

# Looking at Mean-Payoff through Foggy Windows

Paul Hunter\*      Guillermo A. Pérez†

Jean-François Raskin\*

{phunter, gperezme, jraskin}@ulb.ac.be

Université Libre de Bruxelles – Brussels, Belgium

August 8, 2018

## Abstract

Mean-payoff games (MPGs) are infinite duration two-player zero-sum games played on weighted graphs. Under the hypothesis of perfect information, they admit memoryless optimal strategies for both players and can be solved in  $\text{NP} \cap \text{coNP}$ . MPGs are suitable quantitative models for open reactive systems. However, in this context the assumption of perfect information is not always realistic. For the partial-observation case, the problem that asks if the first player has an observation-based winning strategy that enforces a given threshold on the mean-payoff, is undecidable. In this paper, we study the window mean-payoff objectives that were introduced recently as an alternative to the classical mean-payoff objectives. We show that, in sharp contrast to the classical mean-payoff objectives, some of the window mean-payoff objectives are decidable in games with partial-observation.

## 1 Introduction

*Mean-payoff games* (MPGs) [13] are infinite duration, two-player, zero-sum games played on weighted graphs, useful for modelling reactive systems with quantitative objectives and designing algorithms to synthesize controllers for such systems [6]. Like other verification games played on graphs, two players move a token around the graph for an infinite number of steps. One of the players selects a label, after which the second chooses an edge with this label. The token is then moved along the selected edge. This infinite interaction between the two players results in an infinite path in the graph. The objective of Player 1 is to maximize the limiting average payoff of the edges (defined by the weights that annotate them) traversed in this infinite path, while Player 2 tries to minimize this average. It has been shown in [3, 13] that both players in an MPG can play optimally using memoryless strategies, and as a consequence, those games are known to be solvable in  $\text{NP} \cap \text{coNP}$ , and although pseudo-polynomial-time algorithms to solve these games are known [5, 22], the question of whether they can be solved in polynomial-time is an important open question.

In the version of MPG described above, the game is of *perfect information*: both players have complete knowledge of the history of the play up to the current position of the token. For many applications such as controller synthesis, it is often more natural to assume that players have only partial knowledge of the current state of the game. In practice, players may model processes with private variables that other players (processes) may not see, or controllers that acquire information about their environment using sensors with bounded precision, etc. Unfortunately, it has been shown in [10] that MPGs games with *partial-observation* are undecidable.

*Window mean-payoff objectives* (WMP objectives) were recently introduced in [9] as an alternative to the classical MP objectives. In a WMP objective instead of considering the long-run average along the whole play, payoffs are considered over a local bounded window *sliding* along the play. The objective is

---

\*Authors supported by the ERC inVEST (279499) project.

†Author supported by F.R.S.-FNRS fellowship.

then to make sure that the average payoff is at least zero over every window. The WMP objectives enjoy several nice properties. First, in contrast to classical MP games, we have a polynomial-time algorithm for determining WMP games. Second, they can be considered as approximation of the classical MP objectives in the following sense: (i) they are a *strengthening* of the MP objective, i.e. winning for the WMP objective implies winning for the MP objective, (ii) if a (finite memory) strategy is forcing a MP with value  $\varepsilon > 0$  then that strategy is also enforcing the WMP objective for a window size that is bounded by a function of the game and strategy memory sizes.

In this paper we consider the extension of WMP objectives to games with partial-observation. We show that, in sharp contrast with classical MP objectives, some of the WMP objectives are decidable for such games. As in [9], we consider several variants of the window MP objectives. For all objectives, we provide complete complexity results and optimal algorithms. More precisely, our main contributions are as follows:

- First, we consider a definition in which the window size is fixed and the sliding window is started at the initial move of the game, this is called the *direct window* objective. For this definition we give an optimal EXP-time algorithm (Theorem 3) in the form of a reduction to a *safety game*. Additionally, we show that this safety game has a nice structure that induces a natural partial order on game positions. In turn this partial order can be used to obtain a *symbolic algorithm* based on the antichain approach [12]. This shows that WMP objectives allow us not only to recover decidability but they also lead to games that have the potential to be solved efficiently in practice. The antichain approach has already been applied and implemented with success for LTL synthesis [4], omega-regular games with partial observation [2], and language inclusion between non-deterministic Büchi automata [11].
- Second, we consider two natural prefix-independent definitions for the window objectives, the (*uniform*) *fixed window* objectives. We also give optimal EXP-time algorithms for these two definitions (Theorem 5 and Theorem 6), when weights are polynomially bounded in the size of the game arena. For these objectives, we show that the sets of good abstract plays (sequences of observations) of the game with partial observation, for the two definitions, form regular languages whose complements can be recognized by non-deterministic Büchi automata of pseudo-polynomial size (Proposition 2 and Proposition 3). These automata can then be turned into deterministic parity automata that can be used as observers to transform the game of partial-observation into a game of perfect information with a parity objective.
- Finally, we show that, when the size of the window is not fixed but rather left as a parameter, then for all the objectives that we consider the decision problems are undecidable (Theorem 2).

## 2 Preliminaries

**Weighted game arenas.** A *weighted game arena with partial-observation* (WGA, for brevity) is a tuple  $G = \langle Q, q_I, \Sigma, \Delta, w, Obs \rangle$ , where  $Q$  is a finite set of states,  $q_I \in Q$  is the initial state,  $\Sigma$  is a finite set of actions,  $\Delta \subseteq Q \times \Sigma \times Q$  is the transition relation,  $w : \Delta \rightarrow \mathbb{Z}$  is the weight function, and  $Obs \in \text{Partition}(Q)$  is a set of observations containing  $\{q_I\}$ . Let  $W = \max\{|w(t)| : t \in \Delta\}$ . We assume  $\Delta$  is total, i.e. for every  $(q, \sigma) \in Q \times \Sigma$  there exists  $q' \in Q$  such that  $(q, \sigma, q') \in \Delta$ . If every element of  $Obs$  is a singleton, then we say  $G$  is a *WGA with perfect information* and if  $|Obs| = 1$  we say  $G$  is *blind*. For simplicity, we denote by  $\text{post}_\sigma(s) = \{q' \in Q \mid \exists q \in s : (q, \sigma, q') \in \Delta\}$  the set of  $\sigma$ -successors of a set of states  $s \subseteq Q$ .

In this work, unless explicitly stated otherwise, we depict states from a WGA as circles and transitions as arrows labelled by an action-weight pair:  $\sigma, x \in \Sigma \times \{-W, \dots, W\}$ . Observations are represented by dashed boxes and colors, where states with the same color correspond to the same observation.

**Abstract and concrete paths.** A *concrete path* in an WGA is a sequence  $q_0\sigma_0q_1\sigma_1\dots$  where for all  $i \geq 0$  we have  $q_i \in Q$ ,  $\sigma_i \in \Sigma$  and  $(q_i, \sigma_i, q_{i+1}) \in \Delta$ . An *abstract path* is a sequence  $o_0\sigma_0o_1\sigma_1\dots$  where  $o_i \in Obs$ ,  $\sigma_i \in \Sigma$  and such that there is a concrete path  $q_0\sigma_0q_1\sigma_1\dots$  for which  $q_i \in o_i$ , for all  $i$ . Given an abstract path  $\psi$ , let  $\gamma(\psi)$  be the set of concrete paths that agree with the observation and action sequence. Formally

$\gamma(\psi) = \{q_0\sigma_0q_1\sigma_1\dots \mid \forall i \geq 0 : q_i \in o_i \text{ and } (q_i, \sigma, q_{i+1}) \in \Delta\}$ . Also, given abstract (respectively concrete) path  $\rho = o_0\sigma_0\dots$  and integers  $k, l$  we define  $\pi[k..l] = o_k\dots o_l$ ,  $\pi[..k] = \pi[0..k]$ , and  $\pi[l..] = o_l\sigma_l o_{l+1}\dots$ .

Given a concrete path  $\pi = q_0\sigma_0q_1\sigma_1\dots$ , the *payoff* up to the  $(n+1)$ -th element is given by

$$w(\pi[..n]) = \sum_{i=0}^{n-1} w(q_i, \sigma_i, q_{i+1}).$$

If  $\pi$  is infinite, we define two *mean-payoff* values  $\underline{\text{MP}}$  and  $\overline{\text{MP}}$  as:

$$\underline{\text{MP}}(\pi) = \liminf_{n \rightarrow \infty} \frac{1}{n} w(\pi[..n]) \qquad \overline{\text{MP}}(\pi) = \limsup_{n \rightarrow \infty} \frac{1}{n} w(\pi[..n])$$

**Plays and strategies.** A play in a WGA  $G$  is an infinite abstract path starting at  $o_I \in \text{Obs}$  where  $q_I \in o_I$ . Denote by  $\text{Plays}(G)$  the set of all plays and by  $\text{Prefs}(G)$  the set of all finite prefixes of such plays ending in an observation. Let  $\gamma(\text{Plays}(G))$  be the set of concrete paths of all plays in the game, and  $\gamma(\text{Prefs}(G))$  be the set of all finite prefixes of all concrete paths.

An *observation-based strategy for Eve* is a function from finite prefixes of plays to actions, i.e.  $\lambda_{\exists} : \text{Prefs}(G) \rightarrow \Sigma$ . A play  $\psi = o_0\sigma_0o_1\sigma_1\dots$  is *consistent* with  $\lambda_{\exists}$  if  $\sigma_i = \lambda_{\exists}(\psi[..i])$  for all  $i$ . We say an observation-based strategy for Eve  $\lambda_{\exists}$  has *memory*  $\mu$  if there is a set  $M$  with  $|M| = \mu$ , an element  $m_0 \in M$ , and functions  $\alpha_u : M \times \text{Obs} \rightarrow M$  and  $\alpha_o : M \times \text{Obs} \rightarrow \Sigma$  such that for any play prefix  $\rho = o_0\sigma_0\dots o_n$  we have  $\lambda_{\exists}(\rho) = \alpha_o(m_n, o_n)$ , where  $m_n$  is defined inductively by  $m_{i+1} = \alpha_u(m_i, o_i)$  for  $i \geq 0$ .

**Objectives.** An *objective* for a WGA  $G$  is a set  $V_G$  of plays, i.e.  $V_G \subseteq \text{Plays}(G)$ . We say plays in  $V_G$  are *winning for Eve*. Conversely, all plays not in  $V_G$  are *winning for Adam*. We refer to a WGA with a fixed objective as a game. Having fixed a game, we say a strategy  $\lambda$  is *winning* for a player if all plays consistent with  $\lambda$  are winning for that player. Finally, we say that a player *wins* a game if (s)he has a winning strategy. We write  $V$  instead of  $V_G$  if  $G$  is clear from the context.

Given WGA  $G$  and a threshold  $\nu \in \mathbb{Q}$ , the *mean-payoff* (MP) objectives  $\text{MPSup}_G(\nu) = \{\psi \in \text{Plays}(G) \mid \forall \pi \in \gamma(\psi) : \overline{\text{MP}}(\pi) \geq \nu\}$  and  $\text{MPInf}_G(\nu) = \{\psi \in \text{Plays}(G) \mid \forall \pi \in \gamma(\psi) : \underline{\text{MP}}(\pi) \geq \nu\}$  require the mean-payoff value be at least  $\nu$ . We omit the subscript in the objective names when the WGA is clear from the context. Let  $\nu = \frac{a}{b}$ ,  $w'$  be a weight function mapping  $t \in \Delta$  to  $b \cdot w(t) - a$ , for all such  $t$ , and  $G'$  be the WGA resulting from replacing  $w'$  in  $G$  for  $w$ . We note that Eve wins the  $\text{MPSup}_{G'}(0)$  (respectively,  $\text{MPInf}_{G'}(0)$ ) objective if and only if she wins  $\text{MPSup}_G(\nu)$  (resp.,  $\text{MPInf}_G(\nu)$ ).

