

Multiple-Source Multiple-Sink Maximum Flow in Directed Planar Graphs in Near-Linear Time

Glencora Borradaile
Oregon State University

Philip N. Klein
Brown University

Shay Mozes
Brown University

Yahav Nussbaum
Tel-Aviv University

Christian Wulff-Nilsen
Carleton University

May 11, 2011

Abstract

We give an $O(n \log^3 n)$ algorithm that, given an n -node directed planar graph with arc capacities, a set of source nodes, and a set of sink nodes, finds a maximum flow from the sources to the sinks. Previously, the fastest algorithms known for this problem were those for general graphs.

1 Introduction

The maximum flow problem with multiple sources and sinks in a directed graph with arc-capacities is, informally, to find a way to route a single commodity from a given set of sources to a given set of sinks such that the total amount of the commodity that is delivered to the sinks is maximum subject to each arc carrying no more than its capacity. In this paper we study this problem in planar graphs.

The study of maximum flow in planar graphs has a long history. In 1956, Ford and Fulkerson introduced the max st -flow problem, gave a generic augmenting-path algorithm, and also gave a particular augmenting-path algorithm for the case of a planar graph where s and t are on the same face. Researchers have since published many algorithmic results proving running-time bounds on max st -flow for (a) planar graphs where s and t are on the same face, (b) undirected planar graphs where s and t are arbitrary, and (c) directed planar graphs where s and t are arbitrary. The best bounds known are (a) $O(n)$ [14], (b) $O(n \log \log n)$ [17], and (c) $O(n \log n)$ [2], where n is the number of nodes in the graph.

Maximum flow in planar graphs with multiple sources and sinks was studied by Miller and Naor [29]. When it is known how much of the commodity is produced/consumed at each source and each sink, finding a consistent routing of flow that respects arc capacities can be reduced to *negative-length shortest paths*, which we now know can be solved in planar graphs in $O(n \log^2 n / \log \log n)$ time [30]. Otherwise, Miller and Naor gave an $O(n \log^{3/2} n)$ algorithm for the case where all the sinks and the sources are on the boundary of a single face, and generalized it to an $O(k^2 n^{3/2} \log^2 n)$ -time algorithm for the case where the sources and the sinks reside on the boundaries of k different faces.¹

However, the problem of maximum flow with multiple sources and sinks in planar graphs without any additional restrictions remained open. In general (i.e., non-planar) graphs, multiple sources and sinks can be reduced to the single-source single-sink case by introducing an artificial source and sink and connecting them to all the sources and sinks, respectively—but this reduction does not preserve planarity. For more than twenty years since the problem was explicitly stated and considered [29], the fastest known algorithm for computing multiple-source multiple-sink max-flow in a planar graph has been to use this reduction in conjunction with a general maximum-flow algorithm such as that of Sleator and Tarjan [32] which leads to a running time of $O(n^2 \log n)$. For integer capacities less than U , one could instead use the algorithm of Goldberg and Rao [10], which leads to a running time of $O(n^{1.5} \log n \log U)$. No planarity-exploiting algorithm was known for the problem.

Our Result The main result of this paper is an algorithm for the problem that is optimal up to a small poly-logarithmic factor.

Theorem 1.1. *There exists an algorithm that solves the maximum flow problem with multiple sources and sinks in an n -node directed planar graph in $O(n \log^3 n)$ time.*

Application to computer vision problems Multiple-source multiple-sink min-cut arises in addressing a family of problems associated with the terms *metric labeling* (Kleinberg and Tardos, [26]), *Markov Random Fields* [9], and *Potts Model* (see also [4, 15]). In low-level vision problems such as *image restoration*, *segmentation*, *stereo*, and *motion*, the goal is to assign labels from a set to pixels so as to minimize a penalty function. The penalty function is a sum of two parts. One part, the *data component*, has a term for each pixel; the cost depends on the discrepancy between the observed data for the pixel and the label chosen for it. The other part, the *smoothing component*, penalizes neighboring pixels that are assigned different labels.

For the *binary* case (when the set of available labels has size two), finding the optimal solution is reducible to multiple-source multiple-sink min-cut. [12]. For the case of more than two labels, there is a powerful and effective heuristic [4] using very-large-neighborhood [1] local search; the inner loop consists of solving the two-label case. The running time for solving the two-label case is therefore quite important. For this reason, researchers in computer vision have proposed new algorithms for max flow and done experimental studies comparing the run-times of different max-flow algorithms on the instances arising in this context [3, 27]. All of this is evidence for the importance of the problem.

¹The time bound of the first algorithm can be improved to $O(n \log n)$ using the linear-time shortest-path algorithm of Henzinger et al. [14], and the time bound of the second algorithm can be improved to $O(k^2 n \log^2 n)$ using the $O(n \log n)$ -time single-source single-sink maximum flow algorithm of Borradaile and Klein [2].

For the (common) case where the underlying graph of pixels is the two-dimensional grid graph, our result yields a theoretical speed-up for this important computer-vision subroutine.²

Hochbaum [15] describes a special case of the penalty function in which the data component is convex and the smoothing component is linear; in this case, she shows that an optimal solution can be found in time $O(T(m, n) + n \log U)$ where U is the maximum label, and $T(m, n)$ is the time for finding a minimum cut. She mentions specifically image segmentation, for which the graph is planar. For this case, by using our algorithm, the optimal solution can be found in nearly linear time

Application to maximum bipartite matching Consider the problem of maximum matching in a bipartite planar graph. It is well-known how to reduce this problem to multiple-source, multiple-sink maximum flow. Our result is the first planarity-exploiting algorithm for this problem (and the first near-linear one).

Techniques To obtain our result, we employ a wide range of sophisticated algorithmic techniques for planar graphs, some of which we adapted to our needs while others are used unchanged. Our algorithm uses pseudoflows [11, 16] and a divide-and-conquer scheme influenced by that of [19] and that of [29]. We adapt a method for using shortest paths to solve max st -flow when s and t are adjacent [13], and a data structure for implementing Dijkstra in a dense distance graph derived from a planar graph [8]. Among the other techniques we employ are: using cycle separators [28] recursively while keeping the boundary nodes on a constant number of faces [25, 33, 8], an algorithm for single-source single-sink max flow [2, 7], an algorithm for computing multiple-source shortest paths [23, 5], a method for cancelling cycles of flow in a planar graph [20], an algorithm for finding shortest paths in planar directed graphs with negative lengths [24, 30], and a data structure for range queries in a Monge matrix [21].

1.1 Preliminaries

We assume the reader is familiar with the basic definitions of planar embedded graphs and their duals (cf. [2]). Let $G = (V, E)$ be a planar embedded graph with node-set V and arc-set E . We use the term *arc* to emphasize that edges are directed. The term *edge* is used when the direction of an arc is not important. For each arc a in the arc-set E , we define two oppositely directed *darts*, one in the same orientation as a (which we sometimes identify with a) and one in the opposite orientation [2]. We define $\text{rev}(\cdot)$ to be the function that takes each dart to the corresponding dart in the opposite direction. It is notationally convenient to equate the edges, arcs and darts of G with the edges, arcs and darts of the dual G^* .

Let $S \subseteq V$ be a set of nodes called sources, and let $T \subseteq V - S$ be a set of nodes called sinks.

A *flow assignment* $\mathbf{f}(\cdot)$ is a real-valued function on darts satisfying *antisymmetry*:

$$\mathbf{f}(\text{rev}(d)) = -\mathbf{f}(d)$$

A *capacity assignment* $\mathbf{c}(\cdot)$ is a real-valued function on darts.

A flow assignment $\mathbf{f}(\cdot)$ *respects the capacity* of dart d if $\mathbf{f}(d) \leq \mathbf{c}(d)$. $\mathbf{f}(\cdot)$ is called a *pseudoflow* if it respects the capacities of all darts.

For a given flow assignment $\mathbf{f}(\cdot)$, the *net inflow* (or just *inflow*) node v is $\text{inflow}_{\mathbf{f}}(v) = \sum_{\text{dart } d: \text{head}(d)=v} \mathbf{f}(d)$ ³. The *outflow* of v is $\text{outflow}_{\mathbf{f}}(v) = -\text{inflow}_{\mathbf{f}}(v)$. A flow assignment $\mathbf{f}(\cdot)$ is said to *obey conservation* at node v if $\text{inflow}_{\mathbf{f}}(v) = 0$. A *feasible flow* is a pseudoflow that obeys conservation at every node other than the sources and sinks. A *feasible circulation* is a pseudoflow that obeys conservation at all nodes. The *value* of a feasible flow $\mathbf{f}(\cdot)$ is the sum of inflow at the sinks, $\sum_{t \in T} \text{inflow}_{\mathbf{f}}(t)$ or, equivalently, the sum of outflow at the sources. The maximum flow problem is that of finding a feasible flow with maximum value.

For two flow assignments \mathbf{f}, \mathbf{f}' , the *addition* $\mathbf{f} + \mathbf{f}'$ is the flow that assigns $\mathbf{f}(d) + \mathbf{f}'(d)$ to every dart d .

A *residual path* in G is a path whose darts all have strictly positive capacities. For two sets of nodes A, B , $A \xrightarrow{G} B$ is used to denote the existence of some residual a -to- b path in G for some nodes $a \in A$ and $b \in B$. Conversely, $A \not\xrightarrow{G} B$ is used to denote that no such path exists. We will omit the graph G when it is clear from the context.

²Note that the single-source, single-sink max-flow algorithm of [2] was implemented by computer-vision researchers [31] and found to be useful in computer vision and to be faster than the competitors.

³An equivalent definition, in terms of arcs, is $\text{inflow}_{\mathbf{f}}(v) = \sum_{a \in E: \text{head}(a)=v} \mathbf{f}(a) - \sum_{a \in E: \text{tail}(a)=v} \mathbf{f}(a)$.

The *residual graph* of G with respect to a flow assignment $\mathbf{f}(\cdot)$ is the graph $G_{\mathbf{f}}$ with the same arc-set, node-set, sources and sinks, and with capacity assignment $\mathbf{c}_{\mathbf{f}}(\cdot)$ such that for every dart d , $\mathbf{c}_{\mathbf{f}}(d) = \mathbf{c}(d) - \mathbf{f}(d)$. It is well known that a feasible flow \mathbf{f} in G is maximum if and only if $S \stackrel{G_{\mathbf{f}}}{\not\rightarrow} T$.

