Doomsday Equilibria for Omega-Regular Games

Krishnendu Chatterjee[1]*, Laurent Doyen[2], Emmanuel Filiot[3]**, and
Jean-François Raskin[3]***

[1]IST Austria, [2]LSV, ENS Cachan & CNRS, [3]CS-Université Libre de Bruxelles – U.L.B.

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, and 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 [31, 29, 14]. One key application area is synthesis of reactive systems [5, 30, 28]. 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 have been studied in great depth in literature [22, 14, 17].

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

---

* Supported by Austrian Science Fund (FWF) Grant No P23499-N23, FWF NFN Grant No S11407-N23 (RiSE), ERC Start grant (279307: Graph Games), and Microsoft faculty fellows award.

** Supported by the Belgian National Fund for Scientific Research.

*** Supported by ERC Start grant (279499: inVEST).

on graphs has become an active area of research in graph games and reactive synthesis [1, 16, 32]. 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 [25]. 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 (internal criteria), 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 [10] (that captures the adversarial external criteria), rational synthesis [16], and led to several new logics where the non-zero-sum equilibria can be expressed [11, 13, 24, 33, 23]. The complexity of Nash equilibria [32], secure equilibria [10], rational synthesis [16], and of the new logics has been studied recently [11, 13, 24, 33].

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 [10], and of fair-exchange protocols [20, 7] (a key protocol in the area of security for exchange of digital signatures). 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 [28]. However, the model of imperfect-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 [19]. 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 equilibria 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 [10] where objectives of both players are satisfied.

Example 1. Let us consider the two trees of Fig. 1. They model the possible behaviors of two entities Alice and Bob that have the objective of exchanging messages: m_{AB} from Alice to Bob, and m_{BA} from Bob to Alice. Assume for the sake of illustration that m_{AB} models the transfer of property of an house from Alice to Bob, while m_{BA} models the payment of the price of the house from Bob to Alice.

Having that interpretation in mind, let us consider the left tree. On the one hand, Alice has as primary objective (internal criterion) to reach either state 2 or state 4, states in which she has obtained the money, and she has a slight preference for 2 as in that case she received the money while not transferring the property of her house

Fig. 1. A simple example in the domain of Fair Exchange Protocols

to Bob, this corresponds to her adversarial external criterion. On the other hand, Bob would like to reach either state 3 or 4 (with again a slight preference for 3). Also, it should be clear that Alice would hate to reach 3 because she would have transferred the property of her house to Bob but without being paid. Similarly, Bob would hate to reach 2. To summarize, Alice has the following preference order on the final states of the protocol: $2 > 4 > 1 > 3$, while for Bob the order is $3 > 4 > 1 > 2$. Is there a *doomsday threatening equilibrium* in this game ? For such an equilibrium to exist, we must find a pair of strategies that please the two players for their primary objective (internal criterion): reach $\{2, 4\}$ for Alice and reach $\{3, 4\}$ for Bob. Clearly, this is only possible if at the root Alice plays ”send m_{AB}”, as otherwise we would not reach $\{1, 2\}$ violating the primary objective of Bob. But playing that action is not safe for Alice as Bob would then choose ”not send m_{BA}” because he slightly prefers 3 to 4. It can be shown that the only rational way of playing (taking into account both internal and external criteria) is for Alice to play ”not send m_{AB}” and for Bob would to play ”not send m_{BA}”. This way of playing is in fact the only secure equilibrium of the game but this is not what we hope from such a protocol.

The difficulty in this exchange of messages comes from the fact that Alice is starting the protocol by sending her part and this exposes her. To obtain a better behaving protocol, one solution is to add an additional stage after the exchanges of the two messages as depicted in the right tree of Fig. 1. In this new protocol, Alice has the possibility to cancel the exchange of messages (in practice this would be implemented by the intervention of a TTP[1]). For that new game, the preference orderings of the players are as follows: for Alice it is $3 > 7 > 1 = 2 = 4 = 6 = 8 > 5$, and for Bod it is $5 > 7 > 1 = 2 = 4 = 6 = 8 > 3$. Let us now show that there is a doomsday equilibrium in this new game. In the first round, Alice should play ”send m_{AB}” as otherwise the internal objective of Bob would be violated, then Bob should play ”send m_{BA}”, and finally Alice should play “OK” to validate the exchange of messages. Clearly, this profile of strategies satisfies the first property of a doomsday equilibrium: both players have reached their primary objective. Second, let us show that no player has an incentive to deviate from that profile of strategies. First, if Alice deviates then Bob would play ”not send m_{BA}”, and we obtain a doomsday situation as both players have their primary objectives violated. Second, if Bob deviates by playing ”not send m_{BA}”, then Alice would cancel the protocol exchange which again produces a doomsday situation.

