

On the Complexity of Resource-Bounded Logics¹

N. Alechina¹, N. Bulling², S. Demri³ and B. Logan¹

¹ University of Nottingham ² TU Clausthal ³ LSV, CNRS, ENS Paris-Saclay

Abstract

We revisit decidability results for resource-bounded logics and use decision problems on vector addition systems with states (VASS) in order to establish complexity characterisations of (decidable) model checking problems. We show that the model checking problem for the logic $\text{RB}\pm\text{ATL}$ is 2EXPTIME -complete by using recent results on alternating VASS (and in EXPTIME when the number of resources is bounded). Moreover, we establish that the model checking problem for RBTL is EXSPACE -complete. The problem is decidable and of the same complexity for RBTL^* , proving a new decidability result as a by-product of the approach. When the number of resources is bounded, the problem is in PSPACE . We also establish that the model checking problem for $\text{RB}\pm\text{ATL}^*$, the extension of $\text{RB}\pm\text{ATL}$ with arbitrary path formulae, is decidable by a reduction into parity games for single-sided VASS (a variant of alternating VASS). Furthermore, we are able to synthesise values for resource parameters. Hence, the paper establishes formal correspondences between model checking problems for resource-bounded logics advocated in the AI literature and decision problems on alternating VASS, paving the way for more applications and cross-fertilizations.

1 Introduction

Resource-bounded logics. Alternating-time temporal logics such as the logics ATL and ATL^* [5] extend the temporal logics CTL and CTL^* respectively, by interpreting the formulae on concurrent game structures, a sophisticated extension of labelled transition systems, and by allowing modalities to quantify over strategies for a given coalition of agents. ATL significantly extends CTL but the computational properties of the model checking problem remain quite stable. For instance, in [5], the labeling algorithm for model checking CTL is extended to ATL establishing the P -completeness result of the model checking problem for ATL [5, Theorem 5.2]. By contrast, the model checking problem for ATL^* is 2EXPTIME -complete [5, Theorem 5.6] whereas the problem for CTL^* is only PSPACE -complete, see e.g. [22]. The logics ATL and ATL^* are well-established formalisms to reason about multi-agent transition systems, and many variants have been proposed over the years, see e.g. [34, 4]. We focus here on resource-bounded variants.

Resource-bounded logics [13, 12, 38, 3, 2, 14] extend alternating-time temporal logics such as ATL [5] by adding transitions that produce and consume resources to the models. As shown in [2], the introduction of implicit counters in the models (i.e., variables interpreted by natural numbers) and the ability to quantify over strategies for a given set of agents can lead to undecidability, or decidability with a very high worst-case upper bound on the complexity of the model checking problem. The nature of the strategy modalities means that reasoning about resources has similarities to the analysis of runs of vector addition systems with states (a.k.a. VASS) [33],

and more specifically to games on VASS, see e.g. [11]. VASS, and more generally counter machines, are well-known infinite-state systems with many applications in formal verification, see e.g. [10].

Model checking and games on VASS. In our work, we show how existing results on VASS can be used to analyse the model checking problem for resource-bounded logics. As we recall below, model checking problems on VASS based on temporal logics and games are not always decidable; when they are decidable, they are quite difficult to solve, and complexity characterisations often exist. We briefly recapitulate some of these results below.

Temporal logics on VASS often lead to undecidable model checking problems, see e.g., [23, 24], and this is even more true with branching-time temporal logics such as CTL [24], or when the atomic formulae can state properties about the counter values [28]. However there are exceptions. For instance, CTL model checking on one-counter VASS is PSPACE-complete [43, 25] (see also [46]). The control state repeated reachability problem for VASS is shown to be decidable in [29]; this result is generalised to full LTL (for which the atomic formulae are exactly control states), and the model checking problem for LTL on VASS is shown to be EXPSpace-complete in [27]. In [29], a strict fragment of LTL restricted to the “infinitely often” temporal operator GF and atomic formulae stating properties on counter values is also shown decidable by reduction to the reachability problem for VASS.

As far as games for VASS are concerned, the situation is even less encouraging. Indeed, two-player games on VASS in which each player can freely update the counter values are undecidable [11], even with simple winning conditions such as the reachability of a given control state. However, asymmetric VASS games in which at most one player can freely update the counter values and the winning conditions are simple are decidable [40]. For instance, the game on asymmetric VASS (a slight variant of single-sided VASS in [7, 1] or alternating VASS in [16]) with reachability of a control state is shown to be 2EXPTIME-complete in [16] and decidable with parity conditions in [1, 30]². The non-termination problem for games on asymmetric VASS is also 2EXPTIME-complete (the upper bound is from [32] and the lower bound is from [16]).

Our motivation. Our main goal in this paper is to establish formal relationships between model checking problems for resource-bounded logics and decision problems for VASS, so that new decidability results can be established for logical problems or new complexity characterisations can be inherited from problems on counter machines. Of course, this is not surprising; resource values and counter values are similar objects, and logics based on concurrent game structures inherently have games in their semantics. Moreover, earlier work has already explored the connections with counter machines, either to obtain undecidability results or to get complexity lower bounds, see e.g. [2]. In this paper, we extend these results to give optimal complexity upper bounds and new decidability results, even for resource-bounded logics with enriched path formulae such as those in CTL* [22] (see also [19]).

Our contributions. As explained above, our approach is to use results on decision problems for alternating VASS (or for its variants with single-sided VASS) in order to decide optimally model checking problems on resource-bounded logics. So far, the reductions were rather in the other direction to establish undecidability (for instance, by reducing the halting problem for Minsky machines).

²Very recently, a 2EXPTIME upper bound was shown in [35].

- The model checking problem for $\text{RB}\pm\text{ATL}$ is shown to be 2EXPTIME -complete (see Theorem 2 and Theorem 3). The restriction to a bounded number of resources is also shown to be in EXPTIME . The 2EXPTIME lower bound is obtained by a reduction from the state reachability problem for alternating VASS (AVASS) [16], where the upper bound is shown by a reduction to the state reachability and the non-termination problems for AVASS. We need to consider both target problems in order to reduce our logical problem to questions on AVASS. So far, the best known result was decidability established in [3] by taking advantage of the well-quasi-ordering (\mathbb{N}^r, \preceq) .
- The results for $\text{RB}\pm\text{ATL}$ are obtained by using formal relationships between strategies in resource-bounded concurrent game structures and proofs in alternating VASS (the fact that only asymmetric VASS are needed here is the key observation). These relationships are also used to show that the model checking problem for $\text{RB}\pm\text{ATL}^*$ (a new logic extending $\text{RB}\pm\text{ATL}$ as ATL^* extends ATL [5]), is decidable, by a reduction to the parity game problem on single-sided VASS [1]. Note that the complexity characterisation of the parity game problem on single-sided VASS was left open in [1, 16, 32]. More importantly, we show that resource parameters can be effectively computed in the parameterised version of $\text{RB}\pm\text{ATL}^*$ thanks to the fact that the Pareto frontier for any parity game on single-sided VASS is computable [1, Theorem 4]. To the best of our knowledge, this is the first time that resource values have been synthesised in resource-bounded logics (see also [31]), and this is done for the rich new logic, $\text{RB}\pm\text{ATL}^*$.
- The model checking problem for RBTL [12] is shown to be EXPSpace -complete. The restriction to a bounded number of resources is also shown to be in PSPACE . The model checking problem for RBTL^* is shown to be decidable (a new result), and also EXPSpace -complete (see Theorem 5).

In addition, we also provide complexity characterisations for various logical fragments which, e.g., bound the number the resources or the number of agents. For example, the model checking problem for $\text{RB}\pm\text{ATL}$ restricted to a single agent is shown to be EXPSpace -complete (Theorem 4).

2 Logical Preliminaries

We write \mathbb{N} [resp. \mathbb{Z}] for the set of natural numbers [resp. integers] and $[m, m']$ with $m, m' \in \mathbb{Z}$ to denote the set $\{j \in \mathbb{Z} : m \leq j \leq m'\}$. Given a dimension $r \geq 1$ and $a \in \mathbb{Z}$, we write $\vec{a} \in \mathbb{Z}^r$ to denote the vector with all values equal to a . For each $\vec{x} \in \mathbb{Z}^r$, we write $\vec{x}(1), \dots, \vec{x}(r)$ for the entries of \vec{x} . For all $\vec{x}, \vec{y} \in \mathbb{Z}^r$, $\vec{x} \leq \vec{y} \stackrel{\text{def}}{\iff}$ for every $i \in [1, r]$, we have $\vec{x}(i) \leq \vec{y}(i)$. We also write $\vec{x} < \vec{y}$ when $\vec{x} \leq \vec{y}$ and $\vec{x} \neq \vec{y}$.

2.1 The logic $\text{RB}\pm\text{ATL}$ and its variants

We consider the logics $\text{RB}\pm\text{ATL}$ and $\text{RB}\pm\text{ATL}^*$. The logic $\text{RB}\pm\text{ATL}$ was introduced in [3, 4], and extends ATL [5] with resources. $\text{RB}\pm\text{ATL}^*$ extends $\text{RB}\pm\text{ATL}$ to allow path formulae to be any LTL-like formula (see Section 5 for a complete formal definition).

Let PROP be a countably infinite set of atomic propositions. The models for the logics $\text{RB}\pm\text{ATL}$ and $\text{RB}\pm\text{ATL}^*$ are the structures introduced in Definition 1 below. These are concurrent game structures for the logics ATL or ATL^* (see e.g. [5]) but enriched with a cost function

that specifies how resources are produced or consumed. Intuitively, the concurrent game structure is equipped with r counters and the transitions update their values with increments or decrements.

Definition 1. A resource-bounded concurrent game structure \mathfrak{M} is a tuple

$$\mathfrak{M} = (\text{Agt}, S, \text{Act}, r, \text{act}, \text{cost}, \delta, L)$$

such that:

- Agt is a non-empty finite set of agents (by default $\text{Agt} = [1, k]$ for some $k \geq 1$);
- S is a non-empty set of states;
- Act is a non-empty set of actions with a distinguished action idle ;
- $r \geq 1$ is the number of resources;
- $\text{act} : S \times \text{Agt} \rightarrow \mathcal{P}(\text{Act}) \setminus \{\emptyset\}$ is the action manager function, such that for all s and a , $\text{act}(s, a)$ is non-empty and furthermore we have $\text{idle} \in \text{act}(a, s)$;
- $\text{cost} : S \times \text{Agt} \times \text{Act} \rightarrow \mathbb{Z}^r$ is the (partial) cost function; that is, $\text{cost}(s, a, a)$ is defined when $a \in \text{act}(a, s)$,³ and moreover, we stipulate $\text{cost}(s, a, \text{idle}) = \vec{0}$;
- $\delta : S \times (\text{Agt} \rightarrow \text{Act}) \rightarrow S$ is the (partial) transition function such that δ is defined for a state s and joint action \mathfrak{f} whenever for all agents $a \in \text{Agt}$, we have $\mathfrak{f}(a) \in \text{act}(a, s)$;
- $L : \text{PROP} \rightarrow \mathcal{P}(S)$ is a labelling (the definition can be adapted when finite subsets of PROP are involved).

The map δ is also viewed as a deterministic transition relation with transitions of the form $s \xrightarrow{(a_1, \dots, a_k)} s'$ where $\delta(s, \mathfrak{f}) = s'$ and for all $i \in [1, k] = \text{Agt}$, we have $\mathfrak{f}(i) = a_i$. We say that \mathfrak{M} is finite whenever S and Act are finite sets and L is restricted to a finite subset of PROP . Figure 1 presents a finite concurrent game structure (costs are omitted).

The idle action was introduced in [3, 4], where motivations for requiring a distinguished 0-cost action can be found (in Section 5.5 below, we explain why the idle action is not essential for decidability). Given a coalition $A \subseteq \text{Agt}$ and a state s , a joint action by A is a map $\mathfrak{f} : A \rightarrow \text{Act}$ such that for all agents $a \in A$, we have $\mathfrak{f}(a) \in \text{act}(s, a)$. The set of joint actions by A is denoted $D_A(s)$. Given a state s , the set of joint actions by Agt is simply denoted $D(s)$ (instead of $D_{\text{Agt}}(s)$) and the map δ is defined only for such joint actions. We write $\mathfrak{f} \sqsubseteq g$ whenever g is a conservative extension of \mathfrak{f} .

Given a joint action $\mathfrak{f} \in D_A(s)$, we write $\text{out}(s, \mathfrak{f})$ to denote the set below:

$$\text{out}(s, \mathfrak{f}) \stackrel{\text{def}}{=} \{s' \in S \mid \text{there is } g \in D(s) \text{ such that } \mathfrak{f} \sqsubseteq g \text{ and } s' = \delta(s, g)\}.$$

For instance, $\text{out}(s, \mathfrak{f})$ is a singleton set when $\mathfrak{f} \in D(s)$, i.e. if an action is specified for each agent, since δ is a map and not a relation. Given a joint action $\mathfrak{f} \in D_A(s)$ and a state s , the cost of a transition from s by \mathfrak{f} (restricted to A by definition) is defined as follows:

$$\text{cost}_A(s, \mathfrak{f}) \stackrel{\text{def}}{=} \sum_{a \in A} \text{cost}(s, a, \mathfrak{f}(a)).$$

³Unlike in [3], we adopt the convention that positive costs correspond to resource production, and negative costs to resource consumption.

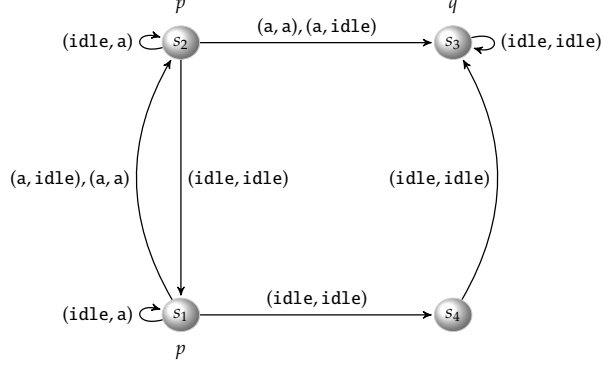


Figure 1: A concurrent game structure (costs are omitted)

Note that the value $\text{cost}_A(s, \mathbf{f})$ does not depend on the costs of the actions by the agents in the opponent coalition ($\text{Agt} \setminus A$) (or equivalently, the cost of actions by agents in ($\text{Agt} \setminus A$) is zero). More generally, given $\mathbf{g} \in D(s)$, we have

$$\text{cost}_A(s, \mathbf{g}) \stackrel{\text{def}}{=} \sum_{a \in A} \text{cost}(s, a, \mathbf{g}(a)).$$

A *computation* λ is a finite sequence or an ω -sequence of the form $s_0 \xrightarrow{\mathbf{f}_0} s_1 \xrightarrow{\mathbf{f}_1} s_2 \dots$ such that for all $0 \leq i < |\lambda| - 1$, we have $s_{i+1} \in \delta(s_i, \mathbf{f}_i)$.⁴ Here, $|\lambda|$ denotes the length of λ , each s_i is a state and each \mathbf{f}_i belongs to $D(s_i)$. For instance, $|s_0 \xrightarrow{\mathbf{f}_0} s_1 \dots \xrightarrow{\mathbf{f}_{n-1}} s_n| = n+1$ and $|s_0 \xrightarrow{\mathbf{f}_0} s_1 \dots \xrightarrow{\mathbf{f}_{n-1}} \dots| = \omega$ for any infinite computation. A *strategy* F_A for the coalition A is a map from the set of finite computations to the set of joint actions of A such that

$$F_A(s_0 \xrightarrow{\mathbf{f}_0} s_1 \dots \xrightarrow{\mathbf{f}_{n-1}} s_n) \in D_A(s_n).$$

A computation $\lambda = s_0 \xrightarrow{\mathbf{f}_0} s_1 \xrightarrow{\mathbf{f}_1} s_2 \dots$ respects the strategy F_A iff for all $i < |\lambda|$, we have, $s_{i+1} \in \text{out}(s_i, F_A(s_0 \xrightarrow{\mathbf{f}_0} s_1 \dots \xrightarrow{\mathbf{f}_{i-1}} s_i))$. A computation λ that respects F_A is *maximal* whenever it cannot be extended further while respecting the strategy. Note that maximal computations respecting F_A are infinite. The set of all maximal computations that respect the strategy F_A that start at the state s is denoted by $\text{Comp}(s, F_A)$. So far, no resource value has been involved in computations. Below, we shall quantify over maximal computations that respect a strategy, and therefore for defining a strategy we can restrict ourselves to finite computations that respect it so far.

Given a bound $\vec{b} \in (\mathbb{N} \cup \{\omega\})^r$ and a computation $\lambda = s_0 \xrightarrow{\mathbf{f}_0} s_1 \xrightarrow{\mathbf{f}_1} s_2 \dots$ in $\text{Comp}(s, F_A)$, let the resource availability at step $i \leq |\lambda|$ be defined as follows: $\vec{v}_0 \stackrel{\text{def}}{=} \vec{b}$ and for all $i < |\lambda| - 1$,

⁴Each transition between two successive states is labelled by a joint action: this is not strictly necessary for the development below, but it provides a more general notion that might be used in other contexts (for example, if the winning condition of strategies depends on the actions of all the agents and not only on those for the agents in A or on the visited states).

$\vec{v}_{i+1} \stackrel{\text{def}}{=} \text{cost}_A(s_i, \vec{f}_i) + \vec{v}_i$ (assuming that $n + \omega = \omega$ for any $n \in \mathbb{Z}$). Then, λ is \vec{b} -consistent iff for all $i < |\lambda|$, $\vec{v}_i \in (\mathbb{N} \cup \{\omega\})^r$. If $\vec{b}(i) = \omega$, we have an infinite supply of the i th resource and effectively disregard what happens on the i th resource. Since the resource availability of the sequence depends only on the agents in A , we call this the *proponent restriction condition*. This condition is very similar to that found in runs of VASS with the sequence of update vectors $\text{cost}_A(s_0, \vec{f}_0), \text{cost}_A(s_1, \vec{f}_1), \dots$. Note also that the above condition is slightly different from the one in [4] but equivalent. We have decided to use our notation in order to more easily show the relationships with VASS decision problems.

The set of all the \vec{b} -consistent (infinite) computations is denoted by $\text{Comp}(s, F_A, \vec{b})$. A \vec{b} -strategy F_A with respect to s is a strategy such that $\text{Comp}(s, F_A) = \text{Comp}(s, F_A, \vec{b})$. This definition also differs slightly from that given in [4]; the notion of \vec{b} -strategy in [4] is not relative to a state and therefore the equality should hold for all states.

With the main definitions of resource-bounded concurrent game structures and strategies in hand, we can now present the logic RB \pm ATL. Given a set of agents $\text{Agt} = \{1, \dots, k\}$ and $r \geq 1$, we write RB \pm ATL(Agt, r) to denote the resource-bounded logic with k agents and r resources whose models are resource-bounded concurrent game structures with the same parameters. Formulae of RB \pm ATL(Agt, r) are defined according to the grammar below:

$$\phi ::= p \mid \neg\phi \mid \phi \wedge \phi \mid \langle\langle A^{\vec{b}} \rangle\rangle X\phi \mid \langle\langle A^{\vec{b}} \rangle\rangle G\phi \mid \langle\langle A^{\vec{b}} \rangle\rangle \phi U \phi,$$

where $p \in \text{PROP}$, $A \subseteq \text{Agt}$ and $\vec{b} \in (\mathbb{N} \cup \{\omega\})^r$. The size of a formula is computed from a DAG representation and the integers are encoded in binary. Note that forthcoming hardness results do not use the conciseness of the DAG representation (with respect to the tree representation). The satisfaction relation \models is defined inductively as follows assuming that \mathfrak{M} is an RB \pm ATL(Agt, r) model (we omit the obvious cases for the Boolean connectives):

$$\begin{array}{ll} \mathfrak{M}, s \models p & \stackrel{\text{def}}{\iff} s \in L(p) \\ \mathfrak{M}, s \models \langle\langle A^{\vec{b}} \rangle\rangle X\phi & \stackrel{\text{def}}{\iff} \text{there is a } \vec{b}\text{-strategy } F_A \text{ w.r.t. } s \text{ such that} \\ & \text{for all } s_0 \xrightarrow{\vec{f}_0} s_1 \dots \in \text{Comp}(s, F_A), \text{ we have } \mathfrak{M}, s_1 \models \phi \\ \mathfrak{M}, s \models \langle\langle A^{\vec{b}} \rangle\rangle G\phi & \stackrel{\text{def}}{\iff} \text{there is a } \vec{b}\text{-strategy } F_A \text{ w.r.t. } s \text{ such that} \\ & \text{for all } \lambda = s_0 \xrightarrow{\vec{f}_0} s_1 \dots \in \text{Comp}(s, F_A), \text{ for all } i < |\lambda|, \\ & \text{we have } \mathfrak{M}, s_i \models \phi \\ \mathfrak{M}, s \models \langle\langle A^{\vec{b}} \rangle\rangle \phi_1 U \phi_2 & \stackrel{\text{def}}{\iff} \text{there is a } \vec{b}\text{-strategy } F_A \text{ w.r.t. } s \text{ such that for all} \\ & \lambda = s_0 \xrightarrow{\vec{f}_0} s_1 \dots \in \text{Comp}(s, F_A), \text{ there is some } i < |\lambda| \\ & \text{such that } \mathfrak{M}, s_i \models \phi_2 \text{ and} \\ & \text{for all } j \in [0, i-1], \text{ we have } \mathfrak{M}, s_j \models \phi_1. \end{array}$$

Standard semantics for temporal operators. It is worth noting that since all the maximal computations are infinite, the index i involved for clauses related to $\langle\langle A^{\vec{b}} \rangle\rangle G$ or $\langle\langle A^{\vec{b}} \rangle\rangle U$ can take any value in \mathbb{N} . The temporal operators X , G and U have their standard meaning from linear-time temporal logic LTL. CTL formulae can be expressed by RB \pm ATL formulae (as it is also classically the case with ATL) if resources are omitted from resource-bounded concurrent game structures thanks to correspondences of the form $\langle\langle \text{Agt}^{t^\omega} \rangle\rangle Gp \approx \text{EG}p$, $\langle\langle \emptyset^{\vec{\omega}} \rangle\rangle p_1 U p_2 \approx \text{A}(p_1 U p_2)$ and $\langle\langle \emptyset^{\vec{\omega}} \rangle\rangle Gp \approx \text{AG}p$, etc.

Ability to safely extend any finite computation. The presence of the idle action allows a (partially defined) strategy to be extended to an infinite strategy as soon as a formula is satisfied along the computations. For instance, $\mathfrak{M}, s \models \langle\langle A^{\vec{b}} \rangle\rangle X\phi$ is equivalent to the existence of $\mathfrak{f} \in D_A(s)$ such that for all $g \sqsupseteq \mathfrak{f}$, we have $\mathfrak{M}, s' \models \phi$ with $\delta(s, g) = s'$ and $\vec{b} + \text{cost}_A(s, \mathfrak{f}) \geq \vec{0}$.

