

Synthesis of Equilibria in Infinite-Duration Games on Graphs

Véronique Bruyère¹, University of Mons (UMONS)

In this survey, we propose a comprehensive introduction to game theory applied to computer-aided synthesis. We study multi-player turn-based infinite-duration games played on a finite directed graph such that each player aims at maximizing a payoff function. We present the well-known notions of Nash equilibrium and subgame perfect equilibrium, as well as interesting strategy profiles of players as response to the strategy announced by a specific player. We provide classical and recent results about the related threshold synthesis problem.

1. INTRODUCTION

Game theory is a well-developed branch of mathematics applied to various domains like economics, biology, computer science, etc. It is the study of mathematical models of interaction and conflict between individuals and the understanding of their decisions assuming that they act rationally [von Neumann and Morgenstern 1944; Osborne and Rubinstein 1994]. The last decades have seen a lot of research on algorithmic questions in game theory motivated by problems from *computer-aided synthesis*.

One important line of research is concerned with *reactive systems* that must continuously react to the events produced by the environment in which they evolve. Such systems are nowadays part of our daily life: think about common yet critical applications like engine control units in automotive, plane autopilots, medical devices, etc. Clearly, any flaw in such critical systems can have catastrophic consequences. A scientifically challenging goal, called *synthesis*, is to propose techniques (models, algorithms and tools) that, given a specification for a system and a model of its environment, compute (synthesize) a *controller* of the system that enforces the specification no matter how the environment behaves. To this end, researchers have advocated the use of *two-player games played on a graph*: the vertices of the graph model the possible configurations, the system and the environment are the two players, the infinite paths in the graph model their continuous interactions (reactive systems are usually not assumed to terminate). As the objectives of the two players are *antagonistic*, we speak of *zero-sum* games. Checking whether there exists a controller for the system reduces to the existence of a *winning strategy* in the corresponding game, and building a controller reduces to computing such a strategy [Grädel et al. 2002].

A lot of research has been done about *Boolean* objectives, in particular about the class of ω -*regular* objectives, like avoiding a deadlock or always granting a request [Grädel et al. 2002]. An infinite path in the game graph is either winning or losing depending on whether the objective is satisfied or not. To allow richer objectives, such as minimizing the energy consumption or guaranteeing a limited response time to a request, existing models have been recently enriched with *quantitative* aspects in a way to associate a payoff (or a cost) to all paths in the graph [Chatterjee et al. 2010a]. The most studied payoffs are liminf/limsup of the weights seen along the path, their mean-payoff or their discounted-sum. In this setting, the *constrained synthesis* problem is to decide whether there exists a winning strategy for the system that ensures a payoff satisfying some given constraints. For instance we would like an energy consumption lying within a certain given interval.

In practical situations, neither the system nor the environment are monolithic, and their objectives are not necessarily antagonistic: systems are composed of several parts

¹Work partially supported by the PDR project *Subgame perfection in graph games* (F.R.S.-FNRS) and the COST Action 16228 *GAMENET* (European Cooperation in Science and Technology).

that must be designed individually, and whose individual objectives must all be taken into account. For such more complex situation, it is advocated to use the model of *multi-player non-zero-sum games played on graphs*: the components are the different players, each of them aiming at satisfying their own objective. The synthesis problem is here different: winning strategies are no longer appropriate and are replaced by the concept of *equilibrium* [Grädel and Ummels 2008]. An equilibrium can be seen as a contract (strategy profile) between the players that makes each player satisfied with respect to his objective and discourages him to break this contract. Moreover, as a game may have several equilibria, one might be interested in one that fulfils certain constraints. For instance, in case of Boolean objectives, we would look for an equilibrium where some players are winning and some others are losing. Different kinds of equilibria have been studied among which the famous notion of *Nash equilibrium* (NE) [Nash 1950] from game theory. A strategy profile is an NE if no player has an incentive to unilaterally deviate from it. In the context of games played on graphs which are sequential by nature, it is well-known that NEs present a serious drawback: they allow for *non-credible threats*, i.e., irrational decisions taken by a player in a subgame to threaten another player and oblige him to follow a given behavior, see e.g. [Osborne and Rubinstein 1994]. To avoid non-credible threats, the notion of NE has been strengthened into the notion of *subgame perfect equilibrium* (SPE) [Selten 1965] also well-studied in game theory: a strategy profile is an SPE if it is an NE in each subgame of the original game. This notion behaves much better for sequential games and excludes non-credible threats.

The (constrained) synthesis problem is rather well *understood for NEs*. For classical ω -regular objectives (like reachability, Bchi, parity, etc) and classical quantitative objectives (like liminf, mean-payoff, etc), the same general approach can be used that works as follows (see e.g. the survey [Bruyère 2017]). Under some very general hypothesis, the plays that are outcomes of an NE can be *characterized* thanks to certain properties of the n two-player zero-sum games where one player (among the n players) is opposed to the coalition of the other players. From this characterization, it follows that there always exists an NE and that the constrained synthesis problem is decidable (with known complexity class). Only the constrained synthesis problem for discounted-sum games is unsolved. It is related to the challenging open problem of target discounted-sum itself related to many open questions in mathematics and computer science [Boker et al. 2015].

Whereas NEs are much studied, SPEs have received less attention and several questions are still unsolved for this concept of equilibrium. The most popular result in game theory is probably Kuhn's theorem that states the existence of SPE in games in finite extensive form (i.e., played on finite trees) [Kuhn 1953]. Unfortunately Kuhn's theorem is not applicable in computer-aided synthesis as we need to consider infinite-duration games. Some preliminary results have been obtained for graph games with ω -regular objectives [Grädel and Ummels 2008]: it is known that there always exists an SPE in these games, however the constrained synthesis problem is only partially solved. Little is known when we shift to quantitative objectives: very recently it has been proved that the constrained synthesis problem is PSPACE-complete for quantitative reachability games [Brihaye et al. 2020] and that it is decidable for mean-payoff games [Brice et al. 2021]. Those two results rely on an adequate new characterization of SPE outcomes like the one known for NEs.

As explained in the beginning of this introduction, the synthesis of controllers for reactive systems is usually studied with the model of two-player zero-sum games played on graphs such that the environment is a player opposed to the system player. A fully antagonistic environment is most often a bold abstraction of reality and several more suitable refinements have been proposed, see for instance the survey [Brenquier et al.

2016]. One such refinement is the framework of *Stackelberg games* [von Stackelberg 1937], a richer non-zero-sum setting from game theory, in which the system called *leader* announces his strategy and then the environment called *follower* rationally responds with a strategy that is optimal with respect to his own objective. The goal of the leader is therefore to announce a strategy that guarantees him a payoff satisfying certain constraints, whatever the optimal response of the follower.

This concept of leader and follower has been developed and extended in the context of computer-aided synthesis. The framework of *rational synthesis* [Fisman et al. 2010; Kupferman et al. 2016] considers several followers composing the environment (instead of one), each of them with their own objective. In that case, rationality of the followers is modeled by assuming that the environment settles to an equilibrium, like an NE or an SPE, where each component is considered to be an *independent selfish* individual, excluding cooperation scenarios between components or the possibility of coordinated rational multiple deviations. Two scenarios are investigated. Either the rational synthesis is *cooperative*: the environment cooperates with the system, it agrees to play an equilibrium that is satisfactory for the system (if it exists). Or it is *adversarial*: the environment can follow any equilibrium, and one has to synthesize a strategy (if it exists) for the system that is satisfactory against all these equilibria. Only partial results are known about the complexity of rational synthesis; let us mention results in [Condurache et al. 2016] for games with ω -regular objectives and in [Gupta and Schewe 2014; Filiot et al. 2020] for mean-payoff and discounted-sum games.

Very recently the framework of *Stackelberg-Pareto synthesis* has been proposed in [Bruyère et al. 2021a]. It is an alternative to rational synthesis where the environment is modeled as a single follower that has several objectives that he wants to satisfy. After responding to the leader with his strategy, the follower receives a tuple of payoffs in the corresponding outcome. Rationality is encoded by the fact that the follower only responds in such a way to receive a *Pareto-optimal* tuple of payoffs, given the strategy announced by the leader. This setting encompasses scenarios where, for instance, several components can collaborate and agree on trade-offs. The goal of the leader is therefore to announce a strategy that guarantees him to obtain a satisfactory own payoff, whatever the response of the follower which ensures him a Pareto-optimal tuple of payoffs. The complexity of Stackelberg-Pareto synthesis is studied in [Bruyère et al. 2021a] for some ω -regular objectives.

In this survey, we propose a *comprehensive introduction* to computer-aided synthesis in the different frameworks of non-zero-sum games presented above. The reader is referred to [Grädel and Ummels 2008; Brenguier et al. 2016; Bruyère 2017] for additional readings. We do not intend to present an exhaustive survey: the considered games are turn-based, deterministic, and with perfect information, and we focus on certain solution concepts and some related results that we find important. Additional references to related work are given in a dedicated section for the interested reader.

This article is structured in the following way. In Section 2, we introduce the concepts of non-zero-sum game, strategy, and Boolean and quantitative objectives (with a focus on parity and mean-payoff games). Section 3 is devoted to the famous NE concept and the related threshold synthesis problem that asks whether there exists an NE that guarantees to each player a payoff at least equal to some given threshold. We recall the characterization of NE outcomes and explain how it is used to solve the threshold synthesis problem. In Section 4, we explain the presence of non-credible threats in NEs and how the concept of SPE excludes this undesirable behavior for games played on graphs. We provide some known solutions to the threshold synthesis problem for SPEs and focus on a recent promising characterization of SPE outcomes [Flesch and Predtetchinski 2017; Brice et al. 2021] that generalizes to SPEs

the previous NE characterization. In Section 5, we study the concept of Stackelberg games, a framework well-suited for the synthesis of reactive systems. We present the cooperative/adversarial rational synthesis introduced in [Fisman et al. 2010; Kupferman et al. 2016] as well as the very recent setting of Stackelberg-Pareto synthesis proposed in [Bruyère et al. 2021a]. Section 6 gathers many references to related work. Finally, we provide a short conclusion in Section 7.

2. PRELIMINARIES

This section introduces the useful notions of multi-player turn-based game played on a graph, Boolean and quantitative payoff function, strategy and strategy profile.

2.1. Games

We consider multi-player turn-based games played on a finite directed graph [Grädel and Ummels 2008].

Definition 2.1. An arena is a tuple $A = (\Pi, V, (V_i)_{i \in \Pi}, E)$ where:

- Π is a finite set of *players*,
- V is a finite set of *vertices* and $E \subseteq V \times V$ is a set of *edges*, such that each vertex has at least one successor²,
- $(V_i)_{i \in \Pi}$ is a partition of V , where V_i is the set of vertices owned³ by player $i \in \Pi$.

A *play* is an infinite sequence $\rho = \rho_0 \rho_1 \dots \in V^\omega$ of vertices such that $(\rho_k, \rho_{k+1}) \in E$ for all $k \in \mathbb{N}$. *Histories* are non-empty finite sequences $h = h_0 \dots h_n \in V^+$ defined in the same way. We often use notation $h\nu$ to mention the last vertex $\nu \in V$ of the history. The set of plays is denoted by Plays and the set of histories (resp. histories ending with a vertex in V_i) by Hist (resp. by Hist _{i}). A *prefix* (resp. *suffix*) of a play $\rho = \rho_0 \rho_1 \dots$ is a finite sequence $\rho_{\leq n} = \rho_0 \dots \rho_n$ (resp. infinite sequence $\rho_{\geq n} = \rho_n \rho_{n+1} \dots$). We use notation $h\rho$ for a play of which history h is prefix. Given a play ρ , we denote by $\text{Inf}(\rho)$ the set of vertices visited infinitely often in ρ . We say that ρ is a *lasso* if it is equal to hg^ω with h, g being two histories.

Definition 2.2. A game G is an arena $A = (\Pi, V, (V_i)_{i \in \Pi}, E)$ such that each player i has a *payoff function* $\text{pay}_i : \text{Plays} \rightarrow \mathbb{R}$ that assigns a *payoff* to each play.

