

Incremental Approximate Maximum Flow via Residual Graph Sparsification

GRAMOZ GORANCI, Faculty of Computer Science, University of Vienna, Vienna, Austria

MONIKA HENZINGER, Institute of Science and Technology Austria, Klosterneuburg, Austria

HARALD RÄCKE, Technical University Munich, Munich, Germany

A. R. SRICHARAN, Doctoral School Computer Science DoCS, Faculty of Computer Science, University of Vienna, Vienna, Austria

We give an algorithm that, with high probability, maintains a $(1 - \epsilon)$ -approximate s - t maximum flow in undirected, uncapacitated n -vertex graphs undergoing m edge insertions in $\tilde{O}(m + nF^*/\epsilon)$ total update time, where F^* is the maximum flow on the final graph. This is the first algorithm to achieve polylogarithmic amortized update time for dense graphs ($m = \Omega(n^2)$), and more generally, for graphs where $F^* = \tilde{O}(m/n)$.

At the heart of our incremental algorithm is the residual graph sparsification technique of Karger and Levine [STOC '02, SICOMP '15], originally designed for computing exact maximum flows in the static setting. Our main contributions are (i) showing how to maintain such sparsifiers for approximate maximum flows in the incremental setting and (ii) generalizing the cut sparsification framework of Fung et al. [STOC '11, SICOMP '19] from undirected graphs to balanced directed graphs.

CCS Concepts: • **Theory of computation** → Dynamic graph algorithms;

Additional Key Words and Phrases: incremental flow, sparsification, approximate flow

ACM Reference format:

Gramoz Goranci, Monika Henzinger, Harald Räcke, and A. R. Sricharan. 2026. Incremental Approximate Maximum Flow via Residual Graph Sparsification. *ACM Trans. Algor.* 22, 3, Article 31 (July 2026), 24 pages. <https://doi.org/10.1145/3816252>

1 Introduction

The maximum s - t flow problem has been at the forefront of research in theoretical computer science and combinatorial optimization. Algorithms developed for maximum s - t flow and its dual problem,

M. Henzinger: This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (MoDynStruct, No. 101019564)  and the Austrian Science Fund (FWF) grant DOI 10.55776/Z422, grant DOI 10.55776/I5982, and grant DOI 10.55776/P33775 with additional funding from the netidee SCIENCE Stiftung, 2020–2024. Views and opinions expressed are those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

Authors' Contact Information: Gramoz Goranci, Faculty of Computer Science, University of Vienna, Vienna, Austria; e-mail: gramoz.goranci@univie.ac.at; Monika Henzinger, Institute of Science and Technology Austria, Klosterneuburg, Austria; e-mail: monika.henzinger@ist.ac.at; Harald Räcke, Technical University Munich, Munich, Germany; e-mail: raecke@in.tum.de; A. R. Sricharan (corresponding author), Doctoral School Computer Science DoCS, Faculty of Computer Science, University of Vienna, Vienna, Austria; e-mail: sricharan.arunapuram@univie.ac.at.



This work is licensed under Creative Commons Attribution International 4.0.

© 2026 Copyright held by the owner/author(s).

ACM 1549-6333/2026/7-ART31

<https://doi.org/10.1145/3816252>

minimum s - t cut, have been highly influential due to their wide applicability [1] and their use as subroutines in other algorithms [2, 28]. A long line of work has improved our understanding of how efficiently maximum flow can be solved in the static setting. Two landmark combinatorial results are the deterministic algorithm by Goldberg and Rao [17], which achieves a running time of $O(\min\{n^{2/3}, m^{1/2}\} \cdot m)$, and the randomized algorithm by Karger and Levine [24] which has a running time of $\tilde{O}(m + nF^*)$,¹ where F^* is the value of the maximum flow. More recently, efforts building upon a novel blend of continuous optimization techniques, the Laplacian paradigm [33] and graph-based data structures have led to even faster algorithms; notably, a randomized $\tilde{O}(m/\epsilon)$ time approximate maximum flow algorithm on undirected graphs [25, 27, 29, 30], and a deterministic $\tilde{O}(m)$ time *exact* algorithm for min-cost flow (and thus maximum flow) even on directed graphs [9, 10], which constitutes a major algorithmic breakthrough.

Recently, there has been growing interest in solving the maximum s - t flow problem in the challenging *dynamic* setting, where the goal is to maintain a flow (or its value) under edge insertions and deletions and answer queries about the maintained flow efficiently. For directed graphs, there are strong conditional hardness results [11, 21] showing a lower bound of $\Omega(n)$ and $\Omega(\sqrt{m})$ amortized update time for maintaining *exact* maximum flow, assuming the **Online Matrix vector (OMv)** conjecture. These strong polynomial lower bounds have motivated a shift in focus toward maintaining *approximate* maximum flows.

Research on maintaining approximate maximum flows in the dynamic setting can be categorized into the following three lines of work. The first line of work [4, 8, 19] deals with the fully dynamic setting, supporting both edge insertions and deletions, and is based on dynamically maintaining tree-based cut approximations. However, these algorithms require at least a logarithmic loss in the quality of the maintained flow. The second line of work [18, 20, 22] is purely combinatorial and works by repeatedly finding augmenting paths in an incremental residual graph together with a lazy rebuilding technique. While these algorithms achieve sub-linear update time, it is not clear how to use them to go beyond the $\sqrt{n}$ barrier on the update time. To address this challenge, the third line of work [4–6, 9] leverages continuous optimization techniques to maintain a $(1 - \epsilon)$ -approximate max flow in the incremental (only edge insertions) and decremental (only edge deletions) settings in $m^{o(1)}$ update time. Table 1 offers a detailed summary of the state-of-the-art results on dynamic maximum flow algorithms.

Focusing on the $(1 - \epsilon)$ -approximation regime, the recent partially dynamic maximum flow algorithms suffer from the following two main drawbacks: (i) they depend on the powerful yet intricate machinery of continuous optimization methods, such as interior point methods, monotone multiplicative weight updates, and dynamic min-ratio cycle/cut problems, (ii) all existing algorithms incur an $m^{o(1)}$ factor in the update time, a limitation that arises in many modern dynamic graph-based data structures and appears difficult to overcome. This leads to the following fundamental question:

Is there an incremental maximum flow algorithm that achieves $(1 - \epsilon)$ approximation with polylogarithmic update time?

We answer this question in the affirmative for dense graphs that are undirected and uncapacitated, as summarized in the theorem below.

THEOREM 1.1. *Given any $\epsilon \in (0, 1)$, there is an incremental randomized algorithm against output-adaptive adversaries that maintains a $(1 - \epsilon)$ -approximate maximum s - t flow f under edge insertions*

¹In this article, we use $\tilde{O}(X)$ to denote $O(X \text{ polylog}(X))$ and $\hat{O}(X)$ to denote $O(X^{1+o(1)})$.

Table 1. Results on Dynamic Max Flow

Setting	Apx. Factor	Flow/Value	Directed	Weighted	Update Time	Reference
Incremental	1	Flow	Yes	Yes	$O(F^*)$	[20, 22]
	1	Flow	No	No	$\tilde{O}(n^{2.5}m^{-1})$	[18]
	$1 + \varepsilon$	Flow	Yes	No	$\tilde{O}(m^{0.5}\varepsilon^{-0.5})$	[18]
	$1 + \varepsilon$	Flow	Yes	Yes	$\tilde{O}(n^{0.5}\varepsilon^{-1})$	[6]
	$1 + \varepsilon$	Flow	No	Yes	$O(m^{o(1)}\varepsilon^{-3})$	[5]
	$1 + \varepsilon$	Flow	Yes	Yes	$O(m^{o(1)}\varepsilon^{-1})$	[9]
	$1 + \varepsilon$	Flow	No	No	$\tilde{O}(nF^*m^{-1}\varepsilon^{-1})$	Theorem 1.1
Decremental	$1 + \varepsilon$	Value	Yes	Yes	$O(m^{o(1)}\varepsilon^{-1})$	[4]
Fully Dynamic	$\tilde{O}(\log n)$	Value	No	Yes	$\tilde{O}(n^{0.667})$	[8]
	$m^{o(1)}$	Value	No	No	$m^{o(1)}$	[19]
	$m^{o(1)}$	Value	No	Yes	$m^{o(1)}$	[4]

The “Flow/Value” column indicates whether the algorithm maintains an actual flow or just the value of a flow. The update time stated is the amortized time over all updates.

on an undirected uncapacitated n -vertex graph G with high probability in total time $O(m \log(n) \alpha(n) + nF^* \log^3(n)/\varepsilon)$, where m is the number of edge insertions, F^* is the value of the max flow after m insertions, and $\alpha(n)$ is the inverse Ackermann function.

Note that in addition to dense graphs ($m = \Omega(n^2)$), our algorithm achieves polylogarithmic amortized update time even for graphs with small flow value $F^* = \tilde{O}(m/n)$, regardless of their density. An important feature of our algorithm is that it builds upon arguably simple and classic combinatorial techniques, such as residual graph sparsification and cut sparsification. Our algorithm can also be extended to allow non-simple graphs, which admit parallel edges between the same vertices.

1.1 Technical Overview

The main idea underlying our result is to dynamize the static algorithm by Karger and Levine [24] (abbrv. KL algorithm) that computes an *exact* max flow in $\tilde{O}(m + nF^*)$ time. We start by reviewing their static construction and then identify the challenges in the dynamic setting, along with our approach to overcome them. We present the algorithm and its analysis in full in Section 3.

The KL algorithm proceeds by repeatedly sampling $\rho = O(n \log n)$ edges from the residual graph, searching for an augmenting path in this sample, and then augmenting along this path before resampling again. When the search fails to find an augmenting path, the number of sampled edges is doubled to 2ρ , and this process is repeated until the sampled graph becomes as large as the original graph. Denoting the value of the maximum flow by F^* , they show that at least $F^*/2$ of the augmentations are performed on the sparsest graph with ρ edges, at least $F^*/4$ of the augmentations are performed on the graph with 2ρ edges and so on. After spending $\tilde{O}(m)$ on pre-processing to compute the sampling probabilities, the total running time thus adds up to $\approx \sum_i (F^*/2^i) \cdot 2^{i-1} \rho$, which gives the required $\tilde{O}(m + nF^*)$ bound. Crucially, the last augmentation from $F^* - 1$ to F^* is performed on a graph of size $\Omega(m)$.

This final augmentation, which requires a sample of size $\Omega(m)$, effectively invalidates any attempts to use the KL approach for obtaining an incremental *exact* max flow algorithm. To understand why, observe that after every edge insertion, the algorithm needs to check if this edge

insertion creates a new $s \rightarrow t$ augmenting path in the sampled graph. As we discuss below, this can be done in two ways, but both lead to algorithmic dead ends.

We could try to use an incremental single-source reachability data structure (e.g., the one by Italiano [23]) to detect augmenting paths in the sampled graph, but since the actual augmentation reverses the direction of edges, we would need to reinitialize the incremental data structure on the sample after each augmentation. Since F^* augmentations are performed on samples of size $\Omega(m)$, this leads to a time bound of $\Omega(mF^*)$. As a concrete example, consider an initially empty bipartite graph $G = (S \cup T, \emptyset)$, with $s \in S$ and $t \in T$. In the first phase, insert all edges between $S \times T \setminus \{t\}$, then in the second phase, insert all the remaining edges of the type $\{(v, t)\}_{v \in S}$. Each edge insertion in the second phase increases the max flow by 1, and any algorithm based on a KL-type approach with an incremental reachability data structure needs to check an $\Omega(m)$ -edge graph to perform each augmentation.

The second approach would be to use a *fully dynamic* directed $s \rightarrow t$ reachability data structure. This would solve the above problem, as an augmentation in the residual graph can be simulated using $O(n)$ directed edge deletions and insertions. However, fully dynamic $s \rightarrow t$ reachability is known to admit strong conditional lower bounds—for any $\eta > 0$, no algorithm can achieve $O(n^{1-\eta})$ worst-case update time and $O(n^{2-\eta})$ worst-case query time, assuming the OMv conjecture [21].

