



## ACSI – Artifact-Centric Service Interoperation



### Deliverable D2.4.1

## Techniques and Tools for KAB, to Manage Action Linkage with Artifact Layer - Iteration 1

Project Acronym	ACSI	
Project Title	Artifact-Centric Service Interoperation	
Project Number	257593	
Workpackage	WP2 – Formal Based Techniques and Tools	
Lead Beneficiary	FUB	
Editor(s)	Diego Calvanese	FUB
	Giuseppe De Giacomo	UNIROMA1
	Babak Bagheri Hariri	FUB
	Riccardo De Masellis	UNIROMA1
	Domenico Lembo	UNIROMA1
	Marco Montali	FUB
Contributor(s)	Francesco Belardinelli	Imperial
	Pieter De Leenheer	Collibra
	Alin Deutsch	Univ. of California, San Diego
	Marlon Dumas	UT
	Dirk Fahland	TU/e
	Paolo Felli	UNIROMA1
	Lior Limonad	IBM Haifa
	Fabio Patrizi	UNIROMA1
Reviewer(s)	Richard Hull	IBM T.J. Watson
	Alessio Lomuscio	Imperial
Dissemination Level	PU	
Contractual Delivery Date	31/05/2012	
Actual Delivery Date	31/05/2012	
Version	1.3	

### Abstract

In this deliverable, we address the key issues for actually realizing the Knowledge and Action Base (KAB), by marrying the descriptions of the static and the dynamic aspects of the domain of interest. First, we investigate the setting of Data Centric Dynamic Systems (DCDSs), corresponding to KABs in which we have complete information on the data. We then move to a setting of incomplete information, and investigate so-called Semantic Artifacts, where the knowledge component of the KAB is realized through a description logics ontology, and where we adopt a functional view of the ontology component when dealing with the dynamics of the KAB. In both settings of DCDSs and of Semantic Artifacts, the evolution of the system is specified by means of unbounded processes, that may acquire new information from the external environment. We are interested in checking temporal properties over such systems, which are undecidable in general and turn out to be infinite state, hence not amenable to traditional model checking techniques. We study suitable restrictions on the form of the process specification that allow us to prove decidability of verification of temporal properties expressed in variants of first-order  $\mu$ -calculus, under various assumptions on the semantics of the studied systems.

## Document History

<b>Version</b>	<b>Date</b>	<b>Comments</b>
V1.0	16/4/2012	First draft with contributions from publications
V1.1	3/5/2012	Restructuring of first draft
V1.2	15/5/2012	Version for internal review
V1.3	31/5/2012	Final version, incorporating internal review comments

# Table of Contents

<b>Document History</b>	<b>3</b>
<b>List of Figures</b>	<b>6</b>
<b>I Introduction</b>	<b>7</b>
<b>II Related Work</b>	<b>11</b>
<b>1 Artifact-centric processes with no database</b>	<b>11</b>
<b>2 Artifact-centric processes with underlying database</b>	<b>11</b>
<b>3 Infinite-state systems</b>	<b>12</b>
<b>4 Processes over ontologies</b>	<b>13</b>
<b>III Data Centric Dynamic Systems</b>	<b>14</b>
<b>5 Data-Centric Dynamic Systems</b>	<b>14</b>
5.1 Data Layer . . . . .	14
5.2 Process Layer . . . . .	15
5.3 Semantics via Transition System . . . . .	15
<b>6 Verification</b>	<b>16</b>
6.1 History Preserving Mu-Calculus . . . . .	18
6.2 Persistence Preserving Mu-Calculus . . . . .	19
<b>7 Deterministic Services</b>	<b>20</b>
7.1 Semantics . . . . .	20
7.2 Run-Bounded Systems . . . . .	22
7.3 Weakly Acyclic DCDSs . . . . .	24
<b>8 Nondeterministic Services</b>	<b>26</b>
8.1 Semantics . . . . .	28
8.2 State-Bounded Systems . . . . .	28
8.3 GR-Acyclic DCDSs . . . . .	30
<b>9 Example DCDS: Travel Reimbursement System</b>	<b>34</b>
<b>10 Discussion</b>	<b>39</b>
<b>IV Semantic Artifacts</b>	<b>43</b>
<b>11 Knowledge Base Formalism</b>	<b>43</b>
<b>12 Formalizing Knowledge and Action Bases</b>	<b>44</b>
12.1 Knowledge Component of a KAB . . . . .	44
12.2 Action Component of a KAB . . . . .	44
12.3 KAB Semantics . . . . .	47

<b>13 Verification</b>	<b>48</b>
13.1 Weakly Acyclic KABs . . . . .	50
<b>14 Discussion</b>	<b>52</b>
<b>V References</b>	<b>53</b>
<b>VI Appendix</b>	<b>58</b>
<b>A Proofs for DCDSs</b>	<b>58</b>
A.1 Bisimulation Invariance for $\mu$ -calculi . . . . .	58
A.2 Deterministic Services . . . . .	61
A.3 Nondeterministic Services . . . . .	66
<b>B Proofs for Semantic Artifacts</b>	<b>75</b>
B.1 Simplified Transition System . . . . .	76
B.2 Positive Dominant . . . . .	77

## List of Figures

1	Semantics of $\mu\mathcal{L}$ . . . . .	17
2	Concrete and abstract transition systems of the DCDS with deterministic services described in Example 7.2; special relations that store the service calls results are represented using a $call \mapsto value$ notation . . . . .	22
3	Concrete and abstract transition systems of the DCDS with deterministic services described in Example 7.1; special relations that store the service calls results are represented using a $call \mapsto value$ notation . . . . .	23
4	Concrete and abstract transition systems of the run-unbounded DCDS with deterministic services $\mathcal{S}$ described in Example 7.3; special relations that store the service calls results are represented using a $call \mapsto value$ notation . . . . .	27
5	Examples of dataflow graphs for DCDSs with deterministic services; special edges are decorated with $*$ . . . . .	27
6	Concrete and abstract transition systems of the state-unbounded DCDS with nondeterministic services of Example 8.2 . . . . .	31
7	Concrete and abstract transition systems obtained when the DCDS described in Example 7.3 has nondeterministic services . . . . .	32
8	Examples of dependency graphs for DCDSs with nondeterministic services; special edges are decorated with $*$ . . . . .	32
9	Dataflow graph for request system . . . . .	36
10	Weakly-acyclic dependency graph of the audit system . . . . .	39
11	The TBox of the KAB for Example 11 . . . . .	46
12	Semantics of $\mu\mathcal{L}_A$ . . . . .	49
13	Weakly acyclic dependency graph for Example 11 . . . . .	51

