

Doomsday Equilibria for Omega-Regular Games

Krishnendu Chatterjee¹, Laurent Doyen², Emmanuel Filiot³, and
Jean-François Raskin⁴

¹IST Austria, ²LSV-ENS de Cachan, ³Université Paris-Est Créteil, ⁴CS-Université
Libre de Bruxelles

Abstract. Two-player games on graphs provide the theoretical framework for many important problems such as reactive synthesis. While the traditional study of two-player zero-sum games has been extended to multi-player games with several notions of equilibria, they are decidable only for perfect-information games, whereas several applications require imperfect-information games.

In this paper we propose a new notion of equilibria, called doomsday equilibria, which is a strategy profile such that all players satisfy their own objective, and if any coalition of players deviates and violates even one of the players objective, then the objective of every player is violated. We present algorithms and complexity results for deciding the existence of doomsday equilibria for various classes of ω -regular objectives, both for imperfect-information games, as well as for perfect-information games. We provide optimal complexity bounds for imperfect-information games, and in most cases for perfect-information games.

1 Introduction

Two-player games on finite-state graphs with ω -regular objectives provide the framework to study many important problems in computer science [Sha53,Rab69,EJ91]. One key application area is synthesis of reactive systems [BL69,RW87,PR89a]. Traditionally, the reactive synthesis problem is reduced to two-player zero-sum games, where vertices of the graph represent states of the system, edges represent transitions, one player represents a component of the system to synthesize, and the other player represents the purely adversarial coalition of all the other components. Since the coalition is adversarial, the game is zero-sum, i.e., the objectives of the two players are complementary. Two-player zero-sum games on graphs have been studied in great depth in the literature [Mar75,EJ91,GTW02].

Instead of considering all the other components as purely adversarial, a more realistic model is to consider them as individual players each with their own objective, as in protocol synthesis where the rational behavior of the agents is to first satisfy their own objective in the protocol before trying to be adversarial to the other agents. Hence, inspired by recent applications in protocol synthesis, the model of multi-player games on graphs has become an active area of research in graph games and reactive synthesis [AHK02,FKL10,UW11]. In a multi-player setting, the games are not necessarily zero-sum (i.e., objectives

are not necessarily conflicting) and the classical notion of rational behavior is formalized as Nash equilibria [Nas50]. Nash equilibria perfectly capture the notion of rational behavior in the absence of external criteria, i.e., the players are concerned only about their own payoff, and they are indifferent to the payoff of the other players. In the setting of synthesis, the more appropriate notion is the adversarial external criteria, where the players are as harmful as possible to the other players without sabotaging with their own objectives. This has inspired the study of refinements of Nash equilibria, such as secure equilibria [CHJ06] (that captures the adversarial external criteria), rational synthesis [FKL10], and led to several new logics where the non-zero-sum equilibria can be expressed [CHP10,DLM10,MMV10,WHY11,MMPV12]. The complexity of Nash equilibria [UW11], secure equilibria [CHJ06], rational synthesis [FKL10], and of the new logics has been studied recently [CHP10,DLM10,MMV10,WHY11].

Along with the theoretical study of refinements of equilibria, applications have also been developed in the synthesis of protocols. In particular, the notion of secure equilibria has been useful in the synthesis of mutual-exclusion protocol [CHJ06], and of fair-exchange protocols [CR12] (a key protocol in the area of security for exchange of digital signatures [MGK02]). One major drawback that all the notions of equilibria suffer is that the basic decision questions related to them are decidable only in the setting of perfect-information games (in a perfect-information games the players perfectly know the state and history of the game, whereas in imperfect-information games each player has only a partial view of the state space of the game), and in the setting of multi-player imperfect-information games they are undecidable [PR89a]. However, the model of partial-information games is very natural because every component of a system has private variables not accessible to other components, and recent works have demonstrated that imperfect-information games are required in synthesis of fair-exchange protocols [JMM12]. In this paper, we provide the first decidable framework that can model them.

We propose a new notion of equilibria which we call *doomsday-threatening equilibria* (for short, doomsday equilibria). A doomsday equilibrium is a strategy profile such that all players satisfy their own objective, and if any coalition of players deviates and violates even one of the players objective, then doomsday follows (every player objective is violated). Note that in contrast to other notions of equilibria, doomsday equilibria consider deviation by an arbitrary set of players, rather than individual players. Moreover, in case of two-player non-zero-sum games they coincide with the secure equilibria where objectives of both players are satisfied. We present algorithms and complexity bounds for deciding the existence of doomsday equilibria for various classes of ω -regular objectives both in perfect-information games, as well as in imperfect-information games, and in most cases with optimal complexity. Our contribution is summarized in Table 1. More specifically:

1. (*Perfect-information games*). We show that deciding the existence of doomsday equilibria in multi-player perfect-information games is (i) PTIME-complete for reachability, Büchi, and coBüchi objectives; (ii) PSPACE-complete for

objectives \ games	safety	reachability	Büchi	co-Büchi	parity
perfect information	PSPACE-C	PTIME-C	PTIME-C	PTIME-C	PSPACE NP-HARD CoNP-HARD
imperfect information	EXPTIME-C	EXPTIME-C	EXPTIME-C	EXPTIME-C	EXPTIME-C

Table 1. Summary of the results

safety objectives; and (iii) in PSPACE and both NP-hard and coNP-hard for parity objectives.

2. (*Imperfect-information games*). We show that deciding the existence of doomsday equilibria in multi-player imperfect-information games is EXPTIME-complete for reachability, safety, Büchi, coBüchi, and parity objectives.

We also show (for both perfect and imperfect information settings) that when the objectives are given by LTL formula, the DE existence problem is 2ExpTime-complete, and we present a Safraless procedure [KV05] which can be optimized using antichain data structures [FJR11a].

The area of multi-player games and various notion of equilibria is an active area of research, but notions that lead to decidability for imperfect-information games and has applications in synthesis, has largely been unexplored. Our work is a contribution in this direction.

2 Doomsday Equilibria for Perfect Information Games

In this section, we define game arena with perfect information, ω -regular objectives, and doomsday equilibria.

Game Arena An n -player game arena G with perfect information is defined as a tuple $(S, \mathcal{P}, s_{\text{init}}, \Sigma, \Delta)$ such that S is a nonempty finite set of *states*, $\mathcal{P} = \{S_1, S_2, \dots, S_n\}$ is a partition of S into n classes of states, one for each player respectively, $s_{\text{init}} \in S$ is the initial state, Σ is a finite set of actions, and $\Delta : S \times \Sigma \rightarrow S$ is the transition function.

Plays in n -player game arena G are constructed as follows. They start in the initial state s_{init} , and then an ω number of rounds are played as follows: the player that owns the current state s chooses a letter $\sigma \in \Sigma$ and the game evolves to the position $s' = \Delta(s, \sigma)$, then a new round starts from s' . So formally, a *play* in G is an infinite sequence $s_0 s_1 \dots s_n \dots$ such that (i) $s_0 = s_{\text{init}}$ and (ii) for all $i \geq 0$, there exists $\sigma \in \Sigma$ such that $s_{i+1} = \Delta(s_i, \sigma)$. The set of plays in G is denoted by $\text{Plays}(G)$, and the set of finite prefixes of plays by $\text{PrefPlays}(G)$. We denote by $\rho, \rho_1, \rho_i, \dots$ plays in G , by $\rho(0..j)$ the prefix of the play ρ up to position j and by $\rho(j)$ the position j in the play ρ . We also use π, π_1, π_2, \dots to denote prefixes of plays. Let $i \in \{1, 2, \dots, n\}$, a prefix π belongs to Player i if $\text{last}(\pi)$, the last state of π , belongs to Player i , i.e. $\text{last}(\pi) \in S_i$. We denote by $\text{PrefPlays}_i(G)$ the set of prefixes of plays in G that belongs to Player i .

Strategies and strategy profiles A *strategy* for Player i , for $i \in \{1, 2, \dots, n\}$, is a mapping $\lambda_i : \text{PrefPlays}_i(G) \rightarrow \Sigma$ from prefixes of plays to actions. A *strategy profile* $\Lambda = (\lambda_1, \lambda_2, \dots, \lambda_n)$ is a tuple of strategies such that λ_i is a strategy of Player i . The strategy of Player i in Λ is denoted by Λ_i , and the tuple of the remaining strategies $(\lambda_1, \dots, \lambda_{i-1}, \lambda_{i+1}, \dots, \lambda_n)$ by Λ_{-i} . For a strategy λ_i of Player i , we define its *outcome* as the set of plays that are consistent with λ_i : formally, $\text{outcome}_i(\lambda_i)$ is the set of $\rho \in \text{Plays}(G)$ such that for all $j \geq 0$, if $\rho(0..j) \in \text{PrefPlays}_i(G)$, then $\rho(j+1) = \Delta(\rho(j), \lambda_i(\rho(0..j)))$. Similarly, we define the *outcome of a strategy profile* $\Lambda = (\lambda_1, \lambda_2, \dots, \lambda_n)$, as the unique play $\rho \in \text{Plays}(G)$ such that for all positions j , for all $i \in \{1, 2, \dots, n\}$, if $\rho(j) \in \text{PrefPlays}_i(G)$ then $\rho(j+1) = \Delta(\rho(j), \lambda_i(\rho(0..j)))$. Finally, given a state $s \in S$ of the game, we denote by G_s the game G whose initial state is replaced by s .

Winning objectives A *winning objective* (or an *objective* for short) φ_i for Player $i \in \{1, 2, \dots, n\}$ is a set of infinite sequences of states, i.e. $\varphi_i \subseteq S^\omega$. A strategy λ_i is *winning* for Player i (against all other players) w.r.t. an objective φ_i if $\text{outcome}_i(\lambda_i) \subseteq \varphi_i$.

Given an infinite sequence of states $\rho \in S^\omega$, we denote by $\text{visit}(\rho)$ the set of states that appear at least once along ρ , i.e. $\text{visit}(\rho) = \{s \in S \mid \exists i \geq 0 \cdot \rho(i) = s\}$, and $\text{inf}(\rho)$ the set of states that appear infinitely often along ρ , i.e. $\text{inf}(\rho) = \{s \in S \mid \forall i \geq 0 \cdot \exists j \geq i \cdot \rho(j) = s\}$. We consider the following types of winning objectives:

- a *safety objective* is defined by a subset of states $T \subseteq S$ that has to be never left: $\text{safe}(T) = \{\rho \in S^\omega \mid \text{visit}(\rho) \subseteq T\}$;
- a *reachability objective* is defined by a subset of states $T \subseteq S$ that has to be reached: $\text{reach}(T) = \{\rho \in S^\omega \mid \text{visit}(\rho) \cap T \neq \emptyset\}$;
- a *Büchi objective* is defined by a subset of states $T \subseteq S$ that has to be visited infinitely often: $\text{Büchi}(T) = \{\rho \in S^\omega \mid \text{inf}(\rho) \cap T \neq \emptyset\}$;
- a *co-Büchi objective* is defined by a subset of states $T \subseteq S$ that has to be reached eventually and never be left: $\text{coBüchi}(T) = \{\rho \in S^\omega \mid \text{inf}(\rho) \subseteq T\}$;
- let $d \in \mathbb{N}$, a *parity objective with d priorities* is defined by a parity function $p : S \rightarrow \{0, 1, \dots, d\}$ as the set of plays such that the smallest priority visited infinitely often is even: $\text{parity}(p) = \{\rho \in S^\omega \mid \min\{p(s) \mid s \in \text{inf}(\rho)\} \text{ is even}\}$.

Büchi, co-Büchi and parity objectives φ are called *tail objectives* because they enjoy the following closure property: for all $\rho \in \varphi$ and all $\pi \in S^*$, $\rho \in \varphi$ iff $\pi \cdot \rho \in \varphi$.

Finally, given an objective $\varphi \subseteq S^\omega$ and a subset $P \subseteq \{1, \dots, n\}$, we write $\langle\langle P \rangle\rangle\varphi$ to denote the set of states s from which the players from P can cooperate to enforce φ when they start playing in s . Formally, $\langle\langle P \rangle\rangle\varphi$ is the set of states s such that there exists a set of strategies $\{\lambda_i \mid i \in P\}$ in G_s , one for each player in P , such that $\bigcap_{i \in P} \text{outcome}_i(\lambda_i) \subseteq \varphi$.

Doomsday Equilibria A strategy profile $\Lambda = (\lambda_1, \lambda_2, \dots, \lambda_n)$ is a *doomsday-threatening equilibrium* (doomsday equilibrium or DE for short) if:




Fig. 1. Examples

1. it is winning for all the players, i.e. $\text{outcome}(\mathcal{A}) \in \bigcap_i \varphi_i$;
2. each player is able to retaliate in case of deviation: for all $1 \leq i \leq n$, for all $\rho \in \text{outcome}_i(\lambda_i)$, if $\rho \notin \varphi_i$, then $\rho \in \bigcap_{j=1}^{j=n} \overline{\varphi_j}$ (doomsday), where $\overline{\varphi_j}$ denotes the complement of φ_j in S^ω .

In other words, when all players stick to their strategies then they all win, and if any arbitrary coalition of players deviates and makes even just one other player lose then this player can retaliate and ensures a doomsday, i.e. all players lose.

Relation with Secure Equilibria In two-player games, the doomsday equilibria coincide with the notion of secure equilibrium [CHJ06] where both players satisfy their objectives. In secure equilibria, for all $i \in \{1, 2\}$, any deviation of Player i that does not decrease her payoff does not decrease the payoff of Player $3-i$ either. In other words, if a deviation of Player i decreases (strictly) the payoff of Player $3-i$, i.e. φ_{3-i} is not satisfied, then it also decreases her own payoff, i.e. φ_i is not satisfied. A two-player secure equilibrium where both players satisfy their objectives is therefore a doomsday equilibrium.

Examples Fig. 1 gives two examples of games with safety and Büchi objectives respectively. Actions are in bijection with edges so they are not represented.