Upward closure. Observe also that a strategy modality $\langle\langle A^{\vec{b}} \rangle\rangle$ reduces the impact of the function cost in two ways. If the i th component of \vec{b} is equal to ω , then there are no constraints on the i th resource along the computation. Moreover, the restriction of cost to proponent agents in A means that the actions of the opponents cost nothing and are always available. In addition, it is worth noting that $\langle\langle A^{\vec{b}} \rangle\rangle \phi \cup \psi \Rightarrow \langle\langle A^{\vec{b}' } \rangle\rangle \phi \cup \psi$ is valid whenever $\vec{b} \leq \vec{b}'$, and therefore whenever $\mathfrak{M}, s \models \langle\langle A^{\vec{b}} \rangle\rangle \phi \cup \psi$ there is a finite set of minimal elements $\vec{m} \in (\mathbb{N} \cup \{\omega\})^r$ (with respect to \leq) such that $\mathfrak{M}, s \models \langle\langle A^{\vec{m}} \rangle\rangle \phi \cup \psi$ (by Dickson's Lemma [21] every upward closed set of $(\mathbb{N} \cup \{\omega\})^r$ admits a finite basis of minimal elements; see also the notion of *Pareto frontier* in Section 5).

Alternative semantics. In the definition of the satisfaction relation \models for RB \pm ATL, in the clauses for a strategy modality followed by a temporal formula, there is an existential quantification over a \vec{b} -strategy F_A followed by a universal quantification over all the computations in $\text{Comp}(s, F_A)$. By definition, $\text{Comp}(s, F_A) = \text{Comp}(s, F_A, \vec{b})$, and therefore all the computations involved in the universal quantification are maximal and infinite, and, of course, all the underlying resource availabilities along the computations are non-negative. Alternative definitions have been considered in the literature that separate maximality from infinity, leading sometimes to different decidability results. For instance, with the *infinite semantics*, the existential quantification is over a strategy (that is not necessarily a \vec{b} -strategy), and the universal quantification is made only over *infinite* computations that respect the strategy (typically nothing is required on maximal and finite computations). Under certain assumptions, this may lead to undecidability, see e.g. [4, Section 6]. Similarly, with the *finite semantics*, the existential quantification is over a strategy (that is not necessarily a \vec{b} -strategy), and the universal quantification is made only over *maximal* (either finite or infinite) computations that respect the strategy. In this paper, we shall not investigate logics with these alternative semantics, as from our technical developments of alternating VASS, we can easily derive new decision problems on alternating VASS that correspond to such logical variants. Another way to define alternative semantics is to change the notion of resource-bounded concurrent game structures; for example, by assuming that there is no distinguished idle action, or requiring that the action manager function is of the form $\text{act} : S \times \text{Agt} \rightarrow \mathcal{P}(\text{Act})$, i.e. an agent may be unable to choose an action from a given state (because the action manager returns an empty set of actions in that state). In what follows, we shall investigate these variants by simply adapting the techniques for RB \pm ATL with the standard semantics defined above.

The *model checking problem for RB \pm ATL* is defined as follows:

Input: $k, r \geq 1$ (in unary), a formula ϕ in RB \pm ATL($[1, k], r$), a finite RB \pm ATL($[1, k], r$) model \mathfrak{M} and a state s ,

Question: $\mathfrak{M}, s \models \phi$?

The encoding of the values in k and r in unary is not essential here, since, if transitions are represented explicitly, the size of \mathfrak{M} is greater than $k + r$.

Proposition 1. [3, Theorem 1] *The model checking problem for RB \pm ATL is decidable.*

A key contribution of this paper is characterising the computational complexity of the model checking problem for $\text{RB}\pm\text{ATL}$. Obviously, $\text{RB}\pm\text{ATL}$ is a quantitative extension of ATL , and whereas the satisfaction of ATL formulae can be restricted to positional strategies (i.e., actions are chosen based on the current state rather than histories), the satisfaction of $\text{RB}\pm\text{ATL}$ formulae may require non-positional strategies in order to keep the amount of each resource above zero.

3 Problems on Vector Addition Systems with States (VASS)

In this section, we recall known complexity/decidability results for model checking and games on VASS, and state necessary complexity characterisations that will be used in the sequel. We then show that using optimal decision procedures for problems such as black boxes leads to optimal decision procedures for model checking problems of resource-bounded logics.

3.1 Alternating VASS

A binary tree \mathfrak{T} , which may contain nodes with (only) one child, is a non-empty subset of $\{1, 2\}^*$ such that, for all $n \in \{1, 2\}^*$ and $i \in \{1, 2\}$, $n \cdot i \in \mathfrak{T}$ implies $n \in \mathfrak{T}$ and, $n \cdot 2 \in \mathfrak{T}$ implies $n \cdot 1 \in \mathfrak{T}$. The nodes of \mathfrak{T} are its *elements*. The root of \mathfrak{T} is ε , the empty word. All notions such as parent, first child, second child, subtree and leaf, have their standard meanings. The height of \mathfrak{T} is the length, i.e. the number of nodes, of the longest simple path from the root to a leaf. An *alternating VASS* (AVASS) [16] is a tuple $\mathcal{A} = (Q, r, R_1, R_2)$ such that:

- Q is a non-empty finite set of *locations* (a.k.a. *control states*) and $r \geq 0$ is the number of resource values;
- R_1 is a finite subset of $Q \times \mathbb{Z}^r \times Q$ (*unary rules*);
- R_2 is a (finite) subset of Q^3 (*fork rules*).

A *derivation skeleton* of \mathcal{A} is a labelling $\mathcal{D} : \mathfrak{T} \rightarrow (R_1 \cup R_2 \cup \{\perp\})$ such that:

- \mathfrak{T} is a binary tree;
- if n has one child in \mathfrak{T} , then $\mathcal{D}(n) \in R_1$;
- if n has two children in \mathfrak{T} , then $\mathcal{D}(n) \in R_2$;
- if n is a leaf in \mathfrak{T} , then $\mathcal{D}(n) = \perp$;
- if $\mathcal{D}(n) = (q, \vec{u}, q')$ and $\mathcal{D}(n \cdot 1) \in R_1 \cup R_2$, then the first location of $\mathcal{D}(n \cdot 1)$ is q' ;
- if $\mathcal{D}(n) = (q, q_1, q_2)$ and $\mathcal{D}(n \cdot i) \in R_1 \cup R_2$ for some $i \in \{1, 2\}$, then the first location of $\mathcal{D}(n \cdot i)$ is q_i .

A *derivation* of \mathcal{A} based on \mathcal{D} is a labelling $\hat{\mathcal{D}} : \mathfrak{T} \rightarrow Q \times \mathbb{Z}^r$ such that:

- if n has one child n' in \mathfrak{T} , $\mathcal{D}(n) = (q, \vec{u}, q')$ and $\hat{\mathcal{D}}(n) = (q, \vec{v})$, then $\hat{\mathcal{D}}(n') = (q', \vec{u} + \vec{v})$;
- if n has two children n' and n'' in \mathfrak{T} , $\mathcal{D}(n) = (q, q_1, q_2)$ and $\hat{\mathcal{D}}(n) = (q, \vec{v})$, then $\hat{\mathcal{D}}(n') = (q_1, \vec{v})$ and $\hat{\mathcal{D}}(n'') = (q_2, \vec{v})$.

Note that fork rules do not update the resources, and therefore there is an asymmetry between unary rules and fork rules. This will be a very useful feature later, when dealing with the propo- nent restriction condition in RB±ATL. Unlike branching VASS (see e.g., [45, 20]), the fork rules have no effect on the counter values.

A derivation $\hat{\mathcal{D}}$ based on \mathcal{D} is *admissible* whenever $\hat{\mathcal{D}} : \mathfrak{T} \rightarrow Q \times \mathbb{N}^r$, i.e., only natural numbers occur in it. An admissible derivation is also called a *proof*. Above, we introduced the primitive notion of computations and their restriction to \vec{b} -consistent computations. Similarly, the primitive notion of derivations can be restricted to proofs (a kind of “ $\vec{0}$ -consistency”).

As an illustration, we present a proof from an alternating VASS having at least the unary rules $r_1 = q_1 \xrightarrow{(-1,+3)} q_0$ and $r_3 = q_2 \xrightarrow{(+3,+3)} q_3$, and the fork rule $r_2 = q_0 \rightarrow q_1, q_2$.

$$\frac{\frac{\frac{\vdots}{(q_3, (4, 8))} \quad (r_3) \quad \frac{\frac{\vdots}{(q_0, (0, 8))} \quad (r_1)}{(q_1, (1, 5))} \quad (r_2)}{(q_0, (1, 5))} \quad (r_1)}{(q_1, (2, 2))} \quad (r_1)$$

Before presenting the decision problems on AVASS, we state a simple property that will be used in the sequel.

Lemma 1. *Given a derivation skeleton $\mathcal{D} : \mathfrak{T} \rightarrow (R_1 \cup R_2 \cup \{\perp\})$ such that $\mathcal{D}(\varepsilon)$ is a rule whose first location is q and $(q, \vec{b}) \in Q \times \mathbb{Z}^r$, there is a unique derivation $\hat{\mathcal{D}}$ of \mathcal{A} based on \mathcal{D} such that $\hat{\mathcal{D}}(\varepsilon) = (q, \vec{b})$.*

Indeed, once the rules are provided by \mathcal{D} , the root value (q, \vec{b}) determines all the values of the derivation since the way $\hat{\mathcal{D}}(n)$ is defined remains essentially deterministic. The *state reachability problem* for AVASS is defined as follows:

Input: An alternating VASS \mathcal{A} and control states q_0 and q_f .

Question: Is there a finite proof of AVASS whose root is equal to $(q_0, \vec{0})$ and each leaf belongs to $\{q_f\} \times \mathbb{N}^r$?

When \mathcal{A} has no fork rules, \mathcal{A} is essentially a VASS [33] and the above problem is an instance of the coverability problem known to be EXPSPACE-complete [37, 39] (see also [8, 18]). The *non-termination problem* for AVASS is defined as follows:

Input: An alternating VASS \mathcal{A} and a control state q_0 .

Question: Is there a proof whose root is equal to $(q_0, \vec{0})$ and all the maximal branches are infinite?

Proposition 2. [16, 32] *The state reachability and non-termination problems for AVASS are 2EXPTIME-complete.*

The decidability of these problems was first established in [40] by using monotonicity of the games. The 2EXPTIME upper bound is preserved if we assume that the root is labelled by (q_0, \vec{b}) with $\vec{b} \in \mathbb{N}^r$ encoded with a binary representation (see Lemma 7 below).

In the sequel, we shall also admit fork rules of any arity $\alpha \geq 1$ and therefore in such slightly extended AVASS, the set of fork rules R_2 is a finite subset of $\bigcup_{\beta \geq 2} Q^\beta$. The notions of derivation skeleton, derivation and proof are also changed to refer to general trees $\mathfrak{T} \subseteq (\mathbb{N} \setminus \{0\})^*$. The set of finite words $\mathfrak{T} \subseteq (\mathbb{N} \setminus \{0\})^*$ is a (not necessarily binary) *tree* iff for all $n \in (\mathbb{N} \setminus \{0\})^*$ and $i \in (\mathbb{N} \setminus \{0\})$, $n \cdot i \in \mathfrak{T}$ implies $n \in \mathfrak{T}$, and $n \cdot i \in \mathfrak{T}$ and $i > 1$ imply $n \cdot (i-1) \in \mathfrak{T}$. In the remainder of the paper, by AVASS we mean such an extended AVASS with fork rules of arbitrary arity.

3.2 Model checking problems

A VASS can be defined as an alternating VASS without any fork rules, and therefore we write it $\mathcal{V} = (Q, r, R)$ where R is a finite set of unary rules. Given a VASS \mathcal{V} , its transition system $\mathfrak{T}\mathfrak{S}(\mathcal{V}) \stackrel{\text{def}}{=} (\mathfrak{B}, \rightarrow, L)$ is such that:

- $\mathfrak{B} \stackrel{\text{def}}{=} Q \times \mathbb{N}^r$;
- L is a labelling with elements of Q also understood as propositional variables and $L(q) \stackrel{\text{def}}{=} \{q\} \times \mathbb{N}^r$;
- \rightarrow is a binary relation on \mathfrak{B} such that $(q, \vec{v}) \rightarrow (q', \vec{v}')$ iff there is a unary rule (q, \vec{u}, q') in R such that $\vec{v}' = \vec{u} + \vec{v}$ where ‘+’ is the component-wise addition on \mathbb{N}^r .

As usual, we also write \rightarrow^* to denote the reflexive and transitive closure of \rightarrow . Since $\mathfrak{T}\mathfrak{S}(\mathcal{V})$ is a Kripke-style structure, it can be used to interpret modal or temporal formulae where the atomic formulae refer to locations, e.g., formulae of the temporal logics LTL or CTL (see also [17]). Since alternating-time temporal logics such as ATL or ATL* strictly extend CTL or CTL* respectively, complexity hardness results for temporal logics can be lifted to alternating-time logics. We adopt this approach.

We first recall some results that will be useful in the sequel.

Proposition 3. *The model checking problem for LTL on VASS is EXPSPACE-complete (the atomic formulae are control states) and it is PSPACE-complete for a fixed number of resources [27].*

EXPSPACE-hardness of model checking on VASS already follows from EXPSPACE-hardness of the state reachability problem for VASS [37], as state reachability is a subproblem of the model checking problem for VASS (consider the LTL formula Fq_f).

3.3 The logic RBTL* and its variants

The models of the logic RBTL* are structures of the form (Q, r, R, L) where (Q, r, R) is a VASS and L is a labelling built on elements of Q understood as propositional variables, so that $L(q) = \{q\}$ (see e.g., [12, Section 3]). For consistency with standard terminology, an infinite proof in (Q, r, R) is called a *path* or *run*, and is represented by $\lambda = (q_0, \vec{v}_0) \rightarrow (q_1, \vec{v}_1) \dots$. We write $\lambda(i)$ to denote the i th configuration (q_i, \vec{v}_i) , and $\lambda[+i, +\infty)$ to denote the suffix of λ starting from (q_i, \vec{v}_i) .

The *state formulae* ϕ and the *path formulae* Φ of RBTL* are defined mutually recursively by the following grammar (relative to a set of locations Q and number of resources r , which is not a significant restriction since we are only interested in model checking)

$$\begin{aligned} \phi &::= q \mid \neg\phi \mid (\phi \wedge \phi) \mid \langle \vec{b} \rangle \Phi \\ \Phi &::= \phi \mid \neg\Phi \mid (\Phi \wedge \Phi) \mid \mathbf{X}\Phi \mid (\Phi\mathbf{U}\Phi) \mid \mathbf{G}\Phi, \end{aligned}$$

where $q \in Q$. Syntactically, every state formula is also a path formula according to this grammar, reflecting the fact that a path uniquely identifies a location in which a formula is interpreted (its starting location).

In presenting the semantics of RBTL*, we make an explicit distinction between state formulae and path formulae. The two satisfaction relations \models_s and \models_p are defined as follows (standard clauses for the Boolean connectives are omitted):

$$\begin{array}{ll}
\mathfrak{M}, q \models_s q' & \stackrel{\text{def}}{\Leftrightarrow} q' = q \\
\mathfrak{M}, q \models_s \langle \vec{b} \rangle \Phi & \stackrel{\text{def}}{\Leftrightarrow} \text{there is an infinite run } \lambda \text{ starting at } (q, \vec{b}) \\
& \text{such that } \mathfrak{M}, \lambda \models_p \Phi \\
\mathfrak{M}, \lambda \models_p \phi & \stackrel{\text{def}}{\Leftrightarrow} \mathfrak{M}, \lambda(0) \models_s \phi \text{ for state formulae } \phi \\
\mathfrak{M}, \lambda \models_p X\Phi & \stackrel{\text{def}}{\Leftrightarrow} \mathfrak{M}, \lambda[1, +\infty) \models_p \Phi \\
\mathfrak{M}, \lambda \models_p \Phi U \Psi & \stackrel{\text{def}}{\Leftrightarrow} \text{there is } i \geq 0 \text{ such that } \mathfrak{M}, \lambda[i, +\infty) \models_p \Psi \text{ and} \\
& \text{for every } j \in [0, i-1], \text{ we have } \mathfrak{M}, \lambda[j, +\infty) \models_p \Phi.
\end{array}$$

As usual, we write $[\vec{b}]\phi$ to denote the formula $\neg\langle \vec{b} \rangle\neg\phi$, and therefore $\mathfrak{M}, q \models_s [\vec{b}]\Phi$ iff for all the infinite runs λ starting at (q, \vec{b}) , we have $\mathfrak{M}, \lambda \models_p \Phi$.

The model checking problem for RBTL* is defined as follows:

Input: a model $\mathfrak{M} = (Q, r, R, L)$, a control state q and a state formula ϕ .

Question: $\mathfrak{M}, q \models_s \phi$?

As CTL is a syntactic fragment of CTL*, RBTL is defined as the syntactic fragment of RBTL* in which any subformula with an outermost connective in $\{U, X, G\}$ is preceded by a modality of the form either $\langle \vec{b} \rangle$ or $[\vec{b}]$. Observe that the model checking problem for RBTL is already EXPSpace-hard, since the state reachability problem for VASS can be reduced to a question of the form $\mathfrak{M}, q_0 \models \langle \vec{0} \rangle q_f$. We consider the computational complexity of the model checking problems for RBTL* and RBTL in Section 4.4.

4 On the Complexity of RB \pm ATL

In this section, we show how to solve the model checking problem for RB \pm ATL by solving instances of decision problems for alternating VASS using a labelling algorithm. The size of the AVASS problem instance is linear in the input resource-bounded concurrent game structure, and the number of calls is also linear in the size of the input formulae. As far as worst-case complexity bounds are concerned, this is probably the best we can hope for. At a high level, our results relate model checking problems for resource-bounded logics in AI and verification games. At a technical level, this reduces to three key correspondences. First, the proponent restriction condition in RB \pm ATL corresponds to the fact that, in AVASS, only unary rules can update the counter values. This is crucial to our results, but alone is not sufficient. Second, each \vec{b} -strategy F_A generates a set of computations that can be represented as a finitely branching tree with infinite branches, which corresponds precisely to the proofs in AVASS (see e.g., Theorem 1 below). Third, roughly speaking, temporal formulae in the scope of a strategy modality correspond to acceptance conditions on branches of the proofs (admissible computations) extracted from the AVASS. In the remainder of this section, we develop these correspondences in detail and these can be summarised in the table below.

RB±ATL	Alternating VASS
Logic in AI	Verification games
<i>proponent restriction condition</i>	updates in R_1 / no update in R_2
computation tree for F_A	proof
formulae in the scope of $\langle\langle A^{\vec{b}} \rangle\rangle$	monotone objectives

4.1 Structural analysis of strategies and proofs

We first establish the necessary formal relationships between strategies in resource-bounded concurrent game structures, and proofs in alternating VASS. The technical developments are not conceptually difficult, but they allow us to derive results that are helpful in solving the model checking problem for RB±ATL using decision procedures on AVASS. Our approach also allows us to pose and solve new model checking problems (see e.g. Section 5 and Section 5.6).

Let \mathfrak{M} be a finite resource-bounded concurrent game structure, $A \subseteq \text{Agt}$ be a coalition, F_A be a strategy and s be a state. We construct an alternating VASS $\mathcal{A}_{\mathfrak{M},A,s}$ such that the set of computations starting in s and respecting F_A corresponds precisely to a derivation skeleton whose root is labelled by a unary rule with first state s . Moreover, if F_A is a \vec{b} -strategy w.r.t. s , then the derivation skeleton can be turned into a proof whose root is labelled by (s, \vec{b}) . This implies that fork rules can have any arity greater than one, and components can have the value ω , where ω is a value that remains constant (i.e., we assume that $n + \omega = n$ for all $n \in \mathbb{Z}$).

Given $\mathfrak{M} = (\text{Agt}, S, \text{Act}, r, \text{act}, \text{cost}, \delta, L)$, the AVASS $\mathcal{A}_{\mathfrak{M},A,s} \stackrel{\text{def}}{=} (Q, r, R_1, R_2)$ is built as follows:

$$Q \stackrel{\text{def}}{=} \{s\} \cup \{(s', \vec{f}) \mid s' \in S, \vec{f} \in D_A(s')\} \cup \{(g, s') \mid s', s'' \in S, g \in D(s''), \delta(s'', g) = s'\}.$$

- The set of unary rules R_1 contains the following elements:
 - For all $\vec{f} \in D_A(s)$, $(s, \text{cost}_A(s, \vec{f}), (s, \vec{f}))$.
 - For all $(g, s') \in Q$, for all $\vec{f} \in D_A(s')$, $((g, s'), \text{cost}_A(s', \vec{f}), (s', \vec{f}))$.
- The set of fork rules R_2 contains the following elements.
 - For all $(s', \vec{f}) \in Q$, let $\{(g_1, s_1), \dots, (g_\alpha, s_\alpha)\} = \{(g, s'') \in S \mid s'' \in \delta(s', g), g \in D(s'), \vec{f} \sqsubseteq g\}$. This set is non-empty because an action manager always returns a non-empty set of actions. We add the α -ary fork rule

$$((s', \vec{f}), (g_1, s_1), \dots, (g_\alpha, s_\alpha)).$$

In order to define the rule unambiguously, we assume an arbitrary linear ordering on the set Q .

It is worth noting that:

- s has a special status in Q simply because any proof whose root configuration contains s has no predecessor configuration.
- By construction, any derivation skeleton from $\mathcal{A}_{\mathfrak{M},A,s}$ has to alternate the rules in R_1 and the rules in R_2 . This property will be used to slightly simplify developments below.

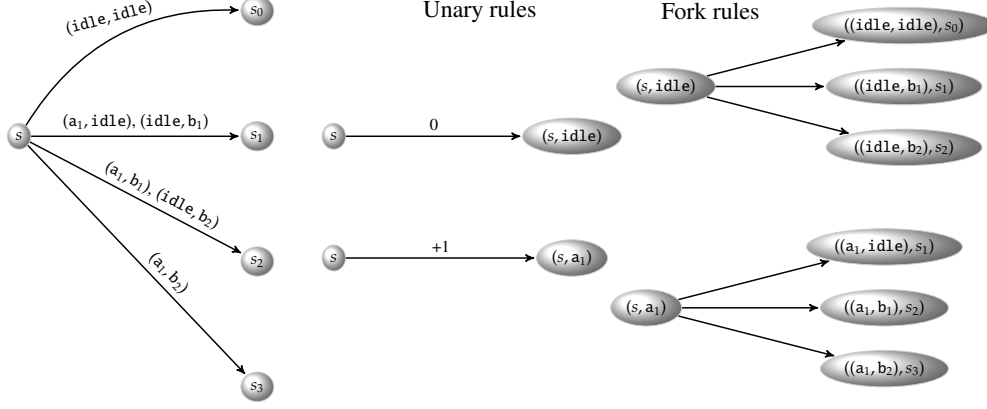


Figure 2: Transitions and its associated unary and fork rules.