Interestingly, we show that the first approach can be extended to the incremental setting if we relax our algorithm to maintain a $(1 - \varepsilon)$ -approximate s - t max flow instead of an exact one. The high-level reason why is as follows: To show that at least $F^*/2^i$ of the augmentations are found in the sampled graph of size $2^{i-1}\rho$, Karger and Levine prove that when at least $F^*/2^i$ of the max flow still remains to be augmented, a sampled graph of size $2^{i-1}\rho$ suffices to preserve the existence of an augmenting path with high probability.² For our setting, if we want to maintain a $(1 - \varepsilon)$ -approximate max flow, the above claim with $i = \log(1/\varepsilon)$ tells us that it suffices to maintain a sparsifier with just $O(n \log(n)/\varepsilon)$ edges incrementally. This then removes the need to augment on an $\Omega(m)$ -edge graph, which lets us circumvent the above barrier. While incrementally maintaining the sparsifier used by Karger and Levine would be highly non-trivial, we show that we can both efficiently maintain a different sparsifier incrementally, and that the new sparsifier is powerful enough to recover the guarantees required for the KL approach to work (at the cost of a $\log n$ factor in the sparsifier size).

The challenge of incrementally maintaining the KL sparsifier is due to its fragility under edge insertions. The KL sampling depends on a parameter called the *strong connectivity* of an edge [3], and in the incremental setting, a single edge insertion could modify the strong connectivities (and thus the sampling probabilities) of *all* other edges of the graph. As a result, no existing algorithm can even just maintain the sampling probabilities, let alone maintain an actual sample from the corresponding distribution. The same issue arises for other well-known connectivity measures used for sampling, such as edge-connectivity [15] and effective resistances [31], and even maintaining approximations to them is quite involved [8, 12, 13, 16]. Instead, our key observation is to use **Nagamochi-Ibaraki (NI) indices** [26] for our sampling because of their resilience to incremental updates.

While the other connectivity parameters mentioned are unique for an edge, the NI index is determined by the forest packing constructed, and different packings give rise to different indices. This flexibility in choosing the index allows us to maintain the same value even under edge updates, which we crucially use to maintain the sampling probabilities. The resilience property we use is as follows: when inserting an edge, the NI index of all other edges remains the same, and our task is

²For us, high probability refers to the failure probability being $\leq 4/n$.

to only determine the NI index of the newly inserted edge efficiently (which then does not change in the future). Efficient maintenance of NI indices is discussed in Section 4.

While we now maintain a set of sampling probabilities using NI indices, it is unclear how they can be used to replace strong connectivity in the KL algorithm. We extend the general result of Fung et al. [15] on using a wide variety of connectivity parameters for cut sparsification on *undirected* graphs to showing that the same parameters also lead to cut sparsification on balanced directed graphs,³ which cover residual graphs that arise in our algorithm as a special case. Their results sample each edge independently, which in our setting leads to a prohibitive $\Omega(m)$ time for sampling. We adapt this to repeated sampling from a probability distribution, where we only spend $O(\log n)$ time per edge for a total $O(n \log^2(n)/\epsilon)$ edges sampled. We present this result in full detail in Section 5.

Finally, we simplify our algorithm by not maintaining an exact sample from the distribution at each timestep, but an oversample. Specifically, we obtain a sample of the correct size ($O(n \log^2(n)/\epsilon)$) right after an augmentation, but between two augmentations, we add edges directly to the sample instead of trying to maintain the distribution. This could blow up the number of edges in the sample to $\Omega(m)$ at some time steps, but since each edge is added this way to at most one sample throughout the algorithm, the total time bound of $\tilde{O}(m + nF^*/\epsilon)$ still holds. This also helps us obtain our guarantee against adversaries that know the current flow maintained by the algorithm, since the sample remains independent of the flow: any time the flow is updated, we resample the graph based on NI indices (which are independent of the flow), and between two augmentations, all inserted edges are directly added to the sample.

2 Preliminaries

2.1 Graphs and Flows

Flows. Our algorithmic results are on undirected uncapacitated graphs $G = (V, E)$. We denote $n = |V|$ and $m = |E|$. For two vertices $s, t \in V$, an s - t flow $f \in \mathbb{R}^m$ assigns a value f_e to each edge such that $\sum_{e \sim v} f_e = 0$ for all $v \neq s, t$ with $|f_e| \leq 1 \forall e$. The value of a flow f is $F(f) = \sum_{e \sim s} f_e$, and the maximum flow $f^* = \operatorname{argmax}_{s,t} F(f)$, with value $F^* = F(f^*)$. We use F without the argument f when the flow is clear from context. For any $\alpha \leq 1$, an s - t flow f is an α -approximate flow if $F \geq \alpha \cdot F^*$. If $\alpha > 1$, then the statement holds with $\alpha' = 1/\alpha$.

Model. In the incremental model, we start with an empty graph with no edges, and the graph is revealed as a sequence of edge insertions $e_1, e_2, \dots$, leading to graphs $G_1, G_2, \dots$. Our goal is to maintain, after each edge insertion e_i , an s - t flow f that is a $(1 - \epsilon)$ approximation to the max s - t flow f_i^* in the current graph G_i . We use F_i^* to denote the value of the flow f_i^* , and we use F^* to denote the final max s - t flow value after all edge insertions.

Directed Graphs. We denote by $\vec{G} = (V, \vec{E})$ a directed (multi-)graph without edge capacities. An edge of integer capacity $x \in \mathbb{Z}^+$ is represented using x parallel edges. For any subset of vertices $S \subseteq V$, $\bar{S} = V \setminus S$ is the complement of S , and $\partial_{\vec{G}}(S)$ denotes the set of edges leaving S , which is also denoted by $\vec{C}(S) = \partial_{\vec{G}}(S)$. Similarly, we use $\overleftarrow{C}(S) = \partial_{\vec{G}}(\bar{S})$ to denote the set of edges entering S . We use $\vec{C}$ and $\overleftarrow{C}$ without the argument S when the subset is clear from context. The capacity $u(\vec{C}) = |\vec{C}|$ of the cut $\vec{C}$ is the number of edges in $\vec{C}$.

Underlying Undirected Graph. For any directed (multi-)graph $\vec{G}$, we denote by G the underlying undirected (multi-)graph obtained from $\vec{G}$ by forgetting edge orientations, and $\partial_G(S)$ is the set

³In the static setting, Cen et al. [7] have shown that sampling according to strong connectivity leads to cut sparsification of balanced directed graphs. Since we work in a dynamic setting, we require the guarantee to hold for a larger class of connectivity parameters, including NI indices.

of edges leaving S , which is also denoted by $C(S) = \partial_G(S)$. For a directed cut $\vec{C} = \partial_{\vec{G}}(S)$, we use $C = \partial_G(S)$ to denote the underlying undirected cut in G , and $u(C) = |C|$ to denote the number of edges in C .

Residual Graph. For any undirected, uncapacitated graph G , the corresponding residual graph for the zero flow $f = 0^m$ is given by G_{0^m} where each edge (u, v) in G is replaced by two directed edges $\vec{uv}$ and $\vec{vu}$. For any G and a non-zero flow f , the residual graph G_f is constructed the following way: start with the graph G_{0^m} as above, and for every edge e that carries flow in the direction $u \rightarrow v$, replace the edge $\vec{uv}$ in G_{0^m} with the edge $\vec{vu}$. In that case there are two edges $\vec{vu}$ in G_f . A flow f^* is a maximum flow if and only if G_{f^*} contains no $s \rightarrow t$ paths.

2.2 Sampling Parameters

We define our importance parameter for sampling edges next. Since the residual graph has multiple edges between the same vertex pairs, we consider multi-graphs (which allow parallel edges) for the following definition.

Definition 2.1 (NI Index [26]). Given an undirected, unweighted (multi-)graph G , an ordered sequence of edge-disjoint spanning forests $T_1, T_2, \dots$ is said to be an NI forest packing of G if each T_i is a maximal spanning forest on $G \setminus \cup_{1 \leq j < i} \{T_j\}$, and $\cup_i \{T_i\}$ is a partition of all the edges of G . The NI index ℓ_e of an edge e is the index i of the forest T_i that e belongs to.

While the NI indices are the connectivity parameters that we will use for our sampling scheme due to their resilience to incremental graph updates, the following connectivity parameter is required for our directed sparsification proofs and cut-counting.

Definition 2.2 (Edge-Connectivity). Given an undirected unweighted (multi-)graph G and two vertices u and v , the edge connectivity between u and v is defined to be the value of the minimum cut (and thus also the maximum flow) between u and v in G . For an edge $e = (u, v)$, we define the edge-connectivity of e to be the edge connectivity between u and v .

As we perform *non-uniform* sampling of the edges, we maintain connectivity parameters λ_e that influence the sampling probabilities. The NI index of an edge and the edge-connectivity of an edge are examples of such connectivity parameters. We will show concentration bounds for classes of edges whose λ_e s are close to each other, which motivates the following definition.

Definition 2.3 (Connectivity Class). Given an undirected graph G and connectivity parameters λ_e for each edge, taking $\Lambda = \max_e \lceil \log \lambda_e \rceil + 1$, the connectivity classes $\mathcal{F} = \{F_i : 1 \leq i \leq \Lambda\}$ are given by

$$F_i = \{e : 2^{i-1} \leq \lambda_e \leq 2^i - 1\}.$$

For any connectivity class F_i , the graph (V, F_i) does not have the same edge-connectivity properties that $G = (V, E)$ did. An edge $e \in F_i$ could have edge-connectivity $\Omega(n)$ in G and 1 in (V, F_i) , e.g., a graph with a single (s, t) edge and $\Omega(n)$ parallel paths from s to t of length two. The following definitions from [15] are used to construct a subgraph of G that is slightly larger than F_i that certify connectivity properties of edges in F_i .

Definition 2.4 (Π -Connected Decomposition [15]). Given an undirected graph G and connectivity parameters λ_e , a decomposition $\mathcal{G} = (G_i = (V, E_i))_{1 \leq i \leq \Lambda}$ is a sequence of subgraphs of G that satisfy $F_i \subseteq E_i$ for all i with $\Lambda = \max_e \lceil \log \lambda_e \rceil + 1$. Further, for a sequence of parameters $\Pi = (\pi_i)_{1 \leq i \leq \Lambda}$, the decomposition satisfies Π -connectivity if all edges $e \in F_i$ have edge-connectivity at least π_i in G_i for all i .

Note that the same edge is allowed to appear in multiple sets E_i , and the decomposition is *not* edge-disjoint. This property is crucial in our proofs.

Definition 2.5 (γ -Overlap [15]). For any $\gamma \geq 1$, an undirected graph G and its Π -connected decomposition $\mathcal{G}$ satisfies γ -overlap if for all i and for all cuts $C = \partial_G(S)$,

$$\sum_{i=0}^{\Lambda} u(C \cap E_i) \cdot \frac{2^{i-1}}{\pi_i} \leq \gamma \cdot u(C).$$

2.3 Balance and Cut Counting

Imbalance of Directed Graphs. Preserving various graph properties by subsampling is difficult on general directed graphs. However, sampling algorithms can be successfully applied when the directed graph exhibits some nice structure, for example, they have a similar number of edges crossing a cut in both directions. Eulerian graphs have exactly the same number of edges crossing any cut in both directions, and more pertinently, the residual graph of a flow that is not a $(1 - \varepsilon)$ approximate max flow has at least an $O(\varepsilon)$ fraction of edges crossing any cut in one direction over the other.