---

[1] TTP stands for *Trusted Third Party*.

objectives	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

So, no player has an incentive to deviate from the equilibrium and the outcome of the protocol is the desired one: the two messages have been fairly exchanged. So, we see that the threat of a doomsday that brings the action ”Cancel” has a beneficial influence on the behavior of the two players. □

It should now be clear that multi-player games with doomsday equilibria provide a suitable framework to model various problems in protocol synthesis. In addition to the definition of doomsday equilibria, our main contributions are to present algorithms and complexity bounds for deciding the existence of such equilibria for various classes of ω-regular objectives both in the perfect-information and in the imperfect-information cases. In all cases but one, we establish the exact complexity. Our technical contributions are 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 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.

The area of multi-player games and various notion of equilibria is an active area of research, but notions that lead to decidability in the imperfect-information setting and has applications in synthesis has largely been an unexplored area. Our work is a step towards it.

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 $(S, \mathcal{P}, s_{\mathsf{init}}, \Sigma, \Delta)$ 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, $s_{\mathsf{init}} \in S$ is the initial state, Σ is a finite set of actions, and $\Delta : S \times \Sigma \to S$ 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) $s_0 = s_{\mathsf{init}}$ and (i) 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 $\mathsf{Plays}(G)$, and the set of finite prefixes of plays by $\mathsf{PrefPlays}(G)$. 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, ... to denote prefixes of plays. Let $i \in \{1, 2, \dots, n\}$, a prefix π belongs to Player i if $\mathsf{last}(\pi)$, the last state of π, belongs to Player i, i.e. $\mathsf{last}(\pi) \in S_i$. We denote by $\mathsf{PrefPlays}_i(G)$ 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 $\lambda_i : \mathsf{PrefPlays}_i(G) \to \Sigma$ 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 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, $\mathsf{outcome}_i(\lambda_i)$ is the set of $\rho \in \mathsf{Plays}(G)$ such that for all $j \geq 0$, if $\rho(0..j) \in \mathsf{PrefPlays}_i(G)$, 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 $\rho \in \mathsf{Plays}(G)$ such that for all positions j, for all $i \in \{1, 2, \dots, n\}$, if $\rho(j) \in \mathsf{PrefPlays}_i(G)$ 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 $\mathsf{outcome}_i(\lambda_i) \subseteq \varphi_i$.

Given an infinite sequence of states $\rho \in S^\omega$, we denote by $\mathsf{visit}(\rho)$ the set of states that appear at least once along ρ, i.e. $\mathsf{visit}(\rho) = \{s \in S | \exists i \geq 0 \cdot \rho(i) = s\}$, and $\mathsf{inf}(\rho)$ the set of states that appear infinitely often along ρ, i.e. $\mathsf{inf}(\rho) = \{s \in S | \forall i \geq 0 \cdot \exists j \geq i \cdot \rho(i) = s\}$. 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: $\mathsf{safe}(T) = \{\rho \in S^\omega \mid \mathsf{visit}(\rho) \subseteq T\}$;
- a *reachability objective* is defined by a subset of states $T \subseteq S$ that has to be reached: $\mathsf{reach}(T) = \{\rho \in S^\omega \mid \mathsf{visit}(\rho) \cap T \neq \emptyset\}$;
- a *Büchi objective* is defined by a subset of states $T \subseteq S$ that has to be visited infinitely often: $\mathsf{Büchi}(T) = \{\rho \in S^\omega \mid \mathsf{inf}(\rho) \cap T \neq \emptyset\}$;
- a *co-Büchi objective* is defined by a subset of states $T \subseteq S$ that has to be reached eventually and never be left: $\mathsf{coBüchi}(T) = \{\rho \in S^\omega \mid \mathsf{inf}(\rho) \subseteq T\}$;
- let $d \in \mathbb{N}$, a *parity objective with d priorities* is defined by a priority function $p : S \to \{0, 1, \dots, d\}$ as the set of plays such that the smallest priority visited infinitely often is even: $\mathsf{parity}(p) = \{\rho \in S^\omega | \min\{p(s) \mid s \in \mathsf{inf}(\rho)\}$ 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 $\bigcap_{i \in P} \mathsf{outcome}_i(\lambda_i) \subseteq \varphi$.

