



# Liveness in broadcast networks

Peter Chini<sup>1</sup> · Roland Meyer<sup>1</sup> · Prakash Saivasan<sup>2</sup>

Received: 30 June 2020 / Accepted: 21 July 2021 / Published online: 10 August 2021

© The Author(s) 2021

## Abstract

We study liveness and model checking problems for broadcast networks, a system model of identical clients communicating via message passing. The first problem that we consider is *Liveness Verification*. It asks whether there is a computation such that one client visits a final state infinitely often. The complexity of the problem has been open. It was shown to be P-hard but in EXPSpace. We close the gap by a polynomial-time algorithm. The latter relies on a characterization of live computations in terms of paths in a suitable graph, combined with a fixed-point iteration to efficiently check the existence of such paths. The second problem is *Fair Liveness Verification*. It asks for a computation where all participating clients visit a final state infinitely often. We adjust the algorithm to also solve fair liveness in polynomial time. Both problems can be instrumented to answer model checking questions for broadcast networks against linear time temporal logic specifications. The first problem in this context is *Fair Model Checking*. It demands that for all computations of a broadcast network, all participating clients satisfy the specification. We solve the problem via the Vardi–Wolper construction and a reduction to *Liveness Verification*. The second problem is *Sparse Model Checking*. It asks whether each computation has a participating client that satisfies the specification. We reduce the problem to *Fair Liveness Verification*.

**Keywords** Liveness verification · Broadcast networks · Model checking

**Mathematics Subject Classification** 68Q17 · 68Q45

---

P. Saivasan: (A part of the work was done when the author was at TU Braunschweig).

---

✉ Peter Chini  
p.chini@tu-bs.de  
Roland Meyer  
roland.meyer@tu-bs.de  
Prakash Saivasan  
prakashs@imsc.res.in

<sup>1</sup> TU Braunschweig, Braunschweig, Germany

<sup>2</sup> The Institute of Mathematical Sciences, Chennai, India

## 1 Introduction

Parameterized systems consist of an arbitrary number of identical clients that communicate via some mechanism like shared memory or message passing [4]. Parameterized systems appear in various applications. In distributed algorithms, a group of clients has to form a consensus [35]. In cache-coherence protocols, coherence has to be guaranteed for data shared among threads [15]. Developing parameterized systems is difficult. The desired functionality has to be achieved not only for a single system instance but for an arbitrary number of clients that is not known a priori. The proposed solutions are generally tricky and sometimes buggy [3], which has led to substantial interest in parameterized verification [8], verification algorithms for parameterized systems.

Broadcast networks are a successful model in parameterized verification [5,7,9,10,16,17,22,25,26,29,43]. A broadcast network consists of an arbitrary number of identical finite-state automata communicating via passing messages. We call these clients since they reflect the interaction of a single client in the parameterized system with its environment. When a client sends a message (by taking a send transition), at the same time some clients receive the message (by taking a corresponding receive transition). A client ready to receive a message may ignore it, and it may be the case that no one receives it.

What makes broadcast networks interesting is the surprisingly low complexity of their verification problems. Earlier works have concentrated on safety verification. In the coverability problem, the question is whether at least one participating client can reach an unsafe state. The problem has been shown to be solvable in polynomial time [16]. In the synchronization problem, all clients need to visit a final state at the same time. Although seemingly harder than coverability, it turned out to be solvable in polynomial time as well [28]. Both problems remain in  $\mathbb{P}$  if the communication topology is slightly restricted [5], a strengthening that usually leads to undecidability results [5,17].

The focus of our work is on liveness verification and model checking. Liveness properties formulate good events that should happen during a computation. To give an example, one would state that every request has to be followed by a response. In the setting of broadcast networks, liveness verification was studied in [17]. The problem generalizes coverability in that at least one client needs to visit a final state infinitely many times. The problem was shown to be solvable in  $\text{EXPSpace}$  by a reduction to repeated coverability in Petri Nets [23,27]. The only known lower bound, however, is  $\mathbb{P}$ -hardness [16].

Our main contribution is an algorithm that solves the liveness verification problem in polynomial time. It closes the gap. We also address a fair variant of liveness verification where all clients participating infinitely often have to see a final state infinitely often, a requirement known as compassion [40]. We give an instrumentation that compiles away compassion and reduces the problem to finding cycles. By our results, safety and liveness verification have the same complexity, a phenomenon also observed in other models [22,24,29,30].

At the heart of our liveness verification algorithm is a fixed-point iteration terminating in polynomial time. It relies on an efficient representation of computations. We first characterize live computations in terms of paths in a suitable graph. Since the

graph is of exponential size, we cannot immediately apply a path finding algorithm. Instead, we show that a path exists if and only if there is a path in some normal form. Paths in normal form can then be found efficiently by the fixed-point iteration.

Our results yield efficient algorithms for model checking broadcast networks against linear time temporal logic (LTL) specifications [39]. Formally, we consider two variants of model checking, like for liveness verification. The first variant incorporates a notion of fairness. Given a broadcast network and a specification, it demands that in each computation, all clients satisfy the specification. The second variant asks for each computation having at least one client satisfying it. We solve both problems by employing the Vardi–Wolper construction [44] and by applying the aforementioned fixed-point iteration for liveness verification. The results show that the given broadcast network only contributes a polynomial factor to the running time needed for model checking.

The paper at hand is an extension of the conference version [12]. It features a new section, namely Sect. 5, showing how model checking problems for broadcast networks can be solved by applying the fixed-point iteration for liveness verification. More precise, the section introduces two model checking problems and presents algorithms along with their time analyses. Moreover, we provide a full version of the paper [13], including all missing proofs.

## 1.1 Related work

Broadcast networks [17,25,43] were introduced to verify ad hoc networks [34,42]. Ad hoc networks are reconfigurable in that the number of clients as well as their communication topology may change during the computation. If the transition relation is compatible with the topology, safety verification has been shown to be decidable [32]. Related studies do not assume compatibility but restrict the topology [31]. If the dependencies among clients are bounded [37], safety verification is decidable independent of the transition relation [45,46]. Verification tools turn these decision procedures into practice [20,38]. D’Osualdo and Ong suggested a typing discipline for the communication topology [21]. In [5], decidability and undecidability results for reachability problems were proven for a locally changing topology. The case when communication is fixed along a given graph was studied in [1]. Topologies with bounded diameter were considered in [18]. Perfect communication where a sent message is received by all clients was studied in [25]. Networks with communication failures were considered in [19]. In [7], a variant of broadcast networks was considered where the clients follow a local strategy. Probabilistic broadcast networks were studied in [6]. By a closer observation, one can find that membership in  $\mathcal{P}$  for the liveness verification problem can also be deduced from a result given in [6]. The corresponding result concerns the complexity of a parity game on broadcast networks. The idea is to construct a suitable Vector Addition Systems with states (VASS) and to employ an algorithm of Kosaraju [36] to find a cycle with zero effect. While this requires some non-trivial machinery like VASS and linear programming [36], our algorithm is comparably simple and tailored to the problem. Algorithmically, we rely on a fixed-point iteration which is easier to implement into model-checking tools and has an improved time-complexity. Moreover, our

abstraction via paths and corresponding normalforms might be of independent interest in (liveness) verification for further parameterized systems. Our fixed-point iteration is inspired by the algorithm for the synchronization problem given in [28]. One can consider our algorithm as a generalization of this result.

Broadcast networks are related to the leader-contributor model. It has a fixed leader and an arbitrary number of identical contributors communicating via shared memory. The model was introduced in [29]. When leader and contributors are finite-state automata, the corresponding reachability problem is NP-complete [26]. In [10,14], the authors took a parameterized complexity look at this problem and proved it fixed-parameter tractable. Liveness verification for the model was studied in [22]. The authors show that repeated reachability is NP-complete. The parameterized complexity of the problem was considered in [11]. Networks with shared memory and randomized scheduler were studied in [9]. For a survey of parameterized verification we refer to [2,8].

## 2 Broadcast networks

We introduce the model of broadcast networks of interest in this paper. Our presentation avoids an explicit characterization of the communication topology in terms of graphs. A *broadcast network* is a concurrent system consisting of an arbitrary but finite number of identical clients that communicate by passing messages to each other. Formally, it is a pair  $\mathcal{N} = (D, P)$ . The *domain*  $D$  is a finite set of messages that can be used for communication. A message  $a \in D$  can either be sent,  $!a$ , or received,  $?a$ . The set  $Ops(D) = \{!a, ?a \mid a \in D\}$  captures the communication operations a client can perform. For modeling the identical clients, we abstract away the internal behavior and focus on the communication with others via  $Ops(D)$ . With this, the clients are given in the form of a finite state automaton  $P = (Q, I, \delta)$ , where  $Q$  is a finite set of states,  $I \subseteq Q$  is a set of initial states, and  $\delta \subseteq Q \times Ops(D) \times Q$  is the transition relation. We extend  $\delta$  to words in  $Ops(D)^*$  and write  $q \xrightarrow{w} q'$  instead of  $(q, w, q') \in \delta$ .

