The Expander Hierarchy and its Applications to Dynamic Graph Algorithms

Gramoz Goranci¹, Harald Räcke², Thatchaphol Saranurak³, and Zihan Tan⁴

¹University of Toronto, Canada ²TU Munich, Germany ³Toyota Technological Institute at Chicago, USA ⁴University of Chicago, USA

Abstract

We introduce a notion for hierarchical graph clustering which we call the *expander* hierarchy and show a fully dynamic algorithm for maintaining such a hierarchy on a graph with n vertices undergoing edge insertions and deletions using $n^{o(1)}$ update time. An expander hierarchy is a tree representation of graphs that faithfully captures the cut-flow structure and consequently our dynamic algorithm almost immediately implies several results including:

- 1. The first fully dynamic algorithm with $n^{o(1)}$ worst-case update time that allows querying $n^{o(1)}$ -approximate conductance, *s*-*t* maximum flows, and *s*-*t* minimum cuts for any given (s, t) in $O(\log^{1/6} n)$ time. Our results are deterministic and extend to multi-commodity cuts and flows. All previous fully dynamic (or even decremental) algorithms for any of these problems take $\Omega(n)$ update or query time. The key idea behind these results is a fully dynamic algorithm for maintaining a *tree flow sparsifier*, a notion introduced by Räcke [FOCS'02] for constructing competitive oblivious routing schemes.
- 2. A deterministic fully dynamic connectivity algorithm with $n^{o(1)}$ worst-case update time. This significantly simplifies the recent algorithm by Chuzhoy et al. that uses the framework of Nanongkai, Saranurak, and Wulff-Nilsen [FOCS'17].
- 3. A deterministic fully dynamic treewidth decomposition algorithm on constant-degree graphs with $n^{o(1)}$ worst-case update time that maintains a treewidth decomposition of width $\operatorname{tw}(G) \cdot n^{o(1)}$ where $\operatorname{tw}(G)$ denotes the treewidth of the current graph. This is the first non-trivial dynamic algorithm for this problem.

Our technique is based on a new stronger notion of the expander decomposition, called the *boundary-linked expander decomposition*. This decomposition is more robust against updates and better captures clustering structure of graphs. Given that the expander decomposition has proved extremely useful in many fields, including approximation, sketching, distributed, and dynamic algorithms, we expect that our new notion will find more future applications.

Contents

1	Introduction	1	
	1.1 Our Results: The Dynamic Expander Hierarchy	1	
	1.2 Applications	3	
	1.3 Comparison of Techniques	9	
2	Technical Overview	6	
	2.1 Tree Flow Sparsifiers	6	
	2.2 Robustness Against Updates	8	
3	Preliminaries 12		
4	Boundary-Linked Expander Decomposition and Hierarchy	12	
	4.1 The Key Subroutine	14	
	4.2 The Algorithm	17	
	4.3 The Analysis	18	
5	Tree Flow Sparsifier From Expander Hierarchy		
6	Fully Dynamic Expander Pruning2		
7	Fully Dynamic Expander Hierarchy	28	
	7.1 Fully Dynamic Expander Decomposition	29	
8	Derandomization and Deamortization	36	
	8.1 De-randomization	36	
	8.2 De-amortization	37	
9	Applications		
	9.1 Dynamic Tree Flow Sparsifier	40	
	9.2 Dynamic Vertex Flow Sparsifiers, Maximum Flow, Multi-commodity Flow, Multi-		
	Way Cut and Multicut	41	
	9.3 Dynamic Sparsest Cut and Lowest Conductance Cut	43	
	9.4 Dynamic Connectivity	44	
	9.5 Treewidth decomposition	44	
\mathbf{A}	Proof of Lemma 5.1	46	

1 Introduction

Computation on *trees* is usually significantly easier than on *general graphs*. Hence, one of the universal themes in graph algorithms is to compute tree representations that faithfully preserve fundamental properties of a given graph. Examples include spanning forests (preserving connectivity), shortest path trees (preserving distances from a source), Gomory-Hu trees (preserving pairwise minimum cuts), low stretch spanning trees and tree embedding (preserving average distances between pairs of vertices), and treewidth decomposition (preserving "tree-like" structure). Among all known approaches for representing a graph with a tree, the tree flow sparsifier introduced by Räcke [Räc02] is astonishingly strong. Roughly speaking, it is a tree Tthat approximately preserves the values of all cuts of a graph G (see the formal definition in Section 1.2). The existence of such trees, which is far from obvious, already enables competitive oblivious routing schemes with both theoretical [Räc02, Räc08] and practical impact [AC03]. Its polynomial-time construction [HHR03, BKR03] also leads to polynomial-time approximation algorithms for many fundamental problems including minimum bisection, min-max partitioning, k-multicut, etc (see e.g. [AGG⁺09, BFK⁺14, CKS13, Räc08, RS14]). More recently, the almostlinear time construction was shown [RST14] and played a key role in obtaining the celebrated result of approximating maximum flows in near-linear time [She13, KLOS14, Pen16]. Given that the construction for static graphs are now well understood, we raise the challenging question of whether it is possible to maintain tree flow sparsifiers in *dynamic graphs* that undergo a sequence of edge insertions and deletions without recomputing from scratch after each update.

In this paper, we answer this question in affirmative by introducing a new notion for hierarchical graph clustering which we call the *expander hierarchy*. We state a precise definition later in Section 1.1. We show that the expander hierarchy is a tree representation of a graph that is strong enough to imply tree flow sparsifiers (and much more), and yet robust against updates in the sense that it admits fully dynamic algorithms on an *n*-vertex graph maintaining the hierarchy in $n^{o(1)}$ update time.

The fact that tree flow sparsifiers can be maintained efficiently immediately allows us to efficiently compute approximate solutions to a wide range of flow/cut-based problems, including max flows, multi-commodity flows, minimum cuts, multi cuts, multi-way cuts, and conductance. Specifically, for all these problems, this gives the first sub-linear time fully dynamic algorithms with $n^{o(1)}$ worst-case update and query time. The power of the expander hierarchy is not limited to tree flow sparsifiers. It also gives an algorithm for the deterministic dynamic connectivity with $n^{o(1)}$ worst-case update time, that significantly simplifies the recent breakthrough result in [CGL⁺19, NSW17] on this problem. Moreover, it gives the first algorithm for maintaining an approximate treewidth decomposition, a central object in the field of fixed-parameter tractable algorithms. We discuss these applications in detail in Section 1.2.

In short, we introduce the expander hierarchy as a clean combinatorial object that is very robust against adversarial updates, yet strong enough to imply many new results and simplify previous important development. It is likely that future development on dynamic graph algorithms can build on such a hierarchy.

1.1 Our Results: The Dynamic Expander Hierarchy

First, we recall definitions related to expanders. Let G = (V, E) be an *n*-vertex *m*-edge unweighted graph. For any set $S, T \subseteq V$, let $E_G(S, T)$ denote a set of edges between S and T. The volume of S is $\operatorname{vol}_G(S) = \sum_{u \in S} \deg_G(u)$ and we write $\operatorname{vol}(G) = \operatorname{vol}_G(V)$. The conductance of a cut $(S, V \setminus S)$ is $\Phi_G(S) = \frac{|E_G(S, V \setminus S)|}{\min\{\operatorname{vol}_G(S), \operatorname{vol}_G(V \setminus S)\}}$ and the conductance of G is denoted by $\Phi_G = \min_{\emptyset \neq S \subset V} \Phi_G(S)$. We say that G is a ϕ -expander iff $\Phi_G \ge \phi$. We need the following generalized notation for induced subgraphs in order to define our new decomposition. Recall that G[S] denotes the subgraph of G induced by S. For any $w \ge 0$, let $G[S]^w$ be obtained from G[S] by adding $\lceil w \rceil$ self-loops to each vertex $v \in S$ for every boundary edge $(v, x), x \notin S$. Note that $G[S]^0 = G[S]$ and vertices in $G[S]^1$ have the same degree as in the original graph G (each self-loop contributes 1 to the degree of the node that it is incident to).

Stronger Expander Decomposition. The core of this paper is to identify a stronger notion of the well-known *expander decomposition* [KVV04] which states that, given any graph G = (V, E) and any parameter $\phi > 0$, there is a partition $\mathcal{U} = (U_1, \ldots, U_k)$ of V into *clusters*, such that:

- 1. $\sum_{i} |E(U_i, V \setminus U_i)| = \tilde{O}(\phi m).$
- 2. For all $i, G[U_i]$ is a ϕ -expander.

Basically, the decomposition says that one can remove $\tilde{O}(\phi)$ -fraction of edges so that each connected component in the remaining graph is a ϕ -expander. As expanders have many algorithmfriendly properties such as, having low diameter, small mixing time, etc., this decomposition has found numerous applications across areas, including property testing [GR98, KSS18], approximation algorithms [KR96, Tre05], fast graph algorithms [ST11, She13, KLOS14], distributed algorithms [CHKM17, CPZ19, CS19, EFF⁺19], and dynamic algorithms [NS17, Wul17, NSW17].

In this paper, we propose a stronger notion of the expander decomposition. For parameters $\alpha, \phi > 0$, we define an (α, ϕ) -boundary-linked expander decomposition of G as a partition $\mathcal{U} = (U_1, \ldots, U_k)$ of V, together with $\phi_1, \ldots, \phi_k \ge \phi$, such that:

- 1. $\sum_{i} |E(U_i, V \setminus U_i)| = \tilde{O}(\phi m).$
- 2. For all i, $G[U_i]^{\alpha/\phi_i}$ is a ϕ_i -expander.
- 3. For all i, $|E(U_i, V \setminus U_i)| \leq \tilde{O}(\phi_i \operatorname{vol}_G(U_i))$.

Compared to the previous definition in [KVV04], we strengthen Property 2 and additionally require Property 3. Before discussing the power of the new decomposition, we start with intuitive observations how it more faithfully captures the clustering structure of graphs. It is instructive to think of $\alpha = 1/\text{polylog}(n)$ and $\phi = 1/n^{o(1)} \ll \alpha$.

Intuitively, in a good graph clustering, vertices in a cluster are better connected to the inside of the cluster than to its outside. Observe that the stronger form of Property 2 implies that, for every vertex $v \in U_i$, $\deg_{G[U_i]}(v) \ge \alpha \cdot \deg_{G[V \setminus U_i]}(v)$. Without this strengthening, there could be a vertex $v \in U_i$ where $\deg_{G[U_i]}(v) \ll \deg_{G[V \setminus U_i]}(v)$, which is counter-intuitive. Moreover, Property 3 additionally implies that $\deg_{G[V - U_i]}(v) \le \tilde{O}(\phi_i) \cdot \deg_{G[U_i]}(v)$ for most $v \in U_i$. That is, most vertices also have few connection to the outside of the cluster which again matches our intuitive understanding of a good graph clustering.

We say that the decomposition \mathcal{U} has slack $s \geq 1$ if we relax Property 2 as follows: for all i, $G[U_i]^{\alpha/\phi_i}$ is a (ϕ_i/s) -expander. We say that \mathcal{U} has no slack if s = 1.

The Expander Hierarchy. Let \mathcal{U} be an (α, ϕ) -boundary-linked expander decomposition of G. Suppose that we contract each cluster $U_i \in \mathcal{U}$ into a vertex where we keep parallel edges and removes self-loops. Let $G_{\mathcal{U}}$ denote the contracted graph. Observe that $\operatorname{vol}(G_{\mathcal{U}}) =$ $\sum_i |E_G(U_i, V \setminus U_i)| = \tilde{O}(\phi m)$. So $\operatorname{vol}(G_{\mathcal{U}}) \ll \operatorname{vol}(G)$ for small enough ϕ . Repeating the process of decomposition and contraction leads to the following definition. A sequence of graphs (G^0, \ldots, G^t) is an (α, ϕ) -expander decomposition sequence of G if $G^0 = G$, G^t has no edges, and for each $i \geq 0$, there is an (α, ϕ) -boundary-linked expander decomposition \mathcal{U}_i where $G^{i+1} = G^i_{\mathcal{U}_i}$.

Now, observe that the sequence (G^0, \ldots, G^t) naturally corresponds to a tree T where the set of vertices at level *i* of T corresponds to vertices of G^i (i.e. $V_i(T) = V(G^i)$) and if a vertex

Applications	How to obtain from the expander hierarchy T
Tree flow sparsifier	Return T itself.
Vertex cut sparsifier w.r.t. a	Return the union of root-to-leaf paths of T over all vertices
terminal set C	$u \in C$. Denoted it by T_C .
$s-t \max$ flow, $s-t \min$ cut,	Solve the problem on $T_{\{s,t\}}$ or T_C , i.e. the vertex cut sparsifier
multi-commodity cut/flow with	defined in the line above.
a terminal set C	
Conductance and sparsest cut	Implement the Top Tree data structure on T .
Pairwise connectivity	Given u and v , check if the roots of u and of v in T are the
	same.
Treewidth decomposition on	Return T itself. For each level-i node $x \in V_i(T)$ that
constant degree graphs	corresponds to a vertex $u_i \in V(G^i)$, the bag B_x contains the
	original endpoints in G of the edges incident to u_i in G^i .

Table 1: Applications of an expander hierarchy T of depth dep(T), that originates from the expander decomposition sequence $(G^0, \ldots, G^{dep(T)})$ of G. All problems are on unweighted graphs.

 $u_i \in V(G^i)$ is contracted to a super-vertex $u_{i+1} \in V(G^{i+1})$, then we add an edge $(u_i, u_{i+1}) \in T$ with weight $\deg_{G^i}(u_i)$. We call this tree an (α, ϕ) -expander hierarchy, which is the central object of this paper. We say that T has slack s if each \mathcal{U}_i from the (α, ϕ) -expander decomposition sequence has slack at most s.

Our main result shows that an expander hierarchy is robust enough to be maintained under edges updates in subpolynomial update time.

Theorem 1.1. There is an algorithm that, given an n-vertex unweighted graph G undergoing edge insertions and deletions, explicitly maintains, with high probability, a $(1/\operatorname{polylog}(n), 2^{-O(\log^{3/4} n)})$ -expander hierarchy of G with depth $O(\log^{1/4} n)$ and slack $2^{-O(\log^{1/2} n)}$ in $2^{O(\log^{3/4} n)}$ amortized update time. The algorithm works against an **adaptive** adversary.

The algorithm in Theorem 1.1 can be both *derandomized* and *deamortized* with essentially the same guarantee up to subpolynomial factors. We note however that the deamortized algorithm does not explicitly maintain the hierarchy, but supports queries of the following form: given a vertex u of G, return a leaf-to-root path of u in the hierarchy in $O(\log^{1/4} n)$ time.¹

1.2 Applications

The dynamic algorithm for maintaining an expander hierarchy in Theorem 1.1 and its derandomized and deamortized counterpart immediately imply a number of applications in dynamic graph algorithms. Table 1 shows that the high-level algorithm for each application can be described in only one or two sentences.

Below, we discuss the contribution of each application. We say a dynamic algorithm is *fully dynamic* if it handles both edge insertions and deletions. Otherwise, it is *incremental* or *decremental*, meaning that it handles only insertions or only deletions of the edges, respectively.

Tree Flow Sparsifiers. Our first application is the *first* non-trivial fully-dynamic (or even decremental) algorithm for *tree flow sparsifiers*, which will be used to obtain many other applications in the paper. Intuitively, a tree flow sparsifier is a tree that approximately captures the flow/cut structure in a graph. As the formal definition of tree flow sparsifiers is a bit hard

¹This kind of guarantee is similar to the dynamic matching algorithm by [BFH19] with worst-case update time.

to digest, here we define a simpler and an almost equivalent notion of *tree cut sparsifiers*. A *tree cut sparsifier* T of a graph G = (V, E) with quality q is a weighted rooted tree such that (i) the leaves of T corresponds to the vertex set V of G, and (2) for any pair A, B of disjoint subsets of V, mincut_T $(A, B) \leq \text{mincut}_G(A, B) \leq q \cdot \text{mincut}_T(A, B)$ where, for any graph H, mincut_H(A, B) denotes the value of a minimum cut separating A from B in H.

Tree flow sparsifiers have been extensively studied in the static setting [BL99, Räc02, HHR03, BKR03, RS14, RST14] and have found many applications in approximation algorithms [AGG⁺09, BFK⁺14, CKS13, Räc08, RS14] and fast algorithms for computing max-flow [She13, KLOS14]. Currently, the fastest algorithm for computing a tree flow sparsifier takes $\tilde{O}(m)$ time to produce a sparsifier of quality $O(\log^4 n)$ with high probability [RST14, Pen16]. However, only little progress has been obtained in the dynamic setting. Very recently, Goranci, Henzinger, and Saranurak [GHS19] show that, by calling the static algorithm of [RST14] in a blackbox manner, they can obtain an *incremental* algorithm² for maintaining tree flow sparsifiers with $\log^{O(\ell)} n$ quality in $\tilde{O}(n^{1/\ell})$ worst-case update time, for any $\ell > 1$, but their technique inherently could not handle edge deletions. The major reason for this lack of progress in dynamic algorithms is the fact that all existing static constructions for tree flow sparsifiers work in a *top-down* manner, which is difficult to dynamize.

We show that an expander hierarchy is itself a tree flow sparsifier. In particular, this hierarchy implies the first *static* tree flow sparsifier based on a *bottom-up* clustering algorithm and is arguably the conceptually simplest of all known constructions. The key feature of this construction is that it can be maintained dynamically, as summarized in the following corollary.

Corollary 1.2. There is a fully dynamic deterministic algorithm on an unweighted n-vertex graph that explicitly maintains a tree flow sparsifier with quality $n^{o(1)}$ in $n^{o(1)}$ amortized update time.

There has been recent interest in designing dynamic algorithms for maintaining trees that preserve important features or graphs, e.g., distances. One example is the work on dynamic low-stretch spanning trees that achieves sub-polynomial stretch [FG19, CZ20], while in the static there are constructions that give nearly logarithmic stretch [AN12]. Driven by this, our work can be thought as a first step in understanding dynamic algorithms for maintaining trees that preserve the cut/flow structure of graphs.

Flow/Cut-based Problems. Using the above theorem, we improve upon the previous results on a wide-range of dynamic cut and flow problems whose previous fully dynamic (and even decremental) algorithms either require $\Omega(n)$ update time³ or $\Omega(n)$ query time⁴, as summarized in the following corollary.

Corollary 1.3. There is a fully dynamic deterministic algorithm on an unweighted n-vertex graph with $n^{o(1)}$ update time that can return an $n^{o(1)}$ -approximation to queries of the following problems:

- 1. s-t maximum flow, s-t minimum cut;
- 2. lowest conductance cut, sparsest cut; and
- 3. multi-commodity flow, multi-cut, multi-way cut, and vertex cut sparsifiers.

For problems in 1 and 2, the query time is $O(\log^{1/6} n)$, while for problems in 3, it is $O(|C| \log^{1/6} n)$ where C is the terminal set of the respective problem. For problems in 1 and 3, the update time is worst-case, while it is amortized for problems in 2.

 $^{^{2}}$ They also show the same trade-off for the weaker *offline* fully dynamic setting where the whole sequence of updates and queries is given from the beginning.

 $^{^3 \}rm{This}$ includes the incremental exact max flow algorithm by [GK18] and the dynamic conductance algorithm by [vdBNS19].

⁴This is by using dynamic graph sparsifiers [ADK⁺16] and running static algorithms on top of the sparsifier.

Previous sub-linear time algorithms are known only for the incremental setting. There are incremental $\log^{O(\ell)} n$ -approximation algorithms with $O(n^{1/\ell})$ update time for all the problems in Corollary 1.3 [GHS19].

Although our approximation ratio of $n^{o(1)}$ is moderately high, we believe that our results might serve as an efficient building block to $(1+\epsilon)$ -approximation dynamic max flow and minimum cut algorithms, analogous to the previous development in the static setting: A fast static $n^{o(1)}$ approximate (multi-commodity) max flow algorithm was first shown by Madry [Mad10], and later, the $(1 + \epsilon)$ -approximate algorithms were devised [She13, KLOS14, Pen16, She17] by combining Madry's technique with the gradient-descent-based method. Although the gradient-descent-based method in the dynamic setting is currently unexplored, we hope that our result will motivate further investigation in this interesting direction.

Connectivity: Bypassing the NSW Framework. Very recently, Chuzhoy et al. [CGL⁺19] combine their new balanced cut algorithm with the framework of Nanongkai, Saranurak, and Wulff-Nilsen (NSW) [NSW17], and obtain a deterministic dynamic connectivity algorithm with $n^{o(1)}$ worst-case update time, answering a major open problem of the field.

Here, we show a significantly simplified algorithm which completely bypasses the complicated framework by [NSW17]. Our algorithm simply follows from the observation that a graph G is connected iff the top level of an expander hierarchy T of G contains only one vertex. Also, two vertices u and v are connected iff the roots of u and v in T are the same. Interestingly, our algorithm is the first algorithm for dynamic connectivity problem that does *not* explicitly maintain a spanning forest.

Corollary 1.4. There is a fully dynamic deterministic algorithm on an n-vertex graph G that maintains connectivity of G using $n^{o(1)}$ worst-case update time and also supports pairwise connectivity in $O(\log^{1/6} n)$ time.

Treewidth. Computing a *treewidth decomposition* with approximately minimum width is a core problem in the area of fixed-parameter tractable algorithms [RS95, Bod96, Ree92, FM06, Ami10, BDD⁺16, FLS⁺18]. We observe that, on constant degree graphs, an expander hierarchy itself gives a treewidth decomposition. Hence, we obtain the first non-trivial dynamic algorithm for this problem.⁵ Our result is summarized in the following corollary, where tw(G) is the treewidth of a graph G (i.e. the minimum width over all tree decomposition of G).

Corollary 1.5. There is a fully dynamic deterministic algorithm on a constant degree n-vertex graph G that explicitly maintains a treewidth decomposition of width $tw(G) \cdot n^{o(1)}$ in $n^{o(1)}$ amortized update time.

1.3 Comparison of Techniques

The expander hierarchy is strictly stronger than the so-called *low-diameter hierarchy* appeared in the algorithms for constructing low-stretch spanning trees in both the static [AKPW95] and dynamic setting [FG19, CZ20]. The low-diameter hierarchy is similar to the expander hierarchy, except that each cluster is only guaranteed to have low diameter. Structurally, the expander hierarchy is strictly stronger since every ϕ -expander automatically has low diameter $\tilde{O}(1/\phi)$, but some low-diameter graph has very bad expansion. This is why the low-diameter hierarchy could

⁵Dvorak, Kupec, and Tuma [DKT13] show a fully dynamic algorithm for maintaining a *treedepth decomposition*, which is closely related to a treewidth decomposition. Let td(G) denote the *treedepth* of a graph G. It is known that $tw(G) \leq td(G) \leq O(tw(G) \log n)$ [BGHK95]. However, the update time of the algorithm [DKT13] is proportional to a tower of height td(G), which is super-linear when $td(G) = \omega(\log^* n)$. So this algorithm might take super-linear time even when tw(G) = O(1).

not be applied to cut/flow-based problems. Algorithmically, previous approaches for maintaining the low-diameter hierarchy [FG19, CZ20] inherently have amortized update time guarantee and assume an *oblivious* adversary, while our algorithm using the expander hierarchy can be made worst-case and deterministic.

Expander decomposition has almost never been used in a hierarchical manner. Many algorithms perform the decomposition only once [Tre05, CKS05, CKS13] and some recursively decompose the graph by removing all edges inside clusters, instead of contracting each cluster (e.g. [ST11, ACK⁺16, JS18, CPZ19, EFF⁺19] for static algorithms and [BvdBG⁺20, CK19] for dynamic ones). An exception is the sensitivity connectivity oracle by Patrascu and Thorup [PT07], which decomposes a graph into a certain hierarchy of expanders. Unfortunately, their hierarchy is not robust and inherently can handle only a single batch of updates, so it does not work in the standard dynamic setting.

The dynamic connectivity and minimum spanning forest algorithm by Nanongkai et al. [NSW17] repeatedly applies the expander decomposition and has a bottom-up flavor as in ours, but their underlying structure does *not* actually yield a hierarchy. More specifically, while each cluster in our expander hierarchy is contracted into a single vertex in the next level, their cluster can only be "compressed" into a smaller set, which might even be cut through in the next level. This leads to a much more complicated structure and requires an ad hoc treatment. The similar structure appears in the very recent work on dynamic *c*-edge connectivity by Jin and Sun [JS20]. We bypass such complication via the boundary-linked expander decomposition and obtain the simplified dynamic connectivity algorithms and other applications. We expect that our clean hierarchy will be easy to work with and lead to more interesting applications in the future.

Very recently, the concurrent work of Chen et al. $[CGH^+20]$ shows how to dynamize several known constructions of vertex sparsifiers for various problems. One of their applications is a fully dynamic algorithm for *s*-*t* maximum flow and minimum cuts. Their algorithm works against an oblivious adversary, has $\tilde{O}(n^{2/3})$ amortized update time and, given a query, returns an $O(\log n(\log \log n)^{O(1)})$ -approximation in $\tilde{O}(n^{2/3})$ time. They also extend the algorithm to work against an adaptive adversary while supporting updates and queries in $\tilde{O}(m^{3/4})$ time. Comparing with Item 1 from Corollary 1.3, our algorithm has $n^{o(1)}$ worst-case update time and $O(\log^{1/6} n)$ query time and is deterministic, but our approximation factor is worse.

2 Technical Overview

This overview is divided into two parts. In Section 2.1 we show that an expander hierarchy is itself a tree flow sparsifier and faithfully captures the cut/flow structure of a graph. In Section 2.2 we show how to maintain an expander hierarchy under dynamic edge updates. Below, we will write (α, ϕ) -decomposition as a shorthand for (α, ϕ) -boundary-linked expander decomposition.

2.1 Tree Flow Sparsifiers

We start by showing how to construct a tree flow sparsifier for an expander, a very special case. Along the way to generalize the idea to general graphs, we will see how expander hierarchies arise naturally. We will also explain why the approaches based on the standard expander decomposition or even the slightly stronger decomposition from [SW19] fail.

Tree flow sparsifiers can be informally defined as follows. Let G = (V, E) be an *n*-vertex *m*-edge graph. Let $D: V \times V \to \mathbb{R}_{\geq 0}$ be a demand (for multi-commodity flow) between pairs of vertices in *G*. We say that *D* can be routed with congestion η in *G* if there is a multi-commodity flow that routes *D* with congestion η . If $\eta = 1$, we say that *D* is *routable* in *G*. A tree *T* is a

tree flow sparsifier of G with quality q, iff (i) any routable demand in G is routable in T, and (ii) any routable demand in T can be routed with congestion q in G.

