

## FLOW LOGIC

ORNA KUPFERMAN AND GAL VARDI

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

*e-mail address:* gal.vardi@mail.huji.ac.il

---

**ABSTRACT.** A flow network is a directed graph in which each edge has a capacity, bounding the amount of flow that can travel through it. Flow networks have attracted a lot of research in computer science. Indeed, many questions in numerous application areas can be reduced to questions about flow networks. This includes direct applications, namely a search for a maximal flow in networks, as well as less direct applications, like maximal matching or optimal scheduling. Many of these applications would benefit from a framework in which one can formally reason about properties of flow networks that go beyond their maximal flow.

We introduce *Flow Logics*: modal logics that treat flow functions as explicit first-order objects and enable the specification of rich properties of flow networks. The syntax of our logic BFL\* (Branching Flow Logic) is similar to the syntax of the temporal logic CTL\*, except that atomic assertions may be *flow propositions*, like  $> \gamma$  or  $\geq \gamma$ , for  $\gamma \in \mathbb{N}$ , which refer to the value of the flow in a vertex, and that first-order quantification can be applied both to paths and to flow functions. For example, the BFL\* formula  $\mathcal{E}((\geq 100) \wedge AG(low \rightarrow (\leq 20)))$  states that there is a legal flow function in which the flow is above 100 and in all paths, the amount of flow that travels through vertices with low security is at most 20.

We present an exhaustive study of the theoretical and practical aspects of BFL\*, as well as extensions and fragments of it. Our extensions include flow quantifications that range over non-integral flow functions or over maximal flow functions, path quantification that ranges over paths along which non-zero flow travels, past operators, and first-order quantification of flow values. We focus on the model-checking problem and show that it is PSPACE-complete, as it is for CTL\*. Handling of flow quantifiers, however, increases the complexity in terms of the network to  $P^{NP}$ , even for the LFL and BFL fragments, which are the flow-counterparts of LTL and CTL. We are still able to point to a useful fragment of BFL\* for which the model-checking problem can be solved in polynomial time. Finally, we introduce and study the query-checking problem for BFL\*, where under-specified BFL\* formulas are used for network exploration.

---

*Key words and phrases:* Flow Network, Temporal Logic.

This research has received funding from the European Research Council under the EU's 7-th Framework Programme (FP7/2007–2013) / ERC grant agreement no 278410.

## 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 [CLR90, GTT89]. The problem was first formulated and solved in the 1950's [FF56, FF62]. It has attracted much research on improved algorithms [EK72, Din70, GT88, Mad16] and applications [AMO93].

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, currents in an electrical circuit, packets in a communication network, and many more [AMO93]. 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 [AMO93] or elimination in partially completed tournaments [Sch66]. 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 [CLR90, AMO93]. Variants of the maximum-flow problem can accommodate further settings, like circulation problems, where there are no sink and target vertices, yet there is a lower bound on the flow that needs to be traversed along each edge [Tar85], networks with multiple source and target vertices, networks with costs for unit flows, networks with multiple commodities, and more [EIS76].

All the above applications reduce the problem at hand to the problem of finding a maximal flow in a network. Often, however, one would like to reason about properties of flow networks that go beyond their maximal flow. This is especially true when the vertices or edges of the network attain information to which the properties can refer. For example, the vertices of a network may be labeled by their security level, and we may want to check whether all legal flow functions are such that the flow in every low-security vertex is at most 20, or check whether there is a flow function in which more than 100 units of flow reach the target and still the flow in every low-security vertex is at most 20. As another example, assume that each vertex in the network is labeled by the service provider that owns it, and we want to find a maximal flow under the constraint that flow travels through vertices owned by at most two providers.

The challenge of reasoning about properties of systems has been extensively studied in the context of formal verification. In particular, in *temporal-logic model checking* [CE81, QS82], we check whether a system has a desired property by translating the system into a labeled state-transition graph, translating the property into a temporal-logic formula, and deciding whether the graph satisfies the formula. Model checking is one of the notable success stories of theoretical computer science, with exciting theoretical research that is being transformed into industrial applications [CGP99, CHVB18]. By viewing networks as labeled state-transition graphs, we can use existing model-checking algorithms and tools in order to reason about the structural properties of networks. We can check, for example, that every path from the source to the target eventually visits a *check-sum* vertex. Most interesting properties of flow networks, however, refer to flows and their values, and not just to the structural properties of the network. Traditional temporal logics do not support the specification and verification of such properties.

We introduce and study *Flow Logics*: modal logics that treat flow functions as explicit first-order objects and enable the specification of rich properties of flow networks. The syntax of our logic BFL<sup>\*</sup> (Branching Flow Logic) is similar to the syntax of the temporal logic CTL<sup>\*</sup>, except that atomic assertions are built from both atomic propositions and *flow propositions*, like  $> \gamma$  or  $\geq \gamma$ , for  $\gamma \in \mathbb{N}$ , which refer to the value of the flow in a vertex, and that first-order quantification can be applied both to paths and to flow functions. Thus, in addition to the path quantifiers  $A$  (“for all paths”) and  $E$  (“there exists a path”) that range over paths, states formulas may contain the flow quantifiers  $\mathcal{A}$  (“for all flow functions”) and  $\mathcal{E}$  (“there exists a flow function”). For example, the BFL<sup>\*</sup> formula  $\mathcal{E}((\geq 100) \wedge AG(low \rightarrow (\leq 20)))$  states the property discussed above, namely that there is a flow function in which the value of the flow is at least 100, and in all paths, the value of flow that travels through in vertices with low security is at most 20.

We study the theoretical aspects of BFL<sup>\*</sup> as well as extensions and fragments of it. We demonstrate their applications in reasoning about flow networks, and we examine the complexity of their model-checking problem. Below we briefly survey our results. (1) We show that while maximal flow can always be achieved by integral flows [FF56], in the richer setting of flow logic, restricting attention to integral flows may change the satisfaction value of formulas. Accordingly, our semantics for BFL<sup>\*</sup> considers two types of flow quantification: one over integral flows and another over *non-integral* ones. (2) We prove that bisimulation [Mil71] is not a suitable equivalence relation for flow logics, which are sensitive to unwinding. We relate this to the usefulness of *past operators* in flow logic, and we study additional aspects of the expressive power of BFL<sup>\*</sup>. (3) We consider extensions of BFL<sup>\*</sup> by path quantifiers that range over paths on which flow travels (rather than over all paths in the network), and by *first-order quantification on flow values*. (4) We study the model-checking complexity of BFL<sup>\*</sup>, its extensions, and some natural fragments. We show that algorithms for temporal-logic model-checking can be extended to handle flow logics, and that the complexity of the BFL<sup>\*</sup> model-checking problem is PSPACE-complete. We study also the *network complexity* of the problem, namely the complexity in terms of the network, assuming that the formula is fixed [LP85, KVV00], and point to a fragment of BFL<sup>\*</sup> for which the model-checking problem can be solved in polynomial time.

One of the concepts that has emerged in the context of formal verification is that of *model exploration*. The idea, as first noted by Chan in [Cha00], is that, in practice, model checking is often used for understanding the system rather than for verifying its correctness. Chan suggested to formalize model exploration by means of *query checking*. The input to the query-checking problem is a system  $S$  and a query  $\varphi$ , where a query is a temporal-logic formula in which some proposition is replaced by the place-holder “?”. A solution to the query is a propositional assertion that, when replaces the place-holder, results in a formula that is satisfied in  $S$ . For example, solutions to the query  $AG(? \rightarrow AX grant)$  are assertions  $\psi$  for which  $S \models AG(\psi \rightarrow AX grant)$ , thus assertions that trigger a grant in all successive possible futures. A query checker should return the strongest solutions to the query (strongest in the sense that they are not implied by other solutions). The work of Chan was followed by further work on query checking, studying its complexity, cases in which only a single strongest solution exists, the case of multiple (possibly related) place-holders, queries on semantic graphs, understanding business processes through query checking and more [BG01, GC02, CG03, SV03, CGG07, GRN12, RCM<sup>+</sup>14]. In the context of flow networks, the notion of model explorations is of special interest, and there is ongoing research on the development of automatic tools for reasoning about networks and

flow-forwarding strategies (c.f., [BDVV18, ZST<sup>+</sup>10]). We develop a theory of query checking for BFL<sup>\*</sup>. We distinguish between *propositional queries*, which, as in the case of CTL<sup>\*</sup>, seek strongest propositional assertions that may replace a ? that serves as a state formula, and *value queries*, which seek strongest values that may replace a ? that serves as a bound in a flow proposition. There, the strength of a solution depends on the tightness of the bound it imposes. We show that while in the propositional case there may be several partially ordered strongest solutions, in the value case the solutions are linearly ordered and thus there is at most one strongest solution. The existence of a single strongest solution has some nice practical applications. In particular, we show that while for propositional queries, the problem of finding all strongest solutions is very complex, for value queries its complexity coincides with that of BFL<sup>\*</sup> model checking.

**Related work.** There are three types of related works: (1) efforts to generalize the maximal-flow problem to richer settings, (2) extensions of temporal logics by new elements, in particular first-order quantification over new types, and (3) works on logical aspects of networks and their use in formal methods. Below we briefly survey them and their relation to our work.

As discussed early in this section, numerous extensions to the classical maximal-flow problems have been considered. In particular, some works that add constraints on the maximal flow, like capacities on vertices, or lower bounds on the flow along edges. Closest to flow logics are works that refer to labeled flow networks. For example, [GCSR13] considers flow networks in which edges are labeled, and the problem of finding a maximal flow with a minimum number of labels. Then, the maximal utilization problem of *capacitated automata* [KT14] amounts to finding maximal flow in a labeled flow network where flow is constrained to travel only along paths that belong to a given regular language. Our work suggests a formalism that embodies all these extensions, as well as a framework for formally reasoning about many more extensions and settings.

The competence of temporal-logic model checking initiated numerous extensions of temporal logics, aiming to capture richer settings. For example, *real-time* temporal logics include clocks with a real-time domain [AH94], *epistemic temporal logics* include knowledge operators [HM90], and *alternating temporal logics* include game modalities [AHK02]. Closest to our work is *strategy logic* [CHP07], where temporal logic is enriched by first-order quantification of strategies in a game. Beyond the theoretical interest in strategy logic, it was proven useful in synthesizing strategies in multi-agent systems and in the solution of rational synthesis [FKL10].

Finally, network verification is an increasingly important topic in the context of protocol verification [KVM12]. Tools that allow verifying properties of network protocols have been developed [WBLS09, LBG<sup>+</sup>15]. These tools support verification of network protocols in the design phase as well as runtime verification [KCZ<sup>+</sup>13]. Some of these tools use a query language called *Network Datalog* in order to specify network protocols [LCG<sup>+</sup>06]. Verification of *Software Defined Networks* has been studied widely, for example in [KZZ<sup>+</sup>13, KPC<sup>+</sup>12, CVP<sup>+</sup>12]. Verification of safety properties in networks with finite-state middleboxes was studied in [VAP<sup>+</sup>16]. Network protocols describe forwarding policies for packets, and are thus related to specific flow functions. However, the way traffic is transmitted in these protocols does not correspond to the way flow travels in a flow network. Thus, properties verified in this line of work are different from these we can reason about with flow logic.

## 2. THE FLOW LOGIC BFL\*

A *flow network* is  $N = \langle AP, V, E, c, \rho, s, T \rangle$ , where  $AP$  is a set of atomic propositions,  $V$  is a set of vertices,  $s \in V$  is a source vertex,  $T \subseteq V \setminus \{s\}$  is a set of target vertices,  $E \subseteq (V \setminus T) \times (V \setminus \{s\})$  is a set of directed edges,  $c : E \rightarrow \mathbb{N}$  is a capacity function, assigning to each edge an integral amount of flow that the edge can transfer, and  $\rho : V \rightarrow 2^{AP}$  assigns each vertex  $v \in V$  to the set of atomic propositions that are valid in  $v$ . Note that no edge enters the source vertex or leaves a target vertex. We assume that all vertices  $t \in T$  are reachable from  $s$  and that each vertex has at least one target vertex reachable from it. For a vertex  $u \in V$ , let  $E^u$  and  $E_u$  be the sets of incoming and outgoing edges to and from  $u$ , respectively. That is,  $E^u = (V \times \{u\}) \cap E$  and  $E_u = (\{u\} \times V) \cap E$ .

