Solving Parity Games on Integer Vectors ⋆

Parosh Aziz Abdulla[1⋆⋆], Richard Mayr[2⋆⋆⋆], Arnaud Sangnier[3], and Jeremy Sproston[4†]

[1] Uppsala University, Sweden
[2] University of Edinburgh, UK
[3] LIAFA, Univ Paris Diderot, Sorbonne Paris Cité, CNRS, France
[4] University of Turin, Italy

Abstract. We consider parity games on infinite graphs where configurations are represented by control-states and integer vectors. This framework subsumes two classic game problems: parity games on vector addition systems with states (VASS) and multidimensional energy parity games. We show that the multidimensional energy parity game problem is inter-reducible with a subclass of single-sided parity games on VASS where just one player can modify the integer counters and the opponent can only change control-states. Our main result is that the minimal elements of the upward-closed winning set of these single-sided parity games on VASS are computable. This implies that the Pareto frontier of the minimal initial credit needed to win multidimensional energy parity games is also computable, solving an open question from the literature. Moreover, our main result implies the decidability of weak simulation preorder/equivalence between finite-state systems and VASS, and the decidability of model checking VASS with a large fragment of the modal μ-calculus.

1 Introduction

In this paper, we consider *integer games*: two-player turn-based games where a color (natural number) is associated to each state, and where the transitions allow incrementing and decrementing the values of a finite set of integer-valued counters by constants. We refer to the players as Player 0 and Player 1.

We consider the classical parity condition, together with two different semantics for integer games: the *energy semantics* and the VASS *semantics*. The former corresponds to *multidimensional energy parity games* [7], and the latter to parity games on VASS (a model essentially equivalent to Petri nets [8]). In energy parity games, the winning objective for Player 0 combines a qualitative property, the classical *parity condition*, with a quantitative property, namely the *energy condition*. The latter means that the values of all counters stay above a finite threshold during the entire run of the game. In VASS

⋆ Technical Report EDI-INF-RR-1417 of the School of Informatics at the University of Edinburgh, UK. (http://www.inf.ed.ac.uk/publications/report/). Full version (including proofs) of material presented at CONCUR 2013 (Buenos Aires, Argentina). arXiv.org - CC BY 3.0.
⋆⋆ Supported by Uppsala Programming for Multicore Architectures Research Center (UpMarc).
⋆⋆⋆ Supported by Royal Society grant IE110996.
† Supported by the project AMALFI (University of Turin/Compagnia di San Paolo) and the MIUR-PRIN project CINA.

parity games, the counter values are restricted to *natural numbers*, and in particular any transition that may decrease the value of a counter below zero is disabled (unlike in energy games where such a transition would be immediately winning for Player 1). So for VASS games, the objective consists only of a parity condition, since the energy condition is trivially satisfied.

We formulate and solve our problems using a generalized notion of game configurations, namely *partial configurations*, in which only a subset C of the counters may be defined. A partial configuration γ denotes a (possibly infinite) set of concrete configurations that are called *instantiations* of γ. A configuration γ' is an instantiation of γ if γ' agrees with γ on the values of the counters in C while the values of counters outside C can be chosen freely in γ'. We declare a partial configuration to be *winning* (for Player 0) if it has an instantiation that is winning. For each decision problem and each set of counters C, we will consider the *C-version* of the problem where we reason about configurations in which the counters in C are defined.

Previous Work. Two special cases of the general C-version are the *abstract* version in which no counters are defined, and the *concrete* version in which all counters are defined. In the energy semantics, the abstract version corresponds to the *unknown initial credit problem* for multidimensional energy parity games, which is coNP-complete [6, 7]. The concrete version corresponds to the *fixed initial credit problem*. For energy games without the parity condition, the fixed initial credit problem was solved in [4] (although it does not explicitly mention energy games but instead formulates the problem as a *zero-reachability* objective for Player 1). It follows from [4] that the fixed initial credit problem for d-dimensional energy games can be solved in d-EXPTIME (resp. $(d-1)$-EXPTIME for offsets encoded in unary) and even the upward-closed winning sets can be computed. An EXPSPACE lower bound is derived by a reduction from Petri net coverability. The subcase of one-dimensional energy parity games was considered in [5], where both the unknown and fixed initial credit problems are decidable, and the winning sets (i.e., the minimal required initial energy) can be computed. The assumption of having just one dimension is an important restriction that significantly simplifies the problem. This case is solved using an algorithm which is a generalization of the classical algorithms of McNaughton [13] and Zielonka [16].

However, for general multidimensional energy parity games, computing the winning sets was an open problem, mentioned, e.g., in [6].

In contrast, under the VASS semantics, all these integer game problems are shown to be undecidable for dimensions ≥ 2 in [2], even for simple safety/coverability objectives. (The one-dimensional case is a special case of parity games on one-counter machines, which is PSPACE-complete). A special subcase are *single-sided* VASS games, where just Player 0 can modify counters while Player 1 can only change control-states. This restriction makes the winning set for Player 0 upward-closed, unlike in general VASS games. The paper [14] shows decidability of coverability objectives for single-sided VASS games, using a standard backward fixpoint computation.

Our Contribution. First we show how instances of the single-sided VASS parity game can be reduced to the multidimensional energy parity game, and vice-versa. I.e., energy games correspond to the single-sided subcase of VASS games. Notice that, since parity conditions are closed under complement, it is merely a convention that Player 0 (and not Player 1) is the one that can change the counters.

Our main result is the decidability of single-sided VASS parity games for general partial configurations, and thus in particular for the concrete and abstract versions described above. The winning set for Player 0 is upward-closed (wrt. the natural multiset ordering on configurations), and it can be computed (i.e., its finitely many minimal elements). Our algorithm uses the Valk-Jantzen construction [15] and a technique similar to Karp-Miller graphs, and finally reduces the problem to instances of the abstract parity problem under the energy semantics, i.e., to the unknown initial credit problem in multidimensional energy parity games, which is decidable by [7].

From the above connection between single-sided VASS parity games and multidimensional energy parity games, it follows that the winning sets of multidimensional energy parity games are also computable. I.e., one can compute the *Pareto frontier* of the minimal initial energy credit vectors required to win the energy parity game. This solves the problem left open in [6, 7].

Our results imply further decidability results in the following two areas: semantic equivalence checking and model-checking. Weak simulation preorder between a finite-state system and a general VASS can be reduced to a parity game on a single-sided VASS, and is therefore decidable. Combined with the previously known decidability of the reverse direction [3], this implies decidability of weak simulation equivalence. This contrasts with the undecidability of weak *bisimulation* equivalence between VASS and finite-state systems [11]. The model-checking problem for VASS is decidable for many linear-time temporal logics [10], but undecidable even for very restricted branching-time logics [8]. We show the decidability of model-checking for a restricted class of VASS with a large fragment of the modal μ-calculus. Namely we consider VASS where some states do not perform any updates on the counters, and these states are used to guard the for-all-successors modal operators in this fragment of the μ-calculus, allowing us to reduce the model-checking problem to a parity game on single-sided VASS.

2 Integer Games

Preliminaries. We use $\mathbb{N}$ and $\mathbb{Z}$ to denote the sets of natural numbers (including 0) and integers respectively. For a set A, we define $|A|$ to be the cardinality of A. For a function $f : A \mapsto B$ from a set A to a set B, we use $f[a \leftarrow b]$ to denote the function f' such that $f(a) = b$ and $f'(a') = f(a')$ if $a' \neq a$. If f is partial, then $f(a) = \perp$ means that f is undefined for a. In particular $f[a \leftarrow \perp]$ makes the value of a undefined. We define $dom(f) := \{a | \, f(a) \neq \perp\}$.

Model. We assume a finite set $\mathcal{C}$ of *counters*. An *integer game* is a tuple $\mathcal{G} = \langle Q, T, \kappa \rangle$ where Q is a finite set of *states*, T is a finite set of *transitions*, and $\kappa : Q \mapsto \{0, 1, 2, \ldots, k\}$ is a *coloring* function that assigns to each $q \in Q$ a natural number in the interval $[0..k]$ for some pre-defined k. The set Q is partitioned into two sets Q_0 (states of Player 0) and Q_1 (states of Player 1). A transition $t \in T$ is a triple $\langle q_1, op, q_2 \rangle$ where $q_1, q_2 \in Q$ are states and op is an operation of one of the following three forms (where $c \in \mathcal{C}$ is a counter): (i) c++ increments the value of c by one; (ii) c- - decrements the value of c by one; (iii) *nop* does not change the value of any counter. We define $\texttt{source}(t) = q_1$,

$\texttt{target}(t) = q_2$, and $\texttt{op}(t) = op$. We say that $\mathcal{G}$ is *single-sided* in case $op = nop$ for all transitions $t \in T$ with $\texttt{source}(t) \in Q_1$. In other words, in a single-sided game, Player 1 is not allowed to changes the values of the counters, but only the state.

Partial Configurations. A *partial counter valuation* $\vartheta : \mathcal{C} \mapsto \mathbb{Z}$ is a partial function from the set of counters to $\mathbb{Z}$. We also write $\vartheta(c) = \perp$ if $c \notin dom(\vartheta)$. A *partial configuration* γ is a pair $\langle q, \vartheta \rangle$ where $q \in Q$ is a state and ϑ is a partial counter valuation. We will also consider *nonnegative partial configurations*, where the partial counter valuation takes values in $\mathbb{N}$ instead of $\mathbb{Z}$. We define $\texttt{state}(\gamma) := q$, $\texttt{val}(\gamma) := \vartheta$, and $\kappa(\gamma) := \kappa(\texttt{state}(\gamma))$. We generalize assignments from counter valuations to configurations by defining $\langle q, \vartheta \rangle [c \leftarrow x] = \langle q, \vartheta[c \leftarrow x] \rangle$. Similarly, for a configuration γ and $c \in \mathcal{C}$ we let $\gamma(c) := \texttt{val}(\gamma)(c)$, $dom(\gamma) := dom(\texttt{val}(\gamma))$ and $|\gamma| := |dom(\gamma)|$. For a set of counters $C \subseteq \mathcal{C}$, we define $\Theta^C := \{\gamma |\ dom(\gamma) = C\}$, i.e., it is the set of configurations in which the defined counters are exactly those in C. We use Γ^C to denote the restriction of Θ^C to nonnegative partial configurations. We partition Θ^C into two sets Θ^C_0 (configurations belonging to Player 0) and Θ^C_1 (configurations belonging to Player 1), such that $\gamma \in \Theta^C_i$ iff $dom(\gamma) = C$ and $\texttt{state}(\gamma) \in Q_i$ for $i \in \{0,1\}$. A configuration is *concrete* if $dom(\gamma) = \mathcal{C}$, i.e., $\gamma \in \Theta^{\mathcal{C}}$ (the counter valuation $\texttt{val}(\gamma)$ is defined for all counters); and it is *abstract* if $dom(\gamma) = \emptyset$, i.e., $\gamma \in \Theta^{\emptyset}$ (the counter valuation $\texttt{val}(\gamma)$ is not defined for any counter). In the sequel, we occasionally write Θ instead of $\Theta^{\mathcal{C}}$, and Θ_i instead of $\Theta^{\mathcal{C}}_i$ for $i \in \{0,1\}$. The same notations are defined over nonnegative partial configurations with Γ, and Γ^C_i and Γ_i for $i \in \{0,1\}$. For a nonnegative partial configuration $\gamma = \langle q, \vartheta \rangle \in \Gamma$, and set of counters $C \subseteq \mathcal{C}$ we define the restriction of γ to C by $\gamma' = \gamma | C = \langle q', \vartheta' \rangle$ where $q' = q$ and $\vartheta'(c) = \vartheta(c)$ if $c \in C$ and $\vartheta'(c) = \perp$ otherwise.

Energy Semantics. Under the energy semantics, an integer game induces a transition relation $\longrightarrow_{\mathcal{E}}$ on the set of partial configurations as follows. For partial configurations $\gamma_1 = \langle q_1, \vartheta_1 \rangle$, $\gamma_2 = \langle q_2, \vartheta_2 \rangle$, and a transition $t = \langle q_1, op, q_2 \rangle \in T$, we have $\gamma_1 \xrightarrow{t}_{\mathcal{E}} \gamma_2$ if one of the following three cases is satisfied: (i) $op = c{+}{+}$ and either both $\vartheta_1(c) = \perp$ and $\vartheta_2(c) = \perp$ or $\vartheta_1(c) \neq \perp$, $\vartheta_2(c) \neq \perp$ and $\vartheta_2 = \vartheta_1[c \leftarrow \vartheta_1(c) + 1]$; (ii) $op = c{-}{-}$, and either both $\vartheta_1(c) = \perp$ and $\vartheta_2(c) = \perp$ or $\vartheta_1(c) \neq \perp$, $\vartheta_2(c) \neq \perp$ and $\vartheta_2 = \vartheta_1[c \leftarrow \vartheta_1(c) - 1]$; (iii) $op = nop$ and $\vartheta_2 = \vartheta_1$. Hence we apply the operation of the transition only if the relevant counter value is defined (otherwise, the counter remains undefined). Notice that, for a partial configuration γ_1 and a transition t, there is at most one γ_2 with $\gamma_1 \xrightarrow{t}_{\mathcal{E}} \gamma_2$. If such a γ_2 exists, we define $t(\gamma_1) := \gamma_2$; otherwise we define $t(\gamma_1) := \perp$. We say that t is *enabled* at γ if $t(\gamma) \neq \perp$. We observe that, in the case of energy semantics, t is not enabled only if $\texttt{state}(\gamma) \neq \texttt{source}(t)$.

VASS Semantics. The difference between the energy and VASS semantics is that counters in the case of VASS range over the natural numbers (rather than the integers), i.e. the VASS semantics will be interpreted over nonnegative partial configurations. Thus, the transition relation $\longrightarrow_{\mathcal{V}}$ induced by an integer game $\mathcal{G} = \langle Q, T, \kappa \rangle$ under the VASS semantics differs from the one induced by the energy semantics in the sense that counters are not allowed to assume negative values. Hence $\longrightarrow_{\mathcal{V}}$ is the restriction of $\longrightarrow_{\mathcal{E}}$ to nonnegative partial configurations. Here, a transition $t = \langle q_1, c{-}{-}, q_2 \rangle \in T$ is enabled from $\gamma_1 = \langle q_1, \vartheta_1 \rangle$ only if $\vartheta_1(c) > 0$ or $\vartheta_1(c) = \perp$. We assume without restriction that at least one transition is enabled from each partial configuration (i.e., there are no deadlocks) in the VASS semantics (and hence also in the energy semantics). Below, we use $\mathtt{sem} \in \{\mathcal{E}, \mathcal{V}\}$ to distinguish the energy and VASS semantics.

Runs. A *run* ρ in semantics $\mathtt{sem}$ is an infinite sequence $\gamma_0 \xrightarrow{t_1}_{\mathtt{sem}} \gamma_1 \xrightarrow{t_2}_{\mathtt{sem}} \cdots$ of transitions between concrete configurations. A *path* π in $\mathtt{sem}$ is a finite sequence $\gamma_0 \xrightarrow{t_1}_{\mathtt{sem}} \gamma_1 \xrightarrow{t_2}_{\mathtt{sem}} \cdots \gamma_n$ of transitions between concrete configurations. We say that ρ (resp. π) is a γ-run (resp. γ-path) if $\gamma_0 = \gamma$. We define $\rho(i) := \gamma_i$ and $\pi(i) := \gamma_i$. We assume familiarity with the logic LTL. For an LTL formula ϕ we write $\rho \models_{\mathcal{G}} \phi$ to denote that the run ρ in $\mathcal{G}$ satisfies ϕ. For instance, given a set β of concrete configurations, we write $\rho \models_{\mathcal{G}} \Diamond\beta$ to denote that there is an i with $\gamma_i \in \beta$ (i.e., a member of β eventually occurs along ρ); and write $\rho \models_{\mathcal{G}} \Box\Diamond\beta$ to denote that there are infinitely many i with $\gamma_i \in \beta$ (i.e., members of β occur infinitely often along ρ).

Strategies. A *strategy* of Player $i \in \{0,1\}$ in $\mathtt{sem}$ (or simply an i-strategy in $\mathtt{sem}$) σ_i is a mapping that assigns to each path $\pi = \gamma_0 \xrightarrow{t_1}_{\mathtt{sem}} \gamma_1 \xrightarrow{t_2}_{\mathtt{sem}} \cdots \gamma_n$ with $\mathtt{state}(\gamma_n) \in Q_i$, a transition $t = \sigma_i(\pi)$ with $t(\gamma_n) \neq \perp$ in $\mathtt{sem}$. We use $\Sigma_i^{\mathtt{sem}}$ to denote the sets of i-strategies in $\mathtt{sem}$. Given a concrete configuration γ, $\sigma_0 \in \Sigma_0^{\mathtt{sem}}$, and $\sigma_1 \in \Sigma_1^{\mathtt{sem}}$, we define $\mathtt{run}(\gamma, \sigma_0, \sigma_1)$ to be the unique run $\gamma_0 \xrightarrow{t_1}_{\mathtt{sem}} \gamma_1 \xrightarrow{t_2}_{\mathtt{sem}} \cdots$ such that (i) $\gamma_0 = \gamma$, (ii) $t_{i+1} = \sigma_0(\gamma_0 \xrightarrow{t_1}_{\mathtt{sem}} \gamma_1 \xrightarrow{t_2}_{\mathtt{sem}} \cdots \gamma_i)$ if $\mathtt{state}(\gamma_i) \in Q_0$, and (iii) $t_{i+1} = \sigma_1(\gamma_0 \xrightarrow{t_1}_{\mathtt{sem}} \gamma_1 \xrightarrow{t_2}_{\mathtt{sem}} \cdots \gamma_i)$ if $\mathtt{state}(\gamma_i) \in Q_1$. For $\sigma_i \in \Sigma_i^{\mathtt{sem}}$, we write $[i, \sigma_i, \mathtt{sem}] : \gamma \models_{\mathcal{G}} \phi$ to denote that $\mathtt{run}(\gamma, \sigma_i, \sigma_{1-i}) \models_{\mathcal{G}} \phi$ for all $\sigma_{1-i} \in \Sigma_{1-i}^{\mathtt{sem}}$. In other words, Player i has a *winning strategy*, namely σ_i, which ensures that ϕ will be satisfied regardless of the strategy chosen by Player $1-i$. We write $[i, \mathtt{sem}] : \gamma \models_{\mathcal{G}} \phi$ to denote that $[i, \sigma_i, \mathtt{sem}] : \gamma \models_{\mathcal{G}} \phi$ for some $\sigma_i \in \Sigma_i^{\mathtt{sem}}$.