Let \mathbf{f} be a pseudoflow in a planar graph G . Let V^+ denote the set of nodes $\{v \in V - (S \cup T) : \text{inflow}_{\mathbf{f}}(v) > 0\}$. Similarly, let V^- denote the set of nodes $\{v \in V - (S \cup T) : \text{inflow}_{\mathbf{f}}(v) < 0\}$. Suppose $S \cup V^+ \stackrel{G_{\mathbf{f}}}{\not\rightarrow} T \cup V^-$. For a graph with n nodes and m edges, there exists an $O(m \log n)$ -time algorithm that converts the pseudoflow \mathbf{f} into a maximum feasible flow \mathbf{f}' [19, 32]. In planar graphs, this can be done in linear time by first canceling flow cycles using the technique of Kaplan and Nussbaum [20], and then by sending back flow from V^+ and into V^- in topological sort order. See Appendix C for details.

1.2 Overview of the algorithm

Consider the following recursive approach for finding a maximum multiple-source multiple-sink flow: split the input graph G in two using a simple cycle separator C [28] and recursively solve the max flow problem in the two subgraphs. When the two recursive calls have been executed, in each of the two subgraphs there is no residual path from any source to any sink. If we further make sure that in each of the two subgraphs there is no residual path from any source to C and from C to any sink, then, since C is a separator, there is no residual path from any source to any sink in G .

We therefore solve a slightly more general problem recursively in the two subgraphs: roughly speaking, find a flow such that there is no residual path from a source to a sink or to C and no residual path from C to a sink (Section 2). After the two recursive calls there is no residual path from any source to any sink in G . However, the requirement that there is no residual path from any source to C and from C to any sink cannot be achieved by a feasible flow but rather by a pseudoflow in which there might be excess inflow or excess outflow on nodes of C . We deal with this by solving a new max flow problem where nodes of C are treated as sources and sinks, limited in supply/demand by their excess inflow/outflow (Section 3).

We exploit a relation between primal circulations and dual shortest paths to maintain a succinct representation of the flow during critical phases of the algorithm, using the fact that there are only $O(\sqrt{n})$ sources and sinks, all cyclically ordered on C (Section 3.2). Even though our representation does not explicitly store the flow on nearly any arc in the graph, we can augment it efficiently towards optimality while maintaining feasibility. An important tool we use is Fakcharoenphol and Rao's efficient implementation of Dijkstra's algorithm [8], which we adapt to our needs.

The resulting pseudoflow can then be turned into a max flow in linear time using existing techniques (Appendix C). This leads to an $O(n \log^3 n)$ time algorithm for max flow.

2 The Algorithm

Now we describe the algorithm, referred to as MULTIPLESOURCEMULTIPLESINKMAXFLOW, in more detail. In order to treat nodes of the cycle separator both as sources and as sinks in recursive calls, we introduce a new node set A . At the top recursion level, $A = \emptyset$. In general, A has constant size; more precisely $|A| \leq 6$.

MULTIPLESOURCEMULTIPLESINKMAXFLOW(G, \mathbf{c}, S, T, A)

Input: a directed planar graph G with non-negative capacities \mathbf{c} , a set S of source nodes, a set T of sink nodes, a set A of at most six nodes

Output: a pseudoflow \mathbf{f} obeying conservation everywhere but S, T, A and s.t. $S \stackrel{G_{\mathbf{f}}}{\not\rightarrow} T$, $S \stackrel{G_{\mathbf{f}}}{\not\rightarrow} A$, $A \stackrel{G_{\mathbf{f}}}{\not\rightarrow} T$.

The algorithm finds a simple cycle separator C [28] and contracts all edges of C except one. This, essentially, merges all the nodes of C into a single supernode v , and turns C into a self loop. For simplicity of presentation we assume that no sources or sinks lie on C .⁴ The algorithm then recursively solves the problem on the subgraphs enclosed and not enclosed by that self loop (the self loop itself need not be included in any of the subgraphs), adding the supernode v that represents C to the set A . In order to keep

⁴This does not lose generality since if a node $u \in C$ is a source, one can introduce a new node u' and an arc $u'u$ whose capacity equals the sum of capacities of arcs outgoing from u . Now consider u' as a source instead of u . Since the separator has just $O(\sqrt{n})$ nodes, this will not affect the running time. Sinks can be handled in a similar fashion.

the cardinality of A at most six, the algorithm alternately applies the cycle separator theorem with weights uniformly distributed on all nodes and uniformly distributed on only the nodes of A . This technique is similar to that used in [25, 33, 8].

After the recursive calls, the algorithm uncontracts the edges of C . At this stage there are no residual paths between sources and sinks in the entire graph, but there might be excess inflow (positive or negative) at the nodes of C . The algorithm then calls the procedure `FIXCONSERVATIONONPATH` that pushes flow between the nodes of C so that there are no residual paths between nodes of C with positive inflow and nodes of C with negative inflow (the path in the name of the procedure is the cycle C without one edge). This procedure is discussed in Section 3; the interface is:

`FIXCONSERVATIONONPATH`($G, P, \mathbf{c}, \mathbf{f}_0$)

Input: a directed planar graph G , simple path P , capacity function \mathbf{c} , and a pseudoflow \mathbf{f}_0

Output: a pseudoflow \mathbf{f} s.t. (i) $\mathbf{f} - \mathbf{f}_0$ satisfies conservation everywhere but P , and

(ii) $\{v \in P : \text{inflow}(v) > 0\} \stackrel{G_{\mathbf{f}}}{\not\rightarrow} \{v \in P : \text{inflow}(v) < 0\}$.

Running Time: $O(n \log^2 n / \log \log n + |P|^2 \log^2 n)$

Next, the algorithm iterates over the nodes a_i of A . The algorithm calls the procedure `CYCLETOSINGLESLIMITEDMAXFLOW` that, roughly speaking, pushes as much excess flow as possible from C to a_i . If C_i^+ is the set of nodes of C that are reachable via residual paths from some node of C with positive inflow at the beginning of iteration i , `CYCLETOSINGLESLIMITEDMAXFLOW` pushes flow among the nodes of C_i^+ and from the nodes of C_i^+ to a_i . The result is that remaining inflow at nodes of C_i^+ is non-negative and there are no residual paths from nodes of C with positive inflow to a_i . See Section 4; the interface is:

`CYCLETOSINGLESLIMITEDMAXFLOW`($G, \mathbf{c}, \mathbf{f}_0, C, a_i$)

Input: a directed planar graph G with capacities \mathbf{c} , a pseudoflow \mathbf{f}_0 , a simple cycle C , a sink a_i .

Assumes: $\forall v \in C^+, \text{inflow}_{\mathbf{f}_0}(v) \geq 0$, where $C^+ = \{v \in C : \{x \in C : \text{inflow}_{\mathbf{f}_0}(x) > 0\} \stackrel{G_{\mathbf{f}_0}}{\rightarrow} v\}$.

Output: a pseudoflow \mathbf{f} s.t. (i) $\mathbf{f} - \mathbf{f}_0$ obeys conservation everywhere but $C^+ \cup \{t\}$, (ii) $\forall v \in C^+, \text{inflow}_{\mathbf{f}}(v) \geq 0$, (iii) $\{v \in C : \text{inflow}_{\mathbf{f}}(v) > 0\} \stackrel{G_{\mathbf{f}}}{\not\rightarrow} a_i$.

Running Time: $O(n \log^2 n / \log \log n + |C|^2 \log^2 n)$.

A similar procedure `SINGLESOURCECYCLELIMITEDMAXFLOW` is called to push flow from a_i to C to eliminate as much negative inflow as possible (using a similarly defined set C_i^-).

Finally, the algorithm pushes back flow from any nodes of C with positive inflow to S and pushes flow back from T into any nodes of C with negative inflow.

Algorithm 1 `MULTIPLESOURCEMULTIPLESINKMAXFLOW`(G, \mathbf{c}, S, T, A)

Input: a directed planar graph G with non-negative capacities \mathbf{c} , a set S of source nodes, a set T of sink nodes, a set A consisting of a constant number of nodes $A = \{a_i\}_{i=1}^k$.

Output: a pseudoflow \mathbf{f} obeying conservation everywhere except S, T, A and s.t. $S \stackrel{G_{\mathbf{f}}}{\not\rightarrow} T$, $S \stackrel{G_{\mathbf{f}}}{\not\rightarrow} A$, $A \stackrel{G_{\mathbf{f}}}{\not\rightarrow} T$.

- 1: add zero capacity arcs to triangulate and 2-connect G (required for simple cycle separators)
 - 2: find a balanced (w.r.t. $|G|$ and $|A|$, alternately) cycle separator C in G disjoint from S and T
 - 3: let P be a path comprising of all of C 's edges except one edge e
 - 4: contract all the edges of P , turning e into a self loop incident to the only remaining node v of C
 - 5: let G_1 and G_2 be the subgraph of G enclosed and not enclosed by e , respectively
 - 6: $\mathbf{f} := \text{MULTIPLESOURCEMULTIPLESINKMAXFLOW}(G_1, \mathbf{c}_{|G_1}, S \cap G_1, T \cap G_1, (A \cap G_1) \cup \{v\})$
 - 7: $\mathbf{f} := \mathbf{f} + \text{MULTIPLESOURCEMULTIPLESINKMAXFLOW}(G_2, \mathbf{c}_{|G_2}, S \cap G_2, T \cap G_2, (A \cap G_2) \cup \{v\})$
 - 8: uncontract the edges of P
 - 9: $\mathbf{f} := \text{FIXCONSERVATIONONPATH}(G, P, \mathbf{c}, \mathbf{f})$
 - 10: **for** $i = 1, 2, \dots, k$
 - 11: $\mathbf{f} := \text{CYCLETOSINGLESLIMITEDMAXFLOW}(G, \mathbf{c}, \mathbf{f}, C, a_i)$
 - 12: $\mathbf{f} := \text{SINGLESOURCECYCLELIMITEDMAXFLOW}(G, \mathbf{c}, \mathbf{f}, a_i, C)$
 - 13: push positive excess from C to S and negative excess to C from T
 - 14: **return** \mathbf{f}
-

Correctness of Algorithm 1 The correctness of Algorithm 1 is proved in detail in Appendix A. The proof consists of a sequence of lemmas that prove that each step of the algorithm eliminates some undesired residual paths without reintroducing undesired residual paths. The arguments used are elementary.