A *flow* is a function  $f : E \rightarrow \mathbb{N}$  that describes how flow is directed in  $N$ . The capacity of an edge bounds the flow in it, thus for every edge  $e \in E$ , we have  $f(e) \leq c(e)$ . All incoming flow must exit a vertex, thus for every vertex  $v \in V \setminus (\{s\} \cup T)$ , we have  $\sum_{e \in E^v} f(e) = \sum_{e \in E_v} f(e)$ . We extend  $f$  to vertices and use  $f(v)$  to denote the flow that travels through  $v$ . Thus, for  $v \in V \setminus (\{s\} \cup T)$ , we define  $f(v) = \sum_{e \in E^v} f(e) = \sum_{e \in E_v} f(e)$ , for the source vertex  $s$ , we define  $f(s) = \sum_{e \in E_s} f(e)$ , and for a target vertex  $t \in T$ , we define  $f(t) = \sum_{e \in E^t} f(e)$ . Note that the preservation of flow in the internal vertices guarantees that  $f(s) = \sum_{t \in T} f(t)$ , which is the amount of flow that travels from  $s$  to all the target vertices together. We say that a flow function  $f$  is *maximal* if for every flow function  $f'$ , we have  $f'(s) \leq f(s)$ . A maximal flow function can be found in polynomial time [FF62]. The maximal flow for  $N$  is then  $f(s)$  for some maximal flow function  $f$ .

The logic BFL\* is a *Branching Flow Logic* that can specify properties of networks and flows in them. As in CTL\*, there are two types of formulas in BFL\*: *state formulas*, which describe vertices in a network, and *path formulas*, which describe paths. In addition to the operators in CTL\*, the logic BFL\* has *flow propositions*, with which one can specify the flow in vertices, and *flow quantifiers*, with which one can quantify flow functions universally or existentially. When flow is not quantified, satisfaction is defined with respect to both a network and a flow function. Formally, given a set  $AP$  of atomic propositions, a BFL\* state formula is one of the following:

- (S1): An atomic proposition  $p \in AP$ .
- (S2): A flow proposition  $> \gamma$  or  $\geq \gamma$ , for an integer  $\gamma \in \mathbb{N}$ .
- (S3):  $\neg \varphi_1$  or  $\varphi_1 \vee \varphi_2$ , for BFL\* state formulas  $\varphi_1$  and  $\varphi_2$ .
- (S4):  $A\psi$ , for a BFL\* path formula  $\psi$  (and  $A$  is termed a *path quantifier*).
- (S5):  $\mathcal{A}\varphi$ , for a BFL\* state formula  $\varphi$  (and  $\mathcal{A}$  is termed a *flow quantifier*).

A BFL\* path formula is one of the following:

- (P1): A BFL\* state formula.
- (P2):  $\neg \psi_1$  or  $\psi_1 \vee \psi_2$ , for BFL\* path formulas  $\psi_1$  and  $\psi_2$ .
- (P3):  $X\psi_1$  or  $\psi_1 U \psi_2$ , for BFL\* path formulas  $\psi_1$  and  $\psi_2$ .

We say that a BFL\* formula  $\varphi$  is *closed* if all flow propositions appear in the scope of a flow quantifier. The logic BFL\* consists of the set of closed BFL\* state formulas. We refer to state formula of the form  $\mathcal{A}\varphi$  as a *flow state formula*.

The semantics of BFL\* is defined with respect to vertices in a flow network. Before we define the semantics, we need some more definitions and notations. Let  $N = \langle AP, V, E, c, \rho, s, T \rangle$ . For two vertices  $u$  and  $w$  in  $V$ , a finite sequence  $\pi = v_0, v_1, \dots, v_k \in V^*$  of vertices is a  $(u, w)$ -*path* in  $N$  if  $v_0 = u$ ,  $v_k = w$ , and  $\langle v_i, v_{i+1} \rangle \in E$  for all  $0 \leq i < k$ . If  $w \in T$ , then  $\pi$  is a *target  $u$ -path*.

State formulas are interpreted with respect to a vertex  $v$  in  $N$  and a flow function  $f : E \rightarrow \mathbb{N}$ . When the formula is closed, satisfaction is independent of the function  $f$  and we omit it. We use  $v, f \models \varphi$  to indicate that the vertex  $v$  satisfies the state formula  $\varphi$  when the flow function is  $f$ . The relation  $\models$  is defined inductively as follows.

- (S1): For an atomic proposition  $p \in AP$ , we have that  $v, f \models p$  iff  $p \in \rho(v)$ .
- (S2): For  $\gamma \in \mathbb{N}$ , we have  $v, f \models > \gamma$  iff  $f(v) > \gamma$  and  $v, f \models \geq \gamma$  iff  $f(v) \geq \gamma$ .
- (S3a):  $v, f \models \neg\varphi_1$  iff  $v, f \not\models \varphi_1$ .
- (S3b):  $v, f \models \varphi_1 \vee \varphi_2$  iff  $v, f \models \varphi_1$  or  $v, f \models \varphi_2$ .
- (S4):  $v, f \models A\psi$  iff for all target  $v$ -paths  $\pi$ , we have that  $\pi, f \models \psi$ .
- (S5):  $v, f \models \mathcal{A}\varphi$  iff for all flow functions  $f'$ , we have  $v, f' \models \varphi$ .

Path formulas are interpreted with respect to a finite path  $\pi$  in  $N$  and a flow function  $f : E \rightarrow \mathbb{N}$ . We use  $\pi, f \models \varphi$  to indicate that the path  $\pi$  satisfies the path-flow formula  $\varphi$  when the flow function is  $f$ . The relation  $\models$  is defined inductively as follows. Let  $\pi = v_0, v_1, \dots, v_k$ . For  $0 \leq i \leq k$ , we use  $\pi^i$  to denote the suffix of  $\pi$  that starts at  $v_i$ , thus  $\pi^i = v_i, v_{i+1}, \dots, v_k$ .

- (P1): For a state formula  $\varphi$ , we have that  $\pi, f \models \varphi$  iff  $v_0, f \models \varphi$ .
- (P2a):  $\pi, f \models \neg\psi$  iff  $\pi, f \not\models \psi$ .
- (P2b):  $\pi, f \models \psi_1 \vee \psi_2$  iff  $\pi, f \models \psi_1$  or  $\pi, f \models \psi_2$ .
- (P3a):  $\pi, f \models X\psi_1$  iff  $k > 0$  and  $\pi^1, f \models \psi_1$ .
- (P3b):  $\pi, f \models \psi_1 U \psi_2$  iff there is  $j \leq k$  such that  $\pi^j, f \models \psi_2$ , and for all  $0 \leq i < j$ , we have  $\pi^i, f \models \psi_1$ .

For a network  $N$  and a closed BFL\* formula  $\varphi$ , we say that  $N$  satisfies  $\varphi$ , denoted  $N \models \varphi$ , iff  $s \models \varphi$  (note that since  $\varphi$  is closed, we do not specify a flow function).

Additional Boolean connectives and modal operators are defined from  $\neg$ ,  $\vee$ ,  $X$ , and  $U$  in the usual manner; in particular,  $F\psi = \mathbf{true}U\psi$  and  $G\psi = \neg F\neg\psi$ . We also define dual and abbreviated flow propositions:  $(< \gamma) = \neg(\geq \gamma)$ ,  $(\leq \gamma) = \neg(> \gamma)$ , and  $(\gamma) = (= \gamma) = (\leq \gamma) \wedge (\geq \gamma)$ , a dual path quantifier:  $E\psi = \neg A\neg\psi$ , and a dual flow quantifier:  $\mathcal{E}\varphi = \neg \mathcal{A}\neg\varphi$ .

**Example 2.1.** Consider a network  $N$  in which target vertices are labeled by an atomic proposition *target*, and low-security vertices are labeled *red*. The BFL\* formula  $\mathcal{E}EF(\mathit{target} \wedge (= 20))$  states that there is a flow in which 20 units reach a target vertex, and the BFL\* formula  $\mathcal{A}((\geq 20) \rightarrow AX(\geq 4))$  states that in all flow functions in which the flow at the source is at least 20, all the successors must have flow of at least 4. Finally,  $\mathcal{E}((\geq 100) \wedge AG(\mathit{red} \rightarrow (\leq 20)))$  states that there is a flow of at least 100 in which the flow in every low-security vertex is at most 20, whereas  $\mathcal{A}((> 200) \rightarrow EF(\mathit{red} \wedge (> 20)))$  states that when the flow is above 200, then there must exist a low-security vertex in which the flow is above 20. As an example to a BFL\* formula with an alternating nesting of flow quantifiers, consider the formula  $\mathcal{E}AG(< 10 \rightarrow \mathcal{A} < 15)$ , stating that there is a flow such that wherever the flow is below 10, then in every flow it would be below 15.  $\square$

**Remark 2.2.** Note that while the semantics of CTL\* and LTL is defined with respect to infinite trees and paths, path quantification in BFL\* ranges over finite paths. We are still going to use techniques and results known for CTL\* and LTL in our study. Indeed, for upper bounds, the transition to finite computations only makes the setting simpler. Also, lower-bound proofs for CTL\* and LTL are based on an encoding of finite runs of Turing machines, and apply also to finite paths.

Specifying finite paths, we have a choice between a weak and a strong semantics for the  $X$  operator. In the weak semantics, the last vertex in a path satisfies  $X\psi$ , for all  $\psi$ .

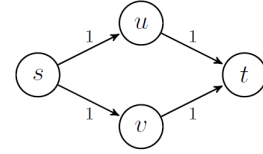
In particular, it is the only vertex that satisfies  $X\mathbf{false}$ . In the strong semantics, the last vertex does not satisfy  $X\psi$ , for all  $\psi$ . In particular, it does not satisfy  $X\mathbf{true}$ . We use the strong semantics.

### 3. PROPERTIES OF BFL<sup>\*</sup>

**3.1. Integral vs. non-integral flow functions.** Our semantics of BFL<sup>\*</sup> considers *integral flow functions*: vertices receive integral incoming flow and partition it to integral flows in the outgoing edges. Integral-flow functions arise naturally in settings in which the objects we transfer along the network cannot be partitioned into fractions, as is the case with cars, packets, and more. Sometimes, however, as in the case of liquids, flow can be partitioned arbitrarily. In the traditional maximum-flow problem, it is well known that the maximum flow can be achieved by integral flows [FF56]. We show that, interestingly, in the richer setting of flow logic, restricting attention to integral flows may change the satisfaction value of formulas.

**Proposition 3.1.** *Allowing the quantified flow functions in BFL<sup>\*</sup> to get values in  $\mathbb{R}$  changes its semantics.*

*Proof.* Consider the network on the right. The BFL<sup>\*</sup> formula  $\varphi = \mathcal{E}(1 \wedge AX(> 0))$  states that there is a flow function in which the flow that leaves the source is 1 and the flow of both its successors is strictly positive. It is easy to see that while no integral flow function satisfies the requirement in  $\varphi$ , a flow function in which 1 unit of flow in  $s$  is partitioned between  $u$  and  $v$  does satisfy it.  $\square$



Proposition 3.1 suggests that quantification of flow functions that allow non-integral flows may be of interest. In Section 4.3 we discuss such an extension.

**3.2. Sensitivity to unwinding.** For a network  $N = \langle AP, V, E, c, \rho, s, T \rangle$ , let  $N_t$  be the unwinding of  $N$  into a tree. Formally,  $N_t = \langle AP, V', E', \rho', s, T' \rangle$ , where  $V' \subseteq V^*$  is the smallest set such that  $s \in V'$ , and for all  $w \cdot v \in V'$  with  $w \in V^*$  and  $v \in V \setminus T$ , and all  $u \in V$  such that  $E(v, u)$ , we have that  $w \cdot v \cdot u \in V'$ , with  $\rho'(w \cdot v \cdot u) = \rho(u)$ . Also,  $\langle w \cdot v, w \cdot v \cdot u \rangle \in E'$ , with  $c'(\langle w \cdot v, w \cdot v \cdot u \rangle) = c(\langle v, u \rangle)$ . Finally,  $T' = V' \cap (V^* \cdot T)$ . Note that  $N_t$  may be infinite. Indeed, a cycle in  $N$  induces infinitely many vertices in  $N_t$ .

