

Uniform strategies

Bastien Maubert and Sophie Pinchinat
S4, IRISA, Campus de Beaulieu
Rennes, 35042, France

Abstract

We consider turn-based game arenas for which we explore a notion of *uniform strategies*, which are strategies submitted to *uniformity constraints*. Such constraints aim at capturing properties of strategies that are not μ -calculus definable as they involve bundles of plays. Typical examples are knowledge-based strategies in imperfect-information games, and the so-called “uniform strategies” for the game semantics of Dependence Logic. We propose a formal language to specify uniform strategies and demonstrate the relevance of the concept by capturing various issues arising from the literature.

1 Introduction

Properties of strategies in, say, extensive (finite or infinite) games are central objects to describe, e.g. when a given strategy is winning. The standard approach starts from a winning condition on plays and winning strategies are those whose induced plays are all winning. By seeing strategies as trees, being winning is therefore measured “vertically” in the tree. We investigate properties that are rather “horizontal”, as they correlate different branches of the tree. This kind of properties on strategies are not μ -calculus definable in general, and yet, they fulfill a true need in game theory, as we illustrate in the present paper.

We develop a mathematical setting which provides a way to define bundles of plays and to describe properties of individuals in the bundle. More concretely, the bundles we consider are equivalence classes of finite and infinite plays, but our setting might also be extended to deal with more general binary relations.

This setting involves a language with an epistemic temporal logic flavor (as ETL[11]), but the bundles of plays need not arise from the epistemic accessibility relation of any player of the game. Additionally, the language possesses an original \exists operator, which allows to quantify over equivalence classes of infinite plays, and the usefulness of which is illustrated.

The need for describing properties of higher-order has already been investigated for example in [6] with the notion of *hyperproperties*, which are properties of sets of system traces. But, as opposed to what we expose here, hyperproperties deal with fixed sets of traces, and moreover, to our knowledge, no logic has ever been developed.

We have chosen to tell our story in a simple framework where games are described by arenas in which all information is put inside the positions, and not on the edges. However, the entire theory can be adapted to more sophisticated models, e.g. with labels on edges, concurrent games, ...

As announced earlier, we illustrate the suitability of our notion by borrowing many frameworks from the literature. These are imperfect-information games with their observation-based and knowledge-based strategies, games with opacity conditions, the non-interference properties of computing systems,

diagnosability of discrete-event systems (with a proposal for a formal definition of *prognosis*), and finally the imperfect-information semantics of Dependence Logic. Proofs of Section 3 are omitted due to lack of space, but they are quite simple. There are even more instances of uniform strategies in the literature, but the numerous examples we give here are already convincing enough to justify the relevance of the notion.

The paper is organized as follows: we start by introducing the game models in Section 2, and the formal notion of uniform strategies. Section 3 is dedicated to the illustrations. In Section 4 we address various expressiveness results. In Section 5 we expose some computational aspects, and we discuss future work in Section 6.

2 Definitions

Let \mathcal{G} be a multi-player turn-based game arena given as a structure $(V_1, \dots, V_N, E, v_0, Prop, \ell)$ where V_k ($k = 1, \dots, N$) is the set of positions of Player k , $V := \bigcup_{k=1, \dots, N} V_k$ is the set of all positions, $E \subseteq V \times V$, and $v_0 \in V$ is the initial position. For $v, v' \in V$, $vE v'$ denotes $(v, v') \in E$. $Prop$ is a countable set of atomic propositions, and we decorate the standard game structure with a valuation $\ell : V \rightarrow 2^{Prop}$. We define $Plays_\omega \subseteq v_0 V^\omega$ to be the set of *infinite plays*, *i.e.* the set of infinite sequences of positions π that start in v_0 and follow the edges of the arena. Similarly, $Plays_* \subseteq v_0 V^*$ is the set of finite plays ρ . For an infinite play $\pi = v_0 v_1 \dots$ and $i \in \mathbb{N}$, $\pi[i] := v_i$, and $\pi[0, i] := \pi[0] \dots \pi[i]$. If $\rho = v_0 v_1 \dots v_n$ is a finite play, $last(\rho) := v_n$ is its last position and $|\rho| = n + 1$ is its length. A strategy for Player k is a partial function $\sigma : Plays_* \rightarrow V$ that defines which position to choose in each finite play ρ in which it is Player k 's turn to play (*i.e.* $last(\rho) \in V_k$). $Outcome(\sigma) \subseteq Plays_\omega$ is the set of infinite plays in which Player k follows σ .

The specification language we use to define uniformity constraints is quite close, both in syntax and semantics, to linear temporal logic with knowledge [11], except the semantics of both our “knowledge-like” operators are based on equivalence relations on plays that do not necessarily derive from the observational power of a player.

The syntax of the language \mathcal{L} is the following :

$$\mathcal{L} : \quad \varphi, \psi ::= p \mid \neg\varphi \mid \varphi \wedge \psi \mid \odot\varphi \mid \varphi \mathcal{U} \psi \mid \mathbb{K}\varphi \mid \mathbb{M}\varphi$$

where p is in $Prop$. As usual we will use the following notations : $true := p \vee \neg p$, $false := \neg true$, $\mathbf{F}\varphi := true \mathcal{U} \varphi$, $\mathbf{G}\varphi := \neg \mathbf{F}\neg\varphi$, and $\varphi \mathcal{W} \psi := \varphi \mathcal{U} \psi \vee \mathbf{G}\varphi$. In addition we will use the two following notations : $\langle \mathbb{K} \rangle \varphi := \neg \mathbb{K} \neg \varphi$ and $\langle \mathbb{M} \rangle \varphi := \neg \mathbb{M} \neg \varphi$.

We use \mathbb{K} instead of the usual knowledge operator K to emphasize that though it has a strong epistemic flavour, notably in various application instances we present here, it need not be interpreted in terms of knowledge in general, but merely as a way to state properties of bundles of plays.

The interpretation of our language requires an arena \mathcal{G} , enriched with an equivalence relation over the set of finite plays $\simeq \subseteq Plays_*^2$ and an equivalence relation over the set of infinite plays $\simeq \subseteq Plays_\omega^2$.

It seems natural to relate \simeq and \simeq , which can be done in several standard ways (we refer to [14]). For example, \simeq can be obtained as the *limit extension* \simeq^{lim} of \simeq : $\pi \simeq^{lim} \pi'$ if there exist two increasing sequences $(\rho_i)_{i \in \mathbb{N}}$ and $(\rho'_i)_{i \in \mathbb{N}}$ of prefixes of π and π' respectively such that for all $i \in \mathbb{N}$, $\rho_i \simeq \rho'_i$. Alternatively, if \simeq can be defined on *suffixes* of finite plays λ , *i.e.* finite plays starting in a position $v \neq v_0$, (which is the case in all our examples), \simeq can also be obtained as the *piecewise*

extension \smile^ω of \smile : $\pi \smile^\omega \pi'$ if there exist two decompositions $\pi = \lambda_1 \lambda_2 \lambda_3 \dots$ and $\pi' = \lambda'_1 \lambda'_2 \lambda'_3 \dots$ such that $\lambda_i \smile \lambda'_i$ for all $i = 1, 2, 3, \dots$. Whether the two definitions coincide depends on the properties of \smile . We can state the following lemmata (which are easy to prove).

