

On First-Order μ -Calculus over Situation Calculus Action Theories

Diego Calvanese

Free Univ. of Bozen-Bolzano, Italy
calvanese@inf.unibz.it

Giuseppe De Giacomo

Sapienza Univ. di Roma, Italy
degiacomo@dis.uniroma1.it

Marco Montali

Free Univ. of Bozen-Bolzano, Italy
lastname@inf.unibz.it

Fabio Patrizi

Abstract

In this paper we study verification of situation calculus action theories against first-order μ -calculus with quantification across situations. Specifically, we consider $\mu\mathcal{L}_a$ and $\mu\mathcal{L}_p$, the two variants of μ -calculus introduced in the literature for verification of data-aware processes. The former requires that quantification ranges over objects in the current active domain, while the latter additionally requires that objects assigned to variables persist across situations. Each of these two logics has a distinct corresponding notion of bisimulation. In spite of the differences we show that the two notions of bisimulation collapse for dynamic systems that are *generic*, which include all those systems specified through a situation calculus action theory. Then, by exploiting this result, we show that for *bounded* situation calculus action theories, $\mu\mathcal{L}_a$ and $\mu\mathcal{L}_p$ have exactly the same expressive power. Finally, we prove decidability of verification of $\mu\mathcal{L}_a$ properties over *bounded* action theories, using finite faithful abstractions. Differently from the $\mu\mathcal{L}_p$ case, these abstractions must depend on the number of quantified variables in the $\mu\mathcal{L}_a$ formula.

1 Introduction

In this paper we study verification of first-order μ -calculus with quantification across situations as a verification language for situation calculus action theories (McCarthy and Hayes 1969; Reiter 2001). Such theories can be seen as one of the most prominent examples in AI of data-aware processes, i.e., dynamic systems in which a rich (first-order) description of the current state is married with a description of how such state evolves through actions (Bhattacharya et al. 2007; Deutsch et al. 2009; Bagheri Hariri et al. 2013a).

After the seminal work by De Giacomo, Ternovskaia, and Reiter (1997) and especially by Claßen and Lake-meyer (2008), there has been an increasing interest in verification in the situation calculus, and recently many important results have been devised regarding sound, complete, and terminating verification, including (Belardinelli, Lomuscio, and Patrizi 2012; De Giacomo, Lesperance, and Patrizi 2012; Bagheri Hariri et al. 2013a; 2013b; Zariw and Claßen 2014; Belardinelli, Lomuscio, and Patrizi 2014; De Giacomo, Lesperance, and Patrizi 2016; De Giacomo et al. 2016).

These results are concerned with verification logics that are variants of those studied in the area of model checking of finite-state transition systems, like LTL, CTL, or modal μ -calculus, which subsumes all of them in the propositional setting (Clarke, Grumberg, and Peled 1999; Baier and Katoen 2008). Obviously, to use them in the context of formalisms with first-order state description, such logics need to be extended with the ability of querying the state in first-order logic. However in most proposals, e.g., (De Giacomo, Lesperance, and Patrizi 2012; De Giacomo et al. 2014), such ability is limited to the use of first-order sentences (closed formulas) instead of propositions, without the possibility of quantifying across states/situations. *Quantification across* (states/situations) refers to the possibility of using variables quantified in the current situation also in future situations. Without quantification across, these first-order temporal logics remain quite similar to their propositional variants (though with infinitely many propositions corresponding to first-order sentences, instead of the usual finite ones). In particular, notions like bisimulation and bisimulation invariance remain essentially those known for the propositional case. Only very few papers study verification logics with quantification across (Belardinelli, Lomuscio, and Patrizi 2012; Bagheri Hariri et al. 2013a; Belardinelli, Lomuscio, and Patrizi 2014; De Giacomo, Lesperance, and Patrizi 2016).

In this paper, we study in depth first-order μ -calculus with quantification across. In particular, we consider the two basic μ -calculus variants proposed in literature, $\mu\mathcal{L}_a$ and $\mu\mathcal{L}_p$, which are characterized by different restrictions on how quantification across is controlled. The logic $\mu\mathcal{L}_a$ requires quantification to range over objects in the active domain, i.e., in the extension of some fluent in the current situation (in situation calculus terms). The logic $\mu\mathcal{L}_a$ was studied by Bagheri Hariri et al. (2013a) in the context of data-centric dynamic systems (DCDS) and in its CTL fragment by Belardinelli, Lomuscio, and Patrizi (2012; 2014). The logic $\mu\mathcal{L}_p$ is a restriction of $\mu\mathcal{L}_a$ in which it is further required that the objects assigned to the quantified variables must persist across the states traversed while checking the formula. The logic $\mu\mathcal{L}_p$ was also studied by Bagheri Hariri et al. (2013a), and then in the context of situation calculus action theories (De Giacomo, Lesperance, and Patrizi 2016).

As shown by Bagheri Hariri et al. (2013a), these two logics can be characterized by two distinct notions of bisimulation over transition systems: history preserving bisimulation (or *a-bisimulation*) for $\mu\mathcal{L}_a$, and persistence preserving bisimulation (or *p-bisimulation*) for $\mu\mathcal{L}_p$. Specifically, $\mu\mathcal{L}_a$ is invariant wrt a-bisimulation while $\mu\mathcal{L}_p$ is invariant wrt p-bisimulation, where bisimulation invariance means that two bisimilar states satisfy the same formulas.

In addition, decidability results for verification have been devised. A crucial notion to get decidability of verification in dynamic system formalisms that allow for first-order state descriptions is that of *state-boundedness* (Belardinelli, Lomuscio, and Patrizi 2012; De Giacomo, Lesperance, and Patrizi 2012; Bagheri Hariri et al. 2013a). In particular, De Giacomo, Lesperance, and Patrizi (2012) show that verification of first-order μ -calculus without quantification across over *bounded action theories* in the situation calculus is decidable. Such theories have an infinite object domain, but the number of object tuples that belong to fluents in each situation remains bounded. Nonetheless, an agent may deal with an infinite number of objects over the course of an infinite execution.

These results are extended to deal with quantification across by De Giacomo, Lesperance, and Patrizi (2016), who show that models of bounded situation calculus action theories can be faithfully abstracted into p-bisimilar finite-state transition systems, thus getting decidability of verification for $\mu\mathcal{L}_p$. Also De Giacomo, Lesperance, and Patrizi (2016) for situation calculus, and Bagheri Hariri et al. (2013a) for DCDS, show that, differently from the $\mu\mathcal{L}_p$ case, in the $\mu\mathcal{L}_a$ case no faithful finite abstraction can exist that is independent from the formula to check. Interestingly, Belardinelli, Lomuscio, and Patrizi (2012; 2014) show that, for their *bounded* first-order defined transition systems, a faithful abstraction depending on the number of variables in the formula exists for the CTL fragment of $\mu\mathcal{L}_a$. However, it remained open till now whether their result extends to full $\mu\mathcal{L}_a$.

Here we investigate thoroughly the two logics $\mu\mathcal{L}_p$ and $\mu\mathcal{L}_a$ and the bisimulation notions associated to them. We establish quite surprising results wrt the expressive power of the two logics, and we establish decidability of verification for $\mu\mathcal{L}_a$ against bounded situation calculus action theories.

Specifically, we present the following results.

- For transition systems that are “generic”, such as those generated by logical theories, and in particular situation calculus theories, the notions of p-bisimilarity and a-bisimilarity collapse, as long as we keep the object domain infinite (a natural assumption in situation calculus).
- Moreover, for generic transition systems with the additional condition that the active domain of each state is finite (though not necessarily smaller than any given bound), $\mu\mathcal{L}_a$ and $\mu\mathcal{L}_p$ have exactly the same expressive power, in the sense that if there is a $\mu\mathcal{L}_a$ formula that is able to distinguish two states/situations, then there is one expressible in $\mu\mathcal{L}_p$, and viceversa.
- As a consequence of the equivalence between p-bisimilarity and a-bisimilarity, we get that if two generic transition systems with infinite object domains are p-

bisimilar, they satisfy the same $\mu\mathcal{L}_a$ formulas. Then we strengthen this result by showing that, if one of the transition systems has a finite object domain that is *large enough*, then it preserves all $\mu\mathcal{L}_a$ formulas that use only a predefined number of variables.

- We further show that, for bounded generic transition systems, and for a set of variables, it is always possible to define a faithful finite-state abstraction that preserves $\mu\mathcal{L}_a$ formulas whose variables belong to that set. This in particular applies to models of bounded situation calculus action theories.
- Finally, we show that given a bounded situation calculus action theory (including those with incomplete information), and a set of variables, we can effectively construct a new situation calculus action theory with finite domain that preserves $\mu\mathcal{L}_a$ formulas whose variables belong to that set. In this way, we get decidability of verification of $\mu\mathcal{L}_a$ formulas over bounded situation calculus action theories.

The impact of these results is quite strong also for another reason. Bagheri Hariri et al. (2013a) have shown that verification of first-order LTL with quantification across time points ranging on active domain is undecidable for trivial bounded-state transition systems. Then, using the folk assumption that μ -calculus can capture LTL also in the first-order case, e.g., (Okamoto 2010), it was concluded that $\mu\mathcal{L}_a$ verification is undecidable for bounded transition systems (hence, including models of bounded situation calculus action theories). Here, we show that this is not true, and that $\mu\mathcal{L}_a$ verification is indeed decidable over bounded situation calculus action theories. This has the consequence that *it is not true that first-order mu-calculus can capture first-order LTL in general!* In other words, once we allow for quantification across, the ability of LTL of talking about single traces cannot be mimicked anymore by μ -calculus. To the best of our knowledge this is the very first formal proof of this notable fact.