### 3 Window Mean-payoff Objectives

In what follows we recall the definitions of the *window mean-payoff* (WMP) objectives introduced in [9] and adapt them to the partial-observation setting. For the classical MP objectives Eve is required to ensure the long-run average of all concretizations of the play is at least  $\nu$ . WMP objectives correspond to conditions which are sufficient for this to be the case. All of them use as a main ingredient the concept of concrete paths being “good”. We formalize this notion below.

Given  $i \geq 0$  and *window size bound*  $l_{\max} \in \mathbb{N}_0 = \mathbb{N} \setminus \{0\}$ , let the set of concrete paths  $\chi$  with the *good window* property be

$$\text{GW}(\nu, i, l_{\max}) = \{\chi \mid \exists j \leq l_{\max} : \frac{w(\chi[i..(i+j)])}{j} \geq \nu\}.$$

As in [9], we assume that the value of  $l_{\max}$  is polynomially bounded by the size of the arena.

For the first of the WMP objectives Eve is required to ensure that all suffixes of all concretizations of the play can be split into concrete paths of length at most  $l_{\max}$  and average weight at least  $\nu$ . The MP objectives are known to be prefix-independent, therefore a prefix-independent version of this first objective is a natural objective to consider as well. We study two such candidates. One which asks of Eve that there is some  $i$  such that all suffixes – after  $i$  – of all concretizations of the play can be split in the same way as before. This




Figure 1: Implications among the objectives




Figure 2: Blind WGA where, for any  $l_{\max} \in \mathbb{N}_0$ , the only possible abstract play is in  $\text{Fix}(l_{\max})$  but not in  $\text{UFix}(l_{\max})$ .




Figure 3: Perfect information WGA where Eve wins both MP objectives but none of the FWMP or BWMP objectives.

is quite restrictive since the  $i$  is *uniform* for all concretizations of the play. The second prefix-independent version of the objective we consider allows for non-uniformity.

Formally, the *fixed window mean-payoff* (FWMP) objectives for a given WGA and threshold  $\nu \in \mathbb{Q}$  are defined below. For convenience we denote by  $\psi$  plays from  $\text{Plays}(G)$  and concrete plays by  $\pi$ , i.e. elements of  $\gamma(\text{Plays}(G))$ .

$$\begin{aligned}
\text{DirFix}(\nu, l_{\max}) &= \{\psi \mid \forall \pi \in \gamma(\psi), \forall i \geq 0 : \pi \in \text{GW}(\nu, i, l_{\max})\} \\
\text{UFix}(\nu, l_{\max}) &= \{\psi \mid \exists i \geq 0, \forall \pi \in \gamma(\psi), \forall j \geq i : \pi \in \text{GW}(\nu, j, l_{\max})\} \\
\text{Fix}(\nu, l_{\max}) &= \{\psi \mid \forall \pi \in \gamma(\psi), \exists i \geq 0, \forall j \geq i : \pi \in \text{GW}(\nu, j, l_{\max})\}
\end{aligned}$$

For the FWMP objectives, we consider  $l_{\max}$  to be a value that is given as input. Another natural question that arises is whether we can remove this input and consider an even weaker objective in which one asks if there exists an  $l_{\max}$ . This is captured in the definition of the *bounded window mean-payoff* (BWMP) objectives which are defined for a given threshold  $\nu \in \mathbb{Q}$ .

$$\begin{aligned}
\text{UDirBnd}(\nu) &= \{\psi \mid \exists l_{\max} \in \mathbb{N}_0, \forall \pi \in \gamma(\psi), \forall i \geq 0 : \pi \in \text{GW}(\nu, i, l_{\max})\} \\
\text{DirBnd}(\nu) &= \{\psi \mid \forall \pi \in \gamma(\psi), \exists l_{\max} \in \mathbb{N}_0, \forall i \geq 0 : \pi \in \text{GW}(\nu, i, l_{\max})\} \\
\text{UBnd}(\nu) &= \{\psi \mid \exists l_{\max} \in \mathbb{N}_0, \exists i \geq 0, \forall \pi \in \gamma(\psi), \forall j \geq i : \pi \in \text{GW}(\nu, j, l_{\max})\} \\
\text{Bnd}(\nu) &= \{\psi \mid \forall \pi \in \gamma(\psi), \exists l_{\max} \in \mathbb{N}_0, \exists i \geq 0, \forall j \geq i : \pi \in \text{GW}(\nu, j, l_{\max})\}
\end{aligned}$$

As with the mean-payoff objectives we can assume, without loss of generality, that  $\nu = 0$ . Henceforth, we omit  $\nu$ .

### 3.1 Relations among objectives

Figure 1 gives an overview of the relative strengths of each of the objectives and how they relate to the mean-payoff objective. The strictness, in general, of most inclusions was established in [9], and Figure 2 provides an example for the remaining case between  $\text{Fix}$  and  $\text{UFix}$ .

In general the mean-payoff objective is not sufficient for the FWMP or BWMP objectives, e.g. see Figure 3. Our first result shows that if, however, Eve has a *finite memory* winning strategy for a *strictly positive* threshold, then this strategy is also winning for any of the FWMP or BWMP objectives. A specific subcase of this was first observed in Lemma 2 of [9].




Figure 4: Gadget which forces Eve to play infinitely many #.




Figure 5: Gadget which, given that Eve will play # infinitely often, forces her to play # in intervals of bounded length.

**Theorem 1.** *Given WGA  $G$ , if Eve has a finite memory winning strategy for the  $MPInf(\varepsilon)$  (or  $MPSup(\varepsilon)$ ) objective, for  $\varepsilon > 0$ , then the same strategy is winning for her in the  $DirFix(\mu)_G$  game – where  $\mu$  is bounded by the memory used by the strategy.*

### 3.2 Lower bounds

In [9] it was shown that in multiple dimensions, with arbitrary window size, solving games with the (direct) fixed window objective was complete for EXP-time. We now show that in our more general setting this hardness result holds, even when the window size is a fixed constant and the weight function is given in unary.

**Lemma 1.** *Let  $l_{max} \in \mathbb{N}_0$  be a fixed constant. Given WGA  $G$ , determining if Eve has a winning strategy for the  $DirFix(l_{max})$ ,  $UFix(l_{max})$  or the  $Fix(l_{max})$  objectives is EXP-hard, even for unary weights.*

*Proof.* We give a reduction from the problem of determining the winner of a safety game with imperfect information, shown in [7] to be EXP-complete.

A safety game with imperfect information is played on a non-weighted game arena with partial-observation  $G = \langle Q, q_I, \Sigma, \Delta, Obs \rangle$ . A play of  $G$  is winning for Eve if and only if it never visits the *unsafe state* set  $\mathcal{U} \subseteq Q$ . Without loss of generality, we assume unsafe states are *trapping*, that is  $(u, \sigma, q) \in \Delta$  and  $u \in \mathcal{U}$  imply that  $u = q$ .

Let  $w$  be the transition weight function mapping  $(u, \sigma, q) \in \Delta$  to  $-1$  if  $u \in \mathcal{U}$  and all other  $t \in \Delta$  to  $0$ . Denote by  $G_w$  the resulting WGA from adding  $w$  to  $G$ . It should be clear that Eve wins the safety game  $G$  if and only if she wins  $MPInf_{G_w}(0)$ ,  $DirFix_{G_w}(l_{max})$ ,  $UFix_{G_w}(l_{max})$ , and  $Fix_{G_w}(l_{max})$  – for any  $l_{max}$ . That is, all objectives are equivalent for  $G_w$ . □ □

In [9] the authors show that determining if Eve has a winning strategy in the  $k$ -dimensional version of the  $UDirBnd$  and  $UBnd$  objectives with perfect information is non-primitive recursive hard. We show that, in our setting, these decision problems are undecidable.

**Theorem 2.** *Given WGA  $G$ , determining if Eve has a winning strategy for any of the BWMP objectives is undecidable, even if  $G$  is blind.*

*Sketch.* Our proof is by a reduction from the universality problem of weighted finite automata which is undecidable [1]. A *weighted finite automaton* is a tuple  $\mathcal{N} = \langle Q, \Sigma, q_I, \Delta, w, F \rangle$  where  $F \subseteq Q$  is a set of final accepting states. A (*accepting*) *run* of the automaton on a word  $x = \sigma_0 \sigma_1 \dots \sigma_n \in \Sigma^*$  is a sequence  $r = q_0 q_1 \dots q_n \in Q^+$  such that  $q_n \in F$  and  $(q_i, \sigma_i, q_{i+1}) \in \Delta$  for all  $0 \leq i < n$ . The cost of the run  $r$  is  $w(r) = \sum_{i=0}^{n-1} w(q_i, \sigma_i, q_{i+1})$ . If the automaton is non-deterministic, it may have several runs on  $x$ . In that case, the *cost of  $x$  in  $\mathcal{N}$*  (denoted by  $\mathcal{N}(x)$ ) is defined as the minimum of the costs of all its accepting runs on  $x$ .

The *universality* problem for weighted automata is to decide whether, for a given automaton  $\mathcal{N}$ , the following holds:  $\forall x \in \Sigma^* : \mathcal{N}(x) < 0$ . If so,  $\mathcal{N}$  is *universal*. Our reduction constructs a blind WGA,  $G_{\mathcal{N}}$ , so that:




Figure 6: Blind gadget to simulate the weighted automaton  $\mathcal{N}$ .

- if  $\mathcal{N}$  is universal, then Eve has a winning strategy for the objective  $\text{UDirBnd}$ ,
- if  $\mathcal{N}$  is not universal, then Adam has a winning strategy for the complement of the objective  $\text{Bnd}$ .

As shown in Fig. 1,  $\text{UDirBnd} \subseteq \text{Bnd}$  and all the other BWMP objectives lie in between those two. So, our reduction establishes the undecidability of all BWMP objectives at once.

Our reduction follows the gadgets given in Figs. 4–6. When the game starts, Adam chooses to play from one of the three gadgets. As the game is blind for Eve, she does not know what is the choice of Adam and so she must be prepared for all possibilities. Note also that as Eve is blind, her strategy can be formalized by an infinite word  $w \in \Sigma \cup \{\#\}^\omega$ . The first two gadgets force Eve to play a word  $w$  such that:  $(C_1)$  there are infinitely many  $\#$  in  $w$ , and  $(C_2)$  there exists a bound  $b \in \mathbb{N}$  such that the distance between two consecutive  $\#$  in  $w$  is bounded by  $b$ . Note that although  $(C_1)$  holds with under a mean-payoff objective, the  $\text{Bnd}$  objective is necessary to ensure  $(C_2)$ .

