

Relating paths in transition systems: the fall of the modal mu-calculus*

Cătălin Dima¹, Bastien Maubert², and Sophie Pinchinat³

¹ Université Paris Est, LACL, UPEC, Créteil, France.
dima@u-pec.fr

² LORIA - CNRS / Université de Lorraine, Nancy, France.
bastien.maubert@gmail.com

³ Université de Rennes 1, IRISA, Rennes, France.
sophie.pinchinat@irisa.fr

Abstract. We revisit Janin and Walukiewicz’s classic result on the expressive completeness of the modal mu-calculus w.r.t. MSO, when transition systems are equipped with a binary relation over paths. We obtain two natural extensions of MSO and the mu-calculus: *MSO with path relation* and the *jumping mu-calculus*. While “bounded-memory” binary relations bring about no extra expressivity to either of the two logics, “unbounded-memory” binary relations make the bisimulation-invariant fragment of MSO with path relation more expressive than the jumping mu-calculus: the existence of winning strategies in games with imperfect-information inhabits the gap.

1 Introduction

Monadic second-order logic (MSO) is a standard for comparing expressiveness of other logics of programs. Ground-breaking expressiveness results on MSO were obtained first on “freely-generated” structures (words, trees, tree-like structures, etc.) [26, 29], then on “non-free” structures like grids [18] or infinite graphs generated by regularity-preserving transformations [10, 8]. Some attention has also been brought to the study of enrichments of MSO with unary predicate symbols or with the “equal level” binary predicate (MSO^{eql}) [11, 25].

Many of these expressiveness results relate MSO with automata and modal logics, among which Janin and Walukiewicz’s seminal result [17] showing that the bisimulation-invariant fragment of MSO interpreted over transition systems is captured by the μ -calculus. Notable exceptions to the classical trilogy between MSO, modal logics and automata are MSO on infinite partial orders (see [23] for partial results) and MSO^{eql} (partial results can be found in [25]).

On the other hand, more recently there has been an increased interest in the expressiveness and decidability of logics defined on structures in which two “orthogonal” relations are considered: the so-called temporal epistemic (multi-agent) logics [12], which combine time-passage relations and epistemic relations on the histories of the system [12, 15]. A natural question that arises is whether

* Catalin Dima acknowledges financial support from the ANR project EQINOCS. Bastien Maubert acknowledges financial support from the ERC project EPS 313360.

there exists a natural extension of MSO, of the μ -calculus, and of tree automata for the temporal epistemic framework, and how would they compare?

Note that appropriate extensions of MSO, of the μ -calculus and of tree automata would rely on two sorts of binary relations: those related to dynamic behaviour and those related to epistemic features. While the temporal part of these logics naturally refers to a tree-like structure, the epistemic part requires, in order to model *e.g.* powerful agents with “perfect recall”, to consider binary relations defined on histories. Such models neither are tree-like structures, nor grid-like structures, nor graphs within the Caucal hierarchy. The only proposals in this direction that we know about are [27, 24, 1] and [19]. [27] mentions an encoding of LTL with knowledge into Chain Logic with equal-level predicate, a fragment of MSO^{eq} . [24] introduces the epistemic μ -calculus and studies its model-checking problem. [1] studies temporal logics over tree models with “jump-edges” which capture observational indistinguishability. Each one considers only one particular relation on histories, all three akin to synchronous perfect recall. [19] proposes a generalization of tree automata, called jumping tree automata, applying them to the study of temporal epistemic logics.

In this paper, we propose natural extensions of MSO and the μ -calculus, respectively called *MSO with path relation* ($\text{MSO}^\curvearrowright$) and the *jumping μ -calculus* ($\text{L}_\mu^\curvearrowright$). $\text{MSO}^\curvearrowright$ is MSO with an additional binary predicate, \curvearrowright , interpreted “transversely” on the tree structure, according to the path relation. $\text{L}_\mu^\curvearrowright$ is a generalization of the epistemic μ -calculus defined in [24]: it features a *jumping modality* whose semantics relies on the path relation, generalizing the knowledge operator K . The path relations that we consider are arbitrary and thus cover all variants of indistinguishability relations of, *e.g.* [16], as well as settings not particularly related to any epistemic interpretation [13]. We investigate whether $\text{L}_\mu^\curvearrowright$ is expressive complete with regards to the bisimulation-invariant fragment of $\text{MSO}^\curvearrowright$. For the class of *recognizable* path relations (finite memory), we observe (Theorem 1) that they add no expressivity w.r.t. the classical case hence the expressive completeness holds. We show however that the case of *regular* path relations breaks this completeness: the class of reachability games with imperfect information and *synchronous perfect recall* [5] where the first player wins cannot be defined in the jumping μ -calculus, while being closed under bisimulation and definable in our extension of MSO (Theorem 2). As an intermediate result, we show that the *jumping tree automata* defined in [19] are equivalent to the jumping μ -calculus.

The paper runs as follows: in Section 2, we develop the framework of our study on transition systems. We introduce in Section 3 the two extensions of MSO and the μ -calculus, and state our results on expressive completeness (Theorem 1 and 2). In Section 4 we establish Theorem 1 and prove that jumping automata are equivalent to the jumping μ -calculus, thanks to which we establish succinctness and complexity results on the jumping μ -calculus with recognizable path relation. In Section 5 we establish our main result (Theorem 2) by proving that winning in reachability games with imperfect information and synchronous perfect recall is not definable in the jumping μ -calculus, unless the winning con-

ditions are observable (Theorem 3). We conclude, comment on the impacts of our results, and give some perspectives in Section 6.