Special Case: Expanders. Suppose G is a ϕ -expander. We can construct a very simple tree flow sparsifier T with quality $q_1 = O(\log(m)/\phi)$ as follows. Let T be a star with a root r that connects to each vertex $v \in V$ with an edge (r, v) of capacity $\deg_G(v)$. Observe that any routable demand in G is routable in T because of the way we set the capacities of edges in T. On the other hand, if a demand D is routable in T, then the total demand on each vertex v is at most $\deg(v)$. But it is well-known from the multi-commodity max-flow/min-cut theorem [LR99] that, on ϕ -expanders, any demand D with such a property can be routed with congestion $O(\log(m)/\phi)$.

Intermediate Case: Two Levels of Expanders. Next, we suppose that G satisfies the following: there is a partition $\mathcal{U} = (U_1, \ldots, U_k)$ of V such that $G[U_i]^1$ is a ϕ -expander for each $U_i \in \mathcal{U}$ and the contracted graph $G_{\mathcal{U}}$ is also a ϕ -expander. We write $V(G_{\mathcal{U}}) = \{u_1, \ldots, u_k\}$. We can naturally construct a tree T corresponding to the partition \mathcal{U} as follows. The set of level-0 vertices of T is $V_0(T) = V(G)$. The set of level-1 vertices is $V_1(T) = V(G_{\mathcal{U}})$. The level-2 contains only the root r of T. For each pair $v \in V_0(T)$ and $u_i \in V_1(T)$ such that $v \in U_i$, we add an edge (v, u_i) with capacity $\deg_G(v)$. And each $u_i \in V_1(T)$, we add an edge (u_i, r) with capacity $\deg_{G_{\mathcal{U}}}(u_i) = |E(U_i, V \setminus U_i)|$. We claim that T has quality $q_2 = O((\log(m)/\phi)^2)$.

By the choice of edge capacity in T, any routable demand in G is routable in T. For the other direction, suppose that a demand D is routable in T, then we will show a multi-commodity flow that route D in G with congestion $O((\log(m)/\phi)^2)$. The idea is to first consider the projected demand $D_{\mathcal{U}}$ where $D_{\mathcal{U}}(u_i, u_j) := \sum_{x \in U_i, y \in U_j} D(x, y)$. Using the argument from the first case, there is a multi-commodity flow $F_{\mathcal{U}}$ that routes $D_{\mathcal{U}}$ in $G_{\mathcal{U}}$ with congestion q_1 . Our goal is to extend this flow $F_{\mathcal{U}}$ to another flow that routes D in G with congestion $O((\log(m)/\phi)^2)$.

For each vertex $u_i \in V(G_{\mathcal{U}})$, consider the flow paths in $F_{\mathcal{U}}$ going through u_i in $G_{\mathcal{U}}$. Observe that these paths corresponds to a demand D_{U_i} between *boundary edges* of U_i (i.e. $E(U_i, V \setminus U_i)$). Our main task now is to route D_{U_i} within $G[U_i]$. Once we have obtain the flow F_i within $G[U_i]$ that routes D_{U_i} for all U_i , this would extend $F_{\mathcal{U}}$ to a flow in G as desired, and we are essentially done (some details are omitted here).

As $G[U_i]^1$ is a ϕ -expander, the max-flow/min-cut theorem implies that any demand D'between the boundary edges of U_i can be routed within $G[U_i]$ with congestion $O(\log(m)/\phi)$ as long as the total demand of D' on each edge is at most 1. However, the total demand of D_i on each boundary edge of U_i can be as large as q_1 (since the flow $F_{\mathcal{U}}$ causes congestion q_1). Therefore, D_i can be routed in $G[U_i]$ with congestion $q_1 \cdot O(\log(m)/\phi) = O((\log(m)/\phi)^2)$. As this holds for all U_i , the tree T has quality $O((\log(m)/\phi)^2)$. Note that we crucially exploit the conductance bound on $G[U_i]^1$ and not just on $G[U_i]$. From the above discussion, the standard expander decomposition cannot give the partition \mathcal{U} as we need.

An important observation is that, this quality can be improved if we are further promised that, for each $U_i \in \mathcal{U}$, $G[U_i]^w$ is a ϕ -expander, for some w > 1. This promise implies that we could route a demand D' between the boundary edges of U_i within $G[U_i]$ with congestion $O(\log(m)/\phi)$ as long as the total demand of D' on each edge is at most w (instead of 1), so the demand D_i can be routed in $G[U_i]$ with congestion $q_1 \cdot O(\frac{\log(m)/\phi}{w})$. Therefore, if \mathcal{U} is an (α, ϕ) -decomposition which guarantees that $G[U_i]^{\alpha/\phi}$ is a ϕ -expander, the quality of T will be $q_1 \cdot O(\frac{\log(m)}{\alpha})$.⁶ That is, we lose only a factor of $O(\log(m)/\alpha)$ per level, instead of $O(\log(m)/\phi)$. **General Case.** Now, we are ready to consider an arbitrary graph G. Let (G^0, \ldots, G^t) be such that (i) $G^0 = G$; (ii) $E(G^t) = \emptyset$; and (iii) $G^{i+1} = G^i_{\mathcal{U}}$ for some (α, ϕ) -decomposition \mathcal{U}_i of

⁶Actually, we have a guarantee that $G[U_i]^{\alpha/\phi_i}$ is a ϕ_i -expander for some $\phi_i \ge \phi$. This implies the same bound of $q_1 \cdot \frac{O(\log m/\phi_i)}{w} = q_1 \cdot O(\frac{\log(m)}{\alpha})$.

 G^i . Observe that if we define a tree from (G^0, \ldots, G^t) using the same idea as above, we would exactly obtain an (α, ϕ) -expander hierarchy T of G. Let t be the depth of T. We can argue inductively that T is a tree flow sparsifier of G with quality $O(\frac{\log m}{\phi}) \cdot O(\frac{\log m}{\phi})^{t-1}$.

Note that that $t = O(\log_{1/\phi} m)$ by Property 1 of the (α, ϕ) -decomposition. From Theorem 4.5, we can compute an (α, ϕ) -expander hierarchy where $\alpha = 1/\operatorname{polylog}(n)$ and $\phi = 2^{-O(\log^{1/2} m)}$, this implies that T has quality $O(\frac{\log m}{\phi}) \cdot O(\frac{\log m}{\alpha})^{t-1} = 2^{O(\log^{1/2} m \log \log m)} = n^{o(1)}$. Note that we need $\phi \ll \alpha$ to obtain the quality of $n^{o(1)}$. For example, if $\alpha = \phi$, the quality we obtain would be $(\frac{\log m}{\phi})^t = \Omega(m)$. This is the reason why we cannot use the expander decomposition algorithm by [SW19], because their algorithm only returns a (weaker version of) (ϕ, ϕ) -decomposition.

In the dynamic setting, our algorithm from Theorem 1.1 maintains an (α, ϕ) -expander hierarchy T that has small slack s where $\alpha = 1/\operatorname{polylog}(n)$, $\phi = 1/2^{O(\log^{3/4} m)}$ and $s = 2^{O(\log^{1/2} m)}$. Following the same analysis, the final quality of T degrades slightly to $O(\frac{s\log m}{\alpha})^{t-1} \cdot O(\frac{s\log m}{\phi}) = 2^{O(\log^{3/4} m)} = n^{o(1)}$.

2.2 Robustness Against Updates

In this section, we show how an (α, ϕ) -expander hierarchy T with small slack s can be maintained in $n^{o(1)}$ update time, where $\alpha = 1/\operatorname{polylog}(n)$, $\phi = 2^{-O(\log^{3/4} n)}$, and $s = 2^{O(\log^{1/2} n)}$. Recall from above that T is a tree flow sparsifier with quality $2^{O(\log^{3/4} m)}$.

Our goal is to illustrate that, because of the right kind of guarantees from the (α, ϕ) -decomposition, the algorithm for maintaining the expander hierarchy can be obtained quite naturally, especially for people familiar with standard techniques in dynamic algorithms.

Reduction to One Level. An (α, ϕ) -expander hierarchy T corresponds to an (α, ϕ) -expander decomposition sequence (G^0, \ldots, G^t) . In particular, for each $i \geq 0$, $G^{i+1} = G^i_{\mathcal{U}_i}$ is obtained from G^i by contracting each cluster of an (α, ϕ) -decomposition \mathcal{U}_i . Therefore, the problem of maintaining an expander hierarchy reduces to maintaining an (α, ϕ) -decomposition \mathcal{U} and $G_{\mathcal{U}}$ on a dynamic graph G. There are two important measures:

- Update Time: The time for computing the updated \mathcal{U} and $G_{\mathcal{U}}$.
- **Recourse:** The number of edge updates to $G_{\mathcal{U}}$.

Suppose that there is an algorithm with τ (amortized) update time and ρ (amortized) recourse. This would imply an algorithm for maintaining an (α, ϕ) -expander decomposition sequence (G^0, \ldots, G^t) with $O(\rho^t \cdot \tau)$ (amortized) update time, because the number of updates can be multiplied by ρ per level. We note that the depth $t = O(\log_{1/\phi} m)$ so we need $\rho = (1/\phi)^{o(1)}$ and $\tau = n^{o(1)}$ to conclude that the final update time is $n^{o(1)}$.

Two Key Tools. From now, we focus on a dynamic graph G and how to maintain an (α, ϕ) -decomposition \mathcal{U} of G with small slack. To do this, we need two algorithmic tools. First, Theorem 4.5 gives a static algorithm for computing an (α, ϕ) -decomposition \mathcal{U} of a graph G with no slack in time $\tilde{O}(m/\phi)$. This algorithm strengthens the previous expander decomposition by Saranurak and Wang [SW19].

Our second tool is the new expander pruning algorithm from Theorem 6.1 with the following guarantee. Suppose $G[U]^w$ is a ϕ -expander where $w \leq 1/(10\phi)$. Suppose there is a sequence of $k \leq \phi \operatorname{vol}(U)/2000$ edge updates to U (i.e. these edges have at least one endpoint in U). Then, the algorithm maintains a small set $P \subseteq U$ such that $G[U \setminus P]^w$ is still a $(\phi/38)$ -expander in total time $\tilde{O}(k/\phi^2)$. More precisely, after the *i*-th update to U, we have $\operatorname{vol}_G(P) = O(i/\phi)$ and $|E_G(P, U \setminus P)| = O(i)$. Theorem 6.1 generalizes the previous expander pruning algorithm by [SW19] that works only when w = 1. A Simple Algorithm with Too Large Recourse. Both tools above suggest the following simple approach. First, we compute an (α, ϕ) -decomposition $\mathcal{U} = \{U_1, \ldots, U_k\}$ of G with no slack where each $G[U_i]^{\alpha/\phi_i}$ is a ϕ_i -expander. Next, given an edge update to U_i , we maintain a pruned set $P_i \subseteq U_i$ so that $G[U_i \setminus P_i]^{\alpha/\phi_i}$ is a $(\phi_i/38)$ -expander. We update \mathcal{U} by adding the singleton sets $\{\{u\}\}_{u \in P_i}$ and replacing U_i by $U_i \setminus P_i$.

For simplicity, we assume that all edge updates have at least one endpoint in the cluster $U_1 \in \mathcal{U}$. Furthermore, assume that there are less than $k_1 \ll \phi^2 \operatorname{vol}_G(U_1)$ updates to U_1 . With this assumption, the updated \mathcal{U} is an (α, ϕ) -decomposition of the updated G with slack 38. To see this, observe that the number of new inter-cluster edges is at most $\operatorname{vol}_G(P_1)$, and so the total number of inter-cluster edges becomes $\tilde{O}(\phi m) + \operatorname{vol}_G(P_1) = \tilde{O}(\phi m)$ satisfying Property 1. Let $U'_1 = U_1 \setminus P_1$. We have that $G[U'_1]^{\alpha/\phi_1}$ is a $(\phi_i/38)$ -expander by the expander pruning algorithm. This satisfies Property 2. For Property 3, observe that $|E_G(U'_1, V \setminus U'_1)| \leq |E_G(U_1, V \setminus U_1)| + |E_G(P_1, U_1 \setminus P_1)| = \tilde{O}(\phi_i \operatorname{vol}_G(U_1))$ because $|E_G(P_1, U_1 \setminus P_1)| \leq O(k_1)$. Note that all singleton clusters $\{\{u\}\}_{u \in P_1}$ satisfies Properties 1 and 3 vacuously. Lastly, the recourse on $G_{\mathcal{U}}$ is $O(\operatorname{vol}_G(P_1)) = O(k_1/\phi_1)$.

Therefore, we have that the amortized recourse and update time are $O(1/\phi)$ and $\tilde{O}(1/\phi^2)$ respectively. While both of them seem small for maintaining a one-level decomposition, the recourse is in fact too large if the expander hierarchy has many levels. After composing the algorithm for t levels using the reduction in the beginning of this section, the update time is at least $\Omega(1/\phi)^t = m^{\Omega(1)}$, which is too large for us.

Previous dynamic algorithms with hierarchical structure have faced the same issue. This is why the dynamic low-stretch spanning trees algorithm of [FG19] has $O(\sqrt{n})$ update time. Chechik and Zhang [CZ20] fixed this issue and improved the update time of [FG19] to $n^{o(1)}$. However, they only require each cluster $U \in \mathcal{U}$ to have a small diameter of $\tilde{O}(1/\phi)$, which is a much weaker guarantee than being a $\Omega(\phi)$ -expander. Unfortunately, their technique is specific to this weaker guarantee (and is also inherently amortized). In the dynamic minimum spanning forests algorithm by Nanongkai et al. [NSW17], they require each cluster to be an expander like us, and can only guarantee $\Omega(1/\phi)$ recourse per update. As we mentioned in Section 1.3, they fix the issue using an ad-hoc and complicated tool tailored to (minimum) spanning forests. Below, we will see how the (α, ϕ) -decomposition allows us to bound the recourse in a simple way.

One-batch Updates. First, let us simplify the situation even more by assuming that all k_1 updates are simultaneously given to U_1 in *one batch*. We will also need a slight generalization of (α, ϕ) -decomposition defined on a subset of vertices here, instead of the whole graph. For a set $P \subseteq V$, an (α, ϕ) -decomposition of P in G is a partition $\mathcal{U}' = \{U'_1, \ldots, U'_k\}$ of P that satisfies the properties of the (α, ϕ) -decomposition, except that Property 1 is now $\sum_{i=1}^{k} |E(U'_i, V \setminus U'_i)| \leq O(|E(P, V \setminus P)|) + \tilde{O}(\phi \operatorname{vol}_G(P))$. Note that the term $\tilde{O}(\phi \operatorname{vol}_G(P)) = \tilde{O}(\phi m)$ as before when P = V and the term $O(|E(P, V \setminus P)|)$ is unavoidable.

Now, we describe the algorithm. Given the batch of k_1 updates, we compute the pruned set $P_1 \subseteq U_1$. Then, we compute an (α, ϕ) -decomposition of P_1 in G and obtain a partition $\mathcal{U}' = \{U'_1, \ldots, U'_{k'}\}$ of P_1 . Finally, we replace U_1 in \mathcal{U} with $\{U_1 \setminus P_1, U'_1, \ldots, U'_{k'}\}$. It follows from the description that the updated \mathcal{U} is an (α, ϕ) -decomposition of G with slack at most 38.

The key step is to bound the total recourse, which is at most

$$\sum_{i=1}^{k} |E(U'_i, V \setminus U'_i)| \le O(|E(P_1, V \setminus P_1)|) + \tilde{O}(\phi \operatorname{vol}_G(P_1)).$$

Recall that, by the pruning algorithm, $\operatorname{vol}_G(P_1) \leq O(k_1/\phi)$ and $|E_G(P_1, U \setminus P_1)| \leq O(k_1)$. So it remains to bound $|E_G(P_1, V \setminus U_1)|$. This is where we exploit Property 2 of the (α, ϕ) - decomposition. Before any update, observe that

$$|E_G(P_1, U_1 \setminus P_1)| \ge \phi_1 \operatorname{vol}_{G^{\alpha/\phi_1}[U_1]}(P_1) \ge \phi_1 \cdot \frac{\alpha}{\phi_1} |E_G(P_1, V \setminus U_1)|$$

where the first inequality is because $G[U_1]^{\alpha/\phi_1}$ was a ϕ_1 -expander and we assume $\operatorname{vol}_{G[U_1]^{\alpha/\phi_1}}(P_1) \leq \operatorname{vol}_{G[U_1]^{\alpha/\phi_1}}(U_1 \setminus P_1)$ (as k_1 is small enough), and the second inequality is by the definition of $G[U_1]^{\alpha/\phi_1}$. So we have $|E_G(P_1, V \setminus U_1)| \leq O(k_1/\alpha)$ after the updates, because $|E_G(P_1, U \setminus P_1)| \leq O(k_1)$ and there are k_1 updates. This implies that the total recourse is $O(k_1/\alpha) + \tilde{O}(k_1)$, which is $\tilde{O}(1/\alpha)$ amortized.

Theorem 4.5 shows that the decomposition \mathcal{U}' of P_1 can be computed in $\tilde{O}(|E_G(P_1, V \setminus P_1)|/\phi^2 + \operatorname{vol}_G(P)/\phi) = \tilde{O}(k_1/(\alpha\phi^2))$ time. Also, the total time for pruning is $\tilde{O}(\operatorname{vol}_G(P)/\phi) = \tilde{O}(k_1/\phi^2)$. So the amortized update time is $\tilde{O}(1/(\alpha\phi^2))$. Plugging these bounds into the reduction, the update time for maintaining an (α, ϕ) -expander hierarchy is $\tilde{O}(1/\alpha)^t \cdot \tilde{O}(1/(\alpha\phi^2)) = n^{o(1)}$.

Removing Assumptions. In fact, in the analysis above, we did not require a strong upper bound on $k_1 \ll \phi^2 \operatorname{vol}_G(U_1)$, but we did require $k_1 \leq \phi_1 \operatorname{vol}_G(U_1)/2000$ because the expander pruning algorithm can handle that many updates and we also need $\operatorname{vol}_{G[U_1]^{\alpha/\phi_1}}(P_1) \leq \operatorname{vol}_{G[U_1]^{\alpha/\phi_1}}(U_1 \setminus P_1)$. This requirement can be removed as follows. If $k_1 > \phi_1 \operatorname{vol}_G(U_1)/2000$, then we just "reset" the cluster U_1 by computing an (α, ϕ) -decomposition \mathcal{U}' of U_1 in G. Then, we remove U_1 from \mathcal{U} and add the new clusters in \mathcal{U}' to \mathcal{U} . The key point is again to bound the recourse which is $\sum_{i=1}^k |E_G(U'_i, V \setminus U'_i)| \leq O(|E_G(U_1, V \setminus U_1)|) + \tilde{O}(\phi \operatorname{vol}_G(U_1)) = \tilde{O}(k_1)$. This is where we exploit Property 3 of the (α, ϕ) -decomposition, which says $|E_G(U_1, V \setminus U_1)| = \tilde{O}(\phi_1 \operatorname{vol}_G(U_1)) = \tilde{O}(k_1)$. So the amortized recourse is $\tilde{O}(1)$ in this case.

To remove the assumption that all updates have an endpoint in U_1 , we simply perform the same algorithm on each cluster $U_i \in \mathcal{U}$. If the number of updates is larger than $\phi \operatorname{vol}(G)$, we just compute the (α, ϕ) -decomposition of the updated graph in $\tilde{O}(\operatorname{vol}(G)/\phi)$ time so that Property 1 of the (α, ϕ) -decomposition is satisfied. In this case, the amortized recourse and update time is $\frac{O(\phi \operatorname{vol}(G))}{\phi \operatorname{vol}(G)} = \tilde{O}(1)$ and $\frac{\tilde{O}(\operatorname{vol}(G)/\phi)}{\phi \operatorname{vol}(G)} = \tilde{O}(1/\phi^2)$ respectively.

From all cases above, we conclude that, given an (α, ϕ) -decomposition \mathcal{U} of G with no slack and one batch of updates, we obtain an (α, ϕ) -decomposition \mathcal{U} of G with slack 38 with $\rho = \tilde{O}(1/\alpha)$ amortized recourse and $\tau = \tilde{O}(1/(\alpha\phi^2))$ amortized update time. This implies an algorithm for an (α, ϕ) -expander hierarchy with slack 38 and $O(\rho^t \tau) = 2^{O(\log^{3/4} m)}$ amortized update time.

Sequence of Updates. Lastly, we remove the final assumption that the updates are given in one batch. The main issue arises from the expander pruning algorithms. Recall that, to bound the recourse, it was enough to bound $O(|E_G(P, V \setminus P)|) + \tilde{O}(\phi \operatorname{vol}_G(P))$ where P is the pruned set of some cluster. However, given a sequence of updates to Theorem 6.1, the set P is a dynamic set that changes through time. Although we can bound $O(|E_G(P, V \setminus P)|) + \tilde{O}(\phi \operatorname{vol}_G(P))$ at any point of time, the total recourse throughout the algorithm can become much larger. To fix this, we use the known trick from [NSW17] (Section 5.2.1) so that the resulting pruned set changes in a much more controlled way.

At a very high level, let $\psi = 2^{O(\log^{1/2} m)}$ and $h = \log_{\psi} m = O(\log^{1/2} m)$. Our algorithm, called Multi-level Pruning from Section 7, will partition $P = P_{h-1} \cup \cdots \cup P_0$ into h parts such that, for each i, $\operatorname{vol}_G(P_i) \leq \operatorname{vol}_G(P_{i+1})/\psi$ and P_i can change only every ψ^i updates. In words, the bigger the part, the less often it changes. At the end, Multi-level Pruning has the amortized recourse $\rho = O(2^{O(\log^{1/2} m)}/\alpha)$ and update time $\tau = O(2^{O(\log^{1/2} m)}/\alpha\phi^2)$. This bound still implies a dynamic (α, ϕ) -expander hierarchy with update time $O(\rho^t \tau) = 2^{O(\log^{3/4} m)}$. However, with this technique, the slack of the maintained (α, ϕ) -decomposition become $s = 38^h = 2^{O(\log^{1/2} m)}$, which gives Theorem 1.1. **Derandomization & Deamortization.** The only randomized component in this paper is the cut-matching game [KRV09, RST14] that is used in our static algorithm for computing an (α, ϕ) -decomposition. By plugging in the new deterministic balanced cut algorithm by Chuzhoy et al. [CGL⁺19] into our framework, this immediately derandomizes the whole algorithm.

We can also make our update time to be worst-case using the standard "building in the background" technique (although this technique prevents us from explicitly maintaining the expander hierarchy). The reason we are allowed to do this is as follows. The only component that is inherently amortized is the expander pruning algorithm from Theorem 6.1. However, Theorem 6.1 is only called by Multi-level Pruning, and we can apply the "building in the background" technique to each level of the algorithm, so that the input to Theorem 6.1 is always in one-batch (not a sequence of updates) and so the running time is worst-case. For other parts of the algorithms, it is clear when the algorithm needs to spend a lot of time to reset or re-preprocess the graph, so we can apply "building in the background" in a straight-forward manner.

3 Preliminaries

By default, all logarithms are to the base of 2. Normally we use n to denote the number of nodes of a graph, and use m to denote the number of edges of a graph. Even when we allow parallel edges and self-loops, we will assume in this paper that m = poly(n). We use $\tilde{O}(\cdot)$ to hide polylog(n) factors.

General Notation. Let G = (V, E) be an unweighted graph. For a vertex $v \in V$, we denote $\deg_G(v)$ as the number of edges incident to v in G. For two subsets $A, B \subseteq V$ of vertices, we denote by $E_G(A, B)$ the set of edges with one endpoint in A and the other endpoint in B. For a subset $S \subseteq V$, we denote by $E_G(S)$ the subset of edges of E with both endpoints in S.

To reduce notational clutter we sometimes use the following shorthand notation for the cardinality of certain edge-sets: the cardinality of edges incident to $S \subseteq V$ is denoted with $\operatorname{out}_G(S) := |E_G(S, V \setminus S)|$; the border of $S \subseteq U$ w.r.t. U is denoted with $\operatorname{border}_{G,U}(S) := |E_G(S, V \setminus U)|$; the cut of $S \subseteq U$ w.r.t. U is denoted with $\operatorname{cut}_{G,U}(S) := |E_G(S, U \setminus S)|$. We drop the subscript G if the graph is clear from the context and we write, e.g., $\operatorname{out}_G(v)$ instead of $\operatorname{out}_G(\{v\})$, i.e., we drop the brackets if the respective set contains just a single vertex.

For an unweighted graph G and a cluster $S \subseteq V$ we use $G[S]^w$ to denote the subgraph of G induced by the vertex set S where we add [w] self-loops to a vertex $v \in S$ for every boundary edge $(v, x), x \notin S$ that is incident to v in G. Note that G[S] is just the standard notion of an induced subgraph and that in the graph $G[U]^1$ the degree of all vertices is the same as in the original graph G (each self-loop contributes 1 to the degree of the node that it is incident to).

Let T be a tree and denote r as its root. We denote by L(T) the set of leaves of T. For each $i \ge 0$, we say that a node $v \in V(T)$ is at the *i*th level of T if the length of the unique path connecting v to r in T is i. So the root r is at the 0th level of T, and all its children are at the 1st level of T, and so on. For each $i \ge 0$, we let $V_i(T)$ be the set of all nodes that lie on the *i*th level of the tree T.

Conductance and Expander. For a weighted graph G and a subset $S \subseteq V$ of its vertices, we define the *volume* of S in G to be $\operatorname{vol}_G(S) = \sum_{v \in S} \deg_G(v)$. We refer to a bi-partition (S, \overline{S}) of V by a *cut* of G if both S and \overline{S} are not \emptyset , and we define the capacity of the cut to be $|E(S,\overline{S})|$. The *conductance* of a cut (S,\overline{S}) in G is defined to be $\Phi_G(S) = \frac{|E(S,\overline{S})|}{\min\{\operatorname{vol}_G(S),\operatorname{vol}_G(\overline{S})\}}$. The conductance of a graph G is defined to be $\Phi_G = \min_{S \subseteq V, S \neq \emptyset} \Phi_G(S)$. For a real number $\phi > 0$, we say that G is a ϕ -expander if $\Phi_G \ge \phi$. We will omit the subsript G in the notations above if the graph is clear from the context.

Definition 3.1 (Near Expander). Given an unweighted graph G = (V, E) and a subset of vertices $A \subseteq V$, we say that A is a *near* ϕ -expander in G for some real number $\phi > 0$ if for all $S \subseteq A$ such that $\operatorname{vol}(S) \leq \operatorname{vol}(A)/2$, we have $|E(S, V \setminus S)| \geq \phi \cdot \operatorname{vol}(S)$.