Running Time Analysis The number of nodes of the separator cycle C used to partition G into G_1 and G_2 is $O(\sqrt{|G|})$. Therefore, each invocation of `FIXCONSERVATIONONPATH`, `CYCLETOSINKLIMITEDMAXFLOW` and `SINGLESTOUCYCLELIMITEDMAXFLOW` in G takes $O(|G| \log^2 |G| / \log \log |G| + |C|^2 \log^2 |C|) = O(|G| \log^2 |G|)$ time.

The way we recursively partition into subgraphs is very similar to that of Fakcharoenphol and Rao [8]. In their algorithm, they spend $O(|G'| \log^2 |G'|)$ time on each subgraph G' in the recursive decomposition and prove a total time bound of $O(n \log^3 n)$. By the same arguments, our algorithm runs in $O(n \log^3 n)$ time.

3 Eliminating Residual Paths Between Nodes with Positive Inflow and Nodes with Negative Inflow on a Path

In this section we present an efficient implementation of the fixing procedure which, roughly speaking, given a path with nodes having positive, negative, or zero inflow, pushes flow between the nodes of the path so that eventually there are no residual paths from nodes with positive inflow to nodes with negative inflow.

We begin by describing an abstract algorithm for the fixing procedure. The abstract algorithm is given as Algorithm 2. It is similar to a technique used by Venkatesan and Johnson [19]. Let M be the sum of capacities of all of the darts of G . The algorithm first increases the capacities of darts of the path P and their reverses by M . Let p_1, p_2, \dots, p_{k+1} be the nodes of P . The algorithm processes the nodes of P one after the other. Processing p_i consists of decreasing the capacities of $d_i = p_i p_{i+1}$ and $\text{rev}(d_i)$ by M (i.e., back to their original capacities), and trying to eliminate positive inflow x at p_i by pushing at most x units of flow from p_i to p_{i+1} . The intuition for doing so is that the flow after the push either obeys conservation at p_i or there are no residual paths from p_i to any of the other nodes of P (this is where we use the large capacities on the darts between unprocessed nodes). See appendix B for a formal proof of correctness. Negative inflow at p_i is handled in a similar manner by pushing flow from p_{i+1} to p_i .

Algorithm 2 `ABSTRACTFIXCONSERVATIONONPATH($G, P, \mathbf{c}, \mathbf{f}_0$)`

Input: directed planar graph G , simple path $P = d_1 d_2 \dots d_k$, capacity function \mathbf{c} , and pseudoflow \mathbf{f}_0

Output: a pseudoflow \mathbf{f}' s.t. (i) $\mathbf{f}' - \mathbf{f}_0$ satisfies conservation at nodes not on P , and (ii) with respect to \mathbf{f}' , there are no residual paths from nodes of P with positive inflow to nodes of P with negative inflow.

```

1:  $\mathbf{f}' = \mathbf{f}_0$ 
2:  $\mathbf{c}[d] = \mathbf{c}[d] + M$  for all darts  $d$  of  $P \cup \text{rev}(P)$ 
3: for  $i = 1, 2, \dots, k$ 
4:   let  $p_i$  and  $p_{i+1}$  be the tail and head of  $d_i$ , respectively
      % reduce the capacities of  $d$  and  $\text{rev}(d)$  by  $M$  and adjust the flow appropriately
5:   for  $d \in \{d_i, \text{rev}(d_i)\}$ 
6:      $\mathbf{c}[d] := \mathbf{c}[d] - M$ 
7:      $\mathbf{f}'[d] := \min\{\mathbf{f}'[d], \mathbf{c}[d]\}$ 
8:      $\mathbf{f}'[\text{rev}(d)] := -\mathbf{f}'[d]$ 
9:    $\text{excess} := \text{inflow at } p_i$ 
10:  if  $\text{excess} > 0$  then  $d := d_i$  else  $d := \text{rev}(d_i)$  find in which direction flow should be pushed
11:  add to  $\mathbf{f}'$  a maximum feasible flow from tail( $d$ ) to head( $d$ ) with limit  $\text{excess}$ 
12: return  $\mathbf{f}'$ 

```

3.1 An Inefficient Implementation

In this section, we give an *inefficient* implementation of line 11 of the abstract algorithm. This will facilitate the explanation of the efficient procedure in the next section. We first review the necessary ideas and tools.

3.1.1 Hassin's algorithm for maximum st -planar flow

An st -planar graph is a planar graph in which nodes s and t are incident to the same face. Hassin [13] gave an algorithm for computing a maximum flow from s to t in an st -planar graph. We briefly describe this algorithm here since we use it in implementing line 11 of Algorithm 2.

Hassin's algorithm starts by adding to G an artificial infinite capacity arc a from t to s . Let d be the dart that corresponds to a and whose head is t . Let t^* be the head in G^* of the dual of d . Compute in the dual G^* a shortest path tree rooted at t^* , where the length of a dual dart is defined as the capacity of the primal dart. Let $\phi[\cdot]$ denote the shortest path distances from t^* in G^* . Consider the flow

$$\rho[d'] = \phi[\text{head}_{G^*}(d')] - \phi[\text{tail}_{G^*}(d')] \text{ for all darts } d' \quad (1)$$

After removing the artificial arc a from G , ρ is a maximum feasible flow from s to t in G . We say that ϕ is a *face potential* vector that induces ρ .

In our algorithm we are interested in a *max flow with limit x* from s to t rather than a maximum flow, i.e., a flow whose value is at most a given number x but is otherwise maximal. It is not difficult to see that setting the capacity of the artificial arc a to x instead of infinity results in the desired limited max flow [22].

In our implementation, instead of using an artificial arc from t to s , we use an existing arc whose endpoints are s and t as the arc a above. In order for this to work, the capacity of the dart d that corresponds to a and whose head is t must be zero (as is indeed the case if a is an artificial arc from t to s). This can always be achieved by first pushing flow on d to saturate it. Also note that in this case, we do not remove a from G . Hence, ρ is a feasible circulation, rather than a maximum flow, since flow is being pushed back from t to s along a . To convert it into a maximum flow one just has to undo the flow on a . If we define \mathbf{f} by

$$\mathbf{f}[d'] = \begin{cases} -\rho[d'] & \text{if } d' \text{ corresponds to } a \\ 0 & \text{otherwise} \end{cases}, \quad (2)$$

then $\rho + \mathbf{f}$ is a maximum feasible st -flow. We will use the fact that this flow can be represented implicitly by the face potential vector ϕ , and the flow values $\mathbf{f}[d']$ for the two darts corresponding to a .

3.1.2 The Inefficient Implementation

Recall that \mathbf{f}_0 is the flow at the beginning of the procedure. Observe that the change to the flow in iteration i of the abstract algorithm (line 11) is a flow between the endpoints of the dart d_i . As discussed in Section 3.1.1, this flow can be represented as the sum of (i) a circulation ρ_i and (ii) a flow on d_i and $\text{rev}(d_i)$. Summing over the first i iterations, the flow \mathbf{f}' at that time can be represented as the sum

$$\mathbf{f}' = \rho + \mathbf{f} \quad (3)$$

where $\rho = \sum_{j=1}^i \rho_j$ is a circulation and \mathbf{f} is a flow assignment that differs from \mathbf{f}_0 only on the darts of $\{d_j\}_{j=1}^i$ and their reverses.

The inefficient implementation of line 11 of the abstract algorithm appears as Algorithm 3. We now describe it. The total flow \mathbf{f}' is maintained by representing \mathbf{f} and the circulation ρ as in Eq. (3). \mathbf{f} is represented explicitly, but, in preparation for the efficient implementation, the circulations ρ_j are represented implicitly by the face-potentials ϕ_j . By linearity of Eq. (1), the sum $\phi = \sum_{j=1}^i \phi_j$ is the face potential vector that induces the circulation ρ .

Recall that d is the dart of C such that flow needs to be sent from $\text{tail}(d)$ to $\text{head}(d)$ (line 10 of Algorithm 2). In lines 1 – 3, the procedure pushes as much as possible on d itself. Consequently, either d is saturated or conservation at p_i is achieved.

Next, an implementation of Hassin's algorithm pushes a maximum flow with limit $|\text{inflow}(p_i)|$ from $\text{tail}(d)$ to $\text{head}(d)$. The procedure first sets the length of darts to their residual capacities (line 4) and sets the length

of $\text{rev}(d)$ to be the flow limit $|\text{inflow}(p_i)|$ (line 5). Since the flow maintained is feasible, all residual capacities are non-negative. The procedure then computes all the $\text{head}_{G^*}(d)$ -to- f distances $\phi_i[f]$ in G^* using Dijkstra's algorithm (line 6). Let ρ_i be the circulation corresponding to the face-potential vector ϕ_i . The procedure sets val equal to $\rho[d]$ in line 7, then subtracts val from $\mathbf{f}[d]$ and adds it to $\mathbf{f}[\text{rev}(d)]$. Finally, in the last line, the current circulation is added to the accumulated circulation by adding the potential ϕ_i to ϕ .