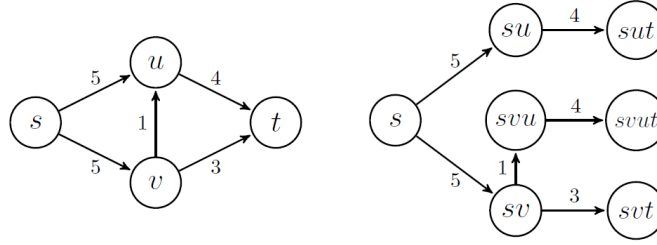
The temporal logic CTL<sup>\*</sup> is insensitive to unwinding. Indeed,  $N$  and  $N_t$  are bisimilar, and for every CTL<sup>\*</sup> formula  $\varphi$ , we have  $N \models \varphi$  iff  $N_t \models \varphi$  [Mil71]. We show that this is not the case for BFL<sup>\*</sup>.

**Proposition 3.2.** *The value of the maximal flow is sensitive to unwinding.*

*Proof.* Consider the network  $N$  appearing in Figure 1 (with the target  $T = \{t\}$ ), and its unwinding  $N_t$  which appears in its right. It is easy to see that the value of the maximal flow in  $N$  is 7, and the value of the maximal flow from  $s$  to  $T'$  in  $N_t$  is 8.  $\square$

**Corollary 3.3.** *The logic BFL<sup>\*</sup> is sensitive to unwinding.*

The sensitivity of BFL<sup>\*</sup> to unwinding suggests that extending BFL<sup>\*</sup> with past operators can increase its expressive power. In Section 4.4, we discuss such an extension.

Figure 1: The flow network  $N$  and its unwinding  $N_t$ 

#### 4. EXTENSIONS AND FRAGMENTS OF BFL\*

In this section we discuss useful extensions and variants of BFL\*, as well as fragments of it. As we shall show in the sequel, while the extensions come with no computational price, their model checking requires additional techniques.

**4.1. Positive path quantification.** Consider a network  $N = \langle AP, V, E, c, \rho, s, T \rangle$  and a flow function  $f : E \rightarrow \mathbb{N}$ . We say that a path  $\pi = v_0, v_1, \dots, v_k$  is positive if the flow along all the edges in  $\pi$  is positive. Formally,  $f(v_i, v_{i+1}) > 0$ , for all  $0 \leq i < k$ . Note that it may be that  $f(v_i) > 0$  for all  $0 \leq i < k$  and still  $\pi$  is not positive. It is sometimes desirable to restrict the range of path quantification to paths along which flow travels. This is the task of the *positive path quantifier*  $A^+$ , with the following semantics.

- $v, f \models A^+\psi$  iff for all positive target  $v$ -paths  $\pi$ , we have that  $\pi, f \models \psi$ .

Dually,  $E^+\psi = \neg A^+\neg\psi$ .

**Example 4.1.** Let  $W$  be a set of workers and  $J$  be a set of jobs. Each worker  $w \in W$  can be assigned to perform jobs from a subset  $J_w \subseteq J$ . It is required to perform all jobs by assigning exactly one worker to each job and at most two jobs to each worker. This problem can be solved using the flow network appearing in Figure 2. The vertices on the left column correspond to the workers, these on the right column correspond to the jobs, and there is an edge between a worker  $w$  and a job  $j$  iff  $j \in J_w$ . The capacity of all edges is 1, except for these from the source  $s$  to the worker vertices, which have capacity 2. It is easy to see that a flow of  $k$  units in the network described in Figure 2 corresponds to a legal assignment in which  $k$  jobs are performed. Assume now that some jobs should be processed in Location  $a$  and the others in Location  $b$ . A worker can process jobs only in a single location,  $a$  or  $b$ . By labeling the job-vertices by their location, the existence of a legal assignment in which  $k$  jobs are processed can be expressed by the BFL\* formula

$$\mathcal{E}(k \wedge AX(A^+Xa \vee A^+Xb)),$$

which uses positive path quantification. □

**4.2. Maximal flow quantification.** It is sometimes desirable to restrict the range of flow quantification to maximal flow functions. This is the task of the *maximal-flow quantifier*  $\mathcal{A}^{max}$ , with the following semantics.

- $v, f \models \mathcal{A}^{max}\varphi$  iff for all maximal-flow functions  $f'$ , we have that  $v, f' \models \varphi$ .



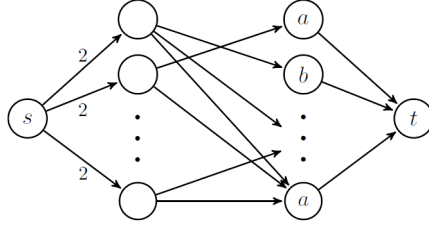


Figure 2: Assigning workers to jobs.

Dually,  $\mathcal{E}^{max}\varphi = \neg\mathcal{A}^{max}\neg\varphi$ .

In a similar manner, it is sometimes helpful to relate to the maximal flow in the network. The max-flow constant  $\gamma_{max} \in \mathbb{N}$  maintains the value of the maximal flow from  $s$  to  $T$ . We also allow arithmetic operations on  $\gamma_{max}$ .

**Example 4.2.** Recall the job-assignment problem from Example 4.1. The BFL<sup>\*</sup> formula

$$\mathcal{E}^{max}(AX(A^+Xa \vee A^+Xb))$$

states that the requirements about the locations do not reduce the number of jobs assigned without this requirement. Then, the formula

$$\mathcal{E}((\geq \gamma_{max} - 4) \wedge AX(A^+Xa \vee A^+Xb))$$

states that the requirements about the locations may reduce the number of jobs performed by at most 4.  $\square$

**4.3. Non-integral flow quantification.** As discussed in Section 3.1, letting flow quantification range over non-integral flow functions may change the satisfaction value of a BFL<sup>\*</sup> formula.

Such a quantification is sometimes desirable, and we extend BFL<sup>\*</sup> with a *non-integral flow quantifier*  $\mathcal{A}^{\mathbb{R}}$ , with the following semantics.

- $v, f \models \mathcal{A}^{\mathbb{R}}\varphi$  iff for all flow functions  $f' : E \rightarrow \mathbb{R}^+$ , we have that  $v, f' \models \varphi$ .

Dually,  $\mathcal{E}^{\mathbb{R}}\varphi = \neg\mathcal{A}^{\mathbb{R}}\neg\varphi$ .

**4.4. Past operators.** As discussed in Section 3.2, while temporal logics are insensitive to unwinding, this is not the case for BFL<sup>\*</sup>. Intuitively, this follows from the fact that the flow in a vertex depends on the flow it gets from all its predecessors. This dependency suggests that an explicit reference to predecessors is useful, and motivates the extension of BFL<sup>\*</sup> by past operators.

Adding past to a branching logic, one can choose between a linear-past semantics — one in which past is unique (technically, the semantics is with respect to an unwinding of the network), and a branching-past semantics — one in which all the possible behaviors that lead to present are taken into an account (technically, the semantics is dual to that of future operators, and is defined with respect to the network) [KPV12]. For flow logics, the branching-past approach is the suitable one, and is defined as follows.

For a path  $\pi = v_0, v_1, \dots, v_k \in V^*$ , a vertex  $v \in V$ , and index  $0 \leq i \leq k$ , we say that  $\pi$  is a source-target  $(v, i)$ -path if  $v_0 = s$ ,  $v_i = v$ , and  $v_k \in T$ . We add to BFL<sup>\*</sup> two past modal

operators,  $Y$  (“Yesterday”) and  $S$  (“Since”), and adjust the semantics as follows. Defining the semantics of logics that refer to the past, the semantics of path formulas is defined with respect to a path and an index in it. We use  $\pi, i, f \models \psi$  to indicate that the path  $\pi$  satisfies the path formula  $\psi$  from position  $i$  when the flow function is  $f$ . For state formulas, we adjust the semantics as follows.

**(S4):**  $v, f \models A\psi$  iff for all  $i$  and for all source-target  $(v, i)$ -paths  $\pi$ , we have that  $\pi, i, f \models \psi$ . Then, for path formulas, we have the following (the adjustment to refer to the index  $i$  in all other modalities is similar).

- $\pi, i, f \models Y\psi_1$  iff  $i > 0$  and  $\pi, i - 1, f \models \psi_1$ .
- $\pi, i, f \models \psi_1 S\psi_2$  iff there is  $0 \leq j < i$  such that  $\pi, j, f \models \psi_2$ , and for all  $j + 1 \leq l \leq i$ , we have  $\pi, l, f \models \psi_1$ .

**Example 4.3.** Recall the job-assignment problem from Example 4.1. Assume that some of the workers have cars, which is indicated by an atomic proposition  $car$  that may label worker vertices. Also, say that a job is a *transit job* if it can be assigned only to workers with cars. Assume that we want to apply the restriction about the single location only to non-transit jobs. Thus, if all the predecessors of a job-vertex are labeled by  $car$ , then this job can be served in either locations. Using past operators, we can specify this property by the formula

$$\mathcal{E}(k \wedge AX(A^+X(a \vee AY car) \vee A^+X(b \vee AY car))).$$

□

**4.5. First-Order quantification on flow values.** The flow propositions in  $BFL^*$  include constants. This makes it impossible to relate the flow in different vertices other than specifying all possible constants that satisfy the relation. In *BFL\* with quantified flow values* we add flow variables  $X = \{x_1, \dots, x_n\}$  that can be quantified universally or existentially and specify such relations conveniently. We also allow the logic to apply arithmetic operations on the values of variables in  $X$ .

For a set of arithmetic operators  $O$  (that is,  $O$  may include  $+$ ,  $*$ , etc.), let  $BFL^*(O)$  be  $BFL^*$  in which Rule S2 is extended to allow expressions with variables in  $X$  constructed by operators in  $O$ , and we also allow quantification on the variables in  $X$ . Formally, we have the following:

- (S2)** A flow proposition  $> g(x_1, \dots, x_k)$  or  $\geq g(x_1, \dots, x_k)$ , where  $x_1, \dots, x_k$  are variables in  $X$  and  $g$  is an expression obtained from  $x_1, \dots, x_k$  by applying operators in  $O$ , possibly using constants in  $\mathbb{N}$ . We assume that  $g : \mathbb{N}^k \rightarrow \mathbb{N}$ . That is,  $g$  leaves us in the domain  $\mathbb{N}$ .

- (S6)**  $\forall x\varphi$ , for  $x \in X$  and a  $BFL^*(O)$  formula  $\varphi$  in which  $x$  is free.

For a  $BFL^*(O)$  formula  $\varphi$  in which  $x$  is a free variable, and a constant  $\gamma \in \mathbb{N}$ , let  $\varphi[x \leftarrow \gamma]$  be the formula obtained by assigning  $\gamma$  to  $x$  and replacing expressions by their evaluation. Then,  $v, f \models \forall x\varphi$  iff for all  $\gamma \in \mathbb{N}$ , we have that  $v, f \models \varphi[x \leftarrow \gamma]$ .

**Example 4.4.** The logic  $BFL^*(\emptyset)$  includes the formula  $\mathcal{E}AG\forall x((split \wedge x \wedge > 0) \rightarrow EX(> 0 \wedge < x))$ , stating that there is a flow in which all vertices that are labeled *split* and with a positive flow  $x$  have a successor in which the flow is positive but strictly smaller than  $x$ . Then,  $BFL^*(div)$  includes the formula  $\mathcal{E}AG\forall x(x \rightarrow EX(\geq x \text{ div } 2))$ , stating that there is a flow in which all vertices have a successor that has at least half of their flow.

Finally,  $\text{BFL}^*(+)$  includes the formula  $\exists x \exists y \mathcal{E}^{\max} AG(\neg(\text{source} \vee \text{target}) \rightarrow x \vee y \vee (x+y))$ , stating that there are values  $x$  and  $y$ , such that it is possible to attain the maximal flow by assigning to all vertices, except maybe source and target vertices, values in  $\{x, y, x+y\}$ .  $\square$

**4.6. Fragments of  $\text{BFL}^*$ .** For the temporal logic  $\text{CTL}^*$ , researchers have studied several fragments, most notably  $\text{LTL}$  and  $\text{CTL}$ . In this section we define interesting fragments of  $\text{BFL}^*$ .