Contracted Graph. Given an unweighted graph G = (V, E) and a partition $\mathcal{U} = (U_1, \ldots, U_r)$ of its vertices, such that the subgraph $G[U_i]$ is connected for each $1 \leq i \leq r$, we define the graph $G_{\mathcal{U}}$ to be the *contracted graph* of G by contracting each cluster U_i to a single vertex, while keeping the parallel edges that connect vertices from the same pair of subsets in \mathcal{U} .

(Single-commodity) Flow Notation. A flow problem $\Pi = (\Delta, T, c)$ on a graph G = (V, E) consists of (i) a source function $\Delta : V \to \mathbb{R}^{\geq 0}$, (ii) a sink capacity function $T : V \to \mathbb{R}^{\geq 0}$, and (iii) an edge capacity function $c : E \to \mathbb{R}^{\geq 0}$. Specifically, for each node $v \in V$, we denote $\Delta(v)$ to be the amount of mass that is placed on v, and we denote T(v) to be the capacity of v as a sink. For an edge e, the capacity c(e) limits how much flow can be routed along e in both directions.

Given a single-commodity flow f on G, we define its edge-formulation by a function f: $V \times V \to \mathbb{R}$, such that for any pair (u, v) of nodes with $(u, v) \in E$, f(u, v) equals the total amount of flow sent from u to v along the edge (u, v) minus the total amount of flow sent from vto u along the edge (u, v). Note that for all pairs (u, v) such that $(u, v) \in E$, f(u, v) = -f(v, u), and f(u, v) = 0 for all pairs (u, v) with $(u, v) \notin E$. We will also refer to an edge-formulation f by a flow. Given a flow problem $\Pi = (\Delta, T, c)$ and a flow f on G, for each node $v \in V$, we define $f_{\Delta}(v) = \Delta(v) + \sum_{u} f(u, v)$ to be the amount of mass ending at v after routing the flow ffrom the initial source function Δ . We say that f is a feasible flow of Π if $|f(u, v)| \leq c(u, v)$ for each edge $(u, v) \in E$, $\sum_{u} f(v, u) \leq \Delta(v)$ for each $v \in V$, and $0 \leq f_{\Delta}(v) \leq T(v)$ for each $v \in V$.

The following subroutine is implicit in [SW19]. They use it as a key subroutine for implementing expander trimming and pruning.

Lemma 3.2 (Incremental Flow). Given an m-edge graph G = (V, E), and a flow problem (Δ, T, c) on G where (i) each edge has integral capacity $1 \le c_e \le c_{\max}$, and (ii) each vertex v can absorb $T(v) = \deg(v)$ mass of flow, there is a deterministic algorithm that maintains an incremental, initially empty set $P \subseteq V$ (i.e., vertices can only join P through time) under a sequence of source-injecting operations of the following form: given $v \in V$, increase $\Delta(v)$.

At any time, as long as $\sum_{v \in V} \Delta(v) \leq \operatorname{vol}(V)/3$, the algorithm guarantees that

- 1. the flow problem (Δ', T', c') on $G[V \setminus P]^1$ is feasible, where $\Delta'(v) = \Delta(v) + c(E(\{v\}, P))$ for all $v \in V$, and T', c' are T, c restricted to $V \setminus P$, respectively, and
- 2. $\operatorname{vol}(P) \le 2 \sum_{v \in V} \Delta(v)$ and $|E(P, V \setminus P)| \le \frac{2 \sum_{v \in V} \Delta(v)}{\min_e \{c_e\}}$.

The total update time is $O(c_{\max} \sum_{v \in V} \Delta(v) \log m)$.

Let us give some intuition about this subroutine. We are given a graph G that undergoes a sequence of "injecting" mass operations, after some time the total mass will not be routable (i.e. the flow problem is not feasible) and get stuck, the above subroutine will maintain a growing set P such that, the mass in the remaining part $G[V \setminus P]^1$ is routable. Moreover, this remains feasible even if we inject additional mass through the cut edges $E(P, V \setminus P)$ at full capacity.

4 Boundary-Linked Expander Decomposition and Hierarchy

In this section we formally introduce the notion of a boundary-linked expander decomposition, which is the main concept of this paper. **Definition 4.1.** (Boundary-Linkedness) For a graph G = (V, E) and parameters $\alpha, \phi \in (0, 1)$ we say that a cluster $U \subseteq V$ is (α, ϕ) -boundary-linked in G if the graph $G[U]^{\alpha/\phi}$ is a ϕ -expander.

Intuitively, the conductance $\Phi_{G[U]^1}$ (i.e., when we choose $\alpha = \phi$) measures how well the edges of the cluster U (including the boundary edges $\Gamma_G(U)$) are connected *inside* the cluster. $\Phi_{G[U]^1} \geq \phi$ means that we can solve an all-to-all multicommodity flow problem between the edges of $E_G(U, V)$ (i.e. edges incident to U) inside G[U] with congestion at most $\tilde{O}(1/\phi)$.⁷ Boundary linkedness with a parameter $\alpha \gg \phi$, means that the boundary edges themselves have higher connectivity. We can solve an all-to-all multicommodity flow problem between boundary-edges with congestion $\tilde{O}(1/\alpha)$ inside U.

Next, we define the notion of a boundary-linked expander decomposition and that of an expander hierarchy. These are the central definitions in this paper.

Definition 4.2 (Boundary-Linked Expander Decomposition). Let G = (V, E) be a graph and $\alpha, \phi \in (0, 1)$ be parameters. Let $U \subseteq V$ be a cluster in G.

An (α, ϕ) -boundary-linked expander decomposition of U in G with slack $s \ge 1$ consists of a partition $\mathcal{U} = \{U_1, \ldots, U_k\}$ of U together with a conductance-bound $\phi_i \ge \phi$ for every $i \in 1, \ldots, k$ such that the following holds:

- 1. $\sum_{i=1}^{k} \operatorname{out}_{G}(U_{i}) \leq O(\operatorname{out}_{G}(U)) + \tilde{O}(\phi \operatorname{vol}_{G}(U)).$
- 2. For all *i*: $G[U_i]^{\alpha/\phi_i}$ is a (ϕ_i/s) -expander.
- 3. For all i: $\operatorname{out}_G(U_i) \leq \tilde{O}(\phi_i \operatorname{vol}_G(U_i))$.

When we have an expander-decomposition with slack 1, we will usually not mention the slack and just call it an (α, ϕ) -boundary linked expander decomposition. The notion of slack will not be important for our static constructions but only becomes important for maintaining boundary-linked expander decompositions dynamically. Instead of writing " (α, ϕ) -boundary-linked expander decomposition", we sometimes write " (α, ϕ) -expander decomposition" or just " (α, ϕ) -ED". If U = V, then we say that \mathcal{U} is an (α, ϕ) -ED of G.

Definition 4.3 (Expander Decomposition Sequence). Let G = (V, E) be a graph with m edges and $\alpha, \phi \in (0, 1)$ be parameters. We say that a sequence of graphs (G^0, G^1, \ldots, G^t) is an (α, ϕ) -expander decomposition sequence of G with slack s or (α, ϕ) -ED-sequence of G if (1) $G^0 = G$, (2) G^t has no edge, and (3) $G^{i+1} = G^i_{\mathcal{U}^i}$ is the contracted graph of G^i w.r.t. to \mathcal{U}^i where \mathcal{U}^i is an (α, ϕ) -ED of G^i with slack s.

Definition 4.4 (Expander Hierarchy). An (α, ϕ) -ED sequence (G^0, G^1, \ldots, G^t) naturally corresponds to a tree T where (1) the set of nodes at level i of T is $V_i(T) = V(G^i)$ and (2) a node $u_i \in V_i(T)$ has a parent $u_{i+1} \in V_{i+1}(T)$ if $u_i \in V(G^i)$ is contracted into the supervertex $u_{i+1} \in V(G^{i+1})$. The edge (u_i, u_{i+1}) is assigned a capacity of $\deg_{G^i}(u_i)$. We call T an (α, ϕ) -expander hierarchy or (α, ϕ) -EH (with slack s).

The next theorem is the main result in this section. Throughout this section, we define $\gamma_{\text{CMP}} = O(\log^2 m)$. This is a value derived from the approximation guarantee for sparsest cut of the cut-matching-game [KRV09] on an *m*-edge graph.

Theorem 4.5. There is a randomized algorithm that, given a graph G = (V, E), a cluster $U \subseteq V$ with $\operatorname{vol}_G(U) = m$ and $\operatorname{out}_G(U) = b$, and parameters α, ϕ , with $\alpha \leq 1/(4\gamma_{CMP} \log_2 m)$ computes an (α, ϕ) -ED of U in $\tilde{O}(b/\phi^2 + m/\phi)$ time with high probability. In particular

⁷In an all-to-all multicommodity flow problem between a subset of edges E' in G[U], there is a weight $w(v) := \sum_{e \in E'} |\{v\} \cap e|$ assigned to every vertex $v \in U$. Then the demand between two vertices $u, v \in U$ is w(u)w(v)/w(U).

- 1. $\sum_{i=1}^{k} \operatorname{out}_{G}(U_{i}) \leq 4 \operatorname{out}_{G}(U) + O(\log^{3} m \cdot \phi \operatorname{vol}_{G}(U)).$
- 2. For all i: $G[U_i]^{\alpha/\phi_i}$ is a ϕ_i -expander.
- 3. For all i: $\operatorname{out}_G(U_i) \leq O(\log^6 m \cdot \phi_i \operatorname{vol}_G(U_i)).$

As an (α, ϕ) -ED-sequence and its corresponding (α, ϕ) -expander hierarchy can be naturally computed bottom up given the above algorithm, we immediately get the following corollary.

Corollary 4.6. There is a randomized algorithm that, given a graph G with m edges and parameters α, ϕ , with $\alpha \leq 1/(2\gamma_{\text{CMP}} \log_2 m)$ computes an (α, ϕ) -expander decomposition sequence of G and its corresponding expander hierarchy in $\tilde{O}(m/\phi)$ time with high probability.

The remainder of this section is devoted to proving Theorem 4.5.

4.1 The Key Subroutine

The major building block for our algorithm is the following sub-routine that when applied to a cluster either (1) finds a sparse balanced cut, or (2) finds a sparse unbalanced cut such that the large side of the cut forms a cluster with good expansion. The sub-routine uses the cut matching game due to Khandekar, Rao, and Vazirani[KRV09] and adds a pruning step ([SW19]) for the case that the cut-matching step returns a very unbalanced cut. The Pruning step is the same as in [SW19] but here we give a different analysis that shows a stronger guarantee.

Lemma 4.7 (Cut-Matching + Trimming). Given an unweighted graph G = (V, E) with m edges and parameters ϕ , w with $w < 1/(8\phi)$, a cut-matching+trimming step runs in time $O(m \log m/\phi)$ and must end in one of the following two cases:

- 1. We find a cut (A, \bar{A}) of G with $cut_G(A, \bar{A}) \leq \gamma_{CMP} \cdot \phi \min\{\operatorname{vol}_G(A), \operatorname{vol}_G(\bar{A})\}$, and $\operatorname{vol}_G(A)$, $\operatorname{vol}_G(\bar{A})$ are both $\Omega(m/\log^2 m)$, i.e., we find a relatively balanced low conductance cut.
- 2. We find a cut (A, \overline{A}) , with $\operatorname{cut}_G(A, \overline{A}) \leq \gamma_{\text{CMP}} \cdot \phi \min\{\operatorname{vol}_G(A), \operatorname{vol}_G(\overline{A})\}$, and $\operatorname{vol}_G(\overline{A}) = m/10$. Moreover, we conclude that $G[A]^w$ is a ϕ -expander. This conclusion may be wrong with probability $o(m^{-10})$.

In the remainder of this section we prove the above theorem.

We use a standard adaptation by Saranurak and Wang [SW19] of the *cut-matching framework*, which was originally proposed by Khandekar, Rao and Vazirani [KRV09]. The following *cut-matching step* was proved in [SW19].

Lemma 4.8 (Adapted Statement of Theorem 2.2 in [SW19]). Given an unweighted graph G = (V, E) with m edges and a parameter $\phi > 0$, the cut-matching step takes $O(m \log m/\phi)$ time and must end with one of three cases:

- 1. We conclude that G has conductance $\Phi_G \ge 8\phi$. This conclusion is wrong with probability $o(m^{-10})$.
- 2. We find a cut (A, A) of G with conductance $\Phi_G(A) \leq \gamma_{\text{KRV}}\phi$, and $\text{vol}_G(A), \text{vol}_G(A)$ are both at least $m/(100\gamma_{\text{KRV}})$, i.e., we find a relatively balanced low conductance cut.
- 3. We find a cut (A, \overline{A}) , such that $\Phi_G(A) \leq \gamma_{\text{KRV}}\phi$ and $\operatorname{vol}_G(\overline{A}) \leq m/(100\gamma_{\text{KRV}})$. Moreover, we conclude that A is a near 8ϕ -expander. This conclusion may be wrong with probability $o(m^{-10})$.

Here, $\gamma_{\text{KRV}} = O(\log^2 m)$.

Recall the definition of near expanders from Definition 3.1. We remark that this is the only building block that is randomized in our algorithms. Once we derandomize it, all our algorithms become deterministic. In fact, in a recent paper [CGL⁺19], a deterministic counterpart of the cut-matching step was constructed. We will use their results and roughly show how to make our algorithms deterministic in Section 8.

In order to obtain Lemma 4.7, we proceed as follows. We run the cut-matching step from Lemma 4.8 on the graph G. If we are in Case 1 of Lemma 4.8 we obtain a valid set A for Case 2 in Lemma 4.7, where \overline{A} is the empty set. If we are in Case 2 we obtain valids set A, \overline{A} for Case 1 in Lemma 4.7. If we are in Case 3 we perform a *trimming operation* on the set A to obtain a set A'. We will need to prove that the set A' fulfills all properties required for Lemma 4.7.

The trimming operation (stated below in Lemma 4.9) is algorithmically exactly the same as in [SW19]. The only difference is in the analysis; we open their black-box and state the guarantee about flow explicitly. Then, we give a new analysis and conclude a stronger statement than the one in [SW19]. More precisely, we show that $G[A']^w$ is a ϕ -expander while they only show that $G[A']^1$ is a ϕ -expander.

Lemma 4.9 (Trimming). We can compute a pruned set $P \subset A$ in time $O(\log m |E_G(A, \overline{A})|/\phi^2)$ with the following properties:

- 1. $\operatorname{vol}_G(P) \leq \frac{4}{\phi} |E_G(A, \bar{A})|$
- 2. $|E_G(A', \bar{A}')| \le 2|E_G(A, \bar{A})|$

where $A' = A \setminus P$. In addition the following flow problem is feasible in G[A'].

- $\Delta(v) = \frac{2}{\phi} |E_G(\{v\}, V \setminus A')|$
- $T(v) = \operatorname{vol}_G(v)$
- $c(e) = 2/\phi$ for every edge in G[A'].

Proof. We run the algorithm from Lemma 3.2 on $G[A]^1$ with $c(e) = 2/\phi$ for every edge. Then we increase $\Delta(v)$ by $2/\phi$ for every edge in $E_G(A, \overline{A})$.

The resulting pruned set P fulfills the properties. Property 1 follows as

$$\operatorname{vol}_{G[A]^1}(P) = \operatorname{vol}_G(P) \le 2\sum_v \Delta(v) = \frac{4}{\phi} |E_G(A, \bar{A})|$$
 .

We have to argue that $A' = A \setminus P$ fulfills all requirements of set A in Case 2 of Lemma 4.7.

• $\Phi_G(A') \leq \gamma_{\text{CMP}}\phi$.

The conductance of the cut is $|E_G(A', \bar{A}')| / \operatorname{vol}_G(\bar{A}')$. We have

$$\operatorname{vol}_G(\bar{A}') \ge \operatorname{vol}_G(\bar{A}) \ge \frac{1}{\gamma_{\operatorname{KRV}}\phi} |E_G(A,\bar{A})| \ge \frac{1}{2\gamma_{\operatorname{KRV}}\phi} |E_G(A',\bar{A}')|$$

Hence, setting $\gamma_{\text{CMP}} = 2\gamma_{\text{KRV}} = O(\log^2 n)$ is sufficient.

• $\operatorname{vol}_G(\bar{A}') \le m/10.$

$$\begin{aligned} \operatorname{vol}_G(A') &\leq \operatorname{vol}_G(P) + \operatorname{vol}_G(A) \leq \frac{4}{\phi} |E_G(A, A)| + \operatorname{vol}_G(A) \leq \frac{4}{\phi} \gamma_{\operatorname{KRV}} \phi \operatorname{vol}_G(A) + \operatorname{vol}_G(A) \\ &\leq 5 \gamma_{\operatorname{CMP}} \operatorname{vol}_G(\bar{A}) = m/10 \end{aligned}$$

The final property is given by the following lemma.

Lemma 4.10. If $w \leq 1/(8\phi)$ then $G[A']^w$ is a ϕ -expander.

Proof. Directly from the guarantee of the cut-matching step from Lemma 4.8, we get that

$$|E_G(A,\bar{A})| \le \gamma_{\rm KRV}\phi \operatorname{vol}_G(\bar{A}) \le \phi m/16 \le \phi \operatorname{vol}_G(A)/16.$$
(1)

Now, consider a subset $S \subseteq A'$ such that $\operatorname{vol}_{G[A']^w}(S) \leq \operatorname{vol}_{G[A']^w}(A')/2$. We first show a helpful claim:

Claim 4.11. $\operatorname{vol}_G(S) \leq \frac{2}{3} \operatorname{vol}_G(A')$.

Proof. By the guarantee of the trimming operation from Lemma 4.9 and Equation (1), $\operatorname{vol}_G(P) \leq \frac{4}{\phi} \cdot |E_G(A,\bar{A})| \leq \frac{4}{\phi} \cdot \frac{\phi}{16} \operatorname{vol}_G(A) \leq \operatorname{vol}_G(A)/4$. So $\operatorname{vol}_G(A') = \operatorname{vol}_G(A) - \operatorname{vol}_G(P) \geq 3 \operatorname{vol}_G(A)/4$. Again by Lemma 4.9 and Equation (1), we have $|E_G(A',\bar{A'})| \leq 2|E_G(A,\bar{A})| \leq 2 \cdot \frac{\phi}{16} \operatorname{vol}_G(A) \leq \frac{\phi}{6} \operatorname{vol}_G(A')$. We get

$$\operatorname{vol}_{G}(S) \leq \operatorname{vol}_{G[A']^{w}}(S) \leq \operatorname{vol}_{G[A']^{w}}(A')/2 = \frac{1}{2} (\operatorname{vol}_{G}(A') + (w-1)|E_{G}(A',A')|)$$

$$\leq \frac{1}{2} (\operatorname{vol}_{G}(A') + \frac{1}{8\phi} \cdot \frac{\phi}{6} \operatorname{vol}_{G}(A')) \leq \frac{2}{3} \operatorname{vol}_{G}(A') .$$

The equality holds because edges between A' and $\bar{A}' = V \setminus A'$ are turned into self-loops of weight w in $G[A']^w$ while having weight 1 in G. Hence, the degree of a vertex in A' incident to such an edge increases by w - 1. The following inequality uses $w - 1 \leq 1/(8\phi)$ and our previous bound on $|E_G(A', \bar{A}')|$.

Recall that $\operatorname{border}_{A'}(S) := |E_G(S, V \setminus A')|$ and $\operatorname{cut}_{A'}(S) := |E_G(S, A' \setminus S)|$. We have to show that $\operatorname{cut}_{A'}(S) \ge \phi \cdot \operatorname{vol}_{G[A']^w}(S)$. From $\operatorname{vol}_G(S) \le \frac{2}{3} \operatorname{vol}_G(A')$ we get $\operatorname{vol}_G(A' \setminus S) = \operatorname{vol}_G(A') - \operatorname{vol}_G(S) \ge \frac{3}{2} \operatorname{vol}_G(S) - \operatorname{vol}_G(S) \ge \operatorname{vol}_G(S)/2$. The fact that A is a near 8 ϕ -expander in G gives that

$$\operatorname{cut}_{A'}(S) + \operatorname{border}_{A'}(S) \ge 8\phi \cdot \min\{\operatorname{vol}_G(S), \operatorname{vol}_G(A \setminus S)\} \\ \ge 8\phi \cdot \min\{\operatorname{vol}_G(S), \operatorname{vol}_G(A' \setminus S)\} \\ \ge 4\phi \cdot \operatorname{vol}_G(S) .$$

$$(2)$$

By the feasibility of the flow problem for G[A'] we obtain

$$\begin{aligned} \frac{2}{\phi} \cdot \operatorname{border}_{A'}(S) &\leq \Delta(S) \leq T(S) + \frac{2}{\phi} \operatorname{cut}_{A'}(S) \\ &= \operatorname{vol}_G(S) + \frac{2}{\phi} \operatorname{cut}_{A'}(S) \\ &\leq \frac{1}{4\phi} \operatorname{border}_{A'}(S) + \left(\frac{1}{4\phi} + \frac{2}{\phi}\right) \operatorname{cut}_{A'}(S) \ , \end{aligned}$$

which yields $\operatorname{border}_{A'}(S) \leq 9/7 \cdot \operatorname{cut}_{A'}(S) \leq 2 \operatorname{cut}_{A'}(S)$. Here, the first inequality is due to the fact that the flow problem injects $2/\phi$ units of flow for every border edge. The second inequality follows because the total flow that can be absorbed at the vertices of S is at most T(S) and the flow that can be send to $A' \setminus S$ is at most $\frac{2}{\phi} \operatorname{cut}_{A'}(S)$ as each edge has capacity $2/\phi$. The final step uses Equation 2.

Finally, we obtain

$$\frac{1}{\phi}\operatorname{cut}_{A'}(S) \geq \frac{1}{4\phi}(\operatorname{cut}_{A'}(S) + 2\operatorname{cut}_{A'}(S)) + \frac{1}{4\phi}\operatorname{cut}_{A'}(S)$$
$$\geq \frac{1}{4\phi}(\operatorname{cut}_{A'}(S) + \operatorname{border}_{A'}(S)) + \frac{1}{8\phi}\operatorname{border}_{A'}(S)$$
$$\geq \operatorname{vol}_G(S) + w\operatorname{border}_{A'}(S) \geq \operatorname{vol}_{G[A']^w}(S) ,$$

as desired.

Running time. The running time of the cut-matching step from Lemma 4.8 is $O(m \log m/\phi)$. The running time of the trimming step from Lemma 4.9 is $O(\log m |E_G(A, \bar{A})|/\phi^2) = O(m \log m/\phi)$ by Equation (1). Hence, the total running time is $O(m \log m/\phi)$. **input** :graph G = (V, E), cluster $U \subseteq V$, parameters α, ϕ **output**: partition $\mathcal{U} = (U_1, \ldots, U_k)$ of U, expansion bounds ϕ_1, \ldots, ϕ_k define \mathcal{U} to contain only U as an active cluster; while \exists active sub-cluster in \mathcal{U} do $\varphi \leftarrow \max\left\{\frac{1}{8\gamma_{\text{CMP}}\log_2^2 m} \cdot \sum_{\text{act. } i} \operatorname{out}(U_i) / \sum_{\text{act. } i} \operatorname{vol}(U_i) , \phi\right\};$ for $U_i \in \mathcal{U}$ do expands $(U_i, \varphi) \leftarrow$ false; while \exists active cluster U_i with expands $(U_i, \varphi) =$ false do apply cut-matching + trimming from Lemma 4.7 to $G[U_i]^{\alpha/\varphi}$; **case 1**: replace U_i by active sets A and \overline{A} in \mathcal{U} ; **case 2**: expands $(A, \varphi) \leftarrow$ true; // A is φ -expanding, w.h.p. replace U_i by active sets A and \overline{A} in \mathcal{U} ; end for every active set $U_i \in \mathcal{U}$ do $\phi_i \leftarrow \varphi;$ // Set expansion ... deactivate $U_i;$ // U_i fulfills Property 2 (w.h.p.) and 3 end end end

Algorithm 1: An algorithm to compute an (α, ϕ) -expander decomposition.

4.2 The Algorithm

The algorithm maintains an expansion parameter φ and a partitioning \mathcal{U} that initially just contains the set U (i.e., the cluster $U \subseteq V$ on which we startet the algorithm) as an active cluster. Recall that $\operatorname{vol}_G(U) = m$. Then the algorithm proceeds in rounds, where a round is an iteration of the outer while loop. During a round the algorithm tries to certify for all active clusters U_i in \mathcal{U} that $G[U_i]^{\alpha/\varphi}$ is φ -expanding. For this it uses the cut-matching+trimming algorithm from Lemma 4.7 with parameter φ on the graph $G[U_i]^{\alpha/\varphi}$. From Lemma 4.7, there are two possible outcomes:

Case 1. The framework finds a sparse fairly balanced cut (A, \overline{A}) . Then the algorithm just replaces U_i by A and \overline{A} in \mathcal{U} .

Case 2. The framework finds an unbalanced cut (A, \overline{A}) and concludes that the larger part A forms a φ -expander. Then the algorithm replaces U_i by A and \overline{A} and remembers the conclusion that A is expanding, i.e., the algorithm will not work on A again during a round.

After the algorithm has determined that w.h.p. all active clusters in \mathcal{U} are φ -expanding it checks for every cluster whether Property 3 from Definition 4.2 of the boundary-linked expanderdecomposition holds. If this is the case for a cluster U_i the algorithm sets ϕ_i to the current value of φ and deactivates the cluster.

The algorithm then proceeds to the next round (possibly increasing φ) and continues until no active clusters are left. Algorithm 1 gives an overview of the algorithm.

4.3 The Analysis

In the following we assume that all conclusions by the algorithm that are correct with high probability are indeed correct.

It is clear that when the algorithm terminates all clusters fulfill Property 2 and Property 3, i.e., we only have to prove that the partition \mathcal{U} fulfills Property 1 and that the algorithm indeed terminates.

Let for a subset $X \operatorname{ivol}_G(X) := \sum_{x \in X} \operatorname{out} \{x\}$ denote the *internal volume* of the set, i.e., the part of the volume that is due to the edges for which both endpoints are in X. In the following the notation $\operatorname{cut}_U(A)$, $\operatorname{border}_U(A)$, and $\operatorname{out}(A)$ are always w.r.t. the graph G. Further we use $Z := \log_2(\operatorname{vol}(U))$ as a shorthand notation. For $x \in V$ we use $\operatorname{cluster}(x)$ to denote the cluster from the partion \mathcal{U} that x is contained in. If $x \notin U$ then this evaluates to undefined.