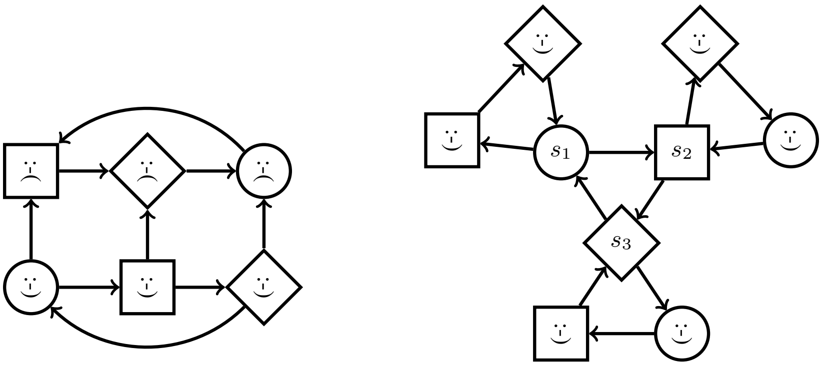

(a) Doomsday (Safety)

(b) Büchi objectives

Fig. 2. Examples of doomsday equilibria for Safety and Büchi objectives

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

1. it is winning for all the players, i.e. $\mathsf{outcome}(\Lambda) \in \bigcap_i \varphi_i$;
2. each player is able to retaliate in case of deviation: for all $1 \leq i \leq n$, for all $\rho \in \mathsf{outcome}_i(\lambda_i)$, 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 retaliates 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 [10] 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.

Example 2. Fig. 2 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. 2(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. 2(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 $i\ mod\ 3 + 1$ and $i\ mod\ 3 + 2$. 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 $s_{(i\ mod\ 3)+1}$. 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 $s_{(i\ mod\ 3)+1}$. 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 in the future to go to the next s state instead of taking their respective loops. 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 retialate 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 doosmday equilibrium in this example. $\square$

3 Complexity of DE for Perfect Information Games

In this section, we prove the following results:

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, hardness already holds for 2-player game arenas,*
- NP-HARD, CONP-HARD *and in* PSPACE *if $(\varphi_i)_{1 \leq i \leq n}$ are parity objectives, hardness already holds for 2-player game arenas,*
- PSPACE-C *if $(\varphi_i)_{1 \leq i \leq n}$ are safety objectives, and* PTIME-C *for game arenas with a fixed number of players.*

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 $G = (S, \mathcal{P}, s_{\mathsf{init}}, \Sigma, \Delta)$ be an n-player game arena with n tail objectives $(\varphi_i)_{1 \le i \le n}$. 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) belongs to $\bigcap_{i=1}^{i=n} \varphi_i$ and (2) stays within the set of states $\bigcap_{i=1}^{i=n} R_i$.*

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 $\Lambda = (\lambda_1, \lambda_2, \ldots, \lambda_n)$. 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}) $\mathsf{outcome}_i(\lambda_i) \subseteq \varphi_i \cup \bigcap_{j=1}^{j=n} \overline{\varphi_j}$. 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 $\rho = \mathsf{outcome}(\Lambda)$ 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, $\mathsf{outcome}(\Lambda_i) \subseteq \varphi_i \cup \bigcap_{j=1}^{j=n} \overline{\varphi_j}$. 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. □

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 $i \in \{1, \ldots, n\}$. 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. [27]. 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 presented in [15]. $\Box$

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 2-player game arenas.*

Proof. The hardness for Büchi and co-Büchi objectives holds already for 2 players. We describe the reduction for Büchi and it is similar for co-Büchi. We reduce the problem of deciding the winner in a two-player zero-sum game arena G with a Büchi objective (known as a PTime-Hard problem [18]) to the existence of a DE for Büchi objectives with two players. Consider a copy G' of the game arena G and the following two objectives: Player 1 has the same Büchi objective as Player 1 in G, and Player 2 has a trivial Büchi objective (i.e. all states are Büchi states). Then clearly there exists a DE in G' iff Player 1 has a winning strategy in G. Details are given in the long version of this paper [8].

