Incremental Approximate Maximum Flow on Undirected Graphs in Subpolynomial Update Time

Jan van den Brand Georgia Tech vdbrand@gatech.edu Li Chen^{*} Georgia Tech lichen@gatech.edu Rasmus Kyng[†] ETH Zurich kyng@inf.ethz.ch

Maximilian Probst Gutenberg[†]

ETH Zurich

maxprobst@ethz.ch

Yang P. Liu[‡] Stanford University yangpliu@stanford.edu Richard Peng[§] University of Waterloo y5peng@uwaterloo.ca

Sushant Sachdeva[¶] University of Toronto sachdeva@cs.toronto.edu Aaron Sidford[∥] Stanford University sidford@stanford.edu

November 7, 2023

Abstract

We provide an algorithm which, with high probability, maintains a $(1 - \epsilon)$ -approximate maximum flow on an undirected graph undergoing *m*-edge additions in amortized $m^{o(1)}\epsilon^{-3}$ time per update. To obtain this result, we provide a more general algorithm that solves what we call the *incremental*, *thresholded*, *p*-norm flow problem that asks to determine the first edgeinsertion in an undirected graph that causes the minimum ℓ_p -norm flow to decrease below a given threshold in value. Since we solve this thresholded problem, our data structure succeeds against an adaptive adversary that can only see the data structure's output. Furthermore, since our algorithm holds for p = 2, we obtain improved algorithms for dynamically maintaining the effective resistance between a pair of vertices in an undirected graph undergoing edge insertions.

Our algorithm builds upon previous dynamic algorithms for approximately solving the minimum-ratio cycle problem that underlie previous advances on the maximum flow problem [Chen-Kyng-Liu-Peng-Probst Gutenberg-Sachdeva, FOCS '22] as well as recent dynamic maximum flow algorithms [v.d.Brand-Liu-Sidford, STOC '23]. Instead of using interior point methods, which were a key component of these recent advances, our algorithm uses an optimization method based on ℓ_p -norm iterative refinement and the multiplicative weight update method. This ensures a monotonicity property in the minimum-ratio cycle subproblems that allows us to apply known data structures and bypass issues arising from adaptive queries.

^{*}Li Chen was supported by NSF Grant CCF-2106444.

[†]The research leading to these results has received funding from the grant "Algorithms and complexity for highaccuracy flows and convex optimization" (no. 200021 204787) of the Swiss National Science Foundation.

[‡]Yang P. Liu was supported by NSF CAREER Award CCF-1844855, NSF Grant CCF-1955039, and a Google Research Fellowship.

[§]Richard Peng was partially supported by the Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery Grant RGPIN-2022-03207

[¶]Sushant Sachdeva was supported by an NSERC Discovery Grant RGPIN-2018-06398, an Ontario Early Researcher Award (ERA) ER21-16-284, and a Sloan Research Fellowship.

^{||}Aaron Sidford was supported by a Microsoft Research Faculty Fellowship, NSF CAREER Award CCF-1844855, NSF Grant CCF-1955039, a PayPal research award, and a Sloan Research Fellowship.

Contents

1	Introduction 1			
	1.1 Results	3		
	1.2 Additional Related Work	5		
	1.3 Paper Organization	5		
2	Preliminaries	6		
3	Technical Overview	6		
	3.1 Previous Tools and Obstacles	7		
	3.2 Adaptive Adversaries and Monotonic Dynamic Min-Ratio Cycle.	8		
	3.3 From IPMs to Iterative Refinement and MWU	8		
	3.4 Putting it Altogether	9		
4	Incremental <i>p</i> -Norm Iterative Refinement	9		
5	Incremental Multiplicative Weight Updates	12		
6	Dynamic Min-Ratio Cycle Data Structures 10			
R	References 17			

1 Introduction

The design and analysis of dynamic graph algorithms is a rich, well-studied research area. Researchers have studied dynamic variations of fundamental graph problems such as minimum spanning tree, shortest path, and bipartite matching. Recent progress on dynamic matching has been widely celebrated [Beh23; BKSW23], and dynamic graph algorithms play a key role in recent advances in static graph algorithms, yielding the first nearly-linear time algorithms for bipartite matching and maximum flow (maxflow) in dense graphs [BLNPSSSW20; BLLSSSW21] and almostlinear time algorithm for maxflow [CKLPPS22]. (See Section 1.2, for additional related work.)

Despite the many successes of dynamic graph algorithms, dynamic maxflow, another class of core problems, has seen relatively little progress. This is perhaps due in part to strong conditional hardness results. In the incremental and decremental settings, where, respectively, edges are only added or only removed, exactly maintaining the *s*-*t* flow value requires $\Omega(n)$ time per update, even in directed unit capacity graphs, assuming the Online-Matrix-vector (OMv)-conjecture. This was shown by [Dah16], building on earlier work by [HKNS15] that introduced the OMv-conjecture and showed a $O(\sqrt{m})$ time per update lower bound for the same problems.

Moving to approximate solutions opens the possibility of faster algorithms. Recently, [GH23] gave an $(1 - \epsilon)$ -approximation algorithm for incremental unit capacity maxflow with amortized update time $\hat{O}(\epsilon^{-1/2}\sqrt{m})^1$, an improvement for sparse graphs in the low accuracy regime. Additionally, [BLS23] recently gave an $(1 - \epsilon)$ -approximation algorithm for incremental maxflow and minimum-cost flow² with amortized update time $\hat{O}(\epsilon^{-1}\sqrt{n})$ by dynamizing and building upon the recent almost-linear time algorithm for maxflow of [CKLPPS22]. For constant ϵ , this runtime is faster than the conditional lower bound for exact incremental maxflow.

In this work, we ask whether we can build upon this recent progress and give a subpolynomial amortized update time algorithm for the dynamic maxflow problem? We answer this in the affirmative by developing such an algorithm for approximate incremental undirected maxflow.

Dynamizing Static Maxflow. To obtain our result, we build upon the recent advance of [BLS23] and the work that underlies it, as well as distinct lines of research related to approximate undirected maxflow. [BLS23] essentially dynamizes the almost linear-time algorithm for maxflow [CKLPPS22]. [CKLPPS22] came at the end of a long line of research that focused on solving flow problems by combining graph theoretic tools with interior point methods (IPMs), a class of continuous optimization methods which obtain high-accuracy solutions to convex optimization problems [DS08; Mąd13; LS14; Mąd16; CMSV17; AMV20; BLNPSSSW20; KLS20; LS20; AMV21; BLLSSSW21; GLP21; BGJLLPS22; CKLPPS22; DGGLPSY22]. The [CKLPPS22] IPM relies on solving an ℓ_1 flow update problem known as "(undirected) min-ratio cycle."³ This problem is solved $m^{1+o(1)}$ times using a data structure with amortized $m^{o(1)}$ time per update. [BLS23], showed how to dynamize this IPM, but obtained an $n^{1/2+o(1)}$ amortized time per update due to the cost of adapting the min-ratio cycle data structures to this setting.

There are two key ideas in [BLS23]. The first is that the ℓ_1 -IPM of [CKLPPS22] can be naturally extended to the incremental setting. In particular, they showed how the IPM can be used to solve a threshold variant of the incremental maximum flow problem, i.e., detecting the first update that causes the maximum flow value to increase above a threshold. However, this extension creates a challenge: The [CKLPPS22] data structure for min-ratio cycle does *not* work against an adaptive adversary. Instead, [CKLPPS22] crucially leverages stability properties of their ℓ_1 IPM

¹In this paper, we use $\widehat{O}(\cdot)$ to suppress subpolynomial $n^{o(1)}$ factors.

²In both cases, this is assuming polynomially bounded capacities; we make this assumption throughout the paper. ³In this paper we refer to this problem as *min-ratio cycle*, omitting the term "undirected." Elsewhere in the literature, but never in this paper, *min-ratio cycle* may refer to a variant with edge-direction constraints.

that ultimately determines the update problems, and uses these stability properties to guarantee that while their data structure may occasionally fail, this will occur infrequently. In the incremental setting, [BLS23] cannot leverage this guarantee to ensure their data structure works. This leads to the second central idea of [BLS23]. They develop a new version of the [CKLPPS22] data structure that works against adaptive adversaries, at the expense of increasing the amortized time per update from $m^{o(1)}$ to $n^{1/2+o(1)}$.

At a high level, our approach is motivated by [BLS23], but ultimately we develop a fundamentally different optimization approach, which yields more tractable update problems. A key observation enabling our algorithms, is that the min-ratio cycle problem succeeds against adaptive adversaries provided that there is a particular monotonicity in the updates. We change the optimization framework to yield subproblems which can be solved by such monotonic updates. Because of this, we manage to show that in our setting, the data structure of [CKLPPS22] can directly solve the update problems in the incremental setting, with only $m^{o(1)}$ amortized time per update.

It is worth mentioning that the authors of this work recently gave a deterministic min-cost flow algorithm [BCKLPGSS23]. However, [BCKLPGSS23] still uses a version of the data structure of [CKLPPS22] that does not work against adaptive adversaries, and critically still uses stability properties of the update sequence to argue correctness.

 ℓ_p -norm Flow and Approximate Undirected Maxflow. To leverage that the min-ratio cycle data structure succeeds against monotonic adversaries, we turn to an approach motivated by lines of research on algorithms for (static) ℓ_p -norm flow and approximate undirected maxflow. One important line of research yielded undirected approximate maxflow in $\tilde{O}(\epsilon^{-1}m)$ time [CKMST11; She13; KLOS14; Pen16; She17]. This sequence of works combined first-order continuous optimization methods with graph theoretic tools. A second line of work focused on a broader class of flow problems, namely ℓ_p -norm flows [AKPS19; APS19; KPSW19; AS20; ABKS21], and developed iterative optimization methods tailored to ℓ_p -norm objectives. The problem of ℓ_p -norm flows asks for a flow \boldsymbol{f} that routes a given demand and minimizes its ℓ_p -norm, $\|\boldsymbol{f}\|_p$. This problem is maxflow for $p = \infty$ and setting $p = O(\epsilon^{-1} \log m)$ yields $(1 - \epsilon)$ -approximate undirected maxflow.

Our Approach to Incremental Flow Problems. We show that the desired monotonicity properties for the subproblems can be achieved in the setting of ℓ_p -norm flows by adapting the optimization framework. Even though ℓ_p -norm flow is less general than directed maxflow, it has interesting consequences including approximate undirected maxflow and effective resistances [SS08; CKMST11].

Switching to ℓ_p -norm flows allows us to use the ℓ_p iterative refinement framework developed in [AKPS19; APS19; AS20], and study the smoothed ℓ_p -norm flow problems of [KPSW19]. The iterative refinement framework shows that smoothed ℓ_p -norm flow computation can be accomplished by a small number of iterations of a *refinement step*. In our context $\hat{O}(p)$ steps suffice.

To solve each refinement step problem, we use a ℓ_1 multiplicative weight update method (MWU). Our method lets us solve smoothed ℓ_p -norm flow to $m^{o(1)}$ accuracy by solving a min-ratio cycle problem roughly $m^{1+o(1)}$ times. Crucially, we show that our MWU induces a *monotonicity* property in our min-ratio cycle problems. Concretely, our sequence of approximate min-ratio cycle problems only change by (1) edge insertions and (2) edge lengths increases (see Problem 5.1). In this way, we create a more tractable data structure problem than those in [BLS23], and we show that this problem can be solved using data structures from [CKLPPS22] with $m^{o(1)}$ amortized update time.

Finally, combining our monotonic ℓ_1 -MWU for computing refinement steps with iterative refinement, we obtain an incremental algorithm for (a decision version of) smoothed ℓ_p -norm flows. From this, we derive an algorithm for approximate incremental undirected maxflow and for incremental electrical flow. Using the incremental electrical flow algorithm, for a fixed pair of vertices s, t, we can detect in an incremental graph the first time the effective resistance between s and t drops below a given threshold.

If we were only focusing on approximate incremental maxflow instead of the more general ℓ_p norm flows, a similar monotonic data structure problem could also be obtained by using an ℓ_1 -oracle
MWU to compute each update step of a first-order optimization method such as one used by [She13;
KLOS14; Pen16; She17]. Wrapping this inside a first-order ℓ_{∞} optimization approach would then
yield a similar algorithm for approximate incremental maxflow.

1.1 Results

Here we present the main results of the paper. This section leverages a variety of notation, in particular graph theory conventions, all provided later in Section 2.

A central result of this paper is an algorithm with subpolynomial update time for the following undirected incremental approximate maxflow problem (which we abbreviate as incremental maxflow in the remainder of the paper). Incremental maxflow is the dynamic data structure problem of maintaining a $(1+\epsilon)$ -(multiplicative) approximate maximum flow in an undirected graph undergoing edge additions (hence the term incremental).

Problem 1.1 (Undirected Incremental Approximate Maxflow Problem). In the undirected incremental approximate maxflow problem (incremental maxflow) we are given a finite set of n vertices V, a distinct pair of elements s and t, and a parameter $\epsilon > 0$. There are then a sequence of $m \leq \text{poly}(n)$ edge insertions where starting from $E = \emptyset$ an undirected edge e is added to E with integral capacity $\mathbf{u}_e \in [1, U]$. The algorithm must maintain a $(1 + \epsilon)$ -approximate maximum flow in the capacitated graph $G = (V, E, \mathbf{u})$ before and after each edge addition.

