

A Note on Game Theory and Verification

Patricia Bouyer

► **To cite this version:**

Patricia Bouyer. A Note on Game Theory and Verification. 17th International Symposium Automated Technology for Verification and Analysis (ATVA'19), Oct 2019, Taipei, Taiwan. pp.3-22, 10.1007/978-3-030-31784-3_1 . hal-02352524

HAL Id: hal-02352524

<https://hal.archives-ouvertes.fr/hal-02352524>

Submitted on 6 Nov 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Note on Game Theory and Verification

Patricia Bouyer*

LSV, CNRS, ENS Paris-Saclay, Université Paris-Saclay, France

Abstract. We present some basics of game theory, focusing on matrix games. We then present the model of multiplayer stochastic concurrent games (with an underlying graph), which extends standard finite-state models used in verification in a multiplayer and concurrent setting; we explain why the basic theory cannot apply to that general model. We then focus on a very simple setting, and explain and give intuitions for the computation of Nash equilibria. We then give a number of undecidability results, giving limits to the approach. Finally we describe the suspect game construction, which (we believe) captures and explains well Nash equilibria and allow to compute them in many cases.

1 Introduction

Multiplayer concurrent games over graphs allow to model rich interactions between players. Those games are played as follows. In a state, each player chooses privately and independently an action, defining globally a move (one action per player); the next state of the game is then defined as the successor (on the graph) of the current state using that move; players continue playing from that new state, and form a(n infinite) play. Each player then gets a reward given by a payoff function (one function per player). In particular, objectives of the players may not be contradictory: those games are non-zero-sum games, contrary to two-player games used for controller or reactive synthesis [31, 24].

Using solution concepts borrowed from game theory, one can describe the interactions between the players, and in particular describe their rational behaviours. One of the most basic solution concepts is that of Nash equilibria [27]. A Nash equilibrium is a strategy profile where no player can improve her payoff by unilaterally changing her strategy. The outcome of a Nash equilibrium can therefore be seen as a rational behaviour of the system. While very much studied by game theoretists (e.g. over matrix games), such a concept (and variants thereof) has been only rather recently studied over games on graphs. Probably the first works in that direction are [17, 15, 32, 33].

Computing Nash equilibria requires to (i) find a good behaviour of the system; (ii) detect deviations from that behaviour, and identify deviating players (called deviators); (iii) punish them. Variants of Nash equilibria (like subgame-perfect equilibria, robust equilibria, *etc*) require slightly different ingredients, but they are mostly of a similar vein.

* Work supported by ERC project EQualIS between 2013 and 2019.

In this note, we first recall some basics of game theory over matrix games. Those games are not sufficient in a verification context: indeed, explicit states are very useful when modelling systems or programs, but are missing in matrix games. However stability notions like Nash equilibria or other solution concepts borrowed from game theory, are very relevant. We thus present the model of concurrent multiplayer games (played on graphs), which extends in a natural way standard models used in verification with multiplayer interactions. We explain how Nash equilibria can be characterized and computed in such general games. The ambition of this note is not to be a full survey of existing results, but rather to give simple explanations and intuitions; it gives formal tools to characterize and compute them and should help understanding simple interactions like Nash equilibria in rich games played on graphs.

Related notes are [31], which discussed the use of two-player zero-sum games in verification, and [22], which discussed solution concepts in multiplayer turn-based games on graphs.

Notations. If Σ is a finite alphabet, then Σ^+ (resp. Σ^ω) denotes the non-empty finite words (resp. infinite words) over Σ . If Γ is a finite set, then we note $\mathcal{D}(\Gamma)$ the set of probability distributions over Γ . We write \mathbb{R} for the set of real numbers.

2 Basics of game theory

In this section we present basic notions from game theory, which will be useful for our purpose. We refer the interested reader to the textbook [26].

A *matrix game* (aka *game in strategic form*) is a tuple $\mathcal{G} = (\mathcal{P}, \Sigma, (\text{payoff}_A)_{A \in \mathcal{P}})$ where $\mathcal{P} = \{A_1, \dots, A_k\}$ is a finite set of players, Σ is a finite set of actions, and for every $A \in \mathcal{P}$, $\text{payoff}_A: \Sigma^{\mathcal{P}} \rightarrow \mathbb{R}$ is a *payoff* (or *utility*) function for player A . In a deterministic setting, such a game is played as followed: independently and simultaneously, each player selects an action, resulting in a *move* (an element of $\Sigma^{\mathcal{P}}$), and each player gets the payoff specified in the game for that move. In a stochastic setting, each player selects a distribution over the actions, resulting in a distribution over the set of moves and an expected value for the payoff.

A *pure strategy* for player $A \in \mathcal{P}$ is the choice of an action $\sigma_A \in \Sigma$, while a *mixed strategy* for player A is a distribution $\sigma_A \in \mathcal{D}(\Sigma)$ over the set of possible actions. Obviously, a pure strategy is a specific case of a mixed strategy where only Dirac probability distributions can be used. We let $\sigma = (\sigma_A)_{A \in \mathcal{P}}$ be a (pure or mixed) *strategy profile* (that is, for every $A \in \mathcal{P}$, σ_A is (pure or mixed) strategy for player A). The probability of a move $m = (a_A)_{A \in \mathcal{P}} \in \Sigma^{\mathcal{P}}$ is written $\sigma(m)$ and defined by:

$$\sigma(m) = \prod_{A \in \mathcal{P}} \sigma_A(a_A)$$

Then, given a player $B \in \mathcal{P}$, the payoff of player B is given by the expected value of payoff_B under σ , that is:

$$\mathbb{E}^\sigma(\text{payoff}_B) = \sum_{m \in \Sigma^{\mathcal{P}}} \sigma(m) \cdot \text{payoff}_B(m)$$

Example 1 (The prisoner's dilemma). Two individuals have committed a crime and are apprehended. The prosecution lacks sufficient evidence to convict the two individuals on the principal charge, but they have enough to convict both on a lesser charge. The prosecutors offer each prisoner a bargain: without any communication between them, the two individuals are offered the opportunity to Betray the other by testifying that the other committed the crime and get (partly) immunity, or to stay Silent. The payoff of both players is summarized in the table below, where the higher is the payoff the shorter is the jail penalty:

		A_2	
		S	B
A_1	S	2, 2	0, 3
	B	3, 0	1, 1

In each cell of the table, the pair ' α_1, α_2 ' represents payoff α_1 (resp. α_2) for player A_1 (resp. A_2). The table can then be read as follows: if both players stay Silent (resp. Betray), then they both get payoff 2 (resp. 1). If only one prisoner Betrays, then he gets payoff 3 while the other prisoner gets payoff 0.

Example 2. We consider the following game (taken from [26, Example 4.34]) with two players, where payoffs are given in the next table:

		A_2	
		L	R
A_1	T	0, 0	2, 1
	B	3, 2	1, 2

Note that in this game (and in several examples in the note), for more readability, we take w.l.o.g. different alphabets for the players.

Example 3 (Matching penny game). The game is a two-player game, where each player has two actions, **a** and **b**. This is a zero-sum game (that is, the sum of the two payoffs in each situation is 0): the first player wins (payoff +1) if the two chosen actions are matching, whereas the second player wins if the two actions are different. The payoffs are summarized below:

		A_2	
		a	b
A_1	a	+1, -1	-1, +1
	b	-1, +1	+1, -1

The study of multiplayer games is to understand the rational behaviours of the players, assumed to be selfish. For instance, if for a player A , one of her strategy σ_A dominates another strategy σ'_A (in the sense that for all strategies of the other players, the payoff is larger using σ_A than using σ'_A), then there is no situation where player A should play σ'_A .

