

Making Weighted Containment Feasible: A Heuristic Based on Simulation and Abstraction

Guy Avni and Orna Kupferman

School of Computer Science and Engineering, Hebrew University, Israel

Abstract. Weighted automata map input words to real numbers and are useful in reasoning about quantitative systems and specifications. The containment problem for weighted automata asks, given two weighted automata \mathcal{A} and \mathcal{B} , whether for all words w , the value that \mathcal{A} assigns to w is less than or equal to the value \mathcal{B} assigns to w . The problem is of great practical interest, yet is known to be undecidable. Efforts to approximate weighted containment by weighted variants of the simulation pre-order still have to cope with large state spaces. One of the leading approaches for coping with large state spaces is abstraction. We introduce an abstraction-refinement paradigm for weighted automata and show that it nicely combines with weighted simulation, giving rise to a feasible approach for the containment problem. The weighted-simulation pre-order we define is based on a quantitative two-player game, and the technical challenge in the setting originates from the fact that the automata assign to words are unbounded. The abstraction-refinement paradigm is based on under- and over-approximation of the automata, where approximation, and hence also the refinement steps, refer not only to the languages of the automata but also to the values they assign to words.

1 Introduction

Traditional automata accept or reject their input, and are therefore Boolean. A weighted finite automaton (WFA, for short) has real-valued weights on its transitions and it maps each word to a real value. Applications of weighted automata include formal verification, where they are used for the verification of quantitative properties [10,11,17,25,32], as well as text, speech, and image processing, where the weights of the automaton are used in order to account for the variability of the data and to rank alternative hypotheses [15,30].

Technically, each transition in a WFA is associated with a weight, the value of a run is the sum of the weights of the transitions traversed along the run, and the value of a word is the value of the maximal run on it.¹ The rich structure of weighted automata makes them intriguing mathematical objects. Fundamental problems that have been solved decades ago for Boolean automata are still open or known to be undecidable in the weighted setting [29]. For example, while in the Boolean setting, nondeterminism does not add to the expressive power of the automata, not all weighted automata can be

¹ The above semantics, which we are going to follow in the paper, is a special case of the general setting, where each weighted automaton is defined with respect to an algebraic semiring.

determinized, and the problem of deciding whether a given nondeterministic weighted automaton can be determinized is still open, in the sense we do not even know whether it is decidable.

A problem of great interest in the context of automata is the *containment* problem. In the Boolean setting, the containment problem asks, given two automata \mathcal{A} and \mathcal{B} , whether all the words in Σ^* that are accepted by \mathcal{A} are also accepted by \mathcal{B} . In the weighted setting, the “goal” of words is not just to get accepted, but also to do so with a maximal value. Accordingly, the containment problem for WFAs asks, given two WFAs \mathcal{A} and \mathcal{B} , whether every word accepted by \mathcal{A} is also accepted by \mathcal{B} , and its value in \mathcal{A} is less than or equal to its value in \mathcal{B} . We then say that \mathcal{B} contains \mathcal{A} , denoted $\mathcal{A} \subseteq \mathcal{B}$. In the Boolean setting, the containment problem is PSPACE-complete [31]. In the weighted setting, the problem is in general undecidable [1,24]. The problem is indeed of great interest: In the automata-theoretic approach to reasoning about systems and their specifications, containment amounts to correctness of systems with respect to their specifications. The same motivation applies for weighted systems, with the specifications being quantitative [10].

Even in the Boolean setting, where the containment problem is decidable, its PSPACE complexity is an obstacle in practice and researchers have suggested two orthogonal methods for coping with it. One is to replace containment by a pre-order that is easier to check, with the leading such pre-order being the *simulation* preorder [28]. Simulation can be checked in polynomial time and symbolically [20,28], and several heuristics for checking containment by variants of simulation have been studied and used in practice [23,26]. A second method, useful also in other paradigms for reasoning about the huge, and possibly infinite, state space of systems is *abstraction* [3,8]. Essentially, in abstraction we hide some of the information about the system. This enables us to reason about systems that are much smaller, yet it gives rise to a 3-valued solution: yes, no, and unknown [5]. In the latter case, the common practice is to refine the abstraction, aiming to add the minimal information that would lead to a definite solution. In particular, in the context of model checking, the method of counterexample guided abstraction-refinement (CEGAR) has proven to be very effective [13].

In this paper we study a combination of the above two methods in the setting of weighted automata. Abstraction frameworks in the 3-valued Boolean semantics are typically based on *modal transition systems* (MTS) [27]. Such systems have two types of transitions: *may* transitions, which over-approximate the transitions of the concrete system, and *must* transitions, which under-approximate them. The over and under approximation refer to the ability of the automaton to take transitions, and hence to its language. In our weighted setting, we combine this with the weights of the transitions: may transitions over-approximate the actual weight and must transitions under-approximate it. This is achieved by defining the weight of may and must transitions according to the maximal and minimal weight, respectively, of the transitions that induce them.

The simulation preorder in the Boolean setting has a game-theoretic characterization. We extend this approach to the weighted setting and define weighted simulation between two WFAs \mathcal{A} and \mathcal{B} by means of a game played between an antagonist, who iteratively generates a word w and an accepting run r of \mathcal{A} on it, and a protagonist, who replies with a run r' of \mathcal{B} on w . The goal of the antagonist is to generate w and r so that

either r' is not accepting, or its value is smaller than the value of r . The goal of the protagonist is to continue the game forever without the antagonist reaching his goal. We say that \mathcal{A} is simulated by \mathcal{B} , denoted $\mathcal{A} \leq \mathcal{B}$ iff the protagonist has a winning strategy. The above definition is similar to the definition of quantitative simulation in [12,14], and has the flavor of the *energy games* in [4]. In these works, however, the winning condition in the game refers only to the weight along the traversed edges. This corresponds to the case the WFAs in question are such that all states are accepting. Even richer than our setting are *energy parity games* [9]. Both energy games and parity energy games can be decided in $\text{NP} \cap \text{co-NP}$. Our main challenge then is to develop an algorithm that would maintain the simplicity of the algorithm in [4] in the richer setting, which is simpler than the one of parity games. We do this by performing a preprocessing on the arena of the game, one after which we can perform only local changes in the algorithm of [4]. This is not easy, as like in parity energy games a winning strategy in the simulation game need not be memoryless. Our main contribution, however, is not the study of simulation games and their solution – the main ideas here are similar to these in [4,9], but the ability to combine simulation with abstraction and refinement, which we see as our main contribution.

Having defined over- and under-approximations of WFAs and the weighted simulation relation, we suggest the following heuristic for checking whether $\mathcal{A} \subseteq \mathcal{B}$. For a WFA \mathcal{U} and an abstraction function α , let $\mathcal{U}_\downarrow^\alpha$ and $\mathcal{U}_\uparrow^\alpha$ be the weighted under and over approximations of \mathcal{U} according to α . Let α and β be approximation functions for \mathcal{A} and \mathcal{B} , respectively. It is not hard to see that if $\mathcal{A}_\uparrow^\alpha \subseteq \mathcal{B}_\downarrow^\beta$, then $\mathcal{A} \subseteq \mathcal{B}$, and that if $\mathcal{A}_\downarrow^\alpha \not\subseteq \mathcal{B}_\uparrow^\beta$, then $\mathcal{A} \not\subseteq \mathcal{B}$. We show that the above is valid not just of containment but also for our weighted-simulation relation. This gives rise to the following heuristics. We start by checking $\mathcal{A}_\uparrow^\alpha \leq \mathcal{B}_\downarrow^\beta$ and $\mathcal{A}_\downarrow^\alpha \not\leq \mathcal{B}_\uparrow^\beta$, for some (typically coarse) initial abstraction functions α and β . As we prove in the paper, if we are lucky and one of them holds, we are done. Otherwise, the winning strategies of the antagonist in the first case and the protagonist in the second case suggest a way to refine α and β , and we repeat the process with the refined abstractions. While refinement in the Boolean case only splits abstract states in order to close the gap between may and must transitions, here we also have refinement steps that tighten the weights along transitions. Note that while repeated refinement can only get us to a solution to the $\mathcal{A} \leq \mathcal{B}$ problem, they also make it more likely that one of our checks returns an answer that would imply a definite solution to the undecidable $\mathcal{A} \subseteq \mathcal{B}$ problem.

Note that our abstraction-refinement procedure combines two games. The first, which corresponds to $\mathcal{A}_\uparrow^\alpha \leq \mathcal{B}_\downarrow^\beta$, approximates the simulation question $\mathcal{A} \leq \mathcal{B}$ from below. The second, which corresponds to $\mathcal{A}_\downarrow^\alpha \leq \mathcal{B}_\uparrow^\beta$, approximates it from above. Such dual approximations have proven useful also in the Boolean setting [6,16,18,21,22], where games combine may and must transitions, and also in settings in which games that are determined are approximated by means other than abstraction. For example, when LTL realizability is done by checking the realizability of approximations of both the specification and its negation [7].

2 Weighted Automata and their Abstraction

A nondeterministic finite weighted automaton on finite words (WFA, for short) is a tuple $\mathcal{A} = \langle \Sigma, Q, \Delta, Q_0, \tau, F \rangle$, where Σ is an alphabet, Q is a set of states, $\Delta \subseteq Q \times \Sigma \times Q$ is a transition relation, $Q_0 \subseteq Q$ is a set of initial states, $\tau : \Delta \rightarrow \mathbb{R}$ is a function that maps each transition to a real value in \mathbb{R} , and $F \subseteq Q$ is a set of accepting states. We assume that there are no *dead-end* states in \mathcal{A} . That is, for every $q \in Q$ there is a letter $\sigma \in \Sigma$ and state $q' \in Q$ such that $\Delta(q, \sigma, q')$.

A run of \mathcal{A} on a word $u = u_1, \dots, u_n \in \Sigma^*$ is a sequence of states $r = r_0, r_1, \dots, r_n$ such that $r_0 \in Q_0$ and for every $0 \leq i < n$ we have $\Delta(r_i, u_{i+1}, r_{i+1})$. The run r is accepting iff $r_n \in F$. The value of the run, denoted $val(r, u)$, is the sum of transitions it traverses. That is, $val(r, u) = \sum_{0 \leq i < n} \tau(\langle r_i, u_{i+1}, r_{i+1} \rangle)$. Since \mathcal{A} is non-deterministic, there can be more than one run on a single word. We define the value that \mathcal{A} assigns to $u \in \Sigma^*$, denoted $val(\mathcal{A}, u)$, as the value of the maximal-valued accepting run of \mathcal{A} on u . That is, $val(\mathcal{A}, u) = \max\{val(r, u) : r \text{ is an accepting run of } \mathcal{A} \text{ on } u\}$. As in NFAs, the language of \mathcal{A} , denoted $L(\mathcal{A})$, is the set of words in Σ^* that \mathcal{A} accepts.

We say that \mathcal{A} is *deterministic* if $|Q_0| = 1$ and for every $q \in Q$ and $\sigma \in \Sigma$, there is at most one state $q' \in Q$ such that $\Delta(q, \sigma, q')$.

An *abstraction function* for a WFA \mathcal{A} is a function $\alpha : Q \rightarrow A$, for a set A , which we assume to be smaller than Q . We refer to the members of Q as the *concrete states* and to these of A as the *abstract states*. The function α induces a partition of Q , and we sometimes refer to abstract states as sets of concrete states. In particular, for a concrete state $c \in Q$ and an abstract state $a \in A$, we use the notation $c \in a$ to indicate that $\alpha(c) = a$.

Consider a WFA \mathcal{A} and an abstraction function α . For parameters $\beta \in \{may, must\}$ and $\gamma \in \{max, min\}$, the abstraction of \mathcal{A} according to α, β , and γ is the WFA $\mathcal{A}_\beta^\gamma[\alpha] = \langle \Sigma, A, \Delta_\beta, A_0, \tau_\gamma, F_\beta \rangle$, where $A_0 = \{\alpha(q_0) : q_0 \in Q_0\}$, and $\Delta_\beta, \tau_\gamma$, and F_β are defined as follows:

- Consider $a, a' \in A$ and $\sigma \in \Sigma$. We define $\Delta_{must} \subseteq A \times \Sigma \times A$ so that $\Delta_{must}(a, \sigma, a')$ iff for every $c \in a$ there is $c' \in a'$ such that $\Delta(c, \sigma, c')$. We define $\Delta_{may} \subseteq A \times \Sigma \times A$ so that $\Delta_{may}(a, \sigma, a')$ iff there exists $c \in a$ and $c' \in a'$ such that $\Delta(c, \sigma, c')$.
- We define the minimum-value function, denoted τ_{min} , of an abstract transition to be the minimum over the values of concrete transitions that induce it. Formally, for $\langle a, \sigma, a' \rangle \in \Delta_\beta$, we define $\tau_{min}(\langle a, \sigma, a' \rangle) = \min\{\tau(\langle c, \sigma, c' \rangle) : c \in a, c' \in a', \text{ and } \Delta(c, \sigma, c')\}$. Similarly, we define the maximal-value function as τ_{max} , with $\tau_{max}(\langle a, \sigma, a' \rangle) = \max\{\tau(\langle c, \sigma, c' \rangle) : c \in a, c' \in a', \text{ and } \Delta(c, \sigma, c')\}$.
- We define $F_{may} = \{a \in A : a \cap F \neq \emptyset\}$ and $F_{must} = \{a \in A : a \subseteq F\}$.