Player i prefers play ρ to play ρ' if $\text{pay}_i(\rho) > \text{pay}_i(\rho')$, that is, he wants to *maximize* his payoff. A payoff function pay_i is *prefix-independent* if $\text{pay}_i(h\rho) = \text{pay}_i(\rho)$ for all $h\rho \in \text{Plays}$. When an initial vertex $v_0 \in V$ is fixed, we call (G, v_0) an *initialized game*. In this case, plays and histories are assumed to start in v_0 , and we then use notations $\text{Plays}(v_0)$, $\text{Hist}(v_0)$, and $\text{Hist}_i(v_0)$ (instead of Plays, Hist, and Hist _{i}).

2.2. Classical Payoff Functions

A particular class of games G are those equipped with *Boolean* functions $\text{pay}_i : \text{Plays} \rightarrow \{0, 1\}$, $i \in \Pi$, such that a Boolean payoff is assigned to each play. The *objective* $\Omega_i = \{\rho \in \text{Plays} \mid \text{pay}_i(\rho) = 1\}$ of player i is composed of his most preferred plays. A play ρ such that $\text{pay}_i(\rho) = 1$ is said to *satisfy* the objective Ω_i of player i . Classical objectives Ω_i are *ω -regular* ones [Grädel et al. 2002; Grädel and Ummels 2008]. In this article, we focus on *parity* objectives as arbitrary ω -regular objectives can be reduced to them. Let $c_i : V \rightarrow \mathbb{N}$, $i \in \Pi$, be a *priority function* that labels the vertices of the arena with integers. Then the parity objective for player i is $\Omega_i = \{\rho \in \text{Plays} \mid \min_{v \in \text{Inf}(\rho)}(c(v)) \text{ is even}\}$. The payoff function associated with a parity objective is prefix-independent. There are

²This condition guarantees that there is no deadlock.

³We also say that player i *controls* the vertices of V_i .




Fig. 1. A two-player game with parity objectives for both players.

many other ω -regular objectives, such as reachability, safety, Büchi, co-Büchi, Streett, Rabin, Muller, etc. They are not studied in this survey.

Example 2.3. Consider the initialized two-player game (G, v_0) in Figure 1 such that player 1 controls all the vertices except vertex v_2 that is controlled by player 2.⁴ Both players have a parity objective, indicated in the figure under each vertex (the i -th component is the priority for player i). For the play $\rho = v_0v_1(v_0v_2v_3)^\omega$, we have payoff $(\text{pay}_1(\rho), \text{pay}_2(\rho)) = (1, 0)$ since the minimum priority visited infinitely often along ρ is equal to 0 for player 1 and 1 for player 2.

Other classical payoff functions are *quantitative* functions $\text{pay}_i : \text{Plays} \rightarrow \mathbb{R}$ defined from a *weight function* $w_i : E \rightarrow \mathbb{Q}$, $i \in \Pi$ [Chatterjee et al. 2010a]. Each edge of the game arena is thus labeled by a $|\Pi|$ -tuple of weights. In this article, we focus on the *mean-payoff* function MP_i . Let $\rho = \rho_0\rho_1 \dots \in \text{Plays}$, then

$$\text{MP}_i(\rho) = \liminf_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} w_i(\rho_k, \rho_{k+1}),$$

that is, $\text{MP}_i(\rho)$ is liminf of running averages of weights seen along the play (with respect to the weight function w_i). In case of a lasso $\rho = hg^\omega$, $\text{MP}_i(\rho)$ is equal to the average weight of the cycle g . The mean-payoff function is prefix-independent. There are several other quantitative functions like the discounted-sum function that is much studied in game theory as is the mean-payoff function [Filar and Vrieze 1997]. Let us also mention the supremum or limsup (resp. infimum, liminf) of the weights seen along a play [Chatterjee et al. 2010a]. Finally, the objective of quantitative reachability means reaching a given target set as quickly as possible (when counting the number of traversed edges). In this case, it is rather a cost that a player wants to minimize (instead of a payoff that he wants to maximize). Only the mean-payoff function is studied in details in this survey.

Example 2.4. Consider the initialized two-player game (G, v_0) in Figure 2 where the arena is equipped with two weight functions w_1, w_2 and the payoff function of each player is the mean-payoff function. As this function is prefix-independent, there are essentially three pairs of payoffs $(\text{pay}_1(\rho), \text{pay}_2(\rho))$ assigned to plays $\rho \in \text{Plays}(v_0)$: $(1, 1)$ (resp. $(0, 3)$, $(2, 2)$) assigned to the left lasso v_2^ω (resp. central lasso $(v_0v_1)^\omega$, right lasso v_3^ω).

2.3. Strategies

Let us now introduce the concept of strategy [Grädel et al. 2002; Grädel and Ummels 2008].

⁴In all examples of this article, there are two players such that circle (resp. square) vertices are controlled by player 1 (resp. player 2).




Fig. 2. A two-player game with mean-payoff function for both players.

Definition 2.5. Let (G, v_0) be an initialized game. A *strategy* σ_i for player i in (G, v_0) is a function $\sigma_i : \text{Hist}_i(v_0) \rightarrow V$ assigning to each history $hv \in \text{Hist}_i(v_0)$ a vertex $v' = \sigma_i(hv)$ such that $(v, v') \in E$. The set of strategies for player i from vertex v_0 is denoted by $\Gamma_i(v_0)$.

Thus $\sigma_i(hv)$ is the next vertex chosen by player i (who controls vertex v) after history hv has been played. A play $\rho \in \text{Plays}(v_0)$ is *compatible* with σ_i if $\rho_{n+1} = \sigma_i(\rho_{\leq n})$ for all n such that $\rho_n \in V_i$.

A strategy σ_i for player i is *memoryless* if it only depends on the last vertex of the history, i.e., $\sigma_i(hv) = \sigma_i(h'v)$ for all $hv, h'v \in \text{Hist}_i(v_0)$. A memoryless strategy can be depicted directly on the game arena by fixing one successor among all successors v' of the vertices v owned by player i . More generally, a strategy is *finite-memory* if it can be encoded by a finite automaton. In this case, the strategy chooses the next vertex depending on the current state of this automaton and the current vertex in the game.⁵

A *strategy profile* is a tuple $\sigma = (\sigma_i)_{i \in \Pi}$ of strategies, where each σ_i is a strategy for player i . Given an initial vertex v_0 , such a strategy profile determines a unique play of (G, v_0) that is compatible with all strategies σ_i . This play is called the *outcome* of σ in (G, v_0) and is denoted by $\langle \sigma \rangle_{v_0}$. We sometimes denote a strategy profile σ as (σ_i, σ_{-i}) where $-i$ denotes the set $\Pi \setminus \{i\}$ of all players except player i and σ_{-i} denotes the strategy profile for those players.

3. NASH EQUILIBRIA

In this section, we consider multi-player games such that each player has his own payoff function that he wants to maximize. We are interested in strategy profiles, called *solution profiles*, that provide payoffs satisfactory to all players.

3.1. Nash Equilibria Concept

A classical solution profile is the well-known notion of *Nash equilibrium* [Nash 1950] from game theory. Informally, a strategy profile is an NE if no player has an incentive to deviate (with respect to his payoff) when the other players stick to their own strategies. In other words, an NE can be seen as a *contract* that makes every player satisfied in the sense that nobody wants to break the contract if the others follow it. In this setting, each player rationally acts as an independent *selfish* individual.

Definition 3.1. Given an initialized game (G, v_0) , a strategy profile $\sigma = (\sigma_i)_{i \in \Pi}$ is a *Nash equilibrium (NE)* if $\text{pay}_i(\langle \sigma \rangle_{v_0}) \geq \text{pay}_i(\langle \sigma'_i, \sigma_{-i} \rangle_{v_0})$ for all players $i \in \Pi$ and all strategies $\sigma'_i \in \Gamma_i(v_0)$.

In this definition, all players stick to their own strategy except player i who switches from strategy σ_i to strategy σ'_i . We say that σ'_i is a *deviating* strategy from σ_i . By using σ'_i instead σ_i , player i is not capable to strictly increase his payoff.

⁵This informal definition is enough for this survey.




Fig. 3. A simple game with two NEs.

Example 3.2. Let us consider the initialized two-player game (G, v_0) of Figure 3 with four plays starting from the initial vertex v_0 . The payoff of each player for all these plays is indicated below the play. Let us consider the strategy profile $\sigma = (\sigma_1, \sigma_2)$ composed of the following memoryless strategies⁶:

- for player 1: $v_1 \rightarrow v_3, v_2 \rightarrow v_5$,
- for player 2: $v_0 \rightarrow v_1$.

This strategy profile is an NE with outcome $\langle \sigma \rangle_{v_0} = v_0 v_1 v_3^\omega$. Indeed, each player has no incentive to deviate since his payoff is maximal (it equals 2). There exists another NE defined by:

- for player 1: $v_1 \rightarrow v_4, v_2 \rightarrow v_5$,
- for player 2: $v_0 \rightarrow v_2$,

such that its outcome equals $v_0 v_2 v_5^\omega$ and has payoffs (1, 1). For instance, if player 2 uses the deviating strategy σ'_2 such that $v_0 \rightarrow v_1$, the resulting outcome $\langle \sigma_1, \sigma'_2 \rangle_{v_0} = v_0 v_1 v_4^\omega$ decreases by 1 the payoff for player 2. Clearly the first NE is more interesting for both players as its outcome has payoffs (2, 2) (instead of (1, 1) for the second NE).

3.2. Threshold Synthesis Problem

Example 3.2 shows that a game may have several NEs. It is thus natural to ask whether there exists an NE that fulfils certain requirements. This motivation leads to the following decision problem [Grädel and Ummels 2008; Bruyère 2017] that we study in details in this survey for different kinds of solution profiles.

Definition 3.3. Let (G, v_0) be an initialized game and $\mu_i \in \mathbb{Q}$ be a threshold for each $i \in \Pi$. The *threshold synthesis problem* is to decide whether there exists a solution profile $\sigma = (\sigma_i)_{i \in \Pi}$ such that $\mu_i \leq \text{pay}_i(\langle \sigma \rangle_{v_0})$ for all players $i \in \Pi$.

Notice that when the given bounds impose no constraint (for instance $\mu_i = 0$, for all $i \in \Pi$, in case of Boolean payoff functions), the decision problem can be rephrased as the *existence synthesis problem*: “decide whether there exists a solution profile”. We will see that for certain classes of games and solution concepts, this problem does not need to be solved because there always exists a solution profile for those games.

Let us here present some results about the existence synthesis problem and the threshold synthesis problem in the case of NEs. They are given for the games on which

⁶As the strategies are memoryless, we simply indicate the chosen successor for each vertex, except for the vertices with only one successor.




Fig. 4. A one-player game with no NE.

we focus in this survey: games with parity objectives and games with mean-payoff functions.

THEOREM 3.4.

- Games with parity objectives always admit an NE [Chatterjee et al. 2004].
- Games with mean-payoff functions always admit an NE [Brihaye et al. 2013].
- Moreover, the threshold synthesis problem is NP-complete for all these games [Ummels 2008; Ummels and Wojtczak 2011a].

NE existence holds for many classes of games. There always exists an NE in all games with ω -regular objectives⁷ [Chatterjee et al. 2004; Grädel and Ummels 2008] and in a large class of games with quantitative payoff functions including those with mean-payoff functions [Brihaye et al. 2013]. Moreover there always exists an NE σ such that each σ_i , $i \in \Pi$, is a finite-memory strategy. The threshold synthesis problem is solved in [Ummels 2008; Condurache et al. 2016] for different types of ω -regular objectives including parity. The case of games with mean-payoff functions is solved in [Ummels and Wojtczak 2011a] and with quantitative reachability objectives in [Brihaye et al. 2019].

Nevertheless there exist games that admit no NE as shown in the next simple example [Brihaye et al. 2016].

Example 3.5. Consider the initialized one-player game (G, v_0) in Figure 4 with the following payoff function $\text{pay}_1 : \text{Plays} \rightarrow \mathbb{N}$ for the unique player:

- $\text{pay}_1(v_0^k v_1^\omega) = k$, for all $k \geq 1$,
- $\text{pay}_1(v_0^\omega) = 0$.