2 Preliminary notions

We first fix a few basic notations. Given two words w and w' over some alphabet Σ , we write $w \preceq w'$ if w is a prefix of w' ; if $w = a_0a_1\dots \in \Sigma^\omega$ is an infinite word we let, for each $i \geq 0$, $w[i] := a_i$ and $w[0, i] := a_0a_1\dots a_i$. For a finite word $w = a_0\dots a_{n-1} \in \Sigma^*$, its *length* is $|w| := n$. Also, given a binary relation $R \subseteq A \times B$ for every $a \in A$, we let $R(a) := \{b \mid (a, b) \in R\}$. In the rest of the paper, we fix $\mathcal{AP} = \{p, p', \dots\}$ a countable set of *atomic propositions* and $\text{Act} = \{a, a', \dots\}$ a countable set of *actions*.

Definition 1. A transition system (over \mathcal{AP} and Act) is a structure $\mathcal{S} = (S, s_i, \{a^S\}_{a \in \text{Act}}, \{p^S\}_{p \in \mathcal{AP}})$, where S is a countable set of states, s_i is an initial state, each a^S is a binary relation over S and each p^S is a subset of S .

The logics that we aim at are concerned with *paths*⁴ in transition systems, that is, on their *tree unfoldings*. To ease the presentation, our trees have at most countable branching degree, but, unless otherwise stated, our results still hold for arbitrary degree. A *tree* is a nonempty, prefix-closed set $\tau \subseteq \mathbb{N}^*$. An element $x \in \tau$ is a *node*, and the empty word ϵ is the *root* of the tree. If $x \cdot i \in \tau$, $x \cdot i$ is a *child* of x . A node with no child is a *leaf*. A *branch* is a sequence of nodes in τ (either finite or infinite) in which each node but the first one is a child of the previous one; a branch is *maximal* if it is infinite or it ends up in a leaf. We write $x \preceq y$ if y can be found on some branch that starts in x , and we let $[\tau]_x = \{y \mid x \preceq y\}$ denote the subtree of τ rooted in x . A *marked tree* (over \mathcal{AP} and Act) is a pair $t = (\tau, m)$, where τ is a tree and $m : \tau \rightarrow (\text{Act} \times 2^{\mathcal{AP}})$ is a *marking* of the nodes, where $m(x) = (a, \ell)$ means that x was reached through action a and ℓ is the set of atomic propositions that hold in x ; we may use notation a^x and ℓ^x for a and ℓ when $m(x) = (a, \ell)$. Node y is an a -child of a node x if y is a child of x and $a^y = a$. The *word* of a node x is $w(x) := m(\epsilon)m(x_1)\dots m(x_n)$, where $\epsilon x_1\dots x_n (= x)$ is the (unique) branch from the root to x . For a finite subset $AP \subset \mathcal{AP}$, an *AP-tree* is a labeled tree $t = (\tau, m)$ such that $\ell^x \subseteq AP$, for every node $x \in t$ (i.e. $x \in \tau$).

Definition 2. Let $\mathcal{S} = (S, s_i, \{a^S\}_{a \in \text{Act}}, \{p^S\}_{p \in \mathcal{AP}})$ be a transition system. The unfolding $t_{\mathcal{S}}$ of \mathcal{S} is the marked tree (τ, m) with least tree τ such that: ϵ is associated⁵ to s_i and $\ell^\epsilon = \{p \mid s_i \in p^S\}$, and for each node $x \in \tau$ associated to state s , if $\langle s \xrightarrow{a_i} s_i \rangle_{i \in I}$ is an enumeration of the outgoing transitions from s (with $I \subseteq \mathbb{N}$), then for each $i \in I$ we have $x \cdot i \in \tau$, $x \cdot i$ is associated to s_i and $m(x \cdot i) = (a_i, \{p \in \mathcal{AP} \mid s_i \in p^S\})$.

⁴ i.e. finite sequences of states and actions that start in the initial state and follow the binary relations

⁵ The notion of “associated state” is only used to define unfoldings and is left informal.

Because in the following only actions and atomic propositions matter, the ordering of children nodes in trees is irrelevant, and the unfolding $t_{\mathcal{S}}$ is therefore uniquely defined up to isomorphism.

Definition 3. We call path relation a binary relation over $(Act \times 2^{\mathcal{AP}})^*$.

A path relation links finite paths of transition systems over \mathcal{AP} and Act . It also induces a binary relation between nodes of marked trees (over \mathcal{AP} and Act) in a natural way by relating nodes x and y whenever their words $w(x)$ and $w(y)$ are related. We use notation \rightsquigarrow for path relations.

Finally, we recall the classic notion of *bisimulation* [20].

Definition 4. A bisimulation between transition systems \mathcal{S} and \mathcal{S}' is a binary relation $\mathcal{Z} \subseteq S \times S'$ such that, for all $(s, s') \in \mathcal{Z}$, for all $p \in \mathcal{AP}$ and $a \in Act$:

1. $s \in p^{\mathcal{S}}$ iff $s' \in p^{\mathcal{S}'}$;
2. for all $r \in a^{\mathcal{S}}(s)$, there is $r' \in a^{\mathcal{S}'}(s')$ such that $(r, r') \in \mathcal{Z}$;
3. and vice-versa.

We write $\mathcal{S} \rightleftharpoons \mathcal{S}'$ whenever there is a bisimulation \mathcal{Z} between \mathcal{S} and \mathcal{S}' such that $(s_i, s'_i) \in \mathcal{Z}$. A class \mathcal{C} of transition systems is *closed under bisimulation*, or *bisimulation closed* if $\mathcal{S} \in \mathcal{C}$ and $\mathcal{S} \rightleftharpoons \mathcal{S}'$ imply $\mathcal{S}' \in \mathcal{C}$, for all \mathcal{S} and \mathcal{S}' .

3 Expressive Completeness Issues

We fix a countable set of *second order variables* $Var = \{X, Y, \dots\}$. Given a marked tree $t = (\tau, m)$, a *valuation* is a mapping $V : Var \rightarrow 2^{\tau}$. For $X \in Var$ and $T \subseteq \tau$, we let $V[T/X]$ be the valuation that maps X to T , and which coincides with V on all other variables.