Definition 2.6 (Balance [14]). Given a directed (multi-)graph $\vec{G} = (V, \vec{E})$, the (im)balance of a cut $\vec{C} = \partial_{\vec{G}}(S)$ is defined as

$$\beta(\vec{C}) = \frac{u(\vec{C})}{u(\overleftarrow{C})},$$

where $\overleftarrow{C} = \partial_{\vec{G}}(\bar{S})$ is the set of edges in the other direction of the cut.

Counting Cut Projections. Karger shows that there are at most $n^{2\alpha}$ distinct cuts whose size is $\leq \alpha\lambda$, where λ is the global mincut. We, like Fung et al. [15], look at larger cuts in the graph and require a tighter bound than the one given by Karger, which requires the following definition.

Definition 2.7 (Directed k -Projection). Let $\vec{G} = (V, \vec{E})$ be a directed (multi-)graph and let G be the corresponding undirected graph. The k -projection of any cut $\vec{C} = \partial(S)$ in $\vec{G}$ is $\vec{C} \cap H_k$, where

$$H_k = \left\{ \vec{e} \in \vec{E} : \text{the edge-connectivity of } e \text{ in } G \text{ is at least } k \right\}.$$

We use this definition to extend the following statement to directed graphs, which was originally shown by Fung et al. [15] for undirected graphs.

LEMMA 2.8 ([15]). *Let $G = (V, E)$ be an undirected (multi-)graph. Let $k \geq \lambda$ be any real number, where λ is the value of the global minimum cut in G . Consider all the cuts $C = \partial_G(S)$ in G such that $u(C) \leq \alpha k$, and let $\mathcal{K}_\alpha^{1k}$ be the set of all undirected k -projections of these cuts, where the undirected k -projection of a cut $C = \partial_G(S)$ is $C \cap H_k$ and H_k is the set of edges with edge-connectivity at least k . Then $|\mathcal{K}_\alpha^{1k}| \leq n^{2\alpha}$.*

2.4 Data Structures

We also need three classic data structures for our results, which we produce here.

FACT 1 (UNIONFIND [32]). *There is a data structure that supports the following operations*

— **ADD(u):** Adds $\{u\}$ to the set of sets in $O(1)$ time.

- FIND(u): Finds the representative of the set that u belongs to in amortized $O(\alpha(n))$ time.
- UNION(u, v): Replaces the sets that u and v belong to with their union in $O(1)$ time.

where n is the number of elements present at that time, and $\alpha(n)$ is the inverse Ackermann function.

FACT 2 (SINGLESOURCEREACHABILITY [23]). There is a data structure that supports the following operations on a graph $\vec{G}$ with n vertices

- INITIALIZE($\vec{G}, s$): Initializes a data structure on $\vec{G}$ with special vertex s in $O(m + n)$ time.
- INSERT($e = \vec{uv}$): Inserts the directed edge $\{e = \vec{uv}\}$ into $\vec{G}$ in amortized $O(1)$ time.
- REACHABLE(t): Returns “True” if t is reachable from s in $\vec{G}$ in amortized $O(1)$ time.
- GETPATH(t): Returns a directed $s \rightarrow t$ path in $\vec{G}$ in time proportional to the length of the path.

FACT 3 (BINARYSEARCHTREE). There is a data structure that supports the following operations

- INITIALIZE: Initializes the binary search tree data structure.
- INSERT(e, x): Inserts the key e with value x in amortized $O(\log n)$ time.
- SEARCH(x): Returns the largest key e with value $\leq x$ in amortized $O(\log n)$ time.

where n is the number of items inserted into the data structure at the time of operation.

High Probability Bounds. Finally, we say that an algorithm satisfies a guarantee with high probability if the guarantee holds with probability at least $1 - 4/n$.

Adversaries. An adversary is said to be *oblivious* if it fixes the update sequence before the algorithm picks its random bits. An adversary is *output-adaptive* if it is allowed to choose the next update/query based on the past updates and the sequence of answers given by the algorithm to past queries. An adversary is *state-adaptive* if it is output-adaptive and also allowed access to the internal state of the algorithm before choosing the next update.

3 Incremental Maximum Flow

We assume that the algorithm is started on the empty graph. Our main result in this section is the following.

THEOREM 1.1. *Given any $\epsilon \in (0, 1)$, there is an incremental randomized algorithm against output-adaptive adversaries that maintains a $(1 - \epsilon)$ -approximate maximum s - t flow f under edge insertions on an undirected uncapacitated n -vertex graph G with high probability in total time $O(m \log(n) \alpha(n) + nF^* \log^3(n)/\epsilon)$, where m is the number of edge insertions, F^* is the value of the max flow after m insertions, and $\alpha(n)$ is the inverse Ackermann function.*

We present the algorithm achieving the guarantees in Algorithm 1. The algorithm works in phases. At the start of a new phase, we sample $O(n \log^2(n)/\epsilon)$ edges from the residual graph using a specific probability distribution, which we will discuss later, and add them to the *sampled graph* H . On top of H , we run an incremental directed single-source reachability algorithm, called D , starting from s . At the beginning, and after each edge insertion, we query D if there exists an $s \rightarrow t$ path in the sample. As long as D does not return “Yes,” for each edge insertion $e = (u, v)$, we add edges $\vec{uv}$ and $\vec{vu}$ to D . When D returns “Yes,” this corresponds to an $s \rightarrow t$ path in the residual graph, and we retrieve this path from D , augment along this path in G_f , and start a new phase.

We need to argue about the time bounds and the correctness. For correctness, we show that whenever an edge insertion increases the max flow to a value larger than a $(1 - \epsilon)^{-1}$ factor over the current flow maintained by the algorithm, the existence of an $s \rightarrow t$ path in the sampled graph H is guaranteed with high probability. Since H is a subsample of the residual graph G_f ,

Algorithm 1: Algorithm for Incremental Approximate Maximum Flow on Undirected, Uncapacitated Graphs

```

1 Function INITIALIZE( $\epsilon$ ):
2    $\rho \leftarrow 5390 \cdot n \log^2(n)/\epsilon$ 
3    $f \leftarrow 0, F \leftarrow 0, G_f \leftarrow (V, \emptyset), H \leftarrow G_f$ 
4    $D \leftarrow \text{SINGLESOURCEREACHABILITY.INITIALIZE}(H, s)$ 
5    $K \leftarrow \text{INCNISAMPLE.INITIALIZE}$  ▷ Algorithm 4
6 Function INSERT( $e = (u, v)$ ):
7   Add  $\vec{uv}$  and  $\vec{vu}$  to  $H$  and  $D$ 
8    $K.\text{INSERT}(e)$ 
9   if  $D.\text{REACHABLE}(t)$  then
10     $p \leftarrow D.\text{GETPATH}(t)$ 
11    update  $f$  and  $G_f$  by augmenting along  $p$ , and update  $F \leftarrow F + 1$ 
12    reset  $H \leftarrow (V, \emptyset)$ 
13    for  $i = 1, 2, \dots, \rho$  do
14       $(u, v) \leftarrow K.\text{SAMPLE}$ 
15      for  $\vec{ab} \in \{\vec{uv}, \vec{vu}\}$  do
16        if  $\vec{ab} \in G_f$  then
17          Add  $\vec{ab}$  to  $H$ 
18    reinitialize  $D \leftarrow \text{SINGLESOURCEREACHABILITY.INITIALIZE}(H, s)$ 
19 Function QUERY(flow/value):
20   if query = flow then return  $f$  else return  $F$ 

```

this provides us with an augmenting path in G_f along which we augment the flow. With a union bound over the at most n times when this can happen (since the max flow value is at most n and each augmentation increases the flow by 1), this will imply that at the end of every timestep, the algorithm maintains a valid $(1 - \epsilon)$ -approximate s - t flow with high probability. To show this claim, we first show that an extension of Fung et al.'s high probability result [15] on independent sampling for cut sparsification on undirected graphs also extends to repeated sampling for cut sparsification on balanced directed graphs. When combined with the fact that the residual graph of a flow that is not a $(1 - \epsilon)$ -approximate max flow is $\Omega(\epsilon)$ -balanced, this shows that sampling approximately $n \log^2(n)/\epsilon$ edges based on the probability distribution discussed below will preserve directed cuts, and thus also the existence of an augmenting path in residual graphs of non- $(1 - \epsilon)$ -approximate flows. Our initial sample at the beginning of a phase already satisfies the size requirement above, and since we add further edges directly to the maintained sample, we always oversample the rate required for the augmenting path preservation guarantee.

For the time bounds, assuming that the sampling can be done in time approximately $O(\log n)$ per sample and that the reachability data structure runs in total time proportional to the number of edges in D , we will show that the algorithm runs in $\tilde{O}(m + nF^*/\epsilon)$ time. Note that the graph H is resampled at most F^* times, since each resample occurs after an augmenting path has been found in the residual graph. Each resample adds $\tilde{O}(n/\epsilon)$ edges to D , for a total of $\tilde{O}(nF^*/\epsilon)$ sampled edges inserted to D throughout the algorithm. This is in addition to the m edges which are inserted to D directly on insertion to G , which gives the claimed running time bound.

Finally, we maintain NI indices incrementally for the probability distribution, which we discuss in Section 4. We use the following definition in our proofs.

Algorithm 2: Algorithm for Directed Balanced Sparsification

Input: Directed graph $\vec{G} = (V, \vec{E})$, connectivity parameters $\{\lambda_e\}_{e \in E}$, parameters $\beta, \gamma \geq 1$, and $\varepsilon < 1$
Output: Sparsified graph $\vec{H} = (V, \vec{F}, w)$

- 1 $\rho \leftarrow \frac{128\gamma(\beta+1)\ln n}{0.38\varepsilon^2} \cdot \sum_{e \in E} \lambda_e^{-1}$
- 2 $\vec{H} \leftarrow (V, \emptyset, w)$
- 3 **for** $i \in 1, 2, \dots, \rho$ **do**
- 4 Sample an edge e from the probability distribution $\{p_e = \lambda_e^{-1} / \sum_e \lambda_e^{-1}\}$
- 5 Insert e into $\vec{H}$ with weight $w_e = (\rho \cdot p_e)^{-1}$ (additively increasing its weight if e already exists in $\vec{H}$)

Definition 3.1 (Phase). Let $t_0 = 0$, and let t_i be the timestep when the i th augmenting path is found. Let m_k be the number of edges inserted in time steps $(t_{k-1}, t_k]$. Phase k refers to the computations performed after the $(k-1)$ th augmentation finishes until the k th augmentation ends.

We also use the following two statements directly, whose proofs can be found in Section 4 and Section 5, respectively.

LEMMA 3.2. *There is an incremental data structure that deterministically maintains an NI forest packing (and thus the NI index ℓ_e of each edge) of an undirected graph under edge insertions in total time $O(m \cdot \log m \cdot \alpha(n))$ where $\alpha(n)$ is the inverse Ackermann function and m is the number of edges in the final graph. At any point of time, it also allows querying for an edge sampled from the probability distribution $\{\ell_e^{-1}/L\}_{e \in E}$, where $L = \sum_e \ell_e^{-1}$ in amortized time $O(\log n)$ for each sample.*

THEOREM 3.3. *Let $\vec{G} = (V, \vec{E})$ be any directed (multi-)graph, and G be the underlying undirected graph. Let $\{\lambda_e\}_{e \in E}$ be some connectivity parameters in G , and $\beta, \gamma \geq 1$ and $\varepsilon < 1$ be input parameters. Let $\vec{H} = (V, \vec{F}, w)$ be the random graph with $O(\gamma\beta \ln(n) \cdot \sum_{e \in E} \lambda_e^{-1}/\varepsilon^2)$ edges generated as in Algorithm 2.*

If there exists a Π -connected decomposition $\mathcal{G} = \{G_i = (V, E_i) : 1 \leq i \leq \Lambda\}$ of G that satisfies γ -overlap, then for all cuts $\vec{C} = \partial_{\vec{G}}(S)$ of balance at least β^{-1} in $\vec{G}$, it holds simultaneously with probability $\geq 1 - 4/n^2$ that

$$(1 - \varepsilon) \cdot u(\vec{C}) \leq w(\vec{C}) \leq (1 + \varepsilon) \cdot u(\vec{C}).$$

Remark. This theorem guarantees sparsification for all balanced cuts, even in potentially unbalanced graphs. When restricted to only using strong connectivity, this statement was shown by Cen et al. [7] with stronger parameters, where the quality of sparsification depends on the balance of the cut, who then used it to give a clearer proof of Karger and Levine [24] static result.

