

Doomsday Equilibria for Omega-Regular Games*

Krishnendu Chatterjee[†]

IST Austria

Laurent Doyen

LSV, ENS Cachan & CNRS

Emmanuel Filiot[‡] Jean-François Raskin[§]

Université Libre de Bruxelles – U.L.B.

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 [22, 20, 9]. One key application area is synthesis of reactive systems [2, 21, 19]. 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 [15, 9, 11].

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 [1, 10, 23]. 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 [18]. 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

*This work was published in the proceedings of the conference VMCAI'14 [7], and a long version is available online [6].

[†]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).

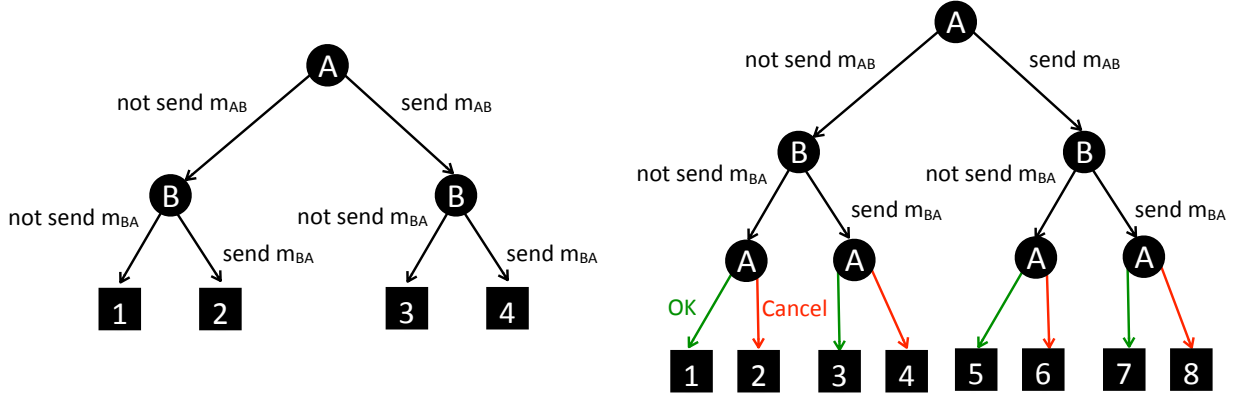


Figure 1: A simple example in the domain of Fair Exchange Protocols

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 [4] (that captures the adversarial external criteria), rational synthesis [10], and led to several new logics where the non-zero-sum equilibria can be expressed [5, 8, 17, 24, 16]. The complexity of Nash equilibria [23], secure equilibria [4], rational synthesis [10], and of the new logics has been studied recently [5, 8, 17, 24].

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 [4], and of fair-exchange protocols [13, 3] (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 [19]. 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 [12]. 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). Given n objectives $\varphi_1, \dots, \varphi_n$ and n strategies $\Lambda_1, \dots, \Lambda_n$ for each of the n players respectively, the strategy profile $\Lambda = (\Lambda_1, \dots, \Lambda_n)$ is a doomsday equilibrium if:

- (a) all players satisfy their own objectives, that is $outcome(\Lambda) \in \varphi_i$ for all $1 \leq i \leq n$ (where $outcome(\Lambda)$ is the path obtained according to the strategies in the profile), and
- (b) if any coalition of players deviates and violates even one of the players objective, then doomsday follows (every player objective is violated), that is for all $1 \leq i \leq n$, for all strategy profiles $\Lambda' = (\Lambda'_1, \dots, \Lambda'_n)$ such that $\Lambda'_i = \Lambda_i$, if $outcome(\Lambda') \notin \varphi_i$, then $outcome(\Lambda') \notin \varphi_j$ for all $1 \leq j \leq n$.

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 [4] where objectives of both players are satisfied.

Example 1. Consider the two trees of Figure 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 a 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. She has a slight preference for 2 as in that case she received the money while not transferring the property of her house to Bob, this corresponds to her adversarial external criterion. On the other hand, Bob would like to reach either state 3 or 4 (similarly with 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 $\{3, 4\}$ 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 to play "not send m_{BA} ". This profile 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 extra stage after the exchanges of the two messages as shown in the right tree of Figure 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 TTP¹). 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 Bob it is $5 > 7 > 1 = 2 = 4 = 6 = 8 > 3$. Now let us 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. This profile of strategies satisfies the first property of a doomsday equilibrium: both players have reached their primary objective, and no player has an incentive to deviate. Indeed, 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. If Bob deviates by playing "not send m_{BA} ", then Alice would cancel the protocol exchange which again produces a doomsday. 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 brought by the action "Cancel" has a beneficial influence on the behavior of the two players. \square

Example 2. Figure 2 gives two examples of games with safety and Büchi objectives respectively.

(Safety) Consider the 3-player game arena with perfect information of Figure 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.

¹TTP stands for Trusted Third Party.

