

MODULE CHECKING OF PUSHDOWN MULTI-AGENT SYSTEMS *

LAURA BOZZELLI  ^a, ANIELLO MURANO  ^b, AND ADRIANO PERON  ^c

^a University of Napoli “Federico II”, Napoli, Italy
e-mail address: laura.bozzelli@unina.it

^b University of Napoli “Federico II”, Napoli, Italy
e-mail address: aniello.murano@unina.it

^c University of Napoli “Federico II”, Napoli, Italy
e-mail address: adrperon@unina.it

ABSTRACT. In this paper, we investigate the module-checking problem of pushdown multi-agent systems (PMS) against ATL and ATL* specifications. We establish that for ATL, module checking of PMS is 2EXPTIME-complete, which is the same complexity as pushdown module-checking for CTL. On the other hand, we show that ATL* module-checking of PMS turns out to be 4EXPTIME-complete, hence exponentially harder than both CTL* pushdown module-checking and ATL* model-checking of PMS. Our result for ATL* provides a rare example of a natural decision problem that is elementary yet but with a complexity that is higher than triply exponential-time.

1. INTRODUCTION

Model checking is a well-established formal-method technique to automatically check for global correctness of systems [CE81, QS82]. Early use of model checking mainly considered *finite-state closed systems*, modelled as labelled state-transition graphs (Kripke structures) equipped with some internal degree of nondeterminism, and specifications given in terms of standard temporal logics such as the linear-time temporal logic LTL [Pnu77] and the branching-time temporal logics CTL and CTL* [EH86]. In the last two decades, model-checking techniques have been extended to the analysis of reactive and distributed component-based systems, where the behavior of a component depends on assumptions on its environment (the other components). One of the first approaches to model check *finite-state open systems* is *module checking* [KV96], a framework for handling the interaction between a system and an external unpredictable environment. In this setting, the system is modeled as a *module* that is a finite-state Kripke structure whose states are partitioned into those controlled by the system and those controlled by the environment. The latter ones intrinsically carry an additional source of nondeterminism describing the possibility that the computation, from these states, can continue with any subset of its possible successor states. This means that while in model checking, we have only one computation tree representing the possible evolution

Key words and phrases: Infinite-state multi-agent systems, Module checking, Pushdown systems, Logics for strategic reasoning, Alternating-time temporal logics.

* The paper is an extended version completed with proofs of the conference paper [BMP20].

	Model Checking	Model Checking (fixed formula)	Module Checking	Module Checking (fixed formula)
CTL	PTIME [EH86]	NLOGSPACE [BVW94]	EXPTIME [KV96]	PTIME [KV96]
CTL*	PSPACE [EH86]	NLOGSPACE [BVW94]	2EXPTIME [KV96]	PTIME [KV96]
ATL	PTIME [AHK02]	PTIME [AHK02]	EXPTIME [BM17]	PTIME [BM17]
ATL*	2EXPTIME [AHK02]	PTIME [AHK02]	3EXPTIME [BM17]	PTIME [BM17]

Table 1: Complexity results on finite-state model checking and finite-state module checking

of the system, in module checking we have an infinite number of trees to handle, one for each possible behavior of the environment. Deciding whether a system satisfies a property amounts to check that all such trees satisfy the property. This makes module checking harder to deal with. Classically, module checking has been investigated with respect to CTL and CTL* [KV96, KV97, BRS07] specifications and for μ -calculus specifications [FMP08]. An extension of module checking has been also used to reason about three-valued abstractions in [dAGJ04, God03]. Other approaches to the verification of multi-component finite-state systems (*multi-agent systems*) are based on the game paradigm: the system is modeled by a multi-player finite-state concurrent game, where at each step, the next state is determined by considering the “intersection” between the choices made simultaneously and independently by all the players (the agents). In this setting, properties are specified in logics for strategic reasoning such as the alternating-time temporal logics ATL and ATL* [AHK02], the latter ones being well-known extensions of CTL and CTL*, respectively, which allow to express cooperation and competition among agents in order to achieve certain goals. In particular, they can express selective quantification over those paths that are the result of the infinite game between a given coalition and the rest of the agents.

For a long time, there has been a common believe that module checking of CTL/CTL* is a special case of model checking of ATL/ATL*. The belief has been refuted in [JM14] where it is proved that module checking includes two features inherently absent in the semantics of ATL/ATL*, namely irrevocability and nondeterminism of strategies. On the other hand, branching-time temporal logics like CTL and CTL* do not accommodate strategic reasoning. These facts have motivated the extension of module checking to a finite-state multi-agent setting for handling specifications in ATL* [JM15, BM17], which turns out to be more expressive than both CTL* module checking and ATL* model checking [JM14, JM15]. Table 1 summarizes known results about the complexity of finite-state model checking and finite-state module checking. All the complexities in Table 1 denote tight bounds.

Verification of pushdown systems. An active field of research is model checking of pushdown systems. These represent an infinite-state formalism suitable to capture the control flow of procedure calls and returns in programs. Model checking of (closed) pushdown systems against standard regular temporal logics (such as LTL, CTL, CTL*, and the modal μ -calculus) is decidable and it has been intensively studied leading to efficient verification algorithms and tools (see [Wal96, BEM97, BR00, AKM12, AMM14]). The verification of

	Pushdown Model Checking	Pushdown Model Checking (fixed formula)	Pushdown Module Checking	Pushdown Module Checking (fixed formula)
CTL	EXPTIME [Wal00]	EXPTIME [Boz06]	2EXPTIME [BMP10]	EXPTIME [BMP10]
CTL*	2EXPTIME [Boz06]	EXPTIME [Boz06]	3EXPTIME [BMP10]	EXPTIME [BMP10]
ATL	EXPTIME [CSW16]	EXPTIME [CSW16]	2EXPTIME Corollary 3.4	EXPTIME Corollary 3.4
ATL*	3EXPTIME [CSW16]	EXPTIME [CSW16]	4EXPTIME Cor. 3.4 & Theorem 4.1	EXPTIME Corollary 3.4

Table 2: Complexity results on pushdown model checking and pushdown module checking

open pushdown systems in a two-player turn-based setting has been investigated in many works (e.g. see [LMS04, HO09]). Open pushdown systems along with the module-checking paradigm have been considered in [BMP10]. As in the case of finite-state systems, for the logic CTL (resp., CTL*), pushdown module-checking is singly exponentially harder than pushdown model-checking, being precisely 2EXPTIME-complete (resp., 3EXPTIME-complete), although with the same program complexity as pushdown model-checking (that is EXPTIME-complete). Pushdown module-checking has been investigated under several restrictions [ALM⁺13, Boz11, MNP08], including the imperfect-information setting case, where the latter variant is in general undecidable [ALM⁺13]. In [MP15, CSW16], the verification of open pushdown systems has been extended to a concurrent game setting (*pushdown multi-agent systems*) by considering specifications in ATL* and the alternating-time modal μ -calculus. In particular, model checking of PMS against ATL* has the same complexity as pushdown module-checking against CTL* [CSW16].

Our contribution. In this paper, we extend the module-checking framework to the verification of multi-agent pushdown systems (PMS) by addressing the module-checking problem of PMS against ATL and ATL* specifications. By [JM14], the considered setting for ATL (reps., ATL*) is strictly more expressive than both pushdown module checking for CTL (resp., CTL*) and ATL (reps., ATL*) model-checking of PMS. We establish that ATL module-checking for PMS has the same complexity as pushdown module-checking for CTL, that is 2EXPTIME-complete. On the other hand, we show that ATL* module-checking of PMS has a very high complexity: it turns out to be exponentially harder than ATL* model-checking of PMS and pushdown module-checking for CTL*, being, precisely, 4EXPTIME-complete with an EXPTIME-complete complexity for a fixed-size formula. The upper bounds are obtained by an automata-theoretic approach. The matching lower bound for ATL* is shown by a technically non-trivial reduction from the acceptance problem for 3EXPSPACE-bounded alternating Turing Machines. Our result for ATL* provides a rare example of a natural decision problem that is elementary yet but with a complexity that is higher than triply exponential-time. To the best of our knowledge, the unique known characterization of the class 4EXPTIME concerns validity of CTL* on alternating automata with bounded cooperative concurrency [HRV90].

Our results confirm that pushdown module checking is exponentially harder than finite-state module checking. Indeed, like the logics CTL and CTL^* , pushdown module checking against ATL (resp., ATL^*) turns out to be exponentially harder than finite-state module checking against ATL (resp., ATL^*) even for a fixed formula. This is illustrated in Tables 1 and 2, where all the complexities denote tight bounds.

The rest of the paper is organized as follows. In Section 2, we recall the concurrent game setting, the class of multi-agent pushdown systems (PMS), and the logics ATL and ATL^* . Moreover, we introduce the PMS module-checking framework for ATL and ATL^* specifications. In Section 3, we describe the proposed automata-theoretic approach for solving the module-checking problem of PMS against ATL and ATL^* , and in Section 4, we show that for the logic ATL^* , the considered problem is 4EXPTIME-hard. Finally Section 5 provides an assessment of the work done, and outlines future research directions.

2. PRELIMINARIES

We fix the following notations. Let AP be a finite nonempty set of atomic propositions, Ag be a finite nonempty set of agents, and Ac be a finite nonempty set of actions that can be made by agents. For a set $A \subseteq Ag$ of agents, an A -decision d_A is an element in Ac^A assigning to each agent $a \in A$ an action $d_A(a)$. For $A, A' \subseteq Ag$ with $A \cap A' = \emptyset$, an A -decision d_A and A' -decision $d_{A'}$, $d_A \cup d_{A'}$ denotes the $(A \cup A')$ -decision defined in the obvious way. Let $Dc = Ac^{Ag}$ be the set of *full decisions* of all the agents in Ag .

Let \mathbb{N} be the set of natural numbers. For an infinite word w over an alphabet Σ and $i \geq 0$, $w(i)$ denotes the $(i+1)^{th}$ letter of w and $w_{\geq i}$ the suffix of w starting from the $(i+1)^{th}$ letter of w , i.e., the infinite word $w(i)w(i+1)\dots$. For a finite word w over Σ , $|w|$ is the length of w .

Given a set Υ of directions, an (*infinite*) Υ -*tree* T is a non-empty prefix closed subset of Υ^* such that for all $\nu \in T$, $\nu \cdot \gamma \in T$ for some $\gamma \in \Upsilon$. Elements of T are called nodes and ε is the root of T . For $\nu \in T$, a child of ν in T is a node of the form $\nu \cdot \gamma$ for some $\gamma \in \Upsilon$. An (*infinite*) path of T is an infinite sequence π of nodes such that $\pi(i+1)$ is a child in T of $\pi(i)$ for all $i \geq 0$. For an alphabet Σ , a Σ -labeled Υ -tree is a pair $\langle T, \text{Lab} \rangle$ consisting of a Υ -tree and a labelling $\text{Lab} : T \mapsto \Sigma$ assigning to each node in T a symbol in Σ . We extend the labeling Lab to paths π in the obvious way, i.e. $\text{Lab}(\pi)$ is the infinite word over Σ given by $\text{Lab}(\pi(0))\text{Lab}(\pi(1))\dots$. The labeled tree $\langle T, \text{Lab} \rangle$ is *complete* if $T = \Upsilon^*$. Given $k \in \mathbb{N} \setminus \{0\}$, a k -ary tree is a $\{1, \dots, k\}$ -tree.

Concurrent game structures (CGS). CGS [AHK02] extend Kripke structures to a setting involving multiple agents. They can be viewed as multi-player games in which players perform concurrent actions, chosen strategically as a function of the history of the game.

Definition 2.1 (CGS). A CGS (over AP , Ag , and Ac) is a tuple $\mathcal{G} = \langle S, s_0, \text{Lab}, \tau \rangle$, where S is a set of states, $s_0 \in S$ is the initial state, $\text{Lab} : S \mapsto 2^{AP}$ maps each state to a set of atomic propositions, and $\tau : S \times Dc \mapsto S \cup \{\dashv\}$ is a transition function that maps a state and a full decision either to a state or to the special symbol \dashv (\dashv is for ‘undefined’) such that for all states s , there exists $d \in Dc$ so that $\tau(s, d) \neq \dashv$. Given a set $A \subseteq Ag$ of agents, an A -decision d_A , and a state s , we say that d_A is *available at state s* if there exists an $(Ag \setminus A)$ -decision $d_{Ag \setminus A}$ such that $\tau(s, d_A \cup d_{Ag \setminus A}) \in S$.

For a state s and an agent a , *state s is controlled by a* if there is a unique $(Ag \setminus \{a\})$ -decision available at state s . Agent a is *passive in s* if there is a unique $\{a\}$ -decision available

at state s . A *multi-agent turn-based game* is a CGS where each state is controlled by an agent.

Note that in modelling independent agents, usually one assumes that at each state s , each agent a has a set $Ac_{a,s} \subseteq Ac$ of actions which are enabled at the state s . This is reflected in the transition function τ by requiring that the set of full decisions d such that $\tau(s, d) \neq \perp$ corresponds to $(Ac_{a,s})_{a \in Ag}$.

We now recall the notion of strategy in a CGS $\mathcal{G} = \langle S, s_0, Lab, \tau \rangle$. Here, we consider *perfect recall* strategies where an agent decides the next action by using all the available information up to the current round. A *play* is an infinite sequence of states $s_1 s_2 \dots$ such that for all $i \geq 1$, s_{i+1} is a *successor* of s_i , i.e. $s_{i+1} = \tau(s_i, d)$ for some full decision d . A *track* (or *history*) ν is a nonempty prefix of some play. Given a set $A \subseteq Ag$ of agents, a *strategy for A* is a mapping f_A assigning to each track ν (representing the history the agents saw so far) an A -decision available at the last state, denoted $lst(\nu)$, of ν . The *outcome* function $out(s, f_A)$ for a state s and the strategy f_A returns the set of all the plays starting at state s that can occur when agents A execute strategy f_A from state s on. Formally, $out(s, f_A)$ is the set of plays $\pi = s_1 s_2 \dots$ such that $s_1 = s$ and for all $i \geq 1$, there is $d \in Ac^{Ag \setminus A}$ so that $s_{i+1} = \tau(s_i, f_A(s_1 \dots s_i) \cup d)$.

Definition 2.2. For a set Υ of directions, a *Concurrent Game Υ -Tree* (Υ -CGT) is a CGS $\langle T, \varepsilon, Lab, \tau \rangle$, where $\langle T, Lab \rangle$ is a 2^{AP} -labeled Υ -tree, and for each node $x \in T$, the successors of x correspond to the children of x in T . Every CGS $\mathcal{G} = \langle S, s_0, Lab, \tau \rangle$ induces a S -CGT, denoted by $Unw(\mathcal{G})$, obtained by unwinding \mathcal{G} from the initial state in the usual way. Formally, $Unw(\mathcal{G}) = \langle T, \varepsilon, Lab', \tau' \rangle$, where $\nu \in T$ iff $s_0 \cdot \nu$ is a track of \mathcal{G} , and for all $\nu \in T$ and $d \in Dc$, $Lab'(\nu) = Lab(lst(\nu))$ and $\tau'(\nu, d) = \nu \cdot \tau(lst(\nu), d)$, with $lst(\varepsilon) = s_0$.

Pushdown multi-agent systems (PMS). PMS, introduced in [MP15], generalize standard pushdown systems to a concurrent multi-player setting.

Definition 2.3. A PMS (over AP , Ag , and Ac) is a tuple $\mathcal{S} = \langle Q, \Gamma \cup \{\gamma_0\}, q_0, Lab, \Delta \rangle$, where Q is a finite set of (control) states, $\Gamma \cup \{\gamma_0\}$ is a finite stack alphabet ($\gamma_0 \notin \Gamma$ is the special *stack bottom symbol*), $q_0 \in Q$ is the initial state, $Lab : Q \mapsto 2^{AP}$ maps each state to a set of atomic propositions, and $\Delta : Q \times (\Gamma \cup \{\gamma_0\}) \times Dc \mapsto (Q \times \Gamma^*) \cup \{\perp\}$ is a transition function (\perp is for ‘undefined’) such that for all pairs $(q, \gamma) \in Q \times (\Gamma \cup \{\gamma_0\})$, there is $d \in Dc$ so that $\Delta(q, \gamma, d) \neq \perp$.

The size $|\Delta|$ of the transition function Δ is given by $|\Delta| = \sum_{(q', \beta) \in Ran(\Delta)} |\beta|$, where $Ran(\Delta)$ is the set of pairs $(q', \beta) \in Q \times \Gamma^*$ such that $(q', \beta) = \Delta(q, \gamma, d)$ for some $(q, \gamma, d) \in Q \times (\Gamma \cup \{\gamma_0\}) \times Dc$. A *configuration* of the PMS \mathcal{S} is a pair (q, β) where q is a (control) state and $\beta \in \Gamma^* \cdot \gamma_0$ is a stack content. Intuitively, when the PMS \mathcal{S} is in state q , the stack top symbol is γ and the agents take a full decision d available at the current configuration, i.e. such that $\Delta(q, \gamma, d) = (q', \beta)$ for some $(q', \beta) \in Q \times \Gamma^*$, then \mathcal{S} moves to the configuration with state q' and stack content obtained by removing γ and pushing β (if $\gamma = \gamma_0$ then γ is not removed). Formally, the PMS $\mathcal{S} = \langle Q, \Gamma \cup \{\gamma_0\}, q_0, Lab, \Delta \rangle$ induces the infinite-state CGS $\mathcal{G}(\mathcal{S}) = \langle S, s_0, Lab', \tau \rangle$, where S is the set of configurations of \mathcal{S} , $s_0 = (q_0, \gamma_0)$ (initially, the stack contains just the bottom symbol γ_0), $Lab'((q, \beta)) = Lab(q)$ for each configuration (q, β) , and the transition function τ is defined as follows for all $((q, \gamma \cdot \beta), d) \in S \times Dc$, where $\gamma \in \Gamma \cup \{\gamma_0\}$:

- either $\Delta(q, \gamma, d) = \perp$ and $\tau((q, \gamma \cdot \beta), d) = \perp$,

- or $\gamma \in \Gamma$, $\Delta(q, \gamma, d) = (q', \beta')$, and $\tau((q, \gamma \cdot \beta), d) = (q', \beta' \cdot \beta)$,
- or $\gamma = \gamma_0$ (hence, $\beta = \varepsilon$), $\Delta(q, \gamma, d) = (q', \beta')$, and $\tau((q, \gamma \cdot \beta), d) = (q', \beta' \cdot \gamma_0)$.

2.1. The logics ATL^* and ATL . We recall the alternating-temporal logics ATL^* and ATL proposed by Alur et al. [AHK02] as extensions of the standard branching-time temporal logics CTL^* and CTL (respectively) [EH86], where the path quantifiers are replaced by more general parameterized quantifiers which allow for reasoning about the strategic capability of groups of agents. For the given sets AP and Ag of atomic propositions and agents, ATL^* formulas φ are defined by the following grammar:

$$\varphi ::= \text{true} \mid p \mid \neg\varphi \mid \varphi \vee \varphi \mid X\varphi \mid \varphi U \varphi \mid \langle\langle A \rangle\rangle \varphi$$

where $p \in AP$, $A \subseteq Ag$, X and U are the standard ‘‘next’’ and ‘‘until’’ temporal modalities, and $\langle\langle A \rangle\rangle$ is the ‘‘existential strategic quantifier’’ parameterized by a set A of agents. Formula $\langle\langle A \rangle\rangle \varphi$ expresses the property that the group of agents A has a collective strategy to enforce property φ . In addition, we use standard shorthands: the ‘‘eventually’’ temporal modality $F\varphi := \text{true} U \varphi$, the ‘‘release’’ temporal modality $\varphi_1 R \varphi_2 := \neg(\neg\varphi_1 U \neg\varphi_2)$, and the ‘‘always’’ temporal modality $G\varphi := \neg\text{true} R \varphi$.

A *state formula* is a formula where each temporal modality is in the scope of a strategic quantifier. The logic ATL is the fragment of ATL^* where each temporal modality is immediately preceded by a strategic quantifier. Formally, the set of ATL formulas are defined by the following grammar:

$$\varphi ::= \text{true} \mid p \mid \neg\varphi \mid \varphi \vee \varphi \mid \langle\langle A \rangle\rangle X\varphi \mid \langle\langle A \rangle\rangle(\varphi U \varphi) \mid \langle\langle A \rangle\rangle(\varphi R \varphi)$$

Note that CTL^* (resp., CTL) corresponds to the fragment of ATL^* (resp., ATL), where only the strategic modalities $\langle\langle Ag \rangle\rangle$ and $\langle\langle \emptyset \rangle\rangle$ (equivalent to the existential and universal path quantifiers E and A , respectively) are allowed.

Given a CGS \mathcal{G} with labeling Lab and a play π of \mathcal{G} , the satisfaction relation $\mathcal{G}, \pi \models \varphi$ for ATL^* is defined as follows (Boolean connectives are treated as usual):

$$\begin{aligned} \mathcal{G}, \pi \models p &\Leftrightarrow p \in Lab(\pi(0)) \\ \mathcal{G}, \pi \models X\varphi &\Leftrightarrow \mathcal{G}, \pi_{\geq 1} \models \varphi \\ \mathcal{G}, \pi \models \varphi_1 U \varphi_2 &\Leftrightarrow \text{there is } j \geq 0 : \mathcal{G}, \pi_{\geq j} \models \varphi_2 \text{ and } \mathcal{G}, \pi_{\geq k} \models \varphi_1 \text{ for all } 0 \leq k < j \\ \mathcal{G}, \pi \models \langle\langle A \rangle\rangle \varphi &\Leftrightarrow \text{for some strategy } f_A \text{ for } A, \mathcal{G}, \pi' \models \varphi \text{ for all } \pi' \in out(\pi(0), f_A). \end{aligned}$$

For a state s of \mathcal{G} , $\mathcal{G}, s \models \varphi$ if there is a play π starting from s such that $\mathcal{G}, \pi \models \varphi$. Note that if φ is a state formula, then for all plays π and π' from s , $\mathcal{G}, \pi \models \varphi$ iff $\mathcal{G}, \pi' \models \varphi$. \mathcal{G} is a model of φ , denoted $\mathcal{G} \models \varphi$, if for the initial state s_0 , $\mathcal{G}, s_0 \models \varphi$. Note that $\mathcal{G} \models \varphi$ iff $Unw(\mathcal{G}) \models \varphi$.