This is for instance the case in the prisoner's dilemma (Example 1), where action B dominates action S. Hence, the only rational issue in this example is

that both players play action B, yielding a payoff of 1 for each. One realizes however that it would be much better for them to both play S, but the threat that the other betrays (plays action B) makes that solution unsafe.

In the game of Example 2, action R (weakly) dominates action L for player A_2 (in the sense, it is better than or equally good), hence playing R for player A_2 is safe; knowing that, player A_1 will play action T; hence, *a priori*, the only rational issue of this game should be the profile (T, R) with payoff (2, 1). However, one also realizes that the profile (B, L) would be much better for both players, so only looking at dominating strategies might be too restrictive.

Finally, there might be no dominating strategies in a game, like in the matching penny game (Example 3), so other solution concepts have to be considered.

One of the most famous solution concepts for rationality is that of *Nash equilibrium* [27]. Let σ be a strategy profile. If $A \in \mathcal{P}$ is a player, and σ'_A is a strategy for A (called a *deviation*), then $\sigma[A/\sigma'_A]$ is the strategy profile such that A plays according to σ'_A and each other player $B \in \mathcal{P} \setminus \{A\}$ plays according to σ_B . Later, we write $\langle -A \rangle$ for the coalition of all the players except player A , that is, $\langle -A \rangle = \mathcal{P} \setminus \{A\}$.

A *mixed (resp. pure) Nash equilibrium* in game \mathcal{G} is a mixed (resp. pure) strategy profile $\sigma^* = (\sigma_A^*)_{A \in \mathcal{P}}$ such that for every $A \in \mathcal{P}$, for every player- A mixed (resp. pure) strategy σ_A ,

$$\mathbb{E}^{\sigma^*[A/\sigma_A]}(\text{payoff}_A) \leq \mathbb{E}^{\sigma^*}(\text{payoff}_A)$$

Note that even for mixed profiles, it is sufficient to look for pure deviations (if a mixed deviation improves the payoff, then so will do a pure deviation). Let σ^* be a strategy profile and σ_A be a deviation for player A such that $\mathbb{E}^{\sigma^*[A/\sigma_A]}(\text{payoff}_A) > \mathbb{E}^{\sigma^*}(\text{payoff}_A)$, then it is a *profitable* deviation for player A w.r.t. σ^* . If such a profitable deviation exists, then the profile is not a Nash equilibrium.

Coming back to the prisoner's dilemma (Example 1), the pair of pure dominating strategies (B, B) is a pure Nash equilibria, whereas the pair (S, S), which would yield a better payoff for both players, is *not* a Nash equilibrium.

In the matching penny game (Example 3), it is not difficult to check that none of the pure strategy profiles can be a Nash equilibrium since in each case, one of the players would benefit from switching to the other action. Also, one can argue that there is a unique Nash equilibrium, where each player plays each action uniformly at random, yielding an expected payoff of 0 for both.

Finally in Example 2, the two profiles (T, R) and (B, L) are the two Nash equilibria of the game. So there might be several Nash equilibria in a game, yielding possibly different payoffs.

A Nash equilibrium expresses a notion of stability. Indeed, it can be seen that a Nash equilibrium $\sigma = (\sigma_A)_{A \in \mathcal{P}}$ is such that each strategy σ_A is the *best-response* to the strategies $(\sigma_B)_{B \in \langle -A \rangle}$ of her adversaries. Formally, let \mathbb{S} (resp. \mathbb{S}_A , $\mathbb{S}_{\langle -A \rangle}$) be the set of mixed strategy profiles (resp. strategies for player A , strategies for coalition $\langle -A \rangle$). For every $\sigma \in \mathbb{S}$, let

$$\text{BR}(\sigma) = \left\{ \sigma' \in \mathbb{S} \mid \forall A \in \mathcal{P}, \sigma'_A \in \operatorname{argmax}_{\sigma''_A \in \mathbb{S}_A} \mathbb{E}^{\sigma[A/\sigma''_A]}(\text{payoff}_A) \right\}$$

be the set of best-response strategy profiles for σ . Then, σ is a Nash equilibrium if and only if $\sigma \in \text{BR}(\sigma)$.

We state now the famous Nash theorem [27], which is one of the important milestones in the game theory domain.

Theorem 1 (Nash theorem). *Every matrix game has a (Nash) equilibrium in mixed strategies.*

The original proof of Nash uses Brouwer’s fixed point theorem (see below). However it can also be seen that it is a consequence of Kakutani’s fixed point theorem (see below), by taking BR as function f (since the set mixed strategy profiles can be seen as a convex subset of $[0, 1]^{|\mathcal{P}| \cdot |\Sigma|}$).

Theorem 2 (Brouwer’s fixed point theorem). *Let $X \subseteq \mathbb{R}^n$ be a convex, compact and nonempty set. Then every continuous function $f: X \rightarrow X$ has a fixed point.*

Theorem 3 (Kakutani’s fixed point theorem). *Let X be a non-empty, compact and convex subset of \mathbb{R}^n . Let $f: X \rightarrow 2^X$ be a set-valued function on X with a closed graph and the property that $f(x)$ is non-empty and convex for all $x \in X$. Then f has a fixed point.*

As a final remark, let us define the *minmax value* of player $A \in \mathcal{P}$ as

$$\bar{v}_A = \min_{(\sigma_B)_{B \in \{-A\}} \in \mathbb{S}_{\{-A\}}} \max_{\sigma_A \in \mathbb{S}_A} \mathbb{E}^\sigma(\text{payoff}_A)$$

where $\sigma = (\sigma_B)_{B \in \mathcal{P}}$. This is the best player A can achieve, when she does not know how the other players will play. We will not discuss the minmax value, the maxmin value and the value of a game, but we notice that for every Nash equilibrium $\sigma \in \mathbb{M}_{\mathcal{P}}$, $\mathbb{E}^\sigma(\text{payoff}_A) \geq \bar{v}_A$ (since otherwise the strategy giving the minmax value will be a profitable deviation).

Conclusion. Game theory is a very rich field of research, of which we have only given few hints on the basic concepts, which will be relevant for the use in verification. We refer again to the textbook [26] for an entry point to this research domain.

Matrix games represent a “one-shot” interaction between the players. In system or program verification, players may represent components or controllers; it is usually useful to allow models with states for such systems, and to consider temporal behaviour of such systems. Hence the interaction is the result of a dynamic process, and not of a one-shot interaction like in matrix games. This is not specific to verification, and towards that goal, more complex interactions have been studied under the names of extensive-form games (games are then played on a tree), or repeated games (a given matrix games is a large number of times). There are many elegant results on these systems, but this note is not sufficient for this purpose.

3 Multiplayer games on graphs in verification

Matrix games and extensions like repeated games are not adapted to study interaction between players in a verification context. Indeed, to represent systems or programs, it is very useful to have models with explicit states. We will therefore first present the model of games on graphs that we will consider, and then argue why those games cannot be solved using the standard well-understood theory that we have recalled. We will then give some results and ideas for the computation of Nash equilibria in such games.

3.1 Definition of the general model and of the problems of interest

We consider the model of concurrent multi-player games, based on the two-player model of [1], and extended with probabilities. The deterministic version of this model was used for instance in [4].

Definition 1. A multiplayer stochastic concurrent game is a tuple

$$\mathcal{G} = (V, v_{\text{init}}, \mathcal{P}, \Sigma, \delta, (\text{payoff}_A)_{A \in \mathcal{P}})$$

where V is a finite set of vertices, $v_{\text{init}} \in V$ is the initial vertex, \mathcal{P} is a finite set of players, Σ is a finite set of actions, $\delta: V \times \Sigma^{\mathcal{P}} \rightarrow \text{Dist}(V)$ associates, with a given vertex and a given action tuple (called move) a distribution over the possible target vertices, and for every $A \in \mathcal{P}$, $\text{payoff}_A: V^\omega \rightarrow \mathbb{R}$ is a payoff function.

