

# Near-Optimal Algorithm for Directed Expander Decompositions

Aurelio L. Sulser ✉

ETH Zurich, Switzerland

Maximilian Probst Gutenberg ✉ 

ETH Zurich, Switzerland

---

## Abstract

In this work, we present the first algorithm to compute expander decompositions in an  $m$ -edge *directed* graph with near-optimal time  $\tilde{O}(m)^1$ . Further, our algorithm can maintain such a decomposition in a dynamic graph and again obtains near-optimal update times. Our result improves over previous algorithms [2, 18] that only obtained algorithms optimal up to subpolynomial factors.

In order to obtain our new algorithm, we present a new push-pull-relabel flow framework that generalizes the classic push-relabel flow algorithm [14] which was later dynamized for computing expander decompositions in *undirected* graphs [16, 31]. We then show that the flow problems formulated in recent work [18] to decompose directed graphs can be solved much more efficiently in the push-pull-relabel flow framework.

Recently, our algorithm has already been employed to obtain the currently fastest algorithm to compute min-cost flows [33]. We further believe that our algorithm can be used to speed-up and simplify recent breakthroughs in combinatorial graph algorithms towards fast maximum flow algorithms [11, 12, 1].

**2012 ACM Subject Classification** Theory of computation → Dynamic graph algorithms

**Keywords and phrases** Directed Expander Decomposition, Push-Pull-Re-label Algorithm

**Digital Object Identifier** 10.4230/LIPIcs.ICALP.2025.132

**Category** Track A: Algorithms, Complexity and Games

**Funding** The research leading to these results has received funding from the starting grant “A New Paradigm for Flow and Cut Algorithms” (no. TMSGI2\_218022) of the Swiss National Science Foundation.

**Acknowledgements** The authors are thankful for many fruitful discussions and feedback on an earlier version of the paper from Rasmus Kyng, Simon Meierhans and Thatchaphol Saranurak.

## 1 Introduction

Over the past two decades, expanders and expander decompositions have been pivotal in advancing on fundamental algorithmic graph problems. The development and application of the first fast algorithm to compute near-expander decompositions was given in the development of the first near-linear time Laplacian solvers [32], a breakthrough in modern graph algorithms. Subsequently, a line of research [16, 36, 28, 29] has focused on strengthening this result by developing fast flow-based pruning techniques that refine near-expander decompositions into expander decompositions. This line of research culminated in [31] where a new, faster, simpler, and more user-friendly expander decomposition framework was presented. This advancement has catalyzed the widespread use of expander decompositions

---

<sup>1</sup> In this article, we use  $\tilde{O}(\cdot)$  notation to suppress factors logarithmic in  $m$ , i.e.  $O(m \log^c m) = \tilde{O}(m)$  for every constant  $c > 0$ .



as a tool in graph algorithms and was instrumental in the recent surge of applications of expander decompositions in both static and dynamic graph settings for various cut, flow, and shortest path problems [32, 21, 16, 36, 28, 29, 10, 2, 26, 5, 35, 30, 6, 25, 13, 15, 9, 3, 3, 24, 19, 34, 23, 7, 20, 12, 1, 33]. In this work, we study the problem of computing and maintaining expander decompositions in *directed* graphs, defined as follows.

► **Definition 1** (Directed Expander Decomposition). *Given an  $m$ -edge directed graph  $G$ , we say a partition  $\mathcal{X}$  of the vertex set of  $G$  and a subset of edges  $E^r \subseteq E$  forms an  $(\beta, \phi)$ -expander decomposition if*

1.  $\forall X \in \mathcal{X}$ ,  $G[X]$  is a  $\phi$ -expander meaning for all cuts  $(S, \bar{S})$ :  $\frac{\min\{e_G(S, \bar{S}), e_G(\bar{S}, S)\}}{\min\{\text{vol}_G(S), \text{vol}_G(\bar{S})\}} \geq \phi$ , and
2.  $|E^r| \leq \beta \cdot \phi \cdot m$ , and
3. the graph  $(G \setminus E^r)/\mathcal{X}$ , that is the graph  $G$  minus the edges in  $E^r$  where expander components in  $\mathcal{X}$  are contracted into supernodes, is a directed acyclic graph (DAG).

In our algorithm, we implicitly maintain an ordering of the partition sets in  $\mathcal{X}$  and let  $E^r$  be the edges that go 'backward' in this ordering of expander components. Note that we can only obtain a meaningful bound on the number of such 'backward' edges since a bound on *all* edges between expander components cannot be achieved as can be seen from any graph  $G$  that is acyclic (which implies that  $\mathcal{X}$  has to be a collection of singletons by Item 1 forcing all edges to be between components). Directed expanders and expander decompositions have been introduced in [2] in an attempt to derandomize algorithms to maintain strongly connected components and single-source shortest paths in directed graphs undergoing edge deletions. Recently, [18] gave an alternative algorithm to compute and maintain directed expander decomposition that refines the framework from [2] and heavily improves subpolynomial factors. Besides working on directed graphs, this algorithm also yields additional properties for expander decompositions that cannot be achieved with existing techniques - even in undirected graphs. In this article, we further refine these techniques to obtain an algorithm that is optimal up to logarithmic factors in  $m$  - as opposed to subpolynomial factors. Since their invention, directed expander decompositions and the techniques to maintain such decompositions have been pivotal in the design of fast dynamic graph data structures and in ongoing research for fast 'combinatorial' maximum flow algorithms. We discuss in Section 1.2 these recent lines of research and how our algorithm benefits recent breakthrough results.

## 1.1 Our Contribution

In this article, we finally give a simple algorithm that generalizes the algorithm from [31] in a clean way. Further, our algorithm is the first to obtain near-optimal runtimes for both static and dynamic expander decompositions in directed graphs. Our result is summarized in the theorem below.

► **Theorem 2.** *Given a parameter  $\phi \leq c/\log^{12} m$  for a fixed constant  $c > 0$ , and a directed  $m$ -edge graph  $G$  undergoing a sequence of edge deletions, there is a randomized data structure that constructs and maintains a  $(O(\log^{19} m), \Omega(\phi/\log^{12} m), O(1/\log^8 m))$ -expander decomposition  $^2 (\mathcal{X}, E^r)$  of  $G$ . The initialization of the data structure takes time  $O(m \log^{20}(m)/\phi)$  and the amortized time to process each edge deletion is  $O(\log^{28}(m)/\phi^2)$ .*

Further, our algorithm has the property that it is refining for up to  $\tilde{O}(\phi \cdot m)$  edge deletions meaning that  $\mathcal{X}$  is a refinement of its earlier versions (every expander component in  $\mathcal{X}$  is a subset of an expander component in any earlier expander decomposition) and the size of  $E^r$

<sup>2</sup> Here we use the augmented expander decomposition definition 7.

does never exceed  $\tilde{O}(\phi m)$ . Our algorithms are deterministic, however, they rely on calling a fast randomized algorithm to find balanced sparse cuts or certify that no such cut exists (see [22, 27]). Our new techniques are much simpler and more accessible than previous work, besides also being much faster. We hope that by giving a simpler algorithm for directed expander decompositions, we can help to make this tool more accessible to other researchers in the field with the hope that this can further accelerate recent advances in dynamic and static graph algorithms.

## 1.2 Applications

Our new algorithms have direct applications to the currently fastest approaches for bipartite matching/ maximum flow/ min-cost flows and the data structures that are employed to obtain these results:

1. Our algorithm is already used in the fastest min-cost flow algorithm that is known to-date [33] which achieves runtime  $m \cdot e^{O(\log^{3/4}(m) \log \log(m))}$  yielding the first improvement over the recent breakthrough in [8] achieving the first near-linear time algorithm for min-cost flows. In the framework of [33], our algorithmic techniques are used to maintain an expander decomposition of an *undirected* graph where it is heavily exploited that our algorithm maintains the expander decomposition such that it refines over time. This guarantee is pivotal in the construction of a fully-dynamic algorithm to maintain dynamic expander hierarchies which is the key data structure in the paper. While [33] could also have relied solely on the techniques [18] to obtain such a data structure with subpolynomial update and query times, these subpolynomial factors would have been substantially larger and would thus not have yielded a faster min-cost flow algorithm overall.
2. In [2], directed expander decompositions for decremental graphs were used to obtain the first algorithm to maintain  $(1 + \epsilon)$ -approximate Single-Source Shortest-Paths (SSSPs) in a decremental graphs in time  $o(mn)$ , though only for the special case of dense graphs. This problem is well-motivated as a simple reduction based on the Multiplicative Weights Update (MWU) framework implies that the maximum flow problem can be solved approximately by solving approximate decremental SSSP. By standard refinement of flows this yields an exact maximum flow algorithm. Very recently, Chuzhoy and Khanna [11, 12] showed that for the update sequence generated by the MWU to solve the bipartite matching problem - a special case of the maximum flow problem - decremental SSSP can be maintained in  $n^{2+o(1)}$  time by refining the techniques from [4, 2]. This yields the first near-optimal 'combinatorial'<sup>3</sup> algorithm for the bipartite matching problem for very dense graphs. While the above algorithms are 'combinatorial', they are still very intricate. Most of these complications stem from the maintenance of the directed expander decomposition used internally by the decremental SSSP data structure. We hope that our technique can help to simplify and speed-up these components to yield a simpler algorithm overall.
3. In independent work [1], an alternative 'combinatorial' maximum flow that runs in  $n^{2+o(1)}$  time was given. This algorithm cleverly extends push-relabel algorithms to run more efficiently when given an ordering of vertices that roughly aligns with the topological order induced by the acyclic graph formed from the support of an optimal maximum flow solution. To obtain this approximate ordering they compute a static expander hierarchy of the directed input graph. This generalizes the notion of directed expander

---

<sup>3</sup> We refrain from defining the scope of combinatorial algorithms here and refer the reader to [11, 12] for a discussion.

decompositions further. In their work, they heavily build on the techniques from [18] to obtain the expander hierarchy. Unfortunately, the algorithm to obtain this hierarchy is very involved. We hope that our new techniques can help to simplify and speed-up their algorithms.

### 1.3 Our Techniques

#### High-Level Strategy

We obtain our result by following the high-level strategy of [31] for undirected graphs: we draw on existing literature (specifically [22, 27]) for an algorithm that either outputs a balanced sparse cut which allows us to recurse on both sides; or outputs a witness that no such cut exists. This witness can be represented as an expander graph  $W$  that embeds into  $G \cup F$  with low congestion where  $F$  is a set of few *fake* edges. In the second case, we set up a flow problem to extract a large expander (the first algorithm only finds balanced sparse cuts, so many unbalanced sparse cuts might remain) which suffices to again recurse efficiently.