## Part I

# Introduction

In this deliverable, we address the key issues for actually realizing the Knowledge and Action Base (KAB). Indeed we marry the description of the static aspects of the domain of interest (the data) with the description of its dynamic aspects (the processes). We devise techniques and tools to formally verify and reason on how the semantic state of the system and the associated data evolve during the execution of a process. First, we investigate the setting of Data Centric Dynamic Systems (DCDSs), corresponding to KABs in which we have complete information on the data, and so the semantic data component can be seen as a relational database with constraints (i.e., a finite model of first-order logic). We then move to a setting of incomplete information, and investigate so-called Semantic Artifacts, where the knowledge component of the KAB is realized through a description logics ontology. In that setting, we adopt a functional view of the ontology component when dealing with the dynamics of the KAB. This view greatly simplifies the problem of combining static aspects with dynamic ones: enabling to base the first one on logical implication and the second one on model checking. In both settings of DCDSs and of Semantic Artifacts, the evolution of the system is specified by means of unbounded processes, that may acquire new information from the external environment. We are interested in checking temporal properties over such systems, which are undecidable in general and turn out to be infinite state, hence not amenable to traditional model checking techniques. We study suitable restrictions on the form of the process specification that make it possible to adopt forms of abstraction resulting in a finite state system. Under such restrictions, and under various assumptions on the semantics of the studied systems, we are able to prove decidability of verification of temporal properties expressed in variants of first-order  $\mu$ -calculus, a very expressive temporal logic subsuming many formalisms known from the literature.

## Data Centric Dynamic Systems

Data-centric dynamic systems (DCDSs) are systems where both the process controlling the dynamics and the manipulated data are equally central. Recently such kinds of systems are increasingly attracting the interest of the scientific community. In particular, the so called artifact-centric approach to modeling business processes has emerged, with the fundamental characteristic of considering both data and processes as first-class citizens in service design and analysis [56, 40, 31, 23, 62, 1]. This holistic view of dealing with data and processes together promises to avoid the notorious discrepancy between data modeling and process modeling of more traditional approaches that consider these two aspects separately [15, 14].

DCDSs are constituted by (i) a *data layer*, which is used to hold the relevant information to be manipulated by the system, and (ii) a *process layer* formed by the invocable (*atomic*) *actions* and a process based on them. Such a process characterizes the dynamic behavior of the system. Executing an action has effects on the data manipulated by the system, on the process state, and on the information exchanged with the external world.

DCDSs deeply challenge formal verification by requiring simultaneous attention to both data and processes: indeed, on the one hand they deal with full-fledged processes and require analysis in terms of sophisticated temporal properties [30]; on the other hand, the presence of possibly unbounded data makes the usual analysis based on model checking of finite-state systems impossible in general, since, when data evolution is taken into account, the whole system becomes infinite-state.

In Part III of this report, we study *relational* DCDSs, where data are represented by a full-fledged relational database, and the process is described in terms of atomic actions that evolve the database. The execution of such actions may involve calls to external services, providing fresh data inserted into the system. As a result such systems are infinite-state in general. In particular, actions are characterized using conditional effects. Effects are specified using first-order (FO) queries to extract from the current database the objects we want to persist in the

next state, and using conjunctive queries on these objects to generate the facts that are true in the next state. In addition, to finalize the next state we call external services (function calls) that provide new information and objects coming from the external world.

On top of such a framework, we introduce powerful verification logics, which are FO variants of  $\mu$ -calculus [51, 57, 35, 21].  $\mu$ -calculus is well known to be more expressive than virtually all temporal logics used in verification, including CTL, LTL, CTL\*, PDL, and many others. Our approach is remarkably robust: while it is common to use simpler logics like CTL and LTL towards verification decidability, our decidability results hold for significantly more expressive  $\mu$ -calculus variants, and thus carry over to all these other logics. Our variants of  $\mu$ -calculus are based on first-order queries over data in the states of the DCDS, and allow for first-order quantification across states (within and across runs), though in a controlled way. No limitations whatsoever are instead put on the fixpoint formulae, which are the key element of the  $\mu$ -calculus.

In particular we consider two variants of  $\mu$ -calculus. The first variant is called  $\mu\mathcal{L}_A$ , and requires that first-order quantification across states be always bounded to the active domain of the state where the quantification is evaluated. This quantification mechanism indirectly preserves, at any point, knowledge of objects that appeared in the history so far, even if they disappeared in the meantime. The second variant, called  $\mu\mathcal{L}_P$ , restricts the first-order quantification in  $\mu\mathcal{L}_A$  by requiring that only quantified object that are still present in the current domain are of interest as we move from one state to the next. That is, knowledge of objects is preserved only if they are continuously present. For these two logics we define novel notions of bisimulation, which we exploit to prove our results.

We show that verification of both  $\mu\mathcal{L}_A$  and  $\mu\mathcal{L}_P$  is undecidable in general. In fact we get undecidability even ruling out first-order quantification and branching time. However we isolate two notable decidable cases. Specifically we start by considering service calls that return values deterministically (depending only on passed parameters). We show that verification of  $\mu\mathcal{L}_A$  properties is decidable under the assumption that the cardinality of fresh data introduced along each run is bounded (*run-bounded* DCDSs), though it need not be bounded across runs. Decidability is therefore not obvious, given that the logic permits quantification over values occurring across (potentially infinitely many) branching run continuations. Run-boundedness is a semantic property which we show is undecidable to check, but for which we propose a sufficient syntactic condition related to the notion of weak acyclicity studied in data exchange [36]. Then, we move to nondeterministic services where same-argument service calls possibly return different values at different moments in time. To exploit the results on run-bounded DCDSs in this case we would have to limit the number of service calls that can be invoked during the execution, which would be a too restrictive condition on the form of DCDSs. So we focus on the above  $\mu\mathcal{L}_P$  fragment of  $\mu\mathcal{L}_A$ . We show that if infinitely many values occur in a run but do not accumulate in the same state (our system is then called *state-bounded*) then  $\mu\mathcal{L}_P$  verification is decidable. This comes as a pleasant surprise, given that when compared to run-boundedness, state-boundedness permits an additional kind of data unboundedness (*within* the run, as opposed to only *across* runs). State-boundedness is a semantic property as well, and we show that checking it is undecidable. We then give a novel syntactic condition, “generate-recall acyclicity”, which suffices to enforce that if a service generates new values by being called an unbounded number of times, then these values cannot be accumulated (“recalled”) indefinitely.