During a communication phase in  $\mathcal{N}$ , one client sends a message that is received by a number of other clients. This induces a change of the current state in each client participating in the communication. We use *configurations* to display the current states of the clients. A configuration is a tuple  $c = (q_1, \dots, q_k) \in Q^k, k \in \mathbb{N}$ . We use  $Set(c)$  to denote the set of client states occurring in  $c$ . To access the components of  $c$ , we use  $c[i] = q_i$ . As the number of clients in the system is arbitrary but fixed, we define the set of all configurations to be  $CF = \bigcup_{k \in \mathbb{N}} Q^k$ . The set of *initial configurations* is given by  $CF_0 = \bigcup_{k \in \mathbb{N}} I^k$ . The communication is modeled by a transition relation among configurations. Let  $c' = (q'_1, \dots, q'_k)$  be another configuration with  $k$  clients and  $a \in D$  a message. We have a transition  $c \xrightarrow{a}_{\mathcal{N}} c'$  if the following conditions hold: (1) there is a sender, an  $i \in [1..k]$  such that  $q_i \xrightarrow{!a} q'_i$ , (2) there is a number of receivers, a set  $R \subseteq [1..k] \setminus \{i\}$  such that  $q_j \xrightarrow{?a} q'_j$  for each  $j \in R$ , and (3) all other clients stay idle, for all  $j \notin R \cup \{i\}$  we have  $q_j = q'_j$ . We use  $idx(c \xrightarrow{a}_{\mathcal{N}} c') = R \cup \{i\}$  to denote the indices of clients that contributed to the transition. Note that it may well

be the case that  $R = \emptyset$ . This means that all clients, except from the sender, ignore the sent message. For  $i \in \text{idx}(c \xrightarrow{a} \mathcal{N} c')$ , we let  $\text{pr}_i(c \xrightarrow{a} \mathcal{N} c') = c[i] \xrightarrow{!a/?a} c'[i]$  be the contribution of client  $i$ . If  $i \notin \text{idx}(c \xrightarrow{a} \mathcal{N} c')$ , we set  $\text{pr}_i(c \xrightarrow{a} \mathcal{N} c') = \epsilon$ .

We extend the transition relation to words  $w \in D^*$  and write  $c \xrightarrow{w} \mathcal{N} c'$ . Such a sequence of consecutive transitions is called a *computation* of  $\mathcal{N}$ . Note that all configurations appearing in a computation have the same number of clients. We write  $c \xrightarrow{*} \mathcal{N} c'$  if there is a word  $w \in D^*$  with  $c \xrightarrow{w} \mathcal{N} c'$ . If  $|w| \geq 1$ , we also use  $c \xrightarrow{+} \mathcal{N} c'$ . Where appropriate, we skip  $\mathcal{N}$  in the index. The definition of  $\text{idx}$  can be extended to computations by combining all clients that contribute to one of the transitions. We also extend the definition of  $\text{pr}_i$  to computations by appending the individual contributions to transitions.

We are interested in infinite computations, sequences  $\pi = c_0 \rightarrow c_1 \rightarrow \dots$  of infinitely many consecutive transitions. Such a computation is called *initialized*, if  $c_0 \in CF_0$ . We use  $\text{Inf}(\pi) = \{i \in \mathbb{N} \mid \exists^\infty j : i \in \text{idx}(c_j \rightarrow c_{j+1})\}$  to denote the set of clients that participate in the computation infinitely often. Let  $F \subseteq Q$  be a set of final states. Then we let  $\text{Fin}(\pi) = \{i \in \mathbb{N} \mid \exists^\infty j : c_j[i] \in F\}$  represent the set of clients that visit final states infinitely often.

### 3 Liveness

We consider the liveness verification problem for broadcast networks. Given a broadcast network  $\mathcal{N} = (D, P)$  with  $P = (Q, I, \delta)$  and a set of final states  $F \subseteq Q$ , the problem asks whether there is an infinite initialized computation  $\pi$  in which at least one client visits a state from  $F$  infinitely often,  $\text{Fin}(\pi) \neq \emptyset$ .

*Liveness Verification*

**Input:** A broadcast network  $\mathcal{N} = (D, P)$  and final states  $F \subseteq Q$ .

**Question:** Is there an initialized computation  $\pi$  with  $\text{Fin}(\pi) \neq \emptyset$ ?

*Liveness Verification* was introduced as *Repeated Coverability* in [17]. The problem is known to be P-hard, due to [16]. Our main contribution is a matching polynomial-time algorithm which allows for deducing completeness.

**Theorem 1** *Liveness Verification is P-complete.*

Key to our algorithm is the following lemma which relates the existence of an infinite computation to the existence of a finite one.

**Lemma 1** *There is a computation  $c_0 \rightarrow c_1 \rightarrow \dots$  visiting states in  $F$  infinitely often if and only if there is a computation  $c_0 \xrightarrow{*} c \xrightarrow{+} c$  with  $\text{Set}(c) \cap F \neq \emptyset$ .*

If there is a computation of the form  $c_0 \xrightarrow{*} c \xrightarrow{+} c$  with  $\text{Set}(c) \cap F \neq \emptyset$ , then  $c \xrightarrow{+} c$  can be iterated infinitely often to obtain an infinite computation visiting  $F$  infinitely often. In turn, in any infinite sequence from  $Q^k$  one can find a repeating configuration (pigeon hole principle). This in particular holds for the infinite sequence of configurations containing final states.

Our polynomial-time algorithm for *Liveness Verification* looks for an appropriate reachable configuration  $c$  that can be iterated. The difficulty is that we have a parameterized system, and therefore the number of configurations is not finite. Our approach is to devise a finite graph in which we search for a cycle that mimics the cycle on  $c$ . While the graph yields a decision procedure, it will be of exponential size and a naive search for a cycle requires exponential time. We show in a second step how to find a cycle in polynomial time.

The graph underlying our algorithm is inspired by the powerset construction for the determinization of finite state automata [41]. The vertices keep track of sets of states  $S$  that a client may be in. Different from finite-state automata, however, there is not only one client in a state  $s \in S$  but arbitrarily (but finitely) many. As a consequence, a transition from  $s$  to  $s'$  may have two effects. Some of the clients in  $s$  change their state to  $s'$  while others stay in  $s$ . In that case, the set of states is updated to  $S' = S \cup \{s'\}$ . Alternatively, all clients may change their state to  $s'$ , in which case we get  $S' = (S \setminus \{s\}) \cup \{s'\}$ .

Formally, the graph of interest is  $G = (V, \rightarrow_G)$ . Vertices are tuples of sets of states,  $V = \bigcup_{k \leq |Q|} \mathcal{P}(Q)^k$ . The parameter  $k$  will become clear in a moment. To define the edges, we need some more notation. For  $S \subseteq Q$  and  $a \in D$ , let

$$post_{\gamma_a}(S) = \{r' \in Q \mid \exists r \in S : r \xrightarrow{a} r'\}$$

denote the set of successors of  $S$  under transitions receiving  $a$ . The set of states in  $S$  where receives of  $a$  are *enabled* is denoted by

$$enabled_{\gamma_a}(S) = \{r \in S \mid post_{\gamma_a}(\{r\}) \neq \emptyset\}.$$

There is a directed edge  $V_1 \rightarrow_G V_2$  from vertex  $V_1 = (S_1, \dots, S_k)$  to vertex  $V_2 = (S'_1, \dots, S'_k)$  if the following three conditions are satisfied: (1) there is an index  $j \in [1..k]$ , states  $s \in S_j$  and  $s' \in S'_j$ , and an element  $a$  from the domain  $D$  such that  $s \xrightarrow{a} s'$  is a send transition. (2) For each  $i \in [1..k]$  there are sets of states  $Gen_i \subseteq post_{\gamma_a}(S_i)$  and  $Kill_i \subseteq enabled_{\gamma_a}(S_i)$  such that

$$S'_i = \begin{cases} (S_i \setminus Kill_i) \cup Gen_i, & \text{for } i \neq j, \\ (U_j \setminus Kill_j) \cup Gen_j \cup \{s'\}, & \text{for } i = j \end{cases}$$

where  $U_j$  is either  $S_j$  or  $S_j \setminus \{s\}$ . (3) For each index  $i \in [1..k]$  and state  $q \in Kill_i$ , the intersection  $post_{\gamma_a}(q) \cap Gen_i$  is non-empty.

Intuitively, an edge in the graph mimics a transition in the broadcast network without making explicit the configurations. Condition (1) requires a sender, a component  $j$  capable of sending a message  $a$ . Clients receiving this message are represented by (2). The set  $Gen_i$  consists of those states that are reached by clients performing a corresponding receive transition. These states are added to  $S_i$ . As mentioned above, states can get killed. If, during a receive transition, all clients move to the target state, the original state will not be present anymore. We capture those states in the set  $Kill_i$

and remove them from  $S_j$ . Condition (3) is needed to guarantee that each killed state is replaced by a target state. Note that for component  $j$  we add  $s'$  due to the send transition. Moreover, we need to distinguish whether state  $s$  gets killed or not.

The following lemma relates a cycle in the constructed graph with a cyclic computation of the form  $c \rightarrow^+ c$ . It is crucial for our result.

**Lemma 2** *There is a cycle  $(\{s_1\}, \dots, \{s_m\}) \rightarrow_G^+ (\{s_1\}, \dots, \{s_m\})$  in  $G$  if and only if there is a configuration  $c$  with  $\text{Set}(c) = \{s_1, \dots, s_m\}$  and  $c \rightarrow^+ c$ .*

The lemma explains the restriction of the nodes in the graph to  $k$ -tuples of sets of states, with  $k \leq |Q|$ . We explore the transitions for every possible state in  $c$ , and there are at most  $|Q|$  different states that have to be considered. We have to keep the sets of states separately to make sure that, for every starting state, the corresponding clients perform a cyclic computation.

**Proof** We fix some notations that we use throughout the proof. Let  $c \in Q^n$  be any configuration and  $s \in \text{Set}(c)$ . By  $\text{Pos}_c(s) = \{i \in [1..n] \mid c[i] = s\}$  we denote the positions of  $c$  storing state  $s$ . Given a second configuration  $d \in Q^n$ , we use the set  $\text{Target}_c(s, d) = \{d[i] \mid i \in \text{Pos}_c(s)\}$  to represent those states that occur in  $d$  at the positions  $\text{Pos}_c(s)$ . Intuitively, if there is a sequence of transitions from  $c$  to  $d$ , these are the target states of those positions of  $c$  storing  $s$ .