(Safety) Consider the 3-player game arena with perfect information of Fig. 1(a) and safety objectives. Unsafe states for each player are given by the respective nodes of the upper part. Assume that the initial state is one of the safe states. This example models a situation where three countries are in peace until one of the countries, say country i , decides to attack country j . This attack will then necessarily be followed by a doomsday situation: country j has a strategy to punish all other countries. The doomsday equilibrium in this example is to play safe for all players.

(Büchi) Consider the 3-player game arena with perfect information of Fig. 1(b) with Büchi objectives for each player: Player i wants to visit infinitely often one of its “happy” states. The position of the initial state does not matter. To make things more concrete, let us use this game to model a protocol where 3 players want to share in each round a piece of information made of three parts: for all $i \in \{1, 2, 3\}$, Player i knows information $i \bmod 3 + 1$ and $i \bmod 3 + 2$. Player i can send or not these informations to the other players. This is modeled by the fact that Player i can decide to visit the happy states of the other players, or move directly to $s_{(i \bmod 3)+1}$. The objective of each player is to have an infinite number of successful rounds where they get all information.

There are several doomsday equilibria. As a first one, let us consider the situation where for all $i \in \{1, 2, 3\}$, if Player i is in state s_i , first it visits the happy states, and when the play comes back in s_i , it moves to $s_{(i \bmod 3)+1}$. This defines an infinite play that visits all the states infinitely often. Whenever some player deviates from this play, the other players retaliate by always choosing to go to the next s state. Clearly, if all players follow their respective strategies all happy states are visited infinitely often. Now consider the strategy of Player i against two strategies of the other players that makes him lose. Clearly, the only way Player i loses is when the two other players eventually never take their states, but then all the players lose.

As a second one, consider the strategies where Player 2 and Player 3 always take their loops but Player 1 never takes his loop, and such that whenever the play deviates, Player 2 and 3 retaliate by never taking their loops. For the same reasons as before this strategy profile is a doomsday equilibrium.

Note that the first equilibrium requires one bit of memory for each player, to remember if they visit their s state for the first or second times. In the second equilibrium, only Player 2 and 3 needs a bit of memory. An exhaustive analysis shows that there is no memoryless doomsday equilibrium in this example.

3 Complexity of DE for Perfect Information Games

In this section, we prove the following result:

Theorem 1. *The problem of deciding the existence of a doomsday equilibrium in an n -player perfect information game arena and n objectives $(\varphi_i)_{1 \leq i \leq n}$ is:*

- PTIME-C if the objectives $(\varphi_i)_{1 \leq i \leq n}$ are either all Büchi, all co-Büchi or all reachability objectives,
- NP-HARD, CONP-HARD and in PSPACE if $(\varphi_i)_{1 \leq i \leq n}$ are parity objectives,
- PSPACE-C if $(\varphi_i)_{1 \leq i \leq n}$ are safety objectives.

In the sequel, game arena with perfect information are just called game arena.

Tail objectives We first present a generic algorithm that works for any tail objective and then analyze its complexity for the different cases. Then we establish the lower bounds. Let us consider the following algorithm:

- compute the retaliation region of each player: $R_i = \langle\langle i \rangle\rangle(\varphi_i \cup \bigcap_{j=1}^{j=n} \overline{\varphi_j})$;
- check for the existence of a play within $\bigcap_{i=1}^{i=n} R_i$ that satisfies all the objectives φ_i .

The correctness of this generic procedure is formalized in the following lemma:

Lemma 1. *Let $G = (S, \mathcal{P}, s_{\text{init}}, \Sigma, \Delta)$ be an n -player game arena with n tail objectives $(\varphi_i)_{1 \leq i \leq n}$. Let $R_i = \langle\langle i \rangle\rangle(\varphi_i \cup \bigcap_{j=1}^{j=n} \overline{\varphi_j})$ be the retaliation region for Player i . There is a doomsday equilibrium in G iff there exists an infinite play that (1) which belongs to $\bigcap_{i=1}^{i=n} \varphi_i$ and (2) stays within the set of states $\bigcap_{i=1}^{i=n} R_i$ and*

Proof. First, assume that there exists an infinite play ρ such that $\rho \in \bigcap_i (\varphi_i \cap R_i^\omega)$. From ρ , and the retaliating strategies that exist in all states of R_i for each player, we show the existence of DE $\Lambda = (\lambda_1, \lambda_2, \dots, \lambda_n)$. Player i plays strategy λ_i as follows: he plays according to the choices made in ρ as long as all the other players do so, and as soon as the play deviates from ρ , Player i plays his retaliating strategy (when it is his turn to play).

First, let us show that if Player j , for some $j \neq i$, deviates and the turn comes back to Player i in a state s then $s \in R_i$. Assume that Player j deviates when he is in some $s' \in S_j$. As before there was no deviation, by definition of ρ , s' belongs to R_i . But no matter what the adversary are doing in a state that belongs to R_i , the next state must be a state that belongs to R_i (there is only the possibility to leave R_i when Player i plays). So, by induction on the length of the segment of play that separates s' and s , we can conclude that s belongs to R_i . From s , Player i plays a retaliating strategy and so all the outcomes from s are in $\varphi_i \cup \bigcap_{j=1}^{j=n} \overline{\varphi_j}$, and since the objective are tails, the prefix up to s is not important and we get (from s_{init}) $\text{outcome}_i(\lambda_i) \subseteq \varphi_i \cup \bigcap_{j=1}^{j=n} \overline{\varphi_j}$. Therefore the second property of the definition of doomsday equilibria is satisfied. Hence Λ is a DE.

Let us now consider the other direction. Assume that Λ is a DE. Then let us show that $\rho = \text{outcome}(\Lambda)$ satisfies properties (1) and (2). By definition of DE, we know that ρ is winning for all the players, so (1) is satisfied. Again by definition of DE, $\text{outcome}(\Lambda_i) \subseteq \varphi_i \cup \bigcap_{j=1}^{j=n} \overline{\varphi_j}$. Let s be a state of ρ and π the prefix of ρ up to s . For all outcomes ρ' of Λ_i in G_s , we have $\pi\rho' \in \varphi_i \cup \bigcap_{j=1}^{j=n} \overline{\varphi_j}$, and since the objectives are tail, we get $\rho' \in \varphi_i \cup \bigcap_{j=1}^{j=n} \overline{\varphi_j}$. Hence $s \in R_i$. Since this property holds for all i , we get $s \in \bigcap_i R_i$, and (2) is satisfied. \square

Accordingly, we obtain the following upper-bounds:

Lemma 2. *The problem of deciding the existence of a doomsday equilibrium in an n -player game arena can be decided in PTIME for Büchi and co-Büchi objectives, and in PSPACE for parity objectives.*

Proof. By Lemma 1 one first needs to compute the retaliation regions R_i for all $i \in \{1, \dots, n\}$. Once the sets R_i have been computed, it is clear that the existence

of a play winning for all players is decidable in PTIME for all the three types of objectives. For the Büchi and the co-Büchi cases, let us show how to compute the retaliation regions R_i . We start with Büchi and we assume that each player wants to visit a set of states T_i infinitely often. Computing the sets R_i boils down to computing the set of states s from which Player i has a strategy to enforce the objective (in LTL syntax) $\Box\Diamond T_i \vee \bigwedge_{j=1}^{j=n} \Diamond\Box\overline{T_j}$, which is equivalent to the formula $\Box\Diamond T_i \vee \Diamond\Box\bigcap_{j=1}^{j=n} \overline{T_j}$. This is equivalent to a disjunction of a Büchi and a co-Büchi objective, which is thus equivalent to a Streett objective with one Streett pair and can be solved in PTime with a classical algorithm, e.g. [PP06]. Similarly, for co-Büchi objectives, one can reduce the computation of the regions R_i in polynomial time to the disjunction of a Büchi objective and a co-Büchi objective.

For the parity case, the winning objectives for the retaliation sets can be encoded compactly as Muller objectives defined by a propositional formula using one proposition per state. Then they can be solved in PSPACE using the algorithm of Emerson and Lei defined in [EL85]. \square

Let us now establish the lower bounds.

Lemma 3. *The problem of deciding the existence of a DE in an n -player game arena is PTIME-HARD for Büchi and co-Büchi objectives, NP-HARD and CONP-HARD for parity objectives. All the hardness results hold even for a fixed number of players.*

The proof of Lemma 3 is decomposed into several lemmas.

Lemma 4. *The problem of deciding the existence of a doomsday equilibrium in a 2-player game arena is PTIME-hard both for Büchi and co-Büchi winning objectives.*

Proof. We explain the result for Büchi objectives (the proof for co-Büchi objectives is similar). To establish this result, we show how to reduce the problem of deciding the winner in a two-player zero-sum game with a Büchi objective (for Player 1), a PTIME-C problem [Imm81], can be reduced to the existence of a doomsday equilibrium in a two-player game arena with Büchi objectives. Let G be the two-player game, S its set of states, and T the set of states that Player 1 wants to visit infinitely often. We reduce the problem of deciding the existence of such a strategy to the existence of a doomsday equilibrium in the same game arena, where the objective of Player 1 is the original Büchi objective, i.e. $\text{Büchi}(T)$, and the objective of Player 2 is trivial: $\text{Büchi}(S)$. Clearly, as Player 2 will always satisfy his objective, Player 1 must have a winning strategy for $\text{Büchi}(T)$ if a doomsday equilibrium exists (and vice versa) otherwise condition 2 would be violated. \square

We now turn to the proof of lower bounds for parity objectives. We build our proof for these lower bounds on the hardness of generalized parity games [CHP07]: in a two-player (called Player A and Player B) zero-sum game, and an objective given by:

- the *disjunction* of two parity objectives, it is NP-HARD to decide if Player A has a winning strategy,
- the *conjunction* of two parity objectives, it is CONP-HARD to decide if Player A has a winning strategy.

We next show that these decision problems can be reduced to the problem of the existence of a doomsday equilibrium in an n player game with parity objectives.

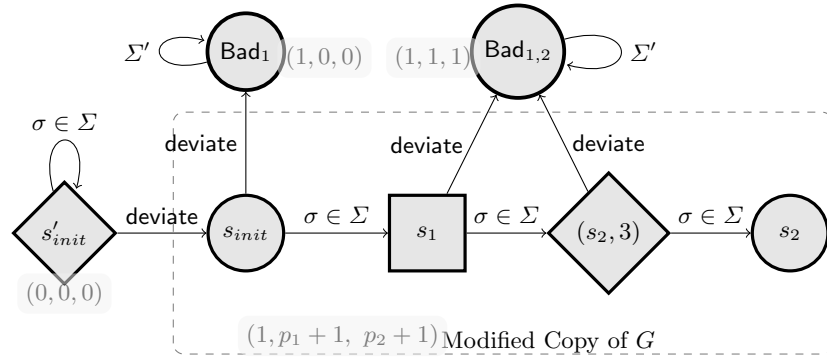


Fig. 2. Structure of the reduction from generalized parity game with a conjunction of two parity objectives to the existence of a doomsday equilibrium with parity objectives.

Lemma 5 (CONP-HARDNESS). *The problem of deciding the existence of a doomsday equilibrium in a 3-player game arena with parity objectives is CONP-HARD.*

Proof. Let $G = (S, \{S_A, S_B\}, s_{\text{init}}, \Sigma, \Delta)$ be a two-player game and a conjunction of two parity objectives defined by the functions p_1 and p_2 that Player A wants to enforce, i.e. the objective of Player A is to ensure an outcome that satisfies the two parity objectives, while the objective of Player B is to ensure an outcome that violates at least one of the two parity objectives. W.l.o.g., we assume that $s_{\text{init}} \in S_A$ and the turns of A and B alternate.

From G , we construct a 3-player game arena $G' = (S', \{S'_1, S'_2, S'_3\}, s'_{\text{init}}, \Sigma', \Delta')$ (depicted in Fig. 2), with:

- the set of states $S' = \{s'_{\text{init}}, \text{Bad}_1, \text{Bad}_{1,2,3}\} \cup S_A \cup S_B \cup (S_A \times \{3\})$, this set is partitioned as follows: $S'_1 = S_A \cup \{\text{Bad}_1\}$, $S'_2 = S_B$, $S'_3 = (S_A \times \{3\}) \cup \{s'_{\text{init}}, \text{Bad}_{1,2,3}\}$.
- the initial state is s'_{init} ,
- the alphabet of actions is $\Sigma' = \Sigma \cup \{\text{deviate}\}$,
- and the transitions of the game G' are defined as follows:
 - For the state s'_{init} , for all $\sigma \in \Sigma$, $\Delta'(s'_{\text{init}}, \sigma) = s'_{\text{init}}$, and $\Delta'(s'_{\text{init}}, \sigma) = s_{\text{init}}$; i.e., the play stay in s'_{init} , unless Player 3 plays *deviate* in which case the play goes to s_{init} that is the copy of the initial state of the game arena G .

- For all states $s \in S_A$, for all $\sigma \in \Sigma$, $\Delta'(s, \sigma) = \Delta(s, \sigma)$, and $\Delta'(s, \text{deviate}) = \text{Bad}_1$, so the transition function on the copy of G behaves from states owned by Player 1 as in the original game and it sends the game to Bad_1 if Player 1 plays the action **deviate**.
 - For all states $s \in S_B$, for all $\sigma \in \Sigma$, $\Delta'(s, \sigma) = (\Delta(s, \sigma), 3)$ and $\Delta'(s, \text{deviate}) = \text{Bad}_{1,2,3}$; i.e., if Player 2 plays an action from the game G , the effect is to send the game to the Player 3 copy of the same state as in the original game, if he deviates, the game reaches the sink state $\text{Bad}_{1,2,3}$.
 - For all states $s \in S_A \times \{3\}$, for all $\sigma \in \Sigma$, $\Delta'((s, 3), \sigma) = s$ and $\Delta'((s, 3), \text{deviate}) = \text{Bad}_{1,2,3}$. I.e. if Player 3 plays an action $\sigma \in \Sigma$, he gives back the turn to Player 1, otherwise he sends the game to $\text{Bad}_{1,2,3}$.
 - The states Bad_1 and $\text{Bad}_{1,2,3}$ are absorbing.
- The parity functions $(p'_i)_{i=1,2,3}$ for the three players are defined to satisfy the following condition:
- first, $p'_i(s'_{\text{init}})$ is even for all $i = 1, 2, 3$ (so if the game stays there for ever, all the players satisfy their objectives).
 - second, in Bad_1 the parity functions return an even number for Player 2 and Player 3 but an odd number for Player 1, this ensures that Player 1 should never play the action **deviate** when the game is in the copy of G ,
 - third, in $\text{Bad}_{1,2,3}$ the parity functions are odd for all the Players. So whenever Player 2 and 3 play **deviate** all players loose,
 - finally, in the copy of G , the parity function is always odd for Player 1, and for all states $q \in S_A \cup S_B \cup (S_A \times \{3\})$, $p'_2(q) = p_1(s) + 1$ and $p'_3(q) = p_2(s) + 1$, where $s = q$ if $q \in S_A \cup S_B$, and s is such that $q = (s, 3)$ if $q \in S_A \times \{3\}$.