- For every (s', \mathfrak{f}) in Q , there is a unique fork rule starting from (s', \mathfrak{f}) .
- The construction also applies in degenerate cases, i.e., when $A = \text{Agt}$ or when $A = \emptyset$ (assuming that $\text{cost}(s', \mathfrak{f}) = \vec{0}$ for the unique $\mathfrak{f} \in D_\emptyset(s')$).

In Figure 2, we illustrate how transitions from the state s are turned into unary rules and fork rules (in the example, $\text{Agt} = \{1, 2\}$, $A = \{1\}$, $\text{act}(s, 1) = \{\text{idle}, a_1\}$, $\text{act}(s, 2) = \{\text{idle}, b_1, b_2\}$, and $\text{cost}(s, 1, a_1) = +1$).

Given an infinite computation $\lambda = s_0 \xrightarrow{g_1} s_1 \xrightarrow{g_2} s_2 \dots$ starting in $s = s_0$ and respecting F_A , we can associate it with an infinite sequence (which we call an *extended computation*)

$$\text{ext}(\lambda, F_A) \stackrel{\text{def}}{=} s_0 \xrightarrow{\vec{u}_0} (s_0, \mathfrak{f}_0) \rightarrow (g_1, s_1) \xrightarrow{\vec{u}_1} (s_1, \mathfrak{f}_1) \rightarrow (g_2, s_2) \xrightarrow{\vec{u}_2} (s_2, \mathfrak{f}_2) \rightarrow (g_3, s_3) \dots$$

where $s_0 = s$, and for all $n \geq 0$, $F_A(s_0 \xrightarrow{g_1} s_1 \dots \xrightarrow{g_n} s_n) = \mathfrak{f}_n$ and $\text{cost}_A(s_n, \mathfrak{f}_n) = \vec{u}_n$. That is, every step $s_i \xrightarrow{g_{i+1}} s_{i+1}$ in the computation λ is decomposed into two parts: $s_i \xrightarrow{\vec{u}_i} (s_i, \mathfrak{f}_i) \rightarrow (g_{i+1}, s_{i+1})$.

The computations in $\text{Comp}(s, F_A)$ can be organised as an infinite tree that corresponds to a derivation skeleton for $\mathcal{A}_{\text{M}, A, s}$. Below we define an infinite tree \mathfrak{T}_{F_A} , a labelling function $\mathfrak{Q} : \mathfrak{T}_{F_A} \rightarrow S$ and a partial map $\mathfrak{R} : \mathfrak{T}_{F_A} \times \mathfrak{T}_{F_A} \rightarrow (\bigcup_{s' \in S} D(s'))$.

- $\mathfrak{T}_{F_A}(\varepsilon) \stackrel{\text{def}}{=} s_0$.
- For all finite words $w = k_1 \dots k_\beta$ in \mathfrak{T}_{F_A} such that $\mathfrak{Q}(w)$ is already defined, we add to \mathfrak{T}_{F_A} the values $k_1 \dots k_\beta \cdot 1, \dots, k_1 \dots k_\beta \cdot \alpha$ such that

- $F_A(\mathfrak{Q}(\varepsilon) \xrightarrow{\mathfrak{R}(\varepsilon, k_1)} \mathfrak{Q}(k_1) \xrightarrow{\mathfrak{R}(k_1, k_1 k_2)} \mathfrak{Q}(k_1 k_2) \dots \xrightarrow{\mathfrak{R}(k_1 \dots k_{\beta-1}, k_1 \dots k_\beta)} \mathfrak{Q}(w)) = \mathfrak{f}$
- $\{(g_1, s_1), \dots, (g_\alpha, s_\alpha)\} = \{(g, s'') \in S \mid s'' \in \delta(s', g), g \in D(s'), \mathfrak{f} \sqsubseteq g\}$ with $s' = \mathfrak{Q}(w)$.
- For all $j \in [1, \alpha]$, we have $\mathfrak{Q}(k_1 \dots k_\beta \cdot j) \stackrel{\text{def}}{=} s_j$ and $\mathfrak{R}(w, w \cdot j) \stackrel{\text{def}}{=} g_j$.

The tree \mathfrak{T}_{F_A} is defined by saturation of the above rules, and the maps \mathfrak{Q} and \mathfrak{R} are defined accordingly. The structure $(\mathfrak{T}_{F_A}, \mathfrak{R}, \mathfrak{Q})$ is a labelled transition system with a tree-like structure encoding all the infinite computations respecting the strategy F_A . A maximal branch w of $(\mathfrak{T}_{F_A}, \mathfrak{R}, \mathfrak{Q})$ is understood as an element of $(\mathbb{N} \setminus \{0\})^\omega$, such that any (strict) finite prefix of w belongs to \mathfrak{T}_{F_A} . The label of w , written $\mathbf{lab}(w)$, is defined as follows:

$$\mathbf{lab}(w) \stackrel{\text{def}}{=} \mathfrak{Q}(\varepsilon) \xrightarrow{\mathfrak{R}(\varepsilon, k_1)} \mathfrak{Q}(k_1) \xrightarrow{\mathfrak{R}(k_1, k_1 k_2)} \mathfrak{Q}(k_1 k_2) \cdots \xrightarrow{\mathfrak{R}(k_1 \cdots k_{\beta-1}, k_1 \cdots k_\beta)} \mathfrak{Q}(k_1 \cdots k_\beta) \cdots$$

where $w = k_1 k_2 k_3 \cdots$. By construction, $\mathbf{lab}(w)$ is a maximal computation.

Lemma 2.

- (I) For every maximal computation λ starting at s and respecting F_A , there is a maximal branch w in $(\mathfrak{T}_{F_A}, \mathfrak{R}, \mathfrak{Q})$ such that $\lambda = \mathbf{lab}(w)$.
- (II) For every maximal branch w in $(\mathfrak{T}_{F_A}, \mathfrak{R}, \mathfrak{Q})$, there is a maximal computation λ starting at s and respecting F_A such that $\mathbf{lab}(w) = \lambda$.

The proof reflects that $(\mathfrak{T}_{F_A}, \mathfrak{R}, \mathfrak{Q})$ contains all the computations from s that respect the strategy F_A .

We build a derivation skeleton $\mathcal{D} : \mathfrak{T}_{F_A} \rightarrow (R_1 \cup R_2)$ as follows, where all the maximal branches of \mathfrak{T}_{F_A} are infinite (we therefore do not need to include \perp in the range of \mathcal{D}).

- $\mathcal{D}(\varepsilon) = (s_0, \text{cost}_A(s_0, F_A(s_0)), (s_0, F_A(s_0)))$.
- $\mathcal{D}(1) = ((s_0, F_A(s_0)), (g_1, s_1), \dots, (g_\alpha, s_\alpha))$ where $1, \dots, \alpha \in \mathfrak{T}_{F_A}$ (but $\alpha + 1 \notin \mathfrak{T}_{F_A}$), for all $j \in [1, \alpha]$, $\mathfrak{T}_{F_A}(j) = s_j$ and $\mathfrak{R}(\varepsilon, j) = g_j$. By construction of $\mathcal{A}_{\mathfrak{M}, A, s}$, $\mathcal{D}(1)$ is the unique fork rule starting from $(s_0, F_A(s_0))$.
- Let $n = 1k_1 1 \cdots 1k_\beta 1$ with $k_1, \dots, k_\beta \geq 1$ and such that

$$\mathcal{D}(1k_1 1 \cdots 1k_\beta) = ((g, s'), \text{cost}_A(s', \mathfrak{f}), (s', \mathfrak{f})) \in R_1.$$

Then, $\mathcal{D}(n) = ((s', \mathfrak{f}), (g_1, s_1), \dots, (g_\alpha, s_\alpha))$ where $k_1 \cdots k_\beta 1, \dots, k_1 \cdots k_\beta \alpha \in \mathfrak{T}_{F_A}$ (but $k_1 \cdots k_\beta (\alpha + 1) \notin \mathfrak{T}_{F_A}$), for all $j \in [1, \alpha]$,

- $\mathfrak{T}_{F_A}(k_1 \cdots k_\beta \cdot j) = s_j$ and,
- $\mathfrak{R}(k_1 \cdots k_\beta, k_1 \cdots k_\beta j) = g_j$.

By construction of $\mathcal{A}_{\mathfrak{M}, A, s}$, $\mathcal{D}(n)$ is the unique fork rule starting from (s', \mathfrak{f}) .

- Let $n = 1k_1 1 \cdots 1k_\beta$. By construction we can assume that we already have that $\mathcal{D}(1k_1 1 \cdots 1k_{\beta-1} 1) = ((s', \mathfrak{f}'), (g_1, s_1), \dots, (g_\alpha, s_\alpha))$. Let g' be equal to $\mathfrak{R}(k_1 \cdots k_\beta, k_1 \cdots k_\beta \cdot 1)$ and \mathfrak{f} be the restriction of g' to A (so $\mathfrak{f} \sqsubseteq g'$). Then,

$$\mathcal{D}(n) = ((g_{k_\beta}, s_{k_\beta}), \text{cost}_A(s_{k_\beta}, \mathfrak{f}), (s_{k_\beta}, \mathfrak{f})).$$

Note that $\mathcal{D}(n)$ is indeed a valid unary rule.

Given an infinite branch w of \mathcal{D} , say $w = 1k_1 1k_2 1k_3 \cdots \in \mathbb{N}^\omega$, we define the extended computation $\text{ext}(w, F_A)$ as follows. Suppose that the label of such a branch is characterised by the values below:

- $\mathcal{D}(\varepsilon) = (s_0, \vec{u}_0, (s_0, \hat{f}_0)); \mathcal{D}(1) = ((s_0, \hat{f}_0), (g_1^1, s_1^1), \dots, (g_{\alpha_1}^1, s_{\alpha_1}^1)).$
- ...
- $\mathcal{D}(1k_11 \dots k_i) = ((g_{k_i}^i, s_{k_i}^i), \vec{u}_i, (s_{k_i}^i, \hat{f}_i)).$
- $\mathcal{D}(1k_11 \dots k_i1) = ((s_{k_i}^i, \hat{f}_i), (g_1^{i+1}, s_1^{i+1}), \dots, (g_{\alpha_{i+1}}^{i+1}, s_{\alpha_{i+1}}^{i+1})).$
- ...
- $\mathcal{D}(1k_11 \dots k_{\beta-1}) = ((g_{k_i}^{\beta-1}, s_{k_{\beta-1}}^{\beta-1}), \vec{u}_{\beta-1}, (s_{k_{\beta-1}}^{\beta-1}, \hat{f}_{\beta-1})).$
- $\mathcal{D}(1k_11 \dots k_{\beta-1}1) = ((s_{k_{\beta-1}}^{\beta-1}, \hat{f}_{\beta-1}), (g_1^\beta, s_1^\beta), \dots, (g_{\alpha_\beta}^\beta, s_{\alpha_\beta}^\beta)).$
- ...

Then,

$$\text{ext}(w, F_A) \stackrel{\text{def}}{=} s_0 \xrightarrow{\vec{u}_0} (s_0, \hat{f}_0) \rightarrow (g_{k_1}^1, s_{k_1}^1) \xrightarrow{\vec{u}_1} (s_{k_1}^1, \hat{f}_1) \rightarrow (g_{k_2}^2, s_{k_2}^2) \xrightarrow{\vec{u}_2} (s_{k_2}^2, \hat{f}_2) \rightarrow (g_{k_3}^3, s_{k_3}^3) \dots$$

Lemma 3.

- (I) For every maximal computation λ starting at s and respecting F_A , there is maximal branch w in \mathcal{D} such that $\text{ext}(\lambda, F_A) = \text{ext}(w, F_A)$.
- (II) For every maximal branch w in \mathcal{D} , there is a maximal computation λ starting at s and respecting F_A such that $\text{ext}(w, F_A) = \text{ext}(\lambda, F_A)$.

The reduction can be easily verified. Note, however, that the proponent restriction is essential for its correctness.

Theorem 1. There is a \vec{b} -strategy w.r.t. s in \mathfrak{M} iff there is a proof in $\mathcal{A}_{\mathfrak{M}, A, s}$ whose root is labelled by (s, \vec{b}) and every maximal branch is infinite.

Proof. First suppose that there is a \vec{b} -strategy F_A w.r.t. s in \mathfrak{M} . Let us consider the structures $(\mathfrak{X}_{F_A}, \mathfrak{R}, \mathfrak{V})$ and $\mathcal{D} : \mathfrak{X} \rightarrow (R_1 \cup R_2)$ as defined above. By Lemma 1, there is a unique derivation $\hat{\mathcal{D}}$ of $\mathcal{A}_{\mathfrak{M}, A, s}$ based on \mathcal{D} such that $\hat{\mathcal{D}}(\varepsilon) = (s, \vec{b})$. It remains to show that $\hat{\mathcal{D}}$ is indeed a proof. Let $w = 1k_11k_21 \dots$ be a maximal branch of $\hat{\mathcal{D}}$. We have:

- $\hat{\mathcal{D}}(\varepsilon) = (s_0, \vec{b}).$
- $\hat{\mathcal{D}}(1) = ((s_0, \hat{f}_0), \vec{u}_0 + \vec{b})$ with $\mathcal{D}(\varepsilon) = (s_0, \vec{u}_0, (s_0, \hat{f}_0)).$
- ...
- $\hat{\mathcal{D}}(1k_11 \dots k_i) = ((g_{k_i}^i, s_{k_i}^i), \sum_{j=1}^{i-1} \vec{u}_j + \vec{b})$ with

$$\mathcal{D}(1k_11 \dots k_{i-1}1) = ((s_{k_{i-1}}^{i-1}, \hat{f}_{i-1}), (g_1^i, s_1^i), \dots, (g_{\alpha_i}^i, s_{\alpha_i}^i)).$$
- $\hat{\mathcal{D}}(1k_11 \dots k_i1) = ((s_{k_i}^i, \hat{f}_i), \sum_{j=1}^i \vec{u}_j + \vec{b})$ with $\mathcal{D}(1k_11 \dots k_i) = ((g_{k_i}^i, s_{k_i}^i), \vec{u}_i, (s_{k_i}^i, \hat{f}_i)).$

• ...

By Lemma 3, there is a maximal computation λ starting at s and respecting F_A such that $\text{ext}(w, F_A) = \text{ext}(\lambda, F_A)$. Since F_A is a \vec{b} -strategy, λ is \vec{b} -consistent and therefore for all $i \geq 0$, we have $\vec{0} \leq \sum_{j=1}^{i-1} \vec{u}_j + \vec{b}$, which implies that $\hat{\mathcal{D}}$ is a proof.

For the proof of the other direction, assuming that there is a proof $\hat{\mathcal{D}}$ whose root is labelled by (s, \vec{b}) and every maximal branch is infinite, we can extract from the underlying derivation \mathcal{D} a strategy F_A (see the similar construction in the proof of Theorem 3 below). Lemma 3 and the fact that $\hat{\mathcal{D}}$ is admissible entail that F_A is a \vec{b} -strategy w.r.t. s (details are omitted). \square

Transitions in \mathfrak{M} can be defined as triples (s', g, s'') such that $\delta(s', g) = s''$. A transition is also denoted by the expression $s' \xrightarrow{g} s''$. The set of transitions of \mathfrak{M} is denoted by $\Sigma_{\mathfrak{M}}$. It is interpreted as a finite alphabet when \mathfrak{M} is finite. An infinite computation $\lambda = s_0 \xrightarrow{g_1} s_1 \xrightarrow{g_2} s_2 \dots$ can be equivalently represented by the ω -word in $\Sigma_{\mathfrak{M}}^\omega$ (with contiguous transitions)

$$(s_0 \xrightarrow{g_1} s_1) \cdot (s_1 \xrightarrow{g_2} s_2) \cdot (s_2 \xrightarrow{g_3} s_3) \dots$$

An ω -word $w \in \Sigma_{\mathfrak{M}}^\omega$ is said to be with contiguous transitions whenever at any position, the second state of the transition is equal to the first state of the next position.

Given an infinite branch of the proof corresponding to the extended computation

$$s \xrightarrow{u_0} (s, \bar{f}_0) \rightarrow (g_{k_1}^1, s_{k_1}^1) \xrightarrow{u_1} (s_{k_1}^1, \bar{f}_1) \rightarrow (g_{k_2}^2, s_{k_2}^2) \xrightarrow{u_2} (s_{k_2}^2, \bar{f}_2) \rightarrow (g_{k_3}^3, s_{k_3}^3) \dots$$

its $\Sigma_{\mathfrak{M}}$ -projection is defined as the sequence

$$(s \xrightarrow{g_{k_1}^1} s_{k_1}^1) \cdot (s_{k_1}^1 \xrightarrow{g_{k_2}^2} s_{k_2}^2) \cdot (s_{k_2}^2 \xrightarrow{g_{k_3}^3} s_{k_3}^3) \dots$$

Lemma 4. Let $L \subseteq \Sigma_{\mathfrak{M}}^\omega$ and $\vec{b} \in (\mathbb{N} \cup \{\omega\})^r$. The statements below are equivalent.

1. There is a \vec{b} -strategy F_A w.r.t. s in \mathfrak{M} such that the set of computations $\text{Comp}(s, F_A)$ is included in L .
2. There is a proof in $\mathcal{A}_{\mathfrak{M}, A, s}$ whose root is labelled by (s, \vec{b}) , every maximal branch is infinite and its $\Sigma_{\mathfrak{M}}$ -projection belongs to L .

Lemma 4 is a consequence of Theorem 1 and Lemma 3 is key to establishing formal correspondences between \mathfrak{M} and $\mathcal{A}_{\mathfrak{M}, A, s}$. The main challenge is to determine classes of languages for which decidability can be obtained by using only the decidability (and complexity characterisation) of the state reachability and non-termination problems for AVASS.

Given $S' \subseteq S$ with $s \in S'$, let $L_{S'}$ be the set of all ω -words such that the transitions use only states in S' (see the formal definition below). Let us define $\mathcal{A}_{\mathfrak{M}, A, s}^{S'}$ as a restriction of $\mathcal{A}_{\mathfrak{M}, A, s}$ in which the opponent has no way to go out of S' . Alternatively, $\mathcal{A}_{\mathfrak{M}, A, s}^{S'}$ can be understood as the restriction of $\mathcal{A}_{\mathfrak{M}, A, s}$ to rules that only involve states in S' (assuming that s is already in S'). We define $\mathcal{A}_{\mathfrak{M}, A, s}^{S'} \stackrel{\text{def}}{=} (Q, r, R_1, R_2)$ as follows:

$$Q \stackrel{\text{def}}{=} \{s\} \cup \{(s', \bar{f}) \mid s' \in S', \bar{f} \in D_A(s')\} \cup \{(g, s') \mid s', s'' \in S', g \in D(s''), \delta(s'', g) = s'\}.$$

- The set of unary rules R_1 contains the following elements.
 - For all $\bar{f} \in D_A(s)$, $(s, \text{cost}_A(s, \bar{f}), (s, \bar{f}))$.
 - For all $(g, s') \in Q$, for all $\bar{f} \in D_A(s')$, $((g, s'), \text{cost}_A(s', \bar{f}), (s', \bar{f}))$.
- The set of fork rules R_2 contains the following elements.
 - For all $(s', \bar{f}) \in Q$, let $\{(g_1, s_1), \dots, (g_\alpha, s_\alpha)\} = \{(g, s'') \mid s'' \in \delta(s', g), g \in D(s'), \bar{f} \sqsubseteq g\}$.
 - If $\{s_1, \dots, s_\alpha\} \subseteq S'$, then we add the α -ary fork rule

$$((s', \bar{f}), (g_1, s_1), \dots, (g_\alpha, s_\alpha)).$$

(Otherwise, nothing is added.)

So, there is at most one fork rule starting from (s', \bar{f}) (possibly zero).

Lemma 5. *Assuming that $s \in S'$, the statements below are equivalent.*

1. *There is a \vec{b} -strategy F_A w.r.t. s in \mathfrak{M} such that the set of computations $\text{Comp}(s, F_A)$ only visit states in S' .*
2. *There is a proof in $\mathcal{A}_{\mathfrak{M}, A, s}^{S'}$ whose root is labelled by (s, \vec{b}) and every maximal branch is infinite (a positive instance of the non-termination problem for AVASS).*

Note that the way \vec{b} -strategies are defined, in (1), F_A generates maximal and infinite computations in which only states in S' are visited. Similarly, the proof in (2) contains only maximal and infinite branches and its root is precisely (s, \vec{b}) .

Proof. Let $\mathcal{A}_{\mathfrak{M}, A, s} = (Q, r, R_1, R_2)$ and $\mathcal{A}_{\mathfrak{M}, A, s}^{S'} = (Q', r, R'_1, R'_2)$. By construction, we have $Q' \subseteq Q$, $R'_1 \subseteq R_1$ and $R'_2 \subseteq R_2$. We write $\Sigma'_{\mathfrak{M}}$ to denote the alphabet $\{s_1 \xrightarrow{a} s_2 : s_1, s_2 \in S' \text{ and } \delta(s_1, g) = s_2\}$ and L to denote the ω -regular language in $(\Sigma'_{\mathfrak{M}})^\omega$ made of infinite sequences of contiguous transitions such that only states in S' can occur.

(1) \rightarrow (2). Suppose there is a \vec{b} -strategy F_A w.r.t. s in \mathfrak{M} such that the set of computations $\text{Comp}(s, F_A)$ only visit states in S' , which amounts to having the set of computations included in L . By Lemma 4, there is a proof in $\mathcal{A}_{\mathfrak{M}, A, s}$ whose root is labelled by (s, \vec{b}) , every maximal branch is infinite and it belongs to L . Since $s \in S'$, all the rules in $R'_1 \cup R'_2$ only involve states in S' and the above proof only visits such states, we have that there is a proof in $\mathcal{A}_{\mathfrak{M}, A, s}^{S'}$ whose root is labelled by (s, \vec{b}) and every maximal branch is infinite.