For parity games, we can reduce zero-sum two-player games with a conjunction of parity objectives (known to be coNP-Hard [12]) to the existence of a DE in a three player game with parity objectives. Similarly, we can reduce the problem of deciding the winner in a two-player zero-sum game with a disjunction of parity objectives (known to be NP-Hard [12]) to the existence of a DE in a two-player game with parity objectives. The main idea in the two cases is to construct a game arena where one of the players can retaliate iff Player 1 in the original two-player zero-sum game has a winning strategy. Details are given in the long version of this paper [8]. $\Box$

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 $G = (S, \mathcal{P}, s_{\mathsf{init}}, \Sigma, \Delta)$ be a game arena, and $(T_i)_{1 \le i \le n}$ be n subsets of S. Let Λ be a doomsday equilibrium in G for the reachability objectives $(\mathsf{Reach}(T_i))_{1 \le i \le n}$. Let s the first state in $\mathsf{outcome}(\Lambda)$ 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, as no strategy from s to reach his target set T_2, then necessarily $s \notin T_2$ and by determinancy 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. ☐

Lemma 4. *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 $\Diamond T_i \vee (\bigwedge_{j=1}^{j=n} \Box \overline{T_j})$. 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 $\mathsf{outcome}_i(\lambda_i) \models \Box((\bigcap_{j=1}^{j=n} \overline{T_j}) \vee \langle\langle i \rangle\rangle \Diamond T_i)\}$, 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 $\bigcap_{j=1}^{j=n} R_j$ before reaching a state s such that $s \in T_i$ and $s \in \bigcap_{j=1}^{j=n} \langle\langle j \rangle\rangle \Diamond T_j$.

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 $\bigcap_{j=1}^{j=n} R_j$. The other direction follows from Proposition 1. ☐

Lemma 5. *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 [18]: 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. ☐

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 6 (EASYNESS). *The existence of a doomsday equilibrium in an n-player game with safety objectives can be decided in* PSPACE*, and in* PTIME *for game arenas with a fixed number of players.*

Proof. We start with the general case where the number of players is not fixed and is part of the input. Let us consider an n-player game arena $G = (S, \mathcal{P}, s_{\mathsf{init}}, \Sigma, \Delta)$ and n safety objectives $\mathsf{safe}(T_1), \ldots, \mathsf{safe}(T_n)$ for $T_1 \subseteq S, \ldots, T_n \subseteq S$. 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 $\mathsf{outcome}_i(\lambda_i)$ 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 [2]) 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 $\bigcap_{i=1}^{i=n}(\mathsf{safe}(T_i) \cap R_i)$.

Now, let us establish the correctness of this algorithm. Assume that an infinite path exists in $\bigcap_{i=1}^{i=n}(\mathsf{safe}(T_i) \cap R_i)$. 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 $\Lambda = (\lambda_1, \lambda_2, \ldots, \lambda_n)$ 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 $\Lambda = (\lambda_1, \lambda_2, \ldots, \lambda_n)$ is a DE in the n-player game G for the safety objectives $(\mathsf{safe}(T_i))_{1 \leq i \leq n}$. Let $\rho = \mathsf{outcome}(\lambda_1, \lambda_2, \ldots, \lambda_n)$. 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, $i \in \{1, \ldots, n\}$, and π the finite prefix of ρ up to s. By definition of DE we have $\mathsf{outcome}(\lambda_i) \models \Box T_i \vee \bigwedge_{j=1}^{j=n} \Diamond\overline{T_j}$. 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 $i \in \{1, \ldots, n\}$, we get $s \in \bigcap_{i=1}^{i=n} R_i$.

