

A Faster Algorithm for Maximum Flow in Directed Planar Graphs with Vertex Capacities

Julian Enoch

Department of Computer Science, University of Texas at Dallas, United States of America

Kyle Fox

Department of Computer Science, University of Texas at Dallas, United States of America

Dor Mesica

Efi Arazi School of Computer Science, The Interdisciplinary Center Herzliya, Israel

Shay Mozes

Efi Arazi School of Computer Science, The Interdisciplinary Center Herzliya, Israel

Abstract

We give an $O(k^3 n \log n \min(k, \log^2 n) \log^2(nC))$ -time algorithm for computing maximum integer flows in planar graphs with integer arc *and vertex* capacities bounded by C , and k sources and sinks. This improves by a factor of $\max(k^2, k \log^2 n)$ over the fastest algorithm previously known for this problem [Wang, SODA 2019].

The speedup is obtained by two independent ideas. First we replace an iterative procedure of Wang that uses $O(k)$ invocations of an $O(k^3 n \log^3 n)$ -time algorithm for maximum flow algorithm in a planar graph with k apices [Borradaile et al., FOCS 2012, SICOMP 2017], by an alternative procedure that only makes one invocation of the algorithm of Borradaile et al. Second, we show two alternatives for computing flows in the k -apex graphs that arise in our modification of Wang's procedure faster than the algorithm of Borradaile et al. In doing so, we introduce and analyze a sequential implementation of the parallel highest-distance push-relabel algorithm of Goldberg and Tarjan [JACM 1988].

2012 ACM Subject Classification Theory of computation → Network flows

Keywords and phrases flow, planar graphs, vertex capacities

Digital Object Identifier 10.4230/LIPIcs.CVIT.2016.23

1 Introduction

The maximum flow problem has been extensively studied in many different settings and variations. This work concerns two related variants of the maximum flow problem in planar graphs. The first variant is the problem of computing a maximum flow in a directed planar network with integer arc *and vertex* capacities, and k sources and sinks. The second variant, which is used in algorithms for the first variant, is the problem of computing a maximum flow in a directed network that is nearly planar; there is a set of k vertices, called apices, whose removal turns the graph planar.

The problem of maximum flow in a planar graph with vertex capacities has been studied in several works since the 1990s [9, 14, 7, 2, 13]. For a more detailed survey of the history of this problem and other relevant results see [13] and references therein. Vertex capacities pose a challenge in planar graphs because the standard reduction from a flow network with vertex capacities to a flow network with only arc capacities does not preserve planarity. The problem can be solved by algorithms for maximum flow in sparse graphs (i.e., graphs with n vertices and $O(n)$ edges that are not necessarily planar). The fastest such algorithms currently known are an $O(n^2/\log n)$ -time algorithm [11] for sparse graphs, and an $O(n^{4/3+o(1)}C^{1/3})$ -time algorithm for sparse graphs with integer capacities bounded by C [8]. Until recently, there was no planarity exploiting algorithm for the case of more than a single source and a single

45 sink. Significant progress on this problem was recently made by Wang [13]. Wang developed
 46 an $O(k^5 n \log^3 n \log^2(nC))$ -time algorithm, where k is the number of sources and sinks, and
 47 C is the largest capacity of a single vertex. This is faster than using the two algorithms for
 48 general sparse graphs mentioned above when $k = \tilde{O}(n^{1/5}/\log^2 C + (nC)^{1/15})$.

49 Wang's algorithm uses multiple calls to an algorithm of Borradaile et al. [2] for computing
 50 a maximum flow in a k -apex graph with only arc capacities. The algorithm of Borradaile
 51 et al. [2] is based on an approach originally suggested by Hochstein and Weihe [6] for a
 52 slightly more restricted problem. In Borradaile et al.'s approach, a maximum flow in a k -apex
 53 graph with n vertices is computed by simulating the Push-Relabel algorithm of Goldberg and
 54 Tarjan [4] on a complete graph with k vertices, corresponding to the k apices of the input
 55 graph. Whenever the Push-Relabel algorithm pushes flow on an arc of the complete graph,
 56 the push operation is simulated by sending flow between the two corresponding apices in the
 57 input k -apex graph. This can be done efficiently using an $O(n \log^3 n)$ time multiple-source
 58 multiple-sink (MSMS) maximum flow algorithm in planar graphs, which is the main result of
 59 the paper of Borradaile et al. [2]. Overall, their algorithm for maximum flow in k -apex graphs
 60 takes $O(k^3 n \log^3 n)$ time. Flow in k -apex graphs can also be computed using the algorithms
 61 for sparse graphs mentioned above. The $O(k^3 n \log^3 n)$ -time algorithm of Borradaile et al. is
 62 faster than these algorithms when $k = \tilde{O}(n^{1/3}/\log^2 C + (nC)^{1/9})$.

63 1.1 Our results and techniques

64 We improve the running time of Wang's algorithm to $O(k^3 n \log n \min(k, \log^2 n) \log^2(nC))$.
 65 This is faster than Wang's result by a factor of $\max(k^2, k \log^2 n)$, extending the range of
 66 values of k for which the planarity exploiting algorithm is the fastest known algorithm for
 67 the problem to $k = \tilde{O}(n^{1/3}/\log^2 C + (nC)^{1/9})$. The improvement is achieved by two main
 68 ideas. At the heart of Wang's algorithm is an iterative procedure for eliminating excess
 69 flow from vertices violating the capacity constraints. Each iteration consists of computing a
 70 circulation with some desired properties. Wang computes this circulation using $O(k)$ calls
 71 to the algorithm of Borradaile et al. for maximum flow in k -apex graphs. We show how to
 72 compute this circulation using a constant number of invocations of the algorithm for k -apex
 73 graphs. This idea alone improves on Wang's algorithm by a factor of k .

74 To further improve the running time, we modify the algorithm of Borradaile et al. for
 75 maximum flow in k -apex graphs [2]. The algorithm of Borradaile et al. uses the Push-Relabel
 76 algorithm of Goldberg and Tarjan [4]. We introduce a sequential implementation of the
 77 parallel highest-distance Push-Relabel algorithm. In this algorithm, which we call batch-
 78 highest-distance, a single operation, Bulk-Push, pushes flow on multiple arcs simultaneously,
 79 instead of just on a single arc as in Goldberg and Tarjan's Push operation. More specifically,
 80 we simultaneously push flow on all admissible arcs whose tails have maximum height (see
 81 Section 3). This is reminiscent of parallel and distributed Push-Relabel algorithms [4, 3],
 82 but our algorithm is sequential, not parallel. We prove that the total number of Bulk-Push
 83 operations performed by the batch-highest-distance algorithm is $O(k^2)$ (this should be
 84 compared to $O(k^3)$ Push operations for the FIFO or highest-distance Push-Relabel algorithms).
 85 We then show that, in the case of the k -apex graphs that show up in Wang's algorithm, we
 86 can implement each Bulk-Push operation using a single invocation of the $O(n \log^3 n)$ -time
 87 MSMS maximum flow algorithm for planar graphs [2]. Hence, we can find a maximum flow
 88 in such k -apex graphs in $O(k^2 n \log^3 n)$ time, which is faster by a factor of k than the time
 89 required by the algorithm of Borradaile et al.

90 We also give another way to modify the algorithm of Borradaile et al. for maximum
 91 flow in k -apex graphs; the second way is better when $k = o(\log^2 n)$. We observe that the

92 structure of the k -apex graphs that arise in Wang's algorithm allows us to implement each
 93 of the $O(k^3)$ Push operations of the FIFO Push-Relabel algorithm used by Borradaile et al.
 94 in just $O(n \log n)$ time. This is done using a procedure due to Miller and Naor [10] for the
 95 case when all sources and sinks lie on a single face.

96 **Roadmap.** In Section 2 we provide preliminary background and notations. Section 3 de-
 97 scribes the sequential implementation of the parallel highest-distance Push-Relabel algorithm,
 98 and its use in an algorithm for finding maximum flow in k -apex graphs. In Section 4 we
 99 describe how to use this Push-Relabel variant to obtain an improved algorithm for computing
 100 maximum flow in planar graphs with vertex capacities.

101 2 Preliminaries

102 All the graphs we consider in this paper are directed. For a graph G we use $V(G)$ and $E(G)$
 103 to denote the vertex set and arc set of G , respectively. For any vertex $v \in V(G)$, let $\deg(v)$
 104 denote the degree of v in G .

105 For a path P we denote by $P[u, v]$ the subpath of P that starts at u and ends at v . We
 106 denote by $P \circ Q$ the concatenation of two paths P, Q such that the first vertex of Q is the
 107 last vertex of P .