Algorithm 3 Inefficient Implementation of line 11 of ABSTRACTFIXCONSERVATIONONPATH (Algorithm 2)

```

% push flow on  $d$  to make its residual capacity zero as required for Hassin's algorithm
1:  $residual\ capacity := \mathbf{c}[d] - \mathbf{f}[d] - \phi[\text{head}_{G^*}(d)] + \phi[\text{tail}_{G^*}(d)]$ 
2:  $val := \min\{residual\ capacity, |\text{inflow}(p_i)|\}$  amount of flow to push on  $d$ 
3:  $\mathbf{f}[d] := \mathbf{f}[d] + val$  ;  $\mathbf{f}[\text{rev}(d)] := -\mathbf{f}[d]$ 

% push excess inflow from  $\text{tail}(d)$  to  $\text{head}(d)$  using Hassin's algorithm
4: let  $\ell[d'] := \mathbf{c}[d'] - \mathbf{f}[d'] - \phi[\text{head}_{G^*}(d')] + \phi[\text{tail}_{G^*}(d')]$  for all darts  $d' \in G$  lengths are residual capacities
5:  $\ell[\text{rev}(d)] := |\text{inflow}(p_i)|$  set the limit on the residual capacity of  $\text{rev}(d)$ 
6:  $\phi_i(\cdot) := \text{DIJKSTRA}(G^*, \ell, \text{head}_{G^*}(d))$  face potential are distances in  $G^*$  from  $\text{head}_{G^*}(d)$  w.r.t. residual capacities
7:  $val := \phi_i[\text{head}_{G^*}(d)] - \phi_i[\text{tail}_{G^*}(d)]$  the amount of flow assigned to  $d$  by the circulation corresponding to  $\phi_i$ 
8:  $\mathbf{f}[d] := \mathbf{f}[d] - val$  ;  $\mathbf{f}[\text{rev}(d)] := -\mathbf{f}[d]$  do not push the circulation on  $d$  and  $\text{rev}(d)$ 
9:  $\phi = \phi + \phi_i$  accumulate the current circulation

```

3.2 An Efficient Implementation

In this section we give an *efficient* implementation of Algorithm 2. We first review the necessary tools.

3.2.1 Fakcharoenphol and Rao's Efficient Dijkstra Implementation

Let X be a set of nodes. Let H be a planar graph in which the nodes of X lie a single face. Let x_1, x_2, \dots be the clockwise order of the nodes of X on that face. Let P be a set of darts not necessarily in the graph H whose endpoints are nodes in X . Fakcharoenphol and Rao [8] described a data structure that can be used in a procedure that efficiently implements Dijkstra's algorithm in $H \cup P$. The procedure takes as input a table D such that $D[i, j]$ stores the distance between x_i and x_j in H , an array ℓ that stores the lengths of the darts in P , and a node $v \in X$. It is assumed that the lengths in D and in ℓ are non-negative. The procedure outputs the distances of the nodes of X from v in $H \cup P$ in $O(|X| \log^2 |X| + |P| \log |X|)$ -time.

We mention a technical issue whose importance will become apparent in the sequel. The procedure partitions the table D into several subtables $\{D_\alpha\}_\alpha$ that correspond to distances between pairs of disjoint sets of nodes of X , where each set consists of nodes that are consecutive in X . It is assumed in [8, footnote on p. 884] that for each such subtable D_α , a data structure that supports range minimum queries of the form $\min_{j_1 \leq j \leq j_2} \{D_\alpha[i, j]\}$ for every i, j_1, j_2 is given. Fakcharoenphol and Rao note that such a data structure can be easily implemented by using a range-search tree [6] for every row i of D_α . The time required to construct all of the range-search trees for D_α is proportional to the size of D_α .

3.2.2 Price Functions, Reduced Lengths and FR-Dijkstra

For a directed graph G with dart-lengths $\ell(\cdot)$, a *price function* is a function ϕ from the nodes of G to the reals. For a dart d , the *reduced length with respect to ϕ* is $\ell_\phi(d) = \ell(d) + \phi(\text{tail}(d)) - \phi(\text{head}(d))$. A *feasible price function* is a price function that induces nonnegative reduced lengths on *all* darts of G (see [18]).

Single-source distances form a feasible price function. Suppose that, for some node r of G , for every node v of G , $\phi(v)$ is the r -to- v distance in G with respect to $\ell(\cdot)$. Then for every arc uw , $\phi(v) \leq \phi(u) + \ell(uw)$, so $\ell_\phi(uw) \geq 0$. Here we assume, without loss of generality, that all distances are finite (i.e., that all nodes are reachable from r) since we can always add arcs with sufficiently large lengths to make all nodes reachable without affecting the shortest paths in the graph.

We will use the following variant of Fakcharoenphol and Rao's efficient Dijkstra implementation. The procedure $\text{FR}(D, \ell, \phi^X, v)$ takes as input the table D and the array ℓ as described above, as well as a feasible price function ϕ^X on the nodes of X and a node $v \in X$. It outputs the distances of the nodes of X from v

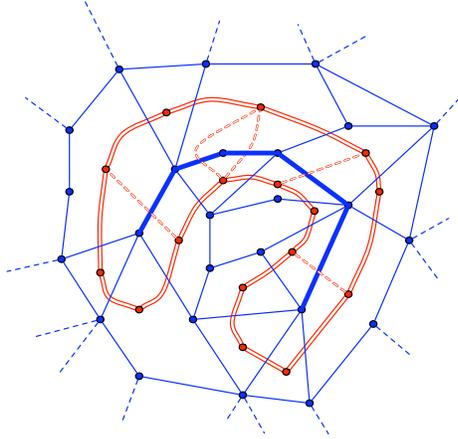


Figure 1: An example illustrating that the nodes of X^* are on the boundary of a single face of H^* . The diagram shows part of the graph G and some edges of its dual G^* . Edges of G are solid blue. Edges of P are bold. Dual edges are double lined red. The dual edges of P are in double lined dashed red.

in $H \cup P$ w.r.t. the reduced lengths w.r.t. ϕ^X . We stress that lengths in D and in ℓ may be negative, but the reduced lengths are all non-negative. The computation takes $O(|X| \log^2 |X| + |P| \log |X|)$ time.

This differs from the procedure described in Section 3.2.1 only in the existence of the price function ϕ^X . We cannot afford to compute the entire table of reduced lengths since that would dominate the running time of the algorithm in [8]. Instead, whenever their algorithm requires some specific reduced length, we can compute it in constant time from D . This, however, does not suffice. Recall that the algorithm in [8] requires that, for each of the subtables D_α , range-search trees that support range minimum queries of the form $\min_{j_1 \leq j \leq j_2} \{D_\alpha[i, j] + \phi^X[x_i] - \phi^X[x_j]\}$ for every i, j_1, j_2 are given. Note that the results of such queries may be different for different price functions. Computing the range-search trees would take $O(|X|^2)$ which will dominate the running time of the entire procedure. To overcome this difficulty we use Monge range-query data structures, due to Kaplan and Sharir [21], which can be constructed from the table D in $O(|X| \log |X|)$ time, and answer queries of the desired form in $O(\log |X|)$ time.

3.2.3 The Procedure

Finally, we describe the efficient implementation. The procedure keeps track of the inflow at each node of P in an array $v[\cdot]$. As in the inefficient implementation, the procedure will maintain the total flow as the sum of a circulation ρ and a flow assignment f that differs from f_0 only on the darts of $P \cup \text{rev}(P)$. Initially f is set equal to f_0 . The circulation ρ will be represented by a face-potential vector ϕ . However, we will show that it suffices to maintain just those entries of ϕ that correspond to faces incident to P .

Define each dart d 's length by $\ell(d) = c[d] - f[d]$. Let X^* be the set of endpoints in the planar dual G^* of the darts of P (i.e. the primal faces incident to P). Let H^* be the graph obtained from G^* by removing the darts of P . Note that in H^* , all the nodes of X^* that did not disappear (i.e., that have degree greater than zero) are on the boundary of a single face, see Figure 1.

The procedure uses a multiple-source shortest paths (MSSP) algorithm [23, 5] to compute a table $D^*[\cdot, \cdot]$ of X^* -to- X^* distances in H^* with respect to the lengths $\ell(\cdot)$.

The implementation of lines 1– 10 of the abstract algorithm using the chosen representation of f' is straightforward. We therefore focus on the implementation of line 11 of Algorithm 2, given as Algorithm 4.

Consider iteration i of the algorithm. The main difference between the inefficient implementation and the efficient one is in implementing the Dijkstra step for computing shortest paths in the dual. Instead of computing the entire shortest-path tree, the procedure computes just the distance labels of nodes in X^* . This is done using the FR data structure, whose initialization requires the X^* -to- X^* distances in H^* with respect to the residual capacities. We now explain how these distances can be obtained.

The flow on a dart d in the primal is

$$f[d] + \phi[\text{head}_{G^*}(d)] - \phi[\text{tail}_{G^*}(d)]. \quad (4)$$

Therefore the residual capacity of d is

$$(c[d] - f[d]) - \phi[\text{head}_{G^*}(d)] + \phi[\text{tail}_{G^*}(d)] \quad (5)$$

which is its *reduced* length $\ell_\phi[d]$ with respect to the length $\ell(\cdot)$ and price function ϕ .

Suppose that d belongs to H^* , i.e. d is not one of the darts of $P \cup \text{rev}(P)$. The procedure never changes $\mathbf{f}[d]$, so $\mathbf{f}[d] = \mathbf{f}_0[d]$. Therefore $\ell(d) = \mathbf{c}[d] - \mathbf{f}_0[d]$. These lengths are known at the beginning of the procedure's execution. The reduced length of an X^* -to- X^* path $Q = d'_1, d'_2, \dots, d'_j$ in H^* is

$$\sum_{i=1}^j (\ell(d'_i) - \phi[\text{head}_{G^*}(d'_i)] + \phi[\text{tail}_{G^*}(d'_i)]) = \left(\sum_{i=1}^j \ell(d'_i) \right) - \phi[\text{end}(Q)] + \phi[\text{start}(Q)]. \quad (6)$$

Therefore, for any nodes $x, y \in X^*$, the x -to- y distance in H^* w.r.t. the residual capacity is given by $D^*[x, y] - \phi[y] + \phi[x]$. Since the procedure maintains the restriction of ϕ to faces of X^* , this distance can be obtained in constant time. Adapting FR's data structure to handle reduced distances w.r.t. a price function ϕ was discussed in Section 3.2.1.

