



Games on Graphs with a Public Signal Monitoring

Patricia Bouyer^(✉) 

LSV, CNRS, ENS Paris-Saclay, Université Paris-Saclay, Cachan, France
bouyer@lsv.fr

Abstract. We study pure Nash equilibria in games on graphs with an imperfect monitoring based on a public signal. In such games, deviations and players responsible for those deviations can be hard to detect and track. We propose a generic epistemic game abstraction, which conveniently allows to represent the knowledge of the players about these deviations, and give a characterization of Nash equilibria in terms of winning strategies in the abstraction. We then use the abstraction to develop algorithms for some payoff functions.

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 (an 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 [23, 30].

The problem of distributed synthesis [25] can be formulated using multiplayer concurrent games. In this setting, there is a global objective Φ , and one particular player called Nature. The question then is whether the so-called grand coalition (all players except Nature) can enforce Φ , whatever Nature does. While the players (except Nature) cooperate (and can initially coordinate), their choice of actions (or strategy) can only depend on what they see from the play so far. When modelling distributed synthesis as concurrent games, information players receive is given via a partial observation function of the states of the game. When the players have perfect monitoring of the play, the distributed synthesis problem reduces to a standard two-player zero-sum game. Distributed synthesis is a fairly hot topic, both using the formalization via concurrent games we have already described and using the formalization via an architecture of processes [26]. The most general decidability results in the concurrent game setting are under the

This work has been supported by ERC project EQUALLIS (FP7-308087).

assumption of hierarchical observation [6, 36] (information received by the players is ordered) or more recently under recurring common knowledge [5].

While distributed synthesis involves several players, this remains nevertheless a zero-sum question. Using solution concepts borrowed from game theory, one can go a bit further in describing the interactions between the players, and in particular in describing rational behaviours of selfish players. One of the most basic solution concepts is that of Nash equilibria [24]. 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 [15, 17, 32, 33]. Several series of works have followed. To roughly give an idea of the existing results, pure Nash equilibria always exist in turn-based games for ω -regular objectives [35] but not in concurrent games; they can nevertheless be computed for large classes of objectives [9, 11, 35]. The problem becomes harder with mixed (that is, stochastic) Nash equilibria, for which we often cannot decide the existence [10, 34].

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. This simple characterization of Nash equilibria is made explicit in [18]. Variants of Nash equilibria require slightly different ingredients, but they are mostly of a similar vein.

In (almost) all these works though, perfect monitoring is implicitly assumed: in all cases, players get full information on the states which are visited; a slight imperfect monitoring is assumed in some works on concurrent games (like [9]), where actions which have been selected are not made available to all the players (we speak of hidden actions). This can yield some uncertainties for detecting deviators but not on states the game can be in, which is rather limited and can actually be handled.

In this work, we integrate imperfect monitoring into the problem of deciding the existence of pure Nash equilibria and computing witnesses. We choose to model imperfect monitoring via the notion of signal, which, given a joint decision of the players together with the next state the play will be in, gives some information to the players. To take further decisions, players get information from the signals they received, and have perfect recall about the past (their own actions and the signals they received). We believe this is a meaningful framework. Let us give an example of a wireless network in which several devices try to send data: each device can modulate its transmission power, in order to maximise its bandwidth and reduce energy consumption as much as possible. However there might be a degradation of the bandwidth due to other devices, and the satisfaction of each device is measured as a compromise between energy consumption and allocated bandwidth, and is given by a quantitative payoff

function.¹ In such a problem, it is natural to assume that a device only gets a global information about the load of the network, and not about each other device which is connected to the network. This can be expressed using imperfect monitoring via public signals.

Following [31] in the framework of repeated matrix games, we put forward a notion of *public signal* (inspired by [31]). A signal will be said public whenever it is common to all players. That is, after each move, all the players get the same information (their own action remains of course private). We will also distinguish several kinds of payoff functions, depending on whether they are publicly visible (they only depend on the public signal), or privately visible (they depend on the public signal and on private actions: the corresponding player knows his payoff!), or invisible (players may not even be sure of their payoff).

The payoff functions we will focus on in this paper are Boolean ω -regular payoff functions and mean payoff functions. Some of the decidability results can be extended in various directions, which we will mention along the way.

As initial contributions of the paper, we show some undecidability results, and in particular that the hypothesis of public signal solely is not sufficient to enjoy all nice decidability results: for mean payoff functions, which are privately visible, one cannot decide the constrained existence of a Nash equilibrium. Constrained existence of a Nash equilibrium asks for the existence of a Nash equilibrium whose payoff satisfies some given constraint.

The main contribution of the paper is the construction of a so-called *epistemic game abstraction*. This abstraction is a two-player turn-based game in which we show that winning strategies of one of the players (Eve) actually correspond to Nash equilibria in the original game. The winning condition for Eve is rather complex, but can be simplified in the case of publicly visible payoff functions. The epistemic game abstraction is inspired by both the epistemic unfolding of [4] used for distributed synthesis, and the suspect game abstraction of [9] used to compute Nash equilibria in concurrent games with hidden actions. In our abstraction, we nevertheless not fully formalize epistemic unfoldings, and concentrate on the structure of the knowledge which is useful under the assumption of public signals; we show that several subset constructions (as done initially in [27], and since then used in various occasions, see e.g. [14, 19, 20, 22]) made in parallel, are sufficient to represent the knowledge of all the players. The framework of [9] happens to be a special case of the public signal monitoring framework of the current paper. This construction can therefore be seen as an extension of the suspect game abstraction.

This generic construction can be applied to several frameworks with publicly visible payoff functions. We give two such applications, one with Boolean ω -regular payoff functions and one with mean payoff functions.

¹ This can be expressed by $\text{payoff}_{\text{player } i} = \frac{R}{\text{power}_i} \left(1 - e^{-0.5\gamma_i}\right)^L$ where γ_i is the signal-to-interference-and-noise ratio for player i , R is the rate at which the wireless system transmits the information and L is the size of the packets [29].

Further Related Works. We have already discussed several kinds of related works. Let us give some final remarks on related works.

We have mentioned earlier that one of the problems for computing Nash equilibria is to detect deviations and players who deviated. Somehow, the epistemic game abstraction tracks the potential deviators, and even though players might not know who exactly is responsible for the deviation (there might be several suspects), they can try to punish all potential suspects. And that what we do here. Very recently, [7] discusses the detection of deviators, and give some conditions for them to become common knowledge of the other players. In our framework, even though deviators may not become fully common knowledge, we can design mechanisms to punish the relevant ones.

Recently imperfect information has also been introduced in the setting of multi-agent temporal logics [2, 3, 20, 21], and the main decidability results assume hierarchical information. However, while those logics allow to express rich interactions, it can somehow only consider qualitative properties. Furthermore, no tight complexity bounds are provided.