Note that without weights, our definition coincides with the standard over- and under-approximations studied in the Boolean case. In the weighted setting, the abstraction approximates, in addition to the transitions, the value that the concrete WFA assigns to words. The two interesting combinations are then the under-approximating WFA $\mathcal{A}_\downarrow^\alpha = \mathcal{A}_{must}^{min}[\alpha]$ and the over-approximating WFA $\mathcal{A}_\uparrow^\alpha = \mathcal{A}_{may}^{max}[\alpha]$. When α is not important or clear from the context, we omit it.

We refer to runs of \mathcal{A}_\downarrow as *must-runs*, runs of \mathcal{A}_\uparrow as *may-runs*, and runs of \mathcal{A} as *concrete-runs*. Note that for every must-run $r = r_0, \dots, r_n$ of \mathcal{A}_\downarrow on some word $u \in \Sigma^*$, there is a matching run $r' = r'_0, \dots, r'_n$ of \mathcal{A} on u such that, for every $0 \leq i \leq n$, we have $r'_i \in r_i$. Similarly, for every run $r = r_0, \dots, r_n$ of \mathcal{A} on some word u , the sequence $r' = \alpha(r_0), \dots, \alpha(r_n)$ is a run of \mathcal{A}_\uparrow on u .

In the Boolean setting, *language containment* refers to words accepted by the automaton. That is, for two NFAs \mathcal{A} and \mathcal{B} , we say that $\mathcal{A} \subseteq \mathcal{B}$ iff $L(\mathcal{A}) \subseteq L(\mathcal{B})$. In the weighted setting, language containment is more involved, as we also have a requirement on the values the automata assign to words. For two WFAs, we say that $\mathcal{A} \subseteq \mathcal{B}$ iff $L(\mathcal{A}) \subseteq L(\mathcal{B})$ and for every $w \in L(\mathcal{A})$ we have $val(\mathcal{A}, w) \leq val(\mathcal{B}, w)$.

The *containment problem* gets as input two automata \mathcal{A} and \mathcal{B} , and decides whether $\mathcal{A} \subseteq \mathcal{B}$. The problem is known to be PSPACE-complete in the Boolean setting [33] and undecidable in the weighted setting [1,24].

Since, in practice, WFAs are typically very large, we would like to reason on their abstractions. As Theorem 1 below shows, \mathcal{A}_\downarrow and \mathcal{A}_\uparrow under- and over-approximates \mathcal{A} , making such a reasoning possible.

Theorem 1. *Consider a WFA \mathcal{A} and an abstraction function α . Then, $\mathcal{A}_\downarrow^\alpha \subseteq \mathcal{A} \subseteq \mathcal{A}_\uparrow^\alpha$.*

Proof: We start by proving that $\mathcal{A}_\downarrow \subseteq \mathcal{A}$. Consider a word $u = u_1, \dots, u_n \in L(\mathcal{A}_\downarrow)$. We prove that $u \in L(\mathcal{A})$ and $val(\mathcal{A}_\downarrow, u) \leq val(\mathcal{A}, u)$. Let $r = a_0, \dots, a_n \in A^*$ be an accepting run of \mathcal{A}_\downarrow on u . Since $a_0 \in A_0$, there is a concrete state $c_0 \in (a_0 \cap Q_0)$. Since r is a must-run, there is a concrete run $r' = c_0, \dots, c_n$ of \mathcal{A} on u such that, for $1 \leq i \leq n$, we have $c_i \in a_i$. Since r is accepting, $r_n \subseteq F$, implying that $c_n \in F$. Thus, r' is accepting and $u \in L(\mathcal{A})$.

It is left to prove that $val(\mathcal{A}_\downarrow, u) \leq val(\mathcal{A}, u)$. We show that for every accepting run r of \mathcal{A}_\downarrow on u and every concrete run r' that corresponds to it, $val(r, u) \leq val(r', u)$. Indeed, by the definition of τ_{min} , we have $val(r, u) = \sum_{0 \leq i \leq n} \tau_{min}(r_i, u_{i+1}, r_{i+1}) \leq \sum_{0 \leq i \leq n} \tau(r_i, u_{i+1}, r_{i+1}) = val(r', u)$, so we are done.

We proceed to the second claim, namely that $\mathcal{A} \subseteq \mathcal{A}_\uparrow$. Consider a word $u = u_1, \dots, u_n \in L(\mathcal{A})$. Let $r = c_0, \dots, c_n \in Q$ be an accepting run of \mathcal{A} on u with maximal value. That is, $val(\mathcal{A}, u) = val(r, u)$. We show that $r' = \alpha(c_0), \dots, \alpha(c_n)$ is an accepting run of \mathcal{A}_\uparrow on u with $val(r, u) \leq val(r', u)$. Thus, we conclude that $val(\mathcal{A}, u) \leq val(\mathcal{A}_\uparrow, u)$.

We show that r' is an accepting run. Since $c_0 \in Q_0$ we have $\alpha(c_0) \in A_0$. Since r is a run of \mathcal{A} , for every $0 \leq i < n$, we have $\Delta(c_i, u_{i+1}, c_{i+1})$. Thus, by the definition of Δ_{may} , we have $\Delta_{may}(\alpha(c_i), u_{i+1}, \alpha(c_{i+1}))$. Since r is accepting, $c_n \in F$. Thus, $\alpha(c_n) \in F_{may}$. We conclude that r' is an accepting run of \mathcal{A}_\uparrow on u .

To conclude the proof, we prove that $val(r, u) \leq val(r', u)$. By the definition of τ_{max} , for every $0 \leq i < n$, we have $\tau(c_i, u_{i+1}, c_{i+1}) \leq \tau_{max}(\alpha(c_i), u_{i+1}, \alpha(c_{i+1}))$. Thus, $val(r, u) = \sum_{0 \leq i < n} \tau(c_i, u_{i+1}, c_{i+1}) \leq \sum_{0 \leq i < n} \tau_{max}(\alpha(c_i), u_{i+1}, \alpha(c_{i+1})) = val(r', u)$, and we are done. \square

3 Weighted Simulation

As discussed in Section 1, the pre-order of simulation [28] is used in the Boolean setting as a heuristic for checking containment. In this section we define weighted simulation

and show that it enjoys the appealing properties of simulation in the Boolean setting. In Section 4, we show that weighted simulation can be checked by reasoning about abstractions of the WFAs in question.

3.1 Defining the Weighted Simulation Relation

Given two WFAs \mathcal{A} and \mathcal{B} , deciding whether $\mathcal{A} \subseteq \mathcal{B}$ can be thought of as a two-player game of one round: Player 1, the Player whose goal it is to show that there is no containment, chooses a word w and a run r_1 of \mathcal{A} on w . Player 2 then replies by choosing a run r_2 of \mathcal{B} on w . Player 1 wins if r_1 is accepting and r_2 is not or if $val(r_1, w) > val(r_2, w)$. While this game clearly captures containment, it does not lead to interesting insights or algorithmic ideas about checking containment. A useful way to view simulation is as a “step-wise” version of the above game in which in each round the players proceed according to a single transition of the WFAs.

We continue to describe the simulation game formally. A game between Player 1 and Player 2 is a pair $\langle G, \Gamma \rangle$, for an arena G and an objective Γ for Player 1. Consider two WFAs \mathcal{A} and \mathcal{B} , where for $\gamma \in \{\mathcal{A}, \mathcal{B}\}$, let $\gamma = \langle \Sigma, Q_\gamma, \Delta_\gamma, q_0^\gamma, F_\gamma, \tau_\gamma \rangle$. For simplicity, we assume that the WFAs are full, in the sense that each state and letter have at least one successor.

The arena of the game that corresponds to $\mathcal{A} \leq \mathcal{B}$ is $G = \langle V, E, v_0, \tau \rangle$. The set V of vertices is partitioned into two disjoint sets: $V_1 = Q_{\mathcal{A}} \times Q_{\mathcal{B}}$ are vertices from which Player 1 proceeds, and $V_2 = Q_{\mathcal{A}} \times \Sigma \times Q_{\mathcal{B}}$ are vertices from which Player 2 proceeds. The players alternate moves, thus $E \subseteq (V_1 \times V_2) \cup (V_2 \times V_1)$. Each play starts in the initial vertex $v_0 = \langle q_0^{\mathcal{A}}, q_0^{\mathcal{B}} \rangle \in V_1$, and $\tau : E \rightarrow \mathbb{R}$ is the weight function. We define $E = E_1 \cup E_2$ and τ as follows:

- $E_1 = \{ \langle \langle p, q \rangle, \langle p', \sigma, q \rangle \rangle : \langle p, \sigma, p' \rangle \in \Delta_{\mathcal{A}} \text{ and } q \in Q_{\mathcal{B}} \}$.
- $E_2 = \{ \langle \langle p, \sigma, q \rangle, \langle p, q' \rangle \rangle : \langle q, \sigma, q' \rangle \in \Delta_{\mathcal{B}} \text{ and } p \in Q_{\mathcal{A}} \}$.
- For $e_1 = \langle \langle p, q \rangle, \langle p', \sigma, q \rangle \rangle \in E_1$, we define $\tau(e_1) = \tau_{\mathcal{A}}(\langle p, \sigma, p' \rangle)$.
- For $e_2 = \langle \langle p, \sigma, q \rangle, \langle p, q' \rangle \rangle \in E_2$, we define $\tau(e_2) = -\tau_{\mathcal{B}}(\langle q, \sigma, q' \rangle)$.

Thus, edges in E_1 leave vertices in V_1 and correspond to Player 1 choosing a letter and a transition in \mathcal{A} . Edges in E_2 leave vertices in V_2 and correspond to Player 2 choosing a transition in \mathcal{B} .

A play of the game is a (possibly infinite) sequence of vertices $\pi = \pi_0, \pi_1, \dots$, where $\pi_0 = v_0$, and for every $i \geq 0$ we have $E(v_i, v_{i+1})$. Every finite play has a value, denoted $val(\pi)$, which is the sum of the edges that are traversed along it: i.e., $val(\pi) = \sum_{0 \leq i < |\pi|} \tau(\langle \pi_i, \pi_{i+1} \rangle)$. We use $\pi[i : j]$, for $i < j$, to refer to the sub-play π_i, \dots, π_j .

A strategy for player $i \in \{1, 2\}$ is a function $\rho_i : V^* \cdot V_i \rightarrow V$. Let \mathcal{S}_i be the set of all strategies for player i . Two strategies $\rho_1 \in \mathcal{S}_1$ and $\rho_2 \in \mathcal{S}_2$, induce a single play obtained when both players follow their strategies. Formally, the outcome of ρ_1 and ρ_2 , denoted $out(\rho_1, \rho_2)$, is the infinite play $\pi = \pi_1, \pi_2, \dots$, where for every $i \geq 0$, we have $\pi_{2i+1} = \rho_1(\pi[0 : 2i])$ and $\pi_{2i+2} = \rho_2(\pi[0 : 2i + 1])$. We say that a strategy ρ_i is *memoryless* if it depends only on the current vertex. Formally, $\rho_i(u_1 \cdot v) = \rho_i(u_2 \cdot v)$ for all $u_1, u_2 \in V^*$ and $v \in V_i$.

It is left to define the objective of the game. A finite play π is winning for Player 1 if the last vertex of π is in $F_{\mathcal{A}} \times (Q_{\mathcal{B}} \setminus F_{\mathcal{B}})$ or the last vertex of π is in $F_{\mathcal{A}} \times F_{\mathcal{B}}$ and $val(\pi) > 0$. The objective $\Gamma \subseteq V^\omega$ of Player 1, namely the set of plays that are winning for Player 1 is defined so that an infinite play π is in Γ iff it has a finite prefix that is winning according to the definition above. Note that for an infinite play π , if $\pi \notin \Gamma$, then it is winning for Player 2. Thus, the objective of Player 2 is $\bar{\Gamma} = V^\omega \setminus \Gamma$. Also note that once the play has a prefix that is winning for Player 1, there is no actual need for the play to continue. A *winning strategy* for Player 1 is a strategy $\rho_1 \in \mathcal{S}_1$ such that for every strategy $\rho_2 \in \mathcal{S}_2$, the play $out(\rho_1, \rho_2)$ is in Γ . A winning strategy for Player 2 is defined symmetrically. We define the simulation relation so that $\mathcal{A} \leq \mathcal{B}$ iff Player 2 has a winning strategy in G .