#### The (Dynamic) Flow Problem in [31]


To outline our algorithm, we first sketch the techniques of [31]. In [31], the following sequence of flow problems is formulated: initially, we add  $\frac{1}{2\phi}$  units of source commodity to each endpoint of an edge in  $F$  and then ask to route the commodity in  $G$  where each vertex  $v$  is a sink of value  $\deg_G(v)$  and each edge has capacity  $\frac{1}{2\phi}$ . It then runs an (approximate) maximum flow algorithm on the flow problem. Whenever the algorithm detects that no feasible flow exists<sup>4</sup>, it finds a cut  $(A, \bar{A})$  where  $A$  is the smaller side of the cut and then poses the same problem for the network  $G[\bar{A}]$  where this time the source commodity is assigned for each edge in  $E_G(A, \bar{A}) \cup F$ . The algorithm terminates once the flow problem can be solved and outputs the final induced graph. In [31], it is shown that once a feasible flow exists then the (induced) graph is a  $\Omega(\phi)$ -expander. Further, it is shown that in the sequence of flow problems, each problem can be warm-started by re-using the flow computed in the previous instance to detect a cut induced on the remaining vertex set. This result is obtained by two main insights:

1. if the flow  $f$  to find the cut  $(A, \bar{A})$  is a *pre-flow*, that is a flow that respects capacities and has no negative excess at any vertex (i.e. it does not route away more flow from a vertex than is inputted by the source), then injecting additional source flow for any edge  $E_G(A, \bar{A})$  guarantees that the induced flow  $f|_{\bar{A}}$  is a pre-flow in the flow problem formulated for  $G[\bar{A}]$ . That is because the amount of flow that was routed via such a cut edge is at most  $2/\phi$  and thus placing  $2/\phi$  new source commodity at the endpoint ensures that no negative excess exists in the induced flow  $f|_{\bar{A}}$ ,
2. the classic push-relabel framework can naturally be extended to warm-start on such a flow  $f|_{\bar{A}}$  as it is built to just further refine pre-flows at every step.

This dynamization of the push-relabel framework allows to bound the cost of computation of *all* flow problems linearly in the amount of source commodity which in term is bounded by the number of edges that appear in either  $F$  or one of the identified min-cuts. Finally, [31] shows that the amount of source commodity remaining in  $\bar{A}$  decreases over the sequence of flow problems proportional to the volume of the set  $A$  of vertices that are removed at each

---

<sup>4</sup> Technically, the algorithm might already output cuts when some cut has capacity less than a constant times the amount of flow that is required to be routed through the cut.



(a) In directed graphs, cuts are asymmetric. While  $(A, \bar{A})$  might be a sparse cut, the cut  $(\bar{A}, A)$  might contain many edges. A straightforward extension of [31] would inject  $2/\phi$  units of commodity to each endpoint of  $E_G(A, \bar{A}) \cup F$ . However, it is not clear with this approach how to bound the total amount of flow injected throughout the algorithm.

(b) We inject  $\Theta(1/\psi)$  units of commodity at the end points of any witness embedding path (green) going through an edge of  $E_G(A, \bar{A}) \cup F$ . We can bound the total amount by  $O(\text{vol}(A)/\text{poly}(\psi))$ . But the injection might well be in the interior of  $A, \bar{A}$  possibly leaving negative excess at the endpoints of the cut edges.

■ **Figure 1** Injection of Commodity due to edges in  $E_G(A, \bar{A}) \cup F$ .

step. Indeed, they observe that at each vertex  $v$  in  $A$  at least  $\deg(v)$  many commodity units are absorbed and that the total amount of source injected due to the cut edges  $E_G(A, \bar{A})$  is bounded by  $\frac{1}{2\phi} \cdot e_G(A, \bar{A}) \leq \text{vol}_G(A)/2$ . This yields that the final induced graph is still large. Thus, the final graph outputted is a large expander, as desired.

### The (Dynamic) Flow Problem in Directed Graphs

In directed graphs, while the above flow problem upon becoming feasible also certifies that the remaining graph is a  $\Omega(\phi)$ -expander, the argument that the sequence of flow problems terminates does not work: the asymmetry of cuts might force us for a small cut  $E_G(A, \bar{A})$  to induce on  $\bar{A}$  while having many edges in  $E_G(\bar{A}, A)$  each of which would add  $2/\phi$  source flow to the new flow problem (see Figure 1a). Hence, we might end up injecting up to  $\Omega(\text{vol}_G(A)/\phi)$  more flow due to the cut edges. This makes it seemingly impossible to argue that the amount of source commodity in the next flow problem is smaller. To recover the argument that the sequence of flow problems terminates (with the remaining expander graph being large) both [2, 18] suggest setting up the flow problems more carefully such that each cut  $(A, \bar{A})$  that is found in this sequence and induced upon is a *sparse cut*. Here, we only describe the less lossy flow problem formulation developed in [18]. To ensure that each cut  $(A, \bar{A})$  that is found is a sparse cut, [18] proposes a slightly different flow problem: instead of adding source commodity  $\frac{1}{2\phi}$  per endpoint of an edge that is fake or not fully contained in the induced graph, it tailors the amount of new source commodity using the witness graph  $W$ , possibly injecting much less source commodity in the process. Concretely, we have that  $W$  is a  $\psi$ -expander over the same vertex set as  $G$  with degrees similar to degrees in  $G$  up to a factor of  $\Theta(\text{poly}(\psi))$  and an embedding  $\Pi$  into  $G \cup F$  with congestion  $O(\phi/\psi)$ , for  $\phi = \tilde{\Theta}(1)$ . To set-up the flow problem, we inject  $\Theta(1/\psi)$  units at the endpoints of any edge  $e$  in the witness  $W$  whose embedding path  $\Pi(e)$  goes through an edge in  $E_G(A, \bar{A}) \cup F$ . We note that any such witness embedding path  $\Pi(e)$  either crosses the sparse cut  $(A, \bar{A})$  or has an endpoint in  $A$ . This allows to bound the additional source injected by  $O(\text{vol}_G(A)/\text{poly}(\psi)) = \tilde{O}(\text{vol}_G(A))$  where we use that  $\psi = \tilde{\Theta}(1)$ . But while correctness and termination of the flow problem sequence are now ensured, this leaves a significant problem: the current flow  $f$  that was used to find the cut  $(A, \bar{A})$  no longer has the property that  $f|_{\bar{A}}$  is a *pre-flow* in the flow problem

formulated on network  $G[\bar{A}]$  even if  $\mathbf{f}$  is a pre-flow. While capacity constraints are still enforced, i.e.  $\mathbf{f}|_{\bar{A}}$  still is a pseudo-flow (see [17] for reference on pseudo-flow), some vertices might now have negative excess since the additional commodity might be injected in the interior of  $\bar{A}$  far from the cut  $(A, \bar{A})$  (see Figure 1b). Indeed, for an edge  $(u, v)$  in  $E_G(A, \bar{A})$  a lot of flow might have been routed through  $(u, v)$  but no additional commodity might be injected at  $v$  since all witness embedding paths passing through  $(u, v)$  might not end in  $v$ . Thus, dynamizing the push-relabel framework does not appear natural for this sequence of problems as it crucially requires that the maintained flow is a pre-flow at all times. In [18], an involved batching technique is used instead (based on the technique in [29]) that does not use dynamic flow problems but instead reduces to few static flow problems, however, at the loss of quality and runtime by subpolynomial factors.

### The Push-Pull-Relabel Framework

The main technical contribution of this paper is a new framework that refines pseudo-flows as efficiently as the push-relabel framework refines pre-flows. Thus, we give a generalization of the latter widely-used and well-studied framework that we believe might have applications well beyond our problem. Recall that the classic push-relabel framework maintains labels  $\ell$  for all vertices and a pre-flow  $\mathbf{f}$ . In each iteration, it ‘pushes’ positive excess flow at a vertex  $v$  to a vertex at a lower label (to be precise to a vertex at level  $\ell(v) - 1$ ); or if no ‘push’ is possible, it increases the labels of some vertices: it ‘relabels’. Using a clever potential-based analysis, one can show that it suffices to only increase the labels to a certain threshold before all flow is settled. In our framework, we allow deleting edges, without compensating by adding source commodity at the endpoints, which might create negative excess leaving  $\mathbf{f}$  a pseudo-flow (instead of a pre-flow). Now, while our framework applies the same strategy for ‘pushes’ and ‘relabels’, we also need a new operation ‘pull’. Intuitively, our algorithm tries to ‘pull’ back the source commodity that now causes the negative excess (this unit of commodity was ‘pushed’ earlier to some other vertices). To do so, a vertex  $v$  with negative excess can ‘pull’ commodity from vertices at a higher level (again it can only pull from a vertex at level  $\ell(v) + 1$ ). But it is not difficult to construct an example where this strategy does not suffice: therefore, we also need to sometimes decrease the label of a vertex to ensure correctness. However, the latter change to the ‘relabel’ operation breaks the property that labels are non-decreasing over time. A property that is crucial in the existing efficiency analysis. Instead, we give a much more careful argument to analyze the potentials that bound the number of push, pull, and relabeling operations that deal with the non-monotonicity of the levels over time. The argument is sketched below. Combining this framework with the above-discussed set-up of dynamic flow problems as proposed in [18] then yields the first near-optimal algorithm to compute an expander decomposition in a directed graph. Further, our technique extends seamlessly to also deal with edge deletions to  $G$ , yielding an algorithm to prune expander graphs that undergo edge deletions.

### A Sketch of the Runtime of Push-Pull-Relabel

The run-time analysis in standard push-relabel considers the run-time contribution of the push and relabel operations separately. Using a potential argument, one can then relate the contribution of the push operations to the relabeling operations. A similar argument albeit much more delicate also allows to relate the contribution of the push and the pull operations to the contribution of the relabeling operations in the extended push-pull-relabel algorithm. What might seem more daunting is to bound the run-time contribution of the relabelings

given that the level function  $\ell$  is no longer point-wise non-decreasing. Let us revisit the argument to bound the run-time contribution of the relabelings in the push-relabel of [31]. Any relabeling of  $v$ , incurs a cost of  $O(\deg(v))$  as each incident edge has to be checked before a relabeling. At such a relabeling of  $v$ , the sink of  $v$ , which has by choice capacity  $\deg_G(v)$ , must be full and any commodity unit in it remains there till termination. Since the label of  $v$  does not decrease and is bounded by  $h$ , we may thus charge for the run-time contribution of the relabelings of  $v$  each unit in the sink of  $v$  exactly  $h$ . Over all vertices, we conclude that any commodity unit was charged at most  $h$  units. In the push-pull-relabel algorithm, a commodity unit might end up in various sinks. So a more clever charging argument needs to be devised. To guide intuition an analogy to flows of protons and electrons is useful. The protons correspond to commodity units and the electrons to the lack of commodity units. We are now moving protons and electrons around in the network and whenever a proton and an electron meet at a vertex they form a neutron which stays put at the vertex indefinitely. The neutrons will provide a means to charge work in the analysis. The sink of each vertex  $v$  can absorb  $\deg(v)$  commodity units or put differently  $\deg(v)$  protons. Electrons only form when we delete an edge  $(u, v)$ , i.e. exactly  $f^+(u, v)$  new electrons form at  $v$ . By choice, we now say that if we perform a push from  $u$  to  $v$  of one unit then exactly one proton moves from  $u$  to  $v$ , while if we perform a pull from  $u$  to  $v$  then exactly one electron moves from  $v$  to  $u$ . Since we only perform a push away from  $v$  if there are more than  $\deg_G(v)$  free protons (not part of a neutron) at  $v$  and a pull towards  $v$  if there are free electrons at  $v$ , we find that  $\min(\mathbf{n}(v) + \deg_G(v)/2, \mathbf{p}(v))$ , where  $\mathbf{p}(v)$  denotes the number of protons at  $v$  (including the ones in a neutron) and  $\mathbf{n}(v)$  denotes the number of electrons at  $v$ , is non-decreasing. This fact allows us to conclude that whenever we change from increasing  $\ell(v)$  to decreasing  $\ell(v)$  at least  $\deg(v)/2$  new neutrons have formed at  $v$ . Since any neutron at  $v$  stays put indefinitely, we can charge the run-time contributions of the relabelings of  $v$  to the number of neutrons at  $v$ .