We later write $v \xrightarrow{m} v'$ whenever $\delta(v, m)(v') > 0$.

As before, we assume an explicit order on $\mathcal{P} = \{A_1, \dots, A_k\}$. Also, given a player $A \in \mathcal{P}$, we write $\langle -A \rangle$ for the coalition $\mathcal{P} \setminus \{A\}$. An element $m = (m_A)_{A \in \mathcal{P}} \in \Sigma^{\mathcal{P}}$ is called a move, and we may write it as $(m_{A_1}, \dots, m_{A_k})$. If $m \in \Sigma^{\mathcal{P}}$ and $A \in \mathcal{P}$, we write $m(A)$ for the A -component of m and $m(\langle -A \rangle)$ for all but the A components of m . In particular, we write $m(-A) = m'(-A)$ whenever $m(B) = m'(B)$ for every $B \in \langle -A \rangle$. Also, if $m \in \Sigma^{\mathcal{P}}$, $B \in \mathcal{P}$ and $a \in \Sigma$, then $m[B/a]$ denotes the move m' such that $m'(-B) = m(-B)$ and $m'(B) = a$.

A *history* π in \mathcal{G} is a finite non-empty sequence $v_0 v_1 \dots v_h \in V^+$ such that for every $1 \leq i \leq h$, there is $m_i \in \Sigma^{\mathcal{P}}$ with $v_{i-1} \xrightarrow{m_i} v_i$. We write $\text{last}(\pi)$ for the last vertex of π (i.e., v_h). If $i \leq h$, we also write $\pi_{\leq i}$ for the prefix $v_0 v_1 \dots v_i$. We write $\text{Hist}(v_0)$ for the set of histories in \mathcal{G} that start at v_0 . Notice that histories do not record moves used along a history.

We extend above notions to infinite sequences in a straightforward way and to the notion of play. We write $\text{Plays}(v_0)$ for the set of full plays that start at v_0 .

Let $A \in \mathcal{P}$ be a player. A *randomized (or mixed) strategy*¹ for player A from v_0 is a mapping $\sigma_A: \text{Hist}(v_0) \rightarrow \text{Dist}(\Sigma)$. An *outcome* of σ_A is a (n infinite) play

¹ This is the terminology used in the verification community, which might nevertheless be confusing with that used in the game theory community.

$\rho = v_0 v_1 \dots$ such that for every $i \geq 0$, writing $m_i(A) = \sigma_A(\rho_{\leq i})$, $v_i \xrightarrow{m_i} v_{i+1}$. We write $\text{out}(\sigma_A, v_0)$ for the set of outcomes of σ_A from v_0 . A *pure (or deterministic) strategy* for player A is a mixed strategy σ_A such that for every history h , $\sigma_A(h)$ is a Dirac probability measure (that is, it associates to some vertex v a probability 1, and to other vertices a probability 0).

A *mixed (resp. pure) strategy profile* is a tuple $\sigma = (\sigma_A)_{A \in \mathcal{P}}$, where, for every player $A \in \mathcal{P}$, σ_A is a mixed (resp. pure) strategy for player A . We write $\text{out}(\sigma, v_0)$ for the set of plays from v_0 , which are outcomes of all strategies part of σ . Note that if σ is pure, then $\text{out}(\sigma, v_0)$ has a single element, hence we may abusively speak of the outcome $\text{out}(\sigma, v_0)$.

Note that strategies, as defined above, can only observe the sequence of visited states along the history, but they may not depend on the exact distributions chosen by the players along the history, nor on the actual sequence of actions played by the players. Notice that this model is more general than the model where actions are visible, which are sometimes considered in the literature—see for instance [33] and [3, Section 6] or [14] for discussions—and the results presented here are valid (though actually simpler) when considering visible actions.

When σ is a strategy profile and σ'_A a player- A strategy, we write $\sigma[A/\sigma'_A]$ for the strategy profile where A plays according to σ'_A , and each other player B plays according to σ_B . The strategy σ'_A is a *deviation* of player A , or an *A-deviation*.

Once a strategy profile $\sigma = (\sigma_A)_{A \in \mathcal{P}}$ is fixed, for every $v_0 \in V$ it standardly induces a probability measure $\mathbb{P}_{v_0}^\sigma$ over the set of plays from v_0 in the game \mathcal{G} , by defining probability of cylinders as described below, and by extending it in a unique way to the generated σ -algebra. For every history $\pi = v_0 v_1 \dots v_h \in \text{Hist}(v_0)$, we let $\text{Cyl}(\pi) = \{\rho \in \text{Plays}(v_0) \mid \pi \text{ is a prefix of } \rho\}$ and we define $\mathbb{P}_{v_0}^\sigma(\text{Cyl}(v_0)) = 1$, and then inductively

$$\mathbb{P}_{v_0}^\sigma(\text{Cyl}(\pi v_{h+1})) = \mathbb{P}_{v_0}^\sigma(\text{Cyl}(\pi)) \cdot \left(\sum_{\substack{m \in \Sigma^{\mathcal{P}} \\ v_h \xrightarrow{m} v_{h+1}}} \sigma(\pi)(m) \cdot \delta(v_h, m)(v_{h+1}) \right)$$

where $\sigma(\pi)(m) = \prod_{A \in \mathcal{P}} \sigma_A(\pi)(m_A)$ is the probability that move m is selected by strategy profile σ .

Let f be a measurable function in the σ -algebra generated by the cylinders above. Then we define its expected value w.r.t. $\mathbb{P}_{v_0}^\sigma$ in a standard way, and denote it $\mathbb{E}_{v_0}^\sigma(f)$. We will therefore assume that payoff functions payoff_A ($A \in \mathcal{P}$) are all measurable!

The notion of Nash equilibrium that we have defined on matrix games extends naturally to games over graphs.

Definition 2. A Nash equilibrium from v_{init} is a strategy profile σ^* such that for every $A \in \mathcal{P}$, for every player- A deviation σ_A ,

$$\mathbb{E}_{v_{\text{init}}}^{\sigma^*[A/\sigma_A]}(\text{payoff}_A) \leq \mathbb{E}_{v_{\text{init}}}^{\sigma^*}(\text{payoff}_A)$$

Note that if σ is a pure profile, then $\mathbb{E}_{v_{\text{init}}}^{\sigma}(\text{payoff}_A) = \text{payoff}_A(\text{out}(\sigma, v_{\text{init}}))$. Also in this case, $\text{out}(\sigma, v_{\text{init}})$ is called the *main outcome* of equilibrium defined by σ .

As in matrix games, given a profile σ^* , a deviation σ_A for player A such that $\mathbb{E}_{v_{\text{init}}}^{\sigma^* [A/\sigma_A]}(\text{payoff}_A) > \mathbb{E}_{v_{\text{init}}}^{\sigma^*}(\text{payoff}_A)$ is called a *profitable deviation* for player A .

Payoff functions. A property ϕ over V^ω is said *prefix-independent* whenever for every ρ , $\rho \models \phi$ if and only if for every suffix ρ' of ρ , $\rho' \models \phi$.

We say that a payoff function $\text{payoff}: V^\omega \rightarrow \mathbb{R}$ is given by a Boolean property ϕ over V^ω whenever $\text{payoff}(\rho) = 1$ if $\rho \models \phi$, and $\text{payoff}(\rho) = 0$ if $\rho \not\models \phi$. Usually, ϕ will be some specific types of properties, like reachability, safety. We then abusively say payoff is a reachability (resp. safety, ...) objective. In a stochastic game, the expected value of such a payoff function is the probability to satisfy the property ϕ .