Now, we assume that Eve plays a word  $w = \#w_1\#w_2\#w_3\#\dots\#w_n\#\dots$  that respects conditions  $C_1$  and  $C_2$ , and we consider what happens when Adam plays in the third gadget (Fig. 6). The best strategy for Adam is to simulate the accepting runs of  $\mathcal{N}$ . If  $\mathcal{N}$  is not universal then Eve chooses  $w_1 \in \Sigma^*$  such that  $\mathcal{N}(w_1) \geq 0$ , and then she plays the strategy  $(\#w_1)^\omega$  and she wins the objective  $\text{UDirBnd}$ . Indeed Eve plays a finite memory strategy that forces a mean-payoff which is at least  $\frac{0.5}{b} > 0$  as each new  $\#$  brings  $+\frac{1}{2}$  and we know that  $\mathcal{N}(w_1) \geq 0$ . The result then follows from Theorem 1. If  $\mathcal{N}$  is universal, no matter which word  $w$  is played by Eve, the mean-payoff of the outcome of the game will be at most  $\frac{-0.5}{b}$  as any word  $w_1 \in \Sigma^*$  played between two consecutive  $\#$  is such that  $\mathcal{N}(w_1) \leq -1$ . Thus, Adam wins for the complement of the mean-payoff objective. From the relation between the mean-payoff objective and  $\text{Bnd}$  (see Fig. 1) it follows that Adam wins for the complement of the objective  $\text{Bnd}$ .  $\square$   $\square$

## 4 Solving DirFix games

In this section we establish an upper bound to match our lower bound of Section 3.2 for determining the winner of DirFix games. We first observe that for WGAs with perfect information the  $\text{DirFix}(l_{\max})$  objective has the flavor of a safety objective. Intuitively, a play  $\pi$  is winning for Eve if every suffix of  $\pi$  has a prefix of length at most  $l_{\max}$  with average weight of at least 0. As soon as the play reaches a point for which this does not hold, Eve loses the play. In WGAs with partial-observation we need to make sure the former holds for all concretizations of an abstract play.

We construct a non-weighted game arena with perfect information  $G'$  from  $G$ . Eve's objective in  $G'$  will consist in ensuring the play never reaches locations in which there is an open window of length  $l_{\max}$ , for some state. This corresponds to a safety objective. Whether Eve wins the new game can be determined in time linear w.r.t. the size of the new game (see, e.g. [21]). The game will be played on a set of functions  $\mathcal{F}$  which is described in detail below. We then show how to transfer winning strategies of Eve from  $G'$  to  $G$

and vice versa in Lemmas 5 and 6. Hence, this yields an algorithm to determine if Eve wins the  $\text{DirFix}(l_{\max})$  objective which runs in exponential time.

**Theorem 3.** *Given WGA  $G$ , determining if Eve has a winning strategy for the  $\text{DirFix}(l_{\max})$  objective is EXP-complete.*

Let us define the functions which will be used as the state space of the game. Intuitively, we keep track of the *belief* of Eve as well as the windows with the minimal weight open at every state of the belief.<sup>1</sup>

For the rest of this section let us fix a WGA with partial-observation  $G$  and a window size bound  $l_{\max} \in \mathbb{N}_0$ . We begin by defining the set of functions  $\mathcal{F}$  as the set of all functions  $f : Q \rightarrow (\{1, \dots, l_{\max}\} \rightarrow \{-W \cdot l_{\max}, \dots, 0\}) \cup \{\perp\}$ . Denote by  $\text{supp}(f)$  the *support* of  $f$ , that is the set of states  $q \in Q$  such that  $f(q) \neq \perp$ . For  $q \in \text{supp}(f)$ , we denote by  $f(q)_i$  the value  $f(q)(i)$ . The function  $f_I \in \mathcal{F}$  is such that  $f_I(q)_l = 0$ , for all  $1 \leq l \leq l_{\max}$ , and  $f_I(q) = \perp$  for all  $q \in Q \setminus \{q_I\}$ . Given  $f_1 \in \mathcal{F}$  and  $\sigma \in \Sigma$ , we say  $f_2 \in \mathcal{F}$  is a  $\sigma$ -successor of  $f_1$  if

- $\text{supp}(f_2) = \text{post}_\sigma(\text{supp}(f_1)) \cap o$  for some  $o \in \text{Obs}$ ;
- for all  $q \in \text{supp}(f_2)$  and all  $1 \leq j \leq l_{\max}$  we have that  $f_2(q)_j$  maps to  $\max\{-W \cdot l_{\max}, \min\{0, \zeta(q)\}\}$ , where  $\zeta(q)$  is defined as follows

$$\zeta(q) = \begin{cases} \min_{\substack{p \in \text{supp}(f_1), \\ (p, \sigma, q) \in \Delta, \\ f_1(p)_{j-1} < 0}} f_1(p)_{j-1} + w(p, \sigma, q) & \text{if } j \geq 2 \\ \min_{\substack{p \in \text{supp}(f_1), \\ (p, \sigma, q) \in \Delta}} w(p, \sigma, q) & \text{otherwise.} \end{cases}$$

**Lemma 2.** *The number of elements in  $\mathcal{F}$  is not greater than  $2^{|Q| \cdot l_{\max} \cdot \log(W \cdot l_{\max})}$ .*

We extend the  $\text{supp}$  operator to finite sequences of functions and actions. In other words, given  $\rho' = f_0 \sigma_0 f_1 \sigma_1 \in (\mathcal{F} \cdot \Sigma)^*$ ,  $\text{supp}(\rho') \mapsto s_0 \sigma_0 s_1 \sigma_1 \dots$  where  $s_i = \text{supp}(f_i)$  for all  $i \geq 0$ . In an abuse of notation, we define the function  $\text{supp}^{-1} : (\text{Obs} \cdot \Sigma)^* \times \mathcal{F} \rightarrow (\mathcal{F} \cdot \Sigma)^*$  which maps abstract paths to function-action sequences. Formally, given  $\rho = o_0 \sigma_0 o_1 \sigma_1 \dots \in \text{Prefs}(G)$  and  $\varphi \in \mathcal{F}$  with  $\text{supp}(\varphi) \subseteq o_0$ ,  $\text{supp}^{-1}(\rho, \varphi) \mapsto f_0 \sigma_0 f_1 \sigma_1 \dots$  where  $f_0 = \varphi$  and for all  $i \geq 0$  we have that  $f_{i+1}$  is the  $\sigma_i$ -successor of  $f_i$  such that  $\text{supp}(f_{i+1}) \subseteq o_{i+1}$ . Both  $\text{supp}$  and  $\text{supp}^{-1}$  are extended to infinite sequences in the obvious manner.

The following two results enunciate the key properties of sequences of the form  $(\mathcal{F} \cdot \Sigma)^*$ . Intuitively, the set of all those sequences corresponds to the windowed, weighted unfolding of  $G$  with information about reachable states as well as open windows.

**Lemma 3.** *Let  $\rho = o_0 \sigma_0 \dots o_n$  be an abstract path,  $\varphi \in \mathcal{F}$  such that  $\text{supp}(\varphi) \subseteq o_0$  and  $\text{supp}^{-1}(\rho, \varphi) = f_0 \sigma_0 \dots f_n \in (\mathcal{F} \cdot \Sigma)^*$ . A state  $q \in Q$  is reachable from some state  $q_0 \in \text{supp}(\varphi)$  through a concrete path  $q_0 \sigma_0 \dots q_n \in \gamma(\rho)$  if and only if  $q \in \text{supp}(f_n)$ .*

Consider an abstract path  $\psi$  and a positive integer  $n$ . We say a window of length  $l$  is open at  $q \in \gamma(\psi[n])$  if there is some concretization  $\chi$  of  $\psi[..n]$  with  $q = \chi[n]$  such that  $\chi \notin \text{GW}(n-l, l)$ .

**Lemma 4.** *Let  $\rho = o_0 \sigma_0 \dots o_n$  be an abstract path,  $\varphi \in \mathcal{F}$  such that  $\text{supp}(\varphi) \subseteq o_0$  and  $\text{supp}^{-1}(\rho, \varphi) = f_0 \sigma_0 \dots f_n \in (\mathcal{F} \cdot \Sigma)^*$ . Given state  $p \in \text{supp}(f_n)$  and  $1 \leq l \leq l_{\max}$  such that  $l \leq n$ , then there is a window of length  $l$  open at  $p$  if and only if  $f_n(p)_l < 0$ .*

Formally, the arena  $G' = \langle \mathcal{F}, f_I, \Sigma, \Delta' \rangle$ . The transition relation  $\Delta'$  contains the transition  $(f_1, \sigma, f_2)$  if  $f_2$  is the  $\sigma$ -successor of  $f_1$ . Eve, in  $G'$ , is required to avoid states  $\mathcal{U} = \{f \in \mathcal{F} \mid \exists q \in \text{supp}(f) : f(q)_{l_{\max}} < 0\}$ .

<sup>1</sup>The terms belief and knowledge are used to denote a state from any variation of the classic ‘‘Reif construction’’ [18] to turn a game with partial-observation into a game with perfect information. Other names for similar constructions include ‘‘knowledge-based subset construction’’ (see e.g. [10]).

**Lemma 5.** *If Eve wins the safety objective in  $G'$ , then she also wins the  $\text{DirFix}(l_{\max})$  objective in  $G$ .*

*Proof.* Assume  $\lambda'$  is a winning strategy for Eve in  $G'$ . We define a strategy  $\lambda$  for her in  $G$  as follows:

$$\lambda(\rho) \mapsto \lambda'(\text{supp}^{-1}(\rho, f_I))$$

for all  $\rho \in \text{Prefs}(G)$ . We claim that  $\lambda$  is winning for her in  $G$ . Towards a contradiction, assume  $\psi \in \text{Plays}(G)$  is consistent with  $\lambda$  and that  $\psi \notin \text{DirFix}(l_{\max})$ . Recall that this implies there are  $n \in \mathbb{N}, q \in Q$  such that there is a window of length  $l_{\max}$  open at  $q \in \gamma(\psi[n])$ . By Lemma 4 we then get that  $f_n$  from  $\text{supp}^{-1}(\psi, f_I) = f_0\sigma_0f_1\sigma_1\dots$  is in  $\mathcal{U}$ . As  $\text{supp}^{-1}(\psi, f_I)$  is consistent with  $\lambda'$ , this contradicts the assumption that  $\lambda'$  was winning.  $\square$   $\square$

**Lemma 6.** *If Eve wins the  $\text{DirFix}(l_{\max})$  objective in  $G$ , then she also wins the safety objective in  $G'$ .*

*Proof.* Assume  $\lambda$  is a winning strategy for Eve in  $G$ . We define a strategy  $\lambda'$  for her in  $G'$  as follows:

$$\lambda'(\rho') \mapsto \lambda \circ \text{obs} \circ \text{supp}(\rho')$$

for all  $\rho' \in \text{Prefs}(G')$ . We claim that  $\lambda'$  is winning for her in  $G'$ . Again, towards a contradiction, assume  $\psi' \in \text{Plays}(G')$  is consistent with  $\lambda'$  and that  $\psi'$  visits some  $f \in \mathcal{U}$ . This implies, by Lemma 4, that there is a window of length  $l_{\max}$  open at some  $q \in \text{supp}(f)$  in  $\psi = \text{obs}(\text{supp}(\psi'))$ . As  $\Delta$  is total, for any  $\sigma \in \Sigma$  Eve plays then there is valid  $\sigma$ -successor of  $q$  that Adam can choose as the next state. Hence there is some  $\chi \in \text{Plays}(G)$  consistent with  $\lambda$  such that  $\chi$  and  $\psi$  have the same prefix up  $i_q$ , where  $q \in \gamma(\chi[i_q])$ , and there is a concretization  $\pi$  of  $\chi$  such that  $\pi[i_q] = q$ . As  $\chi$  is consistent with  $\lambda$  and  $\chi \notin \text{DirFix}(l_{\max})$ , this contradicts the fact that it was a winning strategy.  $\square$   $\square$

## 4.1 A symbolic algorithm for DirFix games

We note that state space of the construction  $G'$  presented in Section 4 admits an order such that if a state is smaller than another state, according to said order, and Eve has a strategy to win from the latter, then she has a strategy to win from the former. In this section we formalize this notion by defining the order and, in line with [4, 8], propose an *antichain*-based algorithm to solve the safety game on  $G'$ .