2.2. ATL^* and ATL Pushdown Module-checking. The module-checking framework was proposed in [KV96] for the verification of finite open systems, that is systems that interact with an environment whose behavior cannot be determined in advance. In such a framework, the system is modeled by a *module* corresponding to a two-player turn-based game between the system and the environment. Thus, in a module, the set of states is partitioned into a set of system states (controlled by the system) and a set of environment states (controlled by the environment).

The module-checking problem takes two inputs: a module M and a branching-time temporal formula ψ . The idea is that the open system should satisfy the specification ψ no

matter how the environment behaves. Let us consider the unwinding $Unw(M)$ of M into an infinite tree. Checking whether $Unw(M)$ satisfies ψ is the usual model-checking problem. On the other hand, for an open system, $Unw(M)$ describes the interaction of the system with a maximal environment, i.e. an environment that enables all the external nondeterministic choices. In order to take into account all the possible behaviors of the environment, we have to consider all the trees T obtained from $Unw(M)$ by pruning subtrees whose root is a successor of an environment state (pruning these subtrees corresponds to disabling possible environment choices). Therefore, a module M satisfies ψ if all these trees T satisfy ψ .

It has been proved in [JM14] that module checking of CTL/CTL^* includes two features inherently absent in the semantics of ATL/ATL^* , namely irrevocability of strategies and nondeterminism of strategies. Intuitively, unlike the standard ATL^* semantics, in module checking, a formula is evaluated over a restricted behaviour of the full execution tree which corresponds to the strategy tree induced by a possible environment behaviour (*irrevocability of the environment's strategies*). On the other hand, temporal logics like CTL and CTL^* do not accommodate strategic reasoning. These facts have motivated the extension of module checking to a multi-agent setting for handling specifications in ATL^* [JM15], which turns out to be more expressive than both CTL^* module checking and ATL^* model checking [JM14, JM15].

In this section, we first recall the ATL^* module-checking framework. Then, we generalize this setting to pushdown multi-agent systems. In the multi-agent module-checking setting, one considers CGS with a distinguished agent (the *environment*).

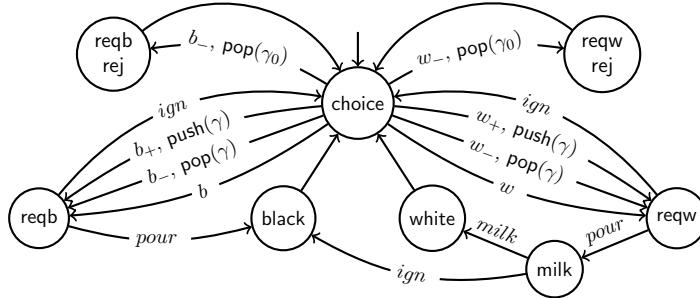
Definition 2.4 (Open CGS). An open CGS is a CGS $\mathcal{G} = \langle S, s_0, Lab, \tau \rangle$ containing a special agent called “the environment” ($env \in Ag$). Moreover, for every state s , either s is controlled by the environment (*environment state*) or the environment is passive in s (*system state*).

For an open CGS $\mathcal{G} = \langle S, s_0, Lab, \tau \rangle$, the set of *environment strategy trees of \mathcal{G}* , denoted $exec(\mathcal{G})$, is the set of S -CGT obtained from $Unw(\mathcal{G})$ by possibly pruning some environment transitions. Formally, $exec(\mathcal{G})$ is the set of S -CGT $\mathcal{T} = \langle T, \varepsilon, Lab', \tau' \rangle$ such that T is a prefix closed subset of the set of $Unw(\mathcal{G})$ -nodes and for all $\nu \in T$ and $d \in Dc$, $Lab'(\nu) = Lab(lst(\nu))$, and $\tau'(\nu, d) = \nu \cdot \tau(lst(\nu), d)$ if $\nu \cdot \tau(lst(\nu), d) \in T$, and $\tau'(\nu, d) = \neg$ otherwise, where $lst(\varepsilon) = s_0$. Moreover, for all $\nu \in T$, the following holds:

- if $lst(\nu)$ is a system state, then for each successor s of $lst(\nu)$ in \mathcal{G} , $\nu \cdot s \in T$;
- if $lst(\nu)$ is an environment state, then there is a nonempty subset $\{s_1, \dots, s_n\}$ of the set of $lst(\nu)$ -successors such that the set of children of ν in T is $\{\nu \cdot s_1, \dots, \nu \cdot s_n\}$.

Intuitively, when \mathcal{G} is in a system state s , then all the transitions from s are enabled. When \mathcal{G} is instead in an environment state, the set of enabled transitions from s depend on the current environment. Since the behavior of the environment is nondeterministic, we have to consider all the possible subsets of the set of s -successors. The only constraint, since we consider environments that cannot block the system, is that not all the transitions from s can be disabled. Note that $Unw(\mathcal{G}) \in exec(\mathcal{G})$ ($Unw(\mathcal{G})$ corresponds to the maximal environment that never restricts the set of its next states).

It is worth noting that the choices made by the environment along an environment strategy tree describe a strategy of the environment which is nondeterministic. This is in contrast with the given notion of strategy for a coalition A of agents which is instead deterministic (at each round, the coalition A selects exactly one A -decision available at the current state).

Figure 1: Multi-agent pushdown coffee machine \mathcal{S}_{cof}

For an open CGS \mathcal{G} and an ATL* formula φ , \mathcal{G} *reactively satisfies* φ , denoted $\mathcal{G} \models^r \varphi$, if for all environment strategy trees $\mathcal{T} \in \text{exec}(\mathcal{G})$, $\mathcal{T} \models \varphi$. Note that $\mathcal{G} \models^r \varphi$ implies $\mathcal{G} \models \varphi$ (since $\text{Unw}(\mathcal{G}) \in \text{exec}(\mathcal{G})$), but the converse in general does not hold. Moreover, $\mathcal{G} \not\models^r \varphi$ is not equivalent to $\mathcal{G} \models^r \neg \varphi$. Indeed, $\mathcal{G} \not\models^r \varphi$ just states that there is some $\mathcal{T} \in \text{exec}(\mathcal{G})$ satisfying $\neg \varphi$.

Pushdown Module-checking. An *open PMS* is a PMS \mathcal{S} such that the induced CGS $\mathcal{G}(\mathcal{S})$ is open. Note that for an open PMS, the property of a configuration of being an environment or system configuration depends only on the control state and the symbol on the top of the stack. The *pushdown module-checking problem against ATL* (resp., ATL*) is checking for a given open PMS \mathcal{S} and an ATL formula (resp., ATL* state formula) φ whether $\mathcal{G}(\mathcal{S}) \models^r \varphi$.

Example 2.5. Consider a coffee machine that allows customers (acting the role of the environment) to choose between the following actions:

- ordering and paying a black or white coffee (actions b or w);
- the same as in the previous point but, additionally, paying a “suspended” coffee (a prepaid coffee) for the benefit of any unknown needy customer claiming it in the future (actions b_+ or w_+);
- asking for an available prepaid (black or white) coffee (actions b_- or w_-).

The coffee machine is modeled by a turn-based open PMS \mathcal{S}_{cof} with three agents: the environment, the brewer br whose function is to pour coffee into the cup (action $pour$), and the milk provider who can add milk (action $milk$). The two system agents can be faulty and ignore the request from the environment (action ign). The stack is exploited for keeping track of the number of prepaid coffees: a request for a prepaid coffee can be accepted only if the stack is not empty. After the completion of a request, the machine waits for further selections. The PMS \mathcal{S}_{cof} is represented as a graph in Figure 1 where each node (control state) is labeled by the propositions holding at it: the state labeled by $choice$ is controlled by the environment, the states labeled by $reqb$ or $reqw$ are controlled by the brewer br , while the state labeled by $milk$ is controlled by the milk provider. The notation $push(\gamma)$ denotes a push stack operation (pushing the symbol $\gamma \neq \gamma_0$), while $pop(\gamma)$ (resp., $pop(\gamma_0)$) denotes a pop operation onto a non-empty (resp., empty) stack. The set of propositions is $\{reqw, reqb, rej, black, white\}$.

In module checking, we can condition the property to be achieved on the behaviour of the environment. For instance, users who never order white coffee and whose request is never

rejected can be served by the brewer alone: $\mathcal{G}(\mathcal{S}_{cof}) \models^r \text{AG}(\neg \text{reqw} \wedge \neg \text{rej}) \rightarrow \langle\langle br \rangle\rangle \text{F black}$. In model checking, the same formula does not express any interesting property since $\mathcal{G}(\mathcal{S}_{cof}) \not\models \text{AG}(\neg \text{reqw} \wedge \neg \text{rej})$. Likewise $\mathcal{G}(\mathcal{S}_{cof}) \models \text{AG} \neg \text{reqw} \rightarrow \langle\langle br \rangle\rangle \text{F black}$, whereas module checking gives a different and more intuitive answer: $\mathcal{G}(\mathcal{S}_{cof}) \not\models^r \text{AG} \neg \text{reqw} \rightarrow \langle\langle br \rangle\rangle \text{F black}$ (there are environments where requests for a prepaid coffee are always rejected).

3. DECISION PROCEDURES

In this section, we provide an automata-theoretic framework for solving the pushdown module-checking problem against ATL and ATL* which is based on the use of *parity alternating automata for CGS* (parity ACG) [SF06] and *parity Nondeterministic Pushdown Tree Automata* (parity NPTA) [KPV02]. The proposed approach (which is proved to be asymptotically optimal in Section 4) consists of two steps. For the given open PMS \mathcal{S} and ATL formula (resp., ATL* state formula) φ , by exploiting known results, we first build in linear-time (resp., double exponential time) a parity ACG $\mathcal{A}_{\neg\varphi}$ accepting the set of CGT which satisfy $\neg\varphi$. Then in the second step, we show how to construct a parity NPTA \mathcal{P} accepting suitable encodings of the environment strategy trees of $\mathcal{G}(\mathcal{S})$ accepted by $\mathcal{A}_{\neg\varphi}$. Hence, $\mathcal{G}(\mathcal{S}) \models^r \varphi$ iff the language accepted by \mathcal{P} is empty.

In the following, we first recall the frameworks of parity NPTA and parity ACG, and the known translations of ATL* and ATL formulas into equivalent parity ACG. Then, in Subsection 3.1, by exploiting parity NPTA, we show that given an open PMS \mathcal{S} and a parity ACG \mathcal{A} , checking that no environment strategy tree of $\mathcal{G}(\mathcal{S})$ is accepted by \mathcal{A} can be done in time double exponential in the size of \mathcal{A} and singly exponential in the size of \mathcal{S} .

Parity NPTA [KPV02]. Here, we describe parity NPTA (without ε -transitions) over labeled complete k -ary trees for a given $k \geq 1$, which are tuples $\mathcal{P} = \langle \Sigma, Q, \Gamma \cup \{\gamma_0\}, q_0, \rho, \Omega \rangle$, where Σ is a finite input alphabet, Q is a finite set of (control) states, $\Gamma \cup \{\gamma_0\}$ is a finite stack alphabet ($\gamma_0 \notin \Gamma$ is the special bottom symbol), $q_0 \in Q$ is an initial state, $\rho : Q \times \Sigma \times (\Gamma \cup \{\gamma_0\}) \rightarrow 2^{(Q \times \Gamma^*)^k}$ is a transition function, and $\Omega : Q \mapsto \mathbb{N}$ is a *parity acceptance condition* over Q assigning to each state a natural number called *color*. The *index* of \mathcal{P} is the number of colors in Ω , i.e., the cardinality of $\Omega(Q)$.

Intuitively, when the automaton is in state q , reading an input node x labeled by $\sigma \in \Sigma$, and the stack contains a word $\gamma \cdot \beta$ in $\Gamma^* \cdot \gamma_0$, then the automaton chooses a tuple $\langle (q_1, \beta_1), \dots, (q_k, \beta_k) \rangle \in \rho(q, \sigma, \gamma)$ and splits in k copies such that for each $1 \leq i \leq k$, a copy in state q_i , and stack content obtained by removing γ and pushing β_i , is sent to the node $x \cdot i$ in the input tree.

Formally, a run of the NPTA \mathcal{P} on a Σ -labeled complete k -ary tree $\langle T, \text{Lab} \rangle$ (with $T = \{1, \dots, k\}^*$) is a $(Q \times \Gamma^* \cdot \gamma_0)$ -labeled tree $r = \langle T, \text{Lab}_r \rangle$ such that $\text{Lab}_r(\varepsilon) = (q_0, \gamma_0)$ (initially, the stack contains just the bottom symbol γ_0) and for each $x \in T$ with $\text{Lab}_r(x) = (q, \gamma \cdot \beta)$, there is $\langle (q_1, \beta_1), \dots, (q_k, \beta_k) \rangle \in \rho(q, \text{Lab}(x), \gamma)$ such that for all $1 \leq i \leq k$, $\text{Lab}_r(x \cdot i) = (q_i, \beta_i \cdot \beta)$ if $\gamma \neq \gamma_0$, and $\text{Lab}_r(x \cdot i) = (q_i, \beta_i \cdot \gamma_0)$ otherwise (note that in this case $\beta = \varepsilon$). The run $r = \langle T, \text{Lab}_r \rangle$ is accepting if for all infinite paths π starting from the root, the highest color $\Omega(q)$ of the states q appearing infinitely often along $\text{Lab}_r(\pi)$ is even. The language $\mathcal{L}(\mathcal{P})$ accepted by \mathcal{P} consists of the Σ -labeled complete k -ary trees $\langle T, \text{Lab} \rangle$ such that there is an accepting run of \mathcal{P} over $\langle T, \text{Lab} \rangle$.

For complexity analysis, we consider the following two parameters:

- the size $|\rho|$ of ρ given by

$$|\rho| := \sum_{(q, \sigma, \gamma) \in Q \times \Sigma \times (\Gamma \cup \{\gamma_0\})} \sum_{\langle (q_1, \beta_1), \dots, (q_k, \beta_k) \rangle \in \rho(q, \sigma, \gamma)} |\beta_1| + \dots + |\beta_k|,$$

- and the smaller parameter $||\rho||$ given by $||\rho|| := \sum_{\beta \in \rho_0} |\beta|$ where ρ_0 is the set of words $\beta \in \Gamma^* \cdot \gamma_0$ occurring in ρ .

It is well-known [KPV02] that emptiness of parity NPTA can be solved in single exponential time by a polynomial time reduction to emptiness of standard two-way alternating tree automata [Var98]. In particular, the following holds (see [KPV02, BMP10]).

Proposition 3.1 [KPV02, BMP10]. *The emptiness problem for a parity NPTA of index m with n states and transition function ρ can be solved in time $O(|\rho| \cdot 2^{O(||\rho||^2 \cdot n^2 \cdot m^2 \log m)})$.*

Parity alternating automata for CGS (parity ACG) [SF06]. ACG generalize alternating automata by branching universally or existentially over all successors that result from the decisions of agents. Formally, for a set X , let $\mathbb{B}^+(X)$ be the set of *positive* Boolean formulas over X , i.e. Boolean formulas built from elements in X using \vee and \wedge . A subset Y of X is a *model* of $\theta \in \mathbb{B}^+(X)$ if the truth assignment that assigns true (resp., false) to the elements in Y (resp., $X \setminus Y$) satisfies θ .

A parity ACG over 2^{AP} and Ag is a tuple $\mathcal{A} = \langle Q, q_0, \delta, \Omega \rangle$, where Q , q_0 , and Ω are defined as for NPTA, while δ is a transition function of the form $\delta : Q \times 2^{AP} \rightarrow \mathbb{B}^+(Q \times \{\square, \diamond\} \times 2^{Ag})$. The transition function δ maps a state and an input letter to a positive Boolean combination of universal atoms (q, \square, A) and existential atoms (q, \diamond, A) . Intuitively, a universal atom (q, \square, A) prescribes that for some A -decision d_A available at the current state s of the input CGS, copies of the automaton in state q are sent to *all* the successors of s which are consistent with d_A . Dually, an existential atom (q, \diamond, A) prescribes that for all A -decisions d_A available at the current state s of the input CGS, a copy of the automaton in state q is sent to *some* successor of s which is consistent with d_A .

The size $|\mathcal{A}|$ of \mathcal{A} is $|Q| + |\text{Atoms}(\mathcal{A})|$, where $\text{Atoms}(\mathcal{A})$ is the set of atoms of \mathcal{A} , i.e. the set of tuples in $Q \times \{\square, \diamond\} \times 2^{Ag}$ occurring in the transition function δ .

We interpret the parity ACG \mathcal{A} over CGT. Note that since the set of full decisions of all agents is finite, a CGT, which is a special CGS, is *finitely-branching*, i.e., each state has a finite number of successors. Thus, given a CGT $\mathcal{T} = \langle T, \varepsilon, \text{Lab}, \tau \rangle$ over AP and Ag , a possible behaviour of \mathcal{A} over the input \mathcal{T} can be formalized by an \mathbb{N} -tree (i.e., a tree with the *countable* set of directions \mathbb{N}) whose nodes are labeled by pairs $(q, \nu) \in Q \times T$ describing a copy of the automaton that is in the state q and reads the node ν of T . Formally, a run of \mathcal{A} over the input \mathcal{T} is a $(Q \times T)$ -labeled \mathbb{N} -tree $r = \langle T_r, \text{Lab}_r \rangle$ such that

- $\text{Lab}_r(\varepsilon) = (q_0, \varepsilon)$ (initially, the automaton is in state q_0 reading the root node of \mathcal{T}),
- for each $y \in T_r$ with $\text{Lab}_r(y) = (q, \nu)$, there is a set $H \subseteq Q \times \{\square, \diamond\} \times 2^{Ag}$ such that H is a model of $\delta(q, \text{Lab}(\nu))$ and the set L of labels associated with the children of y in T_r satisfies the following conditions:
 - for all universal atoms $(q', \square, A) \in H$, there is an available A -decision d_A in the node ν of \mathcal{T} such that for all the children ν' of ν which are consistent with d_A , $(q', \nu') \in L$;
 - for all existential atoms $(q', \diamond, A) \in H$ and for all available A -decisions d_A in the node ν of \mathcal{T} , there is some child ν' of ν which is consistent with d_A such that $(q', \nu') \in L$.

The run r is accepting if for all infinite paths π starting from the root, the highest color of the states appearing infinitely often along $Lab_r(\pi)$ is even. The language $\mathcal{L}(\mathcal{A})$ accepted by \mathcal{A} consists of the CGT \mathcal{T} on AP and Ag such that there is an accepting run of \mathcal{A} over \mathcal{T} .

From ATL* and ATL to parity ACG. In the following we shall exploit a known translation of ATL* state formulas (resp., ATL formulas) into equivalent parity ACG which has been provided in [BM17]. To this end we recall that, for a finite set B disjunct from AP and a CGT $\mathcal{T} = \langle T, \varepsilon, Lab, \tau \rangle$ over AP , a B -labeling extension of \mathcal{T} is a CGT over $AP \cup B$ of the form $\langle T, \varepsilon, Lab', \tau \rangle$, where $Lab'(\nu) \cap AP = Lab(\nu)$ for all $\nu \in T$. A *basic formula* of ATL* is a state formula of ATL* having the form $\langle\langle A \rangle\rangle\varphi$. The result exploited in the following, which corresponds to Theorem 1 in [BM17], is summarized as follows.

Theorem 3.2 [BM17]. *For an ATL* state formula (resp., ATL formula) φ over AP , one can construct in doubly exponential time (resp., linear time) a parity ACG \mathcal{A}_φ over $2^{AP \cup B_\varphi}$, where B_φ is the set of basic subformulas of φ , such that for all CGT \mathcal{T} over AP , \mathcal{T} is a model of φ iff there exists a B_φ -labeling extension of \mathcal{T} which is accepted by \mathcal{A}_φ . Moreover, \mathcal{A}_φ has size $O(2^{2^{O(|\varphi| \cdot \log(|\varphi|))}})$ and index $2^{O(|\varphi|)}$ (resp., size $O(|\varphi|)$ and index 2).*

Note that for the proof of Theorem 3.2, one exploits the fact that for a given CGT \mathcal{T} and ATL* state formula φ , there exists a *unique* B_φ -labeling extension of \mathcal{T} which is *well-formed*, i.e., such that for each node ν of \mathcal{T} , the B_φ -labeling of ν coincides with the set of basic subformulas of φ which hold at node ν of \mathcal{T} . Thus, the automaton ACG \mathcal{A}_φ of Theorem 3.2, by crucially exploiting alternation, recursively checks that the given B_φ -labeling extension of \mathcal{T} is well-formed. Hence, \mathcal{A}_φ accepts a B_φ -labeling extension of \mathcal{T} *only if* it is well-formed.

It is worth noting that while the well-known translation of CTL* formulas into alternating automata involves just a single exponential blow-up, by Theorem 3.2, the translation of ATL* formulas in alternating automata for CGS entails a double exponential blow-up. This seems in contrast with the automata-theoretic approach used in [Sch08] for solving satisfiability of ATL* (recall that ATL* satisfiability has the same complexity as CTL* satisfiability, i.e., it is 2EXPTIME-complete [Sch08]). In particular, given an ATL* state formula φ , one can construct in singly exponential time a parity ACG accepting the set of CGT satisfying some special requirements which provide a necessary and sufficient condition for ensuring the existence of some model of φ [Sch08]. These requirements are based on an equivalent representation of the models of a formula obtained by a sort of widening operation. However, when applied to the environment strategy trees of a CGS, such an encoding is not regular since one has to require that for all nodes in the encoding which are copies of the same environment node in the given environment strategy tree, the associated subtrees are isomorphic. Hence, the approach used in [Sch08] cannot be applied to the module-checking setting.

3.1. Upper bounds for ATL and ATL* pushdown module-checking. Let \mathcal{S} be an open PMS, φ an ATL* (resp., ATL) formula, and $\mathcal{A}_{\neg\varphi}$ the parity ACG over $2^{AP \cup B_\varphi}$ (B_φ is the set of basic subformulas of φ) of Theorem 3.2 associated with the negation of φ . By Theorem 3.2, checking that $\mathcal{G}(\mathcal{S}) \models^r \varphi$ reduces to checking that there are no B_φ -labeling extensions of the environment strategy trees of $\mathcal{G}(\mathcal{S})$ accepted by $\mathcal{A}_{\neg\varphi}$. In this section, we provide an algorithm for checking this last condition. In particular, we establish the following result.

Theorem 3.3. *Given an open PMS \mathcal{S} over AP , a finite set B disjoint from AP , and a parity ACG \mathcal{A} over $2^{AP \cup B}$, checking that there are no B -labeling extensions of the environment strategy trees of $\mathcal{G}(\mathcal{S})$ accepted by \mathcal{A} can be done in time doubly exponential in the size of \mathcal{A} and singly exponential in the size of \mathcal{S} .*

Thus, by Theorem 3.2 and Theorem 3.3, and since the pushdown module-checking problem against **CTL** is already 2EXPTIME-complete, and EXPTIME-complete for a fixed **CTL** formula [BMP10], we obtain the following corollary.