3.1 Running Time

LEMMA 3.4. *Phase k takes total time $O(m_k \log(n)\alpha(n) + n \log^3(n)/\varepsilon)$, where m_k is the number of edges inserted in Phase k and $\alpha(n)$ is the inverse Ackermann function.*

PROOF. Phase k starts with sampling $O(n \log^2(n)/\varepsilon)$ edges from the sampler, where each sample takes time $O(\log n)$ by Lemma 3.2. D has a total of $m_k + O(n \log^2(n)/\varepsilon)$ many edges added to it in its entire lifetime, which gives a total update time of $O(m_k + n \log^2(n)/\varepsilon)$. Inserting each edge into the sampler takes amortized time $O(\log m \cdot \alpha(n))$ by Lemma 3.2. Augmenting along a path takes time proportional to the length of the path, which is at most n . $\square$

LEMMA 3.5. *The algorithm takes total time $O(m \log(n)\alpha(n) + n \log^3(n)F^*/\epsilon)$ where F^* is the value of the final max flow after all edge insertions.*

PROOF. Let K be the total number of phases in the algorithm. Then $K \leq F^*$ since each phase increases the value of the flow by 1. Since the set of edges inserted in each phase is disjoint, $\sum_k m_k = m$. By Lemma 3.4, phase k takes time $O(m_k \log(n)\alpha(n) + n \log^3(n)/\epsilon)$. Together, all phases take time

$$\sum_{k=1}^K O(m_k \log(n)\alpha(n) + n \log^3(n)/\epsilon) = O(m \log(n)\alpha(n) + n \log^3(n)F^*/\epsilon). \quad \square$$

3.2 Correctness

We first show that the residual graph of a non- $(1 - \epsilon)$ -approximate max flow is $\epsilon/2$ -balanced. This then allows us to use our result on sparsifying balanced directed graphs.

LEMMA 3.6. *Let f be any s - t flow in an undirected graph G of value F , and let G_f be its corresponding residual graph. Suppose $F \leq (1 - \epsilon)F^*$, where F^* is the value of the maximum flow. Then every cut in G_f is at least $\epsilon/2$ -balanced.*

PROOF. For any cut $\vec{C} = \partial(S)$ that is not an s - t cut (or a t - s cut), the net flow out of S in G_f is zero, and thus $\vec{C}$ is 1-balanced. All t - s cuts are ≥ 1 -balanced since the residual graph has more flow in the $s \rightarrow t$ direction and thus more capacity in the $t \rightarrow s$ direction. Consider an s - t cut $\vec{C}$. Let c be the capacity of the corresponding undirected cut C in the graph G . Since F units of flow cross C in the forward direction $s \rightarrow t$, the forward capacity across C in G_f is $c - F$, and the backward capacity is thus $c + F$ since the total capacity is $2c$. Using (in order) $c \geq F$, $c \geq F^*$, and $F \leq (1 - \epsilon)F^*$, we get

$$\beta(\vec{C}) = \frac{u(\vec{C})}{u(\overleftarrow{C})} = \frac{c - F}{c + F} \geq \frac{c - F}{2c} = \frac{1}{2} \cdot \left(1 - \frac{F}{c}\right) \geq \frac{1}{2} \cdot \left(1 - \frac{F}{F^*}\right) \geq \frac{\epsilon}{2}. \quad \square$$

We then show that the sum of inverse NI indices is not too large.⁴

LEMMA 3.7. *Let $\{\ell_e\}_{e \in E}$ be a set of NI indices for an undirected (multi-)graph G with m edges. Then $\sum_e \ell_e^{-1} \leq 2n \log m$.*

PROOF. Consider the NI forest packing $T_1, T_2, \dots$ corresponding to the NI indices ℓ_e . Since each forest contains at least one edge, the total number of forests is at most m . As each edge e in T_i has $\ell_e = i$, the sum of NI indices of all edges in forest i contributes at most $\frac{n-1}{i}$ to the total sum. Thus,

$$\sum_{e \in E} \ell_e^{-1} = \sum_{1 \leq i \leq m} \sum_{e \in T_i} \ell_e^{-1} \leq \sum_{1 \leq i \leq m} \frac{n-1}{i} \leq (n-1) \cdot (2 \log m) \leq 2n \log m,$$

where $1 + \frac{1}{2} + \frac{1}{3} + \dots + \frac{1}{m}$ is the m th harmonic number, which is upper bounded by $2 \log m$. $\square$

Since we work with the residual graph, which has two copies of each edge of G , we prove the following lemma to show that an NI forest packing of G can be easily transformed into an NI forest packing of the undirected residual graph.

⁴We focus on NI indices as the connectivity parameter of choice for our algorithm due to the reasons mentioned earlier. In principle, one could use any of the parameters that work with Theorem 3.3 in this section.

LEMMA 3.8. *Let $\{\ell_e\}_{e \in E}$ be a set of NI indices for an undirected, simple graph $G = (V, E)$. Let E' be a copy of the edges in E . Then, for the multigraph $H = (V, E \cup E')$ which doubles the edges in E , assigning $h_e = 2\ell_e - 1$ and $h_{e'} = 2\ell_e$ is a set of NI indices for H .*

PROOF. If $T_1, T_2, \dots$ forms an NI forest packing of G , then $T_1, T_1, T_2, T_2, \dots$ forms an NI forest packing of H . Letting e be the edge in the first copy of T_i and e' be the edge in the second copy gives the lemma. $\square$

We use the result of [15] that shows that an NI forest packing can be used to obtain a decomposition that satisfies Π -connectivity and γ -overlap.

LEMMA 3.9 ([15]). *Let $T_1, T_2, \dots$ be an NI forest packing of an undirected graph G leading to NI indices ℓ_e , and let $\mathcal{F} = \{F_i\}$ be the corresponding connectivity classes. Then $E_1 = F_1$ and $E_i = F_{i-1} \cup F_i$ for $i \geq 2$ is a decomposition that satisfies Π -connectivity and γ -overlap for $\pi_i = 2^{i-1}$ and $\gamma = 2$, i.e., for all i and for all cuts $C = \partial_G(S)$,*

$$\sum_{i=0}^{\Lambda} u(C \cap E_i) \leq 2 \cdot u(C).$$

Using these, we will first show that sampling $\Omega(n \log^2(n)/\varepsilon)$ edges suffices to sample at least one edge across all balanced cuts.

LEMMA 3.10. *Consider a simple undirected, uncapacitated graph G with some s - t flow f of value F . Let ℓ_e be a NI index of edge e . Let $H' = (V, \vec{F})$ be a sample of $5,390 \cdot n \log^2(n)/\varepsilon$ edges chosen i.i.d. with probability $\{\ell_e^{-1}/\sum_e \ell_e^{-1}\}_{e \in G_f}$ from the residual graph G_f . Then it holds with probability $1 - 4/n^2$ that for every directed cut $\vec{C} = \partial(S)$ in G_f of balance at least $\varepsilon/2$, there is an edge $e \in \vec{C}$ that is sampled in H' . In particular, if F^* is the value of the max flow in G and $F \leq (1 - \varepsilon)F^*$, then there exists an augmenting path in the sampled graph H' with probability $1 - 4/n^2$.*

PROOF. We will first argue that sampling the mentioned number of edges and reweighting them as in Algorithm 2 gives a (weighted) cut sparsifier $\vec{H} = (V, \vec{F}, w)$ of G_f . This implies positive weight across every balanced cut in $\vec{H}$. Then, we note that the property of the presence of an edge across the cut holds regardless of the presence of weights. Since the sample H' chosen in the lemma is an unweighted version of $\vec{H}$, the lemma follows.

Since there are at most $\binom{n}{2}$ edges in a simple graph G (and thus at most n^2 edges in a residual graph), Lemma 3.7 gives that $\sum_{e \in E} \ell_e^{-1} \leq 4n \log n$. By Lemma 3.9, there is a decomposition for NI indices that satisfies Π -connectivity and γ -overlap for $\pi_i = 2^{i-1}$ and $\gamma = 2$. Thus, if one runs Algorithm 2 with $\rho = 5,390 \cdot n \log^2(n)/\varepsilon$ edges, then by Theorem 3.3 one obtains a weighted $(1 + \varepsilon')$ -sparsifier $\vec{H} = (V, \vec{F}, w)$ of the residual graph G_f where the parameters are set as $\varepsilon' = 1/2$, $\gamma = 2$, $\beta = 2/\varepsilon$, and $\sum_e \ell_e^{-1} \leq 4n \log n$. Thus, for all cuts $\vec{C}$ in G_f of balance at least $\varepsilon/2$ (Lemma 3.6), we have that the weighted sparsifier $\vec{H} = (V, \vec{F}, w)$ from Theorem 3.3 satisfies

$$w(\vec{C}) \geq \frac{u(\vec{C})}{2} > 0$$

simultaneously with probability $1 - 4/n^2$. In particular, this also means that at least one edge from S to $\bar{S}$ is sampled in $\vec{H}$ for every directed s - t cut $\partial(S)$.

Defining $H' = (V, \vec{F})$ to be the unweighted graph obtained from $\vec{H}$ by dropping the edge weights, the existence of an edge crossing S to $\bar{S}$ holds in H' as well for all cuts $\partial(S)$ with balance

at least $\varepsilon/2$ simultaneously, with probability $1 - 4/n^2$. Since the sampled graph stated in the lemma is generated exactly as in H' , the lemma follows. $\square$

Remark. While we discuss our algorithm purely for simple graphs, this is not an intrinsic barrier. We use the simpleness to apply Lemmas 3.7–3.10, where we concretely use $2 \log n$ as an upper bound on $\log m$, and for a union bound in Lemma 5.1. If we are given a different upper bound on m beforehand, then we could use this value instead in the algorithm to obtain the same guarantees.

This lemma gives us that right after a sample happens, and before any future edge insertions, the augmenting path guarantee holds with high probability.

LEMMA 3.11. *At the beginning of any phase k , if $F \leq (1 - \varepsilon)F_{t_{k-1}}^*$, then there exists an augmenting path in H with probability $1 - 4/n^2$. Here, F_i^* is the max flow at time i , and t_{k-1} is the timestep when the $(k - 1)$ -st augmentation ends. This guarantee holds against output-adaptive adversaries.*

PROOF. Recall that at the beginning of a phase, the graph H is sampled according to NI indices in the undirected graph G . We show that the conditions of Lemma 3.10 hold, even with an output-adaptive adversary. Note that the sampling probabilities are maintained deterministically and are known even to an oblivious adversary. While an output adaptive adversary knows the flow maintained by the algorithm (and thus the current residual graph as well), the sampling according to NI indices is *independent* of the flow maintained by the algorithm and its previous outputs. Lemma 3.10 shows the existence of an augmenting path for the residual graph of *any* flow, so no matter which updates were performed in the past and what information the output-adaptive adversary has. Since we assume in the lemma that we are at the beginning of a phase, there have been no edge insertions after the sampling of H . Since we sample an undirected edge and add both directions of that edge into H (lines 15–17), this is an oversample of the rate required by Lemma 3.10, which only requires adding one direction of the edge. Thus, the claim follows from the guarantees of Lemma 3.10. $\square$

To extend this guarantee to all time steps, we show that Lemma 3.10 holds regardless of the future sequence of edge insertions that happen in the incremental algorithm. Thus, an adversary who cannot see the internal randomness of the sample does not get any power by choosing future updates, since the guarantee holds regardless of the choice of future edge insertions.