(2) \rightarrow (1). Suppose that there is a proof in $\mathcal{A}_{\mathfrak{M}, A, s}^{S'}$ whose root is labelled by (s, \vec{b}) and every maximal branch is infinite. Since $\mathcal{A}_{\mathfrak{M}, A, s}^{S'}$ is defined as a restriction of $\mathcal{A}_{\mathfrak{M}, A, s}$ with $Q' \subseteq Q$, $R'_1 \subseteq R_1$ and $R'_2 \subseteq R_2$, there is also a proof in $\mathcal{A}_{\mathfrak{M}, A, s}$ whose root is labelled by (s, \vec{b}) , every maximal branch is infinite and only states in S' are visited. Equivalently, there is a proof in $\mathcal{A}_{\mathfrak{M}, A, s}$ whose root is labelled by (s, \vec{b}) , every maximal branch is infinite and it belongs to L . By Lemma 4, there is a \vec{b} -strategy w.r.t. s in \mathfrak{M} such that the set of computations $\text{Comp}(s, F_A)$ is included in L , hence only states in S' are visited. \square

Lemma 5 is useful to handle formulae of the form $\langle\langle A^{\vec{b}} \rangle\rangle \mathbf{G}\phi$. Let us consider a similar treatment that will be useful to handle formulae of the form $\langle\langle A^{\vec{b}} \rangle\rangle \phi_1 \mathbf{U} \phi_2$. Given $S_1, S_2 \subseteq S$ with

$s \in S_1 \cup S_2$, let L_{S_1, S_2} be the set of all ω -words with contiguous transitions such that the projection under S belongs to $S_1^* \cdot S_2 \cdot S^\omega$. Typically, the projection of $(s_0 \xrightarrow{a_1} s_1) \cdot (s_1 \xrightarrow{a_2} s_2) \cdot (s_2 \xrightarrow{a_3} s_3) \cdots$ under S is precisely $s_0 s_1 s_2 s_3 \cdots$.

Lemma 6. *The statements below are equivalent:*

1. *There is a \vec{b} -strategy w.r.t. s in \mathfrak{M} such that $\text{Comp}(s, F_A) \subseteq L_{S_1, S_2}$.*
2. *There is a finite proof in $\mathcal{A}_{\mathfrak{M}, A, s}^{S_1 \cup S_2}$ whose root is labelled by (s, \vec{b}) and each leaf contains a control state in $\{(g, s') \in Q \mid s' \in S_2\} \cup \{s' \in Q \mid s' = s, s \in S_2\}$ (a positive instance of the state reachability problem for AVASS).*

The set $\{s' \in Q \mid s' = s, s \in S_2\}$ may appear somewhat unintuitive: basically it corresponds to $\{s\}$ if $s \in S_2$ and \emptyset otherwise. The proof of Lemma 6 below relies on the fact that, in resource-bounded concurrent game structures, $\text{idle} \in \text{act}(s, a)$ for all agents a and states s .

Proof. The proof is similar to the proof of Lemma 5 except that we need to relate finite proofs to infinite ones, and in doing so we take advantage of the presence of the idle action in concurrent game structures.

Let $\mathcal{A}_{\mathfrak{M}, A, s} = (Q, r, R_1, R_2)$ and $\mathcal{A}_{\mathfrak{M}, A, s}^{S_1 \cup S_2} = (Q', r, R'_1, R'_2)$. By construction, we have $Q' \subseteq Q$, $R'_1 \subseteq R_1$ and $R'_2 \subseteq R_2$.

(1) \rightarrow (2). Suppose there is a \vec{b} -strategy F_A w.r.t. s in \mathfrak{M} such that the computations in $\text{Comp}(s, F_A)$ visit a state in S_1 until a state in S_2 is visited, which amounts to having the set of computations included in L_{S_1, S_2} . By Lemma 4, there is a proof $\hat{\mathcal{D}}$ in $\mathcal{A}_{\mathfrak{M}, A, s}$ whose root is labelled by (s, \vec{b}) , every maximal branch is infinite and it belongs to L_{S_1, S_2} . Let $\hat{\mathcal{D}}'$ be the finite proof obtained from $\hat{\mathcal{D}}$ by pruning any subtree as soon as a node is labelled by a control state in S_2 . Existence of such a finite proof is guaranteed by König's Lemma. It is easy to check that $\hat{\mathcal{D}}'$ is a finite proof in $\mathcal{A}_{\mathfrak{M}, A, s}^{S_1 \cup S_2}$ whose root is labelled by (s, \vec{b}) and each leaf contains a control state in $\{(g, s') \in Q \mid s' \in S_2\} \cup \{s' \in Q \mid s' = s, s \in S_2\}$.

(2) \rightarrow (1). Suppose that there is a finite proof $\hat{\mathcal{D}}$ in $\mathcal{A}_{\mathfrak{M}, A, s}^{S_1 \cup S_2}$ whose root is labelled by (s, \vec{b}) and each leaf contains a control state in $\{(g, s') \in Q \mid s' \in S_2\} \cup \{s' \in Q \mid s' = s, s \in S_2\}$. One can extend $\hat{\mathcal{D}}$ in order to obtain an infinite proof $\hat{\mathcal{D}}'$ such that every maximal branch is infinite and it belongs to L_{S_1, S_2} . Any leaf labelled by the control state (g, s') is further extended by application of the unary rule $(g, s') \xrightarrow{\vec{0}} (s', \bar{f})$ where \bar{f} is the idle joint action (with the control state s , a similar method applies). Similarly, any leaf labelled by the control state (s', \bar{f}) is further extended by application of the unique fork rule starting by (s', \bar{f}) . It is easy to check that this not only leads to a derivation but also to a proof, because the extension only deals with the update vector $\vec{0}$. By Lemma 4, there is a \vec{b} -strategy w.r.t. s in \mathfrak{M} such that the set of computations $\text{Comp}(s, F_A)$ is included in L_{S_1, S_2} . \square

4.2 2EXPTIME upper bound

The upper bound is established by giving a labelling algorithm as done in [4] or for standard temporal logics such as CTL and CTL*. The main difference with [4] rests on the treatment of the cases with strategy modalities is not performed in an ad-hoc fashion using the fact that (\mathbb{N}^r, \leq) is a well-quasi-ordering by Dickson's Lemma [21] but rather we explicitly call subroutines that solve decision problems on AVASS. The existence of such subroutines is due to [40] for

monotonic games, and their complexity upper bounds are due to [32, Theorem 3.4] and [16, Theorem 3.1]. Once more, the proof of the 2EXPTIME upper bound is divided into three main steps:

1. we introduce a slight extension of AVASS such that the decision problems remain in 2EXPTIME;
2. we show that the cases for the strategy modalities can be faithfully reduced to subroutines for problems on such extended AVASS (a consequence of developments from Section 4.1);
3. finally, we design a labelling algorithm and establish the complexity upper bound from it.

First, let us introduce a slight extension of decision problems for AVASS.

Lemma 7. *In the following extension of AVASS the state reachability and non-termination problems remain in 2EXPTIME:*

- Fork rules can be α -ary for any $\alpha \geq 1$ (but there is only a finite number of them).
- Reachability is related to a subset $S_f \subseteq S$ (instead of a singleton set).
- The initial configuration is (q_0, \vec{b}) with $\vec{b} \in \mathbb{N}^r$ instead of the fixed tuple $\vec{0}$.
- The value ω in \vec{b} is allowed and absorbs any other value in \mathbb{Z} (a means to ignore components. i.e. to reduce the dimension).

The proof is fairly standard, and consists in using Proposition 2 by simulating a non-binary fork by a linear-size gadget made of unary and binary forking rules, and by adding binary forking rules from states in S_f to a new single final state.

Proof. The lemma states four ways to extend the decision problems on AVASS, either by slightly extending the notion of AVASS, or by considering more general inputs for the problems. For each extension, we show how this can be encoded into the state reachability and the non-termination problems on AVASS using only polynomial-time reductions. The proof of the lemma is then obtained by composition of the reductions (polynomial-time reductions are also known to be closed under compositions) and by invoking Proposition 2 to get the 2EXPTIME upper bound.

- Let $\mathcal{A} = (Q, r, R_1, R_2)$ be an alternating VASS, $q_0, q_f \in Q$, and $\mathfrak{r} = (q_1, \dots, q_{\alpha+1})$ be an (extended) α -ary rule. If $\alpha = 1$, the rule can be treated as a unary rule with the update vector $\vec{0}$, whereas if $\alpha = 2$, it can be treated as a standard binary fork rule. So assume $\alpha \geq 3$. Let R'_2 be the following set of binary fork rules derived from \mathfrak{r} where $q'_1, \dots, q'_{\alpha-1}$ are new control states:

$$R'_2 = \{(q_1, q_2, q'_1)\} \cup \{(q'_j, q_{j+1}, q'_{j+1}) \mid j \in [1, \alpha - 2]\} \cup \{(q'_{\alpha-1}, q_\alpha, q_{\alpha+1})\}.$$

It is easy to show that the statements below are equivalent:

- there is a finite proof in $(Q, r, R_1, R_2 \uplus \mathfrak{r})$ whose root is equal to $(q_0, \vec{0})$ and each leaf belongs to $\{q_f\} \times \mathbb{N}^r$;
- there is a finite proof in $(Q \uplus \{q'_j \mid j \in [1, \alpha - 1]\}, r, R_1, R_2 \uplus R'_2)$ whose root is equal to $(q_0, \vec{0})$ and each leaf belongs to $\{q_f\} \times \mathbb{N}^r$.

If there is more than one extended fork rule, we apply the above reduction as many times as necessary, leading eventually to a reduction to an instance of the state reachability problem for AVASS. The same reduction also works for the non-termination problem.

- Let $\mathcal{A} = (Q, r, R_1, R_2)$, $q_0 \in Q$ and $S_f = \{q_1, \dots, q_\beta\}$. Let $\mathcal{A}' = (Q \uplus \{q_f^{\text{new}}\}, r, R'_1, R_2)$ be defined from \mathcal{A} such that $R'_1 \stackrel{\text{def}}{=} R_1 \uplus \{q_i \xrightarrow{\vec{0}} q_f^{\text{new}} \mid i \in [1, \beta]\}$. It is easy to show that the statements below are equivalent:
 - there is a finite proof in \mathcal{A} whose root is equal to $(q_0, \vec{0})$ and each leaf belongs to $S_f \times \mathbb{N}^r$;
 - there is a finite proof in \mathcal{A}' whose root is equal to $(q_0, \vec{0})$ and each leaf belongs to $\{q_f^{\text{new}}\} \times \mathbb{N}^r$.
- Let $\mathcal{A} = (Q, r, R_1, R_2)$, $q_0, q_f \in Q$ and $\vec{b} \in \mathbb{N}^r$ and $\mathcal{A}' = (Q \uplus \{q'_0\}, r, R_1 \uplus \{q'_0 \xrightarrow{\vec{b}} q_0\}, R_2)$. It is easy to show that the statements below are equivalent:
 - there is a finite proof in \mathcal{A} whose root is equal to (q_0, \vec{b}) and each leaf belongs to $\{q_f\} \times \mathbb{N}^r$;
 - there is a finite proof in \mathcal{A}' whose root is equal to $(q'_0, \vec{0})$ and each leaf belongs to $\{q_f\} \times \mathbb{N}^r$.

Similarly, the statements below are equivalent:

- there is a proof in \mathcal{A} whose root is equal to (q_0, \vec{b}) and all the maximal branches are infinite;
 - there is a proof in \mathcal{A}' whose root is equal to $(q'_0, \vec{0})$ and all the maximal branches are infinite.
- Let $\mathcal{A} = (Q, r, R_1, R_2)$, $q_0, q_f \in Q$ and $\vec{b} \in (\mathbb{N} \cup \{\omega\})^r$. We are looking for proofs whose root is labelled by (q_0, \vec{b}) and any occurrence of ω in \vec{b} remains for all the descendant nodes, which amounts to ignoring some components.

Suppose that ω occurs at least once in \vec{b} and let $\{i_1, \dots, i_\beta\} \subseteq [1, r]$ be the set of positions where ω occurs in \vec{b} . Let $r' = r - \beta$ and let $j_1 < \dots < j_{r'}$ be the indices in $[1, r] \setminus \{i_1, \dots, i_\beta\}$. Let $\bar{f} : [1, r'] \rightarrow \{j_1, \dots, j_{r'}\}$ be the bijection such that $\bar{f}(n) \stackrel{\text{def}}{=} j_n$. We define the alternating VASS $\mathcal{A}' = (Q, r', R'_1, R_2)$ obtained from \mathcal{A} by removing the components in positions in $\{i_1, \dots, i_\beta\}$. The map \bar{f} is extended to $\bar{f} : (\mathbb{Z} \cup \{\omega\})^r \rightarrow \mathbb{Z}^{r'}$ such that for all $\vec{u} \in \mathbb{Z}^{r'}$, $n \in [1, r']$ we have $\bar{f}(\vec{u})(n) \stackrel{\text{def}}{=} \vec{u}(j_n)$. The set of unary rules R'_1 is defined from R_1 as follows:

$$R'_1 \stackrel{\text{def}}{=} \{q \xrightarrow{\bar{f}(\vec{u})} q' \mid q \xrightarrow{\vec{u}} q' \in R_1\}.$$

It is easy to show that the statements below are equivalent:

- there is a finite proof in \mathcal{A} whose root is equal to (q_0, \vec{b}) and each leaf belongs to $\{q_f\} \times \mathbb{N}^r$;

- there is a finite proof in \mathcal{A}' whose root is equal to (q_0, \vec{b}) and each leaf belongs to $\{q_f\} \times \mathbb{N}^r$.

The same reduction also works for the non-termination problem.

□

Lemma 8 relates the satisfaction of a formula with the outermost connective $\langle\langle A^{\vec{b}} \rangle\rangle U$ and the state reachability problem for AVASS.

Lemma 8. *The statements below are equivalent:*

- (I) $\mathfrak{M}, s \models \langle\langle A^{\vec{b}} \rangle\rangle \phi_1 U \phi_2$.
- (II) *there is a finite proof in $\mathcal{A}_{\mathfrak{M}, A, s}^{S_1 \cup S_2}$ whose root is equal to (s, \vec{b}) and each leaf has a control state in $\{(g, s') \in Q \mid s' \in S_2\} \cup \{s' \in Q \mid s' = s, s \in S_2\}$ with $S_i = \{s \mid \mathfrak{M}, s \models \phi_i\}$, $i \in \{1, 2\}$.*

This is a direct consequence of Lemma 6.

Lemma 9 relates to the satisfaction of a formula with outermost connective $\langle\langle A^{\vec{b}} \rangle\rangle G$ and the non-termination problem for AVASS.

Lemma 9. *Assuming that $\mathfrak{M}, s \models \phi_1$, the statements below are equivalent:*

- (I) $\mathfrak{M}, s \models \langle\langle A^{\vec{b}} \rangle\rangle G \phi_1$.
- (II) *there is a proof in $\mathcal{A}_{\mathfrak{M}, A, s}^{S'}$ with $S' = \{s'' \mid \mathfrak{M}, s'' \models \phi_1\}$ whose root is equal to (s, \vec{b}) and every maximal branch is infinite.*

This is a direct consequence of Lemma 5.

Theorem 2. *The model checking problem for $RB \pm ATL$ is in 2EXPTIME.*

Proof. Algorithm 1 is a global model checking algorithm that takes as input a resource-bounded concurrent game structure \mathfrak{M} and a formula ϕ (both built on the same set of agents and with the same number of resources) and returns the set of states that satisfies the formula. By structural induction, one can show that $GMC(\mathfrak{M}, \psi) = \{s \mid \mathfrak{M}, s \models \psi\}$ by using Lemma 8 and Lemma 9. We use the fact that the state reachability and the non-termination problems for extended AVASS are decidable by [40] and by Lemma 7 (for the extension). Let us show how the proof by induction works.

Case $\psi = \langle\langle A^{\vec{b}} \rangle\rangle G \psi'$

The statements below are equivalent:

- $\mathfrak{M}, s \models \psi$.
- there is a \vec{b} -strategy F_A w.r.t. s such that the computations in $\text{Comp}(s, F_A)$ only visit states in $S'_1 = \{s' \mid \mathfrak{M}, s' \models \psi'\}$ (by definition of \models).
- $s \in S'_1$ and there is a proof in $\mathcal{A}_{\mathfrak{M}, A, s}^{S'_1}$ whose root is equal to (s, \vec{b}) and every maximal branch is infinite (by Lemma 9).

Algorithm 1 – RB±ATL model checking –

```

1: procedure GMC( $\mathfrak{M}, \phi$ )
2:   case  $\phi$  of
3:      $p$ : return  $\{s \in S \mid s \in L(p)\}$ 
4:      $\neg\psi$ : return  $S \setminus \text{GMC}(\mathfrak{M}, \psi)$ 
5:      $\psi_1 \wedge \psi_2$ : return  $\text{GMC}(\mathfrak{M}, \psi_1) \cap \text{GMC}(\mathfrak{M}, \psi_2)$ 
6:      $\langle\langle A^{\vec{b}} \rangle\rangle X\psi$ : return  $\{s \mid \exists \hat{f} \in D_A(s), \vec{0} \leq \text{cost}_A(s, \hat{f}) + \vec{b}, \text{ for all } \hat{f} \sqsubseteq g \in D(s), \delta(s, g) \in \text{GMC}(\mathfrak{M}, \psi)\}$ 
7:      $\langle\langle A^{\vec{b}} \rangle\rangle G\psi$ :  $S_1 := \text{GMC}(\mathfrak{M}, \psi)$ ;
       return  $\{s \in S_1 \mid \mathcal{A}_{\mathfrak{M}, A, s'}^{S_1}(s, \vec{b}) \text{ is non-terminating}\}$ 
8:      $\langle\langle A^{\vec{b}} \rangle\rangle \psi_1 U \psi_2$ : return  $\{s \mid \mathcal{A}_{\mathfrak{M}, A, s'}^{S_1 \cup S_2}(s, \vec{b}), S'_2 \text{ is a positive instance of state reachability}\}$ 
       with  $S_1 = \text{GMC}(\mathfrak{M}, \psi_1)$ ,  $S_2 = \text{GMC}(\mathfrak{M}, \psi_2)$ ,  $S'_2 = \{(g, s') \in Q \mid s' \in S_2\} \cup \{s' \in Q \mid s' = s, s \in S_2\}$ 
9:   end case
10: end procedure

```

- $s \in S_1$ and there is a proof in $\mathcal{A}_{\mathfrak{M}, A, s'}^{S_1}$ whose root is equal to (s, \vec{b}) and every maximal branch is infinite with $S_1 = \text{GMC}(\mathfrak{M}, \psi')$ (by induction hypothesis).
- $s \in S_1$ and $\mathcal{A}_{\mathfrak{M}, A, s'}^{S_1}(s, \vec{b})$ is a positive instance of the non-terminating problem for AVASS (by definition).

Case $\psi = \langle\langle A^{\vec{b}} \rangle\rangle \psi_1 U \psi_2$. The statements below are equivalent:

- $\mathfrak{M}, s \models \psi$.
- there is a \vec{b} -strategy F_A w.r.t. s such that for all $\lambda = s_0 \xrightarrow{f_0} s_1 \dots \in \text{Comp}(s, F_A)$, there is some $i < |\lambda|$ such that $\mathfrak{M}, s_i \models \psi_2$ and for all $j \in [0, i-1]$, we have $\mathfrak{M}, s_j \models \psi_1$.
- there is a finite proof in $\mathcal{A}_{\mathfrak{M}, A, s'}^{S_1 \cup S_2}$ whose root is equal to (s, \vec{b}) and each leaf has a control state in $\{(g, s') \in Q \mid s' \in S_2\} \cup \{s' \in Q \mid s' = s, s \in S_2\}$ with $S_i = \{s \mid \mathfrak{M}, s \models \psi_i\}$, $i \in \{1, 2\}$ (by Lemma 8).
- there is a finite proof in $\mathcal{A}_{\mathfrak{M}, A, s'}^{S_1 \cup S_2}$ whose root is equal to (s, \vec{b}) and each leaf has a control state in $\{(g, s') \in Q \mid s' \in S_2\} \cup \{s' \in Q \mid s' = s, s \in S_2\}$ with $S_i = \text{GMC}(\mathfrak{M}, \psi_i)$ $i \in \{1, 2\}$ (by the induction hypothesis).
- $\mathcal{A}_{\mathfrak{M}, A, s'}^{S_1 \cup S_2}(s, \vec{b})$, S'_2 is a positive instance of the state reachability problem for AVASS with $S_i = \text{GMC}(\mathfrak{M}, \psi_i)$ $i \in \{1, 2\}$ and $S'_2 = \{(g, s') \in Q \mid s' \in S_2\} \cup \{s' \in Q \mid s' = s, s \in S_2\}$.

Case $\psi = \langle\langle A^{\vec{b}} \rangle\rangle X\psi'$

The statements below are equivalent:

- $\mathfrak{M}, s \models \psi$.
- there is a \vec{b} -strategy F_A w.r.t. s such that for all $s_0 \xrightarrow{g_0} s_1 \dots \in \text{Comp}(s, F_A)$, we have $\mathfrak{M}, s_1 \models \psi'$ (by definition of \models).

- there is a \vec{b} -strategy F_A w.r.t. s such that for all $s_0 \xrightarrow{q_0} s_1 \dots \in \text{Comp}(s, F_A)$, we have $\mathfrak{M}, s_1 \models \psi'$, and for any finite computation extending $s_0 \xrightarrow{q_0} s_1$, F_A returns the constant map `idle` (thanks to the properties of the action `idle`).
- there is $\mathfrak{f} \in D_A(s)$ such that for all $s_1 \in \text{out}(s, \mathfrak{f})$, we have $\mathfrak{M}, s_1 \models \psi'$ and $\vec{0} \leq \text{cost}_A(s, \mathfrak{f}) + \vec{b}$.
- there is $\mathfrak{f} \in D_A(s)$ such that for all $g \sqsupseteq \mathfrak{f}$, we have $\mathfrak{M}, \delta(s, g) \models \psi'$ and $\vec{0} \leq \text{cost}_A(s, \mathfrak{f}) + \vec{b}$.
- there is $\mathfrak{f} \in D_A(s)$ such that for all $g \sqsupseteq \mathfrak{f}$, we have $\delta(s, g) \in \text{GMC}(\mathfrak{M}, \psi')$ and $\vec{0} \leq \text{cost}_A(s, \mathfrak{f}) + \vec{b}$ (by induction hypothesis).

As far as complexity is concerned, $\text{GMC}(\mathfrak{M}, \psi)$ can be solved by using a recursion depth that is linear in the size of ψ , and the state reachability and the non-termination problems for AVASS can be solved in 2EXPTIME by [32, Theorem 3.4] and [16, Theorem 3.1]. Note also the instances of such problems can be built in polynomial time in the respective sizes of \mathfrak{M} and ϕ . Consequently, the model checking problem for RB \pm ATL is in 2EXPTIME. \square

4.3 2EXPTIME-hardness

In this section, we show a 2EXPTIME-hardness result by reduction from the state reachability problem for AVASS. This improves the EXPSPACE-hardness result in [4].

Theorem 3. *The model checking problem for RB \pm ATL is 2EXPTIME-hard.*

Proof. The proof is by reduction from the state reachability problem for AVASS (see Proposition 2 or [16, Theorem 4.1]). It is divided into three main parts.