**Flow- $\text{CTL}^*$  and Flow- $\text{LTL}$ .** The logics  $\text{Flow-CTL}^*$  and  $\text{Flow-LTL}$  are extensions of  $\text{CTL}^*$  and  $\text{LTL}$  in which atomic state formulas may be, in addition to  $\text{APs}$ , also the flow propositions  $> \gamma$  or  $\geq \gamma$ , for an integer  $\gamma \in \mathbb{N}$ . Thus, no quantification on flow is allowed, but atomic formulas may refer to flow. The semantics of  $\text{Flow-CTL}^*$  is defined with respect to a network and a flow function, and that of  $\text{Flow-LTL}$  is defined with respect to a path in a network and a flow function.

**Linear Flow Logic.** The logic  $\text{LFL}$  is the fragment of  $\text{BFL}^*$  in which only one external universal path quantification is allowed. Thus, an  $\text{LFL}$  formula is a  $\text{BFL}^*$  formula of the form  $A\psi$ , where  $\psi$  is generated without Rule S4.

Note that while the temporal logic  $\text{LTL}$  is a “pure linear” logic, in the sense that satisfaction of an  $\text{LTL}$  formula in a computation of a system is independent of the structure of the system, the semantics of  $\text{LFL}$  mixes linear and branching semantics. Indeed, while all the paths in  $N$  have to satisfy  $\psi$ , the context of the system is important. To see this, consider the  $\text{LFL}$  formula  $\varphi = A\mathcal{A}(\geq 10) \rightarrow X(\geq 4)$ . The formula states that in all paths, all flow functions in which the flow at the first vertex in the path is at least 10, are such that the flow at the second vertex in the path is at least 4. In order to evaluate the path formula  $\mathcal{A}(\geq 10) \rightarrow X(\geq 4)$  in a path  $\pi$  of a network  $N$  we need to know the capacity of all the edges from the source of  $N$ , and not only the capacity of the first edge in  $\pi$ . For example,  $\varphi$  is satisfied in networks in which the source  $s$  has two successors, each connected to  $s$  by an edge with capacity 4, 5, or 6. Consider now the  $\text{LFL}$  formula  $\varphi' = A\mathcal{E}(10 \wedge X(\geq 4))$ . Note that  $\varphi'$  is not equal to the  $\text{BFL}^*$  formula  $\theta = \mathcal{E}(10 \wedge AX(\geq 4))$ . Indeed, in the latter, the same flow function should satisfy the path formula  $X(\geq 4)$  in all paths.

**No nesting of flow quantifiers.** The logic  $\text{BFL}_1^*$  contains formulas that are Boolean combinations of formulas of the form  $\mathcal{E}\varphi$  and  $\mathcal{A}\varphi$ , for a  $\text{Flow-CTL}^*$  formula  $\varphi$ . Of special interest are the following fragments of  $\text{BFL}_1^*$ :

- $\exists\text{BFL}_1^*$  and  $\forall\text{BFL}_1^*$ , where formulas are of the form  $\mathcal{E}\varphi$  and  $\mathcal{A}\varphi$ , respectively, for a  $\text{Flow-CTL}^*$  formula  $\varphi$ .
- $\exists\text{LFL}_1$  and  $\forall\text{LFL}_1$ , where formulas are of the form  $\mathcal{E}A\psi$  and  $\mathcal{A}A\psi$ , respectively, for a  $\text{Flow-LTL}$  formula  $\psi$ , and  $\text{LFL}_1$ , where a formula is a Boolean combination of  $\exists\text{LFL}_1$  and  $\forall\text{LFL}_1$  formulas.

**Conjunctive- $\text{BFL}^*$ .** The fragment  $\text{Conjunctive-BFL}^*$  ( $\text{CBFL}^*$ , for short) contains  $\text{BFL}^*$  formulas whose flow state sub-formulas restrict the quantified flow in a conjunctive way. That is, when we “prune” a  $\text{CBFL}^*$  formula into requirements on the network, atomic flow propositions are only conjunctively related. This would have a computational significance in solving the model-checking problem.

Consider an  $\exists\text{BFL}_1^*$  formula  $\varphi = \mathcal{E}\theta$ . We say that an operator  $g \in \{\vee, \wedge, E, A, X, F, G, U\}$  has a *positive polarity* in  $\varphi$  if all the occurrences of  $g$  in  $\theta$  are in a scope of an even number of negations. Dually,  $g$  has a *negative polarity* in  $\varphi$  if all its occurrences in  $\theta$  are in a scope of an odd number of negations.

In order to define  $\text{CBFL}^*$ , let us first define the fragments  $\exists\text{CBFL}_1^*$  of  $\exists\text{BFL}_1^*$  and  $\forall\text{CBFL}_1^*$  of  $\forall\text{BFL}_1^*$ , which constitute the inner level of  $\text{CBFL}^*$ . The logic  $\exists\text{CBFL}_1^*$  is a fragment of  $\exists\text{BFL}_1^*$  in which the only operators with a positive polarity are  $\wedge, A, X,$  and  $G$ , and the only operators with a negative polarity are  $\vee, E, X,$  and  $F$ . Note that  $U$  is not allowed, as its semantics involves both conjunctions and disjunctions. Note that by pushing negations inside, we make all operators of a positive polarity; that is, we are left only with  $\wedge, A, X,$  and  $G$ . Then, since all requirements are universal and conjunctively related, we can push conjunctions outside so that path formulas do not have internal conjunctions — for example, transform  $AX(\xi_1 \wedge \xi_2)$  into  $AX\xi_1 \wedge AX\xi_2$ , and can get rid of universal path quantification that is nested inside another universal path quantification — for example, transform  $AXAX\xi_1$  into  $AXX\xi$ . Finally, since we use the strong semantics to  $X$ , we can replace formulas that have  $X$  nested inside  $G$  by **false**. The logic  $\forall\text{CBFL}_1^*$  is the dual fragment of  $\forall\text{BFL}_1^*$ . In other words,  $\mathcal{A}\theta$  is in  $\forall\text{CBFL}_1^*$  iff  $\mathcal{E}\neg\theta$  is in  $\exists\text{CBFL}_1^*$ .

Intuitively, the logic  $\text{CBFL}^*$  is obtained by going up a hierarchy in which formulas of lower levels serve as atomic propositions in higher levels. We now define the syntax of  $\text{CBFL}^*$  formally. For simplicity, we define it in a normal form, obtained by applying the rules described above. A  $\text{CBFL}_0^*$  formula is a Boolean assertion over  $AP$ . For  $i \geq 0$ , a  $\text{CBFL}_{i+1}^*$  formula is a Boolean assertion over  $\text{CBFL}_i^*$  formulas and formulas of the form  $\mathcal{E}(A\psi_1 \wedge \cdots \wedge A\psi_n)$ , where  $\psi_j$  is of the form  $X^{k_j}\xi_j$  or  $X^{k_j}G\xi_j$ , where  $k_j \geq 0$  and  $\xi_j$  is a  $\text{CBFL}_i^*$  formula or a flow proposition (that is,  $> \gamma, < \gamma, \geq \gamma,$  or  $\leq \gamma$ , for an integer  $\gamma \in \mathbb{N}$ ). Then, a  $\text{CBFL}^*$  formula is a  $\text{CBFL}_i^*$  formula for some  $i \geq 0$ . Note that both  $\exists\text{CBFL}_1^*$  and  $\forall\text{CBFL}_1^*$  are contained in  $\text{CBFL}_1^*$ .

**Example 4.5.** Recall the job-assignment problem from Example 4.1. The  $\text{CBFL}_1^*$  formula  $\mathcal{A}(< 10 \rightarrow EX \leq 0) \wedge \mathcal{E}(15 \wedge AX \geq 1)$  states that if less than 10 jobs are processed, then at least one worker is unemployed, but it is possible to process 15 jobs and let every worker process at least one job.

**BFL.** The logic BFL is the fragment of  $\text{BFL}^*$  in which every modal operator ( $X, U$ ) is immediately preceded by a path quantifier. That is, it is the flow counterpart of CTL. As we are going to show, while in temporal logic, the transition from  $\text{CTL}^*$  to CTL significantly reduces model-checking complexity, this is not the case for BFL.

**Remark 4.6.** As demonstrated in the examples, the extensions of  $\text{BFL}^*$  simplifies the specification of natural properties. We do not, however, study the expressive power of the extensions and fragments. For some, results about expressive power follow from known results about  $\text{CTL}^*$ . For example, as proven in [KPV12], adding past to  $\text{CTL}^*$  with a branching-past semantics strictly increases its expressive power. The same arguments can be used in order to show that  $\text{BFL}^*$  with past operators is strictly more expressive than  $\text{BFL}^*$ .<sup>1</sup> Likewise, it is easy to show that  $\text{BFL}^*$  is strictly more expressive than BFL. On the other hand, it is

<sup>1</sup>We note that the result from [KPV12] does not immediately imply the analogous result for  $\text{BFL}^*$ , as it might be the case that the flow operators somehow cover up for the extra power of the branching-past temporal operator. In particular, the proof in [KPV12] is based on the fact that only  $\text{CTL}^*$  with past operators is sensitive to unwinding and, as we prove in Section 3.2,  $\text{BFL}^*$  is sensitive to unwinding. Still,

not hard to see that quantification over the maximum flow can be expressed by first-order quantification on flow values. Indeed,  $\mathcal{E}^{max}\psi$  is equivalent to  $\exists x\mathcal{E}((= x) \wedge \psi \wedge \mathcal{A}(\leq x))$ . Likewise, it may well be that every BFL<sup>\*</sup> formula with non-integral flow quantifiers has an equivalent one without such quantifiers, or that positive path quantification can be expressed by first-order quantification on flow values. We leave the study of the expressive power and relative succinctness of the various logics to future research.

## 5. MODEL CHECKING

In this section we study the model-checking problem for BFL<sup>\*</sup>. The problem is to decide, given a flow network  $N$  and a BFL<sup>\*</sup> formula  $\varphi$ , whether  $N \models \varphi$ .

**Theorem 5.1.** *BFL<sup>\*</sup> model checking is PSPACE-complete.*

*Proof.* Consider a network  $N$  and a BFL<sup>\*</sup> formula  $\varphi$ . The idea behind our model-checking procedure is similar to the one that recursively employs LTL model checking in the process of CTL<sup>\*</sup> model checking [EL87]. Here, however, the setting is more complicated. Indeed, the path formulas in BFL<sup>\*</sup> are not “purely linear”, as the flow quantification in them refers to flow in the (branching) network. In addition, while the search for witness paths is restricted to paths in the network, which can be guessed on-the-fly in the case of LTL, here we also search for witness flow functions, which have to be guessed in a global manner.

Let  $\{\varphi_1, \dots, \varphi_k\}$  be the set of flow state formulas in  $\varphi$ . Assume that  $\varphi_1, \dots, \varphi_k$  are ordered so that for all  $1 \leq i \leq k$ , all the subformulas of  $\varphi_i$  have indices in  $\{1, \dots, i\}$ . Our model-checking procedure labels  $N$  by new atomic propositions  $q_1, \dots, q_k$  so that for all vertices  $v$  and  $1 \leq i \leq k$ , we have that  $v \models q_i$  iff  $v \models \varphi_i$  (note that since  $\varphi_i$  is closed, satisfaction is independent of a flow function).

Starting with  $i = 1$ , we model check  $\varphi_i$ , label  $N$  with  $q_i$ , and replace the subformula  $\varphi_i$  in  $\varphi$  by  $q_i$ . Accordingly, when we handle  $\varphi_i$ , it is an  $\exists\text{BFL}_1^*$  or a  $\forall\text{BFL}_1^*$  formula. That is, it is of the form  $\mathcal{E}\xi$  or  $\mathcal{A}\xi$ , for a Flow-CTL<sup>\*</sup> formula  $\xi$ . Assume that  $\varphi_i = \mathcal{E}\xi$ . We guess a flow function  $f : E \rightarrow \mathbb{N}$ , and perform CTL<sup>\*</sup> model-checking on  $\xi$ , evaluating the flow propositions in  $\xi$  according to  $f$ . Since guessing  $f$  requires polynomial space, and CTL<sup>\*</sup> model checking is in PSPACE, so is handling of  $\varphi_i$  and of all the subformulas.