Monadic second order logic with path relation ($\text{MSO}^{\rightsquigarrow}$) is an extension of MSO interpreted over transition systems with a path relation. Its syntax is as follows:

$$\psi ::= sr(X) \mid p(X) \mid succ(X, Y) \mid X \subseteq Y \mid \neg\psi \mid \psi \vee \psi' \mid \exists X.\psi(X) \mid X \rightsquigarrow Y$$

where $p \in \mathcal{AP}$ and $X, Y \in Var$.

An $\text{MSO}^{\rightsquigarrow}$ formula ψ is interpreted over a marked tree $t = (\tau, m)$ with a valuation V and a fixed path relation \rightsquigarrow ; the fact that t with valuation V satisfies ψ is written $t, V \models^{\rightsquigarrow} \psi$, defined inductively as follows:

$$\begin{aligned} t, V \models^{\rightsquigarrow} sr(X) & \text{ if } V(X) = \{\epsilon\} \\ t, V \models^{\rightsquigarrow} p(X) & \text{ if for all } x \in V(X), p \in \ell^x \\ t, V \models^{\rightsquigarrow} succ(X, Y) & \text{ if } V(X) = \{x\}, V(Y) = \{y\}, \text{ and } y \text{ is a child of } x \\ t, V \models^{\rightsquigarrow} X \subseteq Y & \text{ if } V(X) \subseteq V(Y) \\ t, V \models^{\rightsquigarrow} \neg\psi & \text{ if } t, V \not\models^{\rightsquigarrow} \psi \text{ and } t, V \models^{\rightsquigarrow} \psi \vee \psi' \text{ if } t, V \models^{\rightsquigarrow} \psi \text{ or } t, V \models^{\rightsquigarrow} \psi' \\ t, V \models^{\rightsquigarrow} \exists X.\psi(X) & \text{ if there is } T \subseteq \tau \text{ s.t. } t, V[T/X] \models^{\rightsquigarrow} \psi(X) \\ t, V \models^{\rightsquigarrow} X \rightsquigarrow Y & \text{ if } V(X) = \{x\}, V(Y) = \{y\}, \text{ and } x \rightsquigarrow y \end{aligned}$$

If $\psi \in \text{MSO}^{\rightsquigarrow}$ has no free variable, we simply write $t \models^{\rightsquigarrow} \psi$, and $\mathcal{S} \models^{\rightsquigarrow} \psi$ whenever $t_{\mathcal{S}} \models^{\rightsquigarrow} \psi$, for any transition system \mathcal{S} . Let $\mathcal{L}(\psi, \rightsquigarrow) := \{\mathcal{S} \mid \mathcal{S} \models^{\rightsquigarrow} \psi\}$.

The jumping μ -calculus. The syntax of the \curvearrowright -jumping μ -calculus L_μ^{\curvearrowright} is:

$$\varphi ::= X \mid p \mid \neg\varphi \mid \varphi \vee \varphi \mid \lozenge\varphi \mid \heartsuit\varphi \mid \mu X.\varphi(X)$$

where $X \in \text{Var}$, $p \in \mathcal{AP}$, $a \in \text{Act}$, and in the last rule, X appears only under an even number of negations (i.e. positively) in $\varphi(X)$. We classically define the dual operators \heartsuit , \boxplus , \boxminus and ν , e.g. $\boxplus\varphi := \neg\lozenge\neg\varphi$, and $\nu X.\varphi := \neg\mu X.\neg\varphi$.

Given a path relation \curvearrowright , formulas of L_μ^{\curvearrowright} are interpreted over a marked tree with a valuation $V : \text{Var} \rightarrow 2^\tau$. We inductively define $\llbracket \varphi \rrbracket_{\curvearrowright}^{t,V} \subseteq \tau$ with $t = (\tau, m)$:

$$\begin{aligned} \llbracket X \rrbracket_{\curvearrowright}^{t,V} &= V(X) & \llbracket p \rrbracket_{\curvearrowright}^{t,V} &= \{x \in t \mid p \in \ell^x\} & \llbracket \neg\varphi \rrbracket_{\curvearrowright}^{t,V} &= t \setminus \llbracket \varphi \rrbracket_{\curvearrowright}^{t,V} \\ \llbracket \varphi \vee \varphi' \rrbracket_{\curvearrowright}^{t,V} &= \llbracket \varphi \rrbracket_{\curvearrowright}^{t,V} \cup \llbracket \varphi' \rrbracket_{\curvearrowright}^{t,V} & \llbracket \lozenge\varphi \rrbracket_{\curvearrowright}^{t,V} &= \{x \in t \mid x \text{ has an } a\text{-child in } \llbracket \varphi \rrbracket_{\curvearrowright}^{t,V}\} \\ \llbracket \heartsuit\varphi \rrbracket_{\curvearrowright}^{t,V} &= \{x \in t \mid \text{there exists } y \in \llbracket \varphi \rrbracket_{\curvearrowright}^{t,V} \text{ such that } x \curvearrowright y\} \\ \llbracket \mu X.\varphi(X) \rrbracket_{\curvearrowright}^{t,V} &= \bigcap \{T \subseteq t \mid \llbracket \varphi(X) \rrbracket_{\curvearrowright}^{t,V[T/X]} \subseteq T\} \end{aligned}$$

Note that, for each formula $\mu X.\varphi(X)$, function $T \mapsto \llbracket \varphi(X) \rrbracket_{\curvearrowright}^{t,V[T/X]}$ is monotone, and hence has a least fixpoint, namely $\llbracket \mu X.\varphi(X) \rrbracket_{\curvearrowright}^{t,V}$. If $\varphi \in L_\mu^{\curvearrowright}$ has no free variables, we write $t \models^{\curvearrowright} \varphi$ whenever $\epsilon \in \llbracket \varphi \rrbracket_{\curvearrowright}^t$, and $\mathcal{S} \models^{\curvearrowright} \varphi$ whenever $t_{\mathcal{S}} \models^{\curvearrowright} \varphi$, for any transition system \mathcal{S} . We let $\mathcal{L}(\varphi, \curvearrowright) := \{\mathcal{S} \mid \mathcal{S} \models^{\curvearrowright} \varphi\}$.