In [11], a deviator game abstraction is proposed. It twists the suspect game abstraction [9] to allow for more general solution concepts (so-called robust equilibria), but it assumes visibility of actions (hence remove any kind of uncertainties). Relying on results of [13], this deviator game abstraction allows to compute equilibria with mean payoff functions. Our algorithms for mean payoff functions will also rely on the polyhedron problem of [13].

A full version of this paper with all proofs is available as [8]. In this extended abstract, we made the choice to focus on the construction of the epistemic game abstraction and to be more sketchy on algorithms to compute Nash equilibria. We indeed believe the structure of the knowledge represented by the abstraction is the most important contribution, and that algorithms are more standard. However we believe it is important to be able to apply the abstract construction for algorithmic purpose.

2 Definitions

Throughout the paper, if $\mathbb{S} \subseteq \mathbb{R}$, we write $\bar{\mathbb{S}}$ for $\mathbb{S} \cup \{-\infty, +\infty\}$.

2.1 Concurrent Multiplayer Games with Signals

We consider the model of concurrent multi-player games, based on the two-player model of [1]. This model of games was used for instance in [9]. We equip games with *signals*, which will give information to the players.

Definition 1. A concurrent game with signals is a tuple

$$\mathcal{G} = \langle V, v_{\text{init}}, \mathcal{P}, \text{Act}, \Sigma, \text{Allow}, \text{Tab}, (\ell_A)_{A \in \mathcal{P}}, (\text{payoff}_A)_{A \in \mathcal{P}} \rangle$$

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, Act is a finite set of actions, Σ is a finite alphabet, $\text{Allow}: V \times \mathcal{P} \rightarrow$

$2^{\text{Act}} \setminus \{\emptyset\}$ is a mapping indicating the actions available to a given player in a given vertex, $\text{Tab}: V \times \text{Act}^{\mathcal{P}} \rightarrow V$ associates, with a given vertex and a given action tuple the target vertex, for every $A \in \mathcal{P}$, $\ell_A: (\text{Act}^{\mathcal{P}} \times V) \rightarrow \Sigma$ is a signal, and $\text{payoff}_A: V \cdot (\text{Act}^{\mathcal{P}} \cdot V)^{\omega} \rightarrow \mathbb{D}$ is a payoff function with values in a domain $\mathbb{D} \subseteq \overline{\mathbb{R}}$. We say that the game has public signal if there is $\ell: (\text{Act}^{\mathcal{P}} \times V) \rightarrow \Sigma$ such that for every $A \in \mathcal{P}$, $\ell_A = \ell$.

The signals will help the players monitor the game: for taking decisions, a player will have the information given by her signal and the action she played earlier. A public signal will be a common information given to all the players. Our notion of public signal is inspired by [31] and encompasses the model of [9] where only action names were hidden to the players. Note that monitoring by public signal does not mean that all the players have the same information: they have private information implied by their own actions.

An element of $\text{Act}^{\mathcal{P}}$ is called a move. When an explicit order is given on the set of players $\mathcal{P} = \{A_1, \dots, A_{|\mathcal{P}|}\}$, we will write a move $m = (m_A)_{A \in \mathcal{P}}$ as $\langle m_{A_1}, \dots, m_{A_{|\mathcal{P}|}} \rangle$. If $m \in \text{Act}^{\mathcal{P}}$ and $A \in \mathcal{P}$, we write $m(A)$ for the A -component of m and $m(-A)$ for all but the A components of m . In particular, we write $m(-A) = m'(-A)$ whenever $m(B) = m'(B)$ for every $B \in \mathcal{P} \setminus \{A\}$.

A full history h in \mathcal{G} is a finite sequence

$$v_0 \cdot m_0 \cdot v_1 \cdot m_1 \dots m_{k-1} \cdot v_k \in V \cdot (\text{Act}^{\mathcal{P}} \cdot V)^*$$

such that for every $0 \leq i < k$, $m_i \in \text{Allow}(v_i)$ and $v_{i+1} = \text{Tab}(v_i, m_i)$. For readability we will also write h as $v_0 \xrightarrow{m_0} v_1 \xrightarrow{m_1} \dots \xrightarrow{m_{k-1}} v_k$.

We write $\text{last}(h)$ for the last vertex of h (i.e., v_k). If $i \leq k$, we also write $h_{\leq i}$ for the prefix $v_0 \cdot m_0 \cdot v_1 \cdot m_1 \dots m_{i-1} \cdot v_i$. We write $\text{Hist}_{\mathcal{G}}(v_0)$ (or simply $\text{Hist}(v_0)$ if \mathcal{G} is clear in the context) for the set of full histories in \mathcal{G} that start at v_0 .

Let $A \in \mathcal{P}$. The projection of h for A is denoted $\pi_A(h)$ and is defined as:

$$v_0 \cdot (m_0(A), \ell_A(m_0, v_1)) \dots (m_{k-1}(A), \ell_A(m_{k-1}, v_k)) \in V \cdot (\text{Act} \times \Sigma)^*$$

This will be the information available to player A : it contains both the actions she played so far and the signal she received. Note that we assume perfect recall, that is, while playing, A will remember all her past knowledge, that is, all of $\pi_A(h)$ if h has been played so far. We define the *undistinguishability relation* \sim_A as the equivalence relation over full histories induced by π_A : for two histories h and h' , $h \sim_A h'$ iff $\pi_A(h) = \pi_A(h')$. While playing, if $h \sim_A h'$, A will not be able to know whether h or h' has been played. We also define the A -label of h as $\ell_A(h) = \ell_A(m_0, v_1) \cdot \ell_A(m_1, v_2) \dots \ell_A(m_{k-1}, v_k)$.

We extend all the above notions to infinite sequences in a straightforward way and to the notion of *full play*. We write $\text{Plays}_{\mathcal{G}}(v_0)$ (or simply $\text{Plays}(v_0)$ if \mathcal{G} is clear in the context) for the set of full plays in \mathcal{G} that start at v_0 .

We will say that the game \mathcal{G} has *publicly (resp. privately) visible payoffs* if for every $A \in \mathcal{P}$, for every $v_0 \in V$, for every $\rho, \rho' \in \text{Plays}(v_0)$, $\ell_A(\rho) = \ell_A(\rho')$ (resp. $\rho \sim_A \rho'$) implies $\text{payoff}_A(\rho) = \text{payoff}_A(\rho')$. Otherwise they are said *invisible*.

Private visibility of payoffs, while not always assumed (see for instance [3, 19]), are reasonable assumptions: using only her knowledge, a player knows her payoff. Public visibility is more restrictive, but will be required for some of the results.