Consider a computation  $\pi = c \rightarrow^+ c$  with  $\text{Set}(c) = \{s_1, \dots, s_m\}$ . We show that there is a cycle  $(\{s_1\}, \dots, \{s_m\}) \rightarrow_G^+ (\{s_1\}, \dots, \{s_m\})$  in  $G$ . To this end, assume  $\pi$  is of the form  $\pi = c \rightarrow c_1 \rightarrow \dots \rightarrow c_\ell \rightarrow c$ . Since  $c \rightarrow c_1$  is a transition in the broadcast network, there is an edge

$$(\{s_1\}, \dots, \{s_m\}) \rightarrow_G (\text{Target}_c(s_1, c_1), \dots, \text{Target}_c(s_m, c_1))$$

in  $G$  where each state  $s_i$  gets replaced by the set of target states in  $c_1$ . Applying this argument inductively, we get a path in the graph:

$$\begin{aligned} (\{s_1\}, \dots, \{s_m\}) &\rightarrow_G (\text{Target}_c(s_1, c_1), \dots, \text{Target}_c(s_m, c_1)) \\ &\rightarrow_G (\text{Target}_c(s_1, c_2), \dots, \text{Target}_c(s_m, c_2)) \\ &\rightarrow_G \dots \\ &\rightarrow_G (\text{Target}_c(s_1, c), \dots, \text{Target}_c(s_m, c)). \end{aligned}$$

Since  $\text{Target}_c(s_i, c) = \{s_i\}$ , we found the desired cycle.

For the other direction, let a cycle  $\sigma = (\{s_1\}, \dots, \{s_m\}) \rightarrow_G^+ (\{s_1\}, \dots, \{s_m\})$  be given. We construct from  $\sigma$  a computation  $\pi = c \rightarrow^+ c$  in the broadcast network such that  $\text{Set}(c) = \{s_1, \dots, s_m\}$ . The difficulty in constructing  $\pi$  is to ensure that at any point in time there are enough clients in appropriate states. For instance, if a transition  $s \xrightarrow{!a} s'$  occurs, we need to decide on how many clients to move to  $s'$ . Having too few clients in  $s'$  may stall the computation at a later point: there may be a number of sends required that can only be obtained by transitions from  $s'$ . If there are too few clients in  $s'$ , we cannot guarantee the sends. The solution is to start with *enough* clients in any state. With invariants we guarantee that at any point in time, the number of clients in the needed states suffices.



Let cycle  $\sigma$  be  $V_0 \rightarrow_G V_1 \rightarrow_G \dots \rightarrow_G V_\ell$  with  $V_0 = V_\ell = (\{s_1\}, \dots, \{s_m\})$ . Further, let  $V_j = (S_j^1, \dots, S_j^m)$ . We will construct the computation  $\pi$  over configurations in  $Q^n$  where  $n = m \cdot |Q|^\ell$ . The idea is to have  $|Q|^\ell$  clients for each of the  $m$  components of the vertices  $V_i$  occurring in  $\sigma$ . To access the clients belonging to a particular component, we split up configurations in  $Q^n$  into *blocks*, intervals  $I(i) = [(i - 1) \cdot |Q|^\ell + 1 .. i \cdot |Q|^\ell]$  for each  $i \in [1..m]$ . Let  $d \in Q^n$  be arbitrary. For  $i \in [1..m]$ , let  $B_d(i) = \{d[t] \mid t \in I(i)\}$  be the set of states occurring in the  $i$ th block of  $d$ . Moreover, we blockwise collect clients that are currently in a particular state  $s \in Q$ . Let the set  $\text{Pos}_d(i, s) = \{t \in I(i) \mid d[t] = s\}$  be those positions of  $d$  in the  $i$ th block that store state  $s$ .

We fix the configuration  $c \in Q^n$ . For each component  $i \in [1..m]$ , in the  $i$ th block it contains  $|Q|^\ell$  copies of the state  $s_i$ . Formally,  $B_c(i) = \{s_i\}$ . Our goal is to construct the computation  $\pi = c_0 \rightarrow^+ c_1 \rightarrow^+ \dots \rightarrow^+ c_\ell$  with  $c_0 = c_\ell = c$  such that the following two invariants are satisfied. (1) For each  $j \in [0..\ell]$  and  $i \in [1..m]$  we have  $B_{c_j}(i) \subseteq S_j^i$ . (2) For any state  $s$  in a set  $S_j^i$  we have  $|\text{Pos}_{c_j}(i, s)| \geq |Q|^{\ell-j}$ . Intuitively, (1) means that during the computation  $\pi$  we visit at most those states that occur in the cycle  $\sigma$ . Invariant (2) guarantees that at each configuration  $c_j$  there are *enough* clients available in these states.

We construct  $\pi$  inductively. The base case is given by configuration  $c_0 = c$  which satisfies Invariants (1) and (2) by definition. For the induction step, assume  $c_j$  is already constructed such that (1) and (2) hold for the configuration. Our first goal is to construct a configuration  $d$  such that  $c_j \rightarrow^+ d$  and  $d$  satisfies Invariant (2). In a second step we construct a computation  $d \rightarrow^* c_{j+1}$ .

In the cycle  $\sigma$  there is an edge  $V_j \rightarrow_G V_{j+1}$ . From the definition of  $\rightarrow_G$  we get a component  $t \in [1..m]$ , states  $s \in S_j^t$  and  $s' \in S_{j+1}^t$ , and an  $a \in D$  such that there is a send transition  $s \xrightarrow{!a} s'$ . Moreover, there are sets  $\text{Gen}_t \subseteq \text{post}_{\gamma_a}(S_j^t)$  and  $\text{Kill}_t \subseteq \text{enabled}_{\gamma_a}(S_j^t)$  such that the following holds:

$$S_{j+1}^t = (U_t \setminus \text{Kill}_t) \cup \text{Gen}_t \cup \{s'\}.$$

Here,  $U_t$  is either  $S_j^t$  or  $S_j^t \setminus \{s\}$ . We focus on  $t$  and take care of other components later. We apply a case distinction for the states in  $S_{j+1}^t$ .

Let  $q$  be a state in  $S_{j+1}^t \setminus \{s'\}$ . If  $q \in \text{Gen}_t$ , there exists a  $p \in S_j^t$  such that  $p \xrightarrow{?a} q$ . We apply this transition to  $|Q|^{\ell-(j+1)}$  many clients in the  $t$ th block of configuration  $c_j$ . If  $q \in U_t \setminus \text{Kill}_t$  and  $q$  not in  $\text{Gen}_t$ , then certainly  $q \in U_t \subseteq S_j^t$ . In this case, we let  $|Q|^{\ell-(j+1)}$  many clients of block  $t$  stay idle in state  $q$ . For state  $s'$ , we apply a sequence of sends. More precisely, we apply the transition  $s \xrightarrow{!a} s'$  to  $|Q|^{\ell-(j+1)}$  many clients in block  $t$  of  $c_j$ . The first of these sends synchronizes with the previously described receive transitions. The other sends do not have any receivers. For components different from  $t$ , we apply the same procedure. Since there are only receive transitions, we also let them synchronize with the first send of  $a$ . This leads to a computation  $\tau$

$$c_j \xrightarrow{a} d^1 \xrightarrow{a} d^2 \xrightarrow{a} \dots \xrightarrow{a} d^{|Q|^{\ell-(j+1)}} = d.$$



We argue that  $\tau$  is *valid*: there are enough clients in  $c_j$  such that  $\tau$  can be performed. Focus on component  $t$ , the reasoning for other components is similar. Let  $p \in \text{Set}(c_j) = S_j^t$ . Note that the equality is due to Invariants (1) and (2). We count the clients of  $c_j$  in state  $p$  (in block  $t$ ) needed to perform  $\tau$ . We need

$$|Q|^{\ell-(j+1)} \cdot |\text{post}_{\gamma_a}(p) \cup \{p, s'\}| \leq |Q|^{\ell-(j+1)} \cdot |Q| = |Q|^{\ell-j}$$

of these clients. The set  $\text{post}_{\gamma_a}(p) \cup \{p, s'\}$  appears as a consequence of the case distinction above: there may be transitions mapping  $p$  to a state in  $\text{post}_{\gamma_a}(p)$ , it may happen that clients stay idle in  $p$ , and in the case  $p = s$ , we need to add  $s'$  for the send transition. Since  $|\text{Pos}_{c_j}(t, p)| \geq |Q|^{\ell-j}$  by Invariant (2), we get that  $\tau$  is a valid computation. Moreover, note that configuration  $d$  satisfies Invariant (2) for  $j + 1$ : for each state  $q \in S_{j+1}^t$ , the computation  $\tau$  was constructed such that  $|\text{Pos}_d(t, q)| \geq |Q|^{\ell-(j+1)}$ .

To satisfy Invariant (1), we need to erase states that are present in  $d$  but not in  $S_{j+1}^t$ . To this end, we reconsider the set  $\text{Kill}_t \subseteq \text{enabled}_{\gamma_a}(S_j^t)$ . For each state  $p \in \text{Kill}_t$ , we know by the definition of  $\rightarrow_G$  that  $\text{post}_{\gamma_a}(p) \cap \text{Gen}_t \neq \emptyset$ . Hence, there is a  $q \in S_{j+1}^t$  such that  $p \xrightarrow{\gamma_a} q$ . We apply this transition to all clients in  $d$  currently in state  $p$  that were not active in the computation  $\tau$ . In case  $U_t = S_j^t \setminus \{s\}$ , we apply the send  $s \xrightarrow{1a} s'$  to all clients that are still in  $s$  and were not active in  $\tau$ . Altogether, this leads to a computation  $\eta = d \rightarrow^* c_{j+1}$ .