Corollary 3.4. *Pushdown module-checking for ATL^* is in 4EXPTIME while pushdown module-checking for ATL is 2EXPTIME-complete. Moreover, for a fixed ATL^* state formula (resp., ATL formula), the pushdown module-checking problem is EXPTIME-complete.*

In Section 4, we provide a lower bound for ATL^* matching the upper bound in the corollary above. We present now the proof of Theorem 3.3 which is based on a reduction to the emptiness problem of parity **NPTA**. Given an open PMS \mathcal{S} over AP and a parity ACG \mathcal{A} over $2^{AP \cup B}$, we construct in single exponential time a parity **NPTA** \mathcal{P} over $2^{AP \cup B}$ accepting the B -labeling extensions of suitable encodings of the environment strategy trees of $\mathcal{G}(\mathcal{S})$ which are accepted by \mathcal{A} . Since the set B just occurs in the input alphabet $2^{AP \cup B}$ and the behaviour of $\mathcal{G}(\mathcal{S})$ does not depend on B , for simplicity and without loss of generality, we assume that the set B in the statement of Theorem 3.3 is empty.

Encoding of environment strategy trees of open PMS. Let us fix an open PMS $\mathcal{S} = \langle Q, \Gamma \cup \{\gamma_0\}, q_0, Lab, \Delta \rangle$ over AP , and let $\mathcal{G}(\mathcal{S}) = \langle S, s_0, Lab_S, \tau \rangle$. For all pairs $(q, \gamma) \in Q \times (\Gamma \cup \{\gamma_0\})$, we denote by $next_{\mathcal{S}}(q, \gamma)$ the finite set of pairs $(q', \beta) \in Q \times \Gamma^*$ such that there is a full decision d so that $\Delta(q, \gamma, d) = (q', \beta)$. We fix an ordering on the set $next_{\mathcal{S}}(q, \gamma)$ which induces an ordering on the finite set of successors of all the configurations of the form $(q, \gamma \cdot \alpha)$. Moreover, we consider the parameter $k_{\mathcal{S}} = \max\{|next_{\mathcal{S}}(q, \gamma)| \mid (q, \gamma) \in Q \times (\Gamma \cup \{\gamma_0\})\}$ which represents the finite branching degree of $Unw(\mathcal{G}(\mathcal{S}))$. Thus, we can encode each track $\nu = s_0, s_1, \dots, s_n$ of $\mathcal{G}(\mathcal{S})$ starting from the initial state, by the finite word i_1, \dots, i_n over $\{1, \dots, k_{\mathcal{S}}\}$ of length n where for all $1 \leq h \leq n$, i_h represents the index of state s_h in the ordered set of successors of state s_{h-1} . Now, we observe that the transition function τ' of an environment strategy tree $\mathcal{T} = \langle T, \varepsilon, Lab', \tau' \rangle$ of $\mathcal{G}(\mathcal{S})$ is completely determined by T and the transition function τ of $\mathcal{G}(\mathcal{S})$. Hence, for the fixed open CGS $\mathcal{G}(\mathcal{S})$, \mathcal{T} can be simply specified by the underlying 2^{AP} -labeled S -tree $\langle T, Lab' \rangle$.

We consider an equivalent representation of $\langle T, Lab' \rangle$ by a $(2^{AP} \cup \{\perp\})$ -labeled *complete* $k_{\mathcal{S}}$ -tree $\langle \{1, \dots, k_{\mathcal{S}}\}^*, Lab_{\perp} \rangle$, called the \perp -completion encoding of \mathcal{T} (\perp is a fresh proposition), where the labeling Lab_{\perp} is defined as follows for each node $x \in \{1, \dots, k_{\mathcal{S}}\}^*$:

- if x encodes a track $s_0 \cdot \nu$ such that ν is a node of T , then $Lab_{\perp}(x) = Lab'(\nu)$ (*concrete nodes*);
- otherwise, $Lab_{\perp}(x) = \{\perp\}$ (*completion nodes*).

In this way, all the labeled trees encoding environment strategy trees \mathcal{T} of $\mathcal{G}(\mathcal{S})$ have the same structure (they all coincide with $\{1, \dots, k_{\mathcal{S}}\}^*$), and they differ only in their labeling. Thus, the proposition \perp is used to denote both “completion” nodes and nodes in $Unw(\mathcal{G}(\mathcal{S}))$ which are absent in \mathcal{T} (corresponding to possible disabling of environment choices).

Proof of Theorem 3.3. We now prove Theorem 3.3 for the case $B = \emptyset$. We establish by Theorem 3.5 below that given an open PMS \mathcal{S} and a parity ACG $\mathcal{A} = \langle Q_{\mathcal{A}}, q_{\mathcal{A}}^0, \delta, \Omega \rangle$ over 2^{AP} , one can build a parity NPTA \mathcal{P} accepting the \perp -completion encodings of the environment strategy trees of $\mathcal{G}(\mathcal{S})$ which are accepted by \mathcal{A} . Thus, checking that there are no environment strategy trees of $\mathcal{G}(\mathcal{S})$ accepted by \mathcal{A} reduces to emptiness of the language accepted by \mathcal{P} . Moreover, the size of \mathcal{P} is polynomial in the size of \mathcal{S} and singly exponential in the size of \mathcal{A} . Hence, by Proposition 3.1, Theorem 3.3 for the case $B = \emptyset$ directly follows.

Theorem 3.5. *Given an open PMS $\mathcal{S} = \langle Q, \Gamma \cup \{\gamma_0\}, q_0, \text{Lab}, \Delta \rangle$ over AP and a parity ACG $\mathcal{A} = \langle Q_{\mathcal{A}}, q_{\mathcal{A}}^0, \delta, \Omega \rangle$ over 2^{AP} with index h , one can build in single exponential time, a parity NPTA \mathcal{P} accepting the set of $2^{AP} \cup \{\perp\}$ -labeled complete $k_{\mathcal{S}}$ -trees which are the \perp -completion encodings of the environment strategy trees of $\mathcal{G}(\mathcal{S})$ which are accepted by \mathcal{A} . Moreover, \mathcal{P} has index $O(h|\mathcal{A}|^2)$, number of states $O(|Q| \cdot (h|\mathcal{A}|^2)^{O(h|\mathcal{A}|^2)})$, and transition function ρ such that $||\rho|| = O(|\Delta| \cdot (h|\mathcal{A}|^2)^{O(h|\mathcal{A}|^2)})$.*

Proof. First, we observe that for the given parity ACG \mathcal{A} and an input CGT \mathcal{T} , we can associate in a standard way to \mathcal{A} and \mathcal{T} an infinite-state two player parity game, where player 0 plays for acceptance, while player 1 plays for rejection. Winning strategies of player 0 correspond to accepting runs of \mathcal{A} over \mathcal{T} . Thus, since the existence of a winning strategy in parity games implies the existence of a memoryless one, we can restrict ourselves to consider only memoryless runs of \mathcal{A} , i.e. runs $r = \langle T_r, \text{Lab}_r \rangle$ where the behavior of \mathcal{A} along r depends only on the current input node and current state. Formally, r is memoryless if for all nodes y and y' of r having the same label, the subtrees rooted at the nodes y and y' of r are isomorphic. We now provide a representation of the memoryless runs of \mathcal{A} over the environment strategy trees of the open CGS $\mathcal{G}(\mathcal{S})$ induced by the given open PMS \mathcal{S} .

Fix an environment strategy tree $\mathcal{T} = \langle T, \varepsilon, \text{Lab}_{\mathcal{T}}, \tau \rangle$ of $\mathcal{G}(\mathcal{S})$ and let $\langle \{1, \dots, k_{\mathcal{S}}\}^*, \text{Lab}_{\perp} \rangle$ be the \perp -completion encoding of \mathcal{T} . Recall that $\text{Atoms}(\mathcal{A})$ is the set of atoms of \mathcal{A} , i.e., the set of tuples in $Q_{\mathcal{A}} \times \{\square, \diamond\} \times 2^{Ag}$ occurring in the transition function δ of \mathcal{A} .

Let $\text{Ann} := 2^{Q_{\mathcal{A}} \times \text{Atoms}(\mathcal{A})}$ be the finite set of annotations and $\Sigma := (2^{AP} \times \text{Ann} \times \text{Ann}) \cup \{\perp\}$. In the following, a *move* is an element in $Q_{\mathcal{A}} \times \text{Atoms}(\mathcal{A})$. For an annotation $an \in \text{Ann}$, we define the following finite sets:

- $\text{Dom}(an)$ is the set of \mathcal{A} -states q such that $(q, atom) \in an$ for some $atom \in \text{Atoms}(\mathcal{A})$;
- $\text{Cod}(an)$ is the set of \mathcal{A} -states occurring in the atoms of an ;
- For each state $q \in Q_{\mathcal{A}}$, $\text{Atoms}(q, an)$ the set of atoms $atom$ such that $(q, atom) \in an$.

For example, if $an = \{(q_1, (q'_1, \diamond, A_1)), (q_2, (q'_2, \square, A_2))\}$, then $\text{Dom}(an) = \{q_1, q_2\}$, $\text{Cod}(an) = \{q'_1, q'_2\}$, and $\text{Atoms}(q_1, an) = \{(q'_1, \diamond, A_1)\}$.

We represent memoryless runs r of \mathcal{A} over \mathcal{T} as *annotated extensions* of the \perp -completion encoding $\langle \{1, \dots, k_{\mathcal{S}}\}^*, \text{Lab}_{\perp} \rangle$ of \mathcal{T} , i.e., Σ -labeled complete $k_{\mathcal{S}}$ -trees $\langle \{1, \dots, k_{\mathcal{S}}\}^*, \text{Lab}_{\Sigma} \rangle$, where:

(R1) for every concrete node $x \in \{1, \dots, k_{\mathcal{S}}\}^*$ encoding a node ν_x of \mathcal{T} , $\text{Lab}_{\Sigma}(x)$ is of the form $(\text{Lab}_{\perp}(x), an, an')$ (recall that $\text{Lab}_{\perp}(x) = \text{Lab}_{\mathcal{T}}(\nu_x)$), and for every completion node x , $\text{Lab}_{\Sigma}(x) = \text{Lab}_{\perp}(x) = \{\perp\}$.

Intuitively, the meaning of the first annotation an and the second annotation an' in the label of a concrete node x is as follows:

- $\text{Dom}(an)$ represents the set of \mathcal{A} -states q associated with the copies of \mathcal{A} in the run r which read the input node ν_x of \mathcal{T} , while for each $q \in \text{Dom}(an)$, $\text{Atoms}(q, an)$ represents

the model of $\delta(q, \text{Lab}_T(\nu_x))$ selected by \mathcal{A} in r on reading node ν_x in state q . Note that $\text{Cod}(an)$ represents the set of target states of the moves in an .

- Additionally, the second annotation an' in the labeling of node x keeps tracks, in case x is not the root, of the subset of the moves in the first annotation of the parent ν' of ν_x in T for which, starting from ν' , a copy of \mathcal{A} is sent to the current node ν_x along r .

Formally, for the concrete node x with label $(\text{Lab}_\perp(x), an, an')$, we require that the following requirements hold, where x_1, \dots, x_N denote the concrete children of node x in $\langle \{1, \dots, k_S\}^*, \text{Lab}_\perp \rangle$ encoding the children ν_1, \dots, ν_N of node ν_x in T , and $\text{Lab}_\Sigma(x_i) = (\text{Lab}_\perp(x_i), an_i, an'_i)$ for each $1 \leq i \leq N$:

- (R2) for each $q \in \text{Dom}(an)$, $\text{Atoms}(q, an)$ is a model of $\delta(q, \text{Lab}_T(\nu_x))$;
- (R3) the two annotations an and an' are consistent, i.e., $an' = \emptyset$ if x is the root and $\text{Cod}(an') = \text{Dom}(an)$ otherwise; moreover, if x is the root, then $\text{Dom}(an) = \{q_A^0\}$ (q_A^0 is the initial state of \mathcal{A});
- (R4) for all $(q, (q', m, A)) \in an$, the following holds:
 - case $m = \square$: there is an available A -decision d_A in the node ν_x of T such that for all $i \in \{1, \dots, N\}$ so that ν_i is a child of ν_x in T consistent with the A -decision d_A , it holds that $(q, (q', m, A)) \in an'_i$;
 - case $m = \diamond$: for all available A -decisions d_A in the node ν_x of T , there is some $i \in \{1, \dots, N\}$ such that the child ν_i of ν_x in T is consistent with the A -decision d_A and $(q, (q', m, A)) \in an'_i$.
- (R5) $an = \bigcup_{i=1}^{i=N} an'_i$.

An *annotated extension* $\langle \{1, \dots, k_S\}^*, \text{Lab}_\Sigma \rangle$ of $\langle \{1, \dots, k_S\}^*, \text{Lab}_\perp \rangle$ is *well-formed* if it satisfies requirements R1–R5 expressed above. We deduce the following result.

Claim 1. One can construct in singly exponential time a parity NPTA \mathcal{P}_{wf} over Σ -labeled complete k_S -trees accepting the set of *well-formed* annotated extensions of the \perp -completion encodings of the environment strategy trees of $\mathcal{G}(\mathcal{S})$. Moreover, \mathcal{P}_{wf} has number of states $O(|Q| \cdot 2^{O(|Q_A| \cdot |\text{Atoms}(\mathcal{A})|)})$, index 1, and transition function ρ such that $||\rho|| = O(|\Delta|)$.

The proof of Claim 1 is postponed at the end of the proof of Theorem 3.5.

Note that the well-formedness requirement just ensures that the annotated extension $\langle \{1, \dots, k_S\}^*, \text{Lab}_\Sigma \rangle$ of $\langle \{1, \dots, k_S\}^*, \text{Lab}_\perp \rangle$ encodes a memoryless run r of the ACG \mathcal{A} over the input T . In order to ensure that $\langle \{1, \dots, k_S\}^*, \text{Lab}_\Sigma \rangle$ encodes a run r which is also accepting, we need to enforce additional *global* requirements on the annotated extension $\langle \{1, \dots, k_S\}^*, \text{Lab}_\Sigma \rangle$.

Let π be an infinite path of $\langle \{1, \dots, k_S\}^*, \text{Lab}_\Sigma \rangle$ from the root which does not visit \perp -labeled nodes. Then, $\text{Lab}_\Sigma(\pi)$ “collects” all the infinite sequences ν of states in Q_A along the run r associated with the input path of the environment strategy tree T encoded by π . In order to check the acceptance condition on the individual parallel paths ν , the infinite sequence of annotations $\text{Lab}_\Sigma(\pi)$ must allow to distinguish the individual infinite paths over Q_A grouped by $\text{Lab}_\Sigma(\pi)$. This is because we exploit the second annotation an' in the labeling $(\text{Lab}_\perp(x), an, an')$ of a concrete node x . In particular, the individual paths over Q_A grouped by $\text{Lab}_\Sigma(\pi)$ correspond to the so-called Q_A -paths of $\text{Lab}_\Sigma(\pi)$ which are defined as follows.

For all $i \geq 0$, let $Lab_{\Sigma}(\pi(i)) = (\sigma_i, an_i, an'_i)$. Then, a $Q_{\mathcal{A}}$ -path of $Lab_{\Sigma}(\pi)$ is an infinite sequence $q_0q_1\dots$ of $Q_{\mathcal{A}}$ -states such that for all $i \geq 0$, $q_i \in Dom(an_i)$ and $(q_i, (q_{i+1}, m, A)) \in an_i \cap an'_{i+1}$ for some $m \in \{\square, \diamond\}$ and set A of agents.

We need to check that all these $Q_{\mathcal{A}}$ -paths satisfy the acceptance parity condition of \mathcal{A} . To this end, we construct a standard parity nondeterministic tree automaton (parity NTA) \mathcal{A}_{acc} over Σ -labeled complete $k_{\mathcal{S}}$ -trees which accepts an input tree if all the $Q_{\mathcal{A}}$ -paths associated with the infinite paths of the input tree starting at the root satisfy the acceptance parity condition of \mathcal{A} . In order to construct \mathcal{A}_{acc} , we proceed as follows.

We first easily construct a co-parity nondeterministic word automaton \mathcal{B} over Σ with $O(|Q_{\mathcal{A}}| \cdot |Atoms(\mathcal{A})|)$ states and index h (the index of \mathcal{A}) which accepts an infinite word over Σ iff it contains a $Q_{\mathcal{A}}$ -path that does not satisfy the parity acceptance condition of \mathcal{A} . Recall that a co-parity condition over a set Q' of states is defined as a parity condition over Q' , i.e., a mapping of the form $Q' \mapsto \mathbb{N}$, but for acceptance of an infinite sequence of states ρ , we require that the highest color of the states appearing infinitely often along ρ is *odd*. Formally, the co-parity word automaton \mathcal{B} with input alphabet Σ is given by $\mathcal{A} = \langle Q_{\mathcal{B}}, q_{\mathcal{B}}^0, \delta_{\mathcal{B}}, \Omega_{\mathcal{B}} \rangle$, where

- $Q_{\mathcal{B}} = \{q_{\mathcal{B}}^0\} \cup (Q_{\mathcal{A}} \times Atoms(\mathcal{A}))$.
- The transition function $\delta_{\mathcal{B}} : Q_{\mathcal{B}} \times \Sigma \mapsto 2^{Q_{\mathcal{B}}}$ is defined as follows for all $(\sigma, an, an') \in \Sigma$:
 - $\delta_{\mathcal{B}}(q_{\mathcal{B}}^0, (\sigma, an, an')) = an$;
 - for all $(q, (q', m, A)) \in Q_{\mathcal{A}} \times Atoms(\mathcal{A})$, we have that $\delta_{\mathcal{B}}((q, (q', m, A)), (\sigma, an, an')) = \emptyset$ if $(q, (q', m, A)) \notin an'$, and $\delta_{\mathcal{B}}((q, (q', m, A)), (\sigma, an, an')) = an$ otherwise.
- $\Omega_{\mathcal{B}}(q_{\mathcal{B}}^0) = 0$ and $\Omega_{\mathcal{B}}((q, (q', m, A))) = \Omega(q)$ for all $(q, (q', m, A)) \in Q_{\mathcal{A}} \times Atoms(\mathcal{A})$.

By construction, given an input $w \in \Sigma^{\omega}$ with $w = (\sigma_0, an_0, an'_0)(\sigma_1, an_1, an'_1)\dots$, \mathcal{B} accepts w iff there exists a run of \mathcal{B} of the form

$$q_{\mathcal{B}}^0 \cdot (q_0, (q'_0, m_0, A_0)) \cdot (q_1, (q'_1, m_1, A_1)) \dots$$

such that $(q_i, (q'_i, m_i, A_i)) \in an_i \cap an'_{i+1}$ for all $i \geq 0$, and the infinite sequence of \mathcal{A} -states $q_0q_1\dots$ does not satisfy the parity acceptance condition of \mathcal{A} . Hence, \mathcal{B} accepts w iff w contains a $Q_{\mathcal{A}}$ -path that does not satisfy the parity acceptance condition of \mathcal{A} . We now co-determinize the co-parity nondeterministic word automaton \mathcal{B} , i.e., determinize it and complement it in a singly-exponential construction [Saf88] to obtain a *deterministic* parity word automaton \mathcal{B}' that rejects violating $Q_{\mathcal{A}}$ -paths. By [Saf88], \mathcal{B}' has $(nh)^{O(nh)}$ states and index $O(nh)$, where $n = |Q_{\mathcal{A}}| \cdot |Atoms(\mathcal{A})|$. Then the parity NTA \mathcal{A}_{acc} is obtained from \mathcal{B}' by simply running \mathcal{B}' in parallel over all the branches of the input which do not visit a \perp -labeled node. Note that \mathcal{A}_{acc} has $(nh)^{O(nh)}$ states and index $O(nh)$.

Finally, the parity NPTA \mathcal{P} satisfying Theorem 3.5 is obtained by projecting out the annotation components of the input trees accepted by the intersection of the NPTA \mathcal{P}_{wf} of index 1 in Claim 1 with the parity NTA \mathcal{A}_{acc} (recall that parity NPTA are effectively and polynomial-time closed under projection and intersection with nondeterministic tree automata [KPV02]).

In order to conclude the proof of Theorem 3.5 it remains to prove Claim 1.

Proof of Claim 1. In order to define the NPTA \mathcal{P}_{wf} satisfying Claim 1, we need additional definitions which allow to express requirements R4 and R5 in terms of the transition function of the PMS $\mathcal{S} = \langle Q, \Gamma \cup \{\gamma_0\}, q_0, Lab, \Delta \rangle$.

Let $(p, \gamma) \in Q \times (\Gamma \cup \{\gamma_0\})$ with $\text{next}_{\mathcal{S}}(p, \gamma) = \{(p_1, \beta_1), \dots, (p_k, \beta_k)\}$ for some $1 \leq k \leq k_{\mathcal{S}}$. For an annotation an and a tuple $\langle an_1, \dots, an_k \rangle$ of k annotations, we say that $\langle an_1, \dots, an_k \rangle$ is consistent with annotation an and the pair (p, γ) if the following conditions are fulfilled:

- $an = \bigcup_{i=1}^{i=k} an_i$;
- for each move $\eta = (q, (q', m, A)) \in an$, let X_{η} be the subset of $\{(p_1, \beta_1), \dots, (p_k, \beta_k)\}$ consisting of the pairs (p_i, β_i) such that $\eta \in an_i$. Then, $X_{\eta} \neq \emptyset$ and the following holds:
 - case $m = \square$: there is an A -decision d_A available in (p, γ) (i.e., $\Delta(p, \gamma, d) \neq \perp$ for some full decision d consistent with d_A) such that for each full decision d consistent with d_A , either $\Delta(p, \gamma, d) = \perp$ or $\Delta(p, \gamma, d) \in X_{\eta}$;
 - case $m = \Diamond$: for all available A -decisions d_A available in (p, γ) , there is some full decision d consistent with d_A so that $\Delta(p, \gamma, d) \in X_{\eta}$.

We denote by $Cons(p, \gamma, an)$ the set of tuples $\langle an_1, \dots, an_k \rangle$ of k annotations which are consistent with the annotation an and the pair (p, γ) .