LEMMA 3.12. *Consider an undirected, uncapacitated graph $G = (V, E)$ with some s - t flow f of value F . Let $H' = (V, \vec{F})$ be the graph sampled from G_f as in Lemma 3.10, and let $\mathcal{E} = \{X \subseteq (V \times V) \mid X \cap E = \emptyset\}$ be the set of all possible future updates. For any $X \in \mathcal{E}$, define F_X^* to be the max flow in the graph $G_X = (V, E \cup X)$, and define $H'_X = (V, \vec{F} \cup X)$. Then the following guarantee holds with probability at least $1 - 4/n^2$: for every $X \in \mathcal{E}$, if $F \leq (1 - \varepsilon)F_X^*$, then there exists an augmenting path in $H'_X = (V, \vec{F} \cup X)$.*

PROOF. By Lemma 3.10, we have that with probability at least $1 - 4/n^2$, the initially sampled graph H' of G_f satisfies the following: for any directed s - t cut $\partial(S)$ of balance at least $\varepsilon/2$, there is at least one edge sampled from S to $\bar{S}$. We assume in what follows that H' satisfies this guarantee, thereby conditioning on an event of probability at least $1 - 4/n^2$. Fix an arbitrary $X \in \mathcal{E}$ such that $F \leq (1 - \varepsilon)F_X^*$. We will show that the augmenting path guarantee holds in H'_X , which completes the proof.

Note that the augmenting path guarantee is equivalent to showing that for all directed s - t cuts $\partial(S)$ in H'_X , there is at least one edge from S to $\bar{S}$. Let $\vec{C} = \partial(S)$ be an arbitrary directed s - t cut. By

Lemma 3.6, $\vec{C}$ has balance at least $\varepsilon/2$ in $(G \cup X)_f$. If $\vec{C}$ had balance at least $\varepsilon/2$ even in G_f , then there is at least one edge sampled across this cut in H' by the conditioning above. If not, then the balance of $\vec{C}$ changed between G_f and $(G \cup X)_f$, which means that an edge was inserted across $\vec{C}$ in X and, thus, the undirected edge is also added to H'_X . In either case, $\vec{C}$ has an edge crossing the cut in H'_X . Since the choice of cut was arbitrary, this proves the claim. $\square$

LEMMA 3.13. *Fix a phase k . Let i be the first timestep within the phase when $F < (1 - \varepsilon)F_i^*$, where F_i^* is the value of the max flow at time i . Then there exists an augmenting path in H with probability $1 - 4/n^2$. The guarantee holds even if an output-adaptive adversary inserts edges in this phase.*

PROOF. For phase k , we use Lemma 3.12 to argue that the guarantee holds in this phase, for any sequence of future edge insertions. In particular, let i be the first timestep within the phase when $F < (1 - \varepsilon)F_i^*$. Let X be the sequence of edge insertions in this phase, the augmenting path guarantee holds in the graph H'_X with probability $1 - 4/n^2$ regardless of the identity of X . Since the graph H'_X from Lemma 3.12 is the same as the sampled graph H maintained by Algorithm 1 after inserting the edges in X in this phase, the lemma holds. $\square$

Finally, we take a simple union bound to extend the guarantee to all timesteps.

LEMMA 3.14. *Algorithm 1 maintains a $(1 - \varepsilon)$ -approximate s - t max flow at all times with probability $1 - 4/n$ against an output-adaptive adversary.*

PROOF. In any phase k , Lemma 3.13 ensures that an augmenting path exists in the sample at the first timestep i when $F < (1 - \varepsilon)F_i^*$ with probability $1 - 4/n^2$. Each time an augmenting path is found, F increases by 1. Thus, we need to invoke Lemma 3.13 for at most $F^* \leq n$ phases. Taking a union bound over all the phases, we get that the approximation guarantee holds throughout the entire insertion sequence with probability $1 - 4/n$ as required. $\square$

4 Incremental NI Indices

In this section, we show how to incrementally maintain (and sample from) NI indices ℓ_e for each edge e . The main result of this section is the following.

LEMMA 3.2. *There is an incremental data structure that deterministically maintains an NI forest packing (and thus the NI index ℓ_e of each edge) of an undirected graph under edge insertions in total time $O(m \cdot \log m \cdot \alpha(n))$ where $\alpha(n)$ is the inverse Ackermann function and m is the number of edges in the final graph. At any point of time, it also allows querying for an edge sampled from the probability distribution $\{\ell_e^{-1}/L\}_{e \in E}$ where $L = \sum_e \ell_e^{-1}$ in amortized time $O(\log n)$ for each sample.*

We maintain NI indices instead of other sampling parameters because of their stability under incremental updates. When an edge $e = (u, v)$ is inserted into a graph G , it could change other connectivity parameters (such as edge-connectivity, strong-connectivity, or the effective resistance) of every existing edge in the graph. However, the NI index of every other edge remains the same on edge insertion, which is a very desirable property for dynamic algorithms. This happens because while other graph properties are unique for an edge in a given graph, the NI index depends heavily on the particular forest packing used. It could vary wildly for the same edge, and could range from 1 up to n depending on the packing.

The incremental algorithm involves maintaining a collection of forests, where each forest is a union-find data structure on the vertices currently in that forest. On an edge insertion, we binary search across the forests to find the first one where its endpoints are disconnected, and insert

Algorithm 3: Maintaining Incremental NI Indices (INCNIINDEX)

```

1 Function INITIALIZE:
2    $k \leftarrow 0$                                 ▶ Number of forests maintained
3    $tree \leftarrow []$                             ▶ Collection of forests
4    $last\_tree[v] \leftarrow 0$  for all  $v \in V$       ▶ last forest that  $v$  belongs to
5 Function INSERT( $e = (u, v)$ ):
6    $i \leftarrow \text{FINDTREE}(u, v, \min(last\_tree(u), last\_tree(v)))$   ▶ find correct forest for  $e$ 
7   if  $i > k$  then  $k \leftarrow k + 1$ ,  $tree[i] \leftarrow \text{UNIONFIND.INITIALIZE}$   ▶ add new forest
8   for  $x \in \{u, v\}$  do
9     if  $i > last\_tree(x)$  then  $last\_tree[x] \leftarrow i$ ,  $tree[i].\text{ADD}(x)$   ▶ add  $x$  to forest  $i$ 
10     $tree[i].\text{UNION}(u, v)$   ▶ connect  $u$  and  $v$  in forest  $i$ 
11  return  $i$ 
12 Function FINDTREE( $u, v, upper\_bound$ ):
13    $L \leftarrow 0$ ,  $R \leftarrow upper\_bound$ 
14   while  $L \leq R$  do
15      $M \leftarrow \lfloor (L + R) / 2 \rfloor$   ▶ binary search for forest
16     if  $\text{ISCONNECTED}(u, v, M)$  then  $L \leftarrow M + 1$  else  $R \leftarrow M - 1$ 
17   return  $L$ 
18 Function ISCONNECTED( $u, v, i$ ):
19   if  $i = 0$  then return TRUE else return  $tree[i].\text{FIND}(u) = tree[i].\text{FIND}(v)$ 

```

the edge into that forest (and also adding the endpoints to the forest or initialize a new forest if necessary).

Due to the subtleties of implementing this, we present the data structure in full detail below. In particular, if we naïvely “initialized” a new forest by inserting *all* the vertices into the forest, then this would lead to a running time of $\Omega(n)$ per forest, which could be prohibitive. For the extreme case where n parallel edges are inserted between the same two vertices u and v , this necessitates initializing $\Omega(n)$ forests, which leads to $\Omega(n^2)$ total time for the naïve implementation, as opposed to adding a vertex to a forest only when necessary (as done below with the *last_tree* variable).

LEMMA 4.1 *Algorithm 3 maintains the NI forest packing at all times, and runs in total time $O(m \cdot \log m \cdot \alpha(n))$.*

PROOF. We first argue about correctness by induction. It is clearly true before any edge insertions. Suppose it was true until k insertions. Let $e = (u, v)$ be the $(k + 1)$ -th inserted edge. If i is the index of the tree returned by FINDTREE, then it satisfies the property that for all $j < i$, vertices u and v are connected in forest $tree[j]$, and that they are not connected in forest $tree[i]$. Thus connecting u and v in forest $tree[i]$ and setting the index of the edge to i maintains correctness after the $(k + 1)$ -th edge as well.

Next, we argue about the time taken by the algorithm. If there are m edges in the graph currently, then the algorithm maintains $\leq m$ forests. Thus, the binary search takes $O(\log m)$ iterations. In each iteration, it queries a union-find data structure if two vertices are connected, which takes $O(\alpha(n))$ amortized time. Union of two elements in, insertion of new elements into, and initialization of a union-find data structure takes constant time. Thus the algorithm runs in total time $O(m \cdot \log m \cdot \alpha(n))$. $\square$

We quickly recall an example of a binary search tree insertion sequence before we prove the next lemma. In a binary search tree, if there are three consecutive insertions of (key, value) pairs

Algorithm 4: Incremental NI Sampling (INCNISAMPLE)

```

1 Function INITIALIZE:
2    $X \leftarrow \text{INCNIINDEX.INITIALIZE}$ 
3    $Y \leftarrow \text{BINARYSEARCHTREE.INITIALIZE}$ 
4    $L \leftarrow 0$ 
5 Function INSERT( $e = (u, v)$ ):
6    $\ell_e \leftarrow X.\text{INSERT}(e)$ 
7    $Y.\text{INSERT}(e, L)$ 
8    $L \leftarrow L + \ell_e^{-1}$ 
9 Function SAMPLE:
10   $z \leftarrow \text{uniform random number in } [0, L]$ 
11   $e \leftarrow Y.\text{SEARCH}(z)$ 
12  return  $e$ 

```

$(a, 0)$, $(b, 1/3)$, and $(c, 1/2)$, then searching for any $z \in [0, 1/3)$ returns a , searching for $z \in [1/3, 1/2)$ returns b , and searching for $z \in [1/2, \infty)$ returns c .

LEMMA 4.2. *Algorithm 4 maintains a sample from the correct distribution at all times, each insertion takes time $O(\log(m)\alpha(n))$ and each sample takes time $O(\log n)$.*

PROOF. Let $e_1, \dots, e_m$ be the sequence of edge insertions performed, and let ℓ_{e_i} be their corresponding NI indices obtain from Algorithm 3. The ℓ_e values obtained from Algorithm 3 are correct by the guarantees of Lemma 4.1. Define $s_1 = 0$ and $s_i = \sum_{1 \leq j < i} \ell_{e_j}^{-1}$ be the sum of inverse NI indices of all edges inserted until e_{i-1} . We will show that at any time k , for all $i \in \{1, \dots, k\}$, searching for any $z \in [s_i, s_{i+1}] \subseteq [0, s_{k+1}]$ returns key e_i in the binary search tree. This then gives correctness of the algorithm since, after the k th insertion, the edge e_i is sampled with probability $(s_{i+1} - s_i)/L = \ell_{e_i}^{-1}/s_{k+1}$ since $L = s_{k+1}$ after k edge insertions.

The statement is clearly true after the first edge insertion. Suppose it was true until $k-1$ insertions. Since the key e_{k-1} had value s_{k-1} , searching for any $z \in [s_{k-1}, \infty)$ would have given key e_{k-1} . Let e_k be the k th inserted edge, which is inserted with value s_k . Thus, for all $z \in [s_k, \infty)$, the search returns key e_k , and for $z \in [s_{k-1}, s_k)$, the search returns key e_{k-1} . As $s_k - s_{k-1} = \ell_{e_{k-1}}^{-1}$, $s_{k+1} - s_k = \ell_{e_k}^{-1}$, and the other edges e_j for $j < k-1$ are not affected by this insertion, the claim follows.

