

Deviator Detection under Imperfect Monitoring

Extended Abstract

Dietmar Berwanger

CNRS, ENS Paris-Saclay, France

dwb@lsv.fr

R. Ramanujam

The Institute of Mathematical Sciences, India

jam@imsc.res.in

Grim-trigger strategies are a fundamental mechanism for sustaining equilibria in iterated games: The players cooperate along an agreed path, and as soon as one player deviates, the others form a coalition to play him down to his minmax level. A precondition to triggering such a strategy is that the identity of the deviating player becomes common knowledge among the other players. This can be difficult or impossible to attain in games where the information structure allows only imperfect monitoring of the played actions or of the global state.

We study the problem of synthesising finite-state strategies for detecting the deviator from an agreed strategy profile in games played on finite graphs with different information structures. We show that the problem is undecidable in the general case where the global state cannot be monitored. On the other hand, we prove that under perfect monitoring of the global state and imperfect monitoring of actions, the problem becomes decidable, and we present an effective synthesis procedure that covers infinitely repeated games with private monitoring.

1 Introduction

In social situations, a *queue* acts in a self-stabilising manner: when anyone tries to jump the queue, the others give him a dirty look, and often this suffices to enforce the rule. Distributed protocols are often designed to be self-stabilising in a similar sense: when a fault occurs, it is detected and isolated, and the system recovers to a state in such a way that the computation can proceed.

In general, the design of such systems assumes that all processes cooperate and coordinate their actions towards achieving the system goals. When a process deviates from the protocol (perhaps due to faults), the system needs to detect it and recover from the situation. Usually, if such deviation occurs in an isolated one-off manner, it may be hard to detect, but when they occur repeatedly, there is the possibility of detecting the culprit(s). This idea has been applied to intruder detection in security theory.

There are several interesting variations on this theme. One is to ask for protocols or solutions that do not demand detection of deviators but to merely provide resilience. Typical solutions in distributed computing assume a bound on the number of faulty processes and provide solutions that can tolerate that many failures. Another consideration relates to the observational limitations of processes. In a distributed system, processes have only a partial view of the system state, and often cannot observe all the moves of other processes, and these may help the deviator to evade detection successfully.

When distributed computing meets game theory, we have further interesting possibilities to consider. A player may act selfishly if it maximises her payoff, even if this means a deviation. In general the coalition may have no way to hold a member of the coalition to act according to prior commitments (dismissively labelled “cheap talk”). However, in the case of repeated play, there can be threats and punishment to ensure that members do not deviate. In this context, a variety of mechanisms are studied in game theory, notable among them the *grim-trigger strategies*: start out co-operating, and when any

player deviates, punish for ever after. The importance of such strategies is that they play a central role in what are referred to as Folk theorems in game theory.

Nash equilibria provide a robust way to predict how rational players would act offering their best responses to their beliefs about how others might act, in situations where players differ in their knowledge of what others might be doing. At Nash equilibrium, no player has an incentive to deviate unilaterally and shift to a different strategy. Folk theorems then assert that for any Nash equilibrium outcome vector (o_1, \dots, o_n) in an n -player infinitely repeated game (with, say, limit average rewards), then each player i can force the outcome o_i . Conversely every “feasible” and enforceable outcome is the outcome of some Nash equilibrium.

Interestingly, these (and many related) results depend on the ability of players to perfectly monitor the actions of the deviator. In cases where players’ actions may not be directly observable, special solutions are needed, and game theorists have developed a powerful set of techniques for many subclasses of games [9, 12, 1, 14]. Note that this situation pertains more directly to distributed systems where processes are limited in their ability to observe the global state and to record other processes’ actions.

In this paper, we focus on computational questions related to deviator detection under imperfect monitoring. Can we have algorithms that determine, in a game, whether the deviator’s identity be rendered common knowledge? If yes, we could like to construct strategies (for the other players) to achieve this. This problem is related to solving consensus problems on graphs with different communication topologies.