This concludes the reduction.

Clearly, since Player A and B always alternate their moves, in the copy of G , any play will eventually reach a state of Player 2 and a state of Player 3, so that they are always able to retaliate by playing the action **deviate**.

A doomsday equilibrium exists in G' iff Player 1 is also able to retaliate when the game enters the copy of G . But clearly, it is possible if and only if he has a strategy to ensure parity(p'_1) and parity(p'_2), or equivalently iff he has a strategy to ensure parity(p_1) and parity(p_2), iff Player A has a winning strategy in the game G for the conjunction of parity objectives p_1 and p_2 . \square

Lemma 6 (NP-HARDNESS). *The problem of deciding the existence of a doomsday equilibrium in a 2-player game arena with parity objectives is NP-HARD.*

Proof. For this part, we need to show how to reduce the problem of deciding if Player A has a winning strategy in a two-player zero-sum game whose objective is defined by the disjunction of two parity objectives. We only sketch the construction as it is based on the main ideas used in the CONP-HARDNESS result. Let $G = (S, s_{\text{init}}, \Sigma, \Delta)$ be a two-player game with a *disjunction* of two parity objectives defined by the functions p_1 and p_2 . The objective of Player A is to ensure an outcome that satisfies at least one of the two parity objectives (while

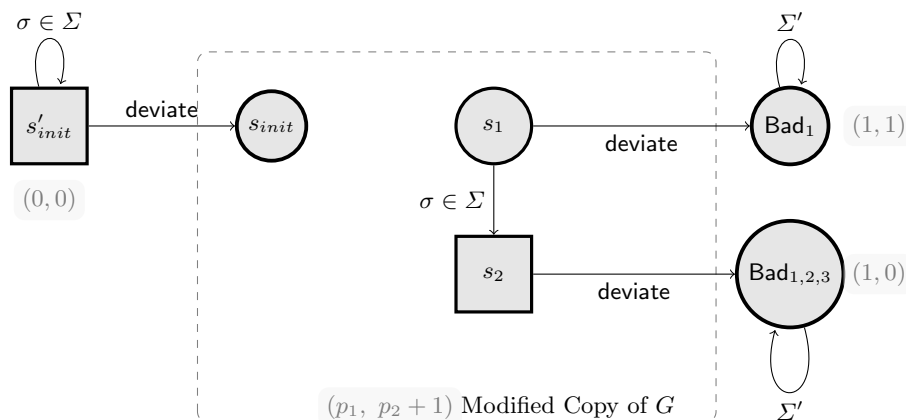


Fig. 3. Structure of the reduction from generalized parity game with a disjunction of two parity objectives to doomsday equilibrium with parity objectives.

the objective of Player B is to ensure an outcome that violates both parity objectives.)

From G , we construct a two-player game G' with parity objectives $(p'_i)_{i=1,2}$ (see Fig. 3). The game arena G' contains a copy of G plus three states s'_{init} (the initial state), Bad_1 and $Bad_{1,2}$. The alphabet of actions is $\Sigma \cup \{\text{deviate}\}$. The partition of the state space is as follows: $S_1 = S_A$ and $S_2 = S_B \cup \{s'_{init}\} \cup \{Bad_1, Bad_{1,2}\}$. The transitions are as follows: if Player 2 plays $\sigma \in \Sigma$ in s'_{init} then the game stays there, if he plays **deviate** then the game enters the copy of G . There the transition function for $\sigma \in \Sigma$ is defined as in G , and if Player 1 plays **deviate** then the game goes to Bad_1 , and if Player 2 plays **deviate** then the game goes to $Bad_{1,2}$. The parity functions $(p'_i)_{i=1,2}$ are defined as follows: p'_1 returns an even number in s'_{init} , is equal to p_1 in the copy of G , returns an odd number in Bad_1 and $Bad_{1,2}$. The function p'_2 returns an even number in s'_{init} , is equal to $p_2 + 1$ in the copy of G (so p'_2 is the complement of p_2), returns an even number in Bad_1 and an odd number in $Bad_{1,2}$. This definition of p'_1 and p'_2 ensures that the two players meet their parity objectives when the game always stay in s'_{init} and when the game enters the copy of G , Player 2 can always retaliate while Player 1 can retaliate if and only if Player A has a winning strategy in the original game. \square

As a corollary of this result, deciding the existence of a secure equilibrium in a 2-player game such that both players satisfy their parity objectives is NP-HARD.

Reachability objectives We now establish the complexity of deciding the existence of a doomsday equilibria in an n -player game with reachability objectives. We first establish an important property for reachability objectives:

Proposition 1. *Let $G = (S, \mathcal{P}, s_{init}, \Sigma, \Delta)$ be a game arena, and $(T_i)_{1 \leq i \leq n}$ be n subsets of S . Let A be a doomsday equilibrium in G for the reachability objectives*

$(\text{Reach}(T_i))_{1 \leq i \leq n}$. Let s the first state in $\text{outcome}(\Lambda)$ such that $s \in \bigcup_i T_i$. Then every player has a strategy from s , against all the other players, to reach his target set.

Proof. W.l.o.g. we can assume that $s \in T_1$. If some player, say Player 2, has no strategy from s to reach his target set T_2 , then necessarily $s \notin T_2$ and by determinacy the other players have a strategy from s to make Player 2 lose. This contradicts the fact that Λ is a doomsday equilibrium as it means that Λ_2 is not a retaliating strategy. \square

Lemma 7. *The problem of deciding the existence of a doomsday equilibrium in an n -player game with reachability objectives is in PTIME.*

Proof. The algorithm consists in:

(1) computing the sets R_i from which player i can retaliate, i.e. the set of states s from which Player i has a strategy to force, against all other players, an outcome such that $\diamond T_i \vee (\bigwedge_{j=1}^{j \neq i} \square \overline{T_j})$. This set can be obtained by first computing the set of states $\langle\langle i \rangle\rangle \diamond T_i$ from which Player i can force to reach T_i . It is done in PTIME by solving a classical two-player reachability game. Then the set of states where Player i has a strategy λ_i such that $\text{outcome}_i(\lambda_i) \models \square((\bigwedge_{j=1}^{j \neq i} \overline{T_j}) \vee \langle\langle i \rangle\rangle \diamond T_i)$, that is to confine the plays in states that do not satisfy the reachability objectives of the adversaries or from where Player i can force its own reachability objective. Again this can be done in PTIME by solving a classical two-player safety game.

(2) then, checking the existence of some $i \in \{1, \dots, n\}$ and some finite path π starting from s_{init} and that stays within $\bigcap_{j=1}^{j \neq i} R_j$ before reaching a state s such that $s \in T_i$ and $s \in \bigcap_{j=1}^{j \neq i} \langle\langle j \rangle\rangle \diamond T_j$.

Let us now prove the correctness of our algorithm. From its output, we can construct the strategy profile Λ where each Λ_j ($j = 1, \dots, n$) is as follows: follow π up to the point where either another player deviates and then play the retaliating strategy available in R_i , or to the point where s is visited for the first time and then play according to a strategy (from s) that force a visit to T_i no matter how the other players are playing. Clearly, Λ witnesses a DE. Indeed, if s is reached, then all players have a strategy to reach their target set (including Player i since $s \in T_i$). By playing so they will all eventually reach it. Before reaching s , if some of them deviate, the other have a strategy to retaliate as π stays in $\bigcap_{j=1}^{j \neq i} R_j$. The other direction follows from Proposition 1. \square

Lemma 8. *The problem of deciding the existence of a DE in a 2-player game with reachability objectives is PTIME-HARD.*

Proof. It is proved by an easy reduction from the And-Or graph reachability problem [Imm81]: if reachability is trivial for one of the two players, the existence of a doomsday equilibrium is equivalent to the existence of a winning strategy for the other player in a two-player zero sum reachability game. \square

Safety Objectives We establish the complexity of deciding the existence of a doomsday equilibrium in an n -player game with perfect information and safety objectives.

Lemma 9 (PSPACE-EASINESS). *The existence of a doomsday equilibrium in an n -player game with safety objectives can be decided in PSPACE.*

Proof. Let us consider an n -player game arena $G = (S, \mathcal{P}, s_{\text{init}}, \Sigma, \Delta)$ and n safety objectives $\text{safe}(T_1), \dots, \text{safe}(T_n)$ for $T_1 \subseteq S, \dots, T_n \subseteq S$. The algorithm is composed of the following two steps:

(1) For each Player i , compute the set of states $s \in S$ in the game such that Player i can retaliate whenever necessary, i.e. the set of states s from where there exists a strategy λ_i for Player i such that $\text{outcome}_i(\lambda_i)$ satisfies $\neg(\Box T_i) \rightarrow \bigwedge_{j=1}^{j=n} \neg \Box T_j$, or equivalently $\neg(\Diamond \overline{T_i}) \vee \bigwedge_{j=1}^{j=n} \Diamond \overline{T_j}$. This can be done in PSPACE using a result by Alur et al. (Theorem 5.4 of [AT04]) on solving two-player games whose Player 1's objective is defined by Boolean combinations of LTL formulas that use only \Diamond and \wedge . We denote by R_i the set of states in G where Player i has a strategy to retaliate.

(2) then, verify whether there exists an infinite path in $\bigcap_{i=1}^{i=n} (\text{safe}(T_i) \cap R_i)$. Now, let us establish the correctness of this algorithm. Assume that an infinite path exists in $\bigcap_{i=1}^{i=n} (\text{safe}(T_i) \cap R_i)$. The strategies λ_i for each Player i are defined as follows: play the moves that are prescribed as long as every other players do so, and as soon as the play deviates from the infinite path, play the retaliating strategy.

It is easy to see that the profile of strategies $\Lambda = (\lambda_1, \lambda_2, \dots, \lambda_n)$ is a DE. Indeed, the states are all safe for all players as long as they play their strategies. Moreover, as before deviation the play is within $\bigcap_{i=1}^{i=n} R_i$, if Player j deviates, we know that the state that is reached after deviation is still in $\bigcap_{j=1}^{j=n} R_j$ and therefore the other players can retaliate.

Second, assume that $\Lambda = (\lambda_1, \lambda_2, \dots, \lambda_n)$ is a DE in the n -player game G for the safety objectives $(\text{safe}(T_i))_{1 \leq i \leq n}$. Let $\rho = \text{outcome}(\lambda_1, \lambda_2, \dots, \lambda_n)$. By definition of doomsday equilibrium, we know that all states appearing in ρ satisfy all the safety objectives, i.e. $\rho \models \bigwedge_{i=1}^{i=n} \Box T_i$. Let us show that the play also remains within $\bigcap_{i=1}^{i=n} R_i$. Let s be a state of ρ , $i \in \{1, \dots, n\}$, and π the finite prefix of ρ up to s . By definition of DE we have $\text{outcome}(\lambda_i) \models \Box T_i \vee \bigwedge_{j=1}^{j=n} \Diamond \overline{T_j}$. Therefore for all outcomes ρ' of λ_i in G_s , $\pi \rho' \models \Box T_i \vee \bigwedge_{j=1}^{j=n} \Diamond \overline{T_j}$. Moreover, $\pi \models \bigwedge_{j=1}^{j=n} \Box T_j$ since it is a prefix of ρ . Therefore $\rho' \models \Box T_i \vee \bigwedge_{j=1}^{j=n} \Diamond \overline{T_j}$ and $s \in R_i$. Since it holds for all $i \in \{1, \dots, n\}$, we get $s \in \bigcap_{i=1}^{i=n} R_i$. \square

Lemma 10 (PSPACE-HARDNESS). *The problem of deciding the existence of a doomsday equilibrium in an n -player game with safety objectives is PSPACE-HARD.*

Proof. We reduce the two-player *multi-reachability* problem to our problem, PSPACE-HARDNESS follows. Let $G = (S^G, \{S_A^G, S_B^G\}, s_{\text{init}}^G, \Sigma^G, \Delta^G)$ be a two-player (Player A and Player B) game arena. Let $\mathcal{T} = \{T_1, T_2, \dots, T_k\}$ be a

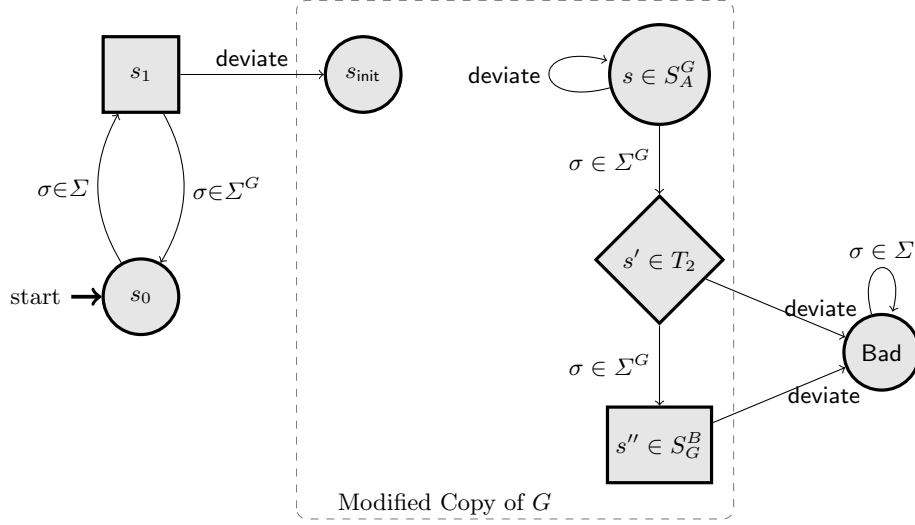


Fig. 4. Structure of the reduction from multi-reachability game to doomsday equilibrium with safety objectives. Round nodes denote Player 0's states and rectangular nodes denote Player 1's states.

family of subsets of S^G supposed to be pairwise disjoint (w.l.o.g.). Also wlog we assume that $T_i \subseteq S_B^G$ for all $i \in \{1, \dots, k\}$. In a multi-reachability game, the objective of Player A is to visit each T in \mathcal{T} , while Player B tries to avoid at least one of the subsets in \mathcal{T} . So, multi-reachability games are two-player zero sum games where the winning plays for Player A are

$$\{\rho = s_0 s_1 \dots s_n \dots \in \text{Plays}(G) \mid \forall i \cdot 1 \leq i \leq k \cdot \exists j \geq 0 \cdot s_j \in T_i\}.$$

It has been shown that the multi-reachability problem for two-player games is PSPACE-C [ATM03, FH10].