There is a subtlety in the definition of  $\eta$ . There may be no send transition for the receivers to synchronize with since  $s$  may not need to be erased. In this case, we synchronize the receive transitions of  $\eta$  with the last send of  $\tau$ .

Computation  $\eta$  substitutes the states in  $\text{Kill}_t$  and state  $s$ , depending on  $U_t$ , by states in  $S_{j+1}^t$ . But this means that in the  $t$ th block of  $c_{j+1}$ , there are only states of  $S_{j+1}^t$  left. Hence,  $B_{c_{j+1}}(t) \subseteq S_{j+1}^t$ , and Invariant (1) holds.

After the construction of  $\pi = c \rightarrow^+ c_\ell$ , it is left to argue that  $c_\ell = c$ . But this is due to the fact that (1) holds for  $c_\ell$  and  $S_\ell^t = (\{s_1\}, \dots, \{s_m\})$ . □

The graph  $G$  is of exponential size. To obtain a polynomial-time procedure, we cannot just search it for a cycle as required by Lemma 2. Instead, we now show that if such a cycle exists, then there is a cycle in a certain normal form. Hence, it suffices to look for a normal-form cycle. As we will show, this can be done in polynomial time. We define the normal form more generally for paths.

A path is in *normal form*, if it takes the shape  $V_1 \xrightarrow{*}_G V_m \xrightarrow{*}_G V_n$  such that the following conditions hold. In the prefix  $V_1 \xrightarrow{*}_G V_m$  the sets of states increase monotonically,  $V_i \sqsubseteq V_{i+1}$  for all  $i \in [1..m - 1]$ . Here,  $\sqsubseteq$  denotes the componentwise inclusion. In the suffix  $V_m \xrightarrow{*}_G V_n$ , the sets of states decrease monotonically,  $V_i \supseteq V_{i+1}$  for all  $i \in [m..n - 1]$ . The following lemma states that if there is a path in the graph, then there is also a path in normal form. The intuition is that the variants of the transitions that decrease the sets of states can be postponed towards the end of the computation.

**Lemma 3** *There is a path from  $V_1$  to  $V_2$  in  $G$  if and only if there is a path in normal form from  $V_1$  to  $V_2$ .*

**Proof** If  $V_1 \rightarrow_G^* V_2$  is a path in normal form, there is nothing to prove. For the other direction, let  $\sigma = V_1 \rightarrow_G^* V_2$  be an arbitrary path. To get a path in normal form, we first simulate the edges of  $\sigma$  in such a way that no states are deleted. In a second step, we erase the states that should have been deleted. We have to respect a particular deletion order to obtain a valid path.

Let  $\sigma = U_1 \rightarrow_G U_2 \rightarrow_G \dots \rightarrow_G U_\ell$  with  $U_1 = V_1$  and  $U_\ell = V_2$ . We inductively construct an increasing path  $\sigma_{\text{inc}} = U'_1 \rightarrow_G \dots \rightarrow_G U'_\ell$  with  $U'_j \supseteq U_j$  for all  $i \leq j$ . For the base case, we set  $U'_1 = U_1$ . Now assume  $\sigma_{\text{inc}}$  has already been constructed up to vertex  $U'_j$ . There is an edge  $e = U_j \rightarrow_G U_{j+1}$  in  $\sigma$ . Since  $U'_j \supseteq U_j$ , we can simulate  $e$  on  $U'_j$ : all states needed for execution are present in  $U'_j$ . Moreover, we mimic  $e$  such that no state gets deleted. This is achieved by setting the corresponding Kill-sets to be empty. Hence, we get  $U'_j \rightarrow U'_{j+1}$  with  $U'_{j+1} \supseteq U'_j$  (no deletion) and  $U'_{j+1} \supseteq U_{j+1}$  (simulation of  $e$ ).

The states in  $V'_2 = U'_\ell$  that are not in  $V_2$  are those states that were deleted along  $\sigma$ . We construct a decreasing path  $\sigma_{\text{dec}} = V'_2 \rightarrow_G^* V_2$ , deleting all these states. To this end, let  $V'_2 = (T_1, \dots, T_m)$  and  $V_2 = (S_1, \dots, S_m)$ . An edge in  $\sigma$  deletes sets of states in each component  $i \in [1..m]$ . Hence, to mimic the deletion, we need to consider subsets of  $Del = \bigcup_{i \in [1..m]} (T_i \setminus S_i) \times \{i\}$ . Note that the index  $i$  in a tuple  $(s, i)$  displays the component the state  $s$  is in.

Consider the equivalence relation  $\sim$  over  $Del$  defined by  $(x, i) \sim (y, t)$  if and only if the last occurrence of  $x$  in component  $i$  and  $y$  in component  $t$  in the path  $\sigma$  coincide. Intuitively, two elements are equivalent if they get deleted at the same time and do not appear again in  $\sigma$ . We introduce an order on the equivalence classes:  $[(x, i)]_\sim < [(y, t)]_\sim$  if and only if the last occurrence of  $(x, i)$  was before the last occurrence of  $(y, t)$ . Since the order is total, we get a partition of  $Del$  into equivalence classes  $P_1, \dots, P_n$  such that  $P_j < P_{j+1}$ .

We construct  $\sigma_{\text{dec}} = K_0 \rightarrow_G \dots \rightarrow_G K_n$  with  $K_0 = V'_2$  and  $K_n = V_2$  as follows. During each edge  $K_{j-1} \rightarrow_G K_j$ , we delete precisely the elements in  $P_j$  and do not add further states. Deleting  $P_j$  is due to an edge  $e = U_k \rightarrow_G U_{k+1}$  of  $\sigma$ . We mimic  $e$  in such a way that no state gets added and set the corresponding Gen sets to the empty set. Since we respect the order  $<$  with the deletions, the simulation of  $e$  is possible. Suppose, we need a state  $s$  in component  $t$  to simulate  $e$  but the state is not available in component  $t$  of  $K_{j-1}$ . Then it was deleted before,  $(s, t) \in P_1 \cup \dots \cup P_{j-1}$ . But this contradicts that  $s$  is present in  $U_k$ . Hence, all the needed states are available. Since after the last edge of  $\sigma_{\text{dec}}$  we have deleted all elements of  $Del$ , we get  $K_n = V_2$ .  $\square$

Using the normal-form result in Lemma 3, we now give a polynomial-time algorithm to check whether  $(\{s_1\}, \dots, \{s_m\}) \rightarrow_G^+ (\{s_1\}, \dots, \{s_m\})$ . The idea is to mimic the monotonically increasing prefix of the path by a suitable post operator, the monotonically decreasing suffix by a suitable pre operator, and intersect the two. The difficulty in computing an appropriate post operator is to ensure that the receive operations are enabled by sends leading to a state in the intersection, and similar for pre. The solution is to use a greatest fixed-point computation. In a first Kleene iteration step, we

determine the ordinary  $post^+$  of  $(\{s_1\}, \dots, \{s_m\})$  and intersect it with the  $pre^*$ . In the next step, we constrain the  $post^+$  and the  $pre^*$  computations to visiting only states in the previous intersection. The results are intersected again, which may remove further states. Hence, the computation is repeated relative to the new intersection. The thing to note is that we do not work with standard post and pre operators but with operators that are constrained by (tuples of) sets of states.

For the definition of the operators, consider  $C = (C_1, \dots, C_m) \in \mathcal{P}(Q)^m$  for an  $m \leq |Q|$ . Given a sequence of sets of states  $X_1, \dots, X_m$ , we define  $post_C(X_1, \dots, X_m) = (X'_1, \dots, X'_m)$  with

$$X'_i = \{s' \in Q \mid \exists s \in X_i : s \xrightarrow{!a} P \downarrow_{C_i} s'\} \\ \cup \{s' \in Q \mid \exists s_1, s_2 \in X_\ell : \exists s \in X_i : s_1 \xrightarrow{!a} P \downarrow_{C_\ell} s_2 \wedge s \xrightarrow{?a} P \downarrow_{C_i} s'\}.$$

Here,  $P \downarrow_{C_i}$  denotes the automaton obtained from  $P$  by restricting it to the states  $C_i$ . Similarly, we define  $pre_C(X_1, \dots, X_m) = (X'_1, \dots, X'_m)$  with

$$X'_i = \{s \in Q \mid \exists s' \in X_i : s \xrightarrow{!a} P \downarrow_{C_i} s'\} \\ \cup \{s \in Q \mid \exists s_1, s_2 \in X_\ell : \exists s' \in X_i : s_1 \xrightarrow{!a} P \downarrow_{C_\ell} s_2 \wedge s \xrightarrow{?a} P \downarrow_{C_i} s'\}.$$

The next lemma shows that the (reflexive) transitive closures of these operators can be computed in polynomial time.

**Lemma 4** *The closures  $post^+_C(X_1, \dots, X_m)$  and  $pre^*_C(X_1, \dots, X_m)$  can be computed in polynomial time.*

**Proof** Both closures can be computed by a saturation. For  $post^+_C(X_1, \dots, X_m)$ , we keep sets  $R_1, \dots, R_m$ , each being the post of a component. Initially, we set  $R_i = X_i$ . The defining equation of  $X'_i$  in  $post^+_C(X_1, \dots, X_m)$  gives the saturation. One substitutes  $X_i$  by  $R_i$  and  $X_\ell$  by  $R_\ell$  on the right side. The resulting set of states is added to  $R_i$ . This process is applied to each component and repeated until the  $R_i$  do not change anymore, the fixed point is reached.

The saturation terminates in polynomial time. After updating  $R_i$  in each component, we either already terminated or added at least one new state to a set  $R_i$ . Since there are  $m \leq |Q|$  of these sets and each one is a subset of  $Q$ , we need to update the sets  $R_i$  at most  $|Q|^2$  many times. For a single of these updates, the dominant time factor comes from finding appropriate send and receive transitions. This can be achieved in  $\mathcal{O}(|\delta|^2)$  time.