Instantiations. Two nonnegative partial configurations γ_1, γ_2 are said to be *disjoint* if (i) $\mathtt{state}(\gamma_1) = \mathtt{state}(\gamma_2)$, and (ii) $dom(\gamma_1) \cap dom(\gamma_2) = \emptyset$ (notice that we require the states to be equal). For a set of counters $C \subseteq \mathcal{C}$, and disjoint partial configurations γ_1, γ_2, we say that γ_2 is a *C-complement* of γ_1 if $dom(\gamma_1) \cup dom(\gamma_2) = C$, i.e., $dom(\gamma_1)$ and $dom(\gamma_2)$ form a partitioning of the set C. If γ_1 and γ_2 are disjoint then we define $\gamma_1 \oplus \gamma_2$ to be the nonnegative partial configuration $\gamma := \langle q, \vartheta \rangle$ such that $q := \mathtt{state}(\gamma_1) = \mathtt{state}(\gamma_2)$, $\vartheta(c) := \mathtt{val}(\gamma_1)(c)$ if $\mathtt{val}(\gamma_1)(c) \neq \perp$, $\vartheta(c) := \mathtt{val}(\gamma_2)(c)$ if $\mathtt{val}(\gamma_2)(c) \neq \perp$, and $\vartheta(c) := \perp$ if both $\mathtt{val}(\gamma_1)(c) = \perp$ and $\mathtt{val}(\gamma_2)(c) = \perp$. In such a case, we say that γ is a *C-instantiation* of γ_1. For a nonnegative partial configuration γ we write $[\![\gamma]\!]_C$ to denote the set of C-instantiations of γ. We will consider the special case where $C = \mathcal{C}$. In particular, we say that γ_2 is a *complement* of γ_1 if γ_2 is a $\mathcal{C}$-complement of γ_1, i.e., $\mathtt{state}(\gamma_2) = \mathtt{state}(\gamma_1)$ and $dom(\gamma_1) = \mathcal{C} - dom(\gamma_2)$. We use $\bar{\gamma}$ to denote the set of complements of γ. If $\gamma_2 \in \overline{\gamma_1}$, we say that $\gamma = \gamma_1 \oplus \gamma_2$ is an *instantiation* of γ_1. Notice that γ in such a case is concrete. For a nonnegative partial configuration γ we write $[\![\gamma]\!]$ to denote the set of instantiations of γ. We observe that $[\![\gamma]\!] = [\![\gamma]\!]_{\mathcal{C}}$ and that $[\![\gamma]\!] = \{\gamma\}$ for any concrete nonnegative configuration γ.

Ordering. For nonnegative partial configurations γ_1, γ_2, we write $\gamma_1 \sim \gamma_2$ if $\mathtt{state}(\gamma_1) = \mathtt{state}(\gamma_2)$ and $dom(\gamma_1) = dom(\gamma_2)$. We write $\gamma_1 \sqsubseteq \gamma_2$ if $\mathtt{state}(\gamma_1) = \mathtt{state}(\gamma_2)$ and $dom(\gamma_1) \subseteq dom(\gamma_2)$. For nonnegative partial configurations $\gamma_1 \sim \gamma_2$, we write

$\gamma_1 \preceq \gamma_2$ to denote that $\texttt{state}(\gamma_1) = \texttt{state}(\gamma_2)$ and $\texttt{val}(\gamma_1)(c) \leq \texttt{val}(\gamma_2)(c)$ for all $c \in dom(\gamma_1) = dom(\gamma_2)$. For a nonnegative partial configuration γ, we define $\gamma\uparrow := \{\gamma' |\, \gamma \preceq \gamma'\}$ to be the upward closure of γ, and define $\gamma\downarrow := \{\gamma' |\, \gamma' \preceq \gamma\}$ to be the downward closure of γ. Notice that $\gamma\uparrow = \gamma\downarrow = \{\gamma\}$ for any abstract configuration γ. For a set $\beta \subseteq \Gamma^C$ of nonnegative partial configurations, let $\beta\uparrow := \cup_{\gamma\in\beta}\gamma\uparrow$. We say that β is upward-closed if $\beta\uparrow = \beta$. For an upward-closed set $\beta \subseteq \Gamma^C$, we use $min(\beta)$ to denote the (by Dickson's Lemma unique and finite) set of minimal elements of β.

Winning Sets of Partial Configurations. For a nonnegative partial configuration γ, we write $[i,\texttt{sem}] : \gamma \models_{\mathcal{G}} \phi$ to denote that $\exists\gamma' \in [\![\gamma]\!] . [i,\texttt{sem}] : \gamma' \models_{\mathcal{G}} \phi$, i.e., Player i is winning from some instantiation γ' of γ. For a set $C \subseteq \mathcal{C}$ of counters, we define $\mathcal{W}[\mathcal{G},\texttt{sem},i,C](\phi) := \left\{\gamma \in \Gamma^C |\, [\texttt{sem},i] : \gamma \models_{\mathcal{G}} \phi\right\}$. If $\mathcal{W}[\mathcal{G},\texttt{sem},i,C](\phi)$ is upward-closed, we define the *Pareto frontier* as $\texttt{Pareto}[\mathcal{G},\texttt{sem},i,C](\phi) := min(\mathcal{W}[\mathcal{G},\texttt{sem},i,C](\phi))$.

Properties. We show some useful properties of the ordering on nonnegative partial configurations. Note that for nonnegative partial configurations, we will not make distinctions between the energy semantics and the VASS semantics; this is due to the fact that in nonnegative partial configurations and in their instantiations we only consider positive values for the counters. For the energy semantics, as we shall see, this will not be a problem since we will consider winning runs where the counter never goes below 0. We now show *monotonicity* and (under some conditions) "reverse monotonicity" of the transition relation wrt. $\preceq$. We write $\gamma_1 \longrightarrow_{\texttt{sem}} \gamma_2$ if there exists t such that $\gamma_1 \xrightarrow{t}_{\texttt{sem}} \gamma_2$.

Lemma 1. *Let γ_1, γ_2, and γ_3 be nonnegative partial configurations. If (i) $\gamma_1 \longrightarrow_{\mathcal{V}} \gamma_2$, and (ii) $\gamma_1 \preceq \gamma_3$, then there is a γ_4 such that $\gamma_3 \longrightarrow_{\mathcal{V}} \gamma_4$ and $\gamma_2 \preceq \gamma_4$. Furthermore, if (i) $\gamma_1 \longrightarrow_{\mathcal{V}} \gamma_2$, and (ii) $\gamma_3 \preceq \gamma_1$, and (iii) $\mathcal{G}$ is single-sided and (iv) $\gamma_1 \in \Gamma_1$, then there is a γ_4 such that $\gamma_3 \longrightarrow_{\mathcal{V}} \gamma_4$ and $\gamma_4 \preceq \gamma_2$.*

We consider a version of the *Valk-Jantzen* lemma [15], expressed in our terminology.

Lemma 2. *[15] Let $C \subseteq \mathcal{C}$ and let $U \subseteq \Gamma^C$ be upward-closed. Then, $min(U)$ is computable if and only if, for any nonnegative partial configuration γ with $dom(\gamma) \subseteq C$, we can decide whether $[\![\gamma]\!]_C \cap U \neq \emptyset$.*

3 Game Problems

Here we consider the parity winning condition for the integer games defined in the previous section. First we establish a correspondence between the VASS semantics when the underlying integer game is single-sided, and the energy semantics in the general case. We will show how instances of the single-sided VASS parity game can be reduced to the energy parity game, and vice-versa. Figure 1 depicts a summary of our results. For either semantics, an instance of the problem consists of an integer game $\mathcal{G}$ and a partial configuration γ. For a given set of counters $C \subseteq \mathcal{C}$, we will consider the *C-version* of the problem where we assume that $dom(\gamma) = C$. In particular, we will consider two special cases: (i) the *abstract* version in which we assume that γ is abstract (i.e., $dom(\gamma) = \emptyset$), and (ii) the *concrete* version in which we assume that γ is concrete (i.e., $dom(\gamma) = \mathcal{C}$).

The abstract version of a problem corresponds to the *unknown initial credit problem* [6, 7], while the concrete one corresponds to deciding if a given initial credit is sufficient or, more generally, computing the Pareto frontier (left open in [6, 7]).

Winning Conditions. Assume an integer game $\mathcal{G} = \langle Q, T, \kappa \rangle$ where $\kappa : Q \mapsto \{0, 1, 2, \ldots, k\}$. For a partial configuration γ and $i : 0 \leq i \leq k$, the relation $\gamma \models_{\mathcal{G}} (\texttt{color} = i)$ holds if $\kappa(\texttt{state}(\gamma)) = i$. The formula simply checks the color of the state of γ. The formula $\gamma \models_{\mathcal{G}} \overline{\texttt{neg}}$ holds if $\texttt{val}(\gamma)(c) \geq 0$ for all $c \in dom(\gamma)$. The formula states that the values of all counters are nonnegative in γ. For $i : 0 \leq i \leq k$, the predicate $even(i)$ holds if i is even. Define the path formula $\texttt{Parity} := \bigvee_{(0 \leq i \leq k) \wedge even(i)} \left((\Box\Diamond(\texttt{color} = i)) \wedge \left(\bigwedge_{i<j\leq k} \Diamond\Box\neg(\texttt{color} = j) \right) \right)$. The formula states that the highest color that appears infinitely often along the path is even.

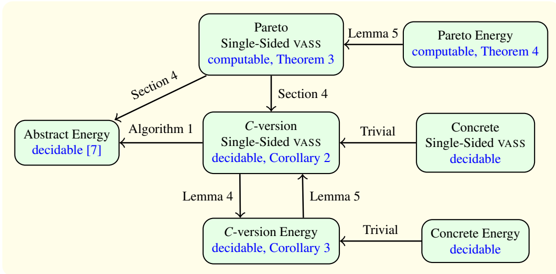

Fig. 1. Problems considered in the paper and their relations. For each property, we state the lemma that show its decidability/computability. The arrows show the reductions of problem instances that we show in the paper.

Energy Parity. Given an integer game $\mathcal{G}$ and a partial configuration γ, we ask whether $[0, \mathcal{E}] : \gamma \models_{\mathcal{G}} \texttt{Parity} \wedge (\Box\overline{\texttt{neg}})$, i.e., whether Player 0 can force a run in the energy semantics where the parity condition is satisfied and at the same time the counters remain nonnegative. The abstract version of this problem is equivalent to the unknown initial credit problem in classical energy parity games [6, 7], since it amounts to asking for the *existence* of a threshold for the initial counter values from which Player 0 can win. The nonnegativity objective $(\Box\overline{\texttt{neg}})$ justifies our restriction to nonnegative partial configurations in our definition of the instantiations and hence of the winning sets.

Theorem 1. *[7] The abstract energy parity problem is decidable.*

The winning set $\mathcal{W}[\mathcal{G}, \mathcal{E}, 0, C](\texttt{Parity} \wedge \Box\overline{\texttt{neg}})$ is upward-closed for $C \subseteq \mathcal{C}$. Intuitively, if Player 0 can win the game with a certain value for the counters, then any higher value for these counters also allows him to win the game with the same strategy. This is because both the possible moves of Player 1 and the colors of configurations depend only on the control-states.

Lemma 3. *For any $C \subseteq \mathcal{C}$, the set $\mathcal{W}[\mathcal{G}, \mathcal{E}, 0, C](\texttt{Parity} \wedge \Box\overline{\texttt{neg}})$ is upward-closed.*

Since this winning set is upward-closed, it follows from Dickson's Lemma that it has finitely many minimal elements. These minimal elements describe the Pareto frontier of the minimal initial credit needed to win the game. In the sequel we will show how to compute this set $\texttt{Pareto}[\mathcal{G}, \mathcal{E}, 0, C](\texttt{Parity} \wedge \Box\overline{\texttt{neg}}) := min(\mathcal{W}[\mathcal{G}, \mathcal{E}, 0, C](\texttt{Parity} \wedge \Box\overline{\texttt{neg}}))$; cf. Theorem 4.

VASS Parity. Given an integer game $\mathcal{G}$ and a nonnegative partial configuration γ, we ask whether $[0,\mathcal{V}] : \gamma \models_{\mathcal{G}} \texttt{Parity}$, i.e., whether Player 0 can force a run in the VASS semantics where the parity condition is satisfied. (The condition $\Box\overline{\texttt{neg}}$ is always trivially satisfied in VASS.) In general, this problem is undecidable as shown in [2], even for simple coverability objectives instead of parity objectives.

Theorem 2. *[2] The VASS Parity Problem is undecidable.*

We will show that decidability of the VASS parity problem is regained under the assumption that $\mathcal{G}$ is single-sided. In [14] it was already shown that, for a single-sided VASS game with reachability objectives, it is possible to compute the set of winning configurations. However, the proof for parity objectives is much more involved.

Correspondence of Single-Sided VASS *Games and Energy Games.* We show that single-sided VASS parity games can be reduced to energy parity games, and vice-versa. The following lemma shows the direction from VASS to energy.

Lemma 4. *Let $\mathcal{G}$ be a single-sided integer game and let γ be a nonnegative partial configuration. Then $[0,\mathcal{V}] : \gamma \models_{\mathcal{G}} \texttt{Parity}$ iff $[0,\mathcal{E}] : \gamma \models_{\mathcal{G}} \texttt{Parity} \wedge \Box\overline{\texttt{neg}}$.*

Hence for a single-sided $\mathcal{G}$ and any set $C \subseteq \mathcal{C}$, we have $\mathcal{W}[\mathcal{G},\mathcal{V},0,C](\texttt{Parity}) = \mathcal{W}[\mathcal{G},\mathcal{E},0,C](\texttt{Parity} \wedge \Box\overline{\texttt{neg}})$. Consequently, using Lemma 3 and Theorem 1, we obtain the following corollary.

Corollary 1. *Let $\mathcal{G}$ be single-sided and $C \subseteq \mathcal{C}$.*

1. *$\mathcal{W}[\mathcal{G},\mathcal{V},0,C](\texttt{Parity})$ is upward-closed.*
2. *The C-version single-sided* VASS *parity problem is reducible to the C-version energy parity problem.*
3. *The* abstract *single-sided* VASS *parity problem (i.e., where $C = \emptyset$) is decidable.*

The following lemma shows the reverse reduction from energy parity games to single-sided VASS parity games.

Lemma 5. *Given an integer game $\mathcal{G} = \langle Q,T,\kappa\rangle$, one can construct a single-sided integer game $\mathcal{G}' = \langle Q',T',\kappa'\rangle$ with $Q \subseteq Q'$ such that $[0,\mathcal{E}] : \gamma \models_{\mathcal{G}} \texttt{Parity} \wedge \Box\overline{\texttt{neg}}$ iff $[0,\mathcal{V}] : \gamma \models_{\mathcal{G}'} \texttt{Parity}$ for every nonnegative partial configuration γ of $\mathcal{G}$.*

Proof sketch. Since $\mathcal{G}'$ needs to be single-sided, Player 1 cannot change the counters. Thus the construction forces Player 0 to simulate the moves of Player 1. Whenever a counter drops below zero in $\mathcal{G}$ (and thus Player 0 loses), Player 0 cannot perform this simulation in $\mathcal{G}'$ and is forced to go to a losing state instead. $\square$

Computability Results. The following theorem (shown in Section 4) states our main computability result. For single-sided VASS parity games, the minimal elements of the winning set $\mathcal{W}[\mathcal{G},\mathcal{V},0,C](\texttt{Parity})$ (i.e., the Pareto frontier) are computable.

Theorem 3. *If $\mathcal{G}$ is single-sided then $\texttt{Pareto}[\mathcal{G},\mathcal{V},0,C](\texttt{Parity})$ is computable.*

In particular, this implies decidability.

Corollary 2. *For any set of counters $C \subseteq \mathcal{C}$, the C-version single-sided* VASS *parity problem is decidable.*

From Theorem 3 and Lemma 5 we obtain the computability of the Pareto frontier of the minimal initial credit needed to win general energy parity games.

Theorem 4. $\mathtt{Pareto}[\mathcal{G}, \mathcal{E}, 0, C](\mathtt{Parity} \wedge \Box\overline{\mathtt{neg}})$ *is computable for any game* $\mathcal{G}$.