We propose a general technique based on methods from automated synthesis towards achieving an epistemic objective (involving common knowledge). With imperfect information, the synthesis problem is undecidable already with simple reachability objectives. So one expects only negative results for deviator detection as well. And indeed, deviator detection under imperfect information of game state is in general undecidable.

However, it is interesting to note that the essence of the deviator detection problem lies in monitoring of player actions and not necessarily game state. Indeed, in repeated games we often have no states at all! We show in this paper that in such as case, the problem is tractable. The main idea is that the problem can then be studied in the setting of coordination games of incomplete information with finitely many states of nature. These are systems with perfect monitoring of state, and uncertainty for a player comes only from unobserved actions of other players. But then, these are games with bounded initial uncertainty that can only reduce as play progresses, and this observation leads us to an algorithmic solution.

While the main result of the paper is the assertion that deviator detection is decidable under imperfect monitoring of actions (with only a finite amount of uncertainty about the state), we see the contribution of the paper as twofold: to highlight a setting in game theory that is of interest to distributed computing; and to illustrate the use of epistemic objectives, that is of interest to games as well as distributed systems. We also suggest that methods from automated synthesis may offer new ways of describing sets of equilibrium solutions and for constructing equilibria, which could be of technical interest for bridging game theory and distributed systems [11].

2 Games

We model distributed systems as infinite games with finitely many states. There is a finite set I of players. We refer to a list $x = (x^i)_{i \in I}$ with one element x^i for every player i as a *profile*. For any such profile, we write x^{-i} to denote the list $(x^j)_{j \in I, j \neq i}$ where the component of Player i is omitted; for element x^i and a list x^{-i} , we denote by (x^i, x^{-i}) the full profile $(x^i)_{i \in I}$.

Game structure. To describe the game dynamics, we fix, for each player $i \in I$, a set A^i of *actions*, a set B^i of *observations*, and a set V^i of *local states* — these are finite sets. We denote by A , B , and V , the set of all profiles of actions, observations, and local states; a profile of local states is also called *global state*. Now, the game form is described by its transition function $\gamma: V \times A \rightarrow V \times B$.

The game is played in stages over infinitely many periods starting from a designated initial state $v_0 \in V$ known to all players. In each period $t \geq 1$, starting in a state v_{t-1} , every player i chooses an action a_t^i . Then the transition $\gamma(v_{t-1}, a_t) = (b_t, v_t)$ determines the observation b_t^i received privately by each player i , and the the global successor state v_t , from which the play proceeds to period $t + 1$.

Thus, a *play* is an infinite sequence $\pi = v_0 a_1 v_1 a_2 v_2 \dots \in V(AV)^\omega$ following the transitions $\gamma(v_{t-1}, a_t) = (v_t, b_t)$ for all $t \geq 1$. A *history* is a finite prefix $v_0 a_1 v_1 \dots a_t v_t \in V(AV)^*$ of a play. We refer to the number of stages played up to period t as the *length* of the history. The sequence of observations $b_1^i b_2^i \dots b_t^i$ received by player i along a history π is denoted by $\beta^i(\pi)$. We assume that each player i always knows his local state v^i and the action a^i she is playing, that is, these data are included in her private observation received in each round. However, she is not perfectly informed about the local states or the actions of the other players, therefore we speak of *imperfect monitoring*. The monitoring function β^i induces an *indistinguishability* relation between histories and plays: $\pi \sim^i \pi'$ if, and only if, $\beta^i(\pi) = \beta^i(\pi')$. This is an equivalence relation between game histories; its classes are called the *information sets* of player i .

A *strategy* for player i is a mapping $s^i: (B^i)^* \rightarrow A^i$ that prescribes an action for every observation sequence. Again, we denote by S the set of all strategy profiles. We say that a history or a play π *follows* a strategy s^i , if $a_{t+1}^i = s^i(\beta^i(\pi_t))$, for all histories π_t of length $t \geq 0$ in π . Likewise, a history or play follows a profile $s \in S$, if it follows the strategy s^i of each player i . The *outcome* $\text{out}(s)$ of a strategy profile s is the unique play that follows it.