Algorithm 4 Efficient Implementation of line 11 of ABSTRACTFIXCONSERVATIONONPATH (Algorithm 2)

```

% push flow on  $d$  to make its residual capacity zero as required for Hassin's algorithm
1:  $\text{residual\_capacity} := \mathbf{c}[d] - \mathbf{f}[d] - \phi^X[\text{head}_{G^*}(d)] + \phi^X[\text{tail}_{G^*}(d)]$ 
2:  $\text{val} := \min\{\text{residual\_capacity}, |\mathbf{v}[p_i]|\}$  amount of flow pushed on  $d$ 
3:  $\mathbf{f}[d] = \mathbf{f}[d] + \text{val}$  ;  $\mathbf{f}[\text{rev}(d)] := -\mathbf{f}[d]$ 
4:  $\mathbf{v}[\text{tail}(d)] = \mathbf{v}[\text{tail}(d)] - \text{val}$  ;  $\mathbf{v}[\text{head}(d)] = \mathbf{v}[\text{head}(d)] + \text{val}$  update the inflow at  $p_i$  and  $p_{i+1}$ 

% push excess inflow from  $\text{tail}(d)$  to  $\text{head}(d)$  using Hassin's algorithm
5: let  $\ell[d'] := \mathbf{c}[d'] - \mathbf{f}[d']$  for all darts  $d' \in P \cup \text{rev}(P)$  lengths of explicit darts are residual capacities  
(not including circulation component)
6:  $\ell[\text{rev}(d)] := |\mathbf{v}[p_i]| - \phi^X[\text{tail}_{G^*}(\text{rev}(d))] + \phi^X[\text{head}_{G^*}(\text{rev}(d))]$  set the limit on the residual capacity of  $\text{rev}(d)$   
(adjusted by circulation component)
7:  $\phi_i^X(\cdot) := FR(D^*, \ell, \phi^X, \text{head}_{G^*}(d))$  face potential are distances in  $G^*$  from  $\text{head}_{G^*}(d)$  w.r.t. the  
reduced lengths induced by  $\phi^X$ 
8:  $\text{val} := \phi_i^X[\text{head}(d)] - \phi_i^X[\text{tail}(d)]$  the amount of flow assigned to  $d$  by the circulation corresponding to  $\phi_i^X$ 
9:  $\mathbf{f}[d] := \mathbf{f}[d] - \text{val}$  ;  $\mathbf{f}[\text{rev}(d)] := -\mathbf{f}[(d)]$  do not push the circulation on  $d$  and  $\text{rev}(d)$ 
10:  $\phi^X = \phi^X + \phi_i^X$  accumulate the current circulation
11:  $\mathbf{v}[\text{tail}(d)] = \mathbf{v}[\text{tail}(d)] - \text{val}$  ;  $\mathbf{v}[\text{head}(d)] = \mathbf{v}[\text{head}(d)] + \text{val}$  update the inflow at  $p_i$  and  $p_{i+1}$ 

```

It follows from the above discussion that, after executing line 11 of Algorithm 2 using the efficient implementation (Algorithm 4) for the last time, the flow assignment \mathbf{f} maintained by the efficient implementation is the same as the one that would have been computed if the inefficient implementation (Algorithm 3) were used. Moreover, the potential function ϕ^X computed by the efficient implementation is the restriction to X^* of the potential function ϕ that would have been computed by the inefficient implementation.

In order to output the flow \mathbf{f}' in the entire graph G , we need to know the potential function ϕ rather than just its restriction ϕ^X . Observe, however, that any pseudoflow that differs from \mathbf{f} by a circulation is a valid output of FIXCONSERVATIONONPATH since a circulation does not change the inflow at any node, nor does it introduce residual paths between nodes that are not connected by a residual path in $G_{\mathbf{f}}$. It therefore suffices to find any feasible circulation χ in $G_{\mathbf{f}}$. This can be done by computing shortest paths in G^* w.r.t. the lengths $\ell := \mathbf{c} - \mathbf{f}$ from an arbitrary node $x \in G^*$. Note, however, that for darts of $P \cup \text{rev}(P)$ these lengths might be negative. We therefore use the $O(n \log^2 n / \log \log n)$ -time algorithm for shortest paths with negative lengths in planar graphs [30] to compute a feasible circulation χ .⁵

The pseudocode for the efficient implementation of FIXCONSERVATIONONPATH is given in appendix D.

Running Time Analysis Let n and k be the number of nodes in G and in P , respectively. The initialization time is dominated by the $O(n \log n + k^2 \log n)$ time for MSSP. The execution of each iteration of the main loop is dominated by the call to FR, which takes $O(k \log^2 k)$ time. The number of iterations is $k - 1$, leading to a total of $O(k^2 \log^2 k)$ time for execution of the main loop. Computing the circulation χ requires one shortest path computation, which takes $O(n \log^2 n / \log \log n)$ time. Thus total running time of the efficient implementation of FIXCONSERVATIONONPATH is $O(n \log^2 n / \log \log n + k^2 \log^2 k)$.

⁵In appendix E, we show how to use one more call to FR followed by a call to Dijkstra's algorithm to compute these distances. While we think doing so is more elegant and simpler to implement, it does not change the asymptotic running time of the algorithm.

4 Pushing Excess Inflow from a Cycle

In this section we describe the procedure `CYCLETOSINGLE-SINKLIMITEDMAXFLOW` that pushes excess inflow from a cycle to a node not on the cycle. The procedure is given as Algorithm 5. The procedure `SINGLE-SOURCETOCYCLELIMITEDMAXFLOW` is very similar. We omit its description.

Algorithm 5 `CYCLETOSINGLE-SINKLIMITEDMAXFLOW`($G, \mathbf{c}, \mathbf{f}_0, C, t$)

Input: a directed planar graph G with capacities \mathbf{c} , a pseudoflow \mathbf{f}_0 , a simple cycle C , a sink a_i .

Assumes: $\forall v \in C^+, \text{inflow}_{\mathbf{f}_0}(v) \geq 0$, where $C^+ = \{v \in C : \{x \in C : \text{inflow}_{\mathbf{f}_0}(x) > 0\} \xrightarrow{G, \mathbf{f}_0} v\}$.

Output: a pseudoflow \mathbf{f} s.t. (i) $\mathbf{f} - \mathbf{f}_0$ obeys conservation everywhere but $C^+ \cup \{t\}$,

(ii) $\forall v \in C^+, \text{inflow}_{\mathbf{f}}(v) \geq 0$, (iii) $\{v \in C : \text{inflow}_{\mathbf{f}}(v) > 0\} \not\xrightarrow{G, \mathbf{f}} a_i$.

- 1: let $C^+ := \{v \in C : \text{there exists a residual path to } v \text{ from a node } x \in C \text{ with } \text{inflow}_{\mathbf{f}_0}(x) > 0\}$
 - 2: delete the nodes of $C - C^+$
 - 3: let v_1, v_2, \dots, v_ℓ be the nodes of C^+ , labeled according to their cyclic order on C
 - 4: add artificial arcs $v_i v_{i+1}$ for $1 \leq i < \ell$
 - 5: let P be the v_1 -to- v_ℓ path of artificial arcs
 - 6: contract all the edges of P , collapsing C^+ into the single node v_1
 - 7: $\mathbf{f} := \text{SINGLE-SOURCE-SINGLE-SINK-MAXFLOW}(G, \mathbf{c} - \mathbf{f}_0, v_1, t)$
 - 8: undo the contraction of the edges of P
 - 9: $\mathbf{f} := \text{FIX-CONSERVATION-ON-PATH}(G, P, \mathbf{c}, \mathbf{f}_0 + \mathbf{f}) - \mathbf{f}_0$
 - 10: modify \mathbf{f} to push back flow to nodes of C^+ whose inflow w.r.t $\mathbf{f}_0 + \mathbf{f}$ is negative
 - 11: $\mathbf{f} := \mathbf{f}_0 + \mathbf{f}$
 - 12: **return** \mathbf{f}
-

To compute C^+ in Line 1, consider the residual graph of G w.r.t. \mathbf{f}_0 . Add a node v and non-zero capacity arcs vw for every node w whose inflow w.r.t. \mathbf{f}_0 is positive (these arcs may not preserve planarity). In $O(|G|)$ time, find the set X of nodes that are reachable from v via darts with non-zero capacity. Then $C^+ = C \cap X$.

Since C^+ consists of all nodes of C reachable via residual paths from the nodes of C with positive inflow, the flow computed by the procedure involves no darts incident to nodes in $C - C^+$. Thus, restricting the computation to the graph obtained from G by deleting the nodes in $C - C^+$ (line 2) does not restrict the computed flow. After deletion, adding artificial arcs between consecutive nodes of C^+ (line 4) will not violate planarity. Contracting the artificial arcs effectively turns C^+ into a single node v_1 . Next, the procedure computes a maximum flow \mathbf{f} from C^+ to t w.r.t. residual capacities $\mathbf{c} - \mathbf{f}_0$. This is done by invoking a single-source single-sink maximum flow algorithm [2] with source v_1 and sink t (line 7). Uncontracting the artificial arcs turns \mathbf{f} into a maximum C^+ -to- t flow in G w.r.t. the residual capacities $\mathbf{c} - \mathbf{f}_0$. However, some of the nodes of C^+ may have negative inflow w.r.t. $\mathbf{f}_0 + \mathbf{f}$. In line 9, the procedure calls `FIX-CONSERVATION-ON-PATH` to reroute the flow \mathbf{f} among the nodes of C^+ so that, w.r.t. $\mathbf{f}_0 + \mathbf{f}$, there are no residual paths from nodes of C^+ with positive inflow to nodes of C^+ with negative inflow. This implies that any node of C^+ that still has negative inflow has pushed too much flow. Line 10 modifies \mathbf{f} to push back such excess flow so that no node of C^+ has negative inflow w.r.t. $\mathbf{f}_0 + \mathbf{f}$. This is done using the procedure of Section 1.1 described in the appendix. Finally, the procedure returns $\mathbf{f}_0 + \mathbf{f}$. See Appendix F for a formal proof of correctness.

We next analyze the running time of this procedure on an n -node graph G and a k -node cycle C . The st -maximum flow computation in line 7 takes $O(n \log n)$ time using the algorithm of Borradaile and Klein [2]. The running time of the procedure is therefore dominated by the call to `FIX-CONSERVATION-ON-PATH` in line 9 which takes $O(n \log^2 n / \log \log n + k^2 \log^2 k)$ time.

Acknowledgements

We thank Haim Kaplan and Micha Sharir for discussions of their unpublished data structure [21]. PNK and SM thank Robert Tarjan for encouraging us to work on this problem.

