

The Unfortunate-Flow Problem

Orna Kupferman

School of Computer Science and Engineering, The Hebrew University, Israel

Gal Vardi

School of Computer Science and Engineering, The Hebrew University, Israel

Abstract

In the traditional maximum-flow problem, the goal is to transfer maximum flow in a network by directing, in each vertex in the network, incoming flow into outgoing edges. The problem is one of the most fundamental problems in TCS, with application in numerous domains. The fact a maximal-flow algorithm directs the flow in all the vertices of the network corresponds to a setting in which the authority has control in all vertices. Many applications in which the maximal-flow problem is applied involve an adversarial setting, where the authority does not have such a control.

We introduce and study the *unfortunate flow problem*, which studies the flow that is guaranteed to reach the target when the edges that leave the source are saturated, yet the most unfortunate decisions are taken in the vertices. When the incoming flow to a vertex is greater than the outgoing capacity, flow is lost. The problem models evacuation scenarios where traffic is stuck due to jams in junctions and communication networks where packets are dropped in overloaded routers.

We study the theoretical properties of unfortunate flows, show that the unfortunate-flow problem is co-NP-complete and point to polynomial fragments. We introduce and study interesting variants of the problem: *integral unfortunate flow*, where the flow along edges must be integral, *controlled unfortunate flow*, where the edges from the source need not be saturated and may be controlled, and *no-loss controlled unfortunate flow*, where the controlled flow must not be lost.

2012 ACM Subject Classification Mathematics of computing → Graph theory → Network flows

Keywords and phrases Flow Network, Graph Algorithms, Games

Digital Object Identifier 10.4230/LIPIcs.ICALP.2018.125

1 Introduction

A *flow network* is a directed graph in which each edge has a capacity, bounding the amount of flow that can travel through it. The amount of flow that enters a vertex equals the amount of flow that leaves it, unless the vertex is a *source*, which has only outgoing flow, or a *target*, which has only incoming flow. The fundamental *maximum-flow problem* gets as input a flow network and searches for a maximal flow from the source to the target [4, 10]. The problem was first formulated and solved in the 1950's [8, 9]. It has attracted much research on improved algorithms, variants, and applications [6, 5, 11, 15].

The maximum-flow problem can be applied in many settings in which something travels along a network. This covers numerous application domains, including traffic in road or rail systems, fluids in pipes, packets in a communication network, and many more [1]. Less obvious applications involve flow networks that are constructed in order to model settings with an abstract network, as in the case of scheduling with constraints [1]. In addition, several classical graph-theory problems can be reduced to the maximum-flow problem. This includes the problem of finding a maximum bipartite matching, minimum path cover, maximum edge-disjoint or vertex-disjoint path, and many more [4, 1]. Variants of the maximum-flow problem can accommodate further settings, like circulation problems [18], multiple source and target vertices, costs for unit flows, multiple commodities, and more [7].



© Orna Kupferman and Gal Vardi;

licensed under Creative Commons License CC-BY

45th International Colloquium on Automata, Languages, and Programming (ICALP 2018).

Editors: Ioannis Chatzigiannakis, Christos Kaklamanis, Dániel Marx, and Don Sannella; Article No. 125; pp. 125:1–125:18

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany




125:2 The Unfortunate-Flow Problem

44 Studies of flow networks so far assume that the vertices in the network are controlled by a central
 45 authority. Indeed, maximum-flow algorithms directs the flow in all vertices of the network. In many
 46 applications of flow networks, however, vertices of the network may not be controlled. Thus, every
 47 vertex may make autonomous and independent decisions regarding how to direct incoming flow to
 48 outgoing edges.

49 Consider, for example, a road network of a city, where the source s models the center of the city
 50 and the target t models the area outside the city. In order to *evacuate* the center of the city, drivers
 51 navigate from s to t . In each vertex, every incoming driver chooses an arbitrary outgoing edge with
 52 free capacity. If the outgoing capacity from a vertex is less than the incoming flow, then a traffic jam
 53 occurs, and flow is lost. As another example, consider a communication network in which packets
 54 are sent from a source router s and should reach a target router t . Whenever an internal router
 55 receives a packet it forwards it to an arbitrary neighbor router. If the outgoing capacity from a vertex
 56 is less than the incoming flow, then packets are dropped, and flow is lost.

57 In both examples, we want to find the flow that is guaranteed to reach the target in the worst
 58 scenario. We now formalize this intuition. Let $\mathcal{G} = \langle V, E, c, s, t \rangle$ be a flow network, where $\langle V, E \rangle$
 59 is a directed graph, $c : E \rightarrow \mathbb{N}$ assigns a capacity for each edge, and s, t are the source and target
 60 vertices. A *preflow* is a function $f : E \rightarrow \mathbb{R}$ that assigns to each edge $e \in E$, a flow in $[0, c(e)]$ such
 61 that the incoming flow to each vertex is greater or equal to its outgoing flow. A *saturating preflow*
 62 is a preflow in which all outgoing edges from s are saturated, and for every vertex $v \in V \setminus \{s, t\}$,
 63 the outgoing flow from v is the minimum between the incoming flow to v and the outgoing capacity
 64 from v . That is, in a saturating preflow, flow loss occurs in a vertex v if and only if the incoming
 65 flow to v is greater than the capacity of the edges outgoing from v . The *unfortunate flow value* of
 66 \mathcal{G} is the minimal flow that reaches t in a saturating preflow. Thus, it is the flow that is guaranteed
 67 to reach t when the edges that leave s are saturated, yet the most unfortunate routing decisions are
 68 taken in junctions. In the *unfortunate-flow problem*, we want to find the unfortunate flow value of \mathcal{G} .

69 ► **Example 1.** Consider the flow network \mathcal{G} appearing in Figure 1 (a). A maximum flow in \mathcal{G} has
 70 value 8, attained, for example, with the preflow in (b). A saturating preflow in \mathcal{G} appears in (c), and
 71 has value 5. While the edges leaving s are saturated, the routing of 7 flow units to the vertex at the
 72 bottom leads to a loss of 4 flow units in this vertex. ◀



■ **Figure 1** A flow network \mathcal{G} , and preflows that attain its maximum-flow and unfortunate-flow values.

73 We introduce the unfortunate-flow problem, study the theoretical properties of saturating pre-
 74 flows, and study the complexity of the problem. We also introduce and study interesting variants
 75 of the problem: *integral unfortunate flow*, where the flow along edges must be integral, *controlled*
 76 *unfortunate flow*, where the edges from the source need not be saturated and may be controlled, and
 77 *no-loss controlled unfortunate flow*, where the controlled flow must avoid loss.

78 Before we describe our contribution, let us review *flow games* and their connection to our contri-
 79 bution here. In flow games [14], the vertices of a flow network are partitioned between two players.
 80 Each player controls how incoming flow is partitioned among edges outgoing from her vertices.
 81 Then, one player aims at maximizing the flow that reaches t and the other player aims at minimizing
 82 it. It is shown in [14] that when the players are restricted to *integral strategies*, thus when the flows
 83 along the edges are integers, then the problem of finding the maximal flow that the maximizer player

84 can guarantee is Σ_2^P -complete. The restriction to integral strategies is crucial. Indeed, unlike the
 85 case of the traditional maximum-flow problem, non-integral strategies may be better than integral
 86 ones. In fact, the problem of finding a maximal flow for the maximizer in a setting with non-integral
 87 strategies was left open in [14]. The unfortunate-flow problem can be viewed as a special case of
 88 flow games, in which the maximizer player controls no vertex.

89 We start with the complexity of the unfortunate-flow problem. We consider the decision-problem
 90 variant, where we are given a threshold $\gamma > 0$ and decide whether the unfortunate flow value is at
 91 least γ . In the case of maximal flow, the problem can be solved in polynomial time [9], and so
 92 are many variants of it. We first show that, quite surprisingly, the unfortunate-flow problem is co-
 93 NP-hard and that it is NP-hard to approximate within any multiplicative factor. We then point to a
 94 polynomial fragment. Intuitively, the fragment restricts the number of vertices in which flow may be
 95 lost, which we pinpoint as the computational bottleneck. Formally, we say that a vertex is a *funnel*
 96 if its incoming capacity is greater than its outgoing capacity. We show that the unfortunate-flow
 97 problem can be solved in time $O(2^{|H|} \cdot (|E|^2 \log|V| + |E||V|\log^2|V|))$, where $H \subseteq V$ is the set
 98 of funnels in \mathcal{G} . In particular, the problem can be solved in strongly-polynomial time if the network
 99 has a logarithmic number of funnels. Our solution reduces the problem to a sequence of *min-cost*
 100 *max-flow* problems [1], implying the desirable *integral flow property*: the unfortunate-flow value can
 101 always be attained by an integral flow. The integral flow property implies a matching co-NP upper
 102 bound, thus the unfortunate-flow problem is co-NP-complete.

103 In some scenarios, we have some initial control on the flow. For example, in evacuation scenar-
 104 ios, as in the example of traffic leaving the city, police may direct cars at the center of the city, but
 105 has no control on them once they leave the center. Likewise, when entering or evacuating stadiums,
 106 police may direct the crowd to different gates, but has no control on how people proceed once they
 107 pass the gates [13]. We study the *controlled unfortunate-flow problem*, where the outgoing flow from
 108 s is bounded and controlled. Formally, there is an integer $\alpha \geq 0$ such that the total outgoing flow
 109 from s is bounded by α , and it is possible to control how this outgoing flow is partitioned among
 110 the edges that leave s . Our goal is to control this flow so that the flow that reaches t in the most
 111 unfortunate case is maximized.¹ We show that the integral-flow property no longer holds in this
 112 setting. Thus, there are networks in which an optimal strategy is to partition the α units of flow that
 113 leave s into non-integers. A troublesome implication of this is that an algorithm that guesses the
 114 strategy has to go over unboundedly many possibilities. This challenge is what has left flow games
 115 undecidable [14]. We show that we can still reduce the controlled unfortunate-flow problem into the
 116 second alternation level of the theory of real numbers under addition and order [17]. The reduction
 117 implies membership in Σ_2^P , and we show a matching lower bound. Thus, the controlled unfortunate-
 118 flow problem is Σ_2^P -complete. We also study a generalization of the problem, where control can be
 119 placed in a subset of the vertices.²

120 Finally, in some scenarios it is crucial for flow not to get lost. For example, in evacuation
 121 scenarios, we may prefer to give up an evacuation attempt under a loss risk, and in communication
 122 networks, we may tolerate low traffic and not tolerate dropping of packets. We say that a flow
 123 network \mathcal{G} is *safe* if all saturating preflows have no loss. For example, networks with no funnels are
 124 clearly safe. It is easy to see that \mathcal{G} is safe if its unfortunate flow value is equal to the maximal flow
 125 the source can generate, thus to the capacity of the edges outgoing from the source. This gives a
 126 co-NP algorithm for deciding the safety of a network. We show one can do better and reduce the

¹ We note that this is different from work done in *evacuation planning*, where the goal is to find routes and schedules of evacuees (for a survey, see [16]).