1. We consider a restriction of the state reachability problem for AVASS that remains 2EXPTIME-hard but that simplifies the definitions in the second part of the proof. Roughly speaking, the set of control states is divided in two disjoint sets, one from which unary rules start, and the other one from which fork rules start.
2. We then define the reduction from the restriction, taking care of the details of the resource-bounded concurrent game structures; essentially, we follow ideas similar to those in the proof of [4, Lemma 6].
3. Finally, we establish the correctness of the reduction.

(1) Given an instance $\mathcal{A} = (Q, r, R_1, R_2)$, q_0, q_f of the state reachability problem, we further assume that there is a partition $Q = Q_1 \uplus Q_2$ such that

- $q_0, q_f \in Q_1$,
- $R_1 \subseteq Q_1 \times \mathbb{Z}^r \times Q_2$ and $R_2 \subseteq Q_2 \times Q_1 \times Q_1$,
- there is no rule starting from q_f and there is at least one unary rule starting from q_0 .

The strict alternation between the control states in Q_1 and those in Q_2 can be obtained by duplicating the control states (in case a control state can start both a unary rule and a fork rule), and by adding new intermediate rules to enforce the alternation. In order to have no rule from q_f , it is sufficient to duplicate it and to reach q_f only for the last visit. The details follow.

Let $\mathcal{A} = (Q, r, R_1, R_2)$ be an alternating VASS. Without loss of generality, we can assume that no rule starts from q_f . Otherwise, we can introduce a new control state q_f^{new} that behaves almost as q_f : copy all the rules where q_f occurs in second or in third position by replacing q_f by q_f^{new} but no rule starting from q_f is copied (details are omitted). We can guarantee that (\star) there is a finite proof whose root is equal to $(q_0, \vec{0})$ and each leaf belongs to $\{q_f\} \times \mathbb{N}^r$ iff with the new AVASS there is a finite proof whose root is equal to $(q_0, \vec{0})$ and each leaf belongs to $\{q_f^{\text{new}}\} \times \mathbb{N}^r$. Similarly, without loss of generality, we can assume that there is a rule in R_1 that starts from q_0 . Otherwise, we add the dummy unary rule $(q_0, \vec{0}, q_0)$.

Let us build the alternating VASS $\mathcal{A}' = (Q', r, R'_1, R'_2)$ verifying the above conditions with $Q' = Q'_1 \uplus Q'_2$ such that (\star) iff there is a finite proof whose root is equal to $((q_0, 1), \vec{0})$ and each leaf belongs to $\{(q_f, 1)\} \times \mathbb{N}^r$. The set Q' is a subset of $Q \times \{1, 2\}$ defined by the clauses below plus auxiliary states introduced with the definition for rules:

- $(q, 1) \in Q' \stackrel{\text{def}}{\Leftrightarrow}$ there is a rule in R_1 that starts from q or $q = q_f$. So $(q_0, 1), (q_f, 1) \in Q'$.
- $(q, 2) \in Q' \stackrel{\text{def}}{\Leftrightarrow}$ there is a rule in R_2 that starts from q .
- $Q'_1 \supseteq \{(q, i) \in Q' \mid i = 1\}$ and $Q'_2 \supseteq \{(q, i) \in Q' \mid i = 2\}$. Obviously, $(q_0, 1), (q_f, 1) \in Q'_1$.

We now define the sets of rules R'_1 and R'_2 .

- For all $q \xrightarrow{\vec{u}} q' \in R_1$ such that $(q', 2) \in Q'$, we add the rule $(q, 1) \xrightarrow{\vec{u}} (q', 2)$ to R'_1 .
- For all $(q_1, q_2, q_3) \in R_2$ such that $(q_2, 1), (q_3, 1) \in Q'$, we add the fork rule

$$((q_1, 2), (q_2, 1), (q_3, 1))$$

to R'_2 .

- For all $r = q \xrightarrow{\vec{u}} q' \in R_1$ such that $(q', 1) \in Q'$, we add the rules below: $(q, 1) \xrightarrow{\vec{u}} q^{\text{new}} \in R'_1$ (q^{new} is new and depends on r) and $(q^{\text{new}}, (q', 1), (q', 1)) \in R'_2$. Moreover, $q^{\text{new}} \in Q'_2$. This amounts to adding an intermediate fork rule leading twice to $(q', 1)$ in order to guarantee that $R'_1 \subseteq Q'_1 \times \mathbb{Z}^r \times Q'_2$.
- For all $r = (q_1, q_2, q_3) \in R_2$ such that either $(q_2, 2) \in Q'$ or $(q_3, 2) \in Q'$. If $(q_2, 2), (q_3, 2) \in Q'$, then we add the rules below:
 - $((q_1, 2), q_2^{\text{new}}, q_3^{\text{new}}) \in R'_2$ where q_2^{new} and q_3^{new} are new and depend on the fork rule r . Moreover, these two new control states belong to Q'_1 .
 - $q_2^{\text{new}} \xrightarrow{\vec{0}} (q_2, 2)$ and $q_3^{\text{new}} \xrightarrow{\vec{0}} (q_3, 2)$ belong to R'_1 .

Again, we add intermediate unary rules in order to guarantee that $R'_2 \subseteq Q'_2 \times Q'_1 \times Q'_1$. If $(q_2, 2), (q_3, 1)$ belong to Q' or if $(q_2, 1), (q_3, 2)$ belong to Q' , the above construction can be easily adapted.




Figure 3: Partial description of \mathfrak{M}

One can show that $\mathcal{A}' = (Q', r, R'_1, R'_2)$ satisfies the above assumption and (\star) iff there is a finite proof whose root is equal to $((q_0, 1), \vec{0})$ and each leaf belongs to $\{(q_f, 1)\} \times \mathbb{N}^r$.

(2) Given an instance $\mathcal{A} = (Q, r, R_1, R_2)$, q_0 and q_f with the restriction above, we build the game structure $\mathfrak{M} = (Agt, S, Act, r, act, cost, \delta, L)$ with $Q_1 \subseteq S$ such that $\mathfrak{M}, q_0 \models \langle\langle \{1\}^{\vec{0}} \rangle\rangle \top \cup q_f$ iff there is a finite proof of AVASS whose root is equal to $(q_0, \vec{0})$ and each leaf belongs to $\{q_f\} \times \mathbb{N}^r$. Here, q_f is also understood as a propositional variable.

Before providing a formal definition of \mathfrak{M} , we illustrate the construction using a simple example. Assume that \mathcal{A} contains the unary rule $r_1 = q_1 \xrightarrow{(-1, +3)} q_0$ and the binary rule $r_2 = q_0 \rightarrow q_1, q_2$. \mathfrak{M} contains two agents, and an inference with r_1 followed by an inference with r_2 is simulated using the transitions shown in Figure 3 below, where the cost of the action (r_1, r_2) is precisely equal to $(-1, +3)$ and both *idle* and its twin action *idle'* have no cost. Note also the presence of a bad state that forbids the choice of the idle action by the first agent, assuming that the objective is to reach the state q_f .

The complete definition of \mathfrak{M} is as follows.

- $Agt \stackrel{\text{def}}{=} \{1, 2\}$.
- A pair of rules $(r_1, r_2) \in R_1 \times R_2$ is *connected* iff the last control state of r_1 is equal to the first control state of r_2 . The set of actions *Act* is equal to the set of connected pairs of rules plus the action *idle* and its twin action *idle'*.
- $S = Q_1 \uplus \{\text{bad}\}$.
- For each control state q in Q_1 , $act(q, 1)$ is the set of connected pairs of rules whose unary rule starts from q plus the idle action. So agent 1 can choose a unary rule immediately followed by a fork rule. $act(q_f, 1)$ is restricted to $\{\text{idle}\}$, because no rule starts from q_f in \mathcal{A} .
- As far as agent 2 is concerned, for all $q \in Q_1$, $act(q, 2) \stackrel{\text{def}}{=} \{\text{idle}, \text{idle}'\}$. So agent 2 can perform two actions that have no effect on resources, which amounts to simulating the effects of fork rules.
- Only the idle action can be performed from the bad state *bad*: $act(\text{bad}, 1) \stackrel{\text{def}}{=} act(\text{bad}, 2) \stackrel{\text{def}}{=} \{\text{idle}\}$.

- The cost of the action (r_1, r_2) is simply the update vector of the unary rule r_1 . Formally, for all $q \in Q_1$, we have $\text{cost}(q, 1, (r_1, r_2)) \stackrel{\text{def}}{=} \vec{u}$ when $r_1 = (q, \vec{u}, q')$ for some q' . Furthermore, $\text{cost}(q, a, \text{idle}) \stackrel{\text{def}}{=} \text{cost}(\text{bad}, a, \text{idle}) \stackrel{\text{def}}{=} \vec{0}$ for all $a \in \{1, 2\}$, and $\text{cost}(q, 2, \text{idle}') \stackrel{\text{def}}{=} \vec{0}$.
- When δ can be defined, we have
 - $\delta(q_f, \bar{f}) \stackrel{\text{def}}{=} q_f$; $\delta(\text{bad}, \bar{f}) \stackrel{\text{def}}{=} \text{bad}$ (\bar{f} is the constant map idle).
 - Whenever $q \in (Q_1 \setminus \{q_f\})$, $\bar{f}(1) = (r_1, r_2)$ with r_1 starting from q and $r_2 = (q_{\text{inter}}, q', q'')$,

$$\delta(q, \bar{f}) \stackrel{\text{def}}{=} \begin{cases} q' & \text{if } \bar{f}(2) = \text{idle} \\ q'' & \text{otherwise, i.e. } \bar{f}(2) = \text{idle}' \end{cases}$$

- $\delta(q, \bar{f}) = \text{bad}$ whenever $\bar{f}(1) = \text{idle}$ for all $q \in Q_1 \setminus \{q_f\}$.

- There is a unique propositional variable q_f and $L(q_f) \stackrel{\text{def}}{=} \{q_f\}$.

(3) We now establish the correctness of the construction.

Without loss of generality, we can assume $q_0 \neq q_f$.

First, suppose that $\mathfrak{M}, q_0 \models \langle\langle\{1\}^{\vec{0}}\rangle\rangle \top \cup q_f$. So, there exists a $\vec{0}$ -strategy $F_{\{1\}}$ such that, for all $\lambda = s_0 \xrightarrow{a_0} s_1 \dots \in \text{Comp}(q_0, F_{\{1\}})$, there is $i \geq 0$ such that $q_i = q_f$ (so by construction of \mathfrak{M} , for all $j \geq i$, we have $q_j = q_f$ and bad does not occur in λ). Since $\{1\}$ is a singleton set, we assume below that $F_{\{1\}}$ returns an action for agent 1 (instead of returning a joint action with respect to the unique agent 1) and $D_{\{1\}}(q)$ is viewed as an action.

From \mathfrak{M} , q_0 and $F_{\{1\}}$, let $(\mathfrak{T}_{F_{\{1\}}}, \mathfrak{R}, \mathfrak{Q})$ be the labelled transition system defined in Section 4.1. We have $\mathfrak{T}_{F_{\{1\}}} \subseteq \{1, 2\}^*$, $\mathfrak{Q} : \mathfrak{T}_{F_{\{1\}}} \rightarrow Q_1$ (because bad does not occur in computation from $\text{Comp}(q_0, F_{\{1\}})$), and \mathfrak{R} is partial map $\mathfrak{T}_{F_{\{1\}}} \times \mathfrak{T}_{F_{\{1\}}} \rightarrow \bigcup_{q \in Q_1} D(q)$. Note the specific structure of $(\mathfrak{T}_{F_{\{1\}}}, \mathfrak{R}, \mathfrak{Q})$:

- For all $w \in \mathfrak{T}_{F_{\{1\}}}$ such that $\mathfrak{Q}(w) = q_f$, $w \cdot 1$ is the unique successor of w , $\mathfrak{Q}(w \cdot 1) = q_f$ and $\mathfrak{R}(w, w \cdot 1)$ is the constant joint action equal to idle .
- For all $w \in \mathfrak{T}_{F_{\{1\}}}$ such that w has exactly two successors $w \cdot 1$ and $w \cdot 2$, there are $q_{\text{inter}} \in Q$ and $\vec{u} \in \mathbb{Z}^r$ such that

$$\begin{aligned} \mathfrak{R}(w, w \cdot 1) &= \overbrace{(((\mathfrak{Q}(w), \vec{u}, q_{\text{inter}}), (q_{\text{inter}}, \mathfrak{Q}(w \cdot 1), \mathfrak{Q}(w \cdot 2))), \text{idle})}^{r_1} \\ \mathfrak{R}(w, w \cdot 2) &= \overbrace{(((\mathfrak{Q}(w), \vec{u}, q_{\text{inter}}), (q_{\text{inter}}, \mathfrak{Q}(w \cdot 1), \mathfrak{Q}(w \cdot 2))), \text{idle}')^{r_2} \end{aligned}$$

and (r_1, r_2) is connected.

Now, let us build a finite derivation skeleton $\mathcal{D} : \mathfrak{T} \rightarrow (R_1 \cup R_2 \cup \{\perp\})$ such that its unique derivation with root labelled by $(q_0, \vec{0})$ is a finite proof with leaves labelled by q_f . The finite tree \mathfrak{T} can be uniquely defined from $\mathfrak{T}_{F_{\{1\}}}$ by using the map $c : \{1, 2\}^* \rightarrow \{1, 2\}^*$, where $c(w)$ is obtained from w by simultaneously replacing every occurrence of 1 by 11 and every occurrence of 2 by 12. So, for instance, $c(\varepsilon) \stackrel{\text{def}}{=} \varepsilon$ and $c(12) \stackrel{\text{def}}{=} 1112$. We stipulate that \mathfrak{T} is the set of words of the form $c(w)$ or $c(w) \cdot 1$ where $w \in \{1, 2\}^*$ and there is $v \in \mathfrak{T}_{F_{\{1\}}}$ such that w is a (non necessarily strict) prefix of v , $\mathfrak{Q}(v) = q_f$ and no strict prefix of v is labelled by q_f . The derivation skeleton \mathcal{D} is defined as follows. For all $w \in \mathfrak{T}_{F_{\{1\}}}$:

- If $\mathfrak{L}(w) \neq q_f$ and $\mathfrak{R}(w, w \cdot 1) = ((r_1, r_2), \text{idle})$, then $\mathcal{D}(c(w)) \stackrel{\text{def}}{=} r_1$ and $\mathcal{D}(c(w) \cdot 1) \stackrel{\text{def}}{=} r_2$.
- If $\mathfrak{L}(w) = q_f$ and $c(w) \in \mathfrak{T}$, then $\mathcal{D}(c(w)) = \perp$.

Similarly to Lemma 3, we can show the following properties:

- (I) For every computation λ starting at q_0 , ending at q_f , visiting q_f only once and respecting $F_{\{1\}}$, there is maximal branch w in \mathcal{D} such that $\text{ext}(\lambda, F_{\{1\}}) = \text{ext}(w, F_{\{1\}})$.
- (II) For every maximal branch w in \mathcal{D} , there is a computation λ starting at q_0 , ending at q_f , visiting q_f only once and respecting $F_{\{1\}}$ such that $\text{ext}(w, F_{\{1\}}) = \text{ext}(\lambda, F_{\{1\}})$.

Consequently, the unique derivation based on \mathcal{D} with root labelled by $(q_0, \vec{0})$ is a finite proof with leaves labelled by q_f . Indeed, by construction for all $q \in Q_1$, we have $\text{cost}(q, 1, (r_1, r_2)) = \vec{u}$ when $r_1 = (q, \vec{u}, q')$ for some q' .

For the converse direction, let us assume the existence of a finite proof $\hat{\mathcal{D}}$ based on the derivation skeleton $\mathcal{D} : \mathfrak{T} \rightarrow (R_1 \cup R_2 \cup \{\perp\})$ such that $\hat{\mathcal{D}}(\varepsilon) = (q_0, \vec{0})$ and each leaf is labelled by a pair in $\{q_f\} \times \mathbb{N}^r$. Let us define a strategy $F_{\{1\}}$. First, we require the following properties:

- $F_{\{1\}}(q_0) \stackrel{\text{def}}{=} (\mathcal{D}(\varepsilon), \mathcal{D}(1))$. Since $q_0 \neq q_f$, we know that $1 \in \mathfrak{T}$.
- For all the finite computations λ ending at the state q_f (we have $Q_1 \subseteq S$ and $q_f \in Q_1$), $F_{\{1\}}(\lambda) \stackrel{\text{def}}{=} \text{idle}$.

Let $\lambda = q_0 \xrightarrow{g_0} q_1 \xrightarrow{g_1} q_2 \cdots \xrightarrow{g_{n-1}} q_n$ be a finite computation respecting (so far) $F_{\{1\}}$ and $q_n \neq q_f$ (since this case is already treated above). Below we define $F_{\{1\}}(\lambda)$.

First, we assume that bad does not occur in λ . Each joint action g_j can be written $((r_1^j, r_2^j), b(k^j))$ with $k^j \in \{1, 2\}$ and $b : \{1, 2\} \rightarrow \{\text{idle}, \text{idle}'\}$ with $b(1) = \text{idle}$ and $b(2) = \text{idle}'$. The derivation skeleton \mathcal{D} verifies the properties below:

- $\mathcal{D}(\varepsilon) = r_1^0$ with $r_1^0 = (q_0, \vec{u}_0, q_{\text{inter}}^0)$.
- $\mathcal{D}(1) = r_2^0$ with $r_2^0 = (q_{\text{inter}}^0, q_1^0, q_2^0)$ and $q_{k^0}^0 = q_1$.
- $\mathcal{D}(1k^0) = r_1^1$ with $r_1^1 = (q_{k^0}^0, \vec{u}_1, q_{\text{inter}}^1)$.
- $\mathcal{D}(1k^01) = r_2^1$ with $r_2^1 = (q_{\text{inter}}^1, q_1^1, q_2^1)$ and $q_{k^1}^1 = q_2$.
- ...
- $\mathcal{D}(1k^01 \cdots k^j) = r_1^{j+1}$ with $r_1^{j+1} = (q_{k^j}, \vec{u}_{j+1}, q_{\text{inter}}^{j+1})$.
- $\mathcal{D}(1k^01 \cdots k^j1) = r_2^{j+1}$ with $r_2^{j+1} = (q_{\text{inter}}^{j+1}, q_1^{j+1}, q_2^{j+1})$ and $q_{k^{j+1}}^{j+1} = q_{j+2}$.
- ...
- $\mathcal{D}(1k^01 \cdots k^{n-2}1) = r_2^{n-1}$ with $r_2^{n-1} = (q_{\text{inter}}^{n-1}, q_1^{n-1}, q_2^{n-1})$ and $q_{k^{n-1}}^{n-1} = q_n$.

Since q_n is different from q_f , $1k^01 \dots k^{n-2}1k^{n-1}$ and $1k^01 \dots k^{n-2}1k^{n-1}1$ exist. We stipulate $F_{\{1\}}(\lambda) \stackrel{\text{def}}{=} (\mathcal{D}(1k^01 \dots k^{n-2}1k^{n-1}), \mathcal{D}(1k^01 \dots k^{n-2}1k^{n-1}1))$. Consequently, any extension $q_0 \xrightarrow{g_0} q_1 \xrightarrow{g_1} q_2 \dots \xrightarrow{g_{n-1}} q_n \xrightarrow{g_n} q_{n+1}$ respecting (again so far) $F_{\{1\}}$ verifies the above correspondence with \mathcal{D} and the state bad cannot be visited. So, the strategy $F_{\{1\}}$ can be defined by using the approach above (more formally, an induction hypothesis should be stated and we should prove that after each step, the property is preserved). One can also check that $F_{\{1\}}$ is $\vec{0}$ -consistent w.r.t q_0 , and, for all $\lambda = s_0 \xrightarrow{g_0} s_1 \dots \in \text{Comp}(q_0, F_{\{1\}})$, there is $i \geq 0$ such that $q_i = q_f$. The $\vec{0}$ -consistency is due to the fact that $\hat{\mathcal{D}}$ is a proof and the reachability condition is a consequence of the fact that every leaf of $\hat{\mathcal{D}}$ is labelled by q_f because of the correspondences between computations respecting $F_{\{1\}}$ and nodes in \mathfrak{T} . \square

The hardness proof above would also work if the proponent restriction is not satisfied, or if no distinguished idle action is assumed in the game structures, or if act may return an empty set of actions. It is also worth observing that one propositional variable and two agents are sufficient to get 2EXPTIME-hardness.

Corollary 1. *For any fixed $r \geq 1$, the model checking problem for RB \pm ATL restricted to at most r resources is in EXPTIME. For $r \geq 4$, the problem is EXPTIME-hard.*

Again, we use [32, Theorem 3.4] and [16, Theorem 3.1], which show that for a bounded number of resources, the state reachability and the non-termination problems for AVASS can be solved in EXPTIME. Moreover, if r is fixed but greater than two, then the model checking problem for RB \pm ATL restricted to at most r resources is PSPACE-hard, since the state reachability problem for VASS of dimension two is PSPACE-complete [9]. When $r = 1$, the model checking problem for RB \pm ATL is NP-hard since the state reachability for VASS of dimension one is NP-complete [26]. (Note that the NP-completeness result does not apply because the model checking problem for RB \pm ATL involves not just the reachability problem but also the non-termination problem.)

We have seen that the model checking problem for RB \pm ATL restricted to two agents is 2EXPTIME-hard; below we show that the restriction to a single agent is only EXPSPACE-complete.

Theorem 4. *The model checking problem for RB \pm ATL restricted to a single agent is EXPSPACE-complete.*

Proof. In order to show the EXPSPACE upper bound, we sketch how to solve the model checking problem for RB \pm ATL restricted to a single agent, by solving instances of the model checking problem for LTL on VASS or instances of the model checking problem for CTL, known to be EXPSPACE-complete (see e.g. [27]) and P-complete (see e.g. [42]) respectively. The labelling algorithm has exactly the same form as the algorithm for full RB \pm ATL. The size of the instances of the problems is linear in the size of the input resource-bounded concurrent game structures, and the number of calls is also linear in the size of the input formulae. This leads to the EXPSPACE upper bound.

The following two properties are essential for the proof. Given a path formula Φ of the form Xp , Gp or p_1Up_2 , one can show that $\mathfrak{M}, s \models \langle\langle \emptyset^{\vec{b}} \rangle\rangle \Phi$ iff $\mathfrak{M}, s \models A\Phi$ in CTL, where \mathfrak{M}' is obtained from \mathfrak{M} by removing the costs and actions from the transitions. Note that the empty coalition \emptyset allows us to quantify over all computations, and therefore the value of the bound \vec{b} is irrelevant. Similarly, one can show that $\mathfrak{M}, s \models \langle\langle \{1\}^{\vec{b}} \rangle\rangle \Phi$ iff there is an infinite run from the initial

configuration (s, \vec{b}) in the VASS \mathcal{V} that satisfies the LTL formula Φ , where \mathcal{V} is obtained from \mathfrak{M} by removing the actions while keeping the costs in \mathbb{Z}^r on the transitions. If \vec{b} has components with the value ω , then we reduce the dimension in \vec{b} and in \mathcal{V} so that only the components with finite values in \vec{b} remain.