In order to derive a bound on $\sum_{i=1}^{k} \operatorname{out}(U_i)$ we proceed as follows. We distribute an initial charge to the edges incident to vertices in U. Whenever we *cut* edges, i.e., we partition a subset U_i into A and \overline{A} we redistribute charge to the edges in the cut. In the end we compare the charge on edges leaving sub-clusters to the initial charge and thereby obtain a bound on $\sum_i \operatorname{out}(U_i)$. In addition we will redistribute charge whenever we adjust the value of φ in the beginning of a round. However, importantly we will never increase the total charge, hence, in the end we can derive a bound on the number of cut-edges by comparing the charge on a cut-edge to the total *initial charge*.

We call one iteration of the outer while-loop a round of the algorithm. Let R denote an upper bound on the number of rounds. Later we will show that $R \leq \log_2 m$. For any round $r \leq R$, we maintain the following invariant concerning the distribution of charge on the edges that have at least one end-point from the set U:

border edges

An edge (u, v) for which not both endpoints are in the same sub-cluster of \mathcal{U} is assigned a charge of $f_B(r)(Z + \log_2(\text{ivol}(\text{cluster}(x))))$ for each end-point $x \in \{u, v\} \cap U$. We call a charge *active* if it comes from an endpoint within an active cluster. This means that an edge could be assigned active as well as inactive charge.

internal edges

An edge (u, v) for which both endpoints are in the same sub-cluster U_i of \mathcal{U} is assigned an active charge of $f_I(r) \log_2(ivol(U_i))$ if this cluster is active. Otherwise, it is assigned a charge of 0.

We refer to the charge on border edges as *border charge* and to the charge on internal edges as *internal charge*. The factors $f_I(r)$ and $f_B(r)$ in the above definition depend on the round and are chosen as follows:

$$f_B(r) = 2R - r$$

$$f_I(r) = 4\gamma_{\rm CMP} Z f_B(r)\varphi(r).$$

Recall the parameter $\gamma_{\text{CMP}} = O(\log^2 m)$ from the cut-matching + trimming algorithm in Lemma 4.7. When we call the algorithm from Lemma 4.7 with conductance parameter φ then the non-empty cuts returned in have conductance at most $\gamma_{\text{CMP}}\varphi$. Below, let $\varphi(r)$ denote the value of φ during round r. For technical reasons we also introduce a round r = 0, which is the beginning of the algorithm. We set $\varphi(0) = \phi$. This gives that the total *initial charge* is

$$initial-charge = f_B(0) \cdot b \cdot (Z + \log_2(ivol(U))) + 4\gamma_{CMP}Zf_B(0)\varphi(0) \cdot ivol(U)\log_2(ivol(U))$$

$$\leq 4RZ \cdot b + 8\gamma_{CMP}RZ^2 \cdot \phi m.$$

The following claim gives Property 1 provided that we can establish the above charge distribution without generating new charge during the algorithm.

Claim 4.12. Suppose that no charge is generated during the algorithm. Then at the end of the algorithm $\sum_{i} \operatorname{out}(U_i) \leq 4b + O(\log^3 m) \cdot \phi m$, i.e., Property 1 holds.

Proof. Observe that in the end every inter-cluster edge will have a charge of at least $f_B(R)Z \ge RZ$. Therefore $\sum_i \operatorname{out}(U_i) \le \frac{1}{RZ}$ initial-charge $\le 4b + 8\gamma_{\text{CMP}}Z \cdot \phi m = 4b + O(\log^3 m) \cdot \phi m$. \Box

Observe that the number of rounds performed by our algorithm is not important for the above claim. This is only important for the running time analysis.

No charge increase during a round. The main task of the analysis is to establish the charging scheme and to show that we do not generate charge throughout the algorithm. We first show that we do not generate charge *during* a round.

Suppose Lemma 4.7 finds a cut of conductance at most $\gamma_{\text{CMP}} \cdot \varphi(r)$ within the graph $H := G[U_i]^{\alpha/\varphi(r)}$. This means we have a set $S \subseteq U_i$ with

$$\operatorname{cut}_{U_i}(S) < \gamma_{\operatorname{CMP}} \cdot \varphi(r) \cdot \min\{\operatorname{vol}_H(S), \operatorname{vol}_H(U_i \setminus S)\}$$
.

W.l.o.g. assume that $ivol(S) \leq ivol(U_i \setminus S)$. Then

$$\operatorname{cut}_{U_i}(S) < \gamma_{\operatorname{CMP}} \cdot \varphi(r) \operatorname{vol}_H(S) = \gamma_{\operatorname{CMP}} \cdot \varphi(r) \cdot (\operatorname{ivol}(S) + \frac{\alpha}{\varphi(r)} \operatorname{border}_{U_i}(S))$$

By performing the cut, every edge that contributes to $\operatorname{border}_{U_i}(S)$ reduces its required charge by at least $f_B(r)$ because one of its endpoints reduces the internal volume of its cluster by a factor of 2. A similar argument holds for the edges with both endpoints in S, which reduce their required charge by at least $f_I(r)$. This means we obtain a charge of at least

$$f_B(r) \cdot \operatorname{border}_{U_i}(S) + f_I(r) \cdot \operatorname{ivol}(S)$$
 (3)

that we can redistribute to the edges in the cut so that these fulfill their increased charge requirement. The new charge for the cut edges (i.e., edges in $\Gamma_G(S, U_i \setminus S)$) is at most

$$f_B(r)(Z + \log_2(\operatorname{ivol}(S))) + f_B(r)(Z + \log_2(\operatorname{ivol}(U_i \setminus S))) \le 4f_B(r)Z$$

where $Z = \log_2(\text{vol}(U))$. This means the new required charge is

$$\begin{aligned} 4f_B(r)Z \cdot \operatorname{cut}_{U_i}(S) &< 4\gamma_{\scriptscriptstyle \mathrm{CMP}} f_B(r)Z \cdot (\varphi(r)\operatorname{ivol}(S) + \alpha\operatorname{border}_{U_i}(S)) \\ &\leq 4\gamma_{\scriptscriptstyle \mathrm{CMP}} \alpha Z f_B(r)\operatorname{border}_{U_i}(S) + 4\gamma_{\scriptscriptstyle \mathrm{CMP}} \varphi(r)Z f_B(r)\operatorname{ivol}(S) \\ &\stackrel{!}{\leq} f_B(r) \cdot \operatorname{border}_{U_i}(S) + f_I(r) \cdot \operatorname{ivol}(S) \ , \end{aligned}$$

where we want to ensure the last inequality so that the new charge on cut-edges is at most the charge that we have for redistribution according to Equation 3. We ensure the last inequality by requiring that

$$4\gamma_{\rm CMP}\alpha Z \leq 1$$

as a precondition of the theorem and by setting

$$f_I(r) := 4\gamma_{\rm CMP} f_B(r) Z\varphi(r).$$

This shows that we can redistribute enough charge to border edges and the total charge does not increase. No charge increase between rounds. Let I denote the index set of active clusters at the start of round r. At the beginning of a round all border edges decrease their charge as $f_B(r)$ decreases. If we choose $\varphi(r) = \phi$ then the charge on internal edges does not increase because in the previous round we had $\varphi(r-1) \ge \phi$ and the charge on an internal edge is increasing with φ . Hence, we only need to consider the case if $\varphi(r)$ is chosen larger than ϕ and hence

$$\varphi(r) = \frac{1}{8\gamma_{\text{CMP}}RZ} \cdot \sum_{i \in I} \operatorname{out}(U_i) / \sum_{i \in I} \operatorname{vol}(U_i).$$

For this case we show that the decrease in charge on active border edges is more than the increase in charge on internal edges. This is sufficient as only active internal edges increase their charge. Every edge in the boundary of an active cluster U_i decreases its charge by at least $(f_B(r-1) - f_B(r))Z \ge Z$. This means the border charge decreases by at least $Z \sum_{i \in I} \operatorname{out}(U_i)$.

What is the total internal charge? Every edge inside an active cluster U_i is assigned a charge of $f_I(r) \log_2(ivol(U_i))$. Recall that we set internal charge of inactive cluster to be zero. Therefore, the total *internal charge* is

$$internal-charge = f_{I}(r) \sum_{i \in I} ivol(U_{i}) \log_{2}(ivol(U_{i}))$$

$$\leq 4\gamma_{CMP} f_{B}(r) Z^{2} \cdot \varphi(r) \cdot \sum_{i \in I} ivol(U_{i})$$

$$\leq 8\gamma_{CMP} R Z^{2} \cdot \varphi(r) \cdot \sum_{i \in I} vol(U_{i})$$

$$= Z \sum_{i \in I} out(U_{i})$$
(4)

where the last step follows by the choice of $\varphi(r)$. This means the reduction of charge on border edges is even lower bounded by the *total* internal charge (not just the increase of internal charge). Hence, the overall charge is not increasing.

Bound on the number of rounds. In order to keep the total number of rounds small we guarantee that the active volume, i.e., $\sum_{i \in I} \operatorname{vol}(U_i)$ decreases by a constant factor between two rounds. In order to guarantee this we first show that the choice of $\varphi(r)$ fulfills the following inequality.

$$active-charge \le 40\gamma_{\rm CMP}R^2Z^2 \cdot \varphi(r) \cdot \sum_{i \in I} \operatorname{vol}(U_i) \quad .$$
(5)

The active charge consists of the total internal charge and the active border charge. Equation 4 gives that

internal-charge
$$\leq 8\gamma_{\text{CMP}}RZ^2 \cdot \varphi(r) \cdot \sum_{i \in I} \operatorname{vol}(U_i)$$

The active border charge is

$$active-border-charge = f_B(r) \sum_{i \in I} \operatorname{out}(U_i)(Z + \log_2(\operatorname{ivol}(U_i)))$$

$$\leq 2f_B(r)Z \sum_{i \in I} \operatorname{out}(U_i) \leq 4RZ \sum_{i \in I} \operatorname{out}(U_i)$$

$$\leq 32\gamma_{\text{CMP}} R^2 Z^2 \varphi(r) \cdot \sum_{i \in I} \operatorname{vol}(U_i)$$
(6)

where the last inequality follows as the algorithm chooses $\varphi \geq \frac{1}{8\gamma_{\text{CMP}}RZ} \cdot \sum_{i \in I} \operatorname{out}(U_i) / \sum_{i \in I} \operatorname{vol}(U_i)$. Combining both inequalities gives Inequality 5. Claim 4.13. The term $\sum_{i \in I} \operatorname{vol}(U_i)$ decreases by a factor of at least 1/2 between two rounds of the algorithm. This gives that $R \leq \log_2 m$.

Proof. The active charge on a boundary edge is at least RZ. Since we do not generate charge during a round and we do not redistribute inactive charge we obtain that at the end of the round

$$\sum_{i \in I'} \operatorname{out}(U_i) \leq \frac{1}{RZ} \operatorname{active-charge}' \leq \frac{1}{RZ} \operatorname{active-charge} \leq 40\gamma_{\text{CMP}} RZ\varphi(r) \sum_{i \in I} \operatorname{vol}(U_i)$$
$$= 40\gamma_{\text{CMP}} RZ\varphi(r) \sum_{i \in I} \operatorname{vol}(U_i) ,$$

where I' denotes the set of active cluster after the first inner while-loop (i.e., before we start deactivating clusters). We use *active-charge'* to denote the active charge at this time. The last equation holds because the active volume does not change during the first while-loop.

Now a simple averaging argument gives that the volume in "bad" clusters (i.e. active clusters that have $80\gamma_{\text{CMP}}\varphi(r)RZ \operatorname{vol}(U_i) < \operatorname{out}(U_i)$) is at most half of the active volume. These are the clusters that make it to the next round. Hence, the claim follows.

Running time. We first derive a bound on the running time of a single round. When we apply the cut-matching+trimming algorithm from Lemma 4.7 to a subgraph $G[U_i]^w$ we can charge the running time to the edges in $G[U_i]^w$. We charge $O(\frac{1}{\varphi} \log m) = O(\frac{1}{\phi} \log m)$ to every edge. Whenever we charge an edge e at least one cluster U_i that contains an endpoint of e changed. We either set expands $(U_i, \varphi) \leftarrow$ true for this cluster (and, hence, stop processing this cluster for this round) or $\operatorname{vol}_G(U_i)$ decreases by a $(1 - 1/\log^2 m)$ factor. This implies that an edge can be charged at most $O(\log^3 m)$ times during a round.

It remains to derive a bound on the total number of edges in active clusters. Note that we cannot simply use m as an upper bound because the algorithm acts on sub-cluster $G[U]^w$, i.e., graphs where w self-loops are added for each border-edge.

The total number of border edges during a round is at most

initial-charge/ $RZ \leq O(b + \gamma_{\text{CMP}} Z \phi m)$

because every border-edge has charge at least RZ.

For each such border edge we add $w = \lceil \alpha/\varphi \rceil \leq 1/(\gamma_{\text{KRV}} Z\phi)$ self-loops (where we use $\phi \leq \alpha \leq 1/(\gamma_{\text{CMP}} Z)$). Therefore there are at most $\tilde{O}(b/\phi + m)$ edges in all graphs $G[U_i]^{\alpha/\phi}$. Hence, the cost of a single round is only $O(\log^4 m(b/\phi^2 + m/\phi))$. Since, the number of rounds is logarithmic the running time follows.

5 Tree Flow Sparsifier From Expander Hierarchy

In this section, we show that an expander hierarchy of a graph G is itself a tree flow sparsifier of G. Usually the concept of a flow sparsifier is defined for weighted graph. In order to simplify the notation and keep it consisten throughout the paper our definitions and proofs only consider unweighted (multi-)graphs. However the extension to weighted graphs is straightforward.

Multicommodity Flow. Given an unweighted (multi-)graph G = (V, E), let \mathcal{P} be the set of all paths in G. A multicommodity flow (that is also referred to as a flow) F is an assignment of non-negative values F_P to all paths $P \in \mathcal{P}$. Each path in \mathcal{P} has one of its endpoints being the originating vertex, and the other endpoints being the terminating vertex. When we assign the value F_P to the path P, we are sending F_P unit of flow from its originating vertex to its terminating vertex. For two vertices $v, v' \in V$, we denote by $\mathcal{P}_{v,v'} \subseteq \mathcal{P}$ the set of paths that originate at v and terminate at v', and we say that the amount of flow that F sends from v to v' is $\sum_{P \in \mathcal{P}_{v,v'}} F_P$. The congestion of the flow F is defined to be $\operatorname{cong}_G(F) = \max_{e \in E} \{F(e)\}$, where F(e) is the total amount of flow sent along the edge e. Given a flow F on G and two subsets $A, B \subseteq V(G)$ of vertices, we define F(A, B) to be the total amount of flow of F that is sent along an edge $e \in E(A, B)$ from its endpoint in A to its endpoint in B. Note that, however, for two vertices $v, v' \in V$ such that $(v, v') \in E$, $F(\{v\}, \{v'\})$ can be smaller than the amount of flow that F sends from v to v'.

Cut and Flow Sparsifiers. Given a (multi-)graph G = (V, E) and a subset $S \subseteq V$ of vertices, let H = (V', E') be a (multi-)graph with $S \subseteq V(H)$. We say that the graph H is a *cut sparsifier* of quality q for G with respect to S, if for each partition (A, B) of S such that both A and B are not empty, we have mincut_H $(A, B) \leq \text{mincut}_G(A, B) \leq q \cdot \text{mincut}_H(A, B)$, where mincut_H(A, B)(mincut_G(A, B), resp.) is the capacity of a minimum cut that separates the subsets A and B of vertices in H (G, resp.). If H is a tree, then H is called a *tree cut sparsifier*.

Given a (multi-)graph G = (V, E) and a subset $S \subseteq V$ of vertices, a set D of demands on S is a function $D : S \times S \to \mathbb{R}_{\geq 0}$, that specifies, for each pair $u, v \in V$ of vertices, a demand D(u, v). We say that the set D of demands is γ -restricted, iff for each vertex $v \in S$, $\sum_{u \in S} D(v, u) \leq \gamma \operatorname{out}(v)$ and $\sum_{u \in S} D(u, v) \leq \gamma \operatorname{out}(v)$, i.e., the demand entering or leaving v is at most γ times the number of edges leaving v. We call it γ -boundary restricted (w.r.t. S) if $\sum_{u \in S} D(v, u) \leq \gamma \operatorname{border}_S(v)$ and $\sum_{u \in S} D(u, v) \leq \gamma \operatorname{border}_S(v)$. Given a subset $S \subseteq V$ and a set D of demands on S, a routing of D in G is a flow F on G, where for each pair $u, v \in S$, the amount of flow that F sends from u to v is D(u, v). We define the congestion $\eta(G, D)$ of a set D of demands in G to be the minimum congestion of a flow F that is a routing of D in G. We say that a set D of demands is routable on G if $\eta(G, D) \leq 1$.

We say that a graph H is a flow sparsifier of quality q for G with respect to S, if $S \subseteq V(H)$, and for any set D of demands on S, $\eta(H, D) \leq \eta(G, D) \leq q \cdot \eta(H, D)$. A flow sparsifier H of G w.r.t. subset V(G) is just called a flow sparsifier for G. If H is a tree we call H a tree flow sparsifier.

We will use the following lemma, which is a direct consequence of approximate max-flow mincut ratios for multicommodity flows. The proof appears in Appendix A.

Lemma 5.1. Given a graph G together with a subset $S \subseteq V$ that is (α, ϕ) -linked in G. Then the following two statements hold.

- We can route a γ -restricted set of demands D on S with congestion $O(\frac{\gamma}{\phi} \log m)$ inside G[S].
- We can route a γ -boundary restricted set of demands D on S with congestion $O(\frac{\gamma}{\alpha} \log m)$ inside G[S].

The main theorem of this section is to show the following theorem that a (α, ϕ) -EH of a graph is automatically a tree flow sparsifier. It is well known that flow sparsifiers are a stronger notion than cut-sparsifiers and that the quality of the two version may differ by a logarithmic factor.

Theorem 5.2. The (α, ϕ) -EH of an undirected, connected graph G with m edges forms a tree flow sparsifier for G with quality $O(s \log m)^t \cdot O(\max\{\frac{1}{\alpha}, \frac{1}{\phi}\}/\alpha^{t-1})$, where t denotes the depth and s the slack of the hierarchy.

If we set $\phi = 2^{-\sqrt{\log m}}$ and so $t = O(\sqrt{\log m})$, then together with Corollary 4.6, we immediately obtain the following corollary:

Corollary 5.3. There is an algorithm, that, given any unweighted m-edge graph G, with high probability, computes a tree flow sparsifier for graph G of quality $O(\log n)^{O(\sqrt{\log n})}$ in time $m^{1+o(1)}$.

Observe that stronger results than the above theorem are known because [RST14] gives a polylogarithmic guarantee on the quality with a running time of O(m polylog m). However, our approach here is simpler and because we are able to efficiently maintain an expander hierarchy we also obtain a result for dynamic graphs. The main tool for proving Theorem 5.2 is the following lemma that shows how to construct a flow sparsifier for a graph G if one is given a flow sparsifier for some contraction $G_{\mathcal{U}}$ of G.

Lemma 5.4. Let G be an unweighted graph and $\mathcal{U} = (U_1, \ldots, U_r)$ be an (α, ϕ) -ED of G with slack s. Given a flow sparsifier $H_{\mathcal{U}}$ for the contracted graph $G_{\mathcal{U}}$ we can construct a flow sparsifier H for G as follows:

- 1. Add vertices of V(G) to $V(H_{\mathcal{U}})$.
- 2. Connect a newly added vertex $v \in U_i$ to the vertex $u_i \in V(H_{\mathcal{U}})$ with $\operatorname{out}_G(v)$ parallel edges.

The quality q_H of the resulting flow sparsifier H is $O((\frac{q}{\alpha} + \frac{1}{\phi})s\log m)$, where q denotes the quality of the flow sparsifier $H_{\mathcal{U}}$.

Proof. For a given set D of demands on V(G) we use $D_{\mathcal{U}}$ to denote the projection of D to $V(G_{\mathcal{U}})$, i.e., for two nodes $u_i, u_j \in V(G_{\mathcal{U}})$ we define $D_{\mathcal{U}}(u_i, u_j) := \sum_{x \in U_i, y \in U_j} D(x, y)$.

We first show that for all demands D we have $\eta(H, D) \leq \eta(G, D)$. Fix some demand D and assume w.l.o.g. that the congestion $\eta(G, D)$ for routing D in G is 1 (this can be obtained by scaling).

For routing between two vertices $x \in U_i$ and $v \in U_j$ from H we split their demand into three parts: $x \to u_i, u_i \to u_j$, and $u_j \to y$. Doing this for all demand-pairs gives us three sets of demands: the source demand D_s defined by $D_s(x, u_i) := \sum_{y \in V(G)} D(x, y)$ (where $x \in U_i$), the projected demand $D_{\mathcal{U}}$ and the target demand $D_t(u_j, y) := \sum_{x \in V(G)} D(x, y)$ (where $y \in U_j$). We route theses demands in H as follows.

- The source and target demand can be routed in H via the edges that were added in Step 2. The total traffic that is generated on the edge (x, u_i) is the total demand that leaves or enters vertex x in D. However, the latter is at most $\operatorname{out}_G(x)$ as otherwise the demand could not be routed in G with congestion 1. Hence, the congestion caused by this step in H is at most 1.
- The projection demand $D_{\mathcal{U}}$ can be routed only along edges belonging to $H_{\mathcal{U}}$. Clearly, this demand can be routed in $G_{\mathcal{U}}$ with congestion at most 1, and, hence, it can also be routed in $H_{\mathcal{U}}$ with congestion at most 1 as $H_{\mathcal{U}}$ is a flow sparsifier for $G_{\mathcal{U}}$.

Observe, that the edges used for routing in the above two steps are disjoint. Hence, we can concurrently route demands D_s, D_t , and $D_{\mathcal{U}}$ with congestion 1, and, hence, we can also route D with this congestion.

Now, we show that $\eta(G, D) \leq q_H \cdot \eta(H, D)$. Fix some demand D and assume w.l.o.g. that $\eta(H, D) = 1$. From this it follows that we can route the projected demand $D_{\mathcal{U}}$ in $H_{\mathcal{U}}$ with congestion 1. Since, $H_{\mathcal{U}}$ is a flow sparsifier for $G_{\mathcal{U}}$ (with quality q) this implies that we can also route $D_{\mathcal{U}}$ in $G_{\mathcal{U}}$ with congestion q.

In the following we describe how to extend a routing F for the demand $D_{\mathcal{U}}$ in $G_{\mathcal{U}}$ to a routing of D in G. In a first step we map the non-empty flow paths of F to G (note that the edges of the contracted multigraph $G_{\mathcal{U}}$ also exist in G; we simply map the flow from edges in $G_{\mathcal{U}}$ to the corresponding edge in G). Consider such a flow path $u_i = u_{s_1}, u_{s_2}, \ldots, u_{s_k} = u_j$. In G its edges connect subsets $U_{s_1}, U_{s_2}, \ldots, U_{s_k}$ but they do not form paths. For example we could have an edge $e = (x_1, x_2)$ followed by an edge $e' = (x'_2, x_3)$, with $x_1 \in U_{s_1}, x_2, x'_2 \in U_{s_2}$, and $x_3 \in U_{s_3}$. In order to obtain *paths* in G we have to connect x_2 to x'_2 . Performing this reconnection step for all routing paths from F induces a flow problem for every cluster U_i . The total demand (incoming and outgoing) for a vertex $x \in U_i$ in this flow problem is the total value of all flow-paths that x participates in. But this can be at most q border $U_i(v)$ as each of these flow paths uses an edge incident to x that leaves U_i and the congestion is at most q. Since U_i is $(\alpha/s, \phi/s)$ -linked and this set of demands is q-boundary-restricted, by Lemma 5.1, we can route such a set of demands in $G[U_i]$ with congestion $O(\frac{q}{\alpha}s \log m)$. As all clusters U_i are vertex-disjoint, performing all reconnections results in congestion $O(\frac{q}{\alpha}s \log m)$.

We also map the empty flow paths of F to empty paths in G as follows. A $u_i - u_i$ path in F is mapped to x - x paths in G with $x \in U_i$ such that the total flow that starts at a vertex x (including empty paths) is exactly $\sum_{y \in V(G)} D(x, y)$. This can be obtained because $\sum_{y \in V(G)} D_{\mathcal{U}}(u_i, y) = \sum_{x \in U_i} \sum_{y \in V(G)} D(x, y)$ and because F routes demands $D_{\mathcal{U}}$.

Let D' denote the set of demands routed by the flow system that we have constructed so far. Observe that D' has the same projection as our demand D, i.e., $D'_{\mathcal{U}} = D_{\mathcal{U}}$. The following claim shows that one can extend a routing for D' to a routing for D with a small increase in congestion.

Claim 5.5. Suppose we are given γ -restricted demands D and D' that fulfill $D'_{\mathcal{U}} = D_{\mathcal{U}}$ and assume that D' can be routed with congestion C'. Then we can route D with congestion at most $O(C' + s\gamma/\phi \cdot \log m)$.

Proof. Since, the projection of demands D and D' are equal we know that $\sum_{(x,y)\in U_i\times U_j} D(x,y) = \sum_{(x,y)\in U_i\times U_j} D'(x,y)$. We successively route D using the flow-paths of the routing for D'. For every pair (x,y) that we want to connect in D we find portals $(x',y') \in U_i \times U_j$ that are connected in D'. Then we add flow paths from x to x' and from y' to y. Formally, we use the following algorithm to compute a demand R'' such that D' together with R'' can route D.

 $\begin{array}{l} R \leftarrow D; \\ R' \leftarrow D'; \\ \textbf{while } \exists x \in U_i, y \in U_j \ \text{with } R(x, y) > 0 \ \textbf{do} \\ & | \ \text{choose } x' \in U_i, y' \in U_j \ \text{with } R'(x', y') > 0 \ ; \\ & | \ \text{choose } R(x, y) \ \text{and } R'(x', y') \ \text{by } \epsilon \ ; \\ & | \ \text{increase } R''(x, x') \ \text{and } R''(y', y) \ \text{by } \epsilon; \\ & | \ // \ \texttt{R}" \ \texttt{stores demand for connecting to portals} \\ \textbf{end} \end{array}$

The demands in R'' are just between vertex pairs inside clusters U_i , $i \in \{1, \ldots, s\}$. The total demand that can enter or leave a vertex v (in R'') is at most $2\gamma \operatorname{out}_G(v)$, because each such demand either occurs in R'' because v is used as an original source/target for demand in D or as a portal (i.e., as a source/target of a demand in D'). Since, both D and D' are γ -restricted we get that R'' is 2γ -restricted. Therefore, we can route R'' with congestion $O(s\gamma/\phi \cdot \log m)$ by Lemma 5.1 using the fact that each U_i is $(\alpha/s, \phi/s)$ -linked.