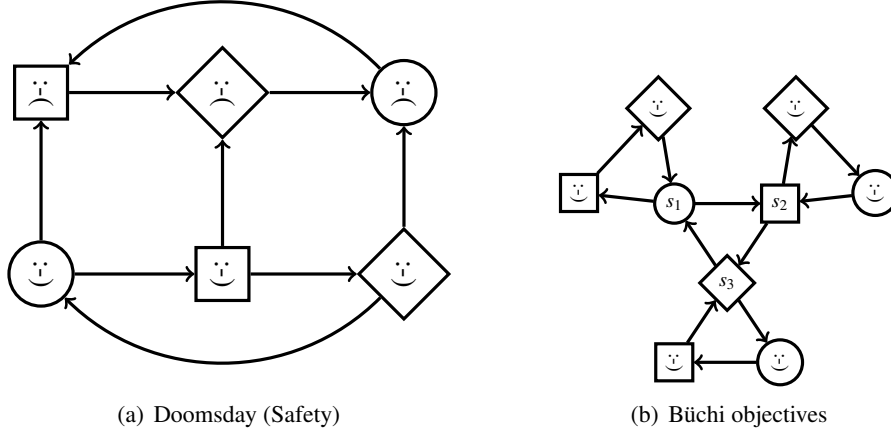


Figure 2: Examples of doomsday equilibria for Safety and Büchi objectives

(Büchi) Consider the 3-player game arena with perfect information of Figure 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 $(i + 1) \bmod 3$ and $(i + 2) \bmod 3$. 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 , then he alternately moves to the happy states and 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 in the future to go to the next s -state instead of visiting the happy states. Clearly, if all players follow their respective strategy, then 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 other two players eventually never visit their happy states anymore, but then all the players lose.

As a second one, consider the strategies where Player 2 and Player 3 always take their loop 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 time. In the second equilibrium, only Player 2 and 3 need a bit of memory. An exhaustive analysis shows that there is no memoryless doomsday equilibrium. \square

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. 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 ob-

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

jectives; (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.

In a long version of this paper [6], we also prove that deciding the existence of a doomsday threatening equilibrium in a game whose objectives are given as LTL formula is 2EXPTIME-complete, but we devise a Safraless procedure [14] suitable to efficient implementation.

The area of multi-player games and various notions 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.

References

- [1] R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49:672–713, 2002.
- [2] J. R. Büchi and L. H. Landweber. Definability in the monadic second-order theory of successor. *J. Symb. Log.*, 34(2):166–170, 1969.
- [3] R. Chadha, S. Kremer, and A. Scedrov. Formal analysis of multiparty contract signing. *J. Autom. Reasoning*, 36(1-2):39–83, 2006.
- [4] K. Chatterjee, T. A. Henzinger, and M. Jurdzinski. Games with secure equilibria. *Theor. Comput. Sci.*, 365(1-2):67–82, 2006.
- [5] K. Chatterjee, T. A. Henzinger, and N. Piterman. Strategy logic. *Inf. Comput.*, 208(6):677–693, 2010.
- [6] Krishnendu Chatterjee, Laurent Doyen, Emmanuel Filiot, and Jean-François Raskin. Doomsday equilibria for omega-regular games. *CoRR*, abs/1311.3238, 2013.
- [7] Krishnendu Chatterjee, Laurent Doyen, Emmanuel Filiot, and Jean-François Raskin. Doomsday equilibria for omega-regular games. In *VMCAI*, pages 78–97, 2014.
- [8] 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.
- [9] E. A. Emerson and C. Jutla. Tree automata, mu-calculus and determinacy. In *FOCS*, pages 368–377. IEEE Comp. Soc., 1991.
- [10] D. Fisman, O. Kupferman, and Y. Lustig. Rational synthesis. In *Proc. of TACAS*, LNCS 6015, pages 190–204. Springer, 2010.
- [11] E. Grädel, W. Thomas, and T. Wilke, editors. *Automata, Logics, and Infinite Games: A Guide to Current Research*, LNCS 2500. Springer, 2002.
- [12] W. Jamroga, S. Mauw, and M. Melissen. Fairness in non-repudiation protocols. In *Proc. of STM*, LNCS 7170, pages 122–139. Springer, 2012.

- [13] 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.
- [14] O. Kupferman and M. Y. Vardi. Safraless decision procedures. In *FOCS*, 2005.
- [15] D. Martin. Borel determinacy. In *Annals of Mathematics*, volume 102, pages 363–371, 1975.
- [16] 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.
- [17] 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.
- [18] J. F. Nash. Equilibrium points in n -person games. *PNAS*, 36:48–49, 1950.
- [19] A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *POPL*, pages 179–190. ACM Press, 1989.
- [20] M. O. Rabin. Decidability of second-order theories and automata on infinite trees. *Trans. Amer. Math. soc.*, 141:1–35, 1969.
- [21] 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.
- [22] L. S. Shapley. Stochastic games. *Proceedings of the National Academy of Sciences*, 39:1095–1100, 1953.
- [23] M. Ummels and D. Wojtczak. The complexity of nash equilibria in stochastic multiplayer games. *Logical Methods in Computer Science*, 7(3), 2011.
- [24] 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.