Lemma 1 *If \smile is a congruence, that is $\lambda_1 \smile \lambda'_1$ and $\lambda_2 \smile \lambda'_2$ implies $\lambda_1 \lambda_2 \smile \lambda'_1 \lambda'_2$, then $\smile^\omega \subseteq \smile^{lim}$.*

Lemma 2 *If \smile is left cancelable, that is $\lambda_1 \smile \lambda'_1$ and $\lambda_1 \lambda_2 \smile \lambda'_1 \lambda'_2$ implies $\lambda_2 \smile \lambda'_2$, then $\smile^\omega \supseteq \smile^{lim}$.*

Notice that epistemic equivalences (accessibility relations in S5) which satisfy the “Uniform No Miracle” assumption of [19] yield congruences over finite histories, thus there may be a canonical way of deriving their extension to infinite histories, if they are also left cancelable. Actually, this is the case for the relevant examples we consider in Section 3. Nevertheless, we want to adopt a general setting where the relation between \smile and \frown need not be given explicitly.

Given an arena \mathcal{G} , and two equivalence relations \smile and \frown on finite and infinite plays respectively, we interpret our specification language \mathcal{L} on triples (Π, π, i) , where $\Pi \subseteq Plays_\omega$, $\pi \in \Pi$ and $i \in \mathbb{N}$. The semantics is given by induction over the structure of the formulas.

$\Pi, \pi, i \models p$	iff	$p \in \ell(\pi[i])$
$\Pi, \pi, i \models \neg\varphi$	iff	$\Pi, \pi, i \not\models \varphi$
$\Pi, \pi, i \models \varphi \wedge \psi$	iff	$\Pi, \pi, i \models \varphi$ and $\Pi, \pi, i \models \psi$
$\Pi, \pi, i \models \odot\varphi$	iff	$\Pi, \pi, i + 1 \models \varphi$
$\Pi, \pi, i \models \varphi \mathcal{U} \psi$	iff	there is $j \geq i$ such that $\Pi, \pi, j \models \psi$ and for all $i \leq k < j$, $\Pi, \pi, k \models \varphi$
$\Pi, \pi, i \models \mathbb{K}\varphi$	iff	for all $\pi' \in \Pi, j \in \mathbb{N}$ such that $\pi[0, i] \smile \pi'[0, j]$, $\Pi, \pi', j \models \varphi$
$\Pi, \pi, i \models \mathbb{X}\varphi$	iff	for all $\pi' \in \Pi, j \in \mathbb{N}$ such that $\pi \frown \pi'$ and $\pi[0, i] \smile \pi'[0, j]$, $\Pi, \pi', j \models \varphi$

The LTL part is classic. $\mathbb{K}\varphi$ is true at some point of a play if φ is true in every equivalent finite play, which is also the classic semantics of epistemic temporal logics. The \mathbb{X} operator is more original: alike \mathbb{K} we consider finite plays that are equivalent to the current one, but the set of finite plays we consider is restricted to prefixes of infinite plays equivalent to the current infinite play. It enables to filter \smile -equivalent plays on the basis of properties of the whole infinite current play. For example, if \smile (and \frown) defines a player’s observational power, \mathbb{X} expresses what her present knowledge would be if she were able to look arbitrarily far away in the future of the current play.

Definition 1 *A uniformity constraint is a formula $\varphi \in \mathcal{L}$. We sometimes write such a constraint $\mu = (\smile, \frown, \varphi)$ in order to make explicit the two equivalences \smile and \frown used in the semantics.*

Now we define two notions of uniform strategies, which differ only in the universe the \mathbb{K} and the \mathbb{X} operators quantify over: $Outcome(\sigma)$ or $Plays_\omega$ (with the latter, equivalent plays not induced by the strategy also count). As we shall see in the examples of the next section, making a nuance is worthwhile.

Definition 2 *Let \mathcal{G} be an arena and $\mu = (\smile, \frown, \varphi)$ be a uniformity constraint. A strategy σ for a Player k is*

- μ -strictly uniform if for all $\pi \in Outcome(\sigma)$, $Outcome(\sigma), \pi, 0 \models \varphi$.
- μ -fully uniform if for all $\pi \in Outcome(\sigma)$, $Plays_\omega, \pi, 0 \models \varphi$.

Remark that the parameter \frown (resp. \smile) plays no role in Definition 2 if φ does not contain any \mathbb{X} (resp. \mathbb{K}) operator, hence it is a mere LTL formula. Notice that in this latter case, some standard ω -regular (winning) conditions can be expressed over plays. The extension to a more powerful logic, such as the full propositional μ -calculus, in order to capture all ω -regular properties is *a priori* possible. However, for the examples considered in Section 3 this full power is not needed.

3 Frameworks from the literature

They are several instances of frameworks where the notion of uniform strategies occurs. This is what this section aims at showing.

In the following, we will talk about *synchronicity* when the equivalence \smile preserves the length, and *asynchronicity* otherwise, but in general we make no assumption on \smile .

3.1 Observation-based and knowledge-based strategies

Games with imperfect information, in the most general sense, are games in which some of the players do not know exactly what is the current position of the game. This can occur in real games, e.g. poker since one does not know what cards her opponents have in hands, but also in situations arising from computer science, like for example a program that observes or controls a system by means of a sub-part of its variables, the interface, while other variables remain hidden. In games with imperfect information, the player's ability to remember what happened so far along a play is a key point to achieve a winning strategy. This is not the case, e.g. for perfect-information parity games, where memoryless strategies are sufficient. It is therefore relevant under an imperfect information assumption to distinguish the *perfect recall* setting, as opposed to the *imperfect recall* one, when the player remembers the whole history of the observation she had of a play, no matter how long it is.

While games with imperfect information and perfect recall have been studied intensively [16, 5], the case of imperfect recall has received much less attention since paradoxes concerning the interpretation of such games were raised [15]. Nonetheless, relevant problems may be modeled with imperfect recall: typically, particular computing resources have very limited memory and cannot remember arbitrarily long histories.

In this presentation, we restrict to the classic synchronous perfect recall setting; however the result would easily adapt to different settings, e.g. if the observational equivalence relation were asynchronous, with imperfect recall...

In two-player imperfect-information games as studied for example in [16, 5, 3], Player 1 only partially observes the positions of the game, such that some positions are indistinguishable to her, while Player 2 has perfect information (the asymmetry is due to the focus being on the existence of strategies for Player 1). Arenas are labelled directed graphs, and in each round, if the position is a node v , Player 1 chooses an label/action a , and Player 2 chooses a next position v' reachable from v through an a -labelled edge.