## Roadmap

In the remainder of the article, we first give preliminaries in Section 2, then present our new push-pull-relabel framework in Section 3 and finally show how to obtain our result in Theorem 2 using the new framework in Section 4.

## 2 Preliminaries

### Graphs

We let  $\deg_G \in \mathbb{R}^{V(G)}$  denote the degree vector of graph  $G$ . For all vertex  $v \in V$ , we have  $\deg_G(v)$  equal to the number of edges incident to  $v$  (both incoming and outgoing are counted). Moreover, for any subset of edges  $D \subseteq E(G)$  we denote by  $G_D$  the graph with vertex set  $V(G)$  and edge set  $D$  and by  $\deg_D$  the degree vector of the subgraph  $G_D$ . We denote by  $\text{vol}_G(S)$  for any  $S \subseteq V$ , the sum of degrees of vertices in  $S$ . We denote by  $E_G(A, B)$  for any  $A, B \subseteq V$  the set of directed edges in  $E(G)$  with tail in  $A$  and head in  $B$ . We define  $e(G) = |E(G)|$  and  $e_G(A, B) = |E_G(A, B)|$ . For any partition  $\mathcal{P}$  of  $V(G)$ , we denote the graph obtained by contracting each partition class to a single vertex by  $G/\mathcal{P}$ . Two vertices in  $G/\mathcal{P}$  are adjacent if there is an edge between the corresponding partition classes in  $G$ . Moreover, for any vertex  $v \in V(G)$  we denote by  $\mathbf{1}_v \in \mathbb{R}^{V(G)}$  the vector with all entries equal to zero apart from the entry at  $v$  equaling one.



### Flows

We call a tuple  $(G, \mathbf{c}, \Delta, \nabla)$  a flow problem, if  $G$  is a directed graph, the capacity function  $\mathbf{c} : V(G) \times V(G) \rightarrow \mathbb{R}^{\geq 0}$  is such that for all  $(u, v) \notin E$  we have  $\mathbf{c}(u, v) = 0$ , and  $\Delta, \nabla : V(G) \rightarrow \mathbb{R}^{\geq 0}$  denote the source and the sink capacities. We denote flows on  $G$  as functions  $\mathbf{f} : V(G) \times V(G) \rightarrow \mathbb{R}$  such that  $\mathbf{f}$  is anti-symmetric, i.e.  $\mathbf{f}(u, v) = -\mathbf{f}(v, u)$ . Given a vertex  $v \in V(G)$  we introduce the notation  $\mathbf{f}(v) = \sum_u \mathbf{f}(v, u)$  and likewise  $\mathbf{c}(v) = \sum_u \mathbf{c}(v, u) + \mathbf{c}(u, v)$ . Moreover, we write  $\mathbf{f}^+(u, v) = \max(\mathbf{f}(u, v), 0)$ . Given a flow  $\mathbf{f}$ , we say a vertex  $v \in V(G)$  has  $\Gamma(v) = \Delta(v) - \mathbf{f}(v)$  excess. We say that it has positive excess if  $\Gamma(v) > \nabla(v)$  and  $v$  has negative excess if  $\Gamma(v) < \nabla(v)/2$ . For a subset  $\tilde{V} \subseteq V(G)$ , we induce the flow  $\mathbf{f}$ , and the sink  $\nabla$ , source  $\Delta$  and edge capacities  $\mathbf{c}$  onto  $G[\tilde{V}]$  in a function sense and write  $\mathbf{f}|_{\tilde{V}}, \nabla|_{\tilde{V}}, \Delta|_{\tilde{V}}, \mathbf{c}|_{\tilde{V}}$ . Moreover, for any subset  $E \subseteq E(G)$ , we write  $\mathbf{f}|_E$  for the flow induced by  $\mathbf{f}$  onto the subgraph of  $G$  consisting only of the edges  $E$ . We say that a flow  $\mathbf{f}$  is a *pseudo-flow* if it satisfies the capacity constraints:

$$\forall (u, v) \in V(G) \times V(G) : -\mathbf{c}(v, u) \leq \mathbf{f}(u, v) \leq \mathbf{c}(u, v).$$

We say  $\mathbf{f}$  is a *pre-flow* if  $\mathbf{f}$  is a pseudo-flow and has no negative excess at any vertex. We say a flow  $\mathbf{f}$  is *feasible* if it is a pre-flow and additionally no vertex has positive excess. Moreover, for any subset  $S \subseteq V$  we denote by  $\Delta(S), \nabla(S)$  the sum  $\sum_{v \in S} \Delta(v), \sum_{v \in S} \nabla(v)$  respectively and for any subset  $D \subseteq E(G)$  we denote by  $\mathbf{c}(D)$  the sum  $\sum_{d \in D} \mathbf{c}(d)$ .

### Expanders

Given graph  $G = (V, E)$ , we say a cut  $(S, \bar{S})$  is  $\phi$ -out sparse if  $\text{vol}_G(S) \leq e(G)$  and  $e_G(S, V \setminus S) < \phi \cdot \text{vol}_G(S)$ . We say  $G$  is a  $\phi$ -out expander if it has no  $\phi$ -out-sparse cut. We say  $G$  is a  $\phi$ -expander if  $G$  and  $G^{rev}$ , the graph where all edges of  $G$  are reversed, are both  $\phi$ -out expander. The next lemma that is folklore and crucial in our expander pruning argument.

► **Lemma 3.** *Given a  $\phi$ -expander  $G = (V, E)$ , then take  $S \subseteq V$  and a set of edge deletions  $D$ . We have that  $e_{G \setminus D}(S, V \setminus S) < \frac{\phi}{4} \cdot \text{vol}_G(S)$  implies  $\min(\text{vol}_G(S), \text{vol}_G(V \setminus S)) < \frac{4 \cdot |D|}{3\phi}$ .*

**Proof.** If  $\min(\text{vol}_G(S), \text{vol}_G(V \setminus S)) \geq \frac{4 \cdot |D|}{3\phi}$ , then  $e_{G \setminus D}(S, V \setminus S) \geq \phi \cdot \min(\text{vol}_G(S), \text{vol}_G(V \setminus S)) - |D| \geq \frac{\phi}{4} \min(\text{vol}_G(S), \text{vol}_G(V \setminus S))$ . ◀

### Graph Embeddings

Given two graphs  $H$  and  $G$  over the same vertex set  $V$ , we say that  $\Pi$  is an embedding of  $H$  into  $G$  if for every edge  $e = (u, v) \in E(H)$ ,  $\Pi(e)$  is a simple  $uv$ -path in  $G$ . We define the congestion of an edge  $e \in E(G)$  induced by the embedding  $\Pi$  to be the maximum number of paths in the image of  $\Pi$  that contain  $e$ . We define the congestion of  $\Pi$  to be the maximum congestion achieved by any such edge  $e \in E(G)$ . Moreover, for any edge  $e \in E(G)$  we will denote by  $\Pi^{-1}(e)$  the set of edges  $f \in E(H)$  such that  $e$  is an edge on the path  $\Pi(f)$ . Given an entire set of edges  $D \subseteq E(G)$ , we denote by  $\Pi^{-1}(D)$  the set of edges  $f \in E(H)$  such that some edge of the path  $\Pi(f)$  is in  $D$ . Given two graphs  $H_1, H_2$  on the same vertex set and embeddings  $\Pi_1 : H_1 \rightarrow G, \Pi_2 : H_2 \rightarrow G$  then we denote by  $\Pi_1 \cup \Pi_2 : H_1 \cup H_2 \rightarrow G$  the embedding of the graph  $H_1 \cup H_2 = (V(H_1), E(H_1) \cup E(H_2))$ . We denote by  $G^{rev}$  the graph obtained from  $G$  by reversing all edges. Given an embedding  $\Pi : H \rightarrow G$ , we denote by  $\Pi^{rev} : H^{rev} \rightarrow G^{rev}$  the embedding obtained by reversing the direction of all embedding paths.



### Expander Decompositions with Witnesses

For the rest of the article, we use a definition of expander decompositions that encodes much more structure than given in Definition 1. In particular, the definition below incorporates the use of witness graphs which are instrumental to our algorithm.

► **Definition 4.** *Given a directed graph  $G$ , we say  $(W, \Pi)$  is a  $(\phi, \psi)$ -out-witness for  $G$  if 1)  $W$  is a  $\psi$ -out-expander, and 2)  $\Pi$  embeds  $W$  into  $G$  with congestion at most  $\frac{\psi}{\phi}$ , and 3)  $\forall v \in V(W) : \deg_G(v) \leq \deg_W(v) \leq \frac{\deg_G(v)}{\psi}$ . If  $(W, \Pi)$  is a  $(\phi, \psi)$ -out-witness for  $G$  and  $(W^{rev}, \Pi^{rev})$  for  $G^{rev}$ , then we say that  $(W, \Pi)$  is a  $(\phi, \psi)$ -witness for  $G$ .*

► **Fact 5.** *If  $(W_1, \Pi_1)$  is a  $(\phi, \psi)$ -out-witness for  $G$  and  $(W_2^{rev}, \Pi_2^{rev})$  is a  $(\phi, \psi)$ -out-witness for  $G^{rev}$ , then  $(W_1 \cup W_2^{rev}, \Pi_1 \cup \Pi_2^{rev})$  is a  $(\psi\phi/4, \psi^2/2)$ -witness for  $G$ .*

**Proof.** We observe that for any cut  $(S, V \setminus S)$ :

$$\begin{aligned} e_{W_1 \cup W_2^{rev}}(S, V \setminus S) &\geq \psi \cdot \min(\text{vol}_{W_1}(S), \text{vol}_{W_2^{rev}}(S), \text{vol}_{W_1}(V \setminus S), \text{vol}_{W_2^{rev}}(V \setminus S)) \\ &\geq \frac{\psi^2}{2} \min(\text{vol}_{W_1 \cup W_2^{rev}}(S), \text{vol}_{W_1 \cup W_2^{rev}}(V \setminus S)). \end{aligned}$$

Item 2, 3 are immediate. ◀

The next fact establishes that a  $(\phi, \psi)$ -(out-)witness for  $G$  certifies that  $G$  is a  $\phi$ -(out-)expander, justifying the name witness.