## Knowledge and Action Bases

Recent work in business processes, services and databases is bringing forward the need of considering both data and processes as first-class citizens in process and service design [34, 64]. In particular, the so called artifact-centric approaches, which advocate a sort of middle ground between a conceptual formalization of dynamic systems and their actual implementation, are promising to be effective in practice [31]. The verification of temporal properties in the presence of data represents a significant research challenge, since data makes the system infinite-state, and neither finite-state model checking [30] nor most of the current technique for infinite-state model checking, which mostly tackle recursion [22], apply to this case. Recently, there have



been some advancements on this issue [32, 7, 13], considering a suitably constrained relational database settings.

Here we are interested in modelling the domain of interest in terms of a semantically rich language as those provided by ontological languages based on Description Logics (DLs) [4]. Such languages natively deal with incomplete knowledge in the modeled domain. This additional flexibility comes with an added cost however: in contrast with relational databases, to evaluate queries we need to resort to logical implication. Moreover incomplete information combined with the ability of evolving the system through actions results in a notoriously difficult setting [65]. In particular, due to the nature of DL assertions (which in general are not definitions but constraints on models), we get one of the most difficult kinds of domain descriptions for reasoning about actions [60], which amounts to dealing with complex forms of state constraints [48, 49]. To overcome this difficulty, virtually all effective solutions presented in the literature are based on a so-called “functional view of knowledge bases” [46]: the KB provides the ability of querying based on logical implication, and the ability of progressing it to a “new” KB through forms of updates [5, 27]. Notice that this functional view is tightly related to an epistemic interpretation of the KB [25]. Indeed our work is also related to that on Epistemic Dynamic Logic [63], and, though out of the scope of this paper, the decidability results presented here could find application in the context of that research as well.

Following this functional view of KBs, in this paper, we introduce the so-called *Knowledge and Action Bases* (KABs). A KAB is equipped with a TBox, expressed in a variant of *DL-Lite<sub>A</sub>* without UNA [26, 3], which extends the core of the Web Ontology Language OWL 2 QL<sup>1</sup> and is particularly well suited for data management. Such a TBox also captures intensional information on the domain of interest, similarly to UML class diagrams or other conceptual data models, though as a software component to be used at runtime. The KAB includes also an ABox, which acts as a storage or state. The ABox maintains the data of interest, which are accessed through by relying on query answering based on logical implication (certain answers). As a query language, we use unions of conjunctive queries, possibly composing their certain answers through full FOL constructs. This gives rise to an *epistemic query language* that asks about what is “known” by the current KB [25]. The KAB then contains *actions* whose execution changes the state of the KB, i.e., its ABox. Such actions are specified as sets of conditional effects, where conditions are (epistemic) queries over the KB and effects are expressed in terms of new ABoxes. Actions have no pre-conditions, instead a process is used to specify which actions can be executed at each step. We model such processes as condition/action rules, where the condition is again expressed as a query over the KB.

Moreover, to capture data that are acquired from external users/environments during action execution, the produced ABoxes may contain special individual terms that are represented by *Skolem terms*. These represent new unknown individuals of which some facts are known by the KB. Notice that all previous decidability results on actions over DL KBs assumed that no information is coming from outside of the system, in the sense that no new individual terms are added while executing actions [5, 27]. In this paper, instead, we allow for arbitrary introduction of new individual terms.

One important aspect is that we can infer equalities between individual terms, including Skolem terms, through functionality constraints as well as through direct assertions, possibly as results of actions. In this setting, we address the verification of temporal/dynamic properties expressed in a first-order variant of  $\mu$ -calculus [57, 61], where atomic formulas are queries over the ontology that can refer both to known and unknown individuals, and where a controlled form of quantification across states is allowed. Unsurprisingly, we show that even for very simple KABs and temporal properties, verification is undecidable. However, we also show that for a very rich class of KABs, verification is indeed decidable and reducible to finite-state model checking. To obtain this result, following [7], we rely on recent results in data exchange on the finiteness of the chase of tuple-generating dependencies [36].

<sup>1</sup><http://www.w3.org/TR/owl2-profiles/>

## Structure of the Deliverable

The deliverable is divided into six parts.

1. Part I is this introduction.
2. Part II reports on related works, putting them in correspondence to the results presented here.
3. Part III introduces Data Centric Dynamic Systems, and presents our decidability and undecidability results for the problem of verification of temporal properties expressed in variants of the  $\mu$ -calculus. It includes also a full-fledged example of a DCDS.
4. Part IV presents Semantics Artifacts and the results on verification of temporal properties in that setting.
5. Part V contains the bibliographic references.
6. Finally, Part VI is an appendix that contains the detailed proofs of the results presented in Parts III and IV.

The work reported here has given rise to several scientific papers, including the following ones: [6, 9, 39, 10, 47, 8].

## Part II

# Related Work

As will be discussed in detail in Section 10, the unrestricted artifact-centric and DCDS models have equivalent expressive capabilities. Our work is therefore closely related to work on verification of artifact-centric business processes, and provides foundational results for verification in that setting. A difference is in how each work trades off between restricting the class of business processes versus the class of properties to verify.

## 1 Artifact-centric processes with no database

Work on formal analysis of artifact-based business processes in restricted contexts has been reported in [38, 37, 15]. Properties investigated include reachability [38, 37], general temporal constraints [37], and the existence of complete execution or dead end [15]. For the variants considered in each paper, verification is generally undecidable; decidability results were obtained only under rather severe restrictions, e.g., restricting all pre-conditions to be “true” [38], restricting to bounded domains [37, 15], or restricting the pre- and post-conditions to be propositional, and thus not referring to data values [37]. [23] adopts an artifact model variation with arithmetic operations but no database. It proposes a criterion for comparing the expressiveness of specifications using the notion of *dominance*, based on the input/output pairs of business processes. Decidability relies on restricting runs to bounded length. [66] addresses the problem of the existence of a run that satisfies a temporal property, for a restricted case with no database and only propositional LTL properties. None of these works model an underlying database (and hence integrity constraints).

## 2 Artifact-centric processes with underlying database

More recently, two lines of work have considered artifact-centric processes that also model an underlying relational database. One considers branching time, the other only linear time.