With the above definition of a strategy, we implicitly assume that players have perfect recall, that is, they may record all the information acquired along a play. Nevertheless, in certain cases, we can restrict our attention to strategy functions computable by automata with finite memory. In this case, we speak of *finite-state strategies*.

Strategy synthesis. The task of automated synthesis is to construct finite-state strategies for solving games (presented in a finite way). Depending on the purpose of the model, the notion of solving has different meanings.

One prominent application area in distributed systems is concerned with synthesising coordination strategies for a coalition with common interests against a fixed adversary — the environment, or Nature [13, 17, 8]. For this purpose, we assume the coalition to be the set $I \setminus \{0\}$ excluding a designated player 0. We are interested in win/lose games. The *winning condition* is described as a set W of plays; a basic example are *reachability* winning conditions, which consist of all plays that reach a designated set of global states. Here, a solution is a distributed winning strategy for the coalition: a profile s^{-0} such that $\text{out}(s^{-0}, r^0) \in W$ for all $r^0 \in S^0$.

The *distributed synthesis problem* for coordination strategies asks: for a given game, determine whether there exists a profile of finite-state strategies for $I \setminus \{0\}$ that is winning against player 0. This problem is well known to be undecidable, already for games with reachability winning conditions.

Theorem 2.1 ([16]). *The distributed synthesis problem is undecidable for reachability games for two players with imperfect information against an adversary.*

If we consider non-zero sum games, where we have players with different, possibly overlapping objectives, a standard solution concept is that of Nash equilibrium. A theory of synthesis for such games

is being developed in the last decade [10, 5, 4]. While there are several positive results for games of perfect information, synthesis questions are difficult in the context of imperfect information.

In this paper, we are interested in a specific case that lies between the approaches of distributed coordination and Nash equilibrium: synthesis of strategies for detecting an unknown adversary—if he should arise— under conditions of imperfect monitoring. We now proceed to define and address this problem.

3 Deviator detection

Unlike the traditional setting concerned with *temporal* objectives on actions or states assumed in a play, we are interested here in *epistemic* objectives which refer to attaining knowledge that a certain event has occurred. For an introduction to knowledge in distributed systems, we refer to the book of Fagin, Halpern, Moses, and Vardi [7, Ch. 2].

Let us fix a game γ with the usual notation. An *event* is a subset E of histories in γ . The event E *occurs* at history π if $\pi \in E$. We say that (the occurrence of) E is private knowledge of Player i at history π , if for any $\pi' \sim^i \pi$, it holds that $\pi' \in E$. Further, an event $E \subseteq \Omega$ is *common knowledge* among the players of a coalition $C \subseteq I$ at history π if, for every sequence of histories π_1, \dots, π_k and players $i_1, \dots, i_k \in C$ such that $\pi \sim^{i_1} \pi_1 \sim^{i_2} \dots \sim^{i_k} \pi_k$, it is the case that $\pi_k \in E$.

Specifically, we are interested in the event that a player i has deviated from a given play. To describe this, we define for each play π and for every player $i \in I$, the event $D^i(\pi)$ consisting of all histories τ that disagree with π such that, for the first round t where the prefixes π_t and τ_t disagree, they differ only in the action of player i . (Since the transition function is deterministic and the initial state is fixed, the first difference between two histories can only occur at an action profile). Obviously, the sets $D^i(\pi)$ are suffix closed, that is, for each $\tau \in D^i(\pi)$, any prolongation history τ' belongs to $D^i(\pi)$ as well. Likewise, if a coalition attains common knowledge of $D^i(\pi)$ at a history τ , it attains common knowledge of $D^i(\pi)$ at every prolongation history of τ .

Now, let W be a designated set of plays in γ . A *deviator detection* strategy with respect to W is a strategy profile $s \in S$ such that:

- (i) the outcome $\pi = \text{out}(s)$ belongs to W , and
- (ii) for each player i and every strategy $r^i \in S^i$, if the outcome $\pi' = \text{out}(s^{-i}, r^i)$ disagrees with π , then the coalition $I \setminus \{i\}$ attains common knowledge of $D^i(\pi)$ at some history of π' .