We define the *uncontrollable predecessors* operator  $\text{UPre} : \mathcal{P}(\mathcal{F}) \rightarrow \mathcal{P}(\mathcal{F})$  as follows:

$$\text{UPre}(S) = \{p' \in \mathcal{F} \mid \forall \sigma \in \Sigma, \exists q' \in S : (p', \sigma, q') \in \Delta'\}.$$

For  $S \in \mathcal{P}(\mathcal{F})$ , we denote by  $\mu X.(S \cup \text{UPre}(X))$ , the *least fixpoint* of the function  $F : X \rightarrow S \cup \text{UPre}(X)$  in the  $\mu$ -calculus notation (see [14]). Note that  $F$  is defined on the powerset lattice, which is finite. The following is a well-known result about the relationship between safety games and the  $\text{UPre}$  operator (see e.g. [15]).

**Proposition 1.** *Eve wins a safety game with unsafe state set  $\mathcal{U}$  if and only if the initial state of the game is not contained in  $\mu X.(\mathcal{U} \cup \text{UPre}(X))$ .*

**Definition 1** (The partial order). *Given  $f', g' \in \mathcal{F}$  we say  $f' \preceq g'$  if and only if  $\text{supp}(f') \subseteq \text{supp}(g')$  and*

$$\forall q \in \text{supp}(f'), \forall i \in \{1, \dots, l_{\max}\}, \exists j \in \{i, \dots, l_{\max}\} : f'(q)_i \geq g'(q)_j.$$

An *antichain* is a non-empty set  $S \in \mathcal{P}(\mathcal{F})$  such that for all  $x, y \in S$  we have  $x \not\preceq y$ . We denote by  $\mathfrak{A}$  the set of all antichains. Given  $a, b \in \mathfrak{A}$ , denote by  $a \sqsubseteq b$  the fact that  $\forall x \in b, \exists y \in a : y \preceq x$ . For  $S \in \mathcal{P}(\mathcal{F})$  we denote by  $\lfloor S \rfloor$  the set of minimal elements of  $S$ , that is  $\lfloor S \rfloor = \{x \in S \mid \forall y \in S : y \preceq x \text{ implies } y = x\}$ . Clearly  $\lfloor S \rfloor$  is an antichain.

Given  $S \in \mathcal{P}(\mathcal{F})$  we denote by  $S^\uparrow$  the *upward-closure* of  $S$ , that is  $S^\uparrow = \{t \in \mathcal{F} \mid S \preceq t\}$ . We say a set  $s \in \mathcal{P}(\mathcal{F})$  is *upward-closed* if  $S = S^\uparrow$ . Note that  $\lfloor S \rfloor^\uparrow = S^\uparrow$  and therefore, if  $S$  is upward-closed, the antichain  $\lfloor S \rfloor$  is a succinct representation of  $S$ .

**Lemma 7.** *The following assertions hold.*






Figure 7: For  $n > l_{\max} + 1$  the abstract path  $(o_0 \dots o_n)^\omega$  is winning for the Fix condition but infinitely often visits an unsafe state in the construction from Section 4.

1.  $\mathcal{U}$  is upward-closed.
2. If  $S, T \in \mathcal{P}(\mathcal{F})$  are two upward-closed sets, then  $S \cup T$  is also upward-closed.

The usual way of showing an antichain algorithm works dictates that we now prove the  $\text{UPre}$  operator, when applied to upward-closed sets, outputs an upward-closed set as well. Unfortunately, this is not true in our case. The following example illustrates this difficulty.

**Example 1.** Consider the WGA from Figure 3 and let  $l_{\max} = 2$ . We note that the function  $f$  such that  $f(q_0) = \perp$  and  $f(q_1)_1 = 1, f(q_1)_2 = 0$  is in  $\text{UPre}(\mathcal{U})$ . We also have that for the function  $g$  such that  $g(q_0) = \perp$  and  $g(q_1)_1 = 0, g(q_1)_2 = 1$  we get that  $f \preceq g$ . It is easy to verify that  $g \notin \text{UPre}(\mathcal{U})$ . Hence,  $\text{UPre}(\mathcal{U})$  is not upward-closed.

However, we claim that one can circumvent this issue by ignoring elements from  $\mathcal{U}$ . Thus we are able to prove that, under some conditions,  $\text{UPre}$  does preserve “upward-closedness”.

**Lemma 8.** Given upward-closed set  $S \in \mathcal{P}(\mathcal{F})$  and  $f, g \in \mathcal{F} \setminus \mathcal{U}$ , if  $f \in \text{UPre}(S)$  and  $f \preceq g$ , then  $g \in \text{UPre}(S)$ .

We define a version of the uncontrollable predecessors’ operator which manipulates antichains instead of subsets of  $\mathcal{F}$ .

$$[\text{UPre}](a) = [\{p' \in \mathcal{F} \setminus \mathcal{U} \mid \forall \sigma \in \Sigma, \exists q' \in a, \exists r' \in \mathcal{F} : (p', \sigma, r') \in \Delta' \wedge q' \preceq r'\}]$$

Given  $a, b \in \mathfrak{A}$  we denote by  $a \sqcup b$  the least upper bound of  $a$  and  $b$ , i.e.  $a \sqcup b = [\{q' \in \mathcal{F} \mid q' \in a \text{ or } q' \in b\}]$ . It is easy to check that  $(a \sqcup b)^\uparrow = a^\uparrow \cup b^\uparrow$  for any  $a, b \in \mathfrak{A}$ .

**Theorem 4.** Given WGA  $G$ , Eve wins the  $\text{DirFix}(l_{\max})$  objective if and only if  $\{q'_I\} \not\sqsupseteq \mu X.([\mathcal{U}] \sqcup [\text{UPre}](X))$ .

## 5 Solving Fix games

Since Fix games are a prefix-independent version of  $\text{DirFix}$  games, it seems logical to consider an analogue of the perfect information game from the previous section with a prefix-independent condition. Indeed, the reader might be tempted to extend the approach used to solve  $\text{DirFix}$  games by replacing the safety objective with a *co-Büchi* objective in order to solve  $\text{UFix}$  or  $\text{Fix}$  games. However, we observe that although Eve winning in the resulting game is sufficient for her to win the original  $\text{Fix}$  game, it is not necessary. Indeed, an abstract play visits states from  $\mathcal{U}$  infinitely often if and only if for infinitely many  $i$  there is a concretization of the play prefix up to  $i$  which violates  $\text{GW}(i, l_{\max})$ . Nevertheless, this does not imply there exists one (infinite) concretization of the play which violates  $\text{GW}(i, l_{\max})$  for infinitely many  $i$ . Figure 7 illustrates this phenomenon.

For the reasons stated above, we propose to solve  $\text{Fix}$  games in a different way. We first introduce the notion of *observer*. Let  $\mathcal{A}$  be a deterministic parity automaton.<sup>2</sup> We say  $\mathcal{A}$  is an *observer for the objective*  $V$  if the language of  $\mathcal{A}$  is  $V$ , i.e.  $\mathcal{L}(\mathcal{A}) = V$ . In [7], the authors show that the synchronized product of  $G$

<sup>2</sup>We refer the reader who is not familiar with parity automata or games to [21].

and an observer for  $V$  is a parity game of perfect information which is won by Eve if and only if she wins  $G$ . Thus, it suffices to find an algorithm to construct an observer for  $\text{Fix}(l_{\max})$  to be able to solve Fix games.

For convenience, we start by describing a non-deterministic machine that accepts as its language the complement of  $\text{Fix}(l_{\max})$ . Note that all elements of  $\text{Fix}(l_{\max})$  start with the observation  $\{q_I\}$  so it suffices to describe the machine that accepts any word  $w \in (\Sigma \cdot \text{Obs})^\omega$  such that  $\{q_I\} \cdot w \in \text{Plays}(G) \setminus \text{Fix}(l_{\max})$ . The machine we construct is similar to the one used in [7] to make objectives of imperfect information games visible. Intuitively, at each step of the game and after Adam has revealed the next observation we will guess his actual choice of state using non-determinism. Additionally, we shall guess whether or not a violating window starts at the next step. The state space of the automaton will therefore consist of a single state from  $Q$ , a negative integer to record the weight of the tracked window, and the length of the current open window.

Formally, let  $\mathcal{N}$  be the automaton consisting of the state space  $F = Q \times \{1, \dots, l_{\max}\} \times \{-W \cdot l_{\max}, \dots, -1\} \cup \{\perp\}$ ; initial state  $(q_I, 1, \perp)$ ; input alphabet  $\Sigma' = \Sigma \times \text{Obs}$ ; and  $\Delta'' \subseteq F \times \Sigma' \times F$ . The transition relation  $\Delta''$  has a transition  $((p, i, n), (\sigma, o), (q, j, m))$  if  $(p, \sigma, q) \in \Delta$ ,  $q \in o$ ,

$$m = \begin{cases} w(p, \sigma, q) & \text{if } w(p, \sigma, q) < 0 \\ n + w(p, \sigma, q) & \text{if } n \neq \perp \wedge n + w(p, \sigma, q) < 0 \wedge i < l_{\max} \\ \perp & \text{otherwise,} \end{cases}$$

and

$$j = \begin{cases} i + 1 & \text{if } m = n + w(p, \sigma, q) \\ 1 & \text{otherwise.} \end{cases}$$

We say a state  $(q, i, n) \in F$  is accepting if  $i = l_{\max}$  and  $n \neq \perp$ . The automaton accepts a word  $x$  if and only if it has a run  $(q_0, i_0, n_0) (\sigma_0, o_1) (q_1, i_1, n_1) (\sigma_1, o_2) \dots$  on  $x$  such that for infinitely many  $j$  we have that  $(q_j, i_j, n_j)$  is accepting.

**Proposition 2.** *The non-deterministic Büchi automaton  $\mathcal{N}$  accepts a word  $\psi \in \text{Plays}(G)$  if and only if  $\psi \notin \text{Fix}(l_{\max})$ .*

At this point we determinize  $\mathcal{N}$  and complement it to get a deterministic automaton with state space of size exponential in the size of  $\mathcal{N}$  yet has parity index that is polynomial w.r.t. the size of  $Q$  (see [17, 19, 20]). The synchronized product of  $G$  and the observer yields a parity game with the same size bounds. The desired result follows from the parity games' algorithm and results of [16].

**Theorem 5.** *Given WGA  $G$ , determining if Eve has a winning strategy for the  $\text{Fix}(l_{\max})$  objective can be decided in time exponential in  $W$  and the size of  $G$ .*

**Corollary 1.** *Given WGA  $G$  with unary encoded weights, deciding if Eve has a winning strategy for the  $\text{Fix}(l_{\max})$  objective is EXP-complete.*

## 6 Solving UFix games

In order to determine the winner of UFix games, we proceed as in the previous section by finding a non-deterministic Büchi automaton that recognizes the set of bad abstract plays. However, in this case the situation is more complicated because a bad abstract play might arise from a violation in the *uniformity*, rather than because of a concrete path with infinitely many window violations. Figure 2 illustrates this issue. To overcome this, we first provide an alternative characterization of the bad abstract plays for Eve. Consider some  $\psi \in \text{Plays}(G)$ . We say  $\pi \in \gamma(\psi)$  merges with infinitely many violating paths if for all  $i \geq 0$ , there are  $j \geq i$ ,  $k \geq j + l_{\max}$  and some  $\chi \in \gamma(\psi[..k])$  such that  $\pi[k] = \chi[k]$  and  $\chi \notin \text{GW}(j, l_{\max})$ . We refer to  $j$  as the position of the violation and to  $k$  as the position of the merge. Our next result formally states the relationship between concrete plays merging for multiple violations and UFix games.