² Not to confuse with the problem of finding critical nodes for firefighters [2, 3]. While there the firefighters block the fire, in our setting they direct the evacuation. Thus, there, the goal is to block undesired vulnerabilities in the network, and here the goal is maximize desired traffic.

125:4 The Unfortunate-Flow Problem

127 safety problem to a maximum weighted flow problem, which can be solved in polynomial time. We
128 then turn to study the *no-loss controlled unfortunate-flow problem*, where we control the flow in
129 edges from s , and we want to maximize the flow to t but in a way that flow loss is impossible. We
130 show that the problem is NP-complete.

131 Due to space limitations, some examples and proofs are omitted and can be found in the full
132 version, in the authors' URLs.

2 Preliminaries

134 A *flow network* is $\mathcal{G} = \langle V, E, c, s, t \rangle$, where V is a set of vertices, $E \subseteq V \times V$ is a set of directed
135 edges, $c : E \rightarrow \mathbb{N}$ is a capacity function, and $s, t \in V$ are source and target vertices. The capacity
136 function assigns to each edge $e \in E$ a nonnegative capacity $c(e) \geq 0$. We define the *size* of \mathcal{G} ,
137 denoted $|\mathcal{G}|$ by $|V| + |E| + |c|$, where $|c|$ is the size required for encoding the capacity function c ,
138 thus assuming the capacities are given in binary. For a vertex $v \in V$, let $E^{\rightarrow v}$ and $E^{v \rightarrow}$ be the sets
139 of incoming and outgoing edges to and from v , respectively. That is, $E^{\rightarrow v} = (V \times \{v\}) \cap E$ and
140 $E^{v \rightarrow} = (\{v\} \times V) \cap E$. A *sink* is a vertex v with no outgoing edges, thus $E^{v \rightarrow} = \emptyset$. We assume that
141 t is a sink, it is reachable from s , and $E^{\rightarrow s} = \emptyset$. We also assume that $\langle V, E \rangle$ does not contain parallel
142 edges and self loops. For a vertex $v \in V$, let $c(\rightarrow v) = \sum_{e \in E^{\rightarrow v}} c(e)$ and $c(v \rightarrow) = \sum_{e \in E^{v \rightarrow}} c(e)$ be
143 the sums of capacities of edges that enter and leave v , respectively. We say that a vertex $v \in V$ is a
144 *funnel* if $c(v \rightarrow) < c(\rightarrow v)$. We use C_s to denote the total capacity of edges outgoing from the source,
145 thus $C_s = c(s \rightarrow)$.

146 A *preflow* in \mathcal{G} is a function $f : E \rightarrow \mathbb{R}$ that satisfies the following two properties:

- 147 ■ For every $e \in E$, we have that $0 \leq f(e) \leq c(e)$.
 - 148 ■ For every vertex $v \in V \setminus \{s\}$, the incoming flow to v is greater or equal to its outgoing flow.
- 149 Formally, $\sum_{e \in E^{\rightarrow v}} f(e) \geq \sum_{e \in E^{v \rightarrow}} f(e)$.

150 For a preflow f and an edge $e \in E$, we say that e is *saturated* if $f(e) = c(e)$. We extend f to
151 vertices: for every vertex $v \in V$, let $f(\rightarrow v) = \sum_{e \in E^{\rightarrow v}} f(e)$ and $f(v \rightarrow) = \sum_{e \in E^{v \rightarrow}} f(e)$. For a
152 vertex $v \in V \setminus \{s, t\}$, the *flow loss of f in v* , denoted $l_f(v)$, is the quantity that enters v and does
153 not leave v . Formally, $l_f(v) = f(\rightarrow v) - f(v \rightarrow)$. Then, $L_f = \sum_{v \in V \setminus \{s, t\}} l_f(v)$ is the *flow loss of*
154 f . The value of a preflow f , denoted $val(f)$, is $f(\rightarrow t)$; that is, the incoming flow to t . Note that
155 $val(f) = f(s \rightarrow) - L_f$. A *flow* is a preflow f with $L_f = 0$. A *maximum flow* is a flow with a maximal
156 value.

157 A *saturating preflow* is a preflow in which all edge in $E^{s \rightarrow}$ are saturated, and for every $v \in$
158 $V \setminus \{s, t\}$, we have $f(v \rightarrow) = \min\{f(\rightarrow v), c(v \rightarrow)\}$. That is, in a saturating preflow, flow loss may
159 occur in a vertex v only if the incoming flow to v is bigger than the capacities of the edges outgoing
160 from v .

161 The *unfortunate value* of a flow network \mathcal{G} , denoted $uval(\mathcal{G})$, is the minimal value of a saturating
162 preflow in \mathcal{G} . That is, it is the value that is guaranteed to reach t when the edges that leave s are
163 saturated, yet the most unfortunate routing decisions are taken in junctions. An *unfortunate saturat-*
164 *ing preflow* is a saturating preflow that attains the network's unfortunate value. The *unfortunate flow*
165 *problem* (UF problem, in short) is to decide, given a flow network \mathcal{G} and a threshold $\gamma > 0$, whether
166 $uval(\mathcal{G}) \geq \gamma$.

3 The Complexity of the Unfortunate-Flow Problem

168 In this section we study the complexity of the unfortunate-flow problem. We start with bad news and
169 show that the problem is co-NP-hard, and in fact is NP-hard to approximate within any multiplicative

170 factor. A more precise analysis of the complexity then enables us to point to a polynomial fragment
 171 and to prove an integral-flow property, which implies a matching co-NP upper bound.

172 ► **Theorem 2.** *The UF problem is co-NP-hard.*

Proof. We show a reduction from CNF-SAT to the complement problem, namely deciding whether $uval(\mathcal{G}) < \gamma$ for some $\gamma \in \mathbb{N}$. Let $\psi = C_1 \wedge \dots \wedge C_m$ be a CNF formula over the variables $x_1 \dots x_n$. We assume that every literal in $x_1, \dots, x_n, \bar{x}_1, \dots, \bar{x}_n$ appears in exactly k clauses in ψ . Indeed, every CNF formula can be converted to such a formula in polynomial time and with a polynomial blowup. We construct a flow network $\mathcal{G} = \langle V, E, c, s, t \rangle$ as demonstrated in Figure 2. Let $Z = \{x_1, \dots, x_n, \bar{x}_1, \dots, \bar{x}_n\}$. For a literal $z \in Z$ and a clause C_i , the network \mathcal{G} contains an edge $\langle z, C_i \rangle$ iff C_i contains z . Thus, each vertex in Z has exactly k outgoing edges. The capacities of these edges are 1. Each vertex C_i has two outgoing edges – to t and to the sink u . In Appendix A.1, we prove that ψ is satisfiable iff $uval(\mathcal{G}) < kn - m + 1$.




Figure 2: The flow network \mathcal{G} . The capacities of the edges entering C_1, \dots, C_m are 1.

174 By a simple manipulation of the network \mathcal{G} constructed in the reduction in the proof of Theorem 2, we can obtain, given a CNF-SAT formula ψ , a network \mathcal{G}' such that if ψ is satisfiable, then
 175 $uval(\mathcal{G}') = 0$, and otherwise, $uval(\mathcal{G}') \geq 1$. Hence the following (a detailed proof can be found in
 176 Appendix A.2).
 177

178 ► **Theorem 3.** *It is NP-hard to approximate the UF problem within any multiplicative factor.*

179 Following the hardness of the problem, we turn to analyze its complexity in terms of the different
 180 parameters of the flow network. Our analysis points to a class of networks for which the UF problem
 181 can be solved in polynomial time.

182 Consider a flow network $\mathcal{G} = \langle V, E, c, s, t \rangle$. Let $H \subseteq V \setminus \{s, t\}$ be the set of funnels in \mathcal{G} .
 183 Thus, $H = \{v : c(\rightarrow v) > c(v \rightarrow)\}$. For $L \subseteq V$, let \mathcal{F}_L be a set of saturating preflows in which
 184 edges outgoing from vertices in L are saturated, and flow loss may occur only in vertices in L . Thus,
 185 $f \in \mathcal{F}_L$ iff f is a saturating preflow in \mathcal{G} such that for every $u \in L$, we have $f(u \rightarrow) = c(u \rightarrow)$, and
 186 for every $u \in V \setminus L$, we have $l_f(u) = 0$. By the definition of a saturating flow, flow loss in \mathcal{G} may
 187 occur only in vertices in H . Accordingly, we have the following.

188 ► **Lemma 4.** $\bigcup_{L \subseteq H} \mathcal{F}_L$ contains all the saturating preflows in \mathcal{G} .

189 By Lemma 4, a search for the unfortunate value of \mathcal{G} can be restricted to preflows in \mathcal{F}_L , for
 190 $L \subseteq H$. Accordingly, the UF problem can be solved by solving $2^{|H|}$ optimization problems, solv-
 191 able by either linear programming (Theorem 5) or a reduction to the min-cost max-flow problem
 192 (Theorem 6).