Using the fact that demands D and D' are O(1)-restricted we can route D with congestion at most $q_H := O((\frac{sq}{\alpha} + \frac{s}{\phi}) \log m)$. This gives the bound on the quality of the sparsifier H. \Box

Proof of Theorem 5.2: Let (G^0, G^1, \ldots, G^t) be some (α, ϕ) -expander decomposition sequence with slack s and let T denote the associated (α, ϕ) -expander-hierarchy. Recall that G^{i+1} is the contraction of G^i w.r.t. some $(\alpha/s, \phi/s)$ -linked partition \mathcal{U}_i of G_i , i.e., $G^{i+1} = G^i_{\mathcal{U}_i}$. G^t corresponds to the root of the tree and consists of just a single vertex⁸, while G^0 is identical to G. Let $T_{\geq i}$ denote the subgraph of T that just contains vertices that have at least distance i to the leaf-level, i.e., $T_{\geq t}$ is just the single root vertex and $T_{\geq 0} = T$. Note that the leaf vertices in $T_{\geq i}$ are the vertices on level i, which correspond to the nodes in G^i .

Let c denote the hidden constant in Lemma 5.4 and define $a := \frac{cs}{\alpha} \log m$ and $b := \frac{cs}{\phi} \log m$. This means $q_H \leq aq + b$ in Lemma 5.4. We show by induction that $T_{\geq i}$ is a sparsifier for G^i with quality $b\frac{a^{t-i}-1}{a-1} + a^{t-i}$. This clearly holds for i = t as then both graphs are identical (just a single vertex) and, hence, $T_{\geq t}$ is a sparsifier of quality 1. Now assume that the statement holds for i + 1 > 0. We prove it for i. We want to show that $T_{\geq i}$ is a sparsifier for G^i . We know that $T_{\geq i+1}$ is a sparsifier for $G^{i+1} = G^i_{\mathcal{U}_i}$; in addition $T_{\geq i}$ is obtained from $T_{\geq i+1}$ by adding vertices of $V(G^i)$ and attaching each vertex $v \in V(G^i)$ to the leaf vertex in $T_{\geq i+1}$ that corresponds to the cluster in \mathcal{U}_i that contains v. This means we can apply Lemma 5.4 and obtain that $T_{\geq i}$ is a sparsifier for G^i of quality

$$a \cdot \left(b\frac{a^{t-i-1}-1}{a-1} + a^{t-i-1}\right) + b = b\frac{a^{t-i}-1}{a-1} + a^{t-i}$$

Hence, for i = 0 we obtain that $T = T_{\geq 0}$ is a sparsifier for $G^0 = G$. The quality is $b\frac{a^t-1}{a-1} + a^t = O(c^t s^t \log^t m \max\{\frac{1}{\alpha}, \frac{1}{\phi}\}/\alpha^{t-1})$. This finishes the proof of Theorem 5.2.

6 Fully Dynamic Expander Pruning

In this section we prove the following theorem, which generalizes Theorem 1.3 in [SW19].

Theorem 6.1 (Fully Dynamic Expander Pruning). Let $0 \le \alpha, \phi \le 1$ and $\alpha/\phi \le w \le 3/(5\phi)$. There is a deterministic algorithm that given a graph G = (V, E), a cluster $U \subseteq V$ such that $G[U]^w$ is an ϕ -expander, and an online sequence of $k \le \phi \operatorname{vol}_{G[U]^w}(U)/120$ edge updates, where each update is an edge insertion or deletion for which at least one of the endpoints is contained in U, maintains a pruned set $P \subseteq U$ of vertices such that the following property holds. For each $1 \le i \le k$, let $G_i = (V, E_i)$ be the graph after the *i*th update, and denote by P_i the set P after the *i*-th update. We have

- 1. $P_0 = \emptyset$, and $P_i \subseteq P_{i+1}$.
- 2. $\operatorname{vol}_{G[U]^w}(P_i) \leq 32i/\phi$, and $|E_G(P_i, U \setminus P_i)| \leq 16i$.
- 3. $|E_G(P_i, V \setminus U)| \le 16i/\alpha$.
- 4. The graph $G_i[U \setminus P_i]^w$ is an $(\phi/38)$ -expander.

Moreover, the total running time for updating P_1, \ldots, P_k is $O(k \log m/\phi^2)$.

While the proof of Theorem 6.1 is similar to the proof of Theorem 1.3 in [SW19], there are subtle differences as we need to work with a cluster in a graph and not the whole graph, and more importantly, we need to show that a stronger notion of a graph defined on a cluster remains an expander.

Our algorithmic construction behind Theorem 6.1 uses Incremental Flow algorithm from Lemma 3.2 as a subroutine. Concretely, let G = (V, E) be a graph and let $U \subseteq V$ be a cluster that is $G[U]^w$ is an ϕ -expander. Let $\Pi = (\Delta, T, c)$ be a flow problem defined on $G[U]^w$ with

⁸For simplicity we assume that G is connected; the proof easily generalizes to graphs with several connected components.

 $\Delta(v) = 0$ for all $v \in U$, $T(v) = \deg_{G[U]^w}(v)$ for all $v \in U$ and $c(e) = 2/\phi$ for all $e \in E(G[U]^w)$. We give $G[U]^w$ and Π as inputs to the Incremental Flow algorithm of Lemma 3.2.

We next show how to handle updates in G. Consider the insertion or deletion of an edge e = (u, v) in G for which at least one of the endpoints is contained in U. For each endpoint $w \in \{u, v\}$ of e such that $w \in U$, we add $8/\phi$ unit of source mass at w, i.e., we set $\Delta(w) = \Delta(w) + 8/\phi$, and pass these source injecting operations to Incremental Flow. This completes the description of an iteration and the algorithm.

We next verify that the above algorithm satisfies the properties of Theorem 6.1. To prove the first property, note that from the Incremental Flow, it is clear that the maintained incremental set P satisfies $P_0 = \emptyset$, and $P_i \subseteq P_{i+1}$ for all $1 \leq i \leq k$ and thus $P \subseteq U$ serves as a pruned set in Theorem 6.1. Next, observing that (i) $\sum_{v \in V} \Delta(v) \leq 16i/\phi$ after *i* edge updates and (ii) $\min_e \{c_e\} = 2/\phi$, and using the second guarantee of Incremental Flow in Lemma 3.2, we get that $\operatorname{vol}_{G[U]^w}(P_i) \leq 32i/\phi$ and $|E_G(P_i, U \setminus P_i)| = |E_{G[U]^w}(P_i, U \setminus P_i)| \leq 16i$, thus proving the second property of Theorem 6.1.

The third property, i.e., the bound on the connectivity between the pruned set P_i and $V \setminus U$, is proved in the lemma below. Throughout, recall that $k \leq \phi \operatorname{vol}_{G[U]^w}(U)/120$ from Theorem 6.1.

Lemma 6.2. Let P_i be the pruned set. Then $|E_G(P_i, V \setminus U)| \leq 16i/\alpha$.

Proof. By the second guarantee of Incremental Flow in Lemma 3.2, we have that

$$\operatorname{vol}_{G[U]^w}(P_i) \le 2 \cdot 16k/\phi \le \operatorname{vol}_{G[U]^w}(U)/2,\tag{7}$$

and

$$|E_G(P_i, U \setminus P_i)| = |E_{G[U]^w}(P_i, U \setminus P_i)| \le 2 \cdot 16/\phi \cdot \phi/2 = 16i.$$

$$\tag{8}$$

As $G[U]^w$ is an ϕ -expander, it follows that $\operatorname{vol}_{G[U]^w}(P_i) \leq 1/\phi \cdot |E_G(P_i, U \setminus P_i)|$, and thus $\operatorname{vol}_{G[U]^w}(P_i) \leq 16i/\phi$. Moreover, $\operatorname{vol}_{G[U]^w}(P_i) \geq w \cdot |E_G(P_i, V \setminus U|)$ by definition of $G[U]^w$. Combining these two bounds and since $w \geq \alpha/\phi$, it follows that $|E_G(P_i, V \setminus U)| \leq 16i/(\phi w) \leq 16i/\alpha$.

It remains to show the fourth property, i.e., the graph $G_i[U \setminus P_i]^w$ is an $(\phi/38)$ -expander for all $1 \leq i \leq k$.

Lemma 6.3. Let $\alpha/\phi \leq w \leq 3/(5\phi)$. The graph $G_i[U \setminus P_i]^w$ is an $(\phi/38)$ -expander.

We will prove the above lemma through several steps. We start by bounding the total amount of mass injected in any subset of the cluster $U \setminus P_i$. To this end, let $A := U \setminus P_i$ be the cluster after pruning the set P_i . The Incremental Flow subroutine guarantees that the flow problem (Δ', T', c') is feasible on $G[U]^w[A]^{19}$ where $\Delta'(v) = \Delta(v) + 2/\phi \cdot |\{e \in E_G(P_i, A) \mid v \in e\}|$ (it is crucial to note here that the flow problem is defined on $G[U]^w[A]^1$ and not on $G_i[U]^w[A]^1$), and $T'(v) = \deg_{G[U]^w}(v)$ for all $v \in A$ and $c'(e) = 2/\phi$ for all $e \in E(G[U]^w[A]^1)$. Let $\Delta'(S) :=$ $\sum_{u \in S} \Delta'(u)$ be the total amount of source mass in S. We next prove a proposition, which will be instrumental in proving Lemma 6.3.

Proposition 6.4. For any set $S \subseteq A$, $\Delta'(S) \leq \operatorname{vol}_{G[U]^w}(S) + \frac{2}{\phi} |E_G(S, A \setminus S)|$.

⁹To explain the notation, let $H = G[U]^w$. We have $G[U]^w[A]^1 = H[A]^1$.

Proof. Consider a feasible flow f for the flow problem (Δ', T', c') defined on $G[U]^w[A]^1$. Recall that $f(v) = \Delta'(v) + \sum_{u} f(u, v)$ and f(u, v) = -f(v, u). It follows that

$$\Delta'(S) = \sum_{v \in S} \left[f(v) + \sum_{u} f(v, u) \right]$$

$$\leq \sum_{v \in S} T'(v) + \sum_{e \in E_{G[U]^{w}}(S, A \setminus S)} c(e)$$

$$= \operatorname{vol}_{G[U]^{w}}(S) + \frac{2}{\phi} |E_{G}(S, A \setminus S)|.$$

In order to leverage the ϕ -expansion of the graph $G[U]^w$, the following lemma shows how to relate the volume of a subset defined on $G_i[A]^w$ with the volume of that subset defined on $G[U]^w$.

Lemma 6.5. Let $S \subset A \subset U$. If $\operatorname{vol}_{G_i[A]^w}(S) \leq \frac{1}{2} \operatorname{vol}_{G_i[A]^w}(A)$, then $\operatorname{vol}_{G[U]^w}(S) \leq \frac{3}{5} \operatorname{vol}_{G[U]^w}(A)$.

Proof. From Equation (7), note that $\operatorname{vol}_{G[U]^w}(A) = \operatorname{vol}_{G[U]^w}(U) - \operatorname{vol}_{G[U]^w}(P_i) \ge \operatorname{vol}_{G[U]^w}(U)/2$, which in turn implies that $k \leq \phi \operatorname{vol}_{G[U]^w}(A)/60$. We also have that $|E_{G_i}(P_i, A)| \leq |E_G(P_i, A)| + e^{-\beta i A_i}$ k < 17k by Equation (8). It follows that

$$\begin{aligned} \operatorname{vol}_{G[U]^w}(S) &\leq \operatorname{vol}_{G_i[U]^w}(S) + wk & (|\operatorname{vol}_{G[U]^w}(S) - \operatorname{vol}_{G_i[U]^w}(S)| \leq wk) \\ &\leq \operatorname{vol}_{G_i[A]^w}(S) + wk & (\operatorname{since} A \subset U) \\ &\leq \frac{1}{2} \operatorname{vol}_{G_i[A]^w}(A) + wk & (\operatorname{by assumption of the lemma}) \\ &\leq \frac{1}{2} (\operatorname{vol}_{G_i[U]^w}(A) + w|E_{G_i}(P_i, A)|) + wk & (\operatorname{by definition of } G_i[A]^w) \\ &\leq \frac{1}{2} (\operatorname{vol}_{G[U]^w}(A) + wk + 17wk) + wk & (|\operatorname{vol}_{G_i[U]^w}(A) - \operatorname{vol}_{G[U]^w}(A)| \leq wk) \\ &\leq \frac{1}{2} \operatorname{vol}_{G[U]^w}(A) + 10wk \\ &\leq \frac{1}{2} \operatorname{vol}_{G[U]^w}(A) + \frac{1}{10} \operatorname{vol}_{G[U]^w}(A) & (k \leq \phi \operatorname{vol}_{G[U]^w}(A)/60 \text{ and } w \leq 3/(5\phi)) \\ &\leq \frac{3}{5} \operatorname{vol}_{G[U]^w}(A) & \Box \end{aligned}$$

Before proceeding to the proof of Lemma 6.3, we introduce some useful notation. Fix an arbitrary subset $S \subseteq A$. Let $a := |E_G(S, P_i)|$ and $c := |E_G(S, A \setminus S)|$ be the number boundary edges of S that cross different parts in the original graph G. Similarly, let $a' := |E_{G_i}(S, P_i)|$ and $c' := |E_{G_i}(S, A \setminus S)|$ be the boundary edges of S that cross different parts in the current graph G_i . Let $v_a := \operatorname{vol}_{G[A]^w}(S), v_u := \operatorname{vol}_{G[U]^w}(S)$ and $v'_a := \operatorname{vol}_{G_i[A]^w}(S), v'_u := \operatorname{vol}_{G_i[U]^w}(S)$. Note that by Proposition 6.4, we have that $\Delta'(S) \leq v_u + \frac{2}{\phi}c$.

Proof of Lemma 6.3. Recall that $A = U \setminus P_i$ and consider any $S \subset A$ such that $\operatorname{vol}_{G_i[A]^w}(S) \leq C$ $\operatorname{vol}_{G_i[A]^w}(A)/2$. To prove that $G_i[A]^w$ is an $(\phi/38)$ -expander, we need to show that $|E_{G_i}(S, A \setminus A)|$ $|S| \ge (\phi/38) \cdot \operatorname{vol}_{G_i[A]^w}(S), \text{ i.e., } c' \ge (\phi/38) \cdot v'_a.$

To this end, we first show a useful relation using the ϕ -expansion of $G[U]^w$. As $\operatorname{vol}_{G_i[A]^w}(S) \leq$ $\operatorname{vol}_{G_i[A]^w}(A)/2$ holds, by Lemma 6.5 we have that $\operatorname{vol}_{G[U]^w}(S) \leq \frac{3}{5} \operatorname{vol}_{G[U]^w}(A)$. From the latter we get $\operatorname{vol}_{G[U]^w}(A \setminus S) = \operatorname{vol}_{G[U]^w}(A) - \operatorname{vol}_{G[U]^w}(S) \ge \frac{5}{3} \operatorname{vol}_{G[U]^w}(S) - \operatorname{vol}_{G[U]^w}(S) \ge \frac{2}{3} \operatorname{vol}_{G[U]^w}(S).$ Therefore,

$$\begin{aligned} |E_G(S, P_i)| + |E_G(S, A \setminus S)| &= |E_{G[U]^w}(S, U \setminus S)| \\ &\geq \phi \cdot \min\{\operatorname{vol}_{G[U]^w}(S), \operatorname{vol}_{G[U]^w}(U \setminus S)\} \\ &\geq \phi \cdot \min\{\operatorname{vol}_{G[U]^w}(S), \operatorname{vol}_{G[U]^w}(A \setminus S)\} \\ &\geq \frac{2}{3}\phi \cdot \operatorname{vol}_{G[U]^w}(S), \end{aligned}$$

or $\frac{3}{2\phi}a + \frac{3}{2\phi}c \ge v_u$. The latter, together with Proposition 6.4, implies that the total amount of source mass in S is bounded by

$$\Delta'(S) \le v_u + \frac{2}{\phi}c \le \frac{3}{2\phi}a + \frac{7}{2\phi}c.$$
(9)

Now, by construction, recall that our algorithm increases the source mass of the endpoints in U from the inserted and deleted edges by $8/\phi$. Moreover, by the first property of Lemma 3.2, the flow problem (Δ', T', c') on $G[U]^w[A]^1$ is feasible. These together imply that $\Delta'(S) \geq \frac{2}{\phi}a + \frac{8}{\phi}|v'_u - v_u|$, $\Delta'(S) \geq \frac{2}{\phi}a + \frac{8}{\phi}|a' - a|$, and $\Delta'(S) \geq \frac{2}{\phi}a + \frac{8}{\phi}|c' - c|$. We claim that

$$|v'_u - v_u|, |a' - a|, |c' - c| \le c/2,$$
(10)

for otherwise $\Delta'(S) \geq \frac{2}{\phi}a + \frac{8}{\phi}(c/2) = \frac{2}{\phi}a + \frac{4}{\phi}c$, which contradicts Equation (9). Since $\frac{2}{\phi}a \leq \Delta'(S) \leq \frac{1}{\phi}a + \frac{3}{\phi}c$, we get that $a \leq 3c$. It follows that

$$\begin{aligned} v'_{a} &= \operatorname{vol}_{G_{i}[A]^{w}}(S) \leq \operatorname{vol}_{G_{i}[U]^{w}}(S) + w \cdot |E_{G_{i}}(S, P_{i})| & \text{(by definition of } G_{i}[A]^{w}) \\ &= v'_{u} + wa' \\ &\leq v_{u} + \frac{c}{2} + w \left(a + \frac{c}{2}\right) & \text{(Equation (10))} \\ &\leq \left(\frac{3}{2\phi}a + \frac{7}{2\phi}c\right) + \frac{1}{\phi}\frac{c}{2} + \frac{3}{5\phi}\left(a + \frac{c}{2}\right) & \left(v_{u} \leq \frac{3}{2\phi}a + \frac{7}{2\phi}c \text{ and } w \leq 3/(5\phi)\right) \\ &\leq \frac{19}{\phi}c & (a \leq 7c) \\ &\leq \frac{38}{\phi}c', & \text{(Equation (10))} \end{aligned}$$

what we wanted to show.

Finally, we analyse the running time. Note that over the course of the algorithm, there are at most $k = \phi \operatorname{vol}_{G[U]^w}(U)/120$ iterations and thus the total amount of mass $\sum_{v \in V} \Delta(v)$ injected in the graph $G[U]^w$ is at most $16/\phi \cdot \phi \operatorname{vol}_{G[U]^w}(U)/120 \leq \operatorname{vol}_{G[U]^w}(U)/3$. The latter implies that the condition on the total mass of Lemma 3.2 is met and by the same lemma we get that the running time is bounded by $O(c_{\max} \sum_{v \in V} \Delta(v) \log m) = O(k \log m/\phi^2)$. This completes the proof of Theorem 6.1.

7 Fully Dynamic Expander Hierarchy

In this section we deal with an undirected unweighted (multi-)graph G that undergoes a sequence of fully adaptive vertex and edge updates with the restriction that only isolated vertices may be deleted.

We assume that at any time G contains at most \bar{n} vertices and at most $\bar{m} = \text{poly}(\bar{n})$ edges. Further, we fix the following parameters throughout this section: $\phi = 2^{-\Theta(\log^{3/4} \bar{n})}$, $\psi = 2^{\Theta(\log^{1/2} \bar{n})}$, $\alpha = 1/\text{poly}(\log \bar{n})$ and we let $h = \log_{\psi}(\bar{m}) = \Theta(\log^{1/2} \bar{n})$ and $\rho = 38^{h}\psi/\alpha$. The main result of this section is the following theorem.

Theorem 7.1. There is a randomized algorithm, for maintaining an (α, ϕ) -expander hierarchy of G with slack $2^{\log^{1/2}(\bar{n})}$ in amortized update time $2^{O(\log^{3/4}(\bar{n}))}$.

The above theorem is based on the following theorem that shows that one can efficiently maintain an expander-decomposition.

Theorem 7.2. There is a randomized algorithm that maintains an (α, ϕ) -expander decomposition \mathcal{U} of G with slack 38^h together with its contracted graph $G_{\mathcal{U}}$ with the following properties:

- update time: $\tilde{O}(\psi \cdot 38^{2h}/\phi^2)$
- amortized recourse (number of updates to $G_{\mathcal{U}}$): $\tilde{O}(\rho) = \tilde{O}(38^h \cdot \psi/\alpha)$.

With the help of Theorem 7.2 we obtain Theorem 7.1 almost immediately.

Proof of Theorem 7.1. We maintain an (α, ϕ) -expander decompositions sequence (G^0, \ldots, G^t) with slack 38^h . For this we use algorithms $\tilde{A}_1, \ldots, \tilde{A}_{t_{\max}}$, where $t_{\max} = 2^{O(\log^{1/4} \bar{n})}$ is an upper bound on the depth of the sequence for our choice of ϕ . The algorithm \tilde{A}_i observes the updates for graph G^{i-1} , maintains an (α, ϕ) -expander decomposition \mathcal{U}^{i-1} with slack 38^h on this graph and generates updates for the contracted graph $G^i := G_{\mathcal{U}^{i-1}}$. The graph G^0 corresponds to the input graph G. The depth t of the maintained expander-hierarchy is determined by the first graph G^t in this sequence that does not contain any edges.

Because of the bounded recourse the number of updates that have to be performed for a graph G^i in this sequence is at most $\tilde{O}(\rho)^i k$, where k is the length of the update sequence for $G = G^0$. This results in a total update time of $\tilde{O}(k \sum_i \rho^i \psi 38^{2h} / \phi^2) = k 2^{\tilde{O}(\log^{3/4} \bar{n})}$.

7.1 Fully Dynamic Expander Decomposition

In this section we prove Theorem 7.2. The theorem follows from the following main lemma.

Lemma 7.3 (Main Lemma). Suppose a graph G initially contains m edges and undergoes a sequence of at most $O(\phi m/\rho)$ adaptive updates such that $V(G) \leq \bar{n}$ and $E(G) \leq \bar{m}$ always hold. Then there exists an algorithm that maintains an (α, ϕ) -expander decomposition \mathcal{U} with slack 38^h and its contracted graph $G_{\mathcal{U}}$ with the following properties:

- 1. update time: $\tilde{O}(\psi \cdot 38^{2h}/\phi^2)$
- 2. preprocessing time: $\tilde{O}(m/\phi)$
- 3. initial volume of $G_{\mathcal{U}}$ (after preprocessing): $\tilde{O}(\phi m)$
- 4. amortized recourse (number of updates to $G_{\mathcal{U}}$): $O(\rho) = O(38^h \cdot \psi/\alpha)$.

Proof of Theorem 7.2. We simply restart the algorithm from the above lemma whenever an update appears that would exceed the update limit. This means we have to perform this after $Z = \Theta(\phi m/\rho) + 1$ updates.

We have to analyze how this increases the update time and the recourse. First observe that before a restart the number of edges can be at most m + Z. Thus, the restart requires preprocessing time $\tilde{O}((m+Z)/\phi)$. Amortizing this against the Z updates increases the amortized update time by $\tilde{O}(\frac{m+Z}{\phi Z}) = \tilde{O}(m/(\phi Z) + 1/\phi) = \tilde{O}(\rho/\phi^2) = \tilde{O}(\psi \cdot 38^{2h}/\phi^2)$, where the last step follows because $\alpha = O(38^h)$.

The amortized recourse increases as follows. Observe that before the restart the total number of edges in $G_{\mathcal{U}}$ is at most $\tilde{O}(\phi m + Z\rho)$, because we only experienced Z updates and the amortized recourse is $O(\rho)$. We delete all these edges. Then we perform a preprocessing step. Since we have at most m + Z edges in G, Property 3 from the above lemma guarantees that this step inserts at most $\tilde{O}(\phi(m + Z))$ edges. Overall this increases the amortized recourse by $\tilde{O}((\phi m + Z\rho + \phi m + \phi Z)/Z) = \tilde{O}(\phi m/Z + \rho) = \tilde{O}(\rho)$.

This means the restarts only increase the recourse to $\tilde{O}(\rho)$.

In the remainder of this section we define the details for the ED-process, i.e., the algorithm from Lemma 7.3.

Multi-level Pruning

In order to define the details of the ED-process we first define a different process called Multi-level Pruning. A variant of this process will serve as a sub-routine in the ED-process.

The input for the Multi-level Pruning process is a cluster U that is (α, ϕ') -linked for parameters α, ϕ' that are known to the process. Then the process receives up to $N \leq \phi' \operatorname{vol}_G(U)/\rho$ many updates for G that are relevant for U, i.e., updates of edges for which at least one endpoint is in U. We will refer to N as the *update limit*. The process maintains a collection of pruned sets P^1, \ldots, P^{\hbar} such that $U \setminus \bigcup_s P^s$ is $(\alpha/38^{\hbar}, \phi'/38^{\hbar})$ -linked in G. Here $\hbar = \lceil \log_{\psi}(N) \rceil \leq h$.

The pruned sets are generated by a hierarchy of algorithms A_{\hbar}, \ldots, A_1 . The algorithm A_s maintains a set \tilde{P}^s and from time to time it changes P^s to the current value of \tilde{P}^s . In this respect P^s is a "snapshot" of \tilde{P}^s from an earlier time step. In the following \tilde{P}_t^s and P_t^s denote the sets \tilde{P}^s and P^s right after the *t*-th update.

The precise relationship between \tilde{P}^s and P^s is as follows. For constructing/maintaining its sets the level s algorithm A_s partitions the update sequence into *batches* of length

$$\ell_s := \begin{cases} N & \text{if } s = \hbar \\ \psi^s & \text{otherwise.} \end{cases}$$

each of which is partitioned into sub-batches of length ℓ_{s-1} ($\ell_0 = 1$). The *i*-th batch on level *s* contains updates number $(i-1)\ell_s + 1, \ldots, i\ell_s$. The *j*-th sub-batch of the *i*-th batch contains updates $(i-1)\ell_s + (j-1)\ell_{s-1} + 1, \ldots, (i-1)\ell_s + j\ell_{s-1}$. As in general $N \neq \psi^{\hbar}$ we allow the last batch for an algorithm to be incomplete and contain less than ℓ_s updates.

The algorithm A_s takes a "snapshot" of \tilde{P}^s at the start of every sub-batch. This means we define $P_t^s := \tilde{P}_{\lfloor t/\ell_{s-1} \rfloor \ell_{s-1}}^s$ if t does not start a new batch; otherwise $P_t^s := \emptyset$ as \tilde{P}_t^s is reset at the start of a batch.