In order to get the EXPSPACE lower bound, we reduce the state reachability problem for VASS to the model checking problem for RB \pm ATL restricted to a single agent, and use the EXPSPACE-hardness established in [37]. Let $\mathcal{V} = (Q, r, R)$ be a VASS and q_0, q_f be locations. One can show that there is a run from $(q_0, \vec{0})$ to some configuration of the form (q_f, \vec{x}) for some $\vec{x} \in \mathbb{N}^r$ iff $\mathfrak{M}, q_0 \models \langle\langle 1 \rangle^{\vec{0}}\rangle \top \text{Up}$, where $\mathfrak{M} = (\text{Agt}, S, \text{Act}, r, \text{act}, \text{cost}, \delta, L)$ is defined from \mathcal{V} as follows (for all $q, q' \in Q$):

- $\text{Agt} = \{1\}$, $S = Q \uplus \{\text{bad}\}$ and $\text{Act} = R \uplus \{\text{idle}\}$.
- L is defined so that p holds true exactly on q_f and $\text{cost}(1, q \xrightarrow{\vec{u}} q') = \vec{u}$.
- $\text{act}(q, 1) = \{\text{idle}\} \uplus \{q \xrightarrow{\vec{u}} q' \mid q \xrightarrow{\vec{u}} q' \in R\}$, $\text{act}(\text{bad}, 1) = \{\text{idle}\}$.
- Finally, $\delta(q, q \xrightarrow{\vec{u}} q') = q'$, $\delta(q, \text{idle}) = \text{bad}$ and $\delta(\text{bad}, \text{idle}) = \text{bad}$.

□

4.4 Bounding other resources in resource-bounded alternating-temporal logics

In this section, we study the model checking problem for resource-bounded logics in which the path formulae are arbitrary, i.e., they can be any LTL-like formulae rather than being restricted to path formulae of the form $\text{G}\psi$, $\text{X}\psi$ and $\psi_1 \cup \psi_2$ as in RB \pm ATL. We have already seen that the model checking problem for RBTL is EXPSPACE-hard (see Section 3.3) and therefore the lower bound also applies to RBTL*. Below, we show that the model checking problem for RBTL* is not only decidable (a new result) but also in EXPSPACE. The arguments for establishing the EXPSPACE upper bound for RBTL and RBTL* are identical, and the EXPSPACE lower bound for the model checking problem for RBTL can be matched with the upper bound for RBTL*.

Theorem 5. *The model checking problem for RBTL* is in EXPSPACE.*

In [12], the problem for RBTL is shown to be decidable by reduction to the reachability problem for VASS. However the best known upper bound for the reachability problem is \mathbf{F}_{ω^3} [36]. Hence, the EXPSPACE upper bound is a substantial improvement. Moreover, the decidability of the model checking problem for RBTL* was left open in [12].

Proof. (sketch) The algorithm to obtain the EXPSPACE upper bound first computes the states in which subformulae hold before dealing with larger formulae. The algorithm is a renaming algorithm. However, there is a caveat: when dealing with subformulae of the form $\langle\vec{b}\rangle\Psi$ where Ψ is an LTL formula, we are entitled to use the model checking algorithm for LTL formulae on VASS that is in EXPSPACE [27] (having ω in one component amounts to ignoring that position). However, in order to systematically consider such subformulae $\langle\vec{b}\rangle\Psi$ when the outermost connective is a path quantifier, we need to perform renamings on-the-fly.

Let us provide a simple example with the formula

$$\phi = \langle \vec{b}_0 \rangle \text{GF} \langle \vec{b}_1 \rangle q \text{U} q'$$

and an arbitrary model \mathfrak{M} . Let us consider some innermost state formula prefixed by a path quantifier, say $\langle \vec{b}_1 \rangle (q \text{U} q')$ (actually here there is only one such a subformula). With the help of a decision procedure for solving the LTL model checking problem on VASS, we determine for which control states q'' we have $\mathfrak{M}, q'' \models_s \langle \vec{b}_1 \rangle (q \text{U} q')$. Say, we obtain the set $\{q_1, \dots, q_\alpha\}$. Now, in ϕ , we replace $\langle \vec{b}_1 \rangle (q \text{U} q')$ by $q_1 \vee \dots \vee q_\alpha$, say we get $\phi_1 = \langle \vec{b}_0 \rangle \text{GF} (q_1 \vee \dots \vee q_\alpha)$. So, we have performed a renaming step by replacing a subformula by a disjunction of propositional variables. This process can be repeated until there are no more path quantifiers. To do this, we substitute some innermost state formula prefixed by a path quantifier in ϕ_1 , say $\langle \vec{b}_0 \rangle \text{GF} (q_1 \vee \dots \vee q_\alpha)$, by a new disjunction of locations (possibly empty) with the help of a decision procedure for solving the LTL model checking problem on VASS. We obtain the manageable formula

$$\phi_2 = q'_1 \vee \dots \vee q'_\beta.$$

Since ϕ_2 is a propositional formula, we are done with the renaming process and it is easy to show that for every control state q'' , we have $\mathfrak{M}, q'' \models_s \phi$ iff $\mathfrak{M}, q'' \models q'_1 \vee \dots \vee q'_\beta$. The above example can be easily generalised to any state formula. Note that the number of renamings is bounded by the size of the input formula, and at each step a subroutine is invoked at most $\text{card}(Q)$ times and requires EXPSPACE, whence we obtain the EXPSPACE upper bound. \square

Corollary 2. *For any fixed $r \geq 1$, the model checking problem for RBTL* restricted to at most r resources is in PSPACE.*

The PSPACE upper bound is a consequence of [27]. The model checking problem for LTL on VASS is PSPACE-complete when the number of counters is bounded [27, Theorem 4.1], and the renaming algorithm briefly described in the proof of Theorem 5 makes only a linear number of calls to the model checking problem for LTL on VASS and the number of counters is preserved when such calls are performed. If r is fixed but greater than two, then the model checking problem for RBTL* restricted to at most r resources is PSPACE-hard, since the state reachability problem for VASS of dimension two is PSPACE-complete [9]. When $r = 1$, the model checking problem for RBTL* is NP-hard, since the state reachability problem for VASS of dimension one is NP-complete [26].

5 The Logic RB \pm ATL* and its Parameterised Variant

We have seen that the model checking problem for RB \pm ATL is 2EXPTIME-complete. Below, we show that the model checking problem for RB \pm ATL* is also decidable. The arguments for establishing the respective decidability of RB \pm ATL and RB \pm ATL* both rest on the decidability of decision problems for alternating VASS. However for the model checking problem for RB \pm ATL* we need to invoke the decidability of parity games on alternating VASS, which is stronger than the decidability of the state reachability and non-termination problems for AVASS. This more complex reduction, which uses ingredients such as the standard equivalence of expressive power of Büchi automata and deterministic parity automata on ω -words, is nevertheless rewarding, as it

allows us to synthesise concrete values for resource parameters, something which has heretofore not been possible for resource-bounded logics.

Below, we introduce $\text{RB}\pm\text{ATL}^*$, an extension of $\text{RB}\pm\text{ATL}$ in which the path formulae are unconstrained, i.e. they can be any LTL-like formula. Although $\text{RB}\pm\text{ATL}^*$ is a new logic, its definition follows a standard schema for branching-time temporal logics.

5.1 Definition

Given a set of agents $\text{Agt} = \{a_1, \dots, a_k\}$ and $r \geq 1$ resources, we write $\text{RB}\pm\text{ATL}^*(\text{Agt}, r)$ to denote the resource-bounded logic with k agents and r resources whose models are resource-bounded concurrent game structures with the same parameters. Formulae of $\text{RB}\pm\text{ATL}^*(\text{Agt}, r)$ are defined according to the grammar below (as in CTL^* or for RBTL^* , we distinguish between state formulae ϕ and path formulae Φ)

$$\begin{aligned}\phi &::= p \mid \neg\phi \mid \phi \wedge \phi \mid \langle\langle A^{\vec{b}} \rangle\rangle \Phi \\ \Phi &::= \phi \mid \neg\Phi \mid (\Phi \wedge \Phi) \mid \text{X}\Phi \mid (\Phi \text{U}\Phi) \mid \text{G}\Phi,\end{aligned}$$

where $p \in \text{PROP}$, $A \subseteq \text{Agt}$ and $\vec{b} \in (\mathbb{N} \cup \{\omega\})^r$. The set of state formulae ϕ for $\text{RB}\pm\text{ATL}^*(\text{Agt}, r)$ extends the set of formulae for $\text{RB}\pm\text{ATL}(\text{Agt}, r)$.⁵ In presenting the semantics for $\text{RB}\pm\text{ATL}^*$, we make an explicit distinction between state formulae and path formulae. The two satisfaction relations \models_s and \models_p are defined as follows.

$$\begin{aligned}\mathfrak{M}, s \models_s p &\stackrel{\text{def}}{\iff} s \in L(p) \\ \mathfrak{M}, s \models_s \langle\langle A^{\vec{b}} \rangle\rangle \Phi &\stackrel{\text{def}}{\iff} \text{there is a } \vec{b}\text{-strategy } F_A \text{ such that} \\ &\quad \text{for all } \lambda = s_0 \xrightarrow{i_0} s_1 \dots \in \text{Comp}(s, F_A), \text{ we have } \mathfrak{M}, \lambda \models_p \Phi \\ \mathfrak{M}, \lambda \models_p \phi &\stackrel{\text{def}}{\iff} \mathfrak{M}, \lambda(0) \models_s \phi \text{ for state formulae } \phi \\ \mathfrak{M}, \lambda \models_p \neg\Phi &\stackrel{\text{def}}{\iff} \mathfrak{M}, \lambda \not\models_p \Phi \\ \mathfrak{M}, \lambda \models_p \Phi \wedge \Phi' &\stackrel{\text{def}}{\iff} \mathfrak{M}, \lambda \models_p \Phi \text{ and } \mathfrak{M}, \lambda \models_p \Phi' \\ \mathfrak{M}, \lambda \models_p \text{X}\Phi &\stackrel{\text{def}}{\iff} \mathfrak{M}, \lambda[1, +\infty) \models_p \Phi \\ \mathfrak{M}, \lambda \models_p \text{G}\Phi &\stackrel{\text{def}}{\iff} \mathfrak{M}, \lambda[i, +\infty) \models_p \Phi \text{ for all } i < |\lambda| \\ \mathfrak{M}, \lambda \models_p \Phi \text{U}\Psi &\stackrel{\text{def}}{\iff} \text{there is } i < |\lambda| \text{ such that } \mathfrak{M}, \lambda[i, +\infty) \models_p \Psi \text{ and} \\ &\quad \text{for every } j \in [0, i-1], \text{ we have } \mathfrak{M}, \lambda[j, +\infty) \models_p \Phi.\end{aligned}$$

Again, all the maximal computations are infinite, i.e., the index i in the clauses for G or U can take any value in \mathbb{N} . The *model checking problem for $\text{RB}\pm\text{ATL}^*$* is defined as follows:

Input: $k, r \geq 1$ (in unary), a state formula ϕ in $\text{RB}\pm\text{ATL}^*([1, k], r)$, a finite model \mathfrak{M} and a state s ,

Question: $\mathfrak{M}, s \models_s \phi$?

Below, we show that the model checking problem for $\text{RB}\pm\text{ATL}^*$ is decidable by reduction to the parity game problem for single-sided VASS [1, Corollary 2]. The latter problem will play a role similar to LTL model checking in CTL^* model checking, see e.g. [42, 19]. In addition, [1, Theorem 4] allows us to synthesise resource bounds. We begin by precisely defining the problem.

⁵ G can be encoded using U and \neg , but we retain it to emphasize that we are dealing with an extension of $\text{RB}\pm\text{ATL}(\text{Agt}, r)$.

5.2 A parameterised variant

We first introduce a *parameterised version* of $\text{RB}\pm\text{ATL}^*$, denoted by $\text{ParRB}\pm\text{ATL}^*$. The formulae of $\text{ParRB}\pm\text{ATL}^*$ are the same as those of $\text{RB}\pm\text{ATL}^*$, except that the concrete values $\vec{b} \in (\mathbb{N} \cup \{\omega\})^r$ are replaced by tuples of variables taken from the set $\text{VAR} = \{\mathbf{x}_1, \mathbf{x}_2, \dots\}$, for example:

$$\langle\langle\{1\}^{(\mathbf{x}_1, \mathbf{x}_2)}\rangle\rangle \top \text{U}q_f \wedge \langle\langle\{2\}^{(\mathbf{x}_2, \mathbf{x}_3)}\rangle\rangle \top \text{U}q'_f.$$

Given a parameterised (state or path) formula ϕ with variables $\mathbf{x}_1, \dots, \mathbf{x}_n$ and a map $v : \{\mathbf{x}_1, \dots, \mathbf{x}_n\} \rightarrow (\mathbb{N} \cup \{\omega\})$, we write $v(\phi)$ to denote the formula in $\text{RB}\pm\text{ATL}^*$ obtained from ϕ by replacing each occurrence of a variable \mathbf{x} by $v(\mathbf{x})$. The *parameterised model checking problem for $\text{ParRB}\pm\text{ATL}^*$* is defined as follows:

Input: $k, r \geq 1$ (in unary), a parameterised state formula ϕ in $\text{ParRB}\pm\text{ATL}^*([1, k], r)$, a finite $\text{RB}\pm\text{ATL}^*([1, k], r)$ model \mathfrak{M} and a state s ,

Question: Compute the set of maps v such that $\mathfrak{M}, s \models_s v(\phi)$.

By ‘compute the set of maps’, we mean being able to characterise the set of maps v such that $\mathfrak{M}, s \models_s v(\phi)$, by using a symbolic representation with nice computational properties. More precisely, we will show that we need only Boolean combinations of atomic formulae of the form $\mathbf{x} \geq k$ where $k \in \mathbb{N}$ and $\mathbf{x} = \omega$. Given such constraints, it is easy to check non-emptiness or to check the satisfaction of $\mathfrak{M}, s \models_s v(\phi)$ for a specific map v . To synthesise such parameters we use a remarkable result from single-sided VASS: the Pareto frontier for any parity game on single-sided VASS is computable [1, Theorem 4].

5.3 Parity acceptance condition

Below, we consider AVASS with a finite set of fork rules included in $\bigcup_{\beta \geq 2} Q^\beta$, and where the proofs are trees with nodes labelled by elements in $Q \times (\mathbb{N} \cup \{\omega\})^r$. Given an AVASS $\mathcal{A} = (Q, r, R_1, R_2)$, a *colouring* col is defined as a map $Q \rightarrow [0, p - 1]$ for some $p \geq 1$ (number of priorities). The *parity game problem* for AVASS is defined as follows:

Input: An alternating VASS \mathcal{A} , a control state $q_0, \vec{b} \in (\mathbb{N} \cup \{\omega\})^r$ and $\text{col} : Q \rightarrow [0, p - 1]$.

Question: Is there a proof whose root is equal to (q_0, \vec{b}) , all the maximal branches are infinite and the maximal colour that appears infinitely often is even (the colour of each configuration is induced by col) ?

Proposition 4. [1, Corollary 2] *The parity game problem for alternating VASS is decidable.*

To be precise, [1, Corollary 2] states the result for single-sided VASS. A single-sided VASS can be viewed as an alternating VASS where the set Q of control states can be partitioned into $Q = Q_1 \uplus Q_2$, unary rules start from states in Q_1 , fork rules start from states in Q_2 and there is at most one fork rule starting from the same control state (necessarily, it belongs to Q_2). The construction of two disjoint sets Q_1 and Q_2 with alternation of unary rules and fork rules can be done as in part (1) in the proof of Theorem 3. However, the colour of the new control states is equal to zero to avoid any new constraint. In order to guarantee the uniqueness of fork rules starting from a given control state, it is sufficient to replace any unary rule $q \xrightarrow{u} q'$ and fork rule $r = (q', q_1, q_2)$ by the unary rule $q \xrightarrow{u} (q', v)$ and the fork rule $((q', v), q_1, q_2)$. The colour of (q', v) is

the colour of q' . With such a transformation, decidability for single-sided VASS can be lifted to alternating VASS (this also implies a reduction for the computation of the Pareto frontier below).

It is not difficult to show that the state reachability and non-termination problems for AVASS are subproblems of the parity game problem, and therefore their decidability also follows from [1]. This decidability result was strengthened in [1] in the following way. Given \mathcal{A} , q_0 and $\text{col} : Q \rightarrow [0, p - 1]$, the set of tuples $\vec{b} \in (\mathbb{N} \cup \{\omega\})^r$ for which there is a proof such that the root is equal to (q_0, \vec{b}) , all the maximal branches are infinite, and for each infinite branch the maximal colour that appears infinitely often is even, is upward closed and computable. This means that it can be represented effectively by a Boolean combination of atomic constraints of the form $x_i \geq k$ where $i \in [1, r]$ and $k \in \mathbb{N}$, and $x_i = \omega$. Since the set is upward closed, by Dickson's Lemma, it has a finite set of minimal elements (with respect to the well-quasi-ordering \leq slightly extended to accommodate the addition of the value ω), allowing the symbolic representation in terms of atomic constraints of the form $x \geq k$ to be easily defined. The *Pareto frontier* of \mathcal{A} , q_0 and $\text{col} : Q \rightarrow [0, p - 1]$ is defined as the set of minimal elements in $(\mathbb{N} \cup \{\omega\})^r$ for which there is a positive solution to the parity game problem.

Proposition 5. [1, Theorem 4] *The Pareto frontier for any parity game on single-sided VASS is computable.*

5.4 A synchronised product

Before defining the reduction from the model checking problem for $\text{RB}\pm\text{ATL}^*$ to the parity game problem, we need to introduce a few more definitions, and in particular a notion of synchronisation that will be useful in the sequel.

Let $\mathfrak{M} = (\text{Agt}, S, \text{Act}, r, \text{act}, \text{cost}, \delta, L)$ be a resource-bounded concurrent game structure. Given the propositional variables p_1, \dots, p_n , we write $\Sigma_n \stackrel{\text{def}}{=} \mathcal{P}(\{p_1, \dots, p_n\})$ to denote the finite alphabet and $L_n(s') \stackrel{\text{def}}{=} \{p_i \mid i \in [1, n], s' \in L(p_i)\}$ for all $s' \in S$. So, by definition $L_n(s') \in \Sigma_n$.

Let $\mathcal{A}_{\mathfrak{M}, A, s} = (Q, r, R_1, R_2)$ be the AVASS defined from \mathfrak{M} , A and s (see Section 4.1), and $\mathbb{A} = (Q', q'_0, \delta : Q' \times \Sigma_n \rightarrow Q', \text{col} : Q' \rightarrow [0, p - 1])$ be a deterministic parity automaton over the alphabet Σ_n . The principle of the *synchronised product* $\mathcal{A}_{\mathfrak{M}, A, s} \otimes \mathbb{A}$ defined below is the following. Any (infinite) branch of a proof of $\mathcal{A}_{\mathfrak{M}, A, s}$ contains control states of the form s , (s', \mathfrak{f}) or (g, s') where s is a distinguished state of \mathfrak{M} , s' is any state, $\mathfrak{f} \in D_A(s')$ and g is joint action in $D(s')$ with $\delta(s', g) = s'$. By construction, (s', \mathfrak{f}) is preceded by a state of the form either (g, s') or s' (if $s' = s$). So an infinite branch of the form $(s_0, \vec{u}_0) ((s_0, \mathfrak{f}_0), \vec{u}_1) ((g_1, s_1), \vec{u}_1) ((s_1, \mathfrak{f}_1), \vec{u}_2) ((g_2, s_2), \vec{u}_2) \dots$ leads to the ω -word $L_n(s_0) L_n(s_1) L_n(s_2) \dots$ that corresponds to a unique run in \mathbb{A} (because \mathbb{A} is deterministic and complete).⁶ The control states of $\mathcal{A}_{\mathfrak{M}, A, s} \otimes \mathbb{A}$ are pairs in $Q \times Q'$ and the second components are therefore control states in Q' for the run on $L_n(s_0) L_n(s_1) L_n(s_2) \dots$.

Let us define the AVASS $\mathcal{A}_{\mathfrak{M}, A, s} \otimes \mathbb{A} \stackrel{\text{def}}{=} (Q'', r, R'_1, R'_2)$ such that:

- $Q'' \stackrel{\text{def}}{=} Q \times Q'$.
- For each unary rule $s \xrightarrow{\vec{u}} (s, \mathfrak{f}) \in R_1$, in R'_1 we have the unary rule $(s, q'_0) \xrightarrow{\vec{u}} ((s, \mathfrak{f}), q'_0)$.
- For each unary rule $(g, s') \xrightarrow{\vec{u}} (s', \mathfrak{f}) \in R_1$ and each $q \in Q'$, in R'_1 we have the unary rule $((g, s'), q) \xrightarrow{\vec{u}} ((s', \mathfrak{f}), q)$. So, firing a unary rule from $\mathcal{A}_{\mathfrak{M}, A, s}$ does not change the second component.

⁶We slightly abuse notation by identifying a branch with its label.

- For each fork rule $((s', \bar{f}), (g_1, s_1), \dots, (g_\alpha, s_\alpha)) \in R_2$, and for each $q \in Q'$, in R'_2 we have the fork rule

$$(((s', \bar{f}), q), ((g_1, s_1), q'), \dots, ((g_\alpha, s_\alpha), q')),$$

with $q' = \delta(q, L_n(s'))$. So, firing a fork rule from $\mathcal{A}_{\mathfrak{M}, \mathbb{A}, s}$ does change the second component in a unique way depending on q and on the letter $L_n(s')$. Again, there is a unique fork rule starting from the control state $((s', \bar{f}), q)$.

Let us define the colouring $\text{col}' : Q'' \rightarrow [0, p-1]$ such that for all $(q, q') \in Q''$, we have $\text{col}'((q, q')) \stackrel{\text{def}}{=} \text{col}(q')$. This is the most natural way to inherit colours from \mathbb{A} to $\mathcal{A}_{\mathfrak{M}, \mathbb{A}, s} \otimes \mathbb{A}$.