No strategy profile is an NE. Indeed if the player chooses to always loop on v_0 , he gets payoff 0 that he can increase by using the deviating strategy σ'_1 such that $\sigma'_1(v_0) = v_1$ in which case he gets payoff 1. If he decides to loop k times on v_0 and then go to v_1 , he gets a payoff that he can increase by 1 by looping once more on v_0 .

3.3. NE Characterization

The proofs of several results about the existence synthesis problem and the threshold synthesis problem for NEs are based on an elegant characterization of NE outcomes that we now present. To this end, we need to introduce some definitions where $\overline{\mathbb{R}}$ denotes the set $\mathbb{R} \cup \{-\infty, +\infty\}$.

Definition 3.6. Let G be a game and $\lambda : V \rightarrow \overline{\mathbb{R}}$ be a function called *requirement function*. A play $\rho = \rho_0 \rho_1 \dots \in \text{Plays}$ is λ -consistent if for all $n \in \mathbb{N}$, we have

$$\text{pay}_i(\rho_{\geq n}) \geq \lambda(\rho_n)$$

such that player i controls vertex ρ_n .

⁷This result holds for the larger class of games with Borel Boolean objectives.

Thus the quantity $\lambda(\rho_n)$ represents the minimal payoff that player i who controls vertex ρ_n requires for the suffix $\rho_{\geq n}$ starting in this vertex. Notice that when the payoff function pay_i is prefix-independent, condition $\text{pay}_i(\rho_{\geq n}) \geq \lambda(\rho_n)$ can be replaced by $\text{pay}_i(\rho) \geq \lambda(\rho_n)$. Notice also that there exists no λ -consistent play if λ is the requirement function such that $\lambda(v) = +\infty$ for all $v \in V$.⁸ On the opposite, each play is λ -consistent if λ is the requirement function such that $\lambda(v) = -\infty$ for all $v \in V$.

Definition 3.7. Let G be a game and $\lambda_{\text{NE}} : V \rightarrow \overline{\mathbb{R}}$ be the requirement function such that for all $i \in \Pi$ and all $v \in V_i$,

$$\lambda_{\text{NE}}(v) = \inf_{\sigma_{-i} \in \Gamma_{-i}(v)} \sup_{\sigma_i \in \Gamma_i(v)} \text{pay}_i(\langle \sigma \rangle_v). \quad (1)$$

In this definition, it is useful to see the players $-i$ forming a *coalition* (as a unique player) opposed to player i : player i tries to maximize his payoff from v against the coalition $-i$ trying to minimize this payoff. The quantity $\lambda_{\text{NE}}(v)$ represents the worst payoff that player i who controls v can hope against the coalition opposed to him.

The announced characterization is based on λ_{NE} -consistency in the following way. The second hypothesis of this theorem asks for the infimum in equation (1) to be achieved.

THEOREM 3.8 ([BRIHAYE ET AL. 2013]). *Let G be a game such that:*

- for each player i , the payoff function pay_i is prefix-independent, and
- for each player i , for each vertex $v \in V_i$, there exists a strategy profile $\sigma_{-i} \in \Gamma_{-i}(v)$ such that $\lambda_{\text{NE}}(v) = \sup_{\sigma_i \in \Gamma_i(v)} \text{pay}_i(\langle \sigma \rangle_v)$.

Then for each initial vertex v_0 , a play $\rho \in \text{Plays}(v_0)$ is an NE outcome in (G, v_0) if and only if ρ is λ_{NE} -consistent.

Let us illustrate all these notions with an example.

Example 3.9. We come back to the game of Example 3.2 for which we suppose the payoff function of each player being prefix-independent. In Figure 5, the value $\lambda_{\text{NE}}(v)$ is indicated to the right of each vertex v of this game.

Consider for instance vertex v_0 such that $\lambda_{\text{NE}}(v_0)$ is equal to 1 and the (memoryless) strategy σ_{-2} from v_0 is depicted in bold in Figure 5 (the coalition -2 is composed of player 1). Indeed either (i) player 2 chooses $v_0 \rightarrow v_1$ and player 1 opposed to him chooses $v_1 \rightarrow v_4$ leading to payoff 0 for player 2, or (ii) player 2 chooses $v_0 \rightarrow v_2$ and player 1 chooses $v_2 \rightarrow v_5$ leading to payoff 1 for player 2 that is the best payoff for him.

With this example, let us provide some intuition on the correctness of the NE characterization. On the one hand, the play $\rho = v_0 v_1 v_4^\omega$ with payoffs $(0, 0)$ is not an NE outcome because it is not λ_{NE} -consistent: $\text{pay}_2(\rho) < \lambda_{\text{NE}}(v_0)$. Indeed due to value $\lambda_{\text{NE}}(v_0) = 1$, player 2 knows that he can guarantee a payoff of at least 1 by deviating from ρ at v_0 with the strategy $v_0 \rightarrow v_2$. This is a better payoff for him than the payoff he gets with ρ . On the other hand, the play $\rho = v_0 v_1 v_3^\omega$ with payoffs $(2, 2)$ is an NE outcome because one can verify that it is λ_{NE} -consistent. One can construct an NE with outcome ρ as follows. It is constructed such as to produce ρ , and as soon as some player i deviates from ρ at vertex $v \in V_i$, it uses the strategy profile σ_{-i} from v of the coalition $-i$ as given in the second hypothesis of Theorem 3.8. As ρ is λ_{NE} -consistent and pay_i is prefix-independent (by the first hypothesis of Theorem 3.8), with this deviation player i receives a payoff $\leq \lambda_{\text{NE}}(v) \leq \text{pay}_i(\rho)$, thus not profitable for him.

⁸Recall that we have $\text{pay}_i : \text{Plays} \rightarrow \mathbb{R}$ by definition.




Fig. 5. Requirement function λ_{NE} for the game of Example 3.2.

3.4. Comments on the NE Characterization

The NE characterization of Theorem 3.8 is stated as given in [Brice et al. 2021]. It appears under various particular forms for instance in [Grädel and Ummels 2008; Ummels and Wojtczak 2011a; Bruyère et al. 2014; Brihaye et al. 2019], and in under more general hypotheses in [Brihaye et al. 2013; Bruyère 2017]. The use of coalitions opposed to some player is also a well-known method in game theory, for instance in the proof of the Folk Theorem in repeated games [Osborne and Rubinstein 1994].

The two hypotheses of Theorem 3.8 are satisfied by many classes of games:

- Concerning the first hypothesis, many studied Boolean and quantitative payoff functions are prefix-independent like the function associated with parity objective and the mean-payoff function.
- The second hypothesis is related to the well-known concepts of value and optimal strategy in *two-player zero-sum* games where the two players i and $-i$ have opposed objectives [Grädel et al. 2002]: player i wants to maximize his payoff whereas player $-i$ wants to minimize this payoff. When it exists, the *value* of a vertex v is the largest (resp. lowest) payoff that player i (resp. player $-i$) can ensure from v , and *optimal strategies* for players i and $-i$ respectively, are strategies realizing this value. The value and optimal strategies for both players exist in many classes of two-player zero-sum games, see e.g. [Gimbert and Zielonka 2005; Bruyère 2017], including games with parity objectives [Emerson and Jutla 1991] and games with mean-payoff functions [Zwick and Paterson 1996]. In equation (1), $\lambda_{\text{NE}}(v)$ is the value of v and in Theorem 3.8, the second hypothesis requires the existence of an optimal strategy for player $-i$.

THEOREM 3.10 ([EMERSON AND JUTLA 1991; ZWICK AND PATERSON 1996]).

For two-player zero-sum games with either parity objectives or mean-payoff functions, the value exists for all vertices v as well as optimal strategies for both players. Moreover, each player can use the same memoryless optimal strategy to ensure the value of all vertices.

Let us come back to Theorem 3.4 and give some comments on its proof:

- The proof of its first two statements follows from the NE characterization given in Theorem 3.8 and from Theorem 3.10 above. For each player i , consider the two-player zero-sum game where he is opposed to player $-i$, and let τ_i, τ_{-i} be the memoryless strategies for those two players as stated in Theorem 3.10. An NE is constructed from $\tau_i, \tau_{-i}, i \in \Pi$, as follows: from the initial vertex play with τ_i for each player i , and as

soon as some player i deviates, punish i by playing τ_{-i} . By construction, this strategy profile produces an outcome that satisfies Theorem 3.8 and is thus an NE.

- Let us give a sketch of proof for the NP-membership of the last statement of Theorem 3.4. Given $\mu_i \in \mathbb{Q}$ for each $i \in \Pi$, one has to decide whether there exists an NE σ such that its outcome satisfies $\mu_i \leq \text{pay}_i(\langle \sigma \rangle_{v_0})$ for all $i \in \Pi$. By Theorem 3.8, it is equivalent to decide whether there exists a play $\rho \in \text{Plays}(v_0)$ such that (i) ρ is λ_{NE} -consistent and (ii) $\mu_i \leq \text{pay}_i(\rho)$ for $i \in \Pi$. For parity objectives (resp. mean-payoff functions) condition (i) is decidable in NP and condition (ii) can be solved in polynomial time [Emerson and Jutla 1991; Ummels 2006] (resp. [Zwick and Paterson 1996; Ummels and Wojtczak 2011b]).

4. SUBGAME PERFECT EQUILIBRIA

In this section, we present another well-known type of solution profile from game theory, the *subgame perfect equilibrium* [Selten 1965], that avoids some NE drawbacks.

4.1. Subgame Perfect Equilibria Concept

NEs do not take into account the sequential nature of games played on graphs. Indeed after any prefix of a play, the players face a new situation and may want to change their strategies. It is well-known that NEs suffer from the problem of *non-credible threat*, see for instance [Osborne and Rubinstein 1994]: the existence of NEs may rely on unrational strategies of some players in subgames. This is illustrated in the next example.

Example 4.1. We come back to Example 3.2 (see Figure 3) and its second NE defined by

- for player 1: $v_1 \rightarrow v_4, v_2 \rightarrow v_5$,
- for player 2: $v_0 \rightarrow v_2$.

Player 2 has no incentive to deviate from v_0 due to the decision of player 1 to play with $v_1 \rightarrow v_4$ in the subgame with initial vertex v_1 . This is a non-credible threat from player 1 in this subgame: it would be more rational for him to play with $v_1 \rightarrow v_3$ in a way to get payoff 2 of instead of 0.

In order to be a subgame perfect equilibrium, a strategy profile is not only required to be an NE from the initial vertex but also after every possible history of the game. Before giving the precise definition of a subgame perfect equilibrium, we need to introduce the following concepts. Let (G, v_0) be an initialized game with payoff functions $\text{pay}_i, i \in \Pi$. Given a history $hv \in \text{Hist}(v_0)$, the *subgame* $(G|_h, v)$ of (G, v_0) is the initialized game with payoff functions $\text{pay}_{i|h}, i \in \Pi$, such that $\text{pay}_{i|h}(\rho) = \text{pay}_i(h\rho)$ for all plays $\rho \in \text{Plays}(v)$. Given a strategy σ_i for player i in (G, v_0) , the strategy $\sigma_{i|h}$ in $(G|_h, v)$ is defined as $\sigma_{i|h}(h') = \sigma_i(hh')$ for all $h' \in \text{Hist}_i(v)$. If $\sigma = (\sigma_i)_{i \in \Pi}$ is a strategy profile in (G, v_0) , we denote by $\sigma|_h$ the strategy profile equal to $(\sigma_{i|h})_{i \in \Pi}$ in $(G|_h, v)$.

Definition 4.2. Given an initialized game (G, v_0) , a strategy profile σ is a *subgame perfect equilibrium (SPE)* if $\sigma|_h$ is an NE in each subgame $(G|_h, v)$ of (G, v_0) with $hv \in \text{Hist}(v_0)$.

By definition, every SPE is an NE. However the converse is false as shown in Example 4.1 where the given strategy profile is an NE but not an SPE.