**Branching time.** Our approach stems from a line of research that has started with [29] and continued with [7] and [13] in the context of artifact-centric processes. The connection between evolution of data-centric dynamic systems and data exchange that we exploit in this paper was first devised in [29]. There the dynamic system transition relation itself is described in terms of TGDs mapping the current state to the next, and the evolution of the system is essentially a form of chase. Under suitable weak acyclicity conditions such a chase terminates, thus making the DCDS transition system finite. A first-order  $\mu$ -calculus without first-order quantification across states is used as the verification formalism for which decidability is shown. Notice the role of getting new objects/values from the external environment, played here by service calls, is played there by nulls. These ideas were further developed in [7], where TGDs were replaced by action rules with the same syntax as here. Semantically however the dynamic system formalism there is deeply different: what we call here service calls are treated there as uninterpreted Skolem terms. This results in an ad-hoc interpretation of equality which sees every Skolem term as equal only to itself (as in the case of nulls [29]). The same first-order  $\mu$ -calculus without first-order quantification across states of [29] is used as the verification formalism, and a form of weak acyclicity is used as a sufficient condition for getting finite-state transition systems and decidability.

In the case of deterministic services, our framework is directly inspired by [7], though here we do interpret service calls. This decision is motivated by our goal of modeling real-life external services, for which two distinct service calls may very well return equal results, even under the

deterministic semantics (for instance if the same service is called with different arguments, or if distinct services are invoked). Interpreting service calls raises a major challenge: even under the run-bounded restriction, the concrete transition system is infinite, because it is infinitely branching. (A service call can be interpreted with any of the constants from the infinite domain.) In contrast to [7], what we show in this case is not that the concrete transition system is finite (it never is), but that it is *bisimilar* to a finite abstract transition system. This leads to a proof technique that is interesting in its own right, being based on novel notions of bisimilarity for the considered  $\mu$ -calculus variants. The reason standard bisimilarity is insufficient is that our logics  $\mu\mathcal{L}_P$  and  $\mu\mathcal{L}_A$  allow first-order quantification across states, so bisimilarity must respect the connection between values appearing both in the current and successor state. Our decision to include first-order quantification across states was motivated by the need to express liveness properties that refer to the same data at various points in time (e.g. “if student  $x$  is enrolled now and continues to be enrolled in the future, then  $x$  will eventually graduate”).

Inspired by [7], [13] builds a similar framework where actions are specified via pre- and post-conditions given as FO formulae interpreted over active domains. The verification logic considered is a first-order variant of CTL with no quantification across states. Thus, it inherits the limitations discussed above on expressibility of liveness properties. In addition, the limited temporal expressivity of CTL precludes expressing certain desirable properties such as fairness. [13] shows that under the assumption that each state has a bounded active domain, one can construct an abstract finite transition system that can be checked instead of the original concrete transition system, which is infinite-state in general. The approach is similar to the one we developed independently for nondeterministic services, however without quantification across states, standard bisimilarity suffices. As opposed to our work, the decidability of checking state-boundedness is not investigated in [13], and no sufficient syntactic conditions are proposed.

**Linear time.** Publication [34] considers an artifact model that has the same expressive capabilities as an unrestricted class of DCDS in which the infinite domain is equipped with a dense linear order, which can be mentioned in pre-, post-conditions, and properties. Runs can receive unbounded external input from an infinite domain, and this input corresponds to nondeterministic services in a DCDS. Verification is decidable even if the input accumulates in states, and runs are neither run-bounded, nor state-bounded. However, this expressive power requires restrictions that render the result incomparable to ours. First, the property language is a first-order extension of LTL, and it is shown that extension to branching time (CTL\*) leads to undecidability. Second, the formulae in pre-, post-conditions and properties access read-only and read-write database relations differently, querying the latter only in limited fashion. In essence, data can be arbitrarily accumulated in read-write relations, but these can be queried only by checking that they contain a given tuple of constants. It is shown that this restriction is tight, as even the ability to check emptiness of a read-write relation leads to undecidability. In addition, no integrity constraints are supported as it is shown that allowing a single functional dependency leads to undecidability. [32] disallows read-write relations entirely (only the artifact variables are writable), but this allows the extension of the decidability result to integrity constraints expressed as embedded dependencies with terminating chase, and to any decidable arithmetic. Again the result is incomparable to ours, as our modeling needs include read-write relations and their unrestricted querying.

### 3 Infinite-state systems

DCDSs are a particular case of infinite-state systems. Research on automatic verification of infinite-state systems has also focused on extending classical model checking techniques (e.g., see [22] for a survey). However, in much of this work the emphasis is on studying recursive control rather than data, which is either ignored or finitely abstracted. More recent work has been focusing specifically on data as a source of infinity. This includes augmenting recursive procedures with integer parameters [18], rewriting systems with data [17], Petri nets with data

associated to tokens [44], automata and logics over infinite alphabets [20, 19, 55, 33, 41, 16, 17], and temporal logics manipulating data [33]. However, the restricted use of data and the particular properties verified have limited applicability to the business process setting we target with the DCDS model.

## 4 Processes over ontologies

Verification of temporal and dynamic properties over ontologies expressed in description logics is notoriously difficult, even for very simple temporal logics. This is due to the combination of incomplete information on the data, typical of ontologies, with the ability of evolving data in the system through actions, which leads immediately to undecidability [65], if not treated with care. In particular, due to the nature of ontologies, which act as sophisticated constraints on possible instantiations, we get one of the most difficult settings for reasoning about actions [60, 48, 49]. To overcome this difficulty, virtually all effective approaches to merge ontologies and actions presented in the literature are based on a so-called “functional view” [46] of ontologies: the ontology provides the ability of querying based on logical deduction, and the ability of progressing it through actions. Such an approach is followed, for example, in the recent work reported in [5, 27]. Our framework here follows this functional view as well. However, differently from the mentioned papers, we allow for new objects to be introduced from the external environment as the process progresses. This gives rise to a potentially infinite number of objects to deal with. Taming this aspect is the crux of our decidability results discussed in Part IV.

## Part III

# Data Centric Dynamic Systems

This part of the report is organized as follows. Section 5 introduces (relational) DCDS's. Section 6 introduces verification of DCDS's and the two variants of  $\mu$ -calculus that we consider. Section 7 focuses on the analysis of DCDS's under the assumption that external service calls behave deterministically. Section 8 consider the case in which external service calls behave non-deterministically. Section 9 presents a full-fledged example of a DCDS, equipped both with deterministic and with nondeterministic service calls. Section 10 discusses the various notions introduced also linking DCDS's to artifact-centric processes.

## 5 Data-Centric Dynamic Systems

In this section, we introduce the notion of (*relational*) *data-centric dynamic system*, or simply DCDS. A DCDS is a pair  $\mathcal{S} = \langle \mathcal{D}, \mathcal{P} \rangle$  formed by two interacting layers: a *data layer*  $\mathcal{D}$  and a *process layer*  $\mathcal{P}$  over it. Intuitively, the data layer keeps all the data of interest, while the process layer modifies and evolves such data. We keep the structure of both layers to the minimum, in particular we do not distinguish between various possible components providing the data, nor those providing the subprocesses running concurrently. Indeed the framework can be further detailed in several directions, while keeping the results obtained here (cf. Section 10).