Corollary 3. *The C-version energy parity problem is decidable.*

4 Solving Single-Sided VASS Parity Games (Proof of Theorem 3)

Consider a single-sided integer game $\mathcal{G} = \langle Q, T, \kappa \rangle$ and a set $C \subseteq \mathcal{C}$ of counters. We will show how to compute the set $\mathtt{Pareto}[\mathcal{G}, \mathcal{V}, 0, C](\mathtt{Parity})$. We reduce the problem of computing the Pareto frontier in the single-sided VASS parity game to solving the abstract energy parity game problem, which is decidable by Theorem 1.

We use induction on $k = |C|$. As we shall see, the base case is straightforward. We perform the induction step in two phases. First we show that, under the induction hypothesis, we can reduce the problem of computing the Pareto frontier to the problem of solving the C-version single-sided VASS parity problem (i.e., we need only to consider individual nonnegative partial configurations in Γ^C). In the second phase, we introduce an algorithm that translates the latter problem to the abstract energy parity problem.

Base Case. Assume that $C = \emptyset$. In this case we are considering the abstract single-sided VASS parity problem. Recall that $\gamma\uparrow = \{\gamma\}$ for any γ with $dom(\gamma) = \emptyset$. Since $C = \emptyset$, it follows that $\mathtt{Pareto}[\mathcal{G}, \mathcal{V}, 0, C](\mathtt{Parity}) = \{\gamma | \, (dom(\gamma) = \emptyset) \wedge ([0, \mathcal{V}] : \gamma \models_{\mathcal{G}} \mathtt{Parity})\}$. In other words, computing the Pareto frontier in this case reduces to solving the abstract single-sided VASS parity problem, which is decidable by Corollary 1.

From Pareto Sets to VASS *Parity.* Assuming the induction hypothesis, we reduce the problem of computing the set $\mathtt{Pareto}[\mathcal{G}, \mathcal{V}, 0, C](\mathtt{Parity})$ to the C-version single-sided VASS parity problem, i.e., the problem of checking whether $[0, \mathcal{V}] : \gamma \models_{\mathcal{G}} \mathtt{Parity}$ for some $\gamma \in \Gamma^C$ when the underlying integer game is single-sided. To do that, we will instantiate the Valk-Jantzen lemma as follows. We instantiate $U \subseteq \Gamma^C$ in Lemma 2 to be $\mathcal{W}[\mathcal{G}, \mathcal{V}, 0, C](\mathtt{Parity})$ (this set is upward-closed by Corollary 1 since $\mathcal{G}$ is single-sided). Take any nonnegative partial configuration γ with $dom(\gamma) \subseteq C$. We consider two cases. First, if $dom(\gamma) = C$, then we are dealing with the C-version single-sided VASS parity game which will show how to solve in the sequel. Second, consider the case where $dom(\gamma) = C' \subset C$. By the induction hypothesis, we can compute the (finite) set $\mathtt{Pareto}[\mathcal{G}, \mathcal{V}, 0, C'](\mathtt{Parity}) = min(\mathcal{W}[\mathcal{G}, \mathcal{V}, 0, C'](\mathtt{Parity}))$. Then to solve this case, we use the following lemma.

Lemma 6. *For all nonnegative partial configurations γ such that $dom(\gamma) = C' \subset C$, we have $[\![\gamma]\!]_C \cap \mathcal{W}[\mathcal{G}, \mathcal{V}, 0, C](\mathtt{Parity}) \neq \emptyset$ iff $\gamma \in \mathcal{W}[\mathcal{G}, \mathcal{V}, 0, C'](\mathtt{Parity})$.*

Hence checking $[\![\gamma]\!]_C \cap \mathcal{W}[\mathcal{G}, \mathcal{V}, 0, C](\mathtt{Parity}) \neq \emptyset$ amounts to simply comparing γ with the elements of the finite set $\mathtt{Pareto}[\mathcal{G}, \mathcal{V}, 0, C'](\mathtt{Parity})$, because $\mathcal{W}[\mathcal{G}, \mathcal{V}, 0, C'](\mathtt{Parity})$ is upward-closed by Corollary 1.

From VASS *Parity to Abstract Energy Parity.* We introduce an algorithm that uses the induction hypothesis to translate an instance of the C-version single-sided VASS parity problem to an equivalent instance of the abstract energy parity problem.

The following definition and lemma formalize some consequences of the induction hypothesis. First we define a relation that allows us to directly classify some nonnegative partial configurations as winning for Player 1 (resp. Player 0).

Definition 1. *Consider a nonnegative partial configuration γ and a set of nonnegative partial configurations β. We write $\beta \lhd \gamma$ if: (i) for each $\hat{\gamma} \in \beta$, $dom(\hat{\gamma}) \subseteq C$ and $|\gamma| = |\hat{\gamma}| + 1$, and (ii) for each $c \in dom(\gamma)$ there is a $\hat{\gamma} \in \beta$ such that $\hat{\gamma} \preceq \gamma[c \leftarrow \perp]$.*

Lemma 7. *Let $\beta = \bigcup_{C' \subseteq C, |C'| = |C|-1} \texttt{Pareto}[\mathcal{G}, \mathcal{V}, 0, C'](\texttt{Parity})$ be the Pareto frontier of minimal Player 0 winning nonnegative partial configurations with one counter in C undefined. Let $\{c_i, \ldots, c_j\} = \mathcal{C} - C$ be the counters outside C.*

1. *For every $\hat{\gamma} \in \beta$ with $\{c\} = C - dom(\hat{\gamma})$ there exists a minimal finite number $v(\hat{\gamma})$ s.t. $[\![\hat{\gamma}[c \leftarrow v(\hat{\gamma})]]\!] \cap \mathcal{W}[\mathcal{G}, \mathcal{V}, 0, \mathcal{C}](\texttt{Parity}) \neq \emptyset$.*
2. *For every $\hat{\gamma} \in \beta$ there is a number $u(\hat{\gamma})$ s.t. $\hat{\gamma}[c \leftarrow v(\hat{\gamma})][c_i \leftarrow u(\hat{\gamma}), \ldots, c_j \leftarrow u(\hat{\gamma})] \in \mathcal{W}[\mathcal{G}, \mathcal{V}, 0, \mathcal{C}](\texttt{Parity})$, i.e., assigning value $u(\hat{\gamma})$ to counters outside C is sufficient to make the nonnegative configuration winning for Player 0.*
3. *If $\gamma \in \Gamma^C$ is a Player 0 winning nonnegative partial configuration, i.e., $[\![\gamma]\!] \cap \mathcal{W}[\mathcal{G}, \mathcal{V}, 0, \mathcal{C}](\texttt{Parity}) \neq \emptyset$, then $\beta \lhd \gamma$.*

The third part of this lemma implies that if $\neg(\beta \lhd \gamma)$ then we can directly conclude that γ is not winning for Player 0 (and thus winning for Player 1) in the parity game.

Now we are ready to present the algorithm (Algorithm 1).

Input and output of the algorithm. The algorithm inputs a single-sided integer game $\mathcal{G} = \langle Q, T, \kappa \rangle$, and a nonnegative partial configuration γ where $dom(\gamma) = C$. To check whether $[0, \mathcal{V}] : \gamma \models_{\mathcal{G}} \texttt{Parity}$, it constructs an instance of the abstract energy parity problem. This instance is defined by a new integer game $\mathcal{G}^{out} = \langle Q_{out}, T_{out}, \kappa_{out} \rangle$ with counters in $\mathcal{C} - C$, and a nonnegative partial configuration γ^{out}. Since we are considering the abstract version of the problem, the configuration γ^{out} is of the form $\gamma^{out} = \langle q^{out}, \vartheta_{out} \rangle$ where $dom(\vartheta_{out}) = \emptyset$. The latter property means that γ^{out} is uniquely determined by the state q^{out} (all counter values are undefined). Lemma 9 relates $\mathcal{G}$ with the newly constructed $\mathcal{G}^{out}$.

Operation of the algorithm. The algorithm performs a forward analysis similar to the classical Karp-Miller algorithm for Petri nets. We start with a given nonnegative partial configuration, explore its successors, create loops when previously visited configurations are repeated and define a special operation for the case when configurations strictly increase. The algorithm builds the graph of the game $\mathcal{G}^{out}$ successively (i.e., the set of states Q^{out}, the set of transitions T^{out}, and the coloring of states κ). Additionally, for bookkeeping purposes inside the algorithm and for reasoning about the correctness of the algorithm, we define a labeling function λ on the set of states and transitions in $\mathcal{G}^{out}$ such that each state in $\mathcal{G}^{out}$ is labeled by a nonnegative partial configuration in Γ^C, and each transition in $\mathcal{G}^{out}$ is labeled by a transition in $\mathcal{G}$.

Algorithm 1: Building an instance of the abstract energy parity problem.

Input: $\mathcal{G} = \langle Q, T, \kappa \rangle$: Single-Sided Integer Game; $\gamma \in \Gamma^{C}$ with $|C| = k > 0$.
Output: $\mathcal{G}^{out} = \langle Q^{out}, T^{out}, \kappa^{out} \rangle$: integer game;
$q^{out} \in Q^{out}$; $\gamma^{out} = \langle q^{out}, \vartheta_{out} \rangle$ where $dom(\vartheta_{out}) = \emptyset$; $\lambda : Q_{out} \cup T_{out} \mapsto \Gamma^{C} \cup T$

1 $\beta \leftarrow \bigcup_{(C' \subseteq C) \wedge |C'| = |C|-1} \texttt{Pareto}[\mathcal{G}, \mathcal{V}, 0, \mathcal{C}'](\texttt{Parity})$;
2 $T^{out} \leftarrow \emptyset$; $\texttt{new}(q^{out})$; $\kappa(q^{out}) \leftarrow \kappa(\gamma)$; $\lambda(q^{out}) \leftarrow \gamma$; $Q_{out} \leftarrow \{q^{out}\}$;
3 **if** $\lambda(q^{out}) \in \Gamma_0$ **then** $Q_0^{out} \leftarrow \{q^{out}\}$; $Q_1^{out} \leftarrow \emptyset$ **else** $Q_1^{out} \leftarrow \{q^{out}\}$; $Q_0^{out} \leftarrow \emptyset$;
4 $\texttt{ToExplore} \leftarrow \{q^{out}\}$;
5 **while** $\texttt{ToExplore} \neq \emptyset$ **do**
6 Pick and remove a $q \in \texttt{ToExplore}$;
7 **if** $\neg(\beta \lhd \lambda(q))$ **then**
8 $\kappa^{out}(q) \leftarrow 1$; $T^{out} \leftarrow T^{out} \cup \{\langle q, nop, q \rangle\}$
9 **else if** $\exists q' . (q', q) \in (T^{out})^* \wedge (\lambda(q') \prec \lambda(q))$ **then**
10 $\kappa^{out}(q) \leftarrow 0$; $T^{out} \leftarrow T^{out} \cup \{\langle q, nop, q \rangle\}$
11 **else for** *each* $t \in T$ *with* $t(\lambda(q)) \neq \perp$ **do**
12 **if** $\exists q' . (q', q) \in (T^{out})^* . \lambda(q') = t(\lambda(q))$ **then**
13 $T^{out} \leftarrow T^{out} \cup \{\langle q, \texttt{op}(t), q' \rangle\}$; $\lambda(\langle q, \texttt{op}(t), q' \rangle) \leftarrow t$
14 **else**
15 $\texttt{new}(q')$; $\kappa(q') \leftarrow \kappa(t(\lambda(q)))$; $\lambda(q') \leftarrow t(\lambda(q))$;
16 **if** $\lambda(q') \in \Gamma_0$ **then** $Q_0^{out} \leftarrow Q_0^{out} \cup \{q'\}$ **else** $Q_1^{out} \leftarrow Q_1^{out} \cup \{q'\}$;
17 $T^{out} \leftarrow T^{out} \cup \{\langle q, \texttt{op}(t), q' \rangle\}$; $\lambda(\langle q, \texttt{op}(t), q' \rangle) \leftarrow t$;
18 $\texttt{ToExplore} \leftarrow \texttt{ToExplore} \cup \{q'\}$;

The algorithm first computes the Pareto frontier $\texttt{Pareto}[\mathcal{G}, \mathcal{V}, 0, C'](\texttt{Parity})$ for all counter sets $C' \subseteq \mathcal{C}$ with $|C'| = |C| - 1$. This is possible by the induction hypothesis. It stores the union of all these sets in β (line 1). At line 2, the algorithms initializes the set of transitions T^{out} to be empty, creates the first state q^{out}, defines its coloring to be the same as that of the state of the input nonnegative partial configuration γ, labels it by γ, and then adds it to the set of states Q^{out}. At line 3 it adds q^{out} to the set of states of Player 0 or Player 1 (depending on where γ belongs), and at line 4 it adds q^{out} to the set `ToExplore`. The latter contains the set of states that have been created but not yet analyzed by the algorithm.

After the initialization phase, the algorithm starts iterating the *while*-loop starting at line 5. During each iteration, it picks and removes a new state q from the set `ToExplore` (line 6). First, it checks two special conditions under which the game is made immediately losing (resp. winning) for Player 0.

Condition 1: If $\neg(\beta \lhd \lambda(q))$ (line 7), then we know by Lemma 7 (item 3) that the nonnegative partial configuration $\lambda(q)$ is not winning for Player 0 in $\mathcal{G}$.

Therefore, we make the state q losing for Player 0 in $\mathcal{G}^{out}$. To do that, we change the color of q to 1 (any odd color will do), and add a self-loop to q. Any continuation of a run from q is then losing for Player 0 in $\mathcal{G}^{out}$.

Condition 2: If Condition 1 did not hold then the algorithm checks (at line 9) whether there is a *predecessor* q' of q in $\mathcal{G}^{out}$ with a label $\lambda(q')$ that is *strictly* smaller than the label $\lambda(q)$ of q, i.e., $\lambda(q') \prec \lambda(q)$. (Note that we are not comparing q to arbitrary other states in $\mathcal{G}^{out}$, but only to predecessors.) If that is the case, then the state q is made winning for Player 0 in $\mathcal{G}^{out}$. To do that, we change the color of q to 0 (any even color will do), and add a self-loop to q. The intuition for making q winning for Player 0 is as follows. Since $\lambda(q') \prec \lambda(q)$, the path from $\lambda(q')$ to $\lambda(q)$ increases the value of at least one of the defined counters (those in C), and will not decrease the other counters in C (though it might have a negative effect on the undefined counters in $\mathcal{C} - C$). Thus, if a run in $\mathcal{G}$ iterates this path sufficiently many times, the value of at least one counter in C will be pumped and becomes sufficiently high to allow Player 0 to win the parity game on $\mathcal{G}$, provided that the counters in $\mathcal{C} - C$ are initially instantiated with sufficiently high values. This follows from the property $\beta \lhd \lambda(q)$ and Lemma 7 (items 1 and 2).

If none of the tests for Condition1/Condition2 at lines 7 and 9 succeeds, the algorithm continues expanding the graph of $\mathcal{G}^{out}$ from q. It generates all successors of q by applying each transition $t \in T$ in $\mathcal{G}$ to the label $\lambda(q)$ of q (line 11). If the result $t(\lambda(q))$ is defined then there are two possible cases. The first case occurs if we have previously encountered (and added to Q^{out}) a state q' whose label equals $t(\lambda(q))$ (line 12). Then we add a transition from q back to q' in $\mathcal{G}^{out}$, where the operation of the new transition is the same operation as that of t, and define the label of the new transition to be t. Otherwise (line 15), we create a new state q', label it with the nonnegative configuration $t(\lambda(q))$ and assign it the same color as $t(\lambda(q))$. At line 16 q^{out} is added to the set of states of Player 0 or Player 1 (depending on where γ belongs). We add a new transition between q and q' with the same operation as t. The new transition is labeled with t. Finally, we add the new state q' to the set of states to be explored.

Lemma 8. *Algorithm 1 will always terminate.*

Lemma 8 implies that the integer game $\mathcal{G}^{out}$ is finite (and hence well-defined). The following lemma shows the relation between the input and output games $\mathcal{G}, \mathcal{G}^{out}$.

Lemma 9. $[0, \mathcal{V}] : \gamma \models_{\mathcal{G}} \texttt{Parity}$ *iff* $[0, \mathcal{E}] : \gamma^{out} \models_{\mathcal{G}^{out}} \texttt{Parity} \wedge \Box\overline{\texttt{neg}}$.

Proof sketch. The left to right implication is easy. Given a Player 0 winning strategy in $\mathcal{G}$, one can construct a winning strategy in $\mathcal{G}^{out}$ that uses the same transitions, modulo the labeling function $\lambda()$. The condition $\Box\overline{\texttt{neg}}$ in $\mathcal{G}^{out}$ is satisfied since the configurations in $\mathcal{G}$ are always nonnegative and the parity condition is satisfied since the colors seen in corresponding plays in $\mathcal{G}^{out}$ and $\mathcal{G}$ are the same.

For the right to left implication we consider a Player 0 winning strategy σ_0 in $\mathcal{G}^{out}$ and construct a winning strategy σ_0' in $\mathcal{G}$. The idea is that a play π in $\mathcal{G}$ induces a play π' in $\mathcal{G}^{out}$ by using the same sequence of transitions, but removing all so-called *pumping sequences*, which are subsequences that end in **Condition 2**. Then σ_0' acts on history π like σ_0 on history π'. For a play according to σ_0' there are two cases. Either it will eventually reach a configuration that is sufficiently large (relative to β) such that a winning strategy is known by induction hypothesis. Otherwise it contains only finitely many pumping sequences and an infinite suffix of it coincides with an infinite suffix of a play according to σ_0 in $\mathcal{G}^{out}$. Thus it sees the same colors and satisfies `Parity`. $\square$

Since γ^{out} is abstract and the abstract energy parity problem is decidable (Theorem 1) we obtain Theorem 3.