Lemma 10. *Let $(s, \vec{b}) \in Q \times (\mathbb{N} \cup \{\omega\})^r$. The statements below are equivalent:*

- (I) $\mathcal{A}_{\mathfrak{M}, \mathbb{A}, s}$ has a proof whose root is equal to (s, \vec{b}) , all the maximal branches are infinite and the L_n -projection of each infinite branch belongs to $L(\mathbb{A})$.
- (II) $\mathcal{A}_{\mathfrak{M}, \mathbb{A}, s} \otimes \mathbb{A}$ has a proof whose root is equal to $((s, q'_0), \vec{b})$, all the maximal branches are infinite and for all infinite branches, the maximal colour that appears infinitely often is even (based on the colouring function col')

Given an infinite branch

$$s_0 \xrightarrow{i_0} (s_0, \bar{f}_0) \rightarrow (g_{k_1}^1, s_{k_1}^1) \xrightarrow{i_1} (s_{k_1}^1, \bar{f}_1) \rightarrow (g_{k_2}^2, s_{k_2}^2) \xrightarrow{i_2} (s_{k_2}^2, \bar{f}_2) \rightarrow (g_{k_3}^3, s_{k_3}^3) \dots$$

in a proof of $\mathcal{A}_{\mathfrak{M}, \mathbb{A}, s}$, its L_n -projection is simply defined as the ω -word in Σ_n^ω below:

$$L_n(s_0) L_n(s_{k_1}^1) L_n(s_{k_2}^2) L_n(s_{k_3}^3) \dots$$

Proof. (I) \rightarrow (II) Let $\hat{\mathcal{D}} : \mathfrak{T} \rightarrow Q \times \mathbb{N}^r$ be a proof of $\mathcal{A}_{\mathfrak{M}, \mathbb{A}, s}$ such that $\hat{\mathcal{D}}(\varepsilon) = (s, \vec{b})$, all the maximal branches are infinite and their L_n -projections belong to $L(\mathbb{A})$. This means that for any label

$$s_0 \xrightarrow{i_0} (s_0, \bar{f}_0) \rightarrow (g_{k_1}^1, s_{k_1}^1) \xrightarrow{i_1} (s_{k_1}^1, \bar{f}_1) \rightarrow (g_{k_2}^2, s_{k_2}^2) \xrightarrow{i_2} (s_{k_2}^2, \bar{f}_2) \rightarrow \dots$$

of an infinite branch, we have $L_n(s_0) L_n(s_{k_1}^1) L_n(s_{k_2}^2) L_n(s_{k_3}^3) \dots \in L(\mathbb{A})$.

Let $\hat{\mathcal{D}}' : \mathfrak{T} \rightarrow (Q \times Q') \times \mathbb{N}^r$ be the map defined below. $\hat{\mathcal{D}}'$ will turn out to be a proof of $\mathcal{A}_{\mathfrak{M}, \mathbb{A}, s} \otimes \mathbb{A}$ built over the same infinite tree \mathfrak{T} . Let $i_1 i_2 i_3 \dots$ be an infinite branch with the label above such that

$$L_n(s_0) L_n(s_{k_1}^1) L_n(s_{k_2}^2) L_n(s_{k_3}^3) \dots \in L(\mathbb{A}).$$

Since \mathbb{A} is deterministic, there is a unique (accepting) run $q'_0 \xrightarrow{L_n(s_0)} q'_1 \xrightarrow{L_n(s_{k_1}^1)} q'_2 \dots$; hence the maximal colour that appears infinitely often is even. For any finite prefix $w \sqsubset i_1 i_2 i_3 \dots$ of length N , we have

$$\hat{\mathcal{D}}'(w) \stackrel{\text{def}}{=} ((q, q'_{\lfloor \frac{N}{2} \rfloor}), \vec{v}) \text{ where } \hat{\mathcal{D}}(w) = (q, \vec{v}).$$

Since \mathbb{A} is deterministic, the map $\hat{\mathcal{D}}'$ can be defined uniquely. Indeed, classically, if \mathbb{A} were nondeterministic, we cannot guarantee that two infinite words in $L(\mathbb{A})$ sharing a common non-empty prefix have the same subrun for that prefix. Since $\mathcal{A}_{\mathfrak{M}, \mathbb{A}, s} \otimes \mathbb{A}$ is also the synchronised product between $\mathcal{A}_{\mathfrak{M}, \mathbb{A}, s}$ and \mathbb{A} , we can check that $\hat{\mathcal{D}}'$ is indeed a proof of $\mathcal{A}_{\mathfrak{M}, \mathbb{A}, s} \otimes \mathbb{A}$, such that

$\hat{\mathcal{D}}'(\varepsilon) = ((s, q'_0), \vec{b})$ and all the maximal branches are infinite. Consider below the label of any infinite branch:

$$(s_0, q'_0) \xrightarrow{i_0} ((s_0, \bar{f}_0), q'_0) \rightarrow ((g_{k_1}^1, s_{k_1}^1), q'_1) \xrightarrow{i_1} ((s_{k_1}^1, \bar{f}_1), q'_1) \rightarrow ((g_{k_2}^2, s_{k_2}^2), q'_2) \xrightarrow{i_2} \dots$$

Since $q'_0 \xrightarrow{L_n(s_0)} q'_1 \xrightarrow{L_n(s_{k_1}^1)} q'_2 \dots$ is an accepting run, the maximal colour that appears infinitely often on the branch is even.

(II) \rightarrow (I) Let $\hat{\mathcal{D}} : \mathfrak{T} \rightarrow (Q \times Q') \times \mathbb{N}^r$ be a proof of $\mathcal{A}_{\mathfrak{M}, A, s} \otimes \mathbb{A}$ whose root is $((s, q'_0), \vec{b})$, all the maximal branches are infinite and the maximal colour that appears infinitely often is even. Let $\hat{\mathcal{D}}' : \mathfrak{T} \rightarrow Q \times \mathbb{N}^r$ be the map defined below. $\hat{\mathcal{D}}'$ will turn out to be a proof of $\mathcal{A}_{\mathfrak{M}, A, s}$, and can be viewed as $\hat{\mathcal{D}}$ where the component in Q' is omitted. For all $w \in \mathfrak{T}$, $\hat{\mathcal{D}}'(w) \stackrel{\text{def}}{=} (q, \vec{v})$ where $\hat{\mathcal{D}}(w) = ((q, q'), \vec{v})$. We have $\hat{\mathcal{D}}'(\varepsilon) = (s, \vec{b})$ and all the maximal branches are infinite. Consider below the label of an infinite branch $i_1 i_2 i_3 \dots$:

$$s_0 \xrightarrow{i_0} (s_0, \bar{f}_0) \rightarrow (g_{k_1}^1, s_{k_1}^1) \xrightarrow{i_1} (s_{k_1}^1, \bar{f}_1) \rightarrow (g_{k_2}^2, s_{k_2}^2) \xrightarrow{i_2} (s_{k_2}^2, \bar{f}_2) \rightarrow \dots$$

Since $\hat{\mathcal{D}}'$ is obtained from $\hat{\mathcal{D}}$ by projection, the label of $i_1 i_2 i_3 \dots$ in $\hat{\mathcal{D}}$ is of the form

$$(s_0, q'_0) \xrightarrow{i_0} ((s_0, \bar{f}_0), q'_0) \rightarrow ((g_{k_1}^1, s_{k_1}^1), q'_1) \xrightarrow{i_1} ((s_{k_1}^1, \bar{f}_1), q'_1) \rightarrow ((g_{k_2}^2, s_{k_2}^2), q'_2) \xrightarrow{i_2} \dots$$

By assumption on $\hat{\mathcal{D}}$, the maximal colour that appears infinitely often is even, and therefore $q'_0 \xrightarrow{L_n(s_0)} q'_1 \xrightarrow{L_n(s_{k_1}^1)} q'_2 \dots$ is an accepting run of \mathbb{A} and

$$L_n(s_0) L_n(s_{k_1}^1) L_n(s_{k_2}^2) L_n(s_{k_3}^3) \dots \in L(\mathbb{A}).$$

□

5.5 Decidability

Theorem 6. *The model checking problem for $RB\pm ATL^*$ is decidable.*

Proof. The model checking problem for $RB\pm ATL^*$ is solved by using the algorithm for the parity game problem for alternating VASS (with a simple renaming technique) as a subroutine. The algorithm uses a dynamic programming approach that first computes in which states the subformulae hold before dealing with larger formulae. However, there is a caveat: when dealing with subformulae of the form $\langle\langle A^{\vec{b}} \rangle\rangle \Phi$ where Φ is a path formula without any strategy modality, we are entitled to use the algorithm to solve the parity game problem for alternating VASS. However, in order to systematically consider such subformulae $\langle\langle A^{\vec{b}} \rangle\rangle \Phi$ when the outermost connective is a strategy modality, we need to perform renamings on-the-fly.

Let ϕ be a formula built over the propositional variables $\{p_1, \dots, p_n\}$ and $\langle\langle A^{\vec{b}} \rangle\rangle \Phi$ be one of its subformulae such that no strategy modality occurs in Φ . Without loss of generality, we can assume that there is an injective map $\text{nom} : S \rightarrow [1, n]$, such that for every $s \in S$, $L(p_{\text{nom}(s)}) = \{s\}$. As a result, each propositional variable $p_{\text{nom}(s)}$ is true in a single state, namely in s . So, even though we assume that ϕ is built over $\{p_1, \dots, p_n\}$, we do not require that all the propositional variables in $\{p_1, \dots, p_n\}$ necessarily occur in ϕ (some of the propositional variables are only used to name states). Given a finite concurrent game structure \mathfrak{M} , it is always possible to enrich it

so that each state can be named by a dedicated propositional variable (also called a nominal in hybrid logics, see e.g. [6]). This can be done in linear time.

Since Φ is an LTL formula built over $\{p_1, \dots, p_n\}$, there is a Büchi automaton \mathbb{A} over the alphabet Σ_n such that $L(\mathbb{A})$ is equal to the set of models of Φ (over the set of propositional variables $\{p_1, \dots, p_n\}$), see e.g. [44]. Say that \mathbb{A} has N states and $N \leq 2^{|\Phi|}$. Since Büchi automata and deterministic parity automata both recognize the set of ω -regular languages, there is deterministic parity automaton \mathbb{B} with initial location q'_0 , $O(N!^2)$ states and $2N$ priorities such that $L(\mathbb{A}) = L(\mathbb{B})$ [41]. The automaton \mathbb{B} can be effectively computed from \mathbb{A} .⁷

Let $X \subseteq S$ be the set of states s such that $\mathcal{A}_{\mathfrak{M}, A, s} \otimes \mathbb{B}$ has a proof whose root is equal to $((s, q'_0), \vec{b})$, all the maximal branches are infinite and the maximal colour that appears infinitely often is even. We update the formula ϕ by replacing every occurrence of $\langle\langle A^{\vec{b}} \rangle\rangle \Phi$ by $\psi = \bigvee_{s' \in X} p_{\text{nom}(s')}$. The set X can be computed thanks to Proposition 4, and this is a correct step thanks to Lemma 10 and Lemma 4. Indeed, for all $s' \in S$, we have $\mathfrak{M}, s' \models_s \langle\langle A^{\vec{b}} \rangle\rangle \Phi$ iff $\mathfrak{M}, s' \models_s \bigvee_{s' \in X} p_{\text{nom}(s')}$, and, therefore, for all $s' \in S$, we have $\mathfrak{M}, s' \models_s \phi$ iff $\mathfrak{M}, s' \models_s \phi[\psi / \langle\langle A^{\vec{b}} \rangle\rangle \Phi]$, where $\phi[\psi / \langle\langle A^{\vec{b}} \rangle\rangle \Phi]$ is obtained from ϕ by substituting every occurrence of $\langle\langle A^{\vec{b}} \rangle\rangle \Phi$ by ψ . We update ϕ until there are no more strategy modalities, and therefore eventually ϕ is a Boolean combination of propositional variables, which is then easy to evaluate on a given state. It is worth noting that the total number of calls to the parity game problem for AVASS is linear in the size of the formula, each instance of the problem has a doubly-exponential number of locations, and the colouring map has an exponential number of priorities in the size of the input formula. \square

The proof of Theorem 6 uses a synchronised product between an alternating VASS and a deterministic parity automaton recognizing ω -words. This is reminiscent of the proof of a 2EXPTIME upper bound for ATL* model checking problem [5, Theorem 5.6]. However, the Rabin tree automata in the proof of [5, Theorem 5.6] are replaced by deterministic parity automata for encoding the LTL formulae, and by alternating VASS (with counters) as outcome of the synchronisation.

Theorem 7. *The model checking problem for $RB\pm ATL^*$ restricted to a single agent is EXPSPACE-complete.*

The proof is similar to the proof of Theorem 4. Again, the upper bound is obtained by solving instances of the model checking problem for LTL on VASS (when the coalition is a singleton set) or instances of the model checking problem for CTL* (when the coalition is the empty set), which are known to be EXPSPACE-complete (see e.g. [27]) and PSPACE-complete (see e.g. [42]) respectively.

Note also that resource-bounded concurrent game structures can be seen as generalisations of VASS and, the logic $RB\pm ATL^*$ is clearly a generalisation of CTL*, by using the correspondences $\langle\langle Agt^{\vec{a}} \rangle\rangle \Phi \approx E \Phi$, and $\langle\langle \emptyset^{\vec{a}} \rangle\rangle \Phi \approx A \Phi$. It may seem surprising that the model checking problem for $RB\pm ATL^*$ is decidable, given that the model checking problem for CTL* on VASS is known to be undecidable. However this can be explained by the different satisfaction relations in the two problems. In the case of $RB\pm ATL^*$, formulae are evaluated on states of a concurrent game structure, not on configurations made of states and counter values, and this makes all the difference.

⁷A similar construction to that for \mathbb{B} was used recently in [15] for model checking pushdown multi-agent systems.

It is also remarkable that the proof of Theorem 6 does *not* use the fact that the idle action is always among the action(s) returned by the action manager. In contrast, the proofs in Section 4 use the idle action in order to extend finite computations to infinite ones, by choosing the idle action for all the agents after the finite part of the computation that, e.g., witnesses the satisfaction of a next or an until formula (see, e.g., Lemma 8). As a consequence, we get the following decidability result (and the results in Section 5.6 below also hold for $\text{RB}\pm\text{ATL}^*$ without idle actions).

Corollary 3. *The model checking problem for the variant of $\text{RB}\pm\text{ATL}^*$ in which no idle action is assumed in the resource-bounded concurrent game structures (and the action manager always returns a non-empty set of actions) is decidable.*

The proof is the same as for Theorem 6. The restriction to a single agent can be also shown EXPSpace -complete, by using analogous arguments from the proof of Theorem 7 (EXPSpace -hardness proof is even simpler).

By combining our results from previous sections and those from [1], we have shown that the model checking problem for $\text{RB}\pm\text{ATL}^*$ is decidable. However, by exploiting techniques for the effective computation of the Pareto frontiers from [1], we can go further, and actually synthesise values for parameters. This is the subject of the next section.

5.6 Symbolic representations for sets of resource values

Let $\mathfrak{M} = (\text{Agt}, S, \text{Act}, r, \text{act}, \text{cost}, \delta, L)$ be a resource-bounded concurrent game structure and ϕ be a $\text{ParRB}\pm\text{ATL}^*$ formula such that the resource variables are among $\mathbf{x}_1, \dots, \mathbf{x}_m$ and the propositional variables are among p_1, \dots, p_n .

We write $\varphi_1 = \langle\langle A_1^{\mathbf{t}_1} \rangle\rangle\phi_1, \dots, \varphi_\alpha = \langle\langle A_\alpha^{\mathbf{t}_\alpha} \rangle\rangle\phi_\alpha$ to denote the subformulae of ϕ whose outermost connective is a strategy modality. The subformulae are arranged in order of increasing size. So, ϕ_1 does not contain a strategy modality, and it can be viewed as an LTL formula built over $\{p_1, \dots, p_n\}$. Each expression \mathbf{t}_i is a tuple of r variables (say $\mathbf{t}_i = (\mathbf{t}_i^1, \dots, \mathbf{t}_i^r)$). A variable can occur more than once in \mathbf{t}_i , and two distinct tuples \mathbf{t}_i and \mathbf{t}_j can share variables. This provides great flexibility in the logical formalism. Below, for each state $s \in S$ and $i \in [1, \alpha]$, we build a *constrained formula* $\hat{\varphi}_i^s$ over $\mathbf{x}_1, \dots, \mathbf{x}_m$ following the grammar:

$$\psi ::= \mathbf{x} \geq k \mid \mathbf{x} = \omega \mid \neg\psi \mid \psi \vee \psi \mid \psi \wedge \psi,$$

where $\mathbf{x} \in \text{VAR}$ and $k \in \mathbb{N}$. Such formulae are interpreted over valuations $\mathbf{v} : \text{VAR} \rightarrow \mathbb{N} \cup \{\omega\}$ with semantics based on the satisfaction relation $\mathbf{v} \models \psi$, and have the following key property: for all \mathbf{v} , we have $\mathbf{v} \models \hat{\varphi}_i^s$ iff $\mathfrak{M}, s \models \mathbf{v}(\langle\langle A_i^{\mathbf{t}_i} \rangle\rangle\phi_i)$. Note that ϕ_i may contain variables in $\mathbf{x}_1, \dots, \mathbf{x}_m$.

Before explaining how to construct the formulae $\hat{\varphi}_i^s$'s, we first explain how to construct from such formulae a constrained formula ψ_s such that for all \mathbf{v} , we have $\mathbf{v} \models \psi_s$ iff $\mathfrak{M}, s \models \mathbf{v}(\phi)$. Let $\varphi_{i_1}, \dots, \varphi_{i_z}$ be the maximal subformulae of ϕ such that their outermost connective is a strategy modality. Given a propositional valuation $\mathfrak{h} : \{\varphi_{i_1}, \dots, \varphi_{i_z}\} \rightarrow \{\perp, \top\}$, we define $\mathfrak{M}, s \models \mathfrak{h}(\phi)$ iff $\mathfrak{h}(\phi)$ obtained from ϕ by replacing simultaneously each φ_{i_j} by $\mathfrak{h}(\varphi_{i_j})$ is true in s ($\mathfrak{h}(\phi)$ may also contain propositional variables). The constrained formula ψ_s is defined as follows:

$$\bigvee_{\mathfrak{h} \text{ s.t. } \mathfrak{M}, s \models \mathfrak{h}(\phi)} \left(\bigwedge_{j \text{ s.t. } \mathfrak{h}(\varphi_{i_j}) = \top} \hat{\varphi}_{i_j}^s \right) \wedge \left(\bigwedge_{j \text{ s.t. } \mathfrak{h}(\varphi_{i_j}) = \perp} \neg \hat{\varphi}_{i_j}^s \right).$$

The generalised disjunction considers all possible valuations \mathfrak{h} that make $\mathfrak{h}(\phi)$ true in s , and the subsequent conjunction ensures that if $\mathfrak{h}(\varphi_\ell) = \top$ then $\hat{\varphi}_\ell^s$ is satisfied, otherwise $\neg\hat{\varphi}_\ell^s$ has to be satisfied.

We now explain now how to build $\hat{\varphi}_i^s$. Let $\varphi_{j_1}, \dots, \varphi_{j_\beta}$ be the maximal subformulae of ϕ_i where the outermost connective is a strategy modality. By assumption, the formulae are arranged in order of increasing size, so we have $\{j_1, \dots, j_\beta\} \subseteq [1, i-1]$, and, possibly, there may be no such subformulae in the case where ϕ_i has no strategy modality (for instance this happens when $i = 1$). Given $S_1, \dots, S_\beta \subseteq S$ and $I \subseteq [1, r]$, let $\psi_{I, S_1, \dots, S_\beta}^s$ be the formula encoding the Pareto frontier of $\mathcal{A}_{\mathfrak{M}, A, s} \otimes \mathbb{A}$ with initial state (s, q'_0) , the ω -components are exactly in I , and associated colouring such that \mathbb{A} is the deterministic parity automaton of L defined below. This latter language is actually an ω -regular language over the alphabet $\mathcal{P}(\{p_1, \dots, p_n\})$ defined by the LTL formula ϕ'_i that is obtained from ϕ_i by replacing every occurrence of φ_j by

$$\left(\bigvee_{s \in S_j} p_{\text{nom}(s)} \right) \wedge \left(\bigwedge_{s \in S \setminus S_j} \neg p_{\text{nom}(s)} \right).$$

The formula $\hat{\varphi}_i^s$ is then defined as the following disjunction:

$$\bigvee_{S_1, \dots, S_\beta \subseteq S, I \subseteq [1, r]} \psi_{I, S_1, \dots, S_\beta}^s \wedge \left(\bigwedge_{\gamma \in [1, \beta]} \left(\bigwedge_{s' \in S_\gamma} \hat{\varphi}_{j_\gamma}^{s'} \right) \wedge \left(\bigwedge_{s' \in S \setminus S_\gamma} \neg \hat{\varphi}_{j_\gamma}^{s'} \right) \right).$$

The proof of Lemma 11 below provides the formal justification for such a construction. However, intuitively, in the generalised disjunction each S_γ allows us to guess where φ_{j_γ} is true, and therefore $s' \in S_\gamma$ should be equivalent to the satisfaction of $\hat{\varphi}_{j_\gamma}^{s'}$. Since $\{j_1, \dots, j_\beta\} \subseteq [1, i-1]$, each formula $\hat{\varphi}_{j_\gamma}^{s'}$ is already defined, and therefore the formula $\hat{\varphi}_i^s$ can be safely built.

Lemma 11. *For all $i \in [1, \alpha]$, for all $s \in S$, for all concretisations \mathfrak{v} , we have $\mathfrak{v} \models \hat{\varphi}_i^s$ iff $\mathfrak{M}, s \models \mathfrak{v}(\langle\langle A_i^{\mathfrak{t}_i} \rangle\rangle \phi_i)$.*

Proof. The proof is by induction on i .

Base case: there is no strategy modality in ϕ_i (this includes the case when $i = 1$).

Recall that $\varphi_i = \langle\langle A_i^{\mathfrak{t}_i} \rangle\rangle \phi_i$ where $A \subseteq \text{Agt}$ and $\mathfrak{t}_i = (\mathfrak{t}_i^1, \dots, \mathfrak{t}_i^r)$. Since there is no strategy modality in ϕ_i , the formula ϕ_i is simply an LTL formula built over propositional variables in $\{p_1, \dots, p_n\}$. Let \mathbb{A} be a deterministic parity automaton such that the models of ϕ_i on Σ_n are precisely $L(\mathbb{A})$. It is known that \mathbb{A} can be effectively computed from ϕ_i such that the number of states in \mathbb{A} is doubly-exponential in the size of ϕ_i , and the number of priorities is (only) exponential in the size of ϕ_i .