### 5.1 Data Layer

The information of interest in the modeled domain is represented in the data layer. It is constituted by a relational schema  $\mathcal{R}$  equipped with equality constraints<sup>2</sup>  $\mathcal{E}$ , e.g., to state keys of relations, and an initial database instance  $\mathcal{I}_0$ , which conforms to the relational schema and the equality constraints. The values stored in this database belong to a predefined, possibly infinite, set  $\mathcal{C}$  of constants. These constants are interpreted as themselves, blurring the distinction between constants and values. We will use the two terms interchangeably.

Given a database instance  $\mathcal{I}$ , its active domain  $\text{ADOM}(\mathcal{I})$  is the subset of  $\mathcal{C}$  such that  $c \in \text{ADOM}(\mathcal{I})$  if and only if  $c$  occurs in  $\mathcal{I}$ .

A *data layer* is a tuple  $\mathcal{D} = \langle \mathcal{C}, \mathcal{R}, \mathcal{E}, \mathcal{I}_0 \rangle$  where:

- $\mathcal{C}$  is a countably infinite set of constants/values.
- $\mathcal{R} = \{R_1, \dots, R_n\}$  is a database schema, constituted by a finite set of relation schemas.
- $\mathcal{E}$  is a finite set  $\{\mathcal{E}_1, \dots, \mathcal{E}_m\}$  of equality constraints. Each  $\mathcal{E}_i$  has the form

$$Q_i \rightarrow \bigwedge_{j=1, \dots, k} z_{ij} = y_{ij},$$

where  $Q_i$  is a domain independent FO query over  $\mathcal{R}$  using constants from the active domain  $\text{ADOM}(\mathcal{I}_0)$  of  $\mathcal{I}_0$  and whose free variables are  $\vec{x}$ , and  $z_{ij}$  and  $y_{ij}$  are either variables in  $\vec{x}$  or constants in  $\text{ADOM}(\mathcal{I}_0)$ .<sup>3</sup>

- $\mathcal{I}_0$  is a database instance that represents the initial state of the data layer, which conforms to the schema  $\mathcal{R}$  and *satisfies* the constraints  $\mathcal{E}$ : namely, for each constraint  $Q_i \rightarrow \bigwedge_{j=1, \dots, k} z_{ij} = y_{ij}$  and for each tuple (i.e., substitution for the free variables)  $\theta \in \text{ans}(Q_i, \mathcal{I})$ , it holds that  $z_{ij}\theta = y_{ij}\theta$ .<sup>4</sup>

<sup>2</sup>Other kinds of constraints can also be included without affecting the results reported here (cf. Section 10).

<sup>3</sup>For convenience, and without loss of generality, we assume that all constants used inside formulae appear in  $\mathcal{I}_0$ .

<sup>4</sup>We use the notation  $t\theta$  (resp.,  $\varphi\theta$ ) to denote the term (resp., the formula) obtained by applying the substitution  $\theta$  to  $t$  (resp.,  $\varphi$ ). Furthermore, given a FO query  $Q$  and a database instance  $\mathcal{I}$ , the *answer*  $\text{ans}(Q, \mathcal{I})$  to  $Q$  over  $\mathcal{I}$  is the set of assignments  $\theta$  from the free variables of  $Q$  to the domain of  $\mathcal{I}$ , such that  $\mathcal{I} \models Q\theta$ . We treat  $Q\theta$  as a boolean query, and with some abuse of notation, we say  $\text{ans}(Q\theta, \mathcal{I}) \equiv \text{true}$  if and only if  $\mathcal{I} \models Q\theta$ .

## 5.2 Process Layer

The process layer constitutes the progression mechanism for the DCDS. We assume that at every time the current instance of the data layer can be arbitrarily queried, and can be updated through action executions, possibly involving external service calls to get new values from the environment. Hence the process layer is composed of three main notions: actions, which are the atomic progression steps for the data layer; external services, which can be called during the execution of actions; and processes, which are essentially nondeterministic programs that use actions as atomic instructions. While we require the execution of actions to be sequential, we do not impose any such constraints on processes, which in principle can be formed by several concurrent branches, including fork, join, and so on. Concurrency is to be interpreted by interleaving and hence reduced to nondeterminism, as often done in formal verification [11, 35]. There can be many ways to provide the control flow specification for processes. Here we adopt a simple rule-based mechanism, but our results can be immediately generalized to any process formalism whose processes control flow is finite-state. Notice that this does not imply that the transition system associated to a process over the data layer is finite-state as well, since the data manipulated in the data layer may grow over time in an unbounded way.

Formally, a process layer  $\mathcal{P}$  over a data layer  $\mathcal{D} = \langle \mathcal{C}, \mathcal{R}, \mathcal{E}, \mathcal{I}_0 \rangle$ , is a tuple  $\mathcal{P} = \langle \mathcal{F}, \mathcal{A}, \varrho \rangle$  where:

- $\mathcal{F}$  is a finite set of *functions*, each representing the interface to an *external service*. Such services can be called, and as a result the function is activated and the answer is produced. How the result is actually computed is *unknown* to the DCDS since the services are indeed external.
- $\mathcal{A}$  is a finite set of *actions*, whose execution progresses the data layer, and may involve external service calls.
- $\varrho$  is a finite set of *condition-action rules* that form the specification of the overall *process*, which tells at any moment which actions can be executed.

An *action*  $\alpha \in \mathcal{A}$  has the form

$$\alpha(p_1, \dots, p_n) : \{e_1, \dots, e_m\},$$

where: (i)  $\alpha(p_1, \dots, p_n)$  is the *signature* of the action, constituted by a name  $\alpha$  and a sequence  $p_1, \dots, p_n$  of *input parameters* that need to be substituted with values for the execution of the action, and (ii)  $\{e_1, \dots, e_m\}$ , also denoted as  $\text{EFFECT}(\alpha)$ , is a set of *effect specifications*, whose specified effects are assumed to take place simultaneously. Each  $e_i$  has the form  $q_i^+ \wedge Q_i^- \rightsquigarrow E_i$ , where:

- $q_i^+ \wedge Q_i^-$  is a query over  $\mathcal{R}$  whose terms are variables  $\vec{x}$ , action parameters, and constants from  $\text{ADOM}(\mathcal{I}_0)$ . The query  $q_i^+$  is a UCQ, and the query  $Q_i^-$  is an arbitrary FO formula whose free variables are included in those of  $q_i^+$ . Intuitively,  $q_i^+$  selects the tuples to instantiate the effect, and  $Q_i^-$  filters away some of them.
- $E_i$  is the effect, i.e., a set of facts for  $\mathcal{R}$ , which includes as terms: terms in  $\text{ADOM}(\mathcal{I}_0)$ , input parameters, free variables of  $q_i^+$ , and in addition Skolem terms formed by applying a function  $f \in \mathcal{F}$  to one of the previous kinds of terms. Such Skolem terms involving functions represent external service calls and are interpreted so as to return a value chosen by an external user/environment when executing the action.