**Lemma 9.** *Given WGA  $G$  and  $\psi \in \text{Plays}(G)$ , there is  $\pi \in \gamma(\psi)$  merging with infinitely many violating paths if and only if  $\psi \notin \text{UFix}(l_{\max})$ .*

*Proof.* ( $\Rightarrow$ ) Assume there is a  $\pi \in \gamma(\psi)$  merging with infinitely many violating paths. We have that there are two infinite sequences of indices  $J = \{j_0, j_1, \dots\}$  and  $K = \{k_0, k_1, \dots\}$  such that  $j_l < j_{l+1}$  and  $j_l + l_{\max} \leq k_l$ , for all  $l \geq 0$ , and for which we know that there is concrete path  $\chi_l \in \gamma(\psi[..k_l])$  such that  $\chi_l[j_l..j_l + l_{\max}]$  realizes an open window of length  $l_{\max}$  and  $\chi_l[k_l] = \pi[k_l]$ , for all  $l \geq 0$ . Observe that for all  $l \geq 0$  we have that  $\chi_l \cdot \pi[k_l..] \in \gamma(\psi)$  and that  $\chi_l \cdot \pi[k_l..] \notin \text{GW}(j_l, l_{\max})$ . In other words,

$$\forall l \geq 0, \exists \alpha \in \gamma(\psi), \exists m \geq l : \alpha \notin \text{GW}(m, l_{\max})$$

which implies that  $\psi \notin \text{UFix}(l_{\max})$ .

( $\Leftarrow$ ) Assume  $\psi \notin \text{UFix}(l_{\max})$ . We have that there is a infinite sequence of indices  $J = \{j_0, j_1, \dots\}$  such that  $j_k < j_{k+1}$ , for all  $k \geq 0$ , and for which we know there is a concrete play  $\pi_k \in \gamma(\psi)$  such that  $\pi_k \notin \text{GW}(j_k - l_{\max}, l_{\max})$ , for all  $k \geq 0$ . Observe that for all  $i \geq 0$  the set  $\gamma(\psi[i])$  is finite and bounded by  $|Q|$ . Thus, by Pigeonhole Principle we have that, for all  $n \geq 0$  there is  $\eta_n \in \{\pi_m \mid 1 \leq m \leq |Q| \cdot n\} \subseteq \gamma(\psi)$  which merges with at least  $n$  violating paths. Consider an arbitrary  $\eta_1$ . If  $\eta_1$  merges with infinitely many violating paths then we are done and the claim holds. Otherwise it only merges with a finite number of violating paths, say  $a_1$ . From the previous argument we know there is an  $\eta_{a'_1} \in \gamma(\psi)$  that merges with at least  $a'_1 = a_1 + 1$ . Clearly  $\eta_1$  and  $\eta_{a'_1}$  are disjoint at every point after  $a'_1$ , lest  $\eta_1$  would merge with a new violating path. We inductively repeat the process, if  $\eta_{a'_i}$  merges with infinitely many violating paths then we are done. Otherwise it only merges with some finite number of violating paths, say  $a_i$ . In that case we turn our attention to  $\eta_{a'_{i+1}}$ . Note that since  $Q$  is finite this process can only be done a finite number of times. Indeed, after having discarded at most  $|Q| - 1$  concrete plays (which are disjoint after some finite point) it must be the case the last remaining possible concrete play has the desired property or we would have a contradiction with our assumptions. Thus, there is some concrete play  $\pi \in \gamma(\psi)$  that merges with infinitely many violating paths.  $\square$

We now construct the non-deterministic Büchi automaton  $\mathcal{N}'$  that recognizes plays that contain a concrete path that merges with infinitely many violating paths. The idea is that we non-deterministically keep track of two paths: one that will eventually witness a violation and then merge with the other, which ultimately serves as the witness for the path that merges with infinitely many violating paths. When the two paths merge, the automaton non-deterministically selects a new path to witness the violation. This is achieved by non-deterministically selecting any state in the belief set of Eve as these states represent the end states of any concrete play consistent with the abstract play so far. To avoid the double exponential associated with taking the Reif construction before determinizing the automaton, we instead compute the belief set on-the-fly using a Moore machine that feeds into our non-deterministic automaton. By transferring the exponential state increase to an exponential increase in the alphabet size, the overall size of the determinized automaton (after composition with the Moore machine) will be at most singly exponential in the size of our game and  $W$ .

More specifically, denote by  $\mathcal{B}$  the machine that, given  $\psi = o_0\sigma_0o_1\sigma_1 \dots \in \text{Plays}(G)$  as its input yields the infinite sequence  $o_0\sigma_0s_0o_1\sigma_1s_1 \dots \in (\text{Obs} \cdot \Sigma \cdot \mathcal{P}(Q))^\omega$  such that  $s_0 = \{q_I\}$  and for all  $i \geq 0$  we have  $s_{i+1} = \text{post}_{\sigma_i}(s_i)$ . One can easily give a definition of  $\mathcal{B}$  – which closely resembles a subset construction – with a state space at most exponential w.r.t.  $G$ . Observe that the machine realizes a continuous function, in the sense that every prefix of length  $i$  of the input uniquely defines the next  $s_{i+1}$  annotation. Thus, the annotation can be done on-the-fly.

Formally,  $\mathcal{N}'$  consists of the state space  $F' = Q \times Q \times \{1, \dots, l_{\max}\} \times \{-W \cdot l_{\max}, \dots, -1\} \cup \{\perp, \top\}$ ; initial state  $(q_I, q_I, 1, \perp)$ ; input alphabet  $\Sigma'' = \Sigma \times \text{Obs} \times \mathcal{P}(Q)$ ; and  $\Delta''' \subseteq F \times \Sigma'' \times F$ . The transition relation

$\Delta'''$  has a transition  $((p, p', i, n), (\sigma, o, s), (q, q', j, m))$  if  $(p', \sigma, q') \in \Delta$ ,  $q \in s$ ,  $q' \in o$ ,

$$m = \begin{cases} w(p, \sigma, q) & \text{if } (p, \sigma, q) \in \Delta \wedge w(p, \sigma, q) < 0 \\ n + w(p, \sigma, q) & \text{if } (p, \sigma, q) \in \Delta \wedge n \neq \perp \wedge n + w(p, \sigma, q) < 0 \wedge i < l_{\max} \\ \top & \text{if } (p, \sigma, q) \in \Delta \wedge (n \neq \top \vee p \neq p') \wedge n \neq \perp \wedge i = l_{\max} \\ \perp & \text{otherwise,} \end{cases}$$

and

$$j = \begin{cases} l_{\max} & \text{if } m = \top \\ i + 1 & \text{if } m = n + w(p, \sigma, q) \\ 1 & \text{otherwise.} \end{cases}$$

We say a state  $(q, q', i, n) \in F'$  is accepting if  $q = q'$ ,  $n = \top$ . The automaton accepts a word  $x$  if and only if it has a run  $(q_0, q'_0, i_0, n_0) (\sigma_0, o_1, s_1) (q_1, q'_1, i_1, n_1) (\sigma_1, o_2, s_2) \dots$  on  $x$  such that for infinitely many  $j$  we have that  $(q_j, q'_j, i_j, n_j)$  is accepting.

**Proposition 3.** *The non-deterministic Büchi automaton  $\mathcal{N}'$  accepts a word  $\alpha = \mathcal{B}(\psi)$ , where  $\psi \in \text{Plays}(G)$ , if and only if  $\psi \notin \text{UFix}(l_{\max})$ .*

We recall that determinizing  $\mathcal{N}'$  and complementing it yields an exponentially bigger deterministic automaton. Its composition with  $\mathcal{B}$ , itself exponentially bigger, accepts the desired set of plays and is still singly exponential in the size of the original arena and  $W$ . Once more, the desired result follows from the algorithm presented in [16].

**Theorem 6.** *Given WGA  $G$ , determining if Eve has a winning strategy for the  $\text{UFix}(l_{\max})$  objective can be decided in time exponential in  $W$  and the size of  $G$ .*

**Corollary 2.** *Given WGA  $G$  with unary encoded weights, deciding if Eve has a winning strategy for the  $\text{UFix}(l_{\max})$  objective is EXP-complete.*

## References

- [1] S. Almagor, U. Boker, and O. Kupferman. What's decidable about weighted automata? In *Automated Technology for Verification and Analysis*, pages 482–491. Springer, 2011.
- [2] D. Berwanger, K. Chatterjee, M. D. Wulf, L. Doyen, and T. A. Henzinger. Alpaqa: A tool for solving parity games with imperfect information. In *TACAS 2009*, volume 5505 of *LNCS*. Springer, 2009.
- [3] H. Björklund, S. Sandberg, and S. Vorobyov. Memoryless determinacy of parity and mean payoff games: a simple proof. *TCS*, 310(1):365–378, 2004.
- [4] A. Bohy, V. Bruyère, E. Filiot, N. Jin, and J.-F. Raskin. Acacia+, a tool for LTL synthesis. In *CAV*, pages 652–657. Springer, 2012.
- [5] L. Brim, J. Chaloupka, L. Doyen, R. Gentilini, and J.-F. Raskin. Faster algorithms for mean-payoff games. *Formal methods in system design*, 38(2):97–118, 2011.
- [6] A. Chakrabarti, L. de Alfaro, T. A. Henzinger, and M. Stoelinga. Resource interfaces. In *EMSOFT 2003*, volume 2855 of *LNCS*. Springer, 2003.
- [7] K. Chatterjee and L. Doyen. The complexity of partial-observation parity games. In *LPAR'10*, pages 1–14. Springer, 2010.
- [8] K. Chatterjee, L. Doyen, T. A. Henzinger, and J.-F. Raskin. Algorithms for omega-regular games with imperfect information. In *CSL*, pages 287–302, 2006.

