Non-terminating processes in the situation calculus

Giuseppe De Giacomo[1] · Eugenia Ternovska[2] · Ray Reiter[3]

Abstract

By their very design, many robot control programs are non-terminating. This paper describes a situation calculus approach to expressing and proving properties of non-terminating programs expressed in GOLOG, a high level logic programming language for modeling and implementing dynamical systems. Because in this approach actions and programs are represented in classical (second-order) logic, it is natural to express and prove properties of GOLOG programs, including non-terminating ones, in the very same logic. This approach to program proofs has the advantage of logical uniformity and the availability of classical proof theory.

Keywords Knowledge representation · Reasoning about actions · Situation calculus · Inductive definitions · Formal verification of GOLOG and CONGOLOG programs

Mathematics Subject Classification (2010) 68T27

1 Introduction

By their very design, many robot control programs are non-terminating. To give a simple example – one we shall use in this paper – an office coffee-delivery robot might be implemented as an infinite loop in which the robot responds to exogenous requests for coffee that are maintained on a queue. Since a future coffee request is always possible, the program never terminates.

As is the case for more conventional programs, we want some reliability assurances for robot controllers. This paper describes an approach for expressing and proving properties of

✉ Giuseppe De Giacomo
degiacomo@diag.uniroma1.it

Eugenia Ternovska
ter@sfu.ca

Ray Reiter
reiter@cs.toronto.edu

[1] Università di Roma "La Sapienza", Rome, Italy

[2] Simon Fraser University, Vancouver, BC, Canada

[3] University of Toronto, Toronto, ON, Canada

non-terminating programs expressed in GOLOG, a high level logic programming language for modeling and implementing dynamic systems. The kinds of properties we have in mind are traditional in computer science: liveness, fairness, etc. We differ from typical approaches in Formal Methods [5, 11, 38, 39] for reasons dictated by the following characteristics of GOLOG:

1. To write a GOLOG program, the programmer first axiomatizes the primitive actions of the application domain, using first-order logic. These actions may also include exogenous events.
2. Next, she describes, in GOLOG, the complex behaviors her robot is to exhibit in this domain. This GOLOG program is interpreted by means of a formula, this time in second-order logic.
3. Finally, a suitable theorem prover executes the program.

Because these features are all represented in classical (second-order) logic, it is natural to express and prove properties of GOLOG programs, including non-terminating ones, in the very same logic. This approach to program proofs has the advantage of logical uniformity and the availability of techniques from classical proof theory. It also provides a very rich language with which to express program properties, as we shall see in this paper. Moreover, it provides for proofs of programs with incomplete initial state, the normal situation in robotics where the agent does not have complete information about the world it inhabits. Finally, this approach gracefully accommodates exogenous event occurrences, and proofs of program properties in their presence.

This paper is a revised version of a manuscript originally presented at the AAAI 1997 Workshop on Robots, Softbots, Immobots: Theories of Action, Planning and Control [24]. The workshop did not have formal proceedings, so the work has remained unpublished until now. However, as an unpublished manuscript, the paper has been cited often, especially lately as effective techniques to actually perform verification of situation calculus-based programs have become available, as we discuss in the final section of the paper.

2 Formal preliminaries

We briefly summarize the main notions of the Situation Calculus and Golog below. We refer to Reiter's Book [44] for a thorough introduction to these.

2.1 The situation calculus

The situation calculus is a second-order language specifically designed for representing dynamically changing worlds. All changes to the world are the result of named *actions*. A possible world history, which is simply a sequence of actions, is represented by a first-order term called a *situation*. The constant S_0 is used to denote the *initial situation*, namely the empty history. There is a distinguished binary function symbol do; $do(\alpha, s)$ denotes the successor situation to s resulting from performing the action α. Actions may be parameterized. For example, $put(x, y)$ might stand for the action of putting object x on object y, in which case $do(put(A, B), s)$ denotes that situation resulting from placing A on B when the history is s. Notice that in the situation calculus, actions are denoted by first-order terms, and situations (world histories) are also first-order terms. For example,

$$do(putdown(A), do(walk(L), do(pickup(A), S_0)))$$

is a situation denoting the world history consisting of the sequence of actions [pickup(A), walk(L), putdown(A)]. Notice that the sequence of actions in a history, in the order in which they occur, is obtained from a situation term by reading off the actions from right to left. The situation calculus has a distinguished predicate symbol *Poss*; the intended meaning of $Poss(a, s)$ is that it is possible to perform the action a in situation s.

Relations (functions) whose truth values (function values) vary from situation to situation are called *relational (functional) fluents*. They are denoted by predicate (function) symbols taking a situation term as their last argument. For example, $hasCoffee(p, s)$ is a relational fluent whose intended meaning is that person p has coffee in situation s; $robotLocation(s)$ is a functional fluent denoting the robot's location in situation s.

When formalizing an application domain, one must specify certain axioms:

– *Action precondition axioms*, one for each primitive action. These characterize the relation *Poss*, and give the preconditions for the performance of an action in a situation. In a robot coffee delivery setting, such an axiom might be:

$$Poss(giveCoffee(person), s) \equiv$$
$$holdingCoffee(s) \wedge$$
$$robotLocation(s) = office(person)$$

This says that the preconditions for the robot to give coffee to person p are that the robot is carrying coffee, and the robot's location is p's office.

– *Successor state axioms*, one for each fluent. These capture the causal laws of the domain, together with a solution to the frame problem [43]. For our coffee delivery robot, the following is an example:

$$Poss(a, s) \supset [holdingCoffee(do(a, s)) \equiv$$
$$a = pickupCoffee \vee$$
$$holdingCoffee(s) \wedge$$
$$\neg(\exists person)a = giveCoffee(person)].$$

In other words, provided the action a is possible, the robot will be holding a cup of coffee after action a is performed iff a is the action of the robot picking up the coffee, or the robot is already holding coffee and a is not the action of the robot giving that coffee to someone.

– Unique names axioms for the primitive actions, stating that different names for actions denote different actions.

– Axioms describing the initial situation – what is true initially, before any actions have occurred. This is any finite set of sentences which mention no situation term, or only the situation term S_0. Examples of axioms for the initial situation for our coffee delivery example are:

$$\neg(\exists p)hasCoffee(p, S_0), \quad robotLocation(S_0) = CM.$$

These have the intended reading that initially, no one has coffee, and the robot is located at the coffee machine (CM).

See [44] for a full description.

2.2 GOLOG

GOLOG [36, 44] is a situation calculus-based logic programming language that allows for defining complex actions using a repertoire of user specified primitive actions.

GOLOG provides the usual kinds of imperative programming language control structures as well as various forms of nondeterminism. Briefly, GOLOG programs are formed by using the following constructs:

1. *Primitive actions: a*. Do action a in the current situation. Actually a is a situation-suppressed action obtained by suppressing the situation argument in each functional fluent (if any) used in object terms instantiating action parameters. The notation $a[s]$ denotes the result of restoring the situation argument s in all functional fluents mentioned in a (see [36, 44]).