Let $A \in \mathcal{P}$ be a player. A *strategy* for player A from v_0 is a mapping $\sigma_A: \text{Hist}(v_0) \rightarrow \text{Act}$ such that for every history $h \in \text{Hist}(v_0)$, $\sigma(h) \in \text{Allow}(\text{last}(h))$. It is said ℓ_A -*compatible* whenever furthermore, for all histories $h, h' \in \text{Hist}(v_0)$, $h \sim_A h'$ implies $\sigma_A(h) = \sigma_A(h')$. An *outcome* of σ_A is a (n infinite) play $\rho = v_0 \cdot m_0 \cdot v_1 \cdot m_1 \dots$ such that for every $i \geq 0$, $\sigma_A(\rho_{\leq i}) = m_i(A)$. We write $\text{out}(\sigma_A, v_0)$ for the set of outcomes of σ_A from v_0 .

A *strategy profile* is a tuple $\sigma_{\mathcal{P}} = (\sigma_A)_{A \in \mathcal{P}}$, where, for every player $A \in \mathcal{P}$, σ_A is a strategy for player A . The strategy profile is said *info-compatible* whenever each σ_A is ℓ_A -compatible. We write $\text{out}(\sigma_{\mathcal{P}}, v_0)$ for the unique full play from v_0 , which is an outcome of all strategies part of $\sigma_{\mathcal{P}}$.

When $\sigma_{\mathcal{P}}$ is a strategy profile and σ'_A a player- A strategy, we write $\sigma_{\mathcal{P}}[A/\sigma'_A]$ for the 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*.

Definition 2. A Nash equilibrium from v_0 is an info-compatible strategy profile σ such that for every $A \in \mathcal{P}$, for every player- A ℓ_A -compatible strategy σ'_A , $\text{payoff}_A(\text{out}(\sigma, v_0)) \geq \text{payoff}_A(\text{out}(\sigma[A/\sigma'_A], v_0))$.

In this definition, deviation σ'_A needs not be ℓ_A -compatible, since the only meaningful part of σ'_A is along $\text{out}(\sigma[A/\sigma'_A], v_0)$, where there are no \sim_A -equivalent histories: any deviation can be made ℓ_A -compatible without affecting the profitability of the resulting outcome. Note also that there might be an A -deviation σ'_A which is not observable by another player B ($\text{out}(\sigma, v_0) \sim_B \text{out}(\sigma[A/\sigma'_A], v_0)$), and there might be two deviations σ'_B (by player B) and σ'_C (by player C) that cannot be distinguished by player A ($\text{out}(\sigma[B/\sigma'_B], v_0) \sim_A \text{out}(\sigma[C/\sigma'_C], v_0)$). Tracking such deviations will be the core of the abstraction we will develop.

Payoff Functions. In the following we will consider various payoff functions. Let Φ be an ω -regular property over some alphabet Γ . The function $\text{pay}_{\Phi}: \Gamma^{\omega} \rightarrow \{0, 1\}$ is defined by, for every $\mathbf{a} \in \Gamma^{\omega}$, $\text{pay}_{\Phi}(\mathbf{a}) = 1$ if and only if $\mathbf{a} \models \Phi$. A publicly (resp. privately) visible payoff function payoff_A for player A is said associated with Φ over Σ (resp. $\text{Act} \times \Sigma$) whenever it is defined by $\text{payoff}_A(\rho) = \text{pay}_{\Phi}(\ell_A(\rho))$ (resp. $\text{payoff}_A(\rho) = \text{pay}_{\Phi}(\pi_A(\rho)_{-v_0})$, where $\pi_A(\rho)_{-v_0}$ crops the first v_0). Such a payoff function is called a Boolean ω -regular payoff function.

Let Γ be a finite alphabet and $w: \Gamma \rightarrow \mathbb{Z}$ be a weight assigning a value to every letter of that alphabet. We define two payoff functions over Γ^{ω} by, for every $\mathbf{a} = (a_i)_{i \geq 1} \in \Gamma^{\omega}$, $\text{pay}_{\underline{\text{MP}}_w}(\mathbf{a}) = \liminf_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n w(a_i)$ and $\text{pay}_{\overline{\text{MP}}_w}(\mathbf{a}) = \limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n w(a_i)$. A publicly visible payoff function payoff_A for player A is said associated with the liminf (resp. limsup) mean payoff of w whenever it is defined by $\text{payoff}_A(\rho) = \text{pay}_{\underline{\text{MP}}_w}(\ell_A(\rho))$ (resp. $\text{pay}_{\overline{\text{MP}}_w}(\ell_A(\rho))$). A privately visible payoff function payoff_A for player A is said associated with the liminf (resp. limsup) mean payoff of w whenever it is defined by $\text{payoff}_A(\rho) = \text{pay}_{\underline{\text{MP}}_w}(\pi_A(\rho)_{-v_0})$ (resp. $\text{pay}_{\overline{\text{MP}}_w}(\pi_A(\rho)_{-v_0})$).




Fig. 1. An example of a concurrent game with public signal (yellow and green: public signal). Edges in red and bold are part of the strategy profile. Dashed edges are the possible deviations. One can notice that none of the deviations is profitable to the deviator, hence the strategy profile is a Nash equilibrium. Convention in the drawing: edges with no label are for complementary labels (for instance the edge from v_5 to $0,0,0$ is labelled by all $\langle a_1, a_2, a_3 \rangle$ not in the set $\{\langle a, a, a \rangle, \langle b, a, a \rangle, \langle b, a, b \rangle\}$ (Color figure online))

Example 1. We now illustrate most notions on the game of Fig. 1. This is a game with three players A_1 , A_2 and A_3 , and which is played basically in two steps, starting at v_0 . Graphically an edge labelled $\langle a_1, a_2, a_3 \rangle$ between two vertices v and v' represents the fact that $a_i \in \text{Allow}(v, A_i)$ for every $i \in \{1, 2, 3\}$ and that $v' = \text{Tab}(v, \langle a_1, a_2, a_3 \rangle)$. As a convention, $*$ stands for both a and b . For readability, bottom vertices explicitly indicate the payoffs of the three players (same order as for actions) if the game ends in that vertex.

After the first step of the game, signal yellow or green is sent to all the players. Histories $v_0 \cdot \langle a, b, a \rangle \cdot v_2$ and $v_0 \cdot \langle a, a, a \rangle \cdot v_1$ are undistinguishable by A_1 and A_3 (same action, same signal), but they can be distinguished by A_2 because of different actions (even if the same signal is observed).

In bold red, we have depicted a strategy profile, which is actually a Nash equilibrium. We analyze the possible deviations in this game to argue for this.