For time bounds, note that each edge insertion into Algorithm 3 takes time $O(\log(m)\alpha(n))$, each insertion into the binary search tree with length ℓ_e^{-1} takes time $O(\log n)$, and each sample takes time $O(\log n)$ as well, given that the random number is generated in time $O(\log n)$. $\square$

Lemma 3.2 now follows from Lemma 4.1 and Lemma 4.2.

5 Balanced Sparsification

While Cen et al. [7] obtain a cut sparsification result for balanced directed graphs, their result only works with strong connectivity and cannot be used with NI indices, which we need for our algorithm. In this section, we show the following result.

THEOREM 3.3. *Let $\vec{G} = (V, \vec{E})$ be any directed (multi-)graph, and G be the underlying undirected graph. Let $\{\lambda_e\}_{e \in E}$ be some connectivity parameters in G , and $\beta, \gamma \geq 1$ and $\varepsilon < 1$ be input parameters. Let $\vec{H} = (V, \vec{F}, w)$ be the random graph with $O(\gamma\beta \ln(n) \cdot \sum_{e \in E} \lambda_e^{-1}/\varepsilon^2)$ edges generated as in Algorithm 2.*

If there exists a Π -connected decomposition $\mathcal{G} = \{G_i = (V, E_i) : 1 \leq i \leq \Lambda\}$ of G that satisfies γ -overlap, then for all cuts $\vec{C} = \partial_{\vec{G}}(S)$ of balance at least β^{-1} in $\vec{G}$, it holds simultaneously with probability $\geq 1 - 4/n^2$ that

$$(1 - \varepsilon) \cdot u(\vec{C}) \leq w(\vec{C}) \leq (1 + \varepsilon) \cdot u(\vec{C}).$$

This is an extension of the result of [15] to balanced directed graphs, with the sampling scheme changed from sampling each edge independently with probability $\approx \text{polylog}(n)/(\lambda_e \varepsilon^2)$ to sampling $\approx O(n \text{polylog}(n)/\varepsilon^2)$ edges with probability distribution proportional to λ_e^{-1} . The proof proceeds by reducing the problem to showing concentration for each connectivity class where the connectivity parameters are close to each other. Inside each connectivity class F_i , the cuts are then collected by the size of the underlying undirected cut in G_i , and concentration is shown for each collection separately. This involves combining the concentration for each cut in the collection, along with a cut-counting argument similar to the one on [15] (which is an extension of Karger's cut-counting). Finally, the concentration of each cut is shown using a standard Chernoff bound.

Concretely, we will show later that the following concentration holds for every connectivity class.

LEMMA 5.1. *It holds for all $i \in [\Lambda]$ and for all cuts $\vec{C} = \partial_{\vec{G}}(S)$ simultaneously with probability at least $1 - 4/n^2$ that*

$$\left| w(\vec{C} \cap F_i) - u(\vec{C} \cap F_i) \right| \leq \frac{\varepsilon}{2} \cdot \left(\frac{u(C \cap E_i) \cdot 2^{i-1}}{\pi_i \cdot \gamma \cdot (\beta + 1)} + u(\vec{C} \cap F_i) \right),$$

where $C = \partial_G(S)$ is the underlying undirected cut of $\vec{C}$ in G .

We first use it to prove Theorem 3.3.

PROOF OF THEOREM 3.3. We group the edges by the probability of being picked, and show concentration for each such connectivity class separately. With probability $\geq 1 - 4/n^2$, we have for all cuts $\vec{C} = \partial_{\vec{G}}(S)$ of balance β^{-1} ,

$$\begin{aligned} \left| w(\vec{C}) - u(\vec{C}) \right| &= \left| \sum_{i=0}^{\Lambda} w(\vec{C} \cap F_i) - \sum_{i=0}^{\Lambda} u(\vec{C} \cap F_i) \right| \\ &\leq \sum_{i=0}^{\Lambda} \left| w(\vec{C} \cap F_i) - u(\vec{C} \cap F_i) \right| \\ &\leq \frac{\varepsilon}{2} \cdot \left(\sum_{i=0}^{\Lambda} \frac{u(C \cap E_i) \cdot 2^{i-1}}{\pi_i \cdot \gamma \cdot (\beta + 1)} + \sum_{i=0}^{\Lambda} u(\vec{C} \cap F_i) \right) && \text{by Lemma 5.1} \\ &\leq \frac{\varepsilon}{2} \cdot \left(\frac{u(C)}{\beta + 1} + u(\vec{C}) \right) \\ &= \frac{\varepsilon}{2} \cdot \left(\frac{u(\vec{C}) + u(\overleftarrow{\vec{C}})}{\beta + 1} + u(\vec{C}) \right) \\ &\leq \varepsilon \cdot u(\vec{C}), && (\text{since } u(\overleftarrow{\vec{C}}) \leq \beta \cdot u(\vec{C})) \end{aligned}$$

where the penultimate inequality follows since

$$\sum_{i=0}^{\Lambda} \frac{u(C \cap E_i) \cdot 2^{i-1}}{\pi_i \cdot \gamma} \leq u(C)$$

by γ -overlap and

$$\sum_{i=0}^{\Lambda} u(\vec{C} \cap F_i) = u(\vec{C})$$

since $\mathcal{F} = \{F_i\}$ is a partition of E . $\square$

To show Lemma 5.1, we partition the cuts in a connectivity class based on the size of the underlying undirected cut, and show concentration in each of them separately. Formally, we will show later that

LEMMA 5.2. *Let C_{ij} be the collection of all cuts $\vec{C} = \partial_{\vec{G}}(S)$ such that for the corresponding undirected cut $C = \partial_G(S)$ in G ,*

$$\pi_i \cdot 2^j \leq u(C \cap E_i) \leq \pi_i \cdot 2^{j+1} - 1.$$

Then for all cuts $\vec{C}$ in C_{ij} simultaneously, it holds with probability $1 - 2/n^{4 \cdot 2^j}$ that

$$\left| w(\vec{C} \cap F_i) - u(\vec{C} \cap F_i) \right| \leq \frac{\varepsilon}{2} \cdot \left(\frac{u(C \cap E_i) \cdot 2^{i-1}}{\pi_i \cdot \gamma \cdot (\beta + 1)} + u(\vec{C} \cap F_i) \right).$$

We first use this to prove Lemma 5.1.

LEMMA 5.1. *It holds for all $i \in [\Lambda]$ and for all cuts $\vec{C} = \partial_{\vec{G}}(S)$ simultaneously with probability at least $1 - 4/n^2$ that*

$$\left| w(\vec{C} \cap F_i) - u(\vec{C} \cap F_i) \right| \leq \frac{\varepsilon}{2} \cdot \left(\frac{u(C \cap E_i) \cdot 2^{i-1}}{\pi_i \cdot \gamma \cdot (\beta + 1)} + u(\vec{C} \cap F_i) \right),$$

where $C = \partial_G(S)$ is the underlying undirected cut of $\vec{C}$ in G .

PROOF. We will show that the statement holds for any fixed i with probability at least $1 - 4/n^4$. Since $\mathcal{F} = \{F_i\}$ is a partition of the edges, and since there are at most n^2 edges in the graph, a union bound over all i with non-empty F_i gives the lemma.

Fix any $i \in [\Lambda]$. For all cuts $\vec{C}$ with no edge in class F_i , the claim holds with probability 1. In what follows, we only consider cuts that have at least one edge in class F_i .

We concentrate on the graph G_i , and only consider cuts that contain at least π_i edges in E_i , since every edge in F_i has connectivity at least π_i in G_i by Π -connectivity. Let C_{ij} be the collection of all cuts $\vec{C} = \partial_{\vec{G}}(S)$ such that for the corresponding undirected cut $C = \partial_G(S)$ in G ,

$$\pi_i \cdot 2^j \leq u(C \cap E_i) \leq \pi_i \cdot 2^{j+1} - 1.$$

Then by Lemma 5.2, for all $\vec{C} \in C_{ij}$, it holds with probability at least $1 - 2/n^{4 \cdot 2^j}$ that

$$\left| w(\vec{C} \cap F_i) - u(\vec{C} \cap F_i) \right| \leq \frac{\varepsilon}{2} \cdot \left(\frac{u(C \cap E_i) \cdot 2^{i-1}}{\pi_i \cdot \gamma \cdot (\beta + 1)} + u(\vec{C} \cap F_i) \right).$$

Thus, the probability that there exists a j for which there exists a cut $\vec{C} \in C_{ij}$ where the bound does not hold is at most

$$\frac{2}{n^4} \cdot \left(\frac{1}{n^{2^0}} + \frac{1}{n^{2^1}} + \dots \right) \leq \frac{4}{n^4}$$

as required. $\square$

To show Lemma 5.2, we show the bound for a single cut in Lemma 5.3 and use this with a directed cut-counting argument as in Lemma 5.4.

LEMMA 5.3. *For any single cut $\vec{C}$ in C_{ij} , it holds with probability $1 - 2/n^{8 \cdot 2^j}$ that*

$$\left| w(\vec{C} \cap F_i) - u(\vec{C} \cap F_i) \right| \leq \frac{\varepsilon}{2} \cdot \left(\frac{u(C \cap E_i) \cdot 2^{i-1}}{\pi_i \cdot \gamma \cdot (\beta + 1)} + u(\vec{C} \cap F_i) \right).$$

LEMMA 5.4. *Let $\vec{G} = (V, \vec{E})$ be a directed (multi-)graph, and G be the underlying undirected graph. Let $k \geq \lambda$ be any real number, where λ is the value of the global minimum cut in G . Consider all the cuts $\vec{C} = \partial_{\vec{G}}(S)$ in $\vec{G}$ such that $u(C) \leq \alpha k$ for the corresponding undirected cut $C = \partial_G(S)$, and let $C_\alpha^{\downarrow k}$ be the set of all directed k -projections of these cuts. Then $|C_\alpha^{\downarrow k}| \leq 2 \cdot n^{2\alpha}$.*

We first use them to show Lemma 5.2.

LEMMA 5.2. *Let C_{ij} be the collection of all cuts $\vec{C} = \partial_{\vec{G}}(S)$ such that for the corresponding undirected cut $C = \partial_G(S)$ in G ,*

$$\pi_i \cdot 2^j \leq u(C \cap E_i) \leq \pi_i \cdot 2^{j+1} - 1.$$

Then for all cuts $\vec{C}$ in C_{ij} simultaneously, it holds with probability $1 - 2/n^{4 \cdot 2^j}$ that

$$\left| w(\vec{C} \cap F_i) - u(\vec{C} \cap F_i) \right| \leq \frac{\varepsilon}{2} \cdot \left(\frac{u(C \cap E_i) \cdot 2^{i-1}}{\pi_i \cdot \gamma \cdot (\beta + 1)} + u(\vec{C} \cap F_i) \right).$$

PROOF. For any single cut $\vec{C}$ in C_{ij} , Lemma 5.3 shows that with probability at least $1 - 2/n^{8 \cdot 2^j}$,

$$\left| w(\vec{C} \cap F_i) - u(\vec{C} \cap F_i) \right| \leq \frac{\varepsilon}{2} \cdot \left(\frac{u(C \cap E_i) \cdot 2^{i-1}}{\pi_i \cdot \gamma \cdot (\beta + 1)} + u(\vec{C} \cap F_i) \right).$$

At this point, we would like to take a union bound over all the cuts in C_{ij} . Ideally, we would like the number of cuts in this set to grow as $n^{c \cdot 2^j}$. However, the number of directed cuts $\vec{C}$ that arise might be much larger than that. We can circumvent this, since the inequality only depends on $\vec{C} \cap F_i$, and not $\vec{C}$. We thus focus on working with the distinct subsets $\vec{C} \cap F_i$ that arise, and use these to prove our claim.