A payoff function payoff over V^ω is said *terminal-reward* if there is some designed subset $\tilde{V} \subseteq V$ such that all vertices of \tilde{V} are sinks in the graph of the game, and a function $w: \tilde{V} \rightarrow \mathbb{R}$ such that for every $\rho \in V^\omega$, $\text{payoff}(\rho) = w(\tilde{v})$ if ρ visits vertex $\tilde{v} \in \tilde{V}$ (which is unique if it exists since it is a sink), and $\text{payoff}(\rho) = 0$ otherwise. A particular case is when the image of w is included in $\{0, 1\}$, in which case we speak of *terminal-reachability*.

Subclasses of games. We use the following subclasses of games. Game \mathcal{G} is said:

- *turn-based* whenever there is a function $J: V \rightarrow \mathcal{P}$ such that for every $v \in V$, for every $m, m' \in \Sigma^{\mathcal{P}}$, $m(J(v)) = m'(J(v))$ implies $\delta(v, m) = \delta(v, m')$;
- *deterministic* whenever for every $v \in V$ and $m \in \Sigma^{\mathcal{P}}$, $\delta(v, m)$ is a Dirac probability measure on some vertex.

The existence and the constrained existence problems. For verification purposes, even if the existence of a Nash equilibrium might be interesting (due to the link with a stability property), we will also be interested in the constrained existence problem, and in the computability of Nash equilibria when they exist.

The *constrained existence problem* asks, given a stochastic multiplayer concurrent game $\mathcal{G} = (V, v_{\text{init}}, \mathcal{P}, \Sigma, \delta, (\text{payoff}_A)_{A \in \mathcal{P}})$ and a predicate P over $\mathbb{R}^{|\mathcal{P}|}$, whether there exists a Nash equilibrium σ such that $\left(\mathbb{E}_{v_0}^{\sigma}(\text{payoff}_A) \right)_{A \in \mathcal{P}} \in P$. Of course, for computability matters, predicates should not be too complicated, but one might think of lower bounds on the expected payoffs, or constraints on the social welfare (that is, the sum of the payoffs of all the players), etc. The *existence problem* is just the same problem when the predicate is $\mathbb{R}^{\mathcal{P}}$.

We add “pure” to the name of the problem if we restrict to pure strategy profiles.

Example 4 (Hide-or-run game). We consider the hide-or-run game represented on Figure 1 (left). There are three vertices and two players A_1 and A_2 . The actions for player A_1 are shoot (the snowball) and wait while the actions for player A_2 are hide and run. Strings **sh**, **wr**, **wh** and **sr** represent all possible moves in the game. In vertex v_0 , if player A_1 plays action **s** and player A_2

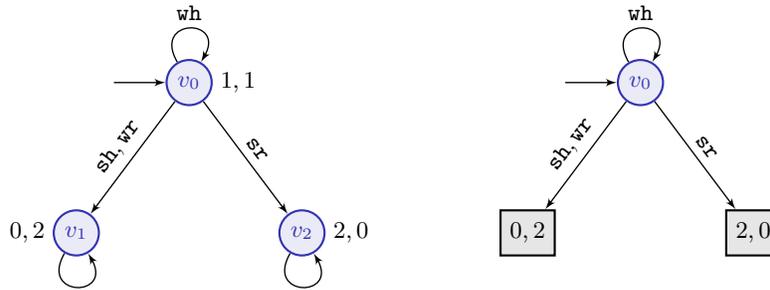


Fig. 1. Hide-or-run game, on the left; and one of its variants, on the right (black squared states indicate payoffs of the two players if the game ends there).

plays action **h**, then the game proceeds to vertex v_1 . Actions from v_1 and v_2 are irrelevant hence omitted in the figure. Pairs of numbers close to each vertex represents a weight for each player. The payoff function for each player will be the mean-payoff along the play of all encountered weights.

In this game, player A_1 wants to hit player A_2 with a (single) snowball. Player A_1 can therefore either **shoot** the ball or **wait**, while Player A_2 can **hide** (in which case she is not hit by the snowball) or **run** (in which case she is hit by the snowball if it is shot at the same time). Payoffs are assigned according to the satisfaction of the two players (payoff 2 when the player is satisfied and 0 otherwise). Note that the pair **wh** forever is only half-satisfactory for both players, hence a payoff of $(1, 1)$.

One realizes that there is no Nash equilibrium in the hide-or-run game: indeed, if the probability of playing **wh** forever from v_0 (resp. of playing **sr** from v_0) is positive, then player A_2 can deviate and get a better payoff; conversely, if the probability of playing **wr** (resp. **sh**) from v_0 is positive, then player A_1 can deviate and get a better payoff.

The game of Figure 1 (on the right) is a slight modification of the previous game, with a terminal-reachability payoff: if the game ends up in the bottom-left vertex (formerly v_1), then the payoff is $(0, 2)$, while it is $(2, 0)$ if the game ends up in the bottom-right vertex of the game (formerly v_2). It is very similar to the first game, the only difference is that playing **wh** forever yields a payoff of $(0, 0)$ instead of $(1, 1)$ previously. This slight modification yields a pure Nash equilibrium in the game, which is to play **sh** from v_0 .

3.2 Why does the standard theory not apply?

While matrix games are obviously special cases of our general model, one may nevertheless wonder why the standard theorems would not apply in this general model. We first realize that Nash theorem (stated as Theorem 1) does not apply: there are indeed potentially infinitely many pure strategies.

As we mentioned earlier, the proof of Nash theorem can be seen as a direct application of Kakutani's fixed point theorem (recalled as Theorem 3), which

is in a much more general setting than its application to Nash theorem. We explain how this theorem can apply in some cases, but why it does not apply in our precise setting. A *stationary* strategy σ is a mixed strategy such that for every $h, h' \in \text{Hist}(v_0)$, $\text{last}(h) = \text{last}(h')$ implies $\sigma(h) = \sigma(h')$. Such a strategy can therefore be viewed as an element \mathbb{R}^N for some integer N (one value for each triple $(v, a, A_i) \in V \times \Sigma \times \mathcal{P}$). The subspace X of \mathbb{R}^N of stationary strategies satisfies the hypotheses of the theorem. As we have already discussed in matrix games, a Nash equilibrium σ is such that each of its components is a best response to the other strategies. When restricted to stationary strategies, the *best-response function* can be defined as (we keep the same notations \mathbb{S} and \mathbb{S}_A):

$$\text{BR}(\sigma) = \left\{ \sigma' \in \mathbb{S} \mid \forall A \in \mathcal{P}, \sigma'_A \in \operatorname{argmax}_{\sigma''_A \in \mathbb{S}_A} \mathbb{E}_{v_0}^{\sigma^{[A/\sigma'_A]}}(\text{payoff}_A) \right\}$$

Nevertheless, over game graphs, continuity of this best-response function is not ensured (hence the graph of BR is not closed). Let us consider for example game of Figure 2 (borrowed from [6]). It is assumed to be turn-based (vertex v_i belongs to player A_i): from v_i , player A_i can either continue or leave the game. A stationary strategy profile σ can be stored as a pair $(\sigma_{A_1}(v_1)(1), \sigma_{A_2}(v_2)(1)) \in [0, 1]^2$, where the first (resp. second) element is the probability that player A_1 (resp. A_2) leaves the game from v_1 (resp. v_2). If one player decides to leave the game with some positive probability, the other player has all incentive to purely continue the game, until eventually reaching the terminal state (with probability 1). Hence $\text{BR}((x, y)) = \{(0, 0)\}$ for every $x, y > 0$. However, if one player purely continues the game, the only way to win some positive payoff $\frac{1}{3}$ is to leave the game with positive probability. Hence $\text{BR}((0, 0)) = \{(x, y) \mid x, y > 0\}$. We conclude that the graph is not closed, so Theorem 3 cannot be applied to the classical BR function. We finally notice that any profile $(x, 0)$ with $x > 0$, or $(0, y)$ with $y > 0$, is a Nash equilibrium.