- First there is an A_2 -deviation to v_1 . This deviation is invisible to both players A_1 and A_3 . For this reason, the strategy out of v_1 for A_1 is to play a (same as out of v_2). On the other hand, even though this would be profitable to her, A_1 cannot deviate from v_1 , since we are in a branch where A_2 has already deviated, and at most one player is allowed to deviate at a time (and anyway A_1 does not know that they are in state v_1).
- There is an A_1 -deviation from v_2 to $0, 1, 0$, which is not profitable to A_1 .
- On the other hand, there is no possible deviation to v_3 , since this would require two players to change their actions simultaneously (A_1 and A_2).
- Then, there is an A_1 -deviation to v_4 and another A_3 -deviation to v_5 ; both activate the green signal. A_2 knows there has been a deviation (because of the green signal), but she doesn't know who has deviated and whether the

game proceeds to v_4 or v_5 (but she knows that if A_1 has deviated, then we are in v_4 , and if A_3 has deviated, we are in v_5). Then, A_2 has to find a way to punish both players, to be safe. On the other hand, both players A_1 and A_3 precisely know what has happened: in case she didn't deviate herself, she knows the other one deviated! And she knows in which state the game is in. Hence in state v_4 , A_3 can help player A_2 punishing A_1 , whereas in state v_5 , A_1 can help player A_2 punishing A_3 . Examples of punishing moves are therefore those depicted in red and bold; and they are part of the global strategy profile. Note that the action of A_2 out of v_5 has to be the same as the one out of v_4 : this is required given the imperfect knowledge of A_2 . On the other hand, the action of A_3 can be different out of v_4 and out of v_5 (which is the case in the given example profile).

Two-Player Turn-Based Game Structures. They are specific cases of the previous model, where at each vertex, at most one player has more than one action in her set of allowed actions. But for convenience, we will give a simplified definition, with only objects that will be useful. A two-player turn-based game structure is a tuple $G = \langle S, S_{\text{Eve}}, S_{\text{Adam}}, s_{\text{init}}, A, \text{Allow}, \text{Tab} \rangle$, where $S = S_{\text{Eve}} \sqcup S_{\text{Adam}}$ is a finite set of states (states in S_{Eve} belong to player Eve whereas states in S_{Adam} belong to player Adam), $s_{\text{init}} \in S$ is the initial state, A is a finite alphabet, $\text{Allow}: S \rightarrow 2^A \setminus \{\emptyset\}$ gives the set of available actions, and $\text{Tab}: S \times A \rightarrow S$ is the next-state function. If $s \in S_{\text{Eve}}$ (resp. S_{Adam}), $\text{Allow}(s)$ is the set of actions allowed to Eve (resp. Adam) in state s .

In this context, strategies will see sequences of states and actions, with full information. Note that we do not include any winning condition or payoff function in the tuple, hence the name structure.

2.2 The Problem

We are interested in the constrained existence of a Nash equilibrium. For simplicity, we define constraints using non-strict thresholds constraints, but could well impose more involved constraints.

Problem 1 (Constrained existence problem). Given a game with signals $\mathcal{G} = \langle V, v_{\text{init}}, \mathcal{P}, \text{Act}, \Sigma, \text{Allow}, \text{Tab}, (\ell_A)_{A \in \mathcal{P}}, (\text{payoff}_A)_{A \in \mathcal{P}} \rangle$ and threshold vectors $(\nu_A)_{A \in \mathcal{P}}, (\nu'_A)_{A \in \mathcal{P}} \in \overline{\mathbb{Q}}^{\mathcal{P}}$, can we decide whether there exists a Nash equilibrium $\sigma_{\mathcal{P}}$ from v_{init} such that for every $A \in \mathcal{P}$, $\nu_A \leq \text{payoff}_A(\text{out}(\sigma_{\mathcal{P}}, v_{\text{init}})) \leq \nu'_A$? If so, compute one. If the constraints on the payoff are trivial (that is, $\nu_A = -\infty$ and $\nu'_A = +\infty$ for every $A \in \mathcal{P}$), we simply speak of the existence problem.

2.3 First Undecidability Results

In this section we state two preliminary undecidability results.

Theorem 1. – *The existence problem in games with signals is undecidable with three players and publicly visible Boolean ω -regular payoff functions.*

- *The constrained existence problem in games with a public signal is undecidable with two players and privately visible mean payoff functions.*

Proofs of these results rely on the distributed synthesis problem [26] for the first one, and on blind two-player mean-payoff games [19] for the second one. While there is no real surprise in the first result since we know that arbitrary partial information yields intrinsic difficulties, the second one suggests restrictions both to public signals and to publicly visible payoff functions.

In the following we will focus on public signals and develop an epistemic game abstraction, which will record and track possible deviations in the game. This will then be applied to get decidability results in two frameworks assuming publicly visible payoff functions.

3 The Epistemic Game Abstraction

Building over [4,9], we construct an epistemic game, which will record possible behaviours of the system, together with possible unilateral deviations. In [4], notions of epistemic Kripke structures are used to really track the precise knowledge of the players. These are mostly useful since undistinguishable states (expressed using signals here) are assumed arbitrary (no hierarchical structure). We could do the same here, but we think that would be overly complex and hide the real structure of knowledge in the framework of public signals. We therefore prefer to stick to simpler subset constructions, which are more commonly used (see e.g. [27] or later [14,19,22]), though it has to be a bit more involved here since also deviations have to be tracked.

Let $\mathcal{G} = \langle V, v_{\text{init}}, \mathcal{P}, \text{Act}, \Sigma, \text{Allow}, \text{Tab}, \ell, (\text{payoff}_A)_{A \in \mathcal{P}} \rangle$ be a concurrent game with public signal. We will first define the epistemic abstraction as a two-player game structure $\mathcal{E}_{\mathcal{G}} = \langle S_{\text{Eve}}, S_{\text{Adam}}, s_{\text{init}}, \Sigma', \text{Allow}', \text{Tab}' \rangle$, and then state the correspondence between \mathcal{G} and $\mathcal{E}_{\mathcal{G}}$. The epistemic abstraction will later be used for decidability and algorithmics purposes. For clarity, we use the terminology “vertices” in \mathcal{G} and “states” (or “epistemic states”) in $\mathcal{E}_{\mathcal{G}}$.

3.1 Construction of the Game Structure $\mathcal{E}_{\mathcal{G}}$

The game $\mathcal{E}_{\mathcal{G}}$ will be played between two players, **Eve** and **Adam**. The aim of **Eve** is to build a suitable Nash equilibrium, whereas the aim of **Adam** is to prove that it is not an equilibrium; in particular, **Adam** will try to find a profitable deviation (to disprove the claim of **Eve** that she is building a Nash equilibrium). Choices available to **Eve** and **Adam** in the abstract game have to reflect partial knowledge of the players in the original game \mathcal{G} . States in the abstract game will therefore store information, which will be sufficient to infer the undistinguishability relation of all the players in the original game. Thanks to the public signal assumption, this information will be simple enough to have a simple structure.

In the following, we set $\mathcal{P}^{\perp} = \mathcal{P} \cup \{\perp\}$, where \perp is a fresh symbol. For convenience, if $m \in \text{Act}^{\mathcal{P}}$, we extend the notation $m(-A)$ when $A \in \mathcal{P}$ to \mathcal{P}^{\perp} by setting $m(-\perp) = m$. We now describe all the components of $\mathcal{E}_{\mathcal{G}}$.