2 Situation Calculus and Boundedness

The *situation calculus* (McCarthy and Hayes 1969; Reiter 2001) is a logical language for representing and reasoning about dynamic worlds with three sorts: objects, actions, and situations. We assume to have countably infinitely many *object constants*, on which we adopt the unique name assumption (UNA). We assume to have a finite number of *action types*, each of which takes a tuple of objects as arguments. A *situation* term denotes a sequence of actions: the constant S_0 denotes the initial situation (no action has yet been done), whereas term $do(a, s)$ denotes the successor situation resulting from performing action a in situation s . We assume to have a finite set \mathcal{F} of *fluents*, i.e., predicates whose extension varies from situation to situation. Fluents take a situation term as their last argument (e.g., *Holding*(x, s)), while the other arguments are of sort object. We assume that there are no functions other than constants and no predicates other than fluents.

Within this language, one can formulate action theories to describe how the world changes as a result of actions. A well studied and popular type of such theories are *basic action*

theories (Reiter 2001). A basic action theory \mathcal{D} is a collection of first-order axioms (plus a second-order characterization of situation terms) conveniently specifying (in terms of size and computational properties): (i) actions' preconditions, characterized through *precondition axioms* involving a special predicate $Poss(a, s)$ capturing when action a is executable in situation s ; (ii) actions' effects and non-effects (i.e., solving the frame problem) by the so-called *successor state axioms*; and (iii) an *initial situation description*, expressed as a first-order logic theory where all fluents are instantiated to S_0 , capturing the world's initial state. We denote by C the set of constants explicitly mentioned in the initial situation description or in precondition or successor state axioms. Actually, for simplicity, w.l.o.g., we assume that all constants in C appear in the initial situation description. Notice that these are the constants we actually predicate on (while on the others we only predicate existence and name uniqueness).

An action theory \mathcal{D} is *bounded*, if for a given natural number B at every executable situation (i.e., reachable through a finite sequence of executable actions), the number of distinct object tuples occurring in the extension of each fluent of \mathcal{D} is at most B . Thus, the interpretation of a fluent at every situation does not use more than B distinct tuples, though these change from situation to situation and collectively are infinitely many (De Giacomo, Lesperance, and Patrizi 2012; 2016). For convenience, with a little abuse of notation, we say that an action theory is *bounded by b* when in each situation the number of objects occurring in the extension of all fluents is at most b . Notice that, when \mathcal{D} is bounded by b , then $B = |\mathcal{F}| \cdot b^k$, where k is the maximal arity of fluents.

A bookshelf is a prototypical example of boundedness.

Example 1 (Avid Reader). An agent is an avid reader and has a bookshelf of a given size. He acquires books, puts them in the bookshelf, reads them, and then puts them back in the bookshelf or gives them away. The available space in the bookshelf is given in units and each book consumes a certain number of units (e.g., one for simplicity). The reader cannot acquire a book if there is not enough space in the bookshelf.

The possible actions are the following:

- *acquire(book)*. Pre: *book* not already in the bookshelf, space available in the bookshelf. Post: *book* in the bookshelf and one less unit available in the bookshelf.
- *read(book)*. Pre: *book* in the bookshelf. Post: *book* in the hand of the avid reader, *book* not in the bookshelf.
- *store(book)*. Pre: *book* in the hand of the avid reader, space available in the bookshelf. Post: *book* in the bookshelf and one less unit available in the bookshelf.
- *discard(book)*. Pre: *book* in the hand of the avid reader. Post: *book* not in the hand of the avid reader and not in the bookshelf.

It is easy to write explicitly precondition and successor state axioms, which we omit for sake of brevity. It is also easy to see that the resulting action theory is indeed bounded. ■

3 Transition Systems

When focussing on verification of temporal properties we do not need to deal directly with full action theory models, since

both actions and situations (both of which do not appear explicitly in formulas to verify) can be essentially disregarded (De Giacomo, Lesperance, and Patrizi 2012; 2016). Specifically, we can focus on *transition systems*.

We denote by $Int_{\Delta}^{\mathcal{F}, C}$, the set of all possible interpretations of the situation-suppressed fluents in \mathcal{F} (i.e., fluents with the situation arguments suppressed) and of the constants in C , over the object domain Δ . A *transition system* (TS) (over the situation-suppressed fluents \mathcal{F} , constants C , and object domain Δ) is a tuple $T = \langle \Delta, Q, q_0, \rightarrow, \mathcal{I} \rangle$, where:

- Q is the *set of states*;
- $q_0 \in Q$ is the *initial state*;
- $\rightarrow \subseteq Q \times Q$ is the *transition relation*; and
- $\mathcal{I} : Q \mapsto Int_{\Delta}^{\mathcal{F}, C}$ is the *labeling function* associating to each state q an interpretation $\mathcal{I}(q) = \langle \Delta, \cdot^{\mathcal{I}(q)} \rangle$ such that the constants in C are interpreted in the same way in all the states over which \mathcal{I} is defined.

We denote by $adom(\mathcal{I}(q))$ the *active domain* of $\mathcal{I}(q)$, i.e., the set of objects occurring in the extension of some (situation-suppressed) fluent in q union the interpretation of constants in C , and by $\tilde{\mathcal{I}}(q)$ the restriction of $\mathcal{I}(q)$ to its active domain.

Among the various TSs, we are interested in those *induced* by models of the situation calculus action theory \mathcal{D} . Consider a model M of \mathcal{D} with object domain Δ^1 and situation domain \mathcal{S} . Given situation s , we can associate to s a first-order interpretation $\mathcal{I}_M(s) \doteq \langle \Delta, \cdot^{\mathcal{I}} \rangle$, where: (i) for every $c \in C$, $c^{\mathcal{I}} = c^M$ and (ii) for every (situation-suppressed) fluent F of \mathcal{D} , $F^{\mathcal{I}} = \{ \vec{d} \mid \langle \vec{d}, s \rangle \in F^M \}$. Then, we can define the TS *induced by M* as the labelled TS $T_M = \langle \Delta, Q, q_0, \mathcal{I}, \rightarrow \rangle$ such that:

- $Q = \mathcal{S}$ is the set of *possible states*, each corresponding to a distinct executable situation in \mathcal{S} ;
- $q_0 = S_0^M \in Q$ is the *initial state*, with S_0^M the initial situation of \mathcal{D} ;
- $\rightarrow \subseteq Q \times Q$ is the *transition relation* such that $q \rightarrow q'$ iff there exists some action a such that $\langle a, q \rangle \in Poss^M$ and $q' = do^M(a, q)$.
- $\mathcal{I} : Q \mapsto Int_{\Delta}^{\mathcal{F}, C}$ is the *labeling function* associating to each state (situation) q the interpretation $\mathcal{I}(q) = \mathcal{I}_M(q)$.

The TS induced by a model M is essentially the tree of executable situations, with each situation labelled by an interpretation of fluents (and constants), corresponding to the interpretation that M associates to that situation. Notice that transitions do not carry any information about the corresponding triggering action.

Generic TS. Next we introduce a key property for TS: *genericity* (Abiteboul, Hull, and Vianu 1995), also called *uniformity* by Belardinelli, Lomuscio, and Patrizi (2014).

We start by recalling the standard notions of *isomorphism* and *isomorphic interpretations*. Two first-order interpretations $\mathcal{I}_1 = \langle \Delta_1, \cdot^{\mathcal{I}_1} \rangle$ and $\mathcal{I}_2 = \langle \Delta_2, \cdot^{\mathcal{I}_2} \rangle$, over the same fluents \mathcal{F} and constants C , are said to be *isomorphic*, written $\mathcal{I}_1 \sim \mathcal{I}_2$, if there exists a bijection (called *isomorphism*)

¹Note that Δ is infinite, since we have assumed that the theories we consider include infinitely many constants with UNA.

$h : \Delta_1 \mapsto \Delta_2$ such that: (i) for every $F \in \mathcal{F}$, $\vec{x} \in F^{\mathcal{I}_1}$ if and only if $h(\vec{x}) \in F^{\mathcal{I}_2}$; (ii) for every $c \in C$, $c^{\mathcal{I}_2} = h(c^{\mathcal{I}_1})$. Intuitively, for two interpretations to be isomorphic, it is required that one can be obtained from the other by renaming the individuals in the interpretation domain. Notice that, necessarily, the interpretation domains of isomorphic interpretations have the same cardinality. When needed, to make it explicit that h is an isomorphism between \mathcal{I}_1 and \mathcal{I}_2 , we write $\mathcal{I}_1 \sim_h \mathcal{I}_2$.

Definition 1 (Generic Transition System). A TS $T = \langle \Delta, Q, q_0, \rightarrow, \mathcal{I} \rangle$ is said to be *generic* if: for every $q_1, q'_1, q_2 \in Q$ and every bijection $h : \Delta \mapsto \Delta$, if $\mathcal{I}(q_1) \sim_h \mathcal{I}(q_2)$ and $q_1 \rightarrow q'_1$, then there exists $q'_2 \in Q$ such that $q_2 \rightarrow q'_2$ and $\mathcal{I}(q'_1) \sim_h \mathcal{I}(q'_2)$.