193 ► **Theorem 5.** *Consider a flow network \mathcal{G} and let H be the set of funnels in \mathcal{G} . The UF problem
 194 for \mathcal{G} can be solved in time $2^{|H|} \cdot poly(|\mathcal{G}|)$.*

195 **Proof.** The algorithm goes over all the subsets of H and for each subset $L \subseteq H$, finds a minimum-
 196 value preflow in \mathcal{F}_L . The latter is done by linear programming. Given L , the linear program for \mathcal{F}_L

125:6 The Unfortunate-Flow Problem

197 is described below. The variable x_e , for every $e \in E$, stands for $f(e)$. The program is of size linear
 198 in $|\mathcal{G}|$, thus the overall complexity is $2^{|H|} \cdot \text{poly}(|\mathcal{G}|)$.

$$\begin{array}{ll}
 \text{minimize} & \sum_{e \in E \rightarrow t} x_e \\
 \text{subject to} & 0 \leq x_e \leq c(e) \quad \text{for each } e \in E \\
 & x_e = c(e) \quad \text{for each } u \in L \cup \{s\}, e \in E^{u \rightarrow} \\
 & \sum_{e \in E^{u \rightarrow}} x_e \leq \sum_{e \in E \rightarrow u} x_e \quad \text{for each } u \in L \\
 & \sum_{e \in E^{u \rightarrow}} x_e = \sum_{e \in E \rightarrow u} x_e \quad \text{for each } u \notin L \cup \{s, t\}
 \end{array}$$

200

201 The complexity of solving each linear program in the algorithm described in the proof of The-
 202 orem 5 is polynomial in $|\mathcal{G}|$, but not *strongly polynomial*. Thus its running time depends (polyno-
 203 mially) on the number of bits required for representing the capacities in \mathcal{G} . We now describe an
 204 alternative algorithm whose complexity depends only on the number of vertices and edges in the
 205 network.

206 Our algorithm reduces the problem of finding a minimal-value preflow in \mathcal{F}_L to the *min-cost*
 207 *max-flow* problem in flow networks with costs [1]. A flow network with costs is $\mathcal{G} = \langle V, E, a, c, s, t \rangle$,
 208 where $\langle V, E, c, s, t \rangle$ is a flow network and $a : E \rightarrow \mathbb{R}$ is a cost function. The cost of a flow f in
 209 \mathcal{G} , denoted $\text{cost}(f)$, is $\sum_{e \in E} a(e) \cdot f(e)$. In the min-cost max-flow problem we are given a flow
 210 network with costs, and find a maximum flow with a minimum cost. By [1], this problem can be
 211 solved in time $O(|E|^2 \log |V| + |E||V| \log^2 |V|)$.

212 ► **Theorem 6.** Consider a flow network $\mathcal{G} = \langle V, E, c, s, t \rangle$ and let H be the set of funnels in \mathcal{G} .
 213 The UF problem for \mathcal{G} can be solved in time $O(2^{|H|} \cdot (|E|^2 \log |V| + |E||V| \log^2 |V|))$.

214 **Proof.** The algorithm finds, for each subset $L \subseteq H$, a minimum-value preflow in \mathcal{F}_L by a reduction
 215 to the min-cost max-flow problem. By Lemma 4, the minimum value found for some $L \subseteq H$ is
 216 $\text{val}(\mathcal{G})$.

217 Consider the flow network with costs $\mathcal{G}' = \langle V', E', a, c', s, t' \rangle$ that is obtained from \mathcal{G} as follows.
 218 We add a new vertex t' and edges $\langle u, t' \rangle$ for every $u \in L \cup \{t\}$, thus $V' = V \cup \{t'\}$ and $E' =$
 219 $E \cup (L \cup \{t\}) \times \{t'\}$. The capacity $c'(e)$ for every new edge $e \in E' \setminus E$ is large (for example, it
 220 may be C_s), and for every $e \in E$, we have $c'(e) = c(e)$. Let $C = \max\{c(e) : e \in E\}$ denote the
 221 maximal capacity in \mathcal{G} . For edges $e \in L \times \{t'\}$, we define $a(e) = -1$; for edges $e \in L \times V$, we
 222 define $a(e) = -C \cdot |V|^2$; and for all the other edges, we define $a(e) = 0$. Intuitively, the costs of
 223 the edges in $L \times V$ are negative and small enough, so that a min-cost max-flow in \mathcal{G}' would have to
 224 saturate them first, and only then try to direct flow to edges in $L \times \{t'\}$.

225 In Appendix A.3, we prove the correctness of the following algorithm: First, find a min-cost
 226 max-flow f' in \mathcal{G}' . If $\text{val}(f') < C_s$ or $\text{cost}(f') > -C \cdot |V|^2 \cdot \sum_{e \in L \times V} c(e)$, then $\mathcal{F}_L = \emptyset$.
 227 Otherwise, the minimal value of a preflow in \mathcal{F}_L is $C_s + \text{cost}(f') + C \cdot |V|^2 \cdot \sum_{e \in L \times V} c(e)$. Since
 228 the min-cost max-flow problem can be solved in time $O(|E|^2 \log |V| + |E||V| \log^2 |V|)$ [1] and there
 229 are $2^{|H|}$ subsets of funnels to check, the required complexity follows. ◀

230 ► **Corollary 7.** The UF problems for networks with a logarithmic number of funnels can be solved
 231 in strongly-polynomial time.

232 We say that a preflow $f : E \rightarrow \mathbb{R}$ is *integral* if $f(e) \in \mathbb{N}$ for all $e \in E$. It is sometimes desirable
 233 to restrict the flow to an integral one, for example in settings in which the objects we transfer along
 234 the network cannot be partitioned into fractions. We now show that the UF problem always has an
 235 integral-flow solution, and that such a solution can be obtained by the algorithm shown in the proof
 236 of Theorem 6. Essentially (see proof in Appendix A.4), it follows from the fact that the min-cost

237 max-flow problem has an integral solution. As we show in Section 4, this *integral flow property* is
 238 not maintained in variants of the UF problem.

239 ► **Theorem 8.** *The UF problem has an integral-flow solution: for every flow network, there exists*
 240 *an integral unfortunate saturating preflow. Moreover, such integral preflow can be found by the*
 241 *algorithm described in the proof of Theorem 6.*

242 The integral-flow property suggests an optimal algorithm for solving the UF problem:

243 ► **Theorem 9.** *The UF problem is co-NP-complete.*

244 **Proof.** Hardness in co-NP is proven in Theorem 2. We prove membership in NP for the comple-
 245 mentary problem: given $\gamma > 0$ and a flow network \mathcal{G} , we need to decide whether $uval(\mathcal{G}) < \gamma$.
 246 According to Theorem 8, it is enough to decide whether there is a saturating preflow f in which for
 247 every $e \in E$, the value $f(e)$ is an integer, and $val(f) < \gamma$. Given a function $f : E \rightarrow \mathbb{N}$, checking
 248 whether f satisfies these requirements can be done in polynomial time, implying membership in
 249 NP. ◀

250 4 The Controlled Unfortunate-Flow Problem

251 In this section we study the controlled unfortunate-flow problem, where the outgoing flow from s is
 252 bounded and controlled. That is, there is $0 \leq \alpha \leq C_s$ such that the total outgoing flow from s is
 253 bounded by α , and it is possible to control how this outgoing flow is partitioned among the edges that
 254 leave s . Our goal is to control this flow so that the flow that reaches t in the worst case is maximized.
 255 As discussed in Section 1, this problem is motivated by scenarios where we have an initial control
 256 on the flow, say by positioning police at the entrance to a stadium or at the center of a city we need
 257 to evacuate, or by transmitting messages we want to send from a router we own.

258 For $\alpha \geq 0$, a *regulator with bound α* is a function $g : E^{s \rightarrow} \rightarrow \mathbb{R}$ that directs α flow units from
 259 s . Formally, for every $e \in E^{s \rightarrow}$, we have $0 \leq g(e) \leq c(e)$, and $\sum_{e \in E^{s \rightarrow}} g(e) \leq \alpha$. A *controlled*
 260 *saturating preflow that respects a regulator g* is a preflow $f : E \rightarrow \mathbb{R}$ such that for every $e \in E^{s \rightarrow}$,
 261 we have $f(e) = g(e)$, and for every $v \in V \setminus \{s, t\}$, we have $f(v \rightarrow) = \min\{f(\rightarrow v), c(v \rightarrow)\}$.
 262 Thus, unlike saturating preflow, here the edges in $E^{s \rightarrow}$ need not be saturated and the flow in them is
 263 induced by g . The *unfortunate g -controlled value* of \mathcal{G} , denoted $cval(\mathcal{G}, g)$, is the minimal value
 264 of a controlled saturating preflow that respects g . Then, the *unfortunate α -controlled value* of \mathcal{G} ,
 265 denoted $cval(\mathcal{G}, \alpha)$ is the maximal unfortunate g -controlled value of \mathcal{G} for some regulator g with
 266 bound α . In the *controlled unfortunate flow problem* (CUF problem, for short), we are given a flow
 267 network \mathcal{G} , a bound $\alpha \geq 0$, and a threshold $\gamma > 0$, and we need to decide whether $cval(\mathcal{G}, \alpha) \geq \gamma$.
 268 Thus, in the CUF problem we need to decide whether there is a regulator g with bound α that ensures
 269 a value of at least γ .

270 For two regulators g and g' , we denote $g \geq g'$ if for every $e \in E^{s \rightarrow}$, we have $g(e) \geq g'(e)$. In the
 271 following theorem we show that the g -controlled unfortunate value is monotonic with respect to g .
 272 Thus, increasing g can only increase the value. In particular, it follows that a maximal $cval(\mathcal{G}, \alpha)$
 273 is obtained with $\alpha = C_s$ and a regulator g in which $g(e) = c(e)$ for every $e \in E^{s \rightarrow}$. Thus, if the
 274 outgoing flow from s is not bounded, then the optimal behavior is to saturate the edges in $E^{s \rightarrow}$.
 275 Essentially (see full proof in Appendix A.5), it follows from the fact that given g and g' such that
 276 $g \geq g'$, and a minimum-value controlled saturating preflow f that respects g , we can construct a
 277 controlled saturating preflow f' that respects g' and such that $val(f) \geq val(f')$.