A state of **Eve** will store a set of vertices of the original game one can be in, together with possible deviators. More precisely, states of **Eve** are defined as $S_{\text{Eve}} = \{s: \mathcal{P}^\perp \rightarrow 2^V \mid |s(\perp)| \leq 1\}$. Let $s \in S_{\text{Eve}}$. If $A \in \mathcal{P}$, vertices of $s(A)$ are those where the game can be in, assuming one has followed the suggestions of **Eve** so far, up to an A -deviation; on the other hand, if $s(\perp) \neq \emptyset$, the single vertex $v \in s(\perp)$ is the one the game is in, assuming one has followed all suggestions by **Eve** so far (in particular, if **Eve** is building a Nash equilibrium, then this vertex belongs to the main outcome of the equilibrium). We define $\text{sit}(s) = \{(v, A) \in V \times \mathcal{P}^\perp \mid v \in s(A)\}$ for the set of *situations* the game can be in at s :

- (a) $(v, \perp) \in \text{sit}(s)$ is the situation where the game has proceeded to vertex v without any deviation;
- (b) $(v, A) \in \text{sit}(s)$ with $A \in \mathcal{P}$ is the situation where the game has proceeded to vertex v benefitting, from an A -deviation.

Structure of state s will allow to infer the undistinguishability relation of all the players in game \mathcal{G} : basically (and we will formalize this later), if she is not responsible for a deviation, player $A \in \mathcal{P}$ will not know in which of the situations of $\text{sit}(s) \setminus V \times \{A\}$ the game has proceeded; if she is responsible for a deviation, player A will know exactly in which vertex $v \in s(A)$ the game has proceeded.

Let $s \in S_{\text{Eve}}$. From state s , **Eve** will suggest a tuple of moves M , one for each possible situation $(v, A) \in \text{sit}(s)$. This tuple of moves has to satisfy the undistinguishability relation: if a player does not distinguish between two situations, her action should be the same in these two situations:

$$\text{Allow}'(s) = \left\{ M \in \prod_{(v,A) \in \text{sit}(s)} \text{Allow}(v) \mid \forall (v_B, B), (v_C, C) \in \text{sit}(s), \forall A \in \mathcal{P} \setminus \{B, C\}, M(v_B, B)(A) = M(v_C, C)(A) \right\}$$

In the above set, the constraint $M(v_B, B)(A) = M(v_C, C)(A)$ expresses the fact that player A should play the same action in the two situations (v_B, B) and (v_C, C) , since she does not distinguish between them. Obviously, we assume Σ' contains all elements of $\text{Allow}'(s)$ above.

States of **Adam** are then copies of states of **Eve** with suggestions given by **Eve**, that is: $S_{\text{Adam}} = \{(s, M) \mid s \in S_{\text{Eve}} \times \text{Allow}'(s)\}$. And naturally, we define $\text{Tab}'(s, M) = (s, M)$ if $M \in \text{Allow}'(s)$.

Let $(s, M) \in S_{\text{Adam}}$. From state (s, M) , **Adam** will choose a signal value which can be activated from some situation allowed in s , after no deviation or a single-player deviation w.r.t. M . From a situation $(v, A) \in \text{sit}(s)$ with $A \in \mathcal{P}$, only A -deviations can be allowed (since we look for unilateral deviations), hence any signal activated by an A -deviation (w.r.t. $M(v, A)$) from v should be allowed. From the situation $(v, \perp) \in \text{sit}(s)$ (if there is one), one can continue without any deviation, or any kind of single-player deviation should be allowed, hence the signal activated by $M(v, \perp)$ from v should be allowed, and any signal activated by some A -deviation (w.r.t. $M(v, \perp)$) from v should be allowed as well. Formally:

$$\text{Allow}'(s, M) = \left\{ \beta \in \Sigma \left| \begin{array}{l} \exists A \in \mathcal{P} \\ \exists v \in s(A) \quad \text{s.t.} \quad \begin{array}{l} \text{(i)} \quad m(-A) = M(v, A)(-A) \\ \text{(ii)} \quad \ell(m, \text{Tab}(v, m)) = \beta \end{array} \\ \exists m \in \text{Act}^{\mathcal{P}} \end{array} \right. \right\} \\ \cup \left\{ \beta \in \Sigma \left| \begin{array}{l} \exists v \in s(\perp) \\ \exists m \in \text{Act}^{\mathcal{P}} \quad \text{s.t.} \quad \begin{array}{l} \text{(i)} \quad m(-A) = M(v, \perp)(-A) \\ \text{(ii)} \quad \ell(m, \text{Tab}(v, m)) = \beta \end{array} \\ \exists A \in \mathcal{P} \end{array} \right. \right\}$$

Note that we implicitly assume that Σ' contains Σ .

It remains to explain how one can compute the next state of some $(s, M) \in S_{\text{Adam}}$ after some signal value $\beta \in \text{Allow}'(s, M)$. The new state has to represent the new knowledge of the players in the original game when they have seen signal β ; this has to take into account all possible deviations that we have already discussed which activate the signal value β . The new state is the result of several simultaneous subset constructions, which we formalize as follows: $s' = \text{Tab}'((s, M), \beta)$, where for every $A \in \mathcal{P}^{\perp}$, $v' \in s'(A)$ if and only if there is $m \in \text{Act}^{\mathcal{P}}$ such that $\beta = \ell(m, v')$, and

1. either there is $v \in s(A)$ such that $m(-A) = M(v, A)(-A)$ and $v' = \text{Tab}(v, m)$;
2. or there is $v \in s(\perp)$ such that $m(-A) = M(v, \perp)(-A)$ and $v' = \text{Tab}(v, m)$.

Note that in case $A = \perp$, the two above cases are redundant.

Before stating properties of $\mathcal{E}_{\mathcal{G}}$, we illustrate the construction.

Example 2. We consider again the example of Fig. 1, and we assume that the public signal when reaching the leaves of the game is uniformly orange. We depict (part of) the epistemic game abstraction of the game on Fig. 2. One can notice that from Eve-states s_1 and s_2 , moves are multi-dimensional, in the sense that there is one move per vertex appearing in the state. There are nevertheless compatibility conditions which should be satisfied (expressed in condition Allow'); for instance, from s_2 , player A_2 does not distinguish between the two options (i) A_1 has deviated and the game is in v_4 , and (ii) A_3 has deviated and the game is in v_5 , hence the action of player A_2 should be the same in the two moves (a in the depicted example, written in red).

3.2 Interpretation of this Abstraction

While we gave an intuitive meaning to the (epistemic) states of $\mathcal{E}_{\mathcal{G}}$, we now need to formalize this. And to do that, we need to explain how full histories and plays in $\mathcal{E}_{\mathcal{G}}$ can be interpreted as full histories and plays in \mathcal{G} .