The synthesis problem for deviator detection is the following: given a game γ with a target set W specified by a finite-state automaton, decide whether there exists a finite-state strategy profile for deviator detection with respect to W and, if so, construct one.

3.1 Deviator detection as a coordination problem

Alternatively, we can cast the deviator detection problem as a more standard problem of distributed synthesis with temporal objectives. Informally, this is done by adding a new player—Nature—that can either remain silent, or at any point take over the identity of an actual player, deviate from his intended action, and continue playing on his behalf. The deviation of Nature takes the game into a fresh copy associated to the corrupted player i where the only way to win for the remaining coalition is by issuing a simultaneous action in which they all expose the identity of i ; however, if the exposure action is not taken in consensus by all players, except for i , the game is lost.

More precisely, we transform the deviator detection game γ into a coordination game $\hat{\gamma}$ against Nature—let us call it exposure game—as follows: First, we add for each player i , actions $j \neq i$ that allow to expose a deviation by player j , that is, we set $\hat{A}^i := A^i \cup \{I \setminus i\}$; we use the shorthand expose^i to denote any action profile where the coalition $I \setminus \{i\}$ chooses i in consensus. Further, we involve a new player N with actions that allow him to stay silent (by choosing ε) or to corrupt the action of any other player, that is $A^N := \{\varepsilon\} \cup \cup_{i \in I} A^i$. His local states are $V^N := \{\varepsilon\} \cup I$. Moreover, we include (global) sink states \oplus and \ominus . The observation sets remain unchanged. The transitions of the new game $\hat{\gamma}$ follow the original transitions as long as Nature stays silent: $\hat{\gamma}((v, \varepsilon), (a, \varepsilon)) = ((v', \varepsilon), b)$ for $(v', b) := \gamma(v, a)$. When Nature decides to deviate from the intended action of a player i , his local state changes from ε to i : $\hat{\gamma}((v, \varepsilon), (a, c^i)) = ((v', i), b)$ for $c^i \neq a^i$ and $(v', b) := \gamma(v, (a^{-i}, c^i))$; for the transitions in the game copy where player i is corrupted, we set $\hat{\gamma}((v, i), (a, c^i)) = ((v', i), (a, c^i))$ for $(v', b) := \gamma(v, (a^{-i}, c^i))$, and $\hat{\gamma}((v, i), (\text{expose}^i, c^i)) = \oplus$; all other moves involving exposure actions lead to \ominus . The (temporal) winning condition \hat{W} of the new game consists of the plays in W where Nature remained silent and of the plays that reach \oplus .

We can analyse the exposure game in terms of the Knowledge of Preconditions principle formulated by Moses in [15]. Once a deviation occurred, the coalition can win only by reaching \oplus via a simultaneous consensus action which requires common knowledge of the identity of the deviator. Conversely, deviator detection strategies in γ can be readily used to win the exposure game.

Lemma 3.1. *Every coordination strategy for the coalition I in the game $(\hat{\gamma}, \hat{W})$ against N corresponds to a deviator-detection strategy with respect to W in γ , and vice versa.*

3.2 General undecidability

The translation of deviator detection games into coordination games shows the problem under a different angle, but it does not bring us closer to an algorithmic solution. In the general setting of imperfect monitoring, we obtain coordination games between multiple players with imperfect information, for which the synthesis problem is undecidable, as pointed out in Theorem 2.1.

Indeed, it turns out that under imperfect monitoring, detecting deviators is no easier than coordinating against an opponent to reach a target set.

Theorem 3.2. *The synthesis problem for deviator detection strategies is undecidable for games with imperfect monitoring.*

Proof. Consider an arbitrary coordination game γ with three players 0, 1, and 2 where the coalition $\{1, 2\}$ seeks to reach a set F of states under imperfect monitoring. We reduce the synthesis problem for this game to one in a deviator detection game γ' among four players 1, 2, X and Y , in which 1, 2 play the same role as in γ whereas X and Y both take the role of player 0. The new game contains two disjoint copies γ_X, γ_Y of γ ; the actions of Y are ignored in the former, and those of X in the latter. The game γ' starts in a fresh state at which it loops with a fixed action profile; the designated set W consists only of this looping play. The actions of player 1 and 2 at this state are perfectly observable to all players, so any deviation from the loop in W is detected instantly. In contrast, the deviations of player X or Y generate the same (fresh) observation to 1 and 2, and they lead to the initial state of γ_X and γ_Y , respectively. These two component games evolve in the same way with the only difference that, when switching to a target state of F in γ_X , the observation X is sent to all players, whereas the observation Y is sent when reaching F in γ_Y .