Intuitively, genericity requires that if two states are isomorphic they induce the “same” transitions (modulo isomorphism). This property is actually always true if the next states are built by a first-order specification involving only the current state and the next one, as long as we do not use predefined domains with special properties that are specified extra-logically (e.g., we allow for natural numbers). In particular it holds for situation calculus specifications (and indeed virtually all first-order based formalisms for reasoning about actions used in AI) (Reiter 2001).

Theorem 2. For every model M of a situation calculus action theory \mathcal{D} , the generated TS T_M is generic.

Proof. By construction of T_M . \square

Bounded-state TS. Next we look at the counterpart of boundedness for action theories in TSs.

Definition 3 (Bounded-State Transition System). A TS $T = \langle \Delta, Q, q_0, \rightarrow, \mathcal{I} \rangle$ is said to be *bounded-state* if for a given b we have that $|\text{adom}(\mathcal{I}(q))| \leq b$ for every $q \in Q$.

That is, we say that T is bounded-state if there is a bound on the number of objects that can be accumulated in the same state. Notice that this does not disallow the possibility of accumulating infinitely many objects along an infinite run (or the entire TS for the matter).

As expected, bounded situation calculus action theories give rise to bounded-state TSs.

Theorem 4. For every model M of a situation calculus action theory \mathcal{D} bounded by b , the generated TS T_M is bounded-state, with each state bounded by b .

Proof. Follows directly from the definition of action theory bounded by b given in Section 2. \square

4 Verification Logics

As a verification logic to specify temporal properties, we focus on modal μ -calculus (Emerson 1996; Stirling 2001; Bradfield and Stirling 2007), one of the most powerful temporal logics for which model checking has been investigated. It is well-known that in the propositional setting μ -calculus is able to capture both linear time logics such as LTL and PSL (Property Specification Language²), and branching time

²<http://www.eda.org/ieee-1850/>

logics such as CTL and CTL* (Clarke, Grumberg, and Peled 1999; Baier and Katoen 2008). The main characteristic of modal μ -calculus is the ability of expressing directly least and greatest fixpoints of (predicate-transformer) operators formed using formulae relating the current state to the next one. By using such fixpoint constructs one can easily express sophisticated temporal properties defined by induction or co-induction.

In the following we consider two first-order variants of modal μ -calculus that have been considered in literature.

The Logic $\mu\mathcal{L}_a$. The first logic is characterized by the assumption that quantification over objects is restricted to those that are present in the current active domain, and was studied by Bagheri Hariri et al. (2013a) and by Belardinelli, Lomuscio, and Patrizi (2014)³. The syntax of $\mu\mathcal{L}_a$ is

$$\Phi ::= \varphi \mid \neg\Phi \mid \Phi_1 \wedge \Phi_2 \mid \exists x. \text{LIVE}(x) \wedge \Phi \mid \langle - \rangle \Phi \mid \mu Z. \Phi,$$

where φ is a first-order formula expressed using situation-suppressed fluents in \mathcal{F} and constants in C , the modal operator $\langle - \rangle \Phi$ denotes the existence of a transition from the current state to a next state where Φ holds, and $\mu Z. \Phi$ denotes the least fixpoint of the formula Φ seen as a predicate transformer wrt Z . We use $\nu Z. \Phi$ as the (standard) abbreviation for $\neg \mu Z. \neg \Phi[Z/\neg Z]$, to denote the greatest fixpoint of Φ . Note that in $\mu\mathcal{L}_a$ quantification across ranges over objects in the current active domain. That is, individuals over which quantification ranges must belong to the active domain of the current situation/state of the TS, as required by $\text{LIVE}(\cdot)$. As usual in μ -calculus, formulas of the form $\mu Z. \Phi$ (and $\nu Z. \Phi$) must obey to the *syntactic monotonicity* of Φ wrt Z , which states that every occurrence of the variable Z in Φ must be within the scope of an even number of negation symbols. This ensures that the semantics of $\mu Z. \Phi$ and $\nu Z. \Phi$ is well defined.

Example 2. The $\mu\mathcal{L}_a$ formula

$$\forall x. \text{Stud}(x) \supset \mu Y. ((\exists y. \text{Grad}(x, y)) \vee \langle - \rangle Y)$$

states that, for each student x in the current state/situation, there exists an evolution that eventually leads to the graduation of x (with some final mark y). \blacksquare

To interpret $\mu\mathcal{L}_a$ formulas over a TS $T = \langle \Delta, Q, q_0, \rightarrow, \mathcal{I} \rangle$, we use valuations (v, V) formed by an individual variable valuation v and a predicate variable valuation V parameterized by v , i.e., which maps each predicate variable Z to a subset $V(v, Z)$ of Q . We define the *extension function* $(\cdot)_{(v, V)}^T$, which maps $\mu\mathcal{L}_a$ formulas to subsets of Q , as shown in Figure 1.⁴

³Actually, Belardinelli, Lomuscio, and Patrizi (2014) consider the CTL fragment of $\mu\mathcal{L}_a$.

⁴By mentioning states/situations explicitly, it is also possible to define the least and greatest fixpoint operators directly in second-order logic as follows (De Giacomo, Ternovskaia, and Reiter 1997):

$$\begin{aligned} \mu Z. \Phi[s] &\equiv \forall \hat{Z}. (\forall \hat{s}. \Phi[\hat{s}] \supset Z(\hat{s})) \supset Z(s) \\ \nu Z. \Phi[s] &\equiv \exists \hat{Z}. (\forall \hat{s}. Z(\hat{s}) \supset \Phi[\hat{s}]) \wedge Z(s) \end{aligned}$$

Note that Φ may contain free individual and predicate variables, and indeed these remain free in $\mu Z. \Phi$ and $\nu Z. \Phi$.

$$\begin{aligned}
(\varphi)_{(v,V)}^T &= \{q \mid q \in Q \text{ and } \mathcal{I}(q), v \models \varphi\} \\
(\neg\Phi)_{(v,V)}^T &= Q \setminus (\Phi)_{(v,V)}^T \\
(\Phi_1 \wedge \Phi_2)_{(v,V)}^T &= (\Phi_1)_{(v,V)}^T \cap (\Phi_2)_{(v,V)}^T \\
(\exists x. \text{LIVE}(x) \wedge \Phi)_{(v,V)}^T &= \{q \mid \exists d \in \text{adom}(\mathcal{I}(q)). \\
&\quad q \in (\Phi)_{(v,V)[x/d]}^T\} \\
(\langle - \rangle \Phi)_{(v,V)}^T &= \{q \mid \exists q'. q \rightarrow q' \text{ and } q' \in (\Phi)_{(v,V)}^T\} \\
(Z)_{(v,V)}^T &= V(Z) \\
(\mu Z. \Phi)_{(v,V)}^T &= \bigcap \{\mathcal{E} \subseteq Q \mid (\Phi)_{(v,V)[Z/\mathcal{E}]}^T \subseteq \mathcal{E}\}
\end{aligned}$$

$(v, V)[x/d]$ stands for (v', V) where v' is as v except that $v'(x) = d$. Similarly $(v, V)[Z/\mathcal{E}]$ stands for (v, V') where V' is as V except that $V'(v, Z) = \mathcal{E}$.

Figure 1: Semantics of $\mu\mathcal{L}_a$

Given a $\mu\mathcal{L}_a$ formula Φ , we say that a TS T satisfies Φ at state q under v and V , written $T, q, (v, V) \models \Phi$, if $q \in (\Phi)_{(v,V)}^T$. When Φ is closed on predicate variables, we omit V , as irrelevant, and write $T, q, v \models \Phi$. If Φ is closed on both individual and predicate variables we simply write $T, q \models \Phi$. For closed formulas, we say that T satisfies Φ , written $T \models \Phi$, if $T, q_0 \models \Phi$.

History-preserving bisimulation. To $\mu\mathcal{L}_a$ corresponds the notion of *history-preserving bisimulation* (or *a-bisimulation*). Given a bijection $h : Q \mapsto Q'$, we denote with $\text{DOM}(h)$ the domain of h , i.e., the set of elements in Q for which h is defined, and with $\text{IMG}(h)$ the image of h , i.e., the set of elements q' such that $q' = h(q)$ for some $q \in Q$. A bijection h' extends h if $\text{DOM}(h) \subseteq \text{DOM}(h')$ and $h'(x) = h(x)$ for all $x \in \text{DOM}(h)$ (or equivalently $\text{IMG}(h) \subseteq \text{IMG}(h')$ and $h'^{-1}(y) = h^{-1}(y)$ for all $y \in \text{IMG}(h)$).