We first show that the number of *distinct directed cuts* $\vec{C} \cap F_i$ that are encountered is still $O(n^{c \cdot 2^j})$ using Lemma 5.4. In particular, we apply Lemma 5.4 on the graph $G_i = (V, E_i)$, setting $k = \pi_i$ and $\alpha k = \pi_i \cdot 2^{j+1} - 1$. The total number of distinct subsets $\vec{C} \cap H_k$ that arise from cuts in C_{ij} is at most

$$2n^{2\alpha} < 2n^{4 \cdot 2^j},$$

where H_k is the subgraph of G_i of all edges with edge-connectivity at least π_i . Since each edge in F_i has edge-connectivity at least π_i by Π -connectivity, we have that $F_i \subseteq H_k$. Every non-empty subset of the form $\vec{C} \cap F_i$ has a corresponding subset of the form $\vec{C} \cap H_k$ that is counted above, and for every $\vec{C} \cap H_k$ counted above, there is at most one non-empty subset of the form $\vec{C} \cap F_i$, which proves the claim.

To simplify notation, we use the following definition for the next claim. Call a subset of directed edges $\vec{R} \subseteq \vec{E}$ to be *bad* for a directed cut $\vec{C} \subseteq \vec{E}$ if $\vec{C} \cap F_i = \vec{R}$ and

$$\left| w(\vec{R}) - u(\vec{R}) \right| > \frac{\varepsilon}{2} \cdot \left(\frac{u(C \cap E_i) \cdot 2^{i-1}}{\pi_i \cdot \gamma \cdot (\beta + 1)} + u(\vec{R}) \right).$$

Note that if $\vec{R}$ is bad for any particular directed cut $\vec{C}$, then it is also bad for

$$\vec{D} = \underset{\vec{C}': (\vec{C}' \cap F_i = \vec{R}) \wedge (\vec{C}' \in C_{ij})}{\operatorname{argmin}} u(C' \cap E_i)$$

since

$$\left| w(\vec{R}) - u(\vec{R}) \right| > \frac{\varepsilon}{2} \cdot \left(\frac{u(C \cap E_i) \cdot 2^{i-1}}{\pi_i \cdot \gamma \cdot (\beta + 1)} + u(\vec{R}) \right) \geq \frac{\varepsilon}{2} \cdot \left(\frac{u(D \cap E_i) \cdot 2^{i-1}}{\pi_i \cdot \gamma \cdot (\beta + 1)} + u(\vec{R}) \right),$$

where the first inequality is because $\vec{R}$ is bad for $\vec{C}$, and the second inequality is by choice of $\vec{D}$.

Thus, the probability we need to bound is

$$\begin{aligned} & \Pr[\exists \text{ a cut } \vec{C} \in C_{ij} \text{ such that } \vec{C} \cap F_i \text{ is bad for } \vec{C}] \\ &= \Pr[\exists \text{ an } \vec{R} \subseteq \vec{E} \text{ and a cut } \vec{C} \in C_{ij} \text{ such that } \vec{R} \text{ is bad for } \vec{C}] \\ &\leq \sum_{\vec{R} \subseteq \vec{E}} \Pr[\exists \text{ a cut } \vec{C} \in C_{ij} \text{ such that } \vec{R} \text{ is bad for } \vec{C}] \\ &= \sum_{\vec{R} \subseteq \vec{E}} \Pr \left[\vec{R} \text{ is bad for } \vec{D} = \underset{\vec{C}': (\vec{C}' \cap F_i = \vec{R}) \wedge (\vec{C}' \in C_{ij})}{\operatorname{argmin}} u(C \cap E_i) \right] \\ &\leq 2n^{4 \cdot 2^j} \cdot \frac{2}{n^{8 \cdot 2^j}} \leq \frac{4}{n^{4 \cdot 2^j}}, \end{aligned}$$

where the penultimate inequality holds because the number of distinct $\vec{R}$ s that can be of the form $\vec{C} \cap F_i$ is at most $2n^{4 \cdot 2^j}$, and for each $\vec{R}$, the probability that $\vec{R}$ is bad for $\vec{D}$ is at most $2/n^{8 \cdot 2^j}$ by Lemma 5.3. $\square$

We will prove a Chernoff bound of the following form in Appendix A.

LEMMA 5.5. *Let $X_1, \dots, X_n$ be n independent random variables that take values in the range $[0, M]$ for some $M > 0$. Then, for any $\varepsilon \in (0, 1)$ and $N \geq E[\sum_i X_i]$, we have*

$$\Pr \left[\left| \sum_{i=1}^n X_i - E \left[\sum_{i=1}^n X_i \right] \right| > \varepsilon N \right] \leq 2 \exp(-0.38 \varepsilon^2 N / M).$$

Using this, we prove Lemma 5.3.

LEMMA 5.3. *For any single cut $\vec{C}$ in C_{ij} , it holds with probability $1 - 2/n^{8 \cdot 2^j}$ that*

$$\left| w(\vec{C} \cap F_i) - u(\vec{C} \cap F_i) \right| \leq \frac{\varepsilon}{2} \cdot \left(\frac{u(C \cap E_i) \cdot 2^{i-1}}{\pi_i \cdot \gamma \cdot (\beta + 1)} + u(\vec{C} \cap F_i) \right).$$

PROOF. For each round $j \in \{1, 2, \dots, \rho\}$, let X_j be the following random variable

$$X_j = \begin{cases} w_e = (\rho \cdot p_e)^{-1} & \text{if an edge } e \in \vec{C} \cap F_i \text{ is chosen in round } j \\ 0 & \text{if an edge } e \notin \vec{C} \cap F_i \text{ is chosen in round } j \end{cases}.$$

Defining $X = \sum_{j=1}^{\rho} X_j$, we have that $X = w(\vec{C} \cap F_i)$ and $E[X] = u(\vec{C} \cap F_i)$. Since $\lambda_e < 2^{i+1}$ for any edge $e \in F_i$, we have that

$$p_e \geq \frac{\lambda_e^{-1}}{\sum_{e \in E} \lambda_e^{-1}} \geq \frac{1}{2^{i+1} \cdot \sum_{e \in E} \lambda_e^{-1}}.$$

This gives an upper bound on the weight as

$$\max_j |X_j| = \max_{e \in \vec{C} \cap F_i} w_e = \max_{e \in \vec{C} \cap F_i} \frac{1}{\rho \cdot p_e} \leq \frac{2^{i+1} \cdot \sum_{e \in E} \lambda_e^{-1}}{\rho}.$$

Recall that in the algorithm, we set

$$\rho = \frac{128\gamma(\beta+1) \ln n}{0.38\epsilon^2} \cdot \sum_{e \in E} \lambda_e^{-1}.$$

Substituting this into the upper bound on the weight, we apply Lemma 5.5 to the random variables X_i using the following explicit values of M and N .

$$M = \left(\frac{128\gamma(\beta+1) \ln n}{0.38 \cdot 2^{i+1}\epsilon^2} \right)^{-1} \quad \text{and} \quad N = u(\vec{C} \cap F_i) + \frac{u(C \cap E_i) \cdot 2^{i-1}}{\pi_i \cdot \gamma \cdot (\beta+1)}.$$

It trivially holds that

$$N = u(\vec{C} \cap F_i) + \frac{u(C \cap E_i) \cdot 2^{i-1}}{\pi_i \cdot \gamma \cdot (\beta+1)} \geq u(\vec{C} \cap F_i) = E[X].$$

Thus,

$$\begin{aligned} & \Pr \left[\left| w(\vec{C} \cap F_i) - u(\vec{C} \cap F_i) \right| > \frac{\epsilon}{2} \cdot \left(\frac{u(C \cap E_i) \cdot 2^{i-1}}{\pi_i \cdot \gamma \cdot (\beta+1)} + u(\vec{C} \cap F_i) \right) \right] \\ & \leq 2 \exp \left(-0.38 \cdot \left(\frac{\epsilon}{2} \right)^2 \cdot \left(\frac{128\gamma(\beta+1) \ln n}{0.38 \cdot 2^{i+1}\epsilon^2} \right) \cdot \left(\frac{u(C \cap E_i) \cdot 2^{i-1}}{\pi_i \cdot \gamma \cdot (\beta+1)} + u(\vec{C} \cap F_i) \right) \right) \\ & \leq 2 \exp \left(-0.38 \cdot \left(\frac{\epsilon}{2} \right)^2 \cdot \left(\frac{128\gamma(\beta+1) \ln n}{0.38 \cdot 2^{i+1}\epsilon^2} \right) \cdot \left(\frac{u(C \cap E_i) \cdot 2^{i-1}}{\pi_i \cdot \gamma \cdot (\beta+1)} \right) \right) \\ & \leq 2 \exp \left(-8 \cdot \frac{u(C \cap E_i) \cdot \ln n}{\pi_i} \right) \\ & \leq 2 \exp(-8 \cdot 2^j \cdot \ln n) \quad (\text{since } \vec{C} \in C_{ij} \implies u(C \cap E_i) \geq \pi_i \cdot 2^j) \end{aligned}$$

as required. $\square$

Finally, we use the undirected projection result of [15] to prove Lemma 5.4.

LEMMA 5.4. Let $\vec{G} = (V, \vec{E})$ be a directed (multi-)graph, and G be the underlying undirected graph. Let $k \geq \lambda$ be any real number, where λ is the value of the global minimum cut in G . Consider all the cuts $\vec{C} = \partial_{\vec{G}}(S)$ in $\vec{G}$ such that $u(C) \leq ak$ for the corresponding undirected cut $C = \partial_G(S)$, and let $C_{\alpha}^{\downarrow k}$ be the set of all directed k -projections of these cuts. Then $|C_{\alpha}^{\downarrow k}| \leq 2 \cdot n^{2\alpha}$.

PROOF. We first apply Lemma 2.8 on the undirected graph G . Consider the set of all undirected cuts $C = \delta(S, \bar{S})$ in G such that $u(C) \leq \alpha k$. Lemma 2.8 gives that there are at most $n^{2\alpha}$ undirected k -projections of such cuts, i.e., $|\mathcal{K}_\alpha^{\downarrow k}| \leq n^{2\alpha}$. We will show that $|C_\alpha^{\downarrow k}| \leq 2 \cdot |\mathcal{K}_\alpha^{\downarrow k}|$, which gives the lemma.

We define a function from the first set to the second set where each image has a pre-image set of size at most two. Recall that H_k is the set of edges in G of edge-connectivity at least k . We map each directed projection $\vec{C} \cap H_k$ in $C_\alpha^{\downarrow k}$ to the corresponding undirected projection $C \cap H_k$. The latter belongs to $\mathcal{K}_\alpha^{\downarrow k}$ since $u(C) \leq \alpha k$ by definition. For each $C \cap H_k \in \mathcal{K}_\alpha^{\downarrow k}$, there are at most two directed projections that map to it since each undirected cut can be directed as either $(S, \bar{S})$ or $(\bar{S}, S)$. $\square$