Let us now turn to the case where the number of players is fixed. Then clearly, in the construction above, all the LTL formulas are of fixed size and so all the associated games can then be solved in polynomial time. □

Lemma 7 (HARDNESS). *The problem of deciding the existence of a doomsday equilibrium in an n-player game with safety objectives is* PSPACE-HARD*, and* PTIME-HARD *when the number of players is fixed.*

Proof. For the general case, we present a reduction from the problem of deciding the winner in a zero-sum two-player game with a conjunction of k reachability objectives (aka generalized reachability games), which is a PSPACE-C problem [3]. The idea of the

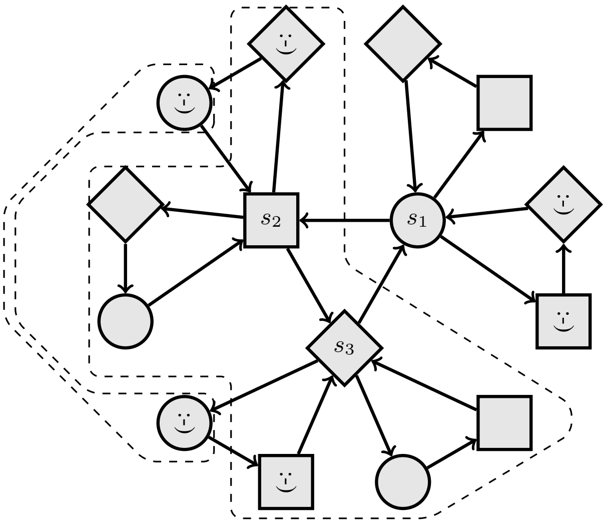

Fig. 3. Game arena with imperfect information and Büchi objectives. Only undistinguishable states of Player 1 (circle) are depicted. Observations are symmetric for the other players.

reduction is to construct a non-zero sum $(k + 1)$-player game where one of the players has a retaliating strategy iff there is a winning strategy in the generalized reachability game.

When the number of players is fixed, PTIME-HARDNESS is proved by an easy reduction from the And-Or graph reachability problem [18]. □

4 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 $G = (S, \mathcal{P}, s_{\mathsf{init}}, \Sigma, \Delta, (O_i)_{1 \leq i \leq n})$ such that $(S, \mathcal{P}, s_{\mathsf{init}}, \Sigma, \Delta)$ 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 order[2]: for all $i \in \{1, \ldots, n\}$, all $q \in S_i$ and all $\sigma \in \Sigma$, $\Delta(q, \sigma) \in S_{(i \ mod \ n)+1}$.

---

[2] This restriction is not necessary to obtain the results presented in this section (e.g. Theorem 2) but it makes some of our notations lighter.

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 $s_0 = s_{\mathsf{init}}$, 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 $\mathsf{Plays}(G)$ and $\mathsf{PrefPlays}(G)$ to denote the set of plays in G and its set of prefixes, and $\mathsf{PrefPlays}_i(G)$ 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 $\mathsf{Obs}_i(\rho) \in (O_i \times (\Sigma \cup \{\tau\})^\omega$ such that for all $j \geq 0$, $\mathsf{Obs}_i(\rho)(j) = (O_i(s_j), \tau)$ if $s_j \notin S_i$, and $\mathsf{Obs}_i(\rho)(j) = (O_i(s_j), \sigma_j)$ 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 $\mathsf{Obs}_i(\pi)$ 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. $\gamma_i(\eta) = \{\rho \in \mathsf{Plays}(G) \mid \mathsf{Obs}_i(\rho) = \eta\}$. 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 $\pi_1, \pi_2 \in \mathsf{PrefPlays}_i(G)$ such that $\mathsf{Obs}_i(\pi_1) = \mathsf{Obs}_i(\pi_2)$, 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 $\pi = s_0, \sigma_0, \dots, s_k \in \mathsf{PrefPlays}_i(G)$, the strategy λ_i^π is defined for all prefixes $\pi' \in \mathsf{PrefPlays}_i(G_{s_k})$ 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 $\mathsf{outcome}_i(\pi, \lambda_i) = \pi \cdot \mathsf{outcome}_i(\lambda_i^\pi)$.

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 3. Fig. 3 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 (i.e. does not visit the happy states). 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 $G = (S, \mathcal{P}, s_{\mathsf{init}}, \Sigma, \Delta, (O_i)_{1 \leq i \leq n})$ 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, $\mathsf{outcome}(\pi, \lambda_i^R) \subseteq \varphi_i \cup \bigcap_{j=1}^{j=n} \overline{\varphi_j}$.