A history-preserving bisimulation relation can be defined as follows. Let $T_1 = \langle \Delta_1, Q_1, q_{10}, \rightarrow_1, \mathcal{I}_1 \rangle$ and $T_2 = \langle \Delta_2, Q_2, q_{20}, \rightarrow_2, \mathcal{I}_2 \rangle$ be two TSs (over the situation-suppressed fluents, and constants of an action theory \mathcal{D}), and let H be the set of all possible bijections $h : D_1 \mapsto D_2$, for $D_1 \subseteq \Delta_1$ and $D_2 \subseteq \Delta_2$. A relation $R \subseteq Q_1 \times H \times Q_2$ is a *history-preserving bisimulation* (or *a-bisimulation*) between T_1 and T_2 , if $\langle q_1, h, q_2 \rangle \in R$ implies that:

1. $\tilde{\mathcal{I}}_1(q_1) \sim_h \tilde{\mathcal{I}}_2(q_2)$, i.e., the interpretations $\mathcal{I}(q_1)$ and $\mathcal{I}(q_2)$, when restricted to their active domains, are isomorphic equivalent according to h ;
2. for each $q'_1 \in Q_1$, if $q_1 \rightarrow_1 q'_1$ then there exists $q'_2 \in Q_2$ such that:
 - (a) $q_2 \rightarrow_2 q'_2$, and
 - (b) there exists a bijection $h' : \text{DOM}(h) \cup \text{adom}(\mathcal{I}_1(q'_1)) \mapsto \text{IMG}(h) \cup \text{adom}(\mathcal{I}_2(q'_2))$ that is an extension of h and such that $\langle q'_1, h', q'_2 \rangle \in R$;
3. for each $q'_2 \in Q_2$, if $q_2 \rightarrow_2 q'_2$ then there exists $q'_1 \in Q_1$ such that:
 - (a) $q_1 \rightarrow_1 q'_1$, and

- (b) there exists a bijection $h' : \text{DOM}(h) \cup \text{adom}(\mathcal{I}_1(q'_1)) \mapsto \text{IMG}(h) \cup \text{adom}(\mathcal{I}_2(q'_2))$ that is an extension of h and such that $\langle q'_1, h', q'_2 \rangle \in R$.

We say that a state $q_1 \in Q_1$ is *history-preserving bisimilar* (or *a-bisimilar*) to $q_2 \in Q_2$, written $q_1 \approx^a q_2$, if there exists an a-bisimulation R between T_1 and T_2 such that $\langle q_1, h, q_2 \rangle \in R$, for some h ; when needed, we also write $q_1 \approx_h^a q_2$, to explicitly name h . Finally, T_1 is said to be *a-bisimilar* to T_2 , written $T_1 \approx^a T_2$, if $q_{10} \approx^a q_{20}$. It is immediate to see that bisimilarity between states and TSs, i.e., the (overloaded) relation \approx^a , is an equivalence relation.

Using the notion of a-bisimilarity, one can prove a suitable version of the classical *bisimulation invariance result* for the μ -calculus, which states that bisimilar TSs satisfy exactly the same μ -calculus formulas, see e.g., (Bradfield and Stirling 2007).

Theorem 5 (Bagheri Hariri et al. 2013a). *Consider two TSs $T_1 = \langle \Delta_1, Q_1, q_{10}, \rightarrow_1, \mathcal{I}_1 \rangle$ and $T_2 = \langle \Delta_2, Q_2, q_{20}, \rightarrow_2, \mathcal{I}_2 \rangle$ with Δ_1 and Δ_2 infinite. If $T_1 \approx^a T_2$, then for every $\mu\mathcal{L}_a$ closed formula Φ , $T_1 \models \Phi$ if and only if $T_2 \models \Phi$.*

The opposite direction of this theorem does not hold in general, but we will show next that it holds for *generic* TSs.

The Logic $\mu\mathcal{L}_p$. Next, we consider a restriction of $\mu\mathcal{L}_a$ called $\mu\mathcal{L}_p$, studied by Bagheri Hariri et al. (2013a) and by De Giacomo, Lesperance, and Patrizi (2016). The syntax of $\mu\mathcal{L}_p$ is

$$\begin{aligned}
\Phi ::= & \varphi \mid \neg\Phi \mid \Phi_1 \wedge \Phi_2 \mid \exists x. \text{LIVE}(x) \wedge \Phi \mid \\
& \text{LIVE}(\bar{x}) \wedge \langle - \rangle \Phi \mid \text{LIVE}(\bar{x}) \wedge [-]\Phi \mid Z \mid \mu Z. \Phi
\end{aligned}$$

Note that in $\mu\mathcal{L}_p$ quantification across ranges over objects in the current active domain that *persist in the extension of some fluents across situations*. This is obtained by forcing through $\text{LIVE}(\bar{x}) \wedge \langle - \rangle \Phi$ and $\text{LIVE}(\bar{x}) \wedge [-]\Phi$ that the variables occurring free in Φ^5 are assigned to objects that are in the active domain of the current situation/state.

Example 3. The following $\mu\mathcal{L}_p$ formula:

$$\forall x. \text{Stud}(x) \supset \mu Y. ((\exists y. \text{Grad}(x, y)) \vee \text{LIVE}(x) \wedge \langle - \rangle Y)$$

states that for each student x in the current state/situation, there exists an evolution where x *remains in the active domain*, and x eventually graduates (with some final mark y).

Persistence-preserving bisimulation. The bisimulation relation that captures $\mu\mathcal{L}_p$ can be defined as follows. Let $T_1 = \langle \Delta_1, Q_1, q_{10}, \rightarrow_1, \mathcal{I}_1 \rangle$ and $T_2 = \langle \Delta_2, Q_2, q_{20}, \rightarrow_2, \mathcal{I}_2 \rangle$ be two TSs over the situation-suppressed fluents \mathcal{F} and constants C , and let H be the set of all possible bijections $h : D_1 \mapsto D_2$, for $D_1 \subseteq \Delta_1$ and $D_2 \subseteq \Delta_2$. A relation $R \subseteq Q_1 \times H \times Q_2$ is a *persistence-preserving bisimulation* (or *p-bisimulation*) between T_1 and T_2 , if $\langle q_1, h, q_2 \rangle \in R$ implies that:

⁵With the proviso that second order variables are substituted by their corresponding fixpoint formula.

1. $\tilde{\mathcal{I}}_1(q_1) \sim_h \tilde{\mathcal{I}}_2(q_2)$;
2. for each $q'_1 \in Q_1$, if $q_1 \rightarrow_1 q'_1$ then there exists $q'_2 \in Q_2$ such that:
 - (a) $q_2 \rightarrow_2 q'_2$, and
 - (b) there exists a bijection $h' : \text{adom}(\mathcal{I}_1(q_1)) \cup \text{adom}(\mathcal{I}_1(q'_1)) \mapsto \text{adom}(\mathcal{I}_2(q_2)) \cup \text{adom}(\mathcal{I}_2(q'_2))$ such that its restriction $h'|_{\text{adom}(\mathcal{I}_1(q_1))}$ coincides with $h|_{\text{adom}(\mathcal{I}_1(q_1))}$ and $\langle q'_1, h'|_{\text{adom}(\mathcal{I}_1(q'_1))}, q'_2 \rangle \in R$;
3. for each $q'_2 \in Q_2$, if $q_2 \rightarrow_2 q'_2$ then there exists $q'_1 \in Q_1$ such that:
 - (a) $q_1 \rightarrow_1 q'_1$, and
 - (b) there exists a bijection $h' : \text{adom}(\mathcal{I}_1(q_1)) \cup \text{adom}(\mathcal{I}_1(q'_1)) \mapsto \text{adom}(\mathcal{I}_2(q_2)) \cup \text{adom}(\mathcal{I}_2(q'_2))$ such that its restriction $h'|_{\text{adom}(\mathcal{I}_1(q_1))}$ coincides with $h|_{\text{adom}(\mathcal{I}_1(q_1))}$ and $\langle q'_1, h'|_{\text{adom}(\mathcal{I}_1(q'_1))}, q'_2 \rangle \in R$.

We say that a state $q_1 \in Q_1$ is *persistence-preserving bisimilar* (or *p-bisimilar*) to $q_2 \in Q_2$, written $q_1 \approx^p q_2$, if there exists a p-bisimulation R between T_1 and T_2 such that $\langle q_1, h, q_2 \rangle \in R$, for some h ; when needed, we also write $q_1 \approx_h^p q_2$, to explicitly name h . Finally, a TS T_1 is said to be *p-bisimilar* to T_2 , written $T_1 \approx^p T_2$, if $q_{10} \approx^p q_{20}$. Again, p-bisimilarity is obviously an equivalence relation.

Theorem 6 (Bagheri Hariri et al. 2013a). *Consider two TSs $T_1 = \langle \Delta_1, Q_1, q_{10}, \rightarrow_1, \mathcal{I}_1 \rangle$ and $T_2 = \langle \Delta_2, Q_2, q_{20}, \rightarrow_2, \mathcal{I}_2 \rangle$ with Δ_1 and Δ_2 infinite. If $T_1 \approx^p T_2$, then for every $\mu\mathcal{L}_p$ closed formula Φ , $T_1 \models \Phi$ if and only if $T_2 \models \Phi$.*

Again the other direction of this theorem does not hold in general, but we show next that it holds for *generic* TSs.

For a bounded situation calculus action theory \mathcal{D} , we can construct finite TS faithful abstractions of the models of \mathcal{D} (De Giacomo, Lesperance, and Patrizi 2012; 2016).

Theorem 7 (De Giacomo, Lesperance, and Patrizi 2016). *Given a model M of a bounded situation calculus action theory, there exists a finite state TS T_M^f that is p-bisimilar to the TS T_M induced by M .*

Putting these two results together, we have that for every model M of \mathcal{D} there exists a finite TS T_M^f which is a faithful abstraction of T_M , i.e., such that for every $\mu\mathcal{L}_p$ closed formula Φ , $T_M \models \Phi$ if and only if $T_M^f \models \Phi$. Hence, we can use T_M^f to model check properties of interest over T_M .