Hardness in PSPACE follows from the hardness of CTL<sup>\*</sup> model checking [SC85].  $\square$

Since BFL<sup>\*</sup> contains CTL<sup>\*</sup>, the lower bound in Theorem 5.1 is immediate. One may wonder whether reasoning about flow networks without atomic propositions, namely when we specify properties of flow only, is simpler. Theorem 5.2 below shows that this is not the case. Essentially, the proof follows from our ability to encode assignments to atomic propositions by values of flow.

**Theorem 5.2.** *BFL<sup>\*</sup> model checking is PSPACE-complete already for  $\forall\text{LFL}_1$  formulas without atomic propositions.*

*Proof.* We describe a reduction from LTL satisfiability, which is PSPACE-hard [SC85]. As discussed in Remark 2.2, PSPACE-hardness applies already to LTL over finite computations.

Consider an LTL formula  $\psi$ . For simplicity, we assume that  $\psi$  is over a single atomic proposition  $p$ . Indeed, LTL satisfiability is PSPACE-hard already for this fragment. We

---

since the specific formula used in [KPV12] in order to prove the sensitivity of CTL<sup>\*</sup> with past to unwinding does not refer to flow, it is easy to see that it has no equivalent BFL<sup>\*</sup> formula.

construct a flow network  $N$  and an  $\forall\text{LFL}_1$  formula  $\varphi$  such that  $N$  does not satisfy  $\varphi$  iff some computation in  $\{p, \neg p\}^*$  satisfies  $\psi$ . The network  $N$  appears in Figure 3. Note that when  $s$  has a flow of 1, then exactly one of  $v_p$  and  $v_{\neg p}$  has a flow of 1. Let  $\psi'$  be the Flow-LTL formula obtained from  $\psi$  by replacing all occurrences of  $p$  by  $\geq 1$ . We define  $\varphi = \mathcal{A}\mathcal{A}(1 \rightarrow X\neg\psi')$ . Thus,  $\neg\varphi = \mathcal{E}\mathcal{E}(1 \wedge X\psi')$ .

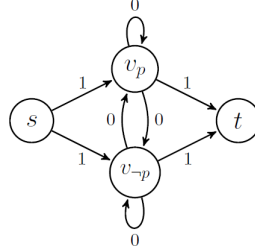


Figure 3: The flow network  $N$ .

We prove that  $N$  satisfies  $\mathcal{E}\mathcal{E}(1 \wedge X\psi')$  iff some computation in  $\{p, \neg p\}^*$  satisfies  $\psi$ . Assume first that some computation  $\pi$  in  $\{p, \neg p\}^*$  satisfies  $\psi$ . Consider a flow  $f$  in  $N$  such that  $f(v_p) = 1$  and  $f(v_{\neg p}) = 0$ , and consider the path in  $N$  that encodes  $\pi$  (starting from the second vertex in the path). For these flow and path, the formula  $X\psi'$  holds, and therefore  $N \models \mathcal{E}\mathcal{E}(1 \wedge X\psi')$ . For the other direction, assume that  $N \models \mathcal{E}\mathcal{E}(1 \wedge X\psi')$ . Let  $f$  and  $\pi$  be the flow function and the path in  $N$  that witness the satisfaction of  $X\psi'$ . The structure of  $N$  guarantees that  $f(v_p) + f(v_{\neg p}) = 1$ . Hence,  $\pi$  with  $f$  encodes a computation in  $\{p, \neg p\}^*$  that satisfies  $\psi$ , and we are done. Note that we can avoid using edges with capacities 0 by changing  $\neg\varphi$  to  $\mathcal{E}\mathcal{E}(1 \wedge X0 \wedge XX\psi')$ .  $\square$

We also note that when the given formula is in BFL, we cannot avoid the need to guess a flow function, yet once the flow function is guessed, we can verify it in polynomial time. Accordingly, the model-checking problem for BFL is in  $\text{P}^{\text{NP}}$ . We discuss this point further below.

In practice, a network is typically much bigger than its specification, and its size is the computational bottleneck. In temporal-logic model checking, researchers have analyzed the *system complexity* of model-checking algorithms, namely the complexity in terms of the system, assuming the specification is of a fixed length. There, the system complexity of LTL and CTL\* is NLOGSPACE-complete [LP85, KVW00]. We prove that, unfortunately, this is not the case of BFL\*. That is, we prove that while the *network complexity* of the model-checking problem, namely the complexity in terms of the network, does not reach PSPACE, it seems to require polynomially many calls to an NP oracle. Essentially, each evaluation of a flow quantifier requires such a call. Formally, we have the following.

**Theorem 5.3.** *The network complexity of BFL\* is in  $\Delta_2^{\text{P}}$  (i.e., in  $\text{P}^{\text{NP}}$ ).*

*Proof.* Fixing the length of the formula in the algorithm described in the proof of Theorem 5.1, we get that  $k$  is fixed, and so is the length of each subformula  $\varphi_i$ . Thus, evaluation of  $\varphi_i$  involves a guess of a flow and then model checking of a fixed size Flow-CTL\* formula, which can be done in time polynomial in the size of the network. Hence, the algorithm from Theorem 5.1 combined with an NP oracle gives the required network complexity.  $\square$

While finding the exact network complexity of model checking  $BFL^*$  and its fragments is interesting from a complexity-theoretical point of view, it does not contribute much to our story. Here, we prove NP and co-NP hardness holds already for very restricted fragments. As good news, in Section 7 we point that for the conjunctive fragment, model checking can be performed in polynomial time.<sup>2</sup>

We first show a simple lemma that is needed for the network-complexity hardness results.

**Lemma 5.4.** *A CNF formula  $\psi$  can be translated to a CNF formula  $\psi'$  in polynomial time, such that  $\psi$  is satisfiable iff  $\psi'$  is satisfiable, and all literals in  $\psi'$  appear exactly the same number of times in  $\psi'$ .*

*Proof.* We consider CNF formulas in which a literal is not allowed to appear more than once in each clause. Let  $|y|_\psi$  denote the number of occurrences of the literal  $y$  in  $\psi$ . Let  $x_1$  and  $x_2$  be two variables in  $\psi$ . We add some ‘true’ clauses to  $\psi$  to obtain  $\psi'$ . First we add clauses of the form  $(x_1 \vee x_2 \vee \neg x_2)$  for  $|\neg x_1|_\psi - |x_1|_\psi$  times if  $|\neg x_1|_\psi > |x_1|_\psi$ , else we add clauses of the form  $(\neg x_1 \vee x_2 \vee \neg x_2)$  for  $|x_1|_\psi - |\neg x_1|_\psi$  times. Thus we obtain a formula  $\psi_1$  with the same number of occurrences for the literals  $x_1$  and  $\neg x_1$ .

Now for every variable  $x_i \neq x_1$ , we add clauses  $(x_i \vee x_1 \vee \neg x_1)$  if the number of occurrences of  $\neg x_i$  in  $\psi_1$  is more than that of  $x_i$ , otherwise we add clauses  $(\neg x_i \vee x_1 \vee \neg x_1)$ , possibly multiple times, until the number of occurrences of  $x_i$  and the number of occurrences of  $\neg x_i$  become the same. Now for each variable  $x$ , both the literals  $x$  and  $\neg x$  appear the same number of times.

As a final step, let  $x_j$  be a variable such that the number of occurrences of the literals  $x_j$  and  $\neg x_j$  is the maximum over all the variables. For every variable  $x_i \neq x_j$ , we add clauses of the form  $(x_i \vee \neg x_i)$  until the number of literals for  $x_i$  and  $x_j$  are the same. We call the resultant formula  $\psi'$ .

Note that all the above steps can be done in polynomial time and the size of  $\psi'$  is polynomial in the size of  $\psi$ .  $\square$

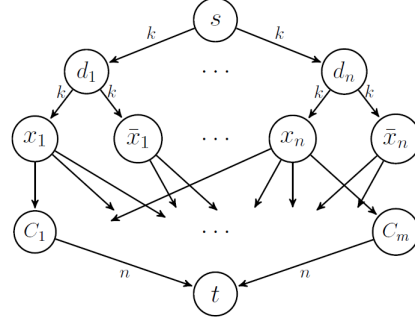
**Theorem 5.5.** *The network complexity of  $\exists BFL_1^*$  and  $\forall BFL_1^*$  is NP-complete and co-NP-complete, respectively. Hardness applies already to  $\exists LFL_1$  and  $\forall LFL_1$  formulas without atomic propositions, and to BFL.*

*Proof.* For the upper bound, it is easy to see that one step in the algorithm described in the proof of Theorem 5.1 (that is, evaluating  $\varphi_i$  once all its flow state subformulas have been evaluated), when applied to  $\varphi_i$  of a fixed length is in NP for  $\varphi_i$  of the form  $\mathcal{E}\xi$  and in co-NP for  $\varphi_i$  of the form  $\mathcal{A}\xi$ .

For the lower bound, we prove NP-hardness for  $\exists LFL_1$ . Co-NP-hardness for  $\forall LFL_1$  follows by dualization. We describe a reduction from CNF-SAT. Let  $\theta = C_1 \wedge \dots \wedge C_m$  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 exactly in  $k$  clauses in  $\theta$ . Indeed, by Lemma 5.4, every CNF formula can be converted to such a formula in polynomial time and with a polynomial blowup.

<sup>2</sup>A possible tightening of our analysis is via the complexity class BH, which is based on a Boolean hierarchy over NP. Essentially, it is the smallest class that contains NP and is closed under union, intersection, and complement. The levels of the hierarchy start with  $BH_1 = NP$ , and each level adds internal intersections as well as intersection with a co-NP (even levels) or an NP (odd levels) language [Wec85]. BH is contained in  $\Delta_2^P$ . It is not hard to prove that the network complexity of the fragment of  $BFL_1^*$  that contains at most  $k$  flow quantifiers is in  $BH_{k+1} \cap co-BH_{k+1}$ . Indeed, the latter contain problems that are decidable in polynomial time with  $k$  parallel queries to an NP oracle [Bei91]. A  $BH_k$  lower bound can also be shown.

We construct a flow network  $N$  and an  $\exists\text{LFL}_1$  formula  $\mathcal{E}A\psi$  such that  $\theta$  is satisfiable iff  $N \models \mathcal{E}A\psi$ . The network  $N$  is constructed as demonstrated on the right. 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  $N$  contains an edge  $\langle z, C_i \rangle$  iff the clause  $C_i$  contains the literal  $z$ . Thus, each vertex in  $Z$  has exactly  $k$  outgoing edges. The capacity of each of these edges is 1. The flow-LTL formula  $\psi = kn \wedge XX(k \vee 0) \wedge XXX(\geq 1)$ . We now prove that  $\theta$  is satisfiable iff  $N \models \mathcal{E}A\psi$ .



Assume first that  $\theta$  is satisfiable. Consider the flow function  $f$  where for every  $1 \leq i \leq n$ , if  $x_i$  holds in the satisfying assignment, then  $f(x_i) = k$  and  $f(\bar{x}_i) = 0$ , and otherwise  $f(x_i) = 0$  and  $f(\bar{x}_i) = k$ . For such  $f$ , all the paths in  $N$  satisfy  $\psi$ .

Assume now that there is a flow function  $f$  with which  $N$  satisfies  $A\psi$ . For every  $1 \leq i \leq n$ , either we have  $f(x_i) = k$  and  $f(\bar{x}_i) = 0$  or we have  $f(x_i) = 0$  and  $f(\bar{x}_i) = k$ . Consider the assignment  $\tau$  to the variables  $x_1, \dots, x_n$  induced by  $f$ , where a literal  $z \in Z$  holds in  $\tau$  iff the flow in the corresponding vertex is positive. Since according to  $\psi$  for every  $i$  the flow in the vertex  $C_i$  is positive, then every clause in  $\theta$  contains at least one literal whose corresponding vertex has a positive flow. Hence, the assignment  $\tau$  satisfies  $\theta$ . Finally, note that  $\psi$  does not contain atomic propositions. Also, the same proof holds with the BFL formula  $\psi = kn \wedge AXEXk \wedge AXAXAX(\geq 1)$ .  $\square$

**5.1. Flow synthesis.** In the flow-synthesis problem, we are given a network  $N$  and an  $\exists\text{BFL}_1^*$  formula  $\mathcal{E}\varphi$ , and we have to return a flow function  $f$  with which  $\varphi$  is satisfied in  $N$ , or declare that no such function exists. The corresponding decision problem, namely, deciding whether there is a flow function with which  $\varphi$  is satisfied in  $N$ , is clearly at least as hard as  $\text{CTL}^*$  model checking. Also, by guessing  $f$ , its complexity does not go beyond  $\text{CTL}^*$  model-checking complexity. The network complexity of the problem coincides with that of  $\exists\text{BFL}_1^*$  model checking. Thus, we have the following.