- [9] K. Chatterjee, L. Doyen, M. Randour, and J.-F. Raskin. Looking at mean-payoff and total-payoff through windows. In *ATVA*, pages 118–132. Springer, 2013.
- [10] A. Degorre, L. Doyen, R. Gentilini, J.-F. Raskin, and S. Toruńczyk. Energy and mean-payoff games with imperfect information. In *CSL*, pages 260–274, 2010.
- [11] L. Doyen and J.-F. Raskin. Antichains for the automata-based approach to model-checking. *Logical Methods in Computer Science*, 5(1), 2009.
- [12] L. Doyen and J.-F. Raskin. Antichain algorithms for finite automata. In *TACAS 2010*, volume 6015 of *LNCS*, pages 2–22. Springer, 2010.
- [13] A. Ehrenfeucht and J. Mycielski. Positional strategies for mean payoff games. *International Journal of Game Theory*, 8:109–113, 1979.
- [14] E. A. Emerson and C. S. Jutla. Tree automata, mu-calculus and determinacy. In *Foundations of Computer Science, 1991. Proceedings., 32nd Annual Symposium on*, pages 368–377. IEEE, 1991.
- [15] E. Grädel. Positional determinacy of infinite games. In *STACS*, pages 4–18. Springer, 2004.
- [16] M. Jurdzinski. Small progress measures for solving parity games. In *STACS*, volume 1770 of *LNCS*, pages 290–301. Springer, 2000.
- [17] N. Piterman. From nondeterministic Büchi and Streett automata to deterministic parity automata. *Logical Methods in Computer Science*, 3(3), 2007.
- [18] J. H. Reif. The complexity of two-player games of incomplete information. *Journal of computer and system sciences*, 29(2):274–301, 1984.
- [19] S. Safra. On the complexity of  $\omega$ -automata. In *Symposium on Foundations of Computer Science*, pages 319–327. IEEE, 1988.
- [20] S. Safra. Exponential determinization for omega-automata with strong-fairness acceptance condition (extended abstract). In *STOC*, pages 275–282, 1992.
- [21] W. Thomas. On the synthesis of strategies in infinite games. In *STACS*, pages 1–13. Springer, 1995.
- [22] U. Zwick and M. Paterson. The complexity of mean payoff games on graphs. *TCS*, 158(1):343–359, 1996.

## A Proof of Theorem 1

*Proof.* In [10] the authors show that if Eve is only allowed to play finite memory strategies then she wins the  $\text{MPInf}(\nu)$  game if and only if she wins the  $\text{MPSup}(\nu)$  game, for any  $\nu \in \mathbb{Q}$ . We show the claim holds for  $\text{MPInf}(\varepsilon)$ . Let  $\lambda_{\exists} = \langle M, m_0, \alpha_u, \alpha_o \rangle$  be the deterministic Moore machine representation of Eve’s finite memory winning strategy. Consider the product of the arena with Eve’s finite memory winning strategy,  $G \times M$ , constructed in the obvious manner, i.e. every path in  $G \times M$  corresponds to a concrete path consistent with her strategy. Clearly all cycles in  $G \times M$  have weight of at least  $\varepsilon$ , otherwise Adam can create a concrete path with mean-payoff value less than  $\varepsilon$  by “pumping” the cycles with value less than  $\varepsilon$ . As any path in  $G \times M$  corresponds to concrete plays consistent with Eve’s strategy, this contradicts the fact that the strategy is winning for her. By Pigeonhole Principle we have that for any path in  $G \times M$ : if a window opens at step  $i$ , then after  $i$  there is a sequence of length at most  $|M||Q| - 1$  that is not involved in any cycle. Now, since every cycle has weight  $\varepsilon > 0$ , after at most

$$\mu = \frac{W \cdot |M||Q|}{\varepsilon} \cdot |M||Q|$$

steps the window will have closed. It follows that for all  $\psi \in \text{Plays}(G)$  consistent with her strategy:

$$\forall \pi \in \gamma(\psi), \forall i \geq 0 : \pi \in \text{GW}(i, \mu)$$

which concludes our argument. □ □

## B Proof of Theorem 2

*Proof.* We provide a reduction from the universality of weighted finite automata which is undecidable [1]. A *weighted finite automaton* is a tuple  $\mathcal{N} = \langle Q, \Sigma, q_I, \Delta, w, F \rangle$  where  $F \subseteq Q$  is a set of final accepting states. A (*accepting*) *run* of the automaton on a word  $x = \sigma_0 \sigma_1 \dots \sigma_n \in \Sigma^*$  is a sequence  $r = q_0 q_1 \dots q_n \in Q^+$  such that  $q_n \in F$  and  $(q_i, \sigma_i, q_{i+1}) \in \Delta$  for all  $0 \leq i < n$ . The cost of the run  $r$  is  $w(r) = \sum_{i=0}^{n-1} w(q_i, \sigma_i, q_{i+1})$ . If the automaton is non-deterministic, it may have several runs on  $x$ . In that case, the *cost of  $x$  in  $\mathcal{N}$*  (denoted by  $\mathcal{N}(x)$ ) is defined as the minimum of the costs of all its accepting runs on  $x$ .

The *universality* problem for weighted automata is to decide whether, for a given automaton  $\mathcal{N}$ , the following holds:

$$\forall x \in \Sigma^* : \mathcal{N}(x) < 0.$$

We construct a blind WGA,  $G_{\mathcal{N}}$ , so that:

- if  $\mathcal{N}$  is universal, then Eve has a winning strategy for the objective  $\text{UDirBnd}$ ,
- if  $\mathcal{N}$  is not universal, then Adam has a winning strategy for the complement of the objective  $\text{Bnd}$ .

As shown in Fig. 1,  $\text{UDirBnd} \subseteq \text{Bnd}$  and all the other BWMP objectives lie in between those two. So, our reduction establishes the undecidability of all BWMP objectives at once.

Our reduction follows the gadgets given in Fig. 4-5-6. When the game starts, Adam chooses to play from one of the three gadgets. As the game is blind for Eve, she does not know what is the choice of Adam and so she must be prepared for all possibilities. Note also that as Eve is blind, her strategy can be formalized by an infinite word  $w \in \Sigma \cup \{\#\}^\omega$ . Let us show first that the two first gadgets force Eve to play a word  $w$  such that:

( $C_1$ ) there are infinitely many  $\#$  in  $w$ , and

( $C_2$ ) there exists a bound  $b \in \mathbb{N}$  such that the distance between two consecutive  $\#$  in  $w$  is bounded by  $b$ .

Assume that Eve plays a word  $w = \#w_1\#w_2\#w_3\#\dots\#w_n\#\dots$  that respects conditions  $C_1$  and  $C_2$ , with each  $w_i \in \Sigma^*$ . First, if Adam decides to play in the first gadget (Fig. 4), then either Adam stays in state  $q_1$  forever, and he does not open any window, or he decides at some point to go from  $q_1$  to  $q_2$ , whereupon he does open a window. However, after at most  $b$  steps Adam has to leave  $q_2$  for  $q_3$  at the next occurrence of the  $\#$  symbol, the bound  $b$  is guaranteed by  $C_2$ . After at most  $b$  additional steps, the open window will be closed as the self loop on  $q_3$  is labeled with the weight  $+1$ . So in this case, Eve wins the objective  $\text{UDirBnd}$ . Second, if Adam decides to play in the second gadget (Fig. 5), then he can go from  $q_4$  to  $q_5$  on the  $\#$  symbols. The windows that open on those transitions will all close within  $b$  steps according to condition  $C_2$  and the game moves back to  $q_4$ . So again, Eve wins for the objective  $\text{UDirBnd}$ .

Now assume that Eve plays a word  $w$  that violates either condition  $C_1$  or condition  $C_2$ . First, if  $w$  violates  $C_1$ , then Adam chooses the first gadget (Fig. 4), and just after Eve has played her last  $\#$ , Adam moves from  $q_1$  to  $q_2$ . As there will be no  $\#$  anymore, Adam can loop on  $q_2$  and the window that he has opened will never close. Hence, Adam wins for the complement of the objective  $\text{Bnd}$ . Second, if  $w$  violates  $C_2$  then there exists an infinite sequence of indices  $i_1 < i_2 < \dots < i_n < \dots$  such that  $|w_{i_1}| < |w_{i_2}| < \dots < |w_{i_n}| < \dots$ . Then Adam can read this sequence of sub-words using runs of the form  $q_4(q_5)^*q_4$ . Each such run will open a window that closes at the end of the sub-word. But as the sequence of lengths of the sub-words is strictly increasing and infinite, Adam wins for the complement of the objective  $\text{Bnd}$ .

Now, we will assume that Eve plays a word  $w = \#w_1\#w_2\#w_3\#\dots\#w_n\#\dots$  that respects conditions  $C_1$  and  $C_2$ , and we consider what happens when Adam plays in the third gadget (Fig. 6). As we will see, in this gadget, the best strategy for Adam is to simulate the accepting runs of  $\mathcal{N}$ .

Assume first that automaton  $\mathcal{N}$  is *non-universal*. Then by definition, there exists a finite word  $w_1 \in \Sigma^*$  such that all accepting runs of  $\mathcal{N}$  on  $w_1$  have a non-negative value, i.e.  $\mathcal{N}(w_1) \geq 0$ . In that case,  $w = (\#w_1)^\omega$  is a finite memory winning strategy for Eve for the objective  $\text{Bnd}$ . Indeed, if Adam simulates accepting runs on  $w$  then the mean-payoff of the outcome is at least  $\frac{0.5}{b} > 0$  as each new  $\#$  brings  $+\frac{1}{2}$  and we know that  $\mathcal{N}(w_1) \geq 0$ . So Eve wins for the objective  $\text{UDirBnd}$  by Theorem 1, as Eve obtains a strictly positive mean-payoff bounded away from zero with a finite memory strategy. Now, if Adam chooses not to follow accepting runs then the game ends up in state  $\perp$  from which the mean-payoff is equal to 1, so we can conclude using similar arguments that Eve wins for the objective  $\text{UDirBnd}$ . Thus in all these cases, Eve wins for the objective  $\text{UDirBnd}$ .

Finally, assume that automaton  $\mathcal{N}$  is universal and let us show then that Adam has a winning strategy for the complement of the  $\text{Bnd}$  objective. Indeed, if Eve plays a word  $w = \#w_1\#w_2\#w_3\#\dots\#w_n\#\dots$  that respects conditions  $C_1$  and  $C_2$ , then we know that  $\mathcal{N}(w_i) < 0$  for each  $i \leq 0$ . On such word, Adam can follow accepting runs in the gadget of Fig. 6. As the length between two consecutive  $\#$  is at most  $b$ , we know that the mean-payoff of the run constructed by Adam is less than or equal to  $\frac{-0.5}{b}$ . It follows that Adam wins the complement of the  $\text{Bnd}$  objective as claimed, as  $\text{Bnd}$  objective implies the mean-payoff objectives (as shown in Fig. 1 ). □ □

## C Missing proofs from Section 4

**Lemma 2.** *The number of elements in  $\mathcal{F}$  is not greater than  $2^{|Q| \cdot l_{\max} \cdot \log(W \cdot l_{\max})}$ .*

*Proof.*

$$\begin{aligned} |\mathcal{F}| &\leq (W \cdot l_{\max})^{|Q| \cdot l_{\max}} \\ &= \left(2^{\log(W \cdot l_{\max})}\right)^{|Q| \cdot l_{\max}} \\ &= 2^{|Q| \cdot l_{\max} \cdot \log(W \cdot l_{\max})}. \end{aligned}$$

□

□

**Lemma 3.** *Let  $\rho = o_0\sigma_0\dots o_n$  be an abstract path,  $\varphi \in \mathcal{F}$  such that  $\text{supp}(\varphi) \subseteq o_0$  and  $\text{supp}^{-1}(\rho, \varphi) = f_0\sigma_0\dots f_n \in (\mathcal{F} \cdot \Sigma)^*$ . A state  $q \in Q$  is reachable from some state  $q_0 \in \text{supp}(\varphi)$  through a concrete path  $q_0\sigma_0\dots q_n \in \gamma(\rho)$  if and only if  $q \in \text{supp}(f_n)$ .*

*Proof.* ( $\Rightarrow$ ) We proceed by induction. We will show that for all  $0 \leq j \leq n$ , for all  $q_j \in \text{supp}(f_j)$  there is a concrete path  $q_0\sigma_0 \dots q_j$  such that  $q_k \in o_k$  for all  $1 \leq k \leq j$  and  $q_0 \in \text{supp}(\varphi)$ . Note that for  $j = 0$  the claim trivially holds. Assume the claim holds for  $j$ . From the definition of  $\sigma$ -successor and  $\text{supp}^{-1}$  we have that  $\text{supp}(f_{j+1}) = \text{post}_{\sigma_j}(\text{supp}(f_j)) \subseteq o_{j+1}$ . This means that for all  $q_{j+1} \in \text{supp}(f_{j+1})$  there must be some  $q_j \in \text{supp}(f_j)$  such that  $(q_j, \sigma_j, q_{j+1}) \in \Delta$ . Hence any  $q_{j+1}$  is reachable from some  $q_j$  via  $\sigma_j$  which, by inductive hypothesis, is in turn reachable from some  $q_0 \in \text{supp}(\varphi)$  via a concrete path of the desired form.