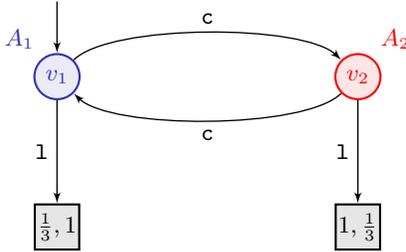


Fig. 2. Turn-based game with terminal rewards (black squared states indicate payoffs of the two players if the game ends there) showing the non-applicability of Kakutani's theorem; the first player who leaves the loop with some positive probability loses.

Though this theorem does not apply in our general context, it can be used in others, for instance for stay-in-a-set games [29], for Nash equilibria with discounted payoffs or ϵ -Nash equilibria [16].

3.3 Discussion on a simple scenario

Let us focus on a simple scenario first. We fix for the rest of this subsection a game $\mathcal{G} = (V, v_{\text{init}}, \mathcal{P}, \Sigma, \delta, (\text{payoff}_A)_{A \in \mathcal{P}})$ which satisfies the following (restricting) assumptions:

- the game is turn-based and deterministic;
- for every $A \in \mathcal{P}$, the payoff function payoff_A is given by a Boolean prefix-independent objective ϕ_A ;

We note (†) the hypotheses of this simple scenario.

For every $A \in \mathcal{P}$, we let $\mathcal{G}[A]$ be the two-player zero-sum game built on the same arena as \mathcal{G} , where A plays against coalition $\langle\!\langle -A \rangle\!\rangle$ (more precisely, all vertices which previously belong to some $B \in \mathcal{P} \setminus \{A\}$ now belongs to $\langle\!\langle -A \rangle\!\rangle$, and all vertices which previously belong to A still belongs to A); the payoff function for A is payoff_A , while the payoff function for $\langle\!\langle -A \rangle\!\rangle$ is $-\text{payoff}_A$ (here, in our simple setting, the objective of player A is ϕ_A while the objective of $\langle\!\langle -A \rangle\!\rangle$ is $\neg\phi_A$). Let $W_{\langle\!\langle -A \rangle\!\rangle}$ (resp. W_A) be the set of winning states for coalition $\langle\!\langle -A \rangle\!\rangle$ (resp. player A) in this game. Since this game is turn-based and the objectives are prefix-independent, the game will be determined, that is, either A has a winning strategy, or the coalition $\langle\!\langle -A \rangle\!\rangle$ has a winning strategy (that is, for every vertex $v \in V$, either $v \in W_A$ or $v \in W_{\langle\!\langle -A \rangle\!\rangle}$). Furthermore, for large classes of objectives, the set $W_{\langle\!\langle -A \rangle\!\rangle}$ (or W_A) can be computed. We report here to the whole literature on the subject, see [21] for an entry point.

One can then characterize pure Nash equilibria by the formula:

$$\Phi_{\text{NE}} = \bigwedge_{A \in \mathcal{P}} \left(\neg\phi_A \Rightarrow \mathbf{G}W_{\langle\!\langle -A \rangle\!\rangle} \right)$$

borrowing notations from the syntax of LTL [28]: that is, Φ_{NE} holds along a play ρ whenever for every $A \in \mathcal{P}$, either ϕ_A holds along the outcome or A cannot enforce winning anywhere along the play (or equivalently, $\langle\!\langle -A \rangle\!\rangle$ can enforce $\neg\phi_A$ in $\mathcal{G}[A]$). Note that the same formula can be used for reachability objectives but that a slightly different one has to be used for safety objectives.

One can show:

Proposition 1. *Assume setting (†). Let $\rho \in \text{Plays}(v_{\text{init}})$. Then, $\rho \models \Phi_{\text{NE}}$ if and only if there is a Nash equilibrium σ from v_{init} such that $\text{out}(\sigma, v_{\text{init}}) = \rho$.*

Proof (Sketch). Indeed, pick a play $\rho \in \text{Plays}(v_{\text{init}})$, and assume that $\rho \models \Phi_{\text{NE}}$. Consider a player $A \in \mathcal{P}$. Such a player may have some interest in deviating only if her objective ϕ_A is not already satisfied by ρ . In that case, she has a profitable deviation after some prefix π of ρ if she is able to ensure winning after π . In particular, if no winning state of A is visited along ρ , then ρ can be completed into a Nash equilibrium as follows:

- all players play along ρ ;

- as soon as a player deviates from ρ , then the coalition $(-A) = \mathcal{P} \setminus \{A\}$ starts playing a counter-strategy to A . Such a strategy is sometimes called a *threat* or a *trigger* strategy.

Conversely assume there is a Nash equilibrium σ from v_{init} such that $\text{out}(\sigma, v_{\text{init}}) = \rho$. Pick a player $A \in \mathcal{P}$ such that $\rho \not\models \phi_A$. Then, since σ is a Nash equilibrium, from every visited vertex v along ρ , A cannot enforce her objective ϕ_A , which means that $v \notin W_A$, hence $v \in W_{(-A)}$. Hence $\rho \models \Phi_{\text{NE}}$. \square

The situation is illustrated on Figure 3. Note that by determinacy, “Player A_1 should lose” can be replaced by “Coalition $\{A_2, A_3\}$ prevents A_1 from winning”.

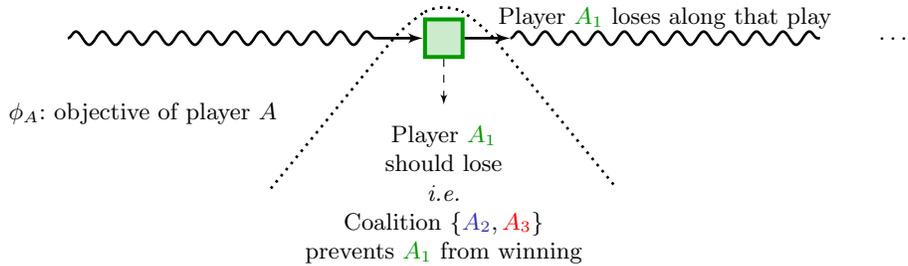


Fig. 3. General shape of a Nash equilibrium in the simple setting (example with three players).

In this simple setting we can also prove the following existence result:

Proposition 2. *Assume setting (\dagger) . There always exists a pure Nash equilibrium from v_{init} .*

Proof (Sketch). In this simple setting, in two-player zero-sum games, there always exists strongly optimal strategies [20], that is, one strategy for each of the two players, say σ and τ , such that each of the two strategies is optimal (for the corresponding player) after any compatible prefix. For every $A \in \mathcal{P}$, we apply this result to each of the games $\mathcal{G}[A]$ from v_{init} , and write σ_A for the corresponding strongly optimal strategy for player A in $\mathcal{G}[A]$.

We argue why the main outcome ρ of $\sigma = (\sigma_A)_{A \in \mathcal{P}}$ satisfies formula Φ_{NE} . Assume that $\rho \not\models \phi_A$. Towards a contradiction assume that one of the visited vertices, say v after prefix π , along ρ , does not belong to $W_{(-A)}$. By the strong determinacy result mentioned at the beginning of the proof, it implies that vertex v belongs to W_A . Since σ_A is strongly optimal, it is also optimal after prefix π : hence it is winning after prefix π . In particular, since ϕ_A is prefix-independent, ρ should be winning as well. Contradiction: $\rho \models \Phi_{\text{NE}}$. \square

Algorithmics issues. By combining the proof of Propositions 1 and 2, one can compute a pure Nash equilibrium from strongly optimal and trigger strategies.