The *process*  $\varrho$  is a finite set of *condition-action rules*, of the form  $Q \mapsto \alpha$ , where  $\alpha$  is an action in  $\mathcal{A}$  and  $Q$  is a FO query over  $\mathcal{R}$  whose free variables are exactly the parameters of  $\alpha$ , and whose other terms can be either quantified variables or constants in  $\text{ADOM}(\mathcal{I}_0)$ .

For a detailed example of a DCDS we refer to Section 9.

## 5.3 Semantics via Transition System

The semantics of a DCDS is defined in terms of a possibly infinite transition system whose states are labeled by databases. Such a transition system represents all possible computations

that the process layer can do on the data layer. A transition system  $\Upsilon$  is a tuple of the form  $\langle \Delta, \mathcal{R}, \Sigma, s_0, db, \Rightarrow \rangle$ , where:

- $\Delta$  is a countably infinite set of values;
- $\mathcal{R}$  is a database schema;
- $\Sigma$  is a set of states;
- $s_0 \in \Sigma$  is the initial state;
- $db$  is a function that, given a state  $s \in \Sigma$ , returns the database associated to  $s$ , which is made up of values in  $\Delta$  and conforms to  $\mathcal{R}$ ;
- $\Rightarrow \subseteq \Sigma \times \Sigma$  is a transition relation between pairs of states.

In order to precisely build the transition system associated to a DCDS, we need to better characterize the behavior of the external services, which are called in the effects of actions. This is done in Sections 7 and 8.

## 6 Verification

To specify dynamic properties over a DCDS, we use  $\mu$ -calculus [35, 61, 21], one of the most powerful temporal logics for which model checking has been investigated in the finite-state setting. Indeed, such a logic is able to express both linear time logics such as LTL and PSL, and branching time logics such as CTL and CTL\* [30]. The main characteristic of  $\mu$ -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 properties defined by induction or co-induction. This is the reason why virtually all logics used in verification can be considered as fragments of  $\mu$ -calculus. From a technical viewpoint,  $\mu$ -calculus separates local properties, i.e., properties asserted on the current state or on states that are immediate successors of the current one, and properties that talk about states that are arbitrarily far away from the current one [21]. The latter are expressed through the use of fixpoints.

In this work, we use a first-order variant of the  $\mu$ -calculus [57], called  $\mu\mathcal{L}$  and defined as follows:

$$\Phi ::= Q \mid \neg\Phi \mid \Phi_1 \wedge \Phi_2 \mid \exists x.\Phi \mid \langle - \rangle\Phi \mid Z \mid \mu Z.\Phi$$

where  $Q$  is a possibly open FO query, and  $Z$  is a second order predicate variable (of arity 0). We make use of the following abbreviations:  $\forall x.\Phi = \neg(\exists x.\neg\Phi)$ ,  $\Phi_1 \vee \Phi_2 = \neg(\neg\Phi_1 \wedge \neg\Phi_2)$ ,  $[-]\Phi = \neg\langle - \rangle\neg\Phi$ , and  $\nu Z.\Phi = \neg\mu Z.\neg\Phi[Z/\neg Z]$ .

As usual in  $\mu$ -calculus, formulae of the form  $\mu Z.\Phi$  (and  $\nu Z.\Phi$ ) must obey to the *syntactic monotonicity* of  $\Phi$  wrt  $Z$ , which states that every occurrence of the variable  $Z$  in  $\Phi$  must be within the scope of an even number of negation symbols. This ensures that the least fixpoint  $\mu Z.\Phi$  (as well as the greatest fixpoint  $\nu Z.\Phi$ ) always exists.

Since  $\mu\mathcal{L}$  also contains formulae with both individual and predicate free variables, given a transition system  $\Upsilon$ , we introduce an individual variable valuation  $v$ , i.e., a mapping from individual variables  $x$  to  $\Delta$ , and a predicate variable valuation  $V$ , i.e., a mapping from predicate variables  $Z$  to subsets of  $\Sigma$ . With these three notions in place, we assign meaning to formulae by associating to  $\Upsilon$ ,  $v$ , and  $V$  an *extension function*  $(\cdot)_{v,V}^{\Upsilon}$ , which maps formulae to subsets of  $\Sigma$ . Formally, the extension function  $(\cdot)_{v,V}^{\Upsilon}$  is defined inductively as shown in Figure 1.

Intuitively, the extension function  $(\cdot)_{v,V}^{\Upsilon}$  assigns to such constructs the following meaning:

- The boolean connectives and quantification of individuals have the expected meaning.



$$\begin{aligned}
(Q)_{v,V}^{\Upsilon} &= \{s \in \Sigma \mid \text{ans}(Qv, \text{db}(s))\} \\
(\neg\Phi)_{v,V}^{\Upsilon} &= \Sigma \setminus (\Phi)_{v,V}^{\Upsilon} \\
(\Phi_1 \wedge \Phi_2)_{v,V}^{\Upsilon} &= (\Phi_1)_{v,V}^{\Upsilon} \cap (\Phi_2)_{v,V}^{\Upsilon} \\
(\exists x.\Phi)_{v,V}^{\Upsilon} &= \{s \in \Sigma \mid \exists t.t \in \Delta \text{ and } s \in (\Phi)_{v[x/t],V}^{\Upsilon}\} \\
(\langle - \rangle\Phi)_{v,V}^{\Upsilon} &= \{s \in \Sigma \mid \exists s'.s \Rightarrow s' \text{ and } s' \in (\Phi)_{v,V}^{\Upsilon}\} \\
(Z)_{v,V}^{\Upsilon} &= V(Z) \\
(\mu Z.\Phi)_{v,V}^{\Upsilon} &= \bigcap \{\mathcal{S} \subseteq \Sigma \mid (\Phi)_{v,V[Z/\mathcal{S}]}^{\Upsilon} \subseteq \mathcal{S}\}
\end{aligned}$$