Unfortunately, in the case of $\mu\mathcal{L}_a$, it is easy to see that no finite TS exists that is a faithful abstraction of T_M *independent from the $\mu\mathcal{L}_a$ formula* to verify. Indeed, assume to have an action that replaces an object in the active domain by one of the objects assigned to its parameters. Then, for every bound n on the number of objects in a candidate finite abstraction, we can write a (fixpoint-free) formula saying that there exists a finite run with more than n distinct objects:

$$\begin{aligned} \exists x_1. \text{LIVE}(x_1) \wedge \langle - \rangle (\exists x_2. \text{LIVE}(x_2) \wedge x_2 \neq x_1 \wedge \\ \langle - \rangle (\exists x_3. \text{LIVE}(x_3) \wedge x_3 \neq x_1 \wedge x_3 \neq x_2 \wedge \\ \dots \\ \langle - \rangle (\exists x_{n+1}. \text{LIVE}(x_{n+1}) \wedge x_{n+1} \neq x_1 \wedge \dots \wedge x_{n+1} \neq x_n))) \end{aligned}$$

This formula is false in the finite abstraction, while true in the original TS, where objects are not “reused” (De Giacomo, Lesperance, and Patrizi 2016).

5 Expressiveness

A notable property of generic TSs is that if the interpretations associated to two states of the same TS are isomorphic wrt the active domain then they are p-bisimilar.

Lemma 8. *If $T = \langle \Delta, Q, q_0, \rightarrow, \mathcal{I} \rangle$ is a generic TS, then for every two states $q, q' \in Q$ and every bijection $h : D \mapsto D$ with $D \subseteq \Delta$ such that $\tilde{\mathcal{I}}_1(q_1) \sim_h \tilde{\mathcal{I}}_2(q_2)$ we have $q \approx_h^p q'$.*

Proof (sketch). By co-induction, we show that the relation $R = \{\langle q_1, h, q_2 \rangle \mid \tilde{\mathcal{I}}_1(q_1) \sim_h \tilde{\mathcal{I}}_2(q_2)\}$ is a p-bisimulation, by exploiting the very definition of generic TS. \square

Observe that for Lemma 8 to hold, the two states need to belong to the *same* generic TS. If the two states belong to different TSs, then we cannot exploit genericity (which relates states of the same TS) and the claim would not hold. Observe also that the opposite direction of Lemma 8 trivially holds, as a consequence of the definition of p-bisimilarity.

Exploiting Lemma 8, we prove the key result of this section: *on generic TSs, p-bisimilarity implies a-bisimilarity.*

Theorem 9. *Consider two generic TSs $T_1 = \langle \Delta_1, Q_1, q_{10}, \rightarrow_1, \mathcal{I}_1 \rangle$ and $T_2 = \langle \Delta_2, Q_2, q_{20}, \rightarrow_2, \mathcal{I}_2 \rangle$ with Δ_1 and Δ_2 infinite. Then, $T_1 \approx^p T_2$ if and only if $T_1 \approx^a T_2$.*

Proof (sketch). The “if” direction is immediate and holds also for TS that are not generic, since \approx^a is stricter than \approx^p .

For the “only-if” direction, we show by co-induction that the relation $R = \{\langle q_1, h, q_2 \rangle \mid q_1 \approx_h^p q_2\}$ is an a-bisimulation.

For the first condition, if $\langle q_1, h, q_2 \rangle \in R$ then, since $q_1 \approx_h^p q_2$, we have that $\tilde{\mathcal{I}}_1(q_1) \sim_h \tilde{\mathcal{I}}_2(q_2)$, i.e., R is closed wrt the first condition of a-bisimulation.

For the second condition consider that, by the second condition of p-bisimulation, we have that for each $q'_1 \in Q_1$, if $q_1 \rightarrow_1 q'_1$ then there exists $q'_2 \in Q_2$ such that: $q_2 \rightarrow_2 q'_2$ and there exists a bijection $h' : \text{adom}(\mathcal{I}_1(q_1)) \cup \text{adom}(\mathcal{I}_1(q'_1)) \mapsto \text{adom}(\mathcal{I}_2(q_2)) \cup \text{adom}(\mathcal{I}_2(q'_2))$ such that its restriction $h'|_{\text{adom}(\mathcal{I}_1(q_1))}$ coincides with h and $\langle q'_1, h'|_{\text{adom}(\mathcal{I}_1(q'_1))}, q'_2 \rangle \in R$.

Now let us consider h'' obtained by *extending* h , as required by a-bisimulation (not only its restriction to $\text{adom}(\mathcal{I}_1(q_1))$, as required by p-bisimulation) such that for all objects d in $\text{adom}(\mathcal{I}_1(q'_1))$ but not in $\text{DOM}(h)$ we have $h''(d) = h'(d)$. Consider $g = h'^{-1} \circ h''$. We have that g can be extended to cover the whole Δ_2 so that $\mathcal{I}(q_2) \sim_g \mathcal{I}(q'_2)$; indeed observe that h'' , h' , and h are identical over the active domain of q_2 , hence, wrt $\mathcal{I}(q_2)$, g is only renaming objects outside the active domain. By genericity, since we have $q_2 \rightarrow_2 q'_2$, there exists a state q''_2 such that $q_2 \rightarrow_2 q''_2$ and $\mathcal{I}(q'_2) \sim_g \mathcal{I}(q''_2)$. By Lemma 8, this implies that $q'_2 \approx_g^p q''_2$. On the other hand $\langle q'_1, h'|_{\text{adom}(\mathcal{I}_1(q'_1))}, q'_2 \rangle \in R$ implies $q'_1 \approx_{h'}^p q'_2$, hence, considering that $h' \circ g = h''$, by composing the two p-bisimulations we have that $q'_1 \approx_{h''}^p q''_2$, i.e., $\langle q'_1, h'', q''_2 \rangle \in R$. Hence we can conclude that R is closed under the second condition of a-bisimulation.

The proof for the third condition of a-bisimulation is analogous. The claim follows. \square

As an immediate consequence we have that if two TSs are p-bisimilar, being also a-bisimilar, by invariance wrt to

a-bisimilarity (Theorem 5), they satisfy the same closed $\mu\mathcal{L}_a$ formulas.

Theorem 10. *Consider two generic TSs $T_1 = \langle \Delta_1, Q_1, q_{10}, \rightarrow_1, \mathcal{I}_1 \rangle$ and $T_2 = \langle \Delta_2, Q_2, q_{20}, \rightarrow_2, \mathcal{I}_2 \rangle$ with Δ_1 and Δ_2 infinite. If $T_1 \approx^p T_2$ then for every $\mu\mathcal{L}_a$ closed formula Φ , $T_1 \models \Phi$ if and only if $T_2 \models \Phi$.*

Next, we study the converse of bisimulation invariance, i.e., of Theorems 5 and 6. In other words, we are interested in understanding for which TSs we have that if two states satisfy exactly the same $\mu\mathcal{L}_p$ (resp., $\mu\mathcal{L}_a$) formulas they are p-bisimilar (resp., a-bisimilar). For that purpose we introduce *generic finite-active-domain TSs*, which are generic TSs with the additional condition that the active domain of every state is finite (though not necessarily bounded by some given b). Obviously, such class of TSs includes generic bounded-state TSs, but also TSs obtained by starting from a database and updating it at each step with a finite number of tuples, as e.g., DCDSs (Bagheri Hariri et al. 2013a).

Theorem 11. *Consider two generic finite-active-domain TSs $T_1 = \langle \Delta_1, Q_1, q_{10}, \rightarrow_1, \mathcal{I}_1 \rangle$ and $T_2 = \langle \Delta_2, Q_2, q_{20}, \rightarrow_2, \mathcal{I}_2 \rangle$ with Δ_1 and Δ_2 infinite. If for every $\mu\mathcal{L}_p$ closed formula Φ , $T_1 \models \Phi$ if and only if $T_2 \models \Phi$ then $T_1 \approx^p T_2$.*

Proof (sketch). We show by co-induction that the relation $R = \{ \langle q_1, h, q_2 \rangle \mid \text{for all } \Phi \in \mu\mathcal{L}_p. T_1, q_1 \models \Phi \text{ iff } T_2, q_2 \models \Phi \text{ and } \mathcal{I}_1(q_1) \sim_h \mathcal{I}_2(q_2) \}$ is a p-bisimulation. R satisfies the first condition of p-bisimulation by definition. Suppose towards contradiction that it does not satisfy the second condition: i.e., there is a tuple $\langle q_1, h, q_2 \rangle$ and q'_1 such that $q_1 \rightarrow_1 q'_1$ but there is no extension h' of h and no q'_2 such that $q_2 \rightarrow_2 q'_2$ and $h'|_{\text{adom}(\mathcal{I}_1(q_1))}$ coincides with $h|_{\text{adom}(\mathcal{I}_1(q_1))}$, $\text{adom}(\mathcal{I}_1(q'_1)) \sim_{h'} \text{adom}(\mathcal{I}_2(q'_2))$, and q'_1 and q'_2 satisfy the same closed $\mu\mathcal{L}_p$ formulas.

Consider the *isomorphism type* of $\mathcal{I}_1(q'_1)$, i.e., the set of interpretations that are isomorphic to $\mathcal{I}_1(q'_1)$. Since $\text{adom}(\mathcal{I}_1(q'_1))$ is finite, there exists a first-order formula with one existentially quantified variable for each object in the active domain that characterizes the isomorphism type (De Giacomo, Lesperance, and Patrizi 2012; 2016), which we call *characteristic formula*. In fact, from such formula we can construct a first-order formula $\Psi(\vec{x})$, which leaves open the variables \vec{x} corresponding to objects already occurring in $\text{adom}(\mathcal{I}_1(q_1))$. In this way, in $\Psi(\vec{x})$, we are parameterizing the characteristic formula on \vec{x} , forcing the objects coming from $\text{adom}(\mathcal{I}_1(q_1))$ to persist.