278 ► **Theorem 10.** *Consider a flow network \mathcal{G} , and let g, g' be two regulators such that $g \geq g'$. Then,*
 279 *$cval(\mathcal{G}, g) \geq cval(\mathcal{G}, g')$.*

280 We now turn to study the complexity of the CUF problem. We first explain why the problem
 281 is challenging. One could expect an algorithm in which, given \mathcal{G} , α , and γ , we guess an *integral*
 282 *regulator* $g : E^{s \rightarrow} \rightarrow \mathbb{N}$ with bound α , and then use an NP oracle in order to check whether
 283 $\text{cval}(\mathcal{G}, g) \geq \gamma$. The problem with the above idea is that it restricts the regulators to integral ones.
 284 In Theorem 11 below we show that in some cases, an optimal regulator must use non-integral values.
 285 Accordingly, an algorithm that guesses a regulator, as has been the case with the guessed flows in
 286 Theorem 9, has to go over unboundedly many possibilities. In fact, when an arbitrary set of vertices
 287 (rather than the source only) may be controlled, the problem is not known to be decidable [14].

288 ► **Theorem 11.** *Integral regulators are not optimal: There is a flow network \mathcal{G} such that for every*
 289 *integral regulator $g : E^{s \rightarrow} \rightarrow \mathbb{N}$ with bound 2 we have $\text{cval}(\mathcal{G}, g) = 1$, but there is a regulator*
 290 *$g' : E^{s \rightarrow} \rightarrow \mathbb{R}$ with bound 2 such that $\text{cval}(\mathcal{G}, g') = 2$.*

Proof. Consider the flow network \mathcal{G} appearing in Figure 3. For every pair $\langle u_i, u_j \rangle$ for $1 \leq i < j \leq 4$, the network \mathcal{G} contains a vertex v_{ij} with incoming edges from u_i and u_j . The capacities of the edges in \mathcal{G} are all 1. It is not hard to see that for every integral regulator g with bound 2 we have $\text{cval}(\mathcal{G}, g) = 1$. Indeed, for such g there is a controlled saturating preflow that respects g , which directs a flow of 2 to some vertex v_{ij} , causing a loss of 1 in v_{ij} . Consider now the regulator g' that assigns a flow of 0.5 to every edge in $E^{s \rightarrow}$. In this case, a flow of more than 1 cannot be directed to any vertex v_{ij} and therefore $\text{cval}(\mathcal{G}, g') = 2$. ◀




Figure 3: The flow network \mathcal{G} . All capacities are 1.

291 We turn to solve the CUF problem. Theorem 11 forces us to consider non-integral regulators.
 292 We do this by a reduction to a problem with a similar challenge, namely *the second alternation*
 293 *level of the theory of real numbers under addition and order*. There, we are given a formula of
 294 the form $\exists x_1 \dots x_n \forall y_1 \dots y_m F(x_1, \dots, x_n, y_1, \dots, y_m)$, where F is a propositional combination
 295 of linear inequalities of the form $a_1 x_1 + \dots + a_n x_n + b_1 y_1 + \dots + b_m y_m \leq d$, for constant integers
 296 $a_1, \dots, a_n, b_1, \dots, b_m$, and d , and we have to decide whether there is an assignment of x_1, \dots, x_n
 297 to real numbers so that F is satisfied for every assignment of y_1, \dots, y_m to real numbers. Even
 298 though the domain of possible solutions is infinite, It is shown in [17] that the problem can be solved
 299 in Σ_2^P , namely the class of problems that can be solved by a nondeterministic polynomial Turing
 300 machine that has an oracle to some NP-complete problem. In [14], a Σ_2^P lower bound is proven
 301 for the problem of finding the value of a flow game, where the outgoing flow of a subset of the
 302 vertices can be controlled. Recall that in the CUF problem, only the flow from the source vertex can
 303 be controlled. While this corresponds to the “exists-forall” nesting of quantifiers that characterizes
 304 reasoning in Σ_2^P , it not clear how to reduce Boolean formulas to unfortunate flows. Indeed, in
 305 the reduction in [14], control in intermediate vertices is used in order to model disjunctions in the
 306 formulas. In the CUF problem, such a control is impossible, as all vertices in the network except for
 307 the source are treated in a conjunctive manner.
 308

309 ► **Theorem 12.** *The CUF problem is Σ_2^P -complete.*

310 **Proof.** We first prove membership in Σ_2^P by a reduction to the second alternation level of the theory
 311 of real numbers under addition and order. Given G , α , and γ , we construct a propositional combin-
 312 ation F of linear inequalities over the variables x_e for every $e \in E^{s \rightarrow}$, and variables y_e , for every
 313 $e \in E$. The formula F states that the values of the variables x_e corresponds to a regulator g with
 314 bound α , and that if the values of the variables y_e correspond to a controlled saturating preflow f

315 that respects g then $val(f) \geq \gamma$. Then, our problem amounts to deciding whether there are real
 316 values x_e such that for every real values y_e the formula F holds.

317 For the lower bound, we describe a reduction from QBF₂, namely satisfiability for quanti-
 318 fied Boolean formulas with one alternation of quantifiers, where the external quantifier is “ex-
 319 exists”. Let ψ be a propositional formula over the variables $x_1, \dots, x_n, y_1, \dots, y_m$, and let $\theta =$
 320 $\exists x_1 \dots \exists x_n \forall y_1 \dots \forall y_m \psi$. Also, let $X = \{x_1, \dots, x_n\}$, $\bar{X} = \{\bar{x}_1, \dots, \bar{x}_n\}$, $Y = \{y_1, \dots, y_m\}$,
 321 $\bar{Y} = \{\bar{y}_1, \dots, \bar{y}_m\}$, $Z = X \cup Y$, and $\bar{Z} = \bar{X} \cup \bar{Y}$. We construct a flow network \mathcal{G}_θ and define α
 322 and γ , such that θ holds iff there is a regulator g with bound α such that $cuvall(\mathcal{G}_\theta, g) \geq \gamma$.

323 We assume that ψ is given in a positive normal form; that is, ψ is constructed from the literals in
 324 $Z \cup \bar{Z}$ using the Boolean operators \vee and \wedge , and that there is $k \geq 1$ such that every literal in $Z \cup \bar{Z}$
 325 appears in ψ exactly k times. Clearly, every Boolean propositional formula can be converted with
 326 only a quadratic blow-up to an equivalent one that satisfies these conditions.

327 We first translate ψ into a Boolean circuit \mathcal{C}_ψ with $k(2n + 2m)$ inputs – one for each occurrence
 328 of a literal in ψ . For example, in Figure 4, on the left, we describe \mathcal{C}_ψ for $\psi = x \vee (\bar{x} \wedge y) \wedge ((x \wedge$
 329 $\bar{y}) \vee (y \vee \bar{y} \vee \bar{x}))$. Each gate in \mathcal{C}_ψ has fan-in 2 and fan-out 1. We say that an input assignment to
 330 \mathcal{C}_ψ is consistent if it corresponds to an assignment to the variables in Z . That is, for each variable
 331 $z \in Z$, there is a value $b \in \{0, 1\}$ such that all the k inputs that correspond to the literal z have value
 332 b and all the k inputs that correspond to the literal \bar{z} have value $1 - b$. If the input to \mathcal{C}_ψ is consistent
 333 then \mathcal{C}_ψ computes the value of ψ for the corresponding assignment.




Figure 4: The Boolean circuit \mathcal{C}_ψ and the external-source flow network \mathcal{G}_ψ .

334 Now, we translate \mathcal{C}_ψ to an external-source flow network $\mathcal{G}_\psi = \langle V, E, c, t \rangle$: a flow network in
 335 which there is no source vertex, and an input flow is given externally. Formally, some of the edges
 336 in E have an unspecified source, to be later connected to edges with an unspecified target. The idea
 337 behind the translation is as follows: The capacities in \mathcal{G}_ψ are all 1. Each OR gate in \mathcal{C}_ψ induces a
 338 vertex v that has in-degree 2 and out-degree 1. Thus, if the incoming flow in each incoming edge to
 339 v is 0 or 1, then its outgoing flow is 1 iff at least one of its incoming edges has flow 1. Then, each
 340 AND gate in \mathcal{C}_ψ induces a vertex v that has in-degree 2 and out-degree 2, yet, one of the two edges
 341 that leaves v leads to a sink. Accordingly, if the incoming flow in each incoming edge to v is 0 or 1,
 342 then the outgoing flow in the edge that does not lead to the sink must be 1 iff both incoming edges
 343 have flow 1. For example, the Boolean circuit \mathcal{C}_ψ from Figure 4 is translated to the external-source
 344 flow network \mathcal{G}_ψ to its right.

345 Given a flow from the external source, we define the unfortunate value of \mathcal{G}_ψ as the minimal
 346 value of a controlled saturating preflow that respects the external flow. The following lemma can be
 347 easily proved by induction on the structure of ψ .

348 ► **Lemma 13.** Consider a Boolean formula ψ and its corresponding external-source flow network
 349 \mathcal{G}_ψ .

- 350 1. Given input flows to \mathcal{G}_ψ , if we increase some input flow, then the new unfortunate value of \mathcal{G}_ψ is
 351 greater than or equal to the original unfortunate value.

352 2. Given input flows in $\{0, 1\}$ to \mathcal{G}_ψ , the unfortunate value of \mathcal{G}_ψ is equal to the output of \mathcal{C}_ψ with
 353 the same input. Thus, if the input flow to \mathcal{G}_ψ corresponds to a consistent input to \mathcal{C}_ψ , then the
 354 unfortunate value of \mathcal{G}_ψ is the value of ψ for the corresponding assignment.

355 We complete the reduction by constructing the flow network \mathcal{G}_θ that uses \mathcal{G}_ψ as a sub-network
 356 as shown in Figure 5. The vertices d_{y_1}, \dots, d_{y_m} are associated with the variables in Y . The vertices
 357 $x_i, \bar{x}_i, y_i, \bar{y}_i$ for every i are associated with the literals in $Z \cup \bar{Z}$. Each outgoing edge from a literal
 358 vertex that enters \mathcal{G}_ψ is connected to an input of \mathcal{G}_ψ that corresponds to this literal. The outgoing
 359 edge from the subnetwork \mathcal{G}_ψ corresponds to an edge from the target vertex of \mathcal{G}_ψ . In Appendix A.6
 360 we describe the network \mathcal{G}_θ for the case $\psi = (x \vee y) \wedge (\bar{x} \vee \bar{y})$.