Another solution consists in computing for every $A \in \mathcal{P}$ the set $W_{(-A)}$ (or equivalently W_A), and to compute an infinite path in the game which satisfies formula Φ_{NE} (which can be done for instance by enumerating the possible set of losing players, and then finding an adequate ultimately periodic play). Obviously, for specific winning conditions, more efficient algorithms can be designed, but this is not the aim of this note. We report e.g. to [34, 4] for more algorithms.

3.4 Back to stochastic concurrent games

By a non-trivial extension of the discussion of Subsection 3.3 (see [36, Section 3] for details), one can show the following existence result:

Theorem 4. *There exists a pure Nash equilibrium in any multiplayer stochastic turn-based game with prefix-independent winning objectives (which we can compute). This also holds in the same setting for any ω -regular objectives. [36, Section 3]*

This result in particular applies to mean-payoff objectives, which are prefix-independent.

Why are we not fully happy with such a result?

- one would like to go from turn-based to concurrent games;
- one would like more general payoff functions;
- one would like to solve the constrained existence problem.

It turns out that those extensions are very intricate, and that we can give a list of (related but incomparable) undecidability results.

Theorem 5. *The following problems are all undecidable:*

1. the **constrained² existence problem** for stochastic multiplayer turn-based games with terminal-reachability objectives. This is true even if we restrict to **pure** strategy profiles. [36, Section 4]
2. the **constrained existence problem** for **deterministic** multiplayer turn-based games with **terminal-reward payoffs**. [35, Section 7]
3. the **constrained existence problem** for **deterministic three-player concurrent** games with **terminal-reachability payoffs**.³ [5]
4. the **existence problem** for **deterministic three-player concurrent** games with **terminal-reward payoffs**. [5]
5. the **constrained existence problem** for **deterministic three-player concurrent** games with **safety objectives**.⁴ [5]

² In the proof, we only impose that a player wins almost-surely.

³ This holds even with a constraint on the social welfare. This result has therefore to be compared with the result of [19], which states that the existence problem is NP-complete in two-player games.

⁴ This result has to be compared with the result of [29], which states that there always exists a Nash equilibrium in a safety game.

3.5 The suspect-game construction [4]

The setting we have chosen here assumes actions are invisible (since only visited vertices are visible along histories). Hence, a deviation from the main outcome can only be detected when the play goes out of the main outcome of the Nash equilibrium. However, even if a deviation occurs, there can be uncertainties for some of the players concerning the identity of the deviator.

Consider for instance the game in Figure 4, with three players. Assume that the main outcome goes through $v_0 \xrightarrow{\text{aaa}} v_1$.

- If the game proceeds to vertex v_2 instead of v_1 , it means that either player A_1 deviated alone (playing **b** instead of **a**), or both players A_1 and A_2 played **b** instead of **a**; the second case cannot occur since Nash equilibria only care of single-player deviations; hence only player A_1 can be the deviator, and all players will therefore know the identity of the deviator.
- If the game proceeds to vertex v_3 , then there are two possible suspects amongst the players: either A_2 or A_3 can be the deviator. In both cases, the two players A_2 and A_3 will know the identity of the deviator, while player A_1 will not know it.

This knowledge about the possible deviators is represented via a suspect function defined as follows:

- $\text{susp}((v_0, v_2), \text{aaa}) = \{A_1\}$
- $\text{susp}((v_0, v_3), \text{aaa}) = \{A_2, A_3\}$

with the meaning that, starting from v_0 , if the game proceeds to v_2 (resp. v_3), then only A_1 (resp. A_2 and A_3) are suspect for the deviation.

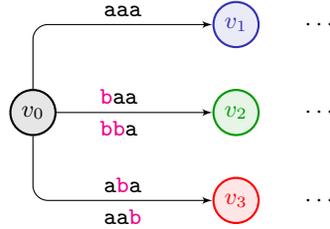


Fig. 4. Several suspect players.

More generally, we consider a game $\mathcal{G} = (V, v_{\text{init}}, \mathcal{P}, \Sigma, \delta, (\text{payoff}_A)_{A \in \mathcal{P}})$, and we define the function $\text{susp}: V^2 \times \Sigma^{\mathcal{P}} \rightarrow 2^{\mathcal{P}}$ as follows:

$$\text{susp}((v_0, v), m) = \{A \in \mathcal{P} \mid \exists b \in \Sigma \text{ s.t. } v_0 \xrightarrow{m[A/b]} v\}$$

Note that in the case no deviation occurred, that is if $v_0 \xrightarrow{m} v$, then the set of suspects is the set of all the players; this is because the set of suspect players becomes only relevant after a deviation has occurred.

The suspect game is now defined as the two-player⁵ turn-based game $\mathcal{S}_G = (S_{\text{Eve}}, S_{\text{Adam}}, s_{\text{init}}, \Gamma, E, (\text{payoff}'_A)_{A \in \mathcal{P}})$ where:

- $S_{\text{Eve}} = V \times 2^{\mathcal{P}}$ is the set of states belonging to **Eve**;
- $S_{\text{Adam}} = S_{\text{Eve}} \times \Sigma^{\mathcal{P}}$ is the set of states belonging to **Adam**;
- $s_{\text{init}} = (v_{\text{init}}, \mathcal{P})$ is the initial state;
- $\Gamma = \Sigma^{\mathcal{P}} \cup V$ is the new alphabet;
- the set of edges is

$$E = \{(v, \text{susp}) \xrightarrow{m} ((v, \text{susp}), m) \mid v \in V, \text{susp} \subseteq \mathcal{P}, m \in \Sigma^{\mathcal{P}}\} \cup$$

$$\{((v, \text{susp}), m) \xrightarrow{v'} (v', \text{susp} \cap \text{susp}((v, v'), m)) \mid \exists A \in \mathcal{P} \exists b \in \Sigma \text{ s.t. } v \xrightarrow{m[A/b]} v'\};$$

– if $\rho = (v_0, \text{susp}_0)(v_0, \text{susp}_0, m_1)(v_1, \text{susp}_1) \dots$, for every $A \in \mathcal{P}$, $\text{payoff}'_A(\rho) = \text{payoff}_A(v_0 v_1 \dots)$.

Given a play $\rho = (v_0, \text{susp}_0)(v_0, \text{susp}_0, m_1)(v_1, \text{susp}_1) \dots$, we define the set of suspect players for ρ as $\text{susp}(\rho) = \bigcap_{i \geq 0} \text{susp}_i$ (this limit is well-defined).

The winning condition for **Eve** is rather non-standard, since it is a condition on the set of outcomes of **Eve**, not on each outcome of the strategy individually. A strategy ζ for **Eve** in \mathcal{S}_G is winning for some $\alpha \in \mathbb{R}^{\mathcal{P}}$ if the unique outcome of ζ where **Adam** complies to **Eve**⁶ has payoff α , and for every other outcome ρ of ζ , for every $A \in \text{susp}(\rho)$, $\text{payoff}'_A(\rho) \leq \alpha_A$.

Example 5. We consider again the small (part of) game depicted on Figure 4 (all missing moves in the figure lead to v_1). The corresponding part of the suspect game is given in Figure 5.