References

- [1] R. Ahuja, Ö. Ergun, J. Orlin, and A. Punnen. A survey of very large scale neighborhood search techniques. *Discrete Applied Mathematics*, 23:75–102, 2002.
- [2] G. Borradaile and P. N. Klein. An $O(n \log n)$ algorithm for maximum st -flow in a directed planar graph. *Journal of the ACM*, 56(2), 2009.
- [3] Y. Boykov and V. Kolmogorov. An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(9):1124–1137, 2004.
- [4] Y. Boykov, O. Veksler, and R. Zabih. Efficient approximate energy minimization via graph cuts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(12):1222–1239, 2001.
- [5] S. Cabello and E. W. Chambers. Multiple source shortest paths in a genus g graph. In *Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 89–97, 2007.
- [6] M. de Berg, M. van Kreveld, M. Overmars, and O. Schwarzkopf. *Computational Geometry: Algorithms and Applications*. Springer-Verlag, 2nd edition edition, 2000. pp. 96–99.
- [7] J. Erickson. Maximum flows and parametric shortest paths in planar graphs. In *Proceedings of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 794–804, 2010.
- [8] J. Fakcharoenphol and S. Rao. Planar graphs, negative weight edges, shortest paths, and near linear time. *Journal of Computer and System Sciences*, 72(5):868–889, 2006.
- [9] S. Geman and D. Geman. Stochastic relaxation, Gibbs distributions, and the Bayesian relation of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6(6):721–742, 1984.
- [10] A. Goldberg and S. Rao. Beyond the flow decomposition barrier. *Journal of the ACM*, 45(5):783–797, 1998.
- [11] A. V. Goldberg and R. E. Tarjan. Solving minimum-cost flow problems by successive approximation. In *Proceedings of the 19th Annual ACM Symposium on Theory of Computing*, pages 7–18, 1987.
- [12] D. Greig, B. Porteous, and A. Seheult. Exact maximum a posteriori estimation for binary images. *Journal of the Royal Statistical Society, Series B*, 51(2):271–279, 1989.
- [13] R. Hassin. Maximum flow in (s, t) planar networks. *Information Processing Letters*, 13(3):107, 1981.
- [14] M. R. Henzinger, P. N. Klein, S. Rao, and S. Subramanian. Faster shortest-path algorithms for planar graphs. *Journal of Computer and System Sciences*, 55(1):3–23, 1997.
- [15] D. S. Hochbaum. An efficient algorithm for image segmentation, Markov random fields and related problems. *Journal of the ACM*, 48(4):686–701, 2001.
- [16] D. S. Hochbaum. The pseudoflow algorithm: A new algorithm for the maximum-flow problem. *Operations Research*, 56(4):992–1009, 2008.
- [17] G. F. Italiano, Y. Nussbaum, P. Sankowski, and C. Wulff-Nilsen. Improved algorithms for min cut and max flow in undirected planar graphs. In *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing*, 2011. To appear.
- [18] D. B. Johnson. Efficient algorithms for shortest paths in sparse networks. *Journal of the ACM*, 24(1):1–13, 1977.
- [19] D. B. Johnson and S. Venkatesan. Using divide and conquer to find flows in directed planar networks in $O(n^{3/2} \log n)$ time. In *Proceedings of the 20th Annual Allerton Conference on Communication, Control, and Computing*, pages 898–905, 1982.

- [20] H. Kaplan and Y. Nussbaum. Maximum flow in directed planar graphs with vertex capacities. In *Proceedings of the 17th European Symposium on Algorithms*, pages 397–407, 2009.
- [21] H. Kaplan and M. Sharir. Finding the maximal empty rectangle containing a query point. manuscript, 2011.
- [22] S. Khuller, J. Naor, and P. Klein. The lattice structure of flow in planar graphs. *SIAM Journal on Discrete Mathematics*, 6(3):477–490, 1993.
- [23] P. N. Klein. Multiple-source shortest paths in planar graphs. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 146–155, 2005.
- [24] P. N. Klein, S. Mozes, and O. Weimann. Shortest paths in directed planar graphs with negative lengths: A linear-space $O(n \log^2 n)$ -time algorithm. *ACM Transactions on Algorithms*, 6(2):1–18, 2010.
- [25] P. N. Klein and S. Subramanian. A fully dynamic approximation scheme for shortest paths in planar graphs. *Algorithmica*, 22(3):235–249, 1998.
- [26] J. Kleinberg and E. Tardos. Approximation algorithms for classification problems with pairwise relationships: Metric labeling and Markov random fields. In *Proceedings of the 40th Annual Symposium on Foundations of Computer Science*, 1999.
- [27] P. Kohli and P. H. S. Torr. Dynamic graph cuts for efficient inference in Markov random fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(12):2079–2088, 2007.
- [28] G. L. Miller. Finding small simple cycle separators for 2-connected planar graphs. *Journal of Computer and System Sciences*, 32(3):265–279, 1986.
- [29] G. L. Miller and J. Naor. Flow in planar graphs with multiple sources and sinks. *SIAM Journal on Computing*, 24(5):1002–1017, 1995. Preliminary version in FOCS 1989.
- [30] S. Mozes and C. Wulff-Nilsen. Shortest paths in planar graphs with real lengths in $O(n \log^2 n / \log \log n)$ time. In *Proceedings of the 18th European Symposium on Algorithms*, pages 206–217, 2010.
- [31] F. R. Schmidt, E. Toeppe, and D. Cremers. Efficient planar graph cuts with applications in computer vision. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 351–356, 2009.
- [32] D. D. Sleator and R. E. Tarjan. A data structure for dynamic trees. *Journal of Computer and System Sciences*, 26(3):362–391, 1983.
- [33] S. Subramanian. *Parallel and Dynamic Shortest-Path Algorithms for Sparse Graphs*. PhD thesis, Brown University, 1995. Available as Brown University Computer Science Technical Report CS-95-04.

A Correctness of Algorithm 1

Lemma A.1. *Let ρ be a circulation. Let u and v be nodes in a graph G . Then*

$$u \xrightarrow{G} v \Rightarrow u \xrightarrow{G_\rho} v.$$

Proof. Pushing a circulation does not change the amount of flow crossing any cut. This implies that if there was no u -to- v residual path before ρ was pushed, then there is none after ρ is pushed. \square

We will use the following two lemmas in the proof of correctness:

Lemma A.2. (*sources lemma*) *Let f be a flow with source set X . Let A, B be two disjoint sets of nodes. Then*

$$A \cup X \xrightarrow{G} B \Rightarrow A \xrightarrow{G_f} B.$$

Proof. The flow f may be decomposed into a cyclic component (a circulation) and an acyclic component. By Lemma A.1, it suffices to show the lemma for an acyclic flow f .

Suppose for the sake of contradiction that there exists a residual a -to- b simple path P after f is pushed for some $a \in A$ and $b \in B$. Let P' be the maximal suffix of P that was residual before the push. Maximality implies that the dart d of P whose head is $\text{start}(P')$ was non-residual before the push, and $f(\text{rev}(d)) > 0$. The fact that $f(\text{rev}(d)) > 0$ implies that before f was pushed there was a residual path Q from some node $x \in X$ to $\text{head}(d)$. Therefore, the concatenation of Q and P' was a residual x -to- b residual path before the push, a contradiction. \square

The proof of the following lemma is symmetric.

Lemma A.3. (*sinks lemma*) *Let f be a flow with sink set X . Let A, B be two disjoint sets of nodes. Then*

$$A \xrightarrow{G} B \cup X \Rightarrow A \xrightarrow{G_f} B.$$

For node sets W, Y, Z , we will use the notation *sources lemma*(W, Y, Z) to indicate the use of the sources lemma to establish that there are no W -to- Y residual paths after a flow with source set Z is pushed. We will use *sinks lemma*(W, Y, Z) in a similar manner.

Lemma A.4. *At any time in the running of the algorithm after the last execution of line 7 and before the execution of line 13, $S \nrightarrow T$, $S \nrightarrow A$, $S \nrightarrow C$, $A \nrightarrow T$, $C \nrightarrow T$.*

Proof. The fact that the properties hold just after line 7 follows from the properties of the recursive calls and from the fact that any residual path from a node in one piece to a node in the other consists of a residual path to C and a residual path from C . Note that, because of the contractions in line 4, at this time the cycle C consists of just the node v . The nonexistence of residual paths with respect to v in the recursive calls implies the nonexistence of residual paths w.r.t. any node of C after the contractions are undone in line 8.

Now, let us prove that the properties are maintained until just before the execution of line 13. By the above argument, there is a cut separating S from $T \cup A \cup C$ which is saturated just after line 8. The procedure in line 9 only pushes flow between vertices of C , and in lines 10–12, flow is only pushed between the nodes of A and C . These sets are all on the same side of the cut which therefore stays saturated. It follows that $S \nrightarrow T$, $S \nrightarrow A$, and $S \nrightarrow C$ at any point after line 7 and before line 13. A similar argument applied to a saturated cut separating $A \cup C$ from T shows $A \nrightarrow T$ and $C \nrightarrow T$. \square

Recall that C_i^+ is the set of nodes of C that are reachable via residual paths from some node of C with positive excess at the beginning of iteration i , and that C_i^- is the set of nodes of C that have residual paths to some node of C with negative excess at the beginning of iteration i .

Lemma A.5. *Just after line 9 of Algorithm 1 is executed, $C_1^+ \nrightarrow C_1^-$.*

Proof. Follows from the definition of `FIXCONSERVATIONONPATH`. \square

Lemma A.6. *For all $i < j$, $C_j^+ \subseteq C_i^+$ and $C_j^- \subseteq C_i^-$*

Proof. It suffices to show that, for all i , $C_{i+1}^+ \subseteq C_i^+$ and $C_{i+1}^- \subseteq C_i^-$.

The flow pushed in iteration i of line 11 can be decomposed into a flow whose sources and sinks are all in C_i^+ and a flow whose sources are in C_i^+ and whose sink is a_i . Therefore, the set X of nodes of C with positive inflow immediately after iteration i of line 11 is a subset of C_i^+ . By definition of SINGLESOURCE-TOCYCLELIMITEDMAXFLOW, the set of nodes of C with positive inflow does not change after line 12 is executed. Therefore, C_{i+1}^+ is the set of nodes reachable from X by a residual path after iteration i of line 12.