( $\Leftarrow$ ) We now show – once more by induction on  $j$  – that for all  $0 \leq j \leq n$ , if there is a concrete path  $q_0\sigma_0 \dots q_j$  such that  $q_0 \in \text{supp}(\varphi)$  and  $q_k \in o_k$  for all  $1 \leq k \leq j$ , then  $q_j \in \text{supp}(f_j)$ . The claim holds for  $j = 0$ . Assume that it holds for some  $j$ . From the assumptions we have that  $(q_j, \sigma_j, q_{j+1}) \in \Delta$  and  $q_{j+1} \in o_{j+1}$ . Further, we know that  $q_j \in \text{supp}(f_j)$  by inductive hypothesis. Hence,  $q_{j+1} \in \text{post}_{\sigma_j}(\text{supp}(f_j)) \subseteq o_{j+1}$  which means that  $q_{j+1} \in \text{supp}(f_{j+1})$ .  $\square$   $\square$

**Lemma 4.** *Let  $\rho = o_0\sigma_0 \dots o_n$  be an abstract path,  $\varphi \in \mathcal{F}$  such that  $\text{supp}(\varphi) \subseteq o_0$  and  $\text{supp}^{-1}(\rho, \varphi) = f_0\sigma_0 \dots f_n \in (\mathcal{F} \cdot \Sigma)^*$ . Given state  $p \in \text{supp}(f_n)$  and  $1 \leq l \leq l_{\max}$  such that  $l \leq n$ , then there is a window of length  $l$  open at  $p$  if and only if  $f_n(p)_l < 0$ .*

Instead of directly providing a proof of Lemma 4, we prove a more general result below. Consider the three conditions stated in Lemma 10. We shall prove that  $\text{C1} \Rightarrow \text{C2} \Rightarrow \text{C3} \Rightarrow \text{C1}$ . Since C1 corresponds to having a window of length  $l$  open at  $p$  from  $\text{supp}(\varphi)$ , the desired result follows from transitivity.

**Lemma 10.** *Let  $\rho = o_0\sigma_0 \dots o_n$  be an abstract path,  $\varphi \in \mathcal{F}$  such that  $\text{supp}(\varphi) \subseteq o_0$  and  $\text{supp}^{-1}(\rho, \varphi) = f_0\sigma_0 \dots f_n \in (\mathcal{F} \cdot \Sigma)^*$ . Given state  $p \in \text{supp}(f_n)$  and  $1 \leq l \leq l_{\max}$  such that  $l \leq n$ , let  $\lambda = n - l$ . The following three statements are equivalent.*

C1. *There is a concrete path  $q_0\sigma_0 \dots q_n \in \gamma(\rho)$  with  $q_n = p$  and  $q_0 \in \text{supp}(\varphi)$  and*

$$\sum_{j=n-l}^m w(q_j, \sigma_j, q_{j+1}) < 0$$

*for all  $n - l \leq m < n$ .*

C2.  *$f_n(p)_l < 0$ .*

C3. *There is a concrete path  $q_0\sigma_0 \dots q_n \in \gamma(\rho)$  with  $q_n = p$  and  $q_0 \in \text{supp}(\varphi)$  such that*

(a)  *$f_j(q_j)_{j-\lambda} < 0$  for all  $\lambda < j \leq n$ , and*

(b)  *$f_k(q_k)_{j-\lambda} + w(q_k, \sigma_k, q_{k+1}) = f_{k+1}(q_{k+1})_{k-\lambda+1}$  for all  $\lambda < k < n$ .*

*Proof.* (C3  $\Rightarrow$  C1) We will apply induction on  $m$ . From the definition of  $\sigma$ -successor we have that  $f_{\lambda+1}(q_{\lambda+1})_1 = \min\{0, w(q_\lambda, \sigma_\lambda, q_{\lambda+1})\}$ . From assumption (a) we know that  $f_{\lambda+1}(q_{\lambda+1})_1 < 0$ . Thus, the claim holds for  $m = \lambda$ . Assume it holds for  $m$ . To conclude the proof, we now show that the claim holds for  $m + 1$  as well.

$$\begin{aligned} \sum_{j=n-l}^{m+1} w(q_j, \sigma_j, q_{j+1}) &= f_m(q_m)_{m-\lambda} + w(q_m, \sigma_m, q_{m+1}) && \text{ind. hyp.} \\ &= f_m(q_m)_{m-\lambda+1} && \text{from (b)} \\ &< 0 && \text{from (a).} \end{aligned}$$

(C1  $\Rightarrow$  C2) We show, by induction on  $m$ , that for all  $\lambda \leq m < n$

$$f_{m+1}(q_{m+1})_{m-\lambda+1} \leq \sum_{j=\lambda}^m w(q_j, \sigma_j, q_{j+1}).$$



The desired result follows. As the base case, consider  $m = \lambda$  and note that by definition of  $\sigma$ -successor we have that

$$\begin{aligned} f_{\lambda+1}(q_{\lambda+1})_1 &= \min(\{0\} \cup \{w(p, \sigma_\lambda, q_{\lambda+1}) \mid p \in \text{supp}(f_\lambda) \wedge (p, \sigma_\lambda, q_{\lambda+1}) \in \Delta\}) \\ &\leq w(q_\lambda, \sigma_\lambda, q_{\lambda+1}). \end{aligned}$$

Thus the claim holds. Assume that the claim is true for  $m$ . From the definition of  $\sigma$ -successor we have that

$$f_{m+2}(q_{m+2})_{\lambda-m+2} \leq f_{m+1}(q_{m+1})_{m-\lambda+1} + w(q_{m+1}, \sigma_{m+1}, q_{m+2}).$$

From the inductive hypothesis we get have that the right hand side of the inequality is equivalent to

$$\sum_{j=\lambda}^{m+1} w(q_j, \sigma_j, q_{j+1}).$$

Thus the claim holds for  $m + 1$  as well.

(C2  $\Rightarrow$  C3) We inductively construct a concrete path  $q_0\sigma_0 \dots q_n\gamma(\rho)$  with  $q_n = p$  and  $q_0 \in \text{supp}(\varphi)$  such that

- (1)  $f_{n-k}(q_{n-k})_{l-k} < 0$  for all  $0 \leq k < l$ , and
- (2)  $f_{n-k}(q_{n-k})_{l-k} = w(q_{n-k+1}, \sigma_{n-k+1}, q_{n-k}) + f_{n-k+1}(q_{n-k+1})_{l-k+1}$  for all  $1 \leq k < l$ .

As these conditions are equivalent to (a)-(b) from C3, the result follows. Note that for  $k = 0$  we have that (1) holds trivially since  $p \in \text{supp}(f_n)$  and  $f_n(p)_l < 0$  by hypothesis. If  $f_{n-k}(q_{n-k})_{l-k} < 0$  then, by definition of  $\sigma$ -successor, it follows that there is some  $q' \in \text{supp}(f_{n-k+1}) \subseteq o_{n-k+1}$  such that  $f_{n-k+1}(q')_{l-k+1} < 0$  and  $f_{n-k}(q_{n-k})_{l-k} = w(q_{n-k+1}, \sigma_{n-k+1}, q_{n-k}) + f_{n-k+1}(q_{n-k+1})_{l-k+1}$ . In other words,  $q'$  is the source of the minimal  $\sigma_{n-k+1}$ -transition of a state from  $\text{supp}(f_{n-k+1})$  to  $q_{n-k}$ . Let  $q_{n-k+1} = q'$ . Continue in this fashion defining every  $q_i$  up to  $q_{n-l}$ . Now, from Lemma 3, we have that  $q_{n-l}$  is reachable from some state in  $\text{supp}(\varphi)$  via a concrete path of the desired form. Any such path is a valid prefix for the sequence  $q_{n-l}\sigma_{n-l} \dots q_n$  we constructed above.  $\square$   $\square$

**Lemma 8.** *Given upward-closed set  $S \in \mathcal{P}(\mathcal{F})$  and  $f, g \in \mathcal{F} \setminus \mathcal{U}$ , if  $f \in \text{UPre}(S)$  and  $f \preceq g$ , then  $g \in \text{UPre}(S)$ .*

*Proof.* We have that for all  $\sigma$ , there is  $h_\sigma \in S$  such that  $(f, \sigma, h_\sigma) \in \Delta'$ . By construction of  $\Delta'$  we also know that there is  $i_\sigma$  such that  $(g, \sigma, i_\sigma) \in \Delta'$ , and furthermore, since  $\text{supp}(f) \subseteq \text{supp}(g)$ , we get that

$$\begin{aligned} \text{supp}(h_\sigma) &= \text{post}_\sigma(\text{supp}(f)) \cap o \\ &\subseteq \text{post}_\sigma(\text{supp}(g)) \cap o \\ &= \text{supp}(i_\sigma) \end{aligned}$$

for some  $o \in \text{Obs}$ . Note that:

- (1) since  $f, g \notin \mathcal{U}$ , then  $f(p)_{l_{\max}} = g(p)_{l_{\max}} = 0$  for all  $p \in \text{supp}(f)$ ; and
- (2)  $i_\sigma(q)_1 = h_\sigma(q)_1$  for all  $q \in \text{supp}(h_\sigma)$ .

From (1) and since  $f \preceq g$ , there is a function  $\alpha : \{1, \dots, l_{\max}\} \rightarrow \{1, \dots, l_{\max} - 1\}$  such that for all  $1 \leq x < l_{\max}$  we have that  $\alpha(x) \geq x$  and  $f(p)_x \geq g(p)_{\alpha(x)}$  holds for all  $p \in \text{supp}(f)$ . Observe that for all  $q \in \text{supp}(h_\sigma)$  and any  $2 \leq x \leq l_{\max}$ , we have that

$$\begin{aligned} h_\sigma(q)_x &= \min_{p \in \text{supp}(f)} (\{0\} \cup \{f(p)_{x-1} + w(p, \sigma, q) \mid f(p)_{x-1} < 0\}) \\ &\geq \min_{p \in \text{supp}(f)} (\{0\} \cup \{g(p)_{\alpha(x-1)} + w(p, \sigma, q) \mid g(p)_{\alpha(x-1)} < 0\}) \\ &\geq i_\sigma(q)_{\alpha(x-1)+1}. \end{aligned}$$

It follows that  $h_\sigma \preceq i_\sigma$  and that, since  $S$  is upward-closed,  $i_\sigma \in S$ . Thus, we have shown that for all  $\sigma$ , there is  $i_\sigma \in S$  such that  $(g, \sigma, i_\sigma) \in \Delta'$ , which implies that  $g \in \text{UPre}(S)$ .  $\square$   $\square$

## D Proof of Theorem 4

**Lemma 11.** *Given upward-closed set  $S \in \mathcal{P}(\mathcal{F})$ ,  $\lfloor \text{UPre} \rfloor(\lfloor S \rfloor) = \lfloor \text{UPre}(S) \setminus \mathcal{U} \rfloor$ .*

*Proof.* We first show that if  $f \in \lfloor \text{UPre} \rfloor(\lfloor S \rfloor)$  then  $f \in \text{UPre}(S) \setminus \mathcal{U}$ . We have that  $f \notin \mathcal{U}$  and  $\forall \sigma \in \Sigma, \exists q' \in \lfloor S \rfloor, \exists r'_\sigma \in \mathcal{F} : (f, \sigma, r'_\sigma) \in \Delta'$  and  $q' \preceq r'_\sigma$ . Since  $S$  is upward-closed and  $q' \preceq r'_\sigma$ , we know that  $r'_\sigma \in S$ . Hence, we get that  $\forall \sigma \in \Sigma, \exists r'_\sigma \in S : (f, \sigma, r'_\sigma) \in \Delta'$ , which implies that  $f \in \text{UPre}(S) \setminus \mathcal{U}$ .