Computing the closure  $pre^*_C(X_1, \dots, X_m)$  is similar. One can apply the above saturation and only needs to reverse the transitions in the client.  $\square$

As argued above, the existence of a cycle reduces to finding a fixed point. The following lemma shows that it can be computed efficiently.

**Lemma 5** *There is a cycle  $(\{s_1\}, \dots, \{s_m\}) \rightarrow_G^+ (\{s_1\}, \dots, \{s_m\})$  if and only if there is a non-trivial solution to the following equation:*

$$C = \text{post}_C^+(\{s_1\}, \dots, \{s_m\}) \cap \text{pre}_C^*(\{s_1\}, \dots, \{s_m\}) .$$

*Such a solution can be found in polynomial time.*

**Proof** We use a Kleene iteration to compute the greatest fixed point. It invokes Lemma 4 as a subroutine. Every step of the Kleene iteration reduces the number of states in  $C$  by at least one, and initially there are at most  $|Q|$  entries with  $|Q|$  states each. Hence, we terminate after quadratically many iterations.

It is left to prove correctness. Let  $(\{s_1\}, \dots, \{s_m\}) \rightarrow_G^+ (\{s_1\}, \dots, \{s_m\})$  be a cycle in  $G$ . By Lemma 3 we can assume it to be in normal form. Let  $(\{s_1\}, \dots, \{s_m\}) \rightarrow_G^+ C$  be the increasing part and  $C \rightarrow_G^* (\{s_1\}, \dots, \{s_m\})$  the decreasing part. Then,  $C$  is a solution to the equation.

For the other direction, let a solution  $C$  be given. Since  $C$  is contained in  $\text{post}_C^+(\{s_1\}, \dots, \{s_m\})$  we can construct a monotonically increasing path  $(\{s_1\}, \dots, \{s_m\}) \rightarrow_G^+ C$ . Similarly, since  $C \subseteq \text{pre}_C^*(\{s_1\}, \dots, \{s_m\})$ , we get a decreasing path  $C \rightarrow_G^* (\{s_1\}, \dots, \{s_m\})$ . Hence, we get the desired cycle.  $\square$

It is yet open on which states  $s_1$  to  $s_m$  to perform the search for a cycle. After all, we need that the corresponding configuration is reachable from an initial configuration. The idea is to use the set of all states reachable from an initial state. Note that there is a live computation if and only if there is one involving all those states. Indeed, if a state is not active during the cycle, the corresponding clients will stop moving after an initial phase. Since the states reachable from an initial state can be computed in polynomial time [16], the proof of Theorem 1 follows. We summarize the algorithm for liveness verification in Algorithm 1.

---

**Algorithm 1** Liveness Verification

---

**Input:** A broadcast network  $\mathcal{N} = (D, P)$  and final states  $F \subseteq Q$ .

**Output:** *True*, if there is an initialized computation  $\pi$  with  $\text{Fin}(\pi) \neq \emptyset$ . *False* otherwise.

- 1: Use the algorithm of [16] to compute the set of reachable states  $\{s_1, \dots, s_m\}$ .
  - 2: **if**  $\{s_1, \dots, s_m\} \cap F = \emptyset$  **then**
  - 3:     **return** *false*. // This check is needed according to Lemma 1.
  - 4: **end if**
  
  - 5: Set  $S = (\{s_1\}, \dots, \{s_m\})$ ,  $C = Q^m$ , and  $C' = \emptyset$ . //  $C$  is the largest element in  $\mathcal{P}(Q)^m$ .
  - 6: **while**  $C \neq C'$  **do** // Computation of the greatest fixed point.
  - 7:     Set  $C' = C$ .
  - 8:     Compute  $C = \text{post}_{C'}^+(S) \cap \text{pre}_{C'}^*(S)$  with the algorithm from Lemma 4.
  - 9:     **if**  $C = \emptyset$  **then**
  - 10:         **return** *false*. // Exclude trivial solution to the equation of Lemma 5.
  - 11:     **end if**
  - 12: **end while**
  
  - 13: **return** *true*. // Largest fixed point is non-trivial. Computation  $\pi$  exists due to Lemma 5.
-

*Liveness Verification* does not take fairness into account. A client may contribute to the live computation (and help the distinguished client reach a final state) without ever making progress towards its own final state.

### 4 Fair liveness

We study the fair liveness verification problem that strengthens the requirement on the computation. Given a broadcast network  $\mathcal{N} = (D, P)$  with clients  $P = (Q, I, \delta)$  and a set of final states  $F \subseteq Q$ , the problem asks whether there is an infinite initialized computation  $\pi$  in which clients that send or receive infinitely often also visit their final states infinitely often,  $\text{Inf}(\pi) \subseteq \text{Fin}(\pi)$ . Computations satisfying the requirement are called *fair*, a notion that dates back to [40] where it was introduced as compassion or strong fairness.

*Fair Liveness Verification*

**Input:** A broadcast network  $\mathcal{N} = (D, P)$  and final states  $F \subseteq Q$ .

**Question:** Is there an initialized computation  $\pi$  with  $\text{Inf}(\pi) \subseteq \text{Fin}(\pi)$ ?

We solve the problem by applying the algorithm from Sect. 3 to an instrumentation of the given network. Formally, an instance  $(\mathcal{N}, F)$  of *Fair Liveness Verification*, is transformed into a new network  $\mathcal{N}_F$ , containing copies of  $Q$ , where  $Q$  are the client states in  $\mathcal{N}$ . The construction ensures that cycles over  $Q$  in  $\mathcal{N}_F$  correspond to cycles in  $\mathcal{N}$  where each participating client sees a final state. Such cycles make up a fair computation. Our main result is:

**Theorem 2** *Fair Liveness Verification is P-complete.*

Hardness follows from [16]. To explain the instrumentation, we need the notion of a good computation. A computation  $c_1 \rightarrow c_2 \rightarrow \dots \rightarrow c_n$  with  $n > 1$  is called *good for F*, denoted  $c_1 \Rightarrow_F c_n$ , if every client that makes a move during the computation sees a final state. Formally, if  $i \in \text{idx}(c_1 \rightarrow^+ c_n)$  then there exists a  $k \in [1..n]$  such that  $c_k[i] \in F$ . The following strengthens Lemma 1.

**Lemma 6** *There is a fair computation from  $c_0$  if and only if  $c_0 \rightarrow^* c \Rightarrow_F c$ .*

**Proof** If there is a computation of the form  $c_0 \rightarrow^* c \Rightarrow_F c$ , then the good cycle  $c \Rightarrow_F c$  can be iterated infinitely often to obtain a fair computation.

For the other direction, let a fair computation  $\pi$  starting in  $c_0$  be given. Since the configurations visited by  $\pi$  are over  $Q^k$  for some  $k$ , there is a configuration  $c$  that repeats infinitely often in  $\pi$ . Hence, we obtain a prefix  $c_0 \rightarrow^* c$ .

Let  $\text{Inf}(\pi) = \{i_1, \dots, i_n\}$  be the clients that participate infinitely often in computation  $\pi$ . By the definition of a fair computation, we have  $\text{Inf}(\pi) \subseteq \text{Fin}(\pi)$ . This means that each of the clients in  $\{i_1, \dots, i_n\}$  visits a state from  $F$  infinitely often along  $\pi$ . Hence, for each  $j \in [1..n]$  we can find a subcomputation  $\pi_j = c \rightarrow^+ c$  of  $\pi$ , in which client  $i_j$  visits a state from  $F$ . Combining all  $\pi_j$  yields desired good cycle  $c \Rightarrow_F c$ .  $\square$

The broadcast network  $\mathcal{N}_F$  is designed to detect good cycles  $c \Rightarrow_F c$ . The idea is to let the clients compute in phases. The original state space  $Q$  is the first phase. As

soon as a client participates in the computation, it moves to a second phase given by a copy  $\hat{Q}$  of  $Q$ . From this copy it enters a third phase  $\tilde{Q}$  upon seeing a final state. From  $\tilde{Q}$  it may return to  $Q$ .

Let the given broadcast network be  $\mathcal{N} = (D, P)$  with  $P = (Q, I, \delta)$ . We define  $\mathcal{N}_F = (D \cup \{n\}, P_F)$  with fresh symbol  $n \notin D$  and extended client

$$P_F = (\bar{Q}, \bar{I}, \bar{\delta}) \quad \text{where} \quad \bar{Q} = Q \cup \hat{Q} \cup \tilde{Q}.$$

Note that the initial phase is  $\tilde{Q}$ . For every transition  $(q, op, q') \in \delta$ , we have  $(q, op, \hat{q}'), (\hat{q}, op, \hat{q}'), (\tilde{q}, op, \tilde{q}') \in \bar{\delta}$ . For every final state  $q \in F$  we have  $(\hat{q}, !n, \tilde{q}) \in \bar{\delta}$ . Finally, for every state  $q \in Q$  we have  $(\tilde{q}, !n, q) \in \bar{\delta}$ . Configuration  $c$  admits a good cycle if and only if there is a cycle at  $c$  in the instrumented network. An initial prefix can be mimicked by computations within  $\tilde{Q}$ .

**Lemma 7**  $c_0 \rightarrow^* c \Rightarrow_F c$  in  $\mathcal{N}$  if and only if  $\tilde{c}_0 \rightarrow^* c \rightarrow^+ c$  in  $\mathcal{N}_F$ .