108 A *flow network* is a directed graph G with a capacity function $c : V(G) \cup E(G) \rightarrow [0, \infty)$
 109 on the vertices and arcs of G , along with two disjoint sets $S, T \subset V(G)$ called *sources* and
 110 *sinks*, respectively. We assume without loss of generality that sources and sinks have infinite
 111 capacities, and that, for any arc $e = (u, v) \in E(G)$, the *reverse* arc (v, u) , denoted $\text{rev}(e)$ is
 112 also in $E(G)$, and has capacity $c(\text{rev}(e)) = 0$.

113 Let $\rho : E(G) \rightarrow [0, \infty)$. To avoid clutter we write $\rho(u, v)$ instead on $\rho((u, v))$. For
 114 each vertex v let $\rho^{in}(v) = \sum_{(u,v) \in E(G)} \rho(u, v)$, and $\rho^{out}(v) = \sum_{(v,u) \in E(G)} \rho(v, u)$. The
 115 function ρ is called a *preflow* if it satisfies the following conservation constraint: for all
 116 $v \in V(G) \setminus (S)$, $\rho^{in}(v) \geq \rho^{out}(v)$. The *excess* of a vertex v with respect to a preflow ρ is
 117 defined by $\text{ex}(\rho, v) = \rho^{in}(v) - \rho^{out}(v)$. A preflow is *feasible on an arc* $e \in E(G)$ if $\rho(e) \leq c(e)$.
 118 It is *feasible on a vertex* $v \in V(G)$ if $\rho^{in}(v) \leq c(v)$. A preflow is said to be *feasible* if, in
 119 addition to the conservation constraint, it is feasible on all arcs and vertices. The value of a
 120 preflow ρ is defined as $|\rho| = \sum_{s \in S} \rho^{out}(s) - \rho^{in}(s)$. A preflow f satisfying $\text{ex}(f, v) = 0$ for
 121 all $v \in V(G) \setminus (S \cup T)$, is called a *flow*. A flow whose value is 0 is called a *circulation*. A
 122 *maximum flow* is a feasible flow whose value is maximum.

123 ► **Remark 1.** The problem of finding a maximum flow in a flow network with multiple sources
 124 and sinks can be reduced to the single-source, single-sink case by adding a super source s and
 125 super sink t , and infinite-capacity arcs (s, s_i) and (t_i, t) for every $s_i \in S$ and $t_i \in T$. If the
 126 original network is planar then this transformation adds two apices to the graph. Throughout
 127 the paper, whenever we refer to the graph G , we mean the graph G after this transformation,
 128 i.e., with a single source, the apex s , and a single sink, the apex t .

129 The *violation of a flow f at a vertex v* is defined by $\text{vio}(f, v) = \max\{0, f^{in}(v) - c(v)\}$.
 130 Thus, if f is a feasible flow then $\text{vio}(f, v) = 0$ for all vertices v . The violation of the flow f is
 131 defined to be $\text{vio}(f) = \max_{v \in V(G)} \text{vio}(f, v)$.¹

132 A preflow ρ is *acyclic* if there is no cycle C such that $\rho(e) > 0$ for every arc $e \in C$. A
 133 preflow *saturates* an arc e if $\rho(e) = c(e)$.

¹ We define violations only with respect to flows (rather than preflows) because we will only discuss preflows in the context of flow networks without finite vertex capacities.

134 The sum of two preflows ρ and η is defined as follows. For every arc $e \in E(G)$, $(\rho+\eta)(e) =$
 135 $\max\{0, \rho(e) + \eta(e) - \rho(\text{rev}(e)) - \eta(\text{rev}(e))\}$. Multiplying the preflow ρ by some constant c to
 136 get the flow $c\rho$ is defined as $(c\rho)(e) = c \cdot \rho(e)$ for all $e \in E(G)$.

137 The *residual capacity* of an arc e with respect to a preflow ρ , denoted by $c_\rho(e)$, is
 138 $c(e) - \rho(e) + \rho(\text{rev}(e))$. The *residual graph* of a flow network G with respect to a preflow ρ is
 139 the graph G where the capacity of every arc $e \in E(G)$ is set to $c_\rho(e)$. It is denoted by G_ρ . A
 140 path of G is called *augmenting* or *residual* (with respect to a preflow ρ) if it is also a path of
 141 G_ρ .

142 Suppose G and H are flow networks such that every arc in G is also an arc in H . If f' is a
 143 (pre)flow in H then the *restriction* of f' to G is the (pre)flow f in G defined by $f(e) = f'(e)$
 144 for all $e \in E(G)$.

145 **3 An algorithm for maximum flow in k -apex graphs**

146 In this section we introduce a sequential implementation of the parallel highest-distance
 147 Push-Relabel algorithm of Goldberg and Tarjan [4], and use it in the algorithm of Borradaile
 148 et al. [2] for maximum flow in k -apex graphs. We first give a high-level description of the
 149 Push-Relabel algorithm.

150 **3.1 The Push-Relabel algorithm [4]**

151 Let H be a flow network (not necessarily planar) with source s and sink t , arc capacities
 152 $c : E(H) \rightarrow \mathbb{R}$, and no finite vertex capacities. The Push-Relabel algorithm maintains a
 153 feasible *preflow* function, ρ , on the arcs of H . A vertex u is called *active* if $\text{ex}(\rho, u) > 0$. The
 154 algorithm starts with a preflow that is zero on all arcs, except for the arcs leaving the source
 155 s , which are saturated. Thus, all the neighbors of s are initially active. When the algorithm
 156 terminates, no vertex is active and the preflow function is guaranteed to be a maximum flow.
 157 The algorithm also maintains a label function h (also known as distance or height function)
 158 over the vertices of H . The label function $h : V(H) \rightarrow \mathbb{N}$ is *valid* if $h(s) = |V(H)|$, $h(t) = 0$
 159 and $h(u) \leq h(v) + 1$ for every residual arc $(u, v) \in E(H_\rho)$.

160 The algorithm progresses by performing two basic operations, **Push** and **Relabel**. A
 161 **Push** operation applies to an arc (u, v) if (u, v) is residual, $\text{ex}(\rho, u) > 0$, and $h(u) =$
 162 $h(v) + 1$. The operation moves excess flow from u to v by increasing the flow on e by
 163 $\min\{\text{ex}(\rho, u), c(e) - \rho(e)\}$.

164 The other basic operation, **Relabel**(u), assigns u the label $h(u) = \min\{h(v) : (u, v) \in$
 165 $E(H_\rho)\} + 1$ and applies to u only if u is active and $h(u)$ is not greater than the label of
 166 any neighbor of u in H_ρ . In other words, **Relabel** applies to an active vertex u only if the
 167 excess flow in u cannot be pushed out of u (because $h(u)$ is not high enough). The algorithm
 168 performs applicable **Push** and **Relabel** operations until no vertex is active.

169 To fit our purposes, we think of the algorithm as one that only maintains explicitly
 170 the excess $\text{ex}(\rho, v)$ and residual capacity $c_\rho(e)$ of each vertex v and arc e of H . The
 171 preflow ρ is implicit. In this view, a **Push**(u, v) operation decreases $\text{ex}(\rho, u)$ and $c_\rho(u, v)$ by
 172 $\min\{\text{ex}(\rho, u), c_\rho(u, v)\}$ and increases $\text{ex}(\rho, v)$ and $c_\rho(v, u)$ by the same amount.

173 We reformulate Goldberg and Tarjan's correctness proof of the generic Push-Relabel
 174 algorithm to fit this view.

175 **► Lemma 2 ([4]).** *Any algorithm that performs applicable Push and Relabel operations in*
 176 *any order satisfies the following properties and invariants:*

- 177 (1) $\text{ex}(\rho, \cdot)$ and $c_\rho(\cdot)$ are non-negative.²
 178 (2) The function h is a valid labeling function.
 179 (3) For all $v \in V$, the value of $h(v)$ never decreases, and strictly increases when $\text{Relabel}(v)$
 180 is called.
 181 (4) $h(v) \leq 2|V(H)| - 1$ for all $v \in V(H)$.
 182 (5) Immediately after $\text{Push}(u, v)$ is performed, either (u, v) is saturated or u is inactive.

183 **Proof.** Properties (1) and (5) are immediate from the definition of Push and the fact that
 184 excess and residual capacities only change during Push operations. Property (2) corresponds
 185 to Lemma 3.1 in [4], Property (3) is proved in Lemma 3.6 in [4], and Property (4) in Lemma
 186 3.7 in [4]. ◀

187 ▶ **Lemma 3.** [4, Lemma 3.3] Properties (1), (2) imply that there is no augmenting path
 188 from s to t at any point of the algorithm.