4.2. Threshold Synthesis Problem

Let us present some results about the threshold synthesis problem and the existence synthesis problem in case of SPEs. Whereas these problems are largely solved for NEs

(see Theorem 3.4 and the comments that follow it), we only have partial answers for SPEs. Moreover, SPEs fail to exist even in simple games like the game with mean-payoff functions of Figure 2 [Solan and Vieille 2003]. Recall that NEs always exist for the class of games with mean-payoff functions as well as for many other classes of games (see previous section).

THEOREM 4.3. *Class of games with parity objectives:*

- *Games with parity objectives always admit an SPE [Ummels 2006; Grädel and Ummels 2008].*
- *The threshold synthesis problem for games with parity objectives is in EXPTIME and is NP-hard [Ummels 2006; Grädel and Ummels 2008].*

Class of games with mean-payoff functions:

- *There exist games with mean-payoff functions that admit no SPE [Solan and Vieille 2003].*
- *The threshold synthesis problem for games with mean-payoff functions is decidable [Brice et al. 2021].*

A well-known result is the existence of an SPE in every game (G, v_0) such that its arena is a tree rooted at v_0 ⁹ [Kuhn 1953]. The example in Figure 3 falls into this class of games. The SPE is easily constructed backwards from the leaves to the initial vertex v_0 . It is proved in [Ummels 2006; Grädel and Ummels 2008] that there exists an SPE in every game with ω -regular objectives¹⁰ and thus with parity objectives (moreover there exists one composed of finite-memory strategies). The existence of an SPE also holds for games with bounded continuous payoff functions [Fudenberg and Levine 1983; Harris 1985] (including the discounting-sum function), and more generally for games where these functions are upper-semicontinuous (resp. lower-semicontinuous) and with finite range [Flesch et al. 2010] (resp. [Purves and Sudderth 2011]). Notice that games with mean-payoff functions do not fall into those classes of games.

The proofs provided in [Ummels 2006; Grädel and Ummels 2008] and in [Flesch et al. 2010] to construct an SPE are based on a non-increasing sequence of sets of plays: the initial set is the set of all plays; then step by step the set loses some plays that definitely are not SPE outcomes; finally one reaches a non-empty fixpoint that exactly contains all SPE outcomes and from which an SPE can be constructed. This fixpoint approach has been further developed for games with reachability objectives [Brihaye et al. 2015; Brihaye et al. 2018] and very recently for games with quantitative reachability objectives [Brihaye et al. 2020] and with mean-payoff functions [Brice et al. 2021]. The latter approach is detailed in the next two sections as it is in the spirit of the NE characterization given in Theorem 3.8.

The proof to establish the EXPTIME-membership of the threshold synthesis problem for games with parity objectives uses an algorithm based on tree automata [Ummels 2006; Grädel and Ummels 2008].

4.3. SPE Characterization

We now present the elegant SPE characterization provided in [Flesch and Predtetchinski 2017; Brice et al. 2021]. It is formulated here as given in [Brice et al. 2021] (it appears under a different formulation in [Flesch and Predtetchinski 2017] for games with finite-range and Borel measurable payoff functions). This characterization uses

⁹In this particular context, plays are finite paths.

¹⁰This is also true for games with Borel Boolean objectives.

an adequate requirement function λ_{SPE} (instead of λ_{NE}) such that SPE outcomes coincide with plays that are λ_{SPE} -consistent. The definition of the requirement function λ_{SPE} is more complex: it is the least fixpoint of some operator on requirement functions. Let us introduce some useful notions.

Definition 4.4. Let (G, v_0) be an initialized game and $\lambda : V \rightarrow \overline{\mathbb{R}}$ be a requirement function.

- Let i be a player and $-i$ be the coalition of the other players. A strategy profile $\sigma_{-i} \in \Gamma_{-i}(v_0)$ is λ -rational if there exists a strategy $\sigma_i \in \Gamma_i(v_0)$ such that for all histories $hv \in \text{Hist}(v_0)$ compatible with σ_{-i} , the play $\langle \sigma|_h \rangle_v$ is λ -consistent.
- If such a strategy profile σ_{-i} exists for some player i , we say that λ is *satisfiable from* v_0 . The set of all λ -rational strategy profiles in (G, v_0) is denoted by $\lambda\text{Rat}(v_0)$.

In this definition, if the coalition $-i$ proposes a λ -rational strategy profile to player i that he accepts to follow (after a finite number of deviations), then the resulting play is λ -consistent.

In the next definition, we introduce an operator that transforms a requirement function λ into a new one that can be seen as the result of *negotiations* between the players: when a player has a requirement to satisfy, another player can hope a better payoff than before and therefore update his own requirement.

Definition 4.5. Let G be a game. The *negotiation function* Nego is an operator that transforms any requirement function λ into the requirement function $\text{Nego}(\lambda)$ defined as follows. For all $i \in \Pi$ and all $v \in V_i$,

$$\text{Nego}(\lambda)(v) = \inf_{\sigma_{-i} \in \lambda\text{Rat}(v)} \sup_{\sigma_i \in \Gamma_i(v)} \text{pay}_i(\langle \sigma \rangle_v), \quad (2)$$

with the convention that $\inf \emptyset = +\infty$. The least fixpoint of Nego is denoted by λ_{SPE} .¹¹

Equation (2) is similar to equation (1) with the exception that each coalition $-i$ plays with strategies σ_{-i} that are λ -rational. The quantity $\text{Nego}(\lambda)(v)$ represents the worst payoff that the player controlling v can hope against the coalition opposed to him while playing with λ -rational strategies. Notice that $\text{Nego}(\lambda)(v) = +\infty$ if and only if λ is not satisfiable from v .

We can now state the announced SPE characterization. Notice that it shares similarities with the NE characterization stated in Theorem 3.8: the first hypothesis is the same, the second one is the second hypothesis of Theorem 3.8 adapted to λ -rational strategy profiles.

THEOREM 4.6 ([FLESCH AND PREDTETCHINSKI 2017; BRICE ET AL. 2021]). *Let G be a game such that:*

- *for each player i , the payoff function pay_i is prefix-independent, and*
- *for each player i , for each vertex v , for each requirement function λ satisfiable from v , there exists a λ -rational strategy profile σ_{-i} from v such that*

$$\inf_{\sigma'_{-i} \in \lambda\text{Rat}(v)} \sup_{\sigma'_i \in \Gamma_i(v)} \text{pay}_i(\langle \sigma' \rangle_v) = \sup_{\sigma_i \in \Gamma_i(v)} \text{pay}_i(\langle \sigma \rangle_v).$$

Then for each initial vertex v_0 , a play $\rho \in \text{Plays}(v_0)$ is an SPE outcome in (G, v_0) if and only if ρ is λ_{SPE} -consistent.

Let us illustrate the computation of λ_{SPE} on an example of game G that admits an SPE (see Example 4.7 below). We start with the *vacuous* requirement function λ_0

¹¹The negotiation function is monotone and the least fixpoint exists by Tarski's fixpoint theorem.




Fig. 6. Computation of the requirement function λ_{SPE} for the game of Example 3.2.

such that $\lambda_0(v) = -\infty$ for all $v \in V$. Notice that all plays are λ_0 -consistent since λ_0 imposes no constraint. We then iteratively apply the negotiation function Nego , i.e., we compute $\lambda_{k+1} = \text{Nego}(\lambda_k)$ for all $k \geq 0$, in a way to update the requirement function until reaching a fixpoint.

Example 4.7. We consider again the game of Example 3.2. In Figure 6, close to each vertex, we indicate the value of λ_0 , λ_1 and $\lambda_2 = \lambda_{\text{SPE}}$ (from left to right).

Notice that λ_1 is equal to λ_{NE} as described in Example 3.9 (see Figure 5). Indeed, when computing λ_1 from λ_0 , both equations (2) and (1) are equal as we have $\sigma_{-i} \in \lambda_0 \text{Rat}(v)$ if and only if $\sigma_{-i} \in \Gamma_{-i}(v)$ (see subscripts of \inf in (1) and (2)).

Let us explain the computation of $\lambda_2(v_0) = 2$. As v_0 is controlled by player 2, the opposed coalition is player 1. The play $v_0 v_1 v_4^\omega$ (which was λ_0 -consistent) is not λ_1 -consistent and the strategy of the coalition -2 such that $v_1 \rightarrow v_4$ is not λ_1 -rational. Therefore the best choice for player 2 from v_0 is $v_0 \rightarrow v_1$ to get payoff 2.

4.4. Comments on the SPE Characterization

We already mentioned the similarities between equations (1) and (2). We have that $\lambda_{\text{NE}} = \text{Nego}(\lambda_0)$ with λ_0 being the vacuous requirement function. Thus the first iteration of the Nego function leads to the NE characterization whereas the least fixpoint of Nego leads to the SPE characterization.

An important result of [Brice et al. 2021] is that games with mean-payoff functions satisfy the second condition of Theorem 4.6. Notice that games with finite-range payoff functions also satisfy this condition. Therefore, the SPE characterization stated in Theorem 4.6 holds for games with parity objectives and for games with mean-payoff functions.

Recall that there exist games with no SPE like the game with mean-payoff objectives of Figure 2 [Solan and Vieille 2003]. Therefore, by Theorem 4.6, for games G with mean-payoff objectives, (G, v_0) admits no SPE if and only if $\lambda_{\text{SPE}}(v_0) = +\infty$.

Example 4.8. We consider again the game G of Figure 2 that admits no SPE from the initial vertex v_0 . In Figure 7, close to each vertex, we indicate the value of λ_0 , λ_1 , λ_2 and $\lambda_3 = \lambda_{\text{SPE}}$ (from left to right). Notice that $\lambda_{\text{SPE}}(v_0) = +\infty$ that confirms the absence of SPE from v_0 .

In Examples 4.7 and 4.8, λ_{SPE} is computed as the limit of the stationary sequence $(\lambda_k)_{k \geq 0}$ such that λ_0 is vacuous requirement and $\lambda_{k+1} = \text{Nego}(\lambda_k)$ for all $k \geq 0$. An




Fig. 7. A game with no SPE.

example of game with mean-payoff functions is given in [Brice et al. 2021] such that this limit is never reached by any requirement function. Nevertheless an algorithm is provided in [Brice et al. 2021] that computes λ_{SPE} for all games with mean-payoff functions, thus showing that the SPE existence synthesis problem is decidable for this class of games (as stated in Theorem 4.3).

5. STACKELBERG GAMES

In this section, we present another concept that is more adequate than the concepts of NE and SPE for the synthesis of reactive systems: *Stackelberg games* [von Stackelberg 1937]. Those games have a specific player called the *leader*, the other players being called *followers*. The leader starts the game by announcing his strategy and the followers respond by playing rationally given that strategy. In case of one follower, his strategy can be an optimal response with respect to his own objective; in case of several followers, they can respond with a strategy profile that is an NE. The goal of the leader is therefore to announce a strategy that guarantees him a payoff at least equal to some given threshold whatever the rational response of the follower(s).

5.1. Several Followers

We begin with the case of several followers such that each of them models one component of the environment. Rationality of the followers is modeled by assuming that the environment settles to an NE: each component is considered to be an independent selfish individual [Fisman et al. 2010; Gupta and Schewe 2014; Kupferman et al. 2016].

Definition 5.1. Let (G, v_0) be an initialized game with a specific player $0 \in \Pi$. Let σ_0 be a strategy for player 0. A σ_0 -Stackelberg profile (σ_0 -SP) is a strategy profile $\sigma = (\sigma_0, (\sigma_i)_{i \in \Pi \setminus \{0\}})$ such that $\text{pay}_i(\langle \sigma \rangle_{v_0}) \geq \text{pay}_i(\langle \sigma'_i, \sigma_{-i} \rangle_{v_0})$ for all players $i \in \Pi \setminus \{0\}$ and all strategies $\sigma'_i \in \Gamma_i(v_0)$.

In this definition, the strategy σ_0 of the leader $i = 0$ is fixed, only deviating strategies σ'_i of the followers, i.e., with $i \neq 0$, are considered. Two variants of the threshold synthesis problem are proposed in the context of Stackelberg games, depending on whether the followers respond to the leader *cooperatively* or *adversarially*.

Definition 5.2. Let (G, v_0) be an initialized game and $\mu \in \mathbb{Q}$ be a threshold.