Figure 5: The flow network \mathcal{G}_θ .

361 In Appendix A.7, we prove that θ holds iff there is a regulator g with bound $(2k+1)m + (k+1)n$
 362 such that for every controlled saturating preflow f that respects g we have $val(f) \geq m + n + 1$.

364 In Theorem 11 we showed that in some cases an optimal regulator must use non-integral values.
 365 Sometimes, however, it is desirable to restrict attention to integral regulators. In the following the-
 366 orem (see proof in Appendix A.8) we show that the Σ_2^P -completeness stays valid also for integral
 367 regulators.

368 ► **Theorem 14.** Let \mathcal{G} be a flow network and let α, γ be integral constants. Deciding whether there
 369 exists an integral regulator $g : E^{s \rightarrow} \rightarrow \mathbb{N}$ with bound α such that $cval(\mathcal{G}, g) \geq \gamma$, is Σ_2^P -complete.

370 ► **Remark. [Bounded Global Control]** In the CUF problem, it is possible to control the flow
 371 leaving the source. This could be generalized by letting an authority control also internal vertices in
 372 the network. In the *bounded global control* problem, we get as input a flow network \mathcal{G} , a number
 373 $k \geq 0$, and a threshold $\gamma > 0$, and we need to decide whether we can guarantee an unfortunate flow
 374 of at least γ by controlling the outgoing flow in at most k vertices. Note that while in the problem of
 375 finding critical nodes for firefighters [2, 3], a firefighter blocks the fire, in our setting the firefighters
 376 direct the evacuation. Thus, there, the goal is to block undesired vulnerabilities in the network, and
 377 here the goal is maximize desired traffic in the network. The formal definition of the bounded global
 378 control problem goes through the flow games of [14], which includes the notion of strategies for
 379 controlling flow. The Σ_2^P algorithm for solving flow games with integral flows can be extended to
 380 solve the bounded global control problem. By making the control on the source vertex essential (say,

381 by adding a transition with a large capacity to a sink), the CUF problem can be reduced to the global
 382 control problem with $k = 1$, implying Σ_2^P completeness.

383 5 Safe Networks and No-Loss Unfortunate Flow

384 In this section we consider settings in which loss must be avoided. We say that a flow network \mathcal{G} is
 385 *safe* if $L_f = 0$ for every saturating preflow f . For example, networks with no funnels are clearly
 386 safe. It is easy to see that \mathcal{G} is safe iff $uval(\mathcal{G}) = C_s$. Together with Theorem 9, this gives a co-NP
 387 algorithm for deciding the safety of a network. We first show that by reducing the safety problem
 388 to the maximum weighted flow problem, we can decide safety in polynomial time. Essentially, the
 389 reduction checks, for every vertex $v \in V$, whether it is possible to direct to v flow that is greater
 390 than its outgoing capacity, and the weights are used in order to filter flow incoming to v . For details,
 391 see Appendix A.9.

392 ► **Theorem 15.** *Deciding whether a flow network is safe can be done in polynomial time.*

393 We now consider the case where the total outgoing flow from s is controlled, and we need to find
 394 an optimal regulator that guarantees no flow loss. Formally, in the *no-loss controlled unfortunate-*
 395 *flow problem* (NLCUF problem, for short), we are given a flow network \mathcal{G} and an integer $\gamma > 0$,
 396 and we need to decide whether there exists a regulator g such that $\sum_{e \in E^{s \rightarrow}} g(e) \geq \gamma$, and for every
 397 controlled saturating preflow that respects g the flow loss is 0 (equivalently, $cuval(\mathcal{G}, g) = \gamma$). That
 398 is, decide whether there is a regulator that ensures no loss and a value of at least γ . We show
 399 that the NLCUF problem is NP-complete. For the upper bound one could expect an algorithm in
 400 which we guess an integral regulator $g : E^{s \rightarrow} \rightarrow \mathbb{N}$ in which the total flow is at least γ , and
 401 then use Theorem 15 in order to check in polynomial time whether flow loss is possible. However,
 402 Theorem 11 shows that in some cases a regulator must use non-integral values in order to ensure that
 403 flow loss is impossible. Consequently, our algorithm is more complicated and uses a result from the
 404 theory of real numbers with addition.

405 ► **Theorem 16.** *The NLCUF problem is NP-complete.*

406 **Proof:** We start with the upper bound. For a rational number q we denote by $\#(q)$ the *length* of q ,
 407 namely, if $q = a/b$ with a, b relatively prime, then $\#(q)$ is the sum of the number of bits in the binary
 408 representations of a and b . Consider a formula $\varphi = \exists x_1, \dots, x_n \forall y_1, \dots, y_m F(x_1, \dots, x_n, y_1, \dots, y_m)$,
 409 where F is a propositional combination of linear inequalities of the form $a_1 x_1 + \dots + a_n x_n + b_1 y_1 +$
 410 $\dots + b_m y_m \leq d$ for integral constants $a_1, \dots, a_n, b_1, \dots, b_m$, and d . The variables $x_1, \dots, x_n, y_1, \dots, y_m$
 411 are real. In [17] (in the proof of Theorem 3.1 there) it is shown that φ holds iff there exists rational
 412 values x_1, \dots, x_n such that for every i the length $\#(x_i)$ is polynomial in the size of φ and for every
 413 real values y_1, \dots, y_m the formula F holds.

414 We construct a propositional combination F of linear inequalities over the variables x_e , for
 415 every $e \in E^{s \rightarrow}$, and y_e , for every $e \in E$. The formula F states that the values of the variables x_e
 416 correspond to a regulator g with bound γ , and that if the values of the variables y_e correspond to a
 417 controlled saturating preflow f that respects g , then $L_f = 0$. Then, our problem amounts to deciding
 418 whether there are real values x_e such that for every real values y_e , the formula F holds. By [17], it
 419 is enough to check whether there are rational values x_e for $e \in E^{s \rightarrow}$ with polynomial lengths such
 420 that for every real values y_e for $e \in E$, the formula F holds. Given values for the variables x_e ,
 421 checking whether for every real values y_e the formula F holds can be done in polynomial time with
 422 the algorithm shown in the proof of Theorem 15. Hence the membership in NP.

125:12 The Unfortunate-Flow Problem

We proceed to the lower bound. We show a reduction from CNF-SAT. Let $\psi = C_1 \wedge \dots \wedge C_m$ be a CNF formula over the variables $x_1 \dots x_n$. We denote $Z = \{x_1, \dots, x_n, \bar{x}_1, \dots, \bar{x}_n\}$. We assume that for every literal $z \in Z$ there is at least one clause in ψ that does not contain z . We construct a flow network $\mathcal{G} = \langle V, E, c, s, t \rangle$ as demonstrated in Figure 6. For a literal $z \in Z$ and a clause C_i , the network \mathcal{G} contains an edge $\langle z, C_i \rangle$ iff the clause C_i does not contain the literal z . Let $\gamma = 2n$. In Appendix A.10, we show that ψ is satisfiable iff there is a regulator that ensures no loss and a value of at least γ .




Figure 6: The flow network \mathcal{G} . Unless stated otherwise, the capacities are 1.

Sometimes it is desirable to restrict attention to integral regulators. As we show below, NP-completeness applies for them too (see Appendix A.11 for proof).

Theorem 17. *Let \mathcal{G} be a flow network and let $\gamma > 0$ be an integer. Deciding whether there exists an integral regulator $g : E^{s \rightarrow} \rightarrow \mathbb{N}$ in \mathcal{G} such that $\sum_{e \in E^{s \rightarrow}} g(e) \geq \gamma$, and for every controlled saturating preflow that respects g the flow loss is 0, is NP-complete.*

6 Discussion

The unfortunate-flow problem captures settings in which the authority has no control on how flow is directed in the vertices of a flow network. For many problems, a transition from a cooperative setting to an adversarial one dualizes the complexity class to which the problem belongs, as in NP for satisfiability vs. co-NP for validity. In the case of flow, the polynomial complexity of the maximum-flow problem is not preserved when we move to the dual unfortunate-flow problem, and we prove that the problem is co-NP-complete.

On the positive side, the integral-flow property of maximal flow is preserved in unfortunate flows. This property, however, is lost once we move to controlled unfortunate flows, where non-integral regulators may be more optimal than integral ones. The need to consider real-valued flows questions the decidability of the controlled unfortunate-flow problem. As we show, the problem is decidable, by a reduction to the second alternation level of the theory of real numbers under addition and order [17]. There, the infinite domain of the real numbers is reduced to a finite one, namely rational numbers of length polynomial in the input. A direct algorithm for the controlled unfortunate-flow problem, thus one that does not rely on [17], is still open. Such a direct algorithm would reduce the real-number domain to a finite one in a tighter manner – one that depends on the network. We see several interesting problems in this direction, in particular finding a *sufficient granularity* that a regulator may need, and *bounding the non-optimality* caused by integral regulators. Similar problems are open in the settings of flow games with two or more players [14, 12].

Finally, the unfortunate-flow problem sets the stage to problems around network design, where the goal is to design networks with maximal unfortunate flows. In particular, in *network repair*, we are given a network and we are asked to modify it in order to increase its unfortunate flow value. Different algorithms correspond to different types of allowed modifications. For example, we may be allowed to change the capacity of a fixed number of edges. Note that unlike the case of maximal flow, here a repair may reduce the capacity of edges. Also, unlike the case of maximal flow, there is no clear theory of minimal cuts that may assist us in such a repair.