Theorem 2. *Simulation is strictly stronger than containment: (1) for all WFAs \mathcal{A} and \mathcal{B} , if $\mathcal{A} \leq \mathcal{B}$, then $\mathcal{A} \subseteq \mathcal{B}$. (2) There are WFAs \mathcal{A} and \mathcal{B} such that $\mathcal{A} \subseteq \mathcal{B}$ and $\mathcal{A} \not\leq \mathcal{B}$.*

Proof: We start with the first claim. Recall that $\mathcal{A} \subseteq \mathcal{B}$ iff $L(\mathcal{A}) \subseteq L(\mathcal{B})$ and for every $u \in L(\mathcal{A})$ we have $val(\mathcal{A}, u) \leq val(\mathcal{B}, u)$. We prove that if $\mathcal{A} \not\subseteq \mathcal{B}$ then Player 1 has a winning strategy. Thus, there is no winning strategy for Player 2 and $\mathcal{A} \not\leq \mathcal{B}$.

Assume that $\mathcal{A} \not\subseteq \mathcal{B}$. That is, there exists a word $u \in \Sigma^*$ such that $u \in L(\mathcal{A}) \setminus L(\mathcal{B})$ or $u \in L(\mathcal{A})$ and $val(\mathcal{A}, u) > val(\mathcal{B}, u)$. Consider the strategy $\rho_1 \in \mathcal{S}_1$ in which Player 1 selects the word u and chooses the run r_1 that maximizes the value of u in \mathcal{A} . We claim that this strategy is winning. That is, for every strategy $\rho_2 \in \mathcal{S}_2$ of Player 2, the play $out(\rho_1, \rho_2)$ is winning for Player 1.

Consider the play $\pi = \pi_1, \dots, \pi_n = out(\rho_1, \rho_2)[0 : 2 \cdot |u|]$. That is, $|u|$ moves for each player. We claim that Player 1 wins π . That is, either the last vertex in π , namely π_n , is in $F_{\mathcal{A}} \times (Q_{\mathcal{B}} \setminus F_{\mathcal{B}})$, or $\pi_n \in F_{\mathcal{A}} \times F_{\mathcal{B}}$ and $val(\pi) > 0$.

Since we define ρ_1 to follow an accepting run of \mathcal{A} on u , then $\pi_n \in F_{\mathcal{A}} \times Q_{\mathcal{B}}$. If $\pi_n \in F_{\mathcal{A}} \times (Q_{\mathcal{B}} \setminus F_{\mathcal{B}})$, we are done. Thus, we assume that $\pi_n \in F_{\mathcal{A}} \times F_{\mathcal{B}}$.

We claim that $val(\pi) > 0$. Assume by way of contradiction that $val(\pi) \leq 0$. Consider the run r_2 of \mathcal{B} on u that is induced by π . Recall that r_1 is the run of \mathcal{A} on u that is followed by ρ_1 . By the definition of the game; $val(\pi) = val(r_1, u) - val(r_2, u)$. Since $val(\pi) \leq 0$, we have that $val(r_2, u) \geq val(r_1, u)$. Now, since $val(\mathcal{B}, u)$ is the value of the maximal accepting run of \mathcal{B} on u , we have that $val(\mathcal{B}, u) \geq val(r_2, u)$. Also, since $val(r_1, u)$ is the maximal value for an accepting run of \mathcal{A} on u , we have that $val(\mathcal{A}, u) = val(r_1, u)$. Combining the two, we have that $val(\mathcal{A}, u) = val(r_1, u) \leq val(r_2, u) \leq val(\mathcal{B}, u)$, which is a contradiction to the fact that $val(\mathcal{A}, u) > val(\mathcal{B}, u)$, and we are done.

We continue to the second claim. Consider the WFAs \mathcal{A} and \mathcal{B} in Figure 1. We prove that $\mathcal{A} \subseteq \mathcal{B}$ but $\mathcal{A} \not\leq \mathcal{B}$. Note that when viewed as NFAs, then \mathcal{B} simulates \mathcal{A} , thus the example shows the weighted aspect of the difference between simulation and containment. It is easy to see that $L(\mathcal{A}) = L(\mathcal{B}) = \{ab, ac\}$. As for the values, $val(\mathcal{A}, ab) = 5 < 10 = val(\mathcal{B}, ab)$ and $val(\mathcal{A}, ac) = val(\mathcal{B}, ac) = 0$. Next, we show that $\mathcal{A} \not\leq \mathcal{B}$. For that, we describe a strategy for Player 1 and show that it is winning: in the first round, Player 1 chooses a and chooses to move to state s_2 in \mathcal{A} . If Player 2 chooses to move to q_4 , then Player 1 chooses the letter c , and Player 2 is stuck (that is, moves to a rejecting sink). If Player 2 chooses to move to q_2 , then Player 1 chooses the letter b , and wins as the value of the run in \mathcal{A} is 5 whereas the value in \mathcal{B} is only 1. \square




Fig. 1. Two WFAs, \mathcal{A} and \mathcal{B} , for which $\mathcal{A} \subseteq \mathcal{B}$ but $\mathcal{A} \not\leq \mathcal{B}$.

As in the Boolean setting, simulation and containment do coincide in case the simulating automaton is deterministic. Indeed, then, there is only one Player 2 strategy, so the “step-wise nature” of simulation does not play a role.

Theorem 3. *If \mathcal{B} is a DWFA, then $\mathcal{A} \subseteq \mathcal{B}$ iff $\mathcal{A} \leq \mathcal{B}$.*

Proof: We show that if \mathcal{B} is a DWFA and $\mathcal{A} \subseteq \mathcal{B}$, then Player 2 wins the game. Since \mathcal{B} is deterministic, there is only one Player 2 strategy, denoted ρ_2 : whenever Player 1 selects a letter, Player 2 replies by choosing the single outgoing transition from the current state. Consider a Player 1 strategy ρ_1 . If $out(\rho_1, \rho_2)$ is winning for Player 1, then there is a word u , namely the word that is selected ρ_1 , such that either $u \in L(\mathcal{A}) \setminus L(\mathcal{B})$ or $val(\mathcal{A}, u) > val(\mathcal{B}, u)$. Thus, $\mathcal{A} \not\leq \mathcal{B}$, which is a contradiction to the assumption. \square

Another property of simulation that stays valid in the weighted setting is its transitivity.

Theorem 4. *For WFAs \mathcal{A} , \mathcal{B} , and \mathcal{C} , if $\mathcal{A} \leq \mathcal{B}$ and $\mathcal{B} \leq \mathcal{C}$, then $\mathcal{A} \leq \mathcal{C}$.*

Proof: Consider WFAs \mathcal{A} , \mathcal{B} , and \mathcal{C} , such that $\mathcal{A} \leq \mathcal{B}$ and $\mathcal{B} \leq \mathcal{C}$. Thus, there is a Player 2 winning strategy ρ_2^1 in the game G_1 that corresponds to \mathcal{A} and \mathcal{B} and a Player 2 winning strategy ρ_2^2 in the game G_2 that corresponds to \mathcal{B} and \mathcal{C} . We prove that $\mathcal{A} \leq \mathcal{C}$ by showing that there is a Player 2 winning strategy ρ_2 in the game G that corresponds to \mathcal{A} and \mathcal{C} . Consider a Player 1 strategy ρ_1 in G . Recall that ρ_1 chooses a word and chooses a run in \mathcal{A} . Intuitively, we play ρ_1 against ρ_2^1 in G_1 . Recall that ρ_2^2 chooses states in \mathcal{B} . Thus, we construct a Player 1 strategy from ρ_2^2 's moves and play it against ρ_2^1 . We define ρ_2 to coincide with ρ_2^2 's moves. It follows from the fact that ρ_2^1 and ρ_2^2 are winning that ρ_2 is also winning.

Formally, assume we define ρ_2 up to round $2k + 1 \geq 1$ in G , and, for $k \geq 0$, let $\pi_{2k+1} = v_0, \dots, v_{2k+1}$ be the outcome of the game in the first $2k + 1$ rounds. Since $v_0 \in V_1$, this is a legitimate assumption. We define $\rho_2(\pi_{2k+1})$ as follows. For $0 \leq i \leq k$, let $v_{2i} = \langle a_i, c_i \rangle$ and let $v_{2i+1} = \langle a_{i+1}, \sigma_{i+1}, c_i \rangle$. Consider the play $\pi_{2k+2}^1 = v_0^1, \dots, v_{2k+2}^1 = out(G^1, \rho_1^1, \rho_2^1)[0 : 2k + 2]$, where ρ_1^1 chooses the same letters and states as ρ_1 . Thus, for $0 \leq j \leq k + 1$, we have $v_{2j}^1 = \langle a_j, b_j \rangle$ and $v_{2j+1}^1 = \langle a_{j+1}, \sigma_{j+1}, b_j \rangle$. Consider the Player 1 strategy ρ_1^2 that chooses the letters $\sigma_1, \dots, \sigma_{k+1}$ and the states b_1, \dots, b_{k+1} . Let $\pi_{2k+2}^2 = v_0^2, \dots, v_{2k+2}^2 = out(\rho_1^2, \rho_2^2)[0 : 2k + 2]$. Thus, $v_{2k+2}^2 = \langle b_{k+1}, c_{k+1} \rangle$. We define $\rho_2(\pi_{2k+1}) = \langle a_{k+1}, c_{k+1} \rangle$, and we claim that this strategy is winning.

We claim that if $a_k \in F_{\mathcal{A}}$, then $c_k \in F_{\mathcal{C}}$. Since ρ_2^1 is winning, $a_k \in F_{\mathcal{A}}$ implies that $b_k \in F_{\mathcal{B}}$, otherwise Player 1 wins in G_1 . Likewise, since ρ_2^2 , if $c_k \notin F_{\mathcal{C}}$, Player 1 wins in G_2 .

Finally, we claim that if $\langle a_k, c_k \rangle \in F_A \times F_C$, then $val(\pi_{2k}) \leq 0$. Recall that $val(\pi_{2k}) = val(a_1, \dots, a_k, \sigma_1, \dots, \sigma_k) - val(c_1, \dots, c_k, \sigma_1, \dots, \sigma_k)$. By the above, $\langle a_k, c_k \rangle \in F_A \times F_C$ implies $\langle a_k, b_k \rangle \in F_A \times F_B$. Since ρ_2^1 is winning, $val(\pi_{2k}^1) = val(a_1, \dots, a_k, \sigma_1, \dots, \sigma_k) - val(b_1, \dots, b_k, \sigma_1, \dots, \sigma_k) \leq 0$. Since $\langle b_k, c_k \rangle \in F_B \times F_C$ and ρ_2^2 is winning, $val(\pi_{2k}^2) = val(b_1, \dots, b_k, \sigma_1, \dots, \sigma_k) - val(c_1, \dots, c_k, \sigma_1, \dots, \sigma_k) \leq 0$. Combining the above, we have $val(\pi_{2k}) = val(\pi_{2k}^1) + val(\pi_{2k}^2) \leq 0$, and we are done. \square

Unlike the Boolean case, here Player 1 need not have a memoryless strategy, as we demonstrate below. The WFAs we use in the example are used also in [9] in order to show that Player 2 has no memoryless winning strategy in energy parity games.

Example 1. We show a family of WFAs $\mathcal{A}_1, \mathcal{A}_2, \dots$ and a WFA \mathcal{B} such that for all $n \geq 1$, Player 1 wins the simulation game corresponding to \mathcal{A}_n and \mathcal{B} , but he has no memoryless winning strategy. Moreover, a winning strategy for Player 1 needs memory of size $\Omega(m \cdot W)$, where m is the size of $\mathcal{A}_n \times \mathcal{B}$ and W is the maximal weight.

Consider the WFAs \mathcal{A}_n and \mathcal{B} in Figure 2. Since $L(\mathcal{B}) = a^*$, then clearly $L(\mathcal{A}_n) \subseteq L(\mathcal{B})$. However, $\mathcal{A}_n \not\subseteq \mathcal{B}$, since for $w = a^n \cdot a^{2Wn+1} a^n$, we have $cost(\mathcal{A}_n, w) = 1 > 0 = cost(\mathcal{B}, w)$.

We claim that in the simulation game (on bottom) that corresponds to the two WFAs, there is a winning Player 1 strategy (i.e., $\mathcal{A}_n \not\subseteq \mathcal{B}$) and that every such winning strategy for Player 1 requires $\Omega(m \cdot W)$ memory. Indeed, a winning Player 1 strategy must proceed to the state (q_n, s_0) and loop there for at least $2Wn + 1$ rounds before returning to the initial state. Thus, a winning strategy must “count” to $2Wn + 1$.