► **Fact 6** (see [18], Claim 2.1). *If  $(W, \Pi)$  is a  $(\phi, \psi)$ -witness for  $G$ , then  $G$  is a  $\phi$ -expander.*

**Proof.** For any set  $S \subseteq V(G)$ , we have in the witness that

$$e_W(S, V \setminus S) \geq \psi \cdot \min(\text{vol}_W(S), \text{vol}_W(V \setminus S)) \geq \psi \cdot \min(\text{vol}_G(S), \text{vol}_G(V \setminus S)).$$

Since for any one of these cut edges there is an embedding path crossing the cut and since the congestion of the embedding is at most  $\frac{\psi}{\phi}$ , we have  $e_G(S, V \setminus S) \geq \frac{\phi}{\psi} \cdot e_W(S, V \setminus S) \geq \phi \cdot \min(\text{vol}_G(S), \text{vol}_G(V \setminus S))$ . ◀

► **Definition 7 (Augmented Expander Decomposition).** *We call a collection  $\mathcal{X}$  and a subset  $E^r \subseteq E$  a  $(\beta, \phi, \psi)$ -expander decomposition of a graph  $G$ , if*

1.  $\forall (X, W, \Pi) \in \mathcal{X}$ ,  $G[X]$  has a  $(\phi, \psi)$ -witness  $(W, \Pi)$ ,
2.  $|E^r| \leq \beta \cdot \phi \cdot e(G)$ ,
3.  $(G \setminus E^r)/\mathcal{P}$  is a DAG, where  $\mathcal{P} = \{X \mid (X, W, \Pi) \in \mathcal{X}\}$ .

*Given two expander decompositions  $(\mathcal{X}_1, E_1^r)$  of the graph  $G$  and  $(\mathcal{X}_2, E_2^r)$  of the graph  $G \setminus \mathcal{D}$ , where  $\mathcal{D} \subseteq E(G)$ , we say  $(\mathcal{X}_2, E_2^r)$  refines  $(\mathcal{X}_1, E_1^r)$  if 1) for all partition classes  $X_2$ , where  $(X_2, W_2, \Pi_2) \in \mathcal{X}_2$ , there is a class  $X_1$ , where  $(X_1, W_1, \Pi_1) \in \mathcal{X}_1$ , such that  $X_2 \subseteq X_1$  and 2)  $E_1^r \subseteq E_2^r \cup \mathcal{D}$ .*

## 3 The Push-Pull-Relabel Framework

In the push-relabel framework as presented in [14], we are trying to compute a feasible flow for a flow problem  $(G, c, \Delta, \nabla)$  by maintaining a pre-flow  $\mathbf{f}$  together with a level function  $\ell$ . The algorithm then runs in iterations terminating once  $\mathbf{f}$  has no positive excess at any vertex. In each iteration of the algorithm, the algorithm identifies a vertex  $v$  that still has positive excess at a vertex  $v$ . This positive excess is then pushed to neighbors on lower levels such that the capacity constraint is still enforced. If this is not possible  $v$  is relabeled

meaning that its label  $\ell(v)$  is increased. In this section, we are extending the push-relabel framework to the dynamic setting, where we allow for increasing the source function  $\Delta$  and the deletion of edges from  $G$ . To do so, we need to introduce the notion of negative excess. Because for a flow  $\mathbf{f}$  of the flow problem, it might happen that once we delete an edge  $(u, v)$  there is more flow leaving the vertex  $v$  than is entering or sourced, i.e.  $\Gamma_t(v) < 0$ . Hence, we need to extend the discussion to include negative excess in the dynamic version. Similar to the standard push-relabel algorithm, we maintain a pseudo-flow  $\mathbf{f}$  and a vertex labeling  $\ell$  in so-called valid states  $(\mathbf{f}, \ell)$ . The only difference is that in the standard push-relabel algorithm,  $\mathbf{f}$  is a pre-flow (not only a pseudo-flow).

► **Definition 8.** *Given a level function  $\ell : V(G) \rightarrow [h]$  and a pseudo-flow  $\mathbf{f}$ , we call a tuple  $(\mathbf{f}, \ell)$  a state for  $(G, \mathbf{c}, \Delta, \nabla)$  if for all edges  $e = (u, v)$  having  $\ell(u) > \ell(v) + 1$  implies  $\mathbf{f}(e) = \mathbf{c}(e)$ .*

*We call the state  $(\mathbf{f}, \ell)$  valid if for all vertices  $v \in V(G)$  : 1)  $\Gamma(v) < \nabla(v)/2$  implies  $\ell(v) = 0$ , 2)  $\nabla(v) < \Gamma(v)$  implies  $\ell(v) = h$ .*

For the remainder of the paper, we have  $h = O\left(\frac{\log(n)}{\phi}\right)$ . To provide the reader with some intuition about the definitions, we remark that the pseudo-flow of a valid state is not feasible in the usual sense. Indeed, some vertices might have positive or negative excess. But these vertices are guaranteed to either be at level  $h$  or at level 0. Moreover, we point out that if a valid state  $(\mathbf{f}, \ell)$  has no vertices at level  $h$ , then  $\mathbf{f}$  might still not be a feasible flow since there might be a vertex  $v$  with  $\Gamma_t(v) < 0$ . But it is straightforward to obtain from  $\mathbf{f}$  a feasible flow: extract from  $\mathbf{f}$  at each vertex  $v$  exactly  $\Delta(v)$  unit flow paths (possibly empty paths starting and ending in  $v$ ). Let us briefly outline how we will use Lemma 9 in the directed expander pruning algorithm 3. In the directed expander pruning algorithm, we start with an out-expander and an adversary performs a number of vertex or edge deletions. We aim to find a small pruning set of vertices such that if we prune away this small pruning set from the remaining graph, then we can certify that the obtained graph is still an expander. The classical way (see [31]) to either certify expansion of the remaining graph or to find a pruning set is by setting up a flow problem, where we inject for any deletion some additional source flow and where each vertex has sink capacity equal to its degree (this is the reason for the condition  $\nabla \geq \deg$  in Lemma 9). We will solve this flow problem using the flow provided by our VALIDSTATE algorithm (see Algorithm 1). We maintain this flow problem using the interface functions INCREASESOURCE, REMOVEEDGES under both adversarial deletions and under necessary pruning deletions. If the valid state computed after any such update has an additional vertex on level  $h$ , this will indicate that more vertices need to be pruned away. If on the other hand all vertices are on levels strictly below  $h$ , we will certify that the remaining graph is an expander.

► **Lemma 9.** *Given a flow problem  $(G = (V, E), \mathbf{c}, \Delta, \nabla)$ , where  $\nabla \geq \deg$ . Then, there is a deterministic data structure VALIDSTATE( $G, \mathbf{c}, \Delta, \nabla$ ) (see Algorithm 1) that initially computes a valid state  $(\mathbf{f}, \ell)$  and after every update of the form*

■ INCREASESOURCE( $\delta$ ) : *where  $\delta \in \mathbb{N}_{\geq 0}^n$ , we set  $\Delta$  to  $\Delta + \delta$ ,*

■ REMOVEEDGES( $D$ ) : *where  $D \subseteq E(V) \setminus \mathcal{D}$ , sets  $\mathcal{D}$  to  $\mathcal{D} \cup D$  (initially  $\mathcal{D} = \emptyset$ ),*

*the algorithm explicitly updates the tuple  $(\mathbf{f}, \ell)$  such that thereafter  $(\mathbf{f}, \ell)$  is a valid state for the current flow instance  $(G \setminus \mathcal{D}, \mathbf{c}|_{E \setminus \mathcal{D}}, \Delta|_{E \setminus \mathcal{D}}, \nabla|_{E \setminus \mathcal{D}})$ . The run-time is  $O(h \cdot (\|\Delta\|_1 + \sum_{d \in \mathcal{D}} (1 + |\mathbf{f}_{t_d}(d)|)))$ , where  $\mathbf{f}_{t_d}(d)$  is equal to  $\mathbf{f}(d)$  at the time  $d$  is deleted and  $\Delta$  is the variable at the end of the algorithm.*

■ **Algorithm 1**  $\text{VALIDSTATE}(G = (V, E), \mathbf{c}, \Delta, \nabla, h)$ .

---

```

def Init():
     $\tilde{E} \leftarrow E(G), (\mathbf{c}, \Delta, \nabla) \leftarrow (\mathbf{c}, \Delta, \nabla), (\mathbf{f}, \ell) \leftarrow (\mathbf{0}, \mathbf{0})$ 
    PushRelabel()
def IncreaseSource( $\delta$ ):
     $\Delta \leftarrow \Delta + \delta$ 
    PushRelabel()
def RemoveEdges( $D$ ):
     $\tilde{E} \leftarrow \tilde{E} \setminus D, (\mathbf{f}, \ell) \leftarrow (\mathbf{f}|_{\tilde{E}}, \ell|_{\tilde{V}})$ 
    PullRelabel()
    PushRelabel()
def PushRelabel():
    while  $\exists v$  where  $\ell(v) < h$  and  $\Gamma(v) > \nabla(v)$  do
        Let  $v$  be a vertex minimizing  $\ell(v)$ 
        if  $\exists (v, u)$  such that  $\mathbf{c}_f(v, u) > 0, \ell(v) = \ell(u) + 1$  then
             $\mathbf{f}(v, u) \leftarrow \mathbf{f}(v, u) + \frac{1}{2}, \mathbf{f}(u, v) \leftarrow -\mathbf{f}(v, u)$  // Sends  $\frac{1}{2}$  units of pos. excess
            from  $v$  to  $u$ 
        else
             $\ell(v) \leftarrow \ell(v) + 1$ 
def PullRelabel():
    while  $\exists v$  where  $\ell(v) > 0$  and  $\Gamma(v) < \nabla(v)/2$  do
        Let  $v$  be a vertex maximizing  $\ell(v)$ 
        if  $\exists (u, v)$  such that  $\mathbf{c}_f(u, v) > 0, \ell(u) = \ell(v) + 1$  then
             $\mathbf{f}(v, u) \leftarrow \mathbf{f}(v, u) - \frac{1}{2}, \mathbf{f}(u, v) \leftarrow -\mathbf{f}(v, u)$  // Sends  $\frac{1}{2}$  units of neg. excess
            from  $v$  to  $u$ 
        else
             $\ell(v) \leftarrow \ell(v) - 1$ 

```

---

The user interface of the data structure `VALIDSTATE` are the functions `INIT`, `INCREASESOURCE` and `REMOVEEDGES`. These functions can be used to initialize or update the flow problem. After any such update the internal state  $(\mathbf{f}, \ell)$  has to be updated to remain valid for the new flow problem. To facilitate these necessary updates to  $(\mathbf{f}, \ell)$  these functions make calls to the internal functions `PUSHRELABEL` and `PULLRELABEL`. As described before, the function `PUSHRELABEL` performs push and relabel operations to handle any positive excess, while the function `PULLRELABEL` performs pull and relabel operations to handle any negative excess. Let us take a closer look at the individual functions: using `INIT` we set up the flow problem. It makes an internal call to `PUSHRELABEL` to compute the first valid state using the standard push-relabel algorithm. Note that at this point no negative excess has been introduced and thus we do not need to resort to any pull operations. The function `INCREASESOURCE` can be used to increase the source of the current flow problem and again this does not introduce any negative excess and hence we can again update the valid state using the standard push-relabel algorithm. Using the function `REMOVEEDGES`, we can restrict the flow problem to a subgraph. The function induces the capacities  $\mathbf{c}, \Delta, \nabla$  and the flow  $\mathbf{f}$  to the subgraph and then makes calls to both `PUSHRELABEL` and `PULLRELABEL` to update the state. The function `PUSHRELABEL` is just the standard push-relabel algorithm. We look for a vertex  $v$  with positive excess  $\Gamma(v) > \nabla(v)$  at level below  $h$ . We pick a vertex on the lowest level possible. We try to push flow along some unsaturated outgoing edge

## 132:12 Near-Optimal Algorithm for Directed Expander Decompositions

to a vertex on a lower level. If it is not possible to push flow, we increase the label of the vertex and repeat. The function `PULLRELABEL` is the analog of the push-relabel algorithm for negative excess. Here, we look for a vertex  $v$  with negative excess  $\Gamma(v) < \nabla(v)/2$  at level above 0. We pick a vertex on the highest level possible. We try to pull flow along some unsaturated incoming edge from a vertex on a higher level. If it is not possible to pull flow then all edges  $(u, v)$  where  $\ell(u) > \ell(v)$  must be saturated and it is safe to decrease the label of the vertex and repeat.