Intuitively, the previous definition allows to express requirements R4 and R5 in case all the possible choices are enabled. More precisely, in case the concrete node ν_x of the environment strategy tree \mathcal{T} in conditions R4 and R5 is associated with a configuration of the form $(p, \gamma \cdot \beta)$ and all the children of ν_x in $Unw(\mathcal{G}(\mathcal{S}))$ are also children of ν_x in \mathcal{T} , then R4 and R5 are equivalent to the following condition: $N = k$ and $\langle an'_1, \dots, an'_k \rangle \in Cons(p, \gamma, an)$.

Now assume that not all the children of node ν_x in $Unw(\mathcal{G}(\mathcal{S}))$ are also children of ν_x in \mathcal{T} . This entails that node ν_x is associated with an environment configuration $(p, \gamma \cdot \beta)$. Since $(p, \gamma \cdot \beta)$ is controlled by the environment, we observe that $\langle an_1, \dots, an_k \rangle \in Cons(p, \gamma, an)$ iff the following conditions hold:

- $an = \bigcup_{i=1}^{i=k} an_i$;
- for each move $\eta = (q, (q', m, A)) \in an$, the following holds:
 - if either $m = \Diamond$ and $env \in A$, or $m = \square$ and $env \notin A$, then $\eta \in an_i$ for some $i \in \{1, \dots, k\}$;
 - otherwise, $\eta \in an_i$ for each $i \in \{1, \dots, k\}$.

This justifies the following definition: an annotation an is *obligation-free* if it does not contain moves of the form $(q, (q', m, A))$ such that either $m = \Diamond$ and $env \in A$, or $m = \square$ and $env \notin A$. Then, in case the concrete node ν_x in \mathcal{T} is associated with an environment configuration $(p, \gamma \cdot \beta)$, requirements R4 and R5 are equivalent to the following condition.

- $N \leq k$ and there are distinct indexes $i_1, \dots, i_N \in \{1, \dots, k\}$ such that the following holds for some $\langle an''_1, \dots, an''_k \rangle \in Cons(p, \gamma, an)$: $an'_j = an''_{i_j}$ for all $1 \leq j \leq N$, and (ii) an''_{ℓ} is obligation-free for all $\ell \in \{1, \dots, k\} \setminus \{i_1, \dots, i_N\}$.

By using the previous definitions and observations, we now define the parity NPTA \mathcal{P}_{wf} of index 1 satisfying Claim 1. Essentially, given a Σ -labeled complete $k_{\mathcal{S}}$ -tree $\langle \{1, \dots, k_{\mathcal{S}}\}^*, Lab_{\Sigma} \rangle$, the automaton \mathcal{P}_{wf} , by simulating the behaviour of the open PMS \mathcal{S} and by exploiting the transition function of the parity ACG \mathcal{A} , checks that the input is a *well-formed* annotated extension of the \perp -completion encoding of some environment strategy tree of $\mathcal{G}(\mathcal{S})$. Formally, the NPTA $\mathcal{P}_{wf} = \langle \Sigma, P, \Gamma \cup \{\gamma_0\}, p_0, \rho, \Omega : p \in P \mapsto \{0\} \rangle$ is defined as follows.

The set P of states consists of the triples (q, an, m) where $q \in Q$ is a state of the PMS \mathcal{S} , $an \in Ann$ is an annotation, and $m \in \{\perp, \top, \vdash\}$ is a state marker such that $an = \emptyset$ if $m = \perp$. Intuitively, whenever the current input node x is a concrete node, then q represents the state of \mathcal{S} associated with node x . The meaning of the state marker m is as follows. When the state marker m is \perp , the NPTA \mathcal{P}_{wf} can read only the letter \perp , while when the

state marker is \top , \mathcal{P}_{wf} can read only letters in $\Sigma \setminus \{\perp\}$. Finally, when \mathcal{P}_{wf} is in states of the form (q, an, \vdash) , then it can read both letters in $\Sigma \setminus \{\perp\}$ and the letter \perp . In this case, it is left to the environment to decide whether the transition to a configuration of the simulated PMS \mathcal{S} of the form (q, β) is enabled. Intuitively, the three types of states are used to ensure that the environment enables all transitions from enabled system configurations, enables at least one transition from each enabled environment configuration, and disables transitions from disabled configurations. Moreover, the annotation an in a control state (q, an, m) of \mathcal{P}_{wf} represents the guessed subset of the moves in the first annotation of the parent x' (if any) of the current *concrete* input node for which, starting from x' , a copy of \mathcal{A} is sent to the current input node: in the transition function, we require that in case the current input symbol σ is not \perp , an coincides with the second annotation of σ .

The initial state p_0 is given by (q_0, \emptyset, \top) . Finally, the transition function $\rho : P \times \Sigma \times (\Gamma \cup \{\gamma_0\}) \rightarrow 2^{(P \times \Gamma^*)^{k_S}}$ is defined as follows. According to the definition of P , the automaton \mathcal{P}_{wf} can be in a state of the form (q, \emptyset, \perp) , (q, an, \top) , or (q, an, \vdash) . Both in the first and the third cases, \mathcal{P}_{wf} can read \perp , which means that the automaton is reading a disabled or a completion node. Thus, independently from the fact that the actual configuration of the automaton is associated with an environment or a system configuration of the open PMS \mathcal{S} , ρ propagates states of the form (q, \emptyset, \perp) to all children of the reading node. Note that for states of the form (q, an, \vdash) and in case \mathcal{P}_{wf} reads \perp , we require that an is obligation-free. If instead the automaton is in a state of the form (q, an, \top) or (q, an, \vdash) and reads a label different from \perp , the possible successor states further depend on the particular kind of the configuration in which the automaton is. If \mathcal{P}_{wf} is in a system configuration of \mathcal{S} , then all the children of the reading node associated with the successors of such a configuration in the CGS $\mathcal{G}(\mathcal{S})$ must not be disabled and so, ρ sends to all of them states with marker \top . If \mathcal{P}_{wf} is in an environment configuration of \mathcal{S} , then all the children of the reading node, but one, associated with the successors of such a configuration in the CGS $\mathcal{G}(\mathcal{S})$ may be disabled and so, ρ sends to all of them states with marker \vdash , except one, to which ρ sends a state with marker \top .

Formally, let $(q, an, m) \in P$, $\sigma \in \Sigma$, and $\gamma \in \Gamma \cup \{\gamma_0\}$ with $next_{\mathcal{S}}(q, \gamma) = \langle (q_1, \beta_1), \dots, (q_k, \beta_k) \rangle$ ($1 \leq k \leq k_S$). Then, $\rho((q, an, m), \sigma, \gamma)$ is defined as follows:

- Case $m \in \{\perp, \vdash\}$, $\sigma = \perp$, and an is obligation-free:

$$\rho((q, an, m), \perp, \gamma) = \{ \langle \underbrace{((q, \emptyset, \perp), \varepsilon), \dots, ((q, \emptyset, \perp), \varepsilon)}_{k_S \text{ pairs}} \rangle \}$$

That is, $\rho((q, an, m), \perp, \gamma)$ contains exactly one k_S -tuple. In this case all the successors of the current \mathcal{S} -configuration are disabled.

- Case $m \in \{\top, \vdash\}$, (q, γ) is associated with *system* \mathcal{S} -configurations, $\sigma = (Lab(q), an', an)$ for some annotation an' such that $Cod(an) = Dom(an')$, and for each $q_{\mathcal{A}} \in Dom(an')$, $Atoms(q_{\mathcal{A}}, an')$ is a model of $\delta(q_{\mathcal{A}}, Lab(q))$:

$$\begin{aligned} \rho((q, an, m), \sigma, \gamma) = \bigcup_{(an_1, \dots, an_k) \in Cons(q, \gamma, an')} & \{ \langle \langle ((q_1, an_1, \top), \beta_1), \dots, ((q_k, an_k, \top), \beta_k), \\ & \underbrace{((q, \emptyset, \perp), \varepsilon), \dots, ((q, \emptyset, \perp), \varepsilon)}_{k_S - k \text{ pairs}} \rangle \} \end{aligned}$$

In this case, all the k successors of the current system \mathcal{S} -configuration are enabled. Moreover, the automaton guesses a tuple $\langle an_1, \dots, an_k \rangle$ of k annotations which are

consistent with the first annotation an' of the input node and the pair (q, γ) , and sends state (q_i, an_i, \top) to the i th child of the current input node for all $1 \leq i \leq k$.

- Case $m \in \{\top, \vdash\}$, (q, γ) is associated with *environment \mathcal{S} -configurations*, $\sigma = (Lab(q), an', an)$ for some annotation an' such that $Cod(an) = Dom(an')$, and for each $q_{\mathcal{A}} \in Dom(an')$, $Atoms(q_{\mathcal{A}}, an')$ is a model of $\delta(q_{\mathcal{A}}, Lab(q))$. In this case $\rho((q, an, m), \sigma, \gamma)$ is defined as follows:

$$\bigcup_{(an_1, \dots, an_k) \in Cons(q, \gamma, an')} \{ \langle \langle (q_1, an_1, \top), \beta_1 \rangle, \langle (q_2, an_2, \vdash), \beta_2 \rangle, \dots, \langle (q_k, an_k, \vdash), \beta_k \rangle, \langle (q, \emptyset, \perp), \varepsilon \rangle, \dots, \langle (q, \emptyset, \perp), \varepsilon \rangle \rangle, \langle \langle (q_1, an_1, \vdash), \beta_1 \rangle, \langle (q_2, an_2, \top), \beta_2 \rangle, \dots, \langle (q_k, an_k, \vdash), \beta_k \rangle, \langle (q, \emptyset, \perp), \varepsilon \rangle, \dots, \langle (q, \emptyset, \perp), \varepsilon \rangle \rangle, \vdots, \langle \langle (q_1, an_1, \vdash), \beta_1 \rangle, \langle (q_2, an_2, \vdash), \beta_2 \rangle, \dots, \langle (q_k, an_k, \top), \beta_k \rangle, \langle (q, \emptyset, \perp), \varepsilon \rangle, \dots, \langle (q, \emptyset, \perp), \varepsilon \rangle \rangle \} \}$$

Thus, the automaton guesses a tuple $\langle an_1, \dots, an_k \rangle$ of k annotations which is consistent with the first annotation an' of the input node and the pair (q, γ) and, additionally, guesses an index $1 \leq i \leq k$. With these choices, the automaton sends state (q_i, an_i, \top) to the i th child of the current input node and, additionally, ensures that the i th successor of the current environment \mathcal{S} -configuration is enabled while all the other successors may be disabled.

- All the other cases: $\rho((q, an, m), \sigma, \gamma) = \emptyset$.

Note that \mathcal{P}_{wf} has $O(|Q| \cdot 2^{O(|Q_{\mathcal{A}}| \cdot |Atoms(\mathcal{A})|)})$ states, $||\rho|| = O(|\Delta|)$, and $|\rho| = O(|\Delta| \cdot 2^{O(k_{\mathcal{S}} \cdot |Q_{\mathcal{A}}| \cdot |Atoms(\mathcal{A})|)})$. By construction, it easily follows that a Σ -labeled complete $k_{\mathcal{S}}$ -tree $\langle \{1, \dots, k_{\mathcal{S}}\}^*, Lab_{\Sigma} \rangle$ is accepted by the NPTA \mathcal{P}_{wf} iff there is an environment strategy tree \mathcal{T} of $\mathcal{G}(\mathcal{S})$ such that $\langle \{1, \dots, k_{\mathcal{S}}\}^*, Lab_{\Sigma} \rangle$ is some *well-formed* annotated extension of the \perp -completion encoding of \mathcal{T} . This concludes the proof of Claim 1. \square

4. 4EXPTIME-HARDNESS OF ATL* PUSHDOWN MODULE-CHECKING

In this section, we establish the following result.

Theorem 4.1. *Pushdown module-checking against ATL* is 4EXPTIME-hard even for two-player turn-based PMS.*

Theorem 4.1 is proved by a polynomial-time reduction from the acceptance problem for 3EXPSPACE-bounded Alternating Turing Machines (ATM, for short) with a binary branching degree. Formally, such a machine is a tuple $\mathcal{M} = \langle \Sigma, Q, Q_{\forall}, Q_{\exists}, q_0, \delta, F \rangle$, where Σ is the input alphabet which contains the blank symbol $\#$, Q is the finite set of states which is partitioned into $Q = Q_{\forall} \cup Q_{\exists}$, Q_{\exists} (resp., Q_{\forall}) is the set of existential (resp., universal) states, q_0 is the initial state, $F \subseteq Q$ is the set of accepting states, and the transition function δ is a mapping $\delta : Q \times \Sigma \rightarrow (Q \times \Sigma \times \{\leftarrow, \rightarrow\})^2$. Note that since \mathcal{M} has a binary branching degree, the transition function δ nondeterministically associates to each pair state/input symbol (q, σ) two possible moves, where each move is represented by a triple (q', σ', d) consisting of a target state q' , the symbol σ' to write in the tape cell currently pointed by the reading head, and a symbol $d \in \{\leftarrow, \rightarrow\}$ encoding the movement of the reading head: \leftarrow (resp., \rightarrow) means that the reading head moves one cell to the left (resp., to the right) of the current cell.

Formally, configurations of \mathcal{M} are words in $\Sigma^* \cdot (Q \times \Sigma) \cdot \Sigma^*$. A configuration $C = \eta \cdot (q, \sigma) \cdot \eta'$ denotes that the tape content is $\eta \cdot \sigma \cdot \eta'$, the current state (resp., current input symbol) is

q (resp., σ), and the reading head is at position $|\eta| + 1$. From a configuration C , the machine \mathcal{M} nondeterministically chooses a triple (q', σ', d) in $\delta(q, \sigma) = \langle (q_l, \sigma_l, d_l), (q_r, \sigma_r, d_r) \rangle$, and then moves to state q' , writes σ' in the current tape cell, and its reading head moves one cell to the left or to the right, according to d . We denote by $\text{succ}_l(C)$ and $\text{succ}_r(C)$ the successors of C obtained by choosing respectively the left and the right triple in $\langle (q_l, \sigma_l, d_l), (q_r, \sigma_r, d_r) \rangle$ (note that the terms ‘left’ and ‘right’ here should not be confused with the movement of the reading head of the ATM). The configuration C is accepting (resp., universal, resp., existential) if the associated state q is in F (resp., in Q_{\forall} , resp., in Q_{\exists}).

Given an input $\alpha \in \Sigma^+$, a (finite) computation tree of \mathcal{M} over α is a finite tree in which each node is labeled by a configuration. The root of the tree is labeled by the initial configuration $(q_0, \alpha(0))\alpha(1)\dots\alpha(n-1)$ associated with α , where $n = |\alpha|$. An *internal* node that is labeled by a universal configuration C has two children, corresponding to $\text{succ}_l(C)$ and $\text{succ}_r(C)$, while an internal node labeled by an existential configuration C has a single child, corresponding to either $\text{succ}_l(C)$ or $\text{succ}_r(C)$. The tree is accepting if each leaf is labeled by an accepting configuration. An input $\alpha \in \Sigma^+$ is *accepted* by \mathcal{M} if there is an accepting computation tree of \mathcal{M} over α .

If the ATM \mathcal{M} is 3EXPSPACE-bounded, then there is a constant $c \geq 1$ such that for each $\alpha \in \Sigma^+$, the space needed by \mathcal{M} on input α is bounded by $\text{Tower}(|\alpha|^c, 3)$, where for all $n, h \in \mathbb{N}$, $\text{Tower}(n, h)$ denotes a tower of exponentials of height h and argument n (i.e, $\text{Tower}(n, 0) = n$ and $\text{Tower}(n, h+1) = 2^{\text{Tower}(n, h)}$). It is well-known [CKS81] that the acceptance problem for 3EXPSPACE-bounded ATM (with a binary branching degree) is 4EXPTIME-complete.

Fix a 3EXPSPACE-bounded ATM \mathcal{M} and an input $\alpha \in \Sigma^+$. Let $n = |\alpha|$. W.l.o.g. we assume that the constant c is 1 and $n > 1$. Hence, any reachable configuration of \mathcal{M} over α can be seen as a word in $\Sigma^* \cdot (Q \times \Sigma) \cdot \Sigma^*$ of length exactly $\text{Tower}(n, 3)$, and the initial configuration $(q_0, \alpha(0))\alpha(1)\dots\alpha(n-1)$ can be represented as the word of length $\text{Tower}(n, 3)$ given by

$$(q_0, \alpha(0))\alpha(1)\dots\alpha(n-1) \cdot (\#)^t$$

where $t = \text{Tower}(n, 3) - n$. Note that for an ATM configuration $C = u_1 u_2 \dots u_{\text{Tower}(n, 3)}$ and for all $i \in [1, \text{Tower}(n, 3)]$ and $\text{dir} \in \{l, r\}$, the value u'_i of the i -th cell of $\text{succ}_{\text{dir}}(C)$ is completely determined by the values u_{i-1} , u_i and u_{i+1} (taking u_{i+1} for $i = \text{Tower}(n, 3)$ and u_{i-1} for $i = 1$ to be some special symbol, say \vdash). Thus, we denote by $\text{next}_{\text{dir}}(u_{i-1}, u_i, u_{i+1})$ the value u'_i of the i -th cell of $\text{succ}_{\text{dir}}(C)$ (note that the function next_{dir} can be trivially obtained from the transition function of \mathcal{M}). According to the previous observation, we use the set Λ of triples of the form (u_p, u, u_s) where $u \in \Sigma \cup (Q \times \Sigma)$, and $u_p, u_s \in \Sigma \cup (Q \times \Sigma) \cup \{\vdash\}$. We prove the following result from which Theorem 4.1 directly follows.

Theorem 4.2. *One can construct, in time polynomial in n and the size of \mathcal{M} , an open turn-based PMS \mathcal{S} and an ATL* formula φ over the set of agents $\text{Ag} = \{\text{sys}, \text{env}\}$ such that \mathcal{M} accepts α iff there is an environment strategy tree in $\text{exec}(\mathcal{G}(\mathcal{S}))$ that satisfies φ iff $\mathcal{G}(\mathcal{S}) \not\models^r \neg\varphi$.*

The rest of this section is devoted to the proof of Theorem 4.2.

4.1. Encoding of ATM configurations. We first define an encoding of the ATM configurations by using the following set Main of atomic propositions:

$$\text{Main} := \Lambda \cup \{0, 1, \forall, \exists, l, r, \text{acc}\} \cup \{\text{s}_1, \text{s}_2, \text{s}_3, \text{e}_1, \text{e}_2, \text{e}_3\}.$$

In the encoding of an ATM configuration, for each ATM cell, we record the content of the cell, the location (cell number) of the cell on the ATM tape, and the contents of the previous and next cell (if any). In order to encode the cell number, which is a natural number in $[0, \text{Tower}(n, 3) - 1]$, for all $1 \leq h \leq 3$, we define the notions of *h-block* and *well-formed h-block*. For $h = 1, 2$, well-formed *h-blocks* encode integers in $[0, \text{Tower}(n, h) - 1]$, while well-formed 3-blocks encode the cells of ATM configurations. In particular, for $h = 2, 3$, a well-formed *h-block* encoding a natural number $m \in [0, \text{Tower}(n, h) - 1]$ is a sequence of $\text{Tower}(n, h - 1)$ well-formed $(h - 1)$ -blocks, where the i^{th} $(h - 1)$ -block encodes both the value and (recursively) the position of the i^{th} -bit in the binary representation of m .

Formally, for each $1 \leq h \leq 3$, an *h-block* bl is a word of the form

$$\{s_h\} \cdot bl_0 \dots bl_t \cdot \{\tau\} \cdot \{e_h\}, \text{ where}$$

- $t \geq 0$,
- $\tau \in \{0, 1\}$ if $h \neq 3$, and $\tau \in \Lambda$ otherwise (we say that τ is the *content* of bl),
- if $h = 1$, then for all $0 \leq i \leq t$, $bl_i \in \{0, 1\}$,
- if $h > 1$, then for all $0 \leq i \leq t$, bl_i is an $(h - 1)$ -block.

Note that the *h-block* bl is enclosed by the start delimiter s_h and the end delimiter e_h . We say that the *h-block* bl is *well-formed* if the following additional conditions hold:

- Case $h = 1$: $t = n - 1$. In this case, the *number* of bl is the natural number in $[0, \text{Tower}(n, 1) - 1]$ whose binary code is given by $bl_0 \dots bl_t$.
- Case $h > 1$: $t = \text{Tower}(n, h - 1) - 1$ and the $(h - 1)$ -block bl_i is well-formed and has *number* i for each $0 \leq i \leq t$. If bl is well-formed, then the *number* of bl is the natural number in $[0, \text{Tower}(n, h) - 1]$ whose binary code is given by $b_0 \dots b_t$ where b_i is the content of the sub-block bl_i for all $0 \leq i \leq t$.

Example 4.3. Let $n = 2$. In this case $\text{Tower}(n, 2) = 16$ and $\text{Tower}(n, 1) = 4$. Thus, we can encode by well-formed 2-blocks all the integers in $[0, 15]$. For example, let us consider the number 14 whose binary code (using $\text{Tower}(n, 1) = 4$ bits) is given by 0111 (assuming that the first bit is the least significant one). For each $b \in \{0, 1\}$, the well-formed 2-block with content b and number 14 is given by

$$\{s_2\} \{s_1\} \{0\} \{0\} \{e_1\} \{s_1\} \{1\} \{0\} \{1\} \{e_1\} \{s_1\} \{0\} \{1\} \{1\} \{e_1\} \{s_1\} \{1\} \{1\} \{1\} \{e_1\} \{b\} \{e_2\}$$

Note that the 1-sub-blocks also encode the position of each bit in the binary code of 14. Now, let us consider $\tau \in \Lambda$ and $\ell \in [0, 2^{16} - 1]$, and let $b_0 \dots b_{15}$ be the binary code of ℓ . Then, the well-formed 3-block with content τ and number ℓ is given by the word $\{s_3\} bl_0, \dots, bl_{15} \{\tau\} \{e_3\}$, where for each $i \in [0, 15]$, bl_i is the well-formed 2-block having content b_i and number i .

ATM configurations $C = u_1 u_2 \dots u_k$ (note that here we do not require that $k = \text{Tower}(n, 3)$) are then encoded by words w_C of the form

$$w_C = tag_1 \cdot bl_1 \cdot \dots \cdot bl_k \cdot tag_2, \text{ where}$$

- $tag_1 \in \{\{l\}, \{r\}\}$,
- for each $i \in [1, k]$, bl_i is a 3-block whose content is (u_{i-1}, u_i, u_{i+1}) (where $u_0 = \vdash$ and $u_{k+1} = \vdash$),
- $tag_2 = \{acc\}$ if C is accepting, $tag_2 = \{\exists\}$ if C is non-accepting and existential, and $tag_2 = \forall$ otherwise.