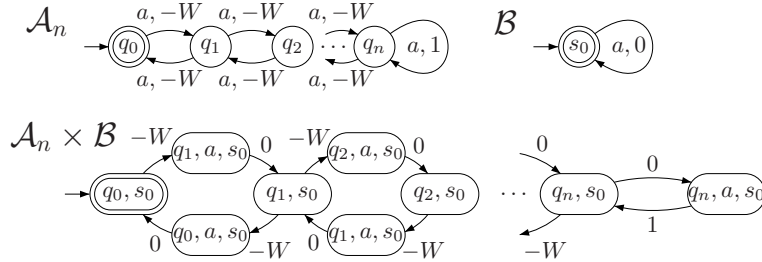


Fig. 2. WFAs \mathcal{A}_n and \mathcal{B} such that $\mathcal{A}_n \not\subseteq \mathcal{B}$ yet Player 1 does not have a memoryless strategy in the corresponding simulation game.

Before we turn to study the solution of the simulation game, observe that the set of infinite runs that are winning for Player 1 is *open*. A set $S \subseteq X$ is open if, intuitively, for every member $x \in S$, there is a non-empty “neighborhood” of x that is contained in S , where the neighborhood of x is defined with respect to some distance function on X . Formally, distances are defined by means of a *Cantor space* as follows. The domain is the set of infinite plays, and two infinite plays are “close” if they have a long common

prefix. In this topology, the set of winning plays for Player 1 is open since it is defined by prefixes. That is, a winning play has many other plays in its neighborhood: all plays that share the same winning prefix. By the Gale-Stewart theorem [19], every game that satisfies this property is determined, hence we have the following.

Theorem 5. *The simulation game is determined. That is, Player 1 or Player 2 has a winning strategy.*

3.2 Solving the Simulation Game

The simulation game stands between the energy games of [4], where the NFAs have no acceptance conditions, and the energy parity games of [9], where the winning condition is richer than the one of WFAs. Both these games are determined and can be decided in $\text{NP} \cap \text{co-NP}$. It is thus not surprising that we are going to show that the same holds for our simulation game. Our main challenge is to develop an algorithm that would maintain the simplicity of the algorithm in [4] in the richer setting of WFAs. The setting is indeed richer, and in particular, as in energy parity games, Player 1 need not have a memoryless winning strategy. We do this by performing a preprocessing on the arena of the game, one after which we can perform only local changes in the algorithm of [4]. Our main contribution, however, is not the study of simulation games and their solution – the main ideas here are similar to these in [4,9]. Rather, it is the combination of these ideas in an abstraction-refinement paradigm, to be described in Section 4.

Reducing $\langle G, \Gamma \rangle$ to a simpler game $\langle G', \Gamma' \rangle$ Consider an arena $G = \langle V, E, v_0, \tau \rangle$.

- Let $W_1 \subseteq V$ be the set of vertices from which Player 1 wins the reachability game with objective $F_A \times (Q_B \setminus F_B)$. That is, $v \in W_1$ iff Player 1 can force the game that starts in v to a vertex in $F_A \times (Q_B \setminus F_B)$.
- Let $W_2 \subseteq V$ be the vertices from which Player 2 wins the safety game with objective $((Q_A \setminus F_A) \times Q_B) \cup V_2$. That is, $v \in W_2$ iff Player 1 can force the game that starts in v to stay in vertices in $((Q_A \setminus F_A) \times Q_B) \cup V_2$.

Lemma 1. $W_1 \cap W_2 = \emptyset$.

Proof: Consider a vertex $v \in W_1$. We prove that $v \notin W_2$. Indeed, Player 1 has a strategy that forces the game from v to reach $F_A \times (Q_B \setminus F_B)$. Since $(F_A \times (Q_B \setminus F_B)) \cap (((Q_A \setminus F_A) \times Q_B) \cup V_2) = \emptyset$, Player 2 cannot force the game to stay in $((Q_A \setminus F_A) \times Q_B) \cup V_2$, and thus $v \notin W_2$. Proving that if $v \in W_2$, then $v \notin W_1$ is similar. \square

It follows from Lemma 1 that we can distinguish between three cases: If $v_0 \in W_1$, then Player 1 wins the game. If $v_0 \in W_2$, Player 2 wins the game. Otherwise, we define a new game, which excludes states from $W_1 \cup W_2$.

We define the new game $\langle G', \Gamma' \rangle$ on the arena $G' = \langle V', E', v_0, \tau' \rangle$. The set V' of vertices are $\{v_{\text{sink}}\} \cup (V \setminus (W_1 \cup W_2))$. The set V'_1 of vertices of Player 1 is $V_1 \cap V'$, and the set V'_2 of vertices of Player 2 is $V_2 \cap V'$. We say that a vertex $v \in V'$ is a

dead-end iff $adj(v) \subseteq (W_1 \cup W_2)$, where $adj(v)$ is defined with respect to E . That is, $adj(v) = \{v' \in V : E(v, v')\}$. The set E' of edges restricts the set E to vertices in V' and includes, in addition, an edge from every dead-end vertex in V' to v_{sink} and an edge from v_{sink} to itself. Recall that we assume that the WFAs on which the simulation game are defined are total, and thus there are no dead-ends in G . Hence, a vertex is a dead-end in G' when all its successors in G are in $W_1 \cup W_2$. Finally, τ' assigns the same value as τ for the edges in E , and assigns 0 to the new edges.

A finite play π is winning for Player 1 in the new game iff the last vertex in π is in $F_A \times F_B$ and $val(\pi) > 0$. As in the original game, Player 2 wins an infinite play iff it does not have a finite prefix that is winning for Player 1. Note that an infinite play π satisfies this condition, i.e., $\pi \in \overline{T'}$, if either one of two conditions: either π is contained in V and $\pi \in \overline{T}$, or π has a finite prefix $\pi[0 : i]$ that ends in a dead-end vertex, i.e., for every $k > i$ we have $\pi_k = v_{sink}$, and for every $j \leq i$, we have that $\pi[0 : j]$ is not winning for Player 1.

The characteristics of the vertices in W_1 and W_2 , as well as the dead-end states enable us to construct, given a winning strategy for Player 1 in G' , a winning strategy for him in G , and similarly for Player 2. Hence, we have the following.

Lemma 2. *Player 1 wins $\langle G, \Gamma \rangle$ iff he wins $\langle G', \Gamma' \rangle$.*

Proof: The claim follows from the determinacy of $\langle G, \Gamma \rangle$ and $\langle G', \Gamma' \rangle$ (the proof of the latter is similar to that of Theorem 5), and from Lemmas 5 and 6 below. \square

Before turning to prove the two directions of Lemma 2 we prove the following two lemmas.

Lemma 3. *Consider the arena G .*

- For a vertex $v_1 \in V'_1$ it holds that $adj(v_1) \cap W_1 = \emptyset$.
- For a vertex $v_2 \in V'_2$ it holds that $adj(v_2) \cap W_2 = \emptyset$.

Proof: Consider a vertex $v_1 \in V'_1$. Assume towards contradiction that there is a vertex $u \in adj(v_1) \cap W_1$. We claim that $v_1 \in W_1$, contradicting the fact that $v_1 \in V'$. Indeed, $u \in W_1$ implies that Player 1 has a strategy ρ_1 that forces the game to $F_A \times (Q_B \setminus F_B)$ from u . A Player 1 strategy from v that forces the game to $F_A \times (Q_B \setminus F_B)$ proceeds from v to u and continues with ρ_1 . Proving the other claim is dual. \square

Lemma 4. *Consider the arena G and a vertex $v \in V'$ that is a dead-end vertex. We claim that $v \in F_A \times F_B$ and that $adj(v) \subseteq W_2$.*

Proof: Consider a vertex $v \in V'$ that is a dead-end vertex. We claim that $v \in V_1$. Assume towards contradiction that $v \in V_2$. Since v is a dead-end vertex $adj(v) \subseteq W_1 \cup W_2$. By Lemma 3 and since we assume $v \in V_2$, we have that $adj(v) \subseteq W_1$. This implies that Player 2, from v , must proceed to a vertex in W_1 . Thus, Player 1 can force the game from v to a vertex in $F_A \times (Q_B \setminus F_B)$, and $v \in W_1$. This is a contradiction to the fact that $v \in V'$. Hence, we have that $v \in V'_1$.

We continue to prove that $v \in F_A \times F_B$. Since $F_A \times (Q_B \setminus F_B) \subseteq W_1$, we have that $v \notin F_A \times (Q_B \setminus F_B)$. We are left to show that $v \notin (Q_A \setminus F_A) \times Q_B$. Indeed, if $v \in (Q_A \setminus F_A) \times Q_B$, then since $adj(v) \subseteq W_2$, we have that $v \in W_2$, which contradicts the fact that $v \in V'$. \square

We continue to prove the two directions of the Lemma 2.

Lemma 5. *If Player 1 wins $\langle G', \Gamma' \rangle$, then he also wins $\langle G, \Gamma \rangle$.*

Proof: Consider a Player 1 winning strategy ρ'_1 in $\langle G', \Gamma' \rangle$. We construct a Player 1 strategy ρ_1 in $\langle G, \Gamma \rangle$ as follows. Recall that from every vertex in W_1 , Player 1 has a strategy that forces the game to $F_A \times (Q_B \setminus F_B)$. We refer to this strategy as the *reachability strategy*. The strategy ρ_1 proceeds in the same manner as ρ'_1 on the vertices $V' \setminus \{v_{sink}\}$ and when the game reaches a vertex in W_1 , it continues with the reachability strategy. For completeness, we define ρ_1 to proceed arbitrarily from vertices in $W_2 \cap V_1$.

We claim that ρ_1 is winning. We consider the ways in which a Player 2 strategy can win against ρ_1 . We claim that outcomes of the game which stay in V' are losing for Player 2. Indeed, since ρ_1 plays the same as ρ'_1 on V' , such outcomes will eventually reach a winning position for Player 1. Thus, the only hope Player 2 has for winning is to “escape” V' .

We distinguish between two ways in which Player 2 can “escape” V' . In the first case, Player 2 can proceed from a vertex $v \in V'_2$ to a vertex in $adj(v) \setminus V'$. However, by Lemma 3, for every $v \in V'_2$, we have that the vertices $adj(v) \setminus V'$ are contained in W_1 . Thus, Player 1 wins from such vertices. For the second case, Player 2 can force the game to a dead-end vertex. However, in such a case, a play that ends in a dead-end vertex is contained in V' . Thus, it is a possible outcome of the game G' when Player 1 plays according to ρ'_1 . Since ρ'_1 is winning, we have that the play is winning for Player 1, and we are done. \square

Lemma 6. *If Player 2 wins $\langle G', \Gamma' \rangle$, then he also wins $\langle G, \Gamma \rangle$.*

Proof: Consider a Player 2 winning strategy ρ'_2 in $\langle G', \Gamma' \rangle$. We construct a Player 2 strategy ρ_2 in $\langle G, \Gamma \rangle$. Recall that from every vertex in W_2 , Player 2 has a strategy, in the game played on G that stays in $((Q_A \setminus F_A) \times Q_B) \cup V_2$. We refer to this strategy as the *safety strategy*. We define ρ_2 to play the same as ρ'_2 on vertices in $V' \setminus \{v_{sink}\}$, and if a play reaches a vertex in W_2 , then it continues with the safety strategy. Note that the safety strategy keeps the game in W_2 . We claim that the game never reaches W_1 . Indeed, ρ'_2 keeps the game in V' and by Lemma 3, only vertices in V'_2 have neighbors in W_1 . Thus, ρ_2 is well defined.

Consider a finite play $\pi = \pi_1, \dots, \pi_n$, that is the outcome in the first n rounds of ρ_2 and some Player 1 strategy. We claim that π is not winning for Player 1. By the above $\pi_n \notin W_1$. Next, we distinguish between two cases. If $\pi_n \in W_2$, then $\pi_n \in ((Q_A \setminus F_A) \times Q_B) \cup V_2$, and thus Player 1 cannot win in π_n . Otherwise, since the safety strategy keeps the game in W_2 , the play π is contained in V' and thus it is an outcome of ρ'_2 against some Player 1 strategy. Since ρ'_2 is winning in G' , the play π is not winning for Player 1, and we are done. \square

It is thus left to show how to solve the game $\langle G', \Gamma' \rangle$.

Solving the game $\langle G', \Gamma' \rangle$ We say that a strategy of Player 1 is an *almost memoryless* strategy if, intuitively, in every play, when visiting a vertex, Player 1 plays in

the same manner except for, possibly, the last visit to the vertex. Formally, consider a Player 1 strategy ρ_1 . Consider a vertex $v \in V_1$. We say that ρ_1 is almost memoryless for v iff there are two vertices $v_1, v_2 \in V_2$ such that for every Player 2 strategy ρ_2 , if $out(\rho_1, \rho_2)[0 : n] = v$, then either $out(\rho_1, \rho_2)[n + 1] = v_1$, or $out(\rho_1, \rho_2)[n + 1] = v_2$ and for every index $n' > n + 1$ we have $out(\rho_1, \rho_2)[n'] \neq v$. We say that ρ_1 is almost memoryless iff it is almost memoryless for every vertex in V_1 .