The role of **Eve** is to search for an equilibrium by suggesting moves to the players, and the role of **Adam** is to check whether there are possible profitable deviations. In particular, winning strategies of **Eve** in the suspect game will coincide with Nash equilibria in the original game:

Proposition 3 (Correctness). *Let $\alpha \in \mathbb{R}^{\mathcal{P}}$. There is a Nash equilibrium in \mathcal{G} with payoff α from v_{init} if and only if **Eve** has a winning strategy for α in \mathcal{S}_G from s_{init} .*

Remark 1. Assume we start with a turn-based game. Then, since the arena of the game is known by the players, as soon as some deviation occurs, then all players will know which player is responsible for the deviation (since this is the player who controls the vertex at which the deviation occurred). In this case, the set of suspects will immediately be a singleton. The winning condition then ensures that, from a vertex controlled by player A_i , if a deviation occurs, then **Eve** plays an optimal strategy for the coalition $(-A_i) = \mathcal{P} \setminus \{A_i\}$. We somehow recover the intuitive explanation we gave in Subsection 3.3.

Also, assume that actions are visible, then similarly, as soon as there is a deviation, the identity of the deviator is known by all the players.

⁵ We call the two players **Eve** and **Adam**.

⁶ That is, from (v, susp, m) , **Adam** chooses to go to (v', susp) where $v \xrightarrow{m} v'$.

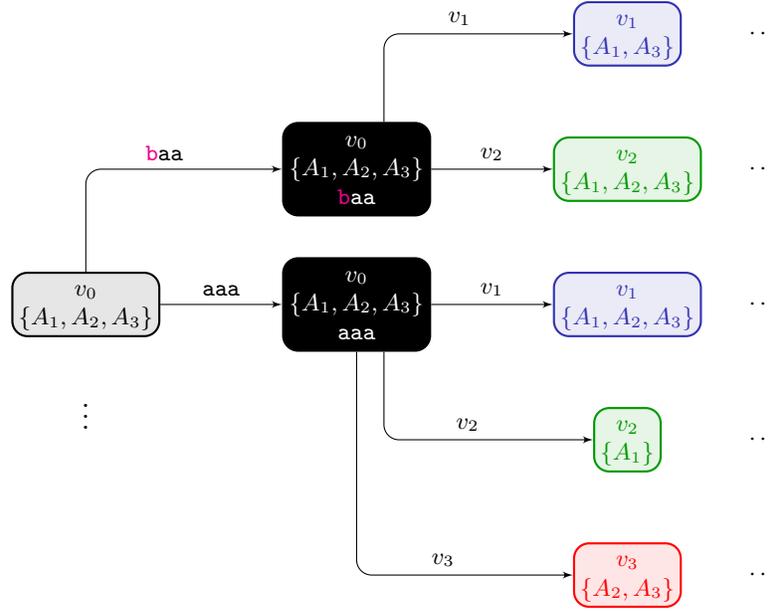


Fig. 5. Illustration of the suspect game construction (states in light colors are Eve’s states while states in dark colors are Adam’s states).

Algorithmics issues. Using the above construction, it is sufficient to solve the suspect game to compute Nash equilibria, since the equivalence of Proposition 3 is constructive. However, the winning condition is non-standard. In [4], many algorithms are designed for specific payoff functions. Complexities obviously depend on the (discrete) payoff functions which are used.

As an illustration, let us look at Figure 6, where each player A has a Boolean objective ϕ_A . We assume players A_1 and A_2 are losing along a play. Then if this play is the main outcome of a Nash equilibrium, it should be the case that from v , A_3 is able to punish both players (with the help of A_1 if A_2 is the deviator, and with the help of A_2 in case A_1 is the deviator); from v' , it is known by everyone that A_1 is the deviator, hence the coalition of both A_2 and A_3 should be able to punish A_1 from there. Algorithmically, it is therefore sufficient to compute states $(v, \{A_i\})$ which are winning for $\neg\phi_{A_i}$ for the coalition $(\neg A_i)$ (or equivalently Adam); and then (in a bottom-up manner) states $(v, \{A_i \mid i \in I\})$ which are winning for Adam for objective $\bigwedge_{i \in I} (A_i \text{ suspect at the limit} \Rightarrow \neg\phi_{A_i})$.

In [35, Section 6], an algorithm for mean-payoff functions is designed (in a setting where actions are visible), which consists in computing values of the various two-player mean-payoff games (A against $(\neg A)$) in each vertex, and then to find a lasso satisfying a given constraint on the payoff.

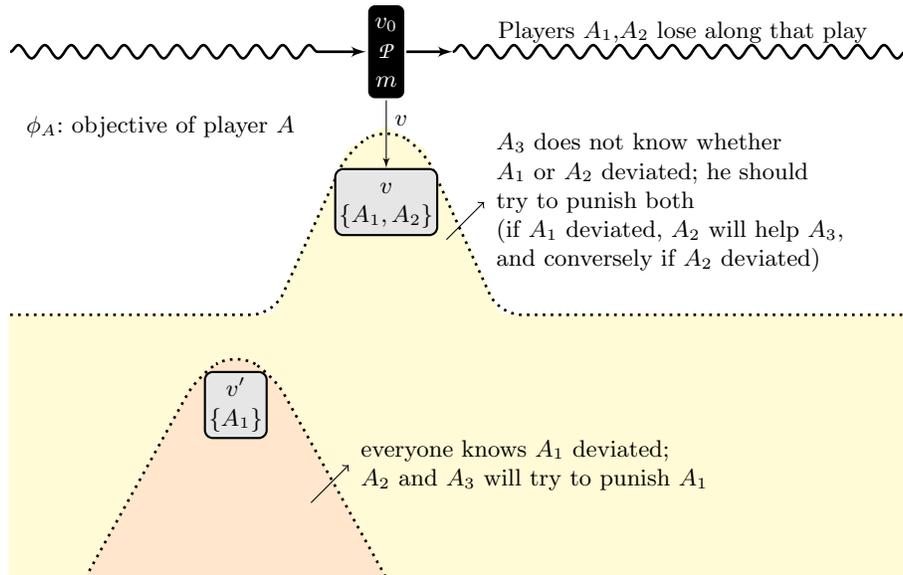


Fig. 6. Overview on the suspect-game construction.

4 Discussion

In this note, we have presented some basics of game theory over matrix games, and discussed how concepts from game theory can be studied in the context of models used in verification. We have discussed in particular a general construction that can be made to compute Nash equilibria in games on graphs, and which gives some general understanding of how interaction between players can be understood. This construction has been refined in several respects (for other solution concepts [9, 18], in some partial information contexts [2, 7]), and might be useful in some more contexts.

Even though there are some known existence results (we have mentioned some of them in Subsection 3.4), for simple payoff functions like terminal reachability payoffs. A related discussion can be found in [23].

In this note, we have not discussed temporal logics for multi-agent systems, even though this is a very rich domain of research (see [25] for some pointers). We have also not discussed domination and admissibility (see [10] among others), nor subgame-perfect equilibria, which have nevertheless been much studied (among others, see [34, 12, 13, 11]).

Acknowledgments. I would like to thank all my co-authors since I started working on multiplayer games played on graphs, that is, Nicolas Markey, Romain Brenguier [8], Daniel Stan [30], Michael Ummels and Nathan Thomasset.