**Proof of Lemma 9.** Since we only decrease the node label  $\ell(v)$  when all edges  $(u, v)$ , where  $\ell(u) > \ell(v)$ , are saturated and since we only increase the node label when all edges  $(v, w)$ , where  $\ell(v) > \ell(w)$ , are saturated, we maintain that  $(\mathbf{f}, \ell)$  is a state. Since we keep on performing push, pull or relabel operations as long as there is a vertex  $v$  with  $\Gamma(v) > \nabla(v)$ ,  $\ell(v) < h$  or  $\Gamma(v) < \nabla(v)/2$ ,  $\ell(v) > 0$ , it is guaranteed that the algorithm computes a valid state eventually. It thus suffices to bound the run-time.

Note that by maintaining at each vertex  $v \in V$ , a linked-list  $L[v]$  containing all non-saturated edges  $(v, u)$  where  $\ell(v) = \ell(u) + 1$ , we can implement a push in time  $O(1)$ . We can maintain such a linked-list for every vertex by spending time  $O(\deg(v))$  every time we relabel a vertex  $v$  plus  $O(1)$  time for each push. We maintain a corresponding list  $L'[v]$  for all pull operations. It thus suffices to bound the contribution of the push, pull, and relabel operations to the run-time. Let us index by  $t$  starting with 0 the push, pull, relabel operations as well as the calls to `INCREASESOURCE`, `REMOVEEDGES` in the order they occur. We denote by  $\tilde{E}_t, \mathbf{f}_t, \ell_t, \Gamma_t$  the state of the variables at index  $t$ . We note that  $\mathbf{f}_t$  is supported on the edges  $\tilde{E}_t$ .

Let us introduce functions  $\mathbf{p}_t, \mathbf{n}_t$  where  $\mathbf{p}_t(v), \mathbf{n}_t(v)$  can be interpreted as the amount of positive/negative units present at vertex  $v \in V$  at time  $t$ . These functions will satisfy

$$\forall v \in V, \Gamma_t(v) = \mathbf{p}_t(v) - \mathbf{n}_t(v). \quad (1)$$

Initially, we have  $\mathbf{p}_0 = \Delta_0, \mathbf{n}_0 = \mathbf{0}$ , where  $\Delta_0$  denotes the variable  $\Delta$  at initialization of the data structure. We describe how  $\mathbf{p}_t, \mathbf{n}_t$  evolve over time. If the  $t$ -th action is

1. a push of  $\frac{1}{2}$  units from  $u$  to  $v$ , then  $\mathbf{p}_{t+1} = \mathbf{p}_t - \frac{1}{2} \cdot \mathbf{1}_u + \frac{1}{2} \cdot \mathbf{1}_v$
2. a pull of  $\frac{1}{2}$  units from  $u$  to  $v$ , then  $\mathbf{n}_{t+1} = \mathbf{n}_t + \frac{1}{2} \cdot \mathbf{1}_u - \frac{1}{2} \cdot \mathbf{1}_v$
3. `INCREASESOURCE`( $\delta$ ), then  $\mathbf{p}_{t+1} = \mathbf{p}_t + \delta$
4. `REMOVEEDGES`( $D$ ), then

$$\forall v \in V : \mathbf{p}_{t+1}(v) = \mathbf{p}_t(v) + \sum_{(v,w) \in D} (\mathbf{f}_t)^+(v, w), \quad \mathbf{n}_{t+1}(v) = \mathbf{n}_t(v) + \sum_{(u,v) \in D} (\mathbf{f}_t)^+(u, v)$$

Clearly  $\|\mathbf{p}_t\|_1$  and  $\|\mathbf{n}_t\|_1$  are non-decreasing, we denote the final values by  $\pi, \nu$ . Let us verify that indeed the transitions above preserve the property in (1). For type 1/2/3 actions, this is immediate. For type 4, we observe that for any vertex  $v \in V$  when inducing the flow  $\mathbf{f}_t$  onto  $\tilde{E}_{t+1}$  every unit of  $\sum_{(v,w) \in D} (\mathbf{f}_t)^+(v, w)$  increases the commodity  $\Gamma_t(v)$  at  $v$  by one and every unit of  $\sum_{(u,v) \in D} (\mathbf{f}_t)^+(u, v)$  decreases the commodity  $\Gamma_t(v)$  at  $v$  by one.

▷ **Claim 10.** We can bound the size of  $\pi, \nu$  by

$$O\left(\|\Delta\|_1 + \sum_{d \in \mathcal{D}} |\mathbf{f}_{t_d}(d)|\right),$$

where  $\Delta$  denotes the final state of the variable.

Proof. Only INCREASESOURCE and REMOVEEDGES can increase  $\|\mathbf{p}_t\|_1, \|\mathbf{n}_t\|_1$ . The bound follows immediately from 3, 4 above.  $\triangleleft$

▷ Claim 11. The contribution to the run-time of the relabeling operations is bounded by

$$O\left(\sum_{v \in V} \nabla(v) \sum_{0 \leq t} |\ell_{t+1}(v) - \ell_t(v)|\right) \leq O(h \cdot \pi).$$

Proof. Since  $\deg(v) \leq \nabla(v)$ , we can bound the run-time contribution of the relabeling operations by  $O\left(\sum_{v \in V} \nabla(v) \sum_{0 \leq t} |\ell_{t+1}(v) - \ell_t(v)|\right)$ . It thus suffices to bound this expression. We treat each  $v \in V(G)$  individually. W.l.o.g.  $\ell_t(v)$  increases at some time  $t$ , otherwise the contribution of  $v$  to the sum is trivially zero. Let us consider the following time stamps  $s_0 < s_1 < s_2 < \dots < s_l$  such that  $s_0$  is the first time  $\Gamma_t(v) \geq \nabla(v)$  and

$$\begin{aligned} \forall i \geq 0 : s_{2i+1} &= \min\{t \geq s_{2i} \mid \Gamma_t(v) \leq \nabla(v)/2\}, \\ \forall i \geq 1 : s_{2i} &= \min\{t \geq s_{2i-1} \mid \Gamma_t(v) \geq \nabla(v)\}. \end{aligned}$$

Additionally, we define the time stamps  $t_0 < t_1 < t_2 < \dots$  such that

$$\begin{aligned} \forall i \geq 0 : t_{2i} &= \max\{t < s_{2i+1} \mid \Gamma_t(v) \geq \nabla(v)\}, \\ \forall i \geq 1 : t_{2i-1} &= \max\{t < s_{2i} \mid \Gamma_t(v) \leq \nabla(v)/2\}. \end{aligned}$$

Let us introduce the vector  $\mathbf{z}_t(v) = \min(\mathbf{p}_t(v), \mathbf{n}_t(v) + \nabla(v)/2)$  (this corresponds to the neutrons introduced in the introduction). We observe that  $\mathbf{z}_t(v)$  is non-decreasing. Indeed, if we push away from  $v$  there must be commodity excess, i.e.  $\Gamma_t(v) \geq \nabla(v) + 1/2$  or equivalently  $\mathbf{p}_t(v) \geq \mathbf{n}_t(v) + \nabla(v) + 1/2$ , and if we pull towards  $v$  there must be negative commodity excess, i.e.  $\Gamma_t(v) \leq \nabla(v)/2 - 1/2$  or equivalently  $\mathbf{p}_t(v) \leq \mathbf{n}_t(v) + \nabla(v)/2 - 1/2$ . We will now establish that  $\mathbf{z}_{s_{2(i+1)}}(v) \geq \mathbf{z}_{s_{2i}}(v) + \nabla(v)/2$ . Consider the time stamps  $t_{2i} \leq t < s_{2i+1}$ : since  $\Gamma_{t+1}(v) < \nabla(v)$  the  $t$ -th operation cannot be a push away from  $v$  and thus  $\mathbf{p}_{t_{2i}}(v) \leq \mathbf{p}_{s_{2i+1}}(v)$ . Combining it with the definition of the time stamp  $s_{2i+1}$ , we find

$$\mathbf{n}_{s_{2i+1}}(v) + \nabla(v)/2 \geq \mathbf{p}_{s_{2i+1}}(v) \geq \mathbf{p}_{t_{2i}}(v).$$

This in turn implies that

$$\mathbf{z}_{s_{2(i+1)}}(v) \geq \mathbf{z}_{s_{2i+1}}(v) \geq \mathbf{p}_{t_{2i}}(v) \geq \mathbf{z}_{t_{2i}}(v) + \nabla(v)/2 \geq \mathbf{z}_{s_{2i}}(v) + \nabla(v)/2,$$

where we used in the first and the last inequality that  $\mathbf{z}_t(v)$  is non-decreasing and in the third the definition of the time stamp  $t_{2i}$ . By induction, we have  $\mathbf{z}_{s_{2i}}(v) \geq \mathbf{z}_{s_0}(v) + i \cdot \nabla(v)/2 \geq \mathbf{z}_{s_0}(v) + (i+1) \cdot \nabla(v)/2 \geq (i+1) \cdot \nabla(v)/2$ . This allows us to bound

$$\begin{aligned} \nabla(v) \cdot \sum_{s_0 \leq t} |\ell_{t+1}(v) - \ell_t(v)| &\leq \nabla(v) \cdot (l+1) \cdot 2 \cdot h \\ &\leq 2 \cdot h \cdot (2\lfloor l/2 \rfloor + 2) \cdot \nabla(v) \leq 4 \cdot h \cdot \mathbf{z}_{s_{2\lfloor l/2 \rfloor}}(v), \end{aligned}$$

where the  $h$  is due to the fact that  $\ell_k(v)$  is non-decreasing between  $s_{2i}, s_{2i+1}$  and non-increasing between  $s_{2i+1}, s_{2(i+1)}$ . Summing over  $v$  and bounding  $\mathbf{z}_t(v) \leq \mathbf{p}_t(v)$  concludes the argument.  $\triangleleft$

## 132:14 Near-Optimal Algorithm for Directed Expander Decompositions

▷ Claim 12. The contribution to the run-time of the push operations and the pull operations is bounded by

$$O\left(h \cdot \left(\|\Delta\|_1 + \sum_{d \in \mathcal{D}} |f_{t_d}(d)|\right)\right).$$