Let $v_0 \in V$, and define $s_0: \mathcal{P}^{\perp} \rightarrow 2^V \in S_{\text{Eve}}$ such that $s_0(\perp) = \{v_0\}$ and $s_0(A) = \emptyset$ for every $A \in \mathcal{P}$. In the following, when $M \in \text{Allow}'(s)$ for some $s \in S_{\text{Eve}}$, if we speak of some $M(v, A)$, we implicitly assume that $(v, A) \in \text{sit}(s)$. Given a full history $H = s_0 \xrightarrow{M_0} (s_0, M_0) \xrightarrow{\beta_0} s_1 \xrightarrow{M_1} (s_1, M_1) \xrightarrow{\beta_1} s_2 \dots (s_{k-1}, M_{k-1}) \xrightarrow{\beta_{k-1}} s_k$ in $\mathcal{E}_{\mathcal{G}}$, we write $\text{concrete}(H)$ for the set of full histories in the original game, which correspond to H , up to a single deviation, that is: $v_0 \xrightarrow{m_0} v_1 \xrightarrow{m_1} v_2 \dots v_{k-1} \xrightarrow{m_{k-1}} v_k \in \text{concrete}(H)$ whenever for every $0 \leq i \leq k-1$, $v_{i+1} = \text{Tab}(v_i, m_i)$ and $\beta_i = \ell(m_i, v_{i+1})$, and:




Fig. 2. Part of the epistemic game corresponding to the game of Fig. 1. For clarity, symbol $-$ is for any choice a or b (the precise choice is meaningless). (Color figure online)

- (a) either $m_i = M_i(v_i, \perp)$ for every $0 \leq i \leq k - 1$;
- (b) or there exist $A \in \mathcal{P}$ and $0 \leq i_0 \leq k - 1$ such that
 - (i) for every $0 \leq i < i_0$, $m_i = M_i(v_i, \perp)$;
 - (ii) $m_{i_0} \neq M_{i_0}(v_{i_0}, \perp)$, but $m_{i_0}(-A) = M_{i_0}(v_{i_0}, \perp)(-A)$;
 - (iii) for every $i_0 < i \leq k - 1$, $m_i(-A) = M_i(v_i, A)(-A)$.

Case (a) corresponds to a concrete history with no deviation (all moves suggested by **Eve** have been followed). Case (b) corresponds to a deviation by player A , and i_0 is the position at which player A has started deviating.

We write $concrete_{\perp}(H)$ for the set of histories of type (a); there is at most one such history, which is the real concrete history suggested by **Eve**. And we write $concrete_A(H)$ for the set of histories of the type (b) with deviator A . The correctness of the approach is obtained thanks to the following characterization of the undistinguishability relations along H : for every $A \in \mathcal{P}$, for every $h_1 \neq h_2 \in concrete(H)$,

$$h_1 \sim_A h_2 \text{ iff } h_1, h_2 \notin concrete_A(H).$$

In particular, a player may not distinguish between deviations by other players, or between a deviation by another player and the real concrete history suggested by **Eve**. But of course, in any case, a player will know that she has deviated!

We extend all these notions to full plays. A full play visiting only **Eve**-states s such that $s(\perp) \neq \emptyset$ is called a \perp -play.

3.3 Winning Condition of Eve

A zero-sum game will be played on the game structure $\mathcal{E}_{\mathcal{G}}$, and the winning condition of **Eve** will be given on the branching structure of the set of outcomes of a strategy for **Eve**, and not individually on each outcome, as standardly in two-player zero-sum games. We write s_{init} for the state of **Eve** such that $s_{\text{init}}(\perp) = \{v_{\text{init}}\}$ and $s_{\text{init}}(A) = \emptyset$ for every $A \in \mathcal{P}$. Let $p = (p_A)_{A \in \mathcal{P}} \in \overline{\mathbb{R}}^{\mathcal{P}}$, and σ_{Eve} be a strategy for **Eve** in $\mathcal{E}_{\mathcal{G}}$; it is said *winning* for p from s_{init} whenever $\text{payoff}(\rho) = p$, where ρ is the unique element of $\text{concrete}_{\perp}(\text{out}_{\perp}(\sigma_{\text{Eve}}, s_{\text{init}}))$ (where we write $\text{out}_{\perp}(\sigma_{\text{Eve}}, s_{\text{init}})$ for the unique outcome of σ_{Eve} from s_{init} which is a \perp -play), and for every $R \in \text{out}(\sigma_{\text{Eve}}, s_{\text{init}})$, for every $A \in \mathcal{P}$, for every $\rho \in \text{concrete}_A(R)$, $\text{payoff}_A(\rho) \leq p_A$.

For every epistemic state $s \in S_{\text{Eve}}$, we define the set of *suspect* players $\text{susp}(s) = \{A \in \mathcal{P} \mid s(A) \neq \emptyset\}$ (this is the set of players that may have deviated).

By extension, if $R = s_0 \xrightarrow{M_0} (s_0, M_0) \xrightarrow{\beta_0} s_1 \dots s_k \xrightarrow{M_k} (s_k, M_k) \xrightarrow{\beta_k} s_{k+1} \dots$, we define $\text{susp}(R) = \lim_{k \rightarrow \infty} \text{susp}(s_k)$. Note that the sequence $(\text{susp}(s_k))_k$ is non-increasing, hence it stabilizes.

Assuming public visibility of the payoff functions in \mathcal{G} , we can define when R is a full play in $\mathcal{E}_{\mathcal{G}}$, and $A \in \mathcal{P}$, $\text{payoff}'_A(R) = \text{payoff}_A(\rho)$, where $\rho \in \text{concrete}(R)$. It is easy to show that payoff'_A is well-defined for every $A \in \mathcal{P}$. Under this assumption, the winning condition of **Eve** can be rewritten as: σ_{Eve} is winning for p from s_{init} whenever $\text{payoff}'(\text{out}_{\perp}(\sigma_{\text{Eve}}, s_{\text{init}})) = p$, and for every $R \in \text{out}(\sigma_{\text{Eve}}, s_{\text{init}})$, for every $A \in \text{susp}(R)$, $\text{payoff}'_A(R) \leq p_A$.

3.4 Correction of the Epistemic Abstraction

The epistemic abstraction tracks everything that is required to detect Nash equilibria in the original game, which we make explicit in the next result. Note that this theorem does not require public visibility of the payoff functions.

Theorem 2. *Let \mathcal{G} be a concurrent game with public signal, and $p \in \overline{\mathbb{R}}^{\mathcal{P}}$. There is a Nash equilibrium in \mathcal{G} with payoff p from v_{init} if and only if **Eve** has a winning strategy for p in $\mathcal{E}_{\mathcal{G}}$ from s_{init} .*

The proof of this theorem highlights a correspondence between Nash equilibria in \mathcal{G} and winning strategies of **Eve** in $\mathcal{E}_{\mathcal{G}}$. In this correspondence, the main outcome of the equilibrium in \mathcal{G} is the unique \perp -concretisation of the unique \perp -play generated by the winning strategy of **Eve**.