The main result of this paper is a randomized algorithm for the incremental maxflow problem that succeeds with high probability in n (whp.) and implements each update in amortized $n^{o(1)}\epsilon^{-3}$ time. This is the first subpolynomial update time for an incremental maxflow problem which achieves even constant approximation.

Before stating our result, we comment on the adversary model. In our algorithms, we assume that the adversary can see the output flow, but not the internal randomness or information stored in the data structure. We call this an *adaptive adversary*. We refer to the stronger adversary which can also see the internal randomness as a *non-oblivious adversary*.

Theorem 1.2 (Incremental Maxflow). There is an algorithm which solves incremental maxflow (Problem 1.1) whp. in amortized $n^{o(1)}\epsilon^{-3}$ time per update against adaptive adversaries.

To obtain this result, we develop dynamic algorithms for the problem of computing smoothed ℓ_p -norm flows [KPSW19] for $p \geq 2$. The smoothed ℓ_p -norm flow problem asks to find a flow routing given vertex demands while minimizing a linear plus quadratic plus p^{th} power objective on the flow. This problem generalizes both the popular and prevalent problems of computing electric flows (and therefore solving Laplacian systems) [ST04; KMP11; CKMPPRX14; JS21a] as well as computing approximate maximum flows on undirected graphs [CKMST11; She13; KLOS14; Pen16; She17].

Problem 1.3 (Smoothed ℓ_p -norm flow). Given an undirected graph G = (V, E), gradient vector $g^G \in \mathbb{R}^E$, edge resistances and weights $\mathbf{r}^G, \mathbf{w}^G \in \mathbb{R}^E_+$, and demand vector $\mathbf{d} \in \mathbb{R}^V$ the smoothed ℓ_p -norm flow problem asks to solve the following optimization problem

$$OPT = \min_{\mathbf{B}^{\top} \boldsymbol{f} = \boldsymbol{d}} \mathcal{E}(\boldsymbol{f}) \quad for \quad \mathcal{E}(\boldsymbol{f}) \stackrel{\text{def}}{=} \langle \boldsymbol{g}^{G}, \boldsymbol{f} \rangle + \left\| \mathbf{R}^{G} \boldsymbol{f} \right\|_{2}^{2} + \left\| \mathbf{W}^{G} \boldsymbol{f} \right\|_{p}^{p}.$$
(1)

 $f \in \mathbb{R}^E$ is said to be feasible or routes the demands if $\mathbf{B}^{\top} \mathbf{f} = \mathbf{d}$, $\mathcal{E}(\mathbf{f})$ is called the energy or smoothed objective value of \mathbf{f} , and a solution to (1) is called a smoothed ℓ_p -norm flow.

For short, we refer to smoothed ℓ_p -norm flows as simply ℓ_p -norm flows throughout. Throughout we also assume that there is a feasible flow $\mathbf{f}^{(0)}$ on the initial graph. This can be ensured by determining the first instance that \mathbf{d} is feasible using a simple union-data data structure. This incurs only an additive $\tilde{O}(1)$ cost in our data structures.

Our main technical result is a high-accuracy algorithm with subpolynomial update time for the following *incremental thresholded* ℓ_p -norm flow problem, or *incremental* ℓ_p -norm flow for short. The incremental ℓ_p -norm flow data structure detects the earliest moment when the optimal value to (1) drops below a given threshold F.

Problem 1.4 (Incremental Thresholded ℓ_p -Norm Flow). Consider a dynamic instance of a ℓ_p -Norm Flow $(G, \mathbf{g}^G, \mathbf{r}^G, \mathbf{w}^G, \mathbf{d})$ that is subject to edge insertion. Let m be the final number of edges in G. Given an objective threshold $F \in \mathbb{R}$, and an error parameter $\epsilon > 0$, the problem of Incremental Threshold ℓ_p -Norm Flow asks, after the initialization or each edge insertion, to either

- 1. correctly certify that OPT > F, or
- 2. output a feasible flow \mathbf{f} with $\mathcal{E}(\mathbf{f}) \leq F + \epsilon$.

Combining an ℓ_1 -MWU with the dynamic min-ratio data structure of [CKLPPS22], we can solve Problem 1.4 in almost linear time whp. against an adaptive adversary.

Theorem 1.5. There is a randomized algorithm for Problem 1.4 that given an initial flow $f^{(0)}$ and inputs ϵ, g^G, r^G, w^G and d with sizes bounded by $\widetilde{O}(1)$ in fixed point arithmetic., runs in $p^2m^{1+o(1)}\log(\frac{\mathcal{E}(f^{(0)})-F}{\epsilon})$ time and succeeds whp. against adaptive adversaries.

As a result, taking p = 2 yields an algorithm for $(1 + \epsilon)$ -approximate incremental electrical flow with subpolynomial update time. This is the first subpolynomial time algorithm for constant accuracy incremental electrical flows.

We show how Theorem 1.5 leads to the incremental maxflow algorithm.

Proof of Theorem 1.2. The algorithm proceeds in about $\tilde{O}(\epsilon^{-1})$ phases. In each phase, we determine when the congestion of the optimal flow has decreased by at least a $(1-\epsilon)$ factor. At the start of such a phase, we find the optimal maxflow \boldsymbol{f} using the almost-linear time algorithm of [CKLPPS22]. Let C be the congestion of \boldsymbol{f} . We wish to determine when the optimal congestion is smaller than $e^{-\epsilon}C$. From the previous discussion, any flow of congestion less than $e^{-\epsilon}C$ must have its ℓ_p -norm at most $m(e^{-\epsilon}C)^p$. Our flow \boldsymbol{f} , on the other hand, has p-norm at most mC^p . So we apply Theorem 1.5 with threshold $F = m(e^{-\epsilon}C)^p$ and error $m(e^{-\epsilon}C)^p$ as well. When the data structure certifies that every flow has ℓ_p -norm at least $m(e^{-\epsilon}C)^p$, we know that \boldsymbol{f} is still a $(1-\epsilon)$ -approximate maxflow. When the data structure outputs a new feasible flow \boldsymbol{f}' s.t. $\mathcal{E}(\boldsymbol{f}') \leq F + m(e^{-\epsilon}C)^p = 2m(e^{-\epsilon}C)^p$, we know the congestion of \boldsymbol{f}' is at most $(2m(e^{-\epsilon}C)^p)^{1/p} \leq e^{-\epsilon/2}C$. This means that the optimal congestion drops at least by a factor of $(1 + \epsilon/2)$ and we restart the whole algorithm with a newly computed maximum flow. It restarts at most $\tilde{O}(1/\epsilon)$ times because the congestion is between $[\exp(-\tilde{O}(1)), \exp(\tilde{O}(1))]$.

Bit Complexity. The number of exact arithmetic operations performed in the algorithm for Theorem 1.5 is only $pm^{1+o(1)}\log(\frac{\mathcal{E}(f^{(0)})-F}{\epsilon})$. The additional p dependency arises as our algorithms manipulate numbers of size $\tilde{O}(p)$ due to the p-th power taken in the objective (1) and we are

using fixed point arithmetic. A potential way to improve the dependency on p is via floating point arithmetic and crude approximations. That is, during the course of the algorithm for Theorem 1.5, we only use $(1 + \exp(-\tilde{O}(1)))$ -multiplicative approximation to the numbers we encountered and use an additional $\tilde{O}(p)$ bits to represent their scales. This way, the numbers encountered would be of size $\tilde{O}(1)$ and we could shave one factor of p from the runtime in Theorem 1.5. This could also improve the amortized update time for incremental approximate maximum flow to $n^{o(1)}\epsilon^{-2}$.

1.2 Additional Related Work

Dynamic Flows and Matching. Exact dynamic maxflow on unit capacity graphs can be maintained in O(m) time per update by performing one augmentation per update, or in O(n) amortized time in the incremental setting [KG03; GK21]. On planar graphs, it can be maintained in the fully dynamic setting in $\tilde{O}(n^{2/3})$ update and query time [INSW11] and in the incremental setting with $\tilde{O}(\sqrt{n})$ update and query time [DGW22]. Beyond that, there is more recent work on the approximate setting, with $(1 - \epsilon)$ -approximate incremental algorithms [BLS23; GH23], as discussed earlier. Finally, in the fully dynamic setting, algorithms with super-constant (i.e. polylog or subpolynomial) approximation ratios and sublinear amortized update time [CGHPS20] and for uncapacitated graphs with subpolynomial worst-case update time [GRST21] are known.

Dynamic bipartite matching (which is a special case of directed maxflow) has also received significant attention in the approximate setting [GP13; BLSZ14; Gup14; BGS15; BS15; BS16; BHN16; PS16; Sol16; ACCSW18; CS18; BDHSS19; BHR19; BK19; CZ19; BLM20; BGS20; Waj20; BFH21; BK21; ABKL22; BK22; GSSU22; Kis22; LMSW22; RSW22; BKS23a; BKSW23; BK23]. In the $(1 - \epsilon)$ -approximate regime, the current state-of-the-art is poly $(1/\epsilon)$ update time for the incremental setting [BK23], poly $(\log(n)/\epsilon)$ for the decremental setting [BKS23b], and $O(\sqrt{m}^{1-\Omega_{\epsilon}(1)})$ for the fully dynamic setting [BKS23a]. For fully dynamic exact bipartite matching, the fastest update time is $O(n^{1.406})$ [San07; BNS19].

Finally, it is worth mentioning that recent works have leveraged numerical optimization methods based on entropy-regularized optimal transport and MWU to design dynamic algorithms for partially dynamic bipartite matching and positive linear programs [JJST22; BKS23b].

Edge Connectivity. The k-edge connectivity between two vertices s, t can be seen as a maxflow of value up to k. Dynamic k-edge st-connectivity has been studied for small constant values of $k \leq 5$ [Fre91; GI91a; GI91b; WT92; DV94; DV95; EGIN97; HK97; DW98; Tho00; HLT01; HRT18]. For super-constant k, [JS21b] presents a fully dynamic algorithm with $n^{o(1)}$ update time for $k = (\log n)^{o(1)}$. [CDKLLPSV21] give an offline fully dynamic algorithm with $\hat{O}(k^{O(k)})$ query time. These results all require small k to be efficient, whereas our result has no such restrictions. We also point out the work in [Tho07] which gives an algorithm to dynamically maintain the (value of) the global min-cut with $\tilde{O}(\sqrt{n})$ worst-case update time.

Dynamic Electric Flows. For p = 2, our incremental ℓ_p -norm flow algorithm can maintain a $(1 - \epsilon)$ -approximate electric flow between two fixed vertices $s, t \in V$ subject to edge insertions. Dynamic electric flows have previously been studied in [GHP17; DGGP19]. Such dynamic electric flow algorithms were also studied for the purpose of accelerating static maxflow and mincost flow algorithms [GLP21; BGJLLPS22]. The closely related concept of dynamic effective resistances has also been studied in the online dynamic [CGHPS20] and offline dynamic setting [LPYZ20].

1.3 Paper Organization

In the remainder of the paper, we provide preliminaries in Section 2 and then give a more technical overview of our approach in Section 3. We then we prove our main result, Theorem 1.5, via

iterative refinement in Section 4. The incremental algorithm for the ℓ_p -norm residual problem is then presented in Section 5. Finally, we argue that the approximate data structure [CKLPPS22] can be used to implement the incremental multiplicative weight method in Section 6.

2 Preliminaries

In this section, we introduce notations we use throughout the paper.

General notation. We denote vectors by boldface lowercase letters. We use uppercase boldface to denote matrices. Often, we use uppercase matrices to denote the diagonal matrices corresponding to lowercase vectors, such as $\mathbf{L} = \operatorname{diag}(\ell)$. For vectors $\boldsymbol{x}, \boldsymbol{y}$ we define the vector $\boldsymbol{x} \circ \boldsymbol{y}$ as the entrywise product, i.e., $(\boldsymbol{x} \circ \boldsymbol{y})_i = \boldsymbol{x}_i \boldsymbol{y}_i$. We also define the entrywise absolute value of a vector $|\boldsymbol{x}|$ as $|\boldsymbol{x}|_i = |\boldsymbol{x}_i|$. We use $\langle \cdot, \cdot \rangle$ as the vector inner product: $\langle \boldsymbol{x}, \boldsymbol{y} \rangle = \boldsymbol{x}^\top \boldsymbol{y} = \sum_i \boldsymbol{x}_i \boldsymbol{y}_i$. We elect to use this notation when $\boldsymbol{x}, \boldsymbol{y}$ have superscripts (such as time indices) to avoid cluttering. For positive real numbers a, b we write $a \approx_{\alpha} b$ for some $\alpha > 1$ if $\alpha^{-1}b \leq a \leq \alpha b$. For positive vectors $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^{[n]}_+$, we say $\boldsymbol{x} \approx_{\alpha} \boldsymbol{y}$ if $\boldsymbol{x}_i \approx_{\alpha} \boldsymbol{y}_i$ for all $i \in [n]$.

Graphs. In this paper, we consider multi-graphs G with edge set E(G) and vertex set V(G). When the graph is clear from context, we use the shorthands E for E(G), V for V(G), m = |E|, and n = |V|. We assume that each edge $e \in E$ has an implicit direction, used to define its edgevertex incidence matrix $\mathbf{B} \in \mathbb{R}^{E \times V}$, i.e., $\mathbf{B}_{e,u} = -1$, $\mathbf{B}_{e,v} = +1$, and zero elsewhere for the row corresponding to the edge e = (u, v). Abusing notation slightly, we often write $e = (u, v) \in E$ where e is an edge in E and u and v are the tail and head of e respectively (note that technically multi-graphs do not allow for edges to be specified by their endpoints).