Expressive Completeness. For a logic \mathcal{L} , a class \mathcal{C} of transition systems is \mathcal{L} -definable if there is a formula of \mathcal{L} whose set of models is exactly \mathcal{C} .

Proposition 1. *For every path relation \curvearrowright , every L_μ^{\curvearrowright} -definable class is closed under bisimulation.*

This result follows from Lemma 1 below.

Let \curvearrowright be a path relation. Note that a marked tree t can be turned into a transition system t^{\curvearrowright} over \mathcal{AP} and $\text{Act}' := \text{Act} \cup \{\check{a}\}$, where \check{a} is a fresh action symbol, by letting $y \in \check{a}^{t^{\curvearrowright}}(x)$ whenever $x \curvearrowright y$. The following lemma states that if two transition systems are bisimilar (w.r.t. the transition relations only), then their unfoldings enriched with a path relation are bisimilar w.r.t. *both* the transition relations and the path relation.

Lemma 1. *Let \mathcal{S} and \mathcal{S}' be two transition systems, and let \curvearrowright be a path relation. If $\mathcal{S} \cong \mathcal{S}'$, then $t_{\mathcal{S}}^{\curvearrowright} \cong t_{\mathcal{S}'}^{\curvearrowright}$.*

Proposition 2. *Every L_μ^{\curvearrowright} -definable class is $\text{MSO}^{\curvearrowright}$ -definable.*

Proposition 2 can easily be established with a straightforward extension of the effective translation of μ -calculus formulas into MSO given, e.g. in [14, Ch14].

We now engage our main concern, the *expressive completeness* of L_μ^{\curvearrowright} with respect to $\text{MSO}^{\curvearrowright}$. As in [17], due to Proposition 1, this question can only be addressed for bisimulation-closed classes of transition systems. We thus seek properties on the path relation \curvearrowright so that L_μ^{\curvearrowright} is expressive complete with respect to $\text{MSO}^{\curvearrowright}$, in the following sense:

Definition 5 (Expressive completeness). L_μ^\smile is expressive complete with respect to MSO^\smile if every bisimulation-closed class of transition systems that is MSO^\smile -definable is also L_μ^\smile -definable.

Note that the results in [17] do not adapt to MSO^\smile since unfoldings of transition systems with a path relation are not tree-like structures.

We briefly recall the notions of recognizable, regular and rational relations, and we refer to [3] for details. Let Σ be a finite alphabet. A binary relation over Σ^* is *rational* if there is a *transducer* (i.e. a finite-state two-tape automaton) that accepts precisely the pairs of words of the relation. A binary relation (over Σ^*) is *regular* if it is accepted by a *synchronous* transducer⁶. The epistemic relation of an agent with asynchronous perfect recall [22] is rational whereas the epistemic relation of an agent with synchronous perfect recall [4, 5] is regular. A binary relation is *recognizable* if there is a finite-state word automaton over $\Sigma \cup \{\#\}$ that accepts words of the form $w\#w'$, whenever w and w' are related. The epistemic relation of a “bounded-memory” agent is recognizable. We say that a path relation \smile (over $(Act \times 2^{AP})^*$) is *rational* (resp. *regular*, *recognizable*) if there are finite subsets $A \subset Act$ and $AP \subset AP$ such that \smile is equal to some rational (resp. regular, recognizable) relation over $(A \times 2^{AP})^*$. Finally, recall that recognizable relations are strictly contained in regular relations, and so are regular relations in rational relations.

Theorem 1. For any recognizable path relation, L_μ^\smile is expressive complete with respect to MSO^\smile .

Theorem 2. There are regular (hence rational) binary relations for which L_μ^\smile is not expressive complete with respect to MSO^\smile .

4 Tree automata for the jumping μ -calculus

We prove that *jumping tree automata* (JTA), introduced in [19], are equivalent to the jumping μ -calculus, which entails complexity and succinctness results for L_μ^\smile with recognizable path relation, and we prove Theorem 1. We assume familiarity with two-player turn-based games.

For a set X , $\mathbb{B}^+(X)$ (with typical elements $\alpha, \beta \dots$) is the set of formulas built with elements of X as atomic propositions using only connectives \vee and \wedge , and with $\top \in \mathbb{B}^+(X)$. Let $Dir = \{\diamond \mid a \in Act\} \cup \{\square \mid a \in Act\} \cup \{\lozenge, \boxtimes\}$ be the set of *automaton directions*.

Definition 6. A jumping tree automaton (JTA) over AP is a structure $\mathcal{A} = (AP, Q, q_\iota, \delta, C)$ where $AP \subset AP$ is a finite set of atomic propositions, Q is a finite set of states, $q_\iota \in Q$ is an initial state, $\delta : Q \times 2^{AP} \rightarrow \mathbb{B}^+(Dir \times Q)$ is a transition function, and $C : Q \rightarrow \mathbb{N}$ is a colouring function.

⁶ i.e. it progresses at the same pace on each tape.