The symbols l and r are used to encode a left and a right ATM successor, respectively. We also use the symbol l to encode the initial configuration. If $k = \text{Tower}(n, 3)$ and for each $i \in [1, k]$, bl_i is a well-formed 3-block having number $i - 1$, then we say that w_C is a *well-formed code* of C . A sequence $w_{C_1} \dots w_{C_p}$ of well-formed ATM configuration codes is *faithful to the evolution* of \mathcal{M} if for each $1 \leq i < p$, either $w_{C_{i+1}}$ starts with the symbol l and $C_{i+1} = \text{succ}_l(C_i)$, or $w_{C_{i+1}}$ starts with the symbol r and $C_{i+1} = \text{succ}_r(C_i)$.

Definition of AP and marked blocks. The set AP of atomic propositions is defined as follows:

$$AP = \text{Main} \cup \{check_1, check_2, check_3, \widehat{check}_3\}$$

where the atomic propositions in $\{check_1, check_2, check_3, \widehat{check}_3\}$ are intuitively used to mark the contents of h -blocks. In particular, in the reduction, we also consider *marked* h -blocks for each $h = 1, 2, 3$. Formally, for each $h = 1, 2$, a marked h -block bl is defined as an h -block but the content of bl is additionally labeled by proposition $check_h$. A marked 3-block bl is defined as a 3-block but the content of bl is additionally labeled either by proposition $check_3$ ($check_3$ -marked 3-block) or by proposition \widehat{check}_3 (\widehat{check}_3 -marked 3-block).

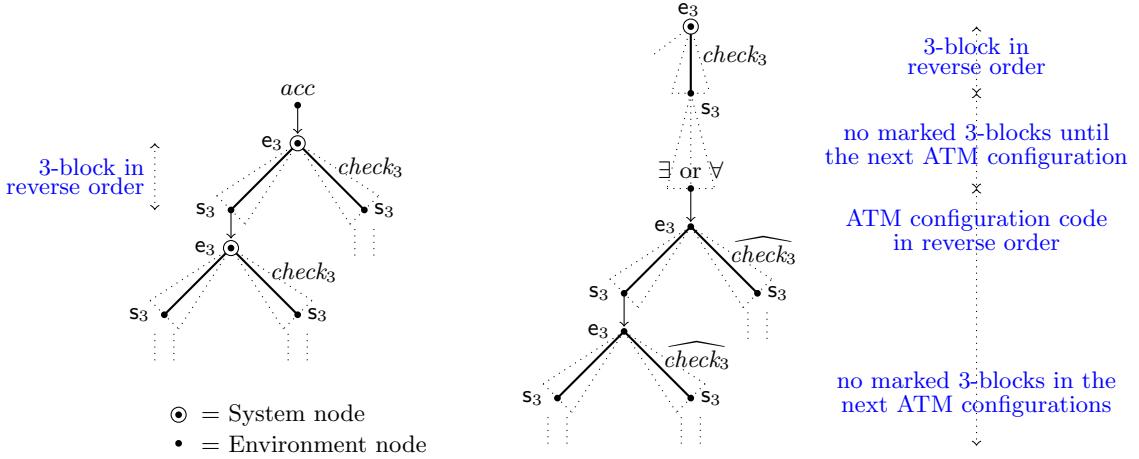
4.2. Construction of the open PMS \mathcal{S} in Theorem 4.2. We now describe the behaviour of the open turn-based PMS \mathcal{S} over $Ag = \{sys, env\}$ in Theorem 4.2. Since \mathcal{S} is turn-based, every configuration is either controlled by the agent *sys* (the *system*) or by the agent *env* (the *environment*). Thus, in the following, for *external* (resp., *internal*) nondeterminism, we mean that at a given configuration, the choices are resolved by the environment (resp., system) agent, i.e., the configuration is controlled by the environment (resp., the system) agent. Intuitively, the PMS generates, for different environment behaviors, all the possible computation trees of \mathcal{M} . External nondeterminism is used in order to produce the actual symbols of each ATM configuration code. Whenever the PMS \mathcal{S} reaches the end of an existential (resp., universal) guessed ATM configuration code w_C , it simulates the existential (resp., universal) choice of \mathcal{M} from C by external (resp., internal) nondeterminism, and, in particular, \mathcal{S} chooses a symbol in $\{l, r\}$ and marks the next guessed ATM configuration with this symbol. This ensures that, once we fix the environment behavior, we really get a tree T where each existential ATM configuration code is followed by (at least) one ATM configuration code marked by a symbol in $\{l, r\}$, and every universal configuration is followed (in different branches) by two ATM configurations codes, one marked by the symbol l and the other one marked by the symbol r .

We have to check that the guessed computation tree T (corresponding to environment choices) corresponds to a legal computation tree of \mathcal{M} over α . To that purpose, we have to check several properties about each computation path π of T , in particular:

- (C1): the ATM configurations codes are well-formed (i.e., the $\text{Tower}(n, 3)$ -bit counter is properly updated),
- (C2): π is faithful to the evolution of \mathcal{M} ,
- (C3): the first ATM configuration of π corresponds to the initial ATM configuration over the input α .

The PMS \mathcal{S} cannot guarantee by itself these requirements. Thus, these checks are performed by a suitable ATL^* formula φ . However, in order to construct an ATL^* formula of size polynomial in n and in the size of the ATM \mathcal{M} , we need to ‘isolate’ the (arbitrary) selected path π from the remaining part of the tree. This is the point where we use the stack of the

Figure 2: Subtree of the unwinding $Unw(\mathcal{G}(\mathcal{S}))$ of the open PMS \mathcal{S} rooted at an *acc*-node (pop-phase)



PMS \mathcal{S} . As the ATM configurations codes are guessed symbol by symbol, they are pushed onto the stack of the PMS \mathcal{S} . This phase is called *push-phase*.

In the push phase, the PMS \mathcal{S} ensures that whenever an *acc*-node x (i.e., a node with label $\{acc\}$) is reached in the unwinding $Unw(\mathcal{G}(\mathcal{S}))$, then the finite path π from the root to node x is labeled by a sequence of ATM configuration codes where the last ATM configuration is accepting: we call such finite paths π of $Unw(\mathcal{G}(\mathcal{S}))$ *accepting* computation paths. Moreover, the stack content associated with node x corresponds to the reverse of the labeling of π . Note that in this phase the unique nodes which are controlled by the system player are the nodes labeled by the proposition \forall , where intuitively the system player simulates the universal choices of the ATM \mathcal{M} from a universal configuration: the \forall -node has two children, one labeled by l (the first symbol of an ATM l -successor) and the other one labeled by r (the first symbol of an ATM r -successor). Thus, \mathcal{M} accepts α iff there is an environment strategy tree \mathcal{T} of $\mathcal{G}(\mathcal{S})$ such that (i) each path of \mathcal{T} from the root visits an *acc*-node and (ii) for each *acc*-node x of \mathcal{T} , the accepting computation path π associated with node x satisfies Conditions C1–C3.

Pop Phase. Whenever an *acc*-node x is reached, the PMS moves to the so called *pop-phase*. Let π be the accepting computation path associated with x (i.e., the finite path of $Unw(\mathcal{G}(\mathcal{S}))$ from the root to x). Recall that \mathcal{S} ensures that the labeling of π is a sequence of ATM configuration codes where the last ATM configuration is accepting and the stack content of x corresponds to the reverse of the π 's labeling. Let us denote by $\langle T_\pi, Lab_\pi \rangle$ the subtree of the unwinding $Unw(\mathcal{G}(\mathcal{S}))$ rooted at the last node of π (i.e., node x).

We now describe the branching behaviour of \mathcal{S} along $\langle T_\pi, Lab_\pi \rangle$ (pop-phase). The structure of $\langle T_\pi, Lab_\pi \rangle$ is also illustrated in Figures 2 and 3. By using both internal and external nondeterminism, the PMS pop the (labeling of the) entire computation path π from the stack. In this way, the PMS \mathcal{S} partitions the sanity checks for π into separate branches. The labelings of these branches correspond to the *reverse* of the π 's labeling but they are augmented with additional information by means of the extra atomic propositions $\widehat{check}_3, \check{check}_3, \check{check}_2, \check{check}_1$. In particular, in the pop-phase, the unique nondeterministic

or branching nodes (i.e., the nodes with at least two children) are end nodes, i.e., nodes labeled by one of the propositions in $\{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3\}$. These nodes have, in particular, a binary branching degree. Moreover:

- the branching behaviour at the branching \mathbf{e}_3 -nodes along $\langle T_\pi, \text{Lab}_\pi \rangle$ is subdivided in two sub-phases. In the first sub-phase, the branching \mathbf{e}_3 -nodes are controlled by the system player, and \mathcal{S} marks by *internal* nondeterminism the Λ -content of exactly one 3-block bl_3 of π with the special proposition check_3 (i.e., the content of bl_3 is additionally labeled by proposition check_3). This means, in particular, that for each 3-block bl_3 of π , there is a path of $\langle T_\pi, \text{Lab}_\pi \rangle$ from the root such that the unique check_3 -marked 3-block corresponds to bl_3 . This is illustrated in the left part of Figure 2. Note that in the first sub-phase a branching \mathbf{e}_3 -node has two children whose labels are of the form $\{\lambda\}$ and $\{\lambda, \text{check}_3\}$, respectively, for some $\lambda \in \Lambda$. After having marked the content of a 3-block with check_3 , \mathcal{S} moves to the second sub-phase, where the branching \mathbf{e}_3 -nodes are controlled by the environment player. In particular, in case the marked 3-block bl_3 does not belong to the first configuration code of π , \mathcal{S} marks by *external* nondeterminism the Λ -content of exactly one 3-block bl'_3 with the special proposition $\widehat{\text{check}}_3$ by ensuring that bl_3 and bl'_3 belong to two consecutive configurations codes along π . Hence, for all 3-blocks bl_3 and bl'_3 of π such that bl_3 and bl'_3 belong to adjacent configurations and bl'_3 follows bl_3 along the reverse of the π 's labeling, there is a path of $\langle T_\pi, \text{Lab}_\pi \rangle$ from the root such that the unique check_3 -marked 3-block corresponds to bl_3 and the unique $\widehat{\text{check}}_3$ -marked 3-block corresponds to bl'_3 . This is illustrated in the right part of Figure 2.
- The branching behaviour at the \mathbf{e}_2 -nodes and \mathbf{e}_1 -nodes along $\langle T_\pi, \text{Lab}_\pi \rangle$, which is illustrated in Figure 3, is as follows. The \mathbf{e}_2 -nodes are controlled by the system player, while the \mathbf{e}_1 -nodes are controlled by the environment nodes. In particular, from each \mathbf{e}_2 -node $x_{\mathbf{e}_2}$ associated with the first symbol of the reverse of some 2-block bl_2 , \mathcal{S} generates by *internal* nondeterminism a tree copy of the reverse of bl_2 . More specifically, node $x_{\mathbf{e}_2}$ has two children y and y_{check_2} whose labels are $\{b\}$ and $\{b, \text{check}_2\}$, respectively, for some $b \in \{0, 1\}$ (see Figures 3(b) and 3(c)). Moreover, the labeled tree (we call *check 2-block-tree*) obtained from the subtree of $\langle T_\pi, \text{Lab}_\pi \rangle$ rooted a node $x_{\mathbf{e}_2}$ by pruning node y and its descendants is structured as follows (see Figure 3(c)):
 - there is an infinite path ρ from the \mathbf{e}_2 -node $x_{\mathbf{e}_2}$ whose labeling consists of a marked copy (of the reverse) of bl_2 (the *content* b of bl_2 is additionally labeled by the special proposition check_2) followed by the suffix \emptyset^ω (see Figures 3(c));
 - there are additional branches chosen by *external* nondeterminism starting at the \mathbf{e}_1 -nodes of the infinite path ρ . As illustrated in Figure 3(c), these additional branches represent marked copies of the (reverse of) 1-sub-blocks bl_1 of bl_2 (the *content* of bl_1 is additionally labeled by the special proposition check_1).

Note that for the \mathbf{e}_1 -nodes of $\langle T_\pi, \text{Lab}_\pi \rangle$, only the ones belonging to check 2-block trees are branching.

Note that for each $h = 1, 2, 3$, in a marked h -block bl , only the content (i.e., the symbol preceding the end-symbol) of bl is marked.

Hence, the subtree $\langle T_\pi, \text{Lab}_\pi \rangle$ of the unwinding $\text{Unw}(\mathcal{G}(\mathcal{S}))$ of $\mathcal{G}(\mathcal{S})$ associated with this pop-phase and the specific accepting computation path π , satisfies the following: the labeling of each *main* path (i.e., a path of $\langle T_\pi, \text{Lab}_\pi \rangle$ starting at the root which does not get trapped into a check 2-block-tree) corresponds to the reverse of the π 's labeling (followed by a suffix with label \emptyset^ω) with the unique difference that exactly one 3-block bl_3 is marked by

$check_3$ and (in case bl_3 does not belong to the first ATM configuration code of π) exactly one 3-block bl'_3 is marked by $\widehat{check_3}$. The PMS \mathcal{S} ensures that bl_3 and bl'_3 belong to two consecutive configurations codes along π (where bl_3 precedes bl'_3 along the main path) and, independently from the environment choices, all the 3-blocks bl_3 of π are checked (i.e., there is a main path whose $check_3$ -marked block corresponds to bl_3).

The additional check 2-block-trees are intuitively used to *isolate* 2-blocks for ensuring by an ATL^* formula φ that the ATM configuration codes along the π 's labeling are well-formed and the π 's labeling is faithful to the evolution of \mathcal{M} . In particular, as detailed in the proof of Lemma 4.7, the ATL^* formula φ requires that the given environment strategy tree of $\mathcal{G}(\mathcal{S})$ satisfies the following:

- all the environment choices in each 2-block check-tree are enabled,
- the environment choices from the $\{e_3\}$ -nodes which are descendants of acc -nodes (i.e., e_3 -nodes associated with the pop-phase) and are controlled by the environment player are *deterministic*. This entails that the subtree rooted at the s_3 -node of a $check_3$ -marked 3-block bl_3 which does not belong to the first ATM configuration code *contains exactly one* $\widehat{check_3}$ -marked 3-block bl'_3 .

Then, by exploiting the previous two requirements, the ATL^* formula φ existentially quantifies over strategies of the system player whose outcomes get trapped into a check 2-block-tree in order to ensure that for the given sequence ν of ATM configuration codes (associated with an accepting computation path π of the push-phase), the following holds:

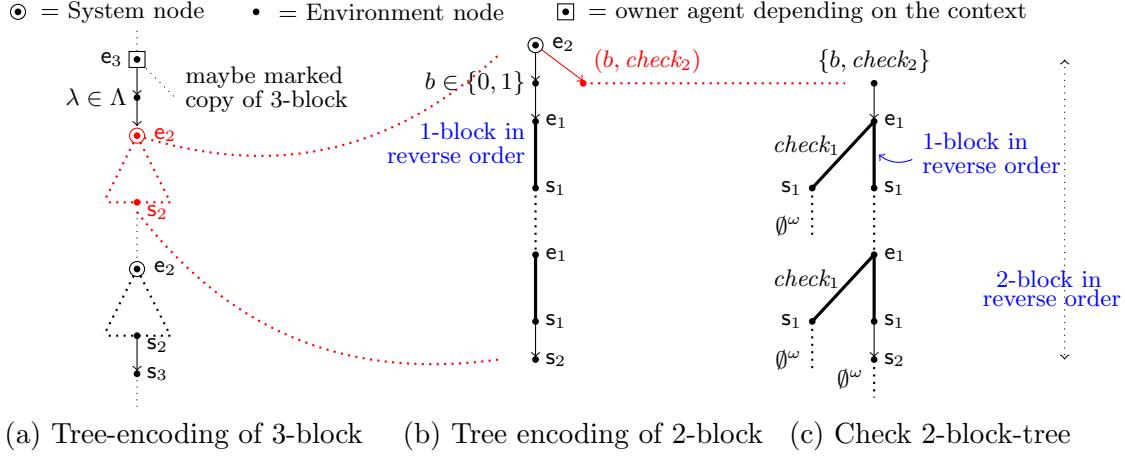
- the configuration codes along ν are well-formed,
- for each $check_3$ -marked 3-block bl_3 which does not belong to the first ATM configuration code of ν , the associated $check_3$ -marked 3-block bl'_3 satisfies the following: bl_3 and bl'_3 have the same number and the Λ -contents of bl_3 and bl'_3 are consistent with the transition function of \mathcal{M} . Since bl_3 and bl'_3 belong to two adjacent configuration codes along ν , the previous conditions ensure that ν is faithful to the evolution of \mathcal{M} . Note that in order to enforce that bl_3 and bl'_3 have the same number, for each 2-sub-block bl_2 of bl_3 , the formula φ requires the existence of a system strategy f starting at the e_2 -node of bl_2 which gets trapped into the check 2-block-tree of a 2-sub-block bl'_2 of bl'_3 such that the copy of bl'_2 in the check 2-block-tree and bl_2 have the same number and the same content. Note that the additional $check_1$ -branches of the check 2-block-tree of bl'_2 are used to check by an LTL formula, asserted at the outcomes of the system strategy f , that bl_2 and the copy of bl'_2 have the same number.

Recall that $AP = \text{Main} \cup \{check_1, check_2, check_3, \widehat{check_3}\}$. We now formally define the AP -labeled trees associated with the *accepting* environment strategy trees of $\mathcal{G}(\mathcal{S})$, i.e. the environment strategy trees where each path from the root visits an $\{acc\}$ -labeled node. In the following, a 2^{AP} -labeled tree is *minimal* if the children of each node have distinct labels. A *branching node* of a tree is a node having at least two distinct children.

Tree-codes. A *tree-code* is a *finite minimal* 2^{AP} -labeled tree $\langle T, \text{Lab} \rangle$ such that

- for each maximal path π from the root, $\text{Lab}(\pi)$ is a sequence of ATM configuration codes;
- a node x is labeled by $\{acc\}$ iff x is a leaf;
- each node labeled by $\{\forall\}$ has two children, one labeled by $\{l\}$ and one labeled by $\{r\}$.

Intuitively, tree-codes correspond to the maximal portions of the *accepting* environment strategy trees of $\mathcal{G}(\mathcal{S})$ where \mathcal{S} performs push operations (push-phase). We now extend a

Figure 3: Marked copies of 2-blocks in the pop-phase of the open PMS \mathcal{S} 

tree-code $\langle T, \text{Lab} \rangle$ with extra nodes in such a way that each leaf x of $\langle T, \text{Lab} \rangle$ is expanded in a tree, called *check-tree* (pop-phase).

Check-trees. The definition of check-trees is based on the notion of *check 2-block-tree* and *simple check-tree*. The structure of a check 2-block-tree for a 2-block bl_2 is illustrated in Figure 3(c). Note that the unique branching nodes are labeled by $\{e_1\}$. In the accepting environment strategy trees of $\mathcal{G}(\mathcal{S})$, these nodes are controlled by the environment. Formally, a check 2-block-tree for a 2-block bl_2 is a *minimal* 2^{AP} -labeled tree $\langle T, \text{Lab} \rangle$ such that:

- there is a path from the root (*main path*) whose labeling is $\rho \cdot \emptyset^\omega$, where ρ is the reverse of the marked copy of bl_2 (the content of bl_2 is additionally labeled by proposition check_2);
- for each $\{e_1\}$ -labeled node x of the main path, there is an infinite path π_s from x (*secondary branch*) such that denoted by bl_1 the 1-subblock of bl_2 associated with node x , the labeling of π_s is $\rho_1 \cdot \emptyset^\omega$, where ρ_1 is the reverse of the marked copy of bl_1 (the content of bl_1 is additionally labeled by proposition check_1);
- each node of $\langle T, \text{Lab} \rangle$ is either a node of the main path or a node of some secondary branch.

A *partial check 2-block-tree* for bl_2 is obtained from the check 2-block-tree for bl_2 by pruning some choices from the $\{e_1\}$ -branching nodes.

Given a sequence ν of ATM configuration codes, a *simple check-tree* for ν is a *minimal* 2^{AP} -labeled tree $\langle T, \text{Lab} \rangle$ such that

- for each path π from the root, $\text{Lab}(\pi)$ corresponds to the *reverse* of ν followed by \emptyset^ω but there is exactly one 3-block bl_3 of ν whose content is additionally marked by proposition check_3 , and in case bl_3 does not belong to the first configuration code of ν , there is exactly one 3-block bl'_3 whose content is marked by proposition $\widehat{\text{check}}_3$; moreover, bl'_3 and bl_3 belong to two consecutive configuration codes, and bl'_3 precedes bl_3 along ν ;
- for each 3-block bl_3 of ν , there is a path π from the root such that the node associated with the content of bl_3 is additionally labeled by proposition check_3 (i.e., all the 3-blocks of ν are checked);

- each branching node x has label $\{\mathbf{e}_3\}$ and two children: one labeled by $\{\lambda\}$ and the other one labeled by $\{\lambda, \text{tag}\}$ for some $\lambda \in \Lambda$ and $\text{tag} \in \{\text{check}_3, \widehat{\text{check}}_3\}$. If $\text{tag} = \text{check}_3$ (resp., $\text{tag} = \widehat{\text{check}}_3$), we say that x is a check_3 -branching (resp., $\widehat{\text{check}}_3$ -branching) node.

Finally, a *check-tree* for ν is a *minimal* 2^{AP} -labeled tree $\langle T, \text{Lab} \rangle$ which is obtained from some simple check-tree $\langle T', \text{Lab}' \rangle$ for ν by adding for each node x of T' with label $\{\mathbf{e}_2\}$ an additional child y and a subtree rooted at y so that the subtree rooted at x obtained by removing all the descendants of x in T' is a partial check 2-block-tree for the 2-block associated with node x in T' .

Thus, in a check-tree, we have four types of branching nodes: check_3 -branching nodes, $\{\mathbf{e}_2\}$ -branching nodes, $\widehat{\text{check}}_3$ -branching nodes, and $\{\mathbf{e}_1\}$ -branching nodes. In the accepting environment strategy trees of $\mathcal{G}(\mathcal{S})$, check_3 -branching nodes and $\{\mathbf{e}_2\}$ -branching nodes are controlled by the system, while $\widehat{\text{check}}_3$ -branching nodes and $\{\mathbf{e}_1\}$ -branching nodes are controlled by the environment.