Lemma 7. *If Player 1 has a winning strategy in $\langle G', \Gamma' \rangle$, then he also has an almost memoryless winning strategy.*

Proof: Assume ρ_1 is a Player 1 winning strategy. Our goal is to construct an almost memoryless winning strategy ρ'_1 from ρ_1 . Intuitively, we divide the Player 1 strategy into two “phases”: in the first phase, which we refer to as the “accumulation phase”, Player 1’s goal is to force the game into accumulating a high value. In the second phase, which we refer to as the “reachability phase”, his goal is to force the game to a winning position, which is a vertex in $F_A \times F_B$. Since all the vertices in V' are not in W_2 , Player 1 can force the game to a winning position from every vertex in V' . Also, since reaching a winning position is done in a memoryless manner, it does not involve a play with cycles, and thus, we bound the maximal value that Player 1 needs to accumulate in the first phase.

We use a similar construction to that in [4]. For every vertex $v \in V'$, we find the Player 2 strategy whose outcome against ρ_1 has the longest play η_v with the minimal value that ends in v . We set the almost memoryless strategy to respond, from every vertex in V'_1 , in the same manner the original winning strategy ρ_1 responds to η_v .

Formally, we define η_v , as the longest play that ends in v with the minimal value when Player 1 plays ρ_1 . That is, for every Player 2 strategy ρ_2 , and every index $m \in \mathbb{N}$, if $out(\rho_1, \rho_2)[0 : m]$ ends in v and does not have a prefix that is winning for Player 1, then $val(\eta_v) \geq val(out(\rho_1, \rho_2)[0 : m])$. Also, if $val(\eta_v) = val(out(\rho_1, \rho_2)[0 : m])$, then $|\eta_v| \geq m$. Since ρ_1 is winning, every outcome of the game is finite, and thus the minimal value and longest plays are well defined.

Recall that from every vertex in $V' \setminus \{v_{sink}\}$, there is a Player 1 strategy that forces the game to a vertex in $F_A \times F_B$. It is well known that such a strategy is memoryless, and thus for every vertex we can use the same reachability strategy. We denote this strategy by ρ_1^{reach} . Since ρ_1^{reach} is memoryless, the play from every vertex in V , that ends in a vertex in $F_A \times F_B$ is of length at most $|V'|$. Thus, its value is at least $\sum\{\tau(e) : e \in E' \text{ and } \tau(e) < 0\}$. Let $M = -\sum\{\tau(e) : e \in E' \text{ and } \tau(e) < 0\}$.

We define the almost memoryless Player 1 strategy ρ'_1 as follows. Consider a play $\pi = \pi_0, \dots, \pi_n$, where $\pi_n \in V'_1$. If there is no index $0 \leq i < |n|$ such that $\pi_i = \pi_n$, then we define $\rho'_1(\pi) = \rho_1(\eta_{\pi_n})$. Otherwise, let i be an index such that $\pi_i = \pi_n$. We distinguish between two cases. If there is an index $0 \leq j \leq |n|$ such that $val(\pi_0, \dots, \pi_j) > M$, Player 1 proceeds to the “reachability phase” and plays according to the reachability strategy. That is, $\rho'_1(\pi) = \rho_1^{reach}(\pi_n)$. Otherwise, Player 1 continues with the “accumulation phase”. That is, $\rho'_1(\pi) = \rho'_1(\pi[0 : i])$.

Clearly, the strategy ρ'_1 is almost memoryless. Also, by playing ρ'_1 , Player 1 wins from a position that is reached after a play with value at least M . Thus, in the following we prove that for every Player 2 strategy, the outcome of the game is either winning for Player 1 or it reaches such a position. We conclude that ρ'_1 is winning.

Claim. Consider a Player 2 strategy ρ_2 . We claim that if a prefix π of the outcome when Player 1 plays ρ'_1 and Player 2 plays ρ_2 , is not winning for Player 1 and if Player 1 has not begun the “reachability phase”, then, assuming π ends in $v \in V'$, we have $val(\pi) \geq val(\eta_v)$. Also, the length $|\pi|$ is at most the length $|\eta_v|$.

Formally, let $\pi = \pi_0, \dots, \pi_n$ be the prefix of length n of $out(\rho'_1, \rho_2)$. If there is no prefix of π that is winning for Player 1 and for every index $0 \leq i \leq n$ we have that $val(\pi_0, \dots, \pi_i) \leq M$, then $val(\pi) \geq val(\eta_{\pi_n})$ and $|\eta_{\pi_n}| \geq n$.

We prove the claim by an induction on the length of π . For the base case, the play of length 0 is a play that ends in v_0 , and thus $0 = val(v_0) \geq val(\eta_{v_0})$, and clearly, $|\eta_{v_0}| \geq 0$.

We assume the claim is correct for a prefix of length $n - 1$ and we prove for length n . We claim that $\eta_{\pi_{n-1}} \cdot \pi_n$ is a play that is an outcome when Player 1 plays ρ_1 . If $\pi_{n-1} \in V'_1$, then since we assume Player 1 has not begun the “reachability phase”, we have $\eta_{\pi_{n-1}} \cdot \pi_n = \eta_{\pi_{n-1}} \cdot \rho'_1(\pi_0, \dots, \pi_{n-1}) = \eta_{\pi_{n-1}} \cdot \rho_1(\eta_{\pi_{n-1}})$. If $\pi_{n-1} \in V'_2$, then clearly, proceeding from π_{n-1} to π_n is a legitimate Player 2 strategy.

Clearly, $val(\pi_0, \dots, \pi_n) = val(\pi_0, \dots, \pi_{n-1}) + \tau'(\pi_{n-1}, \pi_n)$. By the induction hypothesis, $val(\pi_0, \dots, \pi_{n-1}) \geq val(\eta_{\pi_{n-1}})$, and by the above, $val(\eta_{\pi_{n-1}} \cdot \pi_n) \geq val(\eta_{\pi_n})$. Combining the expressions we have $val(\pi) \geq val(\eta_{\pi_n})$.

Similarly, $|\pi| = |\pi_0, \dots, \pi_{n-1}| + 1$, by the induction hypothesis, $|\pi_0, \dots, \pi_{n-1}| \leq |\eta_{\pi_{n-1}}|$, and by the above, $|\eta_{\pi_{n-1}} \cdot \pi_n| \leq |\eta_{\pi_n}|$. Thus, $|\pi| \leq |\eta_{\pi_n}|$, and we are done.

Claim. The value of the prefix between visits to the same vertex increases.

Formally, let $\pi = \pi_0, \dots, \pi_n = out(\rho'_1, \rho_2)[0 : n]$. For a vertex $v \in V'_1$, let $0 \leq i_1 < \dots < i_k \leq |n|$ be the indices in which π visits v , i.e., for every $1 \leq j \leq k$, we have $\pi_{i_j} = v$. We claim that $val(\pi[0 : i_j]) < val(\pi[0 : i_{j+1}])$. Assume towards contradiction that this is not the case. Thus, wlog., $val(\pi[0 : i_1]) \geq val(\pi[0 : i_2])$. Consider the Player 2 strategy that starts with ρ_2 for i_1 rounds, and continues repeatedly with ρ_2 as it is played between rounds i_1 and i_2 . Thus, the new outcome of the game is $\pi[0 : i_1] \cdot \pi[i_1 : i_2] \cdot \dots \cdot \pi[i_1 : i_2]$. We reach a contradiction to the claim above, since either the value of the play decreases and will eventually be lower than $val(\eta_{\pi[i_1]})$, or the value remains the same but the length increases and will eventually exceed the length $|\eta_{\pi[i_1]}|$, and we are done.

Claim. If a prefix of the outcome of the game is not winning for Player 1, then it does not reach v_{sink} .

Assume towards contradiction that the prefix of the outcome of the game $\pi = \pi_0, \dots, \pi_n = out(\rho'_1, \rho_2)[0 : n]$ is not winning for Player 1 and $\pi_n = v_{sink}$. Thus, by Lemma 4, $\pi_{n-1} \in F_{\mathcal{A}} \times F_{\mathcal{B}}$. If Player 1 has begun the “reachability phase”, then he has begun it in an index which is at least $n - |V'|$, and it follows from the above that $val(\pi_0, \dots, \pi_{n-1}) > 0$, which contradicts our assumption that the prefix is not winning for Player 1. Otherwise Player 1 is still in the “accumulation phase”, and by the above, $0 \geq val(\pi_0, \dots, \pi_{n-1}) \geq val(\eta_{\pi_{n-1}})$. Since $\rho'_1(\pi_0, \dots, \pi_{n-1}) = \rho_1(\eta_{\pi_{n-1}}) = v_{sink}$, we reach a contradiction to the fact that ρ_1 is winning, and we are done.

Let pos be the value of the minimal valued positive simple cycle in G' . Formally, $pos = \min\{val(\pi) : \pi \text{ is a cycle and } val(\pi) > 0\}$.

Claim. Consider a Player 2 strategy ρ_2 and let $\pi = \text{out}(\rho'_1, \rho_2)[0 : K]$, where $K = (\lceil M/pos \rceil + 2) \cdot |V'|$. If π does not have a prefix that is winning for Player 1, then $\text{val}(\pi) > M$.

Assume that π does not have a prefix that is winning for Player 1. Thus, by the third claim, it does not reach the vertex v_{sink} . Hence, π is contained in $V' \setminus \{v_{\text{sink}}\}$. Since the arena is finite, there is a vertex $v \in V \setminus \{v_{\text{sink}}\}$ that appears at least $\lceil M/pos \rceil + 2$ times in π . Let $0 \leq i_1 < \dots < i_k$ be the indices in which v appears in π . By the second claim, the cycles have positive value. That is, for $1 \leq j < k$ we have $\text{val}(\pi[i_j : i_{j+1}]) > 0$. By the definition of pos , we have $\text{val}(\pi[i_j : i_{j+1}]) \geq pos$. Also, since cycles have a positive value, we have $\text{val}(\pi[0 : i_1]) \geq -M$. Combining the expressions we have $\text{val}(\pi) = \text{val}(\pi[0 : i_1]) + \sum_{1 \leq j < k} \text{val}(\pi[i_j : i_{j+1}]) \geq -M + (k-1) \cdot pos > 0$, and we are done. \square

For Player 2, our situation is even better as, intuitively, cycles in the game are either good for Player 2, in which case a strategy for him would always proceed to these cycles, or bad for Player 2, in which case a strategy for him would never enter them.

Lemma 8. *If Player 2 has a winning strategy in $\langle G', \Gamma' \rangle$, then he also has a winning memoryless strategy.*

Proof: Consider a winning Player 2 strategy ρ_2 . We construct a memoryless strategy ρ'_2 , in a similar manner to that in Lemma 7. For every vertex $v \in V'$, let η_v be the play with maximal value that reaches v . Formally, for every Player 1 strategy ρ_2 , if for an index n we have $\text{out}(\rho_1, \rho_2)[n] = v$, then $\text{val}(\text{out}(\rho_1, \rho_2)[0 : n]) \leq \text{val}(\eta_v)$. If there is more than one such play, we chose one arbitrarily.

We claim that η_v is well defined. Indeed, there is a bound on the maximal value $\text{val}(\eta_v)$, namely M , and since the arena is finite the value cannot approach the bound arbitrarily close.

Consider a play π that ends in a vertex $v \in V'_2$. We define $\rho'_2(\pi) = \rho_2(\eta_v)$. Clearly, the strategy is memoryless. We claim that it is winning.

Consider a Player 1 strategy ρ_1 and let $\pi = \pi_0, \dots, \pi_n = \text{out}(\pi_1, \pi'_2)[0 : n]$. We prove by induction on n that $\text{val}(\pi) \leq \text{val}(\eta_{\pi_n})$. The proof is dual to the one in Theorem 7. For the base case, v_0 is a trivial outcome of the game when Player 2 plays ρ_2 . Thus, $\text{val}(\eta_{v_0}) \geq 0$. We assume the claim holds for prefixes of length $n-1$ and we prove for n . Clearly, $\text{val}(\pi) = \text{val}(\pi_0, \dots, \pi_{n-1}) + \tau(\pi_{n-1}, \pi_n)$. By the induction hypothesis, $\text{val}(\pi_0, \dots, \pi_{n-1}) \leq \text{val}(\eta_{\pi_{n-1}})$. We claim that the play $\eta_{\pi_{n-1}} \cdot \pi_n$ is a play where Player 2 plays ρ_2 . If $\pi_{n-1} \in V'_1$, then clearly, proceeding from π_{n-1} to π_n is part of a Player 1 strategy. If $\pi_{n-1} \in V'_2$, then $\eta_{\pi_{n-1}} \cdot \pi_n = \eta_{\pi_{n-1}} \cdot \rho'_2(\pi_0, \dots, \pi_{n-1}) = \eta_{\pi_{n-1}} \cdot \rho_2(\eta_{\pi_{n-1}})$. Since $\eta_{\pi_{n-1}}$ is a prefix of the outcome of ρ_2 against some Player 1 strategy, so is $\eta_{\pi_{n-1}} \cdot \pi_n$. Thus, $\text{val}(\eta_{\pi_n}) \geq \text{val}(\eta_{\pi_{n-1}} \cdot \pi_n)$. Combining the expressions we have that $\text{val}(\pi) \leq \text{val}(\eta_{\pi_n})$, and we are done.