JTAs resemble alternating tree automata [14, Ch. 9]. *Action directions* (i.e. \blacklozenge and \blacksquare) are meant to go down the input tree, whereas the new *jump directions* \blacklozenge and \boxtimes rely on an *a priori* given path relation. A JTA \mathcal{A} with a path relation \rightsquigarrow is written $(\mathcal{A}, \rightsquigarrow)$. Acceptance is defined on a two-player parity game between Eve and Adam. Let $t = (\tau, m)$ be an *AP*-tree, and let $\mathcal{A} = (AP, Q, q_L, \delta, C)$. We define the parity game $\mathcal{G}_{\mathcal{A}, \rightsquigarrow}^t = (V, v_L, E, V_E, V_A, C')$ whose set of positions is $V = \tau \times Q \times \mathbb{B}^+(\text{Dir} \times Q)$, whose initial position is $v_L = (\epsilon, q_L, \delta(q_L, \ell^\epsilon))$, and where a position (x, q, α) belongs to Eve iff α is of the form $\alpha_1 \vee \alpha_2$, $[\blacklozenge, q']$, or $[\boxtimes, q']$. The possible moves in $\mathcal{G}_{\mathcal{A}, \rightsquigarrow}^t$ are the following:

$$\begin{aligned} (x, q, \alpha_1 \dagger \alpha_2) &\rightarrow (x, q, \alpha_i) && \text{where } \dagger \in \{\vee, \wedge\} \text{ and } i \in \{1, 2\} \\ (x, q, [\@, q']) &\rightarrow (y, q', \delta(q', \ell^y)) && \text{where } \@ \in \{\blacklozenge, \blacksquare\} \text{ and } y \text{ is an } a\text{-child of } x \\ (x, q, [\ominus, q']) &\rightarrow (y, q', \delta(q', \ell^y)) && \text{where } \ominus \in \{\boxtimes, \boxminus\} \text{ and } x \rightsquigarrow y \end{aligned}$$

Deadlock positions are winning for Eve if of the form (x, q, \top) , $(x, q, [\blacksquare, q'])$ or $(x, q, [\boxminus, q'])$, and winning for Adam otherwise. The colouring function C' of $\mathcal{G}_{\mathcal{A}, \rightsquigarrow}^t$ is inherited from the one of \mathcal{A} by letting $C'(x, q, \alpha) = C(q)$. We let $\mathcal{L}(\mathcal{A}, \rightsquigarrow) = \{\mathcal{S} \mid \text{Eve has a winning strategy in } \mathcal{G}_{\mathcal{A}, \rightsquigarrow}^{\mathcal{S}}\}$.

Proposition 3.

- (a) For every formula $\varphi \in L_\mu^\rightsquigarrow$, there is a JTA \mathcal{A}_φ such that, for every path relation \rightsquigarrow , $\mathcal{L}(\varphi, \rightsquigarrow) = \mathcal{L}(\mathcal{A}_\varphi, \rightsquigarrow)$,
- (b) for every JTA \mathcal{A} , there is an L_μ^\rightsquigarrow -formula $\varphi_{\mathcal{A}}$ such that, for every path relation \rightsquigarrow , $\mathcal{L}(\mathcal{A}, \rightsquigarrow) = \mathcal{L}(\varphi_{\mathcal{A}}, \rightsquigarrow)$.

Moreover, the translations are effective and linear.

When restricting to recognizable path relations, the folklore fact that recognizable relations are MSO-definable gives the following.

Proposition 4. *MSO^{rightsquigarrow} with recognizable path relation is not more expressive than MSO.*

Theorem 1 is obtained from Proposition 2, Proposition 4 and the expressive completeness of the μ -calculus *w.r.t.* MSO. This collapse of the jumping μ -calculus down to the μ -calculus uses transformations that do not provide accurate complexity bounds regarding the jumping μ -calculus. However, by Proposition 3 and the relationship between two-way alternating automata and classic tree automata [28] we get:

Proposition 5. *The satisfiability problem for the jumping μ -calculus with recognizable path relation over transition systems with bounded branching degree is EXPTIME-complete.*

Proposition 6. *For a fixed recognizable path relation, the jumping μ -calculus with path relation over transition systems with bounded branching degree is at most exponentially more succinct than the μ -calculus.*

5 Games and the jumping μ -calculus

We focus on the property stating the existence of a winning strategy in two-player turned-based reachability games with imperfect information. This property gives us Theorem 2, for being bisimilar-invariant, expressible in MSO^ω , but not expressible in L_μ^ω (Theorem 4).

Two-player games with imperfect information [2, 9]. The players are Eveⁱ and Adamⁱ. Eveⁱ partially observes the positions, while Adamⁱ has perfect information. Let $\mathcal{Obs} = \{o, o', \dots\}$ be a countable set of *observations*. An *imperfect-information game arena* is a tuple $G^i = (V, v_\iota, \{a^{G^i}\}_{a \in \mathcal{Act}}, \{o^{G^i}\}_{o \in \mathcal{Obs}})$, where V is a set of positions, $v_\iota \in V$ is an initial position, each a^{G^i} is a binary relation over V and each o^{G^i} is a subset of V such that $\{o^{G^i}\}_{o \in \mathcal{Obs}}$ forms a partition of V . An action $a \in \mathcal{Act}$ is *available* in $v \in V$ if $a^{G^i}(v) \neq \emptyset$. We assume that some action is available in every position, and that two positions with the same observation share the same available actions. For $v \in V$, write o_v the unique observation such that $v \in o_v$.

Players take turns, starting with Eveⁱ. In current position v , Eveⁱ chooses an available action a and Adamⁱ chooses a new position $v' \in a^{G^i}(v)$. A *play* (resp. *partial play*) is an infinite (resp. finite) sequence $\pi = v_0 a_1 v_1 a_2 \dots$ (resp. $\rho = v_0 a_1 v_1 \dots a_n v_n$) such that $v_0 = v_\iota$ and, for all i , $v_{i+1} \in a_{i+1}^{G^i}(v_i)$. In the *synchronous perfect recall* setting [5], Eveⁱ remembers the whole sequence of observations that she receives, and her actions. The *indistinguishability equivalence* over partial plays is thus defined by: for $\rho = v_0 a_1 \dots a_n v_n$ and $\rho' = v'_0 a'_1 \dots a'_n v'_n$, we let $\rho \sim \rho'$ whenever, for all $1 \leq i \leq n$, $o_{v_i} = o_{v'_i}$ and $a_i = a'_i$.