Furthermore suppose that for each q'_2 there is a closed $\mu\mathcal{L}_p$ formula that is true in q'_1 but false in q'_2 . Notice that all q'_2 belonging to the isomorphism type corresponding to $\Psi(\vec{x})$ are p-bisimilar by genericity (Lemma 8) and hence by p-bisimulation invariance (Theorem 6) satisfy the same $\mu\mathcal{L}_p$ formulas. Hence if such a formula exists it is the same for all such states. Let's denote it by Φ . Then $T_1, q_1 \models \exists \vec{x}. \text{LIVE}(\vec{x}) \wedge \langle - \rangle (\Psi(\vec{x}) \wedge \Phi)$ and $T_2, q_2 \models \forall \vec{x}. \text{LIVE}(\vec{x}) \supset \langle - \rangle (\Psi(\vec{x}) \supset \neg \Phi)$, thus q_1 and q_2 do not satisfy the same $\mu\mathcal{L}_p$ formulas, and we get a contradiction. The third condition can be proven analogously. \square

By considering that $\mu\mathcal{L}_p$ is a subset of $\mu\mathcal{L}_a$, as an immediate consequence of Theorems 11 and 9, we get the analogous

result for $\mu\mathcal{L}_a$.

Theorem 12. *Consider two generic finite-active-domain TSs $T_1 = \langle \Delta_1, Q_1, q_{10}, \rightarrow_1, \mathcal{I}_1 \rangle$ and $T_2 = \langle \Delta_2, Q_2, q_{20}, \rightarrow_2, \mathcal{I}_2 \rangle$ with Δ_1 and Δ_2 infinite. If for every $\mu\mathcal{L}_a$ closed formula Φ , $T_1 \models \Phi$ if and only if $T_2 \models \Phi$ then $T_1 \approx^a T_2$.*

Proof (sketch). The proof exploits the fact that $\mu\mathcal{L}_a$ extends $\mu\mathcal{L}_p$, and that equivalence wrt to $\mu\mathcal{L}_p$ formulas guarantees p-bisimilarity, which in turn implies a-bisimilarity for generic finite-active-domain TSs. \square

Summarizing, given a state q_1 of TS T_1 and a state q_2 of TS T_2 , we have:

- always:

$q_1 \approx^a q_2$	implies	$q_1 \approx^p q_2$
$q_1 \approx^p q_2$	implies	$q_1 \approx^{\mu\mathcal{L}_p} q_2$
$q_1 \approx^a q_2$	implies	$q_1 \approx^{\mu\mathcal{L}_a} q_2$
$q_1 \approx^{\mu\mathcal{L}_a} q_2$	implies	$q_1 \approx^{\mu\mathcal{L}_p} q_2$
- when T_1 and T_2 are generic:

$q_1 \approx^p q_2$	implies	$q_1 \approx^a q_2$
$q_1 \approx^p q_2$	implies	$q_1 \approx^{\mu\mathcal{L}_a} q_2$
- when T_1 and T_2 are generic finite-active-domain:

$q_1 \approx^{\mu\mathcal{L}_p} q_2$	implies	$q_1 \approx^p q_2$
$q_1 \approx^{\mu\mathcal{L}_a} q_2$	implies	$q_1 \approx^a q_2$
$q_1 \approx^{\mu\mathcal{L}_p} q_2$	implies	$q_1 \approx^{\mu\mathcal{L}_a} q_2$

where $q_1 \approx^{\mu\mathcal{L}_p} q_2$ denotes that q_1 and q_2 satisfy the same $\mu\mathcal{L}_p$ formulas (and similarly for $\mu\mathcal{L}_a$).

6 Decidability

In this section, we study verification of $\mu\mathcal{L}_a$ formulas over bounded-state generic TSs and over bounded situation calculus action theories. In particular, as De Giacomo, Lesperance, and Patrizi (2012; 2016), Bagheri Hariri et al. (2013a), Bellardinelli, Lomuscio, and Patrizi (2014), we aim at getting decidability of verification by abstracting infinite TSs into finite-state ones. The general idea is to take advantage of what is shown in the previous section, namely that $\mu\mathcal{L}_a$ is invariant wrt p-bisimulation, i.e., if two generic TSs are p-bisimilar (and hence also a-bisimilar by Theorem 9) they satisfy the same $\mu\mathcal{L}_a$ formulas. This appears to allow us to focus on checking whether there exists a finite generic TS that is p-bisimilar to the one generated by the bounded action theory model of interest. However, the results in the previous section assume infinite object domains, and this, together with genericity, prevents one from building a *finite generic TS* by the very definition of genericity (if there exists a transition, then all, infinitely many, isomorphic transitions must exist, each producing a different successor state). To overcome this we need a stronger version of the invariance of $\mu\mathcal{L}_a$ wrt p-bisimulation, which also takes into account that we cannot have a finite abstraction that preserves $\mu\mathcal{L}_a$ and is independent from the formula to check, as discussed at the end of Section 4. The next result establishes such a stronger version of invariance.

In the statements below, we assume that T_1 and T_2 are over the same set of fluents \mathcal{F} and set of (explicitly mentioned) constants C of a given theory \mathcal{D} .

Theorem 13. Consider a finite set $Vars$ of variables and two generic TSs, $T_1 = \langle \Delta_1, Q_1, q_{10}, \rightarrow_1, \mathcal{I}_1 \rangle$, bounded by b and with infinite Δ_1 , and $T_2 = \langle \Delta_2, Q_2, q_{20}, \rightarrow_2, \mathcal{I}_2 \rangle$, with $|\Delta_2| \geq 2b + |Vars|$, such that $T_1 \approx^p T_2$. Then, for every closed $\mu\mathcal{L}_a$ formula Φ with variables renamed apart and belonging to $Vars$, we have $T_1 \models \Phi$ if and only if $T_2 \models \Phi$.

To prove Theorem 13, we first establish the claim for the simpler logic \mathcal{L}_a , which is $\mu\mathcal{L}_a$ without fixpoint constructs. We then generalize it to the infinitary version of \mathcal{L}_a , which captures $\mu\mathcal{L}_a$, by using a well-known line of reasoning in μ -calculus, see (van Benthem 1983; Bradfield and Stirling 2007) or Lemma 2 in (De Giacomo, Lesperance, and Patrizi 2016). We omit the details of this latter part for brevity, and focus instead on the first part by proving the following stronger version of the theorem, though restricted to \mathcal{L}_a formulas only.

We call $free(\Phi)$ the set of free first-order variables of Φ . Obviously, for closed formulas, $free(\Phi)$ is empty.

Lemma 14. Under the same hypothesis of Theorem 13, let Φ be an open \mathcal{L}_a formula with variables renamed apart and belonging to $Vars$. Consider two states $q_1 \in Q_1$, $q_2 \in Q_2$ s.t., for some h , $q_1 \approx_h^p q_2$, and two individual variable valuations v_1, v_2 , mapping variables in $Vars$ to Δ_1 and Δ_2 , respectively. If there exists a bijection \hat{h} between $adom(\mathcal{I}_1(q_1)) \cup \text{IMG}(v_1|_{free(\Phi)})$ and $adom(\mathcal{I}_2(q_2)) \cup \text{IMG}(v_2|_{free(\Phi)})$, whose restriction $\hat{h}|_{adom(\mathcal{I}_1(q_1))}$ coincides with h and s.t., for every individual variable $x \in free(\Phi)$, $\hat{h}(v_1(x)) = v_2(x)$, then $T_1, q_1, v_1 \models \Phi$ if and only if $T_2, q_2, v_2 \models \Phi$.

Proof (sketch). By induction on the structure of Φ .

For $\Phi = \varphi$ first-order, we use the following generalization of Theorem 5.6.3 by Libkin (2007), whose proof is immediately obtained by inspection of the original theorem's proof: for every first-order formula φ and interpretation $\mathcal{I} = \langle \Delta, \cdot^{\mathcal{I}} \rangle$ s.t. $|\Delta| \geq |adom(\mathcal{I})| + |vars(\varphi)|$, φ can effectively be rewritten as a formula φ' , which we call the *domain-independent version* of φ , with quantified variables ranging only over the active domain, s.t. for every valuation v , we have that $\mathcal{I}, v \models \varphi$ iff $\tilde{\mathcal{I}}^v, v \models \varphi'$, for $\tilde{\mathcal{I}}^v$ the restriction of \mathcal{I} to its active domain union the images, through v , of the variables $free(\varphi)$. Considering the invariance of first-order formulas wrt to isomorphic interpretations, that Δ_1 is infinite, and that $|\Delta_2| \geq 2b + |Vars|$, it follows that, $T_1(q_1), v_1 \models \varphi$ iff $\tilde{\mathcal{I}}_1^{v_1}(q_1), v_1 \models \varphi'$ iff $\tilde{\mathcal{I}}_2^{v_2}(q_2), v_2 \models \varphi'$ iff $T_2(q_2), v_2 \models \varphi$, for φ' the domain-independent version of φ .

The case of boolean connectives is straightforward.