2. *Test actions: ϕ?*. Test the truth value of expression ϕ in the current situation. As for primitive actions, ϕ is a situation-suppressed formula consisting of a formula in the language of the situation calculus, but with all situation arguments omitted. The notation $\phi[s]$ denotes the situation calculus formula obtained from ϕ by restoring situation argument s into all fluent names (relational and functional) mentioned in ϕ (see, again, [36, 44]).

3. *Sequence: δ_1; δ_2*. Execute program δ_1, followed by program δ_2.

4. *Nondeterministic action choice: $\delta_1 \mid \delta_2$*. Execute δ_1 or δ_2.

5. *Nondeterministic choice of arguments: $(\pi z)\delta$*. Nondeterministically pick a value for z, and for that value of z, execute program δ.

6. *Nondeterministic repetition: δ^**. Execute δ a nondeterministic number of times.

7. *While loops:* **while**ϕ **do** δ **endWhile**, which is expressed as $(\phi?; \delta)^*; \neg\phi?)$.

8. *Conditionals:* **if**ϕ **then** δ_1 **else** δ_2, which is expressed as $(\phi?; \delta_1) \mid (\neg\phi?; \delta_2)$.

9. *Procedures*, including recursion: **proc** *ProcName*(**v**) $\delta_{ProcName}$ **endProc**.

3 Single step semantics for GOLOG

In [36, 44], GOLOG programs are interpreted by means of a special relation $Do(\delta, s, s')$ that, given a (generally nondeterministic) program δ and a situation s, returns a possible situation s', resulting by executing δ starting from s. In [36, 44], the relation Do is not denoted by a predicate, but instead it is defined implicitly by using *macros expansion rules*:

$$Do(a, s, s') \stackrel{def}{=} Poss(a[s], s) \wedge s' = do(a[s], s)$$

$$Do(\phi?, s, s') \stackrel{def}{=} \phi[s] \wedge s' = s$$

$$Do(\delta_1; \delta_2, s, s') \stackrel{def}{=} (\exists s'')Do(\delta_1, s, s'') \wedge Do(\delta_2, s'', s')$$

$$Do(\delta_1 \mid \delta_2, s, s') \stackrel{def}{=} Do(\delta_1, s, s') \vee Do(\delta_2, s, s')$$

$$Do((\pi z)\delta, s, s') = (\exists z)Do(\delta, s, s')$$

$$Do(\delta^*, s, s') \stackrel{def}{=} (\forall P)[\ldots \supset P(s, s')]$$

where ... stands for the conjunction of:

$$(\forall s)P(s, s)$$
$$(\forall s, s'', s')P(s, s'') \wedge Do(\delta, s'', s') \supset P(s, s').$$

For simplicity, we skip the macro expansion rules for procedures, we refer the reader to [36, 44] for details. By using such macro expansions rules, the relation $Do(\delta, s, s')$ for the particular program δ is defined by a (generally second-order) formula $\Phi_\delta(s, s')$ not mentioning δ at all. This is very convenient, since it completely avoids the introduction

of programs into the language (they are used only during the macro expansion process to get the formulas $\Phi_\delta(s, s')$ corresponding to $Do(\delta, s, s')$. Observe however that, in this way, programs cannot be quantified over, because they are not terms of the language of the situation calculus.

The kind of semantics *Do* associates to programs, which is based on the complete evaluation of the program, is sometimes called *evaluation semantics* [29]. Such a semantics is not well-suited to interpret non-terminating programs, like infinite loops, since for such programs the evaluation can never be completed and a final situation can never be reached.

For non-terminating programs, one needs to rely on a semantics that allows for interpreting *segments of program executions*. So we adopt a kind of semantics called *computational semantics* [29], which is based on "single steps" of computation, or *transitions*.[1] A step here is either a primitive or a test action. We begin by introducing two special relations, *Final* and *Trans*. *Final*(δ) is intended to say that program δ is in a final state, i.e., it may legally terminate in the current situation.[2] *Trans*(δ, s, δ', s') is intended to say that program δ in situation s may legally execute one step, ending in situation s' with program δ' remaining.

To follow this approach, it is necessary to quantify over programs and so, unlike in [36, 44], we need to encode GOLOG programs as first-order terms, including introducing constants denoting variables, and so on. This is laborious, but quite straightforward [33].[3] We omit all such details here and simply use programs within formulas as if they were already first-order terms.

Final and *Trans* are denoted by predicates defined inductively on the structure of the first argument. It is convenient to include a special "empty" program ε, denoting that nothing of the program remains to be performed.

The definition of *Final* is as follows:

$$(\forall \delta) Final(\delta) \equiv (\forall F)[\ldots \supset F(\delta)]$$

where ... stands for the conjunction of the universal closure of the following clauses:

$$F(\varepsilon)$$
$$F(\delta_1) \wedge F(\delta_2) \supset F(\delta_1; \delta_2)$$
$$F(\delta_1) \vee F(\delta_2) \supset F(\delta_1 \mid \delta_2)$$
$$F(\delta) \supset F((\pi z)\delta))$$
$$F(\delta^*)$$
$$F(\delta_{ProcName}) \supset F(ProcName(\mathbf{x}))$$

Observe that being final is a syntactic property of programs: programs of a certain form are considered to be in a final state. Moreover being final does not depend on the objects the program deals with, indeed $Final((\pi z)\delta)$ and $Final(ProcName(\mathbf{x}))$ depend only on δ and $\delta_{ProcName}$ and not on the particular values of z and $\mathbf{x}$ respectively. Observe that from the above definition we get that primitive and test actions are never final: for all primitive actions a $Final(a) \equiv$ **False** and for all tests ϕ? $Final(\phi?) \equiv$ **False**.

The definition of *Trans* is as follows:

$$(\forall \delta, s, \delta', s') Trans(\delta, s, \delta', s') \equiv (\forall T)[\ldots \supset T(\delta, s, \delta', s')]$$

---

[1]Both types of semantics belong to the family of structural operational semantics introduced in [42].

[2]Notice that, *Final* means that the program *may* be considered terminated, but not that all its possible executions are necessarily terminated.