189 ▶ **Lemma 4.** [4, Lemma 3.8] Properties (3), (4) imply that the number of Relabel operations
 190 is at most $2|V(H)| - 1$ per vertex and at most $2|V(H)|^2$ overall.

191 ▶ **Lemma 5.** [4, Lemmas 3.9, 3.10] Properties (1)-(5) imply that the number of Push
 192 operations is $O(|V(H)|^2|E(H)|)$.

193 By Lemma 4 and Lemma 5, the algorithm terminates. Since upon termination no vertex is
 194 active, the implicit preflow ρ is in fact a feasible flow function. By Lemma 3 ρ is a maximum
 195 flow from s to t .

196 Variants of the Push-Relabel algorithm differ in the order in which applicable Push and
 197 Relabel operations are applied. Some variants, such as FIFO , highest-distance, maximal-
 198 excess, etc., guarantee faster termination than the $O(|V(H)|^2|E(H)|)$ guarantee given above.

199 3.2 A sequential implementation of the parallel highest-distance 200 Push-Relabel algorithm

201 We present a sequential implementation of the parallel highest-distance Push-Relabel al-
 202 gorithm, which we call $\text{Batch-Highest-Distance}$. This algorithm attempts to push flow on
 203 multiple edges simultaneously in an operation called Bulk-Push . In that sense, it resembles
 204 the parallel version of the highest-distance Push-Relabel algorithm. It is important to note,
 205 however, that Bulk-Push is a sequential operation and not a parallel/distributed one.

206 We define Bulk-Push , a batched version of the Push operation. $\text{Bulk-Push}(U, W)$ operates
 207 on two sets of vertices, U and W . It is applicable under the following requirements:

- 208 (i) $\text{ex}(u) > 0$ for all $u \in U$.
 209 (ii) There exists an integer h such that $h(u) = h$ and $h(w) = h - 1$ for all $u \in U$ and
 210 $w \in W$.
 211 (iii) There is a residual arc (u, w) for some $u \in U$ and $w \in W$.

212 Note that in a regular Push-Relabel algorithm, conditions (i) and (ii) imply that $\text{Push}(u, w)$
 213 is applicable to any residual arc (u, w) with $u \in U$ and $w \in W$. Condition (iii) guarantees
 214 there is at least one such arc. Bulk-Push pushes as much excess flow as possible from vertices
 215 in U to vertices in W so that after Bulk-Push the following property holds:

- 216 (5*) Immediately after $\text{Bulk-Push}(U, W)$ is called, for all $u \in U$ and $w \in W$, either (u, w) is
 217 saturated or u is inactive.

² This corresponds to the function ρ being a feasible preflow.

218 We replace property (5) with the more general property (5*) in Lemma 2 and Lemma 5.
 219 With this modification, Lemmas 2, 3, 4 and 5 apply to our sequential implementation. The
 220 proofs from [4] need no change except replacing Push with Bulk-Push. Hence, our variant
 221 terminates correctly with a maximum flow from s to t .

222 ► Remark. One may think of Bulk-Push(U, W) as performing in parallel all Push operations
 223 on arcs whose tail is in U and whose head is in W . However, not every maximum flow
 224 with sources U and sinks W can be achieved as the sum of flows pushed by multiple Push
 225 operations. For example, consider the case where U consists of a single vertex u , with
 226 $\text{ex}(u) = 2$, $W = \{w_1, w_2\}$, and the residual capacities of (u, w_1) and (u, w_2) are both 2.
 227 Bulk-Push(U, W) may push one unit of excess flow from u on each of (u, w_1) and (u, w_2) , but
 228 Push(u, w_i) would push 2 units of flow on (u, w_i) , and no flow on the other arc. Therefore, the
 229 correctness of this variant cannot be argued just by simulating Bulk-Push by multiple Push
 230 operations. Instead we chose to argue correctness by stating the generalized property (5*).

231 We now discuss a concrete policy for choosing which Bulk-Push and Relabel operations
 232 to perform in the above algorithm. This policy is similar, but not identical, to the highest-
 233 distance Push-Relabel algorithm [4, 3]. As long as there is an active vertex, the algorithm
 234 repeatedly executes the following two steps, which together are called a *pulse*. Let h_{max}
 235 be the maximum label of an active vertex. That is, $h_{max} = \max\{h(v) : \text{ex}(\rho, v) > 0\}$. Let
 236 H_{max} be the set of all the active vertices whose height is h_{max} . In the first step of the pulse,
 237 the algorithm invokes Bulk-Push(H_{max}, W) where W is the set of all vertices $w \in V$ such
 238 that $h(w) = h_{max} - 1$.³ In the second step of the pulse, the algorithm applies the Relabel
 239 operation to all remaining active vertices in H_{max} in arbitrary order.

■ **Algorithm 1** Batch-Highest-Distance(G, c)

1: Initialize $h(\cdot)$, $c_\rho(\cdot)$ and $\text{ex}(\cdot)$
 2: **while** there exists an active vertex **do**
 3: $h_{max} \leftarrow \max\{h(v) : \text{ex}(\rho, v) > 0\}$
 4: $H_{max} \leftarrow \{v \in V(H) : \text{ex}(\rho, v) > 0, h(v) = h_{max}\}$
 5: $W \leftarrow \{w \in V(H) : h(w) = h_{max} - 1\}$
 6: Bulk-Push(H_{max}, W)
 7: Relabel all active vertices in H_{max} in arbitrary order
 8: **end while**

240 ► Remark. The crucial difference between this policy and the highest-distance Push-Relabel
 241 algorithm [4, 3] is that in the highest-distance algorithm a vertex u with height h_{max} is
 242 relabeled as soon as no more Push operations can be applied to u . In contrast, our variant
 243 first pushes flow from all vertices with height h_{max} and only then relabels all of them.

244 We partition the pulses into two types according to whether any vertices are relabeled
 245 in the relabel step of the pulse. A pulse in which at least one vertex is relabeled is called
 246 *saturating*. All other pulses are called *non-saturating*.⁴

³ Formally it may be that Bulk-Push(H_{max}, W) is not applicable because condition (iii) is not satisfied, e.g., when $W = \emptyset$. In such cases Bulk-Push does not push any flow. Condition (iii) is essential for the termination of the generic generalized algorithm, which may repeat such empty calls to Bulk-Push indefinitely. However, we prove in Lemma 6 that in our specific policy there are $O(|V(H)|^2)$ pulses, regardless of the flow pushed (or not pushed) by Bulk-Push in each pulse.

⁴ This is a generalization of the notions of saturating and non-saturating Push operations in [4].

247 By Lemma 4, the total number of Relabel operations executed by the batch-highest-
 248 distance algorithm is $O(|V(H)|^2)$. We now prove the same bound on the number of Bulk-Push
 249 operations.

250 ► **Lemma 6.** *The number of pulses (and hence also the number of calls to Bulk-Push) executed*
 251 *by the batch-highest-distance algorithm is $O(|V(H)|^2)$.*

252 **Proof.** Note that the Relabel step of a saturating pulse consists of at least one call to Relabel
 253 which strictly increases the height of an active vertex v whose height (before the increase)
 254 was h_{max} . Hence, a saturating pulse strictly increases the value of h_{max} . The fact that
 255 the height of each vertex never decreases and is bound by $2|V(H)|$ implies that (i) there
 256 are $O(|V(H)|^2)$ saturating pulses, and (ii) the total increase in h_{max} over all saturating
 257 Bulk-Push operations is $O(|V(H)|^2)$.

258 As for non-saturating pulses, note that since excess flow is always pushed to a vertex with
 259 lower height, the push step of a pulse does not create excess in any vertex with height greater
 260 than or equal to h_{max} , so all vertices with height greater than h_{max} remain inactive during
 261 the pulse. By property (5*), for every $u \in H_{max}$ and $w \in W$, either (u, w) is saturated, or u
 262 is inactive. Since the pulse is non-saturating, it follows that all the vertices in H_{max} become
 263 inactive during the pulse. Hence, the value of h_{max} strictly decreases during a non-saturating
 264 pulse. Since $h_{max} \geq 0$, the total decrease in h_{max} is also $O(|V(H)|^2)$, so there are $O(|V(H)|^2)$
 265 non-saturating pulses. ◀

266 Note that we do not claim that implementing the Bulk-Push operation by applying
 267 applicable Push(u, w) operations for $u \in U$, $w \in W$ until no more such operations can be
 268 applied would result in fewer Push operations than the $O(|V(H)|^2|E(H)|)$ bound of Lemma 5
 269 for the generic Push-Relabel algorithm. However, in Section 4 we will show a situation
 270 where each call to Bulk-Push can be efficiently implemented using a single invocation of a
 271 multiple-source multiple-sink algorithm in a planar graph.