For $\Phi = \exists y. \text{LIVE}(y) \wedge \Phi'$, suppose first that $T_1, q_1, v_1 \models \Phi$. Then there exists an object $d_1 \in adom(\mathcal{I}_1(q_1))$ s.t. $T_1, q_1, v_1[y/d_1] \models \Phi'$. Notice that d_1 is mapped, through h and thus \hat{h} , to some $d_2 \in adom(\mathcal{I}_2(q_2))$. The claim follows by inductive hypothesis, using $v_1[y/d_1]$, $v_2[y/d_2]$, and the bijection \hat{h} . The other direction is proven analogously.

For $\Phi = \langle - \rangle \Phi'$, suppose that $T_1, q_1, v_1 \models \langle - \rangle \Phi'$. Then, there exists a transition $q_1 \rightarrow_1 q'_1$ such that $T_1, q'_1, v_1 \models \Phi'$. Since $q_1 \approx_h^p q_2$, there exist a transition $q_2 \rightarrow_2 q'_2$, and a bijection $h' : adom(\mathcal{I}_1(q_1)) \cup adom(\mathcal{I}_1(q'_1)) \mapsto$

$adom(\mathcal{I}_2(q_2)) \cup adom(\mathcal{I}_2(q'_2))$ s.t. $h'|_{adom(\mathcal{I}_1(q_1))}$ coincides with $h|_{adom(\mathcal{I}_1(q_1))}$, and $q'_1 \approx_{h'}^p|_{adom(\mathcal{I}_1(q'_1))} q'_2$.

We would like apply the induction hypothesis using Φ' , $q'_1, q'_2, h'|_{adom(\mathcal{I}_1(q'_1))}, v_1, v_2$, and a suitable bijection \hat{h}' that extends $h'|_{adom(\mathcal{I}_1(q'_1))}$. Unfortunately, for q'_1 and q'_2 , an \hat{h}' satisfying the conditions of the lemma may not exist, in general. However, we can show that there exist another state $q''_2 \in Q$ bisimilar to q'_1 and s.t. $q_2 \rightarrow_2 q''_2$, and a bijection \hat{h}' s.t. the inductive hypothesis applies to $\Phi', q'_1, q''_2, \hat{h}', v_1, v_2$. This, by induction, implies that $T_2, q''_2, v_2 \models \Phi'$, thus that $T_2, q_2, v_2 \models \Phi$. Actually, to show the existence of q''_2 and \hat{h}' , we need to exploit the genericity of the two TSs and the cardinality constraints on Δ_2 . In particular, the requirement that $|\Delta_2| \geq 2b + |Vars|$ guarantees that, even in the case where $adom(\mathcal{I}(q_1))$ and $adom(\mathcal{I}(q'_1))$ have no objects in common and v_1 maps all variables into objects that are neither in $adom(\mathcal{I}(q_1))$ nor in $adom(\mathcal{I}(q'_1))$, Δ_2 contains enough objects for \hat{h}' to establish a bijection. The other direction is proven in a similar way. \square

We also show constructively that every bounded-state and generic TS can be abstracted into a p-bisimilar finite-state generic TS with a (finite) object domain of a suitable size.

Theorem 15. Consider a transition system $T_1 = \langle \Delta_1, Q_1, q_{10}, \rightarrow_1, \mathcal{I}_1 \rangle$ that is generic, bounded by b , and with infinite Δ_1 . Then, for every $k \geq 0$, there exists a finite-state generic TS $T_2 = \langle \Delta_2, Q_2, q_{20}, \rightarrow_2, \mathcal{I}_2 \rangle$, with $|\Delta_2| = 2b + k$ such that $T_1 \approx^p T_2$.

Proof (sketch). T_2 is defined as follows. The object domain Δ_2 is a subset of Δ_1 s.t. $|\Delta_2| = 2b + k$ and $adom(\mathcal{I}_1(q_{10})) \subseteq \Delta_2$ (notice that $|adom(\mathcal{I}_1(q_{10}))| \leq b$). The set of states is $Q_2 = \text{Int}_{\Delta_2}^{\mathcal{F}, C}$, which is the (finite) set of interpretations of \mathcal{F} and C over Δ_2 . The initial state q_{20} is the interpretation s.t. $\tilde{q}_{20} = \tilde{\mathcal{I}}_1(q_{10})$. The transition relation \rightarrow_2 is s.t. $q_2 \rightarrow_2 q'_2$ iff there exist two states $q_1, q'_1 \in Q_1$ s.t. $q_1 \rightarrow_1 q'_1$, $\tilde{\mathcal{I}}_1(q_1) \sim_h \tilde{q}_2$, and $\tilde{\mathcal{I}}_1(q'_1) \sim_h \tilde{q}'_2$, for some isomorphism h (notice that here genericity comes into play).⁶ Finally, \mathcal{I}_2 is the identity function. Obviously, T_2 is finite. Moreover, it can be shown that T_2 is generic and that for each state $q_1 \in Q_1$ and every state q_2 s.t. $\tilde{\mathcal{I}}_1(q_1) \sim \tilde{q}_2$, including the initial states q_{10} and q_{20} , we have that $q_1 \approx^p q_2$. \square

As a direct consequence of Theorems 13 and 15, we obtain:

Theorem 16. Given a finite set $Vars$ of variables and a generic TS $T_1 = \langle \Delta_1, Q_1, q_{10}, \rightarrow_1, \mathcal{I}_1 \rangle$, bounded by b and with infinite Δ_1 , there exists a TS $T_2 = \langle \Delta_2, Q_2, q_{20}, \rightarrow_2, \mathcal{I}_2 \rangle$, with $|\Delta_2| \geq 2b + |Vars|$, such that $T_1 \approx^p T_2$ and hence such that, for every closed $\mu\mathcal{L}_a$ formula Φ with variables renamed apart and belonging to $Vars$, we have that $T_1 \models \Phi$ if and only if $T_2 \models \Phi$.

We observe that the finite TS T_2 in Theorem 16 is effectively computable (as in the proof of Theorem 15) in the case where the interpretation $adom(\mathcal{I}_1(q_{10}))$ of the initial state of T_1 restricted to the active domain is known, and one can effectively check whether there exist two states q_1 and q'_1 such

⁶Actually, Q_2 can be restricted to the set of states in $\text{Int}_{\Delta_2}^{\mathcal{F}, C}$ reachable through \rightarrow_2 .

that $\tilde{\mathcal{I}}_1(q_1) = \tilde{q}_2$, $\tilde{\mathcal{I}}_1(q'_1) = \tilde{q}'_2$, and $q_1 \rightarrow_2 q'_1$. When T_2 can be effectively computed, Theorem 16 shows decidability of verification of $\mu\mathcal{L}_a$ formulas. This is the case, e.g., for TSs induced by models of bounded action theories.

However for such theories we can show a stronger result.

Theorem 17. *Let \mathcal{D} be a situation calculus action theory bounded by b , \mathcal{D}_{uno} the part of \mathcal{D} stating the existence, with unique name assumption, of infinitely many constants, and n the maximum among the number of variables occurring in the precondition and successor state axioms of \mathcal{D} . Let C' be a finite set of constants such that $C \subseteq C'$ and $|C'| \geq b + n$. Define the theory $\mathcal{D}' = (\mathcal{D} \setminus \mathcal{D}_{\text{uno}}) \cup \mathcal{D}'_{\text{uno}} \cup \mathcal{D}'_{\text{dc}}$, where:*

$$\mathcal{D}'_{\text{uno}} = \{\bigwedge_{c,c' \in C', c,c' \text{ distinct}} c \neq c'\}, \quad \mathcal{D}'_{\text{dc}} = \{\forall x. \bigvee_{c \in C'} x = c\}.$$

Then, for every model M of \mathcal{D} , there is a model M' of \mathcal{D}' , such that $T_M \approx^p T_{M'}$. Similarly, for every model M' of \mathcal{D}' there is a model M of \mathcal{D} , such that $T_M \approx^p T_{M'}$.

Proof (sketch). Let M be a model of \mathcal{D} with (infinite) domain Δ . The model M' can be obtained by fixing a (finite) domain $\Delta' \subset \Delta$, with cardinality $|C'|$, that includes the interpretation of the constants in C , and taking the interpretation of the initial situation so that $\tilde{\mathcal{I}}_M(S_0) = \tilde{\mathcal{I}}_{M'}(S_0)$. Observe that once the interpretation of the initial situation is fixed, M' is fully determined by \mathcal{D}' . Then, consider $T_M = \langle \Delta, Q, q_0, \mathcal{I}, \rightarrow \rangle$ and $T_{M'} = \langle \Delta', Q', q'_0, \mathcal{I}', \rightarrow' \rangle$, and build the relation $R = \{\langle q, h, q' \rangle \mid \tilde{\mathcal{I}}(q) \sim_h \tilde{\mathcal{I}}'(q')\}$. The proof consists in showing that R is a p-bisimulation s.t. $\langle q_0, h_0, q'_0 \rangle \in R$, for h_0 the identity on $\text{adom}(\mathcal{I}(q_0))$. This exploits the fact that T_M and $T_{M'}$ are constructed by successive action executions. The claim follows by observing that under the same actions (up to object renaming), successors of states with isomorphic labelings have isomorphic labelings. The other direction can be proved analogously. \square

As natural, TSs induced by models of finite-state action theories can be made finite.