Consider a finite outcome of the game π . We claim that π is not winning for Player 1. Indeed, if π ends in a vertex in $v \in F_A \times F_B$, then since ρ_2 is winning, $\text{val}(\eta_v) \leq 0$, and since $\text{val}(\pi) \leq \text{val}(\eta_v)$, we have that π is not winning for Player 1, and we are done. \square

Before turning to prove the complexity results, we remind the reader of the Bellman-Ford algorithm. The algorithm gets as input a weighted directed graph $\langle \mathcal{V}, \mathcal{E}, \theta \rangle$, where \mathcal{V} is a set of vertices, $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ is a set of edges, and $\theta : \mathcal{E} \rightarrow \mathbb{R}$ is a weight function. The algorithm also gets a distinguished source vertex $s \in \mathcal{V}$. It outputs a function $C : \mathcal{V} \rightarrow \mathbb{R}$, where for every $v \in \mathcal{V}$, the value $C(v)$ is the value of the shortest path between s and v . If there is a negative cycle connected to s , the algorithm reports that such a cycle exists but it cannot return a correct answer since no shortest path exists.

We continue to prove the complexity that follows.

Lemma 9. *Solving simulation games is in co-NP.*

Proof: We show how to decide in polynomial time, given a memoryless Player 2 strategy, if it is a winning strategy for Player 2, and thus show that the problem of deciding if Player 1 wins in the game is in co-NP. Given a memoryless Player 2 strategy we trim every edge that starts from vertices in V'_2 and does not coincide with the strategy. We run the longest path version of the Bellman-Ford algorithm. That is, given a directed graph and a source vertex, the algorithm returns a function $C : V \rightarrow \mathbb{R}$ that assigns to every vertex the longest path from the source vertex, and reports if there is a positive valued cycle in the graph.

We claim that the strategy is winning iff there is no positive cycle in the game arena, and if every vertex $v \in F_A \times F_B$ has $C(v) \leq 0$.

We assume that there is no positive cycle in the arena and for every $v \in F_A \times F_B$ we have $C(v) \leq 0$. We prove that the strategy is winning. Assume by way of contradiction that there is a Player 1 strategy that wins against the Player 2 strategy. That is, when Player 1 plays his strategy there is a prefix of the outcome that is $\pi = v_0, \dots, v_n$ such that $v_n \in F_A \times F_B$ and $val(\pi) > 0$. Since every cycle in the arena is non-positive we construct an outcome π' by removing every cycle starting with the longest cycle. Thus, $val(\pi') \geq val(\pi)$. Since π' ends in $v_n \in F_A \times F_B$ and it does not have cycles, we have $C(v_n) \geq val(\pi')$. However, $0 < val(\pi')$ and $C(v_n) \leq 0$, which is a contradiction.

For the other direction, we prove that if one of the two properties does not hold, Player 1 has a winning strategy. If there is a vertex $v \in F_A \times F_B$ such that $C(v) > 0$, then by playing the strategy that corresponds with the longest path in the arena, Player 1 wins.

If there is a positive cycle, Player 1 wins by playing a strategy that traverses the cycle. Recall that there is a Player 1 strategy that forces the game to a vertex in $F_A \times F_B$ from every vertex in V' . Specifically from every vertex in the positive cycle. Thus, Player 1 wins by traversing the cycle until a *sufficient* value is accumulated, and then using the strategy that forces the game to a vertex in $F_A \times F_B$. \square

Lemma 10. *Solving simulation games is in NP.*

Proof: Similar to the above, we guess a memoryless Player 1 strategy ρ_1 and trim the arena according to it. We show how to check, in polynomial time, if there is an almost memoryless winning strategy that corresponds to ρ_1 . Intuitively, we check if Player 2 can close a *good lasso*, which is an outcome of the game that is winning for Player 2, or if Player 2 can reach the vertex v_{sink} is a play that is not losin. If he cannot, we show

that ρ_1 either forces the game to a vertex in $F_A \times F_B$ after a positive valued play, or ρ_1 forces the game to close a positive valued cycle.

The algorithm is the same algorithm as in [4] with a slight adjustment, which in turn is an adjustment of the Bellman-Ford algorithm:

```

function INIT( $s, c$ )
  for  $v \in V'$  do
     $C(v) = \infty$ 
   $C(s) = c$  and  $b = false$ .

function SHORTESTPATH( $s$ )
  for  $|V'| - 1$  rounds do
    for  $\langle u, v \rangle \in E'$  do
      if ( $C(v) \geq C(u) + \tau(v, u)$ ) and
      (if  $v \in F_A \times F_B$ , then  $C(u) + \tau(v, u) \leq 0$ ) then
         $C(v) = C(u) + \tau(v, u)$ 
      if  $v = s$  then  $b = true$ 

function MINRMCR( $s, c$ )
  init( $s, c$ )
  ShortestPath( $v_0$ )
   $C' = C$ 
  ShortestPath( $v_0$ )
  for  $v \in V'$  do
    if  $C'(v) < C(v)$  then
       $C(v) = -\infty$ 
  return ( $C, b$ )

function MAIN
  ( $C_0, b_0$ ) = MINRMCR( $v_0, 0$ )
  if ( $\exists v \in V$  with  $C_0(v) = -\infty$ ) or ( $C_0(v_{sink}) \leq 0$ ) then
    return false
  for  $v \in V' \setminus \{v_{sink}\}$  s.t.  $C_0(v) \neq \infty$  do
    ( $C_v, b_v$ ) = MINRMCR( $v, C(v)$ )
    if  $C_v(v) \leq C(v)$  and  $b_v = true$  then return false
  return true

```

We start by proving that if the algorithm returns *false*, then there is no winning almost memoryless strategy that corresponds to the guessed Player 1 memoryless strategy ρ_1 .

Claim. After a call to *ShortestPath*(v), if $C(u) \neq \infty$, then Player 2 has a strategy for which the finite outcome of the game from v reaches u in a play that is not losing, with a value of $C(u)$. Also, every non-losing acyclic outcome of the game that reaches u has a value at least $C(u)$.

The proof is similar to that of the Bellman-Ford algorithm, and it is by induction on the iteration in which $C(u)$ is updated.

It follows that if $C_0(v_{sink}) \leq 0$, then Player 2 has a strategy with an outcome that reaches v_{sink} in a non-losing play. Thus, Player 2 can win against the Player 1 strategy

ρ_1 , and it follows that ρ_1 does not have a corresponding almost memoryless winning strategy.

If after the call to $MINRMCR(v, c)$ a state $u \in V'$ has $C(u) = -\infty$, then Player 2 has a strategy that closes a cycle with negative value in which no prefix is losing. Similarly, if there is a vertex $v \in V'$ with $C_v(v) \leq C(v)$ and $b_v = true$, Player 2 has a strategy that reaches v from v_0 without losing, and from v , Player 2 can close a non-positive cycle.

If one of these cases occurs, clearly, ρ_1 does not force the game to a positive cycle, and thus it does not have a corresponding almost memoryless winning strategy.

For the other direction, assume the algorithm returns *true*. Assume by way of contradiction that there is a Player 2 strategy ρ_2 that does not lose. Consider the outcome of the game $\pi = out(\rho_1, \rho_2)$.

We distinguish between two cases. In the first case, π reaches v_{sink} . If π is acyclic, then since $C_0(v) > 0$, there is a non-losing path in the arena that ends in v_{sink} with a value less than $C_0(v)$, which is a contradiction to the claim above. Otherwise, there is a cycle in π . We construct a path π' by removing every non-negative valued cycle. Clearly $val(\pi') \leq val(\pi)$. If π' has no cycles and $val(\pi') \leq 0$, we reach a contradiction as before. Thus, there is a negative cycle in π' . Let $v = \pi[i] = \pi[j] = v$, where $i < j$, be the first negative cycle. Clearly, $\pi[i : j]$ is a non-losing cycle. Thus, $C(v) > C(v) + val(\pi[i : j]) \geq C_v(v)$, and we reach a contradiction to the fact the algorithm returns *true*.

In the second case, π does not reach v_{sink} . Similarly to the above, consider a cycle in π . Since π is not losing, there must be a cycle with non-positive value. Let $\pi[i : j]$ be the first such cycle. Every cycle before the index i is a positive cycle and it must start in a Player 2 vertex. Otherwise, Player 2 cannot escape the cycle and the play would have been winning for Player 1. Consider the path π' that is identical to $\pi[0 : i]$ only we trim out every positive cycle. Thus, π' reaches the vertex $\pi[i]$ with $val(\pi') \leq val(\pi[0 : i])$. Since π is not losing, π' is also not losing. Thus, $C_0(\pi[i]) \leq val(\pi') \leq 0$. Since $\pi[i : j]$ is a negative cycle, we have, similar to the above, $C_{\pi[i]}(\pi[i]) > C_0(\pi[i])$, and we reach a contradiction to the fact the algorithm returns *true*, and we are done. \square

Combining Lemma 9 and Lemma 10 we have the following theorem:

Theorem 6. *Solving simulation games is in $NP \cap co-NP$.*

4 An Abstraction-Refinement-based Algorithm for Deciding Simulation

In this section we solve the weighted-simulation problem $\mathcal{A} \leq \mathcal{B}$ by reasoning about abstractions of \mathcal{A} and \mathcal{B} . Recall that for every WFA \mathcal{A} , we have that $\mathcal{A}_\downarrow \subseteq \mathcal{A} \subseteq \mathcal{A}_\uparrow$. We first argue that this order is maintained for the simulation relation. We then use this fact in order to check simulation (and hence, also containment) with respect to abstractions.

Theorem 7. *For every WFA \mathcal{A} and abstraction function α , we have $\mathcal{A}_\downarrow^\alpha \leq \mathcal{A} \leq \mathcal{A}_\uparrow^\alpha$.*

Proof: We start with the first claim, namely $\mathcal{A}_\downarrow \leq \mathcal{A}$. We show that Player 2 has a winning strategy in the game that corresponds to \mathcal{A}_\downarrow and \mathcal{A} . Intuitively, whenever Player 1 selects a letter and a must-transition to proceed with in \mathcal{A}_\downarrow , the winning Player 2 strategy selects a matching concrete transition in \mathcal{A} and proceeds with it. Thus, the winning Player 2 strategy maintains the invariant that when the game reaches a vertex $\langle a, c \rangle$, then $c \in a$. Recall that \mathcal{A}_\downarrow under-approximates \mathcal{A} in three ways: the transition relation, the weight function, and the definition of the accepting states. Consequently, all the prefixes of the play are not winning for Player 1.

Formally, we define the Player 2 strategy $\rho_2 \in \mathcal{S}_2$ as follows. Let $\pi_{2k+1} = v_0, \dots, v_{2k+1}$ be the outcome of the game up to round $2k + 1$, when Player 2 plays ρ_2 and Player 1 plays some strategy. Let $v_{2k} = \langle a_k, c_k \rangle$ and $v_{2k+1} = \langle a_{k+1}, \sigma_{k+1}, c_k \rangle$. Thus, by the definition of the arena, $\Delta_{must}(a_k, \sigma_{k+1}, a_{k+1})$. We define ρ_2 to maintain the invariant $c_k \in a_k$. For $k = 0$ the invariant holds by the definition of the initial states in the arena. Assume $c_k \in a_k$. Since the transition $\langle a_k, \sigma_{k+1}, a_{k+1} \rangle$ is a must transition, there must be a state $c_{k+1} \in a_{k+1}$ such that $\Delta_{\mathcal{A}}(c_k, \sigma_{k+1}, c_{k+1})$. We define $\rho_2(\pi_{2k+1})$ to select one of these states arbitrarily and proceed to the vertex $\langle a_{k+1}, c_{k+1} \rangle$.

We claim that for every Player 1 strategy $\rho_1 \in \mathcal{S}_1$, and for every $k \geq 0$, Player 1 does not win the play $\pi_{2k} = v_0, \dots, v_{2k} = out(\rho_1, \rho_2)[0 : 2k]$. We start by showing that $v_{2k} \notin F_{must} \times (Q_{\mathcal{A}} \setminus F_{\mathcal{A}})$. Recall that a state $a \in A$ is accepting in \mathcal{A}_\downarrow iff all the concrete states that are mapped to it are accepting, i.e., $a \subseteq F_{\mathcal{A}}$. Since we defined ρ_2 to maintain the invariant that $c_k \in a_k$, if $a_k \in F_{must}$, then $c \in F_{\mathcal{A}}$.