[3]We assume that the predicates introduced in this section, including *Final* and *Trans*, cannot occur in tests, hence disallowing self-reference.

where ... stands for the conjunction of the universal closure of the following clauses:

$$Poss(a[s], s) \supset T(a, s, \varepsilon, do(a[s], s))$$

$$\phi[s] \supset T(\phi?, s, \varepsilon, s)$$

$$T(\delta_1, s, \delta_1', s') \supset T(\delta_1; \delta_2, s, \delta_1'; \delta_2, s')$$
$$Final(\delta_1) \wedge T(\delta_2, s, \delta_2', s') \supset T(\delta_1; \delta_2, s, \delta_2', s')$$

$$T(\delta_1, s, \delta_1', s') \supset T(\delta_1 \mid \delta_2, s, \delta_1', s')$$
$$T(\delta_2, s, \delta_2', s') \supset T(\delta_1 \mid \delta_2, s, \delta_2', s')$$

$$(\exists y)T(\delta_y^z, s, \delta', s') \supset T((\pi z)\delta, s, \delta', s')$$

$$T(\delta, s, \delta', s') \supset T(\delta^*, s, \delta'; \delta^*, s')$$

$$T((\delta_{ProcName})_{\mathbf{x}}^{\mathbf{v}}, s, \delta', s') \supset T(ProcName(\mathbf{x}), s, \delta', s')$$

Above, δ_y^z stands for the result of the syntactic substitution of z by y in δ; analogously, $(\delta_{ProcName})_{\mathbf{x}}^{\mathbf{v}}$ stands for the result of the syntactic substitution of the parameters $\mathbf{v}$ by $\mathbf{x}$. The clauses defining *Trans* characterize when a *configuration* (δ, s) can evolve (in a single step) to a configuration (δ', s'). Intuitively they can be read as follows:

– (a, s) evolves to $(\varepsilon, do(\alpha[s], s))$, provided $a[s]$ is possible in s. Observe that after having performed a, nothing remains to be performed.
– $(\phi?, s)$ evolves to (ε, s), provided that $\phi[s]$ holds. Otherwise, it cannot proceed. Observe that the situation remains unchanged.
– $(\delta_1; \delta_2, s)$ can evolve to $(\delta_1'; \delta_2, s')$, provided that (δ_1, s) can evolve to (δ_1', s'). Moreover, it can evolve to (δ_2', s'), provided that δ_1 is final and (δ_2, s) can evolve to (δ_2', s').
– $(\delta_1 \mid \delta_2, s)$ can evolve to (δ', s'), provided that either (δ_1, s) or (δ_2, s) can do so.
– $((\pi z)\delta, s)$ can evolve to (δ', s'), provided that there exists a y such that (δ_y^z, s) can evolve to $(\delta', s') - z$ is bound by π in $(\pi z)\delta$ and is typically free in δ.
– (δ^*, s) can evolve to $(\delta'; \delta^*, s')$ provided that (δ, s) can evolve to (δ', s'). Observe that (δ^*, s) can also not evolve at all, since δ^* is final.
– $(ProcName(\mathbf{x}), s)$ can evolve to (δ', s'), provided that the body $\delta_{ProcName}$ of the procedure *ProcName*, with the actual parameters $\mathbf{x}$ substituted for the formal parameters $\mathbf{v}$, can do so.

The possible configurations that can be reached by a program δ starting in a situation s are those obtained by repeatedly following the transition relation denoted by *Trans* starting from (δ, s), i.e., those in the reflexive transitive closure of the transition relation. Such a relation is denoted by the "reflexive-transitive closure" of *Trans*, *Trans** defined as:

$$(\forall \delta, s, \delta', s')Trans^*(\delta, s, \delta', s') \equiv \forall U[\ldots \supset U(\delta, s, \delta', s')]$$

where ... stands for the conjunction of the universal closure of the following clauses:

$$U(\delta, s, \delta, s)$$
$$U(\delta, s, \delta', s') \wedge Trans(\delta', s', \delta'', s'') \supset U(\delta, s, \delta'', s'')$$

Using *Trans** and *Final*, we may denote the relation *Do* as follows:

$$Do(\delta, s, s') \stackrel{def}{=} (\exists \delta') Trans^*(\delta, s, \delta', s') \wedge Final(\delta')$$

In other words, $Do(\delta, s, s')$ holds if it is possible to repeatedly single-step the program δ, obtaining a program δ' and a situation s' such that δ' can legally terminate in s'. Note that this formulation of *Do* is equivalent to the one in [36, 44] (c.f. [29]).

4 Exogenous actions

Exogenous action are primitive actions that are not under the control of the program [44]. They are executed by other agents in an asynchronous way with respect to the program. *Trans* can be easily modified to take into account exogenous actions as well. It suffice to add to the above definition a clause having, as a first approximation, the form:

$$Exo(exo) \wedge Poss(exo, s) \supset T(\delta, s, \delta, do(exo, s))$$

which says that any configuration (δ, s) can evolve, due to the occurrence of an exogenous action *exo*, to $(\delta, do(exo, s))$, where the situation has changed but the program hasn't.

The above clause enables the occurrence of an exogenous action *exo* every time the action preconditions for *exo*, and hence $Poss(exo, s)$, are true. However it is of interest, to restrict further the actual occurrence of *exo* along a sequence of transitions, establishing some sort of *dynamics* for exogenous actions. Such a dynamics has a role similar to that of programs for normal primitive actions although typically it is not strict enough to extract a program that implements it. Rather the dynamics of exogenous actions has to be specified by means of suitable axioms.

A possible way to follow such a strategy is to introduce a special fluent $DynaPoss(exo, s)$ and modify *Trans* by introducing the following refinement of the above clause:

$$Exo(exo) \wedge Poss(exo, s) \wedge DynaPoss(exo, s) \supset \\ T(\delta, s, \delta, do(exo, s)).$$

Then one uses special axioms expressing the dynamics of exogenous actions by specifying in which situations s, along a sequence of transitions, $DynaPoss(exo, s)$ holds. Such axioms may express sophisticated temporal/dynamic laws and typically they are going to be second-order. Observe that *exo* can actually occur only if both $Poss(exo, s)$ and $DynaPoss(exo, s)$ hold in s.

5 Logical representation of inductive definitions and fixpoints

The relations *Trans* and *Final* are defined inductively. *Inductive definitions* [1, 33, 40] are broadly used in mathematical logic for defining sets and have became widely used in computer science [12]. A *rule-based inductive definition* is a set $\mathcal{R}$ of rules of the form $\frac{P}{c}$, where P is the set of premises and c is the conclusion, together with a closure condition: a set Z is $\mathcal{R}$-closed if each rule in $\mathcal{R}$ whose premises are in Z also has its conclusion in Z. A set H, *inductively defined by* $\mathcal{R}$, is given by $H = \bigcap \{Z \mid Z \text{ is } \mathcal{R}\text{-closed}\}$ or by $H = \bigcup \{Z \mid Z \text{ is } \mathcal{R}\text{-closed}\}$. The former is called a *positive inductive* definition of H, the latter is called a *negative inductive* or *coinductive* definition of H. Let U be a set. An

operator induced by an inductive definition is a total mapping $\Gamma \;:\; Pow(U) \mapsto Pow(U)$, such that

$$\Gamma(Z) = \left\{ c \in U \mid \exists P \subseteq Z \;:\; \frac{P}{c} \in \mathcal{R} \right\}$$

That is, Γ is a mapping taking sets to sets.