272 3.3 The algorithm of Borradaile et al. for k -apex graphs [2]

273 The algorithm of Borradaile et al. [2, Section 5] uses the framework of Hochstein and Weihe [6].
 274 Let H be a graph with a set V^\times of k apices. Denote $V_0 = V(H) \setminus V^\times$. The goal is to
 275 compute a maximum flow in H from a source $s \in V(H)$ to a sink $t \in V(H)$. We assume
 276 that s and t are apices. This is without loss of generality since treating s and t as apices
 277 leaves the number of apices in $O(k)$. Let K^\times be a complete graph over V^\times . The algorithm
 278 computes a maximum flow ρ from s to t in H by simulating a maximum flow computation
 279 from s to t in K^\times using the Push-Relabel algorithm. Whenever a Push operation is performed
 280 on an arc (u, v) of K^\times it is implemented by pushing flow from u to v in the graph H_{uv} ,
 281 induced by $V_0 \cup \{u, v\}$ on the residual graph of H with respect to the flow computed so far.
 282 Note that, because no vertex of V_0 is an apex of H , H_{uv} is a 2-apex graph with apices u, v .
 283 Borradaile et al. use this fact to compute a maximum flow from u to v in H_{uv} as follows.
 284 They split u into multiple copies, each incident to a different vertex w for which (u, w) is
 285 an arc of H_{uv} . A similar process is then applied to v . Note that the resulting graph is
 286 planar. A maximum flow from u to v in H_{uv} is equivalent to a maximum flow with sources
 287 the copies of u and sinks the copies of v in the resulting graph. This flow can be computed
 288 by the multiple-source multiple-sink maximum flow algorithm (the main result in [2]) in
 289 $O(|V(H)| \log^3 |V(H)|)$ time.

290 The correctness of implementing the Push-Relabel algorithm on K^\times in this way was
 291 proved by Hochstein and Weihe [6] by proving essentially that the algorithm satisfies the

292 properties in Lemma 2. Borradaile et al. used the FIFO policy of Push-Relabel, which
 293 guarantees that the number of Push operations is $O(k^3)$, so the overall running time of their
 294 algorithm is $O(k^3|V(H)|\log^3|V(H)|)$.

295 3.4 An algorithm for maximum flow in k -apex graphs

296 We use the algorithm of Borradaile et al. for maximum flow in k -apex graphs from the
 297 previous section, but use our new batch-highest-distance Push-Relabel algorithm instead of
 298 the FIFO Push-Relabel algorithm to compute the maximum flow in K^\times . Note that, in order
 299 to implement the batch-highest-distance algorithm on K^\times we only need to maintain the
 300 excess $ex(\rho, v)$ and labels $h(v)$ of each vertex $v \in K^\times$, and to be able to implement Bulk-Push
 301 so that after the execution, property (5*) is fulfilled. We do not define a flow function in
 302 K^\times nor do we explicitly maintain residual capacities of arcs of K^\times . Instead, we maintain a
 303 preflow ρ in H , and define that an arc (u, v) of K^\times is residual if and only if there exists a
 304 residual path from u to v in H_ρ that is internally disjoint from the vertices of V^\times . Under
 305 this definition, there is no path of residual arcs in K^\times starting at s and ending at t if and
 306 only if there is no such path in H . Since K^\times has $O(k)$ vertices, by Lemma 6, the algorithm
 307 performs $O(k^2)$ pulses.

308 We next describe how a Bulk-Push(U, W) operation in K^\times is implemented. Let $A = U \cup W$.
 309 Let H_A be the graph obtained from H_ρ by deleting the vertices $V^\times \setminus A$. Bulk-Push(U, W)
 310 in K^\times is implemented by pushing a maximum flow in H_A with sources the vertices U and
 311 sinks the vertices W , with the additional restriction that the amount of flow leaving each
 312 vertex $u \in U$ is at most the excess of u . The efficiency of the procedure depends on how
 313 fast we can compute the maximum flow in H_A . We denote the time to execute a single
 314 Bulk-Push operation in the graph H_A by T_{BP} . Note that $T_{BP} = \Omega(k)$, as it takes $\Omega(k)$ time
 315 to construct H_A from H_ρ .

316 The proof of correctness is an easy adaptation of the proof of Hochstein and Weihe [6].
 317 We cannot use their proof without change because Hochstein and Weihe considered only Push
 318 operations along a single arc of K^\times rather than the Bulk-Push operations which involves
 319 more than a single pair of vertices of K^\times .


320 ► **Lemma 7.** *Maximum flow in k -apex graphs can be computed in $O(k^2 \cdot T_{BP})$ time.*

321 **Proof.** We first show that the properties (1)-(4) in the statement of Lemma 2, and the
 322 generalized property (5*) from Section 3.2 hold.

323 Property (1) holds since Bulk-Push(U, W) limits the amount of flow pushed from each
 324 vertex $u \in U$ by the excess of u . Properties (3) and (4) hold without change since Relabel is
 325 not changed.

326 To show property (5*) holds, recall that an arc (u, w) of K^\times is residual if there exists,
 327 in the residual graph H_ρ with respect to the current preflow ρ , a residual path from u to
 328 w that is internally disjoint from any vertex of V^\times . With this definition it is immediate
 329 that property (5*) holds, since our implementation of Bulk-Push(U, W) pushes a maximum
 330 flow in H_A from U that is limited by the excess flow in each vertex of U . Hence, after
 331 Bulk-Push(U, W) is executed, for every $u \in U$ and $w \in W$, either there is no residual path
 332 from u to w in H_ρ that is internally disjoint from V^\times , or u is inactive.

333 As for property (2), since we did not change Relabel, h remains valid after calls to Relabel.
 334 It remains to show that h remains a valid labeling after Bulk-Push(U, W). Consider two
 335 vertices $a, b \in V^\times$. We will show that after Bulk-Push(U, W), either the arc (a, b) of K^\times is
 336 saturated (i.e., is no residual path from a to b in H_ρ), or $h(a) \leq h(b) + 1$. The flow pushed



■ **Figure 1** Illustration of property (2) in the proof of Lemma 7. A $\text{Bulk-Push}(U, W)$ operation pushes flow along paths P and P' (dashed red paths). If Q (blue solid path) is residual after the Bulk-Push operation then the dotted black paths were residual before.

337 (in H_A) by the call $\text{Bulk-Push}(U, W)$ can be decomposed into a set \mathcal{P} of flow paths, each of
 338 which starts at a vertex of U and ends at a vertex of W .

339 Assume that after performing $\text{Bulk-Push}(U, W)$ there is an augmenting a -to- b path, Q in
 340 H_ρ . If Q does not intersect any path $P \in \mathcal{P}$ then Q was residual also before $\text{Bulk-Push}(U, W)$
 341 was called, so $h(a) \leq h(b) + 1$ because h was a valid labeling before the call. Otherwise, Q
 342 intersects some path in \mathcal{P} . Let c, d be the first and last vertices of Q that also belong to
 343 paths in \mathcal{P} . Let $P, P' \in \mathcal{P}$ be paths such that $c \in P$ and $d \in P'$. Let $w \in W$ be the last
 344 vertex of P and let $u \in U$ be the first vertex of P' . See Figure 1 for an illustration. Then,
 345 before $\text{Bulk-Push}(U, W)$ was called, $Q[a, c] \circ P[c, w]$ was a residual path from a to w , and
 346 $P'[u, d] \circ Q[d, b]$ was a residual path from u to b . Since h was a valid labeling before the call,
 347 we have

348
$$h(u) \leq h(b) + 1 \quad \text{and} \quad h(a) \leq h(w) + 1.$$

349 Since $h(u) = h(w) + 1$ it follows that

350
$$h(a) \leq h(w) + 1 = h(u) \leq h(b) + 1,$$

351 showing property (2).