A vector $\boldsymbol{d} \in \mathbb{R}^V$ is a demand vector if it is orthogonal to the all-ones vector, i.e., $\sum_{v \in V} \boldsymbol{d}_v = 0$. We say a flow $\boldsymbol{f} \in \mathbb{R}^E$ routes a demand $\boldsymbol{d} \in \mathbb{R}^V$ if $\mathbf{B}^\top \boldsymbol{f} = \boldsymbol{d}$. We say a flow \boldsymbol{f} is a circulation if it routes an all-zeros demand, i.e., each vertex has zero net flow. For an edge $\boldsymbol{e} = (u, v) \in G$ we let $\boldsymbol{b}_e \in \mathbb{R}^V$ denote the demand vector of routing one unit from u to v.

Dynamic Algorithms. We say G is a dynamic graph, if it undergoes batches $U^{(1)}, U^{(2)}, \ldots$ of updates consisting of edge insertions/deletions that are applied to G. We use $|U^{(t)}|$ to denote the number of updates contained in the batch $U^{(t)}$. The results on dynamic graphs in this article often only consider a subset of the update types and we therefore often state for each dynamic graph which updates are allowed. We say that a dynamic graph G is incremental (and decremental) if it only undergoes edge insertions (and edge deletions respectively). Additionally, we say that the graph G, after applying the first t update batches $U^{(1)}, U^{(2)}, \ldots, U^{(t)}$, is at stage t and denote the graph at this stage by $G^{(t)}$. Additionally, when G is clear, we often denote the value of a variable x at the end of stage t of G by $x^{(t)}$, or a vector x at the end of stage t of G by $x^{(t)}$.

3 Technical Overview

Here we provide a technical overview of the approach we take to obtain the results outlined in Section 1.1. In Section 3.1, we briefly review a variety of previous tools which we leverage and obstacles that we overcome to obtain our results. In Section 3.2, we then elaborate on our central insight about dynamic data structures for the minimum ratio cycle problem that fuels our results. In Section 3.3, we then discuss the dynamic optimization frameworks we use to leverage this data structure. Finally, in Section 3.4, we discuss how we put these pieces together to obtain our results.

3.1 Previous Tools and Obstacles.

We begin by describing the key tools of [CKLPPS22] and [BLS23] which solved maximum flow in almost linear time and obtained an $\hat{O}(\epsilon^{-1}\sqrt{n})$ time incremental algorithm for directed maximum flow. We elaborate on these tools and the obstacles for using them to achieve $n^{o(1)}\epsilon^{-O(1)}$ update time for undirected maximum flow.

Minimum Ratio Cycle. Both [CKLPPS22] and [BLS23] are built upon efficient data structures for approximately solving the min-ratio cycle problem on a fully-dynamic graph, i.e., finding a cycle that approximately minimizes

$$\min_{\boldsymbol{c}\in\mathbb{R}^E \text{ is a cycle }} \frac{\langle \boldsymbol{g}, \boldsymbol{c} \rangle}{\|\mathbf{L}\boldsymbol{c}\|_1}$$
(2)

where $\boldsymbol{g} \in \mathbb{R}^{E}$ and $\boldsymbol{\ell} \in \mathbb{R}^{E}_{+}$ are called *edge gradients* and *lengths* respectively.

Both data structures are based on similar ideas. They maintain a *d*-level hierarchy of vertex and edge sparsification data structures. For vertex sparsification, they use *dynamic low stretch decompositions* which is studied in the context of dynamic shortest paths [CGHPS20]. For edge sparsification, [CKLPPS22] proposed a data structure that maintains graph spanners, which are sparse graphs that preserves all-pairs distances, under edge updates as well as vertex splits. The original min-ratio cycle data structure has $m^{o(1)}$ update time and outputs $m^{o(1)}$ -approximate solutions against oblivious adversaries. Additionally, the data structure succeeds against adaptive adversaries, as long as the inputs satisfy some additional "stability properties". We elaborate on this later in Section 6. In [BLS23], they make the data structure adaptive at a higher update time of $n^{1/2+o(1)}$.

Interior Point Methods. The static algorithm of [CKLPPS22] uses an IPM potential $\Phi(f)$ to find the maximum flow. It starts at some initial flow where $\Phi(f) = \tilde{O}(m)$ and iteratively makes progress until the potential is $\Phi(f) < -\tilde{O}(m)$. At that point, an optimal flow is obtained by standard rounding techniques. When the current flow is f, the algorithm finds a $m^{o(1)}$ -approximate cycle c to an instance of min-ratio cycle (2) with $g \approx \nabla \Phi(f)$ and $\ell \approx \sqrt{\nabla^2 \Phi(f)}$. One can show that augmenting f with a multiple of c decreases the potential by at least $m^{-o(1)}$. After $m^{1+o(1)}$ iterations, the algorithm reaches a flow f whose potential value is at most $-\tilde{O}(m)$.

The incremental maximum flow algorithm of [BLS23] dynamizes the potential reduction procedure of [CKLPPS22]. To handle an edge insertions, [BLS23] keeps augmenting the current flow with an $m^{o(1)}$ -approximate min-ratio cycle until the output cycle cannot make enough, $m^{-o(1)}$, progress. The analysis of this leverages that adding an edge does not affect the feasibility of the current flow and only increases the potential by a constant amount. Additionally, once a flow of cost at most the given threshold appears in the graph, any $m^{o(1)}$ -approximate min-ratio cycle decreases the potential by at least $m^{-o(1)}$ as long as our current flow has its cost larger than the threshold by $\exp(-\tilde{O}(1))$. Since the potential value starts at $\tilde{O}(m)$, over the course of the incremental algorithm, there are at most $m^{1+o(1)}$ approximate min-ratio cycle queries. This yields a $n^{1/2+o(1)}$ -update time using their adaptive data structure for answering min-ratio cycle queries.

Obstacles. In the static case, [CKLPPS22] manages to apply an oblivious data structure for iteratively minimizing the IPM potential due to (a) the existance an optimal flow before initializing the data structure and (b) the stability of the gradient and Hessian of the IPM potential (which are the inputs to the min-ratio cycle data structure). Consequently, whenever the data structure fails to find a cycle good enough to make progress, it must be the case that some part of the data structure is broken, and we need to fix it actively. However, (a) does not hold in the incremental case. That is, whenever the oblivious data structure fails, we cannot distinguish between the case

that (1) the data structure fails due to the obliviousness, or (2) the graph does not support the optimal flow because every feasible flow has a large cost. The issue occurs even when all incremental updates come from an oblivious adversary.

Our Approach. To obtain our results we depart from prior work in both how we reason about adaptive adversaries in the minimum ratio cycle problem and in what optimization method we use for this dynamic data structure. We elaborate on each of these in the next Section 3.2 and Section 3.3 and then discuss how they are put together to obtain our results in Section 3.4.

3.2 Adaptive Adversaries and Monotonic Dynamic Min-Ratio Cycle.

We manage to use the oblivious min-ratio cycle data structure by ensuring that edge lengths are mostly increasing. That is, with a different numerical method, we can divide the entire incremental algorithm into $m^{o(1)}$ phases and within each phase, we solve a sequence of slowly changing minratio cycle problems with the same gradient g and monotonically increasing lengths ℓ . In Section 6 which presents the min-ratio cycle data structure, we work with an adversary where we can assume that all incremental updates are determined before the initialization. That is, the adversary does not have access to the internal state of our algorithm when deciding which edges to insert next. Using these facts, we show, in Section 6, that the updates to the dynamic min-ratio cycle data structure satisfy a weaker form of the hidden stable-flow chasing property, which was the critical property that [CKLPPS22] leveraged to show correctness of their min-ratio cycle data structure. While [CKLPPS22] had to periodically rebuild layers of their data structure when the "lengths" of flows at those layers may have decreased, and this prevented its application to incremental directed maxflow in [BLS23], our monotonicity property allows us to avoid this issue. In particular, whenever the data structure fails to output a good cycle, we know for certain that every feasible flow has large congestion.

3.3 From IPMs to Iterative Refinement and MWU

This paper focuses on solving incremental thresholded ℓ_p -norm flow for $p \geq 2$. This immediately gives an incremental $(1 + \epsilon)$ -approximate maximum flow due to the choice of p. Unlike IPM-based maxflow algorithms, ℓ_p -norm flows can be reduced to approximately solving $m^{o(1)}$ smoothed ℓ_p norm flow residual problems to $m^{o(1)}$ -approximation factors. To approximately solve each residual problem, we use a multiplicative weight update method (MWU) to reduce the problem to a sequence of $m^{1+o(1)}$ slowly changing ℓ_1 -regression sub-problems, which are equivalent to min-ratio cycles in our case. The nature of MWU ensures that the ℓ_1 weights are non-decreasing. These constraints on how the sequence of min-ratio cycle instances change enable us to use the data structure of [CKLPPS22] to achieve the almost-linear runtime.

The reduction to residual problem is achieved via the iterative refinement framework of [AKPS19]. Each residual problem asks to find a circulation Δ such that

$$\|\mathbf{R}\boldsymbol{\Delta}\|_{2} \le m^{o(1)}, \|\mathbf{W}\boldsymbol{\Delta}\|_{p} \le m^{o(1)}, \text{ and } \langle \boldsymbol{g}, \boldsymbol{\Delta} \rangle = -1$$
 (3)

where $\boldsymbol{r}, \boldsymbol{w} \in \mathbb{R}^E_+$ are edge weights derived from the current solution to Problem 1.3, and we are guaranteed the existence of a circulation whose both norms are at most 1.

To find a feasible residual solution to Problem (3), we use a MWU to reduce the problem to a sequence of $m^{1+o(1)} \ell_1$ sub-problems of finding a circulation c such that

$$\|\mathbf{L}\boldsymbol{c}\|_{1} \leq m^{o(1)} \|\boldsymbol{\ell}\|_{1} \text{ and } \langle \boldsymbol{g}, \boldsymbol{c} \rangle = -1.$$
(4)

This is equivalent to finding a min-ratio cycle (2) up to scaling the solution so that $\langle \boldsymbol{g}, \boldsymbol{c} \rangle = -1$. If we can solve each ℓ_1 sub-problem to $m^{o(1)}$ -approximation, we obtain a feasible residual solution. Furthermore, MWU ensures that the ℓ_1 weights are non-decreasing across all $m^{1+o(1)}$ instances.

3.4 Putting it Altogether

To dynamize the approach mentioned in Section 3.3, we make the following observation:

If $OPT \leq F$, there is a feasible solution Δ^* to (3) with $\|\mathbf{R}\Delta\|_2, \|\mathbf{W}\Delta\|_p \leq 1$ that satisfies (4), i.e., $\|\mathbf{L}\Delta^*\|_1 \leq \|\boldsymbol{\ell}\|_1$ for any $\boldsymbol{\ell}$ encountered in the algorithm.

We leverage that Δ^* is fixed with respect to r, w and g. Thus, whenever the $m^{o(1)}$ -approximate min-ratio cycle data structure cannot find a solution to (4), we certify that OPT > F. Otherwise, if $OPT \leq F$, the MWU method, as well as the iterative refinement procedure, would proceed as desired and output a flow f of ℓ_p -norm energy at most $\mathcal{E}(f) \leq F + \epsilon$ in $p^2 m^{1+o(1)} \log(1/\epsilon)$ -time.

4 Incremental *p*-Norm Iterative Refinement

In this section, we discuss the iterative refinement approach to solving *p*-norm flows and how reduce them to incremental MWUs. The iterative refinement framework of [AKPS19] reduces *p*-norm flows to approximately solving a small number of residual problems. The original framework requires an estimate on the optimal residual value, which is obtained via binary search. In our setting, a target objective value is given and we can directly use it to estimate the residual value. This requires a slight modification to the convergence analysis via measuring the progress towards the target value.

The goal of this section is to prove the following theorem:

Theorem 1.5. There is a randomized algorithm for Problem 1.4 that given an initial flow $f^{(0)}$ and inputs ϵ, g^G, r^G, w^G and d with sizes bounded by $\widetilde{O}(1)$ in fixed point arithmetic., runs in $p^2 m^{1+o(1)} \log(\frac{\mathcal{E}(f^{(0)})-F}{\epsilon})$ time and succeeds whp. against adaptive adversaries.

We first define the residual problem, $\mathcal{R}_f(\mathbf{x})$, that we consider to approximate $\mathcal{E}(\mathbf{f} + \mathbf{x}) - \mathcal{E}(\mathbf{f})$ (Problem 4.1) and provide a known lemma about its approximation quality (Lemma 4.2).