455 — **References** —

- 456 **1** R.K. Ahuja, T.L. Magnanti, and J.B. Orlin. *Network flows: Theory, algorithms, and applications*.
457 Prentice Hall Englewood Cliffs, 1993.
- 458 **2** C. Bazgan, M. Chopin, M. Cygan, M.R. Fellows, F. V. Fomin, and E. Leeuwen. Parameterized
459 complexity of firefighting. *Journal of Computer and Systems Science*, 80(7):1285–1297, 2014.
- 460 **3** J. Choudhary, A. Dasgupta, N. Misra, and M.S. Ramanujan. Saving critical nodes with firefighters is
461 FPT. In *Proc. 44th Int. Colloq. on Automata, Languages, and Programming*, volume 80 of *LIPICs*,
462 pages 135:1–135:13, 2017.
- 463 **4** T.H. Cormen, C.E. Leiserson, and R.L. Rivest. *Introduction to Algorithms*. MIT Press and
464 McGraw-Hill, 1990.
- 465 **5** E.A. Dinic. Algorithm for solution of a problem of maximum flow in a network with power estim-
466 ation. *Soviet Math. Doll*, 11(5):1277–1280, 1970. English translation by RF. Rinehart.
- 467 **6** J. Edmonds and R.M. Karp. Theoretical improvements in algorithmic efficiency for network flow
468 problems. *Journal of the ACM*, 19(2):248–264, 1972.
- 469 **7** S. Even, A. Itai, and A. Shamir. On the complexity of timetable and multicommodity flow prob-
470 lems. *SIAM Journal on Computing*, 5(4):691–703, 1976.
- 471 **8** L.R. Ford and D.R. Fulkerson. Maximal flow through a network. *Canadian journal of Mathematics*,
472 8(3):399–404, 1956.
- 473 **9** L.R. Ford and D.R. Fulkerson. *Flows in networks*. Princeton Univ. Press, Princeton, 1962.
- 474 **10** A.V. Goldberg, É. Tardos, and R.E. Tarjan. Network flow algorithms. Technical report, DTIC
475 Document, 1989.
- 476 **11** A.V. Goldberg and R.E. Tarjan. A new approach to the maximum-flow problem. *Journal of the*
477 *ACM*, 35(4):921–940, 1988.
- 478 **12** S. Guha, O. Kupferman, and G. Vardi. Multi-player flow games. In *Proc. 17th International*
479 *Conference on Autonomous Agents and Multiagent Systems*, 2018.
- 480 **13** S. Keren, A. Gal, and E. Karpas. Goal recognition design for non optimal agents. In *Proc. 29th*
481 *AAAI conference*, pages 3298–3304, 2015.
- 482 **14** O. Kupferman, G. Vardi, and M.Y. Vardi. Flow games. In *Proc. 37th Conf. on Foundations of*
483 *Software Technology and Theoretical Computer Science*, volume 93 of *Leibniz International Pro-*
484 *ceedings in Informatics (LIPICs)*, pages 38:38–38:16, 2017.
- 485 **15** A. Madry. Computing maximum flow with augmenting electrical flows. In *Foundations of Com-*
486 *puter Science (FOCS), 2016 IEEE 57th Annual Symposium on*, pages 593–602. IEEE, 2016.
- 487 **16** L. Qingsong, G. Betsy, and S. Shashi. Capacity constrained routing algorithms for evacuation
488 planning: A summary of results. In *International Symposium on Spatial and Temporal Databases*,
489 pages 291–307. Springer, 2005.
- 490 **17** E.D. Sontag. Real addition and the polynomial hierarchy. *Information Processing Letters*,
491 20(3):115–120, 1985.
- 492 **18** É. Tardos. A strongly polynomial minimum cost circulation algorithm. *Combinatorica*, 5(3):247–
493 255, 1985.

494 **A** **Examples and Proofs**495 **A.1 A proof of the reduction in the proof of Theorem 2**

496 Assume that ψ is satisfiable and let τ be a satisfying assignment for ψ . Consider the following
497 saturating preflow f : In each vertex d_i , the incoming flow of $2k$ is directed to the vertex x_i if x_i
498 holds in τ and otherwise it is directed to the vertex \bar{x}_i . Then, in these vertices the incoming flow is
499 greater than the outgoing capacity and therefore the outgoing edges are saturated. In every vertex C_i
500 with a positive incoming flow, f directs a flow of 1 to the sink u and the rest is directed to t . Since

501 τ satisfies ψ , every vertex C_i has a positive incoming flow and therefore $val(f) = kn - m$. Hence
 502 $uval(\mathcal{G}) < kn - m + 1$.

503 Assume now that ψ is not satisfiable and let f be a saturating preflow. For each vertex d_i , the
 504 outgoing flow in at least one outgoing edge is greater than or equal to k . We denote by l_i a literal in
 505 $\{x_i, \bar{x}_i\}$ that has an incoming flow of at least k . These literals induce an assignment to x_1, \dots, x_n .
 506 By our assumption, this assignment does not satisfy ψ . Thus, at least one vertex in C_1, \dots, C_m has
 507 an incoming flow of 0 from the vertices l_1, \dots, l_n . Hence, the total outgoing flow from the vertices
 508 l_1, \dots, l_n is kn and from this flow at most $m - 1$ is lost. Therefore, the incoming flow to t is at least
 509 $kn - (m - 1)$ and thus $uval(\mathcal{G}) \geq kn - m + 1$.

510 A.2 A proof of Theorem 3

511 The network \mathcal{G} constructed in the reduction in the proof of Theorem 2 is such that $uval(\mathcal{G}) \leq kn - m$
 512 if the given CNF-SAT formula ψ is satisfiable and $uval(\mathcal{G}) \geq kn - m + 1$ otherwise. Let \mathcal{G}' be
 513 the network obtained from \mathcal{G} by adding a new vertex v , making v the target of \mathcal{G}' , adding an edge
 514 with capacity $kn - m$ from t to a sink, and an edge with capacity 1 from t to v . Thus, if the flow
 515 that reaches t is at most $kn - m$, then all of it can be directed to the sink. Consequently, if ψ is
 516 satisfiable, then $uval(\mathcal{G}') = 0$, and otherwise, $uval(\mathcal{G}') \geq 1$. Since every approximation algorithm
 517 (within any factor) can determine whether the value is positive or 0, then approximation within any
 518 factor is NP-hard.

519 A.3 Correctness proof of the algorithm in the proof of Theorem 6

520 We first show that there exists a preflow in \mathcal{F}_L with flow loss greater than or equal to β for some $\beta \geq$
 521 0 iff there exists a flow in \mathcal{G}' with value C_s and cost less than or equal to $-C \cdot |V|^2 \cdot (\sum_{e \in L \times V} c(e)) -$
 522 β . Assume that there exists a preflow f in \mathcal{F}_L with $L_f \geq \beta$. This preflow induces a flow f'
 523 in \mathcal{G}' where for every $v \in L$ we have $f'(\langle v, t' \rangle) = l_f(v)$. Since f is a saturating preflow, then
 524 $val(f') = C_s$. Also, since $\sum_{u \in L} l_f(u) \geq \beta$ and for every $e \in L \times V$ we have $f(e) = c(e)$, then
 525 $cost(f') \leq -C \cdot |V|^2 \cdot (\sum_{e \in L \times V} c(e)) - \beta$. Conversely, assume that there exists a flow f' in \mathcal{G}'
 526 with $val(f') = C_s$ and $cost(f') \leq -C \cdot |V|^2 \cdot (\sum_{e \in L \times V} c(e)) - \beta$. It is known that for every
 527 network flow with costs there exists an integral min-cost max-flow [1]. Let f'' be an integral min-
 528 cost max-flow in \mathcal{G}' . Thus, $val(f'') = C_s$ and $cost(f'') \leq -C \cdot |V|^2 \cdot (\sum_{e \in L \times V} c(e)) - \beta$. Since
 529 $cost(f'') \leq -C \cdot |V|^2 \cdot \sum_{e \in L \times V} c(e)$ and f'' is integral, then for every $e \in L \times V$ we must have
 530 $f''(e) = c(e)$. Indeed, otherwise $cost(f'') = \sum_{e \in E'} f''(e)a(e) = (-C \cdot |V|^2) \cdot \sum_{e \in L \times V} f''(e) +$
 531 $(-1) \cdot \sum_{e \in L \times \{t'\}} f''(e) \geq (-C \cdot |V|^2) \cdot [(\sum_{e \in L \times V} c(e)) - 1] - |V| \cdot (|V| - 1) \cdot C > -C \cdot |V|^2 \cdot$
 532 $\sum_{e \in L \times V} c(e) + C \cdot |V|^2 - C \cdot |V|^2 = -C \cdot |V|^2 \cdot \sum_{e \in L \times V} c(e)$. The flow f'' induces a preflow f
 533 in \mathcal{F}_L with $L_f \geq \beta$.

534 Since there exists a preflow in \mathcal{F}_L with flow loss greater than or equal to β iff there exists a flow in
 535 \mathcal{G}' with value C_s and cost less than or equal to $-C \cdot |V|^2 \cdot (\sum_{e \in L \times V} c(e)) - \beta$, then we have that $\mathcal{F}_L \neq$
 536 \emptyset iff there exists a flow in \mathcal{G}' with value C_s and cost less than or equal to $-C \cdot |V|^2 \cdot \sum_{e \in L \times V} c(e)$.
 537 Note that a flow in \mathcal{G}' with value C_s is a maximum flow. Assume that $\mathcal{F}_L \neq \emptyset$ and let $\gamma \geq 0$ be
 538 the maximal flow loss in a preflow in \mathcal{F}_L , namely $\gamma = \max\{L_f : f \in \mathcal{F}_L\}$. Note that according
 539 to Theorem 5 there exists a maximum for this set. Let $-C \cdot |V|^2 \cdot (\sum_{e \in L \times V} c(e)) - \beta$ be the
 540 cost of a min-cost max-flow in \mathcal{G}' . Since there exists a preflow in \mathcal{F}_L with flow loss γ then there
 541 exists a maximum flow f' in \mathcal{G}' with $cost(f') \leq -C \cdot |V|^2 \cdot (\sum_{e \in L \times V} c(e)) - \gamma$. Therefore,
 542 $\beta \geq \gamma$. Conversely, since there exists a flow in \mathcal{G}' with value C_s and cost less than or equal to
 543 $-C \cdot |V|^2 \cdot (\sum_{e \in L \times V} c(e)) - \beta$, then there exists a preflow in \mathcal{F}_L with flow loss greater than or
 544 equal to β , and thus $\gamma \geq \beta$. Hence $\beta = \gamma$. Therefore, if the min-cost max-flow in \mathcal{G}' has value C_s
 545 and cost $\alpha = -C \cdot |V|^2 \cdot (\sum_{e \in L \times V} c(e)) - \beta$ for some $\beta \geq 0$, then the maximal flow loss of a

546 preflow in \mathcal{F}_L is $\beta = -\alpha - C \cdot |V|^2 \cdot \sum_{e \in L \times V} c(e)$, and thus the minimal value of a preflow in \mathcal{F}_L
 547 is $C_s - \beta = C_s + \alpha + C \cdot |V|^2 \cdot \sum_{e \in L \times V} c(e)$.