How is a set \tilde{P}_t^s constructed? The construction of the set \tilde{P}_t^s on level *s* depends on the sets $P_t^{s'}$, s' > s. Let $Q_t^s := \bigcup_{s'>s} P_t^{s'}$ and observe that this set does not change during a batch for algorithm A_s . At the beginning of a batch A_s initializes $\tilde{P}^s := \emptyset$ (since this is also the start of a sub-batch it means also $P^s = \emptyset$ at this point). Then it simulates a run of the algorithm for fully dynamic expander pruning (Theorem 6.1) on subset $U \setminus Q_t^s$ for the ℓ_s updates of the batch. For this run it uses parameters $\alpha_s := \alpha/38^{\hbar-s}$ and $\phi'_s := \phi'/38^{\hbar-s}$ and $w := \alpha/\phi'$.

In order for the simulation to be valid we have to make sure that the preconditions of Theorem 6.1 are met. In particular we require that $U \setminus Q_t^s$ is (α_s, ϕ'_s) -linked and that the number of updates in a batch is at most the update limit of the expander pruning algorithm in Theorem 6.1.

Claim 7.4 (Correctness). For $s \in \{0, \ldots, \hbar\}$ the following properties hold.

- 1. $U \setminus Q_t^s$ is (α_s, ϕ'_s) -linked;
- 2. $\ell_s \leq \phi'_s \operatorname{vol}_G(U \setminus Q^s_t)/120 \leq \phi'_s \operatorname{vol}(G[U \setminus Q^s_t]^w)/120;$

Proof. We prove the lemma via induction. For the base case $s = \hbar$ the set Q_t^{\hbar} is empty. Then the above properties directly follow from the precondition of the input cluster U and the fact that $N \leq \phi' \operatorname{vol}_G(U)/\rho$.

Now, suppose that the statement holds for s + 1. We prove it for s. From the fact that the statement holds for s + 1 we are guaranteed that the simulation of the dynamic expander pruning that is performed by algorithm A_{s+1} is valid.

Part 1 follows because $U \setminus Q_t^s$ is the unpruned part that results from the execution of Theorem 6.1 by A_{s+1} . This theorem guarantees that $G[U \setminus Q_t^s]^w$ is a $(\phi_{s+1}/38)$ -expander with $w = \alpha_{s+1}/\phi'_{s+1}$. But this also means that $G[U \setminus Q_t^s]^w$ is a ϕ_s -expander with $w = \alpha_s/\phi'_s$. Theorem 6.1 also gives the following property for P_t^{s+1} :

$$\operatorname{vol}_{G}(P_{t}^{s+1}) \leq \frac{32}{\phi_{s+1}'} \ell_{s+1} \leq \frac{32}{\phi_{s+1}'} \phi_{s+1}' \operatorname{vol}_{G}(U \setminus Q_{t}^{s+1}) / 120 \leq \frac{1}{2} \operatorname{vol}_{G}(U \setminus Q_{t}^{s+1}) ,$$

where Step 1 is due to Theorem 6.1 and Step 2 is due to induction hypothesis. Hence,

$$\operatorname{vol}_G(U \setminus Q_t^s) = \operatorname{vol}_G(U \setminus Q_t^{s+1}) - \operatorname{vol}_G(P_t^{s+1}) \ge \frac{1}{2} \operatorname{vol}_G(U \setminus Q_t^{s+1})$$

We can use this relationship to obtain a bound on ℓ_s , which gives the second part of the claim. We differentiate two cases. If $s = \hbar - 1$ we get

$$\begin{split} \ell_s &\leq N \leq \phi' \operatorname{vol}_G(U) / \rho = 38\phi'_s \operatorname{vol}_G(U \setminus Q_t^{s+1}) / \rho \leq 76\phi'_s \operatorname{vol}_G(U \setminus Q_t^s) / \rho \\ &\leq \phi'_s \operatorname{vol}_G(U \setminus Q_t^s) / 120 \ , \end{split}$$

where the equality uses the fact that $U \setminus Q_t^{s+1} = U$ for $s+1 = \hbar$. If $s < \hbar - 1$ we have

$$\ell_s = \ell_{s+1}/\psi \le \phi'_{s+1} \operatorname{vol}_G(U \setminus Q_t^{s+1})/(120\psi) \le 38\phi'_s 2\operatorname{vol}_G(U \setminus Q_t^s)/(120\psi) \le \phi'_s \operatorname{vol}_G(U \setminus Q_t^s)/120 ,$$

for sufficiently large n as $\psi = \omega(1)$. This gives Part 2 of the claim.

Claim 7.4 guarantees that a pruned set P_t^s is generated by a valid run of the expander pruning algorithm from Theorem 6.1 on cluster $U \setminus Q_t^s$ with parameters ϕ'_s, α_s . Therefore it fulfills the following properties guaranteed by this theorem.

Claim 7.5. A set P_t^s fulfills the following properties.

1. $\operatorname{vol}_G(P_t^s) \leq 32\ell_s/\phi'_s = O(\psi^s/\phi'_s)$ (from Property 2a in Theorem 6.1) 2. $|E_G(P_t^s, U \setminus Q_t^s \setminus P_s^t)| \leq 16\ell_s = O(\psi^s)$ (from Property 2b in Theorem 6.1) 3. $|E_G(P_t^s, V \setminus (U \setminus Q_t^s))| \leq 16\ell_s/\alpha_s = O(\psi^s/\alpha_s)$ (from Property 3 in Theorem 6.1) 4. $\operatorname{out}_G(P_t^s) \leq 32\ell_s/\alpha_s = O(\psi^s/\alpha_s)$

Proof. The first three properties are directed consequences of Theorem 6.1. The last one follows from Property 2 and Property 3 because $\operatorname{out}_G(P_s^t) = |E_G(P_s^t, V \setminus (U \setminus Q_t^s))| + |E_G(P_s^t, U \setminus Q_t^s \setminus P_s^t)|$.

Claim 7.6. At any time, the cluster $U \setminus \bigcup_s P_t^s = U \setminus Q_t^0$ maintained by the multilevel pruning process fulfills the following properties:

- 1. $U \setminus Q_t^0$ is $(\alpha/38^h, \phi'/38^h)$ -linked in G
- 2. $\operatorname{vol}(U \setminus Q_t^0) \ge \frac{1}{2} \operatorname{vol}_G(U)$.
- 3. $\operatorname{cut}_U(Q_t^0) \le 48N$.

Proof. Part 1 directly follows by applying the above Claim 7.4 for s = 0 and using $\hbar \leq h$. For the remaining parts first observe that $\sum_{s=1}^{\hbar} \ell_s = \sum_{s=1}^{\hbar-1} \psi^s + N \leq 3N$. For Part 2 we estimate $\operatorname{vol}_G(Q_t^0)$ by

$$\operatorname{vol}_{G}(Q_{t}^{0}) \leq \sum_{s} \operatorname{vol}_{G}(P_{t}^{s}) \leq \sum_{s} 32\ell_{s}/\phi_{s}' \leq 32 \cdot 38^{h}/\phi' \cdot 3N \leq \operatorname{vol}_{G}(U)/10$$

where the second step uses Property 1 from Claim 7.5, and the last step uses $N \leq \phi' \operatorname{vol}_G(U)/\rho$. This implies $\operatorname{vol}_G(U \setminus Q_t^0) \geq \operatorname{vol}_G(U)/2$. Part 3 follows because

$$\operatorname{cut}_U(Q_t^0) = |E_G(Q_t^0, U \setminus Q_t^0)| = |E_G(\bigcup_s P_t^s, U \setminus Q_t^0)| = \sum_s |E_G(P_t^s, U \setminus Q_t^0)|$$

$$\leq \sum_s |E_G(P_t^s, U \setminus Q_t^{s+1})| \leq \sum_s 16\ell_s \leq 48N ,$$

where the second inequality is due to Claim 7.5 (Part 2).

Expander Decomposition and **Cluster Decomposition**

The Expander Decomposition process (ED-process) for maintaining the expander decomposition is a process that uses a variant of the Multi-level Pruning process as a sub-routine. We refer to this variant as a Cluster Decomposition process (CD-process). The ED-process gets as input a cluster U in a graph G, parameters α, ϕ' and a sequence of updates relevant for U. It first computes an (α, ϕ') -linked expander decomposition \mathcal{U} of U using Theorem 4.5. For each $U_i \in \mathcal{U}$ with expansion parameter ϕ_i it then starts a CD-process on U_i with parameters α and ϕ_i .

A CD-process is a Multi-level Pruning process with a slight tweak. Whenever, the Multi-level Pruning process (with parameters α, ϕ') as described in the previous section changes a set P^s , the CD-process starts an ED-process on this set (with parameters α, ϕ').

There is one further complication in the definition of an ED-process, which concerns the update limits of the CD-processes. An ED-process handles CD-processes for several clusters. It may happen that one of the CD-processes on some cluster U_i reaches its update limit N—we say the CD-process *expires*. In this case if another update for the cluster appears the ED-process does the following: it uses Theorem 4.5 on the cluster U_i with parameters α, ϕ' and starts a new CD-process on each generated sub-cluster U_{ij} (with parameters α, ϕ_j). We call this step a *restart* of cluster U_i .

There is one subtle issue about the above definition. The CD-process is recursive. The non-recursive case happens when $\operatorname{vol}_G(U) < \rho/\phi'$. Then the CD-process has an update limit $N = \lfloor \phi' \operatorname{vol}_G(U)/\rho \rfloor = 0$. This means any update triggers a restart of the CD-process, which results in computing an expander decomposition for U from scratch.

Observation 7.7. The parameter ϕ' passed to a CD-process or an ED-process on any level of the recursion is at least ϕ , where ϕ is the parameter for the root ED-process.

The following claim means that an ED-process automatically fulfills Property 2 and Property 3 of an (α, ϕ) -boundary-linked expander decomposition with slack 38^h .

Claim 7.8. An ED-process with parameters α, ϕ maintains a cluster-partition \mathcal{U} s.t.

- A cluster $U_i \in \mathcal{U}$ is $(\alpha/38^h, \phi_i/38^h)$ -linked, with $\phi_i \ge \phi$
- A cluster $U_i \in \mathcal{U}$ fulfills $\operatorname{out}_G(U_i) \leq \tilde{O}(\phi_i \operatorname{vol}_G(U_i))$.

Proof. The clusters that are maintained by the ED-process are the sets $U \setminus Q_t^0$ that are maintained by the various CD-processes on various levels of the recursion (U being the set on which the process was started).

A CD-process is always started with parameters (α, φ) for a set U that results from the expander-decomposition algorithm of Theorem 4.5 (run with parameters α and $\phi' \ge \phi$). This set is (α, φ) -linked for $\varphi \ge \phi'$ according to this theorem. Then the first part is a direct consequence of Claim 7.6 (Part 1).

For the second part again observe that a CD-process is started with parameters (α, φ) for a set U that results from the expander-decomposition algorithm of Theorem 4.5. This gives that $\operatorname{out}_G(U) \leq \tilde{O}(\varphi \operatorname{vol}_G(U))$ at the start of the CD-process.

The update limit guarantees that at most $O(\varphi \operatorname{vol}_G(U))$ updates are performed. This means while the CD-process is active it fulfills $\operatorname{out}_G(U) \leq \tilde{O}(\varphi \operatorname{vol}_G(U))$. Claim 7.6 (Part 3) guarantees that the additional edges that are added by the pruning process are $O(N) = \tilde{O}(\varphi \operatorname{vol}_G(U))$, as well. This means $\operatorname{out}_G(U \setminus Q_t^0) = \tilde{O}(\operatorname{vol}_G(U))$. Finally, Claim 7.6 (Part 2) gives that $\operatorname{vol}_G(U \setminus Q_t^0) \geq \operatorname{vol}_G(U)/2$.

Amortized Update Time of the CD-process. In the following we analyze the amortized update time of a CD-process. The cost for a CD-process on some cluster U with parameters α, ϕ' consists of the following parts:

- A. The cost for each algorithm A_s , which consists of
 - 1. the cost for simulating the expander pruning algorithm from Theorem 6.1;
 - 2. the cost for starting an ED-process on each generated set P_t^s .
- B. The cost for maintaining pointers from vertices in G to the corresponding cluster vertex in $G_{\mathcal{U}}$.
- C. The cost for performing a restart on the cluster U when the CD-process expires.
- D. The recursive cost incurred by CD-processes on lower levels of the recursion.

Part A. The amortized cost for A_s to simulate a step of the expander pruning algorithm is $\tilde{O}((\phi'_s)^2) = O((38^{\hbar-s}/\phi')^2)$ due to Theorem 6.1. Summing over all s gives that the amortized simulation cost (i.e. cost A.1) over all algorithms A_s is only $O(38^{2\hbar}/\phi'^2)$.

The cost for starting an ED-process on each generated set P_t^s is dominated by the running time for an expander-decomposition algorithm from Theorem 4.5. This has running time $\tilde{O}(\operatorname{out}_G(P_t^s)/\phi'^2 + \operatorname{vol}_G(P_t^s)/\phi')$. From Part 2 and Part 3 of Claim 7.5 we get $\operatorname{out}_G(P_t^s) = O(\psi^s/\alpha_s)$ and $\operatorname{vol}_G(P_t^s) = O(\psi^s/\phi'_s)$. This gives a total running time of

$$\tilde{O}\left(\psi^s (\frac{1}{\alpha_s \phi'^2} + \frac{1}{\phi'_s \phi'})\right) = \tilde{O}\left(\psi^s 38^{\hbar-s} (\frac{1}{\alpha \phi'^2} + \frac{1}{\phi'^2})\right) = \tilde{O}\left(\psi^s 38^{\hbar-s} / (\alpha \phi'^2)\right)$$

We have to perform this operation at the end of every sub-batch, i.e., we can amortize the cost against the $\ell_{s-1} = \psi^{s-1}$ updates in the sub-batch. Therefore, the amortized cost of algorithm A_s to start an ED-process is $\tilde{O}(\psi 38^{\hbar-s}/(\alpha \phi'^2))$. Summing this over all s gives an amortized cost for Part A.2 of $\tilde{O}(\psi 38^{\hbar}/(\alpha \phi'^2))$.

Combining this with the simulation cost gives that the amortized cost for Part A is $\tilde{O}(\psi 38^{\hbar}/\phi'^2 \cdot (38^{\hbar} + 1/\alpha)) = \tilde{O}(\psi 38^{2\hbar}/\phi'^2).$

Part B. Whenever we initialize a CD-process we also initialize a pointer for every vertex in U to point to the cluster node in $G_{\mathcal{U}}$ corresponding to the set U. Later this pointer is changed for the nodes that are pruned as these then belong to different clusters. However, the running time is O(|U|) whenever we initialize a CD-process on some cluster U. As a CD-process is always started on some cluster that results from the static expander-pruning algorithm of Theorem 4.5 we can amortize the cost for maintaining pointers against the cost of the expander-pruning algorithm

Part C. The cost for a restart is dominated by running the expander-decomposition algorithm from Theorem 4.5 for the cluster U. This is $\tilde{O}(\operatorname{out}_G(U)/\phi'^2 + \operatorname{vol}_G(U)/\phi')$ due to Theorem 4.5. Observe that a restart is only performed after $\lfloor \phi' \operatorname{vol}_{G'}(U)/\rho \rfloor + 1$ updates (the first $\lfloor \phi' \operatorname{vol}_{G'}(U)/\rho \rfloor$ to reach the update limit and one further update to trigger a restart). Here, G' refers to the graph at the time that the CD-process on U was started. We use the following claim.

Claim 7.9. Let G and G' denote the current graph and the graph at the time that the CD-process was started, respectively. Then

- $-\operatorname{vol}_G(U) = O(\operatorname{vol}_{G'}(U))$
- $\operatorname{out}_G(U) = \tilde{O}(\phi' \operatorname{vol}_{G'}(U))$

Proof. We have $\operatorname{vol}_G(U) \leq \operatorname{vol}_{G'}(U) + 2\phi' \operatorname{vol}_{G'}(U)/\rho = O(\operatorname{vol}_{G'}(U))$, because an update can increase the volume by at most 2. Further, we have $\operatorname{out}_G(U) \leq \operatorname{out}_{G'}(U) + 2\phi' \operatorname{vol}_{G'}(U)/\rho = \tilde{O}(\phi' \operatorname{vol}_{G'}(U))$, because Property 3 of an expander decomposition gives $\operatorname{out}_G'(U) \leq \tilde{O}(\phi' \operatorname{vol}_{G'}(U))$.

Using these inequalities we get that the cost for the restart is at most $\tilde{O}(\operatorname{out}_G(U)/\phi'^2 + \operatorname{vol}_G(U)/\phi') = \tilde{O}(\operatorname{vol}_{G'}(U)/\phi')$. We can amortize this cost against N + 1 updates. This gives an amortized cost of $\tilde{O}(\rho/\phi'^2)$ for the restarts.

Part D. The following claim shows that incorporating the cost for the recursion only increases the cost by a logarithmic factor.

Claim 7.10. An update that is relevant for a CD-process C on cluster U can be relevant for at most $O(\log(\text{vol}_G(U)))$ sub CD-processes of C.

Proof. At any time the set of subsets on which a CD-process is running forms a laminar family. In addition suppose we have a CD-process running on subset U. A sub CD-process will be started on some subset of a set P_t^s maintained by the CD-process (recall that the CD-process starts an ED-process on P_t^s , which in turn partitions the set and then starts a CD-process on each part of the partition).

According to Claim 7.5 we have $\operatorname{vol}_G(P_t^s) \leq 32\ell_s/\phi'_s \leq 32\operatorname{vol}_G(U \setminus Q_t^s)/120 \leq \operatorname{vol}_G(U)/2$, where the second Step uses Part 2 of Claim 7.4. This implies that the height of the recursion is only $\log(\operatorname{vol}_G(U))$.

An edge upate is relevant for a CD-process on set U if at least one end-point of the edge is contained in U. This can happen for at most $2\log(\operatorname{vol}_G(U))$ CD-processes.

We get the following lemma.

Lemma 7.11. The amortized time for performing an update operation with the CD-process is $\tilde{O}(\psi 38^{2h}/\phi^2)$.

Proof. The lemma simply follows by combining the costs from all parts and using $\phi' \ge \phi$ and $\hbar \le h$.

Lemma 7.12. An ED-process with parameters α, ϕ started on a graph G with m edges has an amortized update time of $\tilde{O}(\psi 38^{2h}/\phi^2)$ and a pre-processing time of $\tilde{O}(m/\phi)$.

Proof. The ED-process triggers an update for at most 2 CD-processes on the top level. Each of these incurs amortized cost $\tilde{O}(\psi 38^{2\hbar}/\phi'^2)$ with $\phi' \geq \phi$ according to Lemma 7.11. The preprocessing consists of executing the algorithm of Theorem 4.5, which has a running time of $\tilde{O}(m/\phi)$.

Amortized Recourse of the ED-process. In the following we derive a bound on the recourse generated by the root ED-process running on the graph G. We have to analyze how many edge insertions or deletions are generated for the contracted graph $G_{\mathcal{U}}$, where \mathcal{U} is the decomposition maintained by the ED-process. Since we care about the amortized recourse we can focus on edge insertions and amortize the deletions against the insertions at a loss of a factor of 2.

The partition maintained by the ED-process changes whenever a CD-process in the recursion hierarchy changes one of its sets P_t^s . Fix a CD-process \mathcal{C} on some subset U and assume there is an update relevant for U.

Let \bar{s} denote the unique level for which the current time-step ends the current sub-batch of algorithm $A_{\bar{s}}$ without also ending the current batch. All sets P_t^s with $s < \bar{s}$ will be reset to \emptyset because the batch ends and all sets P_t^s with $s > \bar{s}$ will not change. We first issue edge-deletions for all edges that are at the border of some partition of the ED-process inside $\bigcup_{s \leq \bar{s}} P_t^s$ (i.e., they contain exactly one vertex of a sub-partition) and also issue vertex deletions for the corresponding vertices of $G_{\mathcal{U}}$ (we do not have to count these deletions because of the amortization described above). Then we run the expander decomposition algorithm on the new set $P_t^{\bar{s}}$. For each edge in G that afterwards is at the border of a subset in the partition of $P_t^{\bar{s}}$ we issue an edge-insertion. This gives that the total number of insertions is $\sum_i \text{out}_G(U_i)$, where U_1, \ldots, U_k are the subsets of the partition. According to Theorem 4.5 this is at most

$$\tilde{O}\big(\operatorname{out}_G(P_t^s) + \phi' \operatorname{vol}_G(P_t^s)\big) \le \tilde{O}\big(\psi^s/\alpha_s + \phi' \cdot \psi^s/\phi'_s\big) \le \tilde{O}(38^{\hbar-s}\psi^s/\alpha)$$

This means that the algorithm A_s of CD-process \mathcal{C} generates on average $\tilde{O}(38^{\hbar-s}\psi/\alpha)$ updates for $G_{\mathcal{U}}$ per relevant update for \mathcal{C} . This holds because the above updates for $G_{\mathcal{U}}$ are only incurred after $\ell_{s-1} = \psi^{s-1}$ relevant updates for \mathcal{C} . Summing this over all s gives that the amortized recourse for \mathcal{C} due to algorithms A_s is only $\tilde{O}(38^{\hbar}\psi/\alpha) = \tilde{O}(\rho)$.

It remains to derive a bound on the recourse that is generated by \mathcal{C} when a restart is triggered. We delete all edges in $G_{\mathcal{U}}$ that are incident to a current sub-cluster and we also delete the vertices corresponding to these sub-clusters. Then we repartition U and start a CD-process on each cluster. We have to insert all edges that are at the border of a sub-set of the partition, i.e., $\sum_{i} \operatorname{out}_{G}(U_{i})$ many edges, where U_{1}, \ldots, U_{k} denote the subsets in the partition. We have

$$\sum_{i} \operatorname{out}_{G}(U_{i}) \leq \tilde{O}(\operatorname{out}_{G}(U) + \phi' \operatorname{vol}_{G}(U)) \leq \tilde{O}(\phi' \operatorname{vol}_{G}(U))$$

where the first inequality is due to Theorem 4.5 and the second due to Claim 7.9. Since, we can amortize these costs over $N + 1 = \lfloor \phi' \operatorname{vol}_{G'}(U)/\rho \rfloor + 1$ many relevant updates we obtain that the amortized recourse due to restarts is only $\tilde{O}(\rho)$.

Since a single update is relevant for at most at most $O(\log(vol(G)))$ CD-processes we obtain the following lemma.

Lemma 7.13. The amortized recourse of the ED-process is $\tilde{O}(38^{\hbar}\psi/\alpha) = \tilde{O}(\rho)$.

Total boundary for the partition of the ED-process. So far we have only shown that the ED-process maintains a partition that fulfills Property 2 and Property 3 for a boundary-linked partition. It remains to show that the total number of edges between subsets in the partition fulfills Property 1.

Claim 7.14. An ED-process with parameters α, ϕ on a graph G with m edges that receives $Z = O(\phi m/\rho)$ updates maintains a cluster-partition \mathcal{U} such that

$$\sum_{U_i \in \mathcal{U}} \operatorname{out}_G(U_i) \leq \tilde{O}(\phi m)$$
.

Proof. First the ED-process performs a boundary-linked expander decomposition and then starts a CD-process on each cluster. At this point the number of edges between sub-clusters (which equals the number of edges in $G_{\mathcal{U}}$) is at most $\tilde{O}(\phi m)$ according to Theorem 4.5. Because of the bounded recourse from Lemma 7.13 the total number of edges between vertices in $G_{\mathcal{U}}$ can be at most $\tilde{O}(\phi m) + Z \cdot \tilde{O}(\rho) = \tilde{O}(\phi m)$, after Z updates.

We are now ready to prove the main lemma.

Lemma 7.3 (Main Lemma). Suppose a graph G initially contains m edges and undergoes a sequence of at most $O(\phi m/\rho)$ adaptive updates such that $V(G) \leq \bar{n}$ and $E(G) \leq \bar{m}$ always hold. Then there exists an algorithm that maintains an (α, ϕ) -expander decomposition \mathcal{U} with slack 38^{h} and its contracted graph $G_{\mathcal{U}}$ with the following properties:

- 1. update time: $\tilde{O}(\psi \cdot 38^{2h}/\phi^2)$
- 2. preprocessing time: $\tilde{O}(m/\phi)$

- 3. initial volume of $G_{\mathcal{U}}$ (after preprocessing): $\tilde{O}(\phi m)$
- 4. amortized recourse (number of updates to $G_{\mathcal{U}}$): $O(\rho) = O(38^h \cdot \psi/\alpha)$.

Proof. The first property follows from Lemma 7.12. The second and third property follow from the fact that in the pre-processing we perform the expander decomposition algorithm from Theorem 4.5, which has a running time of $\tilde{O}(m/\phi)$ and generates an expander decomposition that fulfills $\sum_{i} \operatorname{out}(U_i) \leq \tilde{O}(\phi m)$. The bound on the amortized recourse follows from Lemma 7.13.

Claim 7.8 shows that the maintained partition fulfills Property 2 and Property 3 of an (α, ϕ) -linked expander decomposition with slack 38^h . Finally, Claim 7.14 shows that it also fulfills Property 1.

8 Derandomization and Deamortization

In this section, we show that our algorithm in Theorem 7.1, which maintains an $(n^{o(1)}, n^{o(1)})$ expander hierarchy of a dynamic graph on n vertices in $n^{o(1)}$ time can be de-randomized and de-amortized easily using the results in [CGL⁺19] and [NSW17]. Throughout the section, we use $\bar{O}(\cdot)$ to hide $(\log \log n)^{O(1)}$ factors. The main result of this section can be summarized as the following theorem.

Theorem 8.1. There is a deterministic algorithm, that, given a fully dynamic unweighted graph G on n vertices, maintains a data structure representing a $(2^{-\overline{O}(\log^{2/3} n)}, 2^{-O(\log^{5/6} n)})$ -expander hierarchy with slack $2^{\overline{O}(\log^{1/2} n)}$ of G in $2^{-O(\log^{5/6} n)}$ worst-case update time and the data structure supports the following query: given a vertex $u \in V(G)$, return a leaf-to-root path of u in the hierarchy in $O(\log^{1/6} n)$ time.

Compared with Theorem 7.1, the algorithm in Theorem 8.1 is deterministic, and also gives worst-case update time guarantees. On the flip side, it does not explicitly maintain a single expander hierarchy, but will constantly switch between several expander hierarchies that we maintain in the background, as we will see later. The proof of the main theorem consists of two parts, that we will show in the following subsections: the first part shows how to derandomize the algorithm in Theorem 7.1 using a recent result in [CGL⁺19]; and the second part shows how to de-amortize the algorithm in Theorem 7.1, using similar techniques from [NSW17]. We note that, by directly combining the methods from the two subsections, we immediately obtain an algorithm for Theorem 8.1. We now describe the two parts in more detail.