We equivalently define this framework in a manner that fits our setting by putting Player 1's actions inside the positions. We have two kinds of positions, of the form v and of the form (v, a) . In a position v , when she chooses an action a , Player 1 actually moves to position (v, a) , then Player 2 moves from (v, a) to some v' . So an imperfect-information game arena is a structure $\mathcal{G}_{imp} = (\mathcal{G}, \sim)$ where $\mathcal{G} = (V_1, V_2, E, v_0, Prop, \ell)$ is a two-player game arena with positions in V_1 of the form v and positions in V_2 of the form (v, a) . For a position $(v, a) \in V_2$, $Pos(v, a) = v$ and $Act(v, a) = a$. $E \subseteq V_1 \times V_2 \cup V_2 \times V_1$, $vE(v', a)$ implies $v = v'$, $v_0 \in V_1$, $Prop = \{p_1\} \cup \{p_a \mid \exists v, (v, a) \in V_2\}$ and

$\ell(v) = \{p_1\}$ for $v \in V_1$, $\ell(v, a) = \{p_a\}$ for $(v, a) \in V_2$. Finally, $\sim \subseteq V_1^2$ is an observational equivalence relation on positions, that relates indistinguishable positions for Player 1, and it extends to finite plays as the least relation \sim such that $\rho v \sim \rho' v'$ whenever $\rho \sim \rho'$ and $v \sim v'$, and $\rho(v, a) \sim \rho'(v', a')$ whenever $\rho \sim \rho'$, $v \sim v'$ and $a = a'$.

We add the classic requirement that the same actions must be available in indistinguishable positions: for all $v, v' \in V_1$, if $v \sim v'$ then $vE(v, a)$ if, and only if, $v'E(v', a)$. In other words, a player can distinguish positions where she has different options.

A strategy for Player 1 is a mapping $\sigma : v_0(V_2V_1)^* \rightarrow V_2$ such that for all finite play of the form ρv , $vE\sigma(\rho v)$. Those strategies are required to be *observation based*, i.e. Player 1 must play the same way in \sim -equivalent finite plays. Strategies for Player 1 can also be *knowledge based*, which is more restrictive. The *knowledge* or *information set* of Player 1 after a finite play is the set of positions that she considers possible according to the observation she has. Formally, let $\rho \in Plays_*$ be a finite play with $last(\rho) \in V_1$. The *knowledge* or *information set* of Player 1 after ρ is $I(\rho) := \{last(\rho') \mid \rho' \in Plays_*, \rho \sim \rho'\}$.

Definition 3 A strategy σ for Player 1 is observation-based (resp. knowledge-based) if for all $\rho, \rho' \in v(V_2V_1)^*$, $\rho \sim \rho'$ (resp. $I(\rho) = I(\rho')$) implies $Act(\sigma(\rho)) = Act(\sigma(\rho'))$.

We can characterize observation-based strategies and knowledge-based strategies as uniform strategies. In order to capture the knowledge-based property, we define \simeq as the smallest reflexive relation such that for all $\rho, \rho' \in Plays_*$ with $last(\rho), last(\rho') \in V_1$, $I(\rho) = I(\rho')$ implies $\rho \simeq \rho'$. We define the formula

$$\text{SameAct} := \mathbf{G}(p_1 \rightarrow \bigvee_{p_a \in Prop} \mathbb{K} \odot p_a)$$

which expresses that whenever it is Player 1's turn to play, there is an action a that is played in all equivalent finite play.

Theorem 3

A strategy σ for Player 1 is observation-based if, and only if, it is $(\sim, \sim^\omega, \text{SameAct})$ -strictly uniform. A strategy σ for Player 1 is knowledge-based if, and only if, it is $(\simeq, \simeq^\omega, \text{SameAct})$ -strictly uniform.

Actually, here the relation on infinite plays is irrelevant, since operator \mathbb{X} is not used; we have arbitrarily chosen to use \sim^ω and \simeq^ω .

3.2 Games with epistemic condition

Uniform strategies enable to express winning conditions that have epistemic features, the relevance of which is exemplified by the *games with opacity condition* studied in [13]. In that case, \mathbb{K} can represent a players' knowledge, or distributed knowledge between a group of players, or common knowledge, depending on the winning condition one wants to define.

Games with opacity conditions are based on two-player imperfect-information arenas with a particular winning condition, called the *opacity condition*, which relies on the knowledge of the player with imperfect information. In such games, some positions are “secret” as they reveal a critical information that the imperfect-information player wants to know (in the epistemic sense).

More formally, let $\mathcal{G}_{inf} = (\mathcal{G}, \sim)$ be an imperfect-information arena as described in Section 3.1, with a distinguished set of positions $S \subseteq V_1$ that denotes the secret. Let $\mathcal{G} = (V_1, V_2, E, v_0, \{p_S\}, \ell)$ be the arena with $\ell^{-1}(\{p_S\}) = S$ (positions labeled by p_S are exactly positions $v \in S$). The *opacity*

winning condition is as follows. An infinite play is winning for Player 1 if there exists a finite prefix ρ of this play whose information set is contained in S , *i.e.* $I(\rho) \subseteq S$, otherwise Player 2 wins. Again, strategies for Player 1 are required to be observation-based. It can easily be shown that:

Theorem 4 *A strategy σ for Player 1 is winning if, and only if, σ is $(\sim, \sim^\omega, \mathbf{FKp}_S)$ -strictly uniform. A strategy σ for Player 2 is winning if, and only if, σ is $(\sim, \sim^\omega, \mathbf{G-Kp}_S)$ -fully uniform.*

In Theorem 4, we make use of the “full” uniformity property because we are interested in the point of view of Player 1 who might consider possible some plays that are not induced by the strategy of Player 2.

3.3 Non-interference

Non-interference, as introduced by [10], is a property evaluated on labelled transition systems which handle Boolean variables. Such systems are tuples $(S, \mathcal{I}, \mathcal{O}, \delta, s_0, \text{out})$ where S is the set of states, $\mathcal{I} = H \uplus L$ is a set of Boolean input variables partitioned into *high security variables* H and *low security variables* L , \mathcal{O} is the set of Boolean output variables, $\delta : S \times 2^{\mathcal{I}} \rightarrow S$ is the transition function that maps each pair of state and input variables valuation¹ to a next state, s_0 is the initial state, and $\text{out} : S \rightarrow 2^{\mathcal{O}}$ is the output function that represents a mapping of states onto valuations of the Boolean output variables. We extend the transition function δ to $S \times (2^{\mathcal{I}})^* \rightarrow S$ as expected: $\delta(s, \epsilon) = s$ and $\delta(s, ua) = \delta(\delta(s, u), a)$.

We consider the definition of non-interference on infinite input sequences used in [7]. We define the L -equivalence, \sim_L over $(2^{\mathcal{I}})^*$ by $u \sim_L u'$ whenever u and u' have the same length and they coincide on the values of the input variables, *i.e.* for all $1 \leq i \leq \text{length}(u)$, for all $l \in L$, $l \in u(i) \Leftrightarrow l \in u'(i)$. Notice that since \sim_L is a congruence and is left cancelable, by Lemmata 1 and 2, \sim_L^{lim} and \sim_L^ω coincide, and we shall write this relation \approx_L . Given an infinite sequence of inputs $w \in (2^{\mathcal{I}})^\omega$, we abuse notation by writing $\text{out}(w)$ for the infinite sequence of output variables valuations encountered in the states along the execution of the system on input w . A system $(S, \mathcal{I}, \mathcal{O}, \delta, s_0, \text{out})$ verifies the *non-interference property* if for any two infinite sequences of inputs $w, w' \in (2^{\mathcal{I}})^\omega$, $w \approx_L w'$ implies $\text{out}(w) = \text{out}(w')$. In other words, the valuations of high security variables have no consequence on the observation of the system.