Thus, for any deviator detection strategy s , upon deviation of either X or Y from the initial loop, players 1 and 2 must coordinate to reach the target set F to identify the deviator. Hence, (s^1, s^2) yields

a solution of the coordination problem. Conversely, any solution of the coordination problem leads to a state in F where the deviator is revealed, so it provides a deviator detection strategy. Since the synthesis problem for coordination problems is undecidable, according to Theorem 2.1, it follows that the synthesis problem for deviator detection strategies is also undecidable. \square \square

4 Perfect monitoring of states

As we could see in the previous section, the algorithmic intractability of games where the global state can be hidden from the players over an unbounded duration of time is preserved when we move from coordination to the deviator detection problem. However, our setting bears two sources of uncertainty: the global state *and* the played action. In this section we consider the case where the uncertainty comes only from the actions played by the other players. Indeed, this is a generalisation of the setting of infinitely repeated games, which can be seen as games with only one global state.

As an example, consider the following simple variant of a beeping model [6]. There are n nodes in a network represented by an undirected graph. The nodes can communicate synchronously. In every round, a node can either beep or stay silent. A silent node can observe whether at least one of its neighbour beeped. We assume that the network is commonly known, we are interested in distributed protocols under which some temporal condition is ensured, e.g., no more than a quarter of the nodes beep in the same round, and that are additionally deviator proof, in the sense that whenever a node deviates from the protocol, the protocol followed by the remaining nodes allows to reach a consensus on the identity of the deviator. This question can be represented as a deviator-detection problem among n players, each with two actions – beep or stay silent – and two observations, telling whether any neighbour beeped or not in the current round. As the effect of an action profile is the same in any round, the game has only one global state. Still, the synthesis problem shows some complexity. Partly, this is due to the structure of the observation functions encoded by the network graph. For instance, one can observe that no deviator detection-strategy can exist on networks that are not two-connected: Any deviation has to be detected by at least two witnesses, and every node that is not a direct witness needs to be finally informed via at least two disjoint paths. But the greater challenge comes from the dynamics of communicating the identity of the deviator: In contrast to the more traditional synthesis problems for temporal conditions, whether a play π is successful is not determined by the strategic choices taken along π itself, but also depends on the choices taken on histories connected to π via the player's indistinguishability relations.

Concretely, we consider games that allow *perfect monitoring of the state* in the sense that for every observation sequence $\beta^i := b_1^i, \dots, b_t^i$ received by any player i along a history, there exists precisely one global state v that is reachable by a history with observation β^i . In other words, all histories in an information set of a player end at the same global state. The condition is obviously met if we include the current global state in the observation of each player. Indeed, every finite game with perfect monitoring of the state can be transformed effectively into one where all the players can observe the current state.

Our main technical result establishes that, under perfect monitoring of states, the deviator detection problem is algorithmically tractable in spite of imperfect private monitoring of actions.

Theorem 4.1. *The synthesis problem for deviator-detection strategies is effectively solvable for games with perfect monitoring of the state and imperfect monitoring of actions.*

The proof relies on a more general result which states, informally, that the synthesis problem for coordination games with a bounded amount of uncertainty are algorithmically tractable. To formulate this more precisely, let us fix a set I of players with their sets of actions, observations, and local states; the set includes a designated player Nature. Consider a finite collection $\gamma_1, \dots, \gamma_n$ of games with perfect

monitoring of the state over the fixed action, observation and state spaces, together with a winning condition W common to all these games. We define the *sum game* γ over the collection as a game with a fresh initial node at which Nature chooses the initial node of any of the games in the collection; no information about this move is delivered to the other players in $I \setminus \{\text{Nature}\}$. Note that the sum γ may not allow perfect monitoring of the state. For this sum game γ , we consider the task to synthesise a coordination strategy for the coalition $I \setminus \{\text{Nature}\}$ to ensure either that the outcome is either winning with respect to W or it *reveals* the initial choice of Nature. That is, we require, for every play $\pi \notin W$ which follows the strategy, that at some history in π , the players attain common knowledge about the component game that has been chosen. We call this a revelation game over $\gamma_1, \dots, \gamma_n$ with condition W .