References

1. Rajeev Alur, Thomas A. Henzinger, and Orna Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49:672–713, 2002.
2. Patricia Bouyer. Games on graphs with a public signal monitoring. In *Proc. 21st International Conference on Foundations of Software Science and Computation Structures (FoSSaCS'18)*, volume 10803 of *Lecture Notes in Computer Science*, pages 530–547. Springer, 2018.
3. Patricia Bouyer, Romain Brenguier, Nicolas Markey, and Michael Ummels. Nash equilibria in concurrent games with Büchi objectives. In *Proc. 30th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'11)*, volume 13 of *LIPIcs*, pages 375–386. Leibniz-Zentrum für Informatik, 2011.
4. Patricia Bouyer, Romain Brenguier, Nicolas Markey, and Michael Ummels. Pure Nash equilibria in concurrent games. *Logical Methods in Computer Science*, 11(2:9), 2015.
5. Patricia Bouyer, Nicolas Markey, and Daniel Stan. Mixed Nash equilibria in concurrent games. In *Proc. 33rd Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'14)*, volume 29 of *LIPIcs*, pages 351–363. Leibniz-Zentrum für Informatik, 2014.
6. Patricia Bouyer, Nicolas Markey, and Daniel Stan. Stochastic equilibria under imprecise deviations in terminal-reward concurrent games. In *Proc. 7th International Symposium on Games, Automata, Logics and Formal Verification (GandALF'16)*, volume 226 of *Electronic Proceedings in Theoretical Computer Science*, pages 61–75, 2016.
7. Patricia Bouyer and Nathan Thomasset. Nash equilibria in games over graphs equipped with a communication mechanism. In *Proc. 44th International Symposium on Mathematical Foundations of Computer Science (MFCS'19)*, volume 138 of *LIPIcs*. Leibniz-Zentrum für Informatik, 2019. To appear.
8. Romain Brenguier. *Nash Equilibria in Concurrent Games – Application to Timed Games*. PhD thesis, ENS Cachan, France, 2012.
9. Romain Brenguier. Robust equilibria in mean-payoff games. In *Proc. 19th International Conference on Foundations of Software Science and Computation Structures (FoSSaCS'16)*, volume 9634 of *Lecture Notes in Computer Science*, pages 217–233. Springer, 2016.
10. Romain Brenguier, Arno Pauly, Jean-François Raskin, and Ocan Sankur. Admissibility in games with imperfect information (invited talk). In *Proc. 28th International Conference on Concurrency Theory (CONCUR'17)*, volume 85 of *LIPIcs*, pages 2:1–2:23. Leibniz-Zentrum für Informatik, 2017.
11. Thomas Brihaye, Véronique Bruyère, Aline Goeminne, and Jean-François Raskin. Constrained existence problem for weak subgame perfect equilibria with ω -regular boolean objectives. In *Proc. 9th International Symposium on Games, Automata, Logics and Formal Verification (GandALF'18)*, volume 277 of *Electronic Proceedings in Theoretical Computer Science*, pages 16–29, 2018.
12. Thomas Brihaye, Véronique Bruyère, Julie De Pril, and Hugo Gimbert. On subgame perfection in quantitative reachability games. *Logical Methods in Computer Science*, 9(1):1–32, 2013.
13. Véronique Bruyère, Stéphane Le Roux, Arno Pauly, and Jean-François Raskin. On the existence of weak subgame perfect equilibria. In *Proc. 20th International*

- Conference on Foundations of Software Science and Computation Structures (FoSSaCS'17)*, volume 10203 of *Lecture Notes in Computer Science*, pages 145–161. Springer, 2017.
14. Krishnendu Chatterjee and Laurent Doyen. Partial-observation stochastic games: How to win when belief fails. *ACM Transactions on Computational Logic*, 15(2:16), 2014.
 15. Krishnendu Chatterjee, Thomas A. Henzinger, and Marcin Jurdziński. Games with secure equilibria. *Theoretical Computer Science*, 365(1-2):67–82, 2006.
 16. Krishnendu Chatterjee, Marcin Jurdziński, and Rupak Majumdar. On Nash equilibria in stochastic games. In *Proc. 18th International Workshop on Computer Science Logic (CSL'04)*, volume 3210 of *Lecture Notes in Computer Science*, pages 26–40. Springer, 2004.
 17. Krishnendu Chatterjee, Rupak Majumdar, and Marcin Jurdziński. On Nash equilibria in stochastic games. In *Proc. 18th International Workshop on Computer Science Logic (CSL'04)*, volume 3210 of *Lecture Notes in Computer Science*, pages 26–40. Springer, 2004.
 18. Rodica Condurache, Youssouf Oualhadj, and Nicolas Troquard. The complexity of rational synthesis for concurrent games. In *Proc. 29th International Conference on Concurrency Theory (CONCUR'18)*, LIPIcs, pages 38:1–38:15. Leibniz-Zentrum für Informatik, 2018.
 19. Vincent Conitzer and Tuomas Sandholm. New complexity results about Nash equilibria. *Games and Economic Behavior*, 63(2):621–641, 2008.
 20. Hugo Gimbert and Florian Horn. Solving simple stochastic tail games. In *Proc. 21st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'10)*, pages 847–862. SIAM, 2010.
 21. Erich Grädel, Wolfgang Thomas, and Thomas Wilke, editors. *Automata, Logics, and Infinite Games: A Guide to Current Research*, volume 2500 of *Lecture Notes in Computer Science*. Springer, 2002.
 22. Erich Grädel and Michael Ummels. Solution concepts and algorithms for infinite multiplayer games. In *New Perspectives on Games and Interaction*, volume 4 of *Texts in Logic and Games*, pages 151–178. Amsterdam University Press, 2008.
 23. Kristoffer Arnsfelt Hansen and Mikhail Raskin. A stay-in-a-set game without a stationary equilibrium. In *Proc. 10th International Symposium on Games, Automata, Logics and Formal Verification (GandALF'19)*, 2019. To appear.
 24. Thomas A. Henzinger. Games in system design and verification. In *Proc. 10th Conference on Theoretical Aspects of Rationality and Knowledge (TARK'05)*, pages 1–4, 2005.
 25. Nicolas Markey. Temporal logics for multi-agent systems (invited talk). In *Proc. 42nd International Symposium on Mathematical Foundations of Computer Science (MFCS'17)*, volume 83 of *LIPIcs*, pages 84:1–84:3. Leibniz-Zentrum für Informatik, 2017.
 26. Michael Maschler, Eilon Solan, and Shmuel Zamir. *Game Theory*. Cambridge University Press, 2013.
 27. John F. Nash. Equilibrium points in n -person games. *Proceedings of the National Academy of Sciences of the United States of America*, 36(1):48–49, 1950.
 28. Amir Pnueli. The temporal logic of programs. In *Proc. 18th Annual Symposium on Foundations of Computer Science (FOCS'77)*, pages 46–57. IEEE Computer Society Press, 1977.
 29. Piercesare Secchi and William D. Sudderth. Stay-in-a-set games. *International Journal of Game Theory*, 30:479–490, 2001.

30. Daniel Stan. *Randomized Strategies in Concurrent Games*. PhD thesis, Université Paris-Saclay, France, 2017.
31. Wolfgang Thomas. Infinite games and verification. In *Proc. 14th International Conference on Computer Aided Verification (CAV'02)*, volume 2404 of *Lecture Notes in Computer Science*, pages 58–64. Springer, 2002. Invited Tutorial.
32. Michael Ummels. Rational behaviour and strategy construction in infinite multiplayer games. In *Proc. 26th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'06)*, volume 4337 of *Lecture Notes in Computer Science*, pages 212–223. Springer, 2006.
33. Michael Ummels. The complexity of Nash equilibria in infinite multiplayer games. In *Proc. 11th International Conference on Foundations of Software Science and Computation Structures (FoSSaCS'08)*, volume 4962 of *Lecture Notes in Computer Science*, pages 20–34. Springer, 2008.
34. Michael Ummels. *Stochastic Multiplayer Games – Theory and Algorithms*. PhD thesis, RWTH Aachen, Germany, 2010.
35. Michael Ummels and Dominik Wojtczak. The complexity of Nash equilibria in limit-average games. In *Proc. 22nd International Conference on Concurrency Theory (CONCUR'11)*, volume 6901 of *Lecture Notes in Computer Science*, pages 482–496. Springer, 2011.
36. Michael Ummels and Dominik Wojtczak. The complexity of Nash equilibria in stochastic multiplayer games. *Logical Methods in Computer Science*, 7(3), 2011.