A first natural problem is to decide the non-interference property of a system. A second more general problem is a control problem: we want to decide whether there is a way of restricting the set of input valuations along the executions, or equivalently to control the environment, so that the system is non-interfering. By constraining the applied restriction to be trivial, the former problem is a particular case of the latter. We can encode the control problem in our setting.

Let $Sys = (S, \mathcal{I}, \mathcal{O}, \delta, s_0, \text{out})$ be an instance of the problem, and write Σ for $2^{\mathcal{I}}$ with typical elements a, b, \dots . Without loss of generality, we can assume that Sys is *complete*: every input valuation yields a transition. We define a two-player game arena that simulates the system, in which Player 1 fixes the *environment*, *i.e.* a subset of the possible inputs in the current state, and Player 2 chooses a particular input among those. More formally, let $\mathcal{G}_{Sys} = (V_1, V_2, E, v_0, Prop, \ell)$, with $V_1 = (\Sigma \uplus \{\epsilon\}) \times S$ and $V_2 = S \times 2^\Sigma$. A position $(a, s) \in V_1$ denotes a situation where the system reaches state s by an a -transition, and $(s, A) \in V_2$ denotes a situation where in state s , the set of possible inputs is A . The set of edges E of the arena is the smallest set such that $(a, s)E(s, A)$ for all $s \in S$, $a \in \Sigma$ and $A \subseteq \Sigma$,

¹we classically confuse valuations over a set B of Boolean variables with elements of 2^B .

and $(s, A)E(a, \delta(s, a))$ whenever $s \in S$ and $a \in A$. The initial position of the arena is $v_0 = (\epsilon, s_0)$, and by letting $Prop = \{p_o \mid o \in 2^{\mathcal{O}}\}$, we set $\ell(a, s) = \ell(s, A) = \{\text{out}(s)\}$.

By writing ι for the canonical projection from $V_1 \cup V_2$ onto $2^{\mathcal{I}}$ (that is $\iota(\epsilon, s_0) = \iota(s, A) = \epsilon$ and $\iota(a, s) = a$) and by extending ι to finite plays as expected, we let $\rho \equiv_L \rho'$ hold whenever $\iota(\rho) \sim_L \iota(\rho')$. Again, \equiv_L is both a congruence and left cancelable, hence \equiv_L^ω and \equiv_L^{lim} coincide. We now define the formula

$$\text{SameOutput} := \mathbf{G} \bigwedge_{p_o \in Prop} (p_o \rightarrow \mathbb{X}p_o)$$

which captures the property that the valuations of output variables along \equiv_L^ω -equivalent executions of the system coincide, and we can establish the following.

Theorem 5 *There is a one-to-one correspondence between $(\equiv_L, \equiv_L^\omega, \text{SameOutput})$ -strictly uniform strategies of Player 1 and the controllers which ensure the non-interference property of the system.*

In particular, the trivial strategy of Player 1, where from any position (a, s) she chooses to move to (s, Σ) , is $(\equiv_L, \equiv_L^\omega, \text{SameOutput})$ -strictly uniform if, and only if, the system has the non-interference property.

Notice that in order to make this control problem more realistic, one would seek a maximal permissive strategy/controller so that environments as “large” as possible are computed, but this is out of the scope of the paper.

3.4 Diagnosis and Prognosis

Diagnosis has been intensively studied, in particular by the discrete-event systems community (see for example [17, 21, 4]). Informally, in this setting, a discrete-event system is *diagnosable* if any occurrence of a faulty event during an execution is eventually detected. More formally, *diagnosability* is a property of discrete-event systems which are structures of the form $Sys = (S, \Sigma, \Sigma_o, \Delta, s_0, F)$, with S the set of states, Σ the set of events, $\Sigma_o \subseteq \Sigma$ the *observable* events, $\Delta \subseteq S \times \Sigma \times S$ the transition relation, s_0 the initial state and $F \subseteq S$ the faulty states; we assume that once a faulty state is reached, only faulty states can be reached (the fault is persistent). We can rephrase this problem in our setting, with a single player simulating the system. Notice that since there is only one player, a strategy defines a unique infinite play.

Let $\mathcal{G}_{Sys} = (V, E, v_0, Prop, \ell)$, with $V = (\Sigma \uplus \{\epsilon\}) \times S$, $(a, s)E(b, s')$ whenever $(s, b, s') \in \Delta$, $v_0 = (\epsilon, s_0)$, $Prop = \{f\}$ and $\ell(a, s) = \{f\}$ if $s \in F$, \emptyset otherwise. We write $\rho \equiv_{\Sigma_o} \rho'$ whenever the sequences of observable events underlying ρ and ρ' are the same (these sequences are obtained from the sequences of positions in the play: for each position of the form (a, s) , keep its letter a and delete it if $a \notin \Sigma_o$).

Theorem 6 *Sys is diagnosable if, and only if, every strategy in \mathcal{G}_{Sys} is $(\equiv_{\Sigma_o}, \equiv_{\Sigma_o}^\omega, \mathbf{F}f \rightarrow \mathbf{FK}f)$ -fully uniform.*

Prognosis is a companion of diagnosis, but focuses on the ability to predict that a fault will occur. Prognosability-like properties can be defined in our setting. As an example, we aim at saying that a system is *prognosable* whenever the fact that a fault occurs in a system is known at least one step in advance. By using our framework, we can propose the following definition.

Definition 4 *A system Sys is prognosable if every strategy in \mathcal{G}_{Sys} is $(\equiv_{\Sigma_o}, \equiv_{\Sigma_o}^\omega, (\neg f)\mathcal{W}(\neg f \wedge \mathbf{K}\odot f))$ -fully uniform.*

3.5 Dependence Logic

Dependence Logic is a flourishing topic introduced recently by Väänänen [18]. It extends first order logic by adding atomic dependence formulas $\text{dep}(t_1, \dots, t_n)$, which express functional dependence of the term t_n on the terms t_1, \dots, t_{n-1} . Evaluating a dependence between terms on a single assignment of the free variables is meaningless: in order to tell whether t depends on t' , one must vary the values of t' and see how the values of t are affected. This is why a formula of Dependence Logic is evaluated on a first-order model \mathcal{M} and a *set of assignments* for the free variables, called a *team*. If t is a term, \mathcal{M} a model and s an assignment for the free variables in t , $\llbracket t \rrbracket_s^{\mathcal{M}} \in M$ is the interpretation of t in the model \mathcal{M} with the assignment s .

For this logic, a game semantics is given in [18] that is said to be a game with imperfect information. However the game arena is of perfect information: it is called game with imperfect information because an additional “uniformity requirement” is imposed on strategies.