Inductive definitions are strongly related to *fixpoint properties*. i.e., properties defined as solutions of recursive equations. Specifically, positive inductive definitions are related to least fixpoints, i.e., minimal solution of the recursive equations, whereas negative inductive definitions are related to greatest fixpoints, i.e., maximal solutions of the recursive equations. Dynamic properties are typically fixpoint properties, expressed as the least or greatest solutions of certain recursive logical equations (e.g. see [49]).

Every property definable as an extreme fixpoint must have, by definition:

- its own construction principle, a recursive equation a fixpoint of which is our property;
- an appropriate induction or coinduction principle to guarantee the minimality or maximality of the solution of the recursive equation.

5.1 Construction principle

To define a set Z, here denoted by a predicate $Z(\mathbf{x})$, we need to say what its elements are. The *construction principle* tells us how to obtain these elements recursively.

$$(\forall \mathbf{x}) Z(\mathbf{x}) \equiv \Phi(Z, \mathbf{x}) \tag{1}$$

In this case Φ is called a *constructor* for Z. Any solution of this recursive equation is called a *fixpoint* of the operator Φ. The Knaster-Tarski Theorem [31, 50] guarantees that if the operator Φ is monotone, the (1) has both a least and a greatest solution. A sufficient condition for monotonicity is that all occurrences of Z occur within a even number of negations.[4] This condition is always satisfied in this paper.

5.2 Induction principle: least fixpoints

To guarantee that Z is the smallest solution, we apply the *induction principle:*[5]

$$(\forall P, \mathbf{x})\{[(\forall \mathbf{y})\Phi(P, \mathbf{y}) \supset P(\mathbf{y})] \supset [Z(\mathbf{x}) \supset P(\mathbf{x})]\} \tag{2}$$

i.e., whatever solution P of the recursive specification we take, Z is included in it.

A set Z satisfying construction principle (1) and induction principle (2) is denoted by $\mu_{P,\mathbf{y}}\Phi(P, \mathbf{y})(\mathbf{x})$, and it is called a *least fixpoint* of an operator $\Phi(P, \mathbf{y})$. Note that in $\mu_{P,\mathbf{y}}\Phi(P, \mathbf{y})(\mathbf{x})$ the predicate variable P and the individual variables $\mathbf{y}$ are considered bounded by μ, while the individual variables $\mathbf{x}$ are free. Another view of $\mu_{P,\mathbf{y}}\Phi(P, \mathbf{y})(\mathbf{x})$ is that $\mu_{P,\mathbf{y}}\Phi(P, \mathbf{y})$ is the name of a defined predicate, and $\mathbf{x}$ are its arguments.

We can rewrite the induction principle (2) in the following way

$$(\forall \mathbf{x})\{Z(\mathbf{x}) \supset [(\forall P)[(\forall \mathbf{y})\Phi(P, \mathbf{y}) \supset P(\mathbf{y})] \supset P(\mathbf{x})]\} \tag{3}$$

---

[4]Interpreting $\Phi \supset \Psi$ as an abbreviation for $\neg\Phi \vee \Psi$.

[5]The idea of defining a least fixpoint using two principles, construction and induction, is from [28].

Notice that implication in the opposite direction follows from the construction principle (1). We obtain

$$(\forall \mathbf{x})\{\mu_{P,\mathbf{y}}\Phi(P,\mathbf{y})(\mathbf{x}) \equiv [(\forall P)[(\forall \mathbf{y})\Phi(P,\mathbf{y}) \supset P(\mathbf{y})] \supset P(\mathbf{x})]\} \qquad (4)$$

The last sentence is often considered as a formal definition of a least fixpoint. Observe that it has exactly the form we have used to define *Trans* and *Final* (as well as $Do(\delta^*, s, s')$ in [36]).

5.3 Coinduction principle: greatest fixpoints

To guarantee that Z is the biggest solution of (1), we apply the *coinduction principle:*

$$(\forall P, \mathbf{x})\{[(\forall \mathbf{y})P(\mathbf{y}) \supset \Phi(P,\mathbf{y})] \supset [P(\mathbf{x}) \supset Z(\mathbf{x})]\} \qquad (5)$$

i.e., whatever solution P of the recursive specification we take, Z includes it.

We can rewrite the coinduction principle (5) in the following way

$$(\forall \mathbf{x})\{[(\exists P)[(\forall \mathbf{y})P(\mathbf{y}) \supset \Phi(P,\mathbf{y})] \wedge P(\mathbf{x})] \supset Z(\mathbf{x})\} \qquad (6)$$

An explicit expression for a greatest fixpoint can be obtained in a similar way as was done for a least fixpoint:

$$(\forall \mathbf{x})\{\nu_{P,\mathbf{y}}\Phi(P,\mathbf{y})(\mathbf{x}) \equiv [(\exists P)[(\forall \mathbf{y})P(\mathbf{y}) \supset \Phi(P,\mathbf{y})] \wedge P(\mathbf{x})]\} \qquad (7)$$

The last sentence can be taken as a definition of a greatest fixpoint.

6 Examples of expressible dynamic properties

With *Trans* and *Final* in place, a wide variety of dynamic properties can be expressed by relying on second-order formulae expressing least and greatest fixpoint properties. In particular properties expressible by logics of programs, such as dynamic logics [32], mu-calculus [41, 49], and temporal logics [25], can be rephrased in our setting. Let us present some examples.

1. The formula:

$$Q_1(\delta_0, s_0) \overset{def}{=} \\ \mu_{P,\delta,s}[\psi(\delta, s) \vee (\exists \delta', s')Trans(\delta, s, \delta', s') \wedge \\ P(\delta', s')](\delta_0, s_0)$$

(where δ_0, s_0 are individual variables) defines a predicate $Q_1(\delta_0, s_0)$ that denotes the smallest set of configurations C_1 such that a configuration (δ, s) belongs to this set (the predicate Q_1 is true on (δ, s)) if and only if either ψ is true on (δ, s) or there exists a configuration (δ', s'), reachable in one step by the relation *Trans*, which also belongs to the set C_1.

In this way, the formula expresses that from each configuration (δ_0, s_0) on which the specified predicate is true, there exists an execution path that eventually reaches a configuration (δ, s) on which ψ is true.

As a special case, by taking $\psi(\delta, s)\overset{def}{=}\phi(s) \wedge \mathit{Final}(\delta)$, one can express that there exists a terminating execution of program δ_0 starting from situation s_0 such that ϕ is true in the final situation.