**Proof (Idea)** The reasoning for prefixes is simple. By adding or removing  $!n$ -transitions leading from  $\tilde{Q}$  to  $Q$ , we can turn a prefix  $c_0 \rightarrow^* c$  of  $\mathcal{N}$  to  $\tilde{c}_0 \rightarrow^* c$  of  $\mathcal{N}_F$  and vice versa. Let a cycle  $c \rightarrow^+ c$  in  $\mathcal{N}_F$  be given. Note that in  $c$ , all clients are in states from  $Q$ . As soon as a client participates in the cycle, it will immediately move from  $Q$  to  $\hat{Q}$ . To return to  $Q$ , via  $\tilde{Q}$ , it needs to see a state from  $F$ , resulting in a good cycle  $c \Rightarrow_F c$  in  $\mathcal{N}$ .

For the other direction, let  $c \Rightarrow_F c$  be a good cycle in  $\mathcal{N}$ . If a client participates in it, we simulate its computation on  $\mathcal{N}_F$ . Assume it starts in  $q \in Q$ . Upon its first transition, it moves from  $Q$  to  $\hat{Q}$ . It stays within the states of  $\hat{Q}$  until it sees a state from  $F$ . Note this definitely happens since the assumed cycle is good. After visiting  $F$ , the client moves to  $\tilde{Q}$  via sending  $n$ . It continues its computation in  $\tilde{Q}$  until it reaches  $\tilde{q}$ . Then, it moves to  $q$  by taking an  $!n$ -transition. Altogether, this constitutes a cycle  $c \rightarrow^+ c$  in  $\mathcal{N}_F$ . □

For the proof of Theorem 2, it is left to state the algorithm for finding a computation  $\tilde{c}_0 \rightarrow^* c \rightarrow^+ c$  in  $\mathcal{N}_F$ . We compute the states reachable from an initial state in  $\mathcal{N}_F$ . As we are interested in a configuration  $c$  over  $Q$ , we intersect this set with  $Q$ . Both steps can be done in polynomial time. Let  $s_1, \dots, s_m$  be the states in the intersection. To these we apply the fixed-point iteration from Lemma 5. By Lemma 2, the iteration witnesses the existence of a cycle over a configuration  $c$  of  $\mathcal{N}_F$  that involves only the states  $s_1$  up to  $s_m$ .

### 5 Model checking broadcast networks

We consider model checking for broadcast networks against linear time specifications. Given a specification, described in linear time temporal logic (LTL) [39], we test whether all (infinite) computations of a broadcast network satisfy the specification. Like for liveness verification, we consider two variants of the problem. These differ in when a computation of a broadcast network satisfies a specification. *Fair Model Checking* requires that the individual computations of all clients participating infinitely

often satisfy the specification. *Sparse Model Checking* asks for at least one client that participates infinitely often and satisfies the specification. We show how to solve both problems by reducing them to *Liveness Verification* and *Fair Liveness Verification* via incorporating the Vardi–Wolper construction [44]. But before we consider the problems, we recall syntax and semantics of LTL. Let  $\mathcal{A}$  be a set of atomic propositions. An LTL formula  $\varphi$  over  $\mathcal{A}$  is defined as follows:

$$\varphi ::= p \in \mathcal{A} \mid \neg\varphi \mid \varphi \vee \varphi \mid \mathcal{X}\varphi \mid \varphi \mathcal{U}\varphi.$$

An LTL formula consists of propositions and combinations of the same via negation, union, *next operator*  $\mathcal{X}$ , and *until operator*  $\mathcal{U}$ . Semantics is defined in terms of the *satisfaction relation*  $\models$ . It describes when a formula is fulfilled. We define it inductively along the structure of LTL formulas. Let  $w \in \mathcal{P}(\mathcal{A})^\omega$  be an infinite word consisting of sets of propositions and let  $i \in \mathbb{N}$ . We have

$$\begin{aligned} w, i \models p & \quad \text{iff } p \in w(i), \\ w, i \models \neg\varphi & \quad \text{iff } w, i \not\models \varphi, \\ w, i \models \varphi_1 \vee \varphi_2 & \quad \text{iff } w, i \models \varphi_1 \text{ or } w, i \models \varphi_2, \\ w, i \models \mathcal{X}\varphi & \quad \text{iff } w, i + 1 \models \varphi, \\ w, i \models \varphi_1 \mathcal{U}\varphi_2 & \quad \text{iff } \exists j \geq i : w, j \models \varphi_2 \text{ and } \forall k \in [i..j - 1] : w, k \models \varphi_1, \\ w \models \varphi & \quad \text{iff } w, 0 \models \varphi. \end{aligned}$$

For model checking broadcast networks against LTL specifications, we need to define when a network computation satisfies a formula. Let  $\mathcal{N} = (D, P)$  be a network with clients  $P = (Q, I, \delta)$  and let  $\varphi$  be an LTL formula. Moreover, let  $\lambda : Q \rightarrow \mathcal{P}(\mathcal{A})$  be a map associating to each state of  $P$  a set of atomic proposition. A computation  $\sigma = q_0 \rightarrow q_1 \rightarrow \dots$  of  $P$  with  $q_0 \in I$  *satisfies* formula  $\varphi$  if  $\lambda(\sigma) \models \varphi$ . Here,  $\lambda(\sigma) = \lambda(q_0).\lambda(q_1) \dots \in \mathcal{P}(\mathcal{A})^\omega$  is the infinite word obtained by applying  $\lambda$  to the whole computation. We also write  $\sigma \models \varphi$ .

For network computations, we consider two notions of satisfaction for LTL formulas. Like for liveness, one of the notions incorporates a fairness assumption, the other one focuses on a single client. An initialized infinite computation  $\pi$  of  $\mathcal{N}$  *satisfies*  $\varphi$  *under fairness* if each client  $i \in \text{Inf}(\pi)$  satisfies  $\varphi$  via its contribution,  $pr_i(\pi) \models \varphi$ . We write  $\pi \models_{\text{fair}} \varphi$ . Hence, satisfaction under fairness ensures that each client participating infinitely often satisfies the specification. An initialized infinite computation  $\pi$  of  $\mathcal{N}$  *sparingly satisfies*  $\varphi$  if there is a client  $i \in \text{Inf}(\pi)$  that satisfies  $\varphi$ ,  $pr_i(\pi) \models \varphi$ . We denote it by  $\pi \models_{\text{sparse}} \varphi$ . In contrast to fairness, a single client satisfying the specification suffices.

The two notions of satisfaction yield two decision problems: *Fair Model Checking* and *Sparse Model Checking*. We state both in the subsequent sections and develop algorithms based on the results from Sects. 3 and 4.



## 5.1 Fair model checking

We want to test whether all computations of a broadcast network satisfy a given LTL formula under the fairness assumption. Formally, a broadcast network  $\mathcal{N}$  is said to *satisfy* an LTL formula  $\varphi$  *under fairness*, written  $\mathcal{N} \models_{\text{fair}} \varphi$ , if for each initialized infinite computation  $\pi$  we have that  $\pi \models_{\text{fair}} \varphi$ . With this notion at hand, we can formalize the corresponding decision problem *Fair Model Checking*. Note that we do not explicitly mention the map  $\lambda$  as part of the input. We assume it to be given with the broadcast network.

### *Fair Model Checking*

**Input:** A broadcast network  $\mathcal{N} = (D, P)$  and an LTL formula  $\varphi$ .

**Question:** Does  $\mathcal{N} \models_{\text{fair}} \varphi$  hold?

Our goal is to prove the following theorem by presenting an algorithm for *Fair Model Checking*. Note that the exponential factor in the time estimation only depends on the size of the formula. The *size*  $|\mathcal{N}| = \max\{|D|, |Q|\}$  of the broadcast network only contributes a polynomial factor.

**Theorem 3** *Fair Model Checking can be solved in time  $2^{\mathcal{O}(|\varphi|)} \cdot |\mathcal{N}|^{\mathcal{O}(1)}$ .*

We need to develop an algorithm for checking  $\mathcal{N} \models_{\text{fair}} \varphi$ . A direct iteration over all computations  $\pi$  of  $\mathcal{N}$  along with a test whether  $\pi \models_{\text{fair}} \varphi$  is not tractable since there are infinitely many candidates for  $\pi$ . We rather search for a computation that violates  $\varphi$ . The non-existence of such a computation ensures that  $\mathcal{N}$  satisfies the formula. Phrased differently, we have  $\mathcal{N} \not\models_{\text{fair}} \varphi$  if and only if there is a computation  $\pi$  of  $\mathcal{N}$  with  $\pi \not\models_{\text{fair}} \varphi$ .

The latter can be reformulated. Let  $\pi$  be a computation with  $\pi \not\models_{\text{fair}} \varphi$ . Then there is a client  $i \in \text{Inf}(\pi)$  that violates  $\varphi$ , we have  $pr_i(\pi) \models \neg\varphi$ . But by definition this means  $\pi \models_{\text{sparse}} \neg\varphi$ . Hence, we obtain that  $\pi \not\models_{\text{fair}} \varphi$  if and only if  $\pi \models_{\text{sparse}} \neg\varphi$ . The following lemma summarizes the reasoning.

**Lemma 8** *We have  $\mathcal{N} \not\models_{\text{fair}} \varphi$  if and only if there is an initialized infinite computation  $\pi$  of  $\mathcal{N}$  such that  $\pi \models_{\text{sparse}} \neg\varphi$ .*

Due to Lemma 8 it suffices to find a computation of  $\mathcal{N}$  that sparsely satisfies  $\neg\varphi$ . We develop an algorithm for this task via two steps. (1) We employ the Vardi–Wolper construction [44] to build a new broadcast network  $\mathcal{N}_{\neg\varphi}$  such that computations of  $\mathcal{N}_{\neg\varphi}$  visiting final states infinitely often are computations of  $\mathcal{N}$  that sparsely satisfy  $\neg\varphi$ . (2) We decide the existence of the former with the fixed-point iteration for *Liveness Verification*, developed in Sect. 3.