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 8. *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:*

(F_1) $\rho \in \bigcap_{i=1}^{i=n} \varphi_i$, *i.e. ρ is winning for all the players,*
(F_2) *for all Player i, $1 \leq i \leq n$, for all prefixes κ of $\mathsf{Obs}_i(\rho)$, κ is good for retaliation for Player i,*
(F_3) *for all Player i, $1 \leq i \leq n$, $\mathsf{Obs}_i(\rho)$ is doomsday compatible for Player i.*

Proof. First, assume that conditions (F_1), (F_2) and (F_3) hold and show that there exists a DE in G. We construct a DE $(\lambda_1, \ldots, \lambda_n)$ 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 (F_2)). 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 (F_3) we know that if the outcome does not satisfy φ_i, then it does not satisfies 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 $\kappa = \kappa_1(o_1, \sigma_1)\dots(o_m, \sigma_n m$ 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 (F_2), 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 $\mathsf{obs}_i(\rho)$ is not doomsday compatible for Player i, then by definition, there is a path ρ' in $\mathsf{Plays}(G)$ that is compatible with the observations and actions of player i in ρ (i.e., $\mathsf{obs}_i(\rho) = \mathsf{obs}_i(\rho')$), 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 $\mathsf{obs}_i(\rho)$ 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 $\mathsf{obs}_i(\rho)$ 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. ☐

Theorem 2. *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, all Büchi, all co-Büchi or all parity objectives. Hardness already holds for 2-player game arenas.*

Proof. By Lemma 8, 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 for all the considered objectives [4, 9]. ☐

5 Conclusion

We defined the notion of doomsday threatening equilibria both for perfect and imperfect information n player games with omega-regular objectives. This notion generalizes to n player games the winning secure equilibria [10]. Applications in the analysis of security protocols are envisioned and will be pursued as future works.

We have settled the exact complexity in games of perfect information for almost all omega-regular objectives with complexities ranging from PTIME to PSPACE, the only small gap that remains is for parity objectives where we have a PSPACE algorithm and both NP and coNP-hardness. Surprisingly, the existence of doomsday threatening equilibria in n player games with imperfect information is decidable and more precisely EXPTIME-C for all the objectives that we have considered.

In a long version of this paper [8], we provide a solution in 2EXPTIME for deciding the existence of a doomsday threatening equilibrium in a game whose objectives are given as LTL formula (this solution is optimal as it is easy to show that the classical LTL realizability problem can be reduced to the DE existence problem). We also provide a Safraless solution [21] suitable to efficient implementation.

References

1. R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49:672–713, 2002.
2. R. Alur and S. La Torre. Deterministic generators and games for LTL fragments. *TOCL*, 5, 2004.
3. 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.
4. D. Berwanger and L. Doyen. On the power of imperfect information. In *FSTTCS*, pages 73–82, 2008.
5. J. R. Büchi and L. H. Landweber. Definability in the monadic second-order theory of successor. *J. Symb. Log.*, 34(2):166–170, 1969.
6. 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.
7. R. Chadha, S. Kremer, and A. Scedrov. Formal analysis of multiparty contract signing. *J. Autom. Reasoning*, 36(1-2):39–83, 2006.
8. K. Chatterjee, L. Doyen, E. Filiot, and J.-F. Raskin. Doomsday equilibria for omega-regular games. *CoRR*, abs/1311.3238, 2013.