By definition of C_i^+ , immediately before iteration i of line 11 there are no C_i^+ -to- $\{C - C_i^+\}$ residual paths. By *sources lemma*($C_i^+, C - C_i^+, C_i^+$), there are no C_i^+ -to- $\{C - C_i^+\}$ residual paths immediately after iteration i of line 11 as well. This shows that there are no X -to- $\{C - C_i^+\}$ residual paths at that time.

The flow pushed in line 12 can be decomposed into a flow whose sources and sinks are all in C_i^- and a flow whose source is a_i and whose sinks are all in C_i^- . By *sinks lemma*($C_i^+, C - C_i^+, C_i^-$), there are no C_i^+ -to- $\{C - C_i^+\}$ residual paths immediately after iteration i of line 12. This shows that there are no X -to- $\{C - C_i^+\}$ residual paths at that time. Hence $C_{i+1}^+ \subseteq C_i^+$, as desired.

The proof of the analogous claim for C_{i+1}^- is similar. \square

Lemma A.7. *Just after line 11 of Algorithm 1 is executed in iteration i , $C_i^+ \rightarrow C_i^-$, $C_{i+1}^+ \rightarrow \{a_j\}_{j \leq i}$, $\{a_j\}_{j < i} \rightarrow C_i^-$.*

Proof. The flow pushed in line 11 can be decomposed into a flow whose sources and sinks are all in C_i^+ and a flow whose sources are in C_i^+ and whose sink is a_i .

- $C_i^+ \rightarrow C_i^-$ by *sources lemma*(C_i^+, C_i^-, C_i^+)
- $C_{i+1}^+ \rightarrow \{a_j\}_{j < i}$ by *sources lemma*(C_{i+1}^+, a_j, C_i^+)
- $C_{i+1}^+ \rightarrow \{a_i\}$ by definition of CYCLETOSINGLE-SINKLIMITEDMAXFLOW
- $\{a_j\}_{j < i} \rightarrow C_i^-$ by *sources lemma*(a_j, C_i^-, C_i^+)

\square

Lemma A.8. *Just after line 12 of Algorithm 1 is executed in iteration i , $C_i^+ \rightarrow C_i^-$, $C_{i+1}^+ \rightarrow \{a_j\}_{j \leq i}$, $\{a_j\}_{j \leq i} \rightarrow C_{i+1}^-$.*

Proof. The flow pushed in line 12 can be decomposed into a flow whose sources and sinks are all in C_i^- and a flow whose source is a_i and whose sinks are all in C_i^- .

- $C_i^+ \rightarrow C_i^-$ by *sinks lemma*(C_i^+, C_i^-, C_i^-)
- $C_{i+1}^+ \rightarrow \{a_j\}_{j \leq i}$ by *sinks lemma*(C_{i+1}^+, a_j, C_i^-)
- $\{a_j\}_{j < i} \rightarrow C_{i+1}^-$ by *sinks lemma*(a_j, C_{i+1}^-, C_i^-)
- $a_i \rightarrow C_{i+1}^-$ by definition of SINGLESOURCE-TOCYCLELIMITEDMAXFLOW

\square

Let $C^+(C^-)$ denote the set of nodes in C with positive (negative) inflow just before line 13 is executed.

Corollary A.9. *Just before line 13 of Algorithm 1 is executed, $C^+ \rightarrow C^-$, $C^+ \rightarrow A$, $A \rightarrow C^-$.*

Lemma A.10. *The following are true upon termination:*

1. f is a pseudoflow
2. f obeys conservation everywhere except at S, T, A
3. $S \xrightarrow{G_f} T, S \xrightarrow{G_f} A, A \xrightarrow{G_f} T$.

Proof. Since every addition to \mathbf{f} along the algorithm respects capacities of all darts, \mathbf{f} is a pseudoflow at all times. By induction, the only nodes that do not obey conservations after the recursive calls are those of S, T and A . Subsequent changes to \mathbf{f} only violate conservation on the nodes of C , but any such violation is eliminated in line 13. Therefore upon termination \mathbf{f} obeys conservation everywhere except S, T and A .

Since, by Lemma A.4 and Corollary A.9 before line 13 $C^+ \rightarrow A$ and $C \rightarrow T$, the flow pushed back from C^+ in line 13 must be pushed back to S . Similarly, the flow pushed back to C^- must be pushed back from T .

Let f_+ (f_-) be the flow pushed back from C^+ to S (from T to C^-) in line 13. Consider first pushing back f_+ .

- $S \rightarrow T$ by *sources lemma*(S, T, C^+)
- $S \rightarrow A$ by *sources lemma*(S, A, C^+)
- $A \rightarrow T$ by *sources lemma*(A, T, C^+)
- $S \rightarrow C^-$ by *sources lemma*(S, C^-, C^+)
- $A \rightarrow C^-$ by *sources lemma*(A, C^-, C^+)

Next consider pushing f_-

- $S \rightarrow T$ by *sinks lemma*(S, T, C^-)
- $S \rightarrow A$ by *sinks lemma*(S, A, C^-)
- $A \rightarrow T$ by *sinks lemma*(A, T, C^-)

□

B Correctness of Algorithm 2

Lemma B.1. *The following holds immediately after iteration i of the main loop (line 3).*

1. For $j \leq i, j' > j$, if p_j has positive inflow, there is no residual path from p_j to $p_{j'}$. If p_j has negative inflow, there is no residual path from $p_{j'}$ to p_j .
2. For $j, j' \leq i$, if p_j has positive inflow and $p_{j'}$ has negative inflow then there is no p_j -to- $p_{j'}$ residual path.

Proof. By induction on the number of iterations i of the loop. For $i = 0$ the invariants are trivially satisfied.

First note that the adjustments to capacities and flow in lines 6–8 do not create any new residual paths since capacities are only reduced, and no residual capacity increases. Therefore, it suffices to argue just about the flow pushed in line 11.

Assume the invariants hold up until the beginning of the i^{th} iteration. We show that the invariants hold at the end of the iteration. Suppose that p_i has positive inflow at the end of the i^{th} iteration (the case of negative inflow is similar).

1. Since the flow pushed in line 11 is limited by $|\text{inflow}(p_i)|$, the fact that p_i has positive inflow at the end implies that the flow pushed was a maximum flow from p_i to p_{i+1} . Since the capacities of darts between p_{k+1} and p_k for $k > i$ are sufficiently large, the maximality of the flow implies that there are no p_i -to- p_k residual paths for any $k > i$.

The invariant holds for nodes p_j with positive inflow and $j < i$ by *sinks lemma*($\{p_j\}, \{p_{j'} : j' > j\}, \{p_{i+1}\}$).

2. Invariant 2 holds for $j, j' < i$ by *sinks lemma*($\{p_j : j < i, \text{inflow}(p_j) > 0\}, \{p_j : j < i, \text{inflow}(p_j) < 0\}, \{p_{i+1}\}$).

The invariant holds for p_i by invoking the *sources lemma*($\{p_i\}, \{p_j : j < i, \text{inflow}(p_j) < 0\}, \{p_i\}$).

□

Corollary B.2. *Algorithm 2 is correct*

Proof. The flow $\mathbf{f}' - \mathbf{f}_0$ satisfies conservation everywhere except at nodes of P since the algorithm only pushes flow between nodes of P . By Lemma B.1, with respect to \mathbf{f}' , there are no residual paths from nodes of P with positive inflow to nodes of P with negative inflow. □

C Converting a Pseudoflow into a Maximum Feasible Flow

Let \mathbf{f} be a pseudoflow in a planar graph G with node set V , sources S and sinks T . Let V^+ denote the set of nodes $\{v \in V - (S \cup T) : \text{inflow}_{\mathbf{f}}(v) > 0\}$. Similarly, let V^- denote the set of nodes $\{v \in V - (S \cup T) : \text{inflow}_{\mathbf{f}}(v) < 0\}$. Suppose $S \cup V^+ \xrightarrow{G, \mathbf{f}} T \cup V^-$. In this appendix we show how to convert \mathbf{f} into a maximum feasible flow \mathbf{f}' . This procedure was first described for planar graphs by Johnson and Venkatesan [19]. The original description of the procedure took $O(n \log n)$ time, but using the flow cycles canceling technique of Kaplan and Nussbaum [20] the running time is $O(n)$.

We begin by converting \mathbf{f} to an acyclic pseudoflow in linear time [20]. That is, after the conversion there is no cycle C such that $\mathbf{f}[d] > 0$ for every dart d of C . Since \mathbf{f} is acyclic, the darts with a positive flow induce a topological ordering on the nodes of the graph G . Let v be the last member of V^+ in the topological ordering, and let d be an arbitrary dart that carries flow into v . We set $\mathbf{f}[d] = \max\{\mathbf{f}[d] - \text{inflow}_{\mathbf{f}}(v), 0\}$, and set $\mathbf{f}[\text{rev}(d)]$ accordingly. The flow assignment \mathbf{f} maintains the invariant $S \cup V^+ \xrightarrow{G, \mathbf{f}} T \cup V^-$ by *sinks lemma* ($S \cup V^+, T \cup V^-, \{v\}$). As long as v is in V^+ , there must be a dart d which carries flow to v . By changing the flow on d we cannot add to V^+ a new node that appears later than v in the topological ordering. We repeat this process until V^+ is empty. Since we reduce the flow on each dart at most once, this takes linear time. Next we handle V^- while keeping the invariant $S \cup V^+ \xrightarrow{G, \mathbf{f}} T \cup V^-$ in a symmetric way, by repeatedly fixing the first vertex in V^- in the topological ordering.

The total running time is $O(n)$, and since V^+, V^- are both empty, we get from the invariant $S \cup V^+ \xrightarrow{G, \mathbf{f}} T \cup V^-$ that the resulting pseudoflow is a feasible flow. This is the required flow \mathbf{f}' .