The termination proof in Lemma 8 relies on Dickson's Lemma, and thus there is no elementary upper bound on the complexity of Algorithm 1 or on the size of the constructed energy game $\mathcal{G}^{out}$. The algorithm in [4] for the fixed initial credit problem in pure energy games without the parity condition runs in d-exponential time (resp. $(d-1)$-exponential time for offsets encoded in unary) for dimension d, and is thus not elementary either. As noted in [4], the best known lower bound is EXPSPACE hardness, easily obtained via a reduction from the control-state reachability (i.e., coverability) problem for Petri nets.

5 Applications to Other Problems

5.1 Weak simulation preorder between VASS and finite-state systems

Weak simulation preorder [9] is a semantic preorder on the states of labeled transition graphs, which can be characterized by weak simulation games. A configuration of the game is given by a pair of states (q_1, q_0). In every round of the game, Player 1 chooses a labeled step $q_1 \xrightarrow{a} q_1'$ for some label a. Then Player 0 must respond by a move which is either of the form $q_0 \xrightarrow{\tau^* a \tau^*} q_0'$ if $a \neq \tau$, or of the form $q_0 \xrightarrow{\tau^*} q_0'$ if $a = \tau$ (the special label τ is used to model internal transitions). The game continues from configuration (q_1', q_0'). A player wins if the other player cannot move and Player 0 wins every infinite play. One says that q_0 *weakly simulates* q_1 iff Player 0 has a winning strategy in the weak simulation game from (q_1, q_0). States in different transition systems can be compared by putting them side-by-side and considering them as a single transition system.

We use $\langle Q, T, \Sigma, \lambda \rangle$ to denote a labeled VASS where the states and transitions are defined as in Section 2, Σ is a finite set of labels and $\lambda : T \mapsto \Sigma$ assigns labels to transitions.

It was shown in [3] that it is decidable whether a finite-state labeled transition system weakly simulates a labeled VASS. However, the decidability of the reverse direction was open. (The problem is that the weak $\xrightarrow{\tau^* a \tau^*}$ moves in the VASS make the weak simulation game infinitely branching.) We now show that it is also decidable whether a labeled VASS weakly simulates a finite-state labeled transition system. In particular this implies that weak simulation equivalence between a labeled VASS and a finite-state labeled transition system is decidable. This is in contrast to the undecidability of weak *bisimulation* equivalence between VASS and finite-state systems [11].

Theorem 5. *It is decidable whether a labeled* VASS *weakly simulates a finite-state labeled transition system.*

Proof sketch. Given a labeled VASS and a finite-state labeled transition system, one constructs a single-sided VASS parity game s.t. the VASS weakly simulates the finite system iff Player 0 wins the parity game. The idea is to take a controlled product of the finite system and the VASS s.t. every round of the weak simulation game is encoded by a single move of Player 1 followed by an arbitrarily long sequence of moves by Player 0. The move of Player 1 does not change the counters, since it encodes a move in the finite system, and thus the game is single-sided. Moreover, one enforces that every sequence of consecutive moves by Player 0 is finite (though it can be arbitrarily long), by assigning an odd color to Player 0 states and a higher even color to Player 1 states.

5.2 μ-Calculus model checking VASS

While model checking VASS with linear-time temporal logics (like LTL and linear-time μ-calculus) is decidable [8, 10], model checking VASS with most branching-time logics (like EF, EG, CTL and the modal μ-calculus) is undecidable [8]. However, we show that Theorem 3 yields the decidability of model checking single-sided VASS with a guarded fragment of the modal μ-calculus. We consider a VASS $\langle Q,T\rangle$ where the states, transitions and semantics are defined as in Section 2, and reuse the notion of partial configurations and the transition relation defined for the VASS semantics on integer games. We specify properties on such VASS in the positive μ-calculus L_μ^{pos} whose atomic propositions q refer to control-states $q \in Q$ of the input VASS.

The syntax of the positive μ-calculus L_μ^{pos} is given by the following grammar: $\phi ::= q \mid X \mid \phi\wedge\phi \mid \phi\vee\phi \mid \Diamond\phi \mid \Box\phi \mid \mu X.\phi \mid \nu X.\phi$ where $q\in Q$ and X belongs to a countable set of variables $\mathcal{X}$. The semantics of L_μ^{pos} is defined as usual (see appendix). To each closed formula ϕ in L_μ^{pos} (i.e., without free variables) it assigns a subset of concrete configurations $[\![\phi]\!]$.

The model-checking problem of VASS with L_μ^{pos} can then be defined as follows. Given a VASS $\mathcal{S} = \langle Q,T\rangle$, a closed formula ϕ of L_μ^{pos} and an initial configuration γ_0 of $\mathcal{S}$, do we have $\gamma_0 \in [\![\phi]\!]$? If the answer is yes, we will write $\mathcal{S},\gamma_0 \models \phi$. The more general *global model-checking problem* is to compute the set $[\![\phi]\!]$ of configurations that satisfy the formula. The general unrestricted version of this problem is undecidable.

Theorem 6. *[8] The model-checking problem of* VASS *with L_μ^{pos} is undecidable.*

One way to solve the μ-calculus model-checking problem for a given Kripke structure is to encode the problem into a parity game [12]. The idea is to construct a parity game whose states are pairs, where the first component is a state of the structure and the second component is a subformula of the given μ-calculus formula. States of the form $\langle q,\Box\phi\rangle$ or $\langle q,\phi\wedge\psi\rangle$ belong to Player 1 and the remainder belong to Player 0. The colors are assigned to reflect the nesting of least and greatest fixpoints. We can adapt this construction to our context by building an integer game from a formula in L_μ^{pos} and a VASS $\mathcal{S}$, as stated by the next lemma.

Lemma 10. *Let $\mathcal{S}$ be a* VASS, *γ_0 a concrete configuration of $\mathcal{S}$ and ϕ a closed formula in L_μ^{pos}. One can construct an integer game $\mathcal{G}(\mathcal{S},\phi)$ and an initial concrete configuration γ_0' such that $[0,\mathcal{V}] : \gamma_0' \models_{\mathcal{G}(\mathcal{S},\phi)}$* `Parity` *if and only if $\mathcal{S},\gamma_0 \models \phi$.*

Now we show that, under certain restrictions on the considered VASS and on the formula from L_μ^{pos}, the constructed integer game $\mathcal{G}(\mathcal{S},\phi)$ is single-sided, and hence we obtain the decidability of the model-checking problem from Theorem 3. First, we reuse the notion of single-sided games from Section 2 in the context of VASS, by saying that a VASS $\mathcal{S} = \langle Q,T\rangle$ is single-sided iff there is a partition of the set of states Q into two sets Q_0 and Q_1 such that $op = nop$ for all transitions $t\in T$ with $\texttt{source}(t) \in Q_1$. The guarded fragment L_μ^{sv} of L_μ^{pos} for single-sided VASS is then defined by guarding the $\Box$ operator with a predicate that enforces control-states in Q_1. Formally, the syntax of L_μ^{sv} is given by the following grammar: $\phi ::= q \mid X \mid \phi\wedge\phi \mid \phi\vee\phi \mid \Diamond\phi \mid Q_1 \wedge \Box\phi \mid \mu X.\phi \mid \nu X.\phi$, where Q_1 stands for the formula $\bigvee_{q\in Q_1} q$. By analyzing the construction of Lemma 10 in this restricted case, we obtain the following lemma.

Lemma 11. *If $\mathcal{S}$ is a single-sided* VASS *and $\phi \in L_\mu^{sv}$ then the game $\mathcal{G}(\mathcal{S},\phi)$ is equivalent to a single-sided game.*

By combining the results of the last two lemmas with Corollary 1, Theorem 3 and Corollary 2, we get the following result on model checking single-sided VASS.

Theorem 7.

1. *Model checking L_μ^{sv} over single-sided* VASS *is decidable.*
2. *If $\mathcal{S}$ is a single-sided* VASS *and ϕ is a formula of L_μ^{sv} then $[\![\phi]\!]$ is upward-closed and its set of minimal elements is computable.*

6 Conclusion and Outlook

We have established a connection between multidimensional energy games and single-sided VASS games. Thus our algorithm to compute winning sets in VASS parity games can also be used to compute the minimal initial credit needed to win multidimensional energy parity games, i.e., the Pareto frontier.

It is possible to extend our results to integer parity games with a mixed semantics, where a subset of the counters follow the energy semantics and the rest follow the VASS semantics. If such a mixed parity game is single-sided w.r.t. the VASS counters (but not necessarily w.r.t. the energy counters) then it can be reduced to a single-sided VASS parity game by our construction in Section 3. The winning set of the derived single-sided VASS parity game can then be computed with the algorithm in Section 4.

References

1. P.A. Abdulla, M. Atig, and J. Stenman. The minimal cost reachability problem in priced timed pushdown systems. In *LATA'12*, volume 7183 of *LNCS*, pages 58–69. Springer, 2012.
2. P.A. Abdulla, A. Bouajjani, and J. d'Orso. Deciding monotonic games. In *CSL'03*, volume 2803 of *LNCS*, pages 1–14. Springer, 2003.
3. P.A. Abdulla, K. Čerāns, B. Jonsson, and Y. Tsay. General decidability theorems for infinite-state systems. In *LICS'96*, pages 313–321. IEEE, 1996.
4. T. Brázdil, P. Jančar, and A. Kučera. Reachability games on extended vector addition systems with states. In *ICALP'10*, volume 6199 of *LNCS*, pages 478–489. Springer, 2010.
5. K. Chatterjee and L. Doyen. Energy parity games. *TCS*, 458:49–60, 2012.
6. K. Chatterjee, L. Doyen, T. Henzinger, and J.-F. Raskin. Generalized mean-payoff and energy games. In *FSTTCS'10*, volume 8 of *LIPIcs, Schloss Dagstuhl - LZI*, pages 505–516, 2010.
7. K. Chatterjee, M. Randour, and J.-F. Raskin. Strategy synthesis for multi-dimensional quantitative objectives. In *CONCUR'12*, volume 7454 of *LNCS*, pages 115–131. Springer, 2012.
8. J. Esparza and M. Nielsen. Decibility issues for Petri nets - a survey. *Journal of Informatik Processing and Cybernetics*, 30(3):143–160, 1994.
9. R.J. van Glabbeek. The linear time – branching time spectrum I; the semantics of concrete, sequential processes. In J.A. Bergstra, A. Ponse, and S.A. Smolka, editors, *Handbook of Process Algebra*, chapter 1, pages 3–99. Elsevier, 2001.

10. P. Habermehl. On the complexity of the linear-time mu-calculus for Petri-nets. In *ICATPN'97*, volume 1248 of *LNCS*, pages 102–116. Springer, 1997.
11. P. Jančar, J. Esparza, and F. Moller. Petri nets and regular processes. *J. Comput. Syst. Sci.*, 59(3):476–503, 1999.
12. D. Kirsten. Alternating tree automata and parity games. In E. Grädel, W. Thomas, and T. Wilke, editors, *Automata, Logics, and Infinite Games*, volume 2500 of *LNCS*, pages 153–167. Springer, 2002.
13. R. McNaughton. Infinite games played on finite graphs. *Annals of Pure and Applied Logic*, 65:149–184, 1993.
14. J.-F. Raskin, M. Samuelides, and L. Van Begin. Games for counting abstractions. *Electr. Notes Theor. Comput. Sci.*, 128(6):69–85, 2005.
15. R. Valk and M. Jantzen. The residue of vector sets with applications to decidability problems in Petri nets. *Acta Inf.*, 21:643–674, 1985.
16. W. Zielonka. Infinite games on finitely coloured graphs with applications to automata on infinite trees. *TCS*, 200:135–183, 1998.

Appendix

Proof of Lemma 1

Let $\gamma_1 = \langle q_1, \vartheta_1 \rangle$, $\gamma_2 = \langle q_2, \vartheta_2 \rangle$, and $\gamma_3 = \langle q_3, \vartheta_3 \rangle$ be nonnegative partial configurations and let $t = \langle q'_1, op, q'_2 \rangle$ in T.

Assume that $\gamma_1 \xrightarrow{t}_{\mathcal{V}} \gamma_2$ and that $\gamma_1 \preceq \gamma_3$. From this we know that $q_1 = q'_1 = q_3$, that $q_2 = q'_2$ and also that $dom(\gamma_1) = dom(\gamma_2) = dom(\gamma_3)$. There are several cases for a transition t that can be taken from γ_1. If op is an increment or a *nop* operation then only the control-state matters for taking the transition. If op is a decrement transition then the initial value of the decremented counter has to be either $\perp$ or ≥ 1. Since this is the case for ϑ_1 and since $\gamma_1 \preceq \gamma_3$, we deduce that this also holds for ϑ_3. Then we obtain the nonnegative partial configuration $\gamma_4 = \langle q_4, \vartheta_4 \rangle$ from γ_3 by following the rule of the transition relation $\xrightarrow{t}_{\mathcal{V}}$. Moreover, we can deduce that $\gamma_2 \preceq \gamma_4$, because any operation on the undefined counters leaves the counters undefined, and for the other counters one can easily prove that for all $c' \in dom(\gamma_1)$, $\vartheta_4(c') = \vartheta_2(c') + (\vartheta_3(c') - \vartheta_1(c'))$.

Now suppose that $\gamma_1 \xrightarrow{t}_{\mathcal{V}} \gamma_2$, that $\gamma_3 \preceq \gamma_1$, that $\mathcal{G}$ is single-sided and that $\gamma_1 \in \Gamma_1$. It follows that $q_1 = q'_1 = q_3$, that $q_2 = q'_2$ and also that $dom(\gamma_1) = dom(\gamma_2) = dom(\gamma_3)$. Furthermore, since $\gamma_1 \in \Gamma_1$, we deduce that $q_1 \in Q_1$ and, since $\mathcal{G}$ is single-sided, we have that $op = nop$. Hence, by definition of the transition relation $\xrightarrow{t}_{\mathcal{V}}$, we obtain $\vartheta_1 = \vartheta_2$ and so by choosing $\gamma_4 = \langle q_2, \vartheta_3 \rangle$ we obtain that $\gamma_3 \xrightarrow{t} \gamma_4$. Since $\gamma_3 \preceq \gamma_1$, we have $\vartheta_3(c) \leq \vartheta_1(c)$ for all $c \in dom(\gamma_1)$ and hence $\gamma_4 \preceq \gamma_2$.

Proof of Lemma 2

Usually the Valk and Jantzen Lemma, which allows the computation of the minimal elements of an upward-closed set of vectors of naturals, is stated a bit differently by using vectors of naturals with ω at some indexes to represent any integer values (see for instance in [1]). In our context, the ω are replaced by undefined values for the counters in the considered nonnegative partial configurations, but the idea is the same. The usual way to express the Valk and Jantzen Lemma is as follows: For $C \subseteq \mathcal{C}$ and an upward-closed set $U \subseteq \Gamma^C$, $min(U)$ is computable if and only if for any nonnegative partial configuration γ with $dom(\gamma) \subseteq C$, one can decide whether $[\![\gamma]\!]_C\!\downarrow \cap U \neq \emptyset$. Now we show that this way of formalizing the Valk and Jantzen Lemma is equivalent to the statement of Lemma 2.

First, if we assume that $min(U)$ is computable, then it is obvious that for any nonnegative partial configuration γ with $dom(\gamma) \subseteq C$, we can decide whether $[\![\gamma]\!]_C \cap U \neq \emptyset$. In fact, it suffices to check whether there exists a $\gamma_1 \in min(U)$ such that for all $c \in dom(\gamma)$, we have $\gamma(c) \geq \gamma_1(c)$ (since U is upward-closed). Since $min(U)$ is finite, it is possible check this condition for all nonnegative partial configurations γ_1 in $min(U)$.

Now assume that for any nonnegative partial configuration γ with $dom(\gamma) \subseteq C$, we can decide whether $[\![\gamma]\!]_C \cap U \neq \emptyset$. Consider a configuration γ_1 with $dom(\gamma_1) \subseteq C$. First note that $\gamma_1\!\downarrow$ is a finite set and also that $[\![\gamma_1]\!]_C\!\downarrow = \bigcup_{\gamma_2 \in \gamma_1\downarrow} [\![\gamma_2]\!]_C$. But since $\gamma_1\!\downarrow$ is finite, and since we can decide whether $[\![\gamma_2]\!]_C \cap U \neq \emptyset$ for each $\gamma_2 \in \gamma_1\!\downarrow$, we can decide whether $[\![\gamma_1]\!]_C\!\downarrow \cap U \neq \emptyset$. By the Valk and Jantzen Lemma, $min(U)$ is computable.