**Theorem 5.6.** *The flow-synthesis problem for  $\exists\text{BFL}_1^*$  is PSPACE-complete, and its network complexity is NP-complete.*

## 6. MODEL CHECKING EXTENSIONS OF $\text{BFL}^*$

In Section 4, we defined several extensions of  $\text{BFL}^*$ . In this section we study the model-checking complexity for each of the extensions, and show that they do not require an increase in the complexity. The techniques for handling them are, however, richer: For positive path quantification, we have to refine the network and add a path-predicate that specifies positive flow, in a similar way fairness is handled in temporal logics. For maximal-flow quantification, we have to augment the model-checking algorithm by calls to a procedure that finds the maximal flow. For non-integral flow quantification, we have to reduce the model-checking problem to a solution of a linear-programming system. For past operators, we have to extend the model-checking procedure for  $\text{CTL}^*$  with branching past. Finally,



for first-order quantification over flow values, we have to first bound the range of relevant values, and then apply model checking to all relevant values.

**6.1. Positive path quantification.** Given a network  $N$ , it is easy to generate a network  $N'$  in which we add a vertex in the middle of each edge, and in which the positivity of paths correspond to positive flow in the new intermediate vertices. Formally, assuming that we label the new intermediate vertices by an atomic proposition  $edge$ , then the BFL<sup>\*</sup> path formula  $\xi_{positive} = G(edge \rightarrow > 0)$  characterizes positive paths, and replacing a state formula  $A\psi$  by the formula  $A(\xi_{positive} \rightarrow \psi)$  restricts the range of path quantification to positive paths. Now, given a BFL<sup>\*</sup> formula  $\varphi$ , it is easy to generate a BFL<sup>\*</sup> formula  $\varphi'$  such that  $N \models \varphi$  iff  $N' \models \varphi'$ . Indeed, we only have to (recursively) modify path formulas so that vertices labeled  $edge$  are ignored:  $X\xi$  is replaced by  $XX\xi$ , and  $\xi_1 U \xi_2$  is replaced by  $(\xi_1 \vee edge)U(\xi_2 \wedge \neg edge)$ . Hence, the complexity of model-checking is similar to BFL<sup>\*</sup>.

**6.2. Maximal flow quantification.** The maximal flow  $\gamma_{max}$  in a flow network can be found in polynomial time. Our model-checking algorithm for BFL<sup>\*</sup> described in the proof of Theorem 5.1 handles each flow state subformula  $\mathcal{E}\varphi$  by guessing a flow function  $f : E \rightarrow \mathbb{N}$  with which the Flow-CTL<sup>\*</sup> formula  $\varphi$  holds. For an  $\mathcal{E}^{max}$  quantifier, we can guess only flow functions for which the flow leaving the source vertex is  $\gamma_{max}$ . In addition, after calculating the maximal flow, we can substitute  $\gamma_{max}$ , in formulas that refer to it, by its value. Hence, the complexity of model-checking is similar to that of BFL<sup>\*</sup>.

**6.3. Non-integral flow quantification.** Recall that our BFL<sup>\*</sup> model-checking algorithm handles each flow state subformula  $\mathcal{E}\varphi$  by guessing a flow function  $f : E \rightarrow \mathbb{N}$  with which the Flow-CTL<sup>\*</sup> formula holds. Moving to non-integral flow functions, the guessed function  $f$  should be  $f : E \rightarrow \mathbb{R}$ , where we cannot bound the size or range of guesses.

Accordingly, in the non-integral case, we guess, for every vertex  $v \in V$ , an assignment to the flow propositions that appear in  $\varphi$ . Then, we perform two checks. First, that  $\varphi$  is satisfied with the guessed assignment — this is done by CTL<sup>\*</sup> model checking, as in the case of integral flows. Second, that there is a non-integral flow function that satisfies the flow constraints that appear in the vertices. This can be done in polynomial time by solving a system of linear inequalities [Sch03]. Thus, as in the integral case, handling each flow state formula  $\mathcal{E}\varphi$  can be done in PSPACE, and so is the complexity of the entire algorithm.

**6.4. Past operators.** Recall that our algorithm reduces BFL<sup>\*</sup> model checking to a sequence of calls to a CTL<sup>\*</sup> model-checking procedure. Starting with a BFL<sup>\*</sup> formula with past operators, the required calls are to a model-checking procedure for CTL<sup>\*</sup> with past. By [KPV12], model checking CTL<sup>\*</sup> with branching past is PSPACE-complete, and thus so is the complexity of our algorithm.

**6.5. First-Order quantification on flow values.** For a flow network

$$N = \langle AP, V, E, c, \rho, s, T \rangle,$$

let  $C_N = 1 + \sum_{e \in E} c(e)$ . Thus, for every flow function  $f$  for  $N$  and for every vertex  $v \in V$ , we have  $f(v) < C_N$ . We claim that when we reason about BFL<sup>\*</sup> formulas with quantified flow values, we can restrict attention to values in  $\{0, 1, \dots, C_N\}$ .

**Lemma 6.1.** *Let  $N$  be a flow network and let  $\theta = \forall x_1 \varphi$  be a BFL<sup>\*</sup>( $\{+, *\}$ ) formula over the variables  $X = \{x_1, \dots, x_n\}$ , and without free variables. Then,  $N \models \theta$  iff  $N \models \varphi[x_1 \leftarrow \gamma]$ , for every  $0 \leq \gamma \leq C_N$ .*

*Proof.* We show that for every  $\gamma > C_N$  we have  $N \models \varphi[x_1 \leftarrow \gamma]$  iff  $N \models \varphi[x_1 \leftarrow C_N]$ . Every expression  $g(X)$  in  $\varphi$  is obtained from the variables in  $X$  by applying the operators  $+, *$ , and possibly by using constants in  $\mathbb{N}$ . Hence, every such  $g$  is monotonically increasing for every variable  $x_i$ . Let  $\gamma_2, \dots, \gamma_n$  be an assignment for the variables  $x_2, \dots, x_n$ , respectively. Let  $g_1(x_1) = g(x_1, \gamma_2, \dots, \gamma_n)$ . Let  $\gamma_1$  be an integer bigger than  $C_N$ . Since  $g_1$  is monotonically increasing, we have  $g_1(C_N) \leq g_1(\gamma_1)$ . If  $g_1(C_N) < g_1(\gamma_1)$ , then  $g_1$  is not a constant function, and since it contains only the operators  $+, *$  and constants in  $\mathbb{N}$ , we have  $g_1(C_N) \geq C_N$ . Therefore, if  $g_1(C_N) \neq g_1(\gamma_1)$ , then  $g_1(\gamma_1) > g_1(C_N) \geq C_N$ . Recall that for every flow function  $f$  in  $N$  and for every vertex  $v \in V$ , we have  $f(v) < C_N$ . Therefore, every state subformula  $> g_1(x_1)$  and every state subformula  $\geq g_1(x_1)$  holds for  $\gamma_1$  iff it holds for  $C_N$ . Therefore,  $N \models \varphi[x_1 \leftarrow \gamma]$  iff  $N \models \varphi[x_1 \leftarrow C_N]$ . Hence, we have  $N \models \forall x_1 \varphi$  iff  $N \models \varphi[x_1 \leftarrow \gamma]$  for every  $0 \leq \gamma \leq C_N$ .  $\square$

By Lemma 6.1, the model-checking problem for BFL<sup>\*</sup>( $\{+, *\}$ ) is PSPACE-complete.

## 7. A POLYNOMIAL FRAGMENT

The upper bounds to the complexity of the model-checking problem that are given in Sections 5 and 6 refer to worst-case complexity. As has been the case of CTL<sup>\*</sup>, in practice the complexity is often lower. In particular, there are fragments of CTL<sup>\*</sup>, and hence also BFL<sup>\*</sup> and its extensions for which we can show that the described model-checking algorithms perform better, say by showing that the particular syntax guarantees a bound on the size of the automata associated with path formulas. In this section we show that the model-checking problem for CBFL<sup>\*</sup> (see Section 4.6) can be solved in polynomial time. The result is not based on a tighter complexity analysis for the case the specification is in CBFL<sup>\*</sup> but rather on a different model-checking algorithm for this logic.

Our model-checking algorithm reduces the evaluation of a CBFL<sup>\*</sup> formula into a sequence of solutions to the *vertex-constrained flow problem*. In this problem, we are given a flow network  $N = \langle AP, V, E, c, \rho, s, T \rangle$  in which each vertex  $v \in V$  is attributed by a range  $[\gamma_l, \gamma_u] \in \mathbb{N} \times (\mathbb{N} \cup \{\infty\})$ . The problem is to decide whether there is a flow function  $f : E \rightarrow \mathbb{R}$  such that for all vertices  $v \in V$ , we have  $\gamma_l \leq f(v) \leq \gamma_u$ .

**Lemma 7.1.** *The vertex-constrained flow problem can be solved in polynomial time. If there is a solution that is a non-integral flow function, then there is also a solution that is an integral flow function, and the algorithm returns such a solution.*

*Proof.* Given  $N$ , we construct a new flow network  $N'$  in which the attributions are on the edges. Thus,  $N'$  is a flow network with lower bounds on the flow. We construct  $N'$  by splitting every vertex  $v \in V$  with attribution  $[\gamma_l, \gamma_u]$  into two vertices in  $N'$ :  $v_{in}$  and

$v_{out}$ . The vertices are connected by an edge  $\langle v_{in}, v_{out} \rangle$  with attribution  $[\gamma_l, \gamma_u]$ . For every edge  $\langle u, v \rangle \in E$ , we add to  $N'$  an edge  $\langle u_{out}, v_{in} \rangle$  with attribution  $[0, c(\langle u, v \rangle)]$ . It is well known that deciding whether there exists a feasible flow and finding such a flow in a flow network with lower bounds can be done in polynomial time (see, for example, Chapter 6.7 in [AMO93]). Also, the algorithm for finding such a feasible flow reduces the problem to a maximum-flow problem. Accordingly, if there is a feasible flow then the algorithm returns an integral one.  $\square$

**Theorem 7.2.** *CBFL $^*$  model checking can be solved in polynomial time.*

*Proof.* Let  $N = \langle AP, V, E, c, \rho, s, T \rangle$ , and consider a CBFL $^*$  formula  $\varphi$ . If  $\varphi$  is in CBFL $^*_0$ , we can clearly label in linear time all the vertices in  $N$  by a fresh atomic proposition  $p_\varphi$  that maintains the satisfaction of  $\varphi$ . That is, in all vertices  $v \in V$ , we have that  $p_\varphi \in \rho(v)$  iff  $v \models \varphi$ . Otherwise,  $\varphi$  is a CBFL $^*_{i+1}$  formula for some  $i \geq 0$ . We show how, assuming that the vertices of  $N$  are labeled by atomic propositions that maintain satisfaction of the subformulas of  $\varphi$  that are CBFL $^*_i$  formulas, we can label them, in polynomial time, by a fresh atomic proposition that maintains the satisfaction of  $\varphi$ .

Recall that  $\varphi$  is a Boolean assertion over CBFL $^*_i$  formulas and flow formulas of the form  $\mathcal{E}(A\psi_1 \wedge \dots \wedge A\psi_n)$ , where each  $\psi_j$  is of the form  $X^{k_j}\xi_j$  or  $X^{k_j}G\xi_j$ , where  $k_j \geq 0$  and  $\xi_j$  is a CBFL $^*_i$  formula or a flow proposition (that is,  $> \gamma$ ,  $< \gamma$ ,  $\geq \gamma$ , or  $\leq \gamma$ , for an integer  $\gamma \in \mathbb{N}$ ).