Problem 4.1 (Residual Problem). For feasible flow f in p-Norm flow instance (G, g^G, r^G, w^G, d) , we define its residual problem, $\mathcal{R}_f(x)$, (that approximates) as follows:

$$\mathcal{R}_{oldsymbol{f}}(oldsymbol{x}) \stackrel{ ext{def}}{=} \langle oldsymbol{g},oldsymbol{x}
angle + \|\mathbf{R}oldsymbol{x}\|_2^2 + \|\mathbf{W}oldsymbol{x}\|_p^p,$$

where $\boldsymbol{g} \stackrel{\text{def}}{=} \boldsymbol{g}^G + 2(\mathbf{R}^G)^2 \boldsymbol{f} + p(\mathbf{W}^G)^p |\boldsymbol{f}|^{p-2} \boldsymbol{f}, \boldsymbol{r} \stackrel{\text{def}}{=} \sqrt{(\boldsymbol{r}^G)^2 + 2p^2(\mathbf{W}^G)^p |\boldsymbol{f}|^{p-2}}, \boldsymbol{w} \stackrel{\text{def}}{=} p \boldsymbol{w}^G.$ We often ignore the subscript \boldsymbol{f} when it is clear from the context.

Lemma 4.2 (Iterative Refinement, [AKPS19; APS19; AS20]). For any f and x, we have

$$\mathcal{E}(\boldsymbol{f} + \boldsymbol{x}) - \mathcal{E}(\boldsymbol{f}) \leq \mathcal{R}_f(\boldsymbol{x}), \text{ and}$$

 $\mathcal{E}(\boldsymbol{f} + \lambda \boldsymbol{x}) - \mathcal{E}(\boldsymbol{f}) \geq \lambda \mathcal{R}_f(\boldsymbol{x}), \text{ for some } \lambda = O(p)$

Using Lemma 4.2, we can relate the optimal residual objective value to the gap between the current feasible flow f and the target threshold F.

Lemma 4.3 (Threshold Certification by Residual Value). If feasible flow \mathbf{f} in a p-Norm flow instance $(G, \mathbf{g}^G, \mathbf{r}^G, \mathbf{w}^G, \mathbf{d})$ satisfies $OPT \leq F$, then there is a circulation \mathbf{c}^* with $\mathcal{R}(\mathbf{c}^*) \leq (F - \mathcal{E}(\mathbf{f}))/\lambda$.

Proof. Let f^* be any feasible flow such that $\mathcal{E}(f^*) \leq F$. Lemma 4.2 yields that

$$F - \mathcal{E}(f) \ge \mathcal{E}(f^*) - \mathcal{E}(f) \ge \lambda \mathcal{R}\left(\frac{f^* - f}{\lambda}\right)$$

The conclusion follows because $f^* - f$ is a circulation.

Let $R \stackrel{\text{def}}{=} (\mathcal{E}(f) - F)/\lambda > 0$ be the residual threshold. Our goal now is to find a circulation c such that $\mathcal{R}(c) \leq -R/K$ for some $K = m^{o(1)}$. We show that we can compute this by solving the following *incremental residual problem*.

Problem 4.4 (Incremental K-Approximate Residual Problem). Consider an incremental graph G = (V, E) with at most m edges, a gradient vector $\mathbf{g} \in \mathbb{R}^E$, ℓ_2 and ℓ_p edge weights $\mathbf{r}, \mathbf{w} \in \mathbb{R}^E_+$ and a approximation factor K > 0. The Incremental K-Approximate Residual Problem asks, after the initialization or each edge insertion, to either

- 1. certify that there is no circulation c^* with $\langle g, c^* \rangle = -1$, $\|\mathbf{R}c^*\|_2 \leq 1$, and $\|\mathbf{W}c^*\|_p \leq 1$, or
- 2. output a circulation c such that $\langle g, c \rangle = -1$, $\|\mathbf{R}c\|_2 \leq K$, and $\|\mathbf{W}c\|_p \leq K$.

In Section 5, we will present an almost-linear time algorithm for Problem 4.4.

Lemma 4.5. For some $K = m^{o(1)}$, there is a randomized algorithm for Problem 4.4, denoted as $\mathcal{A}^{(4.5)}$, that runs in $m^{1+o(1)}$ -time and succeeds with high probability against an adaptive adversary.

Here, recall that the adversary is only adaptive against the yes/no output of the algorithm, and cannot see the internal randomness, including the flow that is being stored internally. Obviously, our algorithm can output a flow when we get a "no" output, by computing a high-accuracy *p*-norm minimizing flow, but this is different from the flow being stored internally.

At first glance, the solution guarantee of Problem 4.4 has nothing to do with the residual problem (Problem 4.1) and R. However, since we only need a solution whose objective is better than $m^{-o(1)}(F - \mathcal{E}(f))/\lambda$, we can reduce incremental *p*-norm regression to the form of Problem 4.4 using the following pair of lemmas. The first is a straightforward scaling argument.

Lemma 4.6. If there's some c^* such that $\mathcal{R}(c^*) \leq -R$ for some threshold R > 0 then,

$$\frac{\langle \boldsymbol{g}, \boldsymbol{c}^* \rangle}{\|\mathbf{R}\boldsymbol{c}^*\|_2} \leq -2\sqrt{R} \text{ and } \frac{\langle \boldsymbol{g}, \boldsymbol{c}^* \rangle}{\|\mathbf{W}\boldsymbol{c}^*\|_p} \leq -R^{(p-1)/p}$$

Proof. From the bound that $\mathcal{R}(\boldsymbol{c}^*) \leq -R$, we know

$$\langle \boldsymbol{g}, \boldsymbol{c}^* \rangle + \|\mathbf{R}\boldsymbol{c}^*\|_2^2 \leq -R$$

Thus we know that

$$\frac{\langle \boldsymbol{g}, \boldsymbol{c}^* \rangle}{\left\| \mathbf{R} \boldsymbol{c}^* \right\|_2} \leq -\frac{R}{\left\| \mathbf{R} \boldsymbol{c}^* \right\|_2} - \left\| \mathbf{R} \boldsymbol{c}^* \right\|_2 \leq -2\sqrt{R},$$

by the AM-GM inequality. Similarly,

$$\langle \boldsymbol{g}, \boldsymbol{c}^* \rangle + \| \mathbf{W} \boldsymbol{c}^* \|_p^p \leq -R$$

and hence

$$\frac{\langle \boldsymbol{g}, \boldsymbol{c}^* \rangle}{\|\mathbf{W}\boldsymbol{c}^*\|_p} \leq -\frac{R}{\|\mathbf{W}\boldsymbol{c}^*\|_p} - \|\mathbf{W}\boldsymbol{c}^*\|_p^{p-1} \leq -R^{\frac{p-1}{p}},$$

where we used the (weighted) AM-GM inequality again.

The second, given below, is used to argue that the solution quality of Problem 4.4 translates to a bound on the residual objective value.

Lemma 4.7. Given a residual threshold R > 0 and the current residual problem $\mathcal{R}(\mathbf{c})$ with gradient \mathbf{g} and ℓ_2 and ℓ_p edge weights \mathbf{r} and \mathbf{w} , consider a circulation \mathbf{c} such that $\langle \mathbf{g}, \mathbf{c} \rangle = -1$, $\|2\sqrt{R}\mathbf{R}\mathbf{c}\|_2 \leq K$, and $\|R^{(p-1)/p}\mathbf{W}\mathbf{c}\|_p \leq K$ for some $K \geq 1$. We have $\mathcal{R}((R/(2K^2))\mathbf{c}) \leq -R/(6K^2)$.

Proof. From the assumptions of the lemma, we know that

$$\langle \boldsymbol{g}, \boldsymbol{c} \rangle = -1, \ \|\mathbf{R}\boldsymbol{c}\|_2 \leq \frac{K}{2\sqrt{R}}, \text{ and } \ \|\mathbf{W}\boldsymbol{c}\|_p \leq KR^{-(p-1)/p}$$

Plugging these into the definition of the residual problem yields

$$\begin{aligned} \mathcal{R}\left(\frac{R}{2K^{2}}\boldsymbol{c}\right) &= \frac{R}{2K^{2}}\langle\boldsymbol{g},\boldsymbol{c}\rangle + \frac{R^{2}}{4K^{4}} \|\mathbf{R}\boldsymbol{c}\|_{2}^{2} + \frac{R^{p}}{2^{p}K^{2p}} \|\mathbf{W}\boldsymbol{c}\|_{p}^{p} \\ &\leq -\frac{R}{2K^{2}} + \frac{R^{2}}{4K^{4}} \frac{K^{2}}{4R} + \frac{R^{p}}{2^{p}K^{2p}} K^{p} R^{-(p-1)} \\ &= -\frac{R}{2K^{2}} + \frac{R}{16K^{2}} + \frac{R}{2^{p}K^{p}} \leq \frac{-R}{5K^{2}} \,. \end{aligned}$$

Now, we are ready to state Algorithm 1 which we use to prove Theorem 1.5.

To show correctness, we first prove that each step makes an exponential progress. Leveraging this lemma we then prove Theorem 1.5.

Lemma 4.8 (Iterative Refinement Convergence Rate). At the end of any step t in Algorithm 1,

$$\mathcal{E}(\boldsymbol{f}^{(t+1)}) - F \le \left(1 - \frac{1}{6K^2\lambda}\right) \left(\mathcal{E}(\boldsymbol{f}^{(t)}) - F\right)$$

Proof. At the end of step t, Lemma 4.7 ensures that $\mathcal{R}(\frac{R}{2K^2}c) \leq -\frac{R}{6K^2}$. Lemma 4.2 then yields

$$\mathcal{E}(\boldsymbol{f}^{(t+1)}) - \mathcal{E}(\boldsymbol{f}^{(t)}) \le \mathcal{R}\left(\frac{R}{2K^2}\boldsymbol{c}\right) \le \frac{-1}{6K^2\lambda} \left(\mathcal{E}(\boldsymbol{f}^{(t)}) - F\right)$$

Thus, we have

$$\begin{aligned} \mathcal{E}(\boldsymbol{f}^{(t+1)}) - F &= \mathcal{E}(\boldsymbol{f}^{(t+1)}) - \mathcal{E}(\boldsymbol{f}^{(t)}) + \mathcal{E}(\boldsymbol{f}^{(t)}) - F \\ &\leq \left(1 - \frac{1}{6K^2\lambda}\right) \left(\mathcal{E}(\boldsymbol{f}^{(t)}) - F\right) \end{aligned}$$

Proof of Theorem 1.5. We first show that whenever the algorithm outputs that OPT > F, this is indeed the case. Let t be an arbitrary step in which the algorithm outputs that OPT > F, let $\mathbf{f}^{(t)}$ be the flow maintained at that step, and let $\mathcal{R}(\mathbf{x})$ be the corresponding residual problem instance (Problem 4.1). Correctness of Lemma 4.5 ensures that no circulation \mathbf{c}^* satisfies

$$\langle \boldsymbol{g}^{(t)}, \boldsymbol{c}^* \rangle = -1, \|2\sqrt{R}\mathbf{R}^{(t)}\boldsymbol{c}^*\|_2 \le 1, \|R^{(p-1)/p}\mathbf{W}^{(t)}\boldsymbol{c}^*\|_p \le 1$$

Algorithm 1: Incremental Thresholded *p*-Norm Flow

1 global variables p: the norm considered in Problem 1.4. $\mathbf{2}$ $\lambda \leftarrow O(p)$: the scaling factor in Lemma 4.2. 3 $K \leftarrow m^{o(1)}$: the approximation factor in Lemma 4.5. $\mathbf{4}$ procedure INCREMENTAL PNORM $(G, d, g^G, r^G, w^G, F, \epsilon)$ 5 Initialize $f^{(0)}$ to be the optimal *p*-norm flow on *G*. 6
$$\begin{split} S \stackrel{\text{\tiny def}}{=} O(pK^2\log(\frac{\mathcal{E}(f^{(0)})-F}{\epsilon})) \\ \text{for } step \ t = 0, 1, \dots, S \ \textbf{do} \end{split}$$
7 8 Construct residual problem $(\boldsymbol{g}^{(t)}, \boldsymbol{r}^{(t)}, \boldsymbol{w}^{(t)})$ with respect to $\boldsymbol{f}^{(t)}$. 9 Define residual threshold $R \stackrel{\text{def}}{=} (\mathcal{E}(\mathbf{f}^{(t)}) - F)/\lambda$. 10Initialize and run the incremental approximate residual algorithm 11 while $\mathcal{A}^{(4.5)}$ cannot find a circulation satisfying Lemma 4.5 item 2 do 12Claim that OPT > F. 13 Wait for a new edge e. $\mathbf{14}$ Insert the new edge e to $\mathcal{A}^{(4.5)}$ with $\mathbf{15}$ $\boldsymbol{g}_{e}^{(t)} \stackrel{\text{def}}{=} \boldsymbol{g}_{e}^{G}, \, \boldsymbol{r}_{e}^{(t)} \stackrel{\text{def}}{=} 2\sqrt{R}\boldsymbol{r}_{e}^{G}, \, \text{and} \, \boldsymbol{w}_{e}^{(t)} \stackrel{\text{def}}{=} R^{(p-1)/p} \boldsymbol{w}_{e}^{G}$ $\mathcal{A}^{(4.5)}$ finds a circulation c such that 16 $\langle \boldsymbol{g}^{(t)}, \boldsymbol{c} \rangle = -1, \|2\sqrt{R}\mathbf{R}^{(t)}\boldsymbol{c}\|_{2} \leq K, \text{ and } \|R^{(p-1)/p}\mathbf{W}^{(t)}\boldsymbol{c}\|_{p} \leq K.$ Set $\boldsymbol{f}^{(t+1)} \leftarrow \boldsymbol{f}^{(t)} + \frac{R}{2K^{2}}\boldsymbol{c}$ $\mathbf{17}$ return $f^{(S)}$ 18

This fact and Lemma 4.6 then implies that every circulation c^* must have $\mathcal{R}(c^*) > -R = (F - \mathcal{E}(f^{(t)}))/\lambda$. Lemma 4.3 implies that OPT > F.