Let ϕ be a sentence (formula with no free variable) of Dependence Logic, and let \mathcal{M} be a first order model. $\mathcal{G}^{\mathcal{M}}(\phi)$ is a two player game between Player 1 and Player 2; positions are of the form (φ, n, s, i) , where φ is a subformula of ϕ , n is the position in ϕ of the first symbol of φ , s is an assignment whose domain contains the free variables of φ , and $i \in \{1, 2\}$. The index n is used to decide, given two positions containing the same dependence atom, whether they are the same *syntactic* subformulas of ϕ or not. For a subformula φ , $\text{len}(\varphi)$ is the number of symbols in φ . The game starts in position $(\phi, 1, \emptyset, 2)$ and the rules are as follows:

position $(t_1 = t_2, n, s, i)$:	if $\llbracket t_1 \rrbracket_s^{\mathcal{M}} = \llbracket t_2 \rrbracket_s^{\mathcal{M}}$, Player i wins, otherwise the opponent wins.
position $(Rt_1 \dots t_m, n, s, i)$:	if $\llbracket R \rrbracket^{\mathcal{M}} \llbracket t_1 \rrbracket_s^{\mathcal{M}} \dots \llbracket t_m \rrbracket_s^{\mathcal{M}}$, Player i wins, otherwise the opponent wins.
position $(\text{dep}(t_1, \dots, t_m), n, s, i)$:	Player i wins.
position $(\neg\varphi, n, s, i)$:	move to the position $(\varphi, n + 1, s, i^*)$, where i^* is the opponent of i .
position $(\varphi \vee \psi, n, s, i)$:	i chooses between position $(\varphi, n + 1, s, i)$ and $(\psi, n + 1 + \text{len}(\varphi), s, i)$.
position $(\exists x\varphi, n, s, i)$:	i chooses a value $a \in M$ and moves to $(\varphi, n + 2, s(a/x), i)$

A strategy σ for Player 1 is *uniform* in the sense of [18] if, for every two finite plays $\rho, \rho' \in \text{Outcome}(\sigma)$ such that $\text{last}(\rho) = (\text{dep}(t_1, \dots, t_m), n, s, 1)$ and $\text{last}(\rho') = (\text{dep}(t_1, \dots, t_m), n, s', 1)$ contain the same (syntactically speaking) atomic dependence subformula, if s and s' agree on t_1, \dots, t_{m-1} , they also agree on t_m . This reflects the idea of dependence atoms: the values of the terms t_1, \dots, t_{m-1} determine the value of t_m . A sentence ϕ of Dependence Logic is true in a model \mathcal{M} if Player 1 has a winning uniform strategy in $\mathcal{G}^{\mathcal{M}}(\phi)$.

We characterize uniform strategies in the sense of [18] as uniform strategies in our sense. The game described above easily fits in our setting (we add loops on terminal positions so as to obtain infinite plays). Let ϕ be a sentence of Dependence Logic, and \mathcal{M} be a finite model. For each object $a \in M$ of the domain we use one atomic proposition p_a , and we also use the proposition d to mark positions that contain dependence atoms. So $\text{Prop} = \{p_a \mid a \in M\} \uplus \{d\}$, and the valuation ℓ is as follows :

$$\ell(\text{dep}(t_1, \dots, t_m), n, s, i) = \begin{cases} \{p_a, d\} & \text{if } i = 1, \text{ with } a = \llbracket t_m \rrbracket_s^{\mathcal{M}} \\ \emptyset & \text{otherwise} \end{cases}$$

$$\ell(-, n, s, i) = \emptyset$$

We define the equivalence relation \simeq on finite plays as the smallest reflexive relation such that if there is $\varphi = \text{dep}(t_1, \dots, t_m)$ and n s.t $\text{last}(\rho) = (\varphi, n, s, 1)$, $\text{last}(\rho') = (\varphi, n, s', 1)$, and $\llbracket t_i \rrbracket_s^{\mathcal{M}} = \llbracket t_i \rrbracket_{s'}^{\mathcal{M}}$

for $i = 1, \dots, n - 1$, then $\rho \simeq \rho'$. Now we define the formula

$$\text{AgreeOnLast} := \mathbf{G}(d \rightarrow \bigvee_{a \in M} \mathbb{K}p_a)$$

which expresses that whenever the current position contains a dependence atom $\text{dep}(t_1, \dots, t_m)$, it agrees with all equivalent finite plays on some value a for t_m . Since equivalent plays are those ending in a position that has the same dependence atom and agrees on the first $m - 1$ terms, it is easy to prove:

Theorem 7 *A strategy σ for Player 1 is uniform if, and only if, it is $(\simeq, \simeq^\omega, \text{AgreeOnLast})$ -strictly uniform.*

3.6 Dependence logic and games with imperfect information

As we said, the semantics game for Dependence Logic presented in the previous section is said to be a game with imperfect information. We do not agree, because the difference between games with perfect information and games with imperfect information (at least in the perfect recall setting, it is not as clear otherwise, see [15]) lies in the fact that in the latter, some finite plays are related (indistinguishable), and players must behave the same way in these related situations. Concerning the semantics game for Dependence Logic, the difference with perfect-information games is that some plays are related, those ending in positions bound to the same atomic dependence formula $\text{dep}(t_1, \dots, t_n)$ with valuations agreeing on t_1, \dots, t_{n-1} , and the valuations in these positions must agree on t_n . So it is not that players should behave the same way in indistinguishable situations, but rather that the players should *have behaved* in such a way that the valuations for t_n are the same in indistinguishable situations.

But it is true that there is a similarity between these two constraints on allowed strategies, as shown by looking at the formulas of the uniformity constraints capturing observation-based strategies (**SameAct**) and uniform strategies in the sense of Dependence Logic (**AgreeOnLast**):

$$\text{SameAct} = \mathbf{G}(p_1 \rightarrow \bigvee_{p_a \in Prop} \mathbb{K} \odot p_a) \quad \text{and} \quad \text{AgreeOnLast} = \mathbf{G}(d \rightarrow \bigvee_{a \in M} \mathbb{K}p_a)$$

In the first case, the same thing must *happen* in equivalent situations, whereas in the second case, the same thing must *hold* in equivalent situations.

Neither semantics games for Dependence Logic are games with imperfect information in the classical sense, nor games with imperfect information can be easily described using the uniform strategy notion of [18], but both can be characterized in a very similar way with our notion of uniform strategies.

4 Expressiveness


By looking at its semantics, the language \mathcal{L} clearly has at least the expressive power of ETL [11] for a single agent with relation \simeq over finite histories. In the following, we write $\mathcal{L}^{\mathbb{K}}$ for the syntactic fragment of \mathcal{L} without operator \mathbb{K} .

Already, the language $\mathcal{L}^{\mathbb{K}}$ may lead to specify sets of uniform strategies that do not form a regular tree-language, hence they are not μ -calculus definable. Indeed, for example the property that nodes at the same depth of a labelled-tree share the same label cannot be regular (by a simple Pumping

Lemma argument). As already mentioned in [1], this is what requires, with an appropriate relation \sim , the $\mathcal{L}^{\mathbb{K}}$ -formula $\mathbf{FK}p_S$ of Theorem 4, expressing that a strategy of Player 1 is winning in a game with opacity condition.