In game-theoretic terminology, the sum game constructed above is actually a game of *incomplete* information — the uncertainty about the global state of the game is due to not knowing which of the finitely many component games is being played. Nevertheless, as the component games deliver different observations, the players may be able to recover this missing information. It turns out that this restricted form of imperfect information is algorithmically tractable. The setting is similar to that of multi-environment Markov decision processes studied in [18].

Theorem 4.2. *The synthesis problem is effectively solvable for revelation games on components with perfect monitoring of the state. Moreover, the set of all winning strategies admits a regular representation.*

The idea is to keep track of the knowledge that players have about the index of the actual component game while the play proceeds. This knowledge can be represented by epistemic structures similar to the ones used in [2]. Here it is sufficient to consider epistemic structures on a subset of component indices, with epistemic equivalence relations \sim^i relating indices k, k' whenever player i considers it possible to be in component k if the actual history is in component k' . The tracking construction associates to each history π a structure that is strongly connected via these relations; we call this structure the epistemic state of π .

Intuitively, the construction represents the actions in the original game abstractly by their effect on the uncertainty about the component. In contrast to the concrete actions in the game, which can be monitored only imperfectly, the abstract updates on the epistemic structure can be monitored perfectly; the resulting game is thus solvable with standard methods as one with perfect information, where the winning condition asks to satisfy W or to reach an epistemic structure over a singleton, representing that the players attain common knowledge about the actual component. The perfect-information solution yields a regular representation of the set of winning strategies over a product alphabet of global game states and epistemic states. In the full paper, we show that this abstraction can be maintained by a finite-state construction and that it allows to represent a solution whenever one exists.

To prove Theorem 4.1 using the result of Theorem 4.2, we view the deviator detection problem as a revelation game. Towards this, we adapt the exposure game constructed in Subsection 3.1 for a given deviator detection problem γ with target set W . The exposure game $(\hat{\gamma}, \hat{W})$ is already close to the setting of revelation games, but there is one twist: Besides choosing the deviator, in the exposure game Nature can also choose the period in which to deviate. To account for this, we transform $\hat{\gamma}$ by letting Nature pick a candidate deviator $i \in I$ in the first move; this choice is hidden from the other players. In every later round, Nature can choose to either remain silent or to corrupt the action of player i . As a target condition for the new game, we fix the set of all plays in the target set W where Nature remained silent. The obtained revelation game has the same set of solutions as the deviator detection problem at the outset.

5 Conclusion

Thus what we have here is a building block for constructing equilibria in games based on epistemic states. We are interested primarily in the issue of detecting a deviation from an agreed strategy profile. This task is more specific than constructing equilibria by detecting deviations from the set of distributed strategies that ensure a win, as done for instance, in [3]. The crucial difference relies in the fact that, in the latter case, the deviation events can be described by a (regular) set of game histories, while in our setting, the notion of deviation is relative to a strategy profile that is not fixed within the game structure. To illustrate the difference, consider the example of a beeping model from 4 with the trival target set that contains all possible plays. Obviously, every strategy profile is an equilibrium here, but deviator-detection strategies remain nevertheless intricate.

Our abstraction from imperfect monitoring of actions to games with perfect information is fairly generic. The central clue is that if there is only a finite amount of information hidden in the beginning of a play, and we can decide whether the coalition can recover it. In the context of distributed systems, there is a wide variety of situations that involve only imperfect observation of actions, and where uncertainty about system states may be bounded. Hence we reasonably expect that these techniques will be applicable, not only for deviator detection, but in other algorithmic questions on inferring global information in such systems.