Step (1) relies on the following well-known characterization of LTL formulas in terms of automata. The result is crucial for our development since it allows for representing all words that satisfy a given formula as a language of a Büchi automaton. We assume familiarity with the notion of Büchi automata [33].

**Theorem 4** [44] *For any LTL formula  $\psi$  over  $\mathcal{A}$ , one can construct a Büchi automaton  $B_\psi$  of size at most  $2^{\mathcal{O}(|\psi|)}$  such that for any word  $w \in \mathcal{P}(\mathcal{A})^\omega$*

$$w \models \psi \text{ if and only if } w \in L(B_\psi).$$

Apply Theorem 4 to the formula  $\neg\varphi$  of interest. We obtain an automaton  $B_{\neg\varphi} = (Q_B, I_B, \delta_B, F_B)$  over the alphabet  $\mathcal{P}(\mathcal{A})$  with states  $Q_B$ , initial states  $I_B$ , transitions  $\delta_B$ , and final states  $F_B$ . Further, we assume  $B_{\neg\varphi}$  to be non-blocking: in each state we have an outgoing transition on each letter. The language of  $B_{\neg\varphi}$  consists of exactly those words in  $\mathcal{P}(\mathcal{A})^\omega$  that satisfy  $\neg\varphi$ .

In order to obtain computations of  $\mathcal{N}$  that satisfy  $\neg\varphi$ , we have to build a cross product of  $P$ , the clients in  $\mathcal{N}$ , with  $B_{\neg\varphi}$ . Intuitively, a computation in the cross product is then a computation of a client which at the same time satisfies the formula  $\neg\varphi$ . The construction needs to take into account that  $B_{\neg\varphi}$  is an automaton over a different alphabet than  $Ops(D)$ . Let  $P = (Q, I, \delta)$  be the clients of  $\mathcal{N}$  and  $\lambda : Q \rightarrow \mathcal{P}(\mathcal{A})$  the given map. We define the new client  $P_{\neg\varphi} = (Q_{\neg\varphi}, I_{\neg\varphi}, \delta_{\neg\varphi})$  over  $Ops(D)$ . The states are given by  $Q_{\neg\varphi} = Q \times Q_B$ , the initial states by  $I_{\neg\varphi} = I \times I_B$ , transitions  $\delta_{\neg\varphi}$  are defined by

$$(q, p) \xrightarrow{op} (q', p') \in \delta_{\neg\varphi} \text{ if } q \xrightarrow{op} q' \in \delta \text{ and } p \xrightarrow{\lambda(q)} p' \in \delta_B.$$

We also define a set of final states by  $F_{\neg\varphi} = Q \times F_B$ . Then, a computation of  $P_{\neg\varphi}$  that visits  $F$  infinitely often is a computation of  $P$  that satisfies  $\neg\varphi$ .

The broadcast network of interest is  $\mathcal{N}_{\neg\varphi} = (D, P_{\neg\varphi})$ . It uses  $P_{\neg\varphi}$  as clients and yields the desired result: a computation of  $\mathcal{N}_{\neg\varphi}$  in which a client visits  $F_{\neg\varphi}$  infinitely often is a computation of  $\mathcal{N}$  that sparsely satisfies  $\neg\varphi$ .

**Lemma 9** *There is a computation  $\pi$  of  $\mathcal{N}$  such that  $\pi \models_{sparse} \neg\varphi$  if and only if there is a computation  $\pi'$  of  $\mathcal{N}_{\neg\varphi}$  with  $\text{Fin}(\pi') \neq \emptyset$ .*

Note that the lemma reasons about initialized infinite computations and  $\text{Fin}(\pi')$  is defined via the final states  $F_{\neg\varphi}$ . Putting Lemmas 8 and 9 together yields a reduction to *Liveness Verification*. It is left to test whether there is a computation  $\pi'$  of  $\mathcal{N}_{\neg\varphi}$  in which a client visits  $F_{\neg\varphi}$  infinitely often. Step (2) decides this by applying the fixed-point iteration of Sect. 3 to  $(\mathcal{N}_{\neg\varphi}, F_{\neg\varphi})$ .

Regarding the complexity stated in Theorem 3, consider the following. By Theorem 4,  $B_{\neg\varphi}$  has at most  $2^{\mathcal{O}(|\varphi|)}$  states and can be constructed in time  $2^{\mathcal{O}(|\varphi|)}$ . Hence, the clients  $P_{\neg\varphi}$  have at most  $|Q| \cdot 2^{\mathcal{O}(|\varphi|)}$  states. The size of the alphabet is  $\mathcal{O}(|D|) \cdot |\mathcal{N}|$ . The size of  $\mathcal{N}_{\neg\varphi}$  is at most  $2^{\mathcal{O}(|\varphi|)} \cdot |\mathcal{N}|$ . Since the fixed-point iteration for solving *Liveness verification* runs in polynomial time, we obtain the desired complexity estimation. It is left to prove Lemma 9.

**Proof (Idea)** First, let  $\pi$  be an initialized infinite computation of  $\mathcal{N}$  such that  $\pi \models_{sparse} \neg\varphi$ . By definition there is a client  $i \in \text{Inf}(\pi)$  with  $pr_i(\pi) \models \neg\varphi$ . This means that the word  $\lambda(pr_i(\pi)) \in \mathcal{P}(\mathcal{A})^\omega$  satisfies  $\neg\varphi$  and thus lies in the language  $L(B_{\neg\varphi})$ . Hence, there is an initialized computation  $r_i$  of  $B_{\neg\varphi}$  on the word that accepts it. Let  $b_i \in Q_B^\omega$  be the sequence of states appearing in  $r_i$ . Since  $B_{\neg\varphi}$  is non-blocking, there is also an initialized computation  $r_j$  on the word  $\lambda(pr_j(\pi)) \in \mathcal{P}(\mathcal{A})^\omega$ , for each client  $j \neq i$ . Note that  $r_j$  is not necessarily accepting. We extract the sequence of states  $b_j \in Q_B^\omega$  from each  $r_j$ . Denote by  $b$  the vector of all sequences  $b = (b_1, \dots, b_k)$ , where  $k$  is the number of clients.

We combine  $\pi$  with  $b$  to obtain a computation  $\pi'$  of  $\mathcal{N}_{\neg\varphi}$  that satisfies  $\text{Fin}(\pi') \neq \emptyset$ . Let  $\pi(\ell) \in Q^k$  denote  $\ell$ th configuration of  $\pi$  and let  $\pi(\ell, \ell')$  refer to the subcomputation of  $\pi$  starting in  $\pi(\ell)$  and ending in  $\pi(\ell')$ . The idea is to define  $\pi'$  by induction and to join each client of  $\pi$  with the corresponding sequence of states in  $b$ . Technically, this requires maintaining two invariants during the induction: (1)  $\pi'(0, \ell)$  is an initialized computation of  $\mathcal{N}_{\neg\varphi}$ . (2) When we project a configuration visited by  $\pi'(0, \ell)$  to  $Q$ , we obtain the corresponding configuration of  $\pi$ . We refer to [13] for the detailed construction. Note that client  $i$  visits  $F_{\neg\varphi}$  infinitely often. Hence,  $\pi'$  is the desired computation.

For the other direction, note that the projection of a computation  $\pi'$  of  $\mathcal{N}_{\neg\varphi}$  onto the states of  $\mathcal{N}$  provides a proper computation of  $\mathcal{N}$ . Since  $\text{Fin}(\pi') \neq \emptyset$ , there is an  $i \in \text{Fin}(\pi')$  such that  $\lambda(\text{pr}_i(\pi')) \in L(B_{\neg\varphi})$ , meaning that  $\text{pr}_i(\pi) \models \neg\varphi$ . Thus, client  $i \in \text{Inf}(\pi)$  satisfies  $\neg\varphi$  and hence  $\pi \models_{\text{sparse}} \neg\varphi$ .  $\square$

### 5.2 Sparse model checking

The second model checking problem that we consider is *Sparse Model Checking*. It demands that all computations of a broadcast network sparsely satisfy an LTL specification. We design an algorithm similar to *Fair Model Checking* that invokes the Vardi–Wolper construction to establish a reduction to *Fair Liveness Verification*. The latter can then be solved with the algorithm from Sect. 4.

We state the decision problem. To this end, a broadcast network  $\mathcal{N}$  is said to *sparsely satisfy* an LTL formula  $\varphi$  if for each initialized infinite computation  $\pi$  we have that  $\pi \models_{\text{sparse}} \varphi$ . In this case, we write  $\mathcal{N} \models_{\text{sparse}} \varphi$ .

*Sparse Model Checking*

**Input:** A broadcast network  $\mathcal{N} = (D, P)$  and an LTL formula  $\varphi$ .

**Question:** Does  $\mathcal{N} \models_{\text{sparse}} \varphi$  hold?

The following theorem states the main result. Note that the size of the broadcast network only contributes a polynomial factor.

**Theorem 5** *Sparse Model Checking can be solved in time  $2^{\mathcal{O}(|\varphi|)} \cdot |\mathcal{N}|^{\mathcal{O}(1)}$ .*

Like for *Fair Liveness Verification* it is simpler to find a computation violating the given specification than checking whether all computations satisfy it. Hence, we give an algorithm deciding  $\mathcal{N} \not\models_{\text{sparse}} \varphi$ . If a computation  $\pi$  of  $\mathcal{N}$  does not sparsely satisfy  $\varphi$ , all clients in  $\text{Inf}(\pi)$  will satisfy  $\neg\varphi$ . But this means that  $\pi \models_{\text{fair}} \neg\varphi$ . We obtain the following lemma.