Extended tree-codes. An *extended tree-code* is a minimal 2^{AP} -labeled tree $\langle T_e, \text{Lab}_e \rangle$ such that there is a tree-code $\langle T, \text{Lab} \rangle$ so that $\langle T_e, \text{Lab}_e \rangle$ is obtained from $\langle T, \text{Lab} \rangle$ by replacing each leaf x (recall that x is labeled by $\{\text{acc}\}$) with a check-tree for the sequence of labels associated with the path of $\langle T, \text{Lab} \rangle$ starting at the root and leading to x . By construction and the intuitions given about the PMS \mathcal{S} , we obtain the following result.

Lemma 4.4. *One can build, in time polynomial in the size of the ATM \mathcal{M} , an open turn-based PMS \mathcal{S} over AP and $Ag = \{\text{env}, \text{sys}\}$ such that the following hold, where an environment strategy tree is accepting if each path from the root visits an $\{\text{acc}\}$ -labeled node:*

- the set of 2^{AP} -labeled trees $\langle T, \text{Lab} \rangle$ associated with the accepting environment strategy trees $\langle T, \text{Lab}, \tau \rangle$ in $\text{exec}(\mathcal{G}(\mathcal{S}))$ coincides with the set of extended tree-codes;
- for each accepting environment strategy tree $\langle T, \text{Lab}, \tau \rangle$ in $\text{exec}(\mathcal{G}(\mathcal{S}))$, the unique nodes controlled by the system in a check-subtree of $\langle T, \text{Lab}, \tau \rangle$ are the check_3 -branching nodes and the $\{\mathbf{e}_2\}$ -branching nodes.

Proof. Since the PMS \mathcal{S} is turn-based, each configuration is either controlled by the environment or by the system agent. Thus, in specifying the transition function Δ of \mathcal{S} , we can abstract away from the set of full decisions, and we just specify for each state p and stack symbol γ , the set of pairs (p', β) such that $(p', \beta) \in \Delta(p, \gamma, d)$ for some full decision d . In particular, the transition function of Δ consists of the following types of transitions:

- push transitions $p \xrightarrow{\text{push}(\gamma)} p'$ meaning that for each stack symbol γ' , $(p', \gamma \cdot \gamma') \in \Delta(p, \gamma', d)$ for some full decision d ;
- pop transitions $p \xrightarrow{\text{push}(\gamma)} p'$ meaning that $(p', \varepsilon) \in \Delta(p, \gamma, d)$ for some full decision d ;
- internal transitions $p \xrightarrow{\gamma} p'$ that do not use the stack, meaning that for each stack symbol γ , there is some full decision d such that either $\gamma = \gamma_0$ and $(p', \varepsilon) \in \Delta(p, \gamma, d)$, or $\gamma \neq \gamma_0$ and $(p', \gamma) \in (p, \gamma, d)$.

The initial state of \mathcal{S} is denoted by *in* and the set of states is given by

$$Q_{\text{push}} \cup Q_{\text{pop}} \cup \{q_\emptyset\}$$

where $in \in Q_{\text{push}}$, Q_{push} is the set of states used in the push phase, Q_{pop} is the set of states used in the pop-phase, and the state q_\emptyset is a sink state. The unique transition from state q_\emptyset is the internal transition $q_\emptyset \xrightarrow{\gamma} q_\emptyset$. Moreover, the propositional labeling of q_\emptyset is the empty set. Note that each configuration with control state q_\emptyset is deterministic.

The stack alphabet Γ is defined as follows:

$$\Gamma := \text{Main} \cup \{\text{in}\} \cup (\Lambda \times \{\text{first}, \text{first}_{\text{in}}\})$$

where the symbol in is pushed onto the stack in the first step of the push phase, while a symbol $(\lambda, t) \in \Lambda \times \{\text{first}, \text{first}_{\text{in}}\}$ is pushed onto the stack on generating the content λ of the first 3-block of a guessed ATM configuration code C . In particular, $t = \text{first}_{\text{in}}$ means that C is the first guessed ATM configuration code, while $t = \text{first}$ means it isn't. The flags in $\{\text{first}, \text{first}_{\text{in}}\}$ are used in the pop-phase for ensuring that the generation of the propositions check_3 and $\widehat{\text{check}}_3$ is consistent with the definition of check-tree.

Push Phase. The set Q_{push} of states used in the push phase is defined as follows:

$$\begin{aligned} Q_{\text{push}} &:= \{\text{in}\} \cup ((\Lambda \cup \{\perp\}) \times (Q \cup \{\perp\}) \times \{\text{in}, \perp\} \times \text{Main}') \\ \text{Main}' &:= \text{Main} \cup (\{bl_1, bl_2\} \times \{0, 1\}) \cup (\Lambda \times \{\text{first}, \text{first}_{\text{in}}\}) \end{aligned}$$

The intuitive meaning of a state $(\lambda_{\perp}, q_{\perp}, \text{in}_{\perp}, m) \in Q_{\text{push}}$ is as follows:

- The symbol $\lambda_{\perp} \in \Lambda \cup \{\perp\}$ keeps track of the last Λ -symbol of the prefix of the current guessed ATM configuration code generated so far if such a prefix contains a Λ -symbol; otherwise, $\lambda_{\perp} = \perp$.
- The symbol $q_{\perp} \in Q \cup \{\perp\}$ keeps track of the state associated to the current guessed ATM configuration code C if the prefix of C generated so far contains a Λ -symbol (u_p, u, u_q) where u is of the form (σ, q_{\perp}) ; otherwise, $q_{\perp} = \perp$.
- The symbol $\text{in}_{\perp} \in \{\text{in}, \perp\}$ keeps track whether the current guessed ATM configuration code C is the first one to be generated ($\text{in}_{\perp} = \text{in}$) or not ($\text{in}_{\perp} = \perp$).
- The main symbol m has the following meaning: if $m \in \text{Main}$ (resp., $m = (\lambda, t) \in \Lambda \times \{\text{first}, \text{first}_{\text{in}}\}$), then m (resp., λ) is the symbol currently generated for the current guessed ATM configuration code C . The flag $t \in \{\text{first}, \text{first}_{\text{in}}\}$ means that λ is the content of the first 3-block of C with $t = \text{first}_{\text{in}}$ iff C is the first guessed ATM configuration code. If instead $m = (bl_1, b)$ (resp., $m = (bl_2, b)$) for some $b \in \{0, 1\}$, then b represents the content of a 1-block (resp., 2-block) of C .

The propositional label of state $(\lambda_{\perp}, q_{\perp}, \text{in}_{\perp}, m)$ is $\{m\}$ if $m \in \text{Main}$, is λ if $m = (\lambda, t) \in \Lambda \times \{\text{first}, \text{first}_{\text{in}}\}$, and is $\{b\}$ if $m = (bl_k, b)$ for some $k = 1, 2$ and $b \in \{0, 1\}$. The propositional label of the initial state in is $\{l\}$. Moreover, all the configurations associated with the states in Q_{push} are controlled by the environment with the exception of the configurations associated with the push states of the form $(\lambda_{\perp}, q_{\perp}, \forall)$, which are instead controlled by the system. Transitions from states in Q_{push} are push transitions. These transitions are defined as follows, where $\text{acc} \in \text{Main}$ is also used as control state in Q_{pop} and has propositional labeling $\{\text{acc}\}$.

- Transitions from state in : $\text{in} \xrightarrow{\text{push}(\text{in})} (\perp, \perp, \text{in}, s_3)$.
- Transitions from states $(\lambda_{\perp}, q_{\perp}, \text{in}_{\perp}, l), (\lambda_{\perp}, q_{\perp}, \text{in}_{\perp}, r) \in Q_{\text{push}}$:
 $(\lambda_{\perp}, q_{\perp}, \text{in}_{\perp}, l) \xrightarrow{\text{push}(l)} (\lambda_{\perp}, q_{\perp}, \text{in}_{\perp}, s_3)$ and $(\lambda_{\perp}, q_{\perp}, \text{in}_{\perp}, r) \xrightarrow{\text{push}(r)} (\lambda_{\perp}, q_{\perp}, \text{in}_{\perp}, s_3)$.
- Transitions from states $(\lambda_{\perp}, q_{\perp}, \text{in}_{\perp}, s_k) \in Q_{\text{push}}$, where $k = 1, 2, 3$:
 - $(\lambda_{\perp}, q_{\perp}, \text{in}_{\perp}, s_k) \xrightarrow{\text{push}(s_k)} (\lambda_{\perp}, q_{\perp}, \text{in}_{\perp}, s_{k-1})$ for each $k = 2, 3$.
 - $(\lambda_{\perp}, q_{\perp}, \text{in}_{\perp}, s_1) \xrightarrow{\text{push}(s_1)} (\lambda_{\perp}, q_{\perp}, \text{in}_{\perp}, b)$ for each $b \in \{0, 1\}$.
- Transitions from states $(\lambda_{\perp}, q_{\perp}, \text{in}_{\perp}, b) \in Q_{\text{push}}$, where $b \in \{0, 1\}$:

- $(\lambda_{\perp}, q_{\perp}, in_{\perp}, b) \xrightarrow{\text{push}(b)} (\lambda_{\perp}, q_{\perp}, in_{\perp}, b') \text{ for all } b' \in \{0, 1\};$
 - $(\lambda_{\perp}, q_{\perp}, in_{\perp}, b) \xrightarrow{\text{push}(b)} (\lambda_{\perp}, q_{\perp}, in_{\perp}, (bl_1, b')) \text{ for all } b' \in \{0, 1\}.$
- Transitions from states $(\lambda_{\perp}, q_{\perp}, in_{\perp}, (bl_k, b)) \in Q_{\text{push}}$, where $k = 1, 2$ and $b \in \{0, 1\}$:

$$(\lambda_{\perp}, q_{\perp}, in_{\perp}, (bl_k, b)) \xrightarrow{\text{push}(b)} (\lambda_{\perp}, q_{\perp}, in_{\perp}, e_k).$$
- Transitions from states $(\lambda_{\perp}, q_{\perp}, in_{\perp}, e_1) \in Q_{\text{push}}$:
 - $(\lambda_{\perp}, q_{\perp}, in_{\perp}, e_1) \xrightarrow{\text{push}(e_1)} (\lambda_{\perp}, q_{\perp}, in_{\perp}, s_1);$
 - $(\lambda_{\perp}, q_{\perp}, in_{\perp}, e_1) \xrightarrow{\text{push}(e_1)} (\lambda_{\perp}, q_{\perp}, in_{\perp}, (bl_2, b)) \text{ for all } b \in \{0, 1\}.$
- Transitions from states $(\perp, q_{\perp}, in_{\perp}, e_2) \in Q_{\text{push}}$:
 - $(\perp, q_{\perp}, in_{\perp}, e_2) \xrightarrow{\text{push}(e_2)} (\perp, q_{\perp}, in_{\perp}, s_2);$
 - $(\perp, q_{\perp}, in_{\perp}, e_2) \xrightarrow{\text{push}(e_2)} (\perp, q'_{\perp}, in_{\perp}, (\lambda, t)) \text{ for all } (\lambda, t) \in \Lambda \times \{first, first_{in}\} \text{ and } q'_{\perp} \in Q \cup \{\perp\} \text{ such that the following holds, where } \lambda = (u_p, u, u_s):$
 - * either $u \in \Sigma$ and $q'_{\perp} = \perp$, or u is of the form $(\sigma, q'_{\perp}) \in \Sigma \times Q$;
 - * $u_p = \perp$;
 - * $t = first_{in}$ if $in_{\perp} = in$, and $t = first$ otherwise.
- Transitions from states $(\lambda, q_{\perp}, in_{\perp}, e_2) \in Q_{\text{push}}$ where $\lambda \in \Lambda$:
 - $(\lambda, q_{\perp}, in_{\perp}, e_2) \xrightarrow{\text{push}(e_2)} (\lambda, q_{\perp}, in_{\perp}, s_2);$
 - $(\lambda, q_{\perp}, in_{\perp}, e_2) \xrightarrow{\text{push}(e_2)} (\lambda, q'_{\perp}, in_{\perp}, \lambda') \text{ for all } \lambda' \in \Lambda \text{ and } q'_{\perp} \in Q \cup \{\perp\} \text{ such that the following holds, where } \lambda = (u_p, u, u_s) \text{ and } \lambda' = (u'_p, u', u'_s):$
 - * either $u' \in \Sigma$ and $q'_{\perp} = q_{\perp}$, or u' is of the form $(\sigma, q'_{\perp}) \in \Sigma \times Q$ and $q_{\perp} = \perp$;
 - * $u'_p = u$ and $u' = u_s$.
- Transitions from states $(\lambda_{\perp}, q_{\perp}, in_{\perp}, (\lambda', t)) \in Q_{\text{push}}$, where $(\lambda', t) \in \Lambda \times \{first, first_{in}\}$:

$$(\lambda_{\perp}, q_{\perp}, in_{\perp}, (\lambda', t)) \xrightarrow{\text{push}((\lambda', t))} (\lambda', q_{\perp}, in_{\perp}, e_3)$$
- Transitions from states $(\lambda_{\perp}, q_{\perp}, in_{\perp}, \lambda') \in Q_{\text{push}}$, where $\lambda' \in \Lambda$:

$$(\lambda_{\perp}, q_{\perp}, in_{\perp}, \lambda') \xrightarrow{\text{push}(\lambda')} (\lambda', q_{\perp}, in_{\perp}, e_3)$$
- Transitions from states $(\lambda_{\perp}, q_{\perp}, in_{\perp}, e_3) \in Q_{\text{push}}$:
 - if λ_{\perp} is not of the form (u_p, u, \perp) : $(\lambda_{\perp}, q_{\perp}, in_{\perp}, e_3) \xrightarrow{\text{push}(e_3)} (\lambda_{\perp}, q_{\perp}, in_{\perp}, s_3);$
 - else if $q_{\perp} \neq \perp$ and q_{\perp} is universal and non-accepting:

$$(\lambda_{\perp}, q_{\perp}, in_{\perp}, e_3) \xrightarrow{\text{push}(e_3)} (\lambda_{\perp}, q_{\perp}, in_{\perp}, \forall);$$
 - else if $q_{\perp} \neq \perp$ and q_{\perp} is existential and non-accepting:

$$(\lambda_{\perp}, q_{\perp}, in_{\perp}, e_3) \xrightarrow{\text{push}(e_3)} (\lambda_{\perp}, q_{\perp}, in_{\perp}, \exists);$$
 - else if $q_{\perp} \neq \perp$ and q_{\perp} is accepting: $(\lambda_{\perp}, q_{\perp}, in_{\perp}, e_3) \xrightarrow{\text{push}(e_3)} acc;$
 - else: $(\lambda_{\perp}, q_{\perp}, in_{\perp}, e_3) \xrightarrow{\text{push}(e_3)} q_{\emptyset}.$
- Transitions from states $(\lambda_{\perp}, q_{\perp}, in_{\perp}, \exists), (\lambda_{\perp}, q_{\perp}, in_{\perp}, \forall) \in Q_{\text{push}}$:
 - $(\lambda_{\perp}, q_{\perp}, in_{\perp}, \exists) \xrightarrow{\text{push}(\exists)} (\perp, \perp, \perp, dir) \text{ for all } dir \in \{l, r\};$
 - $(\lambda_{\perp}, q_{\perp}, in_{\perp}, \forall) \xrightarrow{\text{push}(\forall)} (\perp, \perp, \perp, dir) \text{ for all } dir \in \{l, r\}.$

Let ν be a sequence of ATM configuration codes of the form $\nu = \rho \cdot \{acc\}$. Note that each symbol of ρ is of the form $\{p\}$ where $p \in Main$. Thus, ρ corresponds to a word ρ' over $Main$. Let ρ'' obtained from ρ' by replacing the first symbol of ρ' with in and by replacing for each configuration code C along ρ' , the content λ of the first 3-block of C with (λ, t) , where $t = first_{in}$ if C is the first configuration code of ρ' , and $t = first$ otherwise. We denote by $Stack(\nu)$ the stack content given by $(\rho'')^R \cdot \gamma_0$, where $(\rho'')^R$ is the reverse of ρ'' . By construction, the following two claims hold.

Claim 1. Let \mathcal{T} be an accepting environment strategy tree of $\mathcal{G}(\mathcal{S})$. Then, the finite labeled tree obtained from \mathcal{T} by pruning all the subtrees of \mathcal{T} rooted at the children of acc -nodes is a tree-code. Moreover, for each acc -node x of \mathcal{T} , let π be the finite path from the root to node x . Then, the stack content of node x is $Stack(\nu)$ where ν is the labeling of π .

Claim 2. Let $\langle T, Lab \rangle$ be a tree-code. Then there exists an accepting environment strategy tree \mathcal{T} of $\mathcal{G}(\mathcal{S})$ such that the finite labeled tree obtained from \mathcal{T} by pruning all the subtrees of \mathcal{T} rooted at the children of acc -nodes is isomorphic to $\langle T, Lab \rangle$.

Pop Phase. The set Q_{pop} of states used in the pop phase is defined as follows:

$$Q_{pop} := Q_{pop}^1 \cup Q_{pop}^2$$

where Q_{pop}^1 is used to generate the nodes of the *main* paths of a check-tree (i.e., the paths that do not get trapped into check 2-block trees), while the states in Q_{pop}^1 are used to generate check 2-block trees. We first consider the set Q_{pop}^2 which is defined as follows:

$$Q_{pop}^1 := \{acc\} \cup (\{check_3\} \times \Lambda) \cup (\{check_3^0, check_3^1, check_3^2, \widehat{check}_3\} \times Main)$$

Intuitively, for a state $(t, m) \in Q_{pop}^1$, m represents the last symbol which has been popped from the stack. Moreover, the flag $t \in \{check_3^0, check_3^1, check_3^2, \widehat{check}_3\}$ has the following meaning:

- $t = check_3^0$ iff the proposition $check_3$ has not been generated so far.
- $t = check_3$ iff $m \in \Lambda$ (i.e., m is the content of a 3-block) and proposition $check_3$ is currently generated.
- $t = check_3^1$ iff the proposition $check_3$ has been generated and is associated with a 3-block of the current ATM configuration code.
- $t = check_3^2$ iff the proposition $check_3$ has been generated and is associated with a 3-block of an ATM configuration code preceding the current one.
- $t = \widehat{check}_3$ iff the proposition \widehat{check}_3 has been generated.

For each $(t, m) \in Q_{pop}^1$, the propositional labeling of state (t, m) is $\{t, m\}$ if $(t, m) \in \{check_3, \widehat{check}_3\} \times \Lambda$, and $\{m\}$ otherwise. All the configurations associated with the states in Q_{pop}^1 are controlled by the environment with the exception of the configurations associated with the states of the form $(check_3^0, e_3)$, which are instead controlled by the system. Transitions from states in Q_{pop}^1 are pop transitions. They are defined as follows where $(check_2, b) \in Q_{pop}^2$ for each $b \in \{0, 1\}$.

- Transitions from state acc :
 - $acc \xrightarrow{\text{pop}(e_3)} (check_3^0, e_3)$.
 - $acc \xrightarrow{\text{pop}(\gamma)} q_\emptyset$ for all $\gamma \neq e_3$.
- Transitions from states $(check_3, \lambda) \in Q_{pop}^1$ (note that $\lambda \in \Lambda$):

- $(check_3, \lambda) \xrightarrow{\text{pop}(\mathbf{e}_2)} (check_3^1, \mathbf{e}_2)$.
- $(check_3, \lambda) \xrightarrow{\text{pop}(\gamma)} q_\emptyset$ for all $\gamma \neq \mathbf{e}_2$.
- Transitions from states $(t, \mathbf{e}_2) \in Q_{pop}^1$:
 - $(t, \mathbf{e}_2) \xrightarrow{\text{pop}(\gamma)} q_\emptyset$ for all $\gamma \notin \{0, 1\}$.
 - $(t, \mathbf{e}_2) \xrightarrow{\text{pop}(b)} (t, b)$ and $(t, \mathbf{e}_2) \xrightarrow{\text{pop}(b)} (check_2, b)$ for all $b \in \{0, 1\}$.
- Transitions from states $(t, \mathbf{e}_3) \in Q_{pop}^1$:
 - $(t, \mathbf{e}_3) \xrightarrow{\text{pop}(\gamma)} q_\emptyset$ for all $\gamma \notin \Lambda \cup (\Lambda \times \{first, first_{in}\})$.
 - $(t, \mathbf{e}_3) \xrightarrow{\text{pop}(\lambda)} (t, \lambda)$ for all $t \in \{check_3^1, \widehat{check}_3\}$ and $\lambda \in \Lambda$.
 - $(t, \mathbf{e}_3) \xrightarrow{\text{pop}((\lambda, first))} (t, \lambda)$ for all $t \in \{check_3^1, \widehat{check}_3\}$ and $(\lambda, t) \in \Lambda \times \{first, first_{in}\}$.
 - $(check_3^0, \mathbf{e}_3) \xrightarrow{\text{pop}(\lambda)} (check_3^0, \lambda)$ and $(check_3^0, \mathbf{e}_3) \xrightarrow{\text{pop}(\lambda)} (check_3, \lambda)$ for all $\lambda \in \Lambda$.
 - $(check_3^0, \mathbf{e}_3) \xrightarrow{\text{pop}((\lambda, first))} (check_3^0, \lambda)$ and $(check_3^0, \mathbf{e}_3) \xrightarrow{\text{pop}((\lambda, first))} (check_3, \lambda)$ for all $\lambda \in \Lambda$.
 - $(check_3^0, \mathbf{e}_3) \xrightarrow{\text{pop}((\lambda, first_{in}))} (check_3, \lambda)$ for all $\lambda \in \Lambda$.
 - $(check_3^2, \mathbf{e}_3) \xrightarrow{\text{pop}(\lambda)} (check_3^2, \lambda)$ and $(check_3^2, \mathbf{e}_3) \xrightarrow{\text{pop}(\lambda)} (\widehat{check}_3, \lambda)$ for all $\lambda \in \Lambda$.
 - $(check_3^2, \mathbf{e}_3) \xrightarrow{\text{pop}((\lambda, t))} (\widehat{check}_3, \lambda)$ for all $(\lambda, t) \in \Lambda \times \{first, first_{in}\}$.
- Transitions from states $(t, m) \in Q_{pop}^1 \setminus (\{check_3\} \times \Lambda)$ with $m \notin \{\mathbf{e}_2, \mathbf{e}_3\}$:
 - $(t, m) \xrightarrow{\text{pop}(\gamma)} q_\emptyset$ for all $\gamma \notin Main$.
 - $(t, m) \xrightarrow{\text{pop}(\gamma)} (t, m)$, where $\gamma \in Main$, and either $m \notin \{\exists, \forall\}$ or $t \neq check_3^1$.
 - $(check_3^1, m) \xrightarrow{\text{pop}(\gamma)} (check_3^2, m)$ for all $\gamma \in Main$ and $m \in \{\exists, \forall\}$.