From $G = (S^G, \{S_A^G, S_B^G\}, s_{\text{init}}^G, \Sigma^G, \Delta^G)$ and $\mathcal{T} = \{T_1, T_2, \dots, T_k\}$ that define a multi-reachability game, we construct a game arena $G' = (S, \{S_0, S_1, \dots, S_k\}, s_0, \Sigma, \Delta)$ with $k+1$ players and a set of $k+1$ safety objectives $\text{Safe}_0, \dots, \text{Safe}_k$ such that Player A wins the multi-reachability objective defined by G and \mathcal{T} iff there exists a doomsday equilibrium in G' for the safety objectives $\text{Safe}_0, \dots, \text{Safe}_k$.

The structure of the reduction is depicted in Fig. 4. The state space of G' is composed of three parts: an initial part on the left, a modified copy of G , and a part on the right. The set S of states is $\{s_0, s_1\} \cup S_A^G \cup S_B^G \cup \{\text{Bad}\}$. This set of states is partitioned as follows: $S_0 = \{s_0\} \cup S_A^G \cup \{\text{Bad}\}$, $S_1 = S_B^G \setminus \bigcup_{i=2}^k T_i$ and for all i , $2 \leq i \leq k$, $S_i = T_i$.

The sets of safety objectives are defined as follows: $\text{Safe}_0 = \text{safe}(\{s_0, s_1\})$, and for all $i \in \{1, 2, \dots, k\}$, $\text{Safe}_i = \text{safe}(S \setminus (\{\text{Bad}\} \cup T_i))$. The alphabet of actions is $\Sigma = \Sigma^G \cup \{\text{deviate}\}$, and the transition function is defined as follows:

$$- \Delta(s_0, \sigma) = s_1, \text{ for all } \sigma \in \Sigma,$$

- $\Delta(s_1, \sigma) = \begin{cases} s_0 & \text{if } \sigma \in \Sigma^G \\ s_{\text{init}}^G & \text{if } \sigma = \text{deviate} \end{cases}$
- for all $s \in S_A^G \cup S_B^G$:
 - for all $\sigma \in \Sigma^G$, $\Delta(s, \sigma) = \Delta^G(s, \sigma)$
 - for the letter **deviate**: for all states $s \in S_A^G$, $\Delta(s, \text{deviate}) = s$, and for all $s \in S_B^G$, $\Delta(s, \text{deviate}) = \text{Bad}$
- **Bad** is a sink state.

Now, let us justify this construction. First, assume that, in the two-player game arena $G = (S^G, \{S_1^G, S_2^G\}, s_{\text{init}}^G, \Sigma^G, \Delta^G)$ with the multi-reachability objective given by $\mathcal{T} = \{T_1, T_2, \dots, T_k\}$, Player A has a winning strategy. In that case, we show that there exists a doomsday equilibrium in the game G' for the safety objectives $(\text{Safe}_i)_{0 \leq i \leq k}$. To establish the existence of a doomsday equilibrium, we consider the strategy profile $\Lambda = (\lambda_0, \lambda_1, \dots, \lambda_k)$ whose strategies respect the following conditions:

- If all the players follows the strategy profile Λ , the outcome of the game is $(s_0 \cdot s_1)^\omega$, i.e. Player 1 avoids to play **deviate** in s_1 .
- Whenever player 1 plays **deviate** in s_1 , then the game enters the sub game of G' corresponding to G , and the game thus enters an unsafe state for Player 0 (as s_{init} is not part of Safe_0). From there, Player 0 must retaliate by forcing a visit to each set in $\mathcal{T} = \{T_1, T_2, \dots, T_k\}$ to make sure that all the other players lose. By hypothesis, in G , Player A has a winning strategy for the multi-reachability objective, so we know that if the other players play letters that are in Σ^G then all sets in \mathcal{T} will eventually be visited when Player 0 plays according to the winning strategy of Player A in G . On the other hand, if the letter **deviate** is played then the game goes to the state **Bad** where all the safety objectives are violated. So, we have established that Player 0 can retaliate if he plays as Player A in the copy of G . Now, let us consider all the other players. According to the definition of the transition function, Player i has the option to retaliate whenever he enters its unsafe set T_i by choosing the action **deviate** and so force a visit to **Bad**. So, all other players have also the ability to retaliate whenever they enter their unsafe region.

So, we have established that $(\lambda_0, \lambda_1, \dots, \lambda_k)$ witnesses a doomsday equilibrium in G' .

Now, let us consider the other direction. Let $(\lambda_0, \lambda_1, \dots, \lambda_k)$ be a profile of strategies which witnesses a doomsday equilibrium for G' and the safety objectives given by the subsets of plays $(\text{Safe}_i)_{i=0, \dots, k}$. In that case, if we consider a prefix of play that enters for the first time the state s_{init} , we know by definition of doomsday equilibrium that Player 0 has a strategy to retaliate against any strategies of the adversaries. If all the other players chooses their letters in Σ^G then it should be the case that the play visits all the sets in \mathcal{T} . So, this clearly means that Player A has a winning strategy in G for the multi-reachability objective defined by \mathcal{T} , this strategy simply follows the strategy λ_0 in the copy of G . \square

4 Doomsday Equilibria Synthesis from LTL Specifications and Perfect Information

The doomsday equilibria synthesis problem (DE-synthesis problem) asks to automatically generate (whenever possible) a doomsday equilibrium from n objectives given by n LTL formulas $\varphi_1, \dots, \varphi_n$ (one for each player). Its associated decision problem (whether there exists a doomsday equilibrium) is called the DE-realizability problem. In the setting of synthesis from LTL specifications, the game graph is not given explicitly and the players do not play in turns, but concurrently. Each LTL formula φ_i is defined over a set P of atomic propositions, partitioned into n parts P_1, \dots, P_n .

Let $\Sigma = 2^P$ and $\Sigma_i = 2^{P_i}$ for all i . The *DE-realizability problem* can be seen as an n -player game where at each round r , Player i chooses a set of propositions $\sigma_i^r \in \Sigma_i$, for all $i \in \{1, \dots, n\}$. The outcome of the game is the infinite word $w = (\sigma_1^0 \cup \dots \cup \sigma_n^0)(\sigma_1^1 \cup \dots \cup \sigma_n^1) \dots \in \Sigma^\omega$. *Strategies* for Player i are mappings $\lambda_i : \Sigma^* \rightarrow \Sigma_i$. The *outcome* of λ_i , denoted by $\text{outcome}_i(\lambda_i)$ (or just $\text{outcome}(\lambda_i)$), is the set of ω -words $\sigma_0 \sigma_1 \dots \in \Sigma^\omega$ such that for all $r \geq 0$, $\sigma_r \cap \Sigma_i = \lambda_i(\sigma_0 \dots \sigma_{r-1})$. Given a strategy profile A , the outcome of A is defined by $\text{outcome}(A) = \bigcap_i \text{outcome}(A_i)$. Note that $\text{outcome}(A)$ is a singleton and we therefore identify it to an ω -word in Σ^ω . The *DE-realizability problem* asks to decide whether there exists a strategy profile A such that

1. $\text{outcome}(A) \models \bigwedge_{i=1}^n \varphi_i$
2. $\text{outcome}(A_i) \models \neg \varphi_i \rightarrow \bigwedge_{i \neq j} \neg \varphi_j$ for all $i \in \{1, \dots, n\}$.

If such a strategy profile exists, it is called a doomsday equilibrium and we say that $(\varphi_1, \dots, \varphi_n)$ is *DE-realizable*. The *synthesis problem* asks to construct a doomsday equilibrium if one exists.

In this section, we show that DE-realizability is a 2EXPTIME-C problem. As we will show, it generalizes over the classical (two-player) LTL realizability problem [PR89b], which is already 2EXPTIME-C.

Our decision procedure relies on Safra's determinization, which is notoriously difficult to implement efficiently [ATW06]. Therefore, we give a procedure that does not rely on it and is more likely to lead to efficient implementations. In particular, based on previous work on (classical) LTL synthesis [FJR11b, BBF⁺12], we show how to optimize this latter procedure with antichain data structures.

As the sets P_i are disjoint, we identify $\Sigma_1 \times \dots \times \Sigma_n$ with Σ , so we may freely write $(\sigma_1, \dots, \sigma_n) \in \Sigma$ where $\sigma_i \in \Sigma_i$. For all $\sigma \in \Sigma$, $\pi_i(\sigma)$ denotes the projection of σ on Σ_i .

Example 1 (Mutual Exclusion). Let $P = \{a_1, \dots, a_n\}$ and for all $i \in \{1, \dots, n\}$, let $P_i = \{a_i\}$. Consider the following LTL specifications:

$$\varphi_i =_{\text{def}} G(F(a_i)) \wedge G\left(\bigwedge_{i \neq j} \neg a_j \vee \neg a_i\right)$$

Intuitively, there is a common resource that each process i wants to access by asserting a_i infinitely often. However, when process i holds the resource, all the

other processes must not hold the resource. There exists a doomsday equilibrium which consists in asserting a_i every n time units, starting from the i -th tick, and always asserting a_i as soon as some player deviates from her strategy. The outcome of this strategy profile is $w = \{a_1\}\{a_2\} \dots \{a_n\}\{a_1\}\{a_2\} \dots$. Assume that Player i plays its retaliating strategy because some other player has deviated, then for all $j \neq i$, if Player j asserts a_j at some point, then the mutual exclusion condition of φ_j is violated, and if Player j never asserts a_j , the liveness condition of φ_j is then violated. In both cases, φ_j is not satisfied.

Example 2 (Classical LTL Realizability). This example shows that the DE-realizability problem is more general than the classical (0-sum 2-players) LTL realizability problem [PR89b]. Let P be a set of propositions partitioned into P_1 , the output propositions controlled by Player 1, and P_2 , the input propositions controlled by Player 2. Let φ be an LTL formula over P . It is said to be *realizable* if Player 1 has a strategy λ_1 such that for all strategies λ_2 of Player 2, $\text{outcome}(\lambda_1, \lambda_2) \models \varphi$. It is equivalent to say that the pair of formulas (φ, \top) is DE-realizable. Note however that the setting of DE-realizability slightly differs from the one of [PR89b], in which the two players play in a turn-based manner. However, it is easy to see that a formula φ is realizable in the concurrent setting iff it is realizable in the turn-based setting where Player 1 plays first.

4.1 Decidability and Finite-Memory Strategies

In this section, we prove that the DE-realizability problem is 2EXPTIME-C, and that whenever a tuple of objectives is DE-realizable, it is realizable by finite-memory strategies. A *finite-memory strategy* for Player i is defined by a deterministic finite-state machine $M = (Q, q_0, T)$ where Q is a finite set of states with initial state $q_0 \in Q$ and $T : Q \times \Sigma \rightarrow Q$ is a (partial) transition function s.t.

1. *it is complete for all actions of the other players:* for all $q \in Q$, for all $\sigma_{-i} \in \Sigma \setminus \Sigma_i$, there exists $\sigma_i \in \Sigma_i$ such that $T(q, \sigma_{-i} \cup \sigma_i)$ is defined.
2. *Player i 's actions only depend on the current state:* for all $q \in Q$, for all $\sigma, \sigma' \in \Sigma$, if $T(q, \sigma)$ and $T(q, \sigma')$ are defined, then $\Sigma_i \cap \sigma = \Sigma_i \cap \sigma'$. We write $O_M(q)$ (or just $O(q)$) for $\Sigma_i \cap \sigma$.

The language $L(M)$ of M is defined as the set of ω -words $\sigma_0\sigma_1\sigma_2 \dots \in \Sigma^\omega$ such that there exists a sequence of states $p_0p_1 \dots$ such that $p_0 = q_0$ and for all $i \geq 0$, $T(p_i, \sigma_i) = p_{i+1}$. As M is deterministic, for all words $w \in L(M)$ and all prefixes u of w , there exists a unique state denoted by $\text{reach}_M(u)$ reached by M after reading u . The strategy λ_M induced by M is defined only on prefixes u of $L(M)$ as follows: $\lambda_M(u) = O_M(\text{reach}_M(u))$. Clearly the following holds: $\text{outcome}_i(\lambda_M) = L(M)$. The size of λ_M is defined as the number of states of M .

The 2ExpTime upper bound for DE-realizability is obtained via a reduction to the problem of deciding the existence of doomsday equilibrium in an n player game with parity objectives, as LTL objectives can be encoded as parity objectives. However, one cannot reuse directly the results of Section 3 as we now

are in a concurrent setting. Moreover, we need to inspect closely the size of the strategies that witness the existence of doomsday equilibria in parity games. It will be needed to establish the correctness of the Safraless procedure given later. The game arenas considered in this section are *concurrent* (each player chooses his actions concurrently). They are defined as tuples $G = (S, \mathcal{P}, s_{\text{init}}, \Sigma, \Delta)$ where $\Delta : S \times \Sigma \rightarrow S$. The notion of strategies, outcomes and winning conditions carry over to this setting. Given a tuple of parity objectives, the procedure to decide the existence of a doomsday equilibrium in such a game is very similar to the one in the turn-based setting. However in the following lemma, we closely inspect the complexity and the size of winning strategies.