D Pseudocode for FIXCONSERVATIONONPATH

Algorithm 6 FIXCONSERVATIONONPATH($G, P, \mathbf{c}, \mathbf{f}_0$)

Input: a directed planar graph G , simple path $P = d_1 d_2 \dots d_k$, capacity function \mathbf{c} , and a pseudoflow \mathbf{f}_0

Output: a pseudoflow \mathbf{f} s.t. (1) $\mathbf{f} - \mathbf{f}_0$ satisfies conservation at nodes not on P and (2) w.r.t \mathbf{f} , no residual path exists from $\{v \in P : \text{inflow}(v) > 0\}$ to $\{v \in P : \text{inflow}(v) < 0\}$.

```

1: let  $X^*$  be the set of endpoints in the planar dual  $G^*$  of the darts of  $P$ 
2: initialize to zero an array  $\phi^X[\cdot]$  indexed by the nodes of  $X^*$ 
3: let  $H^*$  be the graph obtained from  $G^*$  by removing the darts of  $P$ 
4: compute, using the MSSP algorithm [23, 5], a table  $D^*[\cdot, \cdot]$  of distances in  $H^*$  between nodes of  $X^*$  w.r.t
   the lengths  $\mathbf{c} - \mathbf{f}_0$ 
5:  $\mathbf{f} := \mathbf{f}_0$ 
6:  $\mathbf{c}[d] := \mathbf{c}[d] + M$  for all darts of  $P \cup \text{rev}(P)$ 
7: initialize an array of length  $k$  by  $\mathbf{v}[i] := \text{inflow at } p_i \text{ with respect to } \mathbf{f}$ 
8: for  $i = 1, 2, \dots, k$ 
9:   let  $p_i$  and  $p_{i+1}$  be the tail and head of  $d_i$ , respectively

   % reduce the capacities of  $d$  and  $\text{rev}(d)$  by  $M$  and adjust the flow appropriately
10:  for  $d \in \{d_i, \text{rev}(d_i)\}$ 
11:     $\mathbf{c}[d] := \mathbf{c}[d] - M$ 
12:     $\text{old flow} := \mathbf{f}[d] + \phi^X[\text{head}_{G^*}(d)] - \phi^X[\text{tail}_{G^*}(d)]$ 
13:     $\text{new flow} := \min\{\text{old flow}, \mathbf{c}[d]\}$  adjusted flow must not exceed adjusted capacity
14:     $\mathbf{f}[d] := \text{new flow} - \phi^X[\text{head}_{G^*}(d)] + \phi^X[\text{tail}_{G^*}(d)]$  the explicit flow does not include the circulation component
15:     $\mathbf{f}[\text{rev}(d)] := -\mathbf{f}[d]$ 
16:     $\mathbf{v}[\text{head}(d)] := \mathbf{v}[\text{head}(d)] + \text{new flow} - \text{old flow}$  update the inflow at  $p_i$  and  $p_{i+1}$ 
17:     $\mathbf{v}[\text{tail}(d)] := \mathbf{v}[\text{tail}(d)] - \text{new flow} + \text{old flow}$ 

18:  if  $\mathbf{v}[p_i] > 0$  then  $d := d_i$  else  $d := \text{rev}(d_i)$  find in which direction flow should be pushed

   % push flow on  $d$  to make its residual capacity zero as required for Hassin's algorithm
19:   $\text{residual capacity} := \mathbf{c}[d] - \mathbf{f}[d] - \phi^X[\text{head}_{G^*}(d)] + \phi^X[\text{tail}_{G^*}(d)]$ 
20:   $\text{val} := \min\{\text{residual capacity}, |\mathbf{v}[p_i]|\}$  amount of flow pushed on  $d$ 
21:   $\mathbf{f}[d] = \mathbf{f}[d] + \text{val}$ ;  $\mathbf{f}[\text{rev}(d)] := -\mathbf{f}[d]$ 
22:   $\mathbf{v}[\text{tail}(d)] = \mathbf{v}[\text{tail}(d)] - \text{val}$ ;  $\mathbf{v}[\text{head}(d)] = \mathbf{v}[\text{head}(d)] + \text{val}$  update the inflow at  $p_i$  and  $p_{i+1}$ 

   % push excess inflow from  $\text{tail}(d)$  to  $\text{head}(d)$  using Hassin's algorithm
23:  let  $\ell[d'] := \mathbf{c}[d'] - \mathbf{f}[d']$  for all darts  $d' \in P \cup \text{rev}(P)$  lengths of explicit darts are residual capacities
   (not including circulation component)
24:   $\ell[\text{rev}(d)] := |\mathbf{v}[p_i]| - \phi^X[\text{tail}_{G^*}(\text{rev}(d))] + \phi^X[\text{head}_{G^*}(\text{rev}(d))]$  set the limit on the residual capacity of  $\text{rev}(d)$ 
   (adjusted by circulation component)
25:   $\phi_i^X(\cdot) := FR(D^*, \ell, \phi^X, \text{head}_{G^*}(d))$  face potential are distances in  $G^*$  from  $\text{head}_{G^*}(d)$  w.r.t. the
   reduced lengths induced by  $\phi^X$ 
26:   $\text{val} := \phi_i^X[\text{head}(\text{rev}(d))] - \phi_i^X[\text{tail}(\text{rev}(d))]$  the amount of flow assigned to  $d$  by the circulation corresponding to  $\phi_i^X$ 
27:   $\mathbf{f}[d] := \mathbf{f}[d] + \text{val}$ ;  $\mathbf{f}[\text{rev}(d)] := -\mathbf{f}[d]$  do not push the circulation on  $d$  and  $\text{rev}(d)$ 
28:   $\phi^X = \phi^X + \phi_i^X$  accumulate the current circulation
29:   $\mathbf{v}[\text{tail}(d)] = \mathbf{v}[\text{tail}(d)] - \text{val}$ ;  $\mathbf{v}[\text{head}(d)] = \mathbf{v}[\text{head}(d)] + \text{val}$  update the inflow at  $p_i$  and  $p_{i+1}$ 

   % find a feasible circulation
30:  let  $x$  be an arbitrary node in  $G^*$ 
31:  let  $\ell[d] := \mathbf{c}[d] - \mathbf{f}[d]$  for all darts  $d \in G$ 
32:   $\chi := \text{SINGLESOURCESHORTESTPATHS}(G^*, \ell, x)$  distances of the nodes of  $G^*$  from  $x$ 
33:   $\mathbf{f}'[d] := \mathbf{f}[d] + \chi[\text{head}_{G^*}(d)] - \chi[\text{tail}_{G^*}(d)]$  for every dart  $d \in G$  add explicit flow and implicit circulation
34:  return  $\mathbf{f}'$ 

```

E Alternative to Line 32 of Algorithm 6

In this section we show that it is not necessary to use a generic shortest path algorithm that works with negative lengths to compute a feasible circulation χ in line 32 of Algorithm 6. Instead of choosing x to be an arbitrary node in G^* , let x be an arbitrary node of X^* . Let $\chi(y)$ denote the x -to- y distance in G^* w.r.t. the lengths ℓ . Instead of computing χ using a shortest path algorithm that accepts negative lengths, we will compute it more efficiently in two steps. Pseudocode is given below as Algorithm 7. In the first step we use FR to compute the distances to just the nodes of X^* . In the second step we extend these distances to all other nodes using Dijkstra's algorithm.

In the first step the algorithm computes distances in G^* from x to all nodes of X^* w.r.t. ℓ_{ϕ^X} , the reduced lengths of ℓ induced by the feasible price function ϕ^X . This is done by an additional invocation of FR (line 2). Let ψ^X denote these distance labels. By definition of reduced lengths and a telescoping sum similar to the one in the derivation of Eq. (6),

$$\psi^X[y] = \chi(y) + \phi^X[x] - \phi^X[y]. \quad (7)$$

Since both ϕ^X and ψ^X are known, the algorithm can compute the unreduced distances $\chi[y]$ for all $y \in X^*$ (line 4). Next, the algorithm runs Dijkstra's algorithm on G^* , initializing the label of each node y of X^* to its correct value $\chi[y]$. Since in the dual the darts of $P \cup \text{rev}(P)$ are only incident to nodes of X^* , Dijkstra's algorithm initialized in this manner correctly outputs the distance labels for all nodes of G^* even if some of the darts of $P \cup \text{rev}(P)$ may have negative lengths.

Algorithm 7 Alternative to line 32 of Algorithm 6

- 1: let x be an arbitrary node in X^*
 - 2: $\psi^X = \text{FR}(\text{D}^*, \ell, \phi^X, x)$ find distances of the nodes of X^* from x w.r.t. price function ϕ^X
 - 3: initialize to ∞ an array χ indexed by the nodes of G^*
 - 4: $\chi[y] = \psi^X[y] - \phi^X[x] + \phi^X[y]$ for every $y \in X^*$ distances of the nodes of X^* from x
 - 5: extend ℓ to all darts of G^* by $\ell[d] := \mathbf{c}[d] - \mathbf{f}_0[d]$ for every dart $d \in G$
 - 6: $\chi := \text{DIJKSTRA}(G^*, \ell, x, \chi)$ distances of the nodes of G^* from x
-

F Correctness of Algorithm 5

Lemma F.1. *Algorithm 5 is correct.*

Proof. Any flow that is pushed by the algorithm originates at C^+ and terminates at $C^+ \cup \{t\}$. Therefore, $\mathbf{f} - \mathbf{f}_0$ violates conservation only at $C^+ \cup \{t\}$. In line 10, flow of \mathbf{f} is pushed back so that no node of C^+ has negative inflow w.r.t. $\mathbf{f}_0 + \mathbf{f}$. It is possible to do so by only pushing back flow of \mathbf{f} (rather than flow of $\mathbf{f}_0 + \mathbf{f}$) since by assumption no node of C^+ has negative inflow w.r.t. \mathbf{f}_0 .

By maximality of the flow pushed in line 7, just after line 7 is executed there are no C^+ -to- t residual paths. Clearly, this remains true when the capacities of the artificial darts are set to zero. In line 9 flow is pushed among the nodes of C^+ , so by *sources lemma*(C^+, t, C^+), there are no C^+ -to- t residual paths after line 9 either. Moreover, by definition of `FIXCONSERVATIONONPATH`, there are no $C_{>0}$ -to- $C_{<0}$ residual paths immediately after line 9 is executed, where $C_{>0}$ ($C_{<0}$) is the set of nodes of C^+ with positive (negative) inflow at that time. Line 10 pushes flow into $C_{<0}$, making all the nodes of $C_{<0}$ obey conservation. By *sinks lemma*($C_{>0}, t, C_{<0}$) there are no $C_{>0}$ -to- t residual paths upon termination of the procedure. This completes the proof since $C_{>0}$ is the set of nodes of C with positive inflow upon termination. \square