Now, we analyze the quality of the output, $f^{(S)}$. Applying Lemma 4.8 inductively on steps t yields that

$$\mathcal{E}(\boldsymbol{f}^{(S)}) - F \le \left(1 - \frac{1}{6K^2\lambda}\right)^S \left(\mathcal{E}(\boldsymbol{f}^{(0)}) - F\right) \le \epsilon$$

by definitions of S and λ (Lemma 4.2).

Finally, the runtime comes directly from $S = m^{o(1)} p \log(\frac{1}{\epsilon})$ applications of Lemma 4.5. (Recall that there is an extra factor of p due to bit complexity.)

5 Incremental Multiplicative Weight Updates

In this section, we present a MWU-based algorithm that solves Problem 4.4 and proves Lemma 4.5. We use Algorithm 2 to prove Lemma 4.5. At a high level, the algorithm outputs the final circulation as an average of $T = \tilde{O}(m)$ circulations. At each iteration *i*, we compute a set of ℓ_1 edge weights $\ell \in \mathbb{R}^E_+$ and compute an approximate min ratio cycle that minimizes $\langle \boldsymbol{g}, \boldsymbol{c} \rangle / \|\mathbf{L}\boldsymbol{c}\|_1$. As is too costly to compute each cycle from scratc, we instead apply a dynamic data structure simplified from [CKLPPS22] to compute such cycle efficiently.

Problem 5.1 (Approximate Monotonic Min Ratio Cycle (MRC)). Consider a target ratio $\alpha > 0$ and a dynamic graph G = (V, E) with edge gradients $\mathbf{g} \in \mathbb{R}^{E}$, and lengths $\boldsymbol{\ell} \in \mathbb{R}^{E}_{+}$ under edge insertions and length increases. In addition, there is a hidden circulation \mathbf{c}^{*} not necessarily supported on G, i.e., there is possibly some edge not in G with $\mathbf{c}^{*}_{e} \neq 0$. The problem of κ -approximate monotonic min ratio cycle asks, after the initialization and each edge update,

- 1. If c^* is not supported on G: output any circulation.
- 2. If c^* is supported on G: output a circulation c such that

$$\frac{\langle \boldsymbol{g}, \boldsymbol{c} \rangle}{\|\mathbf{L}\boldsymbol{c}\|_1} \leq \frac{-\alpha}{\kappa}$$

given that $\frac{\langle \boldsymbol{g}, \boldsymbol{c}^* \rangle}{\|\mathbf{L}\boldsymbol{c}^*\|_1} \leq -\alpha$.

We can solve Problem 5.1 in almost-linear time by using a data structure from [CKLPPS22]. The guarantee of this data structure is stated below and the lemma is proved in Section 6.

Lemma 5.2 (Approximate Monotonic MRC Data Structure). For some $\kappa = m^{o(1)}$, there is a randomized algorithm, denoted $\mathcal{A}^{(5,2)}$, that maintains a collection of $s = m^{o(1)}$ spanning trees T_1, T_2, \ldots, T_s and solves the κ -approximate monotonic MRC problem (Problem 5.1) in $(m+Q)m^{o(1)}$ -time where Q is the total number of updates. In particular, it always output a cycle Δ represented by $m^{o(1)}$ off-tree edges and paths on a spanning tree T_i for some $i \in [s]$.

We now state Algorithm 2 which shows Lemma 4.5. The algorithm runs a MWU algorithm with ℓ_1 -norm subproblems for $\tilde{O}(m)$ iterations. This is motivated by the ℓ_1 -IPM of [CKLPPS22] for solving min-cost flow, and contrasts with ℓ_2 -norm MWUs of [CKMST11; AKPS19] for approximate maxflow and *p*-norm regression. These ℓ_1 -norm subproblems are min-ratio cycle problems, which we solve by calling a min-ratio cycle data structure. If the output has sufficiently good ratio, the algorithm proceeds to the next iteration. Otherwise, we conclude that *G* does not support a valid circulation yet, and ask for the next edge insertion.

To analyze Algorithm 2, we keep track of the following two potentials:

$$\Phi^{(i)} \stackrel{\text{def}}{=} \left\| \mathbf{R} \boldsymbol{a}^{(i)} \right\|_2^2 \tag{5}$$

$$\Psi^{(i)} \stackrel{\text{def}}{=} \left\| \mathbf{W} \boldsymbol{b}^{(i)} \right\|_{q}^{q} \tag{6}$$

where \boldsymbol{a} and \boldsymbol{b} are defined as

$$\boldsymbol{a}^{(i)} \stackrel{\text{def}}{=} \frac{K}{\sqrt{m}} \boldsymbol{r}^{-1} + \frac{1}{T} \sum_{j=0}^{i-1} |\boldsymbol{\Delta}^{(j)}| \tag{7}$$

$$\boldsymbol{b}^{(i)} \stackrel{\text{def}}{=} \frac{K}{m^{1/q}} \boldsymbol{w}^{-1} + \frac{1}{T} \sum_{j=0}^{i-1} |\boldsymbol{\Delta}^{(j)}| \tag{8}$$

Because of the definition of $\boldsymbol{a}^{(i)}$ and $\boldsymbol{b}^{(i)}$, we know $\boldsymbol{a}^{(i)}, \boldsymbol{b}^{(i)} \geq |\boldsymbol{c}^{(i)}|$ after any iteration *i* and thus both $\Phi^{(i)} \leq \|\mathbf{R}\boldsymbol{c}^{(i)}\|_2^2$ and $\Psi^{(i)} \leq \|\mathbf{W}\boldsymbol{c}^{(i)}\|_q^q$.

Algorithm 2: Incremental K-Approximate Residual Algorithm 1 global variables p: the norm considered in Problem 4.4. $\mathbf{2}$ $q \leftarrow \min\{\log_2 m, p\}$: the norm which the algorithm handles. 3 $\kappa \leftarrow m^{o(1)}$: the approximation factor in Lemma 5.2. $\mathbf{4}$ $K \leftarrow 100q\kappa$: the approximation factor of this algorithm. $\mathbf{5}$ $\alpha \leftarrow K^{1-q}/(10q)$: the target ratio given in Lemma 5.7. 6 procedure INCREMENTALMWU(G, g, r, w)7 Initialize $c^{(0)} = 0$, $a^{(0)} = Km^{-1/2}r^{-1}$, and $b^{(0)} = Km^{-1/q}w^{-1}$. 8 Initialize edge lengths $\boldsymbol{\ell}^{(0)} \stackrel{\text{def}}{=} K^{q-2} \boldsymbol{r}^2 \boldsymbol{a}^{(0)} + \boldsymbol{w}^q (\boldsymbol{b}^{(0)})^{q-1}.$ 9 Initialize a κ -approximate MRC algorithm $\mathcal{A}^{(5,2)}$ on (G, g, ℓ) and target ratio α . 10 Set $T \leftarrow 100 qm$ 11 for *iteration* $i = 0, 1, \ldots, T$ do $\mathbf{12}$ Define the edge length $\boldsymbol{\ell}^{(i)} \stackrel{\text{def}}{=} K^{q-2} \boldsymbol{r}^2 \boldsymbol{a}^{(i)} + \boldsymbol{w}^q (\boldsymbol{b}^{(i)})^{q-1}.$ $\mathbf{13}$ Maintain an estimation $\widetilde{\ell} \in [\ell^{(i)}, 2\ell^{(i)}].$ $\mathbf{14}$ Feed updates to the estimation $\tilde{\ell}$ to $\mathcal{A}^{(5.2)}$. 15while $\mathcal{A}^{(5,2)}$ outputs a cycle of ratio $> -\alpha/\kappa$ do 16 Claim that there's no circulation c^* s.t. $\mathbf{17}$ $\langle \boldsymbol{g}, \boldsymbol{c}^* \rangle = -1, \|\mathbf{R}\boldsymbol{c}^*\|_2 \leq 1, \text{ and } \|\mathbf{W}\boldsymbol{c}^*\|_n \leq 1$ Wait for a new edge e. 18 Set $\boldsymbol{a}_{e}^{(i)} = Km^{-1/2}\boldsymbol{r}_{e}^{-1}$ and $\boldsymbol{b}_{e}^{(i)} = Km^{-1/q}\boldsymbol{w}_{e}^{-1}$. 19 Insert the new edge e to $\mathcal{A}^{(5,2)}$ with gradient \mathbf{q}_e and length $\mathbf{20}$ $\widetilde{\boldsymbol{\ell}}_e = \boldsymbol{\ell}_e^{(i)} = K^{q-2} \boldsymbol{r}_e^2 \boldsymbol{a}_e^{(i)} + \boldsymbol{w}_e^q (\boldsymbol{b}_e^{(i)})^{q-1}$ $\mathcal{A}^{(5.2)}$ outputs a cycle $\Delta^{(i)}$ s.t. $\langle \boldsymbol{q}, \Delta^{(i)} \rangle / \| \widetilde{\mathbf{L}} \Delta^{(i)} \|_1 < -\alpha/\kappa$. 21 Scale $\mathbf{\Delta}^{(i)}$ such that $\langle \boldsymbol{g}, \mathbf{\Delta}^{(i)} \rangle = -1$ and $\| \widetilde{\mathbf{L}} \mathbf{\Delta}^{(i)} \|_1 \leq \kappa / \alpha \leq K^q$. 22 Update $\boldsymbol{c}^{(i+1)} \stackrel{\text{def}}{=} \boldsymbol{c}^{(i)} + \frac{1}{T} \boldsymbol{\Delta}^{(i)}, \ \boldsymbol{a}^{(i+1)} \stackrel{\text{def}}{=} \boldsymbol{a}^{(i)} + \frac{1}{T} |\boldsymbol{\Delta}^{(i)}|, \text{ and } \boldsymbol{b}^{(i+1)} \stackrel{\text{def}}{=} \boldsymbol{b}^{(i)} + \frac{1}{T} |\boldsymbol{\Delta}^{(i)}|$ $\mathbf{23}$ return $c^{(T)}$ $\mathbf{24}$

Lemma 5.3 (Initial Potential). Initially, $\Phi^{(0)} = K^2$ and $\Psi^{(0)} = K^q$.

Lemma 5.4 (Increase in Φ). At any iteration *i*, we have $\Phi^{(i+1)} \leq \Phi^{(i)} + 3K^2/T$.

Proof. For simplicity in the presentation, we ignore superscripts (i) in the proof. Let Δ be the cycle output in Line 21 such that $\|\tilde{\mathbf{L}}\Delta\|_1 \leq K^q$ and Φ' be the new potential values after updating $\mathbf{a}' \stackrel{\text{def}}{=} \mathbf{a} + \frac{1}{T} |\Delta|$. For any edge e, we know

$$K^{q-2} \boldsymbol{r}_e^2 \boldsymbol{a}_e | \boldsymbol{\Delta}_e | \leq \boldsymbol{\ell}_e | \boldsymbol{\Delta}_e | \leq \| \mathbf{L} \boldsymbol{\Delta} \|_1 \leq \| \widetilde{\mathbf{L}} \boldsymbol{\Delta} \|_1 \leq K^q$$

Rearrangement and the monotonicity of \boldsymbol{a} yields

$$\frac{K^2}{m} \frac{|\boldsymbol{\Delta}_e|}{\boldsymbol{a}_e} \leq \boldsymbol{r}_e^2 \boldsymbol{a}_e^2 \frac{|\boldsymbol{\Delta}_e|}{\boldsymbol{a}_e} = \boldsymbol{r}_e^2 \boldsymbol{a}_e |\boldsymbol{\Delta}_e| \leq K^2$$

and therefore $|\Delta_e| \leq m a_e \leq T a_e/3$. Applying the definition of Φ' yields

$$egin{aligned} \Phi' &= \sum_e r_e^2 (oldsymbol{a}_e + rac{1}{T} |oldsymbol{\Delta}_e|)^2 \ &\leq \sum_e r_e^2 (oldsymbol{a}_e^2 + rac{3}{T} oldsymbol{a}_e |oldsymbol{\Delta}_e|) = \Phi + rac{3}{T} \sum_e r_e^2 oldsymbol{a}_e |oldsymbol{\Delta}_e| \ &\leq \Phi + rac{3}{T} K^{2-q} \, \| \mathbf{L} oldsymbol{\Delta} \|_1 \leq \Phi + rac{3K^2}{T} \end{aligned}$$

where we use $(1+x)^2 \le 1+3x$ for $x \in [0,1]$.

Lemma 5.5 (Increase in Ψ). At any iteration *i*, we have $\Psi^{(i+1)} \leq \Psi^{(i)} + 4qK^q/T$.