3.5 Remarks on the Construction

We did not formalize the epistemic unfolding as it is made in [4]. We believe we do not really learn anything for public signal using it. And the above extended subset construction can much better be understood.

One could argue that this epistemic game gives more information to the players, since **Eve** explicitly gives to everyone the move that should be played.

But in the real game, the players also have that information, which is obtained by an initial coordination of the players (this is required to achieve equilibria).

Finally, notice that the epistemic game constructed here generalizes the suspect game construction of [9], where all players have perfect information on the states of the game, but cannot see the actions that are precisely played. Somehow, games in [9] have a public signal telling the state the game is in (that is, $\ell(m, v) = v$). So, in the suspect game of [9], the sole uncertainty is in the players that may have deviated, not in the set of states that are visited.

Remark 1. Let us analyze the size of the epistemic game abstraction. The size of the alphabet is bounded by $|\Sigma| + |\text{Act}|^{|\mathcal{X}| \cdot |V| \cdot (1 + |\mathcal{X}|)}$. Furthermore, $|\Sigma|$ is bounded by $|V| \cdot |\text{Act}|^{|\mathcal{X}|}$. The number of states is therefore in $O(2^{|\mathcal{X}| \cdot |V|} \cdot |\text{Act}|^{|\mathcal{X}|^2 \cdot |V|})$. The epistemic game is therefore of exponential size w.r.t. the initial game. Note that we could reduce the bounds by using tricks like those in [9, Proposition 4.8], but this would not avoid an exponential blowup.

4 Two Applications with Publicly Visible Payoffs

While the construction of the epistemic game has transformed the computation of Nash equilibria in a concurrent game with public signal to the computation of winning strategies in a two-player zero-sum turn-based game, we cannot apply standard algorithms out-of-the-box, because the winning condition is rather complex. In the following, we present two applications of that approach in the context of publicly visible payoffs, one with Boolean payoff functions, and another with mean payoff functions. Remember that in the latter case, public visibility is required to have decidability (Theorem 1).

The epistemic game has a specific structure, which can be used for algorithmic purpose. The main outcome of a potential Nash equilibrium is given by a \perp -play, that is, a play visiting only epistemic states s with $s(\perp) \neq \emptyset$. There are now two types of deviations:

- (i) those that are invisible to all players (except the deviator): they are tracked along the main \perp -play. Assuming public visibility of the payoff functions, such a deviation cannot be profitable to any of the players (the payoff of all concrete plays along that \perp -play coincides with the payoff of the main outcome), hence no specific punishing strategy has to be played.
- (ii) those that leave the main \perp -play at some point, and visit only epistemic states s such that $s(\perp) = \emptyset$ from that point on: those are the deviations that need to be punished. Note nevertheless that the deviator may not precisely be known by all the players, hence punishing strategies need to take this into account. However, the set of potential deviators along a deviating play is non-increasing, and we can solve subgames with specific subsets of potential deviators separately (e.g. in a bottom-up approach). The winning objectives in those subgames will depend on the payoff functions (and will mostly be conjunctions of constraints on those functions), and also on the value of those payoff functions along the main outcome.

Using such an approach and results of [16] on generalized parity games, we obtain the following result for Boolean ω -regular payoff functions:

Theorem 3. *The constrained existence problem is in EXPSpace and EXPTIME-hard for concurrent games with public signal and publicly visible Boolean payoff functions associated with parity conditions. The lower bound holds even for Büchi conditions and two players.*

The same approach could be used for the ordered objectives of [9], which are finite preference relations over sets of ω -regular properties. Also, we believe we can enrich the epistemic game construction and provide an algorithm to decide the constrained existence problem for Boolean ω -regular invisible payoff functions.

We have also investigated publicly visible mean payoff functions. While we could have used the same bottom-up approach as above and applied results from [12, 13], we adopt an approach similar to that of [11], which consists in transforming the winning condition of Eve in $\mathcal{E}_{\mathcal{G}}$ into a so-called *polyhedron query* in a multi-dimensional mean-payoff game. Given such a game, a polyhedron query asks whether there exists a strategy for Eve which achieves a payoff belonging to some given polyhedron. Using this approach, we get the following result:

Theorem 4. *The constrained existence problem is in $\text{NP}^{\text{NEXPTIME}}$ (hence in EXPSpace) and EXPTIME-hard for concurrent games with public signal and publicly visible mean payoff functions.*

5 Conclusion

In this paper, we have studied concurrent games with imperfect monitoring modelled using signals. We have given some undecidability results, even in the case of public signals, when the payoff functions are not publicly visible. We have then proposed a construction to capture single-player deviations in games with public signals, and reduced the search of Nash equilibria to the synthesis of winning strategies in a two-player turn-based games (with a rather complex winning condition though). We have applied this general framework to two classes of payoff functions, and obtained decidability results.

As further work we wish to understand better if there could be richer communication patterns which would allow representable knowledge structures for Nash equilibria and thereby the synthesis of Nash equilibria under imperfect monitoring. A source of inspiration for further work will be [28].