- The *cooperative threshold synthesis problem* is to decide whether there exists a strategy σ_0 for player 0 and a σ_0 -SP σ such that $\mu \leq \text{pay}_0(\langle \sigma \rangle_{v_0})$.
- The *adversarial threshold synthesis problem* is to decide whether there exists a strategy σ_0 for player 0 such that for all σ_0 -SP σ , we have $\mu \leq \text{pay}_0(\langle \sigma \rangle_{v_0})$.

We illustrate these notions with the next example.




Fig. 8. A cooperative or adversarial follower.

Example 5.3. Let us consider the two-player game (G, v_0) depicted in Figure 8 such that the payoffs are indicated below each of the three plays. The leader is the player that controls the circle vertices and the (unique) follower is the other player. Suppose that the leader announces to play $v_1 \rightarrow v_4$. The follower has two possible responses that are NEs: he can play either $v_0 \rightarrow v_1$ or $v_0 \rightarrow v_2$. A cooperative follower will play $v_0 \rightarrow v_1$ that maximizes the payoff (equal to 1) of the leader. An adversarial follower will play $v_0 \rightarrow v_2$ such that the leader only gets a payoff of 0.

Clearly every NE is a σ_0 -SP. However the next example shows that σ_0 -SPs may produce a strictly better payoff for the leader than all NEs. Thus a solution to the cooperative threshold synthesis problem is not necessarily an NE.

Example 5.4. Let us come back to the game (G, v_0) of Figure 8 such that the payoffs of play $v_0 v_1 v_4^\omega$ are now equal to $(1, 2)$ instead of $(1, 1)$. The only NE is the strategy profile (σ_0, σ_1) such that player 0 plays $v_1 \rightarrow v_3$ and player 1 plays $v_0 \rightarrow v_2$. The outcome of this NE gives payoff 0 to the leader. Suppose now that the leader announces to play $v_1 \rightarrow v_4$. The only NE response of the follower is to play $v_0 \rightarrow v_1$. The resulting strategy profile (σ'_0, σ'_1) is a σ'_0 -SP that is not an NE and whose outcome gives payoff 1 to the leader. Therefore (σ'_0, σ'_1) is solution to the cooperative threshold synthesis problem with threshold $\mu = 1$ whereas there is no solution that is an NE.

In case of Boolean payoff functions, the cooperative/adversarial threshold synthesis problem is interesting only with the threshold μ equal to 1: we ask for the objective Ω_0 of the leader to be satisfied. Notice that a σ_0 -SP such that Ω_0 is satisfied is then an NE because the leader has no incentive to deviate. Therefore in this case the cooperative threshold synthesis problem is nothing else than the threshold synthesis problem studied for NEs in Section 3.2. Hence for the cooperative threshold synthesis problem we focus on quantitative payoff functions only. We have the next result:

THEOREM 5.5 ([GUPTA AND SCHEWE 2014]). *The cooperative threshold synthesis problem is NP-complete for games with mean-payoff functions.*

The techniques used to prove this result are related to some used for NEs, in particular they build on the NE characterization of Theorem 3.8 and on [Ummels and Wojtczak 2011a]. The cooperative threshold synthesis problem is studied for games with discounted-sum functions in [Gupta et al. 2015; Filiot et al. 2020].

We have the following results for the adversarial threshold synthesis problem.

THEOREM 5.6 ([CONDURACHE ET AL. 2016]). *For games with parity objectives, the adversarial threshold synthesis problem (with threshold $\mu = 1$) is in EXPTIME and PSPACE-hard.*

The EXPTIME-membership stated in this theorem uses tree automata techniques. The adversarial threshold synthesis problem is studied in [Filiot et al. 2020; Balachander et al. 2020] for games with mean-payoff functions and with *two players*. An example of game is given in [Filiot et al. 2020] such that the follower has no NE (i.e., no best response, as there is only one follower) to respond to the strategy announced by the leader. Therefore two notions of ϵ -best responses are respectively proposed in [Filiot et al. 2020] and [Balachander et al. 2020] and in both variants the adversarial threshold synthesis problem (with strict inequality $> \mu$) is proved to be in NP.

5.2. One Follower with Several Payoff Functions

We continue with the case of only one follower modeling the environment however with *several payoff functions*, one function for each component of the environment. After responding to the leader with his own strategy, the follower receives a tuple of payoffs in the corresponding outcome. Rationality of the follower is encoded by the fact that he only responds in such a way to receive a *Pareto-optimal* tuple of payoffs, given the strategy announced by the leader. This setting encompasses scenarios where, for instance, several components can collaborate and agree on trade-offs. The goal of the leader is therefore to announce a strategy that guarantees his own payoff to be larger than a given threshold, whatever the response of the follower which ensures him a Pareto-optimal tuple of payoffs [Bruyère et al. 2021a].

In this context, we consider games with *two* players (instead of several players) such that player 0 is the leader with one payoff function pay_0 and player 1 is the follower with *several* payoff functions $(\text{pay}_i)_{i \in \{1, \dots, n\}}$. We denote by \leq the partial order on pairs of n -tuples of payoffs $p = (p_1, \dots, p_n)$, $p' = (p'_1, \dots, p'_n)$ such that $p \leq p'$ if and only if $p_i \leq p'_i$ for all $i \in \{1, \dots, n\}$. Given a strategy σ_0 for player 0, we denote P_{σ_0} the set of n -tuples of payoffs (for player 1) of plays compatible with σ_0 that are *Pareto-optimal* with respect to \leq .¹²

$$P_{\sigma_0} = \sup\{(\text{pay}_i(\rho))_{i \in \{1, \dots, n\}} \in \mathbb{R}^n \mid \rho \text{ is compatible with } \sigma_0\} \quad (3)$$

We say that P_{σ_0} is *achievable* if it is not empty and for each $p \in P_{\sigma_0}$, there exists a play ρ compatible with σ_0 such that $p = (\text{pay}_i(\rho))_{i \in \{1, \dots, n\}}$. Finite-range payoff functions (thus in particular Boolean payoff functions) always yield an achievable set P_{σ_0} . For games with mean-payoff functions, the set P_{σ_0} may be empty as shown in the example from [Filiot et al. 2020] that we mentioned at the end of the previous subsection. Notice that this example uses an infinite-memory strategy σ_0 ; indeed the set P_{σ_0} is always achievable when σ_0 has finite memory [Chatterjee et al. 2010b].

In the context of games with Boolean payoff functions, the following notion of rational response of the follower and the related threshold synthesis problem are introduced in [Bruyère et al. 2021a].

Definition 5.7. Let (G, v_0) be an initialized game with Boolean payoff functions and let σ_0 be a strategy for player 0. A σ_0 -*Pareto-optimal* (σ_0 -*PO*) strategy is a strategy σ_1 for player 1 such that $(\text{pay}_i(\rho))_{i \in \{1, \dots, n\}} \in P_{\sigma_0}$ where $\rho = \langle (\sigma_0, \sigma_1) \rangle_{v_0}$.

Definition 5.8. Let (G, v_0) be an initialized game with Boolean payoff functions and let $\mu \in \mathbb{Q}$ be a threshold. The *Pareto-optimal threshold synthesis problem* is to decide

¹²Notice that P_{σ_0} only refers to n -tuples of payoffs received by player 1 and not to the payoff received by player 0.




Fig. 9. A game where player 1 has $n = 3$ payoff functions.

whether there exists a strategy σ_0 for player 0 such that for all σ_0 -PO strategy σ_1 for player 1, we have $\mu \leq \text{pay}_0(\langle\langle\sigma_0, \sigma_1\rangle\rangle_{v_0})$.

The next example shows that player 0 sometimes needs to play a complex strategy σ_0 , i.e., not memoryless, in order to have a solution to the previous problem [Bruyère et al. 2021a].

Example 5.9. Consider the game depicted in Figure 9 in which player 0 has one Boolean payoff function and player 1 has $n = 3$ Boolean payoff functions. All these payoff functions are prefix-independent such that the tuple of payoffs¹³ of plays eventually looping on vertices v_1 , v_4 , v_6 or v_7 is displayed in the arena next to those vertices, and the tuple of payoffs of plays eventually looping on cycle (v_3v_5) is equal to $(0, (0, 1, 0))$. The leader is again the player that controls the circle vertices. Let us study the Pareto-optimal threshold synthesis problem with bound $\mu = 1$ (player 0 wants to see his objective being satisfied).

Consider the memoryless strategy σ_0 of player 0 such that $v_3 \rightarrow v_5$. The set of triples of payoffs for player 1 of plays compatible with σ_0 is equal to $\{(0, 0, 1), (0, 1, 0), (1, 0, 0), (0, 1, 1)\}$ and the set of those that are Pareto-optimal is $P_{\sigma_0} = \{(1, 0, 0), (0, 1, 1)\}$. Notice that play $\rho = v_0v_2v_4^\omega$ is compatible with σ_0 , has triple of payoffs $(1, 0, 0)$ for player 1 and payoff 0 for player 0. Strategy σ_0 is therefore not a solution to the Pareto-optimal threshold synthesis problem. In this game, there is only one other memoryless strategy for player 0 (such that $v_3 \rightarrow v_7$). One can verify that it is again not a solution to the problem.

We can however define a finite-memory strategy σ_0 such that $\sigma_0(v_0v_2v_3) = v_5$ and $\sigma_0(v_0v_2v_3v_5v_3) = v_7$ and show that it is now a solution to the problem. Indeed, the set of Pareto-optimal triples of payoffs is $P_{\sigma_0} = \{(0, 1, 1), (1, 1, 0)\}$ and player 0 receive payoff 1 for every play compatible with σ_0 whose triples of payoffs for player 1 is in this set.

We have the following result for the Pareto-optimal threshold synthesis problem.

THEOREM 5.10 ([BRUYÈRE ET AL. 2021A]). *For games with parity objectives, the Pareto-optimal threshold synthesis problem (with threshold $\mu = 1$) is NEXPTIME-complete. It is fixed-parameter tractable (in parameters equal to the number of objectives of the follower and the maximal priority used in each priority function).*

The FPT complexity result of this theorem is obtained thanks to a reduction to a so-called Challenger-Prover zero-sum game such that Prover tries to show the existence of a solution to the problem, while Challenger tries to disprove it. Classical tree automata

¹³We denote such a tuple by $(p_0, (p_1, p_2, p_3))$ to distinguish the payoff p_0 of player 0 from the triple of payoffs (p_1, p_2, p_3) of player 1.

techniques seem inadequate in this context. The results of Theorem 5.10 also hold for games with reachability objectives [Bruyère et al. 2021a].

For games with quantitative payoff functions (like mean-payoff), as the set P_{σ_0} given in (3) may be empty, adequate variants of the Pareto-optimal threshold synthesis problem need to be introduced and further studied.

6. RELATED WORK

In this section, we indicate additional pointers to related work.

NEs capture rational behaviors when the players only care about their own payoff and are indifferent to the payoff of the other players. The notion of *secure equilibrium* is introduced in [Chatterjee et al. 2006] for two-player games with lexicographic objectives: each player first tries to maximize his own payoff and then to minimize the payoff of the other player. It is proved in [Chatterjee et al. 2006] that every two-player game with Borel Boolean objectives has a secure equilibrium; this result is generalized to multi-player games in [De Pril et al. 2014]. Concerning games with quantitative objectives, general hypotheses are provided in [De Pril et al. 2014] that guarantee the existence of a secure equilibrium; the threshold synthesis problem for secure equilibria is studied in [Bruyère et al. 2014]. A variant of secure equilibrium, called *Doomsday equilibrium*, is studied in [Chatterjee et al. 2017] for games with ω -regular objectives with both perfect and imperfect information. The concept of *assume-guarantee synthesis* inspired from secure equilibria is introduced in [Chatterjee and Henzinger 2007] and is further studied for digital contract signing in [Chatterjee and Raman 2014] and for concurrent reactive programs with partial information in [Bloem et al. 2015]. NEs with *lexicographic* objectives in concurrent games are investigated in [Gutierrez et al. 2017] such that each player first prefers to satisfy his Büchi objective and then to minimize his mean-payoff.