Proof of Lemma 3

We will show that the set $\mathcal{W}[\mathcal{G},\mathcal{E},0,C](\texttt{Parity} \wedge \Box\overline{\texttt{neg}})$ is upward-closed. Let $\gamma_1, \gamma_2 \in \Gamma^C$ (with $\gamma_1 = \langle q_1, \vartheta_1 \rangle$ and $\gamma_2 = \langle q_1, \vartheta_2 \rangle$) such that $\gamma_1 \in \mathcal{W}[\mathcal{G},\mathcal{E},0,C](\texttt{Parity} \wedge \Box\overline{\texttt{neg}})$ and $\gamma_1 \preceq \gamma_2$. In order to prove that $\mathcal{W}[\mathcal{G},\mathcal{E},0,C](\texttt{Parity} \wedge \Box\overline{\texttt{neg}})$ is upward-closed, we need to show that $\gamma_2 \in \mathcal{W}[\mathcal{G},\mathcal{E},0,C](\texttt{Parity} \wedge \Box\overline{\texttt{neg}})$.

Since $\gamma_1 \in \mathcal{W}[\mathcal{G},\mathcal{E},0,C](\texttt{Parity} \wedge \Box\overline{\texttt{neg}})$, there exists $\gamma_1' \in [\![\gamma_1]\!]$ such that $[0,\mathcal{E}]: \gamma_1' \models_{\mathcal{G}} \texttt{Parity} \wedge \Box\overline{\texttt{neg}}$, i.e., there exists $\gamma_1' = \langle q_1, \vartheta_1' \rangle \in [\![\gamma_1]\!]$ and $\sigma_0 \in \Sigma_0^{\mathcal{E}}$ such that $\texttt{run}(\gamma_1', \sigma_0, \sigma_1) \models_{\mathcal{G}} \texttt{Parity} \wedge \Box\overline{\texttt{neg}}$ for all $\sigma_1 \in \Sigma_1^{\mathcal{E}}$. Let us first define the following concrete configuration $\gamma_2' = \langle q_1, \vartheta_2' \rangle$ with:

$$\vartheta_2'(c) = \begin{cases} \vartheta_2(c) & \text{if } c \in dom(\gamma_2) \\ \vartheta_1'(c) & \text{if } c \notin dom(\gamma_2) \end{cases}$$

By definition we have $\gamma_2' \in [\![\gamma_2]\!]$ and since $\gamma_1 \preceq \gamma_2$, we also have $\gamma_1' \preceq \gamma_2'$ (i.e. $\vartheta_1'(c) \leq \vartheta_2'(c)$ for all $c \in \mathcal{C}$). We want to show that $[0,\mathcal{E}] : \gamma_2' \models_{\mathcal{G}} \texttt{Parity} \wedge \Box\overline{\texttt{neg}}$, i.e., that player 0 has a winning strategy from the concrete configuration γ_2'.

We now show how to build a winning strategy $\sigma_0' \in \Sigma_0^{\mathcal{E}}$ for player 0 from the configuration γ_2'. First to any γ_2'-path $\pi = \gamma_0'' \xrightarrow{t_1}_{\mathcal{E}} \gamma_1'' \xrightarrow{t_2}_{\mathcal{E}} \cdots \gamma_n''$ we associate the γ_1'-path $\alpha(\pi) = \gamma_0''' \xrightarrow{t_1}_{\mathcal{E}} \gamma_1''' \xrightarrow{t_2}_{\mathcal{E}} \cdots \gamma_n'''$ where for all $j \in \{0,\ldots,n\}$, if $\gamma_j'' = \left\langle q_j'', \vartheta_j'' \right\rangle$ then $\gamma_j''' = \left\langle q_j'', \vartheta_j''' \right\rangle$ with $\vartheta_j'''(c) = \vartheta_j''(c) - (\vartheta_2'(c) - \vartheta_1'(c))$ for all $c \in \mathcal{C}$ (i.e. to obtain $\alpha(\pi)$ from π, we decrement each counter valuation by the difference between $\vartheta_2'(c) - \vartheta_1'(c)$). Note that $\alpha(\pi)$ is a valid path since we are considering the energy semantics where the counters can take negative values. Now we define the strategy $\sigma_0' \in \Sigma_0^{\mathcal{E}}$ for player 0 as $\sigma_0'(\pi) = \sigma_0(\alpha(\pi))$ for each γ_2'-path π. Here again the strategy is well defined since in energy games the enabledness of a transition depends only on the control-state and not on the counter valuation.

We will now prove that for all strategies $\sigma_1' \in \Sigma_1^{\mathcal{E}}$, we have $\texttt{run}(\gamma_2', \sigma_0', \sigma_1') \models_{\mathcal{G}} \texttt{Parity} \wedge \Box\overline{\texttt{neg}}$. Let $\sigma_1' \in \Sigma_1^{\mathcal{E}}$. Using σ_1', we will construct another strategy $\sigma_1 \in \Sigma_1^{\mathcal{E}}$ and prove that if $\texttt{run}(\gamma_1', \sigma_0, \sigma_1) \models_{\mathcal{G}} \texttt{Parity} \wedge \Box\overline{\texttt{neg}}$ then $\texttt{run}(\gamma_2', \sigma_0', \sigma_1') \models_{\mathcal{G}} \texttt{Parity} \wedge \Box\overline{\texttt{neg}}$. Before we give the definition of σ_1, we introduce another notation. To any γ_1'-path $\pi = \gamma_0''' \xrightarrow{t_1}_{\mathcal{E}} \gamma_1''' \xrightarrow{t_2}_{\mathcal{E}} \cdots \gamma_n'''$ we associate the γ_2'-path $\overline{\alpha}(\pi) = \gamma_0'' \xrightarrow{t_1}_{\mathcal{E}} \gamma_1'' \xrightarrow{t_2}_{\mathcal{E}} \cdots \gamma_n''$ where for all $j \in \{0,\ldots,n\}$, if $\gamma_j''' = \left\langle q_j'', \vartheta_j''' \right\rangle$ then $\gamma_j'' = \left\langle q_j'', \vartheta_j'' \right\rangle$ with $\vartheta_j''(c) = \vartheta_j'''(c) + (\vartheta_2'(c) - \vartheta_1'(c))$ for all $c \in \mathcal{C}$ (i.e., to obtain $\overline{\alpha}(\pi)$ from π, we increment each counter valuation by the difference between $\vartheta_2'(c) - \vartheta_1'(c)$). Note that $\overline{\alpha}(\pi)$ is a valid path. Now we define the strategy $\sigma_1 \in \Sigma_1^{\mathcal{E}}$ for Player 1 as $\sigma_1(\pi) = \sigma_1'(\overline{\alpha}(\pi))$ for each γ_1'-path π.

We extend in the obvious way the function $\alpha()$ [resp. $\overline{\alpha}()$] to γ_2'-run [resp. to γ_1'-run]. Then one can easily check that we have $\alpha(\texttt{run}(\gamma_2', \sigma_0', \sigma_1')) = \texttt{run}(\gamma_1', \sigma_0, \sigma_1)$ and that $\texttt{run}(\gamma_2', \sigma_0', \sigma_1') = \overline{\alpha}(\texttt{run}(\gamma_1', \sigma_0, \sigma_1))$ by construction of the strategy σ_0' and σ_1. First, remember that σ_0 is a winning strategy for Player 0 from the configuration γ_1'. Thus we have $\texttt{run}(\gamma_1', \sigma_0, \sigma_1) \models_{\mathcal{G}} \texttt{Parity} \wedge \Box\overline{\texttt{neg}}$. Since in $\overline{\alpha}(\texttt{run}(\gamma_1', \sigma_0, \sigma_1))$ the sequence of control-states are the same and all the counter valuations along the path are greater or equal to the ones seen in $\texttt{run}(\gamma_1', \sigma_0, \sigma_1)$ (remember that we add to each configuration, to each counter c the quantity $\vartheta'_2(c) - \vartheta'_1(c) \geq 0$), this allows us to deduce that $\overline{\alpha}(\mathtt{run}(\gamma'_1, \sigma_0, \sigma_1)) \models_{\mathcal{G}} \mathtt{Parity} \wedge \Box\overline{\mathtt{neg}}$. Hence we have $\mathtt{run}(\gamma'_2, \sigma'_0, \sigma'_1) \models_{\mathcal{G}} \mathtt{Parity} \wedge \Box\overline{\mathtt{neg}}$.

Finally we have proved that there exist $\gamma'_2 \in [\![\gamma_2]\!]$ and $\sigma'_0 \in \Sigma_0^{\mathcal{E}}$ such that $\mathtt{run}(\gamma'_2, \sigma'_0, \sigma'_1) \models_{\mathcal{G}} \mathtt{Parity} \wedge \Box\overline{\mathtt{neg}}$ for all $\sigma'_1 \in \Sigma_1^{\mathcal{E}}$. So Player 0 has a winning strategy from an instantiation of the configuration γ_2. Hence $\gamma_2 \in \mathcal{W}[\mathcal{G}, \mathcal{E}, 0, C](\mathtt{Parity} \wedge \Box\overline{\mathtt{neg}})$.

Proof of Lemma 4

Let $\mathcal{G} = \langle Q, T, \kappa \rangle$ be a single-sided integer game and $\gamma \in \Gamma^C$ a nonnegative partial configuration.

First we will assume that $[0, \mathcal{E}] : \gamma \models_{\mathcal{G}} \mathtt{Parity} \wedge \Box\overline{\mathtt{neg}}$. This means that there exists $\gamma' = \langle q, \vartheta \rangle \in [\![\gamma]\!]$ and $\sigma_0 \in \Sigma_0^{\mathcal{E}}$ such that $\mathtt{run}(\gamma', \sigma_0, \sigma_1) \models_{\mathcal{G}} \mathtt{Parity} \wedge \Box\overline{\mathtt{neg}}$ for all $\sigma_1 \in \Sigma_1^{\mathcal{E}}$. The idea we will use here is that since the strategy σ_0 keeps the value of the counters positive, then the same strategy can be followed under the VASS semantics, and furthermore this strategy will be a winning strategy for the VASS parity game. Let us formalize this idea. We build the strategy $\sigma'_0 \in \Sigma_0^{\mathcal{V}}$ as follows: for any path $\pi = \gamma_0 \xrightarrow{t_1}_{\mathcal{V}} \gamma_1 \xrightarrow{t_2}_{\mathcal{V}} \cdots \gamma_n$, we have $\sigma'_0(\pi) = \sigma_0(\pi)$ if $\sigma_0(\pi)(\gamma_n) \neq \bot$ (under the VASS semantics) and otherwise $\sigma'_0(\pi)$ equals any enabled transition. Note that this definition is valid since any path in the VASS semantics is also a path in the energy semantics. We consider now a strategy $\sigma'_1 \in \Sigma_1^{\mathcal{V}}$. This strategy can be easily extended to a strategy $\sigma_1 \in \Sigma_1^{\mathcal{E}}$ for the energy game by playing any transition when the input path is not a path valid under the VASS semantics. First note that since σ_0 is a winning strategy in the energy parity game we have $\mathtt{run}(\gamma', \sigma_0, \sigma_1) \models_{\mathcal{G}} \mathtt{Parity} \wedge \Box\overline{\mathtt{neg}}$. From the way we build the strategies, we deduce that $\mathtt{run}(\gamma', \sigma'_0, \sigma'_1) = \mathtt{run}(\gamma', \sigma_0, \sigma_1)$. Since the colors seen along a run depend only of the control-state, we deduce that $\mathtt{run}(\gamma', \sigma'_0, \sigma'_1) \models_{\mathcal{G}} \mathtt{Parity}$. Hence we have proven that $[0, \mathcal{V}] : \gamma \models_{\mathcal{G}} \mathtt{Parity}$.

We now assume that $[0, \mathcal{V}] : \gamma \models_{\mathcal{G}} \mathtt{Parity}$. This means that there exists $\gamma' = \langle q, \vartheta \rangle \in [\![\gamma]\!]$ and $\sigma_0 \in \Sigma_0^{\mathcal{V}}$ such that $\mathtt{run}(\gamma', \sigma_0, \sigma_1) \models_{\mathcal{G}} \mathtt{Parity}$ for all $\sigma_1 \in \Sigma_1^{\mathcal{V}}$. We build a strategy $\sigma'_0 \in \Sigma_0^{\mathcal{E}}$ as follows: for any path π in the VASS semantics $\sigma'_0(\pi) = \sigma_0(\pi)$; otherwise, if π is not a valid path under the VASS semantics, $\sigma'_0(\pi)$ is equal to any transition enabled in the last configuration of the path. Take now a strategy $\sigma'_1 \in \Sigma_0^{\mathcal{E}}$ for Player 1 in the energy parity game. From σ'_1, we define a strategy $\sigma_1 \in \Sigma_0^{\mathcal{V}}$ as follows: for any path π in the VASS semantics, let $\sigma_1(\pi) = \sigma'_1(\pi)$. Note that since the game is single-sided this strategy is well defined; in fact, in a single-sided game, in the states of Player 1, all the outgoing transitions are enabled in the energy and in the VASS semantics (because in single-sided games, Player 1 does not change the counter values). But then we have $\mathtt{run}(\gamma', \sigma_0, \sigma_1) = \mathtt{run}(\gamma', \sigma'_0, \sigma'_1)$ and since $\mathtt{run}(\gamma', \sigma_0, \sigma_1) \models_{\mathcal{G}} \mathtt{Parity}$ and since it is a valid run under the VASS semantics, we deduce that the values of the counters always remain positive. Consequently we have $\mathtt{run}(\gamma', \sigma'_0, \sigma'_1) \models_{\mathcal{G}} \mathtt{Parity} \wedge \Box\overline{\mathtt{neg}}$. We conclude that $[0, \mathcal{E}] : \gamma \models_{\mathcal{G}} \mathtt{Parity} \wedge \Box\overline{\mathtt{neg}}$.

Proof of Lemma 5

Let $\mathcal{G} = \langle Q, T, \kappa \rangle$ be an integer game. From it we build a single-sided integer game $\mathcal{G}' = \langle Q', T', \kappa' \rangle$ as follows:

- $Q' = Q \uplus \{q_t \mid t \in T\} \uplus \{q_\ell\}$ (where $\uplus$ denotes the disjoint union operator), with $Q'_0 = Q_0 \uplus \{q_t \mid t \in T\} \uplus \{q_\ell\}$ and $Q'_1 = Q_1$;
- T' is the smallest set of transitions such that, for each transition $t = \langle q_1, op, q_2 \rangle$ in T, the following conditions are respected:
  - $\langle q_1, nop, q_t \rangle \in T'$;
  - $\langle q_t, op, q_2 \rangle \in T'$;
  - $\langle q_t, nop, q_\ell \rangle \in T'$;
  - $\langle q_\ell, nop, q_\ell \rangle \in T'$;
- κ' is defined as follows:
  - for all $q \in Q$, $\kappa'(q) = \kappa(q)$;
  - for all $t \in T$, $\kappa'(q_t) = 0$;
  - $\kappa'(q_\ell) = 1$.

By construction $\mathcal{G}'$ is single-sided. Also note that once the system enters the losing state q_ℓ, Player 0 loses the game since the only possible infinite run from this state remains in q_ℓ and the color associated to this state is odd (it is equal to 1). Figure 2 depicts the encoding of transitions of the form $\langle q_1, c\text{--}, q_2 \rangle$.

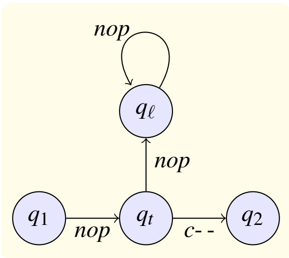

Fig. 2. Translating a transition $\langle q_1, c\text{--}, q_2 \rangle$ from an energy game to a single-sided VASS game. Note that $\kappa(q_\ell)$ is odd.

We will now prove that $\mathcal{W}[\mathcal{G}, \mathcal{E}, 0, C](\texttt{Parity} \wedge \Box\overline{\texttt{neg}}) = \mathcal{W}[\mathcal{G}', \mathcal{V}, 0, C](\texttt{Parity}) \cap \{\gamma \mid \texttt{state}(\gamma) \in Q\}$. First let $\gamma \in \mathcal{W}[\mathcal{G}, \mathcal{E}, 0, C](\texttt{Parity} \wedge \Box\overline{\texttt{neg}})$. This means that there exists $\gamma' \in [\![\gamma]\!]$ and $\sigma_0 \in \Sigma_0^{\mathcal{E}}$ such that $[0, \sigma_0, \mathcal{E}] : \gamma \models_{\mathcal{G}} \texttt{Parity} \wedge \Box\overline{\texttt{neg}}$. From σ_0, we will build a winning strategy $\sigma'_0 \in \Sigma_0^{\mathcal{V}}$ for player 0 in $\mathcal{G}'$. Let us first introduce some notation. To a path in $\mathcal{G}'$, $\pi = \gamma_0 \xrightarrow{t_1}_{\mathcal{V}} \gamma_{t_1} \xrightarrow{t'_1}_{\mathcal{V}} \gamma_1 \xrightarrow{t_2}_{\mathcal{V}} \gamma_{t_2} \xrightarrow{t'_2}_{\mathcal{V}} \gamma_2 \cdots \gamma_n$ with $\texttt{state}(\gamma_n) \in Q$, we associate the path $\beta(\pi) = \gamma_0 \xrightarrow{t1}_{\mathcal{E}} \gamma_1 \xrightarrow{t2}_{\mathcal{E}} \cdots \gamma_n$ in $\mathcal{G}$ (by construction of $\mathcal{G}'$ such a path exists). The strategy σ'_0 is then defined as follows. For all paths $\pi = \gamma_0 \xrightarrow{t_1}_{\mathcal{V}} \gamma_1 \xrightarrow{t_2}_{\mathcal{V}} \gamma_2 \cdots \gamma_n$ in $\mathcal{G}'$:

- if $\mathtt{state}(\gamma_n) \in Q$, then $\sigma_0'(\pi) = \langle \mathtt{state}(\gamma_n), nop, q_t \rangle$ with $t = \sigma_0(\beta(\pi))$;
- if $\mathtt{state}(\gamma_n) = q_t$ for some transition $t = \langle q_1, op, q_2 \rangle \in T$, then if $\langle q_t, op, q_2 \rangle$ is enabled in γ_n, $\sigma_0'(\pi) = \langle q_t, op, q_2 \rangle$, otherwise $\sigma_0' = \langle q_t, nop, q_\ell \rangle$;
- if $\mathtt{state}(\gamma_n) = q_\ell$, then $\sigma_0'(\pi) = \langle q_\ell, nop, q_\ell \rangle$.