Lemma 11. *The existence of a doomsday equilibrium in an n -player (concurrent) game arena G with parity objectives f_1, \dots, f_n can be decided with an algorithm whose time complexity is exponential in n and in the number of different priorities.*

Moreover, if a doomsday equilibrium exists, then it is witnessed by a tuple of finite-memory strategies, each of them of size $m^n d^{nd} (d^2 n)!$ at most, where m is the number of vertices in the game and d is the maximal number of priorities over all parity objectives.

Proof. We first prove that the sets of states from which each player can retaliate (the sets R_i) are computable with an algorithm whose time complexity is exponential only in n (the number of players) and in d the number of different priorities, i.e. $d = \max_i d_i$ where $d_i = |\text{codom}(f_i)|$. Let s be a state of G . We have $s \in R_i$ iff Player i has a strategy λ_i from state s such that

$$\text{outcome}_i(\lambda_i) \models f_i \vee \bigwedge_{i \neq j} \neg f_j \quad \text{or equivalently,} \quad \text{outcome}_i(\lambda_i) \models \bigwedge_{i \neq j} (f_i \vee \neg f_j)$$

We show that we can test whether $s \in R_i$ for any given state s , by solving a generalized parity game G_i with a conjunction of $n - 1$ parity objectives, encoding every objective $\neg f_i \vee f_j$ respectively. For this purpose we can see G with the parity function f_k ($1 \leq k \leq n$) as a deterministic parity automaton over Σ , with initial state s . Let us denote by A_k this automaton: its language is the set of infinite sequences of actions in Σ whose runs are in $\text{parity}(f_k)$.

It is well known that deterministic parity automata are closed under complement, intersection, and union. Given a deterministic parity automaton A with n_A states and two priority functions g_1 and g_2 with at most e priorities each, the union of the languages defined by (A, g_1) and (A, g_2) can be defined by a deterministic parity automaton with at most $n_A \cdot e^e$ states and e^2 different priorities (see for instance [BFK]). It is also well known that deterministic parity automata can be complemented in linear time by shifting the parity function from 1. Therefore for all $j \neq i$, one can construct (from A_i and $\overline{A_j}$) a deterministic parity automaton B_j for the winning objective $f_i \vee \neg f_j$ of size $m \cdot d^d$ at most with d^2 different priorities.

Finally, if g_j denotes the parity function of B_j , in order to test whether $s \in R_i$, it suffices to solve the generalized parity game $G_i := G \otimes B_1 \otimes \dots \otimes B_{i-1} \otimes$

$B_{i+1} \otimes \cdots \otimes B_n$ with the conjunction of parity objectives defined by the parity functions $(s, q_1, \dots, q_{i-1}, q_{i+1}, \dots, q_n) \mapsto g_j(q_j)$. It is known (see for instance [CHP07]) that a generalized parity game of size α with a conjunction of k parity conditions and d' different priorities each can be solved in time complexity $O(\alpha^{\text{poly}(k, d')} \text{exp}(k, d'))$, where *poly* and *exp* are polynomial and exponential functions respectively. The number of vertices of G_i is $m \cdot (m \cdot d^d)^{n-1} = m^n d^{d(n-1)}$ with a conjunction of $n-1$ parity functions with d^2 different priorities each. Therefore, the game G_i can be solved in time complexity

$$O((m^n d^{d(n-1)})^{\text{poly}(n-1, d^2)} \text{exp}(n-1, d^2)) = O(m^{\text{poly}'(n, d)} \text{exp}'(n, d))$$

for some polynomial and exponential functions poly' and exp' respectively.

Once the sets R_i are computed for all $i = 1, \dots, n$, then for each state $s \in R_i$, we compute the set of actions $\sigma_i \in \Sigma_i$ such that for all actions of the other players, the successor state is in R_i . Such actions are called *good actions* and can be computed in PTime. It remains to check the existence of an infinite path that satisfies all parity objectives and is made of good actions only. For this purpose we restrict the vertices of G to $\bigcap_i R_i$ and the actions to good actions, and see the induced subgraph as a Streett automaton A with nd pairs (each parity condition is encoded with d pairs), whose non-emptiness witnesses existence of a doomsday equilibrium. It is well-known that non-emptiness of Streett automata can be decided in PTime.

Let us now establish a bound on the size of finite-memory strategies sufficient to witness the existence of a doomsday equilibrium. Generalized parity games with conjunction of parity objectives can be seen as Streett games, and therefore one can reuse classical results on the size of strategies that are sufficient to win generalized parity games (see [CHP07] for more details). In particular, one gets that the existence of winning strategies in each game G_i is witnessed by finite-memory strategies of size $m^n d^{d(n-1)} (d^2(n-1))!$ at most. This bounds the size of retaliating strategies. It remains to bound the size of (a representation of) an infinite path in G that satisfies all parity conditions.

It is known that the non-emptiness of a Streett automaton with α states and β pairs is witnessed by a lasso path uv^ω such that $|u| + |v| \leq \alpha\beta^2\beta!$. Therefore the non-emptiness of the Streett automaton A is witnessed by a lasso path of length $m(nd)^2(nd)!$ at most (see [Pit07] for instance). Given such an infinite path, remind the strategy of Player i is to play good actions as long as the play follows the path, and its retaliating strategy if the play deviates from it. Therefore, if there exists a doomsday equilibrium, there exists one such that the strategy of Player i , for all i , has finite-memory, of size (at most)

$$m(nd)^2(nd)! + m^n d^{d(n-1)} (d^2(n-1))! \leq m^n d^{nd} (d^2n)!$$

This concludes the proof. \square

Theorem 2. *The DE-realizability problem from LTL objectives is 2EXPTIME-C.*

Proof. Upper bound We reduce the DE-realizability problem from n LTL objectives to the problem of deciding the existence of a doomsday equilibrium in an n -player concurrent game with parity objectives.

Based on several automata constructions among which Safra's determinization, we can construct from each LTL objective φ_i some equivalent deterministic parity automaton A_i whose set of states is denoted by Q_i , transition function by $T_i : Q_i \times \Sigma \rightarrow Q_i$ and parity function by f_i . Moreover, we can assume that A_i has doubly exponentially many states in $|\varphi_i|$ and exponentially many different priorities (see for instance [KV05]). Then the problem reduces to deciding the existence of a doomsday equilibrium in the synchronized product of those automata. The product $A = A_1 \otimes \dots \otimes A_n$ is the classical product automaton: its set of states is $Q = Q_1 \times \dots \times Q_n$ and there exists a transition $((q_1, \dots, q_n), \sigma, (q'_1, \dots, q'_n))$ in A iff $(q_i, \sigma, q'_i) \in T_i$ for all $i = 1, \dots, n$. Since A is deterministic, we can see it as a concurrent game arena with parity objectives given by the functions $(q_1, \dots, q_n) \mapsto f_i(q_i)$ for all players i . There exists a doomsday equilibrium in this game iff $(\varphi_1, \dots, \varphi_n)$ is DE-realizable. By Lemma 11, this game can be solved with an algorithm whose time complexity is exponential only in n and in the number of different priorities. This yields an algorithm for DE-realizability whose time complexity is doubly exponential in $\sum |\varphi_i|$.

Lower bound This follows from Example 2 and the 2-Exptime hardness of (two players) LTL realizability [PR89b]. \square

The following lemma states that finite-memory strategies are sufficient to witness DE-realizability of n LTL objectives, and gives an explicit bound on their memory size.

Lemma 12. *If a tuple of LTL objectives $(\varphi_1, \dots, \varphi_n)$ is DE-realizable, then it is DE-realizable by finite-memory strategies of size $2^{n^2 \cdot 2^{8F}} \cdot (2^{8F} n)!$ at most, where $F = \max_i |\varphi_i|$.*

Proof. As shown in Theorem 2, the problem reduces to test the existence of a doomsday equilibrium in a game arena with n parity objectives. The result then follows from Lemma 11. Let us inspect the construction in the proof of Theorem 2. Each deterministic parity automaton A_i is constructed from the LTL formula φ_i . It is known that A_i can be constructed by the following steps: (i) construct an alternating Büchi automaton with at most $|\varphi_i|$ states, (ii) remove universal transitions with Mihano-Ayashi construction, this yields a non-deterministic Büchi automaton with $2^{2|\varphi_i|}$ states, (iii) transform the latter automaton into a deterministic parity automaton, using for instance the improved Safra's construction of [Pit07], which yields an automaton with

$$(2^{2|\varphi_i|}) \cdot 2^{2^{2|\varphi_i|+2}} \leq 2^{2^{3(|\varphi_i|+1)}}$$

states and $2^{2|\varphi_i|+1} - 1$ priorities. Therefore the n -player game $A_1 \otimes \dots \otimes A_n$ has $2^{\sum_i 2^{3(|\varphi_i|+1)}}$ states and each objective has $2^{2|\varphi_i|+1} - 1$ priorities. By Lemma 11, we get, for $F = \sum_i |\varphi_i|$, that the existence of a doomsday equilibrium is witnessed by finite-memory strategies of size at most

$$(2^{n \cdot 2^{3(F+1)}})^n (2^{2F+1} - 1)^{n \cdot 2^{2F+1} - 1} ((2^{2F+1} - 1)^{2n})! \leq 2^{n^2 \cdot 2^{8F}} \cdot (2^{8F} n)!$$

This concludes the proof. \square

4.2 Safraless Procedure

The solution given in Theorem 2 relies on Safra’s determinization of non-deterministic Büchi automata. In this section, we give a procedure that avoids such a determinization by using universal co-Büchi automata.

Remind that a universal coBüchi automaton A is exactly the dual of a Büchi automaton: it accepts an infinite word w if all the runs of A on w visit a finite number of accepting states. It can be easily constructed from an LTL formula φ by constructing a Büchi automaton for $\neg\varphi$.

We denote by $\llbracket\varphi\rrbracket$ the models of any LTL formula φ . Given a tuple of LTL objectives $(\varphi_1, \dots, \varphi_n)$, a partition $P = \uplus_i P_i$, and parameter $K \in \mathbb{N}$, the Safraless procedure is based on the following steps:

1. for all $i \in \{1, \dots, n\}$, construct a universal co-Büchi automaton A_i equivalent to $\varphi_i \vee \bigwedge_{j \neq i} \neg\varphi_j$.
2. following the Safraless method for two-player LTL realizability [FJR11a,SF07], for all $i \in \{1, \dots, n\}$, construct a two-player safety game $G(A_i, K)$ such that Player 1 (the protagonist) wins $G(A_i, K)$ iff Player i has a strategy A_i (in the two-player LTL realizability game) such that all runs of A_i on words of $\text{outcome}(A_i)$ visit at most K accepting states (this implies $\text{outcome}(A_i) \subseteq L(A_i)$ and therefore $\text{outcome}(A_i) \models \varphi_i \vee \bigwedge_{j \neq i} \neg\varphi_j$).
3. test the existence of a strategy profile $\Lambda = (\Lambda_1, \dots, \Lambda_n)$ such that all strategies A_i is winning in $G(A_i, K)$, and such that $\text{outcome}(\Lambda) \models \bigwedge_i \varphi_i$.

We will show that this algorithm is correct and complete for a sufficiently large parameter K . Before this, let us explain how those steps are done:

(1) For the automata construction, we can apply well-known LTL-to-Büchi automata constructions on the negation of $\varphi_i \vee \bigwedge_{j \neq i} \neg\varphi_j$.

(2) For the second step, we use the reduction of classical two-player LTL realizability to two-player safety games [FJR11a,SF07]. The safety games we obtain are turn-based: in the initial state, Player 1 plays first and chooses some action $\sigma_i \in \Sigma_i$ and then Player 2 chooses some action $\sigma_{-i} \in \Sigma \setminus \Sigma_i$ and the game moves to another state (in a deterministic way), and so on. The transition relation is therefore of type $P \times \Sigma \rightarrow P$ where P are the states of $G(A_i, K)$. The safety condition is a subset of states that Player 1 does not want to leave. Any Player 1’s strategy in this two-player turn-based game can be seen as a strategy of Player i in the DE-realizability setting. However in the latter setting, the players choose their actions concurrently. This does not harm as any Player 1’s strategy in the turn-based safety game is winning iff it is winning in the same safety game where both players choose their actions concurrently (because Player 1 starts first in the turn-based setting).

(3) For all $i = 1, \dots, n$, $G(A_i, K)$ is a safety game, the set of all winning strategies of Player 1 can be compactly represented by the subgame obtained by restricting Player 1’s action to good actions, i.e. those actions from which whatever Player 2 does, the next state is safe. In order to check the existence of a doomsday equilibrium, it suffices to check the existence of an infinite path that satisfies $\bigwedge_i \varphi_i$ in the synchronized product of those subgames. This is a classical

model-checking problem. If there exists an infinite path, the tuple of strategies output by the algorithm is defined for all i as follows: as long as the play is a prefix of this infinite path, Player i plays according to this infinite path, and if the play deviates from this infinite path, Player i applies any strategy winning in $G(A_i, K)$ for Player 1.

Lemma 13. *Let $F = \max_i |\varphi_i|$. The Safrless procedure is correct and complete for any $K \geq 2^{n^2 2^{9F}} \cdot (2^{8F} n)!$.*

Proof. Let us first show that if strategy profile Λ is found, then $(\varphi_1, \dots, \varphi_n)$ is DE-realizable by Λ . The infinite word $\text{outcome}(\Lambda)$ corresponds to an infinite path in the synchronized product of the games $G(A_i, K)$ which by definition satisfies $\bigwedge_i \varphi_i$. Moreover, for all $i \in \{1, \dots, n\}$, Λ_i is a winning strategy in $G(A_i, K)$, therefore $\text{outcome}(\Lambda_i) \in \subseteq L(A_i) = \llbracket \varphi_i \vee \bigwedge_{j \neq i} \neg \varphi_j \rrbracket$, by definition of $G(A_i, K)$.