SPEs are immune of the problem of non-credible threats. Another concept that avoids this problem is studied in [Berwanger 2007] with the concept of *admissible* strategies which are strategies not dominated by any other strategies. Dominated strategies are eliminated for each player, step by step (as strategies that were not dominated may become dominated), until the process stabilizes. The algorithmic properties of this concept is studied in [Brenquier et al. 2014] for games with ω -regular objectives. The *assume-admissible synthesis* related to the first iteration of this elimination procedure is studied in [Brenquier et al. 2016; Basset et al. 2017a; Brenquier et al. 2017]. The notion of admissibility in timed games is investigated in [Basset et al. 2017b].

The notion of *weak SPE* is introduced in [Brihaye et al. 2015] as a useful concept to study the existence SPEs (possibly with constraints) in quantitative reachability games. The exact complexity class of this problem is later settled in [Brihaye et al. 2020]. While an SPE must be resistant to any unilateral deviation of one player, a weak SPE must be resistant to deviations restricted to deviating strategies that differ from the original one on a finite number of histories only. Weak SPEs are easier to characterize and to manipulate algorithmically and they coincide with SPEs for large classes of games, for instance for games played on finite trees [Kuhn 1953] or with lower-semicontinuous payoff functions [Flesch et al. 2010]. In [Bruyère et al. 2021b] the authors study general conditions that guarantee the existence of a weak SPE for quantitative games. The constrained synthesis problem is studied in [Brihaye et al. 2018; Goeminne 2020] for different kinds of ω -regular objectives.

NEs are also investigated in the more general framework of *concurrent* games, where the players make their choices concurrently instead of in a turn-based way. In this setting, the constrained synthesis problem for NEs is undecidable for concurrent deterministic games with terminal-reward payoffs and with randomized strategies [Bouyer

et al. 2014]. When the NEs are composed of pure strategies, this problem is solved in [Klimos et al. 2012] for quantitative reachability games thanks to the effective representation of the outcomes of NEs as the language accepted by some Büchi automaton. It is also solved in [Bouyer et al. 2015] for all classical ω -regular objectives thanks to the concept of *suspect game*. The latter game is a two-player zero-sum game where one player tries to have the players play an NE and the other player tries to disprove this attempt by finding a profitable deviation for one of the original players. The existence of NEs and SPEs in multi-player *timed games* with reachability objectives is studied in [Bouyer et al. 2010; Brihaye and Goeminne 2020]. The existence problem for NEs is solved in [Gutierrez et al. 2020] for concurrent games where the objectives of the players are specified by LTL formulas. The proposed algorithm reduces the problem to solving a collection of parity games and has been implemented in the Equilibrium Verification Environment (EVE) system.

When shifting to *stochastic* games, we are again faced with undecidability. For instance, for turn-based stochastic games with ω -regular objectives, it is undecidable whether there exists an NE where a given player satisfies his objective with probability 1 [Ummels and Wojtczak 2011a]. In [Chatterjee et al. 2004], the authors give an algorithm for computing an NE in stochastic games with ω -regular objectives. The complexity of finding an NE satisfying certain constraints is analyzed in [Ummels and Wojtczak 2011a].

Non-zero-sum games are also examined in the context of *imperfect information*. In [Gutierrez et al. 2018], the authors study the existence problem for NEs in concurrent games with imperfect information and LTL objectives. The notions of admissible strategy in [Brennguier et al. 2017] and of Doomsday equilibrium in [Filiot et al. 2018] are studied in the context of games with imperfect information. Rational synthesis is investigated in [Filiot et al. 2018] in the imperfect information setting. An extension of the suspect game is proposed in the setting of games with imperfect monitoring in [Bouyer 2018].

Popular extensions of temporal logics have been introduced to express the existence of strategies for the components of a system, like ATL [Alur et al. 2002], a well-known extension of CTL. To be able to reason about solution concepts like NEs, *Strategy Logics* (SL) has been introduced in [Chatterjee et al. 2010; Mogavero et al. 2014]. An extension of SL is introduced in [Aminof et al. 2019] for stochastic systems and in [Berthon et al. 2021] for systems with imperfect information. Let us also mention a quantitative extension of SL proposed in [Bouyer et al. 2020]. The complexity of *rational synthesis* is studied thanks to adequate variants of SL in [Fisman et al. 2010; Kupferman et al. 2016] for games with objectives defined by LTL formulas and such that the components of the environment behave according to an NE, an SPE, or using dominant strategies.

Other kinds of constraints on NEs and SPEs are investigated like the *social welfare* equal to the sum of the payoffs for all players, see e.g. in [Conitzer and Sandholm 2003; Bouyer et al. 2014; Brihaye et al. 2019].

There are several works about the existence of solution concepts like NEs and SPEs for any (possibly infinite) number of players and actions. Notice that we are quickly faced with non-existence results. For instance, there is no NE in the simple one-player game where the player has to play once by choosing an action among his infinite action set $\{1, 2, \dots\}$ and action n gives him payoff $1 - \frac{1}{n}$. For general conditions implying NE or SPE existence, see e.g. [Roux 2013; Flesch and Predtetchinski 2016; 2017; Cingiz et al. 2020; Bruyère et al. 2021b; Roux and Pauly 2021].

7. CONCLUSION

In this introductory survey, we presented both classical and recent results about the threshold synthesis problem for multi-player non-zero-sum games played on graphs.

We focused on two important payoff functions: on parity objectives for the Boolean case and on mean-payoff functions for the quantitative case. The threshold synthesis problem has been studied according to different solution concepts. We first recalled the concepts of NE and SPE from game theory as well as the known results to the threshold synthesis problem. In this context we presented a very useful characterization of NE outcomes and the one obtained recently for SPEs. We then presented Stackelberg games well-suited for the synthesis of reactive systems. The environment is modeled either as several players each of them with their own objective or as a single player with several objectives (one for each of his components). Illustrative examples were provided all along the survey and a last section was completely dedicated to related work.

ACKNOWLEDGMENTS

We would like to thank Raphaël Berthon, Thomas Brihaye, Emmanuel Filiot, János Flesch, Aline Goeminne, Jean-François Raskin, Gatan Staquet, and Clément Tamines for their useful discussions and comments that helped us to improve the presentation of this article.

REFERENCES

- Rajeev Alur, Thomas A. Henzinger, and Orna Kupferman. 2002. Alternating-time temporal logic. *J. ACM* 49, 5 (2002), 672–713. DOI: <http://dx.doi.org/10.1145/585265.585270>
- Benjamin Aminof, Marta Kwiatkowska, Bastien Maubert, Aniello Murano, and Sasha Rubin. 2019. Probabilistic Strategy Logic. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, Sarit Kraus (Ed.). ijcai.org, 32–38. DOI: <http://dx.doi.org/10.24963/ijcai.2019/5>
- Mrudula Balachander, Shibashis Guha, and Jean-François Raskin. 2020. Fragility and Robustness in Mean-payoff Adversarial Stackelberg Games. *CoRR* abs/2007.07209 (2020). <https://arxiv.org/abs/2007.07209>
- Nicolas Basset, Gilles Geeraerts, Jean-François Raskin, and Ocan Sankur. 2017a. Admissibility in Concurrent Games. In *44th International Colloquium on Automata, Languages, and Programming, ICALP 2017, July 10-14, 2017, Warsaw, Poland (LIPIcs)*, Ioannis Chatzigiannakis, Piotr Indyk, Fabian Kuhn, and Anca Muscholl (Eds.), Vol. 80. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 123:1–123:14. DOI: <http://dx.doi.org/10.4230/LIPIcs.ICALP.2017.123>
- Nicolas Basset, Jean-François Raskin, and Ocan Sankur. 2017b. Admissible Strategies in Timed Games. In *Models, Algorithms, Logics and Tools - Essays Dedicated to Kim Guldstrand Larsen on the Occasion of His 60th Birthday (Lecture Notes in Computer Science)*, Luca Aceto, Giorgio Bacci, Giovanni Bacci, Anna Ingólfssdóttir, Axel Legay, and Radu Mardare (Eds.), Vol. 10460. Springer, 403–425. DOI: http://dx.doi.org/10.1007/978-3-319-63121-9_20
- Raphaël Berthon, Bastien Maubert, Aniello Murano, Sasha Rubin, and Moshe Y. Vardi. 2021. Strategy Logic with Imperfect Information. *ACM Trans. Comput. Log.* 22, 1 (2021), 5:1–5:51. DOI: <http://dx.doi.org/10.1145/3427955>
- Dietmar Berwanger. 2007. Admissibility in Infinite Games. In *STACS 2007, 24th Annual Symposium on Theoretical Aspects of Computer Science, Aachen, Germany, February 22-24, 2007, Proceedings (Lecture Notes in Computer Science)*, Wolfgang Thomas and Pascal Weil (Eds.), Vol. 4393. Springer, 188–199. DOI: http://dx.doi.org/10.1007/978-3-540-70918-3_17
- Roderick Bloem, Krishnendu Chatterjee, Swen Jacobs, and Robert Könighofer. 2015. Assume-Guarantee Synthesis for Concurrent Reactive Programs with Partial Information. In *Tools and Algorithms for the Construction and Analysis of Systems - 21st International Conference, TACAS 2015, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2015, London, UK, April 11-18, 2015. Proceedings (Lecture Notes in Computer Science)*, Christel Baier and Cesare Tinelli (Eds.), Vol. 9035. Springer, 517–532. DOI: http://dx.doi.org/10.1007/978-3-662-46681-0_50
- Udi Boker, Thomas A. Henzinger, and Jan Otop. 2015. The Target Discounted-Sum Problem. In *30th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2015, Kyoto, Japan, July 6-10, 2015*. IEEE Computer Society, 750–761. DOI: <http://dx.doi.org/10.1109/LICS.2015.74>
- Patricia Bouyer. 2018. Games on Graphs with a Public Signal Monitoring. In *Foundations of Software Science and Computation Structures - 21st International Conference, FOSSACS 2018, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2018, Thessaloniki, Greece, April 14-20, 2018, Proceedings (Lecture Notes in Computer Science)*, Christel Baier and Ugo Dal Lago (Eds.), Vol. 10803. Springer, 530–547. DOI: http://dx.doi.org/10.1007/978-3-319-89366-2_29