Note that regular property checking techniques may be used to verify *observational determinism* [2, 12]. This approach amounts to evaluate a polyadic μ -calculus formula over the self-product of the system. These results are consistent with our observation that \mathcal{L} can express non regular tree-language properties. In fact, observational determinism turns out to be a particular *hyperproperty*, in the sense of [6], that is a property of sets of traces. Observational determinism is a *2-safety hyperproperty*, and as established by [6], checking a k -safety hyperproperty ($k \in \mathbb{N}$) can be reduced to checking a safety property on a product of k copies of the structure. However, the opacity-guarantee property expressed in our language (see Section 3.2), although being a safety hyperproperty, is not a k -safety hyperproperty for any k : indeed, for arbitrary arenas, using k copies of the arena may not be sufficient to guarantee that all reachable information sets (of size possibly greater than the fixed k) are not contained in the secret. As a result, techniques using a polyadic μ -calculus on the product of k copies of the structure do not provide a complete method.

Beyond the fact that operator \mathbb{K} yields non-regular properties of strategies, a gain in expressiveness is reached by using the operator \mathbb{X} , quantifying over a \sim -equivalence class (of infinite plays). This gain is strict as exemplified by the instances of Figure 1: two arenas \mathcal{G} and \mathcal{G}' are considered, and we distinguish the plays π and π' : $\pi = v_0v_1(v_2)^\omega$, $\pi' = v'_0v'_1(v'_2)^\omega$. It can be shown by induction over the structure of the formulas that the logic $\mathcal{L}^{\mathbb{K}}$ cannot separate plays π and π' , whereas $\mathbb{X}p$ holds only in π' .



$$\text{Plays}(\mathcal{G}), v_0v_1(v_2)^\omega, 0 \not\models_{(\sim, \sim^\omega)} \mathbb{X}p \quad \text{and} \quad \text{Plays}(\mathcal{G}'), v'_0v'_1(v'_2)^\omega, 0 \models_{(\sim, \sim^\omega)} \mathbb{X}p.$$

Figure 1

5 Some computational aspects

In this section we will be interested in two problems: the uniform strategy checking problem and the uniform strategy existence problem.

Definition 5 *The strictly (resp. fully) uniform strategy problem is, given a game \mathcal{G} and a uniformity constraint $\mu = (\sphericalangle, \frown, \varphi)$, to decide whether there exists a strategy σ that is μ -strictly (resp. fully) uniform.*

Definition 6 *The strict (resp. full) uniformity problem is to decide whether a given strategy σ is μ -strictly (resp. fully) uniform.*

Theorem 8 *The uniform strategy problem is undecidable.*

We prove this result by reduction of the Post Correspondence Problem (PCP). Let Σ be an alphabet with at least two letters, and let $(\alpha_1, \beta_1) \dots (\alpha_N, \beta_N)$ be a nonempty sequence of ordered couples of words over Σ . We define the game $G = (V_1, V_2, E, v_0, Prop, \ell)$ as follows:

$V_2 = \{v_0, v_f\}$, $V_1 = \{v_\alpha, v_\beta\} \cup \{\overline{\alpha_i[j]} \mid 1 \leq i \leq N, 0 \leq j < |\alpha_i|\} \cup \{\overline{\beta_i[j]} \mid 1 \leq i \leq N, 0 \leq j < |\beta_i|\}$, $v_0 E = \{v_\alpha, v_\beta\}$, $v_\alpha E = \{\overline{\alpha_i[0]} \mid 1 \leq i \leq N\}$, $v_\beta E = \{\overline{\beta_i[0]} \mid 1 \leq i \leq N\}$, $\overline{\alpha_i[j]} E = \{\overline{\alpha_i[j+1]}\}$ for $1 \leq i \leq N$ and $0 \leq j < |\alpha_i| - 1$, $\overline{\beta_i[j]} E = \{\overline{\beta_i[j+1]}\}$ for $1 \leq i \leq N$ and $0 \leq j < |\beta_i| - 1$, $\overline{\alpha_i[|\alpha_i| - 1]} E = \{\overline{\alpha_j[0]} \mid 1 \leq j \leq N\} \cup \{v_f\}$ and $\overline{\beta_i[|\beta_i| - 1]} E = \{\overline{\beta_j[0]} \mid 1 \leq j \leq N\} \cup \{v_f\}$ for $1 \leq i \leq N$ and finally $v_f E = \{v_f\}$. The game starts in v_0 , and Player 2 chooses whether we are going to read the α_i 's or the β_i 's. Next Player 1 chooses the sequence of indexes she proposes as a solution to the PCP. If the first index chosen is i in position v_α (resp. v_β), there is no choice but to read all the letters of α_i (resp. β_i). When arrived in $\overline{\alpha_i[|\alpha_i| - 1]}$ (resp. $\overline{\beta_i[|\beta_i| - 1]}$), Player 1 chooses the next index, say j , and moves to $\overline{\alpha_j[0]}$ (resp. $\overline{\beta_j[0]}$) or decides to go to the final sink position v_f . Positions are decorated with propositions in $Prop = \{p_x \mid x \in \Sigma\} \cup \{p_f\}$. $\ell(v_0) = \ell(v_\alpha) = \ell(v_\beta) = \emptyset$, $\ell(\overline{\alpha_i[j]}) = \{p_{\alpha_i[j]}\}$, $\ell(\overline{\beta_i[j]}) = \{p_{\beta_i[j]}\}$ and $\ell(v_f) = \{p_f\}$.

We define two different equivalence relations on plays, \simeq_1 and \simeq_2 . For \simeq_2 we first define an ‘‘observation’’ function, that is the sequence of indexes chosen by Player 1 in the play.

$$Obs(v_0) = \epsilon \quad \text{and} \quad Obs(\rho v) = \begin{cases} Obs(\rho) i & \text{if } v \in \{\overline{\alpha_i[0]}, \overline{\beta_i[0]}\} \text{ for some } i \\ Obs(\rho) & \text{otherwise} \end{cases}$$

Now, $\rho \simeq_1 \rho'$ if $|\rho| = |\rho'|$, $\rho \simeq_2 \rho'$ if $Obs(\rho) = Obs(\rho')$, and we let

$$\text{GoodWord} := \mathbf{F}p_f \wedge \mathbf{G} \bigwedge_{\{p_x \mid x \in \Sigma\}} p_x \rightarrow \mathbb{A}p_x$$

Lemma 9 *The PCP has a solution if, and only if, Player 1 has a $(\simeq_1, \simeq_2^\omega, \text{GoodWord})$ -fully uniform strategy.*

Proof To prove this lemma, first notice that for any strategy σ for Player 1 there are only two plays in $Outcome(\sigma)$: one for each choice of Player 2 in the initial position. Notice also that each finite (resp. infinite) play ρ (resp. π) defines a word in Σ^* , $w(\rho)$ (resp. $w(\pi)$).