Proof. For simplicity in the presentation, we ignore superscripts (i) in the proof. Let Δ be the cycle output in Line 21 such that $\|\widetilde{\mathbf{L}}\Delta\|_1 \leq K^q$ and Ψ' be the new potential values after updating $\mathbf{b}' \stackrel{\text{def}}{=} \mathbf{b} + \frac{1}{T} |\Delta|$. For any edge e, we know

$$w_e^q b_e^{q-1} |\mathbf{\Delta}_e| \leq \ell_e |\mathbf{\Delta}_e| \leq \|\mathbf{L}\mathbf{\Delta}\|_1 \leq \|\mathbf{\widetilde{L}}\mathbf{\Delta}\|_1 \leq K^q$$

Rearrangement and the monotonicity of \boldsymbol{b} yields

$$\frac{K^q}{m} \frac{|\boldsymbol{\Delta}_e|}{\boldsymbol{b}_e} \leq \boldsymbol{w}_e^q \boldsymbol{b}_e^q \frac{|\boldsymbol{\Delta}_e|}{\boldsymbol{b}_e} = \boldsymbol{w}_e^q \boldsymbol{b}_e^{q-1} |\boldsymbol{\Delta}_e| \leq K^q$$

and therefore $|\mathbf{\Delta}_e| \leq m \mathbf{b}_e \leq \frac{T \mathbf{b}_e}{100q}$. Applying the definition of Ψ' yields

$$\begin{split} \Psi' &= \sum_{e} \boldsymbol{w}_{e}^{q} (\boldsymbol{b}_{e} + \frac{1}{T} |\boldsymbol{\Delta}_{e}|)^{q} \\ &\leq \sum_{e} \boldsymbol{w}_{e}^{q} (\boldsymbol{b}_{e}^{q} + \frac{4q}{T} \boldsymbol{b}_{e}^{q-1} |\boldsymbol{\Delta}_{e}|) = \Psi + \frac{4q}{T} \sum_{e} \boldsymbol{w}_{e}^{q} \boldsymbol{b}_{e}^{q-1} |\boldsymbol{\Delta}_{e}| \\ &\leq \Psi + \frac{4q}{T} \| \mathbf{L} \boldsymbol{\Delta} \|_{1} \leq \Psi + \frac{4qK^{q}}{T} \end{split}$$

where we use $(1+x)^q \le 1 + 4qx$ for $x \in [0, 1/100q]$.

Lemma 5.6 (Final Potential). When Algorithm 2 returns $\mathbf{c}^{(T)}$, we have $\Phi^{(T)} \leq 4K^2$ and $\Psi^{(T)} \leq 5qK^q$. Therefore, we have $\|\mathbf{R}\mathbf{c}^{(T)}\|_2 \leq 2K$ and $\|\mathbf{W}\mathbf{c}^{(T)}\|_p \leq \|\mathbf{W}\mathbf{c}^{(T)}\|_q \leq 2K$.

Proof. This follows directly from Lemma 5.3, Lemma 5.4, and Lemma 5.5.

Lemma 5.7 (Existence of a good ℓ_1 solution). If there is a circulation \mathbf{c}^* such that $\langle \mathbf{g}, \mathbf{c}^* \rangle = -1$, $\|\mathbf{R}\mathbf{c}^*\|_2 \leq 1$, and $\|\mathbf{W}\mathbf{c}^*\|_p \leq 1$, we have $\|\mathbf{L}^{(i)}\mathbf{c}^*\|_1 \leq 20qK^{q-1}$ at any iteration *i*.

Proof. In this proof, we consider an arbitrary iteration i and ignore the superscripts (i) for simplicity. By definition and the bounds on Φ and Ψ (Lemma 5.6), we have by the Cauchy-Schwarz inequality and Hölder inquality that,

$$\begin{split} \|\mathbf{L}\boldsymbol{c}^{*}\|_{1} &= K^{q-2}\sum_{e}\boldsymbol{r}_{e}^{2}\boldsymbol{a}_{e}|\boldsymbol{c}_{e}^{*}| + \sum_{e}\boldsymbol{w}_{e}^{q}\boldsymbol{b}_{e}^{q-1}|\boldsymbol{c}_{e}^{*}| \\ &\leq K^{q-2} \|\mathbf{R}\boldsymbol{a}\|_{2} \|\mathbf{R}\boldsymbol{c}^{*}\|_{2} + \left\|\mathbf{W}^{q-1}\boldsymbol{b}^{q-1}\right\|_{q/(q-1)} \|\mathbf{W}\boldsymbol{c}^{*}\|_{q} \\ &= K^{q-2} \|\mathbf{R}\boldsymbol{a}\|_{2} \|\mathbf{R}\boldsymbol{c}^{*}\|_{2} + \|\mathbf{W}\boldsymbol{b}\|_{q}^{q-1} \|\mathbf{W}\boldsymbol{c}^{*}\|_{q} \\ &\leq K^{q-2}(2K) + (5qK^{q})^{(q-1)/q} \cdot 2 \\ &\leq 2K^{q-1} + 10qK^{q-1} \leq 20qK^{q-1} \end{split}$$

where the second inequality comes from $\|\mathbf{W}\mathbf{c}^*\|_q \le m^{1/q-1/p}\|\mathbf{W}\mathbf{c}^*\|_p \le m^{1/\log_2 m} = 2.$

Proof of Lemma 4.5. We first show that if line 16 of Algorithm 1 occurs, then the claim about c^* in the following line 17 is true. The correctness of Lemma 5.2 says that there's no circulation c^* such that $\langle \boldsymbol{g}, \boldsymbol{c}^* \rangle = -1$ and $\|\tilde{\mathbf{L}}\boldsymbol{c}^*\|_1 \leq 1/\alpha = 20qK^{q-1}$. Correctness then follows from Lemma 5.7.

Next, the quality of the output $c^{(T)}$ is established by Lemma 5.6.

Finally, the runtime follows from Lemma 5.2 and the standard technique of maintaining $\boldsymbol{a}, \boldsymbol{b}$ and $\tilde{\boldsymbol{\ell}}$ using dynamic tree data structures such as link-cut trees. In particular, $\boldsymbol{\ell}$ is monotone and bounded by $m^{O(p)}$ on each edge, so we can maintain a $\tilde{\boldsymbol{\ell}} \approx_2 \boldsymbol{\ell}$ with $\tilde{O}(p)$ changes per change.

6 Dynamic Min-Ratio Cycle Data Structures

In this section, we prove Lemma 5.2. We use the oblivious min-ratio cycle data structure from [CKLPPS22]. In particular, the data structure can handle adversaries stronger than oblivious ones as long as the update sequence satisfies what is called the *hidden stable-flow chasing* property. In the problem of approximate monotonic min-ratio cycles, we show that the update sequences satisfy this property and the monotonicity in edge lengths allows us to use a special case of the data structure in [CKLPPS22].

Definition 6.1 (Hidden Stable-Flow Chasing Updates, Definition 6.1 [CKLPPS22]). Consider a dynamic graph $G^{(t)}$ undergoing batches of updates $U^{(1)}, \ldots, U^{(t)}, \ldots$ consisting of edge insertions and deletions. We say the sequences $\mathbf{g}^{(t)}, \boldsymbol{\ell}^{(t)}$, and $U^{(t)}$ satisfy the hidden stable-flow chasing property if there are hidden dynamic circulations $\mathbf{c}^{(t)}$ and hidden dynamic upper bounds $\mathbf{w}^{(t)}$ such that the following holds at all stages t:

- 1. $\mathbf{c}^{(t)}$ is a circulation: $\mathbf{B}_{C^{(t)}}^{\top} \mathbf{c}^{(t)} = 0.$
- 2. $\boldsymbol{w}^{(t)}$ upper bounds the length of $\boldsymbol{c}^{(t)}$: $|\boldsymbol{\ell}_e^{(t)}\boldsymbol{c}_e^{(t)}| \leq \boldsymbol{w}_e^{(t)}$ for all $e \in E(G^{(t)})$.
- 3. For any edge e in the current graph $G^{(t)}$, and any stage $t' \leq t$, if the edge e was already present in $G^{(t')}$, i.e. $e \in G^{(t)} \setminus \bigcup_{s=t'+1}^{t} U^{(s)}$, then $\boldsymbol{w}_e^{(t)} \leq 2\boldsymbol{w}_e^{(t')}$.

Theorem 6.2 (Simpler version of [CKLPPS22, Theorem 7.1]). Let G = (V, E) be a dynamic graph undergoing τ batches of updates $U^{(1)}, \ldots, U^{(\tau)}$ containing only edge insertions and deletions with edge gradient $\mathbf{g}^{(t)}$ and length $\ell^{(t)}$ such that the update sequence satisfies the hidden stable-flow chasing property (Definition 6.1) with hidden dynamic circulation $\mathbf{c}^{(t)}$ and width $\mathbf{w}^{(t)}$. In addition, $\|\mathbf{w}^{(t)}\|_1$ is non-decreasing in t, i.e., $\|\mathbf{w}^{(t)}\|_1 \leq \|\mathbf{w}^{(t+1)}\|_1$ for all steps t.

There is an algorithm on G that for $d = (\log m)^{1/8}$ maintains a collection of $s = O(\log m)^d$ spanning trees T_1, T_2, \ldots, T_s and after each update outputs a circulation Δ represented by $\exp(O(\log^{7/8} m \log \log m))$ off-tree edges and paths on some $T_i, i \in [s]$. The output circulation Δ satisfies $\mathbf{B}^{\mathsf{T}} \Delta = 0$ and for some $\kappa = \exp(-O(\log^{7/8} m \log \log m))$

$$\frac{\langle \boldsymbol{g}^{(t)}, \boldsymbol{\Delta} \rangle}{\|\mathbf{L}^{(t)}\boldsymbol{\Delta}\|_1} \leq \kappa \frac{\langle \boldsymbol{g}^{(t)}, \boldsymbol{c}^{(t)} \rangle}{(d+1)\|\boldsymbol{w}^{(t)}\|_1}$$

The algorithm succeeds whp. with total runtime $(m+Q)m^{o(1)}$ for $Q \stackrel{\text{def}}{=} \sum_{t=1}^{\tau} \left| U^{(t)} \right| \le \operatorname{poly}(n)$.

This follows from [CKLPPS22, Theorem 7.1] because the terms $\|\boldsymbol{w}^{(\mathsf{prev}_i^{(t)})}\|_1 \leq \|\boldsymbol{w}^{(t)}\|_1$ by the increasing property, and $\mathsf{prev}_i^{(t)} \leq t$.

Proof of Lemma 5.2. We apply the dynamic algorithm of Theorem 6.2 to prove the lemma. In order to do so, we need to show that there is a hidden circulation $c^{(t)}$ and width $w^{(t)}$ that satisfies the hidden stable-flow chasing property (Definition 6.1). At any stage t, if the monotonic MRC problem instance (Problem 5.1) inserts a new edge e with gradient g_e and length ℓ_e , we create an update batch $U^{(t)}$ consists of a single edge insertion (e, g_e, ℓ_e) . Otherwise, if the problem instance increases the length of an edge from ℓ_e to ℓ'_e , we create an update batch $U^{(t)}$ consists of a single edge insertion (e, g_e, ℓ_e) .

At any stage t (including stage 0, the moment after the initialization), we define the hidden circulation $c^{(t)}$ to be c^* if every edge in the circulation c^* appears in $G^{(t)}$. Otherwise, we set $c^{(t)}$ to be the all zero circulation. We also define the hidden width $w^{(t)}$ on each edge e as

$$\boldsymbol{w}_{e}^{(t)} \stackrel{\text{def}}{=} \begin{cases} \left| \boldsymbol{\ell}_{e}^{(t)} \boldsymbol{c}_{e}^{*} \right| &, \text{ if } e \in \operatorname{supp}(\boldsymbol{c}^{*}) \\ 0 &, \text{ otherwise.} \end{cases}$$

Now, we verify each condition of Definition 6.1 one at a time. Item 1 holds at any stage t because both \mathbf{c}^* and $\mathbf{0}$ are circulations. For any edge $e \in \operatorname{supp}(\mathbf{c}^*)$, $\mathbf{c}_e^{(t)}$ is either 0 or \mathbf{c}_e^* and $\mathbf{w}_e^{(t)} \geq |\boldsymbol{\ell}^{(t)} \mathbf{c}_e^{(t)}|$ holds. For any other edge $e \notin \operatorname{supp}(\mathbf{c}^*)$, both $\mathbf{c}_e^{(t)}$ and $\mathbf{w}_e^{(t)}$ are always 0 and Item 2 follows. Item 3 follows from the fact that \mathbf{c}^* is fixed beforehand and that an edge is updated when its length doubles. Also, $\|\mathbf{w}^{(t)}\|_1$ is non-decreasing in t because edge lengths are non-decreasing.

Finally, we need to show that the output after each update is always valid. At stage t, if $\sup(c^*) \not\subseteq G^{(t)}$, Problem 5.1 does not care what we output. If $\sup(c^*) \subseteq G^{(t)}$, we know $g^{(t)} = g$, $c^{(t)} = c^*$ and $\|w^{(t)}\|_1 = \|\mathbf{L}^{(t)}c^*\|_1$, and the output circulation $\boldsymbol{\Delta}$ satisfies

$$\frac{\langle \boldsymbol{g}, \boldsymbol{\Delta} \rangle}{\|\mathbf{L}^{(t)} \boldsymbol{\Delta}\|_1} \leq \kappa \frac{\langle \boldsymbol{g}, \boldsymbol{c}^{(t)} \rangle}{(d+1) \|\boldsymbol{w}^{(t)}\|_1} = \frac{\kappa}{d+1} \frac{\langle \boldsymbol{g}, \boldsymbol{c}^* \rangle}{\|\mathbf{L}^{(t)} \boldsymbol{c}^*\|_1}$$

from Theorem 6.2. The conclusion follows because $d = (\log m)^{1/8}$.

Acknowledgements

We thank the anonymous reviewers for their feedback and suggestions.