548 Accordingly, the algorithm finds a min-cost max-flow f' in \mathcal{G}' . If $val(f') < C_s$ or $cost(f') >$
 549 $-C \cdot |V|^2 \cdot \sum_{e \in L \times V} c(e)$, then $\mathcal{F}_L = \emptyset$. Otherwise, the minimal value of a preflow in \mathcal{F}_L is
 550 $C_s + cost(f') + C \cdot |V|^2 \cdot \sum_{e \in L \times V} c(e)$. Since the min-cost max-flow problem can be solved in
 551 time $O(|E|^2 \log |V| + |E||V| \log^2 |V|)$ [1] and there are $2^{|H|}$ subsets of funnels to check, the required
 552 complexity follows.


553 A.4 A proof of Theorem 8

554 The algorithm shown in the proof of Theorem 6 runs over all the subsets of H and for each subset
 555 $L \subseteq H$ finds a minimum-value preflow in \mathcal{F}_L . In the proof of Theorem 6 it is shown that if
 556 $\mathcal{F}_L \neq \emptyset$ and $\gamma = \max\{L_f : f \in \mathcal{F}_L\}$, then a min-cost max-flow in \mathcal{G}' has value C_s and cost
 557 $-C \cdot |V|^2 \cdot (\sum_{e \in L \times V} c(e)) - \gamma$. Since the min-cost max-flow problem always has a solution in
 558 which all flows are integral [1], let f' be an integral min-cost max-flow in \mathcal{G}' . Thus $val(f') = C_s$
 559 and $cost(f') = -C \cdot |V|^2 \cdot (\sum_{e \in L \times V} c(e)) - \gamma$. In the proof of Theorem 6 it is also shown that an
 560 integral flow in \mathcal{G}' with value C_s and cost $-C \cdot |V|^2 \cdot (\sum_{e \in L \times V} c(e)) - \gamma$ induces a preflow in \mathcal{F}_L
 561 with flow loss γ . Hence, the preflow that f' induces in \mathcal{F}_L is integral and has a maximal flow loss.
 562 Since this is the case for every $L \subseteq H$, then the UF problem has an integral-flow solution.

563 A.5 A proof of Theorem 10

564 Let $\mathcal{G} = \langle V, E, c, s, t \rangle$, and let f be a minimum-value controlled saturating preflow that respects g .
 565 We construct a controlled saturating preflow f' that respects g' and such that $val(f') \leq val(f)$. The
 566 preflow f' is obtained from f by reducing flows in edges according to the following process: We
 567 start with $f' = f$. For every $e \in E^{s \rightarrow}$, if $g'(e) < g(e)$ then we change f' such that $f'(e) = g'(e)$.
 568 Now, for some vertices the incoming flow in f' might be less than the outgoing flow. We fix this
 569 with the following steps. Let $V = \{v_1, \dots, v_n\}$. We choose a vertex $v \in V \setminus \{s, t\}$ such that
 570 $f'(v \rightarrow) - f'(\rightarrow v)$ is maximal. Then, we find the minimal i such that $\langle v, v_i \rangle \in E$ and $f'(\langle v, v_i \rangle) > 0$
 571 and reduce the flow in $\langle v, v_i \rangle$ such that $f'(v \rightarrow) = f'(\rightarrow v)$. If there is not enough flow in $\langle v, v_i \rangle$,
 572 then we reduce the flow there to 0 and then find the next i such that $\langle v, v_i \rangle \in E$ and $f'(\langle v, v_i \rangle) > 0$
 573 and reduce the flow there too. We repeat this until $f'(v \rightarrow) = f'(\rightarrow v)$. Then we find the next vertex
 574 $v \in V \setminus \{s, t\}$ such that $f'(v \rightarrow) - f'(\rightarrow v)$ is maximal and repeat the above steps. We finish this
 575 process when for every $v \in V \setminus \{s, t\}$, we have $f'(v \rightarrow) = \min\{f'(\rightarrow v), c(v \rightarrow)\}$. Since f' is obtained
 576 from f by reducing flow in some edges, then $val(f') \leq val(f)$.

577 We show that the above process terminates after a finite number of iterations. Assume that the
 578 process does not terminate. Since we only reduce flows, there is a finite number of iterations in
 579 which a flow in an edge is reduced to 0. Thus, from some point, for every $v \in V$, there is an edge
 580 $e_v \in E^{v \rightarrow}$ such that a flow reduction in $E^{v \rightarrow}$ occurs only in e_v . Let $U \subseteq V$ be the vertices from
 581 which we reduce flow infinitely many times. Thus, from some point we reduce flow only for edges
 582 e_v for $v \in U$. Every vertex $v \in U$ must have a vertex $u \in U$ such that $e_u = \langle u, v \rangle$. Therefore, the
 583 subgraph induced by the vertices U and the edges $\{e_v : v \in U\}$ consists of disjoint cycles. Hence,
 584 if for some $v \in U$ we have $e_v = \langle v, w \rangle$ then $w \in U$. Note that from some point, for every vertex
 585 $v \in U$ the incoming flow is less than its outgoing capacity. Thus, from some point, if we reduce a
 586 flow of ϵ from an edge $e_u = \langle u, v \rangle$ then $u, v \in U$ and a flow of at least ϵ should be reduced later
 587 from the edge e_v . Since in every iteration we choose a vertex from which the flow reduction that is
 588 needed is maximal, then in the next iteration a flow of at least ϵ will be reduced. Hence, a flow of at
 589 least ϵ will be reduced in every iteration from this point and on. Since the flows are bounded below
 590 by 0 then this process cannot continue for infinitely many iterations.

591 **A.6 An example of the reduction in Theorem 12**Figure 7: The flow network \mathcal{G}_θ for $\theta = \exists x \forall y (x \vee y) \wedge (\bar{x} \vee \bar{y})$.592 **A.7 A proof of the reduction in the proof of Theorem 12**

593 we prove that θ holds iff there is a regulator g with bound $(2k+1)m + (k+1)n$ such that for every
594 controlled saturating preflow f that respects g we have $\text{val}(f) \geq m + n + 1$. Assume first that θ
595 holds. Let π be an assignment for X such that for every assignment for Y , the formula ψ holds.
596 Consider the regulator g where for every vertex $u \in d_{y_1}, \dots, d_{y_m}$ we have $g(\langle s, u \rangle) = 2k + 1$,
597 for every vertex x in X such that $\pi(x) = 1$ we have $g(\langle s, x \rangle) = k + 1$ and $g(\langle s, \bar{x} \rangle) = 0$, and
598 for every vertex x in X such that $\pi(x) = 0$ we have $g(\langle s, x \rangle) = 0$ and $g(\langle s, \bar{x} \rangle) = k + 1$. Note
599 that $\sum_{e \in E^{s \rightarrow}} g(e) = (2k+1)m + (k+1)n$. Thus, the input flows to \mathcal{G}_ψ for the variables X must
600 be consistent with the assignment π and for every i , the vertex u_{x_i} must have an incoming flow of
601 1. Since for each i , the vertex d_{y_i} has an incoming flow of $2k + 1$, then either y_i or \bar{y}_i
602 has an incoming flow of at least $k + 1$. Therefore, for every i , the outgoing edges of either y_i or \bar{y}_i
603 are saturated. Hence, according to Lemma 13 (1), a minimum-value controlled saturating preflow
604 should direct a flow of $2k + 1$ to either y_i or \bar{y}_i . Thus, the input flows to \mathcal{G}_ψ for the variables Y are
605 consistent with some assignment τ to these variables. According to Lemma 13 (2), the unfortunate
606 value of \mathcal{G}_ψ for these input flows is the value of ψ for the assignments π and τ , which is 1. Therefore,
607 for every controlled saturating preflow f that respects g we have $\text{val}(f) = m + n + 1$.

608 Now, assume that θ does not hold. We show that for every regulator g in \mathcal{G}_θ such that $\sum_{e \in E^{s \rightarrow}} g(e) \leq$
609 $(2k+1)m + (k+1)n$, there is a controlled saturating preflow f that respects g such that $\text{val}(f) <$
610 $m + n + 1$. First, in order to ensure an incoming flow of at least 1 to every vertex u_{y_i} , the regulator
611 g must have $g(\langle s, d_{y_i} \rangle) = 2k + 1$. Since $\sum_{e \in E^{s \rightarrow}} g(e) \leq (2k+1)m + (k+1)n$, then in order to
612 ensure also an incoming flow of 1 to every vertex u_{x_i} , the regulator g must have $g(\langle s, x_i \rangle) = k + 1$
613 or $g(\langle s, \bar{x}_i \rangle) = k + 1$. Therefore, the input flows to \mathcal{G}_ψ for the variables X are consistent with some
614 assignment π . According to the assumption, for every assignment to X there is an assignment to Y
615 such that ψ does not hold. Let τ be such an assignment to Y . A preflow that assigns to the vertices
616 $Y \cup \bar{Y}$ flows that respect τ , namely a flow of $2k + 1$ to the literals that hold in τ , will result in input
617 flows to \mathcal{G}_ψ for the variables Y that are consistent with τ . According to Lemma 13 (2), in this case
618 the unfortunate value of \mathcal{G}_ψ is 0 and therefore the incoming flow to t will be less than $m + n + 1$.