2. The formula:

$$Q_2(\delta_0, s_0)\overset{def}{=}$$
$$\mu_{P,\delta,s}\{\psi(\delta, s) \vee [(\exists\delta', s')\mathit{Trans}(\delta, s, \delta', s')] \wedge$$
$$(\forall\delta', s')\mathit{Trans}(\delta, s, \delta', s') \supset P(\delta', s')\}(\delta_0, s_0)$$

defines a predicate $Q_2(\delta_0, s_0)$ that denotes the smallest set of configurations C_2 such that the predicate is true on configuration (δ, s) if and only if either ψ is true on (δ, s) or there exists a configuration (δ', s') reachable in one step by the relation Trans, and on all such configurations the predicate is still true.

In this way, the formula expresses that from each configuration (δ_0, s_0) on which the specified predicate is true, all execution paths eventually reach a configuration (δ, s) on which ψ is true.

3. The formula:

$$Q_3(\delta_0, s_0)\overset{def}{=}$$
$$\nu_{P,\delta,s}[\psi(\delta, s) \wedge$$
$$(\exists\delta', s')\mathit{Trans}(\delta, s, \delta', s') \wedge P(\delta', s')](\delta_0, s_0)$$

defines a predicate $Q_3(\delta_0, s_0)$ that denotes the greatest set of configurations C_3, such that the predicate is true on configuration (δ, s) if and only if both ψ is true on (δ, s) and the predicate is still true on at least one configuration (δ', s') reachable in one step by the relation Trans.

In this way, the formula expresses that from each configuration (δ_0, s_0) on which the specified predicate is true, there exists a non-terminating execution path along which ψ is always true.

As a special case, by $\psi(\delta, s)\overset{def}{=}\mathbf{True}$, one can express that there exists a non-terminating execution path.

4. The formula:

$$Q_4(\delta_0, s_0)\overset{def}{=}$$
$$\nu_{P,\delta,s}[\psi(\delta, s) \wedge$$
$$(\forall\delta', s')\mathit{Trans}(\delta, s, \delta', s') \supset P(\delta', s')](\delta_0, s_0)$$

defines a predicate that denotes the greatest set of configurations C_4 such that the predicate is true on configuration (δ, s) if and only if both ψ is true on (δ, s) and the predicate is still true on each configuration (δ', s') reachable in one step by the relation Trans.

In this way, the formula expresses that from each configuration (δ_0, s_0) on which the specified predicate is true, along all execution paths ψ is always true.

As a special case, by $\psi(\delta, s)\overset{def}{=}\neg\mathit{Final}(\delta) \wedge (\exists\delta', s')\mathit{Trans}(\delta, s, \delta', s')$, one can express that all execution paths are non-terminating and no final state is ever reached.

7 Example: a coffee delivery robot

Here, we describe a robot whose task is to deliver coffee in an office environment. The robot can carry just one cup of coffee at a time, and there is a central coffee machine from which it

gets the coffee. The robot receives *asynchronous* requests for coffee from employees. These requests are put in a queue. The robot continuously takes the first request from the queue and serves coffee to the specified person. The use of the queue guarantees that all requests will in fact be served (implementing a *fair* serving policy).

7.1 Representation of the queue

As usual, to define an abstract data type, we need to specify the *domain of its values*, and its *functions and predicates*.

The domain of values for queues is constructed inductively from the constant *nil* and the functor $cons(\cdot, \cdot)$ as follows:[6]

$$(\forall q)IsQueue(q) \equiv (\forall Q)[\ldots \supset Q(q)]$$

where ... stands for the conjunction of:

$$Q(nil)$$
$$(\forall f, r)Q(r) \supset Q(cons(f, r))$$

The functions and predicates for queues are the usual *first*(·), *dequeue*(·), *enqueue*(·, ·) and *isEmpty*(·). They are defined in our setting as follows:

$$(\forall f, r)first(cons(f, r)) = f \quad \text{(unspecified for \texttt{nil})}$$

$$(\forall f, r)dequeue(cons(f, r)) = r \,\text{(unspecified for \texttt{nil})}$$

$$(\forall p)enqueue(nil, p) = cons(p, nil)$$
$$(\forall p, f, r)enqueue(cons(f, r), p) = cons(f, enqueue(r, p))$$

$$(\forall q)isEmpty(q) \equiv (q = nil)$$

To these, we add the function $length(\cdot)$ that returns the length of the queue, and the predicate $isFull(\cdot)$, since we are going to need queues of a bounded length.

$$length(nil) = 0$$
$$(\forall f, r)length(cons(f, r)) = 1 + length(r)$$

$$(\forall q)isFull(q) \equiv (length(q) = 100)$$

We enforce unique name assumption for terms built from *nil* and $cons(\cdot, \cdot)$, but obviously not for those built with the functions $dequeue(\cdot)$, $enqueue(\cdot, \cdot)$ and $length(\cdot)$.

7.2 Formalization of the example

Primitive actions

– *requestCoffee*(*person*). A request for coffee is received from the employee *person*. This action is an *exogenous* one, i.e., an action not under the control of the robot. $(\forall p)Exo(requestCoffee(p))$ holds.

---

[6]Equivalently, $(\forall q_0)IsQueue(q_0) \equiv \mu_{Q,q}[q = nil \vee (\exists f, r)q = cons(f, r) \wedge Q(r)](q_0)$.

- *selectRequest(person)*. The first request in the queue is selected, and the employee *person* that made that request will be served.
- *pickupCoffee*. The robot picks up a cup of coffee from the coffee machine.
- *giveCoffee(person)*. The robot gives a cup of coffee to *person*.
- *startGo(loc₁, loc₂)*. The robot starts to go from location loc_1 to loc_2.
- *endGo(loc₁, loc₂)*. The robot ends its process of going from location loc_1 to loc_2.

Fluents

- *queue(s)*. A functional fluent denoting the queue of requests in situation s.
- *robotLocation(s)*. A functional fluent denoting the robot's location in situation s.
- *hasCoffee(person, s)*. *person* has coffee in s.
- *going(loc₁, loc₂, s)*. In situation s, the robot is going from loc_1 to loc_2.
- *holdingCoffee(s)*. In situation s, the robot is holding a cup of coffee.

Situation independent predicates and functions

- *office(person)*. Denotes the office of *person*.
- *CM*. Constant denoting coffee machine's location.
- *Sue, Mary, Bill, Joe*. Constants denoting people.