Theorem 18. *Let \mathcal{D}' be a (bounded) situation calculus action theory defined as in Theorem 17, for some finite C' . Then, for every model M' of \mathcal{D}' with (finite) object domain Δ' , the corresponding induced TS $T_{M'}$ is p-bisimilar to a TS T_F that is generic, finite-state, and effectively computable from \mathcal{D}' , $\tilde{\mathcal{I}}_{M'}(S_0)$, and Δ' .*

Proof (sketch). We prove the result by providing an algorithm to compute $T_F = \langle \Delta_F, Q_F, q_{F0}, \rightarrow_F, \mathcal{I}_F \rangle$. We set $\Delta_F = \Delta'$, and \mathcal{I}_F as the identity function, and we initialize $q_{F0} = \mathcal{I}_{M'}(S_0)$, $Q_F = \{q_{F0}\}$, and $\rightarrow_F = \emptyset$. Then, starting with $q = q_{F0}$, we consider all actions a that, in M' , are executable in those situations s s.t. $\mathcal{I}_{M'}(s) = q$. This requires evaluating only the (situation-suppressed) precondition axiom of a against $\mathcal{I}(q)$. Notice that since Δ_F is finite, there are only finitely many actions. For every a , we then compute the interpretation of situation $s' = do^{M'}(a, s)$, for s as above. To this end, it is enough to evaluate the (situation-suppressed) right-hand side of each successor-state axiom against q (i.e., $\mathcal{I}_{M'}(s)$, for s as above), with the action assigned to a , thus producing a new interpretation $q' = \mathcal{I}_{M'}(s')$. Observe that the finiteness of Δ_F guarantees that both precondition and successor-state axioms can be effectively evaluated. Then,

if not already present, we add the obtained q' to Q_F , and let $q \rightarrow_F q'$. Finally, we iterate these steps on the newly added states, until no new states are added. Termination is an obvious consequence of Δ_F 's finiteness, which implies that only finitely many states can be added to Q_F . Genericity is a consequence of the fact that the interpretation of states is obtained by answering first-order queries, which are unable to distinguish objects outside the active domain. \square

With these results in place we can immediately prove that if we are given a model M of \mathcal{D} , then it is decidable to check $T_M \models \Phi$. That is we have decidability in case of complete information. Furthermore we can extend such a result to deal with verification in presence of incomplete information. We write $\mathcal{D} \models \Phi$ if $T_M \models \Phi$, for every model M of \mathcal{D} .

Theorem 19. *Let \mathcal{D} be a situation calculus bounded action theory (with infinite object domain) and Φ a closed $\mu\mathcal{L}_a$ formula with all variables renamed apart and belonging to a finite set Vars . Then, it is decidable to check whether $\mathcal{D} \models \Phi$.*

Proof (sketch). Given \mathcal{D} , let \mathcal{D}' be an action theory as in Theorem 17, with $|C'| = 2b + m$, for m the maximum between $|\text{Vars}|$ and the maximum number of variables occurring in the action precondition and successor-state axioms of \mathcal{D} (n of Theorem 17). By Theorem 17, every model M of \mathcal{D} with infinite object domain Δ , has a corresponding p-bisimilar model M' of \mathcal{D}' with finite object domain Δ' of size $|C'|$, and viceversa. Thus, by Theorem 13, for corresponding M and M' , we have that $T_M \models \Phi$ iff $T_{M'} \models \Phi$. Hence, since by Theorem 17, the models of \mathcal{D}' “cover” those of \mathcal{D} and viceversa, it follows that $\mathcal{D} \models \Phi$ iff $\mathcal{D}' \models \Phi$. Finally, decidability is easily obtained by observing that the models M' of \mathcal{D}' are finitely many, up to object renaming, and that by Theorem 18, checking whether $T_{M'} \models \Phi$ is decidable. \square

7 Conclusions

In this paper we have studied first-order μ -calculus with quantification across, in the two main variants proposed in literature. We have seen that the two corresponding notions of bisimulation collapse for the class of generic transition systems, which includes all transition systems generated by reasoning about actions formalisms based on first-order representation of states, and logical mechanisms to generate from the current state the next one, in particular situation calculus. From this we could derive decidability of verification for $\mu\mathcal{L}_a$ over bounded action theories. This result contrasts with verification for the first-order LTL variant corresponding to $\mu\mathcal{L}_a$, which has been shown to be undecidable.

Acknowledgments. This work has been supported by the Sapienza research project *Immersive Cognitive Environments*; by Ripartizione Diritto allo Studio, Università e Ricerca Scientifica di Provincia Autonoma di Bolzano–Alto Adige, under project VERISYNCoPATED (*Verification and Synthesis from Components of Processes that Manipulate Data*); by the Free University of Bozen-Bolzano project KENDO (*Knowledge-driven ENterprise Distributed cOmputing*); and by the EU IP project Optique (*Scalable End-user Access to Big Data*), grant agreement n. FP7-318338.

References

- Abiteboul, S.; Hull, R.; and Vianu, V. 1995. *Foundations of Databases*. Addison Wesley Publ. Co.
- Bagheri Hariri, B.; Calvanese, D.; De Giacomo, G.; Deutsch, A.; and Montali, M. 2013a. Verification of relational data-centric dynamic systems with external services. In *Proc. of the 32nd ACM SIGACT SIGMOD SIGAI Symp. on Principles of Database Systems (PODS)*, 163–174. Extended version available at <http://arxiv.org/abs/1203.0024>.
- Bagheri Hariri, B.; Calvanese, D.; Montali, M.; De Giacomo, G.; De Masellis, R.; and Felli, P. 2013b. Description logic Knowledge and Action Bases. *J. of Artificial Intelligence Research* 46:651–686.
- Baier, C., and Katoen, J.-P. 2008. *Principles of Model Checking*. The MIT Press.
- Belardinelli, F.; Lomuscio, A.; and Patrizi, F. 2012. An abstraction technique for the verification of artifact-centric systems. In *Proc. of the 13th Int. Conf. on the Principles of Knowledge Representation and Reasoning (KR)*, 319–328.
- Belardinelli, F.; Lomuscio, A.; and Patrizi, F. 2014. Verification of agent-based artifact systems. *J. of Artificial Intelligence Research* 51:333–376.
- Bhattacharya, K.; Caswell, N. S.; Kumaran, S.; Nigam, A.; and Wu, F. Y. 2007. Artifact-centered operational modeling: Lessons from customer engagements. *IBM Systems J.* 46(4):703–721.
- Bradfield, J., and Stirling, C. 2007. Modal mu-calculi. In *Handbook of Modal Logic*, volume 3. Elsevier. 721–756.
- Clarke, E. M.; Grumberg, O.; and Peled, D. A. 1999. *Model Checking*. Cambridge, MA, USA: The MIT Press.
- Claßen, J., and Lakemeyer, G. 2008. A logic for non-terminating Golog programs. In *Proc. of the 11th Int. Conf. on the Principles of Knowledge Representation and Reasoning (KR)*, 589–599.
- De Giacomo, G.; Lespérance, Y.; Patrizi, F.; and Vassos, S. 2014. LTL verification of online executions with sensing in bounded situation calculus. In *Proc. of the 21st Eur. Conf. on Artificial Intelligence (ECAI)*, 369–374.
- De Giacomo, G.; Lesperance, Y.; Patrizi, F.; and Sardina, S. 2016. Verifying ConGolog programs on bounded situation calculus theories. In *Proc. of the 30th AAAI Conf. on Artificial Intelligence (AAAI)*.
- De Giacomo, G.; Lesperance, Y.; and Patrizi, F. 2012. Bounded Situation Calculus action theories and decidable verification. In *Proc. of the 13th Int. Conf. on the Principles of Knowledge Representation and Reasoning (KR)*, 467–477.
- De Giacomo, G.; Lesperance, Y.; and Patrizi, F. 2016. Bounded situation calculus action theories. *Artificial Intelligence*. To appear. Preliminary version available at <http://arxiv.org/abs/1509.02012>.
- De Giacomo, G.; Ternovskaia, E.; and Reiter, R. 1997. Non-terminating processes in the situation calculus. In *Proc. of the AAAI 1997 Workshop on Robots, Softbots, Immobiles: Theories of Action, Planning and Control*, 18–28.
- Deutsch, A.; Hull, R.; Patrizi, F.; and Vianu, V. 2009. Automatic verification of data-centric business processes. In *Proc. of the 12th Int. Conf. on Database Theory (ICDT)*, 252–267.
- Emerson, E. A. 1996. Model checking and the mu-calculus. In *Descriptive Complexity and Finite Models*, 185–214. AMS, DIMACS.
- Libkin, L. 2007. Embedded finite models and constraint databases. In *Finite Model Theory and Its Applications*. Springer. 257–338.
- McCarthy, J., and Hayes, P. J. 1969. Some philosophical problems from the standpoint of artificial intelligence. *Machine Intelligence* 4:463–502.
- Okamoto, K. 2010. Comparing expressiveness of first-order modal μ -calculus and first-order CTL*. *RIMS Kokyuroku* 1708:1–14.
- Reiter, R. 2001. *Knowledge in Action. Logical Foundations for Specifying and Implementing Dynamical Systems*. The MIT Press.
- Stirling, C. 2001. *Modal and Temporal Properties of Processes*. Springer.
- van Benthem, J. 1983. *Modal Logic and Classical Logic*. Bibliopolis, Napoli.
- Zarriß, B., and Claßen, J. 2014. Verifying CTL* properties of Golog programs over local-effect actions. In *Proc. of the 21st Eur. Conf. on Artificial Intelligence (ECAI)*, 939–944.