619 A.8 A proof of Theorem 14

620 According to Theorem 2, given an integral regulator g , deciding whether the minimal value of a
 621 controlled saturating preflow that respects g is at least γ is in co-NP. Thus we can guess an integral
 622 regulator g such that $\sum_{e \in E^{s \rightarrow}} g(e) \leq \alpha$ and then use an NP oracle in order to check whether the
 623 minimal value of a controlled saturating preflow that respects g is at least γ . Hence the Σ_2^P upper
 624 bound.

625 For the lower bound, note that the reduction described in the proof of Theorem 12 holds also for
 626 integral regulators.

627 A.9 A proof of Theorem 15

628 Consider a flow network $\mathcal{G} = \langle V, E, c, s, t \rangle$. For every vertex $v \in V \setminus \{s, t\}$, we construct a flow
 629 network \mathcal{G}_v in which every edge e is associated with a weight $a_v(e)$. The weighted flow network
 630 $\mathcal{G}_v = \langle V_v, E_v, a_v, c_v, s, t_v \rangle$ is obtained by extending \mathcal{G} as follows. We add a new vertex t_v and new
 631 edges from every vertex in V to t_v , thus, $V_v = V \cup \{t_v\}$ and $E_v = E \cup V \times \{t_v\}$. The capacity
 632 c_v for the new edges is large and for every $e \in E$ we have $c_v(e) = c(e)$. For every edge $e \in E^{-v}$
 633 we have $a_v(e) = 1$ and for every $e \notin E^{-v}$ we have $a_v(e) = 0$. We show that flow loss is possible
 634 in \mathcal{G} iff there is a vertex $v \in V$ such that the maximum weighted flow in \mathcal{G}_v has weight greater than
 635 $\sum_{v' \in V} c(\langle v, v' \rangle)$, that is, there is a flow f in \mathcal{G}_v such that $\sum_{e \in E_v} a_v(e) f(e) > \sum_{v' \in V} c(\langle v, v' \rangle)$.
 636 Intuitively, the maximum weighted flow in \mathcal{G}_v has weight greater than $\sum_{v' \in V} c(\langle v, v' \rangle)$ iff an incom-
 637 ing flow of more than $\sum_{v' \in V} c(\langle v, v' \rangle)$ to v is possible in \mathcal{G} iff flow loss is possible in v . Checking
 638 for every $v \in V$ whether the maximum weighted flow in \mathcal{G}_v has weight greater than $\sum_{v' \in V} c(\langle v, v' \rangle)$
 639 can be done in polynomial time by solving a linear program.

640 Assume that flow loss is possible. Let f be a saturating preflow in \mathcal{G} such that $l_f(v) > 0$ for
 641 some vertex $v \in V$, that is, the incoming flow to v is greater than its outgoing capacity. The preflow
 642 f induces a flow f' in \mathcal{G}_v , where the flow losses in f and the flow in t are directed to the vertex
 643 t_v . Thus, for every $u \in V \setminus \{s, t\}$ we have $f'(\langle u, t_v \rangle) = l_f(u)$ and $f'(\langle t, t_v \rangle) = f(\rightarrow t)$. In
 644 the flow f' the incoming flow to v is greater than the outgoing capacity from v in \mathcal{G} and therefore
 645 $\sum_{e \in E_v} a_v(e) f'(e) > \sum_{v' \in V} c(\langle v, v' \rangle)$.

646 Now, assume that there is a vertex $v \in V \setminus \{s, t\}$ such that there is a flow in \mathcal{G}_v with weight
 647 greater than $\sum_{v' \in V} c(\langle v, v' \rangle)$. Hence, there is a flow f in \mathcal{G}_v such that the incoming flow to the
 648 vertex v is greater than $\sum_{v' \in V} c(\langle v, v' \rangle)$. The flow f induces a preflow f' in \mathcal{G} , where for every
 649 $e \in E$ we have $f'(e) = f(e)$. Note that f' is a preflow in \mathcal{G} but it may not be a saturating preflow.
 650 We change f' according to the following steps in order to obtain a saturating preflow in \mathcal{G} such that
 651 the flow in each edge $e \in E$ is greater or equal to $f(e)$. We denote $V = \{v_1, \dots, v_n\}$. First, for
 652 every $e \in E^{s \rightarrow}$ we change f' such that $f'(e) = c(e)$. Then, we choose a vertex $u \in V \setminus \{s, t\}$ such
 653 that $\min\{f'(\rightarrow u), c(u \rightarrow)\} - f'(u \rightarrow)$ is maximal. We find the minimal i such that $\langle u, v_i \rangle \in E$ and
 654 $f'(\langle u, v_i \rangle) < c(\langle u, v_i \rangle)$ and increase the flow in $\langle u, v_i \rangle$ such that $f'(u \rightarrow) = \min\{f'(\rightarrow u), c(u \rightarrow)\}$.
 655 If there is not enough free capacity in $\langle u, v_i \rangle$ then we increase the flow there to $c(\langle u, v_i \rangle)$ and find
 656 the next i such that $\langle u, v_i \rangle \in E$ and $f'(\langle u, v_i \rangle) < c(\langle u, v_i \rangle)$ and increase the flow there also. We
 657 repeat this until $f'(u \rightarrow) = \min\{f'(\rightarrow u), c(u \rightarrow)\}$. Then, we find the next $u \in V \setminus \{s, t\}$ such
 658 that $\min\{f'(\rightarrow u), c(u \rightarrow)\} - f'(u \rightarrow)$ is maximal and repeat the above steps. We finish this process
 659 when for every $u \in V \setminus \{s, t\}$ we have $f'(u \rightarrow) = \min\{f'(\rightarrow u), c(u \rightarrow)\}$. Since in this process
 660 we only increase flows, then we still have that the incoming flow to the vertex v is greater than
 661 $\sum_{v' \in V} c(\langle v, v' \rangle)$ and therefore there is a flow loss in v .

662 We now show that the above process terminates after a finite number of iterations. Assume that
 663 the process does not terminate. Since we only increase flows, there is a finite number of iterations in
 664 which a flow in an edge $e \in E$ is increased to $c(e)$. Thus, from some point, for every $u \in V$ there is

125:18 The Unfortunate-Flow Problem

665 an edge $e_u \in E^{u \rightarrow}$ such that a flow increase in $E^{u \rightarrow}$ occurs only in e_u . Let $U \subseteq V$ be the vertices
666 in which we commit a flow increase infinitely many times. Thus, from some point we add flow only
667 to edges e_u for $u \in U$. Every vertex $w \in U$ must have a vertex $u \in U$ such that $e_u = \langle u, w \rangle$.
668 Therefore, the subgraph induced by the vertices U and the edges $\{e_u : u \in U\}$ consists of disjoint
669 cycles. Hence, if for some $u \in U$ we have $e_u = \langle u, w \rangle$ then $w \in U$. Note that for every vertex
670 $u \in U$ the flow in e_u is always less than $c(e_u)$. Thus, from some point, if we add a flow of ϵ to an
671 edge $e_u = \langle u, w \rangle$ then $u, w \in U$ and a flow of at least ϵ should be added later to the edge e_w . Since
672 in every iteration we choose a vertex to which the flow increase that is needed is maximal, then in
673 the next iteration a flow of at least ϵ will be added. Hence, a flow of at least ϵ will be added in every
674 iteration from this point and on. Since the sum of the flows in the edges of \mathcal{G} is bounded above by
675 $\sum_{e \in E} c(e)$, then this process cannot continue for infinitely many iterations.

676 A.10 A proof of the reduction in the proof of Theorem 16

677 Assume that ψ is satisfiable and let π be a satisfying assignment. Consider the regulator g in which
678 for each literal $z \in Z$ that holds in π we have $g(\langle s, z \rangle) = 2$ and for every other edge $e \in E^{s \rightarrow}$ we
679 have $g(e) = 0$. Thus, the total outgoing flow from s is $2n$. We show that every saturating preflow in
680 \mathcal{G} that respects g has no flow losses. First, a flow loss cannot occur in a vertex u_i because for every
681 i either x_i or \bar{x}_i has an incoming flow of 0. A flow loss also cannot occur in a vertex $z \in Z$ since
682 its outgoing capacity is at least 2. Finally, since π satisfies ψ , each clause C_i contains at least one
683 literal that holds in π and therefore there is at least one literal $z \in Z$ with $g(\langle s, z \rangle) > 0$ such that
684 $\langle z, C_i \rangle \notin E$. Thus, the incoming flow to C_i is at most $n - 1$ and hence a flow loss cannot occur in
685 the vertex C_i .

686 Now assume that ψ is not satisfiable and let g be a regulator with $\sum_{e \in E^{s \rightarrow}} g(e) \geq 2n$. We show
687 that there is a saturating preflow in \mathcal{G} that respects g and has flow losses. First, if there is some i such
688 that $g(\langle s, x_i \rangle) + g(\langle s, \bar{x}_i \rangle) > 2$ then a flow greater than 1 can be directed to u_i and thus a flow loss
689 occurs. Since $\sum_{e \in E^{s \rightarrow}} g(e) \geq 2n$ we assume that for every i we have $g(\langle s, x_i \rangle) + g(\langle s, \bar{x}_i \rangle) = 2$.
690 If for some i we have $g(\langle s, x_i \rangle) > 0$ and $g(\langle s, \bar{x}_i \rangle) > 0$ then a flow greater than 1 can be directed to
691 u_i and thus a flow loss occurs. Otherwise, for every i we have $g(\langle s, x_i \rangle) = 2$ or $g(\langle s, \bar{x}_i \rangle) = 2$ and
692 therefore g induces an assignment to the variables x_1, \dots, x_n . Since ψ is not satisfiable, there is a
693 clause C_i that does not contain literals $z \in Z$ with $g(\langle s, z \rangle) = 2$ and therefore a flow of n can be
694 directed to the vertex C_i , resulting in a flow loss.

695 A.11 A proof of Theorem 17

696 Given an integral regulator, deciding whether a flow loss is possible can be done in polynomial time
697 according to Theorem 15. Hence the NP upper bound.

698 For the lower bound, note that the reduction shown in the proof of Theorem 16 holds also for
699 integral regulators.