Next, we show that if $v_{2k} \in F_{must} \times F_{\mathcal{A}}$, then $val(\pi_{2k}) \leq 0$. We prove this claim by induction on k . For $k = 0$, the value of the “empty-play” is 0, and the claim holds trivially. Assuming that the claim is correct for π_{2k-2} , i.e., $val(\pi_{2k-2}) \leq 0$, we show that it is correct for π_{2k} . We show that $\tau(v_{2k-2}, v_{2k-1}) + \tau(v_{2k-1}, v_{2k}) \leq 0$. Recall that $v_{2k-2} = \langle a_{k-1}, c_{k-1} \rangle$, $v_{2k-1} = \langle a_k, \sigma_k, c_{k-1} \rangle$, and $v_{2k} = \langle a_k, c_k \rangle$. Since \mathcal{A}_\downarrow uses the weight function τ_{min} , and since the concrete transition $\langle c_{k-1}, \sigma_k, c_k \rangle$ matches the must transition $\langle a_{k-1}, \sigma_k, a_k \rangle$, we have $\tau_{min}(a_{k-1}, \sigma_k, a_k) \leq \tau_{\mathcal{A}}(c_{k-1}, \sigma_k, c_k)$. We combine the above: $val(\pi_{2k}) = val(\pi_{2k-2}) + \tau(v_{2k-2}, v_{2k-1}) + \tau(v_{2k-1}, v_{2k}) = val(\pi_{2k-2}) + \tau_{min}(a_{k-1}, \sigma_k, a_k) - \tau_{\mathcal{A}}(c_{k-1}, \sigma_k, c_k) \leq \tau_{min}(a_{k-1}, \sigma_k, a_k) - \tau_{\mathcal{A}}(c_{k-1}, \sigma_k, c_k) \leq \tau_{min}(a_{k-1}, \sigma_k, a_k) - \tau_{min}(a_{k-1}, \sigma_k, a_k) = 0$, and we are done.

We proceed to prove the second claim, namely $\mathcal{A} \leq \mathcal{A}_\uparrow$. Recall that in the game that corresponds to \mathcal{A} and \mathcal{A}_\uparrow , Player 1 chooses the word the automata run on, and the run of \mathcal{A} . Player 2 chooses the run of \mathcal{A}_\uparrow . Intuitively, we construct a Player 2 winning strategy that responds to Player 1 moves by taking the may transition that matches the concrete transition that Player 1 takes. Similar to the previous claim, since \mathcal{A}_\uparrow over-approximates the transitions, the accepting states, and the weight function, Player 1 never wins.

Formally, we define $\rho_2 \in \mathcal{S}_2$ as follows. For $k \geq 0$, let $\pi_{2k+1} = v_0, \dots, v_{2k+1}$ be the outcome of the game in the first $2k + 1$ rounds, where Player 2 plays ρ_2 and Player 1 plays some strategy. Let $v_{2k} = \langle c_k, a_k \rangle$ and $v_{2k+1} = \langle c_{k+1}, \sigma_{k+1}, a_k \rangle$. We define ρ_2 to maintain the invariant $c_k \in a_k$. For the first round, $k = 0$, and we have $v_0 = \langle c_0, a_0 \rangle$, where c_0 is the initial state in \mathcal{A} and a_0 is the initial state in \mathcal{A}_\uparrow . By the definition of \mathcal{A}_\uparrow , we have $c_0 \in a_0$. For $k > 1$, assume $c_{k-1} \in a_{k-1}$. By the definition of the arena of the game, since $E(v_{2k}, v_{2k+1})$, we have $\Delta_{\mathcal{A}}(c_k, \sigma_{k+1}, c_{k+1})$. Consider the abstract

state $a_{k+1} \in A$ for which $c_{k+1} \in a_{k+1}$. By the definition of the may transitions, we have $\Delta_{may}(a_k, \sigma_{k+1}, a_{k+1})$. We define $\rho_2(\pi_{2k+1}) = \langle c_{k+1}, a_{k+1} \rangle$.

We claim that Player 1 never wins. We prove that at round $2k$ both conditions for termination do not hold on the play π_{2k} . We start by showing that $v_{2k} \notin F_A \times (A \setminus F_{may})$. Recall that a state $a \in A$ is accepting in \mathcal{A}_\uparrow iff there is a concrete accepting state that is mapped to it, i.e., $a \cap F_A \neq \emptyset$. Thus, for $v_{2k} = \langle c_k, a_k \rangle$, since we defined ρ_2 to maintain the invariant $c_k \in a_k$, we have that if $c_k \in F_A$, then $a_k \in F_{may}$.

To conclude the proof, we prove by induction on k that $val(\pi_{2k}) \leq 0$. For $k = 0$, the “empty-play” has a value of 0 and the claim is trivial. Assuming $val(\pi_{2k-2}) \leq 0$, we prove that $val(\pi_{2k}) \leq 0$. We show that $\tau(v_{2k-2}, v_{2k-1}) + \tau(v_{2k-1}, v_{2k}) \leq 0$. Recall that $v_{2k-2} = \langle c_{k-1}, a_{k-1} \rangle$, $v_{2k-1} = \langle c_k, \sigma_k, a_{k-1} \rangle$, and $v_{2k} = \langle c_k, a_k \rangle$. By the definition of the weight function in \mathcal{A}_\uparrow , namely τ_{max} , since the concrete transition $\langle c_{k-1}, \sigma_k, c_k \rangle$ matches the may transition $\langle a_{k-1}, \sigma_k, a_k \rangle$, we have $\tau_{max}(a_{k-1}, \sigma_k, a_k) \geq \tau_A(c_{k-1}, \sigma_k, c_k)$. Since $val(\pi_{2k}) = val(\pi_{2k-2}) + \tau(v_{2k-2}, v_{2k-1}) + \tau(v_{2k-1}, v_{2k})$, $\tau(v_{2k-2}, v_{2k-1}) = \tau_A(c_{k-1}, \sigma_k, c_k)$, and $\tau(v_{2k-1}, v_{2k}) = -\tau_{max}(a_{k-1}, \sigma_k, a_k)$, combining the above with the induction hypothesis, we have $val(\pi_{2k}) \leq 0$, and we are done. \square

We note that beyond the use of Theorem 7 in practice, it provides an additional witness to the appropriateness of our definition of weighted simulation.

Recall that our algorithm solves the weighted-simulation problem $\mathcal{A} \leq \mathcal{B}$ by reasoning about abstractions of \mathcal{A} and \mathcal{B} . We first show that indeed we can conclude about the existence of simulation or its nonexistence by reasoning about the abstractions:

Theorem 8. *Consider two WFAs \mathcal{A} and \mathcal{B} and abstraction functions α and β .*

- If $\mathcal{A}_\uparrow^\alpha \leq \mathcal{B}_\downarrow^\beta$, then $\mathcal{A} \leq \mathcal{B}$.
- If $\mathcal{A}_\downarrow^\alpha \not\leq \mathcal{B}_\uparrow^\beta$, then $\mathcal{A} \not\leq \mathcal{B}$.

Proof: We start with the first claim. By Theorem 7, we know that $\mathcal{A} \leq \mathcal{A}_\uparrow^\alpha$ and $\mathcal{B}_\downarrow^\beta \leq \mathcal{B}$. Thus, by Theorem 4 (transitivity of \leq), we know that $\mathcal{A}_\uparrow^\alpha \leq \mathcal{B}_\downarrow^\beta$ implies that $\mathcal{A} \leq \mathcal{B}$. For the second claim, assume, by way of contradiction, that $\mathcal{A} \leq \mathcal{B}$. By Theorem 7 we know that $\mathcal{A}_\downarrow^\alpha \leq \mathcal{A}$ and $\mathcal{B} \leq \mathcal{B}_\uparrow^\beta$. Hence, transitivity implies that $\mathcal{A}_\downarrow^\alpha \leq \mathcal{B}_\uparrow^\beta$, contradicting the claim’s assumption. \square

Our algorithm proceeds as follows. We start by checking whether $\mathcal{A}_\uparrow^\alpha \leq \mathcal{B}_\downarrow^\beta$ and $\mathcal{A}_\downarrow^\alpha \not\leq \mathcal{B}_\uparrow^\beta$, for some (typically coarse) initial abstraction functions α and β . By Theorem /refinement-thm, if we are lucky and one of them holds, we are done. Otherwise, the winning strategies of the Player 2 in the first case and Player 1 in the second case suggest a way to refine α and β , and we repeat the process with the refined abstractions. While refinement in the Boolean case only splits abstract states in order to close the gap between may and must transitions, here we also have refinement steps that tighten the weights along transitions. Below is a formal description of the algorithm.

Input: Two WFAs \mathcal{A} and \mathcal{B} , with abstraction functions α and β

Output: *yes* if $\mathcal{A} \leq \mathcal{B}$ and *no* otherwise

while true do
 if Player 2 wins the game that corresponds to $\mathcal{A}_\uparrow^\alpha$ and $\mathcal{B}_\downarrow^\beta$ **then** return *yes*
 else let ρ_1 be a winning Player 1 strategy in the game
 if Player 1 wins the game that corresponds to $\mathcal{A}_\downarrow^\alpha$ and $\mathcal{B}_\uparrow^\beta$ **then** return *no*
 else let ρ_2 be a winning Player 2 strategy in the game
 $\alpha', \beta' = \text{refine}(\mathcal{A}, \mathcal{B}, \alpha, \beta, \rho_1, \rho_2)$
 set: $\alpha = \alpha'$ and $\beta = \beta'$
end while

By Theorem 8, if the algorithm returns *yes*, then \mathcal{A} simulates \mathcal{B} , and if the algorithm returns *no*, then \mathcal{A} does not simulate \mathcal{B} . If $\mathcal{A}_\uparrow^\alpha \not\leq \mathcal{B}_\downarrow^\beta$ and $\mathcal{A}_\downarrow^\alpha \leq \mathcal{B}_\uparrow^\beta$, then the answer is indefinite and we refine the abstractions. This is done by the procedure *refine*, described below.

Recall that we refine α and β in case both $\mathcal{A}_\uparrow^\alpha \not\leq \mathcal{B}_\downarrow^\beta$ and $\mathcal{A}_\downarrow^\alpha \leq \mathcal{B}_\uparrow^\beta$. When this happens, the algorithm for checking simulation generates a winning strategy ρ_1 of Player 1 in the game corresponding to $\mathcal{A}_\uparrow^\alpha \not\leq \mathcal{B}_\downarrow^\beta$ and a winning strategy ρ_2 of Player 2 in the game corresponding to $\mathcal{A}_\downarrow^\alpha \leq \mathcal{B}_\uparrow^\beta$. The procedure *refine* gets as input the WFAs \mathcal{A} and \mathcal{B} , the abstraction functions α and β , and the winning strategies ρ_1 and ρ_2 . It returns two new abstraction functions α' and β' .

In order to see the idea behind *refine*, assume that Player 1 wins in the concrete game. Then, ρ_2 is winning in a *spurious* manner in the game that corresponds to $\mathcal{A}_\downarrow^\alpha \leq \mathcal{B}_\uparrow^\beta$. Our goal in the refinement process is to remove at least one of the reasons ρ_2 is winning. Also, if Player 1 wins in the concrete game, then refinement in the game corresponding to $\mathcal{A}_\uparrow^\alpha \leq \mathcal{B}_\downarrow^\beta$ should reveal the fact that ρ_1 is winning. The situation is dual if Player 2 wins the concrete game. Since during the refinement process we cannot know which of the players wins the concrete game, we perform refinements to the two games simultaneously until we reach a definite answer. Since we assume that the WFAs are finite, we are guaranteed to eventually terminate. Thus, our procedure is complete (it attempts, however, to solve only the simulation, rather than containment, problem).

Observe that the arenas on which the strategies are defined have the same vertices but different edges. Edges that appear in one game but not the other correspond to may transitions that are not must transitions. Our refinement procedure is based on the algorithm for solving the simulation game as described in Section 3.2. Recall that the algorithm first performs a pre-processing stage in which it removes two sets of vertices: vertices from which Player 1 wins (namely the set W_1), and vertices from which Player 2 wins (namely the set W_2). The first set of vertices are the winning vertices in the un-weighted reachability game with objective $F_{\mathcal{A}} \times (Q_{\mathcal{B}} \setminus F_{\mathcal{B}})$. The second set of vertices are the winning vertices in the un-weighted safety game with objective $((Q_{\mathcal{A}} \setminus F_{\mathcal{A}}) \times Q_{\mathcal{B}}) \cup V_2$.

Since the two winning sets depend on the edges of the game, the sets we remove are not the same in the two games. We refer to the vertices after their removal as $V'_{\mathcal{A}_\uparrow, \mathcal{B}_\downarrow}$ and $V'_{\mathcal{A}_\downarrow, \mathcal{B}_\uparrow}$, and we refine them until the initial vertex is in both sets.

We describe the refinement according to the strategy ρ_1 . Refinement according to ρ_2 is dual.