- Patricia Bouyer, Romain Brenguier, and Nicolas Markey. 2010. Nash Equilibria for Reachability Objectives in Multi-player Timed Games. In *CONCUR 2010 - Concurrency Theory, 21th International Conference, CONCUR 2010, Paris, France, August 31-September 3, 2010. Proceedings (Lecture Notes in Computer Science)*, Paul Gastin and François Laroussinie (Eds.), Vol. 6269. Springer, 192–206. DOI : http://dx.doi.org/10.1007/978-3-642-15375-4_14
- Patricia Bouyer, Romain Brenguier, Nicolas Markey, and Michael Ummels. 2015. Pure Nash Equilibria in Concurrent Deterministic Games. *Log. Methods Comput. Sci.* 11, 2 (2015). DOI : [http://dx.doi.org/10.2168/LMCS-11\(2:9\)2015](http://dx.doi.org/10.2168/LMCS-11(2:9)2015)
- Patricia Bouyer, Orna Kupferman, Nicolas Markey, Bastien Maubert, Aniello Murano, and Giuseppe Perelli. 2020. Reasoning About Quality and Fuzziness of Strategic Behaviours. In *ECAI 2020 - 24th European Conference on Artificial Intelligence, 29 August-8 September 2020, Santiago de Compostela, Spain, August 29 - September 8, 2020 - Including 10th Conference on Prestigious Applications of Artificial Intelligence (PAIS 2020) (Frontiers in Artificial Intelligence and Applications)*, Giuseppe De Giacomo, Alejandro Catalá, Bistra Dilkina, Michela Milano, Senén Barro, Alberto Bugarín, and Jérôme Lang (Eds.), Vol. 325. IOS Press, 2887–2888. DOI : <http://dx.doi.org/10.3233/FAIA200437>
- Patricia Bouyer, Nicolas Markey, and Daniel Stan. 2014. Mixed Nash Equilibria in Concurrent Terminal-Reward Games. In *34th International Conference on Foundation of Software Technology and Theoretical Computer Science, FSTTCS 2014, December 15-17, 2014, New Delhi, India (LIPIcs)*, Venkatesh Raman and S. P. Suresh (Eds.), Vol. 29. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 351–363. DOI : <http://dx.doi.org/10.4230/LIPIcs.FSTTCS.2014.351>
- Romain Brenguier, Lorenzo Clemente, Paul Hunter, Guillermo A. Pérez, Mickael Randour, Jean-François Raskin, Ocan Sankur, and Mathieu Sassolas. 2016. Non-Zero Sum Games for Reactive Synthesis. In *Language and Automata Theory and Applications - 10th International Conference, LATA 2016, Prague, Czech Republic, March 14-18, 2016, Proceedings (Lecture Notes in Computer Science)*, Adrian-Horia Dediu, Jan Janousek, Carlos Martín-Vide, and Bianca Truthe (Eds.), Vol. 9618. Springer, 3–23. DOI : http://dx.doi.org/10.1007/978-3-319-30000-9_1
- Romain Brenguier, Arno Pauly, Jean-François Raskin, and Ocan Sankur. 2017. Admissibility in Games with Imperfect Information (Invited Talk). In *28th International Conference on Concurrency Theory, CONCUR 2017, September 5-8, 2017, Berlin, Germany (LIPIcs)*, Roland Meyer and Uwe Nestmann (Eds.), Vol. 85. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2:1–2:23. DOI : <http://dx.doi.org/10.4230/LIPIcs.CONCUR.2017.2>
- Romain Brenguier, Guillermo A. Pérez, Jean-François Raskin, and Ocan Sankur. 2016. Admissibility in Quantitative Graph Games. In *36th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2016, December 13-15, 2016, Chennai, India (LIPIcs)*, Akash Lal, S. Akshay, Saket Saurabh, and Sandeep Sen (Eds.), Vol. 65. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 42:1–42:14. DOI : <http://dx.doi.org/10.4230/LIPIcs.FSTTCS.2016.42>
- Romain Brenguier, Jean-François Raskin, and Ocan Sankur. 2017. Assume-admissible synthesis. *Acta Informatica* 54, 1 (2017), 41–83. DOI : <http://dx.doi.org/10.1007/s00236-016-0273-2>
- Romain Brenguier, Jean-François Raskin, and Mathieu Sassolas. 2014. The complexity of admissibility in Omega-regular games. In *Joint Meeting of the Twenty-Third EACSL Annual Conference on Computer Science Logic (CSL) and the Twenty-Ninth Annual ACM/IEEE Symposium on Logic in Computer Science (LICS), CSL-LICS '14, Vienna, Austria, July 14 - 18, 2014*, Thomas A. Henzinger and Dale Miller (Eds.). ACM, 23:1–23:10. DOI : <http://dx.doi.org/10.1145/2603088.2603143>
- Léonard Brice, Jean-François Raskin, and Marie van den Bogaard. 2021. Subgame-perfect Equilibria in Mean-payoff Games. *CoRR* abs/2101.10685 (2021). <https://arxiv.org/abs/2101.10685>
- Thomas Brihaye, Véronique Bruyère, Aline Goeminne, and Jean-François Raskin. 2018. Constrained Existence Problem for Weak Subgame Perfect Equilibria with ω -Regular Boolean Objectives. In *Proceedings Ninth International Symposium on Games, Automata, Logics, and Formal Verification, GandALF 2018, Saarbrücken, Germany, 26-28th September 2018 (EPTCS)*, Andrea Orlandini and Martin Zimmermann (Eds.), Vol. 277. 16–29. DOI : <http://dx.doi.org/10.4204/EPTCS.277.2>
- Thomas Brihaye, Véronique Bruyère, Aline Goeminne, Jean-François Raskin, and Marie van den Bogaard. 2020. The Complexity of Subgame Perfect Equilibria in Quantitative Reachability Games. *Log. Methods Comput. Sci.* 16, 4 (2020). <https://lmcs.episciences.org/6883>
- Thomas Brihaye, Véronique Bruyère, Aline Goeminne, and Nathan Thomasset. 2019. On Relevant Equilibria in Reachability Games. In *Reachability Problems - 13th International Conference, RP 2019, Brussels, Belgium, September 11-13, 2019, Proceedings (Lecture Notes in Computer Science)*, Emmanuel Filiot, Raphaël M. Jungers, and Igor Potapov (Eds.), Vol. 11674. Springer, 48–62. DOI : http://dx.doi.org/10.1007/978-3-030-30806-3_5
- Thomas Brihaye, Véronique Bruyère, Noémie Meunier, and Jean-François Raskin. 2015. Weak Subgame Perfect Equilibria and their Application to Quantitative Reachability. In *24th EACSL An-*

- nual Conference on Computer Science Logic, CSL 2015, September 7-10, 2015, Berlin, Germany (LIPIcs), Stephan Kreutzer (Ed.), Vol. 41. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 504–518. DOI : <http://dx.doi.org/10.4230/LIPIcs.CSL.2015.504>
- Thomas Brihaye, Julie De Pril, and Sven Schewe. 2013. Multiplayer Cost Games with Simple Nash Equilibria. In *Logical Foundations of Computer Science, International Symposium, LFCS 2013, San Diego, CA, USA, January 6-8, 2013. Proceedings (Lecture Notes in Computer Science)*, Sergei N. Artëmov and Anil Nerode (Eds.), Vol. 7734. Springer, 59–73. DOI : http://dx.doi.org/10.1007/978-3-642-35722-0_5
- Thomas Brihaye, Amit Kumar Dhar, Gilles Geeraerts, Axel Haddad, and Benjamin Monmege. 2016. Efficient Energy Distribution in a Smart Grid using Multi-Player Games. In *Proceedings Cassting Workshop on Games for the Synthesis of Complex Systems and 3rd International Workshop on Synthesis of Complex Parameters, Cassting/SynCoP 2016, Eindhoven, The Netherlands, April 2-3, 2016 (EPTCS)*, Thomas Brihaye, Benoît Delahaye, Loig Jezequel, Nicolas Markey, and Jiri Srba (Eds.), Vol. 220. 1–12. DOI : <http://dx.doi.org/10.4204/EPTCS.220.1>
- Thomas Brihaye and Aline Goeminne. 2020. On Subgame Perfect Equilibria in Turn-Based Reachability Timed Games. In *Formal Modeling and Analysis of Timed Systems - 18th International Conference, FORMATS 2020, Vienna, Austria, September 1-3, 2020, Proceedings (Lecture Notes in Computer Science)*, Nathalie Bertrand and Nils Jansen (Eds.), Vol. 12288. Springer, 94–110. DOI : http://dx.doi.org/10.1007/978-3-030-57628-8_6
- Véronique Bruyère. 2017. Computer Aided Synthesis: A Game-Theoretic Approach. In *Developments in Language Theory - 21st International Conference, DLT 2017, Liège, Belgium, August 7-11, 2017, Proceedings (Lecture Notes in Computer Science)*, Émilie Charlier, Julien Leroy, and Michel Rigo (Eds.), Vol. 10396. Springer, 3–35. DOI : http://dx.doi.org/10.1007/978-3-319-62809-7_1
- Véronique Bruyère, Noémie Meunier, and Jean-François Raskin. 2014. Secure equilibria in weighted games. In *Joint Meeting of the Twenty-Third EACSL Annual Conference on Computer Science Logic (CSL) and the Twenty-Ninth Annual ACM/IEEE Symposium on Logic in Computer Science (LICS), CSL-LICS '14, Vienna, Austria, July 14 - 18, 2014*, Thomas A. Henzinger and Dale Miller (Eds.). ACM, 26:1–26:26. DOI : <http://dx.doi.org/10.1145/2603088.2603109>
- Véronique Bruyère, Jean-François Raskin, and Clément Tamines. 2021a. Stackelberg-Pareto Synthesis. *CoRR* abs/2102.08925 (2021). <https://arxiv.org/abs/2102.08925>
- Véronique Bruyère, Stéphane Le Roux, Arno Pauly, and Jean-François Raskin. 2021b. On the existence of weak subgame perfect equilibria. *Inf. Comput.* 276 (2021), 104553. DOI : <http://dx.doi.org/10.1016/j.ic.2020.104553>
- Krishnendu Chatterjee, Laurent Doyen, Herbert Edelsbrunner, Thomas A. Henzinger, and Philippe Rannou. 2010b. Mean-Payoff Automaton Expressions. In *CONCUR 2010 - Concurrency Theory, 21th International Conference, CONCUR 2010, Paris, France, August 31-September 3, 2010. Proceedings (Lecture Notes in Computer Science)*, Paul Gastin and François Laroussinie (Eds.), Vol. 6269. Springer, 269–283. DOI : http://dx.doi.org/10.1007/978-3-642-15375-4_19
- Krishnendu Chatterjee, Laurent Doyen, Emmanuel Filiot, and Jean-François Raskin. 2017. Doomsday equilibria for omega-regular games. *Inf. Comput.* 254 (2017), 296–315. DOI : <http://dx.doi.org/10.1016/j.ic.2016.10.012>
- Krishnendu Chatterjee, Laurent Doyen, and Thomas A. Henzinger. 2010a. Quantitative languages. *ACM Trans. Comput. Log.* 11, 4 (2010), 23:1–23:38. DOI : <http://dx.doi.org/10.1145/1805950.1805953>
- Krishnendu Chatterjee and Thomas A. Henzinger. 2007. Assume-Guarantee Synthesis. In *Tools and Algorithms for the Construction and Analysis of Systems, 13th International Conference, TACAS 2007, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2007 Braga, Portugal, March 24 - April 1, 2007, Proceedings (Lecture Notes in Computer Science)*, Orna Grumberg and Michael Huth (Eds.), Vol. 4424. Springer, 261–275. DOI : http://dx.doi.org/10.1007/978-3-540-71209-1_21
- Krishnendu Chatterjee, Thomas A. Henzinger, and Marcin Jurdzinski. 2006. Games with secure equilibria. *Theor. Comput. Sci.* 365, 1-2 (2006), 67–82. DOI : <http://dx.doi.org/10.1016/j.tcs.2006.07.032>
- Krishnendu Chatterjee, Thomas A. Henzinger, and Nir Piterman. 2010. Strategy logic. *Inf. Comput.* 208, 6 (2010), 677–693. DOI : <http://dx.doi.org/10.1016/j.ic.2009.07.004>
- Krishnendu Chatterjee, Rupak Majumdar, and Marcin Jurdzinski. 2004. On Nash Equilibria in Stochastic Games. In *Computer Science Logic, 18th International Workshop, CSL 2004, 13th Annual Conference of the EACSL, Karpacz, Poland, September 20-24, 2004, Proceedings (Lecture Notes in Computer Science)*, Jerzy Marcinkowski and Andrzej Tarlecki (Eds.), Vol. 3210. Springer, 26–40. DOI : http://dx.doi.org/10.1007/978-3-540-30124-0_6
- Krishnendu Chatterjee and Vishwanath Raman. 2014. Assume-guarantee synthesis for digital contract signing. *Formal Aspects Comput.* 26, 4 (2014), 825–859. DOI : <http://dx.doi.org/10.1007/s00165-013-0283-6>