The states in Q_{pop}^2 are used to generate the non-root nodes of check 2-block trees. The set Q_{pop}^2 is defined as follows:

$$Q_{pop}^2 := (\{check_1, check_2\} \times \{0, 1\}) \cup (\{check_1^0\} \times \{s_1, e_1, s_2, 0, 1\}) \cup (\{check_1^1\} \times \{s_1, 0, 1\})$$

Intuitively, for a state $(t, m) \in Q_{pop}^2$, m represents the last symbol which has been popped from the stack. Moreover, the flag $t \in \{check_2, check_1, check_1^0, check_1^1\}$ has the following meaning:

- $check_2$ is associated with the unique root's child x of a check 2-block tree $\langle T, Lab \rangle$. The node x is labeled with the marked content $(check_2, b)$ of the 2-block encoded by $\langle T, Lab \rangle$.
- $t = check_1^0$ is associated with the nodes of the main path of a check 2-block tree $\langle T, Lab \rangle$ whose labels are in $\{s_1, e_1, s_2, 0, 1\}$.
- $t = check_1^1$ (resp., $t = check_1$) is related to the nodes of the secondary branches of a check 2-block tree $\langle T, Lab \rangle$ whose labels are in $\{s_1, 0, 1\}$ (resp., whose labels are of form $(check_1, b)$, i.e., the marked contents of 1-blocks).

For each $(t, m) \in Q_{pop}^2$, the propositional labeling of state (t, m) is $\{t, m\}$ if $t \in \{check_1, check_2\}$, and $\{m\}$ otherwise. All the configurations associated with the states in Q_{pop}^2 are controlled by the system with the exception of the configurations associated with the state $(check_1^0, e_1)$, which are instead controlled by the environment. Transitions from states in Q_{pop}^2 are pop transitions. They are defined as follows. Note that all the configurations associated with states in $Q_{pop}^2 \setminus \{(check_1^0, e_1)\}$ are deterministic.

- Transitions from states $(check_2, b) \in Q_{pop}^2$ (note that $b \in \{0, 1\}$):
 - $(check_2, b) \xrightarrow{\text{pop}(e_1)} (check_1^0, e_1)$.
 - $(check_2, b) \xrightarrow{\text{pop}(\gamma)} q_\emptyset$ for all $\gamma \neq e_1$.
- Transitions from states $(check_1^0, m) \in Q_{pop}^2$ (note that $m \in \{s_1, e_1, s_2, 0, 1\}$):
 - $(check_1^0, m) \xrightarrow{\text{pop}(\gamma)} (check_1^0, \gamma)$ for all $m \in \{s_1, 0, 1\}$ and $\gamma \in \{s_1, e_1, s_2, 0, 1\}$.
 - $(check_1^0, e_1) \xrightarrow{\text{pop}(b)} (check_1^0, b)$ and $(check_1^0, e_1) \xrightarrow{\text{pop}(b)} (check_1, b)$ for all $b \in \{0, 1\}$.
 - $(check_1^0, m) \xrightarrow{\text{pop}(\gamma)} q_\emptyset$ if either $m = s_2$ or $\gamma \notin \{s_1, e_1, s_2, 0, 1\}$.
- Transitions from states $(check_1, b) \in Q_{pop}^2$ (note that $b \in \{0, 1\}$):
 - $(check_1, b) \xrightarrow{\text{pop}(b')} (check_1^1, b')$ for all $b' \in \{0, 1\}$.
 - $(check_1, b) \xrightarrow{\text{pop}(\gamma)} q_\emptyset$ for all $\gamma \notin \{0, 1\}$.
- Transitions from states $(check_1^1, m) \in Q_{pop}^2$ (note that $m \in \{s_1, 0, 1\}$):
 - $(check_1^1, b) \xrightarrow{\text{pop}(\gamma)} (check_1^1, \gamma)$ for all $b \in \{0, 1\}$ and $\gamma \in \{s_1, 0, 1\}$.
 - $(check_1^1, m) \xrightarrow{\text{pop}(\gamma)} q_\emptyset$ if either $m \notin \{0, 1\}$ or $\gamma \notin \{s_1, 0, 1\}$.

By construction, the following claim holds,

Claim 3. Let ν be a sequence of ATM configuration codes of the form $\nu = \rho \cdot \{acc\}$ and \mathcal{T} be the CGT obtained by unwinding $\mathcal{G}(\mathcal{S})$ from configuration $(acc, \text{Stack}(\nu))$. Then, the environments strategy trees of \mathcal{T} correspond to the check trees associated with ν .

By Claims 1–3, it follows that the set of 2^{AP} -labeled trees $\langle T, \text{Lab} \rangle$ associated with the accepting environment strategy trees $\langle T, \text{Lab}, \tau \rangle$ in $\text{exec}(\mathcal{G}(\mathcal{S}))$ coincides with the set of extended tree-codes. Moreover, the unique nodes controlled by the system in a check-subtree of $\langle T, \text{Lab}, \tau \rangle$ are the $check_3$ -branching nodes and the $\{e_2\}$ -branching nodes. This concludes the proof of Lemma 4.4. \square

4.3. Construction of the ATL* formula φ in Theorem 4.2. We now illustrate in detail the construction of the ATL* formula φ in Theorem 4.2. To this end, we need an additional definition and a preliminary result.

Definition 4.5 (Well-formed Check-trees). A check-tree $\langle T, \text{Lab} \rangle$ for a sequence ν of ATM configuration codes is *well-formed* if

- $\langle T, \text{Lab} \rangle$ satisfies the *goodness property*, which means that:
 - there are no \widehat{check}_3 -branching nodes,¹ i.e., the unique branching e_3 -nodes are the $check_3$ -branching nodes.² This entails that the subtree rooted at the $\{s_3\}$ -node of a $check_3$ -marked 3-block contains at most one \widehat{check}_3 -marked 3-block.

¹Recall that a \widehat{check}_3 -branching node is a e_3 -node having two children, one labeled by $\{\widehat{check}_3, \lambda\}$ and one which is not marked and is labeled by $\{\lambda\}$ for some $\lambda \in \Lambda$.

²Recall that the $check_3$ -branching nodes in the check-subtrees of the accepting environment strategy trees of $\mathcal{G}(\mathcal{S})$ are controlled by the system player.

- Each $\{\mathbf{e}_1\}$ -node in a partial check 2-block-tree has two children (i.e., all the choices in the $\{\mathbf{e}_1\}$ -branching nodes are enabled).³
- The ATM configuration codes in ν are well-formed;
- ν starts with the code of the initial configuration for α ;
- *fairness condition*: ν is faithful to the evolution of \mathcal{M} and for each path visiting a (well-formed) $check_3$ -marked 3-block bl_3 and a (well-formed) \widehat{check}_3 -marked 3-block bl'_3 , bl_3 and bl'_3 have the same number.

Next we show the following preliminary result.

Lemma 4.6. *One can construct in time polynomial in n and $|AP|$, three CTL* formulas φ_{good} , φ_{init} and φ_{3bl} over AP satisfying the following for each check-tree $\langle T_c, Lab_c \rangle$, where ν is the sequence of ATM configuration codes associated with $\langle T_c, Lab_c \rangle$:*

- $\langle T_c, Lab_c \rangle$ satisfies φ_{good} iff $\langle T_c, Lab_c \rangle$ satisfies the goodness property in Definition 4.5;
- $\langle T_c, Lab_c \rangle$ satisfies φ_{init} iff the first configuration code of ν is associated with an ATM configuration of the form $(q_0, \alpha(0))\alpha(1)\dots\alpha(n-1) \cdot (\#)^k$ for some $k \geq 0$;
- if $\langle T_c, Lab_c \rangle$ is good, then $\langle T_c, Lab_c \rangle$ satisfies φ_{3bl} iff the 3-blocks along ν are well-formed.

Proof. Fix a check-tree $\langle T_c, Lab_c \rangle$ and let ν be the sequence of ATM configuration codes associated with $\langle T_c, Lab_c \rangle$.

The CTL* formula φ_{good} ensuring the goodness property in Definition 4.5 is defined as follows:

$$\begin{aligned} \varphi_{good} := & \text{AG}(\mathbf{e}_3 \rightarrow \neg(\text{EX} \widehat{check}_3 \wedge \text{EX} \neg \widehat{check}_3)) \wedge \\ & \text{AG}(\widehat{check}_2 \rightarrow \text{AG}(\mathbf{e}_1 \rightarrow (\text{EX} \widehat{check}_1 \wedge \text{EX} \neg \widehat{check}_1))) \end{aligned}$$

where

- the first conjunct ensures that there are no \widehat{check}_3 -branching nodes, i.e., no \mathbf{e}_3 -node of the check-tree has both a child marked by \widehat{check}_3 and a child which is not marked by \widehat{check}_3 ;
- the second conjunct asserts that each \mathbf{e}_1 -node associated with a marked 2-block has exactly two children. Recall that each \mathbf{e}_1 -node associated with a marked 2-block has at most two children, one which is not marked and the other one which is marked by \widehat{check}_1 .

The definition of the CTL* formula φ_{init} is involved but standard.

$$\begin{aligned} \varphi_{init} := & \text{EF} \left((acc \vee \exists \vee \forall) \wedge ((\neg l \wedge \neg r) \cup (l \wedge \neg \text{EX} \bigvee_{p \in AP} p)) \wedge \right. \\ & \left. (\mathbf{e}_3 \rightarrow \text{X} \psi_{\#}) \cup (\mathbf{e}_3 \wedge \text{X}(\psi_n \wedge (\neg \mathbf{e}_3 \cup (\mathbf{e}_3 \wedge \text{X}(\psi_{n-1} \wedge \dots (\neg \mathbf{e}_3 \cup (\mathbf{e}_3 \wedge \text{X}(\psi_1 \wedge \text{X} \neg \mathbf{e}_3))) \dots)))) \right) \end{aligned}$$

$$\begin{aligned} \text{where } \psi_{\#} := & \bigvee_{(u_p, \#, u_s) \in \Lambda} (u_p, \#, u_s), \quad \psi_1 := \bigvee_{(u_p, (q_0, \alpha(0)), u_s) \in \Lambda} (u_p, (q_0, \alpha(0)), u_s), \text{ and for all} \\ 2 \leq i \leq n, \psi_i := & \bigvee_{(u_p, \alpha(i-1), u_s) \in \Lambda} (u_p, \alpha(i-1), u_s). \end{aligned}$$

Recall that the labelings of the paths along the check-tree $\langle T_c, Lab_c \rangle$ are associated to the reverse of ν and the first symbol (resp., the last symbol) of a configuration code is of the form $\{p\}$ where $p \in \{l, r\}$ (resp., $p \in \{acc, \exists, \forall\}$). Moreover, each path of the check-tree $\langle T_c, Lab_c \rangle$ has a suffix labeled by \emptyset^ω . Thus, the previous formula asserts that the last configuration code along the reverse of ν (corresponding to the first configuration code of ν) has the form $(q_0, \alpha(0))\alpha(1)\dots\alpha(n-1) \cdot (\#)^k$ for some $k \geq 0$.

³Recall that the $\{\mathbf{e}_1\}$ -nodes in the check-subtrees of the accepting environment strategy trees of $\mathcal{G}(\mathcal{S})$ are controlled by the environment.

Construction of the CTL^* formula φ_{3bl} . Assuming that the check-tree $\langle T_c, \text{Lab}_c \rangle$ is good, the CTL^* formula φ_{3bl} requires that the 3-blocks along ν are well-formed (hence, the n -bit and 2^n -bit counters in a 3-block are properly updated).

$$\varphi_{3bl} := \varphi_{2bl} \wedge \varphi_{2,first} \wedge \varphi_{2,last} \wedge \varphi_{2,inc}$$

The conjunct φ_{2bl} checks that the 2-blocks are well-formed. Again we recall that the labelings of the paths along the check-tree $\langle T_c, \text{Lab}_c \rangle$ are associated to the reverse of ν .

$$\begin{aligned} \varphi_{2bl} := & \text{AG}\left(e_1 \rightarrow (X^{n+2}s_1 \wedge \bigwedge_{i=1}^{n+1} \bigvee_{b \in \{0,1\}} X^i b)\right) \wedge \text{AG}\left((e_1 \wedge X^{n+3}s_2) \rightarrow \bigwedge_{i=2}^{n+1} X^i 0\right) \wedge \\ & \text{AG}\left((\neg s_1 \wedge X e_1) \rightarrow \bigwedge_{i=3}^{n+2} X^i 1\right) \wedge \text{AG}\left((e_1 \wedge X^{n+3}e_1) \rightarrow \right. \\ & \left. \bigvee_{i=2}^{n+1} \left[(X^i 1 \wedge X^{n+3+i} 0) \wedge \bigwedge_{j=2}^{i-1} \bigvee_{b \in \{0,1\}} (X^j b \wedge X^{n+3+j} b) \wedge \bigwedge_{j=i+1}^{n+1} (X^j 0 \wedge X^{n+3+j} 1) \right] \right) \end{aligned}$$

where:

- the first conjunct in the definition of φ_{2bl} ensures well-formedness of 1-blocks. Recall that the reverse of a well-formed 1-block is of the form $\{e_1\}\{b\}\{b_1\} \dots \{b_n\}\{s_1\}$, where $b, b_1, \dots, b_n \in \{0, 1\}$ and b is the content of the 1-block.
- The second conjunct ensures that the first 1-block bl_1 of a 2-block has number 0, i.e., the reverse of bl_1 has the form $\{e_1\}\{b\}\{0\} \dots \{0\}\{s_1\}$ for some $b \in \{0, 1\}$.
- The third conjunct ensures that the last 1-block bl_1 of a 2-block has number $2^n - 1$, i.e., the reverse of bl_1 has the form $\{e_1\}\{b\}\{1\} \dots \{1\}\{s_1\}$ for some $b \in \{0, 1\}$.
- Finally, the last conjunct ensures that for two adjacent 1-blocks bl_1 and bl'_1 along a 2-block, bl_1 and bl'_1 have consecutive numbers.

The second conjunct $\varphi_{2,first}$ in the definition of φ_{3bl} ensures that the first 2-block bl_2 of a 3-block along ν has number 0, i.e., the content of each 1-sub-block of bl_2 is 0.

$$\varphi_{2,first} := \text{AG}\left([e_2 \wedge X(\neg e_2 \cup s_3)] \rightarrow X[(\neg e_2 \wedge (e_1 \rightarrow X0)) \cup s_3]\right)$$

The second conjunct $\varphi_{2,last}$ guarantees that the last 2-block bl_2 of a 3-block has number $2^{2^n} - 1$, i.e., the content of each 1-sub-block of bl_2 is 1.

$$\varphi_{2,last} := \text{AG}\left([\neg s_2 \wedge X e_2 \wedge F s_2] \rightarrow X[(\neg s_2 \wedge (e_1 \rightarrow X1)) \cup s_2]\right)$$

Finally, the last conjunct $\varphi_{2,inc}$ in the definition of φ_{3bl} guarantees that for all adjacent 2-blocks bl_2 and bl'_2 of a 3-block along ν , bl_2 and bl'_2 have consecutive numbers. For this, assuming that bl'_2 follows bl_2 along the reverse of ν , we need to check that there is a 1-sub-block \bar{bl}_1 of bl_2 whose content is 1 and the following holds:

- the 1-sub-block of bl'_2 with the same number as \bar{bl}_1 has content 0;
- Let bl_1 be a 1-sub-block of bl_2 distinct from \bar{bl}_1 , and bl'_1 be the 1-sub-block of bl'_2 having the same number as bl_1 . Then, bl_1 and bl'_1 have the same content if bl_1 precedes \bar{bl}_1 along the reverse of bl_2 ; otherwise, the content of bl_1 is 0 and the content of bl'_1 is 1.

In order to check these conditions, we exploit the branches of a check 2-block-tree $\langle T', \text{Lab}' \rangle$ in $\langle T_c, \text{Lab}_c \rangle$ associated with (a copy of) bl'_2 which lead to $check_1$ -marked copies of the 1-sub-blocks of bl'_2 (see Figure 3(c)). Note that these branches consist (of the reverse) of a $check_1$ -marked 1-sub-block of bl'_2 followed by the suffix \emptyset^ω . Moreover, since $\langle T_c, \text{Lab}_c \rangle$ is

good all the choices in the e_1 -nodes of $\langle T', Lab' \rangle$ are enabled (i.e, for each 1-sub-block bl'_1 of bl'_2 , there is a branch for the $check_1$ -marked copy of bl'_1). Then, the formula $\varphi_{2,inc}$ is defined as follows.

$$\varphi_{2,inc} := \text{AG} \left((e_2 \wedge X(\neg s_3 \cup e_2)) \longrightarrow X \left[\left\{ \neg e_2 \wedge (e_1 \rightarrow \bigvee_{b \in \{0,1\}} \theta(b, b)) \right\} \cup \left\{ \theta(1, 0) \wedge e_1 \wedge X((\neg e_2 \wedge (e_1 \rightarrow \theta(0, 1))) \cup e_2) \right\} \right] \right)$$

where for all $b, b' \in \{0, 1\}$, the auxiliary subformula $\theta(b, b')$ in the definition of $\varphi_{2,inc}$ requires that for the current 1-sub-block bl_1 of bl_2 and for the path from bl_1 which leads to the $check_1$ -marked copy bl'_1 of the 1-sub-block of bl'_2 having the same number as bl_1 , the following holds: the content of bl_1 is b and the content of bl'_1 is b' .

$$\theta(b, b') := Xb \wedge E \left([\neg e_2 \cup (e_2 \wedge X(check_2 \wedge F(check_1 \wedge b')))] \wedge \bigwedge_{i=1}^n \bigvee_{c \in \{0,1\}} [X^{i+1}c \wedge F(check_1 \wedge X^i c)] \right)$$

This concludes the proof of Lemma 4.6. \square

An extended tree-code $\langle T_e, Lab_e \rangle$ is *well-formed* if each check-tree in $\langle T_e, Lab_e \rangle$ is well-formed. Evidently, there is a well-formed extended tree-code if and only if there is an accepting computation tree of \mathcal{M} over α . By exploiting Lemma 4.6, we now establish the following result that together with Lemma 4.4 provides a proof of Theorem 4.2.

Lemma 4.7. *One can construct in time polynomial in n and $|AP|$, an ATL^* formula φ over AP and $Ag = \{env, sys\}$ such that for each environment strategy tree $\mathcal{T} = \langle T, Lab, \tau \rangle$ in $\text{exec}(\mathcal{G}(\mathcal{S}))$, \mathcal{T} is a model of φ iff $\langle T, Lab \rangle$ is a well-formed extended tree-code.*

Proof. By Lemma 4.4, the set of $2^{|AP|}$ -labeled trees associated with the *accepting* environment strategy trees of $\mathcal{G}(\mathcal{S})$ coincides with the set of extended tree-codes. Let φ_{good} , φ_{init} , and φ_{3bl} be the CTL^* formulas of Lemma 4.6 having size polynomial in n and $|AP|$. Note that since the paths quantifiers of CTL^* correspond to the strategic quantifiers $\langle\langle\emptyset\rangle\rangle$ and $\langle\langle Ag\rangle\rangle$, each CTL^* formula can be seen as an ATL^* formula. Then, the ATL^* formula φ is given by

$$\varphi := \text{AF} acc \wedge \text{AG}(acc \rightarrow (\varphi_{good} \wedge \varphi_{init} \wedge \varphi_{3bl} \wedge \varphi_{conf} \wedge \varphi_{fair}))$$

where for an environment strategy tree $\mathcal{T} = \langle T, Lab, \tau \rangle$ of the PMS \mathcal{S} of Lemma 4.4, the first conjunct ensures that \mathcal{T} is accepting (recall that \mathcal{T} is accepting iff each path from the root visits an $\{acc\}$ -labeled node), while the subformulas φ_{good} , φ_{init} , φ_{3bl} , φ_{conf} , and φ_{fair} ensure that each check-tree $\langle T_c, Lab_c \rangle$ of \mathcal{T} is well-formed. Hence, an environment strategy tree $\mathcal{T} = \langle T, Lab, \tau \rangle$ of $\mathcal{G}(\mathcal{S})$ satisfies φ iff $\langle T, Lab \rangle$ is a well-formed extended tree-code.

Fix a check-tree $\langle T_c, Lab_c \rangle$ of an accepting environment strategy tree of the PMS \mathcal{S} , and let ν be the sequence of ATM configuration codes associated with $\langle T_c, Lab_c \rangle$. By Lemma 4.6, we have that:

- $\langle T_c, Lab_c \rangle$ satisfies φ_{good} iff it satisfies the goodness property in Definition 4.5;
- φ_{init} guarantees that the first configuration code of ν is associated with an ATM configuration of the form $(q_0, \alpha(0))\alpha(1)\dots\alpha(n-1) \cdot (\#)^k$ for some $k \geq 0$;
- φ_{3bl} enforces well-formedness of 3-blocks along ν .

We now consider the conjuncts φ_{conf} and φ_{fair} of φ which ensure the following properties for the given check-tree $\langle T_c, Lab_c \rangle$:

- φ_{conf} requires that the ATM configuration codes along ν are well-formed;

- φ_{fair} ensures that ν satisfies the fairness condition in Definition 4.5.

By means of the formulas φ_{good} and φ_{3bl} , we can assume that the check-tree $\langle T_c, Lab_c \rangle$ is good and all the 3-blocks along ν are well-formed. For defining the ATL^* formulas φ_{conf} and φ_{fair} , we exploit the following pattern: starting from an $\{\mathbf{e}_2\}$ -node x_{bl_2} related to a 2-block bl_2 of the good check-tree $\langle T_c, Lab_c \rangle$, we need to *isolate* another 2-block bl'_2 following bl_2 along the reverse of ν and checking, in particular, that bl_2 and bl'_2 have the same number. Moreover, for the case of the formula φ_{conf} , we require that the 3-block of bl'_2 is adjacent to the 3-block of bl_2 within the same ATM configuration code, while for the case of the formula φ_{fair} , we require that the 3-block of bl_2 (resp., bl'_2) is *check*₃-marked (resp., $\widehat{\text{check}}_3$ -marked) in the considered path of $\langle T_c, Lab_c \rangle$.