352 We have shown that properties (1)-(4) and (5*) hold. Hence, by Lemmas 6 and 4, the
 353 algorithm terminates after performing $O(|V^\times|^2) = O(k^2)$ Bulk-Push and Relabel operations.
 354 Since each Relabel takes $O(k)$ time, and each Bulk-Push takes $\Omega(k)$ time, the total running
 355 time of the algorithm is $O(k^2 \cdot T_{BP})$. By Lemma 3, when the algorithm terminates there is
 356 no residual path from s to t in K^\times . By our definition of residual arcs of K^\times this implies that
 357 there is no residual path from s to t in H_ρ , so ρ is a maximum flow from s to t in H . ◀

358 **4 A faster algorithm for maximum flow with vertex capacities**

359 In this section, we give a faster algorithm for computing a maximum flow in a directed
 360 planar graph with integer arc and vertex capacities bounded by C , parameterized by the
 361 number k of terminal vertices (sources and sinks). The fastest algorithm currently known for
 362 this problem is by Wang [13]. It runs in $O(k^5 n \text{ polylog}(nC))$ time. We first sketch Wang's
 363 algorithm. We only go into details in the parts of the algorithm that will be modified in our
 364 algorithm in Section 4.2.

365 **4.1 Wang's algorithm**

366 Wang's algorithm uses the following two auxiliary graphs. In both these graphs only the
 367 arcs are capacitated. Let G be a planar network. Let k be the total number of sources and
 368 sinks in G . Recall from Remark 1 that we turn G into a 2-apex flow network with a single
 369 super-source s and super-sink t .

370 ► **Definition 8** (The graph G°). *For a flow network G , the network G° is obtained by the
 371 following procedure. For each vertex $v \in V(G)$, replace v with an undirected cycle C_v with
 372 $d = \deg(v)$ vertices v_1, \dots, v_d .⁵ Each arc in C_v has capacity $c(v)/2$. Connect each arc incident
 373 to v with a different vertex v_i , preserving the clockwise order of the arcs so that no new
 374 crossings are introduced.*

375 ► **Definition 9** (The graph G^\times). *Let f be a flow in G . Let X be the set of infeasible vertices,
 376 i.e., vertices $x \in V(G)$ such that $f^{in}(x) > c(x)$. The graph G^\times is defined as follows. Starting
 377 with G° , for each vertex $x \in X$, replace the cycle representing x with two vertices x^{in} , x^{out}
 378 and an arc (x^{in}, x^{out}) of capacity $c(x)$.*

379 *Every arc of capacity $c > 0$ going from a vertex $u \notin C_x$ to a vertex in C_x becomes an arc
 380 (u, x^{in}) of capacity c . Similarly, every arc of capacity $c > 0$ going from a vertex of C_x to a
 381 vertex $u \notin C_x$ becomes an arc (x^{out}, u) with capacity c .*

382 Note that even though $G - \{s, t\}$ is planar, $G^\times - \{s, t\}$ is not. $\{x^{in} : x \in X\} \cup \{x^{out} :$
 383 $x \in X\} \cup \{s, t\}$ is an apex set in G^\times . Thus, G^\times is a $(2|X| + 2)$ -apex graph.

384 Recall that if H and G are two graphs such that every arc of G is also an arc of H ,
 385 then the restriction of a flow f' in H to G is a flow f in G such that $f(e) = f'(e)$ for all
 386 $e \in E(G)$. Thus we can speak of the restriction of a flow f° in G° , to a flow f in G , and of
 387 the restriction of a flow f^\times in G^\times to a flow f in G .

388 Let λ^* be the value of the maximum flow in G . Wang's algorithm uses binary search to
 389 find λ^* . Let λ be the current candidate value for λ^* . The algorithm computes a flow f° with
 390 value λ in the graph G° . Let f be the restriction of f° to G . Wang proves that the set X
 391 of infeasible vertices under f has size at most $k - 2$, and that the sum of the violations of
 392 the vertices in X is at most $(k - 2)C$. As long as $\text{vio}(f) > 2k$, the algorithm improves f .
 393 This improvement phase, which will be described shortly, is the crux of the algorithm. If
 394 $\text{vio}(f) \leq 2k$, then $O(k^2)$ iterations of the classical Ford-Fulkerson algorithm suffice to get rid
 395 of all the remaining violations.

396 The improvement phase of the algorithm is based on finding a circulation g that cancels
 397 the violations on the infeasible vertices and does not create too much violation on other
 398 vertices. It can then be shown that adding $1/k \cdot g$ to the flow f decreases $\text{vio}(f)$ by a
 399 multiplicative factor of roughly $1 - 1/k$. After $O(k \log(kC))$ iterations of the improvement
 400 step, $\text{vio}(f)$ is at most $2k$.

401 Wang proves that in order to find the circulation g , it suffices to compute a circulation
 402 g^\times in G^\times that satisfies the following properties:

- 403 1. $f^\times + g^\times$ is feasible in G^\times .
- 404 2. The restriction of $f^\times + g^\times$ to G has no violations on vertices of X .
- 405 3. The restriction of $f^\times + g^\times$ to G has at most $(k - 2) \cdot \text{vio}(f)$ violation on any vertex in
 406 $V(G) \setminus X$.

⁵ By undirected cycle we mean that there are directed arcs in both directions between every pair of consecutive vertices of the cycle C_v .

407 The desired circulation g is the restriction of g^\times to G . If no such g^\times exists then g does not
 408 exist, which implies that $\lambda > \lambda^*$.

409 Wang essentially shows that any algorithm for finding g^\times in $O(T)$ time, where $T = \Omega(n)$,
 410 yields an algorithm for maximum flow with vertex capacities in $O(kT \log(kC) \log(nC))$
 411 time. The additional terms stem from the $O(k \log(kC))$ iterations of the improvement
 412 step, and the $\log(nC)$ steps of the binary search. Wang shows how to compute g^\times in
 413 $T = O(k^4 n \log^3 n)$ time, by eliminating the violation at each vertex of X one after the other
 414 in an auxiliary graph obtained from G^\times . Thus, the overall running time of his algorithm is
 415 $O(k^5 n \log^3 n \log(kC) \log(nC))$.

416 4.2 A faster algorithm for computing g^\times

417 We propose a faster way of computing the circulation g^\times by eliminating the violations in
 418 all the vertices of X in a single shot. Doing so correctly requires some care in defining
 419 the appropriate capacities in the auxiliary graph, since we only know that for each $x \in X$,
 420 g^\times should eliminate at least $\text{vio}(f, x)$ units of flow from x , but the actual amount of flow
 421 eliminated from x may have to be larger. This issue does not come up when resolving the
 422 violations one vertex at a time as was done by Wang.

423 Define an auxiliary graph H as follows. Starting with $G_{f^\times}^\times$, the residual graph of G^\times
 424 with respect to f^\times ,

- 425 ■ For each $x \in X$, set the capacity of the arc $(x^{\text{in}}, x^{\text{out}})$ to be 0 and the capacity of
 426 $(x^{\text{out}}, x^{\text{in}})$ to be $c(x)$.
- 427 ■ Add a super source s' and arcs (s', x^{in}) with capacity $\text{vio}(f, x)$ for every $x \in X$.
- 428 ■ Add a super sink t' and arcs (x^{out}, t') with capacity $\text{vio}(f, x)$ for every $x \in X$.

429 Note that $\{s, t\} \cup \bigcup_{x \in X} \{x^{\text{in}}, x^{\text{out}}\} \cup \{s', t'\}$ is an apex set of size $O(k)$ in H (recall from
 430 Remark 1 that s and t are the super source and super sink of the original graph G).

431 The circulation g^\times can be found using the following algorithm. Find a maximum flow h'
 432 from s' to t' in H using Lemma 7. Convert h' to an acyclic flow h of the same value using the
 433 algorithm of Sleator and Tarjan [12] (cf. [13, Lemma 2.5]). If h does not saturate every arc
 434 incident to s' and t' , return that the desired circulation g does not exist. Otherwise, h can
 435 be extended to the desired circulation g^\times by setting $g^\times(x^{\text{out}}, x^{\text{in}}) = h(x^{\text{out}}, x^{\text{in}}) + \text{vio}(f, x)$
 436 for every $x \in X$ and $g^\times(e) = h(e)$ for all other arcs.

437 The following lemma shows that any single Bulk-Push operation in the algorithm of
 438 Lemma 7 on H can be implemented by a constant number of calls to the $O(n \log^3 n)$ -time
 439 multiple-source multiple-sink maximum flow algorithm in planar graphs of Borradaile et
 440 al. [2]. There are two challenges that need to be overcome. First, the graph H is $O(k)$ -apex
 441 graph rather than planar. Second, the algorithm of Borradaile et al. computes a maximum
 442 flow from multiple sources to multiple sinks, not a maximum flow under the restriction that
 443 each source sends at most some given limit. This is not a problem in the case of a single
 444 source, or a limit on just the total value of the flow, since then some of the flow pushed can
 445 be "undone". When each of the multiple sources has a different limit, undoing the flow from
 446 one source can create residual paths from another source that did not yet reach its limit.