8.1 De-randomization

In this section we provide the proof of the following theorem.

Theorem 8.2. There is a deterministic algorithm, that, given a fully dynamic unweighted graph G on n vertices, explicitly maintains a $(2^{-\bar{O}(\log^{2/3} n)}, 2^{-O(\log^{5/6} n)})$ -expander hierarchy with slack $2^{\bar{O}(\log^{1/2} n)}$ of G in amortized update time $2^{\bar{O}(\log^{5/6} n)}$.

Recall that the algorithm in Theorem 7.1 utilizes the algorithm in Lemma 7.3 as a subroutine, and upon this, everything is deterministic. Recall also that the algorithm in Lemma 7.3 utilizes as subroutines the algorithm in Theorem 6.1, which is deterministic, and the algorithm in Theorem 4.5 as subroutines, which is randomized, and upon this, everything is deterministic. Observe that the only randomized part in the algorithm in Theorem 7.1 is the subroutine of the cut-matching game. Therefore, the only part that is randomized in the algorithm of Theorem 7.1 is also the cut-matching step in Lemma 4.8.

A recent result by Chuzhoy et al $[CGL^+19]$ gave a deterministic algorithm for the cutmatching step with weaker parameters, that is stated as follows. **Lemma 8.3.** There is a deterministic algorithm, that, given an unweighted graph G = (V, E) with m edges and a parameter $\phi > 0$,

- 1. either certifies that G has conductance $\Phi_G \ge 8\phi$;
- 2. or finds a cut (A, \overline{A}) of G with conductance $\Phi_G(A) \leq \gamma^* \phi$, and $\operatorname{vol}_G(A), \operatorname{vol}_G(\overline{A})$ are both at least $m/(16\gamma^*)$, i.e., we find a relatively balanced low conductance cut;
- 3. or finds a cut (A, \overline{A}) , such that $\Phi_G(A) \leq \gamma^* \phi$ and $\operatorname{vol}_G(\overline{A}) \leq m/(16\gamma^*)$, and A is a near 8ϕ -expander;

with the parameter $\gamma^* = 2^{O(\log^{2/3} m(\log \log n)^{1/3})} = 2^{\overline{O}(\log^{2/3} m)}$. Moreover, the algorithm runs in time $\widetilde{O}(m\gamma^*/\phi)$,

The following corollary is immediately obtained by replacing the randomized cut-matching step in the algorithm of Theorem 4.5 with the algorithm in the above theorem.

Corollary 8.4. There is a deterministic algorithm that, given a graph G = (V, E), a cluster $C \subseteq V$ with $\operatorname{vol}_G(C) = m$ and $|E_G(C, V \setminus C)| \leq b$, and parameters α, ϕ such that $\alpha \leq 2^{-\bar{O}(\log^{2/3} m)}$, computes an (α, ϕ) -expander decomposition of C in $\tilde{O}(b/\phi^2 + m\gamma^*/\phi)$ time, with $\gamma^* = 2^{O(\log^{2/3} m(\log \log n)^{1/3})} = 2^{\bar{O}(\log^{2/3} m)}$.

To obtain the algorithm for Theorem 8.2, we simply replace the randomized cut-matching step with the algorithm in Corollary 8.4. We also change the parameters in Section 7 accordingly as $\alpha = 2^{-\bar{O}(\log^{2/3} m)}$ and $\phi = 2^{-O(\log^{5/6} n)}$. We keep $\psi = 2^{O(\sqrt{\log n})}$, so $h = \Theta(\log_{\psi} n) = \Theta(\sqrt{\log n})$ as before, the depth of the hierarchy is $O(\log^{1/6} n)$ and $\rho = O(\psi \cdot 38^{h}/\alpha) = 2^{\bar{O}(\log^{2/3} m)}$. From the same proof of Section 7 (with distinct parameters), we can show that we can maintain an $(\alpha, \phi) = (2^{-\bar{O}(\log^{2/3} n)}, 2^{-O(\log^{5/6} n)})$ -expander hierarchy with slack $2^{\bar{O}(\log^{1/2} n)}$ of an *n*-vertex fully-dynamic graph in amortized update time $O(1/\phi^2) \cdot \rho^{\operatorname{dep}(T)} = 2^{O(\log^{5/6} n)}$.

8.2 De-amortization

The main result in this section is the following theorem.

Theorem 8.5. There is a randomized algorithm, that, given a fully dynamic unweighted graph G on n vertices undergoing adaptive edge insertions and deletions, maintains a data structure representing a $(2^{-O(\log^{1/2} n)}, 2^{-O(\log^{3/4} n)})$ -expander hierarchy with slack $2^{O(\log^{1/2} n)}$ of G in $2^{O(\log^{3/4} n)}$ worst-case update time and the data structure supports the following query: given a vertex u, return a leaf-to-root path of u in the hierarchy in $O(\log^{1/4} n)$ time with high probability.

In order to construct an algorithm for Theorem 8.5, we first show that we can de-amortize the algorithm in Lemma 7.3, and then we describe how to use this new algorithm of Lemma 7.3 to construct an algorithm for Theorem 8.5.

The crux in de-amortizing the algorithm in Lemma 7.3 is to de-amortize the core subroutine: Multi-level Pruning. Recall that Multi-level Pruning extensively uses the algorithm in Theorem 6.1 for expander pruning, that we denote by \mathcal{B} . However, the algorithm in Theorem 6.1 only guarantees small amortized update time, and cannot be de-amortized. To overcome this issue, the key observation is that, we cannot sequentially feed the algorithm \mathcal{B} up to the current update and force it to produce information with respect to the current graph. Instead, when we feed the algorithm \mathcal{B} with a batch of updates that has already shown up in the update sequence, we have to wait for a certain number of updates that is comparable to the length of the batch that we feed to \mathcal{B} , so that it can distribute the running time for processing the batch that we feed to it evenly to the new updates, thus achieving the worst-case update time guarantee.

In the remainder of this section, we first describe how to de-amortize the Multi-level Pruning process, and then describe how to use it to further de-amortize the algorithm for Lemma 7.3 and eventually provide an algorithm for Theorem 8.5.

De-amortize Multi-level Pruning. Recall that the Multi-level Pruning consists of a hierarchy of algorithms A_1, \ldots, A_{\hbar} , such that, when an higher-level algorithm produces a pruned set, every lower-level algorithm works on the remaining graph where the pruned set is taken out from U. The high-level intuition for de-amortizing Multi-level Pruning is to "delay" the work in each algorithm, so that there are enough updates for the algorithm to distribute their work on. For this to be accomplished, we will have to incur a multiplicative loss of ψ in the update time.

Let U be the cluster that we run the Multi-level Pruning process on. We denote by D the sequence of updates on G that are relevant to U. For a pair of integers $1 \le i < j \le N$, we denote D(i, j) as the subsequence of D from the *i*th update to the *j*th update (including both). We call such a subsequence a *batch*. For a cluster W of vertices that is (α', ϕ') -linked in G upon the (i-1)th update, we denote by $\mathcal{B}(W, i, j, \alpha', \phi')$ to be the run of the algorithm in Theorem 6.1, starting with the cluster W in G with boundary-linkedness parameters α' and ϕ' , handling the updates in D(i, j). Note that, an update on G may be irrelevant of W (i.e., both endpoints of the updated edge are not in W). In this case we simply ignore this update in the run of $\mathcal{B}(W, i, j, \alpha', \phi')$.

Recall that the input for the Multi-level Pruning process is a cluster that is (α', ϕ') -linked in G for parameters α', ϕ' that are known to the algorithm, and the Multi-level Pruning process handles the next $N = \phi' \operatorname{vol}_G(U)/1200$ updates that are relevant to U. Recall that $\psi = 2^{O(\sqrt{\log n})}$ and $\hbar = \log_{\psi} N$.

The new algorithm consists of a hierarchy of $\hbar - 1$ sub-algorithms $A'_{\hbar-1}, A'_{\hbar-2}, \ldots, A'_1$. We first describe the work of sub-algorithm $A'_{\hbar-1}$. Recall that in Section 7, the work of A_{\hbar} is divided into stages with length $\ell_{\hbar-1} = \psi^{\hbar-1}$ each. Similarly, the work of $A'_{\hbar-1}$ is also divided into stages with length $\ell_{\hbar-1}$ each. However, the work in each stage is now completely different. In the first stage, $A'_{\hbar-1}$ does nothing. For each $1 \leq t \leq N/\ell_{\hbar-1}$, note that, at the beginning of the (t+1)th stage, the batch $D(1, t\ell_{\hbar-1})$ of updates has completely shown up. We simply let $A'_{\hbar-1}$ run $\mathcal{B}(U, 1, t\ell_{\hbar-1}, \alpha', \phi')$ in its (t+1)th stage, with the work evenly distributed upon all updates in this stage. And after this stage is finished, $A'_{\hbar-1}$ sends the pruned set $P^{\hbar-1}_{t\ell_{\hbar-1}}$ to $A'_{\hbar-2}$. Intuitively, the sub-algorithm $A'_{\hbar-1}$ processes batches of size $[\psi^{\hbar-1}, \psi^{\hbar}]$, and is always " $\psi^{\hbar-1}$ updates late" compared to the current update. From the above discussion, in a stage of $A'_{\hbar-1}$, at most $N = \psi^{\hbar}$ updates are handled.

We now describe, for each $2 \leq s \leq \hbar - 2$, the work of sub-algorithm A'_s , which is similar to $A'_{\hbar-1}$. The work of sub-algorithm A'_s is divided into stages with length ℓ_s each. In the first stage, it does nothing. For each $1 \leq t \leq 2\psi - 1$, note that at the end of the *t*th stage on A'_s , the batch $D(1, t\ell_s)$ of updates has completely shown up, and A'_s has not received anything from A'_{s+1} . In the (t+1)th stage, we simply let A'_s run $\mathcal{B}(U, 1, t\ell_s, \alpha', \phi')$ in its (t+1)th stage, with the work evenly distributed upon all updates in this stage. And after this stage is finished, A'_s sends the pruned set $P^s_{t\ell_s}$ to A'_{s-1} . Starting from the $(2\psi + 1)$ th stage on A'_s , we call every next ψ stages on A'_s a phase of A'_s . Note that, from the description of the work on A'_{s+1} . We now describe the work of A'_s within this phase. For each $0 \leq t \leq \psi - 1$, in the (t+1) stage within this phase. For each $0 \leq t \leq \psi - 1$, in the (t+1) stage within this phase. For each $0 \leq t \leq \psi - 1$, in the (t+1) stage within this phase. For each $0 \leq t \leq \psi - 1$, in the (t+1) stage within this phase. This completes the work of $A'_s R^{h-s-1}$, $\phi'/3R^{h-s-1}$, $\phi'/3R^{h-s-1}$, $\phi'/3R^{h-s-1}$, which is stage on A'_s is finished, A'_s sends the pruned set $P^s_{t\ell_s}$ to A'_{s-1} . Intuitively, the

subalgorithm A_s process batches of size $[\psi^s, 2\psi^{s+1}]$, and is always " ψ^s updates late" compared with the current update. From the above discussion, in a stage of A'_s , at most $2\psi^{s+1}$ updates are handled.

It remains to describe the work of A'_1 . While the sub-algorithms $A'_2, A'_3, \ldots, A'_{\hbar-1}$ can be one-stage late, the work on A'_1 has to be up-to-date. The work of sub-algorithm A'_1 is also divided into stages with length $\ell_1 = \psi$ each. In the first 2ψ stages, upon the *i*th update, the sub-algorithm A'_1 simply runs $\mathcal{B}(U, 1, i, \alpha'/38^{\hbar-2}, \phi'/38^{\hbar-2})$. Starting from the $(2\psi + 1)$ th stage on A'_1 , we call every next ψ stages on A'_1 a phase of A'_1 . Note that, from the description of the work on A'_1 , for each $t' \ge 0$, at the beginning of the (t'+1)th phase on A'_1 , it receives a set $P^2_{t'\ell_2}$ from A'_2 . We now describe the work of A'_s within this phase. For each $0 \le i \le \psi^2 - 1$, upon the *i*th update stage within this phase, we let it run the entire process $\mathcal{B}(\overline{P^2_{t'\ell_2}}, t'\ell_{s+1} + 1, t'\ell_{s+1} + i, \alpha'/38^{\hbar-2}, \phi'/38^{\hbar-2})$. Put in other words, within the phase, upon each update, the machine M_1 makes an individual run of the algorithm in Theorem 6.1 handling all updates in this phase (from the first update in this phase to the current update). This completes the work on A'_1 . From the above discussion, upon each update, A'_1 handles a batch of at most $2\psi^2$ updates.

From the discussion, the worst-case update time in this de-amortized Multi-level Pruning is at most $O(\psi^2)$ -factor larger than the amortized update time of Multi-level Pruning. Therefore, the worst-case update-time $O(\psi^2 \cdot 38^{2h}/\phi'^2)$.

De-amortize Cluster Decomposition. Recall that the input to Cluster Decomposition process is a cluster that is (α, ϕ') -linked in the current graph G. Also recall that, the CD-process contains a main Multi-level Pruning process, and additionally, for each set \tilde{P}^s in the collection $\{\tilde{P}^1, \ldots, \tilde{P}^{\hbar}\}$ of sets maintained by the Multi-level Pruning process, the CD-process recomputes an (α, ϕ') -expander decomposition on \tilde{P}^s every time it changes (namely, every ψ^s updates on U). For obtaining the expander decomposition in time, we tweak the de-amortized Multi-level Pruning a bit, by letting each sub-algorithm A'_s runs, in each phase, not only a process of \mathcal{B} handling a batch of updates, but also an (α, ϕ') -expander decomposition on the pruned out set of vertices, after completing the process of \mathcal{B} , with the total work of both tasks evenly distributed on all updates in this stage. Note that this increase the worst-case update time by O(1)-factor.

De-amortize the algorithm for Lemma 7.3. Recall that the algorithm for Lemma 7.3 simply first computes an (α, ϕ) -expander decomposition, and then, for each cluster in the (α, ϕ) -expander decomposition, it starts a CD-process on it with respect to the well-linkedness parameter (α, ϕ') of this cluster. We have already shown how to de-amortize the CD-process. However, to completely de-amortize the algorithm for Lemma 7.3, we need one more step. Note that when the CD-process on a cluster U has handled $N = \phi' \operatorname{vol}_G(U)/1200$ updates, the cluster U will be reset. In particular, the algorithm will recompute an (α, ϕ) -expander decomposition. This cluster-resetting step needs to be de-amortized as well.

In order to achieve this, we run three CD-processes on the same cluster U in parallel, each maintaining an $(\alpha/38^h, \phi/38^h)$ -expander decomposition of U. At any time, one of the CD-process is used by the algorithm (that we call *online*), and the others are temporarily not (that we call *in the background*). When the online CD-process terminates, we switch it into the background, and bring online another CD-process that was in the background. We carefully choose the "offset" between these CD-process and schedule their work so that at any time, the online CD-process maintains an available decomposition of the current cluster.

We now describe the algorithm in more detail. We maintain 3 tweaked CD-process in parallel. Each tweaked CD-process has three phases: the preparing phase; the chasing phase; and the working phase; each spans the time of a consecutive N/3 updates (recall that N is the number of updates that can be handled by a CD-process). The offset between each pair of tweaked CD-process is also N/3. Assume the input is a cluster U that is (α, ϕ') -linked in G. Assume that some tweaked CD-process starts at the kth update, and we denote by G_k the graph after the kth update. In the first phase of the tweaked CD-process, the preparing phase, it computes an (α, ϕ') -expander decomposition of U in G_k . In the second phase, the chasing phase, it handles the batch D(k + 1, k + 2N/3) of updates. Note that, before this phase, the tweaked CD-process is N/3 updates behind, and after this phase, the tweaked CD-process manages to maintain an $(\alpha/38^h, \phi'/38^h)$ -expander decomposition of the up-to-date graph. Intuitively, this can be achieved by running a normal CD-process with double speed. In the third phase, the working phase, it runs a normal CD-process to handle the batch D(k + 2N/3 + 1, k + N) of updates. The work in the first and the second phases is evenly distributed over all updates in that phase. Each tweaked CD-process is online only at its working phase. It is not hard to see that, the combination of three tweaked CD-process defined above maintains an $(\alpha/38^h, \phi'/38^h)$ -expander decomposition of the up-to-date graph, and achieves the worst-case update time within a O(1)-factor of the worst-case update time of a normal de-amortized CD-process described above. Therefore, the worst-case update time is $O(\psi^2 \cdot 38^{2h}/\phi'^2)$.

Constructing the algorithm for Theorem 8.5. Recall that the algorithm in Lemma 7.3 maintains an $(\alpha/38^h, \phi'/38^h)$ -expander decomposition, such that the amortized recourse in the contracted graph with respect to the decomposition is $\rho = O(38^h \cdot \psi/\alpha)$. However, to construct an algorithm for Theorem 8.5 using the de-amortized algorithm for Lemma 7.3, we need to ensure that the worst-case recourse is $2^{O(\sqrt{\log n})}$, preferably $O(\rho)$. This can be achieved by further tweaking the de-amortized algorithm for Lemma 7.3 a bit. Specifically, for each *i* and in each stage of G^i , we not only distribute the running time evenly over all updates, but also distribute the recourse that is needed to propagate to the graph G^{i-1} at one-level above. In this way, we ensure that the worst-case recourse for the graph G^i is at most $O(\rho)^i$, thus achieving the worst-case update time $O(\psi^2 \cdot 38^{2h}/\phi'^2) \cdot O(\rho)^i = 2^{-O(\log^{3/4} n)}$.

9 Applications

In this section we show that our dynamic expander hierarchy almost directly leads to a number of applications in dynamic graph algorithms.

9.1 Dynamic Tree Flow Sparsifier

We start by reviewing the notion of tree flow sparsifiers. Given a weighted graph G = (V, E, c)and a subset $S \subseteq V$ of vertices, a set D of demands on S is a function $D: S \times S \to \mathbb{R}_{\geq 0}$, that specifies, for each pair $u, v \in S$ of vertices, a demand D(u, v). Given a subset $S \subseteq V$ and a set D of demands on S, a routing of D in G is a flow F on G, where for each pair $u, v \in S$, the amount of flow that F sends from u to v is D(u, v). We define the congestion $\eta(G, D)$ of a set D of demands in G to be the minimum congestion of a flow F that is a routing of D in G. We say that a tree T is a tree flow sparsifier of quality q for G with respect to S, if $S \subseteq V(T)$, and for any set D of demands on S, $\eta(T, D) \leq \eta(G, D) \leq q \cdot \eta(T, D)$. A tree flow sparsifier H of Gw.r.t. the subset V(G) is just called a tree flow sparsifier for G.

We design an algorithm that explicitly maintains a tree flow sparsifier for a graph G that undergoes edge insertions and deletions, which proceeds as follows: given an unweighted dynamic graph G on n vertices, maintain a $(2^{-\bar{O}(\log^{2/3} n)}, 2^{-O(\log^{5/6} n)})$ -expander hierarchy with slack $2^{\bar{O}(\log^{1/2} n)}$ of G using Theorem 8.2.

We immediately obtain the following result, which proves Corollary 1.2 from the introduction.

Corollary 9.1. There is a deterministic fully dynamic algorithm on a graph G with n vertices that explicitly maintains a tree flow sparsifier for G with quality $2^{O(\log^{5/6} n)}$ and depth $O(\log^{1/6} n)$ using $2^{O(\log^{5/6} n)}$ amortized update time.

Proof. To bound the quality of the tree flow sparsifier, the main observation is that an expander hierarchy of a graph G is itself a tree flow sparsifier for G. Concretely, let $\alpha := 2^{-\bar{O}(\log^{2/3} n)}$, $\phi := 2^{-O(\log^{5/6} n)}$ and $s := 2^{\bar{O}(\log^{1/2} n)}$. By Theorem 8.2, the depth of (α, ϕ) -expander hierarchy we maintain is $t := O(\log^{1/6} n)$. Using Theorem 5.2, it follows that our (α, ϕ) -expander hierarchy of G with slack s and depth t is a tree flow sparsifier for G with quality $O(s \log m)^t \cdot O(\max\{\frac{1}{\alpha}, \frac{1}{\phi}\}/\alpha^{t-1}) = 2^{O(\log^{5/6} n)}$.

Since we can maintain an (α, ϕ) -expander hierarchy with slack s of G in $2^{O(\log^{5/6} n)}$ amortized update time (Theorem 8.2), it follows that the amortized update time for maintaining a tree flow sparsifier for G is also bounded by $2^{O(\log^{5/6} n)}$.

9.2 Dynamic Vertex Flow Sparsifiers, Maximum Flow, Multi-commodity Flow, Multi-Way Cut and Multicut

We show that a dynamic tree flow sparsifier can be used to maintain a tree flow sparsifier with w.r.t. a subset S (also known as *vertex flow sparsifiers*), an approximation to the value of the following problems (i) maximum flow/minimum cut, (ii) maximum concurrent (multi-commodity) flow, (iii) multi-way cut and (iv) multicut.

In the dynamic vertex flow sparsifier¹⁰ problem, the graph G undergoes insertions or deletions of edges and the following queries are supported: given any subset $S \subseteq V(G)$, return a tree flow sparsifier for G w.r.t. S. The main idea behind designing an algorithm for this problem is the observation that given a tree flow sparsifier for G, one can easily extract a tree flow sparsifier for G w.r.t. any subset $S \subseteq V(G)$. Concretely, given an unweighted dynamic graph G on n vertices, let T be the maintained tree flow sparsifier for G from Corollary 9.1. For a vertex pair u, v, let $T_{u,v}$ denote the (unique) path between u and v in T. Upon receiving a query associated with an arbitrary subset $S \subseteq V(G)$, we do the following:

- Construct the subtree $T' := \bigcup_{u \in S} T_{u,r_T}$ that consists of all the paths from vertices in S to the root r_T of T.
- Return T'.

We immediately obtain the following result, which proves the third item of Corollary 1.3 from the introduction.

Corollary 9.2. There is a deterministic fully dynamic algorithm on a graph G with n vertices such that given a query associated with an arbitrary $S \subseteq V(G)$ outputs a tree flow sparsifier with quality $2^{O(\log^{5/6} n)}$ for G w.r.t. S using $2^{O(\log^{5/6} n)}$ amortized update time and $O(|S| \log^{1/6} n)$ query time. Moreover, the update time can be made worst-case while keeping the same quality and running time guarantees.

Proof. We first show that the output tree T' is a tree flow sparsifier with quality 1 for T w.r.t. S. Since T is a tree, every demand among two leaf vertices u, v in T is routed according to the unique path $T_{u,v}$ between u and v in T. If $u, v \in S$, note that $T_{u,v}$ is entirely contained in

¹⁰In general, vertex sparsifiers that preserve the (multi-commodity) flow between terminal vertices are not restricted to tree instances. However, as a byproduct of our techniques, the vertex sparsifiers we consider in this paper are always trees.

the sub-tree $T' = \bigcup_{u \in S} T_{u,rt}$. Therefore, every demand that we route in T between any vertex pair u, v in S, can also be routed in T' with the same congestion. For the other reduction, by construction we have that $T' \subseteq T$, i.e., every demand that we route in T' between any vertex pair in S can be routed in T with the same congestion. Combining the above gives that T' is a tree flow sparsifier with quality 1 for T w.r.t. S. As T is a tree flow sparsifier with quality $2^{O(\log^{5/6} n)}$ for G (Corollary 9.1), by the transitivity property of flow sparsifiers, it follows that T' is a tree flow sparsifier with quality $2^{O(\log^{5/6} n)}$ for G w.r.t. S.

We next analyze the running time. The claimed amortized update time follows directly from Corollary 9.1. For the query time, Corollary 9.1 ensures that at any time the depth of T is $O(\log^{1/6} n)$. The latter guarantees that the length of each path from a leaf vertex to the root in T is $O(\log^{1/6} n)$, which in turn implies that the time to compute T' and its size are both bounded by $O(|S| \log^{1/6} n)$.

To achieve our worst-case update time, we replace the expander hierarchy from Theorem 8.2 with the one from Thereom 8.1, which in turn allows us to query for any given vertex u, the leaf-to-root path of u in the hierarchy. Since we only need such paths for the construction of T', our claim follows.

The above corollary readily implies a fully-dynamic algorithm for the all-pair approximate maximum flow problem: upon receiving a query associated with an arbitrary vertex pair $u, v \in V$ we let $S = \{u, v\}$ and then compute a tree flow sparsifier T' for G w.r.t. S using Corollary 9.2. Finally, we compute the maximum flow from u to v in T' and return its value as an estimate. We have the following result, which proves the first item of Corollary 1.3 from the introduction.

Corollary 9.3. There is a deterministic fully dynamic algorithm on a graph G with n vertices that maintains for every vertex pair $u, v \in V$, an estimate that approximates the maximum flow from u to v in G up to a factor of $2^{O(\log^{5/6} n)}$ using $2^{O(\log^{5/6} n)}$ worst-case update time and $O(\log^{1/6} n)$ query time.

We next show that the same idea extends to the maximum concurrent (multi-commodity) flow problem, which is defined as follows: given an unweighted graph G and k source-sink pairs s_i, t_i , each associated with a non-negative demand D(i), compute the congestion $\eta(G, D)$, i.e., the minimum congestion a flow F that is a routing of D in G, where $D := (D(1), \ldots, D(k))$. We study a dynamic version of the problem, where G undergoes edge updates and the k source-sink pairs are made available only at query time. Our dynamic construction uses Corollary 9.2, and whenever the k source-sink pairs are revealed to us, we define $V_k = \bigcup_i \{s_i, t_i\}$ and then compute a tree flow sparsifier T' for G w.r.t. V_k . Finally, we compute the congestion $\eta(T, D)$ in T and return this value as an estimate. The result below follows from the definition of tree flow sparsifiers and proves the third item of Corollary 1.3 from the introduction.

Corollary 9.4. There is a fully dynamic deterministic algorithm on a graph G with n vertices that maintains for every demand set D defined on k source-sink pairs s_i, t_i , an estimate that approximates $\eta(G, D)$ up to a factor of $2^{O(\log^{5/6} n)}$ using $2^{O(\log^{5/6} n)}$ worst-case update time and $O(k \log^{1/6} n)$ query time.