Given $I \subseteq [1, r]$, we can construct a constrained formula ψ_I^s characterising the Pareto frontier of $\mathcal{A}_{\mathfrak{M}, A, s} \otimes \mathbb{A}$ with initial state (s, q'_0) and the values at the positions in I are equal to ω at the root (see [1, Theorem 4]). Consequently ψ_I^s is equivalent to $\psi_I^s \wedge (\bigwedge_{j \in I} \mathfrak{t}_i^j = \omega) \wedge (\bigwedge_{j \notin I} \mathfrak{t}_i^j \neq \omega)$. This means that for all $\vec{b} \in (\mathbb{N} \cup \{\omega\})^r$ such that (C_I) for all $j \in [1, r]$, $\vec{b}(j) = \omega$ iff $j \in I$, we have $\vec{b} \models \psi_I^s$ (meaning $\mathfrak{v} \models \psi_I^s$ with $\mathfrak{v}(\mathfrak{t}_i^j) \stackrel{\text{def}}{=} \vec{b}(j)$ for all $j \in [1, r]$) iff there is a proof of $\mathcal{A}_{\mathfrak{M}, A, s} \otimes \mathbb{A}$ whose root is $((s, q'_0), \vec{b})$, all the maximal branches are infinite and the maximal colour that appears infinitely often is even.

By Lemma 10, for all $\vec{b} \in (\mathbb{N} \cup \{\omega\})^r$ such that (C_I) , we have $\vec{b} \models \psi_I^s$ iff $\mathcal{A}_{\mathfrak{M}, A_i, s}$ has a proof whose root is (s, \vec{b}) , all the maximal branches are infinite and the L_n -projection of each infinite branch belongs to $L(\mathbb{A})$. By Lemma 4, for all $\vec{b} \in (\mathbb{N} \cup \{\omega\})^r$ such that (C_I) , we have $\vec{b} \models \psi_I^s$ iff there is a \vec{b} -strategy F_{A_i} w.r.t. s in \mathfrak{M} such that the set of computations $\text{Comp}(s, F_{A_i})$ is included in $L(\mathbb{A})$. So, for all $\vec{b} \in (\mathbb{N} \cup \{\omega\})^r$ such that (C_I) , we have $\vec{b} \models \psi_I^s$ iff $\mathfrak{M}, s \models \langle\langle A_i^{t_i} \rangle\rangle \phi_i$. Consequently, for all concretisations v such that for all $j \in [1, r]$ $v(t_j^i) = \omega$ iff $j \in I$, we have $v \models \psi_I^s$ iff $\mathfrak{M}, s \models v(\langle\langle A_i^{t_i} \rangle\rangle \phi_i)$. The formula $\hat{\varphi}_i^s$ is defined as a generalised disjunction parameterised by all the possible values for I , i.e. $\hat{\varphi}_i^s \stackrel{\text{def}}{=} \bigvee_{I \subseteq [1, r]} \psi_I^s$, and it is easy to check that for all v , we have $v \models \hat{\varphi}_i^s$ is equivalent to $\mathfrak{M}, s \models v(\langle\langle A_i^{t_i} \rangle\rangle \phi_i)$.

Induction step. $\varphi_i = \langle\langle A_i^{t_i} \rangle\rangle \phi_i$ with $\varphi_{j_1}, \dots, \varphi_{j_\beta}$, the maximal subformulae of ϕ_i ($\beta \geq 1$) such that its outermost connective is a strategy modality and for all $\gamma \in [1, \beta]$, for all $s' \in S$, for all v , we have $v \models \hat{\varphi}_{j_\gamma}^{s'}$ iff $\mathfrak{M}, s' \models v(\langle\langle A_{j_\gamma}^{t_{j_\gamma}} \rangle\rangle \phi_{j_\gamma})$.

The proof for the induction step is quite similar to the proof for the base case, except that we need to show that the renaming mechanism we use is correct. First, let us state a few basic properties. Given a state formula ϕ in $\text{RB}\pm\text{ATL}^*$, we write $\mathfrak{M} \models \phi \stackrel{\text{def}}{\Leftrightarrow}$ for all $s' \in S$, we have $\mathfrak{M}, s' \models_s \phi$, i.e. ϕ is valid in \mathfrak{M} .

(P1) Let ϕ, ψ, ψ' be state formulae in $\text{RB}\pm\text{ATL}^*$ such that ψ occurs in ϕ and $\mathfrak{M} \models \psi \Leftrightarrow \psi'$. Then $\mathfrak{M} \models \phi \Leftrightarrow \phi[\psi'/\psi]$, where $\phi[\psi'/\psi]$ is defined from ϕ by replacing every occurrence of ψ by ψ' .

(P2) Let $\gamma \in [1, \beta]$ and $S' \subseteq S$. For all v , the statements below are equivalent.

(P2.1) $\mathfrak{M} \models [(\bigvee_{s \in S'} p_{\text{nom}(s)}) \wedge (\bigwedge_{s \notin S'} \neg p_{\text{nom}(s)})] \Leftrightarrow v(\varphi_{j_\gamma})$.

(P2.2) For all $s' \in S$, $v \models \hat{\varphi}_{j_\gamma}^{s'}$ iff $s' \in S'$.

The proof of (P1) is quite standard but it is worth noting that the formulae ϕ , ψ and ψ' need to be *state* formulae. The proof of (P2) uses the induction hypothesis in a straightforward way.

Let $S_1, \dots, S_\beta \subseteq S$ and $I \subseteq [1, r]$. We write ϕ'_i to denote the formula obtained from ϕ_i replacing every occurrence of the formula φ_{j_γ} by

$$\varphi_{j_\gamma}^* \stackrel{\text{def}}{=} \left(\bigvee_{s' \in S_\gamma} p_{\text{nom}(s')} \right) \wedge \left(\bigwedge_{s' \in S \setminus S_\gamma} \neg p_{\text{nom}(s')} \right).$$

So, even though this is not explicit in the notation, the formula ϕ'_i obviously depends on S_1, \dots, S_β . Like the base case, the formula ϕ'_i is an LTL formula built over $\{p_1, \dots, p_n\}$, and one can compute a deterministic parity automaton \mathbb{A} such that the models of ϕ'_i on Σ_n are precisely $L(\mathbb{A})$. We can also construct a constrained formula $\psi_{I, S_1, \dots, S_\beta}^s$ characterising the Pareto frontier of $\mathcal{A}_{\mathfrak{M}, A_i, s} \otimes \mathbb{A}$, with initial state (s, q'_0) and the values at the positions in I are equal to ω at the root, see [1, Theorem 4].

Reasoning in the same way as in the base case (basically replace ϕ_i by ϕ'_i), we can show that, for all concretisations v such that for all $j \in [1, r]$ $v(t_j^i) = \omega$ iff $j \in I$, we have $v \models \psi_{I, S_1, \dots, S_\beta}^s$ iff $\mathfrak{M}, s \models v(\langle\langle A_i^{t_i} \rangle\rangle \phi'_i)$ (+). However, the formula that is important to us is $\langle\langle A_i^{t_i} \rangle\rangle \phi_i$ (rather than $\langle\langle A_i^{t_i} \rangle\rangle \phi'_i$), and this is the place where the induction hypothesis is again invoked.

First, let us introduce an auxiliary notion. A concretisation v is said to be *compatible with* I, S_1, \dots, S_β iff the conditions below hold:

1. $I = \{\gamma \in [1, r] \mid v(\mathbf{t}_i^\gamma) = \omega\}$.
2. For all $\gamma \in [1, \beta]$, $S_\gamma = \{s' \in S \mid \mathfrak{M}, s' \models v(\varphi_{j_\gamma})\}$.

Assuming that v is compatible with I, S_1, \dots, S_β , we have that:

- By induction hypothesis, for all $\gamma \in [1, \beta]$, $S_\gamma = \{s' \in S \mid v \models \hat{\varphi}_{j_\gamma}^{s'}\}$.
- By (P2), for all $\gamma \in [1, \beta]$, $\mathfrak{M} \models [(\bigvee_{s' \in S_\gamma} p_{\text{nom}(s')}) \wedge (\bigwedge_{s' \notin S_\gamma} \neg p_{\text{nom}(s')})] \Leftrightarrow v(\varphi_{j_\gamma})$.
- By (P1) and the newly established property (\dagger),

$$v \models \psi_{I, S_1, \dots, S_\beta}^s \text{ iff } \mathfrak{M}, s \models v(\langle\langle A_i^{\mathbf{t}_i} \rangle\rangle \phi_i)[\varphi_{j_1}/\varphi_{j_1}^*, \dots, \varphi_{j_\beta}/\varphi_{j_\beta}^*].$$

- Since $v(\phi_i)[\varphi_{j_1}/\varphi_{j_1}^*, \dots, \varphi_{j_\beta}/\varphi_{j_\beta}^*] = v(\phi_i)$, $v \models \psi_{I, S_1, \dots, S_\beta}^s$ iff $\mathfrak{M}, s \models v(\langle\langle A_i^{\mathbf{t}_i} \rangle\rangle \phi_i)$.
- By using the compatibility of v , we get that $v \models \psi_{I, S_1, \dots, S_\beta}^s \wedge (\bigwedge_{\gamma \in [1, \beta]} (\bigwedge_{s' \in S_\gamma} \hat{\varphi}_{j_\gamma}^{s'}) \wedge (\bigwedge_{s' \notin S_\gamma} \neg \hat{\varphi}_{j_\gamma}^{s'}))$ iff $\mathfrak{M}, s \models v(\langle\langle A_i^{\mathbf{t}_i} \rangle\rangle \phi_i)$.

The formula $\hat{\varphi}_i^s$ is defined as a generalised disjunction parameterised by all the possible values for I, S_1, \dots, S_β , and it is easy to check that for all v , we have $v \models \hat{\varphi}_i^s$ equivalent to $\mathfrak{M}, s \models v(\langle\langle A_i^{\mathbf{t}_i} \rangle\rangle \phi_i)$. \square

Theorem 8. *The parameterised model checking problem for ParRB \pm ATL* can be solved.*

Proof. Let ϕ be a state formula in ParRB \pm ATL* $([1, k], r)$, \mathfrak{M} be a resource-bounded concurrent game structure and $s \in S$. We write $\varphi_1 = \langle\langle A_1^{\mathbf{t}_1} \rangle\rangle \phi_1, \dots, \varphi_z = \langle\langle A_z^{\mathbf{t}_z} \rangle\rangle \phi_z$ to denote the maximal subformulae of ϕ such that the outermost connective is a strategy modality. By Lemma 11, for all $j \in [1, z]$, for all valuations $v : \text{VAR} \rightarrow \mathbb{N} \cup \{\omega\}$, we have $v \models \hat{\varphi}_j^s$ iff $\mathfrak{M}, s \models v(\langle\langle A_j^{\mathbf{t}_j} \rangle\rangle \phi_j)$, where $\hat{\varphi}_j^s$ is a constrained formula that can be built from ϕ, \mathfrak{M} and s . Let ψ_s be the formula defined by:

$$\bigvee_{\mathfrak{h}: \{\varphi_1, \dots, \varphi_z\} \rightarrow \{\perp, \top\} \text{ s.t. } \mathfrak{M}, s \models \mathfrak{h}(\phi)} \left(\bigwedge_{j \text{ s.t. } \mathfrak{h}(\varphi_j) = \top} \hat{\varphi}_j^s \right) \wedge \left(\bigwedge_{j \text{ s.t. } \mathfrak{h}(\varphi_j) = \perp} \neg \hat{\varphi}_j^s \right).$$

Each expression $\hat{\varphi}_j^s$ and $\neg \hat{\varphi}_j^s$ is a constrained formula, and $\mathfrak{M}, s \models \mathfrak{h}(\phi)$ can be decided for any valuation \mathfrak{h} . Consequently, ψ_s defined above is a constrained formula that can be effectively computed.

Let v be a concretisation such that $\mathfrak{M}, s \models v(\phi)$. There exists $\mathfrak{h}_0 : \{\varphi_1, \dots, \varphi_z\} \rightarrow \{\perp, \top\}$ such that $\mathfrak{M}, s \models \mathfrak{h}_0(\phi)$ and for all $j \in [1, z]$, we have $\mathfrak{h}_0(\varphi_j) = \top$ iff $\mathfrak{M}, s \models v(\varphi_j)$. By Lemma 11, for all $j \in [1, z]$, we have $\mathfrak{h}_0(\varphi_j) = \top$ iff $v \models \hat{\varphi}_j^s$. So,

$$v \models \left(\bigwedge_{j \text{ s.t. } \mathfrak{h}_0(\varphi_j) = \top} \hat{\varphi}_j^s \right) \wedge \left(\bigwedge_{j \text{ s.t. } \mathfrak{h}_0(\varphi_j) = \perp} \neg \hat{\varphi}_j^s \right),$$

which entails $v \models \psi_s$.

Conversely, assuming that $v \models \psi_s$ for some concretisation v , there is a map \mathfrak{h} such that $v \models (\bigwedge_{j \text{ s.t. } \mathfrak{h}(\varphi_j)=\top} \hat{\varphi}_j^s) \wedge (\bigwedge_{j \text{ s.t. } \mathfrak{h}(\varphi_j)=\perp} \neg \hat{\varphi}_j^s)$ and $\mathfrak{M}, s \models \mathfrak{h}(\phi)$. Again by Lemma 11, and by Boolean reasoning we obtain $\mathfrak{M}, s \models v(\phi)$. Consequently, the formula ψ_s can be computed and it is a symbolic representation for all the maps v such that $\mathfrak{M}, s \models v(\phi)$. \square

6 Concluding Remarks

We have related model checking problems for resource-bounded logics and decision problems for alternating VASS, such as state reachability, non-termination, and, more generally, parity game problems. While the existence of such relationships is perhaps not surprising, we have been able to obtain several new complexity and decidability results including:

1. The model checking problem for the logic $\text{RB}\pm\text{ATL}$ introduced in [3, 4] is 2EXPTIME -complete. No complexity upper bound was previously known.
2. We have introduced a new logic $\text{RB}\pm\text{ATL}^*$ that extends $\text{RB}\pm\text{ATL}$ (as ATL^* extends ATL), and we have shown that its model checking problem is decidable. For the parameterised version $\text{ParRB}\pm\text{ATL}^*$, given \mathfrak{M}, s and ϕ in $\text{ParRB}\pm\text{ATL}^*$, we have explained how we can synthesise a formula ψ such that $\mathfrak{M}, s \models v(\phi)$ iff $v \models \psi$ for all interpretations v for the resource parameters. Moreover, ψ is a Boolean combination of atomic constraints of the form $x \geq k$ and $x = \omega$.
3. The model checking problem for RBTL^* introduced in [12] is EXSPACE -complete. The decidability of the model checking problem for RBTL^* and the complexity upper bound for RBTL were not previously known.

We have been also able to provide complexity results for fragments and variants of these resource-bounded logics, and we believe that the simple framework we have proposed may be used to obtain further results for new resource-bounded logics. However this is future work.

References

- [1] P.A. Abdulla, R. Mayr, A. Sangnier, and J. Sproston. Solving Parity Games on Integer Vectors. In *CONCUR'13*, volume 8052 of *Lecture Notes in Computer Science*, pages 106–120. Springer, 2013.
- [2] N. Alechina, N. Bulling, B. Logan, and H.N. Nguyen. On the boundary of (un)decidability: Decidable model-checking for a fragment of resource agent logic. In *IJCAI'15*, pages 1494–1501. AAAI Press, 2015.
- [3] N. Alechina, B. Logan, H.N. Nguyen, and F. Raimondi. Decidable model-checking for a resource logic with production of resources. In *ECAI'14*, pages 9–14, 2014.
- [4] N. Alechina, B. Logan, H.N. Nguyen, and F. Raimondi. Technical report: Model-checking for resource-bounded ATL with production and consumption of resources. *CoRR*, abs/1504.06766, 2015.
- [5] R. Alur, Th. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49(5):672–713, 2002.
- [6] C. Areces, P. Blackburn, and M. Marx. Hybrid logics: characterization, interpolation and complexity. *The Journal of Symbolic Logic*, 66(3):977–1010, 2001.
- [7] B. Bérard, S. Haddad, M. Sassolas, and N. Sznajder. Concurrent games on VASS with inhibition. In *CONCUR'12*, volume 7454 of *Lecture Notes in Computer Science*, pages 39–52. Springer, 2012.
- [8] M. Blockeet and S. Schmitz. Model-checking coverability graphs of vector addition systems. In *MFCS'11*, volume 6907 of *Lecture Notes in Computer Science*, pages 108–119. Springer, 2011.
- [9] M. Blondin, A. Finkel, S. Göller, C. Haase, and P. McKenzie. Reachability in two-dimensional vector addition systems with states is PSPACE-complete. In *LICS'15*, pages 32–43. ACM Press, 2015.
- [10] A. Bouajjani, M. Bozga, P. Habermehl, R. Iosif, P. Moro, and T. Vojnar. Programs with lists are counter automata. In *CAV'06*, volume 4144 of *Lecture Notes in Computer Science*, pages 517–531. Springer, 2006.

- [11] T. Brázdil, P. Jančar, and A. Kucera. Reachability games on extended vector addition systems with states. In *ICALP'10*, volume 6199 of *Lecture Notes in Computer Science*, pages 478–489. Springer, 2010.
- [12] N. Bulling and B. Farwer. Expressing properties of resource-bounded systems: The logics RBTL* and RBTL. In *CLIMA X*, volume 6214 of *Lecture Notes in Computer Science*, pages 22–45. Springer, 2009.
- [13] N. Bulling and B. Farwer. On the (Un-)Decidability of Model-Checking Resource-Bounded Agents. In *ECAI'10*, pages 567–572, 2010.
- [14] N. Bulling and H.N. Nguyen. Model checking resource bounded systems with shared resources via alternating Büchi pushdown systems. In *PRIMA'15*, volume 9387 of *Lecture Notes in Computer Science*, pages 640–649, 2015.
- [15] T. Chen, F. Song, and Z. Wu. Global model checking on pushdown multi-agent systems. In *AAAI'16*, pages 2459–2465. AAAI Press, 2016.
- [16] J.B. Courtois and S. Schmitz. Alternating vector addition systems with states. In *MFCS'14*, volume 8634 of *Lecture Notes in Computer Science*, pages 220–231. Springer, 2014.
- [17] Ph. Darondeau, S. Demri, R. Meyer, and Ch. Morvan. Petri net reachability graphs: Decidability status of FO properties. *Logical Methods in Computer Science*, 8(4), 2012.
- [18] S. Demri. On selective unboundedness of VASS. *Journal of Computer and System Sciences*, 79(5):689–713, 2013.
- [19] S. Demri, V. Goranko, and M. Lange. *Temporal Logics in Computer Science*. Cambridge University Press, 2016.
- [20] S. Demri, M. Jurdziński, O. Lachish, and R. Lazić. The covering and boundedness problems for branching vector addition systems. *Journal of Computer and System Sciences*, 79(1):23–38, 2013.
- [21] L.E. Dickson. Finiteness of the odd perfect and primitive abundant numbers with n distinct prime factors. *American Journal of Mathematics*, pages 413–422, 1913.
- [22] A. Emerson. Temporal and modal logic. In *Handbook of Theoretical Computer Science*, pages 996–1072. Elsevier, 1990.
- [23] J. Esparza. On the decidability of model checking for several μ -calculi and Petri nets. In *ICALP'94*, volume 787 of *Lecture Notes in Computer Science*, pages 115–129. Springer, 1994.
- [24] J. Esparza. Decidability and complexity of Petri net problems — an introduction. In *Advances in Petri Nets 1998*, volume 1491 of *Lecture Notes in Computer Science*, pages 374–428. Springer, 1998.
- [25] S. Göller and M. Lohrey. Branching-time model checking of one-counter processes and timed automata. *SIAM Journal of Computing*, 42(3):884–923, 2013.
- [26] C. Haase. *On the Complexity of Model Checking Counter Automata*. PhD thesis, University of Oxford, 2012.
- [27] P. Habermehl. On the complexity of the linear-time μ -calculus for Petri nets. volume 1248 of *Lecture Notes in Computer Science*, pages 102–116. Springer, 1997.
- [28] R.R. Howell and L.E. Rosier. Problems concerning fairness and temporal logic for conflict-free Petri nets. *Theoretical Computer Science*, 64:305–329, 1989.
- [29] P. Jančar. Decidability of a temporal logic problem for Petri nets. *Theoretical Computer Science*, 74(1):71–93, 1990.
- [30] P. Jančar. On reachability-related games on vector addition systems with states. In *RP'15*, volume 9328 of *Lecture Notes in Computer Science*, pages 50–62. Springer, 2015.
- [31] L. Juhl, K.G. Larsen, and J.-F. Raskin. Optimal bounds for multiweighted and parametrised energy games. In *Theories of Programming and Formal Methods - Essays Dedicated to Jifeng He on the Occasion of His 70th Birthday*, volume 8051 of *Lecture Notes in Computer Science*, pages 244–255. Springer, 2013.
- [32] M. Jurdziński, R. Lazić, and S. Schmitz. Fixed-dimensional energy games are in pseudo-polynomial time. In *ICALP'15*, volume 9135 of *Lecture Notes in Computer Science*, pages 260–272. Springer, 2015.
- [33] R.M. Karp and R.E. Miller. Parallel program schemata. *Journal of Computer and System Sciences*, 3(2):147–195, 1969.
- [34] F. Laroussinie, N. Markey, and G. Oreiby. On the expressiveness and complexity of ATL. *Logical Methods in Computer Science*, 4(2), 2008.
- [35] R. Lazić and S. Schmitz. Private communication, November 2016.
- [36] J. Leroux and S. Schmitz. Demystifying reachability in vector addition systems. In *LICS'15*, pages 56–67. IEEE, 2015.
- [37] R.J. Lipton. The reachability problem requires exponential space. Technical Report 62, Department of Computer Science, Yale University, 1976.
- [38] D. Della Monica, M. Napoli, and M. Parente. On a logic for coalitional games with priced-resource agents. *ENTCS*, 278:215–228, 2011.
- [39] C. Rackoff. The covering and boundedness problems for vector addition systems. *Theoretical Computer Science*, 6(2):223–231, 1978.
- [40] J.F. Raskin, M. Samuelides, and L. Van Begin. Games for counting abstractions. *ENTCS*, 128(6):69–85, 2005.
- [41] S. Schewe. Tighter bounds for the determinisation of Büchi automata. In *FOSSACS'09*, volume 5504 of *Lecture Notes in Computer Science*, pages 167–181. Springer, 2009.

- [42] Ph. Schnoebelen. The complexity of temporal logic model checking. In *AIML'02*, pages 437–459. King's College Publication, 2003.
- [43] O. Serre. Parity games played on transition graphs of one-counter processes. In *FOSSACS'06*, volume 3921 of *Lecture Notes in Computer Science*, pages 337–351. Springer, 2006.
- [44] M. Vardi and P. Wolper. Reasoning about infinite computations. *Information and Computation*, 115:1–37, 1994.
- [45] K.N. Verma and J. Goubault-Larrecq. Karp-Miller Trees for a Branching Extension of VASS. *Discrete Mathematics and Theoretical Computer Science*, 7:217–230, 2005.
- [46] S. Vester. On the complexity of model-checking branching and alternating-time temporal logics in one-counter systems. In *ATVA'15*, volume 9364 of *Lecture Notes in Computer Science*, pages 361–377. Springer, 2015.