Suppose that there exists a $(\simeq_1, \simeq_2^\omega, \text{GoodWord})$ -fully uniform strategy σ for Player 1. Let $\pi \in Outcome(\sigma)$ be the play induced by σ that defines the longest word $w(\pi)$. Let $i_1 \dots i_N = Obs(\pi)$.

Because $Plays_\omega, \pi, 0 \models \mathbf{F}p_f$, we have that $N < 0$. Depending on the choice made by Player 2 in π , $w(\pi) = \alpha_{i_1} \dots \alpha_{i_N}$ or $w(\pi) = \beta_{i_1} \dots \beta_{i_N}$. Without loss of generality assume that $w(\pi) = \alpha_{i_1} \dots \alpha_{i_N}$. Let $\pi' \in Plays_\omega$ be the play in which Player 2 makes the other choice than in π , *i.e.* she chooses v_β , and in which Player 1 chooses the same sequence of indices $i_1 \dots i_N$, so that we have $w(\pi') = \beta_{i_1} \dots \beta_{i_N}$. We write $l = |w(\pi)|$ and $l' = |w(\pi')|$. By assumption $l \geq l'$. We prove that $w(\pi) = w(\pi')$, which concludes. Let $0 \leq i < l$. We have that $Plays_\omega, \pi, i + 2 \models p_{w(\pi)[i]}$ (the first two positions of every play, $v_0 v_\alpha$ or $v_0 v_\beta$, do not correspond to any letter). Because $Plays_\omega, \pi, 0 \models \mathbf{G} \bigwedge_{\{p_x | x \in \Sigma\}} p_x \rightarrow \mathbb{X}p_x$ and $Plays_\omega, \pi, i + 2 \models p_x$, we have that $Plays_\omega, \pi, i + 2 \models \mathbb{X}p_x$. Clearly $Obs(\pi) = Obs(\pi')$, so $\pi \simeq_2^\omega \pi'$, and $\pi[0, i + 2] \simeq \pi'[0, i + 2]$, so $Plays_\omega, \pi', i + 2 \models p_x$, hence $w(\pi')[i] = x$. So for $i = 0 \dots l - 1$ we have that $w(\pi)[i] = w(\pi')[i]$. This implies that $l' \geq l$, and since $l \geq l'$ we have that $l = l'$, hence $\alpha_{i_1} \dots \alpha_{i_N} = \beta_{i_1} \dots \beta_{i_N}$.

Now suppose that $i_1 \dots i_N$ is a solution to the problem. We define the strategy σ for Player 1 (we only define it for relevant finite plays, *i.e.* those that follow the strategy):

$$\sigma(\rho v) = \begin{cases} \overline{\alpha_{i_1}[0]} & \text{if } v = v_\alpha \\ \overline{\beta_{i_1}[0]} & \text{if } v = v_\beta \\ \overline{\alpha_i[j + 1]} & \text{if } v = \overline{\alpha_i[j]} \text{ with } j < |\alpha_i| - 1 \\ \overline{\beta_i[j + 1]} & \text{if } v = \overline{\beta_i[j]} \text{ with } j < |\beta_i| - 1 \\ \overline{\alpha_{i_{k+1}}[0]} & \text{if } v = \overline{\alpha_i[|\alpha_i| - 1]}, Obs(\rho v) = i_1 \dots i_k \text{ and } k < N \\ \overline{\beta_{i_{k+1}}[0]} & \text{if } v = \overline{\beta_i[|\beta_i| - 1]}, Obs(\rho v) = i_1 \dots i_k \text{ and } k < N \\ v_f & \text{if } v = \overline{\alpha_i[|\alpha_i| - 1]} \text{ and } Obs(\rho v) = i_1 \dots i_N \\ v_f & \text{if } v = \overline{\beta_i[|\beta_i| - 1]} \text{ and } Obs(\rho v) = i_1 \dots i_N \end{cases}$$

We prove that σ is ($\simeq_1, \simeq_2^\omega, \mathbf{GoodWord}$)-fully uniform. Let $\pi \in Outcome(\sigma)$. $N < \omega$, so $Plays_\omega, \pi, 0 \models \mathbf{F}p_f$. Now let $i \geq 0$ and suppose that there exists $x \in \Sigma$ such that $Plays_\omega, \pi, i \models p_x$. Necessarily $i \geq 2$ and $x = w(\pi)[i - 2]$. Let π', j such that $\pi \simeq_2^\omega \pi'$ and $\pi, i \simeq_1 \pi', j$. By definition of σ , we have $Obs(\pi) = i_1 \dots i_N$, hence $w(\pi) = \alpha_{i_1} \dots \alpha_{i_N}$ or $w(\pi) = \beta_{i_1} \dots \beta_{i_N}$. Since $\pi \simeq_2^\omega \pi'$, $Obs(\pi') = Obs(\pi) = i_1 \dots i_N$, so $w(\pi') = \alpha_{i_1} \dots \alpha_{i_N}$ or $w(\pi') = \beta_{i_1} \dots \beta_{i_N}$. Since $i_1 \dots i_N$ is a solution to the PCP, we have $\alpha_{i_1} \dots \alpha_{i_N} = \beta_{i_1} \dots \beta_{i_N}$, so $w(\pi) = w(\pi')$. And because $\pi, i \simeq_1 \pi', j$, we have $i = j$. $Plays_\omega, \pi', i \models p_{w(\pi')[i]}$ and $w(\pi)[i] = w(\pi')[i]$, so $Plays_\omega, \pi', i \models p_x$. So $Plays_\omega, \pi, 0 \models \mathbf{GoodWord}$. \square

Now we expose a positive result concerning the uniformity problem, and relate a subclass of our framework to Epistemic Temporal Logic. The subclass of the framework is defined by some further assumptions on the strategies used, the equivalence relations and the language.

First, we consider that strategies can be represented by finite I/O automaton, as done in eg.[8]. Such an automaton representing a strategy for Player k takes positions of the play as inputs, and when it is Player k 's turn to play, it outputs the next position the player should choose. This assumption is basically that the strategies considered require only finite memory, represented by the states of the automaton.

Next, we consider that equivalence relations \simeq are *observation-based, synchronous with perfect-recall*. We mean that for every game \mathcal{G} and relation \simeq , there exists a set of observations \mathcal{O} (corresponding or not to one of the player's observational power) and a mapping $Obs : V \rightarrow \mathcal{O}$ such that for all $\rho, \rho' \in Plays_*$, $\rho \simeq \rho'$ if, and only if, $|\rho| = |\rho'|$ and for all $0 \leq i < |\rho|$, $Obs(\rho[i]) = Obs(\rho'[i])$.

Finally we restrict the language used to $\mathcal{L}^{\mathbf{K}}$, and we consider only strict uniformity.

In these conditions, we have the following result:

Theorem 10 *Considering strategies with finite memory, observation-based synchronous perfect-recall equivalence relations, and the sub-language $\mathcal{L}^{\mathbb{K}}$, the strict uniformity problem is in PSPACE.*

The proof is done by reduction to the framework of Linear Time Logic of Knowledge (LTLK) described in [20, 9].