Proof. Let us introduce the function  $\Phi(t) = \sum_{v \in V} \Gamma_t(v) \cdot \ell_t(v)$ . We remark that any push or pull operation of  $\frac{1}{2}$  units from  $u$  to  $v$ , where  $\ell_t(u) > \ell_t(v)$ , decreases  $\Gamma_t(u)$  by  $\frac{1}{2}$  and increases  $\Gamma_t(v)$  by at most  $\frac{1}{2}$  while preserving all other  $\Gamma_t$  entries. Hence, any such push or pull operation decreases the function  $\Phi(t)$ . This allows us to bound the run-time contribution of the push and pull operations by  $\sum_{0 \leq t} \max(\Phi(t) - \Phi(t+1), 0)$ . Since at termination the state is valid and thus all vertices with negative excess are on level 0, we have  $\lim_t \Phi(t) \geq 0$  for all  $t$ . Together with  $\Phi(0) = 0$ , this implies that it suffices to bound the increases to  $\Phi$ . As argued before push and pull operations can only decrease the potential. So let  $\mathcal{T}$  be the subset of indices that do not correspond to a push or a pull. Let us denote by  $t_1, t_2, \dots, t_k$  all indices at which we perform a call to INCREASESOURCE or REMOVEEDGES. For any  $i \geq 0$ , we denote by  $\mathcal{T}_{t_i, t_{i+1}}$  all time indices  $t$  in  $\mathcal{T}$  where  $t_i \leq t < t_{i+1}$  (where  $t_{k+1} = \infty$  for notational convenience). To ease notation, we write

$$R = h \cdot \left(\|\Delta\|_1 + \sum_{d \in \mathcal{D}} |f_{t_d}(d)|\right).$$

As a first step, we observe

$$\begin{aligned} \sum_{t \in \mathcal{T}_{t_i, t_{i+1}}} \max(\Phi(t+1) - \Phi(t), 0) &= \sum_{t \in \mathcal{T}_{t_i+1, t_{i+1}}} \sum_{v \in V} |\Gamma_t(v)| \cdot |\ell_{t+1}(v) - \ell_t(v)| \\ &\quad + \sum_{v \in V} (\Gamma_{t_{i+1}}(v) - \Gamma_{t_i}(v)) \cdot \ell_{t_{i+1}}(v), \end{aligned}$$

where the first term of the expression is due to the relabeling operations in  $\mathcal{T}_{t_i, t_{i+1}}$  and the second term is due to the function call at index  $t_i$ . Since  $\Gamma_{t_i}(v)$  can only increase if  $\mathbf{p}_{t_i}(v)$  increases, we may thus bound the contribution of a vertex of the second term by

$$(\mathbf{p}_{t_{i+1}}(v) - \mathbf{p}_{t_i}(v)) \cdot h. \quad (2)$$

Combining Claim 10, Claim 11, we find that  $\sum_i \sum_{t \in \mathcal{T}_{t_i+1, t_{i+1}}} \sum_{v \in V} 2\nabla(v) \cdot |\ell_{t+1}(v) - \ell_t(v)| \leq O(R)$ . We may thus subtract  $\sum_{t \in \mathcal{T}_{t_i+1, t_{i+1}}} \sum_{v \in V} 2\nabla(v) \cdot |\ell_{t+1}(v) - \ell_t(v)|$  and reduce the discussion to bounding the expression

$$\sum_{t \in \mathcal{T}_{t_i+1, t_{i+1}}} \sum_{v \in V} (|\Gamma_t(v)| - 2\nabla(v)) \cdot |\ell_{t+1}(v) - \ell_t(v)| + \sum_{v \in V} (\Gamma_{t_{i+1}}(v) - \Gamma_{t_i}(v)) \cdot \ell_{t_{i+1}}(v). \quad (3)$$

We treat each vertex  $v \in V$  individually. Let  $\tau_i(v) \geq 0$  be maximum such that  $\forall t_i \leq t < t_i + \tau_i(v) : \Gamma_t(v) > 2 \cdot \nabla(v)$  or  $\Gamma_t(v) < -2 \cdot \nabla(v)$ . Let us consider the first term of the expression (3). Note that if  $|\Gamma_t(v)| \leq 2\nabla(v)$  then after any push or pull operation at time  $t$ , we still have that  $|\Gamma_{t+1}(v)| \leq 2\nabla(v)$ . This implies that for all  $t \in \mathcal{T}_{t_i+\tau_i, t_{i+1}} : |\Gamma_t(v)| \leq 2\nabla(v)$ . It thus suffices to consider for each  $v \in V$  the expression

$$\begin{aligned} &\sum_{t \in \mathcal{T}_{t_i+1, t_i+\tau_i(v)}} (|\Gamma_t(v)| - 2\nabla(v)) \cdot |\ell_{t+1}(v) - \ell_t(v)| \\ &\leq (|\Gamma_{t_{i+1}}(v)| - 2\nabla(v)) \cdot \sum_{t \in \mathcal{T}_{t_i+1, t_i+\tau_i(v)}} |\ell_{t+1}(v) - \ell_t(v)|. \end{aligned} \quad (4)$$

Let us consider the vertices  $v$  such that the quantity above is strictly positive. Clearly,  $\Gamma_{t_i+1}(v) > 2\nabla(v)$  or  $\Gamma_{t_i+1}(v) < -2\nabla(v)$ . Let us consider the first case. Since at time  $t_i$  the tuple  $(\mathbf{f}_{t_i}, \ell_{t_i})$  is valid, we must have that  $\Gamma_{t_i}(v) < 2\nabla(v)$  otherwise  $\ell_{t_i}(v) = h$  and hence  $\sum_{t \in \mathcal{T}_{t_i+1, t_i+\tau_i(v)}} |\ell_{t+1}(v) - \ell_t(v)| = 0$ . Since  $\Gamma_{t_i}(v)$  can only increase if  $\mathbf{p}_{t_i}(v)$  increases, we can bound

$$(|\Gamma_{t_i+1}(v)| - 2\nabla(v)) \leq (\Gamma_{t_i+1}(v) - \Gamma_{t_i}(v)) \leq \mathbf{p}_{t_i+1}(v) - \mathbf{p}_{t_i}(v).$$

Moreover, the label can only increase during  $\mathcal{T}_{t_i+1, t_i+\tau_i(v)}$  and hence  $\sum_{t \in \mathcal{T}_{t_i+1, t_i+\tau_i(v)}} |\ell_{t+1}(v) - \ell_t(v)| \leq h$ . We may thus bound the contribution of such a vertex to (4) by

$$(\mathbf{p}_{t_i+1}(v) - \mathbf{p}_{t_i}(v)) \cdot h. \quad (5)$$

If on the other hand  $\Gamma_{t_i+1}(v) < -2\nabla(v)$ , then since at time  $t_i$  the tuple  $(\mathbf{f}_{t_i}, \ell_{t_i})$  is valid we must have that  $\Gamma_{t_i}(v) \geq 0$  otherwise  $\ell_{t_i}(v) = 0$  and hence  $\sum_{t \in \mathcal{T}_{t_i+1, t_i+\tau_i(v)}} |\ell_{t+1}(v) - \ell_t(v)| = 0$ . Since  $\Gamma_{t_i}(v)$  can only decrease if  $\mathbf{n}_{t_i}(v)$  increases, we can bound

$$(|\Gamma_{t_i+1}(v)| - 2\nabla(v)) \leq (\Gamma_{t_i}(v) - \Gamma_{t_i+1}(v)) \leq \mathbf{n}_{t_i+1}(v) - \mathbf{n}_{t_i}(v).$$

Moreover, the label can only decrease during  $\mathcal{T}_{t_i+1, t_i+\tau_i(v)}$  and hence  $\sum_{t \in \mathcal{T}_{t_i+1, t_i+\tau_i(v)}} |\ell_{t+1}(v) - \ell_t(v)| \leq h$ . We may thus bound the contribution of such a vertex to (4) by

$$(\mathbf{n}_{t_i+1}(v) - \mathbf{n}_{t_i}(v)) \cdot h. \quad (6)$$

Summing over all time indices  $t_1, t_2, \dots$  and combining Claim 10, Claim 11, we can bound both equation 2, 5, 6 by  $O(R)$ . This yields the conclusion.  $\triangleleft$

#### 4 Directed Expander Decompositions via the Push-Pull-Relabel Framework

In this section, we discuss our main technical result that states that if  $G$  is initially a  $\phi$ -expander, then an expander decomposition can be maintained efficiently. The algorithm behind this theorem heavily relies on the new push-pull-relabel algorithm presented in the previous section.

► **Theorem 13.** *Given a  $\phi$ -expander  $G = (V, E)$  with  $(\phi, \psi)$ -witness  $(W, \Pi)$ , there is a deterministic data structure `BIDIRECTEDEXPANDERPRUNING`( $G, W, \Pi$ ) (see Algorithm 2). After every update of the form*

■ `REMOVEEDGES`( $D$ ): *where  $D \subseteq E$ , sets  $\mathcal{D}$  to  $\mathcal{D} \cup D$  (initially  $\mathcal{D} = \emptyset$ ), the algorithm explicitly updates  $\tilde{V} \subseteq V$  and the tuple  $(E^r, \mathcal{P})$ , where  $E^r \subseteq E(G)$  and  $\mathcal{P}$  is a partition of  $V \setminus \tilde{V}$  (initially  $E^r, \mathcal{P} = \emptyset$ ), such that both  $E^r$  and  $\mathcal{P}$  are non-decreasing and such that after the update*

1. *the graph  $\tilde{G} = (G \setminus \mathcal{D})[\tilde{V}]$  has a  $\left(\frac{\phi \cdot \psi^6}{32000}, \frac{\psi^4}{800}\right)$ -witness  $(\tilde{W}, \Pi)$ ,*
  2.  $|E^r| \leq \frac{\phi}{4} \cdot \sum_{P \in \mathcal{P}} \text{vol}_G(P)$
  3.  $\sum_{P \in \mathcal{P}} \text{vol}_{\tilde{W}}(P) \leq \frac{8}{\psi} \cdot |\Pi^{-1}(\mathcal{D})|$
  4.  $(G \setminus (\mathcal{D} \cup E^r)) / (\mathcal{P} \cup \{\tilde{V}\})$  *is a directed acyclic graph (DAG),*
- provided  $\frac{200}{\psi^2} \cdot |\Pi^{-1}(\mathcal{D})| < e(G)$ . The run-time is*
- $O\left(\frac{h}{\psi^2} \cdot |\Pi^{-1}(\mathcal{D})| + h \cdot \sum_{d \in \mathcal{D} \cup E_{\mathcal{P}}} (1 + |\mathbf{f}_{t_d}(d)|)\right) = O\left(h \cdot \frac{|\mathcal{D}|}{\psi^2 \phi}\right)$ , *where  $\mathcal{D}$  denotes the variable at the end of all updates and  $E_{\mathcal{P}}$  denotes all intercluster edges of the partition  $\mathcal{P} \cup \{\tilde{V}\}$ .*



Following the high-level approach of [31], we obtain our main result Theorem 2 for the special case where  $G$  is static by running a generalization of the cut-matching game to directed graphs [22, 27] and then apply Theorem 13 to handle unbalanced cuts. Combining this result again with Theorem 13, we obtain our main result Theorem 2 in full generality. Both of these reductions have been known from previous work [2, 18] and are omitted.