Primitive action preconditions

$$Poss(requestCoffee(p), s) \equiv$$
$$\neg isFull(queue(s))$$

$$Poss(selectRequest(p), s) \equiv$$
$$\neg isEmpty(queue(s)) \wedge p = first(queue(s))$$

$$Poss(pickupCoffee, s) \equiv$$
$$\neg holdingCoffee(s) \wedge$$
$$robotLocation(s) = CM$$

$$Poss(giveCoffee(person), s) \equiv$$
$$holdingCoffee(s) \wedge$$
$$robotLocation(s) = office(person)$$

$$Poss(startGo(loc_1, loc_2), s) \equiv$$
$$\neg(\exists l, l')going(l, l', s) \wedge loc_1 \neq loc_2 \wedge$$
$$robotLocation(s) = loc_1$$

$$Poss(endGo(loc_1, loc_2), s) \equiv$$
$$going(loc_1, loc_2, s).$$

Successor state axioms

$$Poss(a, s) \supset$$
$$[queue(do(a, s)) = q \equiv$$
$$(\exists p)a = requestCoffee(p) \wedge$$
$$q = enqueue(queue(s), p) \vee$$
$$(\exists p)a = selectRequest(p) \wedge$$
$$q = dequeue(queue(s), p) \vee$$
$$(\forall p)a \neq requestCoffee(p) \wedge$$
$$a \neq selectRequest(p) \wedge q = queue(s)]$$

$$Poss(a, s) \supset$$
$$[hasCoffee(person, do(a, s)) \equiv$$
$$a = giveCoffee(person) \vee hasCoffee(person, s)]$$

$$Poss(a, s) \supset$$
$$[robotLocation(do(a, s)) = loc \equiv$$
$$(\exists loc')a = endGo(loc', loc) \vee$$
$$robotLocation(s) = loc \wedge$$
$$\neg(\exists loc', loc'')a = endGo(loc', loc'')]$$

$$Poss(a, s) \supset$$
$$[going(l, l', do(a, s)) \equiv$$
$$a = startGo(l, l') \vee$$
$$going(l, l', s) \wedge a \neq endGo(l, l')]$$

$$Poss(a, s) \supset$$
$$[holdingCoffee(do(a, s)) \equiv$$
$$a = pickupCoffee \vee$$
$$holdingCoffee(s) \wedge$$
$$\neg(\exists person)a = giveCoffee(person)].$$

Additional axioms

(The following axiom is not strictly necessary, we add it for clarity.)

$$(\forall s)IsQueue(queue(s))$$
(the values of $queue(\cdot)$ are queues)

Unique names axioms stating that the following terms, together with those formed from *nil* and $cons(\cdot, \cdot)$ (see above), are pairwise unequal:

$$Sue, \ Mary, \ Bill, \ Joe, \ CM, \ office(Sue),$$
$$office(Mary), \ office(Bill), \ office(Joe).$$

Initial situation

$$robotLocation(S_0) = CM \wedge \neg holdingCoffee(S_0) \wedge$$
$$\neg(\exists l, l')going(l, l', S_0) \wedge \neg(\exists p)hasCoffee(p, S_0) \wedge$$
$$queue(S_0) = nil$$

Robot's GOLOG program

The robot executes the program *DeliverCoffee* defined as follows (note the suppressed situation argument in primitive and test actions):

proc *DeliverCoffee*
    while True do
        if $\neg isEmpty(queue)$
            then $(\pi p)selectRequest(p); ServeCoffee(p)$
            else True?   (skip)
    endWhile
endProc

proc *ServeCoffee(p)*
    $Goto(CM)$;
    pickupCoffee;
    $Goto(office(p))$;
    $giveCoffee(p)$
endProc

proc *Goto(loc)*
    $startGo(robotLocation, loc)$;
    $endGo(robotLocation, loc)$
endProc

Dynamics of exogenous actions

Along all possible evolutions of any program δ_0, starting from S_0, into any configuration, in a finite number of transitions, a situation s is reached where somebody may request coffee (*DynaPoss* holds) (provided that it is possible to request coffee, i.e., that also *Poss* holds):

$$(\forall \delta_0, \delta, s)Trans^*(\delta_0, S_0, \delta, s) \supset ExoLaws(\delta, s)$$

$$ExoLaws(\delta_1, s_1) \overset{def}{=}$$
$$\mu_{E,\delta,s}\{[(\exists p)DynaPoss(requestCoffee(p), s)] \vee$$
$$[(\forall \delta', s')Trans(\delta, s, \delta', s') \supset E(\delta', s')]\}(\delta_1, s_1)$$

7.3 Reasoning

Next, we show some dynamic properties of the overall system (the program plus the exogenous actions). First, it is easy to see, from its structure, that the program *DeliverCoffee* will never reach a final configuration:

$$(\forall \delta, s)Trans^*(DeliverCoffee, S_0, \delta, s) \supset \neg Final(\delta).$$

It is also possible to show the following more complex property: every request for coffee sooner or later will be served. Formally, the *fairness* property $Fair(DeliverCoffee, S_0)$ holds, where:

$$Fair(\delta_0, s_0) \stackrel{def}{=}$$
$$(\forall p, \delta, s)Trans^*(\delta_0, s_0, \delta, do(requestCoffee(p), s)) \supset$$
$$EventuallyServed(p, \delta, do(requestCoffee(p), s))$$

and

$$EventuallyServed(p, \delta_1, s_1) \stackrel{def}{=}$$
$$\mu_{P,\delta,s}\{[(\exists s'')s = do(selectRequest(p), s'')] \vee$$
$$[((\exists \delta'.s')Trans(\delta, s, \delta', s')) \wedge$$
$$(\forall \delta', s')Trans(\delta, s, \delta', s') \supset P(\delta', s')]\}(\delta_1, s_1)$$

It is also possible to show that there exists an (infinite) execution path where no coffee is ever served:

$$PossiblyAlwaysIdle(DeliverCoffee, S_0)$$

where

$$PossiblyAlwaysIdle(\delta_0, s_0) \stackrel{def}{=}$$
$$\nu_{A,\delta,s}\{[(\forall p, s'')(s \neq do(selectRequest(p), s'')] \wedge$$
$$[(\exists \delta', s')Trans(\delta, s, \delta', s') \wedge A(\delta', s')]\}(\delta_0, s_0).$$

However, by the fairness property above, this means that no requests for coffee were made along that execution path.

8 Discussion

The 1997 manuscript [24], which constitutes the core part of this paper, was the first work addressing verification of possibly non-terminating programs in the situation calculus. Notably, no effective techniques were available at the time, except for the easy case where the object domain is assumed to be finite and known a priori (i.e., essentially the propositional situation calculus). Since then, there has been growing interest in reasoning about and verifying agent programs in the situation calculus and outside of it. Here, we review some of the work that was directly influenced by the original 1997 manuscript.

Several phenomena in reasoning about complex actions have been studied by leveraging on the transition semantics. The semantics itself was originally introduced to deal with concurrency in CONGOLOG [13, 14]. Later it was exploited for capturing: execution monitoring in [23], the distinction between online and offline execution in INDIGOLOG [22, 46], the notions of ability and epistemic feasibility [34, 35], compilation of search over nondeterministic programs into planning problems [6], connections with web service composition [48], etc.. The situation calculus causal laws, the transition semantics and the temporal properties were captured in the unifying framework of inductive/conductive definitions in [51]. Apart from the literature on the situation calculus, the work in this paper has influenced other areas of Artificial Intelligence, in particular research on merging reasoning about actions with description logics [2, 4, 53].