An *interpreted environment* for one agent is a tuple $E = (S, I, \Delta, O, \ell, F)$ where S is a set of states, $I \subseteq S$ is a set of initial states, $\Delta \subseteq S^2$ is a transition relation, $O : S \rightarrow \mathcal{O}$ is an observation function, where \mathcal{O} is a set of observations, $\ell : S \rightarrow 2^{Prop}$ is a labelling function, and $F \subseteq S$ is a Büchi acceptance condition.

Let G be a game, \mathcal{A}_σ an I/O automaton for a strategy σ , \sim an equivalence relation and φ a $\mathcal{L}^{\mathbb{K}}$ -formula. It is not difficult, by means of an appropriate synchronous product of the game arena with \mathcal{A}_σ , to obtain a transition system $S = (S, I, \Delta)$ whose set of traces is $Out(\sigma)$. Then, since \sim is observation-based we can define the observation function O , the labelling function ℓ is defined like in the original game, and we take the trivial acceptance condition $F = S$. This defines the interpreted environment E_σ^G constructible in polynomial time. And the syntax and semantics of LTLK formulas being the same as for $\mathcal{L}^{\mathbb{K}}$, the result follows from the result proved in [9] that the uniformity problem for LTLK with one agent and synchronous perfect-recall is in PSPACE.

6 Discussion

Our notion of uniform strategies is a very general notion the relevance of which is demonstrated by the many examples from the literature that can be captured and which provides a clean mathematical setting to work on. Still, the present proposal deserves further study in many directions. For example, we should investigate whether the hypothesis on the binary relations between plays can be relaxed, just as epistemic accessibility relations need not be equivalences in general. Even if we stick to equivalences, those not satisfying the Uniform No Miracle property of [19] are not necessarily congruences, which may complicate the way to extend it to infinite plays.

Also, results on expressiveness, and answers on decidability and computational complexity questions are burning topics. In its full generality, the setting should lead to numerous undecidability results because it enables to navigate both on the vertical and horizontal dimensions of the trees. A preliminary comparison between our language and Second Order Logic over trees should throw light on these issues.

Note that we could also allow several \mathbb{K}_i operators associated to several relations \sim_i between plays; this would enable us to represent, e.g. the different players' knowledge. Theorem 10 stating that the uniformity problem under certain assumptions is in PSPACE would no longer hold, but the same reduction gives a non-elementary upper bound [20], and if relations are based only on the last position of plays instead of having perfect-recall, then the problem falls back in PSPACE [9].

Finally, uniform strategies are defined via a language. At the moment, the language is based on LTL, but it can be extended to the propositional μ -calculus with the \mathbb{K} and \mathbb{X} operators. Whichever temporal logic is chosen, it would be interesting to have a language-independent definition of uniform strategies, and to establish the completeness of a given language, like the one we propose, with regard to this hypothetical semantic-based definition.

References

- [1] R. Alur, P. Černý, and S. Zdancewic. Preserving secrecy under refinement. *Automata, Languages and Programming*, pages 107–118, 2006.
- [2] G. Barthe, P.R. D’Argenio, and T. Rezk. Secure information flow by self-composition. In *Computer Security Foundations Workshop, 2004. Proceedings. 17th IEEE*, pages 100–114. IEEE, 2004.
- [3] D. Berwanger and L. Doyen. On the power of imperfect information. In *Proc. of FSTTCS*, pages 73–82. Citeseer, 2008.
- [4] C.G. Cassandras and S. Lafortune. Discrete event systems- the state of the art and new directions. *Applied and computational control, signals, and circuits.*, 1:1–65, 1999.
- [5] K. Chatterjee, L. Doyen, T. Henzinger, and J.F. Raskin. Algorithms for omega-regular games with imperfect information. In *Computer Science Logic*, pages 287–302. Springer, 2006.
- [6] M.R. Clarkson and F.B. Schneider. Hyperproperties. *Journal of Computer Security*, 18(6):1157–1210, 2010.
- [7] R. Dimitrova, B. Finkbeiner, M. Kovács, M. Rabe, and H. Seidl. Model checking information flow in reactive systems. In *Verification, Model Checking, and Abstract Interpretation*, pages 169–185. Springer, 2012.
- [8] S. Dziembowski, M. Jurdzinski, and I. Walukiewicz. How much memory is needed to win infinite games? In *Logic in Computer Science, 1997. LICS’97. Proceedings., 12th Annual IEEE Symposium on*, pages 99–110. IEEE, 1997.
- [9] K. Engelhardt, P. Gammie, and R. Van Der Meyden. Model checking knowledge and linear time: Pspace cases. *Logical Foundations of Computer Science*, pages 195–211, 2007.
- [10] J.A. Goguen and J. Meseguer. Security policies and security models. In *IEEE Symposium on Security and Privacy*, volume 12, 1982.
- [11] J.Y. Halpern and M.Y. Vardi. The complexity of reasoning about knowledge and time. 1. lower bounds. *Journal of Computer and System Sciences*, 38(1):195–237, 1989.
- [12] M. Huisman, P. Worah, and K. Sunesen. A temporal logic characterisation of observational determinism. In *Computer Security Foundations Workshop, 2006. 19th IEEE*, pages 13–pp. IEEE, 2006.
- [13] Bastien Maubert, Sophie Pinchinat, and Laura Bozzelli. Opacity issues in games with imperfect information. In Giovanna D’Agostino and Salvatore La Torre, editors, *GandALF*, volume 54 of *EPTCS*, pages 87–101, 2011.
- [14] D. Peled, T. Wilke, and P. Wolper. An algorithmic approach for checking closure properties of temporal logic specifications and [omega]-regular languages. *Theoretical Computer Science*, 195(2):183–203, 1998.
- [15] M. Piccione and A. Rubinstein. The absent-minded driver’s paradox: synthesis and responses. *Games and Economic Behavior*, 20(1):121–130, 1997.

- [16] J.H. Reif. The complexity of two-player games of incomplete information. *Journal of computer and system sciences*, 29(2):274–301, 1984.
- [17] M. Sampath, R. Sengupta, S. Lafortune, K. Sinaamohideen, and D. Teneketzis. Diagnosability of discrete event systems. *IEEE Transactions on Automatic Control*, 40(9):1555–1575, 1995.
- [18] J. Väänänen. Dependence logic, 2007.
- [19] J. Van Benthem, J. Gerbrandy, T. Hoshi, and E. Pacuit. Merging frameworks for interaction. *Journal of Philosophical Logic*, 38(5):491–526, 2009.
- [20] R. van der Meyden and N. Shilov. Model checking knowledge and time in systems with perfect recall. In *Foundations of Software Technology and Theoretical Computer Science*, pages 432–445. Springer, 1999.
- [21] T.S. Yoo and S. Lafortune. Polynomial-time verification of diagnosability of partially observed discrete-event systems. *Automatic Control, IEEE Transactions on*, 47(9):1491–1495, 2002.
- [22] S. Zdancewic and A.C. Myers. Observational determinism for concurrent program security. In *Computer Security Foundations Workshop, 2003. Proceedings. 16th IEEE*, pages 29–43. IEEE, 2003.