#### 4.1 Reduction to Out-Expanders

In this subsection, we show that the task of maintaining an expander decomposition as described in Theorem 13 can be reduced to a simpler problem that only requires maintaining out-expanders (however under edge and vertex deletions). In particular, we reduce to the following statement whose algorithm is deferred to Section 5 and the proof is omitted.

► **Lemma 14.** *For every  $\phi$ -out-expander  $G = (V, E)$  with  $(\phi, \psi)$ -witness  $(W, \Pi)$ , there is a deterministic data structure  $\text{DIRECTEDEXPANDERPRUNING}(G, W, \Pi)$  (see Algorithm 3). After every update of the form (initially  $\tilde{V} = V, \mathcal{S} = \emptyset, E_{\mathcal{S}} = \emptyset, E_{\mathcal{S}}^+ = \emptyset, E_{\mathcal{S}}^- = \emptyset, \mathcal{D} = \emptyset$ )*

- $\text{REMOVEEDGES}(D)$ : where  $D \subseteq E(\tilde{V}) \setminus \mathcal{D}$ , sets  $\mathcal{D}$  to  $\mathcal{D} \cup D$ ,
- $\text{REMOVEVERTICES}(S)$ : where  $S \subseteq \tilde{V}$ , sets  $\tilde{V}$  to  $\tilde{V} \setminus S$ ,  $\mathcal{S}$  to  $\mathcal{S} \cup S$ ,  $E_{\mathcal{S}}^+$  to  $E_{\mathcal{S}}^+ \cup E_G(S, \tilde{V} \setminus S)$ ,  $E_{\mathcal{S}}^-$  to  $E_{\mathcal{S}}^- \cup E_G(\tilde{V} \setminus S, S)$  and  $E_{\mathcal{S}}$  to  $E_{\mathcal{S}}^+ \cup E_{\mathcal{S}}^-$

*the algorithm explicitly updates  $\tilde{V} \subseteq V$  and the tuple  $(E^r, \mathcal{P})$ , where  $E^r \subseteq E$  and  $\mathcal{P}$  is a partition of  $V \setminus (\mathcal{S} \cup \tilde{V})$  (initially  $E^r = \emptyset, \mathcal{P} = \emptyset$ ), such that both  $E^r$  and  $\mathcal{P}$  are non-decreasing and such that after the update*

1. *the graph  $(G \setminus \mathcal{D})[\tilde{V}]$  has a  $\left(\frac{\phi \cdot \psi^4}{400}, \frac{\psi^2}{20}\right)$ -out-witness  $(\tilde{W}, \tilde{\Pi})$ ,*
  2.  $|E^r| \leq \frac{\phi}{4} \cdot \sum_{P \in \mathcal{P}} \text{vol}_G(P)$ ,
  3.  $\sum_{P \in \mathcal{P}} \text{vol}_{\tilde{W}}(P) \leq \frac{4}{3\psi} \cdot |\Pi^{-1}(\mathcal{D} \cup E_{\mathcal{S}}^-)|$ ,
  4.  $E^r = \bigcup_P E_{G \setminus (\mathcal{D} \cup E_{\mathcal{S}}^-)}(P, V_P)$ , where  $V_P$  is equal to  $\tilde{V}$  at the time  $P$  is added to  $\mathcal{P}$ ,
- provided  $\frac{20}{\psi} \cdot |\Pi^{-1}(\mathcal{D} \cup E_{\mathcal{S}} \cup E_{\mathcal{P}})| < e(G)$ , where  $\mathcal{D}$  denotes the variable at the end of all updates and  $E_{\mathcal{P}}$  denotes all intercluster edges of the partition  $\mathcal{P} \cup \{\tilde{V}\}$ . The algorithm runs in total time  $O\left(\frac{h \cdot |\Pi^{-1}(\mathcal{D} \cup E_{\mathcal{S}} \cup E_{\mathcal{P}})|}{\psi^2} + h \cdot \sum_{d \in \mathcal{D} \cup E_{\mathcal{S}} \cup E_{\mathcal{P}}} (1 + |f_{t_d}(d)|)\right)$ .*

■ **Algorithm 2**  $\text{BiDIRECTEDEXPANDERPRUNING}(G, W, \Pi)$ .

---

```

def Init():
     $(\tilde{V}_1, \mathcal{P}_1, E_1^r) \leftarrow \text{DIRECTEDEXPANDERPRUNING}(G, W, \Pi)$ 
     $(\tilde{V}_2, \mathcal{P}_2, E_2^r) \leftarrow \text{DIRECTEDEXPANDERPRUNING}(G^{rev}, W, \Pi)$ 
     $\tilde{V} \leftarrow \tilde{V}_1 \cap \tilde{V}_2$ 
     $\mathcal{P} \leftarrow \mathcal{P}_1 \cup \mathcal{P}_2$ 
     $E^r \leftarrow E_1^r \cup E_2^r$  // dynamically updated

def RemoveEdges(D):
     $(\tilde{V}_1, \mathcal{P}_1, E_1^r).$ REMOVEEDGES( $D$ )
     $(\tilde{V}_2, \mathcal{P}_2, E_2^r).$ REMOVEVERTICES( $\tilde{V}_2 \setminus \tilde{V}_1$ )
     $(\tilde{V}_2, \mathcal{P}_2, E_2^r).$ REMOVEEDGES( $D$ )
    AdjustPartition()

def AdjustPartition():
    while  $\tilde{V}_1 \neq \tilde{V}_2$  do
         $(\tilde{V}_1, \mathcal{P}_1, E_1^r).$ REMOVEVERTICES( $\tilde{V}_1 \setminus \tilde{V}_2$ )
         $(\tilde{V}_2, \mathcal{P}_2, E_2^r).$ REMOVEVERTICES( $\tilde{V}_2 \setminus \tilde{V}_1$ )

```

---

We present **BiDIRECTEDEXPANDERPRUNING** (see Algorithm 2), which implements this reduction. The **INIT** function initializes two out-expander pruning data structure. One for the graph  $G$  with witness  $(W, \Pi)$  with variables  $\tilde{V}_1, \mathcal{P}_1, E_1^r$  and one for the graph  $G^{rev}$  with witness  $(W, \Pi)$  with variables  $\tilde{V}_2, \mathcal{P}_2, E_2^r$ . From the guarantees of the one directional **DIRECTEDEXPANDERPRUNING** algorithm, we will have that after each update  $G[\tilde{V}_1]$  and  $G^{rev}[\tilde{V}_2]$  are out-expanders. To conclude that the  $G[\tilde{V}]$  is in fact an expander in both directions, we will ensure using the function **ADJUSTPARTITION** that  $\tilde{V}_1, \tilde{V}_2$  agree. If they do not agree, we will remove the vertices  $\tilde{V}_1 \setminus \tilde{V}_2$  from  $\tilde{V}_1$  using the function **REMOVEVERTICES** of the algorithm **DIRECTEDEXPANDERPRUNING** (see Algorithm 3). Using the function **REMOVEEDGES** of **BiDIRECTEDEXPANDERPRUNING**, we can remove edges from  $G[\tilde{V}]$ . This is again accomplish by calling the analogous function **REMOVEEDGES** of **DIRECTEDEXPANDERPRUNING** for both directions. Again we have to adjust the partition such that  $\tilde{V}_1, \tilde{V}_2$  agree afterwards. The algorithm and the proof are omitted.

## 5 Maintaining Out-Expanders

It remains to prove Lemma 14. Algorithm 3 implements of **DIRECTEDEXPANDERPRUNING**. The **INIT** function initializes for an out-expander  $(G, W, \Pi)$  the expander pruning variables  $\tilde{V}, \mathcal{P}, E^r, \mathcal{D}$  as well as a first valid state  $(f, \ell)$  of the **VALIDSTATE** algorithm (see Algorithm 1).

■ **Algorithm 3** **DIRECTEDEXPANDERPRUNING**( $G, W, \Pi$ ).

---

```

def Init():
     $\tilde{V} \leftarrow V, \mathcal{P} \leftarrow \emptyset, E_S^+ \leftarrow \emptyset, E_S^- \leftarrow \emptyset, E^r \leftarrow \emptyset, \mathcal{D} \leftarrow \emptyset$ 
     $(f, \ell) \leftarrow \text{VALIDSTATE}(G, \frac{18}{\phi\psi^2} \cdot \mathbf{1}, \mathbf{0}, \deg_W, h)$ 

def RemoveEdges(D):
     $\mathcal{D} \leftarrow \mathcal{D} \cup D$ 
     $(f, \ell).\text{REMOVEEDGES}(D)$ 
     $(f, \ell).\text{INCREASESOURCE}(\frac{4}{\psi} \deg_{\Pi^{-1}(D)})$ 
    AdjustPartition()

def RemoveEdges(D):
     $\tilde{V} \leftarrow \tilde{V} \setminus S, E_S^+ \leftarrow E_S^+ \cup E_{G \setminus \mathcal{D}}(S, \tilde{V}), E_S^- \leftarrow E_S^- \cup E_{G \setminus \mathcal{D}}(\tilde{V}, S)$ 
     $(f, \ell).\text{REMOVEEDGES}(E_{G \setminus \mathcal{D}}(\tilde{V}, S) \cup E_{G \setminus \mathcal{D}}(S, \tilde{V}))$ 
     $(f, \ell).\text{INCREASESOURCE}(\frac{4}{\psi} \deg_{\Pi^{-1}}(E_{G \setminus \mathcal{D}}(\tilde{V}, S) \cup E_{G \setminus \mathcal{D}}(S, \tilde{V})))$ 
    AdjustPartition()

def AdjustPartition():
    while  $\exists v \in \tilde{V}$  with  $\ell(v) = h$  do
         $P \leftarrow \text{Prune}(G \setminus \mathcal{D}, \ell)$ 
         $\mathcal{P} \leftarrow \mathcal{P} \cup \{P\}, \tilde{V} \leftarrow \tilde{V} \setminus P, E^r \leftarrow E^r \cup E_{G \setminus \mathcal{D}}(P, \tilde{V})$ 
         $(f, \ell).\text{REMOVEEDGES}(E_{G \setminus \mathcal{D}}(P, \tilde{V}) \cup E_{G \setminus \mathcal{D}}(\tilde{V}, P))$ 
         $(f, \ell).\text{INCREASESOURCE}(\frac{4}{\psi} \deg_{\Pi^{-1}}(E_{G \setminus \mathcal{D}}(P, \tilde{V}) \cup E_{G \setminus \mathcal{D}}(\tilde{V}, P)))$ 

def Prune(G,  $\ell$ ):
     $S \leftarrow \emptyset, i \leftarrow h$ 
    repeat
         $S \leftarrow S \cup \{v \in \tilde{V} \mid \ell(v) = i\}$ 
         $i \leftarrow i - 1$ 
    until  $\text{vol}_{G \setminus (\mathcal{D} \cup E_S^-)}(\{v \in \tilde{V} \mid \ell(v) \leq i\}) < (1 + \frac{\phi}{72}) \cdot \text{vol}_{G \setminus (\mathcal{D} \cup E_S^-)}(S)$ 
    return the cut  $S$ 