**Figure 1 – Semantics of  $\mu\mathcal{L}$** 

- The extension of  $\langle - \rangle\Phi$  consists of the states  $s$  such that for *some* state  $s'$  with  $s \Rightarrow s'$ , we have that  $\Phi$  holds in  $s'$ , while the extension of  $[-]\Phi$  consists of the states  $s$  such that for *all* states  $s'$  with  $s \Rightarrow s'$ , we have that  $\Phi$  holds in  $s'$ .
- The extension of  $\mu X.\Phi$  is the *smallest subset*  $\mathcal{S}_\mu$  of  $\Sigma$  such that, assigning to  $Z$  the extension  $\mathcal{S}_\mu$ , the resulting extension of  $\Phi$  is contained in  $\mathcal{S}_\mu$ . That is, the extension of  $\mu X.\Phi$  is the *least fixpoint* of the operator  $(\Phi)_{v,V[Z/\mathcal{S}]}^{\Upsilon}$  (here  $V[Z/\mathcal{S}]$  denotes the predicate valuation obtained from  $V$  by forcing the valuation of  $Z$  to be  $\mathcal{S}$ ).
- Similarly, the extension of  $\nu X.\Phi$  is the *greatest subset*  $\mathcal{S}_\nu$  of  $\Sigma$  such that, assigning to  $X$  the extension  $\mathcal{S}_\nu$ , the resulting extension of  $\Phi$  contains  $\mathcal{S}_\nu$ . That is, the extension of  $\nu X.\Phi$  is the *greatest fixpoint* of the operator  $(\Phi)_{v,V[X/\mathcal{S}]}^{\Upsilon}$ . Formally,  $(\nu Z.\Phi)_{v,V}^{\Upsilon} = \bigcup \{\mathcal{S} \subseteq \Sigma \mid \mathcal{S} \subseteq (\Phi)_{v,V[Z/\mathcal{S}]}^{\Upsilon}\}$ .

**Example 6.1.** An example of  $\mu\mathcal{L}$  formula is:

$$\exists x_1, \dots, x_n. \bigwedge_{i \neq j} x_i \neq x_j \wedge \bigwedge_{i \in \{1, \dots, n\}} \mu Z. [\text{Stud}(x_i) \vee \langle - \rangle Z] \quad (1)$$

The formula asserts that there are at least  $n$  distinct objects/values, each of which eventually denotes a student along some execution path. Notice that the formula does not imply that all of these students will be in the same state, nor that they will all occur in a single run. It only says that in the entire transition systems there are (at least)  $n$  distinct students. ■

When  $\Phi$  is a closed formula,  $(\Phi)_{v,V}^{\Upsilon}$  depends neither on  $v$  nor on  $V$ , and we denote the extension of  $\Phi$  simply by  $(\Phi)^{\Upsilon}$ . We say that a closed formula  $\Phi$  holds in a state  $s \in \Sigma$  if  $s \in (\Phi)^{\Upsilon}$ . In this case, we write  $\Upsilon, s \models \Phi$ . We say that a closed formula  $\Phi$  holds in  $\Upsilon$ , denoted by  $\Upsilon \models \Phi$ , if  $\Upsilon, s_0 \models \Phi$ , where  $s_0$  is the initial state of  $\Upsilon$ . We call *model checking* verifying whether  $\Upsilon \models \Phi$  holds.

In particular we are interested in formally verifying properties of a DCDS. Given the transition system  $\Upsilon_{\mathcal{S}}$  of a DCDS  $\mathcal{S}$  and a dynamic property  $\Phi$  expressed in  $\mu\mathcal{L}$ ,<sup>5</sup> we say that  $\mathcal{S}$  *verifies*  $\Phi$  if

$$\Upsilon_{\mathcal{S}} \models \Phi.$$

The challenging point is that  $\Upsilon_{\mathcal{S}}$  is in general infinite state, so we would like to devise a finite-state transition system which is a faithful abstraction of  $\Upsilon_{\mathcal{S}}$ , in the sense that it preserves the truth value of all  $\mu\mathcal{L}$  formulae. Unfortunately, this program is doomed right from the start if we insist on using full  $\mu\mathcal{L}$  as the verification formalism. Indeed formulae of the form (1) defeat any kind of finite-state transition system. So next we introduce two interesting sublogics of  $\mu\mathcal{L}$  that serve better our objective.

<sup>5</sup>We remind the reader that, without loss of generality, we assume that all constants used inside formulae  $\Phi$  appear in the initial database instance of the DCDS.

## 6.1 History Preserving Mu-Calculus

The first fragment of  $\mu\mathcal{L}$  that we consider is  $\mu\mathcal{L}_A$ , which is characterized by the assumption that quantification over individuals is restricted to individuals that are present in the current database. To enforce such a restriction, we introduce a special predicate  $\text{LIVE}(x)$ , which states that  $x$  belongs to the current active domain. The logic  $\mu\mathcal{L}_A$  is defined as follows:

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

We make use of the usual abbreviation, including  $\forall x.\text{LIVE}(x) \rightarrow \Phi = \neg(\exists x.\text{LIVE}(x) \wedge \neg\Phi)$ . Formally, the extension function  $(\cdot)_{v,V}^\Upsilon$  is defined inductively as in Figure 12, with the new special predicate  $\text{LIVE}(x)$  interpreted as follows:

$$(\text{LIVE}(x))_{v,V}^\Upsilon = \{s \in \Sigma \mid x/d \in v \text{ implies } d \in \text{ADOM}(db(s))\}$$

**Example 6.2.** As an example, consider the following  $\mu\mathcal{L}_A$  formula:

$$\begin{aligned} \nu X.(\forall x.\text{LIVE}(x) \wedge \text{Stud}(x) \rightarrow \\ \mu Y.(\exists y.\text{LIVE}(y) \wedge \text{Grad}(x, y) \vee \langle - \rangle Y) \wedge [-]X), \end{aligned}$$

which states that, along every path, it is always true, for each student  $x$ , that there exists an evolution that eventually leads to a graduation of the student (with some final mark  $y$ ). ■

We are going to show that under suitable conditions we can get a faithful finite abstraction for a DCDS that preserves all formulae of  $\mu\mathcal{L}_A$ , and hence enables us in principle to use standard model checking techniques. Towards this goal, we introduce a notion of bisimulation that is suitable for the kind of transition systems we consider here. In particular, we have to take into account that the two transition systems are over different data domains, and hence we have to consider the correspondence between the data in the two transition systems and how such data evolve over time. To do so, we introduce the following notions.