References

[ABKS21]	Deeksha Adil, Brian Bullins, Rasmus Kyng, and Sushant Sachdeva. "Almost-linear-time Weighted ℓ_p -norm Solvers in Slightly Dense Graphs via Sparsification". In: <i>arXiv preprint</i> arXiv:2102.06977 (2021) (cit. on p. 2).
[AKPS19]	Deeksha Adil, Rasmus Kyng, Richard Peng, and Sushant Sachdeva. "Iterative refinement for ℓ_p -norm regression". In: <i>Proceedings of the Thirtieth Annual ACM-SIAM Symposium</i> on Discrete Algorithms. SIAM. 2019, pp. 1405–1424 (cit. on pp. 2, 8, 9, 13).
[APS19]	Deeksha Adil, Richard Peng, and Sushant Sachdeva. "Fast, provably convergent irls algorithm for p-norm linear regression". In: Advances in Neural Information Processing Systems 32 (2019) (cit. on pp. 2, 9).
[AS20]	Deeksha Adil and Sushant Sachdeva. "Faster p-norm minimizing flows, via smoothed q-norm problems". In: <i>Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms</i> . SIAM. 2020, pp. 892–910 (cit. on pp. 2, 9).
[ACCSW18]	Moab Arar, Shiri Chechik, Sarel Cohen, Cliff Stein, and David Wajc. "Dynamic Match- ing: Reducing Integral Algorithms to Approximately-Maximal Fractional Algorithms". In: <i>ICALP</i> . Vol. 107. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018, 7:1–7:16 (cit. on p. 5).

Г		

[ABKL22]	Sepehr Assadi, Soheil Behnezhad, Sanjeev Khanna, and Huan Li. "On Regularity Lemma and Barriers in Streaming and Dynamic Matching". In: <i>CoRR</i> abs/2207.09354 (2022) (cit. on p. 5).
[AMV21]	Kyriakos Axiotis, Aleksander Madry, and Adrian Vladu. "Faster Sparse Minimum Cost Flow by Electrical Flow Localization". In: <i>FOCS</i> . IEEE, 2021, pp. 528–539 (cit. on p. 1).
[AMV20]	Kyriakos Axiotis, Aleksander Mądry, and Adrian Vladu. "Circulation control for faster minimum cost flow in unit-capacity graphs". In: 2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS). IEEE. 2020, pp. 93–104 (cit. on p. 1).
[BGS15]	Surender Baswana, Manoj Gupta, and Sandeep Sen. "Fully Dynamic Maximal Matching in O(log n) Update Time". In: <i>SIAM J. Comput.</i> 44.1 (2015), pp. 88–113 (cit. on p. 5).
[Beh23]	Soheil Behnezhad. "Dynamic Algorithms for Maximum Matching Size". In: <i>Proceedings</i> of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA). SIAM, 2023, pp. 129–162 (cit. on p. 1).
[BDHSS19]	Soheil Behnezhad, Mahsa Derakhshan, MohammadTaghi Hajiaghayi, Cliff Stein, and Madhu Sudan. "Fully Dynamic Maximal Independent Set with Polylogarithmic Update Time". In: <i>FOCS</i> . IEEE Computer Society, 2019, pp. 382–405 (cit. on p. 5).
[BK22]	Soheil Behnezhad and Sanjeev Khanna. "New Trade-Offs for Fully Dynamic Matching via Hierarchical EDCS". In: SODA. SIAM, 2022, pp. 3529–3566 (cit. on p. 5).
[BLM20]	Soheil Behnezhad, Jakub Lacki, and Vahab S. Mirrokni. "Fully Dynamic Matching: Beating 2-Approximation in Δ^{ϵ} Update Time". In: SODA. SIAM, 2020, pp. 2492–2508 (cit. on p. 5).
[BFH21]	Aaron Bernstein, Sebastian Forster, and Monika Henzinger. "A Deamortization Approach for Dynamic Spanner and Dynamic Maximal Matching". In: <i>ACM Trans. Algorithms</i> 17.4 (2021), 29:1–29:51 (cit. on p. 5).
[BGS20]	Aaron Bernstein, Maximilian Probst Gutenberg, and Thatchaphol Saranurak. "Deter- ministic decremental reachability, SCC, and shortest paths via directed expanders and congestion balancing". In: 2020 IEEE 61st Annual Symposium on Foundations of Com- puter Science (FOCS). IEEE. 2020, pp. 1123–1134 (cit. on p. 5).
[BHR19]	Aaron Bernstein, Jacob Holm, and Eva Rotenberg. "Online Bipartite Matching with Amortized $O(\log^2 n)$ Replacements". In: J. ACM 66.5 (2019), 37:1–37:23 (cit. on p. 5).
[BS15]	Aaron Bernstein and Cliff Stein. "Fully Dynamic Matching in Bipartite Graphs". In: <i>ICALP (1).</i> Vol. 9134. Lecture Notes in Computer Science. Springer, 2015, pp. 167–179 (cit. on p. 5).
[BS16]	Aaron Bernstein and Cliff Stein. "Faster Fully Dynamic Matchings with Small Approx- imation Ratios". In: SODA. SIAM, 2016, pp. 692–711 (cit. on p. 5).
[BHN16]	Sayan Bhattacharya, Monika Henzinger, and Danupon Nanongkai. "New deterministic approximation algorithms for fully dynamic matching". In: <i>STOC.</i> ACM, 2016, pp. 398–411 (cit. on p. 5).
[BK21]	Sayan Bhattacharya and Peter Kiss. "Deterministic Rounding of Dynamic Fractional Matchings". In: <i>ICALP</i> . Vol. 198. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021, 27:1–27:14 (cit. on p. 5).
[BKS23a]	Sayan Bhattacharya, Peter Kiss, and Thatchaphol Saranurak. "Dynamic $(1+\epsilon)$ -Approximate Matching Size in Truly Sublinear Update Time". In: <i>arXiv preprint arXiv:2302.05030</i> (2023) (cit. on p. 5).
[BKS23b]	Sayan Bhattacharya, Peter Kiss, and Thatchaphol Saranurak. "Dynamic algorithms for packing-covering lPS via multiplicative weight updates". In: <i>Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)</i> . SIAM. 2023, pp. 1–47 (cit. on p. 5).

[BKSW23]	Sayan Bhattacharya, Peter Kiss, Thatchaphol Saranurak, and David Wajc. "Dynamic Matching with Better-than-2 Approximation in Polylogarithmic Update Time". In: <i>Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)</i> . SIAM, 2023, pp. 100–128 (cit. on pp. 1, 5).
[BK19]	Sayan Bhattacharya and Janardhan Kulkarni. "Deterministically Maintaining a $(2 + \epsilon)$ -Approximate Minimum Vertex Cover in $O(1/\epsilon^2)$ Amortized Update Time". In: SODA. SIAM, 2019, pp. 1872–1885 (cit. on p. 5).
[BK23]	Joakim Blikstad and Peter Kiss. "Incremental (1- ϵ)-approximate dynamic matching in $O(\text{poly}(1/\epsilon))$ update time". In: <i>CoRR</i> abs/2302.08432 (2023) (cit. on p. 5).
[BLSZ14]	Bartlomiej Bosek, Dariusz Leniowski, Piotr Sankowski, and Anna Zych. "Online Bipar- tite Matching in Offline Time". In: <i>FOCS</i> . IEEE Computer Society, 2014, pp. 384–393 (cit. on p. 5).
[BLS23]	Jan van den Brand, Yang P Liu, and Aaron Sidford. "Dynamic Maxflow via Dynamic Interior Point Methods". In: <i>Proceedings of the 55th Annual ACM Symposium on Theory of Computing.</i> 2023, pp. 1215–1228 (cit. on pp. 1, 2, 5, 7, 8).
[BCKLPGSS23]	Jan van den Brand, Li Chen, Rasmus Kyng, Yang P Liu, Richard Peng, Maximilian Probst Gutenberg, Sushant Sachdeva, and Aaron Sidford. "A Deterministic Almost-Linear Time Algorithm for Minimum-Cost Flow". In: 2022 IEEE 64rd Annual Symposium on Foundations of Computer Science (FOCS). IEEE. 2023 (cit. on p. 2).
[BGJLLPS22]	Jan van den Brand, Yu Gao, Arun Jambulapati, Yin Tat Lee, Yang P. Liu, Richard Peng, and Aaron Sidford. "Faster maxflow via improved dynamic spectral vertex sparsifiers". In: <i>STOC.</i> ACM, 2022, pp. 543–556 (cit. on pp. 1, 5).
[BLLSSSW21]	Jan van den Brand, Yin Tat Lee, Yang P. Liu, Thatchaphol Saranurak, Aaron Sidford, Zhao Song, and Di Wang. "Minimum cost flows, MDPs, and ℓ_1 -regression in nearly linear time for dense instances". In: <i>STOC.</i> ACM, 2021, pp. 859–869 (cit. on p. 1).
[BLNPSSSW20]	Jan van den Brand, Yin-Tat Lee, Danupon Nanongkai, Richard Peng, Thatchaphol Saranurak, Aaron Sidford, Zhao Song, and Di Wang. "Bipartite matching in nearly- linear time on moderately dense graphs". In: 2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS). IEEE. 2020, pp. 919–930 (cit. on p. 1).
[BNS19]	Jan van den Brand, Danupon Nanongkai, and Thatchaphol Saranurak. "Dynamic Matrix Inverse: Improved Algorithms and Matching Conditional Lower Bounds". In: <i>FOCS</i> . IEEE Computer Society, 2019, pp. 456–480 (cit. on p. 5).
[CDKLLPSV21]	Parinya Chalermsook, Syamantak Das, Yunbum Kook, Bundit Laekhanukit, Yang P Liu, Richard Peng, Mark Sellke, and Daniel Vaz. "Vertex sparsification for edge connectivity". In: <i>Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)</i> . SIAM. 2021, pp. 1206–1225 (cit. on p. 5).
[CS18]	Moses Charikar and Shay Solomon. "Fully Dynamic Almost-Maximal Matching: Breaking the Polynomial Worst-Case Time Barrier". In: <i>ICALP</i> . Vol. 107. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018, 33:1–33:14 (cit. on p. 5).
[CZ19]	Shiri Chechik and Tianyi Zhang. "Fully Dynamic Maximal Independent Set in Expected Poly-Log Update Time". In: <i>FOCS</i> . IEEE Computer Society, 2019, pp. 370–381 (cit. on p. 5).
[CGHPS20]	Li Chen, Gramoz Goranci, Monika Henzinger, Richard Peng, and Thatchaphol Sara- nurak. "Fast dynamic cuts, distances and effective resistances via vertex sparsifiers". In: 2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS). IEEE. 2020, pp. 1135–1146 (cit. on pp. 5, 7).

[CKLPPS22]	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). https://arxiv.org/abs/2203.00671. IEEE. 2022, pp. 612–623 (cit. on pp. 1, 2, 4, 6–8, 13, 16).
[CKMST11]	Paul Christiano, Jonathan A. Kelner, Aleksander Mądry, Daniel A. Spielman, and Shang-Hua Teng. "Electrical flows, Laplacian systems, and faster approximation of maximum flow in undirected graphs". In: <i>Proceedings of the 43rd ACM Symposium on</i> <i>Theory of Computing, STOC 2011, San Jose, CA, USA, June 6-8 2011.</i> Available at https://arxiv.org/abs/1010.2921. ACM, 2011, pp. 273–282 (cit. on pp. 2, 3, 13).
[CKMPPRX14]	Michael B Cohen, Rasmus Kyng, Gary L Miller, Jakub W Pachocki, Richard Peng, Anup B Rao, and Shen Chen Xu. "Solving SDD linear systems in nearly m log1/2 n time". In: <i>Proceedings of the forty-sixth annual ACM symposium on Theory of computing.</i> 2014, pp. 343–352 (cit. on p. 3).
[CMSV17]	Michael B. Cohen, Aleksander Mądry, Piotr Sankowski, and Adrian Vladu. "Negative-Weight Shortest Paths and Unit Capacity Minimum Cost Flow in $\tilde{O}(m^{10/7})$ Time (Extended Abstract)". In: <i>SODA</i> . SIAM, 2017, pp. 752–771 (cit. on p. 1).
[Dah16]	Soren Dahlgaard. "On the Hardness of Partially Dynamic Graph Problems and Connections to Diameter". In: <i>ICALP</i> . Vol. 55. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016, 48:1–48:14 (cit. on p. 1).
[DS08]	Samuel I Daitch and Daniel A Spielman. "Faster approximate lossy generalized flow via interior point algorithms". In: <i>Proceedings of the fortieth annual ACM symposium on Theory of computing.</i> 2008, pp. 451–460 (cit. on p. 1).
[DGW22]	Debarati Das, Maximilian Probst Gutenberg, and Christian Wulff-Nilsen. "A near- optimal offline algorithm for dynamic all-pairs shortest paths in planar digraphs". In: <i>Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)</i> . SIAM. 2022, pp. 3482–3495 (cit. on p. 5).
[DV94]	Yefim Dinitz and Alek Vainshtein. "The connectivity carcass of a vertex subset in a graph and its incremental maintenance". In: <i>STOC.</i> ACM, 1994, pp. 716–725 (cit. on p. 5).
[DV95]	Yefim Dinitz and Alek Vainshtein. "Locally Orientable Graphs, Cell Structures, and a New Algorithm for the Incremental Maintenance of Connectivity Carcasses". In: <i>SODA</i> . ACM/SIAM, 1995, pp. 302–311 (cit. on p. 5).
[DW98]	Yefim Dinitz and Jeffery R. Westbrook. "Maintaining the Classes of 4-Edge-Connectivity in a Graph On-Line". In: <i>Algorithmica</i> 20.3 (1998), pp. 242–276 (cit. on p. 5).
[DGGLPSY22]	Sally Dong, Yu Gao, Gramoz Goranci, Yin Tat Lee, Richard Peng, Sushant Sachdeva, and Guanghao Ye. "Nested Dissection Meets IPMs: Planar Min-Cost Flow in Nearly-Linear Time". In: <i>Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)</i> . 2022, pp. 124–153. eprint: https://epubs.siam.org/doi/pdf/10.1137/1.9781612 (cit. on p. 1).
[DGGP19]	David Durfee, Yu Gao, Gramoz Goranci, and Richard Peng. "Fully dynamic spectral vertex sparsifiers and applications". In: <i>STOC</i> . ACM, 2019, pp. 914–925 (cit. on p. 5).
[EGIN97]	David Eppstein, Zvi Galil, Giuseppe F. Italiano, and Amnon Nissenzweig. "Sparsification - a technique for speeding up dynamic graph algorithms". In: <i>J. ACM</i> 44.5 (1997), pp. 669–696 (cit. on p. 5).
[Fre91]	Greg N. Frederickson. "Ambivalent Data Structures for Dynamic 2-Edge-Connectivity and k Smallest Spanning Trees". In: <i>FOCS</i> . IEEE Computer Society, 1991, pp. 632–641 (cit. on p. 5).