References

1. Alur, R., Henzinger, T., Kupferman, O.: Alternating-time temporal logic. *J. ACM* **49**, 672–713 (2002)
2. Berthon, R., Maubert, B., Murano, A.: Decidability results for ATL* with imperfect information and perfect recall. In: *Proceedings of 16th Conference Autonomous Agents and Multiagent Systems (AAMAS 2017)*, pp. 1250–1258. ACM (2017)
3. Berthon, R., Maubert, B., Murano, A., Rubin, S., Vardi, M.Y.: Strategy logic with imperfect information. In: *Proceedings of 32nd Annual Symposium Logic in Computer Science (LICS 2017)*, pp. 1–12. IEEE Computer Society Press (2017)
4. Berwanger, D., Kaiser, L., Puchala, B.: Perfect-information construction for coordination in games. In: *Proceedings of 30th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2011)*, LIPIcs, vol. 13, pp. 387–398. LZI (2011)
5. Berwanger, D., Mathew, A.B.: Infinite games with finite knowledge gaps. *Inf. Comput.* **254**, 217–237 (2017)
6. Berwanger, D., Mathew, A.B., Van den Bogaard, M.: Hierarchical information and the synthesis of distributed strategies. *Acta Informatica*, 1–33. Springer, Heidelberg (2017). <https://doi.org/10.1007/s00236-017-0306-5>
7. Berwanger, D., Ramanujam, R.: Deviator detection under imperfect monitoring. In: *Proceedings of 5th International Workshop Strategic Reasoning (SR 2017)* (2017)
8. Bouyer, P.: Games on graphs with a public signal monitoring. Research report, arXiv <https://arxiv.org/abs/1710.07163> (2017)
9. Bouyer, P., Brenguier, R., Markey, N., Ummels, M.: Pure Nash equilibria in concurrent games. *Log. Methods Comput. Sci.* **11**(2:9) (2015). [https://doi.org/10.2168/LMCS-11\(2:9\)2015](https://doi.org/10.2168/LMCS-11(2:9)2015)
10. Bouyer, P., Markey, N., Stan, D.: Mixed Nash equilibria in concurrent games. In: *Proceedings of 33rd Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2014)*, LIPIcs, vol. 29, pp. 351–363. LZI (2014)
11. Brenguier, R.: Robust equilibria in mean-payoff games. In: Jacobs, B., Löding, C. (eds.) *FoSSaCS 2016*. LNCS, vol. 9634, pp. 217–233. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49630-5_13
12. Brenguier, R., Raskin, J.-F.: Optimal values of multidimensional mean-payoff games. Research report hal-00977352, Université Libre de Bruxelles, Belgium (2014). <https://hal.archives-ouvertes.fr/hal-00977352>
13. Brenguier, R., Raskin, J.-F.: Pareto curves of multidimensional mean-payoff games. In: Kroening, D., Păsăreanu, C.S. (eds.) *CAV 2015, Part II*. LNCS, vol. 9207, pp. 251–267. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-21668-3_15
14. Chatterjee, K., Doyen, L., Henzinger, T., Raskin, J.-F.: Algorithms for ω -regular games with imperfect information. *Log. Methods Comput. Sci.* **3**(3) (2007). [https://doi.org/10.2168/LMCS-3\(3:4\)2007](https://doi.org/10.2168/LMCS-3(3:4)2007)
15. Chatterjee, K., Henzinger, T., Jurdziński, M.: Games with secure equilibria. *Theor. Comput. Sci.* **365**(1–2), 67–82 (2006)
16. Chatterjee, K., Henzinger, T.A., Piterman, N.: Generalized parity games. In: Seidl, H. (ed.) *FoSSaCS 2007*. LNCS, vol. 4423, pp. 153–167. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-71389-0_12
17. Chatterjee, K., Majumdar, R., Jurdziński, M.: On Nash equilibria in stochastic games. In: Marcinkowski, J., Tarlecki, A. (eds.) *CSL 2004*. LNCS, vol. 3210, pp. 26–40. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-30124-0_6

18. Condurache, R., Filiot, E., Gentilini, R., Raskin, J.-F.: The complexity of rational synthesis. In: Proceedings of 43rd International Colloquium on Automata, Languages and Programming (ICALP 2016), LIPIcs, vol. 55, pp. 121:1–121:15. Leibniz-Zentrum für Informatik (2016)
19. Degorre, A., Doyen, L., Gentilini, R., Raskin, J.-F., Toruńczyk, S.: Energy and mean-payoff games with imperfect information. In: Dawar, A., Veith, H. (eds.) CSL 2010. LNCS, vol. 6247, pp. 260–274. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-15205-4_22
20. Dima, C., Enea, C., Guelev, D.P.: Model-checking an alternating-time temporal logic with knowledge, imperfect information, perfect recall and communicating coalitions. In: Proceedings of 1st International Symposium Games, Automata, Logics and Formal Verification (GandALF 2010), Electronic Proceedings in Theoretical Computer Science, vol. 25, pp. 103–117 (2010)
21. Dima, C., Tiplea, F.L.: Model-checking ATL under imperfect information and perfect recall semantics is undecidable. Research report arXiv, <http://arxiv.org/abs/1102.4225> (2011)
22. Doyen, L., Raskin, J.-F.: Games with imperfect information: theory and algorithms. Lectures in Game Theory for Computer Scientists, pp. 185–212. Cambridge University Press (2011)
23. Henzinger, T.: Games in system design and verification. In: Proceedings of 10th Conference Theoretical Aspects of Rationality and Knowledge (TARK 2005), pp. 1–4 (2005)
24. Nash, J.F.: Equilibrium points in n -person games. Proc. Nat. Acad. Sci. U.S.A. **36**(1), 48–49 (1950)
25. Peterson, G.L., Reif, J.H.: Multiple-person alternation. In: Proceedings of 20th Annual Symposium on Foundations of Computer Science (FOCS 1979), pp. 348–363. IEEE Computer Society Press (1979)
26. Pnueli, A., Rosner, R.: Distributed reactive systems are hard to synthesize. In: Proceedings of 31st Annual Symposium on Foundations of Computer Science (FOCS 1990), pp. 746–757. IEEE Computer Society Press (1990)
27. Reif, J.H.: The complexity of two-player games of incomplete information. J. Comput. System Sciences **29**(2), 274–301 (1984)
28. Renault, J., Tomala, T.: Repeated proximity games. Int. J. Game Theory **27**(4), 539–559 (1998)
29. Saraydar, C.U., Mandayam, N.B., Goodman, D.J.: Pareto efficiency of pricing-based power control in wireless data networks. In: Proceedings of IEEE Wireless Comm. and Networking Conference (WCNC 1999), pp. 231–235. IEEE Computer Society Press (1999)
30. Thomas, W.: Infinite games and verification. In: Brinksma, E., Larsen, K.G. (eds.) CAV 2002. LNCS, vol. 2404, pp. 58–65. Springer, Heidelberg (2002). https://doi.org/10.1007/3-540-45657-0_5
31. Tomala, T.: Pure equilibria of repeated games with public observation. Int. J. Game Theory **27**(1), 93–109 (1998)
32. Ummels, M.: Rational behaviour and strategy construction in infinite multiplayer games. In: Arun-Kumar, S., Garg, N. (eds.) FSTTCS 2006. LNCS, vol. 4337, pp. 212–223. Springer, Heidelberg (2006). https://doi.org/10.1007/11944836_21
33. Ummels, M.: The complexity of Nash equilibria in infinite multiplayer games. In: Amadio, R. (ed.) FoSSaCS 2008. LNCS, vol. 4962, pp. 20–34. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-78499-9_3

34. Ummels, M., Wojtczak, D.: The complexity of Nash equilibria in limit-average games. In: Katoen, J.-P., König, B. (eds.) CONCUR 2011. LNCS, vol. 6901, pp. 482–496. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-23217-6_32
35. Ummels, M., Wojtczak, D.: The complexity of Nash equilibria in stochastic multiplayer games. *Log. Methods Comput. Sci.* **7**(3) (2011). [https://doi.org/10.2168/LMCS-7\(3:20\)2011](https://doi.org/10.2168/LMCS-7(3:20)2011)
36. van der Meyden, R., Wilke, T.: Synthesis of distributed systems from knowledge-based specifications. In: Abadi, M., de Alfaro, L. (eds.) CONCUR 2005. LNCS, vol. 3653, pp. 562–576. Springer, Heidelberg (2005). https://doi.org/10.1007/11539452_42

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