Next, we show that if  $f \in \lfloor \text{UPre}(S) \setminus \mathcal{U} \rfloor$  then  $f \in \{p' \in \mathcal{F} \setminus \mathcal{U} \mid \forall \sigma \in \Sigma, \exists q' \in a, \exists r' \in \mathcal{F} : (p', \sigma, r') \in \Delta' \text{ and } q' \preceq r'\}$ . We know that  $f \notin \mathcal{U}$  and  $\forall \sigma \in \Sigma, \exists r' \in Q : (f, \sigma, r') \in \Delta'$ . By definition of  $\lfloor S \rfloor$ , we know there is  $q_{r'} \in \lfloor S \rfloor$  such that  $q_{r'} \preceq r'$ . Thus, we get that  $\forall \sigma \in \Sigma, \exists r'_\sigma \in S, \exists q_{r'_\sigma} \in \lfloor S \rfloor : (f, \sigma, r'_\sigma) \in \Delta' \text{ and } q_{r'_\sigma} \preceq r'_\sigma$ .

Finally, we note that if  $f \in \lfloor \text{UPre} \rfloor(\lfloor S \rfloor)$  then not only is it true that  $f \in \text{UPre}(S) \setminus \mathcal{U}$ , but furthermore  $f \in \lfloor \text{UPre}(S) \setminus \mathcal{U} \rfloor$ . Indeed, if this were not the case, then there would be  $g \in \lfloor \text{UPre}(S) \setminus \mathcal{U} \rfloor$  such that  $g \preceq f$  and  $f \neq g$ . Then, by the argument explained in the previous paragraph, this would contradict minimality of  $f$  in  $\lfloor \text{UPre} \rfloor(\lfloor S \rfloor)$ . Similarly, if  $f \in \lfloor \text{UPre}(S) \setminus \mathcal{U} \rfloor$  then  $f \in \lfloor \text{UPre} \rfloor(\lfloor S \rfloor)$ , as otherwise, by the argument from the first paragraph of the proof, minimality in the first set would be contradicted. Thus, the claim holds.  $\square$   $\square$

**Theorem 4.** *Given WGA  $G$ , Eve wins the  $\text{DirFix}(l_{\max})$  objective if and only if  $\{q'_I\} \not\sqsubseteq \mu X.([\mathcal{U}] \sqcup \lfloor \text{UPre} \rfloor(X))$ .*

*Proof.* We note that for any upward-closed set  $S \subseteq \mathcal{F}$  such that  $\mathcal{U} \subseteq S$  we have, from Lemma 8 that  $\mathcal{U} \cup \text{UPre}(S)$  is again upward-closed and a superset of  $\mathcal{U}$ . In fact, it holds that

$$\begin{aligned} \mathcal{U} \cup \text{UPre}(S) &= ([\mathcal{U}] \sqcup \lfloor \text{UPre}(S) \rfloor)^\uparrow \\ &= ([\mathcal{U}] \sqcup \lfloor \text{UPre} \rfloor(\lfloor S \rfloor))^\uparrow \end{aligned} \quad \text{from Lemma 11.}$$

It is easy to show by induction that  $\mu X.(\mathcal{U} \cup \text{UPre}(X)) = (\mu X.([\mathcal{U}] \sqcup \lfloor \text{UPre} \rfloor(\lfloor X \rfloor)))^\uparrow$ . Thus,  $\{q'_I\} \not\sqsubseteq \mu X.([\mathcal{U}] \sqcup \lfloor \text{UPre} \rfloor(\lfloor S \rfloor))$  if and only if  $q_I \notin \mu X.(\mathcal{U} \cup \text{UPre}(S))$ . From Proposition 1 and Lemmas 5 and 6 we know this is the case if and only if Eve has a winning strategy in the safety game in  $G'$  if and only if she wins the  $\text{DirFix}(l_{\max})$  objective in  $G$ .  $\square$   $\square$

## E Proof of Proposition 2

*Proof.* ( $\Rightarrow$ ) Assume  $\mathcal{N}$  accepts  $\psi$ . Let  $r = (q_0, i_0, n_0)(\sigma_0, o_1)(q_1, i_1, n_1)(\sigma_1, o_2) \dots$  be one of the accepting runs of the automaton on  $\psi$ . By construction of  $\mathcal{N}$  we have that  $q_0 \sigma_0 q_1 \sigma_1 \dots \in \gamma(\psi)$ . Let  $\pi_r$  denote this concrete play and  $J = \{j_0, j_1, j_2, \dots\}$  be an infinite set of indices such that  $j_k < j_{k+1}$  and  $(q_{j_k}, i_{j_k}, n_{j_k})$  is accepting for all  $k \geq 0$ . Such a sequence is guaranteed to exist since  $r$  is accepting. One can easily verify by induction on the definition of  $\Delta''$  that for all  $k \geq 0$  it holds that  $\pi_r \notin \text{GW}(i_{j_k} - l_{\max}, l_{\max})$ . It follows that  $\forall m \geq 0, \exists n \geq m : \pi_r \notin \text{GW}(n, l_{\max})$ , which concludes our argument.

( $\Leftarrow$ ) Assume that  $\psi = o_0 \sigma_0 o_1 \sigma_1 \dots \notin \text{Fix}(l_{\max})$ . Let  $\pi = q_0 \sigma_0 q_1 \sigma_1 \in \gamma(\psi)$  be the concrete play such that for infinitely many  $i$  it is the case that  $\pi \notin \text{GW}(i, l_{\max})$ . We describe the infinite run of  $\mathcal{N}$  on  $\psi$  that accepts. Let  $J = \{j_0, j_1, j_2, \dots\}$  be an infinite set of indices such that  $j_k + l_{\max} < j_{k+1}$  and  $\pi \notin \text{GW}(j_k, l_{\max})$  for  $k \geq 0$ . The sequence is guaranteed to exist because of our choice of  $\pi$ . Observe that this implies there is a run  $r = (q_0, i_0, n_0) (\sigma_0, o_1) (q_1, i_1, n_1) (\sigma_1, o_2) \dots$  of the automaton where for all  $k \geq 0$  we have that  $n_{j_k+1} = w(q_{j_k}, \sigma_{j_k}, q_{j_k+1})$  and for all  $1 < l < l_{\max}$  then

$$n_{j_k+l} = n_{j_k+l-1} + w(q_{j_k+l}, \sigma_{j_k+l}, q_{j_k+l+1}).$$

Furthermore, in this run it holds that for all  $k \geq 0$  we have  $i_{j_k+l_{\max}} = l_{\max}$ . Hence, said run is such that for all  $k \geq 0$  the state  $(q_{j_k+l_{\max}}, i_{j_k+l_{\max}}, n_{j_k+l_{\max}})$  is accepting. We conclude that the automaton accepts  $\psi$ .  $\square$   $\square$

## F Proof of Proposition 3

*Proof.* ( $\Rightarrow$ ) Assume  $\mathcal{N}'$  accepts  $\alpha$ . Let  $r = (q_0, q'_0 i_0, n_0) (\sigma_0, o_1, s_1) (q_1, q'_1 i_1, n_1) (\sigma_1, o_2, s_2) \dots$  be one of the accepting runs of the automaton on  $\alpha$ . By construction of  $\mathcal{N}'$  we have that  $q'_0 \sigma_0 q'_1 \sigma_1 \dots \in \gamma(\psi)$ . Let  $\pi_r$  denote this concrete play,  $J = \{j_0, j_1, \dots\}$  and  $K = \{k_0, k_1, \dots\}$  be two infinite sets of indices such that  $j_l < j_{l+1}$  and  $j_l + l_{\max} \leq k_l$ , for all  $l \geq 0$ , and for which we know that

- $(q_{k_l}, q'_{k_l}, i_{k_l}, n_{k_l})$  is accepting for all  $l \geq 0$ , and
- $n_{j_l + l_{\max}} < 0 \wedge i_{j_l + l_{\max}} = l_{\max}$ .

Such sequences are guaranteed to exist since  $r$  is accepting. Assuming the correctness of  $\mathcal{B}$ , one can easily verify by induction on the definition of  $\Delta'''$  that for all  $l \geq 0$  we have that  $\pi_r$ , at  $k_l$  merges with a path having a violation at  $j_l$ . It follows that  $\pi_r$  merges with infinitely many violating paths. From Lemma 9 we get that  $\psi \notin \text{UFix}(l_{\max})$ .

( $\Leftarrow$ ) Assume that  $\psi = o_0 \sigma_0 o_1 \sigma_1 \dots \notin \text{UFix}(l_{\max})$ . Let  $\pi = q_0 \sigma_0 q_1 \sigma_1 \in \gamma(\psi)$  be the concrete play that merges with infinitely many violating paths (see Lemma 9). We describe the infinite run of  $\mathcal{N}$  on  $\alpha = \mathcal{B}(\psi)$  that accepts. Let  $J = \{j_0, j_1, \dots\}$  and  $K = \{k_0, k_1, \dots\}$  be two infinite sets of indices such that  $j_l < j_{l+1}$  and  $j_l + l_{\max} + 1 < k_l$ , for all  $l \geq 0$ , and for which we know that there is some  $\chi_l \in \gamma(\psi[..k_l])$  such that  $\pi[k_l] = \chi_l[k_l]$  and for all  $j_k < m \leq j_k + l_{\max} + 1$  we have  $w(\chi[j_k..m]) < 0$ . The sequences are guaranteed to exist because of our choice of  $\pi$ . Observe that this implies there is a run  $r = (q_0, q'_0 i_0, n_0) (\sigma_0, o_1, s_1) (q_1, q'_1 i_1, n_1) (\sigma_1, o_2, s_2) \dots$  of the automaton where for all  $l \geq 0$  we have that  $n_{j_l+1} = w(q_{j_l}, \sigma_{j_l}, q_{j_l+1})$  and for all  $1 < b < l_{\max}$  then

$$n_{j_l+b} = n_{j_l+b-1} + w(q_{j_l+b}, \sigma_{j_l+b}, q_{j_l+b+1}).$$

Furthermore, we have that  $n_{j_l+b} = \top$  for all  $l_{\max} \leq b \leq k_l$  and  $q_{k_l} = q'_{k_l}$ . Hence, said run is such that for all  $l \geq 0$  the state  $(q_{k_l}, q'_{k_l}, i_{k_l}, n_{k_l})$  is accepting. We conclude that the automaton accepts  $\psi$ .  $\square$   $\square$