Conversely, let us show that if $(\varphi_1, \dots, \varphi_n)$ is DE-realizable, then for a sufficiently large bound K , the algorithm above returns a tuple of strategies. We know by Lemma 12 that if $(\varphi_1, \dots, \varphi_n)$ is DE-realizable, then it is DE-realizable by finite-memory strategies represented by finite-state machines M_i of size $\alpha := 2^{n^2 \cdot 2^{8F}} \cdot (2^{8F} n)!$ at most. By definition of doomsday equilibria, for all $i \in \{1, \dots, n\}$, $L(M_i) \subseteq L(A_i)$. Therefore each run of A_i on any word of $L(M_i)$ visits at most $\alpha \cdot n_i$ accepting states, where n_i is the number of states of A_i (see [FJR11a] for details). Therefore by taking $K = \alpha \cdot (\max_i n_i)$, the strategy λ_{M_i} induced by M_i is winning for Player 1 in $G(A_i, K)$. It remains to explicit the number of states of A_i . The automaton A_i can be obtained by constructing a non-deterministic Büchi automaton for the negated objective $\neg(\varphi_i \vee \bigwedge_{j \neq i} \neg \varphi_j) = \neg \varphi_i \wedge \bigvee_{j \neq i} \varphi_j$ of size $\sum_j |\varphi_j| + n + 1$. It is known that from any LTL formula of size n an equivalent automaton of size 2^{2n} can be constructed. Therefore we can assume that A_i has

$$2^{2 \cdot (\sum_j |\varphi_j| + n + 1)} \leq 2^{5nF}$$

states at most. Now

$$2^{5nF} \cdot \alpha = 2^{5nF} \cdot 2^{n^2 \cdot 2^{8F}} \cdot (2^{8F} n)! = 2^{5nF + n^2 \cdot 2^{8F}} \cdot (2^{8F} n)! \leq 2^{n^2 2^{9F}} \cdot (2^{8F} n)!$$

Therefore we can take $K = 2^{n^2 2^{9F}} \cdot (2^{8F} n)!$. □

4.3 Towards an efficient implementation

We discuss in this section some optimizations that could be used to efficiently implement the Safrless procedure.

Incremental Algorithm The Safrless procedure can be made incremental: check existence of a doomsday equilibrium for increasing values of the parameter K (initialized to 0) until the large bound K of Lemma 13 is reached. For classical (two-player) LTL realizability, it has been noticed that small values of K (up to 5) are sufficient to conclude for realizability in practice [FJR11a].

Antichain Data Structure Step 2 of the Safraless procedure is a classical two-player LTL realizability problem. There exist several efficient implementations of LTL realizability and synthesis, like Lily [JB06], Unbeast [Ehl11] or Acacia+ [BBF⁺12]. The latter tool is an implementation of the bounded synthesis method [FJR11a,SF07] optimized with antichain data structures. It is based on the following observation: the positions of the games $G(A_i, K)$ are functions f from Q_i to $\{-1, 0, \dots, K + 1\}$ (for each state q of Q_i there is a counter that counts the maximal number of accepting states visited so far by the runs ending in q , up to $K + 1$). Therefore they can be partially ordered with the pairwise comparison of vector components, denoted by \preceq_i . Moreover, it is also known that the winning region W_i of $G(A_i, K)$, for all i , are downward-closed sets for \preceq_i , and can therefore be represented by the antichain of their maximal elements, that we denote by L_i , i.e. $W_i = \downarrow L_i$ where \downarrow denotes the downward closure. The antichain that represents W_i is computed by a backward fixpoint algorithm, starting from the safe positions, i.e. the positions f such that for all $q \in Q_i$, $f(q) \leq K$.

Once the antichains L_1, \dots, L_n representing the winning sets W_1, \dots, W_n have been computed, it remains to check whether there exists an infinite sequence of actions $\sigma_0 \sigma_1 \dots \in \Sigma^\omega$ such that for all $i \in \{1, \dots, n\}$, the sequence $(\sigma_0 \cap \Sigma_i)(\sigma_0 \cap (\Sigma \setminus \Sigma_i))(\sigma_1 \cap \Sigma_i)(\sigma_1 \cap (\Sigma \setminus \Sigma_i)) \dots$ induces a play in $G(A_i, K)$ that stays in W_i . This problem can be encoded as a model-checking problem where the transitions of the games $G(A_i, K)$ are compactly represented. This can be done for instance by using NuSMV.

5 Complexity of DE for Imperfect Information Games

In this section, we define n -player game arenas with imperfect information. We adapt to this context the notions of observation, observation of a play, observation-based strategies, and we study the notion of doomsday equilibria when players are restricted to play observation-based strategies.

Game arena with imperfect information An n -player game arena with *imperfect information* is a tuple $G = (S, \mathcal{P}, s_{\text{init}}, \Sigma, \Delta, (O_i)_{1 \leq i \leq n})$ such that $(S, \mathcal{P}, s_{\text{init}}, \Sigma, \Delta)$ is a game arena (of perfect information) and for all i , $1 \leq i \leq n$, $O_i \subseteq 2^S$ is a *partition* of S . Each block in O_i is called an *observation* of Player i . We assume that the players play in a predefined order¹: for all $i \in \{1, \dots, n\}$, all $q \in S_i$ and all $\sigma \in \Sigma$, $\Delta(q, \sigma) \in S_{(i \bmod n) + 1}$.

Observations For all $i \in \{1, \dots, n\}$, we denote by $O_i(s) \subseteq S$ the block in O_i that contains s , that is the observation that Player i has when he is in state s . We say that two states s, s' are *undistinguishable* for Player i if $O_i(s) = O_i(s')$. This defines an equivalence relation on states that we denote by \sim_i . The notions of plays and prefixes of plays are slight variations from the perfect information setting: a play in G is a sequence $\rho = s_0, \sigma_0, s_1, \sigma_1, \dots \in (S \cdot \Sigma)^\omega$ such that $s_0 =$

¹ This restriction is not necessary to obtain the results presented in this section (e.g. Theorem 3) but it makes some of our notations lighter.




Fig. 5. Game arena with imperfect information and Büchi objectives. Only undistinguishable states of Player 1 (Player circle) are depicted. Observations are symmetric for the other players.

s_{init} , and for all $j \geq 0$, we have $s_{j+1} = \Delta(s_j, \sigma_j)$. A prefix of play is a sequence $\pi = s_0, \sigma_0, s_1, \sigma_1, \dots, s_k \in (S \cdot \Sigma)^* \cdot S$ that can be extended into a play. As in the perfect information setting, we use the notations $\text{Plays}(G)$ and $\text{PrefPlays}(G)$ to denote the set of plays in G and its set of prefixes, and $\text{PrefPlays}_i(G)$ for the set of prefixes that end in a state that belongs to Player i . While actions are introduced explicitly in our notion of play and prefix of play, their visibility is limited by the notion of observation. The *observation* of a play $\rho = s_0, \sigma_0, s_1, \sigma_1, \dots$ by Player i is the infinite sequence written $\text{Obs}_i(\rho) \in (O_i \times (\Sigma \cup \{\tau\}))^\omega$ such that for all $j \geq 0$, $\text{Obs}_i(\rho)(j) = (O_i(s_j), \tau)$ if $s_j \notin S_i$, and $\text{Obs}_i(\rho)(j) = (O_i(s_j), \sigma_j)$ if $s_j \in S_i$. Thus, only actions played by Player i are visible along the play, and the actions played by the other players are replaced by τ . The observation $\text{Obs}_i(\pi)$ of a prefix π is defined similarly. Given an infinite sequence of observations $\eta \in (O_i \times (\Sigma \cup \{\tau\}))^\omega$ for Player i , we denote by $\gamma_i(\eta)$ the set of plays in G that are compatible with η , i.e. $\gamma_i(\eta) = \{\rho \in \text{Plays}(G) \mid \text{Obs}_i(\rho) = \eta\}$. The functions γ_i are extended to prefixes of sequences of observations naturally.

Observation-based strategies and doomsday equilibria A strategy λ_i of Player i is *observation-based* if for all prefixes of plays $\pi_1, \pi_2 \in \text{PrefPlays}_i(G)$

such that $\text{Obs}_i(\pi_1) = \text{Obs}_i(\pi_2)$, it holds that $\lambda_i(\pi_1) = \lambda_i(\pi_2)$, i.e. while playing with an observation-based strategy, Player i plays the same action after undistinguishable prefixes. A strategy profile Λ is observation-based if each Λ_i is observation-based. Winning objectives, strategy outcomes and winning strategies are defined as in the perfect information setting. We also define the notion of outcome relative to a prefix of a play. Given an observation-based strategy λ_i for Player i , and a prefix $\pi = s_0, \sigma_0, \dots, s_k \in \text{PrefPlays}_i(G)$, the strategy λ_i^π is defined for all prefixes $\pi' \in \text{PrefPlays}_i(G_{s_k})$ where G_{s_k} is the game arena G with initial state s_k , by $\lambda_i^\pi(\pi') = \lambda_i(\pi \cdot \pi')$. The set of outcomes of the strategy λ_i relative to π is defined by $\text{outcome}_i(\pi, \lambda_i) = \pi \cdot \text{outcome}_i(\lambda_i^\pi)$.

The notion of doomsday equilibrium is defined as for games with perfect information but with the additional requirements that *only* observation-based strategies can be used by the players. Given an n -player game arena with imperfect information G and n winning objectives $(\varphi_i)_{1 \leq i \leq n}$ (defined as in the perfect information setting), we want to solve the problem of deciding the existence of an *observation-based strategy profile* Λ which is a doomsday equilibrium in G for $(\varphi_i)_{1 \leq i \leq n}$.

Example Fig. 5 depicts a variant of the example in the perfect information setting, with imperfect information. In this example let us describe the situation for Player 1. It is symmetric for the other players. Assume that when Player 2 or Player 3 send their information to Player 1 (modeled by a visit to his happy states), Player 1 cannot distinguish which of Player 2 or 3 has sent the information, e.g. because of the usage of a cryptographic primitive. Nevertheless, let us show that there exists doomsday equilibrium. Assume that the three players agree on the following protocol: Player 1 and 2 send their information but not Player 3.

Let us show that this sequence witnesses a doomsday equilibrium and argue that this is the case for Player 1. From the point of view of Player 1, if all players follow this profile of strategies then the outcome is winning for Player 1. Now, let us consider two types of deviation. First, assume that Player 2 does not send his information. In that case Player 1 will observe the deviation and can retaliate by not sending his own information. Therefore all the players are losing. Second, assume that Player 2 does not send his information but Player 3 does. In this case it is easy to verify that Player 1 cannot observe the deviation and so according to his strategy will continue to send his information. This is not problematic because all the plays that are compatible with Player 1's observations are such that: (i) they are winning for Player 1 (note that it would be also acceptable that all the sequence are either winning for Player 1 or losing for all the other players), and (ii) Player 1 is always in position to retaliate along this sequence of observations. In our solution below these two properties are central and will be called *doomsday compatible* and *good for retaliation*.

Generic Algorithm We present a generic algorithm to test the existence of an observation-based doomsday equilibrium in a game of imperfect information. To present this solution, we need two additional notions: sequences of

observations which are *doomsday compatible* and prefixes which are *good for retaliation*. These two notions are defined as follows. In a game arena $G = (S, \mathcal{P}, s_{\text{init}}, \Sigma, \Delta, (O_i)_{1 \leq i \leq n})$ with imperfect information and winning objectives $(\varphi_i)_{1 \leq i \leq n}$,

- a sequence of observations $\eta \in (O_i \times (\Sigma \cup \{\tau\}))^\omega$ is *doomsday compatible* (for Player i) if $\gamma_i(\eta) \subseteq \varphi_i \cup \bigcap_{j=1}^{j=n} \overline{\varphi_j}$, i.e. all plays that are compatible with η are either winning for Player i , or not winning for any other player,
- a prefix $\kappa \in (O_i \times (\Sigma \cup \{\tau\}))^* \cdot O_i$ of a sequence of observations is *good for retaliation* (for Player i) if there exists an observation-based strategy λ_i^R such that for all prefixes $\pi \in \gamma_i(\kappa)$ compatible with κ , $\text{outcome}(\pi, \lambda_i^R) \subseteq \varphi_i \cup \bigcap_{j=1}^{j=n} \overline{\varphi_j}$.

The next lemma shows that the notions of sequences of observations that are doomsday compatible and good for retaliation prefixes are important for studying the existence of doomsday equilibria for imperfect information games.

Lemma 14. *Let G be an n -player game arena with imperfect information and winning objectives φ_i , $1 \leq i \leq n$. There exists a doomsday equilibrium in G if and only if there exists a play ρ in G such that:*

- (F₁) $\rho \in \bigcap_{i=1}^{i=n} \varphi_i$, i.e. ρ is winning for all the players,
- (F₂) for all Player i , $1 \leq i \leq n$, for all prefixes κ of $\text{Obs}_i(\rho)$, κ is good for retaliation for Player i ,
- (F₃) for all Player i , $1 \leq i \leq n$, $\text{Obs}_i(\rho)$ is doomsday compatible for Player i .

Proof. First, assume that conditions (F₁), (F₂) and (F₃) hold and show that there exists a DE in G . We construct a DE $(\lambda_1, \dots, \lambda_n)$ as follows. For each player i , the strategy λ_i plays according to the (observation of the) path ρ in \mathcal{G} , as long as the previous observations follow ρ . If an observation is unexpected for Player i (i.e., differs from the sequence in ρ), then λ_i switches to an observation-based retaliating strategy λ_i^R (we will show that such a strategy exists as a consequence of (F₂)). This is a well-defined profile and a DE because: (1) all strategies are observation-based, and the outcome of the profile is the path ρ that satisfies all objectives; (2) if no deviation from the observation of ρ is detected by Player i , then by condition (F₃) we know that if the outcome does not satisfy φ_i , then it does not satisfy φ_j , for all $1 \leq j \leq n$, (3) if a deviation from the observation of ρ is detected by Player i , then the sequence of observations of Player i so far can be decomposed as $\kappa = \kappa_1(o_1, \sigma_1) \dots (o_m, \sigma_m)$ where (o_1, σ_1) is the first deviation of the observation of ρ , and (o_m, σ_m) is the first time it is Player i 's turn to play after this deviation (so possibly $m = 1$). By condition (F₂), we know that κ_1 is good for retaliation. Clearly, $\kappa_1(o_1, \sigma_1) \dots (o_\ell, \sigma_\ell)$ is retaliation compatible as well for all $\ell \in \{1, \dots, m\}$ since retaliation goodness is preserved by player j 's actions for all j . Therefore κ is good for retaliation and by definition of retaliation goodness there exists an observation-based retaliation strategy λ_i^R for Player i which ensures that that regardless of the strategies of the opponents in coalition, if the outcome does not satisfy φ_i , then for all $j \in \{1, \dots, n\}$, it does not satisfy φ_j either.