One can then easily verify using the definition of $\mathcal{G}'$ and of the strategy σ_0' that since $[0, \sigma_0, \mathcal{E}] : \gamma \models_{\mathcal{G}} \mathtt{Parity} \wedge \Box\overline{\mathtt{neg}}$, we have $[0, \sigma_0', \mathcal{V}] : \gamma \models_{\mathcal{G}} \mathtt{Parity}$ and hence that $\gamma \in \mathcal{W}[\mathcal{G}', \mathcal{V}, 0, C](\mathtt{Parity}) \cap \{\gamma \mid \mathtt{state}(\gamma) \in Q\}$.

The proof that if we take $\gamma \in \mathcal{W}[\mathcal{G}', \mathcal{V}, 0, C](\mathtt{Parity}) \cap \{\gamma \mid \mathtt{state}(\gamma) \in Q\}$ then γ belongs also to $\mathcal{W}[\mathcal{G}, \mathcal{E}, 0, C](\mathtt{Parity} \wedge \Box\overline{\mathtt{neg}})$ is done similarly.

Proof of Lemma 6

Let γ be a nonnegative partial configuration such that $dom(\gamma) = C' \subset C$. Suppose that $[\![\gamma]\!]_C \cap \mathcal{W}[\mathcal{G}, \mathcal{V}, 0, C](\mathtt{Parity}) \neq \emptyset$, i.e., there is a $\gamma_1 \in [\![\gamma]\!]_C$ where $\gamma_1 \in \mathcal{W}[\mathcal{G}, \mathcal{V}, 0, C](\mathtt{Parity})$. Since $\gamma_1 \in \mathcal{W}[\mathcal{G}, \mathcal{V}, 0, C](\mathtt{Parity})$ there is a $\gamma_2 \in [\![\gamma_1]\!]$ with $[0, \mathcal{V}] : \gamma_2 \models_{\mathcal{G}} \mathtt{Parity}$. Notice that $\gamma_2 \in [\![\gamma]\!]$. It follows that $\gamma \in \mathcal{W}[\mathcal{G}, \mathcal{V}, 0, C'](\mathtt{Parity})$.

Now, suppose that $\gamma \in \mathcal{W}[\mathcal{G}, \mathcal{V}, 0, C'](\mathtt{Parity})$. By definition there is a $\gamma_1 \in [\![\gamma]\!]$ such that $[0, \mathcal{V}] : \gamma_1 \models_{\mathcal{G}} \mathtt{Parity}$. Define γ_2 by $\gamma_2(c) := \gamma_1(c)$ for all $c \in C$ and $\gamma_2(c) := \perp$ for all $c \notin C$. Then $\gamma_2 \in [\![\gamma]\!]_C$ and $\gamma_2 \in \mathcal{W}[\mathcal{G}, \mathcal{V}, 0, C](\mathtt{Parity})$, hence $[\![\gamma]\!]_C \cap \mathcal{W}[\mathcal{G}, \mathcal{V}, 0, C](\mathtt{Parity}) \neq \emptyset$.

Proof of Lemma 7

1. Consider a partial nonnegative configuration $\hat{\gamma} \in \beta$ where $c \in C - dom(\hat{\gamma})$.
   Since $[\![\hat{\gamma}]\!] \cap \mathcal{W}[\mathcal{G}, \mathcal{V}, 0, \mathcal{C}](\mathtt{Parity}) \neq \emptyset$, there exists a minimal finite number $v(\hat{\gamma})$ s.t. $\mathcal{W}(\hat{\gamma}) := [\![\hat{\gamma}[c \leftarrow v(\hat{\gamma})]]\!] \cap \mathcal{W}[\mathcal{G}, \mathcal{V}, 0, \mathcal{C}](\mathtt{Parity}) \neq \emptyset$.
2. In particular, $\mathcal{W}(\hat{\gamma})$ is upward-closed w.r.t. the counters in $\mathcal{C} - C$ and $min(\mathcal{W}(\hat{\gamma}))$ is finite. Let $u(\hat{\gamma})$ be the maximal constant appearing in $min(\mathcal{W}(\hat{\gamma}))$. Thus, an instantiation of $\hat{\gamma}[c \leftarrow v(\hat{\gamma})]$ where the counters in $\mathcal{C} - C$ have values $\geq u(\hat{\gamma})$ is certainly winning for Player 0, i.e., in $\mathcal{W}[\mathcal{G}, \mathcal{V}, 0, \mathcal{C}](\mathtt{Parity})$.
3. The first condition of Def. 1 is satisfied by the definition of β. Moreover, since $\gamma \in \Gamma^C$ and $[\![\gamma]\!] \cap \mathcal{W}[\mathcal{G}, \mathcal{V}, 0, \mathcal{C}](\mathtt{Parity}) \neq \emptyset$, for every $c \in C$ we have $[\![\gamma[c \leftarrow \perp]]\!] \cap \mathcal{W}[\mathcal{G}, \mathcal{V}, 0, \mathcal{C}](\mathtt{Parity}) \neq \emptyset$. Since β are by definition the minimal nonnegative configurations (with a domain which is exactly one element smaller than C) that have this property, there must exist some element $\hat{\gamma} \in \beta$ s.t. $\hat{\gamma} \preceq \gamma$. Therefore, also the second condition of Def. 1 is satisfied and we get $\beta \lhd \gamma$. $\square$

Proof of Lemma 8

We assume the contrary and derive a contradiction. If Algorithm 1 does not terminate then, in the graph of the game $\mathcal{G}^{out}$, it will build an infinite sequence of states $q_0, \ldots, q_k, \ldots$ such that, for all $i, j \in \mathbb{N}$, the following properties hold: $i < j$ implies

(a) $(q_i, q_j) \in (T^{out})^*$, and,
(b) $\lambda(q_i) \neq \lambda(q_j)$, and,

(c) $\lambda(q_i) \not\preceq \lambda(q_j)$.

The property (a) comes from the way we build the transition relation when adding new state to the set to `ToExplore` at Line 17 of the algorithm (and from the fact that the VASS is finitely branching and hence so is the graph of the game $\mathcal{G}^{out}$). The property (b) is deduced thanks to the test at Line 12 that necessarily fails infinitely often, otherwise the algorithm would terminate. The property (c) is obtained thanks to the test at Line 9 which must also fail infinitely often if the algorithm does not terminate. Since the number of counters is fixed, the set $(\Gamma^{\mathcal{C}}, \preceq)$ is well-quasi-ordered by Dickson's Lemma. Hence in the infinite sequence of states $q_0, \ldots, q_k, \ldots$ there must appear two states q_i and q_j with $i < j$ such that $\lambda(q_i) \preceq \lambda(q_j)$, which is a contradiction to the conjunction of (b) and (c). This allows us to conclude that the Algorithm 1 necessarily terminates.

Proof of Lemma 9

We show both directions of the equivalence.

Left to right implication. If $[0, \mathcal{V}] : \gamma \models_{\mathcal{G}}$ `Parity` then there exists a concrete nonnegative configuration $\gamma_0 \in [\![\gamma]\!]$ with $\gamma_0 = \gamma \oplus \gamma'$ s.t. $[0, \mathcal{V}] : \gamma_0 \models_{\mathcal{G}}$ `Parity`, i.e., γ' assigns values to the counters in $\mathcal{C} - C$. Moreover, we have $\gamma^{out} = \langle q^{out}, \vartheta_{out} \rangle$ where $\lambda(q^{out}) = \gamma$ and $dom(\vartheta_{out}) = \emptyset$.

Using the winning strategy $\sigma_0 \in \Sigma_0^{\mathcal{V}}$ of Player 0 in $\mathcal{G}$ from γ_0, we will construct a winning strategy $\sigma_0' \in \Sigma_0^{\mathcal{E}}$ of Player 0 from a concrete configuration $\gamma_0' \in [\![\gamma^{out}]\!]$ in $\mathcal{G}^{out}$, where $\gamma_0' = \langle q^{out}, \mathtt{val}(\gamma') \rangle$. We do this by maintaining a correspondence between nonnegative configurations in both games and between the used sequences of transitions. Let $\pi = \gamma_0 \xrightarrow{t_1}_{\mathcal{V}} \gamma_1 \xrightarrow{t_2}_{\mathcal{V}} \ldots \gamma_n$ a partial play in $\mathcal{G}$, and $\pi' = \gamma_0' \xrightarrow{t_1'}_{\mathcal{E}} \gamma_1' \xrightarrow{t_2'}_{\mathcal{E}} \ldots \gamma_n'$ a partial play in $\mathcal{G}^{out}$.

We will define σ_0' to ensure that either the following invariant holds for all $i \geq 0$ or **Condition 2** holds for some γ_n' and the invariant holds for all $i \leq n$.

1. $\lambda(t_i') = t_i$
2. $\lambda(\mathtt{state}(\gamma_i')) = \gamma_i|C$
3. $\mathtt{val}(\gamma_i') = \gamma_i|(\mathcal{C} - C)$
4. $\kappa(\gamma_i') = \kappa(\gamma_i)$

These conditions are satisfied for the initial states at $i = 0$, since $\lambda(\mathtt{state}(\gamma_0')) = \lambda(q^{out}) = \gamma = \gamma_0|C$, $\mathtt{val}(\gamma_0') = \mathtt{val}(\gamma') = \gamma_0|(\mathcal{C} - C)$ and $\kappa(\gamma_0') = \kappa(q^{out}) = \kappa(\gamma) = \kappa(\gamma_0)$.

For the step we choose $\sigma_0'(\pi') := t_{n+1}'$ s.t. $\lambda\left(t_{n+1}'\right) = t_{n+1} = \sigma_0(\pi)$ which maintains the invariant.

It cannot happen that **Condition 1** holds in π'. All visited nonnegative configurations γ_i in the winning play π are also winning for Player 0. By Lemma 7 (item 3), we have $\beta \lhd \gamma_i|C$ and thus $\beta \lhd \gamma_i|C = \lambda(\mathtt{state}(\gamma_i'))$ so that **Condition 1** is false at γ_i'.

Since $\mathcal{G}$ is a VASS-game, we have $\gamma_i \geq 0$ for all $i \geq 0$. Therefore $\lambda(\mathtt{state}(\gamma_i')) = \gamma_i|C \geq 0$ and $\mathtt{val}(\gamma_i') = \gamma_i|(\mathcal{C} - C) \geq 0$. Thus the same transitions are possible in $\mathcal{G}^{out}$ as in $\mathcal{G}$.

In the case where **Condition 2** eventually holds in $\mathcal{G}^{out}$, Player 0 trivially wins the game in $\mathcal{G}^{out}$. Otherwise we have $\mathtt{val}(\gamma_i') = \gamma_i|(\mathcal{C}-C) \geq 0$ for all $i \geq 0$ and thus the nonnegativity condition $\Box\overline{\mathtt{neg}}$ of $\mathcal{G}^{out}$ is satisfied by π'.

Finally, since the parity condition is satisfied by π and $\kappa(\gamma_i) = \kappa(\gamma_i')$, the parity condition is also satisfied by π'. Therefore σ_0' is winning for Player 0 in $\mathcal{G}^{out}$ from $\gamma_0' \in [\![\gamma^{out}]\!]$ and thus we obtain $[0,\mathcal{E}] : \gamma^{out} \models_{\mathcal{G}^{out}} \mathtt{Parity} \wedge \Box\overline{\mathtt{neg}}$ as required.

Right to left implication. If $[0,\mathcal{E}] : \gamma^{out} \models_{\mathcal{G}^{out}} \mathtt{Parity} \wedge \Box\overline{\mathtt{neg}}$ then there exists a concrete nonnegative configuration $\gamma' \in [\![\gamma^{out}]\!]$ s.t. $[0,\mathcal{E}] : \gamma' \models_{\mathcal{G}^{out}} \mathtt{Parity} \wedge \Box\overline{\mathtt{neg}}$. Due to the concreteness of γ' and the $\Box\overline{\mathtt{neg}}$ property, we also have $[0,\mathcal{V}] : \gamma' \models_{\mathcal{G}^{out}} \mathtt{Parity}$. Thus Player 0 has a winning strategy σ_0 in the VASS parity game on $\mathcal{G}^{out}$ from the concrete nonnegative configuration γ'.

Using σ_0, we will construct a winning strategy σ_0' for Player 0 in the VASS parity game on $\mathcal{G}$ from some nonnegative configuration $\gamma_0 \in [\![\gamma]\!]$. Let $\gamma_0 = \gamma \oplus \gamma''$, where γ'' is some yet to be constructed function assigning sufficiently high values to counters in $\mathcal{C}-C$. We only prove the sufficient condition that a winning strategy σ_0' exists, but do not construct a Turing machine that implements it. This is because σ_0' uses the numbers $v(\hat{\gamma})$ and $u(\hat{\gamma})$ from Lemma 7 that are not computed here.

In order to construct σ_0' and γ'', we need some definitions. Consider a sequence of transitions in $\mathcal{G}^{out}$ that leads from q to q' ending with **Condition 2** at line 9 in the algorithm. We call this sequence a *pumping sequence*. Its effect is nonnegative on all counters in C and strictly increasing in at least one of them, although its effect may be negative on counters in $\mathcal{C}-C$. Due to the finiteness of $\mathcal{G}^{out}$ (by Lemma 8), the number of different pumping sequences is bounded by some number p and their maximal length is bonded by some number l. For the given finite $\beta = \bigcup_{C' \subseteq C, |C'|=|C|-1} \mathtt{Pareto}[\mathcal{G},\mathcal{V},0,C'](\mathtt{Parity})$ we use the constants from Lemma 7 to define the following finite upper bounds $v := max(\{v(\hat{\gamma}) \mid \hat{\gamma} \in \beta\})$ and $u := max(\{u(\hat{\gamma}) \mid \hat{\gamma} \in \beta\})$.

Now we define σ_0'. The intuition is as follows. Either the current nonnegative configuration is already known to be winning for Player 0 by induction hypothesis (if the current nonnegative configuration is sufficiently large compared to nonnegative configurations in $[\![\beta]\!]$) in which case he plays according to his known winning strategy from the induction hypothesis. Otherwise, for a given history π in $\mathcal{G}$, Player 0 plays like for a history π' in $\mathcal{G}^{out}$, where π' is derived from π as follows. For π' we first use a sequence of transitions in $\mathcal{G}^{out}$ whose labels (see line 13 of the algorithm) correspond to the sequence of transitions in π, but then we remove all subsequences from π' which are pumping sequences in $\mathcal{G}^{out}$. Thus Player 0 plays from nonnegative configurations in $\mathcal{G}$ that are possibly larger than the corresponding (labels of) nonnegative configurations in $\mathcal{G}^{out}$ on the counters in C. The other counters in $\mathcal{C}-C$ might differ between the games and will have to be chosen sufficiently high by the initial γ'' to stay positive during the game (see below). We show that the history of the winning game in $\mathcal{G}$ will contain only finitely many such pumping sequences, and thus finite initial values (encoded in γ'') for the counters in $\mathcal{C}-C$ will suffice to win the game.

Let $\pi = \gamma_0 \xrightarrow{t_1} \gamma_1 \xrightarrow{t_2} \cdots \gamma_n$ be a path in $\mathcal{G}$, where Player 0 played according to strategy σ_0'. Our strategy σ_0' will maintain the invariant that π induces a sequence of states $\hat{\pi} = q_0, q_1, \ldots, q_n$ in $\mathcal{G}^{out}$. The sequence $\hat{\pi}$ is almost like a path in $\mathcal{G}^{out}$ with transitions whose label is the same as the transitions in π, except that it contains back-jumps to previously visited states whenever a pumping sequence is completed.

Let $q_0 = q^{out}$. For the step from q_i to q_{i+1} there are two cases. For a given transition t_i in $\mathcal{G}$ appearing in π there is a unique transition t'_i in $\mathcal{G}^{out}$ with $\lambda(t'_i) = t_i$. As an auxiliary construction we define the state q'_{i+1}, which is characterized uniquely by $\lambda\left(q'_{i+1}\right) = t'_i(\lambda(q_i))$. If there is a $j \leq i$ s.t. the sequence from q_j to q'_{i+1} is a pumping sequence and q_j is not part of a previously identified pumping sequence (the construction ensures that there can be at most one such j), then let $q_{i+1} := q_j$, i.e., we jump back to the beginning of the pumping sequence. Otherwise, if no pumping sequence is completed at q'_{i+1}, then let $q_{i+1} = q'_{i+1}$, so that we have $q_i \xrightarrow{t'_i} q_{i+1}$. From the sequence $\hat{\pi}$ we obtain a genuine path π' in $\mathcal{G}^{out}$ by deleting all pumping sequences from $\hat{\pi}$.

In the case where γ_n belongs to Player 0 we define $\sigma'_0(\pi)$ by case distinction.