**Lemma 10** *We have  $\mathcal{N} \not\models_{\text{sparse}} \varphi$  if and only if there is an initialized infinite computation  $\pi$  of  $\mathcal{N}$  such that  $\pi \models_{\text{fair}} \neg\varphi$ .*

As above, we construct a broadcast network  $\mathcal{N}_{\neg\varphi}$  with a set of final states  $F_{\neg\varphi}$ . We invoke Theorem 4 on  $\neg\varphi$  and obtain an automaton  $B_{\neg\varphi}$ . Now we construct the clients  $P_{\neg\varphi}$  as before, as a cross product of  $P$  and  $B_{\neg\varphi}$ . The resulting broadcast network is  $\mathcal{N}_{\neg\varphi} = (D, P_{\neg\varphi})$ . The final states are  $F_{\neg\varphi} = Q \times F_B$ . Computations of  $\mathcal{N}_{\neg\varphi}$  in which each client that moves infinitely often also visits  $F_{\neg\varphi}$  infinitely often are precisely those

computations of  $\mathcal{N}$  that satisfy  $\neg\varphi$  under fairness. The following lemma summarizes the statement.

**Lemma 11** *There is a computation  $\pi$  of  $\mathcal{N}$  such that  $\pi \models_{\text{fair}} \neg\varphi$  if and only if there is a computation  $\pi'$  of  $\mathcal{N}_{\neg\varphi}$  with  $\text{Inf}(\pi') \subseteq \text{Fin}(\pi')$ .*

From Lemmas 10 and 11, we obtain a reduction to *Fair Liveness Verification*. Hence, we can apply the algorithm from Sect. 4 to decide the existence of  $\pi'$ . Since the algorithm takes only polynomial time, the complexity estimation given in Theorem 5 follows similar to the estimation of Theorem 3.

**Acknowledgements** We thank Arnaud Sangnier for helping us to identify the connection of our work with [6] and for helpful discussions on the topic.

**Funding** Open Access funding enabled and organized by Projekt DEAL.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

1. Abdulla PA, Atig MF, Rezine O (2013) Verification of directed acyclic ad hoc networks. In: FORTE, volume 7892 of LNCS. Springer, pp 193–208
2. Abdulla PA, Sistla AP, Talupur M (2018) Model checking parameterized systems. In: Handbook of model checking, pp 685–725
3. Akhiani H, Doligez D, Harter P, Lamport L, Scheid J, Tuttle MR, Yu Y (1999) Cache coherence verification with TLA+. In: FM, volume 1709 of LNCS. Springer, pp 1871–1872
4. Apt KR, Kozen D (1986) Limits for automatic verification of finite-state concurrent systems. Inf Process Lett 22(6):307–309
5. Balasubramanian AR, Bertrand N, Markey N (2018) Parameterized verification of synchronization in constrained reconfigurable broadcast networks. In: TACAS, volume 10806 of LNCS. Springer, pp 38–54
6. Bertrand N, Fournier P, Sangnier A (2014) Playing with probabilities in reconfigurable broadcast networks. In: FOSSACS, volume 8412 of LNCS. Springer, pp 134–148
7. Bertrand N, Fournier P, Sangnier A (2015) Distributed local strategies in broadcast networks. In: CONCUR, volume 42 of LIPIcs. Schloss Dagstuhl, pp 44–57
8. Bloem R, Jacobs S, Khalimov A, Konnov I, Rubin S, Veith H, Widder J (2015) Decidability of parameterized verification. Synthesis lectures on distributed computing theory. Morgan and Claypool Publishers
9. Bouyer P, Markey N, Randour M, Sangnier A, Stan D (2016) Reachability in networks of register protocols under stochastic schedulers. In: ICALP, volume 55 of LIPIcs. Schloss Dagstuhl, pp 106:1–106:14
10. Chini P, Meyer R, Saivasan P (2018) Fine-grained complexity of safety verification. In: TACAS, volume 10806 of LNCS. Springer, pp 20–37
11. Chini P, Meyer R, Saivasan P (2019) Complexity of liveness in parameterized systems. In: FSTTCS, volume 150 of LIPIcs. Schloss Dagstuhl, pp 37:1–37:15
12. Chini P, Meyer R, Saivasan P (2019) Liveness in broadcast networks. In: NETYS, volume 11704 of LNCS. Springer, pp 52–66

13. Chini P, Meyer R, Saivasan P (2019) Liveness in broadcast networks. CoRR. [arXiv:1904.00833](https://arxiv.org/abs/1904.00833)
14. Chini P, Meyer R, Saivasan P (2020) Fine-grained complexity of safety verification. *J Autom Reason* 64(7):1419–1444
15. Delzanno G (2000) Automatic verification of parameterized cache coherence protocols. In: CAV, volume 1855 of LNCS. Springer, pp 53–68
16. Delzanno G, Sangnier A, Traverso R, Zavattaro G (2012) On the complexity of parameterized reachability in reconfigurable broadcast networks. In: FSTTCS, volume 18 of LIPIcs. Schloss Dagstuhl, pp 289–300
17. Delzanno G, Sangnier A, Zavattaro G (2010) Parameterized verification of ad hoc networks. In: CONCUR, volume 6269 of LNCS. Springer, pp 313–327
18. Delzanno G, Sangnier A, Zavattaro G (2011) On the power of cliques in the parameterized verification of ad hoc networks. In: FOSSACS, volume 6604 of LNCS. Springer, pp 441–455
19. Delzanno G, Sangnier A, Zavattaro G (2012) Verification of ad hoc networks with node and communication failures. In: FORTE, volume 7273 of LNCS. Springer, pp 235–250
20. D’Osualdo E, Kochems J, Luke Ong C-H (2013) Automatic verification of erlang-style concurrency. In: SAS, volume 7935 of LNCS. Springer, pp 454–476
21. D’Osualdo E, Luke Ong C-H (2016) On hierarchical communication topologies in the  $\pi$ -calculus. In: ESOP, volume 9632 of LNCS. Springer, pp 149–175
22. Durand-Gasselín A, Esparza J, Ganty P, Majumdar R (2015) Model checking parameterized asynchronous shared-memory systems. In: CAV, volume 9206 of LNCS. Springer, pp 67–84
23. Esparza J (2003) Some applications of Petri nets to the analysis of parameterised systems (talk). WISP
24. Esparza J (2014) Keeping a crowd safe: On the complexity of parameterized verification (invited talk). In: STACS, volume 25 of LIPIcs. Schloss Dagstuhl, pp 1–10
25. Esparza J, Finkel A, Mayr R (1999) On the verification of broadcast protocols. In: LICS. IEEE, pp 352–359
26. Esparza J, Ganty P, Majumdar R (2013) Parameterized verification of asynchronous shared-memory systems. In: CAV, volume 8044 of LNCS. Springer, pp 124–140
27. Esparza J, Nielsen M (1994) Decidability issues for Petri nets—a survey. *Bull EATCS* 52:244–262
28. Fournier P (2015) Parameterized verification of networks of many identical processes. Ph.D. thesis, University of Rennes 1
29. Hague M (2011) Parameterised pushdown systems with non-atomic writes. In: FSTTCS, volume 13 of LIPIcs. Schloss Dagstuhl, pp 457–468
30. Hague M, Meyer R, Muskalla S, Zimmermann M (2018) Parity to safety in polynomial time for pushdown and collapsible pushdown systems. In: MFCS, volume 117 of LIPIcs. Schloss Dagstuhl, pp 57:1–57:15
31. Hüchting R, Majumdar R, Meyer R (2014) Bounds on mobility. In: CONCUR, volume 8704 of LNCS. Springer, pp 357–371
32. Joshi S, König B (2008) Applying the graph minor theorem to the verification of graph transformation systems. In: CAV, volume 5123 of LNCS. Springer, pp 214–226
33. Khossainov B, Nerode A (2001) Automata theory and its applications. Birkhauser, Boston
34. König B, Kozioura V (2006) Counterexample-guided abstraction refinement for the analysis of graph transformation systems. In: TACAS, volume 3920 of LNCS. Springer, pp 197–211
35. Konnov IV, Lazic M, Veith H, Widder J (2017) A short counterexample property for safety and liveness verification of fault-tolerant distributed algorithms. In: POPL. ACM, pp 719–734
36. Kosaraju SR, Sullivan GF (1988) Detecting cycles in dynamic graphs in polynomial time (preliminary version). In: STOC. ACM, pp 398–406
37. Meyer R (2008) On boundedness in depth in the  $\pi$ -calculus. In: IFIP TCS, volume 273 of IFIP. Springer, pp 477–489
38. Meyer R, Strazny T (2010) Petruccio: From dynamic networks to nets. In: CAV, volume 6174 of LNCS. Springer, pp 175–179
39. Pnueli A (1977) The temporal logic of programs. In: FOCS. IEEE, pp 46–57
40. Pnueli A, Sa’ar Y (2008) All you need is compassion. In: VMCAI, volume 4905 of LNCS. Springer, pp 233–247
41. Rabin MO, Scott D (1959) Finite automata and their decision problems. *IBM J Res Dev* 3(2):114–125
42. Saksena M, Wibling O, Jonsson B (2008) Graph grammar modeling and verification of ad hoc routing protocols. In: TACAS, volume 4963 of LNCS. Springer, pp 18–32

43. Singh A, Ramakrishnan CR, Smolka SA (2009) Query-based model checking of ad hoc network protocols. In: CONCUR, volume 5710 of LNCS. Springer, pp 603–619
44. Vardi M, Wolper P (1986) An automata-theoretic approach to automatic program verification. In: LICS. IEEE, pp 322–331
45. Wies T, Zuffrey D, Henzinger TA (2010) Forward analysis of depth-bounded processes. In: FoSSaCS, volume 6014 of LNCS. Springer, pp 94–108
46. Zufferey D (2013) Analysis of Dynamic Message Passing Programs (a framework for the analysis of depth-bounded systems). Ph.D. thesis, Institute of Science and Technology

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.