Acknowledgements This work was supported by the Indo-French Joint Research Unit RELAX (UMI CNRS 2000).

References

- [1] Massimiliano Amarante. Recursive structure and equilibria in games with private monitoring. *Economic Theory*, 22(2):353–374, 2003.
- [2] Dietmar Berwanger, Łukasz Kaiser, and Bernd Puchala. A perfect-information construction for coordination in games. In *Proceedings of Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2011)*, volume 13 of *LIPICs*, pages 387–398. Leibniz-Zentrum für Informatik, December 2011.
- [3] Patricia Bouyer, Romain Brenguier, Nicolas Markey, and Michael Ummels. Concurrent games with ordered objectives. In *Foundations of Software Science and Computational Structures - 15th International Conference, FOSSACS 2012. Proceedings*, pages 301–315, 2012.
- [4] Patricia Bouyer, Romain Brenguier, Nicolas Markey, and Michael Ummels. Pure Nash equilibria in concurrent games. *Logical Methods in Computer Science*, 11(2:9), June 2015.
- [5] Krishnendu Chatterjee, Laurent Doyen, Emmanuel Filiot, and Jean-François Raskin. Doomsday equilibria for omega-regular games. In Kenneth L. McMillan and Xavier Rival, editors, *Verification, Model Checking, and Abstract Interpretation*, volume 8318 of *LNCS*, pages 78–97. Springer, 2014.
- [6] Alejandro Cornejo and Fabian Kuhn. Deploying wireless networks with beeps. In Nancy A. Lynch and Alexander A. Shvartsman, editors, *Distributed Computing: 24th International Symposium, DISC 2010. Proceedings*, pages 148–162. Springer, 2010.
- [7] Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi. *Reasoning about Knowledge*. MIT Press, 1995.
- [8] B. Finkbeiner and S. Schewe. Uniform distributed synthesis. In *Proc. of Logic in Computer Science (LICS'05)*, pages 321–330. IEEE, 2005.
- [9] Drew Fudenberg, David I Levine, and Eric Maskin. The Folk Theorem with Imperfect Public Information. *Econometrica*, 62(5):997–1039, September 1994.

- [10] Julian Gutierrez and Michael Wooldridge. Equilibria of concurrent games on event structures. In *Computer Science Logic (CSL) and Logic in Computer Science (LICS)*, CSL-LICS '14, pages 46:1–46:10, New York, NY, USA, 2014. ACM.
- [11] Joseph Y. Halpern. Computer science and game theory: A brief survey. *CoRR*, abs/cs/0703148, 2007.
- [12] Michihiro Kandori and Hitoshi Matsushima. Private observation, communication and collusion. *Econometrica*, 66(3):pp. 627–652, 1998.
- [13] Orna Kupferman and Moshe Y. Vardi. Synthesizing distributed systems. In *Proc. of LICS '01*, pages 389–398. IEEE Computer Society Press, June 2001.
- [14] George Mailath and Larry Samuelson. *Repeated Games and Reputations: Long-Run Relationships*. Oxford University Press, 2006.
- [15] Yoram Moses. Relating knowledge and coordinated action: The knowledge of preconditions principle. In *Theoretical Aspects of Rationality and Knowledge, TARK 2015, Proc.*, volume 215 of *EPTCS*, pages 231–245, 2015.
- [16] Gary L. Peterson and John H. Reif. Multiple-Person Alternation. In *Proc 20th Annual Symposium on Foundations of Computer Science, (FOCS 1979)*, pages 348–363. IEEE, 1979.
- [17] Amir Pnueli and Roni Rosner. Distributed reactive systems are hard to synthesize. In *Proceedings of the 31st Annual Symposium on Foundations of Computer Science, (FoCS '90)*, pages 746–757, 1990.
- [18] Jean-Francois Raskin and Ocan Sankur. Multiple-Environment Markov Decision Processes. In *Foundation of Software Technology and Theoretical Computer Science (FSTTCS 2014)*, volume 29 of *LIPICs*, pages 531–543. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2014.