447 ► **Lemma 10.** *Any single Bulk-Push operation in the execution of the algorithm of Lemma 7*
 448 *on the graph H defined above can be implemented in $O(n \log^3 n)$ time.*

449 **Proof.** Let $V^\times = \{s, t\} \cup \bigcup_{x \in X} \{x^{\text{in}}, x^{\text{out}}\} \cup \{s', t'\}$ be the set of apices of H . Recall that
 450 the algorithm of Lemma 7 invokes the batch-highest-distance Push-Relabel algorithm on
 451 a complete graph K^\times over V^\times , and maintains a corresponding preflow in H . Consider a
 452 single Bulk-Push(U, W) operation from a set of apices U to a set of apices W . Let ρ denote

453 the preflow pushed in H up to this Bulk-Push operation. Let $A = U \cup W$. To correctly
 454 implement $\text{Bulk-Push}(U, W)$, we find a flow ρ' with sources U and sinks W in the graph H ,
 455 which satisfies the following properties:

- 456 (i) For every $u \in U$, $\text{ex}(\rho + \rho', u) \geq 0$, and
- 457 (ii) For every $u \in U$ and $w \in W$, either $\text{ex}(\rho + \rho', u) = 0$ or there is no residual path in
 458 $H_{\rho+\rho'}$ from u to w that is internally disjoint from V^\times .

459 Condition (i) guarantees that ρ' does not push more flow from a vertex $u \in U$ than the
 460 current excess of u . Condition (ii) is condition (5*) from Section 3.2.

461 Let H'' be the graph obtained from H_ρ by deleting the vertices $V^\times \setminus A$. Note that the
 462 absence of residual paths that are internally disjoint from V^\times in H'' is equivalent to the
 463 absence of such paths in H . We will compute ρ' using a constant number of invocations of
 464 the $O(n \log^3 n)$ -time multiple-source multiple-sink maximum flow algorithm in planar graphs
 465 of Borradaile et al. [2]. Instead of invoking this algorithm on H'' , which is not planar, we
 466 shall invoke it on modified versions of H'' which are planar.

467 Starting with H'' , we split each vertex $w \in W$ into $\deg(w)$ copies. Each arc e that was
 468 incident to w before the split is now incident to a distinct copy of w , and is embedded so
 469 that it does not cross any other arc in the graph. Let H' denote the resulting graph, and let
 470 W' denote the set of vertices created as a result of splitting all the vertices of W .

471 The set W' replaces W as the set of sinks of the flow ρ' we need to compute. Note that
 472 U is an apex set in H' . We then build the flow ρ' gradually, by computing the following
 473 steps, each using a single invocation of the multiple-source multiple-sink maximum flow
 474 algorithm of Borradaile et al. in $O(n \log^3 n)$ time. In what follows, when we say that the
 475 flow ρ' satisfies condition (ii) for a subset U' of U we mean that for every $u \in U'$ and $w \in W$,
 476 either $\text{ex}(\rho + \rho', u) = 0$ or there is no residual path in $H_{\rho+\rho'}$ from u to w that is internally
 477 disjoint from V^\times .

478 (1) If $s \in U$, starting with H' , we split s into $\deg(s)$ copies. Each arc e that was incident
 479 to s before the split is now incident to a distinct copy of s , and is embedded so that it
 480 does not cross any other arc in the graph. We also delete all other vertices of U . We
 481 compute in the resulting graph, which is planar, a maximum flow with sources the copies
 482 of s and the sinks W' . Let ρ'_s be the flow computed. If $|\rho'_s| > \text{ex}(\rho, s)$, we decrease $|\rho'_s|$
 483 by pushing $|\rho'_s| - \text{ex}(\rho, s)$ units of flow back from W' to the copies of s . This can be
 484 done in $O(n)$ time in reverse topological order w.r.t. ρ'_s (cf. [2, Section 1.4]). We view ρ'_s
 485 as a flow in H , and set $\rho' = \rho'_s$. By construction ρ' satisfies condition (i), and satisfies
 486 condition (ii) for the subset $\{s\}$.

487 (2) If $t \in U$, starting with $H'_{\rho'}$, we repeat step (1) with t taking the role of s to compute a
 488 flow ρ'_t . Set $\rho' \leftarrow \rho' + \rho'_t$. By construction of ρ'_t , ρ' now satisfies condition (i),
 489 and satisfies condition (ii) for the subset $U \cap \{s, t\}$.

490 (3) Let U^{in} be the set $U \cap \{x^{in} : x \in X\}$. If $U^{in} \neq \emptyset$, starting with $H'_{\rho'}$, we delete all the
 491 vertices of U that are not in U^{in} . Note that, since the resulting graph does not contain
 492 s, t, s', t' , nor any x^{out} for any $x \in X$, and since arcs incident to x^{in} only cross those
 493 incident to x^{out} , the resulting graph is planar. For every $x^{in} \in U^{in}$ we add a vertex x'
 494 and an arc (x', x^{in}) with capacity $\text{ex}(\rho, x^{in})$. The resulting graph is still planar. We
 495 compute a maximum flow ρ'_{in} with sources $\{x' : x^{in} \in U^{in}\}$ and sinks W' . We view ρ'_{in}
 496 as a flow in H , and set $\rho' \leftarrow \rho' + \rho'_{in}$. By construction of ρ'_{in} , ρ' now satisfies condition
 497 (i), and satisfies condition (ii) for the subset $U \cap (\{s, t\} \cup U^{in})$.

498 (4) We repeat step (3) with out taking the role of in to compute a flow ρ'_{out} . By construction
 499 of ρ'_{in} , ρ' now satisfies condition (i), and satisfies condition (ii) for $U \cap (\{s, t\} \cup U^{in} \cup U^{out})$.
 500 Since s' and t' are the source and sink of the flow computed by the Push-Relabel algorithm,

501 they are never active vertices, so they never belong to U . Hence $\{s, t\} \cup U^{in} \cup U^{out} \supseteq U$,
 502 and conditions (i) and (ii) are fully satisfied by ρ' . ◀

503 Combining Lemma 10 and Lemma 7, we get the following lemma.

504 ▶ **Lemma 11.** *The algorithm described above finds a circulation g^\times such that*

- 505 1. $f^\times + g^\times$ is feasible in G^\times .
 - 506 2. The restriction of $f^\times + g^\times$ to G has no violations at vertices of X .
 - 507 3. The restriction of $f^\times + g^\times$ to G has violation at most $(k - 2) \cdot \text{vio}(f)$ at any vertex in
 508 $V(G) \setminus X$.
- 509 in $O(k^2 n \log^3 n)$ time if such a circulation exists.

510 **Proof.** We first analyze the running time. Computing the graph H can be done in $O(n)$
 511 time. Computing the flow h' in H using the algorithm of Lemma 7 takes $O(k^2 \cdot T_{BP})$ time.
 512 By Lemma 10, $T_{BP} = O(n \log^3 n)$ for the graph H . Transforming h' to an acyclic flow h
 513 using the algorithm of Sleator and Tarjan [12] takes $O(n \log n)$ time. Finally, computing g^\times
 514 from h takes $O(n)$ time. Hence, the total time to compute g^\times is $O(k^2 n \log^3 n)$.

515 In order to prove the correctness of the algorithm, we will first prove that there exists a
 516 feasible flow h in H that saturates every arc incident to s' and t' if and only if there exists a
 517 circulation g^\times in G^\times that satisfies conditions (1) and (2) in the statement of the lemma.

518 (\Leftarrow) Assume the circulation g^\times exists in G^\times . Define a flow h in H as follows. For
 519 every arc $e \in E(H)$ not of the form (x^{out}, x^{in}) set $h(e) = g^\times(e)$. For every $x \in X$, set
 520 $h(x^{out}, x^{in}) = g^\times(x^{out}, x^{in}) - \text{vio}(f, x)$, $h(s, x^{in}) = \text{vio}(f, x)$ and $h(x^{out}, t) = \text{vio}(f, x)$. Since
 521 the restriction of $f^\times + g^\times$ to G has no violations on the vertices of x , $g^\times(x^{out}, x^{in}) \geq \text{vio}(f, x)$,
 522 so $h(x^{out}, x^{in}) \geq 0$ and h is a well defined flow. By definition, the flow h saturates every arc
 523 incident to s' and t' .