1. We let $\sigma'_0(\pi) := t_i$ where $\lambda(t_i) = \sigma_0(\pi')$, except when the condition of the following case 2 holds.
2. By $\lambda(q_0) = \gamma$ and $\gamma_0 = \gamma \oplus \gamma'$ we have $\lambda(q_0) \preceq \gamma_0|C$. Since the effects of the sequences of transitions in π and $\hat{\pi}$ are the same, and pumping sequences have a nondecreasing effect on the counters in C, we obtain $\lambda(q_i) \preceq \gamma_i|C$ for all $i \geq 0$.
   Since σ_0 is winning in $\mathcal{G}^{out}$ we have $\beta \lhd \lambda(q_i)$ and thus $\beta \lhd \gamma_i|C$. By Def. 1, there exists some $\hat{\gamma} \in \beta$ and counter $c \notin dom(\hat{\gamma})$ s.t. $\hat{\gamma} \preceq \gamma_i|C[c \leftarrow \perp]$.
   Condition for case 2: If $\gamma_i(c) \geq v(\hat{\gamma})$ and $\gamma_i(c') \geq u(\hat{\gamma})$ for every counter $c' \in \mathcal{C} - C$ then, by Lemma 7 (items 1 and 2) and monotonicity (Lemma 1), Player 0 has a winning strategy σ''_0 from γ_i. In this case σ'_0 henceforth follows this winning strategy σ''_0.

Now we show that σ'_0 is winning for Player 0 in $\mathcal{G}$ from the initial nonnegative configuration $\gamma_0 = \gamma \oplus \gamma'$ for some sufficiently large but finite γ'. We distinguish two cases, depending on whether case 2 above is reached or not.

If Case 2 is reached: Consider the case where condition 2 above holds at some reached game nonnegative configuration γ_n. Every pumping sequence α has nondecreasing effect on all counters in C and strictly increases at least some counter $c_\alpha \in C$. Thus if $\hat{\pi}$ contains the pumping sequence α at least v times, then $\gamma_n(c_\alpha) - \lambda(q_n)(c_\alpha) \geq v$ and in particular $\gamma_n(c_\alpha) \geq v$. If additionally, γ_n is sufficiently large on the counters outside C, i.e., $\gamma_n(c') \geq u$ for every counter $c' \in \mathcal{C} - C$, then case 2 above applies and the winning strategy σ''_0 takes over.

The path π (resp. $\hat{\pi}$) can contain at most $v * p$ pumping sequences of a combined length that is bounded by $v * p * l$ before the first condition $\gamma_n(c_\alpha) \geq v$ becomes true for some pumping sequence α. In this case it is sufficient for σ''_0 to win if the values in the counters in $\mathcal{C} - C$ are $\geq u$ at nonnegative configuration γ_n. How large does a counter $c' \in \mathcal{C} - C$ need to be at the (part of the) initial nonnegative configuration γ'' in order to satisfy this additional condition later at γ_n? Since σ_0 is winning in the VASS game from γ' in $\mathcal{G}^{out}$, an initial value $\gamma'(c')$ is sufficient to keep the counter c' above 0 in the game on $\mathcal{G}^{out}$. Thus an initial value of $\gamma'(c') + u$ is sufficient to keep the counter c' above u in the game on $\mathcal{G}^{out}$. Moreover, the game played according to σ'_0 in $\mathcal{G}$ contains the same transitions (modulo the labeling $\lambda(\ldots)$) as the game played according to σ_0 on $\mathcal{G}^{out}$, except for the $\leq v * p * l$ extra transitions in pumping sequences. Since a single transition can decrease a counter by at most one, an initial counter value of $\gamma''(c') = \gamma'(c') + u + v * p * l$ is sufficient in order to have $c' \geq u$ whenever case 2 applies and then σ_0'' (and thus σ_0') is winning for Player 0. The counters in C are always large enough by construction, since $\lambda(q_i) \preceq \gamma_i|C$ for all $n \geq i \geq 0$. The parity objective is satisfied by σ_0', since it is satisfied by σ_0'' on the infinite suffix of the game.

If Case 2 is not reached: Otherwise, if case 2 is not reached, then the VASS game on $\mathcal{G}$ played according to σ_0' is like the VASS game on $\mathcal{G}^{out}$ played according to σ_0, except for the finitely many interludes of pumping sequences, of which there are at most $p * v$ (with a combined length $\leq v * p * l$). Since σ_0 is winning the VASS game on $\mathcal{G}^{out}$ from γ', this keeps the counters nonnegative. At most $v * p * l$ extra transitions happen in $\mathcal{G}$ (in the pumping sequences) and a single transition can decrement a counter by at most one. Thus it is sufficient for staying nonnegative in $\mathcal{G}$ if $\gamma''(c') \geq \gamma'(c') + v * p * l$ for all $c' \in \mathcal{C} - C$. The counters in $c \in C$ trivially stay nonnegative, since $\lambda(q_i) \preceq \gamma_i|C$ for all $i \geq 0$. The parity objective is satisfied, since the colors of the nonnegative configurations γ_i and q_i in π and $\hat{\pi}$ coincide, the colors of an infinite suffix of $\hat{\pi}$ coincide with the colors of an infinite suffix of π' and π' satisfies the parity objective as σ_0 is winning in $\mathcal{G}^{out}$.

Combination of the cases. While σ_0' might not be able to enforce either of the two cases described above, one of them will certainly hold in any play. We define the (part of the) initial nonnegative configuration γ'' to be sufficiently high to win in either case, by taking the maximum of the requirements for the cases.

We let $\gamma''(c') := \gamma'(c) + u + v * p * l$ for all $c' \in \mathcal{C} - C$ and obtain that σ_0' is a winning strategy for Player 0 in the parity game on $\mathcal{G}$ from the initial nonnegative configuration $\gamma_0 = \gamma \oplus \gamma' \in [\![\gamma]\!]$. Thus $[0, \mathcal{V}] : \gamma \models_{\mathcal{G}} \texttt{Parity}$, as required. $\square$

Proof of Theorem 5

Given a labeled finite-state system $\left\langle S, \xrightarrow{a}, \Sigma \right\rangle$ and a labeled VASS $\langle Q, T, \Sigma, \lambda \rangle$ with initial states s_0 and $\langle q_0, \vartheta \rangle$, respectively, we construct a single-sided integer game $\mathcal{G} = \langle Q_0 \uplus Q_1, T', \kappa \rangle$ with initial configuration $\gamma = \langle \langle s_0, q_0, 1 \rangle, \vartheta \rangle$ s.t. $\langle q_0, \vartheta \rangle$ weakly simulates s_0 if and only if $[0, \mathcal{V}] : \gamma \models_{\mathcal{G}} \texttt{Parity}$. Then decidability follows from Theorem 3.

Let $Q_1 = \{\langle s, q, 1 \rangle \mid s \in S, q \in Q\} \cup \{win_0\}$ and $Q_0 = \{\langle s, q, 0 \rangle \mid s \in S, q \in Q\} \cup \{\langle s, q^a, 0 \rangle \mid s \in S, q \in Q, a \in \Sigma\} \cup \{lose_0\}$. Let $\kappa(Q_1) = 2$ and $\kappa(Q_0) = 1$, i.e., Player 0 wins the parity game iff states belonging to Player 1 are visited infinitely often.

Now we define T'. For every finite-state system transition $s \xrightarrow{a} s'$ and every $q \in Q$, we add a transition $\langle \langle s, q, 1 \rangle, nop, \langle s', q^a, 0 \rangle \rangle$. Here the state q^a encodes the choice of the symbol a by Player 1, which restricts the future moves of Player 0. For every VASS transition $t = \langle q_1, op, q_2 \rangle \in T$ with label $\lambda(t) = \tau$ and every $s \in S, a \in \Sigma$ we add a transition $\langle \langle s, q_1^a, 0 \rangle, op, \langle s, q_2^a, 0 \rangle \rangle$. This encodes the first arbitrarily long sequence of τ-moves in the Player 0 response of the form $\tau^* a \tau^*$. For every VASS transition $t = \langle q_1, op, q_2 \rangle \in T$ with label $\lambda(t) = a \neq \tau$ and $s \in S$ we add a transition $\langle \langle s, q_1^a, 0 \rangle, op, \langle s, q_2, 0 \rangle \rangle$. This encodes the a-step in in the Player 0 response of the form $\tau^* a \tau^*$. Moreover, we add transitions $\langle \langle s, q^\tau, 0 \rangle, nop, \langle s, q, 0 \rangle \rangle$ for all $s \in S, q \in Q$ (since a τ-move in the weak simulation game does not strictly require a response step). For every VASS transition $t = \langle q_1, op, q_2 \rangle \in T$ with label $\lambda(t) = \tau$ and $s \in S$ we add a transition

$\langle\langle s,q_1,0\rangle, op, \langle s,q_2,0\rangle\rangle$. This encodes the second arbitrarily long sequence of τ-moves in the Player 0 response of the form $\tau^* a \tau^*$. Finally, for all $s \in S, q \in Q$ we add transitions $\langle\langle s,q,0\rangle, nop, \langle s,q,1\rangle\rangle$. Here Player 0 switches the control back to Player 1. He cannot win by delaying this switch indefinitely, because the color of the states in Q_0 is odd.

The following transitions encode the property of the simulation game that a player loses if he gets stuck. For every state in $q \in Q_1$ with no outgoing transitions we add a transition $\langle q, nop, win_0\rangle$. In particular this creates a loop at state win_0. Since the color of win_0 is even, this state is winning for Player 0. For every state in $q \in Q_0$ with no outgoing transitions we add a transition $\langle q, nop, lose_0\rangle$. In particular this creates a loop at state $lose_0$. Since the color of $lose_0$ is odd, this state is losing for Player 0.

This construction yields a single-sided integer game, since all transitions from states in Q_1 have operation *nop*.

A round of the weak simulation game is encoded by the moves of the players between successive visits to a state in Q_1. A winning strategy for Player 0 in the weak simulation game directly induces a winning strategy for Player 0 in the parity game $\mathcal{G}$, since the highest color that is infinitely often visited is 2, and thus $[0,\mathcal{V}] : \gamma \models_{\mathcal{G}} \texttt{Parity}$. Conversely, $[0,\mathcal{V}] : \gamma \models_{\mathcal{G}} \texttt{Parity}$ implies a winning strategy for Player 0 in the parity game on $\mathcal{G}$ which ensures that color 2 is seen infinitely often. Therefore, states in Q_1 are visited infinitely often. Thus, either infinitely many rounds of the weak simulation game are simulated or state win_0 is reached in $\mathcal{G}$ and Player 1 gets stuck in the weak simulation game. In either case, Player 0 wins the weak simulation game and $\langle q_0, \vartheta\rangle$ weakly simulates s_0. $\square$

Semantics of L_μ^{pos}

The syntax of the positive μ-calculus L_μ^{pos} is given by the following grammar: $\phi ::= q \mid X \mid \phi \wedge \phi \mid \phi \vee \phi \mid \Diamond\phi \mid \Box\phi \mid \mu X.\phi \mid \nu X.\phi$ where $q \in Q$ and X belongs to a countable set of variables $\mathcal{X}$.

Free and bound occurrences of variables are defined as usual. We assume that no variable has both bound and free occurrences in some ϕ, and that no two fixpoint subterms bind the same variable (this can always be ensured by renaming a bound variable). A formula is closed if it has no free variables. Without restriction, we do not use any negation in our syntax. Negation can be pushed inward by the usual dualities of fixpoints, and the negation of an atomic proposition referring to a control-state can be expressed by a disjunction of propositions referring to all the other control-states.

We now give the interpretation over the VASS $\langle Q,T\rangle$ of a formula of L_μ^{pos} according to an environment $\rho : \mathcal{X} \to 2^\Gamma$ which associates to each variable a subset of concrete configurations. Given ρ, a formula $\phi \in L_\mu^{pos}$ represents a subset of concrete configurations, denoted by $[\![\phi]\!]_\rho$ and defined inductively as follows.

$$
\begin{array}{ll}
[\![q]\!]_\rho & = \{\gamma \in \Gamma \mid \mathtt{state}(\gamma) = q\} \\
[\![X]\!]_\rho & = \rho(X) \\
[\![\phi \wedge \psi]\!]_\rho & = [\![\phi]\!]_\rho \cap [\![\psi]\!]_\rho \\
[\![\phi \vee \psi]\!]_\rho & = [\![\phi]\!]_\rho \cup [\![\psi]\!]_\rho \\
[\![\Diamond\phi]\!]_\rho & = \left\{\gamma \in \Gamma \mid \exists \gamma' \in [\![\phi]\!]_\rho \text{ s.t. } \gamma \longrightarrow_{\mathcal{V}} \gamma'\right\} \\
[\![\Box\phi]\!]_\rho & = \left\{\gamma \in \Gamma \mid \forall \gamma' \in \Gamma, \gamma \longrightarrow_{\mathcal{V}} \gamma' \text{ implies } \gamma' \in [\![\phi]\!]_\rho\right\} \\
[\![\mu X.\phi]\!]_\rho & = \bigcap\left\{\Gamma' \subseteq \Gamma \mid [\![\phi]\!]_{\rho[X\leftarrow\Gamma']} \subseteq \Gamma'\right\} \\
[\![\nu X.\phi]\!]_\rho & = \bigcup\left\{\Gamma' \subseteq \Gamma \mid \Gamma' \subseteq [\![\phi]\!]_{\rho[X\leftarrow\Gamma']}\right\}
\end{array}
$$

where the notation $\rho[X \leftarrow \Gamma']$ is used to define an environment equal to ρ on every variable except on X where it returns Γ'. We recall that $(2^\Gamma, \subseteq)$ is a complete lattice and that, for every $\phi \in L_\mu^{pos}$ and every environment ρ, the function $G : 2^\Gamma \mapsto 2^\Gamma$, which associates to $\Gamma' \subseteq \Gamma$ the set $G(\Gamma') = [\![\phi]\!]_{\rho[X\leftarrow\Gamma']}$, is monotonic. Hence, by the Knaster-Tarski Theorem, the set $[\![\mu X.\phi]\!]_\rho$ (resp. $[\![\nu X.\phi]\!]_\rho$) is the least fixpoint (resp. greatest fixpoint) of G, and it is well-defined. Finally we denote by $[\![\phi]\!]$ the subset of configurations $[\![\phi]\!]_{\rho_0}$ where ρ_0 is the environment which assigns the empty set to each variable.

Proof of Lemma 10

We consider a VASS $\mathcal{S} = \langle Q, T\rangle$ and ϕ a formula in L_μ^{pos}. We will use in this proof the set of subformulae of ϕ, denoted by $sub(\phi)$. For formulae in L_μ^{pos} we assume that no variable is bounded by the same fixpoint. Hence given a formula ϕ and a bounded variable $X \in \mathcal{X}$, we can determine uniquely the subformula of ϕ that bounds the variable X; such a formula will be denoted by ϕ_X. We also denote by $free(\phi)$ the set of free variables in ϕ. The integer game $\mathcal{G}(\mathcal{S}, \phi) = \langle Q', T', \kappa\rangle$ is built as follows:

- $Q' = Q \times sub(\phi)$
- The transition relation T' is the smallest set respecting the following conditions for all the formulae $\psi \in sub(\phi)$:
  - If $\psi = q$ with $q \in Q$, then $\langle\langle q', \psi\rangle, nop, \langle q', \psi\rangle\rangle$ belongs to T' for all states q' in Q;
  - If $\psi = X$ with $X \in \mathcal{X}$ and $X \notin free(\phi)$, then $\langle\langle q, \psi\rangle, nop, \langle q, \phi_X\rangle\rangle$ belongs to T' for all states q in Q;
  - If $\psi = X$ with $X \in \mathcal{X}$ and $X \in free(\phi)$, then $\langle\langle q, \psi\rangle, nop, \langle q, \psi\rangle\rangle$ belongs to T' for all states q in Q;
  - If $\psi = \psi' \wedge \psi''$ or $\psi = \psi' \vee \psi''$ then $\langle\langle q, \psi\rangle, nop, \langle q, \psi'\rangle\rangle$ and $\langle\langle q, \psi\rangle, nop, \langle q, \psi''\rangle\rangle$ belong to T' for all states q in Q;
  - If $\psi = \Diamond\psi'$ or $\psi = \Box\psi'$ then for all states $q \in Q$ and for all transitions $\langle q, op, q'\rangle \in T$, we have $\langle\langle q, \psi\rangle, op, \langle q', \psi'\rangle\rangle$ in T';
  - If $\psi = \mu X.\psi'$ or $\psi = \nu X.\psi'$, then $\langle\langle q, \psi\rangle, nop, \langle q, \psi'\rangle\rangle$ for all states $q \in Q$.
- A state $\langle q, \psi\rangle$ belongs to Q'_0 if and only if:
  - $\psi = q'$ with $q' \in Q$, or,
  - $\psi = X$ with $X \in \mathcal{X}$, or,
  - $\psi = \psi' \vee \psi''$, or,

- $\psi = \Diamond\psi'$, or,
- $\psi = \mu X.\psi''$, or,
- $\psi = \nu X.\psi''$.

– A state $\langle q,\psi\rangle$ belongs to Q'_1 if and only if:
  - $\psi = \psi' \wedge \psi''$, or,
  - $\psi = \Box\psi'$.