Remark 1. Here, Eve remembers the sequence of actions. We point out that if she does not, then Theorem 3 below does not hold.

A (uniform) *strategy* for Eve is a partial function $\sigma : \{v_\iota\}(\mathcal{Act} \cdot V)^* \rightarrow \mathcal{Act}$ such that for two partial plays ρ and ρ' , if $\rho \sim \rho'$, then $\sigma(\rho) = \sigma(\rho')$. We say that a play $\pi = v_0 a_1 v_1 \dots$ *follows* a strategy σ if for all $i \geq 0$, $a_{i+1} = \sigma(v_0 a_1 v_1 \dots a_i v_i)$.

A (*reachability*) *game with imperfect information* \mathcal{G}^i is an imperfect-information arena $G^i = (V, v_\iota, \{a^{G^i}\}_{a \in \mathcal{Act}}, \{o^{G^i}\}_{o \in \mathcal{Obs}})$ together with a reachability winning condition $F \subseteq V$. A strategy for Eveⁱ is *winning* if every play that follows it visits F . F is *observable* if for all positions v and v' , $o_v = o_{v'}$ implies $(v \in F \Leftrightarrow v' \in F)$.

Note that a game with imperfect information is a transition system over $\mathcal{AP} = \mathcal{Obs} \cup \{F\}$ and \mathcal{Act} . Note also that the relation \sim on partial plays induces a path relation, that we shall also write \sim .

Consider the class $\mathcal{R}(A, O)$ of reachability games with imperfect information where actions range over $A \subseteq \mathcal{Act}$ and observations range over $O \subseteq \mathcal{Obs}$. To address logical definability, we restrict to games where A and O are finite sets. Also, because we address μ -calculus definability, we close by bisimulation, so that from now on a *game* is a transition system whose subsystem connected to the initial position is a game as defined above.

Notice that the subclass $\mathcal{R}(A, O)$ where Eveⁱ wins⁷ is bisimulation closed since, according to [5], for any two imperfect-information reachability games \mathcal{G}^i and $\mathcal{G}^{i'}$, if $\mathcal{G}^i \simeq \mathcal{G}^{i'}$, then Eveⁱ wins in \mathcal{G}^i if, and only if, she wins in $\mathcal{G}^{i'}$.

Theorem 3. *The subclass of $\mathcal{R}(A, O)$ of games with observable winning condition where Eveⁱ wins is definable in L_μ^\sim , namely by $\mu X.(F \vee \bigvee_{a \in A} \square \square X)$.*

Actually, a result similar to Theorem 3 can be established for parity conditions.

Theorem 4. *The subclass of $\mathcal{R}(A, O)$ where Eveⁱ wins is not L_μ^\sim -definable.*

Theorem 4 entails Theorem 2 since the subclass of $\mathcal{R}(A, O)$ where Eveⁱ wins is clearly MSO^\sim -definable, closed under bisimulation, and the synchronous perfect-recall relation \sim is regular⁸.

Proof of Theorem 4. We prove that Theorem 4 holds, already when we consider only two actions and one observation, *i.e.* $A = \{a_0, a_1\}$ and $O = \{o\}$.

The proof is dealt with by contradiction: Assume that there is a formula $\Phi_{\text{Win}} \in L_\mu^\sim$ such that for every $\mathcal{G}^i \in \mathcal{R}(A, O)$, $\mathcal{G}^i \models \Phi_{\text{Win}}$ if, and only if, Eveⁱ has a winning strategy in \mathcal{G}^i . By Proposition 3, there is a JTA $\mathcal{A} = (AP, Q, q_l, \delta, C)$ such that $\mathcal{L}(\Phi_{\text{Win}}, \sim) = \mathcal{L}(\mathcal{A}, \sim)$. Let $N := |Q| + 1$.

We describe 2^N (unfoldings of) games in $\mathcal{R}(A, O)$ where Eveⁱ wins. For each one, we exhibit a winning strategy in the (perfect information) acceptance game of \mathcal{A} on this unfolding. We then employ the ‘‘pigeon hole’’ principle to show that at least two of these ‘‘accepting’’ strategies can be combined into a new ‘‘accepting’’ strategy of Eve, entailing acceptance by \mathcal{A} of a game in $\mathcal{R}(A, O)$ where Eveⁱ has no winning strategy, hence the contradiction.

The family of game unfoldings that we consider is depicted in Figure 1. They all share the same unmarked tree τ , and for all $k \in \{1, \dots, 2^N + 2\}$, $[\tau]_{y_k}$ is the full binary tree with action a_0 (resp. a_1) leading to the left (resp. right) child. Let $w_k \in \{0, 1\}^N$ be the binary representation of $k - 1$. Between the different trees t_i , the markings only differ on the leaves of $[\tau]_{y_{2^{N+1}}}$ and $[\tau]_{y_{2^{N+2}}}$: for $1 \leq i \leq 2^N$, in $[t_i]_{y_{2^{N+1}}}$ and $[t_i]_{y_{2^{N+2}}}$, the only nodes in F are $y_{2^{N+1}} \cdot w_i$ and $y_{2^{N+2}} \cdot w_i$. Finally, since Eveⁱ is blind, her strategies are simply described by sequences of actions.

For each $1 \leq i \leq 2^N$, write $\mathcal{G}_i = (V^i, v_l^i, E^i, V_E^i, V_A^i, C^i)$ for $\mathcal{G}_{\mathcal{A}, \sim}^{t_i}$, *i.e.* the acceptance game of \mathcal{A} on t_i with relation \sim .