Second, assume that there exists a DE $(\lambda_1, \dots, \lambda_n)$ in G , and show that $(F_1), (F_2)$ and (F_3) hold. Let ρ be the outcome of the profile $(\lambda_1, \dots, \lambda_n)$. Then ρ satisfies (F_1) by definition of DE. Let us show that it also satisfies (F_3) . By contradiction, if $\text{obs}_i(\rho)$ is not doomsday compatible for Player i , then by definition, there is a path ρ' in $\text{Plays}(G)$ that is compatible with the observations and actions of player i in ρ (i.e., $\text{obs}_i(\rho) = \text{obs}_i(\rho')$), but ρ' does not satisfy φ_i , while it satisfies φ_j for some $j \neq i$. Then, given the strategy λ_i from the profile, the other players in coalition can choose actions to construct the path ρ' (since ρ and ρ' are observationally equivalent for player i , the observation-based strategy λ_i is going to play the same actions as in ρ). This would show that the profile is not a DE, establishing a contradiction. Hence $\text{obs}_i(\rho)$ is doomsday compatible for Player i for all $i = 1, \dots, n$ and (F_3) holds. Let us show that ρ also satisfies (F_2) . Assume that this not true. Assume that κ is a prefix of $\text{obs}_i(\rho)$ such that κ is not good for retaliation for Player i for some i . By definition it means that the other players can make a coalition and enforce an outcome ρ' , from any prefix of play compatible with κ , that is winning for one of players of the coalition, say Player $j, j \neq i$, and losing for Player i . This contradicts the fact that λ_i belongs to a DE. \square

Theorem 3. *The problem of deciding the existence of a doomsday equilibrium in an n -player game arena with imperfect information and n objectives is EXPTIME- C for objectives that are either all reachability, all safety, or all parity objectives.*

Proof. By Lemma 14, we know that we can decide the existence of a doomsday equilibrium by checking the existence of a play ρ in G that respects the conditions $(F_1), (F_2)$, and (F_3) . It can be shown (see Appendix), for all $i \in \{1, \dots, n\}$, that the set of good for retaliation prefixes for Player i is definable by a finite-state automaton C_i , and the set of observation sequences that are doomsday compatible for Player i is definable by a Streett automaton D_i .

From the automata $(D_i)_{1 \leq i \leq n}$ and $(C_i)_{1 \leq i \leq n}$, we construct using a synchronized product a finite transition system T and check for the existence of a path in T that satisfy the winning objectives for each player in G , the Streett acceptance conditions of the $(D_i)_{1 \leq i \leq n}$, and whose all prefixes are accepted by the automata $(C_i)_{1 \leq i \leq n}$. The size of T is exponential in G and the acceptance condition is a conjunction of Streett and safety objectives. The existence of such a path can be established in polynomial time in the size of T , so in exponential time in the size of G . The EXPTIME-hardness is a consequence of the EXPTIME-hardness of two-player games of imperfect information with safety [BD08], reachability and parity objectives [CDHR07]. \square

Doomsday Equilibria Synthesis from LTL Specifications and Imperfect

Information The doomsday equilibria synthesis problem can be extended to an imperfect setting. Remind that in the perfect information setting, LTL formulas are defined over a set of atomic propositions P , such that each player controls a subset P_i of P whose propositions are all visible to the other players. In the imperfect information setting, each set P_i is splitted into two sets P_i^{vis} and P_i^{inv}

of visible and invisible propositions, as defined for instance in [?]. We let $P^{inv} = \bigcup_i P_i^{inv}$ and $P^{vis} = \bigcup_i P_i^{vis}$. Note that $P = P^{inv} \cup P^{vis}$. Similarly to the perfect information setting, a strategy for Player i is a mapping $(2^P)^* \rightarrow 2^{P_i}$. However in the synthesis problem with imperfect information, we restrict strategies to observation-based strategies. The observation of $\sigma \in 2^P$ for Player i is defined by $\text{Obs}_i(\sigma) = \sigma \setminus (P^{inv} \setminus P_i^{inv}) \in 2^{P^{vis} \cup P_i}$, i.e. Player i observes only his propositions and the visible propositions of the other players. Observations are extended to finite and ω - words naturally. A strategy λ_i of Player i is *observation-based* if $\lambda_i(\pi) = \lambda_i(\pi')$ whenever $\text{Obs}_i(\pi) = \text{Obs}_i(\pi')$. Given a strategy $\lambda_i : (2^{P^{vis} \cup P_i})^* \rightarrow 2^{P_i}$, we define its (observation-based) concretization $\eta_i(\lambda_i) : (2^P)^* \rightarrow 2^{P_i}$ for all prefixes $\pi \in (2^P)^*$ by $\eta_i(\lambda_i)(\pi) = \lambda_i(\text{Obs}_i(\pi))$.

The *doomsday equilibria realizability problem with imperfect information* asks, given n LTL objectives $\varphi_1, \dots, \varphi_n$ over the atomic propositions P partitionned into $P_1^{vis}, P_1^{inv}, \dots, P_n^{vis}, P_n^{inv}$, whether there exists a doomsday equilibrium $\Lambda = (\lambda_1, \dots, \lambda_n)$ for the objectives $\varphi_1, \dots, \varphi_n$ such that each strategy λ_i is observation-based. The synthesis problem asks to automatically construct such a strategy profile if it exists.

The Safraless procedure for the perfect information setting can be extended to the imperfect information setting, as we describe now. Remind that in the perfect information setting, each objective $\varphi_i \cup \bigcap_j \overline{\varphi_j}$ is encoded as a universal co-Büchi automaton A_i . In the imperfect information setting, we additionally hide in A_i the propositions that are invisible for Player i , i.e. we construct a universal automaton $\text{Obs}_i(A_i)$ obtained as a copy of A_i where every transition (q, σ, q') is replaced by $(q, \text{Obs}_i(\sigma), q')$. The following proposition is immediate:

Proposition 2. *For all $i \in \{1, \dots, n\}$, for all $\rho \in (2^{P^{vis} \cup P_i})^*$, $\rho \in L(\text{Obs}_i(A_i))$ iff for all $\rho' \in (2^P)^*$ such that $\text{Obs}_i(\rho') = \rho$, we have $\rho' \in L(A_i)$.*

Then, we test the existence of a strategy profile Λ such that

1. for all $i \in \{1, \dots, n\}$, $A_i : (2^{P^{vis} \cup P_i}) \rightarrow 2^{P_i}$
2. for all $i \in \{1, \dots, n\}$, A_i realizes the objective $\text{Obs}_i(A_i)$ in the classical LTL two-player perfect information setting for the output propositions P_i and input propositions P^{vis} .
3. $\text{outcome}(\Lambda_1, \dots, \Lambda_n) \models \bigwedge_i \varphi_i$.

Proposition 3. *Λ satisfies (1), (2) and (3) iff $\eta(\Lambda) = (\eta_1(\Lambda_1), \dots, \eta_n(\Lambda_n))$ is an observation-based DE for $\varphi_1, \dots, \varphi_n$.*

Proof. Suppose that Λ satisfies 1–3. Clearly, $\eta(\Lambda)$ is winning for $\varphi_1, \dots, \varphi_n$. Moreover by definition of $\eta(\cdot)$, $\eta(\Lambda)$ is observation-based. Let us show that for all i , $\text{outcome}_i(\eta_i(\Lambda_i)) \subseteq L(A_i)$. By hypothesis $\text{outcome}_i(\Lambda_i) \subseteq L(\text{Obs}_i(A_i))$. Let $\rho \in \text{outcome}_i(\eta_i(\Lambda_i))$. We have $\text{Obs}_i(\rho) \in \text{outcome}_i(\Lambda_i)$ and by Proposition 2, we get that $\rho \in L(A_i)$.

Conversely, suppose that Λ' is an observation-based DE for $\varphi_1, \dots, \varphi_n$. For all i , define $A_i : (2^{P^{vis} \cup P_i})^* \rightarrow 2^{P_i}$ as $A'_i|_{(2^{P^{vis} \cup P_i})^*}$, and $\Lambda = (\Lambda_i)_{1 \leq i \leq n}$. Then clearly Λ satisfies (1) and (3). Let us show that it satisfies (2) too. By definition

of DE, for all i , $\text{outcome}_i(A'_i) \subseteq L(A_i)$. Let $\rho \in \text{outcome}_i(A_i)$. Clearly $\rho \in \text{outcome}_i(A'_i)$ and therefore $\rho \in L(A_i)$. It remains to show that $\rho \in L(\text{Obs}_i(A_i))$. By Proposition 2, it suffices to show that any $\rho' \in (2^P)^*$ such that $\text{Obs}_i(\rho') = \rho$ satisfies $\rho' \in L(A_i)$. It is indeed the case since all such ρ' are in $\text{outcome}_i(A'_i)$. Therefore A_i satisfies (2). \square

As for perfect information, it is possible to show that if a strategy profile A that satisfies 1–3 exists, then there exists one with finite memory, bounded by some constant K_φ that depends (doubly exponentially) on the size of $\varphi_1, \dots, \varphi_n$. As for perfect information, there is a Safraless approach: check (2) by solving safety games that depend on a parameter K (less than K_φ) as in [FJR11a], and checking (3) is a model-checking problem w.r.t. to the solution of the safety games (one has to search for an infinite path in the product of the safety games where unsafe edges have been removed). This leads to a 2EXPTIME procedure, which is not more than the perfect information setting. Moreover, optimization such as testing bounds K incrementally or antichains still applying in the imperfect information setting.

Theorem 4. *The DE-realizability problem from LTL objectives with imperfect information is 2EXPTIME-C.*

References

- [AHK02] R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49:672–713, 2002.
- [AT04] R. Alur and S. La Torre. Deterministic generators and games for LTL fragments. *TOCL*, 5, 2004.
- [ATM03] R. Alur, S. La Torre, and P. Madhusudan. Playing games with boxes and diamonds. In *CONCUR*, volume 2761 of *LNCS*, pages 127–141. Springer, 2003.
- [ATW06] C. Schulte Althoff, W. Thomas, and N. Wallmeier. Observations on determinization of Büchi automata. *Theor. Comput. Sci.*, 363(2):224–233, 2006.
- [BBF⁺12] A. Bohy, V. Bruyère, E. Filiot, N. Jin, and J.-F. Raskin. Acacia+, a tool for ltl synthesis. In *CAV*, pages 652–657, 2012.
- [BD08] D. Berwanger and L. Doyen. On the power of imperfect information. In *FSTTCS*, pages 73–82, 2008.
- [BFK] R. Ehlers B. Finkbeiner and A. Kupriyanov. Lecture on automata, games and verification. <http://iskren.info/reading/saar/agv/ps/sol05.pdf>.
- [BL69] J. R. Büchi and L. H. Landweber. Definability in the monadic second-order theory of successor. *J. Symb. Log.*, 34(2):166–170, 1969.
- [CDHR07] K. Chatterjee, L. Doyen, T.A. Henzinger, and J.-F. Raskin. Algorithms for omega-regular games with imperfect information. *LMCS*, 3(3), 2007.
- [CHJ06] K. Chatterjee, T. A. Henzinger, and M. Jurdzinski. Games with secure equilibria. *Theor. Comput. Sci.*, 365(1-2):67–82, 2006.
- [CHP07] K. Chatterjee, T.A. Henzinger, and N. Piterman. Generalized parity games. In *FoSSaCS*, volume 4423 of *LNCS*, pages 153–167. Springer, 2007.

- [CHP10] K. Chatterjee, T. A. Henzinger, and N. Piterman. Strategy logic. *Inf. Comput.*, 208(6):677–693, 2010.
- [CR12] K. Chatterjee and V. Raman. Synthesizing protocols for digital contract signing. In *VMCAI*, pages 152–168, 2012.
- [CZL09] Y. Cai, T. Zhang, and H. Luo. An improved lower bound for the complementation of rabin automata. In *LICS*, pages 167–176. IEEE Computer Society, 2009.
- [DLM10] A. Da Costa Lopes, F. Laroussinie, and N. Markey. ATL with strategy contexts: Expressiveness and model checking. In *FSTTCS*, volume 8 of *LIPICs*, pages 120–132, 2010.
- [Ehl11] R. Ehlers. Unbeast: Symbolic bounded synthesis. In *TACAS*, pages 272–275, 2011.
- [EJ91] E. A. Emerson and C. Jutla. Tree automata, mu-calculus and determinacy. In *FOCS*, pages 368–377. IEEE Comp. Soc., 1991.
- [EL85] E.A. Emerson and C.-L. Lei. Modalities for model checking: Branching time strikes back. In *POPL*, pages 84–96, 1985.
- [FH10] N. Fijalkow and F. Horn. The surprising complexity of reachability games. *CoRR*, abs/1010.2420, 2010.
- [FJR11a] E. Filiot, N. Jin, and J.-F. Raskin. Antichains and compositional algorithms for ltl synthesis. In *Journal of Formal Methods in System Design*. Springer, 2011.
- [FJR11b] E. Filiot, N. Jin, and J.-F. Raskin. Antichains and compositional algorithms for LTL synthesis. *Formal Methods in System Design*, 39(3):261–296, 2011.
- [FKL10] D. Fisman, O. Kupferman, and Y. Lustig. Rational synthesis. In *Proc. of TACAS*, LNCS 6015, pages 190–204. Springer, 2010.
- [GTW02] E. Grädel, W. Thomas, and T. Wilke, editors. *Automata, Logics, and Infinite Games: A Guide to Current Research*, LNCS 2500. Springer, 2002.
- [Imm81] N. Immerman. Number of quantifiers is better than number of tape cells. *Journal of Computer and System Sciences*, 22:384–406, 1981.
- [JB06] B. Jobstmann and R. Bloem. Optimizations for LTL synthesis. In *Formal Methods in Computer-Aided Design (FMCAD)*, pages 117–124. IEEE Computer Society, 2006.
- [JMM12] W. Jamroga, S. Mauw, and M. Melissen. Fairness in non-repudiation protocols. In *Proc. of STM*, LNCS 7170, pages 122–139. Springer, 2012.
- [KV05] O. Kupferman and M. Y. Vardi. Safraless decision procedures. In *FOCS*, 2005.
- [Mar75] D. Martin. Borel determinacy. In *Annals of Mathematics*, volume 102, pages 363–371, 1975.
- [MGK02] O. Markowitch, D. Gollmann, and S. Kremer. On fairness in exchange protocols. In *ICISC*, pages 451–464, 2002.
- [MMPV12] F. Mogavero, A. Murano, G. Perelli, and M.Y. Vardi. What makes at1* decidable? a decidable fragment of strategy logic. In *CONCUR*, pages 193–208, 2012.
- [MMV10] F. Mogavero, A. Murano, and M. Y. Vardi. Reasoning about strategies. In *Proc. of FSTTCS*, volume 8 of *LIPICs*, pages 133–144. Schloss Dagstuhl - LZfI, 2010.
- [Nas50] J. F. Nash. Equilibrium points in n -person games. *PNAS*, 36:48–49, 1950.
- [Pit07] N. Piterman. From nondeterministic Büchi and streett automata to deterministic parity automata. *Logical Methods in Computer Science*, 3(3), 2007.