```

---

This valid state  $(\mathbf{f}, \ell)$ , we will use to find the pruning cuts in function PRUNE. Using REMOVEEDGES the user can remove edges  $D$  from  $G[\tilde{V}]$ . Removing edges might force us to prune away some part of the  $G[\tilde{V}]$  and add the pruned part to the collection of pruning sets  $\mathcal{P}$ . To find the pruning set  $P$ , we adopt a similar strategy as in [31]. We inject additional source flow into the flow problem: for any witness edge  $e = (u, v) \in E(W)$ , where an edge on the embedding path  $\Pi(e)$  is in  $D$ , we increase the sources  $\Delta(u), \Delta(v)$  by  $\frac{4}{\psi}$ . Since we updated the flow problem by increasing the source capacities and removing edges, we need to update the valid state  $(\mathbf{f}, \ell)$ . This new valid state  $(\mathbf{f}, \ell)$ , then allows us to locate the pruning set in the function ADJUSTPARTITION. In ADJUSTPARTITION, we check whether there is a vertex in the remaining graph  $G[\tilde{V}]$  on level  $h$ . If there is no such vertex, it will certify that in fact  $G[\tilde{V}]$  has already a good out-witness and we don't need to prune away any subgraph. If on the other hand, there is a vertex on level  $h$  we will find a pruning set  $P$  using the function PRUNE. This set  $P$  is then added to the collection of pruning sets  $\mathcal{P}$  and the vertices in  $P$  are removed from  $\tilde{V}$ . The cut edges  $E_{G \setminus \mathcal{D}}(P, \tilde{V})$  are added to the remove edges  $E^r$ . Since  $\tilde{V}$  has become smaller, we need to update the valid state  $(\mathbf{f}, \ell)$  again. We do this once more by injecting additional source at the endpoints of any witness edge  $e \in E(W)$ , where some edge on the path  $\Pi(e)$  is in  $E_{G \setminus \mathcal{D}}(\tilde{V}, P)$ . Thereafter we again check whether in the new valid state  $(\mathbf{f}, \ell)$  there is a vertex in  $\tilde{V}$  on level  $h$ . We keep on doing this procedure until there no longer is a vertex on level  $h$ . We will prove that the volume of the pruning sets can be related to the size of the set of edges  $D$  initially removed by the user. So far we have not explained how to find the pruning set in function PRUNE. This is accomplished by a standard level cutting procedure. We start with  $P$  being all vertices on level  $h$  and then check whether the vertices on the next level have volume at least  $\phi \cdot \text{vol}_G(P)$ . If it is the case, we add vertices on the next level to  $P$  and otherwise return  $P$ . Through the function REMOVEVERTICES the user can remove vertices  $S$  from  $\tilde{V}$ . Similar to the function REMOVEEDGES, we will have to inject additional source at the endpoints of the witness edges  $e \in E(W)$ , where some edge of  $\Pi(e)$  is in  $E_{G \setminus \mathcal{D}}(\tilde{V}, S)$ , and update the valid state  $(\mathbf{f}, \ell)$  accordingly. This might again leave some vertices on level  $h$  and will again force us to prune some vertices. This is again accomplished by a call to ADJUSTPARTITION. And similar to REMOVEEDGES, we will again prove that the volume of the pruning sets can be related to the volume of the set  $S$  initially removed by the user.

---

## References

- 1 Aaron Bernstein, Joakim Blikstad, Thatchaphol Saranurak, and Ta-Wei Tu. Maximum flow by augmenting paths in  $n^{2+o(1)}$  time. *to appear at FOCS'24*, 2024.
- 2 Aaron Bernstein, Maximilian Probst Gutenberg, and Thatchaphol Saranurak. Deterministic decremental reachability, scc, and shortest paths via directed expanders and congestion balancing. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1123–1134. IEEE, 2020. doi:10.1109/FOCS46700.2020.00108.
- 3 Aaron Bernstein, Maximilian Probst Gutenberg, and Thatchaphol Saranurak. Deterministic decremental sssp and approximate min-cost flow in almost-linear time. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1000–1008. IEEE, 2022.
- 4 Aaron Bernstein, Maximilian Probst Gutenberg, and Christian Wulff-Nilsen. Near-optimal decremental sssp in dense weighted digraphs. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1112–1122. IEEE, 2020. doi:10.1109/FOCS46700.2020.00107.
- 5 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. In *49th International Colloquium on Automata, Languages, and Programming (ICALP 2022)*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022.
- 6 Parinya Chalermsook, Syamantak Das, Yunbum Kook, Bundit Laekhanukit, Yang P Liu, Richard Peng, Mark Sellke, and Daniel Vaz. Vertex sparsification for edge connectivity. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1206–1225. SIAM, 2021. doi:10.1137/1.9781611976465.74.
  - 7 Li Chen, Rasmus Kyng, Yang P Liu, Simon Meierhans, and Maximilian Probst Gutenberg. Almost-linear time algorithms for incremental graphs: Cycle detection, sccs,  $s$ - $t$  shortest path, and minimum-cost flow. *to appear at STOC'24*, 2024.
  - 8 Li Chen, Rasmus Kyng, Yang P Liu, Richard Peng, Maximilian Probst Gutenberg, and Sushant Sachdeva. Maximum flow and minimum-cost flow in almost-linear time. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 612–623. IEEE, 2022. doi:10.1109/FOCS54457.2022.00064.
  - 9 Julia Chuzhoy. Decremental all-pairs shortest paths in deterministic near-linear time. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 626–639, 2021. doi:10.1145/3406325.3451025.
  - 10 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. doi:10.1145/3313276.3316320.
  - 11 Julia Chuzhoy and Sanjeev Khanna. A faster combinatorial algorithm for maximum bipartite matching. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2185–2235. SIAM, 2024. doi:10.1137/1.9781611977912.79.
  - 12 Julia Chuzhoy and Sanjeev Khanna. Maximum bipartite matching in  $n^{2+o(1)}$  time via a combinatorial algorithm. *to appear at STOC'24*, 2024.
  - 13 Julia Chuzhoy and Thatchaphol Saranurak. Deterministic algorithms for decremental shortest paths via layered core decomposition. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2478–2496. SIAM, 2021. doi:10.1137/1.9781611976465.147.
  - 14 Andrew V Goldberg and Robert E Tarjan. A new approach to the maximum-flow problem. *Journal of the ACM (JACM)*, 35(4):921–940, 1988. doi:10.1145/48014.61051.
  - 15 Gramoz Goranci, Harald Räcke, Thatchaphol Saranurak, and Zihan Tan. The expander hierarchy and its applications to dynamic graph algorithms. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2212–2228. SIAM, 2021. doi:10.1137/1.9781611976465.132.
  - 16 Monika Henzinger, Satish Rao, and Di Wang. Local flow partitioning for faster edge connectivity. *SIAM Journal on Computing*, 49(1):1–36, 2020. doi:10.1137/18M1180335.
  - 17 Dorit S Hochbaum. The pseudoflow algorithm: A new algorithm for the maximum-flow problem. *Operations research*, 56(4):992–1009, 2008. doi:10.1287/OPRE.1080.0524.
  - 18 Yiding Hua, Rasmus Kyng, Maximilian Probst Gutenberg, and Zihang Wu. Maintaining expander decompositions via sparse cuts. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 48–69. SIAM, 2023. doi:10.1137/1.9781611977554.CH2.
  - 19 Wenyu Jin and Xiaorui Sun. Fully dynamic  $st$  edge connectivity in subpolynomial time. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 861–872. IEEE, 2022.
  - 20 Wenyu Jin, Xiaorui Sun, and Mikkel Thorup. Fully dynamic min-cut of superconstant size in subpolynomial time. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2999–3026. SIAM, 2024. doi:10.1137/1.9781611977912.107.
  - 21 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*, pages 217–226. SIAM, 2014. doi:10.1137/1.9781611973402.16.
- 22 Rohit Khandekar, Satish Rao, and Umesh Vazirani. Graph partitioning using single commodity flows. *Journal of the ACM (JACM)*, 56(4):1–15, 2009. doi:10.1145/1538902.1538903.
  - 23 Rasmus Kyng, Simon Meierhans, and Maximilian Probst Gutenberg. A dynamic shortest paths toolbox: Low-congestion vertex sparsifiers and their applications. *to appear at STOC’24*, 2024.
  - 24 Rasmus Kyng, Simon Meierhans, and Maximilian Probst. Derandomizing directed random walks in almost-linear time. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 407–418. IEEE, 2022. doi:10.1109/FOCS54457.2022.00046.
  - 25 Jason Li. Deterministic mincut in almost-linear time. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 384–395, 2021. doi:10.1145/3406325.3451114.
  - 26 Yang P Liu. Vertex sparsification for edge connectivity in polynomial time. In *14th Innovations in Theoretical Computer Science Conference (ITCS 2023)*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023.
  - 27 Anand Louis. Cut-matching games on directed graphs. *arXiv preprint arXiv:1010.1047*, 2010. arXiv:1010.1047.
  - 28 Danupon Nanongkai and Thatchaphol Saranurak. Dynamic spanning forest with worst-case update time: adaptive, las vegas, and  $o(n^{1/2-\varepsilon})$ -time. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1122–1129, 2017. doi:10.1145/3055399.3055447.
  - 29 Danupon Nanongkai, Thatchaphol Saranurak, and Christian Wulff-Nilsen. Dynamic minimum spanning forest with subpolynomial worst-case update time. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 950–961. IEEE, 2017. doi:10.1109/FOCS.2017.92.
  - 30 Thatchaphol Saranurak. A simple deterministic algorithm for edge connectivity. In *Symposium on Simplicity in Algorithms (SOSA)*, pages 80–85. SIAM, 2021. doi:10.1137/1.9781611976496.9.
  - 31 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*, pages 2616–2635. SIAM, 2019. doi:10.1137/1.9781611975482.162.
  - 32 Daniel A Spielman and Shang-Hua Teng. Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, pages 81–90, 2004. doi:10.1145/1007352.1007372.
  - 33 Jan Van Den Brand, Li Chen, Rasmus Kyng, Yang P Liu, Simon Meierhans, Maximilian Probst Gutenberg, and Sushant Sachdeva. Almost-linear time algorithms for decremental graphs: Min-cost flow and more via duality. *to appear at FOCS’24*, 2024.
  - 34 Jan Van Den Brand, Li Chen, Richard Peng, Rasmus Kyng, Yang P Liu, Maximilian Probst Gutenberg, Sushant Sachdeva, and Aaron Sidford. A deterministic almost-linear time algorithm for minimum-cost flow. In *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 503–514. IEEE, 2023. doi:10.1109/FOCS57990.2023.00037.
  - 35 Jan Van Den Brand, Yin Tat Lee, Yang P Liu, Thatchaphol Saranurak, Aaron Sidford, Zhao Song, and Di Wang. Minimum cost flows, mdps, and l1-regression in nearly linear time for dense instances. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 859–869, 2021. doi:10.1145/3406325.3451108.
  - 36 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*, pages 1130–1143, 2017. doi:10.1145/3055399.3055415.