References

- [1] Ravindra K. Ahuja, Thomas L. Magnanti, James B. Orlin, and M. R. Reddy. 1995. Applications of network optimization. In *Handbooks in Operations Research and Management Science*. M. O. Ball, T. L. Magnanti, C. L. Monma, and G. L. Nemhauser (Eds.), Vol. 7, Elsevier, 1–83. Retrieved from <https://www.sciencedirect.com/science/article/abs/pii/S0927050705801185>
- [2] Sanjeev Arora, Elad Hazan, and Satyen Kale. 2012. The multiplicative weights update method: A meta-algorithm and applications. *Theory Comput.* 8, 1 (2012), 121–164. DOI: <https://doi.org/10.4086/TOC.2012.V008A006>
- [3] András A. Benczúr and David R. Karger. 2015. Randomized approximation schemes for cuts and flows in capacitated graphs. *SIAM J. Comput.* 44, 2 (2015), 290–319. DOI: <https://doi.org/10.1137/070705970>
- [4] Jan van den Brand, Li Chen, Rasmus Kyng, Yang P. Liu, Simon Meierhans, Maximilian Probst Gutenberg, and Sushant Sachdeva. 2024. Almost-linear time algorithms for decremental graphs: Min-cost flow and more via duality. In *Proceedings of the 2024 IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 2010–2032. DOI: <https://doi.org/10.1109/FOCS61266.2024.00120>
- [5] Jan van den Brand, Li Chen, Rasmus Kyng, Yang P. Liu, Richard Peng, Maximilian Probst Gutenberg, Sushant Sachdeva, and Aaron Sidford. 2024. Incremental approximate maximum flow on undirected graphs in subpolynomial update time. In *Proceedings of the Symposium on Discrete Algorithms (SODA)*. SIAM, 2980–2998. DOI: <https://doi.org/10.1137/1.9781611977912.106>
- [6] Jan van den Brand, Yang P. Liu, and Aaron Sidford. 2023. Dynamic maxflow via dynamic interior point methods. In *Proceedings of the Symposium on Theory of Computing (STOC)*. ACM, New York, NY, 1215–1228. DOI: <https://doi.org/10.1145/3564246.3585135>
- [7] Ruoxu Cen, Yu Cheng, Debmalaya Panigrahi, and Kevin Sun. 2021. Sparsification of directed graphs via cut balance. In *Proceedings of the International Colloquium on Automata, Languages, and Programming (ICALP)*. Schloss Dagstuhl—Leibniz-Zentrum für Informatik, 45:1–45:21. DOI: <https://doi.org/10.4230/LIPICS.ICALP.2021.45>
- [8] Li Chen, Gramoz Goranci, Monika Henzinger, Richard Peng, and Thatchaphol Saranurak. 2020. Fast dynamic cuts, distances and effective resistances via vertex sparsifiers. In *Proceedings of the 2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 1135–1146. DOI: <https://doi.org/10.1109/FOCS46700.2020.00109>
- [9] Li Chen, Rasmus Kyng, Yang P. Liu, Simon Meierhans, and Maximilian Probst Gutenberg. 2024. Almost-linear time algorithms for incremental graphs: Cycle detection, SCCs, s-t shortest path, and minimum-cost flow. In *Proceedings of the Symposium on Theory of Computing (STOC)*. ACM, New York, NY, 1165–1173. DOI: <https://doi.org/10.1145/3618260.3649745>
- [10] Li Chen, Rasmus Kyng, Yang P. Liu, Richard Peng, Maximilian Probst Gutenberg, and Sushant Sachdeva. 2022. Maximum flow and minimum-cost flow in almost-linear time. In *Proceedings of the 2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 612–623. DOI: <https://doi.org/10.1109/FOCS54457.2022.00064>
- [11] Søren Dahlgaard. 2016. On the hardness of partially dynamic graph problems and connections to diameter. In *Proceedings of the 43rd International Colloquium on Automata, Languages, and Programming (ICALP '16)*. Schloss Dagstuhl—Leibniz-Zentrum für Informatik, Article 48, 1–14. DOI: <https://doi.org/10.4230/LIPICS.ICALP.2016.48>
- [12] David Durfee, Yu Gao, Gramoz Goranci, and Richard Peng. 2018. Fully dynamic effective resistances. arXiv:1804.04038. Retrieved from <http://arxiv.org/abs/1804.04038>
- [13] David Durfee, Yu Gao, Gramoz Goranci, and Richard Peng. 2019. Fully dynamic spectral vertex sparsifiers and applications. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing (STOC '19)*, 914–925. DOI: <https://doi.org/10.1145/3313276.3316379>
- [14] Alina Ene, Gary L. Miller, Jakub Pachocki, and Aaron Sidford. 2016. Routing under balance. In *Proceedings of the 48th annual ACM symposium on Theory of Computing (STOC '16)*, 598–611. DOI: <https://doi.org/10.1145/2897518.2897654>

- [15] Wai Shing Fung, Ramesh Hariharan, Nicholas J. A. Harvey, and Debmalya Panigrahi. 2019. A general framework for graph sparsification. *SIAM J. Comput.* 48, 4 (2019), 1196–1223. DOI: <https://doi.org/10.1137/16M1091666>
- [16] Yu Gao, Yang P. Liu, and Richard Peng. 2021. Fully dynamic electrical flows: Sparse maxflow faster than Goldberg-Rao. In *Proceedings of the 2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, 516–527. DOI: <https://doi.org/10.1109/FOCS52979.2021.00058>
- [17] Andrew V. Goldberg and Satish Rao. 1998. Beyond the flow decomposition barrier. *J. ACM* 45, 5 (1998), 783–797. DOI: <https://doi.org/10.1145/290179.290181>
- [18] Gramoz Goranci and Monika Henzinger. 2023. Efficient data structures for incremental exact and approximate maximum flow. In *Proceedings of the 50th International Colloquium on Automata, Languages, and Programming (ICALP '23)*. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Article 69, 1–14. DOI: <https://doi.org/10.4230/LIPICS.ICALP.2023.69>
- [19] Gramoz Goranci, Harald Räcke, Thatchaphol Saranurak, and Zihan Tan. 2021. The expander hierarchy and its applications to dynamic graph algorithms. In *Proceedings of the Symposium on Discrete Algorithms (SODA)*. SIAM, 2212–2228. DOI: <https://doi.org/10.1137/1.9781611976465.132>
- [20] Manoj Gupta and Shahbaz Khan. 2021. Simple dynamic algorithms for maximal independent set, maximum flow and maximum matching. In *Proceedings of the Symposium on Simplicity in Algorithms (SOSA)*. SIAM, 86–91. DOI: <https://doi.org/10.1137/1.9781611976496.10>
- [21] Monika Henzinger, Sebastian Krinninger, Danupon Nanongkai, and Thatchaphol Saranurak. 2015. Unifying and strengthening hardness for dynamic problems via the online matrix-vector multiplication conjecture. In *Proceedings of the Symposium on Theory of Computing (STOC)*. ACM, New York, NY, 21–30. DOI: <https://doi.org/10.1145/2746539.2746609>
- [22] Monika Rauch Henzinger. 1997. A static 2-approximation algorithm for vertex connectivity and incremental approximation algorithms for edge and vertex connectivity. *J. Algorithms* 24, 1 (1997), 194–220. DOI: <https://doi.org/10.1006/JAGM.1997.0855>
- [23] Giuseppe F. Italiano. 1986. Amortized efficiency of a path retrieval data structure. *Theor. Comput. Sci* 48, 3 (1986), 273–281. DOI: [https://doi.org/10.1016/0304-3975\(86\)90098-8](https://doi.org/10.1016/0304-3975(86)90098-8)
- [24] David R. Karger and Matthew S. Levine. 2015. Fast augmenting paths by random sampling from residual graphs. *SIAM J. Comput.* 44, 2 (2015), 320–339. DOI: <https://doi.org/10.1137/070705994>
- [25] Jonathan A. Kelner, Yin Tat Lee, Lorenzo Orecchia, and Aaron Sidford. 2014. An almost-linear-time algorithm for approximate max flow in undirected graphs, and its multicommodity generalizations. In *Proceedings of the Symposium on Discrete Algorithms (SODA)*. SIAM, 217–226. DOI: <https://doi.org/10.1137/1.9781611973402.16>
- [26] Hiroshi Nagamochi and Toshihide Ibaraki. 1992. A linear-time algorithm for finding a sparse k-connected spanning subgraph of a k-connected graph. *Algorithmica* 7, 5–6 (1992), 583–596. DOI: <https://doi.org/10.1007/BF01758778>
- [27] Richard Peng. 2016. Approximate undirected maximum flows in $O(\text{mpolylog}(n))$ time. In *Proceedings of the Symposium on Discrete Algorithms (SODA)*. SIAM, 1862–1867. DOI: <https://doi.org/10.1137/1.9781611974331.CH130>
- [28] Jonah Sherman. 2009. Breaking the multicommodity flow barrier for $O(\text{vlog } n)$ -approximations to sparsest cut. In *Proceedings of the 2009 50th Annual IEEE Symposium on Foundations of Computer Science (FOCS '09)*. IEEE Computer Society, 363–372. DOI: <https://doi.org/10.1109/FOCS.2009.66>
- [29] Jonah Sherman. 2013. Nearly maximum flows in nearly linear time. In *Proceedings of the 2013 IEEE 54th Annual Symposium on Foundations of Computer Science*. IEEE Computer Society, 263–269. DOI: <https://doi.org/10.1109/FOCS.2013.36>
- [30] Jonah Sherman. 2017. Area-convexity, ∞ regularization, and undirected multicommodity flow. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing (STOC '17)*, 452–460. DOI: <https://doi.org/10.1145/3055399.3055501>
- [31] Daniel A. Spielman and Nikhil Srivastava. 2011. Graph sparsification by effective resistances. *SIAM J. Comput.* 40, 6 (2011), 1913–1926. DOI: <https://doi.org/10.1137/080734029>
- [32] Robert Endre Tarjan and Jan van Leeuwen. 1984. Worst-case analysis of set union algorithms. *J. ACM* 31, 2 (1984), 245–281. DOI: <https://doi.org/10.1145/62.2160>
- [33] Shang-Hua Teng. 2010. The Laplacian paradigm: Emerging algorithms for massive graphs. In *Theory and Applications of Models of Computation*. J. Kratochvíl, A. Li, J. Fiala, and P. Kolman (Eds.), Springer, 2–14. DOI: https://doi.org/10.1007/978-3-642-13562-0_2

Appendix

A Chernoff Bounds

We use the following Chernoff bound from [15] to obtain Lemma 5.5.

LEMMA A.1 ([15]). *Let $Y_1, \dots, Y_n$ be n independent random variables each taking values in $[0, 1]$. Let $\mu = E[\sum_i Y_i]$. Then for all $\delta > 0$, we have*

$$\Pr \left[\left| \sum_{i=1}^n Y_i - \mu \right| > \delta \mu \right] \leq 2 \exp(-0.38 \min\{1, \delta\} \delta \mu).$$

Recall the bound we want to prove.

LEMMA 5.5. *Let $X_1, \dots, X_n$ be n independent random variables that take values in the range $[0, M]$ for some $M > 0$. Then, for any $\varepsilon \in (0, 1)$ and $N \geq E[\sum_i X_i]$, we have*

$$\Pr \left[\left| \sum_{i=1}^n X_i - E \left[\sum_{i=1}^n X_i \right] \right| > \varepsilon N \right] \leq 2 \exp(-0.38 \varepsilon^2 N / M).$$

PROOF. Define $Y_i = M^{-1} \cdot X_i$. Then $Y_i \in [0, 1]$ and $\mu = E[\sum_i Y_i] = M^{-1} \cdot E[\sum_i X_i]$. Applying Lemma A.1 for $\delta = \frac{\varepsilon N}{E[\sum_i X_i]}$, we get that

$$\Pr \left[\left| \sum_{i=1}^n X_i - E \left[\sum_{i=1}^n X_i \right] \right| > \varepsilon N \right] \leq 2 \exp(-0.38 \min\{1, \delta\} \delta \mu).$$

If $\delta \leq 1$, then using the fact that $N \geq E[\sum_i X_i]$, the bound becomes

$$2 \exp \left(-0.38 \cdot \left(\frac{\varepsilon N}{E[\sum_i X_i]} \right)^2 \cdot M^{-1} \cdot E \left[\sum_i X_i \right] \right) \leq 2 \exp(-0.38 \varepsilon^2 N / M),$$

and if $\delta > 1$, then the bound becomes

$$2 \exp \left(-0.38 \cdot \left(\frac{\varepsilon N}{E[\sum_i X_i]} \right) \cdot M^{-1} \cdot E \left[\sum_i X_i \right] \right) \leq 2 \exp(-0.38 \varepsilon N / M),$$

where the final bound follows since $\varepsilon^2 < \varepsilon$ for $\varepsilon \in [0, 1]$. □

Received 14 October 2025; revised 8 May 2026; accepted 11 May 2026