Clearly, for each $1 \leq i \leq 2^N$, Eveⁱ wins t_i , and thus Eve wins \mathcal{G}_i . Let σ_i be a winning strategy for Eve in each game \mathcal{G}_i . Let $\text{visit}_{\sigma_i} : \tau \rightarrow 2^Q$ associate to each node of τ the set of states of \mathcal{A} visited by strategy σ_i : formally, $\text{visit}_{\sigma_i}(x) := \{q \mid \exists \pi \in \text{Out}(\mathcal{G}_i, \sigma_i), \exists n \geq 0, \exists \alpha \in \mathcal{B}^+(\text{Dir} \times Q) \text{ s.t. } \pi[n] = (x, q, \alpha)\}$.

Since there are at most $2^{|Q|}$ different such sets of states, and we have 2^N strategies with $N = |Q| + 1$, there exist $i \neq j$ s.t. $\text{visit}_{\sigma_i}(y_{2^{N+1}}) = \text{visit}_{\sigma_j}(y_{2^{N+1}})$. Fix such a pair (i, j) , and define the game unfolding t_0 , obtained from t_i by replacing the subtree $[t_i]_{y_{2^{N+1}}}$ with $[t_j]_{y_{2^{N+1}}}$ (see Figure 1). Note that t_0 is

⁷ *i.e.* has a winning strategy

⁸ a one-state transducer that accepts it can easily be exhibited




Fig. 1. The tree t_i , the tree t_j , and the hybrid tree t_0 .

the unfolding of a game in $\mathcal{R}(A, O)$. Let us write $\mathcal{G}_0 = (V^0, v_\iota^0, E^0, V_E^0, V_A^0, C^0)$ for $\mathcal{G}_{\mathcal{A}, \sim}^{t_0}$. Observe that the three games \mathcal{G}_i , \mathcal{G}_j and \mathcal{G}_0 share the same set of positions: $V^0 = V^i = V^j = \tau \times Q \times \mathbb{B}^+(Dir \times Q)$ that we now write V . Also, for all $1 \leq k \leq 2^N + 2$, $\ell_0^{y_k} = \ell_i^{y_k} = \ell_j^{y_k} (= \{o\})$, that we now write ℓ . We succinctly write v_k^q for positions $(y_k, q, \delta(q, \ell))$, which play an important role.

The following lemma allows us to transfer the existence of winning strategies in positions v_k^q from \mathcal{G}_i and \mathcal{G}_j to \mathcal{G}_0 .

Lemma 2. 1. For all $q \in Q$ and $k \neq 2^N + 1$, $(\mathcal{G}_0, v_k^q) \Leftrightarrow (\mathcal{G}_i, v_k^q)$, and
 2. for all $q \in Q$ and $k \neq 2^N + 2$, $(\mathcal{G}_0, v_k^q) \Leftrightarrow (\mathcal{G}_j, v_k^q)$.

By design of t_0 , Eveⁱ has no winning strategy in t_0 . Therefore, $t_0 \not\stackrel{i}{\sim} \Phi_{\text{Win}}$, and thence $t_0 \notin \mathcal{L}(\mathcal{A}, \sim)$, *i.e.* Eve does not have a winning strategy in the acceptance game \mathcal{G}_0 of JTA \mathcal{A} on t_0 with path relation \sim . We establish Proposition 7 below, which provides a contradiction and terminates the proof of Theorem 4.

Proposition 7. *Eve has a winning strategy in \mathcal{G}_0 .*

Proof sketch. Let us define $\text{Start}_\tau := \{\epsilon, x_1, \dots, x_{2^N+2}\}$, the two first levels of τ , and $\text{Start}_\mathcal{G} := \{(x, q, \alpha) \in V \mid x \in \text{Start}_\tau\}$. Every play in \mathcal{G}_0 starts in $\text{Start}_\mathcal{G}$. Note that from any position of $\text{Start}_\mathcal{G}$, the same moves are available in \mathcal{G}_0 , \mathcal{G}_i and \mathcal{G}_j . In \mathcal{G}_0 , we let Eve follow σ_i as long as the game is in $\text{Start}_\mathcal{G}$. If the game remains in $\text{Start}_\mathcal{G}$ for ever (by jumping infinitely), the obtained play is an outcome of σ_i , which is winning for Eve in \mathcal{G}_i . Because positions have the same colour in all acceptance games, this play is also winning for Eve in \mathcal{G}_0 . Otherwise, the play exits $\text{Start}_\mathcal{G}$ by going down the tree, hence it reaches some position v_k^q . Because v_k^q has been reached by the winning strategy σ_i , it is a winning position for Eve in the perfect-information parity game \mathcal{G}_i . If $k \neq 2^N + 1$, by Point 1 of Lemma 2, *i.e.* $(\mathcal{G}_0, v_k^q) \Leftrightarrow (\mathcal{G}_i, v_k^q)$, Eve also has a winning strategy from v_k^q in \mathcal{G}_0 .

If $k = 2^N + 1$, because $\text{visit}_{\sigma_i}(y_{2^N+1}) = \text{visit}_{\sigma_j}(y_{2^N+1})$, σ_j also visits position $v_{2^N+1}^q$, and therefore $v_{2^N+1}^q$ is a winning position for Eve in \mathcal{G}_j . By Point 2 of Lemma 2, *i.e.* $\mathcal{G}_0, v_k^q \Leftrightarrow \mathcal{G}_j, v_k^q$, Eve also has a winning strategy from v_k^q in \mathcal{G}_0 . \square

6 Conclusion and perspectives

We have considered a general setting where transition systems are equipped with path relations. We have proposed natural extensions of MSO and the μ -calculus in this setting: $\text{MSO}^\curvearrowright$, which is MSO with path relations, and $\text{L}_\mu^\curvearrowright$, which is the jumping μ -calculus. We have studied the question of whether the bisimulation-invariant fragment of $\text{MSO}^\curvearrowright$ and $\text{L}_\mu^\curvearrowright$ have the same expressivity, like in [17]. In the case of recognizable relations, the whole picture collapses to the classic case (Theorem 1). However, for the synchronous perfect recall path relation (a regular binary relation that captures models of agency with time and knowledge), the answer is negative (Theorem 2).