- [PP06] N. Piterman and A. Pnueli. Faster solutions of rabin and streett games. In *LICS*, pages 275–284, 2006.
- [PR89a] A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *POPL*, pages 179–190. ACM Press, 1989.
- [PR89b] A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *ACM Symposium on Principles of Programming Languages (POPL)*. ACM, 1989.
- [Rab69] M. O. Rabin. Decidability of second-order theories and automata on infinite trees. *Trans. Amer. Math. soc.*, 141:1–35, 1969.
- [RW87] P. J. Ramadge and W. M. Wonham. Supervisory control of a class of discrete event processes. *SIAM Journal on Control and Optimization*, 25(1):206–230, 1987.
- [SF07] S. Schewe and B. Finkbeiner. Bounded synthesis. In *ATVA*, volume 4762 of *LNCS*, pages 474–488. Springer-Verlag, 2007.
- [Sha53] L. S. Shapley. Stochastic games. *PNAS*, 39:1095–1100, 1953.
- [UW11] M. Ummels and D. Wojtczak. The complexity of nash equilibria in stochastic multiplayer games. *Logical Methods in Computer Science*, 7(3), 2011.
- [WHY11] F. Wang, C. Huang, and F. Yu. A temporal logic for the interaction of strategies. In *Proc. of CONCUR*, LNCS 6901, pages 466–481. Springer, 2011.

A Complexity of Doomsday Equilibria for Imperfect Information Games

We now present automata construction to recognize sequences of observations that are doomsday compatible and prefixes that are good for retaliation.

Lemma 15. *Given an n -player game G with imperfect information and a set of reachability, safety or parity objectives $(\varphi_i)_{1 \leq i \leq n}$, we can construct for each Player i , in exponential time, a deterministic Streett automaton D_i whose language is exactly the set of sequences of observations $\eta \in (O_i \times (\Sigma \cup \{\tau\}))^\omega$ that are doomsday compatible for Player i , i.e.*

$$L(D_i) = \{\eta \in (O_i \times (\Sigma \cup \{\tau\}))^\omega \mid \forall \rho \in \gamma_i(\eta) \cdot \rho \in \varphi_i \cup \bigcap_{j \neq i} \overline{\varphi_j}\}.$$

For each D_i , the size of its set of states is bounded by $\mathbf{O}(2^{nk \log k})$ and the number of Streett pairs is bounded by $\mathbf{O}(nk^2)$ where k is the number of states in G .

Proof. Let $G = (S, (S_i)_{1 \leq i \leq n}, s_{\text{init}}, \Sigma, \Delta, (O_i)_{1 \leq i \leq n})$, and let us show the constructions for Player i , $1 \leq i \leq n$. We treat the three types of winning conditions as follows.

We start with safety objectives. Assume that the safety objectives are defined implicitly by the following tuple of sets of safe states: (T_1, T_2, \dots, T_n) , i.e. $\varphi_i = \text{safe}(T_i)$. First, we construct the automaton

$$A = (Q^A, q_{\text{init}}^A, (O_i \times (\Sigma \cup \{\tau\})), \delta^A)$$

over the alphabet $O_i \times (\Sigma \cup \{\tau\})$ as follows:

- $Q^A = S$, i.e. the states of A are the states of the game structure G ,
- $q_{\text{init}}^A = s_{\text{init}}$,
- $(q, (o, \sigma), q') \in \delta^A$ if $q \in o$ and there exists $\sigma' \in \Sigma$ such that $\Delta(q, \sigma') = q'$ and such that $\sigma = \tau$ if $q \notin S_i$, and $\sigma = \sigma'$ if $q \in S_i$.

The acceptance condition of A is *universal* and expressed with LTL syntax:

A word w is accepted by A iff *all* runs ρ of A on w satisfy $\rho \models \Box T_i \vee \bigwedge_{j \neq i} \Diamond \overline{T_j}$.

Clearly, the language defined by A is exactly the set of sequences of observations $\eta \in (O_i \times (\Sigma \cup \{\tau\}))^\omega$ that are *doomsday compatible* for Player i , this is because the automaton A checks (using universal nondeterminism) that all plays that are compatible with a sequence of observations are doomsday compatible.

Let us show that we can construct a deterministic Streett automaton D_i that accepts the language of A and whose size is such that: (i) its number of states is at most $\mathbf{O}(2^{(nk \log k)})$ and (ii) its number of Streett pairs is at most $\mathbf{O}(nk)$. We obtain D with the following series of constructions:

- First, note that we can equivalently see A as the intersection of the languages of $n - 1$ universal automata A_j with the acceptance condition $\Box T_i \vee \Diamond \overline{T_j}$, $j \neq i$, $1 \leq j \leq n$.

- Each A_j can be modified so that a violation of T_i is made permanent and a visit to $\overline{T_j}$ is recorded. For this, we use a state space which is equal to $Q^A \times \{0, 1\} \times \{0, 1\}$, the first bit records a visit to $\overline{T_i}$ and the second a visit to $\overline{T_j}$. We denote this automaton by A'_j , and its acceptance condition is now $\Box\Diamond(Q^A \times \{0, 1\} \times \{0\}) \rightarrow \Box\Diamond(Q^A \times \{0\} \times \{0, 1\})$. Clearly, this is a universal Streett automaton with a single Streett pair.
- A'_j , which is a universal Streett automaton, can be complemented (by duality) by interpreting it as a nondeterministic Rabin automaton (with one Rabin pair). This nondeterministic Rabin automaton can be made deterministic using a Safra like procedure, and according to [CZL09] we obtain a deterministic Rabin automaton with $\mathbf{O}(2^{k \log k})$ states and $\mathbf{O}(k)$ Rabin pairs. Let us call this automaton A''_j .
- Now, A''_j can be complemented by considering its Rabin pairs as Streett pairs (by dualization of the acceptance condition): we obtain a deterministic Streett automaton with $\mathbf{O}(k)$ Streett pairs for each A_j .
- Now, we need to take the intersection of the $n - 1$ deterministic automata A''_j (interpreted as Streett automata). Using a classical synchronized product we obtain a single deterministic Streett automaton D_i of size with $\mathbf{O}(2^{nk \log k})$ states and $\mathbf{O}(nk)$ Streett pairs. This finishes our proof for safety objectives.

Let us now consider reachability objectives. Therefore we now assume the states in T_1, \dots, T_n to be target states for each player respectively, i.e. $\varphi_i = \text{reach}(T_i)$. The construction is in the same spirit as the construction for safety. Let $A = (Q^A, q_{\text{init}}^A, O_i \times (\Sigma \cup \{\tau\}), \delta^A)$ be the automaton over $(O_i \times (\Sigma \cup \{\tau\}))$ constructed from G as for safety, with the following (universal) acceptance condition;

A word w is accepted by A iff all runs ρ of A on w satisfy $\rho \models (\bigvee_{j \neq i} \Diamond T_j) \rightarrow \Diamond T_i$.

Clearly, the language defined by A is exactly the set of sequences of observations $\eta \in ((\Sigma \cup \{\tau\}) \times O_i)^\omega$ that are *doomsday compatible* for Player i (w.r.t. the reachability objectives). Let us show that we can construct a deterministic Streett automaton D_i that accepts the language of A and whose size is such that: (i) its number of states is at most $\mathbf{O}(2^{nk \log k})$ and (ii) its number of Streett pairs is at most $\mathbf{O}(nk)$. We obtain D_i with the following series of constructions:

- First, the acceptance condition can be rewritten as $\bigwedge_{j \neq i} (\Diamond T_j \rightarrow \Diamond T_i)$. Then clearly if A_j is a copy of A with acceptance condition $\Diamond T_j \rightarrow \Diamond T_i$ then $L(A) = \bigcap_{j \neq i} L(A_j)$.
- For each A_j , we construct a universal Streett automaton with one Streett pair by memorizing the visits to T_i and T_j and considering the acceptance condition $\Box\Diamond T_j \rightarrow \Box\Diamond T_i$. So, we get a universal automaton with a single Streett pair.
- Then we follow exactly the last three steps (3 to 5) of the construction for safety.

Finally, let us consider parity objectives. The construction is similar to the other cases. Specifically, we can take as acceptance condition for A the universal

condition $\bigwedge_{j \neq i} (\text{parity}_i \vee \overline{\text{parity}_j})$, and treat each condition $\text{parity}_i \vee \overline{\text{parity}_j}$ separately. We dualize the acceptance condition of A , into the nondeterministic condition $\overline{\text{parity}_i} \wedge \text{parity}_j$. This acceptance condition can be equivalently expressed as a Streett condition with at most $\mathbf{O}(k)$ Streett pairs. This automaton accepts exactly the set of observation sequences that are not doomsday compatible for Player i against Player j . Now, using optimal procedure for determinization, we can obtain a deterministic Rabin automaton, with $\mathbf{O}(k^2)$ pairs that accepts the same language [Pit07]. Now, by interpreting the pairs of the acceptance condition as Streett pairs instead of Rabin pairs, we obtain a deterministic Streett automaton A_j that accepts the set of observations sequences that are doomsday compatible for Player i against Player j . Now, it suffices to take the product of the $n - 1$ deterministic Streett automata A_j to obtain the desired automaton A , its size is at most $\mathbf{O}(2^{nk \log k})$ with at most $\mathbf{O}(nk^2)$ Streett pairs. \square

Lemma 16. *Given an n -player game arena G with imperfect information and a set of reachability, safety or parity objectives $(\varphi_i)_{1 \leq i \leq n}$, for each Player i , we can construct a finite-state automaton C_i that accepts exactly the prefixes of observation sequences that are good for retaliation for Player i .*

Proof. Let us show how to construct this finite-state automaton for any Player i , $1 \leq i \leq n$. Our construction follows these steps:

- First, we construct from G , according to lemma 15, a deterministic Streett automaton $D_i = (Q^{D_i}, q_{\text{init}}^{D_i}, (O_i \times (\Sigma \cup \{\tau\}), \delta^{D_i}, \text{St}^{D_i})$ that accepts exactly the set of sequences of observations $\eta \in (O_i \times (\Sigma \cup \{\tau\}))^\omega$ that are *doomsday compatible* for Player i . We know that the number of states in D_i is $\mathbf{O}(2^{|S|^2 \log |S|})$ and the number of Streett pairs is bounded by $\mathbf{O}(|S|^2 \cdot n)$, where $|S|$ is the number of states in G .
- Second, we consider a turn-based game played on D_i by two players, A and B, that move a token from states to states along edges of D_i as follows:
 1. initially, the token is in some state q
 2. then in each round: B chooses an observation $o \in O_i$ in the set $\{o \in O_i \mid \exists(q, (o, \sigma), q') \in \delta^{D_i}\}$. Then A chooses a transition $(q, (o, \sigma), q') \in \delta^{D_i}$ (which is completely determined by σ as D_i is deterministic), and the token is moved to q' where a new round starts.

The objective of A is to enforce from state q an infinite sequence of states, so a run of D_i that starts in q , and which satisfies St^{D_i} the Streett condition of D_i . For each q , this can be decided in time polynomial in the number of states in D_i and exponential in the number of Streett pairs in St^{D_i} , see [PP06] for an algorithm with the best known complexity. Thus, the overall complexity is exponential in the size of the game structure G . We denote by $\text{Win} \subseteq Q^{D_i}$ the set of states q from which A can win the game above.

- Note that if $(o_1, \sigma_1) \dots (o_m, \sigma_m)$ is the trace of a path from q_{init} in D_i to a state $q \in \text{Win}$, then clearly $(o_1, \sigma_1) \dots (o_{n-1}, \sigma_{n-1})o_n$ is good for retaliation. Indeed, the winning strategy of A in q is an observation based retaliating strategy λ_i^R for Player i in G . On the other hand, if a prefix of observations reaches $q \notin \text{Win}$ then by determinacy of Streett games, we know that B has a

winning strategy in q and this winning strategy is a strategy for the coalition (against Player i) in G to enforce a play in which Player i does not win and at least one of the other players wins. So, from D_i and Win , we can construct a finite state automaton C_i which is obtained as a copy of D_i with the following acceptance condition: a prefix $\kappa = (o_0, \sigma_0), (o_1, \sigma_1), \dots, (o_{k-1}, \sigma_{k-1}), o_k$ is accepted by C_i if there exists a path $q_0 q_1 \dots q_k$ in C_i such that q_0 is the initial state of C_i and either there exists a transition labeled (o_k, σ) from q_k to a state of Win . \square