9. K. Chatterjee, L. Doyen, T.A. Henzinger, and J.-F. Raskin. Algorithms for omega-regular games with imperfect information. *LMCS*, 3(3), 2007.
10. K. Chatterjee, T. A. Henzinger, and M. Jurdzinski. Games with secure equilibria. *Theor. Comput. Sci.*, 365(1-2):67–82, 2006.
11. K. Chatterjee, T. A. Henzinger, and N. Piterman. Strategy logic. *Inf. Comput.*, 208(6):677–693, 2010.
12. K. Chatterjee, T.A. Henzinger, and N. Piterman. Generalized parity games. In *FoSSaCS*, volume 4423 of *LNCS*, pages 153–167. Springer, 2007.
13. 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.
14. E. A. Emerson and C. Jutla. Tree automata, mu-calculus and determinacy. In *FOCS*, pages 368–377. IEEE Comp. Soc., 1991.
15. E.A. Emerson and C.-L. Lei. Modalities for model checking: Branching time strikes back. In *POPL*, pages 84–96, 1985.
16. D. Fisman, O. Kupferman, and Y. Lustig. Rational synthesis. In *Proc. of TACAS*, LNCS 6015, pages 190–204. Springer, 2010.
17. E. Grädel, W. Thomas, and T. Wilke, editors. *Automata, Logics, and Infinite Games: A Guide to Current Research*, LNCS 2500. Springer, 2002.
18. N. Immerman. Number of quantifiers is better than number of tape cells. *Journal of Computer and System Sciences*, 22:384–406, 1981.
19. W. Jamroga, S. Mauw, and M. Melissen. Fairness in non-repudiation protocols. In *Proc. of STM*, LNCS 7170, pages 122–139. Springer, 2012.
20. S. Kremer and J.-F. Raskin. A game-based verification of non-repudiation and fair exchange protocols. *Journal of Computer Security*, 11(3):399–430, 2003.
21. O. Kupferman and M. Y. Vardi. Safraless decision procedures. In *FOCS*, 2005.
22. D. Martin. Borel determinacy. In *Annals of Mathematics*, volume 102, pages 363–371, 1975.
23. F. Mogavero, A. Murano, G. Perelli, and M.Y. Vardi. What makes atl* decidable? a decidable fragment of strategy logic. In *CONCUR*, pages 193–208, 2012.
24. 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.
25. J. F. Nash. Equilibrium points in n-person games. *PNAS*, 36:48–49, 1950.
26. N. Piterman. From nondeterministic Büchi and streett automata to deterministic parity automata. *Logical Methods in Computer Science*, 3(3), 2007.
27. N. Piterman and A. Pnueli. Faster solutions of rabin and streett games. In *LICS*, pages 275–284, 2006.
28. A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *POPL*, pages 179–190. ACM Press, 1989.
29. M. O. Rabin. Decidability of second-order theories and automata on infinite trees. *Trans. Amer. Math. soc.*, 141:1–35, 1969.
30. 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.
31. L. S. Shapley. Stochastic games. *PNAS*, 39:1095–1100, 1953.
32. M. Ummels and D. Wojtczak. The complexity of nash equilibria in stochastic multiplayer games. *Logical Methods in Computer Science*, 7(3), 2011.
33. 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 Additional Details – Doomsday Equilibria for Imperfect Information Games

We present automata constructions to recognise sequences of observations that are doomsday compatible and prefixes that are good for retaliation.

Lemma 9. *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 $G = (S, (S_i)_{1\leq i\leq n}, s_{\mathsf{init}}, \Sigma, \Delta, (O_i)_{1\leq i\leq n})$, 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. $\varphi_i = \mathsf{safe}(T_i)$. First, we construct the automaton

$$A = (Q^A, q^A_{\mathsf{init}}, (O_i \times (\Sigma \cup \{\tau\}), \delta^A)$$

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,
- $q^A_{\mathsf{init}} = s_{\mathsf{init}}$,
- $(q, (o, \sigma), q') \in \delta^A$ 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 [6] 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 $T_1, \ldots, T_n$ to be target states for each player respectively, i.e. $\varphi_i = \mathsf{reach}(T_i)$. The construction is in the same spirit as the construction for safety. Let $A = (Q^A, q^A_{\mathsf{init}}, O_i \times (\Sigma \cup \{\tau\}), \delta^A)$ 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 $\bigwedge_{j \neq i}(\mathsf{parity}_i \vee \overline{\mathsf{parity}_j})$, and treat each condition $\mathsf{parity}_i \vee \overline{\mathsf{parity}_j}$ separately. We dualize the acceptance condition of A, into the nondeterministic condition $\overline{\mathsf{parity}_i} \wedge \mathsf{parity}_j$. 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 $\mathsf{O}(k^2)$ pairs that accepts the same language [26]. 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. ☐

Lemma 10. *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 9, a deterministic Streett automaton $D_i = (Q^{D_i}, q^{D_i}_{\mathsf{init}}, (O_i \times (\Sigma \cup \{\tau\}), \delta^{D_i}, \mathsf{St}^{D_i})$ 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. 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 $o \in O_i$ s.t. $\exists (q,(o,\sigma),q') \in \delta^{D_i}$. 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 [27] 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 $\mathsf{Win} \subseteq Q^{D_i}$ 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 $q \in \mathsf{Win}$, 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 $q \notin \mathsf{Win}$ 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 $q_0q_1\dots q_k$ 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. ☐