Focusing on the propositional situation calculus (where fluents have only situation as an argument) decidable verification techniques were devised in [52]. Later these were extended to a one-argument fluents fragment of the situation calculus [27]. Techniques for verification resorting to second-order theorem proving with no decidability guarantees have been studied in [47] where the CASL verification environment for multi-agent CONGOLOG

programs is described. In [9], *characteristic graphs* of programs were introduced to define a form of regression over programs, to be used as a pre-image computation step in (sound) procedures for verifying GOLOG and CONGOLOG programs inspired by model checking. Verification of programs over a two-variable fragment of the situation calculus was shown to be decidable in [10, 54, 55]. The work in [30] established conditions to verify loop invariants and persistence properties. Finally, the work in [21, 45] can be seen as a direct follow-up to the present work in that the authors propose techniques (with model-checking ingredients) to reason about infinite executions of GOLOG and CONGOLOG programs directly based on second-order logic exploiting fixpoint approximates. More recently, [15–17, 19, 20] have shown that one obtains robust decidability results for temporal verification of "bounded situation calculus action theories", i.e., situation calculus action theories such that, in every situation, the number of object tuples forming the extension of each fluent is bounded by a fixed number. Such a research is related to decidable verification for bounded data-aware dynamic systems [3, 7]. These decidability results for bounded situation calculus action theories have been extended to GOLOG and CONGOLOG programs (without recursive procedures), where atomic actions are specified by a bounded situation calculus action theory [18].

Finally, the general results in [8], combined with the result in [18], allow for devising decidability for exactly the fixpoint language over programs presented here, in the case of bounded action theories and programs with tail-recursion only. Indeed, the language studied in this paper corresponds to the mu-calculus with full-fledged quantification across configurations, over infinite transitions systems generated by *Trans* and *Final*. The paper [8] shows that verifying properties expressed in such a logic against any bounded-state generic transition system is reducible to propositional mu-calculus model checking over a finite transition system which is a faithful abstraction of the original one. At the same time, using the construction in [18], one can show that *Trans* and *Final* over bounded action theories and programs with tail-recursion only induce bounded-state generic transition systems.

One of the drawbacks of the transition semantics used in this paper is the need for introducing a special sort for programs which, similarly to situations, needs to be axiomatized in second-order logic. There has been a constant interest by the research community to find ways of compiling away such a program sort as can be done for GOLOG under the *Do*-semantics [36]. In particular, [26], showed how to compile arbitrary programs into Petri nets plus (unbounded) stacks for recursion, and then encode these into a basic action theory. More recently, [37] has demonstrated how to drop the second-order axioms required for representing CONGOLOG programs as terms, by storing all information about programs into a new distinguished situation term and exploiting the standard foundational second-order induction axiom for situations. Both proposals focus on correctness results for the *Do*-semantics, which is concerned with finite traces, hence those results cannot be used for doing temporal verification of non-terminating programs. In the same spirit, [18] also compiles away the program sort, and in fact this is one of the key elements for achieving decidability for temporal verification of non-terminating programs over bounded theories.

Much of the above research explicitly recognizes the unpublished 1997 manuscript, constituting this paper, as one of its main foundations. It is quite amazing, though certainly not an isolated case, to see how an unpublished manuscript could have such a strong impact on the succeeding research.

References

1. Aczel, P.: An introduction to inductive definitions. In: Barwise, J. (ed.) Handbook of Mathematical Logic, pp. 739–782. Elsevier (1977)
2. Baader, F., Liu, H., ul Mehdi, A.: Verifying properties of infinite sequences of description logic actions. In: ECAI, pp. 53–58 (2010)
3. Bagheri Hariri, B., Calvanese, D., De Giacomo, G., Deutsch, A., Montali, M.: Verification of relational data-centric dynamic systems with external services. In: PODS, pp. 163–174 (2013)
4. Bagheri Hariri, B., Calvanese, D., Montali, M., De Giacomo, G., Masellis, R.D., Felli, P.: Description logic knowledge and action bases. J. Artif. Intell. Res. (JAIR) **46**, 651–686 (2013)
5. Baier, C., Katoen, J.P.: Principles of Model Checking. MIT Press (2008)
6. Baier, J.A., Fritz, C., McIlraith, S.A.: Exploiting procedural domain control knowledge in state-of-the-art planners. In: ICAPS, pp. 26–33 (2007)
7. Belardinelli, F., Lomuscio, A., Patrizi, F.: Verification of agent-based artifact systems. J. Artif. Intell. Res. (JAIR) **51**, 333–376 (2014)
8. Calvanese, D., De Giacomo, G., Montali, M., Patrizi, F.: First-order μ-calculus over generic transition systems and applications to the situation calculus. Inf. Comput. **259**(3), 328–347 (2018)
9. Claßen, J., Lakemeyer, G.: A logic for non-terminating Golog programs. In: KR, pp. 589–599 (2008)
10. Claßen, J., Liebenberg, M., Lakemeyer, G., Zarrieß, B.: Exploring the boundaries of decidable verification of non-terminating Golog programs. In: AAAI, pp. 1012–1019 (2014)
11. Cousot, P.: Methods and logics for proving programs. In: Handbook of Theoretical Computer Science, Volume B: Formal Models and Sematics (B), pp. 841–994 (1990)
12. Cousot, P., Cousot, R.: Inductive definitions, semantics and abstract interpretation. In: POPL, pp. 83–94 (1992)
13. De Giacomo, G., Lespérance, Y., Levesque, H.J.: Reasoning about concurrent execution prioritized interrupts, and exogenous actions in the situation calculus. In: IJCAI, pp. 1221–1226 (1997)
14. De Giacomo, G., Lespérance, Y., Levesque, H.J.: ConGolog, a concurrent programming language based on the situation calculus. Artif. Intell. **121**(1–2), 109–169 (2000)
15. De Giacomo, G., Lespérance, Y., Patrizi, F.: Bounded situation calculus action theories and decidable verification. In: KR (2012)
16. De Giacomo, G., Lespérance, Y., Patrizi, F.: Bounded epistemic situation calculus theories. In: IJCAI (2013)
17. De Giacomo, G., Lespérance, Y., Patrizi, F.: Bounded situation calculus action theories. Artif. Intell. **237**, 172–203 (2016)
18. De Giacomo, G., Lespérance, Y., Patrizi, F., Sardiña, S.: Verifying ConGolog programs on bounded situation calculus theories. In: AAAI, pp. 950–956 (2016)
19. De Giacomo, G., Lespérance, Y., Patrizi, F., Vassos, S.: LTL verification of online executions with sensing in bounded situation calculus. In: ECAI, pp. 369–374 (2014)
20. De Giacomo, G., Lespérance, Y., Patrizi, F., Vassos, S.: Progression and verification of situation calculus agents with bounded beliefs. In: AAMAS, pp. 141–148 (2014)
21. De Giacomo, G., Lespérance, Y., Pearce, A.R.: Situation calculus based programs for representing and reasoning about game structures. In: KR (2010)
22. De Giacomo, G., Levesque, H.J.: An incremental interpreter for high-level programs with sensing. In: Levesque, H.J., Pirri, F. (eds.) Logical Foundation for Cognitive Agents: Contributions in Honor of Ray Reiter, pp. 86–102. Springer (1999)
23. De Giacomo, G., Reiter, R., Soutchanski, M.: Execution monitoring of high-level robot programs. In: KR, pp. 453–465 (1998)
24. De Giacomo, G., Ternovskaia, E., Reiter, R.: Non-terminating processes in the situation calculus. In: Proc.of the AAAI'97 Workshop on Robots, Softbots, Immobots: Theories of Action, Planning and Control (1997)
25. Emerson, E.A.: Automated temporal reasoning about reactive systems. In: Logics for Concurrency: Structure versus Automata, no. 1043 in Lecture Notes in Computer Science, pp. 41–101. Springer (1996)
26. Fritz, C., Baier, J.A., McIlraith, S.A.: ConGolog, Sin Trans: Compiling ConGolog into basic action theories for planning and beyond. In: KR, pp. 600–610 (2008)
27. Gu, Y., Kiringa, I.: Model checking meets theorem proving: A situation calculus based approach. In: 11th International Workshop on Nonmonotonic Reasoning, Action, and Change (2006)
28. Hehner, E.C.R.: A Practical Theory of Programming. Texts and Monographs in Computer Science Springer (1993)