[GI91a]	Zvi Galil and Giuseppe F. Italiano. "Fully Dynamic Algorithms for Edge-Connectivity Problems (Extended Abstract)". In: <i>STOC</i> . ACM, 1991, pp. 317–327 (cit. on p. 5).
[GI91b]	Zvi Galil and Giuseppe F. Italiano. "Fully Dynamic Algorithms for Edge-Connectivity Problems (Extended Abstract)". In: <i>STOC</i> . ACM, 1991, pp. 317–327 (cit. on p. 5).
[GLP21]	Yu Gao, Yang P. Liu, and Richard Peng. "Fully Dynamic Electrical Flows: Sparse Maxflow Faster Than Goldberg-Rao". In: <i>FOCS</i> . IEEE, 2021, pp. 516–527 (cit. on pp. 1, 5).
[GH23]	Gramoz Goranci and Monika Henzinger. "Efficient Data Structures for Incremental Ex- act and Approximate Maximum Flow". In: 50th International Colloquium on Automata, Languages, and Programming (ICALP 2023). Schloss Dagstuhl-Leibniz-Zentrum für In- formatik. 2023 (cit. on pp. 1, 5).
[GHP17]	Gramoz Goranci, Monika Henzinger, and Pan Peng. "The Power of Vertex Sparsifiers in Dynamic Graph Algorithms". In: <i>ESA</i> . Vol. 87. LIPIcs. Schloss Dagstuhl - Leibniz- Zentrum für Informatik, 2017, 45:1–45:14 (cit. on p. 5).
[GRST21]	Gramoz Goranci, Harald Räcke, Thatchaphol Saranurak, and Zihan Tan. "The Expander Hierarchy and its Applications to Dynamic Graph Algorithms". In: <i>SODA</i> . SIAM, 2021, pp. 2212–2228 (cit. on p. 5).
[GSSU22]	Fabrizio Grandoni, Chris Schwiegelshohn, Shay Solomon, and Amitai Uzrad. "Maintain- ing an EDCS in General Graphs: Simpler, Density-Sensitive and with Worst-Case Time Bounds". In: SOSA. SIAM, 2022, pp. 12–23 (cit. on p. 5).
[Gup14]	Manoj Gupta. "Maintaining Approximate Maximum Matching in an Incremental Bi- partite Graph in Polylogarithmic Update Time". In: <i>FSTTCS</i> . Vol. 29. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2014, pp. 227–239 (cit. on p. 5).
[GK21]	Manoj Gupta and Shahbaz Khan. "Simple dynamic algorithms for maximal independent set, maximum flow and maximum matching". In: <i>Symposium on Simplicity in Algorithms (SOSA)</i> . SIAM. 2021, pp. 86–91 (cit. on p. 5).
[GP13]	Manoj Gupta and Richard Peng. "Fully Dynamic $(1 + e)$ -Approximate Matchings". In: <i>FOCS</i> . IEEE Computer Society, 2013, pp. 548–557 (cit. on p. 5).
[HKNS15]	Monika Henzinger, Sebastian Krinninger, Danupon Nanongkai, and Thatchaphol Sara- nurak. "Unifying and Strengthening Hardness for Dynamic Problems via the Online Matrix-Vector Multiplication Conjecture". In: <i>STOC.</i> ACM, 2015, pp. 21–30 (cit. on p. 1).
[HK97]	Monika Rauch Henzinger and Valerie King. "Fully dynamic 2-edge connectivity algorithm in polylogarithmic time per operation". In: <i>SRC Technical Note</i> 4 (1997) (cit. on p. 5).
[HLT01]	Jacob Holm, Kristian de Lichtenberg, and Mikkel Thorup. "Poly-logarithmic determin- istic fully-dynamic algorithms for connectivity, minimum spanning tree, 2-edge, and biconnectivity". In: J. ACM 48.4 (2001), pp. 723–760 (cit. on p. 5).
[HRT18]	Jacob Holm, Eva Rotenberg, and Mikkel Thorup. "Dynamic Bridge-Finding in $\tilde{O}(\log^2 n)$ Amortized Time". In: SODA. SIAM, 2018, pp. 35–52 (cit. on p. 5).
[INSW11]	Giuseppe F. Italiano, Yahav Nussbaum, Piotr Sankowski, and Christian Wulff-Nilsen. "Improved algorithms for min cut and max flow in undirected planar graphs". In: <i>STOC.</i> ACM, 2011, pp. 313–322 (cit. on p. 5).
[JJST22]	Arun Jambulapati, Yujia Jin, Aaron Sidford, and Kevin Tian. "Regularized Box-Simplex Games and Dynamic Decremental Bipartite Matching". In: <i>ICALP</i> . Vol. 229. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022, 77:1–77:20 (cit. on p. 5).
[JS21a]	Arun Jambulapati and Aaron Sidford. "Ultrasparse ultrasparsifiers and faster laplacian system solvers". In: <i>Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)</i> . SIAM. 2021, pp. 540–559 (cit. on p. 3).

[JS21b]	Wenyu Jin and Xiaorui Sun. "Fully Dynamic s-t Edge Connectivity in Subpolynomial Time (Extended Abstract)". In: <i>FOCS</i> . IEEE, 2021, pp. 861–872 (cit. on p. 5).
[KLS20]	Tarun Kathuria, Yang P. Liu, and Aaron Sidford. "Unit Capacity Maxflow in Almost $O(m^{4/3})$ Time". In: 61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020. IEEE, 2020, pp. 119–130 (cit. on p. 1).
[KLOS14]	Jonathan A. Kelner, Yin Tat Lee, Lorenzo Orecchia, and Aaron Sidford. "An Almost- Linear-Time Algorithm for Approximate Max Flow in Undirected Graphs, and its Mul- ticommodity Generalizations". In: <i>Proceedings of the Twenty-Fifth Annual ACM-SIAM</i> Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014. Ed. by Chandra Chekuri. SIAM, 2014, pp. 217–226 (cit. on pp. 2, 3).
[Kis22]	Peter Kiss. "Improving update times of dynamic matching algorithms from amortized to worst case". In: <i>Proceedings of the 13th Innovations in Theoretical Computer Science Conference (ITCS)</i> . 2022, p. 94 (cit. on p. 5).
[KMP11]	Ioannis Koutis, Gary L Miller, and Richard Peng. "A nearly-m log n time solver for sdd linear systems". In: 2011 IEEE 52nd Annual Symposium on Foundations of Computer Science. IEEE. 2011, pp. 590–598 (cit. on p. 3).
[KG03]	S. Kumar and P. Gupta. "An Incremental Algorithm for the Maximum Flow Problem". In: J. Math. Model. Algorithms 2.1 (2003), pp. 1–16 (cit. on p. 5).
[KPSW19]	Rasmus Kyng, Richard Peng, Sushant Sachdeva, and Di Wang. "Flows in Almost Linear Time via Adaptive Preconditioning". In: <i>Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing</i> . 2019, pp. 902–913 (cit. on pp. 2, 3).
[LMSW22]	Hung Le, Lazar Milenkovic, Shay Solomon, and Virginia Vassilevska Williams. "Dy- namic Matching Algorithms Under Vertex Updates". In: <i>ITCS</i> . Vol. 215. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022, 96:1–96:24 (cit. on p. 5).
[LS14]	Yin Tat Lee and Aaron Sidford. "Path Finding Methods for Linear Programming: Solv- ing Linear Programs in $\tilde{O}(\sqrt{\text{rank}})$ Iterations and Faster Algorithms for Maximum Flow". In: FOCS. IEEE Computer Society, 2014, pp. 424–433 (cit. on p. 1).
[LPYZ20]	Huan Li, Stacy Patterson, Yuhao Yi, and Zhongzhi Zhang. "Maximizing the Number of Spanning Trees in a Connected Graph". In: <i>IEEE Trans. Inf. Theory</i> 66.2 (2020), pp. 1248–1260 (cit. on p. 5).
[LS20]	Yang P Liu and Aaron Sidford. "Faster energy maximization for faster maximum flow". In: <i>Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing.</i> 2020, pp. 803–814 (cit. on p. 1).
[Mąd13]	Aleksander Mądry. "Navigating central path with electrical flows: From flows to matchings, and back". In: 2013 IEEE 54th Annual Symposium on Foundations of Computer Science. IEEE. 2013, pp. 253–262 (cit. on p. 1).
[Mąd16]	Aleksander Mądry. "Computing Maximum Flow with Augmenting Electrical Flows". In: 57th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2016, 9-11 October 2016, Hyatt Regency, New Brunswick, New Jersey, USA. Available at https://arxiv.org/abs/ IEEE Computer Society, 2016, pp. 593-602 (cit. on p. 1).
[PS16]	David Peleg and Shay Solomon. "Dynamic $(1 + \epsilon)$ -Approximate Matchings: A Density-Sensitive Approach". In: SODA. SIAM, 2016, pp. 712–729 (cit. on p. 5).
[Pen16]	Richard Peng. "Approximate undirected maximum flows in $O(mpolylog(n))$ time". In: Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms. SIAM. 2016, pp. 1862–1867 (cit. on pp. 2, 3).
[RSW22]	Mohammad Roghani, Amin Saberi, and David Wajc. "Beating the Folklore Algorithm for Dynamic Matching". In: <i>ITCS</i> . Vol. 215. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022, 111:1–111:23 (cit. on p. 5).

[San07]	Piotr Sankowski. "Faster dynamic matchings and vertex connectivity". In: SODA. SIAM, 2007, pp. 118–126 (cit. on p. 5).
[She13]	Jonah Sherman. "Nearly Maximum Flows in Nearly Linear Time". In: 54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, 26-29 October, 2013, Berkeley, CA, USA. IEEE Computer Society, 2013, pp. 263–269 (cit. on pp. 2, 3).
[She17]	Jonah Sherman. "Area-convexity, ℓ_{∞} regularization, and undirected multicommodity flow". In: <i>Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Com-</i> <i>puting.</i> 2017, pp. 452–460 (cit. on pp. 2, 3).
[Sol16]	Shay Solomon. "Fully Dynamic Maximal Matching in Constant Update Time". In: <i>FOCS</i> . IEEE Computer Society, 2016, pp. 325–334 (cit. on p. 5).
[SS08]	Daniel A Spielman and Nikhil Srivastava. "Graph sparsification by effective resistances". In: <i>Proceedings of the fortieth annual ACM symposium on Theory of computing.</i> 2008, pp. 563–568 (cit. on p. 2).
[ST04]	Daniel A. Spielman and Shang-Hua Teng. "Nearly-linear time algorithms for graph par- titioning, graph sparsification, and solving linear systems". In: <i>Proceedings of the 36th</i> <i>Annual ACM Symposium on Theory of Computing, STOC 2004, Chicago, IL, USA, June</i> 13-16, 2004. Available at https://arxiv.org/abs/0809.3232, https://arxiv.org/abs/0808.4134, https://arxiv.org/abs/cs/0607105. 2004, pp. 81-90 (cit. on p. 3).
[Tho00]	Mikkel Thorup. "Near-optimal fully-dynamic graph connectivity". In: <i>STOC</i> . ACM, 2000, pp. 343–350 (cit. on p. 5).
[Tho07]	Mikkel Thorup. "Fully-dynamic min-cut". In: <i>Combinatorica</i> 27.1 (2007), pp. 91–127 (cit. on p. 5).
[Waj20]	David Wajc. "Rounding dynamic matchings against an adaptive adversary". In: <i>STOC</i> . ACM, 2020, pp. 194–207 (cit. on p. 5).
[WT92]	Jeffery R. Westbrook and Robert Endre Tarjan. "Maintaining Bridge-Connected and Biconnected Components On-Line". In: <i>Algorithmica</i> 7.5&6 (1992), pp. 433–464 (cit. on p. 5).