Since CBFL $^*_i$  subformulas have already been evaluated, we describe how to evaluate subformulas of the form  $\theta = \mathcal{E}(A\psi_1 \wedge \dots \wedge A\psi_n)$ . Intuitively, since the formulas in  $\theta$  include no disjunctions, they impose constraints on the vertices of  $N$  in a deterministic manner. These constraints can be checked in polynomial time by solving a vertex-constrained flow problem. Recall that for each  $1 \leq j \leq n$ , the formula  $\psi_j$  is of the form  $X^{k_j}\xi_j$  or  $X^{k_j}G\xi_j$ , for some  $k_j \geq 0$ , and a CBFL $^*_i$  formula or a flow proposition  $\xi_j$ . In order to evaluate  $\theta$  in a vertex  $v \in V$ , we proceed as follows. For each  $1 \leq j \leq n$ , the formula  $\xi_j$  imposes either a Boolean constraint (in case  $\xi_j$  is a CBFL $^*_i$  formula) or a flow constraint (in case  $\xi_j$  is a flow proposition) on a finite subset  $V_j^v$  of  $V$ . Indeed, if  $\psi_j = X^{k_j}\xi_j$ , then  $V_j^v$  includes all the vertices reachable from  $v$  by a path of length  $k_j$ , and if  $\psi_j = X^{k_j}G\xi_j$ , then  $V_j^v$  includes all the vertices reachable from  $v$  by a path of length at least  $k_j$ . We attribute each vertex by the constraints imposed on it by all the conjuncts in  $\theta$ . If one of the Boolean constraints does not hold, then  $\theta$  does not hold in  $v$ . Otherwise, we obtain a set of flow constraints for each vertex in  $V$ . For example, if  $\theta = \mathcal{E}(AXXp \wedge AXX > 5 \wedge AG \leq 8)$ , then in order to check whether  $\theta$  holds in  $s$ , we assign the flow constraint  $\leq 8$  to all the vertices reachable from  $s$ , and assign the flow constraint  $> 5$  to all the successors of the successors of  $s$ . If one of these successors of successors does not satisfy  $p$ , we can skip the check for a flow and conclude that  $s$  does not satisfy  $\theta$ . Otherwise, we search for such a flow, as described below.

The flow constraints for a vertex induce a closed, open, or half-closed range. The upper bound in the range may be infinity. For example, the constraints  $> 6$ ,  $< 10$ ,  $\leq 8$  induce the half-closed range  $(6, 8]$ . Note that it may be that the induced range is empty. For example, the constraints  $\leq 6$  and  $> 8$  induce an empty range. Then,  $\theta$  does not hold in  $v$ . Since we are interested in integral flows, we can convert all strict bounds to non-strict ones. For example, the range  $(6, 8]$  can be converted to  $[7, 8]$ . Note that since we are interested in integral flow, a non-empty open range may not be satisfiable, and we refer to it as an empty range. For example, the range  $(6, 7)$  is empty. Hence, the satisfaction of  $\theta$  in  $v$  is reduced to an instance of the vertex-constrained flow problem. By Lemma 7.1, deciding whether there is a flow function that satisfies the constraints can be solved in polynomial time.  $\square$

**Remark 7.3.** Note that the same algorithm can be applied when we consider non-integral flow functions, namely in CBFL<sup>\*</sup> with the  $\mathcal{A}^{\mathbb{R}}$  flow quantifier. There, the induced vertex-constrained flow problem may include open boundaries. The solution need not be integral, but can be found in polynomial time by solving a system of inequalities [Sch03].

## 8. QUERY CHECKING

As discussed in Section 1, query checking is a useful methodology for system exploration. A query is a specification with the place-holder “?”, and the goal is to find replacements to ? with which the specification holds. In this section we extend the methodology to BFL<sup>\*</sup>. We first need some definitions.

A *propositional BFL<sup>\*</sup> query* (propositional query, for short) is a BFL<sup>\*</sup> formula in which a single state formula is “?”. For a propositional query  $\psi$  and a propositional assertion  $\theta$  over  $AP$ , we denote by  $\psi[? \leftarrow \theta]$  the formula obtained from  $\psi$  by replacing ? by  $\theta$ . Given a network  $N$  and a query  $\psi$ , a *solution* to  $\psi$  in  $N$  is a propositional assertion  $\theta$  over  $AP$  such that  $N \models \psi[? \leftarrow \theta]$ . For example, the propositional query  $\mathcal{E}\mathcal{E}((? \wedge \geq 5)Utarget)$  asks which propositional assertions  $\theta$  are such that the network satisfies  $\mathcal{E}\mathcal{E}((\theta \wedge \geq 5)Utarget)$ , namely propositional assertions  $\theta$  such that there is a flow and there is a path to the target in which all vertices satisfy  $\theta$  and the flow in them is at least 5. Note that we do not allow the solution  $\theta$  to include flow propositions. Note also that only a single occurrence of ? is allowed in the query. Richer settings in CTL<sup>\*</sup> query checking allow queries with several place holders, namely  $?_1, ?_2, \dots, ?_m$ , possibly with multiple appearances to each of them [GC02]. The focus on queries with a single ? enables the query checker to refer to the *polarity* of queries. Formally, we say that a propositional query  $\psi$  is *positive* (*negative*) if the single ? in it is in the scope of an even (respectively, odd) number of negations. The polarity of queries implies monotonicity, in the following sense (the proof proceeds by an easy induction on the structure of the query).

**Lemma 8.1.** *Consider a positive (negative) BFL<sup>\*</sup> query  $\psi$ . Let  $\theta$  and  $\theta'$  be propositional formulas such that  $\theta$  implies  $\theta'$  ( $\theta'$  implies  $\theta$ , respectively). Then,  $\psi[? \leftarrow \theta]$  implies  $\psi[? \leftarrow \theta']$ .*

**Corollary 8.2.** *Consider a network  $N$ . Let  $\psi$  be a positive (negative) BFL<sup>\*</sup> query, let  $\theta$  be a solution for  $\psi$  in  $N$ , and let  $\theta'$  be a propositional formula such that  $\theta$  implies  $\theta'$  ( $\theta'$  implies  $\theta$ , respectively). Then,  $\theta'$  is a solution for  $\psi$  in  $N$ .*

It is not hard to see that a propositional BFL<sup>\*</sup> query may have multiple solutions. Corollary 8.2 enables us to partially order them. Then, given a query  $\psi$ , we say that a solution  $\theta$  is *strongest* if there is no solution  $\theta'$  such that  $\theta'$  is not equivalent to  $\theta$  and either  $\psi$  is positive, in which case  $\theta'$  implies  $\theta$ , or  $\psi$  is negative, in which case  $\theta$  implies  $\theta'$ . A query, however, may not only have multiple solutions, but may also have multiple strongest solutions. Accordingly, as is the case with CTL<sup>\*</sup>, a search for all strongest solutions has to examine all possible solutions [BG01]. Consider a network  $N$  and a propositional BFL<sup>\*</sup> query  $\psi$ . Each solution to  $\psi$  in  $N$  corresponds to an assignment to all subsets of atomic propositions, thus there are  $2^{2^{|AP|}}$  assertions to check. For each we need to model check the formula  $\psi[? \leftarrow \theta]$ , for an assertion  $\theta$ . Accordingly, the length of  $\psi[? \leftarrow \theta]$  is  $O(|\psi| + 2^{|AP|})$ . Thus, finding all strongest solutions to a propositional query is a very complex task, involving  $2^{2^{|AP|}}$  executions of the algorithm described in Theorem 5.1 for a formula of length  $O(|\psi| + 2^{|AP|})$ .

Fortunately, in the context of  $\text{BFL}^*$ , we are able to point to a class of interesting queries for which query checking is not more complex than model checking: queries in which the place holder is the value in a flow proposition. Formally, a *value BFL\* query* (value query, for short) is a  $\text{BFL}^*$  formula in which a single flow proposition is of the form  $>?$ ,  $\geq?$ ,  $<?$ , or  $\leq?$ . Each value query  $\psi$  is either a *lower-bound query*, in case  $\psi$  is positive iff the flow proposition with  $?$  is  $>?$  or  $\geq?$  (that is, either  $\psi$  is positive and the flow proposition with  $?$  is  $>?$  or  $\geq?$ , or  $\psi$  is negative and the flow proposition with  $?$  is  $<?$  or  $\leq?$ ), or an *upper-bound query*, in case  $\psi$  is positive iff the flow proposition with  $?$  is  $<?$  or  $\leq?$ . For a value query  $\psi$  and an integer  $\gamma \in \mathbb{N}$ , we denote by  $\psi[? \leftarrow \gamma]$  the  $\text{BFL}^*$  formula obtained from  $\psi$  by replacing  $?$  by  $\gamma$ . Given a network  $N$  and a value query  $\psi$ , a *solution* to  $\psi$  in  $N$  is an integer  $\gamma \in \mathbb{N}$  such that  $N \models \psi[? \leftarrow \gamma]$ . For example, the value query  $\mathcal{E}EF(\text{target} \wedge (\geq?))$  asks which values of flow can reach a target vertex.

For a value query  $\psi$ , we say that a solution  $\gamma$  to  $\psi$  in  $N$  is *strongest* if  $\psi$  is a lower-bound query and there is no solution  $\gamma'$  to  $\psi$  in  $N$  such that  $\gamma' > \gamma$ , or  $\psi$  is an upper-bound query and there is no solution  $\gamma'$  to  $\psi$  in  $N$  such that  $\gamma' < \gamma$ . Since a strongest solution in a value query is either a maximum or a minimum of a set in  $\mathbb{N}$ , there is at most one strongest solution.<sup>3</sup> Hence the following lemma.

**Lemma 8.3.** *Consider a network  $N$ . Let  $\gamma$  be the strongest solution for a value  $\text{BFL}^*$  query  $\psi$ . If  $\psi$  is a lower-bound query, then the set of solutions is  $\mathbb{N} \cap [\gamma, \infty)$ . If  $\psi$  is an upper-bound query then the set of solutions is  $\mathbb{N} \cap [0, \gamma]$ .*

**Remark 8.4.** Note that in our definition of a value query, we do not allow a flow proposition of the form  $=?$ . Indeed, such a flow proposition encodes a conjunction of two flow propositions of different polarities. If we decide to allow such value queries then the set of solutions does not satisfy a property as in Lemma 8.3. For example, the set of solutions for the query  $\mathcal{E}[(=?) \wedge ((\geq 2 \wedge \leq 4) \vee (\geq 6 \wedge \leq 9))]$  and a network with a maximal flow of 9, is  $\{2, 3, 4, 6, 7, 8, 9\}$ .

Thus, unlike the case of propositional queries, in the case of value queries we can talk about the strongest solution, and study the complexity of finding it. The lower bound in the following theorem corresponds to the problem of deciding whether there is a solution to a given value  $\text{BFL}^*$  query.

**Theorem 8.5.** *The problem of finding the strongest solution to a value  $\text{BFL}^*$  query is PSPACE-complete. Furthermore, the network complexity is in  $\Delta_2^P$ .*

*Proof.* Model checking of a  $\text{BFL}^*$  formula  $\psi$  can be reduced to deciding whether there is a solution for the value query  $\psi \wedge (\mathcal{A} \leq?)$ . Hence the PSPACE-hardness.

Lemma 8.3 implies that in order to find the strongest solution for a value query  $\psi$ , we can run a binary search on the values, and for each value  $\gamma$  to check whether  $\psi[? \leftarrow \gamma]$  holds using  $\text{BFL}^*$  model-checking. Thus, the number of calls to the  $\text{BFL}^*$  model checker is logarithmic in  $\sum_{e \in E} c(e)$ , and hence polynomial in  $|N|$ . The PSPACE upper bound then follows from the PSPACE upper bound for  $\text{BFL}^*$  model checking. The  $\Delta_2^P$  upper bound for the network complexity follows from the  $\Delta_2^P$  upper bound for the network complexity of  $\text{BFL}^*$  model checking.  $\square$

<sup>3</sup>Note that if we allow solutions in  $\mathbb{R}$ , then flow propositions with strict inequality do not have unique strongest solutions. In the case of solutions in  $\mathbb{N}$ , however, the flow propositions  $< \gamma$  and  $> \gamma$  are equivalent to  $\leq \gamma - 1$  and  $\geq \gamma + 1$ , respectively.

## 9. DISCUSSION

We introduced the flow logic  $\text{BFL}^*$  and studied its theoretical and practical aspects, as well as extensions and fragments of it. Below we discuss possible directions for future research.

At the more theoretical front, an important aspect that we left open in this work is the *expressive power* of the different extensions and fragment of  $\text{BFL}^*$ . As discussed in Remark 4.6, some cases follow easily from known expressiveness results for  $\text{CTL}^*$ , yet the full picture is open. Also, as has been the case in traditional temporal logics, questions of *succinctness* are of interest too.