Recall that in a good check-tree, the unique nodes controlled by the system are the *check*₃-branching nodes and the $\{\mathbf{e}_2\}$ -nodes, and each unmarked 2-block is associated with a check 2-block-tree (2-*CBT* for short). In particular, in a 2-*CBT*, all the nodes, but the root (which is an $\{\mathbf{e}_2\}$ -node), are controlled by the environment. Moreover, each strategy of the system selects exactly one child for each node controlled by the system. Hence, there is a strategy f_{bl_2} of the player system such that

- (*) each play consistent with the strategy f_{bl_2} starting from the $\{\mathbf{e}_2\}$ -node x_{bl_2} “gets trapped” in the 2-*CBT* of bl'_2 , and
- (**) each path starting from the node x_{bl_2} and leading to some marked 1-block of the 2-*CBT* for bl'_2 is consistent with the strategy f_{bl_2} .

Thus, in order to *isolate* a 2-block bl'_2 , an ATL^* formula “guesses” the strategy f_{bl_2} and check that conditions (*) and (**) are fulfilled by simply requiring that each outcome from the current node x_{bl_2} visits a node marked by proposition *check*₂. Additionally, by exploiting the branches of the 2-*CBT* leading to marked 1-blocks, we can check by a formula of size polynomial in n and the size of \mathcal{M} that bl_2 and bl'_2 have the same number. We now proceed with the technical details about the construction of the ATL^* formulas φ_{conf} and φ_{fair} .

Construction of the ATL^* formula φ_{conf} . The ATL^* formula φ_{conf} is defined as follows.

$$\varphi_{conf} := \varphi_{3,first} \wedge \varphi_{3,last} \wedge \varphi_{3,inc}$$

The conjunct $\varphi_{3,first}$ requires that the first 3-block bl_3 of an ATM configuration code along ν has number 0, i.e., the content of each 2-sub-block of bl_3 is 0.

$$\varphi_{3,first} := \text{AG} \left([\mathbf{e}_3 \wedge \mathbf{X}(\neg \mathbf{e}_3 \cup (l \vee r))] \longrightarrow \mathbf{X}[(\neg \mathbf{e}_3 \wedge (\mathbf{e}_2 \rightarrow \mathbf{X}0)) \cup (l \vee r)] \right)$$

The second conjunct $\varphi_{3,last}$ guarantees that the last 3-block bl_3 of an ATM configuration code has number $\text{Tower}(n, 3) - 1$, i.e., the content of each 2-sub-block of bl_3 is 1).

$$\varphi_{3,last} := \text{AG} \left([\neg s_3 \wedge \mathbf{X}e_3 \wedge \mathbf{F}s_3] \longrightarrow \mathbf{X}[(\neg s_3 \wedge (\mathbf{e}_2 \rightarrow \mathbf{X}1)) \cup s_3] \right)$$

The last conjunct $\varphi_{3,inc}$ in the definition of φ_{conf} checks that for all *adjacent* 3-blocks bl_3 and bl'_3 of an ATM configuration code along ν , bl_3 and bl'_3 have consecutive numbers. For this, assuming that bl'_3 follows bl_3 along the reverse of ν , we need to check that there is a 2-sub-block \bar{bl}_2 of bl_3 whose content is 1 and the following holds:

- the 2-sub-block of bl'_3 with the same number as \bar{bl}_2 has content 0;
- Let bl_2 be a 2-sub-block of bl_3 distinct from \bar{bl}_2 , and bl'_2 be the 2-sub-block of bl'_3 having the same number as bl_2 . Then, bl_2 and bl'_2 have the same content if bl_2 precedes \bar{bl}_2 along the reverse of bl_3 ; otherwise, the content of bl_2 is 0 and the content of bl'_2 is 1.

Formula $\varphi_{3,inc}$ is then defined as follows.

$$\varphi_{3,inc} := \text{AG} \left((e_3 \wedge X((\neg l \wedge \neg r) \cup e_3)) \longrightarrow X \left[\left\{ \neg e_3 \wedge (e_2 \rightarrow \bigvee_{b \in \{0,1\}} \eta(b, b)) \right\} \right. \right. \\ \left. \left. \cup \left\{ \eta(1, 0) \wedge e_2 \wedge X((\neg e_3 \wedge (e_2 \rightarrow \eta(0, 1))) \cup e_3) \right\} \right] \right)$$

where for all $b, b' \in \{0, 1\}$, we exploit the auxiliary formula $\eta(b, b')$ to require from the current e_2 -node x of the current 2-sub-block bl_2 of bl_3 that the content of bl_2 is b and the 2-sub-block bl'_2 of bl'_3 having the same number as bl_2 has content b' . In order to ensure the last condition, the formula $\eta(b, b')$ asserts the existence of a strategy f_x of the player system such that the following two conditions hold:

- (1) each outcome of f_x from the node x visits a node marked by $check_2$ whose parent (e_2 -node) belongs to a 2-block of bl'_3 . This ensures that all the outcomes “get trapped” in the *same* check 2-block-tree associated with some 2-block bl'_2 of bl'_3 . Moreover, bl'_2 has content b' .
- (2) For each outcome π' of f_x from x which leads to a marked 1-sub-block bl'_1 (hence, a marked copy of a 1-sub-block of bl'_2), denoting by bl_1 the 1-sub-block of bl_2 having the same number as bl'_1 , it holds that bl_1 and bl'_1 have the same content. This ensures that bl_2 and bl'_2 have the same number.

The first (resp., second) condition is implemented by the first (resp., second) conjunct in the argument of the strategic quantifier $\langle\langle sys \rangle\rangle$ in the definition of $\eta(b, b')$ below.

$$\eta(b, b') := Xb \wedge \langle\langle sys \rangle\rangle \left(\left[\neg e_3 \cup (e_3 \wedge X(\neg e_3 \cup (check_2 \wedge b'))) \right] \wedge \right. \\ \left. [F check_1 \rightarrow X((\neg e_2 \wedge (e_1 \rightarrow X\eta_1)) \cup s_2)] \right)$$

$$\eta_1 := \left(\bigwedge_{i=1}^{i=n} \bigvee_{b \in \{0,1\}} ((X^i b) \wedge F(check_1 \wedge X^i b)) \right) \longrightarrow \bigvee_{b \in \{0,1\}} (b \wedge F(check_1 \wedge b))$$

Note that for each outcome π' of strategy f_x which leads to a marked 1-sub-block bl'_1 of bl'_2 , the subformula η_1 of $\eta(b, b')$ is asserted at the content node of each 1-sub-block bl_1 of bl_2 . Thus, η_1 requires that whenever bl_1 and bl'_1 have the same number, then bl_1 and bl'_1 have the same content as well.

Construction of the ATL^* formula φ_{fair} . We can assume that the check-tree $\langle T_c, Lab_c \rangle$ is good and all the ATM configuration codes along ν are well-formed. Since $\langle T_c, Lab_c \rangle$ satisfies the goodness property, for each $check_3$ -marked 3-block bl_3 which does not belong to the first configuration code of ν , there is *exactly one* $\widehat{check_3}$ -marked 3-block bl'_3 in the subtree of $\langle T_c, Lab_c \rangle$ rooted at the s_3 -node of bl_3 . Moreover, bl_3 and bl'_3 belong to two adjacent configuration codes along ν . Thus, by construction, in order to ensure that ν is faithful to the evolution of \mathcal{M} , it suffices to require that for each (well-formed) $check_3$ -marked 3-block bl_3 in $\langle T_c, Lab_c \rangle$ which does not belong to the first configuration code of ν , the associated (well-formed) $\widehat{check_3}$ -marked 3-block bl'_3 satisfies the following conditions, where (u_p, u, u_s) (resp., (u'_p, u', u'_s)) is the content of bl_3 (resp., bl'_3)

- bl_3 and bl'_3 have the same number,
- $u = next_l(u'_p, u', u'_s)$ if l marks the ATM configuration code of bl_3 , and $u = next_r(u'_p, u', u'_s)$ otherwise.

Thus, formula φ_{fair} is defined as follows:

$$\varphi_{fair} := \bigwedge_{dir \in \{l, r\}} \mathsf{AG} \left(\left[\mathit{check}_3 \wedge [(\neg l \wedge \neg r) \vee (dir \wedge \mathsf{X}(\exists \vee \forall))] \right] \longrightarrow \left[((\neg e_3 \wedge (e_2 \rightarrow \psi_=_)) \vee s_3) \wedge \bigvee_{(u_p, u, u_s), (u'_p, u', u'_s) \in \Lambda: u = \mathit{next}_{dir}(u'_p, u', u'_s)} ((u_p, u, u_s) \wedge \mathsf{EF}(\widehat{\mathit{check}}_3 \wedge (u'_p, u', u'_s))) \right] \right)$$

where the auxiliary formula $\psi_=_$ in the definition of φ_{fair} requires from the current e_2 -node x of the current 2-sub-block bl_2 of bl_3 that the 2-sub-block bl'_2 of bl'_3 having the same number as bl_2 has the same content as bl_2 too. In order to ensure the last condition, the formula $\psi_=_$ asserts the existence of a strategy f_x of the player system such that the following holds:

- (1) each outcome of f_x from the node x visits a node marked by check_2 whose parent (e_2 -node) belongs to a $\widehat{\mathit{check}}_3$ -marked 3-block. This ensures that all the outcomes “get trapped” in the *same* 2-block check-tree associated with some 2-block bl'_2 of bl'_3 . Moreover, bl_2 and bl'_2 have the same content.
- (2) For each outcome π' of f_x from x which leads to a marked 1-sub-block bl'_1 (hence, a marked copy of a 1-sub-block of bl'_2), denoting by bl_1 the 1-sub-block of bl_2 having the same number as bl'_1 , it holds that bl_1 and bl'_1 have the same content. This ensures that bl_2 and bl'_2 have the same number.

Thus, formula $\psi_=_$ is defined as follows.

$$\psi_=_ := \langle\langle \mathit{sys} \rangle\rangle \left(\bigvee_{b \in \{0, 1\}} [\mathsf{X}b \wedge \mathsf{F}\{\widehat{\mathit{check}}_3 \wedge (\neg e_3 \vee (b \wedge \mathit{check}_2))\}] \wedge [\mathsf{F}\mathit{check}_1 \rightarrow \mathsf{X}((\neg e_2 \wedge (e_1 \rightarrow \mathsf{X}\eta_1)) \vee s_2)] \right)$$

where η_1 corresponds to the homonymous subformula of the auxiliary formula $\eta(b, b')$ used in the definition of $\varphi_{3,inc}$. This concludes the proof of Lemma 4.7. \square

5. CONCLUSION

Module checking is a useful game-theoretic framework to deal with branching-time specifications. The setting is simple and powerful as it allows to capture the essence of the adversarial interaction between an open system (possibly consisting of several independent components) and its unpredictable environment. The work on module checking has brought an important contribution to the strategic reasoning field, both in computer science and AI [AHK02]. It is known [JM14] that CTL/CTL* module checking is incomparable with ATL/ATL* model checking. In particular the former can keep track of all moves made in the past, while the latter cannot. This is a severe limitation in ATL/ATL* and has been studied under the name of irrevocability of strategies in [ÅGJ07]. Remarkably, this feature can be handled with more sophisticated logics such as *Strategy Logics* [CHP10, MMPV14], *ATL with strategy contexts* [LM15], and *quantified CTL* [LM14]. However, for such logics, the relative model checking question for finite-state multi-agent systems (modelled by finite-state concurrent game structures) turns out to be non-elementarily decidable.

In this paper, we have addressed the module-checking problem of multi-agent pushdown systems (PMS) against ATL and ATL* specifications. PMS endow finite-state multi-agent systems with an additional expressive power, the possibility of using a stack to store unbounded information. The stack is the standard low level mechanism which allows to

structure agents in modules and to implement recursive calls and returns of modules. Hence, the considered framework is suitable for formally reasoning on the behaviour of software agents with (recursive) procedural modularity. As a main contribution, we have established the exact computational complexity of pushdown module-checking against ATL and ATL*. While for ATL, the considered problem is 2EXPTIME-complete, which is the same complexity as pushdown module-checking for CTL, for ATL*, pushdown module-checking turns out to be 4EXPTIME-complete, hence exponentially harder than both CTL* pushdown module-checking and ATL* model-checking of PMS. As future work, we aim to investigate the considered problems in the setting of *imperfect information under memoryless strategies*. We recall that this setting is decidable in the finite-state case [AHK02]. However, moving to pushdown systems one has to distinguish whether the missing information relies in the control states, in the pushdown store, or both. We recall that in pushdown module-checking only the former case is decidable for specifications given in CTL and CTL* [ALM⁺13].

Another interesting question to investigate is the exact computational complexity of pushdown module checking against the fragment ATL⁺ of ATL*, where each temporal modality is immediately preceded either by a strategic quantifier or by a Boolean connective. Our results just imply that pushdown module checking against ATL⁺ lies somewhere between 2EXPTIME and 4EXPTIME.

REFERENCES

- [ÅGJ07] T. Ågotnes, V. Goranko, and W. Jamroga. Alternating-time temporal logics with irrevocable strategies. In *Proc. 11th TARK*, pages 15–24, 2007. doi:10.1145/1324249.1324256.
- [AHK02] R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49(5):672–713, 2002. doi:10.1145/585265.585270.
- [AKM12] B. Aminof, O. Kupferman, and A. Murano. Improved model checking of hierarchical systems. *Inf. Comput.*, 210:68–86, 2012. doi:10.1016/j.ic.2011.10.008.
- [ALM⁺13] B. Aminof, A. Legay, A. Murano, O. Serre, and M. Y. Vardi. Pushdown module checking with imperfect information. *Inf. Comput.*, 223(1):1–17, 2013. doi:10.1016/j.ic.2012.11.005.
- [AMM14] B. Aminof, F. Mogavero, and A. Murano. Synthesis of hierarchical systems. *Sci. Comput. Program.*, 83:56–79, 2014. doi:10.1016/j.scico.2013.07.001.
- [BEM97] A. Bouajjani, J. Esparza, and O. Maler. Reachability Analysis of Pushdown Automata: Application to Model-Checking. In *Proc. 8th CONCUR*, LNCS 1243, pages 135–150. Springer, 1997. doi:10.1007/3-540-63141-0__10.
- [BM17] L. Bozzelli and A. Murano. On the complexity of ATL and ATL* module checking. In *Proc. 8th GandALF*, EPTCS 256, pages 268–282, 2017. doi:10.4204/EPTCS.256.19.
- [BMP10] L. Bozzelli, A. Murano, and A. Peron. Pushdown module checking. *Formal Methods in System Design*, 36(1):65–95, 2010. doi:10.1007/s10703-010-0093-x.
- [BMP20] Laura Bozzelli, Aniello Murano, and Adriano Peron. Module checking of pushdown multi-agent systems. In Diego Calvanese, Esra Erdem, and Michael Thielscher, editors, *Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning, KR 2020, Rhodes, Greece, September 12–18, 2020*, pages 162–171, 2020. doi:10.24963/kr.2020/17.
- [Boz06] L. Bozzelli. Complexity results on branching-time pushdown model checking. In *Proc. 7th VMCAI*, LNCS 3855, pages 65–79. Springer, 2006. doi:10.1007/11609773__5.
- [Boz11] L. Bozzelli. New results on pushdown module checking with imperfect information. In *Proc. 2nd GandALF*, EPTCS 54, pages 162–177, 2011. doi:10.4204/EPTCS.54.12.
- [BR00] T. Ball and S. Rajamani. Bebop: a symbolic model checker for boolean programs. In *Proc. 7th SPIN Workshop*, LNCS 1885, pages 113–130. Springer, 2000. doi:10.1007/10722468__7.
- [BRS07] S. Basu, P. S. Roop, and R. Sinha. Local module checking for CTL specifications. *ENTCS* 176, (2):125–141, 2007. doi:10.1016/j.entcs.2006.02.035.

[BVW94] O. Bernholtz, M.Y. Vardi, and P. Wolper. An automata-theoretic approach to branching-time model checking (extended abstract). In *Proc. 6th CAV*, LNCS 818, pages 142–155. Springer, 1994. doi:10.1007/3-540-58179-0_50.

[CE81] E.M. Clarke and E.A. Emerson. Design and synthesis of synchronization skeletons using branching time temporal logic. In *Proc. of Workshop on Logic of Programs*, LNCS 131, pages 52–71. Springer-Verlag, 1981. doi:10.1007/BFb0025774.

[CHP10] K. Chatterjee, T. A. Henzinger, and N. Piterman. Strategy logic. *Inf. Comput.*, 208(6):677–693, 2010. doi:10.1016/j.ic.2009.07.004.

[CKS81] A.K. Chandra, D.C. Kozen, and L.J. Stockmeyer. Alternation. *Journal of the ACM*, 28(1):114–133, 1981. doi:10.1145/322234.322243.

[CSW16] T. Chen, F. Song, and Z. Wu. Global Model Checking on Pushdown Multi-Agent Systems. In *Proc. 30th AAAI*, pages 2459–2465. AAAI Press, 2016. URL: <http://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/view/11984>.

[dAGJ04] L. de Alfaro, P. Godefroid, and R. Jagadeesan. Three-Valued Abstractions of Games: Uncertainty, but with Precision. In *Proc. 19th LICS*, pages 170–179. IEEE Computer Society, 2004. doi:10.1109/LICS.2004.1319611.

[EH86] E.A. Emerson and J.Y. Halpern. "Sometimes" and "Not Never" revisited: on branching versus linear time temporal logic. *Journal of the ACM*, 33(1):151–178, 1986. doi:10.1145/4904.4999.

[FMP08] A. Ferrante, A. Murano, and M. Parente. Enriched μ -calculi module checking. *Logical Methods in Computer Science*, 4(3:1):1–21, 2008. doi:10.2168/LMCS-4(3:1)2008.

[God03] P. Godefroid. Reasoning about Abstract Open Systems with Generalized Module Checking. In *Proc. 3rd EMSOFT*, LNCS 2855, pages 223–240. Springer, 2003. doi:10.1007/978-3-540-45212-6_15.

[HO09] M. Hague and C.-H. Luke Ong. Winning regions of pushdown parity games: A saturation method. In *Proc. 20th CONCUR*, LNCS 5710, pages 384–398. Springer, 2009. doi:10.1007/978-3-642-04081-8_26.

[HRV90] D. Harel, R. Rosner, and M. Y. Vardi. On the power of bounded concurrency III: Reasoning about programs (preliminary report). In *Proc. 5th LICS*, pages 478–488. IEEE Computer Society, 1990. doi:10.1109/LICS.1990.113770.

[JM14] W. Jamroga and A. Murano. On module checking and strategies. In *Proc. AAMAS*, pages 701–708. IFAAMAS/ACM, 2014.

[JM15] W. Jamroga and A. Murano. Module checking of strategic ability. In *Proc. AAMAS*, pages 227–235. ACM, 2015.

[KPV02] O. Kupferman, N. Piterman, and M.Y. Vardi. Pushdown specifications. In *Proc. 9th LPAR*, LNCS 2514, pages 262–277. Springer, 2002. doi:10.1007/3-540-36078-6_18.

[KV96] O. Kupferman and M.Y. Vardi. Module checking. In *Proc. 8th CAV*, LNCS 1102, pages 75–86. Springer, 1996. doi:10.1007/3-540-61474-5_59.

[KV97] O. Kupferman and M.Y. Vardi. Module checking revisited. In *Proc. 9th CAV*, LNCS 1254, pages 36–47. Springer, 1997. doi:10.1007/3-540-63166-6_7.

[LM14] F. Laroussinie and N. Markey. Quantified CTL: Expressiveness and Complexity. *Logical Methods in Computer Science*, 10(4), 2014. doi:10.2168/LMCS-10(4:17)2014.

[LM15] F. Laroussinie and N. Markey. Augmenting ATL with strategy contexts. *Inf. Comput.*, 245:98–123, 2015. doi:10.1016/j.ic.2014.12.020.

[LMS04] C. Löding, P. Madhusudan, and O. Serre. Visibly pushdown games. In *Proc. 24th FSTTCS*, LNCS 3328, pages 408–420. Springer, 2004. doi:10.1007/978-3-540-30538-5_34.

[MMPV14] F. Mogavero, A. Murano, G. Perelli, and M. Y. Vardi. Reasoning About Strategies: On the Model-Checking Problem. *ACM Trans. Comput. Log.*, 15(4):34:1–34:47, 2014. doi:10.1145/2631917.

[MNP08] A. Murano, M. Napoli, and M. Parente. Program complexity in Hierarchical Module Checking. In *Proc. 15th LPAR*, LNCS 5330, pages 318–332. Springer, 2008. doi:10.1007/978-3-540-89439-1_23.

[MP15] A. Murano and G. Perelli. Pushdown Multi-Agent System Verification. In *Proc. 24th IJCAI*, pages 1090–1097. AAAI Press, 2015. URL: <http://ijcai.org/Abstract/15/158>.

[Pnu77] A. Pnueli. The temporal logic of programs. In *Proc. 18th FOCS*, pages 46–57. IEEE, 1977. doi:10.1109/SFCS.1977.32.

- [QS82] J.P. Queille and J. Sifakis. Specification and verification of concurrent systems in CESAR. In *Proc. 5th SP*, LNCS 137, pages 337–351. Springer, 1982. doi:10.1007/3-540-11494-7_22.
- [Saf88] S. Safra. On the complexity of ω -automata. In *Proc. 29th FOCS*, pages 319–327. IEEE Computer Society, 1988. doi:10.1109/SFCS.1988.21948.
- [Sch08] S. Schewe. ATL* Satisfiability is 2EXPTIME-complete. In *Proc. 35th ICALP*, LNCS 5126, pages 373–385. Springer, 2008. doi:10.1007/978-3-540-70583-3_31.
- [SF06] S. Schewe and B. Finkbeiner. Satisfiability and Finite Model Property for the Alternating-Time μ -calculus. In *Proc. 20th CSL*, LNCS 4207, pages 591–605. Springer, 2006. doi:10.1007/11874683_39.
- [Var98] M.Y. Vardi. Reasoning about the past with two-way automata. In *Proc. 25th ICALP*, LNCS 1443, pages 628–641. Springer, 1998. doi:10.1007/BFb0055090.
- [Wal96] I. Walukiewicz. Pushdown processes: Games and Model Checking. In *Proc. 8th CAV*, LNCS 1102, pages 62–74. Springer, 1996. doi:10.1007/3-540-61474-5_58.
- [Wal00] I. Walukiewicz. Model checking CTL properties of pushdown systems. In *Proc. 20th FSTTCS*, LNCS 1974, pages 127–138. Springer, 2000. doi:10.1007/3-540-44450-5_10.