– The coloring function κ is then defined as follows:
  - for all $q,q' \in Q$, if $q' = q$ then $\kappa\langle q,q'\rangle = 0$ and if $q' \neq q$ then $\kappa\langle q,q'\rangle = 1$;
  - for all $q \in Q$ and all $X \in free(\phi)$, $\kappa\langle q,X\rangle = 1$;
  - for all $q \in Q$, for all subformulae $\psi \in sub(\phi)$ if $\psi \neq \mu X.\psi''$ and $\psi \neq \nu X.\psi''$ and $\psi \neq q'$ with $q' \in Q$ and $\psi \neq X$ with $X \in free(\phi)$, then $\kappa\langle q,\psi\rangle = 0$;
  - for all $q \in Q$, for all subformulae $\psi \in sub(\phi)$ such that $\psi \neq \mu X.\psi''$, $\kappa\langle q,\psi\rangle = m$ where m is the smallest odd number greater or equal to the alternation depth of ψ;
  - for all $q \in Q$, for all subformulae $\psi \in sub(\phi)$ such that $\psi \neq \mu X.\psi''$, $\kappa\langle q,\psi\rangle = m$ where m is the smallest even number greater or equal to the alternation depth of ψ;

Before providing the main property of the game $\mathcal{G}(\mathcal{S},\phi)$, we introduce a new winning condition which will be useful in the sequel of the proof. This winning condition uses an environment $\rho : \mathcal{X} \to 2^\Gamma$ and is given by the formula $\texttt{Parity} \vee \bigvee_{X\in free(\phi)} \Diamond(X \wedge \rho(X))$ where $X \wedge \rho(X)$ holds in the configurations of the form $\langle\langle q,X\rangle,\vartheta\rangle$ such that $\langle q,\vartheta\rangle \in \rho(X)$. It states that a run is winning if it respects the parity condition or if at some point it encounters a configuration of the form $\langle\langle q,X\rangle,\vartheta\rangle$ with $X \in free(\phi)$ and $\langle q,\vartheta\rangle \in \rho(X)$. We denote by $\texttt{Cond}(\phi,\rho)$ the formula $\bigvee_{X\in free(\phi)} \Diamond(X \wedge \rho(X))$.

We will now prove the following property: for all formulae ϕ in L^{pos}_μ, for all concrete configurations $\gamma = \langle q,\vartheta\rangle$ of $\mathcal{S}$ and all environments $\rho : \mathcal{X} \to 2^\Gamma$, we have $\gamma \in [\![\phi]\!]_\rho$ iff $[0,\mathcal{V}] : \langle\langle q,\phi\rangle,\vartheta\rangle \models_{\mathcal{G}(\mathcal{S},\phi)} \texttt{Parity} \vee \texttt{Cond}(\phi,\rho)$, i.e., iff $\langle\langle q,\phi\rangle,\vartheta\rangle \in \mathcal{W}[\mathcal{G}(\mathcal{S},\phi),\mathcal{V},0,\mathcal{C}](\texttt{Parity} \vee \texttt{Cond}(\phi,\rho))$.

We reason by induction on the length of ϕ. For the base case with $\phi = q$ with $q \in Q$ or $\phi = X$ with $X \in \mathcal{X}$ the property trivially holds. We then proceed with the induction reasoning. It is easy to prove that the property holds for formulae of the form $\phi' \wedge \phi''$ or $\phi' \vee \phi''$ if the property holds for ϕ' and ϕ'' and the same for formulae of the form $\Diamond\phi'$ and $\Box\phi'$. We consider now a formula ϕ of the form $\mu X.\psi$ and assume that the property holds for the formula ψ. Let $G : 2^\Gamma \mapsto 2^\Gamma$ be the function which associates to any subset of configurations Γ' the set $G(\Gamma') = [\![\psi]\!]_{\rho[X\leftarrow\Gamma']}$. By induction hypothesis we have $\langle q,\vartheta\rangle \in G(\Gamma')$ iff $\langle\langle q,\psi\rangle,\vartheta\rangle \in \mathcal{W}[\mathcal{G}(\mathcal{S},\psi),\mathcal{V},0,\mathcal{C}](\texttt{Parity} \vee \texttt{Cond}(\psi,\rho[X\leftarrow\Gamma']))$. We denote by μG the least fixpoint of G. We want to prove that $\langle q,\vartheta\rangle \in \mu G$ iff $\langle\langle q,\mu X.\psi\rangle,\vartheta\rangle \in \mathcal{W}[\mathcal{G}(\mathcal{S},\mu X.\psi),\mathcal{V},0,\mathcal{C}](\texttt{Parity} \vee \texttt{Cond}(\mu X.\psi,\rho))$. We define the following set of configurations $\Gamma_\mu = \big\{\langle q,\vartheta\rangle \in \Gamma \mid \langle\langle q,\mu X.\psi\rangle,\vartheta\rangle \in \mathcal{W}[\mathcal{G}(\mathcal{S},\mu X.\psi),\mathcal{V},0,\mathcal{C}](\texttt{Parity} \vee \texttt{Cond}(\mu X.\psi,\rho))\big\}$. So finally what we want to prove is that $\mu G = \Gamma_\mu$.

– We begin by proving that $\mu G \subseteq \Gamma_\mu$. By definition $\mu G = \bigcap\{\Gamma' \subseteq \Gamma \mid G(\Gamma') \subseteq \Gamma'\}$. Hence it is enough to prove that $G(\Gamma_\mu) \subseteq \Gamma_\mu$. Let $\langle q,\vartheta\rangle \in G(\Gamma_\mu)$. This means that $\langle\langle q,\psi\rangle,\vartheta\rangle \in \mathcal{W}[\mathcal{G}(\mathcal{S},\psi),\mathcal{V},0,\mathcal{C}](\texttt{Parity} \vee \texttt{Cond}(\psi,\rho[X\leftarrow\Gamma_\mu]))$

by definition of G. We want to prove that $\langle\langle q, \mu X.\psi\rangle, \vartheta\rangle \in \mathcal{W}[\mathcal{G}(\mathcal{S},\psi), \mathcal{V}, 0, \mathcal{C}](\texttt{Parity} \vee \texttt{Cond}(\mu X.\psi, \rho))$. First note that the configuration $\langle\langle q, \mu X.\psi\rangle, \vartheta\rangle$ belongs to Player 0, and from this configuration, Player 0 has a unique choice which is to go to the state $\langle\langle q, \psi\rangle, \vartheta\rangle$. Then from $\langle\langle q, \psi\rangle, \vartheta\rangle$, if Player 0 plays as in the game $\mathcal{G}(\mathcal{S},\psi)$ where it has a winning strategy, there are two options:

1. a control-state of the form $\langle q', X\rangle$ is never encountered and in that case Player 0 wins because it was winning in $\mathcal{G}(\mathcal{S},\psi)$ and the run performed is the same;
2. a control-state of the form $\langle q', X\rangle$ is encountered, but in that case, Player 0 is necessarily in a configuration $\langle\langle q', X\rangle, \vartheta\rangle$ with $\langle q', \vartheta\rangle \in \rho[X \leftarrow \Gamma_\mu](X)$ (by definition of the winning condition in $\mathcal{G}(\mathcal{S},\psi)$), ie with $\langle q', \vartheta\rangle \in \Gamma_\mu$. But this means that from this configuration, Player 0 has a winning strategy for the game $\mathcal{G}(\mathcal{S},\phi)$.

Hence we have shown that $\langle q, \vartheta\rangle \in \Gamma_\mu$ and consequently $G(\Gamma_\mu) \subseteq \Gamma_\mu$. This allows us to deduce that $\mu G \subseteq \Gamma_\mu$.

- We will now prove that $\Gamma_\mu \subseteq \mu G$. For this we will prove that for all $\Gamma' \subseteq \Gamma$ such that $G(\Gamma') = \Gamma'$, we have $\Gamma_\mu \subseteq \Gamma'$. This will in fact imply that $\Gamma_\mu \subseteq \mu G$, since μG is the least fixpoint of the function G. Let $\Gamma' \subseteq \Gamma$ such that $G(\Gamma') = \Gamma'$ and let $\langle q, \vartheta\rangle \in \Gamma_\mu$. We reason *by contradiction* and assume that $\langle q, \vartheta\rangle \notin \Gamma'$. Since $\langle q, \vartheta\rangle \in \Gamma_\mu$, this means that Player 0 has a winning strategy to win in the game $\mathcal{G}(\mathcal{S}, \mu X.\psi))$ from the configuration $\langle\langle q, \mu X.\psi\rangle, \vartheta\rangle$ with the objective $\texttt{Parity} \vee \texttt{Cond}(\mu X.\psi, \rho)$. Since $\langle q, \vartheta\rangle \notin \Gamma' = G(\Gamma')$, this means that there is no winning strategy for Player 0 in the game $\mathcal{G}(\mathcal{S},\psi)$ from configuration $\langle\langle q, \psi\rangle, \vartheta\rangle$ with the objective $\texttt{Parity} \vee \texttt{Cond}(\psi, \rho[X \leftarrow \Gamma'])$. Since Player 0 has a winning strategy to win in the game $\mathcal{G}(\mathcal{S}, \mu X.\psi)$, we can adapt this strategy to the game $\mathcal{G}(\mathcal{S},\psi)$ (by restricting it to the path possible in this game and beginning one step later). But since this strategy is not winning in the game $\mathcal{G}(\mathcal{S},\psi)$ with the objective $\texttt{Parity} \vee \texttt{Cond}(\psi, \rho[X \leftarrow \Gamma'])$, it means that there is a path π_0 in $\mathcal{G}(\mathcal{S},\psi)$ that respects this strategy and this path necessarily terminates in a state of the form $\langle\langle q_1, X\rangle, \vartheta_1\rangle$ with $\langle q_1, \vartheta_1\rangle \notin \rho[X \leftarrow \Gamma'](X)$, i.e., with $\langle q_1, \vartheta_1\rangle \notin \Gamma'$ (otherwise this strategy which is winning in $\mathcal{G}(\mathcal{S}, \mu X.\psi)$ would also be winning in $\mathcal{G}(\mathcal{S},\psi)$). On the other hand, in $\mathcal{G}(\mathcal{S}, \mu X.\psi)$, $\langle\langle q_1, X\rangle, \vartheta_1\rangle$ has a unique successor which is $\langle\langle q_1, \mu X.\psi\rangle, \vartheta_1\rangle$ and from which Player 0 has a winning strategy since we have followed a winning strategy in the game $\mathcal{G}(\mathcal{S}, \mu X.\psi)$ that has lead us to that configuration. Hence we have $\langle q_1, \vartheta_1\rangle \notin \Gamma'$ and $\langle q_1, \vartheta_1\rangle \in \Gamma_\mu$. So from $\langle q_1, \vartheta_1\rangle$ we can perform a similar reasoning following the winning strategy in $\mathcal{G}(\mathcal{S}, \mu X.\psi)$ to reach a configuration $\langle\langle q_2, X\rangle, \vartheta_2\rangle$ such that $\langle q_2, \vartheta_2\rangle \notin \Gamma'$ and $\langle q_2, \vartheta_2\rangle \in \Gamma_\mu$. Finally, by performing the same reasoning we succeed in building an infinite play in $\mathcal{G}(\mathcal{S}, \mu X.\psi)$ which follows a winning strategy and such that the sequence of the visited configurations is of the form:

$$\langle\langle q, \mu X.\psi\rangle, \vartheta\rangle \ldots \langle\langle q_1, \mu X.\psi\rangle, \vartheta_1\rangle \ldots \langle\langle q_2, \mu X.\psi\rangle, \vartheta_2\rangle \ldots$$

Note that for all $i \geq 1$, $\kappa(\langle\langle q_i, \mu X.\psi\rangle, \vartheta_i\rangle)$ is the maximal priority in the game $\mathcal{G}(\mathcal{S}, \mu X.\phi)$ and it is odd by definition of the game. This means that the path we obtain following a winning strategy for Player 0 is losing, which is a contradiction. Hence we have $\langle q, \vartheta\rangle \in \Gamma'$. From this we deduce that $\Gamma_\mu \subseteq \mu G$.

If we consider a formula ϕ of the form $\nu X.\psi$, a reasoning similar to the previous one can be performed in order to show that the property holds.

Thanks to the previous proof, for all formulae ϕ in L_μ^{pos}, for all concrete configurations $\gamma = \langle q, \vartheta \rangle$ of $\mathcal{S}$, we have $\gamma \in [\![\phi]\!]_{\rho_0}$ iff $\langle \langle q, \phi \rangle, \vartheta \rangle \in \mathcal{W}[\mathcal{G}(\mathcal{S},\phi)), \mathcal{V}, 0, \mathcal{C}](\texttt{Parity} \vee \texttt{Cond}(\phi,\rho_0))$ where ρ_0 is the environment which assigns to each variable the empty set. This means that for all formulae ϕ in L_μ^{pos}, for all concrete configurations $\gamma = \langle q, \vartheta \rangle$ of $\mathcal{S}$, we have $\gamma \in [\![\phi]\!]_{\rho_0}$ iff $\langle \langle q, \phi \rangle, \vartheta \rangle \in \mathcal{W}[\mathcal{G}(\mathcal{S},\phi)), \mathcal{V}, 0, \mathcal{C}](\texttt{Parity})$ because $\texttt{Cond}(\phi,\rho_0)$ is equivalent to the formula which is always false. By denoting $\gamma' = \langle \langle q, \phi \rangle, \vartheta \rangle$, we have hence that $[0, \mathcal{V}] : \gamma' \models_{\mathcal{G}(\mathcal{S},\phi)}$ $\texttt{Parity}$ if and only if $\mathcal{S}, \gamma \models \phi$.

Proof of Lemma 11

Let $\mathcal{S} = \langle Q, T \rangle$ be a single-sided VASS and $\phi \in L_\mu^{sv}$. Strictly speaking the construction of the game $\mathcal{G}(\mathcal{S},\phi)$ proposed in the proof of Lemma 10 does not build a single-sided game. However we can adapt this construction in order to build an equivalent single-sided game. In this manner we adapt the construction to the case of $\phi \in L_\mu^{sv}$ by changing the rules for the outgoing transitions for states in the game of the form $\langle q, Q_1 \wedge \Box\psi \rangle$. To achieve this we build a game $\mathcal{G}'(\mathcal{S},\phi) = \langle Q', T', \kappa \rangle$ the same way as $\mathcal{G}(\mathcal{S},\phi)$ except that we perform the following change in the definition of transition relation for states of the form $\langle q, Q_1 \wedge \Box\psi \rangle$:

- If $\psi = Q_1 \wedge \Box\psi'$ then for all states $q \in Q_1$, for all transitions $\langle q, nop, q' \rangle \in T$, we have $\langle \langle q, \psi \rangle, nop, \langle q', \psi' \rangle \rangle$ in T', and, for all states $q \in Q_0$, we have $\langle \langle q, \psi \rangle, nop, \langle q, \psi \rangle \rangle$ in T'.

Then the states of the form $\langle q, Q_1 \wedge \Box\psi \rangle$ will belong to Player 1 and the coloring of such states will be defined as follows:

- for all $q \in Q$, for all subformulae $\psi \in sub(\phi)$, if $\psi = Q \wedge \Box\psi'$ then if $q \in Q_1$, $\kappa(\langle q, \psi \rangle) = 0$ else $\kappa(\langle q, \psi \rangle) = 1$.

Apart from these changes the definition of the game $\mathcal{G}'(\mathcal{S},\phi)$ is equivalent to the one of $\mathcal{G}(\mathcal{S},\phi)$. By construction, since $\mathcal{S}$ is single-sided and by definition of L_μ^{sv}, we have that such an integer game $\mathcal{G}'(\mathcal{S},\phi)$ is single-sided. Furthermore, for any concrete configuration $\gamma = \langle \langle q, \psi \rangle, \vartheta \rangle$, one can easily show that $[0, \mathcal{V}] : \gamma \models_{\mathcal{G}(\mathcal{S},\phi)} \texttt{Parity}$ iff $[0, \mathcal{V}] : \gamma \models_{\mathcal{G}'(\mathcal{S},\phi)} \texttt{Parity}$.

Proof of Theorem 7

Let $\mathcal{S} = \langle Q, T \rangle$ be a single-sided VASS, ϕ be closed formula of L_μ^{sv} and γ_0 be an initial configuration of $\langle Q, T \rangle$. Using Lemma 10 and Lemma 11, we have that $[0, \mathcal{V}] : \gamma_0' \models_{\mathcal{G}'(\mathcal{S},\phi)} \texttt{Parity}$ if and only if $\mathcal{S}, \gamma_0 \models \phi$ where $\mathcal{G}'(\mathcal{S},\phi)$ is a single-sided integer game. Hence, thanks to Corollary 2, we can deduce that the model-checking problem of L_μ^{sv} over single-sided VASS is decidable. Furthermore, by using the result of these two lemmas we have that $\langle q, \vartheta \rangle \in [\![\phi]\!]_{\rho_0}$ iff $\langle \langle q, \phi \rangle, \vartheta \rangle \in \mathcal{W}[\mathcal{G}'(\mathcal{S},\phi), \mathcal{V}, 0, C](\texttt{Parity})$. Hence by Corollary 1, we deduce that $[\![\phi]\!]_{\rho_0}$ is upward-closed and by Theorem 3 that we can compute its set of minimal elements which is equal to $\{\langle q, \vartheta \rangle \mid \langle \langle q, \phi \rangle, \vartheta \rangle \in \texttt{Pareto}[\mathcal{G}'(\mathcal{S},\phi), \mathcal{V}, 0, \mathcal{C}](\texttt{Parity})\}$.