524 To show that h is feasible in H it is enough to show that $h(x^{out}, x^{in}) \leq c(x)$ for every
 525 $x \in X$ (on all other arcs h is feasible because g^\times is feasible in $G_{f^\times}^\times$). Let $x \in X$. Since $f^\times + g^\times$
 526 is feasible in G^\times , $g^\times(x^{out}, x^{in}) \leq f^\times(x^{in}, x^{out})$. Since $h(x^{out}, x^{in}) = g^\times(x^{out}, x^{in}) - \text{vio}(f, x)$,
 527 $h(x^{out}, x^{in}) \leq f^\times(x^{in}, x^{out}) - \text{vio}(f, x) = c(x)$.

528 (\Rightarrow) Assume there exist a feasible flow h in H that saturates every arc incident to s'
 529 and t' , and let g^\times be the circulation obtained from h as described above. We show that
 530 $f^\times + g^\times$ is feasible in G^\times . On all arcs e not of the form (x^{out}, x^{in}) , $g^\times(e) = h(e)$ and the
 531 capacity of e in $G_{f^\times}^\times$ equals the capacity of e in H . Therefore, since h is a feasible flow in
 532 H , g^\times is feasible on e in $G_{f^\times}^\times$, so $f^\times + g^\times$ is feasible on e in G^\times . We now focus on the
 533 arcs (x^{out}, x^{in}) for each $x \in X$. Let $e = (x^{out}, x^{in})$. Observe that $0 \leq h(e) \leq c(x)$. Since
 534 $g^\times(e) = h(e) + \text{vio}(f, x)$ we have that $\text{vio}(f, x) \leq g^\times(e) \leq c(x) + \text{vio}(f, x) = f^\times(e)$. Since
 535 $(f^\times + g^\times)(x^{in}, x^{out}) = f^\times(x^{in}, x^{out}) - g^\times(x^{out}, x^{in})$, we have $0 \leq (f^\times + g^\times)(x^{in}, x^{out}) \leq c(x)$,
 536 so $f^\times + g^\times$ is feasible in G^\times .

537 To finish proving the (\Rightarrow) direction, we show that the restriction of $f^\times + g^\times$ to G has
 538 no violations on the vertices of X . By definition of G^\times and of residual graph, the only arcs
 539 in $G_{f^\times}^\times$ that can carry flow out of x^{in} are the reverses of the arcs that carry flow into x in
 540 f , and the only arcs that can carry flow into x^{out} are the reverses of the arcs that carry
 541 flow out of x in f . We will show that $(f + g)^{in}(x) \leq c(x)$ by considering separately the
 542 contribution of the flow on arcs of G that in G^\times are incident to x^{out} , and arcs of G that in
 543 G^\times are incident to x^{in} .

544 The only arc of f^\times that carries flow into x^{out} is (x^{in}, x^{out}) . Thus, there is no arc e of G
 545 such that $f^\times(e)$ carries flow into x^{out} . Since g^\times only carries flow into x^{out} along the reverses
 546 of arcs that carry flow out of x^{out} in f^\times and since for every such arc e' , $g^\times(e') \leq f^\times(\text{rev}(e'))$,
 547 there is also no arc e of G such that $(f^\times + g^\times)(e)$ carries flow into x^{out} .

548 The total flow that f^\times carries into x^{in} is $c(x) + \text{vio}(f, x)$. Let z denote the total amount
 549 of flow that g^\times carries into x^{in} . Since the only arc incident to x^{in} that carries flow in g^\times
 550 and does not belong to G is (x^{out}, x^{in}) , the total amount of flow that g^\times carries into x^{in} on
 551 arcs that belong to G is $z - g^\times(x^{out}, x^{in})$. On the other hand, g^\times carries z units of flow out
 552 of x^{in} , and all of this flow is pushed along the reverses of arcs that carry flow into x^{in} in f^\times
 553 (and also belong to G). Hence, the total amount of flow that $f^\times + g^\times$ carries into x^{in} on
 554 arcs that belong to G is $c(x) + \text{vio}(f, x) + (z - g^\times(x^{out}, x^{in})) - z$. Therefore,

$$\begin{aligned}
 555 \quad (f + g)^{in}(x) &= c(x) + \text{vio}(f, x) - g^\times(x^{out}, x^{in}) \\
 556 &= c(x) + \text{vio}(f, x) - (h(x^{out}, x^{in}) + \text{vio}(f, x)) \\
 557 &= c(x) - h(x^{out}, x^{in}) \\
 558 &\leq c(x).
 \end{aligned}$$


559 We have thus shown that the algorithm computes a flow g^\times satisfying conditions (1) and
 560 (2) in the statement of the lemma. To see that condition (3) is also satisfied, note that the value
 561 of the flow h is $\sum_{x \in X} \text{vio}(f, x) \leq (k-2) \cdot \text{vio}(f)$. Since h is acyclic, $h^{in}(v) \leq (k-2) \cdot \text{vio}(f)$
 562 for all $v \in H$. Since for all $v \in V(G) \setminus X$, $f^{in}(v) \leq c(v)$, and $h^{in}(v) = (g^\times)^{in}(v)$, it follows
 563 that the violation of $f^\times + g^\times$ at v is at most $(k-2) \cdot \text{vio}(f)$. ◀

564
 565 Using the $O(k^2 n \log^3 n)$ -time algorithm of Lemma 11 in the improvement phase of Wang's
 566 algorithm instead of using Wang's $O(k^4 n \log^3 n)$ -time procedure for this phase results in a
 567 running time of $O(k^3 n \text{polylog}(nC))$ for finding a maximum flow in G .

568 4.3 The case $k = o(\log^2 n)$

569 As previously mentioned, we can use an alternative algorithm to the one of Lemma 11 for
 570 the case of $k = o(\log^2 n)$. This algorithm computes the circulation g^\times in $O(k^3 n \log n)$ time.
 571 We use the same auxiliary graph H as defined above and again compute a maximum flow h'
 572 from s' to t' in H . Let $V^\times = \{s, t\} \cup \bigcup_{x \in X} \{x^{in}, x^{out}\} \cup \{s', t'\} \cup S \cup T$ be the set of apices
 573 of H along with the original sources S and sinks T of G , and let K^\times be the complete graph
 574 on V^\times . As in the strategy of Borradaile et al. [2], we simulate a maximum flow computation
 575 from s' to t' in K^\times using the FIFO Push-Relabel algorithm. Each $\text{Push}(u, v)$ operation can be
 576 performed in only $O(n \log n)$ time. In short, we push along the arc (u, v) directly if u or v is
 577 one of s, t, s' , or t' . If $u \in \{x_1^{in}, x_1^{out}\}$ and $v \in \{x_2^{in}, x_2^{out}\}$ for two distinct vertices $x_1, x_2 \in X$,
 578 then we can directly push from u to v in a planar subgraph of H using the single-source
 579 single-sink maximum flow algorithm in planar graphs of Borradaile and Klein [1]. Finally, if
 580 $u, v \in \{x^{in}, x^{out}\}$ for some $x \in X$, we can split both u and v into several copies that all lie
 581 on a common face. We can then use a divide-and-conquer procedure of Miller and Naor [10]
 582 to perform the Push in $O(n \log n)$ time. See Appendix A for details. By using the better of
 583 the two procedures for computing g^\times , we get our main theorem.

584 ► **Theorem 12.** *A maximum flow in an n -vertex planar flow network G with integer arc and*
 585 *vertex capacities bounded by C can be computed in $O(k^3 n \log n \min(k, \log^2 n) \log(kC) \log(nC))$*
 586 *time.*



■ **Figure 2** Illustration of the auxiliary graphs used in the algorithm of Lemma 10. Only a portion of the graphs around some vertex $x \in X$ is shown. (a) the graph G° . (b) the graph H . Note that the only crossings are between arcs incident to x^{in} and arcs incident to x^{out} . (c) the graph H' in the case that x^{out} belongs to W . x^{out} is split into multiple copies, eliminating all arc crossings.

Appendices

587

A Alternative algorithm for $k = o(\log^2 n)$

588

589 We now provide an alternative algorithm to the one given in the previous section that is
 590 faster for small $k = o(\log^2 n)$. Specifically, we describe an algorithm for computing the
 591 circulation g^\times that runs in $O(k^3 n \log n)$ time instead of the $O(k^2 n \log^3 n)$ time required by
 592 the algorithm of Lemma 11. The final running time for computing a maximum flow with
 593 integer arc and vertex capacities is therefore $O(k^4 n \log n \log(kC) \log(nC))$.