Our results suggest that the adequate logic on transition systems with path relations may lie in between $\text{L}_\mu^\curvearrowright$ and $\text{MSO}^\curvearrowright$: on the one hand, the latter is undecidable for regular path relations as simple as the “Equal Level” relation [25]; on the other hand fundamental decidable properties, such as winning in two-player imperfect-information games with perfect recall, are not captured by the former (Theorem 4). Remark that this property is expressible in Alternating-time Temporal Logic with imperfect information, which is therefore not subsumed by the jumping mu-calculus, but whose model-checking is decidable for one agent with synchronous perfect recall.

In addition, both the epistemic mu-calculus and imperfect-information games are decidable for perfect-recall when the knowledge of the agents is hierarchically ordered [7, 21]. Recent results on games with imperfect information also signal other classes of models with several agents and perfect recall that can be handled [6]. So even for the case of several relations on paths, a logic that would encompass both the (jumping) mu-calculus and games with imperfect information may have good computational properties on interesting classes of systems.

References

1. R. Alur, P. Černý, and S. Chaudhuri. Model checking on trees with path equivalences. In *TACAS’07*, pages 664–678. Springer, 2007.
2. Krzysztof R Apt and Erich Grädel. *Lectures in game theory for computer scientists*. Cambridge University Press, 2011.
3. J. Berstel. *Transductions and context-free languages*, volume 4. Teubner Stuttgart, 1979.
4. D. Berwanger, K. Chatterjee, M. De Wulf, L. Doyen, and Thomas A. Henzinger. Strategy construction for parity games with imperfect information. *Inf. Comput.*, 208(10):1206–1220, 2010.
5. D. Berwanger and L. Kaiser. Information tracking in games on graphs. *Journal of Logic, Language and Information*, 19(4):395–412, 2010.
6. Dietmar Berwanger and Anup Basil Mathew. Games with recurring certainty. In *Proceedings of SR 2014*, pages 91–96, 2014.

7. R. Bozianu, C. Dima, and C. Enea. Model checking an epistemic mu-calculus with synchronous and perfect recall semantics. In *TARK'2013*, 2013.
8. D. Caucal. On infinite transition graphs having a decidable monadic theory. *Theor. Comput. Sci.*, 290(1):79–115, 2003.
9. K. Chatterjee and L. Doyen. The complexity of partial-observation parity games. In *LPAR 17*, pages 1–14. Springer, 2010.
10. B. Courcelle and J. Engelfriet. *Graph Structure and Monadic Second-Order Logic - A Language-Theoretic Approach*. Cambridge University Press, 2012.
11. C. C. Elgot and M. O. Rabin. Decidability and undecidability of extensions of second (first) order theory of (generalized) successor. *J. Symb. Log.*, 31(2):169–181, 1966.
12. R. Fagin, J. Halpern, Y. Moses, and M. Vardi. *Reasoning about knowledge*. The MIT Press, 2004.
13. E. Grädel. Model-checking games for logics of imperfect information. *Theor. Comput. Sci.*, 493:2–14, 2013.
14. E. Grädel, W. Thomas, and Th. Wilke. *Automata, Logics, and Infinite Games, volume 2500 of LNCS*. Springer Verlag, 2002.
15. J. Y. Halpern, R. van der Meyden, and M. Y. Vardi. Complete Axiomatizations for Reasoning about Knowledge and Time. *SIAM J. Comput.*, 33(3):674–703, 2004.
16. J. Y. Halpern and M. Y. Vardi. The complexity of reasoning about knowledge and time. 1. Lower bounds. *J. Comp. Sys. Sci.*, 38(1):195–237, 1989.
17. D. Janin and I. Walukiewicz. On the expressive completeness of the propositional mu-calculus with respect to monadic second order logic. In *Proceedings of CONCUR'96*, pages 263–277. Springer, 1996.
18. O. Matz, N. Schweikardt, and W. Thomas. The monadic quantifier alternation hierarchy over grids and graphs. *Inf. Comput.*, 179(2):356–383, 2002.
19. B. Maubert and S. Pinchinat. Jumping automata for uniform strategies. In *Proceedings of FSTTCS'13*, pages 287–298, 2013.
20. D. Park. Concurrency and automata on infinite sequences. In *Proc. 5th GI Conf. on Th. Comp. Sci., LNCS 104*, pages 167–183. Springer-Verlag, March 1981.
21. A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *POPL'89*, pages 179–190. ACM Press, 1989.
22. B. Puchala. Asynchronous omega-regular games with partial information. In *Proceedings of MFCS'2010*, pages 592–603, 2010.
23. F. Reiter. Distributed graph automata. *CoRR*, abs/1404.6503, 2014.
24. N. V. Shilov and N. O. Garanina. Combining knowledge and fixpoints. Technical Report Preprint n.98, <http://www.iis.nsk.su/files/preprints/098.pdf>, A.P. Ershov Institute of Informatics Systems, Novosibirsk, 2002.
25. W. Thomas. Infinite trees and automaton-definable relations over omega-words. *Theor. Comput. Sci.*, 103(1):143–159, 1992.
26. W. Thomas. Languages, automata, and logic. In G. Rozenberg and A. Salomaa, editors, *Handbook of Formal Languages*, volume 3, Beyond Words, pages 389–455. Springer Verlag, 1997.
27. R. van der Meyden and N. Shilov. Model checking knowledge and time in systems with perfect recall (extended abstract). In *Proceedings of FSTTCS'99*, volume 1738 of LNCS, pages 432–445, 1999.
28. M. Y. Vardi. Reasoning about the past with two-way automata. In *Proceedings of ICALP'98*, volume 1443, pages 628–641, 1998.
29. I. Walukiewicz. Monadic second-order logic on tree-like structures. *Theor. Comput. Sci.*, 275(1-2):311–346, 2002.