We finally consider a dynamic version of the multi-way cut problem, which is defined as follows. Given an unweighted graph G and k distinguished vertices s_1, \ldots, s_k , the goal is to remove a minimum number of edges F such that no pair of distinguished vertices s_i and s_j with $i \neq j$ belong to the same connected component after the removal of F from G. We study a dynamic version of the problem, where G undergoes edge updates and the k distinguished vertices are made available only at query time. Similarly to above, we use Corollary 9.2 and whenever the k distinguished vertices are revealed to us, we define $V_k = \bigcup_i \{s_i\}$ and then compute a tree flow sparsifier T' for G w.r.t. V_k . Finally, we compute an optimal solution to the multi-way cut problem on T' with respect to the queried k distinguished vertices and return this value as an esimate. The result below follows from the definition of tree flow sparsifiers and proves the third item of Corollary 1.3 from the introduction.

Corollary 9.5. There is a fully dynamic deterministic algorithm on a graph G with n vertices that maintains for any k distinguished vertices s_1, \ldots, s_k , an estimate that approximates an optimal solution to the multi-way cut up to a factor of $2^{O(\log^{5/6} n)}$ using $2^{O(\log^{5/6} n)}$ worst-case update time and $O(k \log^{1/6} n)$ query time.

The dynamic multicut essentially follows the same idea and we omit it here for the sake of brevity.

9.3 Dynamic Sparsest Cut and Lowest Conductance Cut

We show that a dynamic tree flow sparsifier can be used to maintain sparsest cuts, multi-cuts and multi-way cuts. Throughout, we only focus on the dynamic sparsest cut problem. An almost identical idea extends to the lowest conductance cut but we omit a detailed description here for the sake of brevity.

Let G = (V, E, c) be a weighted graph. For any cut (S, \overline{S}) such that $|S| \leq |\overline{S}|$, let $c(\delta(S))$ be the sum over capacites of all edges with one endpoint in S and the other in \overline{S} , where $\overline{S} = V \setminus S$. Let $\alpha(G, S) := c(\delta(S))/|S|$ be the sparsity of (S, \overline{S}) . The *sparsest cut* problem asks to find a cut (S, \overline{S}) such that $|S| \leq |\overline{S}|$ with smallest possible sparsity in G, which we denoted by $\alpha(G)$. We study a dynamic version of this problem, where G undergoes edge updates and at query time we need to report the sparsity $\alpha(G)$ of the sparsest cut in the current graph G. To design a dynamic algorithm, we follow a well-known approach used to solve the static version of the problem: given a graph G, (1) compute a tree flow sparsifier T with quality q for G and (2) solve the sparsest cut problem on T. Since a tree flow sparsifier is also a tree cut sparsifier with the same quality, it is easy to verify that $\alpha(T)$ approximates $\alpha(G)$ up to a factor of q. The main advantage of this approach is that computing sparsest cut on trees is much easier.

To see this, consider a (rooted) tree flow sparsifier $T = (V(T), E(T), c^T)$ with quality q and depth t for G such that the leaf nodes of T correspond to the vertices of G. It is known that the sparsest cut on a tree must occur at one of the edges in T. We can also build a data-structure such that given an internal node x in T (except the root), it reports the number of leaf nodes in the sub-tree rooted at x. Using these two observations, an algorithm for computing $\alpha(T)$ works as follows:

- For each edge $e = (x, p(x)) \in E(T)$ (as T is rooted), where p is the parent of x, compute the sparsity of the cut (S, \overline{S}) obtained by removing (x, p(x)) in T using $c^{T}(e)/|S|$, where |S| is precisely the number of leaf nodes in the sub-tree rooted at x.
- Return $\min_{e \in E(T)} c^T(e) / |S|$.

In a similar vein, using Corollary 9.1 we maintain a tree flow sparsifier T for an unweighted dynamic graph G. As T undergoes changes, we additionally update the information about the number of leaf nodes at an internal node and the edge with the smallest sparsity in T. Since these updates can be implemented in time proportional to the time needed to maintain T, we obtain the following result, which proves the second item of Corollary 1.3 from the introduction.

Corollary 9.6. There is a deterministic fully dynamic algorithm on a graph G with n vertices that maintains an estimate that approximates $\alpha(G)$ up to a factor of $2^{O(\log^{5/6} n)}$ using $2^{O(\log^{5/6} n)}$ amortized update time and $O(\log^{1/6} n)$ query time.

9.4 Dynamic Connectivity

We observe that the data-structure representation of the expander hierarchy from Theorem 8.1 leads to a dynamic algorithm for maintaining connectivity information of G. More precisely, a graph G is connected iff the top level our expander hierarchy consists of a single vertex. Moreover, two vertices u and v are connected iff the roots of u and v in the hierarchy are the same. These observations lead to the following result, which proves Corollary 1.4 from the introduction.

Corollary 9.7. There is a deterministic fully dynamic algorithm on a n-vertex graph G that maintains connectivity of G using $2^{-O(\log^{5/6} n)}$ worst-case update time and also supports pairwise connectivity queries in $O(\log^{1/6} n)$ time.

9.5 Treewidth decomposition

A treewidth decomposition T of a graph G = (V, E) is a tree such that each node x in T corresponds to a set $B_x \subseteq V$ of vertices called a *bag*. For each edge $(u, v) \in E$, there must exist a node x whose bag B_x contains both u and v. Moreover, for each vertex $u \in V$, $\{x \mid u \in B_x\}$ must induce a connected subtree of T. A width of T is $\max_x |B_x| - 1$. The treewidth tw(G) of G is the minimum width over all treewidth decomposition of G.

We obtain the first dynamic algorithm for maintaining a tree width decomposition. The main observation behind our construction is that a treewidth decomposition of a graph can be directly derived from an expander hierarchy, which works as follows. Let T be a expander hierarchy of a graph G = (V, E). We simply let T itself be the treewidth decomposition. It remains to define a bag B_x for each node $x \in T$.

To this end, for each node $x \in T$, let U be a cluster from T corresponding to a node x. Recall that $U \subseteq V(G^i)$ for some i. Let $E_{G_i}(U, V(G^i))$ denote the set of edges in G^i incident to a vertex from U. For each $e^i \in E_{G_i}(U, V(G^i))$, there is a corresponding "original" edge e of G. The bag $B_x \subseteq V$ consists of the endpoints of all "original" edges correspond to edges from $E_{G_i}(U, V(G^i))$. See Figure 1 for an example.



Figure 1: An illustration of a bag B_x of a node x in an expander hierarchy.

Lemma 9.8. The expander hierarchy T is a treewith decomposition of G.

Proof. Observe that for each edge $(u, v) \in E$ there is a unique cluster graph C in T containing a edge $e' \in E(C)$ corresponding to e. Let x be the node in T corresponding to C. It is clear that the bag B_x containing both u and v.

Next, suppose towards a contradiction that there is a vertex u where the set $\{x \mid u \in B_x\}$ does not induce a connected subtree of T. Let T_1 and T_2 be two disconnected induced subtrees. Let y be a node in a path connecting T_1 and T_2 such that y is neither in T_1 nor in T_2 . Observe that the bag B_y is a separator that separates vertices in the bags in T_1 and T_2 . More precisely, let $V_1 = \bigcup_{x \in T_1} B_x$ and $V_2 = \bigcup_{x \in T_2} B_x$. Observe that in the graph $G[V \setminus B_y]$, no pair of vertices between $V_1 \setminus B_y$ and $V_2 \setminus B_y$ can be connected. However, we have that $u \in V_1 \setminus B_y$ and $u \in V_2 \setminus B_y$, which is a contradiction.

To bound the width of our treewidth decomposition, we need the notions of well-linkedness and flow-linkedness.

Definition 9.9. A set of $S \subset V(G)$ is γ -well-linked in G iff any cut (A, B) in G, $|E(A, B)| \ge \gamma \cdot \min\{|A \cap S|, |B \cap S|\}$.

Definition 9.10. A set $S \subset V(G)$ is γ -flow-linked in G if given any multi-commodity flow demand D on S where the total demand on each vertex $v \in S$ is at most 1, the congestion for routing D in G is at most $\eta(G, D) \leq 1/\gamma$.

It is easy to see that any γ -flow-linked set in G is a γ -well-linked set in G. The next fact relates well-linkedness and treewidth in a bounded degree graph.

Fact 9.11 (A paraphrase of Corollary 2.1 from [CC13]). Let G = (V, E) be a graph with maximum degree Δ . Let $B \subseteq V$ be a set of vertices that is γ -well-linked in G. Then $\operatorname{tw}(G) \geq \frac{\gamma |B|}{3\Delta} - 1$ or equivalently $|B| = O(\frac{\Delta}{\gamma} \cdot \operatorname{tw}(G))$.

The following key technical lemma relates the notion of tree flow sparsifiers to the notion of treewidth via flow-linkedness.

Lemma 9.12. If T is a tree flow sparsifier with quality q, then each bag B_x is $\Omega(1/q)$ -flow-linked. *Proof.* Let D be any demand D on B_x where the total demand on each vertex $v \in B_x$ is at most 1. It suffices to show D is routable in T, i.e. $\eta(T, D) \leq 1$. This is because T has quality q, so D can be routed in G with congestion q, i.e. $\eta(G, D) \leq q$.

The crucial observation is that all vertices $v \in B_x$ can route one unit of flow in T to x simultaneously without congestion. The latter holds since each $v \in B_x$ is an endpoint of some boundary edge e of a cluster correspond to the node x or the children of x in T. Therefore, the edge e contributes to one unit capacity to every tree-edge in the path from v to x in T.

Now, to route D in T, each vertex $v \in B_x$ just sends flow (equal to its total demand) of at most one unit to x, which causes no congestion. Connecting the all flow paths that meet at x completes the proof of the lemma.

As an expander hierarchy is a good quality tree flow sparsifiers, our construction of treewidth decomposition has small width. This fact is summarized as follows:

Corollary 9.13. Let G be a constant degree graph and let T an (α, ϕ) -expander hierarchy of G with depth t and slack s. Then each bag B_x has size at most $\operatorname{tw}(G) \cdot O(s \log m)^t \cdot O(\max\{\frac{1}{\alpha}, \frac{1}{\phi}\}/\alpha^{t-1})$.

Proof. By Theorem 5.2, T has quality $q = O(s \log m)^t \cdot O(\max\{\frac{1}{\alpha}, \frac{1}{\phi}\} \cdot \frac{1}{\alpha^{t-1}})$, and thus each bag B_x is $\Omega(1/q)$ -flow-linked by Lemma 9.12, and hence also $\Omega(1/q)$ -well-linked. Finally, by Fact 9.11, $|B_x| \leq O(\operatorname{tw}(G) \cdot q)$ as desired.

Our dynamic algorithm for treewidth decomposition proceeds as follows. We maintain a (α, ϕ) -expander hierarchy T with slack $s := 2^{\overline{O}(\log^{1/2} n)}$ and depth $t := O(\log^{1/6} n)$ of G using Theorem 8.2, where $\alpha := 2^{-\overline{O}(\log^{2/3} n)}$, $\phi := 2^{-O(\log^{5/6} n)}$. Using Corollary 9.13 and observing that we can explicitly update all the bags within the same running time guaranteed by Theorem 8.2, we get the following result which proves Corollary 1.5 from the introduction.

Corollary 9.14. There is a deterministic fully dynamic algorithm on a constant degree graph G with n vertices that maintains a treewidth decomposition of G with width $tw(G) \cdot 2^{O(\log^{5/6} n)}$ using $2^{O(\log^{5/6} n)}$ amortized update time.

A Proof of Lemma 5.1

We show this via an approximate maxflow-mincut theorem. It is well known [LR99] that the optimum congestion required for solving a multicommodity flow problem with demands D in an undirected graph G = (V, E) is at most $O(\log n / \operatorname{sparsity}(G, D))$, where $\operatorname{sparsity}(G, D) = \max_{X \subseteq V} |E_G(X, V \setminus X)| / D(X, V \setminus X)$, and $D(X, V \setminus X) = \sum_{(x,y) \in X \times V \setminus X} (D(x, y) + D(y, x))$ is the demand that has to cross the cut X.

With this in mind we prove the lemma by showing that D has low sparsity in G[S]. Fix a subset X. The demand that originates at vertices in X is at most $\sum_x \gamma \deg_G(x)$ because the demand is γ -restricted. The same holds for the demand that ends at vertices in X. This means that the total demand that can cross the cut can be at most $\gamma \min\{\operatorname{vol}_G(X), \operatorname{vol}_G(V \setminus X)\}$. But since $G[S]^{\alpha/\phi}$ is a ϕ -expander we know that $|E_G(X, S \setminus X)| \geq \phi \min\{\operatorname{vol}_{G[S]^{\alpha/\phi}}(X), \operatorname{vol}_{G[S]^{\alpha/\phi}}(S \setminus X)\} \geq \phi \min\{\operatorname{vol}_G(X), \operatorname{vol}_G(S \setminus X)\}$. Using approximate maxflow-mincut gives the first statement.

For the second statement the demand that has to cross the cut can be at most

$$D(X, S \setminus X) \leq \gamma \cdot \min\{\sum_{v \in X} |E_G(\{v\}, V \setminus S)|, \sum_{v \in S \setminus X} |E_G(\{v\}, V \setminus S)|\}$$

$$\leq \gamma \cdot \frac{\phi}{\alpha} \cdot \min\{\operatorname{vol}_{G[S]^{\alpha/\phi}}(X), \operatorname{vol}_{G[S]^{\alpha/\phi}}(S \setminus X)\}$$

$$\leq \gamma \cdot \frac{\phi}{\alpha} \cdot \frac{1}{\phi} |E_G(X, S \setminus X)| ,$$

where the first inequality is due to the γ -boundary restriction, the second due to the reweighting of boundary edges in the graph $G[S]^{\alpha/\phi}$ and the last inequality follows from the ϕ -expansion of $G[S]^{\alpha/\phi}$.

References

- [AC03] David Applegate and Edith Cohen. Making intra-domain routing robust to changing and uncertain traffic demands: understanding fundamental tradeoffs. In Proceedings of the ACM SIGCOMM 2003 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication, August 25-29, 2003, Karlsruhe, Germany, pages 313–324, 2003.
- [ACK⁺16] Alexandr Andoni, Jiecao Chen, Robert Krauthgamer, Bo Qin, David P. Woodruff, and Qin Zhang. On sketching quadratic forms. In Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science, Cambridge, MA, USA, January 14-16, 2016, pages 311–319, 2016.
- [ADK⁺16] Ittai Abraham, David Durfee, Ioannis Koutis, Sebastian Krinninger, and Richard Peng. On fully dynamic graph sparsifiers. In IEEE 57th Annual Symposium on Foundations of Computer Science, FOCS 2016, 9-11 October 2016, Hyatt Regency, New Brunswick, New Jersey, USA, pages 335–344, 2016.
- [AGG⁺09] Konstantin Andreev, Charles Garrod, Daniel Golovin, Bruce M. Maggs, and Adam Meyerson. Simultaneous source location. ACM Trans. Algorithms, 6(1):16:1–16:17, 2009.
- [AKPW95] Noga Alon, Richard M. Karp, David Peleg, and Douglas B. West. A graph-theoretic game and its application to the k-server problem. *SIAM J. Comput.*, 24(1):78–100, 1995.
- [Ami10] Eyal Amir. Approximation algorithms for treewidth. *Algorithmica*, 56(4):448–479, 2010.
- [AN12] Ittai Abraham and Ofer Neiman. Using petal-decompositions to build a low stretch spanning tree. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 395–406, 2012.
- $\begin{array}{lll} [\text{BDD}^+16] & \text{Hans L. Bodlaender, Pål Grønås Drange, Markus S. Dregi, Fedor V. Fomin, Daniel Lokshtanov, and Michal Pilipczuk. A c^k n 5-approximation algorithm for treewidth. \\ SIAM J. Comput., 45(2):317-378, 2016. \end{array}$
- [BFH19] Aaron Bernstein, Sebastian Forster, and Monika Henzinger. A deamortization approach for dynamic spanner and dynamic maximal matching. In Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019, pages 1899–1918, 2019.
- [BFK⁺14] Nikhil Bansal, Uriel Feige, Robert Krauthgamer, Konstantin Makarychev, Viswanath Nagarajan, Joseph Naor, and Roy Schwartz. Min-max graph partitioning and small set expansion. SIAM J. Comput., 43(2):872–904, 2014.
- [BGHK95] Hans L. Bodlaender, John R. Gilbert, Hjálmtyr Hafsteinsson, and Ton Kloks. Approximating treewidth, pathwidth, frontsize, and shortest elimination tree. J. Algorithms, 18(2):238–255, 1995.
- [BKR03] Marcin Bienkowski, Miroslaw Korzeniowski, and Harald Räcke. A practical algorithm for constructing oblivious routing schemes. In SPAA 2003: Proceedings of the Fifteenth Annual ACM Symposium on Parallelism in Algorithms and Architectures,

June 7-9, 2003, San Diego, California, USA (part of FCRC 2003), pages 24–33, 2003.

- [BL99] Yair Bartal and Stefano Leonardi. On-line routing in all-optical networks. *Theor. Comput. Sci.*, 221(1-2):19–39, 1999.
- [Bod96] Hans L. Bodlaender. A linear-time algorithm for finding tree-decompositions of small treewidth. *SIAM J. Comput.*, 25(6):1305–1317, 1996.
- [BvdBG⁺20] Aaron Bernstein, Jan van den Brand, Maximilian Probst Gutenberg, Danupon Nanongkai, Thatchaphol Saranurak, Aaron Sidford, and He Sun. Fully-dynamic graph sparsifiers against an adaptive adversary, 2020.
- [CC13] Chandra Chekuri and Julia Chuzhoy. Large-treewidth graph decompositions and applications. In Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013, pages 291–300, 2013.
- [CGH⁺20] Li Chen, Gramoz Goranci, Monika Henzinger, Richard Peng, and Thatchaphol Saranurak. Fast dynamic cuts, distances and effective resistances via vertex sparsifers, 2020. In submission to FOCS'20.
- [CGL⁺19] 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. *CoRR*, abs/1910.08025, 2019.
- [CHKM17] Keren Censor-Hillel, Bernhard Haeupler, Jonathan A. Kelner, and Petar Maymounkov. Rumor spreading with no dependence on conductance. SIAM J. Comput., 46(1):58–79, 2017.
- [CK19] Julia Chuzhoy and Sanjeev Khanna. A new algorithm for decremental single-source shortest paths with applications to vertex-capacitated flow and cut problems. In Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, pages 389–400, 2019.
- [CKS05] Chandra Chekuri, Sanjeev Khanna, and F. Bruce Shepherd. Multicommodity flow, well-linked terminals, and routing problems. In Proceedings of the 37th Annual ACM Symposium on Theory of Computing, Baltimore, MD, USA, May 22-24, 2005, pages 183–192, 2005.
- [CKS13] Chandra Chekuri, Sanjeev Khanna, and F. Bruce Shepherd. The all-or-nothing multicommodity flow problem. *SIAM J. Comput.*, 42(4):1467–1493, 2013.
- [CPZ19] Yi-Jun Chang, Seth Pettie, and Hengjie Zhang. Distributed triangle detection via expander decomposition. In Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019, pages 821–840, 2019.
- [CS19] Yi-Jun Chang and Thatchaphol Saranurak. Improved distributed expander decomposition and nearly optimal triangle enumeration. In Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing, PODC 2019, Toronto, ON, Canada, July 29 - August 2, 2019., pages 66–73, 2019.

- [CZ20] Shiri Chechik and Tianyi Zhang. Dynamic low-stretch spanning trees in subpolynomial time. 2020. To appear at SODA 2020.
- [DKT13] Zdenek Dvorák, Martin Kupec, and Vojtech Tuma. Dynamic data structure for tree-depth decomposition. *CoRR*, abs/1307.2863, 2013.
- [EFF⁺19] Talya Eden, Nimrod Fiat, Orr Fischer, Fabian Kuhn, and Rotem Oshman. Sublinear-time distributed algorithms for detecting small cliques and even cycles. In 33rd International Symposium on Distributed Computing, DISC 2019, October 14-18, 2019, Budapest, Hungary., pages 15:1–15:16, 2019.
- [FG19] Sebastian Forster and Gramoz Goranci. Dynamic low-stretch trees via dynamic low-diameter decompositions. In Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, Phoenix, AZ, USA, June 23-26, 2019., pages 377–388, 2019.
- [FLS⁺18] Fedor V. Fomin, Daniel Lokshtanov, Saket Saurabh, Michal Pilipczuk, and Marcin Wrochna. Fully polynomial-time parameterized computations for graphs and matrices of low treewidth. ACM Trans. Algorithms, 14(3):34:1–34:45, 2018.
- [FM06] Uriel Feige and Mohammad Mahdian. Finding small balanced separators. In Proceedings of the 38th Annual ACM Symposium on Theory of Computing, Seattle, WA, USA, May 21-23, 2006, pages 375–384, 2006.
- [GHS19] Gramoz Goranci, Monika Henzinger, and Thatchaphol Saranurak. Fast dynamic flows, cuts, distances via vertex sparsifers. 2019. unpublished.
- [GK18] Manoj Gupta and Shahbaz Khan. Simple dynamic algorithms for maximal independent set and other problems. *CoRR*, abs/1804.01823, 2018.
- [GR98] Oded Goldreich and Dana Ron. A sublinear bipartiteness tester for bunded degree graphs. In Proceedings of the Thirtieth Annual ACM Symposium on the Theory of Computing, Dallas, Texas, USA, May 23-26, 1998, pages 289–298, 1998.
- [HHR03] Chris Harrelson, Kirsten Hildrum, and Satish Rao. A polynomial-time tree decomposition to minimize congestion. In SPAA 2003: Proceedings of the Fifteenth Annual ACM Symposium on Parallelism in Algorithms and Architectures, June 7-9, 2003, San Diego, California, USA (part of FCRC 2003), pages 34–43, 2003.
- [JS18] 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, SODA 2018, New Orleans, LA, USA, January 7-10, 2018, pages 2487–2503, 2018.
- [JS20] Wenyu Jin and Xiaorui Sun. Fully dynamic *c*-edge connectivity in subpolynomial time. *arXiv preprint arXiv:2004.07650*, 2020.
- [KLOS14] Jonathan A. Kelner, Yin Tat Lee, Lorenzo Orecchia, and Aaron Sidford. An almost-linear-time algorithm for approximate max flow in undirected graphs, and its multicommodity generalizations. In Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014, pages 217–226, 2014.

- [KR96] Jon M. Kleinberg and Ronitt Rubinfeld. Short paths in expander graphs. In 37th Annual Symposium on Foundations of Computer Science, FOCS '96, Burlington, Vermont, USA, 14-16 October, 1996, pages 86–95, 1996.
- [KRV09] Rohit Khandekar, Satish Rao, and Umesh Vazirani. Graph partitioning using single commodity flows. *Journal of the ACM (JACM)*, 56(4):19, 2009.
- [KSS18] Akash Kumar, C. Seshadhri, and Andrew Stolman. Finding forbidden minors in sublinear time: A n¹/2+o(1)-query one-sided tester for minor closed properties on bounded degree graphs. In 59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7-9, 2018, pages 509–520, 2018.
- [KVV04] Ravi Kannan, Santosh S. Vempala, and Adrian Vetta. On clusterings: Good, bad and spectral. J. ACM, 51(3):497–515, 2004.
- [LR99] Frank Thomson Leighton and Satish Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. J. ACM, 46(6):787– 832, 1999.
- [Mad10] 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, 2010.
- [NS17] Danupon Nanongkai and Thatchaphol Saranurak. Dynamic spanning forest with worst-case update time: adaptive, las vegas, and $o(n^{1/2} \epsilon)$ -time. In Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017, pages 1122–1129, 2017.
- [NSW17] Danupon Nanongkai, Thatchaphol Saranurak, and Christian Wulff-Nilsen. Dynamic minimum spanning forest with subpolynomial worst-case update time. In 58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017, pages 950–961, 2017.
- [Pen16] Richard Peng. Approximate undirected maximum flows in O(mpolylog(n)) time. In Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016, pages 1862– 1867, 2016.
- [PT07] Mihai Patrascu and Mikkel Thorup. Planning for fast connectivity updates. In *FOCS*, pages 263–271. IEEE Computer Society, 2007.
- [Räc02] Harald Räcke. Minimizing congestion in general networks. In 43rd Symposium on Foundations of Computer Science (FOCS 2002), 16-19 November 2002, Vancouver, BC, Canada, Proceedings, pages 43–52, 2002.
- [Räc08] Harald Räcke. Optimal hierarchical decompositions for congestion minimization in networks. In Proceedings of the 40th Annual ACM Symposium on Theory of Computing, Victoria, British Columbia, Canada, May 17-20, 2008, pages 255–264, 2008.

- [Ree92] Bruce A. Reed. Finding approximate separators and computing tree width quickly. In Proceedings of the 24th Annual ACM Symposium on Theory of Computing, May 4-6, 1992, Victoria, British Columbia, Canada, pages 221–228, 1992.
- [RS95] Neil Robertson and Paul D. Seymour. Graph minors .xiii. the disjoint paths problem. J. Comb. Theory, Ser. B, 63(1):65–110, 1995.
- [RS14] Harald Räcke and Chintan Shah. Improved guarantees for tree cut sparsifiers. In Algorithms - ESA 2014 - 22th Annual European Symposium, Wroclaw, Poland, September 8-10, 2014. Proceedings, pages 774–785, 2014.
- [RST14] Harald Räcke, Chintan Shah, and Hanjo Täubig. Computing cut-based hierarchical decompositions in almost linear time. In Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014, pages 227–238, 2014.
- [She13] Jonah Sherman. Nearly maximum flows in nearly linear time. In *FOCS*, pages 263–269. IEEE Computer Society, 2013.
- [She17] Jonah Sherman. Area-convexity, l_{∞} regularization, and undirected multicommodity flow. In Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017, pages 452–460, 2017.
- [ST11] Daniel A. Spielman and Shang-Hua Teng. Spectral sparsification of graphs. *SIAM J. Comput.*, 40(4):981–1025, 2011.
- [SW19] 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, San Diego, California, USA, January 6-9, 2019, pages 2616–2635, 2019.
- [Tre05] Luca Trevisan. Approximation algorithms for unique games. In 46th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2005), 23-25 October 2005, Pittsburgh, PA, USA, Proceedings, pages 197–205, 2005.
- [vdBNS19] Jan van den Brand, Danupon Nanongkai, and Thatchaphol Saranurak. Dynamic matrix inverse: Improved algorithms and matching conditional lower bounds. CoRR, abs/1905.05067, 2019. To appear at FOCS 2019.
- [Wul17] Christian Wulff-Nilsen. Fully-dynamic minimum spanning forest with improved worst-case update time. In Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017, pages 1130–1143, 2017.