- Kutay Cingiz, János Flesch, P. Jean-Jacques Herings, and Arkadi Predtetchinski. 2020. Perfect information games where each player acts only once. *Economic Theory* 69 (2020), 965–985. <https://doi.org/10.1007/s00199-019-01199-3>
- Rodica Condurache, Emmanuel Filiot, Raffaella Gentilini, and Jean-François Raskin. 2016. The Complexity of Rational Synthesis. In *43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016, July 11-15, 2016, Rome, Italy (LIPIcs)*, Ioannis Chatzigiannakis, Michael Mitzenmacher, Yuval Rabani, and Davide Sangiorgi (Eds.), Vol. 55. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 121:1–121:15. DOI : <http://dx.doi.org/10.4230/LIPIcs.ICALP.2016.121>
- Vincent Conitzer and Tuomas Sandholm. 2003. Complexity Results about Nash Equilibria. In *IJCAI-03, Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence, Acapulco, Mexico, August 9-15, 2003*, Georg Gottlob and Toby Walsh (Eds.). Morgan Kaufmann, 765–771. <http://ijcai.org/Proceedings/03/Papers/111.pdf>
- Julie De Pril, János Flesch, Jeroen Kuipers, Gijs Schoenmakers, and Koos Vrieze. 2014. Existence of Secure Equilibrium in Multi-player Games with Perfect Information. In *Mathematical Foundations of Computer Science 2014 - 39th International Symposium, MFCS 2014, Budapest, Hungary, August 25-29, 2014. Proceedings, Part II (Lecture Notes in Computer Science)*, Erzsébet Csuhaj-Varjú, Martin Dietzfelbinger, and Zoltán Ésik (Eds.), Vol. 8635. Springer, 213–225. DOI : http://dx.doi.org/10.1007/978-3-662-44465-8_19
- E. Allen Emerson and Charanjit S. Jutla. 1991. Tree Automata, Mu-Calculus and Determinacy (Extended Abstract). In *32nd Annual Symposium on Foundations of Computer Science, San Juan, Puerto Rico, 1-4 October 1991*. IEEE Computer Society, 368–377. DOI : <http://dx.doi.org/10.1109/SFCS.1991.185392>
- Jerzy Filar and Koos Vrieze. 1997. *Competitive Markov Decision Processes*. Springer.
- Emmanuel Filiot, Raffaella Gentilini, and Jean-François Raskin. 2018. Rational Synthesis Under Imperfect Information. In *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2018, Oxford, UK, July 09-12, 2018*, Anuj Dawar and Erich Grädel (Eds.). ACM, 422–431. DOI : <http://dx.doi.org/10.1145/3209108.3209164>
- Emmanuel Filiot, Raffaella Gentilini, and Jean-François Raskin. 2020. The Adversarial Stackelberg Value in Quantitative Games. In *47th International Colloquium on Automata, Languages, and Programming, ICALP 2020, July 8-11, 2020, Saarbrücken, Germany (Virtual Conference) (LIPIcs)*, Artur Czumaj, Anuj Dawar, and Emanuela Merelli (Eds.), Vol. 168. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 127:1–127:18. DOI : <http://dx.doi.org/10.4230/LIPIcs.ICALP.2020.127>
- Dana Fisman, Orna Kupferman, and Yoav Lustig. 2010. Rational Synthesis. In *Tools and Algorithms for the Construction and Analysis of Systems, 16th International Conference, TACAS 2010, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2010, Paphos, Cyprus, March 20-28, 2010. Proceedings (Lecture Notes in Computer Science)*, Javier Esparza and Rupak Majumdar (Eds.), Vol. 6015. Springer, 190–204. DOI : http://dx.doi.org/10.1007/978-3-642-12002-2_16
- János Flesch, Jeroen Kuipers, Ayala Mashiah-Yaakovi, Gijs Schoenmakers, Eilon Solan, and Koos Vrieze. 2010. Perfect-Information Games with Lower-Semicontinuous Payoffs. *Math. Oper. Res.* 35, 4 (2010), 742–755. DOI : <http://dx.doi.org/10.1287/moor.1100.0469>
- János Flesch and Arkadi Predtetchinski. 2016. Subgame-Perfect ϵ -Equilibria in Perfect Information Games with Common Preferences at the Limit. *Math. Oper. Res.* 41, 4 (2016), 1208–1221. DOI : <http://dx.doi.org/10.1287/moor.2015.0774>
- János Flesch and Arkadi Predtetchinski. 2017. A Characterization of Subgame-Perfect Equilibrium Plays in Borel Games of Perfect Information. *Math. Oper. Res.* 42, 4 (2017), 1162–1179. DOI : <http://dx.doi.org/10.1287/moor.2016.0843>
- Drew Fudenberg and David Levine. 1983. Subgame-perfect equilibria of finite- and infinite-horizon games. *Journal of Economic Theory* 31 (1983), 251–268.
- Hugo Gimbert and Wiesław Zielonka. 2005. Games Where You Can Play Optimally Without Any Memory. In *CONCUR 2005 - Concurrency Theory, 16th International Conference, CONCUR 2005, San Francisco, CA, USA, August 23-26, 2005, Proceedings (Lecture Notes in Computer Science)*, Martín Abadi and Luca de Alfaro (Eds.), Vol. 3653. Springer, 428–442. DOI : http://dx.doi.org/10.1007/11539452_33
- Aline Goeminne. 2020. Constrained existence of weak subgame perfect equilibria in multiplayer Büchi games. *Inf. Process. Lett.* 163 (2020), 105996. DOI : <http://dx.doi.org/10.1016/j.ipl.2020.105996>
- Erich Grädel, Wolfgang Thomas, and Thomas Wilke (Eds.). 2002. *Automata, Logics, and Infinite Games: A Guide to Current Research [outcome of a Dagstuhl seminar, February 2001]*. Lecture Notes in Computer Science, Vol. 2500. Springer. DOI : <http://dx.doi.org/10.1007/3-540-36387-4>
- Erich Grädel and Michael Ummels. 2008. Solution Concepts and Algorithms for Infinite Multiplayer Games. In *New Perspectives on Games and Interaction*, Vol. 4. Amsterdam University Press, 151–178.

- Anshul Gupta and Sven Schewe. 2014. Quantitative Verification in Rational Environments. In *21st International Symposium on Temporal Representation and Reasoning, TIME 2014, Verona, Italy, September 8-10, 2014*, Amedeo Cesta, Carlo Combi, and François Laroussinie (Eds.). IEEE Computer Society, 123–131. DOI: <http://dx.doi.org/10.1109/TIME.2014.9>
- Anshul Gupta, Sven Schewe, and Dominik Wojtczak. 2015. Making the Best of Limited Memory in Multi-Player Discounted Sum Games. In *Proceedings Sixth International Symposium on Games, Automata, Logics and Formal Verification, GandALF 2015, Genoa, Italy, 21-22nd September 2015 (EPTCS)*, Javier Esparza and Enrico Tronci (Eds.), Vol. 193. 16–30. DOI: <http://dx.doi.org/10.4204/EPTCS.193.2>
- Julian Gutierrez, Aniello Murano, Giuseppe Perelli, Sasha Rubin, and Michael J. Wooldridge. 2017. Nash Equilibria in Concurrent Games with Lexicographic Preferences. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, Carles Sierra (Ed.). ijcai.org, 1067–1073. DOI: <http://dx.doi.org/10.24963/ijcai.2017/148>
- Julian Gutierrez, Muhammad Najib, Giuseppe Perelli, and Michael J. Wooldridge. 2020. Automated temporal equilibrium analysis: Verification and synthesis of multi-player games. *Artif. Intell.* 287 (2020), 103353. DOI: <http://dx.doi.org/10.1016/j.artint.2020.103353>
- Julian Gutierrez, Giuseppe Perelli, and Michael J. Wooldridge. 2018. Imperfect information in Reactive Modules games. *Inf. Comput.* 261 (2018), 650–675. DOI: <http://dx.doi.org/10.1016/j.ic.2018.02.023>
- Christopher Harris. 1985. Existence and characterization of perfect equilibrium in games of perfect information. *Econometrica* 53 (1985), 613–628.
- Miroslav Klimos, Kim G. Larsen, Filip Stefanak, and Jeppe Thaarup. 2012. Nash Equilibria in Concurrent Priced Games. In *Language and Automata Theory and Applications - 6th International Conference, LATA 2012, A Coruña, Spain, March 5-9, 2012. Proceedings (Lecture Notes in Computer Science)*, Adrian-Horia Dediu and Carlos Martín-Vide (Eds.), Vol. 7183. Springer, 363–376. DOI: http://dx.doi.org/10.1007/978-3-642-28332-1_31
- Harold W. Kuhn. 1953. Extensive games and the problem of information. *Classics in Game Theory* (1953), 46–68.
- Orna Kupferman, Giuseppe Perelli, and Moshe Y. Vardi. 2016. Synthesis with rational environments. *Ann. Math. Artif. Intell.* 78, 1 (2016), 3–20. DOI: <http://dx.doi.org/10.1007/s10472-016-9508-8>
- Fabio Mogavero, Aniello Murano, Giuseppe Perelli, and Moshe Y. Vardi. 2014. Reasoning About Strategies: On the Model-Checking Problem. *ACM Trans. Comput. Log.* 15, 4 (2014), 34:1–34:47. DOI: <http://dx.doi.org/10.1145/2631917>
- John F. Nash. 1950. Equilibrium points in n -person games. In *PNAS*, Vol. 36. National Academy of Sciences, 48–49.
- Martin J. Osborne and Ariel Rubinstein. 1994. *A course in Game Theory*. MIT Press, Cambridge, MA.
- Roger A. Purves and William D. Sudderth. 2011. Perfect Information Games with Upper Semicontinuous Payoffs. *Math. Oper. Res.* 36, 3 (2011), 468–473. DOI: <http://dx.doi.org/10.1287/moor.1110.0504>
- Stéphane Le Roux. 2013. Infinite sequential Nash equilibrium. *Log. Methods Comput. Sci.* 9, 2 (2013). DOI: [http://dx.doi.org/10.2168/LMCS-9\(2:3\)2013](http://dx.doi.org/10.2168/LMCS-9(2:3)2013)
- Stéphane Le Roux and Arno Pauly. 2021. Equilibria in multi-player multi-outcome infinite sequential games. *Inf. Comput.* 276 (2021), 104557. DOI: <http://dx.doi.org/10.1016/j.ic.2020.104557>
- Reinhard Selten. 1965. Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit. *Zeitschrift für die gesamte Staatswissenschaft* 121 (1965), 301–324 and 667–689.
- Eilon Solan and Nicolas Vieille. 2003. Deterministic multi-player Dynkin games. *Journal of Mathematical Economics* 39 (2003), 911–929.
- Michael Ummels. 2006. Rational Behaviour and Strategy Construction in Infinite Multiplayer Games. In *FSTTCS 2006: Foundations of Software Technology and Theoretical Computer Science, 26th International Conference, Kolkata, India, December 13-15, 2006, Proceedings (Lecture Notes in Computer Science)*, S. Arun-Kumar and Naveen Garg (Eds.), Vol. 4337. Springer, 212–223. DOI: http://dx.doi.org/10.1007/11944836_21
- Michael Ummels. 2008. The Complexity of Nash Equilibria in Infinite Multiplayer Games. In *Foundations of Software Science and Computational Structures, 11th International Conference, FOSSACS 2008, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2008, Budapest, Hungary, March 29 - April 6, 2008. Proceedings (Lecture Notes in Computer Science)*, Roberto M. Amadio (Ed.), Vol. 4962. Springer, 20–34. DOI: http://dx.doi.org/10.1007/978-3-540-78499-9_3
- Michael Ummels and Dominik Wojtczak. 2011a. The Complexity of Nash Equilibria in Limit-Average Games. In *CONCUR 2011 - Concurrency Theory - 22nd International Conference, CONCUR 2011, Aachen, Germany, September 6-9, 2011. Proceedings (Lecture Notes in Computer Science)*, Joost-Pieter Katoen and Barbara König (Eds.), Vol. 6901. Springer, 482–496. DOI: http://dx.doi.org/10.1007/978-3-642-23217-6_32

- Michael Ummels and Dominik Wojtczak. 2011b. The Complexity of Nash Equilibria in Limit-Average Games. *CoRR* abs/1109.6220 (2011). <http://arxiv.org/abs/1109.6220>
- John von Neumann and Oskar Morgenstern. 1944. *Theory of Games and Economic Behavior*. Princeton University Press.
- Heinrich Freiherr von Stackelberg. 1937. *Marktform und Gleichgewicht*. Wien und Berlin, J. Springer.
- Uri Zwick and Mike Paterson. 1996. The Complexity of Mean Payoff Games on Graphs. *Theor. Comput. Sci.* 158, 1&2 (1996), 343–359. DOI : [http://dx.doi.org/10.1016/0304-3975\(95\)00188-3](http://dx.doi.org/10.1016/0304-3975(95)00188-3)