Another natural problem regarding  $\text{BFL}^*$  is the *satisfiability* problem. That is, given a  $\text{BFL}^*$  formula, decide whether there is a flow network that satisfies it. In Section 3.2, we showed that, unlike  $\text{CTL}^*$ , the logic  $\text{BFL}^*$  is sensitive to unwinding. Thus, the standard algorithm for the satisfiability problem of  $\text{CTL}^*$ , which is based on checking emptiness of tree automata, is not useful in the case of  $\text{BFL}^*$ . Moreover, the satisfiability problem is challenging already to the linear fragment LFL of  $\text{BFL}^*$ . Indeed, as discussed in Section 4.6, the semantics of LFL mixes linear and branching semantics, and there is no simple reduction of the satisfiability problem for LFL into the emptiness of word automata.

In the algorithmic side, the relation between maximum flows and minimum cuts has been significant in the context of the traditional maximum-flow problem. It would be interesting to study the relation between cuts and  $\text{BFL}^*$ . In particular, the notion of cuts that satisfy some structural property that can be expressed with  $\text{BFL}^*$  seems interesting.

Finally, many problems in various domains are solved using a reduction to the maximum-flow problem [CLR90, AMO93]. Flow logics allow reasoning about properties of flow networks that go beyond their maximal flow. An interesting direction for future research is to study applications of flow logics for rich variants of such problems. For example, using flow logics we can solve variants of matching or scheduling problems that involve restrictions on the allowed matches or variants of schedules.

## REFERENCES

- [AH94] R. Alur and T.A. Henzinger. A really temporal logic. *Journal of the ACM*, 41(1):181–204, 1994.
- [AHK02] R. Alur, T.A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49(5):672–713, 2002.
- [AMO93] R.K. Ahuja, T.L. Magnanti, and J.B. Orlin. *Network flows: Theory, algorithms, and applications*. Prentice Hall Englewood Cliffs, 1993.
- [BDVV18] R. Birkner, D. Drachler-Cohen, L. Vanbever, and M.T. Vechev. Net2Text: Query-guided summarization of network forwarding behaviors. In *15th USENIX Symposium on Networked Systems Design and Implementation*, pages 609–623, 2018.
- [Bei91] R. Beigel. Bounded queries to sat and the boolean hierarchy. *Theoretical Computer Science*, 84(2):199–223, 1991.
- [BG01] G. Bruns and P. Godefroid. Temporal logic query checking. In *Proc. 16th IEEE Symp. on Logic in Computer Science*, pages 409–420. IEEE Computer Society, 2001.
- [CE81] E.M. Clarke and E.A. Emerson. Design and synthesis of synchronization skeletons using branching time temporal logic. In *Proc. Workshop on Logic of Programs*, volume 131 of *Lecture Notes in Computer Science*, pages 52–71. Springer, 1981.
- [CG03] M. Chechik and A. Gurfinkel. TLQSolver: A temporal logic query checker. In *Proc. 15th Int. Conf. on Computer Aided Verification*, Lecture Notes in Computer Science, pages 210–214. Springer, 2003.
- [CGG07] M. Chechik, M. Gheorghiu, and A. Gurfinkel. Finding state solutions to temporal queries. In *Proc. Integrated Formal Methods*, pages 273–292, 2007.
- [CGP99] E.M. Clarke, O. Grumberg, and D. Peled. *Model Checking*. MIT Press, 1999.

- [Cha00] W. Chan. Temporal-logic queries. In *Proc. 12th Int. Conf. on Computer Aided Verification*, volume 1855 of *Lecture Notes in Computer Science*, pages 450–463. Springer, 2000.
- [CHP07] K. Chatterjee, T. A. Henzinger, and N. Piterman. Strategy logic. In *Proc. 18th Int. Conf. on Concurrency Theory*, pages 59–73, 2007.
- [CHVB18] E.M. Clarke, T.A. Henzinger, H. Veith, and R. Bloem. *Handbook of model checking*. Springer, 2018.
- [CLR90] T.H. Cormen, C.E. Leiserson, and R.L. Rivest. *Introduction to Algorithms*. MIT Press and McGraw-Hill, 1990.
- [CVP<sup>+</sup>12] M. Canini, D. Venzano, P. Perešini, D. Kostić, and J. Rexford. A nice way to test openflow applications. In *9th USENIX Symposium on Networked Systems Design and Implementation (NSDI 12)*, pages 127–140, 2012.
- [Din70] E.A. Dinic. Algorithm for solution of a problem of maximum flow in a network with power estimation. *Soviet Math. Doll*, 11(5):1277–1280, 1970. English translation by RF. Rinehart.
- [EIS76] S. Even, A. Itai, and A. Shamir. On the complexity of timetable and multicommodity flow problems. *SIAM Journal on Computing*, 5(4):691–703, 1976.
- [EK72] J. Edmonds and R.M. Karp. Theoretical improvements in algorithmic efficiency for network flow problems. *Journal of the ACM*, 19(2):248–264, 1972.
- [EL87] E.A. Emerson and C.-L. Lei. Modalities for model checking: Branching time logic strikes back. *Science of Computer Programming*, 8:275–306, 1987.
- [FF56] L.R. Ford and D.R. Fulkerson. Maximal flow through a network. *Canadian journal of Mathematics*, 8(3):399–404, 1956.
- [FF62] L.R. Ford and D.R. Fulkerson. *Flows in networks*. Princeton Univ. Press, Princeton, 1962.
- [FKL10] D. Fisman, O. Kupferman, and Y. Lustig. Rational synthesis. In *Proc. 16th Int. Conf. on Tools and Algorithms for the Construction and Analysis of Systems*, volume 6015 of *Lecture Notes in Computer Science*, pages 190–204. Springer, 2010.
- [GC02] A. Gurfinkel and M. Chechik. Temporal logic query checking through multi-valued model checking. Technical report, CSRG Technical Report 457, University of Toronto, Department of Computer Science, 2002.
- [GCSR13] D. Granata, R. Cerulli, M.G. Scutellá, and A. Raiconi. Maximum flow problems and an NP-complete variant on edge-labeled graphs. In *Handbook of Combinatorial Optimization*, pages 1913–1948. Springer, 2013.
- [GRN12] M. Gueffaz, S. Rampacek, and C. Nicolle. Temporal logic to query semantic graphs using the model checking method. *Journal of Software*, 7(7), 2012.
- [GT88] A.V. Goldberg and R.E. Tarjan. A new approach to the maximum-flow problem. *Journal of the ACM*, 35(4):921–940, 1988.
- [GTT89] A.V. Goldberg, É. Tardos, and R.E. Tarjan. Network flow algorithms. Technical report, DTIC Document, 1989.
- [HM90] J.Y. Halpern and Y. Moses. Knowledge and common knowledge in a distributed environment. *Journal of the ACM*, 37(3):549–587, 1990.
- [KCZ<sup>+</sup>13] P. Kazemian, M. Chang, H. Zeng, G. Varghese, N. McKeown, and S. Whyte. Real time network policy checking using header space analysis. In *10th USENIX Symposium on Networked Systems Design and Implementation (NSDI 13)*, pages 99–111, 2013.
- [KPC<sup>+</sup>12] M. Kuzniar, P. Peresini, M. Canini, D. Venzano, and D. Kostic. A soft way for openflow switch interoperability testing. In *Proceedings of the 8th international conference on Emerging networking experiments and technologies*, pages 265–276. ACM, 2012.
- [KPV12] O. Kupferman, A. Pnueli, and M.Y. Vardi. Once and forall. *Journal of Computer and Systems Science*, 78(3):981–996, 2012.
- [KT14] O. Kupferman and T. Tamir. Properties and utilization of capacitated automata. In *Proc. 34th Conf. on Foundations of Software Technology and Theoretical Computer Science*, volume 29 of *LIPICs*, pages 33–44. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, Germany, 2014.
- [KVM12] P. Kazemian, G. Varghese, and N. McKeown. Header space analysis: Static checking for networks. In *9th USENIX Symposium on Networked Systems Design and Implementation (NSDI 12)*, pages 113–126, 2012.
- [KVW00] O. Kupferman, M.Y. Vardi, and P. Wolper. An automata-theoretic approach to branching-time model checking. *Journal of the ACM*, 47(2):312–360, 2000.

- [KZZ<sup>+</sup>13] A. Khurshid, X. Zou, W. Zhou, M. Caesar, and P.B. Godfrey. Veriflow: verifying network-wide invariants in real time. In *10th USENIX Symposium on Networked Systems Design and Implementation (NSDI 13)*, pages 15–27, 2013.
- [LBG<sup>+</sup>15] N.P. Lopes, N. Bjørner, P. Godefroid, K. Jayaraman, and G. Varghese. Checking beliefs in dynamic networks. In *12th USENIX Symposium on Networked Systems Design and Implementation (NSDI 15)*, pages 499–512, 2015.
- [LCG<sup>+</sup>06] B.T. Loo, T. Condie, M. Garofalakis, D.E. Gay, J.M. Hellerstein, P. Maniatis, R. Ramakrishnan, T. Roscoe, and I. Stoica. Declarative networking: language, execution and optimization. In *Proceedings of the 2006 ACM SIGMOD international conference on Management of data*, pages 97–108. ACM, 2006.
- [LP85] O. Lichtenstein and A. Pnueli. Checking that finite state concurrent programs satisfy their linear specification. In *Proc. 12th ACM Symp. on Principles of Programming Languages*, pages 97–107, 1985.
- [Mad16] A. Madry. Computing maximum flow with augmenting electrical flows. In *Foundations of Computer Science (FOCS), 2016 IEEE 57th Annual Symposium on*, pages 593–602. IEEE, 2016.
- [Mil71] R. Milner. An algebraic definition of simulation between programs. In *Proc. 2nd Int. Joint Conf. on Artificial Intelligence*, pages 481–489. British Computer Society, 1971.
- [QS82] J.P. Queille and J. Sifakis. Specification and verification of concurrent systems in Cesar. In *Proc. 9th ACM Symp. on Principles of Programming Languages*, volume 137 of *Lecture Notes in Computer Science*, pages 337–351. Springer, 1982.
- [RCM<sup>+</sup>14] M. Räim, C. Di Ciccio, F.M. Maggi, M. Mecella, and J. Mendling. Log-based understanding of business processes through temporal logic query checking. In *OTM Confederated International Conferences "On the Move to Meaningful Internet Systems"*, pages 75–92. Springer, 2014.
- [SC85] A.P. Sistla and E.M. Clarke. The complexity of propositional linear temporal logic. *Journal of the ACM*, 32:733–749, 1985.
- [Sch66] B.L. Schwartz. Possible winners in partially completed tournaments. *SIAM Review*, 8(3):302–308, 1966.
- [Sch03] A. Schrijver. *Combinatorial Optimization*. Springer, 2003.
- [SV03] M. Samer and H. Veith. Validity of ctl queries revisited. In *Proc. 12th Annual Conf. of the European Association for Computer Science Logic*, volume 2803, pages 470–483. Springer, 2003.
- [Tar85] É. Tardos. A strongly polynomial minimum cost circulation algorithm. *Combinatorica*, 5(3):247–255, 1985.
- [VAP<sup>+</sup>16] Y. Velner, K. Alpernas, A. Panda, A. Rabinovich, M. Sagiv, S. Shenker, and S. Shoham. Some complexity results for stateful network verification. In *Proc. 23rd Int. Conf. on Tools and Algorithms for the Construction and Analysis of Systems*, pages 811–830. Springer, 2016.
- [WBL09] A. Wang, P. Basu, B.T. Loo, and O. Sokolsky. Declarative network verification. In *International Symposium on Practical Aspects of Declarative Languages*, pages 61–75. Springer, 2009.
- [Wee85] G. Wechsung. On the boolean closure of np. In *Proceedings of the International Conference on Fundamentals of Computation Theory*, volume 199 of *Lecture Notes in Computer Science*, pages 485–493. Springer, 1985.
- [ZST<sup>+</sup>10] W. Zhou, M. Sherr, T. Tao, X. Li, B.T. Loo, and Y. Mao. Efficient querying and maintenance of network provenance at internet-scale. In *Proceedings of the 2010 ACM SIGMOD International Conference on Management of data*, pages 615–626. ACM, 2010.