29. Hennessy, M.: The Semantics of Programming Languages: An Elementary Introduction Using Structural Operational Semantics. Wiley, New York (1990)
30. Kelly, R.F., Pearce, A.R.: Property persistence in the situation calculus. Artif. Intell. **174**(12–13), 865–888 (2010)
31. Knaster, B.: Un thèoréme sur les fonctions d'ensembles. Ann. Soc. Polon. Math. **6**, 133–134 (1928)
32. Kozen, D., Tiuryn, J.: Logics of programs. In: van Leeuwen, J. (ed.) Handbook of Theoretical Computer Science, pp. 790–840. Elsevier (1990)
33. Leivant, D.: Higher order logic. In: Handbook of Logic in Artificial Intelligence and Logic Programming, vol. 2. Clarendon Press, pp. 229–321 (1994)
34. Lespérance, Y.: On the epistemic feasibility of plans in multiagent systems specifications. In: 8th International Workshop on Intelligent Agents VIII, ATAL 2001 Seattle, WA, USA, August 1-3, 2001, Revised Papers, pp. 69–85 (2001)
35. Lespérance, Y., Levesque, H.J., Lin, F., Scherl, R.B.: Ability and knowing how in the situation calculus. Stud. Logica. **66**(1), 165–186 (2000)
36. Levesque, H.J., Reiter, R., Lespérance, Y., Lin, F., Scherl, R.B.: Golog: A logic programming language for dynamic domains. J. Log. Program. **31**(1-3), 59–83 (1997)
37. Lin, F.: A first-order semantics for Golog and ConGolog under a second-order induction axiom for situations. In: KR (2014)
38. Loeckx, J., Sieber, K.: Foundation of Program Verification. Teubner-Wiley, New York (1987)
39. Manna, Z., Pnueli, A.: The Temporal Logic of Reactive and Concurrent Systems, vol.1–2. Springer (1992)
40. Moschovakis, Y.: Elementary Induction on Abstract Structures. Amsterdam (1974)
41. Park, D.: Fixpoint induction and proofs of program properties. In: Machine Intelligence, vol. 5, pp. 59–78. Edinburgh University Press (1970)
42. Plotkin, G.: A structural approach to operational semantics. Tech. Rep. DAIMI-FN-19, Computer Science Dept. Aarhus Univ Denmark (1981)
43. Reiter, R.: The frame problem in the situation calculus: a simple solution (sometimes) and a completeness result for goal regression. In: Lifschitz, V. (ed.) Artificial Intelligence and Mathematical Theory of Computation: Papers in Honor of John McCarthy, pp. 359–380. Academic Press, San Diego (1991)
44. Reiter, R.: Knowledge in Action. Logical Foundations for Specifying and Implementing Dynamical Systems. The MIT Press (2001)
45. Sardiña, S., De Giacomo, G.: Composition of ConGolog programs. In: IJCAI, pp. 904–910 (2009)
46. Sardiña, S., De Giacomo, G., Lespérance, Y., Levesque, H.J.: On the semantics of deliberation in IndiGolog- from theory to implementation. Ann. Math. Artif. Intell. **41**(2–4), 259–299 (2004)
47. Shapiro, S., Lespérance, Y., Levesque, H.J.: The cognitive agents specification language and verification environment for multiagent systems. In: AAMAS, pp. 19–26 (2002)
48. Sohrabi, S., Prokoshyna, N., McIlraith, S.A.: Web service composition via the customization of Golog programs with user preferences. In: Conceptual Modeling: Foundations and Applications - Essays in Honor of John Mylopoulos, pp. 319–334 (2009)
49. Stirling, C.: Modal and temporal logics for processes. In: Logics for Concurrency: Structure versus Automata, no. 1043 in Lecture Notes in Computer Science, pp. 149–237. Springer (1996)
50. Tarski, A.: A lattice-theoretical fixpoint theorem and its applications. Pac. J. Math. **5**(2), 285–309 (1955)
51. Ternovskaia, E.: Inductive definability and the situation calculus. In: Transactions and Change in Logic Databases, International Seminar on Logic Databases and the Meaning of Change, Schloss Dagstuhl, Germany, September 23-27, 1996 and ILPS '97 Post-Conference Workshop on (Trans)Actions and Change in Logic Programming and Deductive Databases, (DYNAMICS'97) Port Jefferson, NY, USA, October 17, 1997, Invited Surveys and Selected Papers, pp. 227–248 (1998)
52. Ternovskaia, E.: Automata theory for reasoning about actions. In: IJCAI, pp. 153–159 (1999)
53. Wang, Y., Chang, L., Li, F., Gu, T.: Verification of branch-time property based on dynamic description logic. In: Intelligent Information Processing VII - 8th IFIP TC 12 International Conference, IIP 2014, Hangzhou, China, October 17-20, 2014, Proceedings, pp. 161–170 (2014)
54. Zarrieß, B., Claßen, J.: Verifying CTL* properties of Golog programs over local-effect actions. In: ECAI, pp. 939–944 (2014)
55. Zarrieß, B., Claßen, J.: Decidable verification of Golog programs over non-local effect actions. In: AAAI, pp. 1109–1115 (2016)