Recall that, by Theorem 7, if Player 1 wins, then he has an almost-memoryless winning strategy. Thus, we assume ρ_1 is almost memoryless. Also recall that ρ_1 is winning in the game played on the vertices $V'_{\mathcal{A}_\uparrow, \mathcal{B}_\downarrow}$, and since ρ_2 is winning in the game played on the vertices $V'_{\mathcal{A}_\downarrow, \mathcal{B}_\uparrow}$, the strategy ρ_1 is not winning in this game.

We proceed as in the algorithm for solving simulation games: we “guess” the strategy ρ_1 and check if (actually, how) Player 2 can win against this strategy. Since ρ_1 is not winning in the game played on the vertices $V'_{\mathcal{A}_\uparrow, \mathcal{B}_\downarrow}$, we find at least one path π that is winning for Player 2. As seen in the algorithm, π is either a path that reaches v_{sink} or is a lasso contained in the vertices $V'_{\mathcal{A}_\uparrow, \mathcal{B}_\downarrow} \setminus \{v_{sink}\}$. More formally, π is of the form $\pi_1 \cdot \pi_2^\omega$. The path π_1 is a simple path that uses vertices from $V'_{\mathcal{A}_\uparrow, \mathcal{B}_\downarrow} \setminus \{v_{sink}\}$ and it (and every prefix of it) is not losing for Player 2. The path π_2 is either the cycle that is the single vertex v_{sink} or it is a cycle contained in $V'_{\mathcal{A}_\uparrow, \mathcal{B}_\downarrow} \setminus \{v_{sink}\}$. In the second case, $val(\pi_2) \leq 0$, and for every $0 \leq i \leq |\pi_2|$, if $\pi_2[i] \in F_{\mathcal{A}} \times F_{\mathcal{B}}$, then $val(\pi_1) + val(\pi_2[0 : i]) \leq 0$.

Since π is not a path in $V'_{\mathcal{A}_\downarrow, \mathcal{B}_\uparrow}$, at least one of the following three cases hold:

- π uses a vertex in $V'_{\mathcal{A}_\uparrow, \mathcal{B}_\downarrow} \setminus V'_{\mathcal{A}_\downarrow, \mathcal{B}_\uparrow}$,
- π traverses an edge that corresponds to a may but not must transition, or
- the sum of the edges traversed in π is larger in the one game than in the other.

In the first case, we refine the vertices $V'_{\mathcal{A}_\uparrow, \mathcal{B}_\downarrow}$ and $V'_{\mathcal{A}_\downarrow, \mathcal{B}_\uparrow}$, as described above. In the second case, the refinement is similar to the one done in the Boolean setting, where we close the gap between may and must transitions. Finally, in the third case, we split states in order to tighten the weights on the transitions. Recall that these weights are defined by taking the minimum or maximum of the corresponding set of transitions. Therefore, splitting of states indeed tightens the weights.

Example 2. Consider the two simulation games \mathcal{G}_1 and \mathcal{G}_2 in Figure 2. The game \mathcal{G}_1 corresponds to \mathcal{A}_\uparrow and \mathcal{B}_\downarrow , and \mathcal{G}_2 corresponds to \mathcal{A}_\downarrow and \mathcal{B}_\uparrow , for some two WFAs \mathcal{A} and \mathcal{B} with abstraction functions. In the figure, we use circle and boxes in order to denote, respectively, the nodes in which Player 1 and Player 2 proceed. In \mathcal{G}_1 , we define $F_{\mathcal{A}} \times (Q_{\mathcal{B}} \setminus F_{\mathcal{B}}) = \{s_4\}$, $F_{\mathcal{A}} \times F_{\mathcal{B}} = \{s_5, s_6\}$, and $(Q_{\mathcal{A}} \setminus F_{\mathcal{A}}) \times Q_{\mathcal{B}} = \{s_4\}$, and the definition is similar in \mathcal{G}_2 . Due to lack of space, we omit the letters from the arenas.

We show that Player 1 wins \mathcal{G}_1 in three different ways and Player 2 wins \mathcal{G}_2 . In the first strategy, in \mathcal{G}_1 , Player 1 proceeds from s_0 to s_1 . Player 2 is then forced to continue to s_4 , which is losing for Player 2 since it is a vertex in $F_{\mathcal{A}} \times (Q_{\mathcal{B}} \setminus F_{\mathcal{B}})$. In the second winning strategy, Player 1 proceeds from s_0 to s_2 . The game continues by alternating between s_2 and s_6 . Since the cycle has a positive value and $s_6 \in F_{\mathcal{A}} \times F_{\mathcal{B}}$, Player 1 wins the prefix $s_0 s_2 s_6 s_2 s_6 s_2 s_6$. Finally, in the third strategy, Player 1 proceeds from s_0 to s_3 . Player 2 is then forced to proceed to s_6 . Since $s_6 \in F_{\mathcal{A}} \times F_{\mathcal{B}}$ and $val(s_0 s_2 s_6) > 0$, Player 1 wins the prefix. Clearly, in \mathcal{G}_2 , the Player 2 strategy that proceeds from s_1 to s_5 is winning.

We proceed to describe the refinement of the abstractions using these strategies. First, note that in \mathcal{G}_1 , by playing the first strategy described above, Player 1 can force the game to a vertex in $F_{\mathcal{A}} \times (Q_{\mathcal{B}} \setminus F_{\mathcal{B}})$ from the initial vertex. Thus, $s_0 \notin V'_{\mathcal{A}_\uparrow, \mathcal{B}_\downarrow}$. We start by refining the set of Player 1 winning vertices in the reachability game with objective $F_{\mathcal{A}} \times (Q_{\mathcal{B}} \setminus F_{\mathcal{B}})$. In this process we refine the vertex s_1 .

Next, we apply the third Player 1 winning strategy on G_2 and see how Player 2 can win against it. Player 2 wins because Player 1 uses the edge $\langle s_0, s_3 \rangle$, which is not in G_2 . We refine s_0 , and after the refinement the strategy is no longer valid for Player 1 in G_1 . After these two refinements, Player 1 can still win in G_1 using the second strategy, and we apply it in G_2 . The outcome of the game against a winning Player 2 strategy is $s_0(s_2s_6)^\omega$. In this path, we find the failure vertex s_2 and refine it in order to tighten the values of the edges.

The two resulting games after these three refinements are G'_1 and G'_2 (see the right side of Figure 2). Player 1 wins in G'_2 by proceeding from s_0 to s_2 , and thus we are done.

Note that since not all the values on the edges are the same in G'_1 and G'_2 , the refinement is not exhausted. That is, the arenas are not $Q_A \times Q_B$. Thus, the abstraction-refinement algorithm successfully decides simulation on a smaller state space than the concrete one. Since, however, we found that $\mathcal{A} \not\subseteq \mathcal{B}$, then by Theorem 2, it might still be the case that $\mathcal{A} \subseteq \mathcal{B}$.




Fig. 3. An example of applying the refinement algorithm on two simulation games.

5 Directions for Future Research

We introduced the notions of abstraction and simulation for weighted automata and argue that they form a useful heuristic for checking containment – a problem of practical interest that is known to be undecidable. In the Boolean setting, researchers have suggested ways for closing the gap between containment and simulation [23,26]. Some, like these that extend the definition of simulation with a look ahead, are easy to extend to the weighted setting. Other ways require special treatment of the accumulated weights and are subject to future research. Finally, the rich weighted setting allows one to measure the differences between systems. For example, we can talk about one WFA t -approximating another WFA, in the sense that the value of a word in the second is at most t times its value in the first [2]. Our weighted simulation corresponds to the special case $t = 1$ and we plan to study approximated weighted simulation.

References

1. S. Almagor, U. Boker, and O. Kupferman. What’s decidable about weighted automata? In *9th ATVA*, LNCS 6996, pages 482–491, 2011.

2. B. Aminof, O. Kupferman, and R. Lampert. Formal analysis of online algorithms. In *9th ATVA*, LNCS 6996, page 213–227 Springer, 2011.
3. T. Ball, E. Bounimova, B. Cook, V. Levin, J. Lichtenberg, C. McGarvey, B. Ondrusek, S.K. Rajamani, and A. Ustuner. Thorough static analysis of device drivers. In *EuroSys*, 2006.
4. K. Larsen N. Markey P. Bouyer, U. Fahrenberg and J. Srba. Infinite runs in weighted timed automata with energy constraints. In *FORMATS*, pages 33–47, 2008.
5. G. Bruns and P. Godefroid. Model checking partial state spaces with 3-valued temporal logics. In *Proc. 11th CAV*, pages 274–287, 1999.
6. T. Ball and O. Kupferman. An abstraction-refinement framework for multi-agent systems. In *Proc. 21st LICS*, 2006.
7. U. Boker and O. Kupferman. Co-ing Büchi made tight and helpful. In *Proc. 24th LICS*, pages 245–254, 2009.
8. P. Cousot and R. Cousot. Abstract interpretation: a unified lattice model for the static analysis of programs by construction or approximation of fixpoints. In *Proc. 4th POPL*, pages 238–252, 1977.
9. K. Chatterjee and L. Doyen. Energy parity games. In *Proc. 37th ICALP*, pages 599–610, 2010.
10. K. Chatterjee, L. Doyen, and T. Henzinger. Quantitative languages. In *Proc. 17th CSL*, pages 385–400, 2008.
11. K. Chatterjee, L. Doyen, and T. Henzinger. Probabilistic weighted automata. In *Proc. 20th CONCUR*, pages 224–258, 2009.
12. K. Chatterjee, L. Doyen, and T. A. Henzinger. Expressiveness and closure properties for quantitative languages. *LMCS*, 6(3), 2010.
13. E. M. Clarke, O. Grumberg, S. Jha, Y. Lu, and H. Veith. Counterexample-guided abstraction refinement for symbolic model checking. *Journal of the ACM*, 50(5):752–794, 2003.
14. P. Černý, T. Henzinger, and A. Radhakrishna. Simulation distances. In *Proc. 21st CONCUR*, pages 253–268, 2010.
15. K. Culik and J. Kari. Digital images and formal languages. *Handbook of formal languages, vol. 3: beyond words*, pages 599–616, 1997.
16. L. de Alfaro and P. Roy. Solving games via three-valued abstraction refinement. In *Proc. 18th CONCUR*, pages 74–89, 2007.
17. M. Droste and P. Gastin. Weighted automata and weighted logics. In *Proc. 32nd ICALP*, pages 513–525, 2005.
18. O. Grumberg, M. Lange, M. Leucker, and S. Shoham. When not losing is better than winning: Abstraction and refinement for the full μ -calculus. *I&C*, 205(8):1130–1148, 2007.
19. D. Gale and F. M. Stewart. Infinite games of perfect information. *Ann. Math. Studies*, 28:245–266, 1953.
20. M.R. Henzinger, T.A. Henzinger, and P.W. Kopke. Computing simulations on finite and infinite graphs. In *Proc. 36th FOCS*, pages 453–462, 1995.
21. T. Henzinger, R. Jhala, and R. Majumdar. Counterexample-guided control. In *Proc. 30th ICALP*, pages 886–902, 2003.
22. T.A. Henzinger, R. Majumdar, F.Y.C. Mang, and J-F Raskin. Abstract interpretation of game properties. In *Proc. 7th SAS*, LNCS 1824, pages 245–252, 2000.
23. Y. Kesten, N. Piterman, and A. Pnueli. Bridging the gap between fair simulation and trace containment. In *Proc. 15th CAV*, LNCS 2725, pages 381–393, 2003.
24. D. Krob. The equality problem for rational series with multiplicities in the tropical semiring is undecidable. *International Journal of Algebra and Computation*, 4(3):405–425, 1994.
25. D. Kuperberg. Linear temporal logic for regular cost functions. In *Proc. 28th STACS*, pages 627–636, 2011.
26. N. A. Lynch and M.R. Tuttle. Hierarchical correctness proofs for distributed algorithms. In *Proc. 6th PODC*, pages 137–151, 1987.

27. K.G. Larsen and G.B. Thomsen. A modal process logic. In *Proc. 3rd LICS*, 1988.
28. R. Milner. An algebraic definition of simulation between programs. In *Proc. 2nd IJCAI*, pages 481–489, 1971.
29. M. Mohri. Finite-state transducers in language and speech processing. *Computational Linguistics*, 23(2):269–311, 1997.
30. M. Mohri, F.C.N. Pereira, and M. Riley. Weighted finite-state transducers in speech recognition. *Computer Speech and Language*, 16(1):69–88, 2002.
31. A.R. Meyer and L.J. Stockmeyer. The equivalence problem for regular expressions with squaring requires exponential time. In *Proc. 13th SWAT*, pages 125–129, 1972.
32. M.P. Schützenberger. On the definition of a family of automata. *Information and Control*, 4(2-3):245–270, 1961.
33. A.P. Sistla, M.Y. Vardi, and P. Wolper. The complementation problem for Büchi automata with applications to temporal logic. *Theoretical Computer Science*, 49:217–237, 1987.