594 We use the same auxiliary graph H as defined above and again compute a maximum
 595 flow h' from s' to t' in H . Let $V^\times = \{s, t\} \cup \bigcup_{x \in X} \{x^{in}, x^{out}\} \cup \{s', t'\} \cup S \cup T$ be the
 596 set of apices of H along with the original sources S and sinks T of G , and let K^\times be the
 597 complete graph on V^\times . Instead of using the batch-highest-distance Push-Relabel algorithm
 598 as in Lemma 7, we more directly follow the strategy of Borradaile et al. [2] by simulating a
 599 maximum flow computation from s to t in K^\times using the FIFO Push-Relabel algorithm. We
 600 do not wish to directly use the multiple-source multiple-sink flow algorithm of Borradaile et
 601 al. [2], because then each of the $O(k^3)$ Push operations would take $O(n \log^3 n)$ time. But as
 602 above, we may take advantage of the structure of H to perform each Push operation more
 603 quickly.

604 ► **Lemma 13.** *Any single (individual arc) Push operation in the graph H defined above can*
 605 *be implemented in $O(n \log n)$ time.*

606 **Proof.** Consider a single Push(u, v) operation where $u, v \in V^\times$. Let ρ denote the preflow
 607 pushed in H by the FIFO Push-Relabel algorithm up to this Push operation. We find a flow ρ'
 608 with source u and sink v in the graph H such that either $\text{ex}(\rho + \rho', u) = 0$ or $\text{ex}(\rho + \rho', u) > 0$
 609 and there is no residual path in $H_{\rho + \rho'}$ from u to v that is internally disjoint from V^\times .

610 Let H' be the graph obtained from H_ρ by deleting the vertices $V^\times \setminus \{u, v\}$. Instead
 611 of invoking the $O(n \log^3 n)$ -time multiple-source multiple-sink maximum flow algorithm of
 612 Borradaile et al. [2], we will compute ρ' as follows. As before, we must consider a few different
 613 cases.

614 ■ If $u = s$ or $v = s$, then $v \in S$ or $u \in S$, respectively. We push up to $\text{ex}(\rho, u)$ units of flow
 615 directly along the arc (u, v) in constant time, either saturating the arc or reducing the
 616 excess flow in u to 0. We may similarly push directly along the arc (u, v) in constant

- 617 time if one of u or v is one of t , s' , or t' instead.
- 618 ■ If $u \in \{x_1^{in}, x_1^{out}\}$ and $v \in \{x_2^{in}, x_2^{out}\}$ for two distinct vertices $x_1, x_2 \in X$, then the graph
 619 H' is planar. We add a vertex u' and an arc (u', u) with capacity $\text{ex}(\rho, u)$ and compute the
 620 maximum flow ρ' with source u' and sink v using the single-source single-sink maximum
 621 flow algorithm in planar graphs of Borradaile and Klein [1].
- 622 ■ If neither of the above cases apply, then $u, v \in \{x^{in}, x^{out}\}$ for some $x \in X$. If arc (u, v)
 623 has positive residual capacity, we push up to $\text{ex}(\rho, u)$ units of flow directly along it in
 624 constant time. Similar to Step 1 in the proof of Lemma 10, starting with H' , we split u
 625 into $\text{deg}(u)$ copies so that each arc that was incident to u is now incident to a distinct
 626 copy of u . Similarly, we split v into $\text{deg}(v)$ copies so each arc that was incident to v is
 627 now incident to a distinct copy of v . The resulting graph is planar, and all copies of u
 628 and v lie on a common face. As mentioned by Borradaile et al. [2, p. 1280], we can then
 629 plug the linear time shortest paths in planar graphs algorithm of Henzinger et al. [5] into
 630 a divide-and-conquer procedure of Miller and Naor [10] to compute a maximum flow ρ'_u
 631 with sources the copies of u and sinks the copies of v in $O(n \log n)$ time. Again, if the
 632 value of this flow is greater than the excess of u , we push the appropriate amount of flow
 633 back to the copies of u in $O(n)$ time. Finally, we view ρ'_u as a flow in H to set $\rho' = \rho'_u$.
 634 ◀

635 As a consequence of the previous lemma, we immediately get a variation of Lemma 11 with
 636 a running time of $O(k^3 n \log n)$. We use our $O(k^3 n \log n)$ time algorithm in the improvement
 637 phase of Wang's algorithm whenever $k = o(\log^2 n)$.

638 ——— **References** ———

- 639 **1** Glencora Borradaile and Philip N. Klein. An $O(n \log n)$ algorithm for maximum st -flow in a
640 directed planar graph. *J. ACM*, 56(2):9:1–9:30, 2009. doi:10.1145/1502793.1502798.
- 641 **2** Glencora Borradaile, Philip N. Klein, Shay Mozes, Yahav Nussbaum, and Christian Wulff-
642 Nilsen. Multiple-source multiple-sink maximum flow in directed planar graphs in near-linear
643 time. *SIAM J. Comput.*, 46(4):1280–1303, 2017. doi:10.1137/15M1042929.
- 644 **3** Joseph Cheriyan and S. N. Maheshwari. Analysis of preflow push algorithms for maximum
645 network flow. *SIAM J. Comput.*, 18(6):1057–1086, 1989. doi:10.1137/0218072.
- 646 **4** Andrew V. Goldberg and Robert Endre Tarjan. A new approach to the maximum-flow problem.
647 *J. ACM*, 35(4):921–940, 1988. doi:10.1145/48014.61051.
- 648 **5** Monika Rauch Henzinger, Philip N. Klein, Satish Rao, and Sairam Subramanian. Faster
649 shortest-path algorithms for planar graphs. *J. Comput. Syst. Sci.*, 55(1):3–23, 1997. doi:
650 10.1006/jcss.1997.1493.
- 651 **6** Jan M. Hochstein and Karsten Weihe. Maximum s - t -flow with k crossings in $O(k^3 n \log n)$
652 time. In Nikhil Bansal, Kirk Pruhs, and Clifford Stein, editors, *Proceedings of the Eighteenth*
653 *Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2007, New Orleans, Louisiana,*
654 *USA, January 7-9, 2007*, pages 843–847. SIAM, 2007. URL: <http://dl.acm.org/citation.cfm?id=1283383.1283473>.
- 655 **7** Haim Kaplan and Yahav Nussbaum. Maximum flow in directed planar graphs with vertex
656 capacities. *Algorithmica*, 61(1):174–189, 2011. doi:10.1007/s00453-010-9436-7.
- 657 **8** Tarun Kathuria, Yang P. Liu, and Aaron Sidford. Unit capacity maxflow in almost $O(m^{4/3})$
658 time. In *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020,*
659 *Durham, NC, USA, November 16-19, 2020*, pages 119–130. IEEE, 2020. doi:10.1109/
660 FOCS46700.2020.00020.
- 661 **9** Samir Khuller and Joseph Naor. Flow in planar graphs with vertex capacities. *Algorithmica*,
662 11(3):200–225, 1994. doi:10.1007/BF01240733.
- 663 **10** Gary L. Miller and Joseph Naor. Flow in planar graphs with multiple sources and sinks. *SIAM*
664 *J. Comput.*, 24(5):1002–1017, 1995. doi:10.1137/S0097539789162997.
- 665 **11** James B. Orlin. Max flows in $O(nm)$ time, or better. In Dan Boneh, Tim Roughgarden, and
666 Joan Feigenbaum, editors, *Symposium on Theory of Computing Conference, STOC'13, Palo*
667 *Alto, CA, USA, June 1-4, 2013*, pages 765–774. ACM, 2013. doi:10.1145/2488608.2488705.
- 668 **12** Daniel Dominic Sleator and Robert Endre Tarjan. A data structure for dynamic trees. *J.*
669 *Comput. Syst. Sci.*, 26(3):362–391, 1983. doi:10.1016/0022-0000(83)90006-5.
- 670 **13** Yipu Wang. Maximum integer flows in directed planar graphs with vertex capacities and
671 multiple sources and sinks. In Timothy M. Chan, editor, *Proceedings of the Thirtieth Annual*
672 *ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA,*
673 *January 6-9, 2019*, pages 554–568. SIAM, 2019. doi:10.1137/1.9781611975482.35.
- 674 **14** Xianchao Zhang, Weifa Liang, and Guoliang Chen. Computing maximum flows in undirected
675 planar networks with both edge and vertex capacities. In Xiaodong Hu and Jie Wang, editors,
676 *Computing and Combinatorics, 14th Annual International Conference, COCOON 2008, Dalian,*
677 *China, June 27-29, 2008, Proceedings*, volume 5092 of *Lecture Notes in Computer Science*,
678 pages 577–586. Springer, 2008. doi:10.1007/978-3-540-69733-6_57.
- 679