Given two domains  $\Delta_1$  and  $\Delta_2$ , a *partial bijection*  $h$  between  $\Delta_1$  and  $\Delta_2$  is a bijection between a subset of  $\Delta_1$  and  $\Delta_2$ . Given a partial function  $f : S \rightarrow S'$ , we denote with  $\text{DOM}(f)$  the domain of  $f$ , i.e., the set of elements in  $S$  on which  $f$  is defined, and with  $\text{IM}(f)$  the image of  $f$ , i.e., the set of elements  $s'$  in  $S'$  such that  $s' = f(s)$  for some  $s \in S$ . A partial bijection  $h'$  *extends*  $h$  if  $\text{DOM}(h) \subseteq \text{DOM}(h')$  (or equivalently  $\text{IM}(h) \subseteq \text{IM}(h')$ ) and  $h'(x) = h(x)$  for all  $x \in \text{DOM}(h)$  (or equivalently  $h'^{-1}(y) = h^{-1}(y)$  for all  $y \in \text{IM}(h)$ ). Let  $db_1$  and  $db_2$  be two databases over two domains  $\Delta_1$  and  $\Delta_2$  respectively, both conforming to the same schema  $\mathcal{R}$ . We say that a partial bijection  $h$  *induces an isomorphism* between  $db_1$  and  $db_2$  if  $\text{ADOM}(db_1) \subseteq \text{DOM}(h)$ ,  $\text{ADOM}(db_2) \subseteq \text{IM}(h)$ , and  $h$  projected on  $\text{ADOM}(db_1)$  is an isomorphism between  $db_1$  and  $db_2$ .

Given two transition systems  $\Upsilon_1 = \langle \Delta_1, \mathcal{R}, \Sigma_1, s_{01}, db_1, \Rightarrow_1 \rangle$  and  $\Upsilon_2 = \langle \Delta_2, \mathcal{R}, \Sigma_2, s_{02}, db_2, \Rightarrow_2 \rangle$ , and the set  $H$  of partial bijections between  $\Delta_1$  and  $\Delta_2$ , a *history preserving bisimulation* between  $\Upsilon_1$  and  $\Upsilon_2$  is a relation  $\mathcal{B} \subseteq \Sigma_1 \times H \times \Sigma_2$  such that  $\langle s_1, h, s_2 \rangle \in \mathcal{B}$  implies that:

1.  $h$  is a partial bijection between  $\Delta_1$  and  $\Delta_2$  that induces an isomorphism between  $db_1(s_1)$  and  $db_2(s_2)$ ;
2. for each  $s'_1$ , if  $s_1 \Rightarrow_1 s'_1$  then there is an  $s'_2$  with  $s_2 \Rightarrow_2 s'_2$  and a bijection  $h'$  that extends  $h$ , such that  $\langle s'_1, h', s'_2 \rangle \in \mathcal{B}$ .
3. for each  $s'_2$ , if  $s_2 \Rightarrow_2 s'_2$  then there is an  $s'_1$  with  $s_1 \Rightarrow_1 s'_1$  and a bijection  $h'$  that extends  $h$ , such that  $\langle s'_1, h', s'_2 \rangle \in \mathcal{B}$ .

A state  $s_1 \in \Sigma_1$  is *history preserving bisimilar* to  $s_2 \in \Sigma_2$  wrt a partial bijection  $h$ , written  $s_1 \approx_h s_2$ , if there exists a history preserving bisimulation  $\mathcal{B}$  between  $\Upsilon_1$  and  $\Upsilon_2$  such that  $\langle s_1, h, s_2 \rangle \in \mathcal{B}$ . A state  $s_1 \in \Sigma_1$  is *history preserving bisimilar* to  $s_2 \in \Sigma_2$ , written  $s_1 \approx s_2$ , if there exists a partial bijection  $h$  and a history preserving bisimulation  $\mathcal{B}$  between  $\Upsilon_1$  and  $\Upsilon_2$

such that  $\langle s_1, h, s_2 \rangle \in \mathcal{B}$ . A transition system  $\Upsilon_1$  is *history preserving bisimilar* to  $\Upsilon_2$ , written  $\Upsilon_1 \approx \Upsilon_2$ , if there exists a partial bijection  $h_0$  and a history preserving bisimulation  $\mathcal{B}$  between  $\Upsilon_1$  and  $\Upsilon_2$  such that  $\langle s_{01}, h_0, s_{02} \rangle \in \mathcal{B}$ . The next theorem gives us the classical invariance result of  $\mu$ -calculus wrt bisimulation, in our setting.

**Theorem 6.1.** *Consider two transition systems  $\Upsilon_1$  and  $\Upsilon_2$  such that  $\Upsilon_1 \approx \Upsilon_2$ . Then for every  $\mu\mathcal{L}_A$  closed formula  $\Phi$ , we have:*

$$\Upsilon_1 \models \Phi \text{ if and only if } \Upsilon_2 \models \Phi.$$

## 6.2 Persistence Preserving Mu-Calculus

The second fragment of  $\mu\mathcal{L}$  that we consider is  $\mu\mathcal{L}_P$ , which further restricts  $\mu\mathcal{L}_A$  by requiring that individuals over which we quantify must continuously persist along the system evolution for the quantification to take effect.

With a slight abuse of notation, in the following we write  $\text{LIVE}(x_1, \dots, x_n) = \bigwedge_{i \in \{1, \dots, n\}} \text{LIVE}(x_i)$ .

The logic  $\mu\mathcal{L}_P$  is defined as follows:

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

where  $Q$  is a possibly open FO query,  $Z$  is a second order predicate variable, and the following assumption holds: in  $\langle - \rangle (\text{LIVE}(\vec{x}) \wedge \Phi)$  and  $[-] (\text{LIVE}(\vec{x}) \wedge \Phi)$ , the variables  $\vec{x}$  are exactly the free variables of  $\Phi$ , with the proviso that we substitute to each bounded predicate variable  $Z$  in  $\Phi$  its bounding formula  $\mu Z. \Phi'$ . We use the usual abbreviations, including:  $\langle - \rangle (\text{LIVE}(\vec{x}) \rightarrow \Phi) = \neg [-] (\text{LIVE}(\vec{x}) \wedge \neg\Phi)$  and  $[-] (\text{LIVE}(\vec{x}) \rightarrow \Phi) = \neg \langle - \rangle (\text{LIVE}(\vec{x}) \wedge \neg\Phi)$ . Intuitively, the use of  $\text{LIVE}(\cdot)$  in  $\mu\mathcal{L}_P$  ensures that individuals are only considered if they persist along the system evolution, while the evaluation of a formula with individuals that are not present in the current database trivially leads to false or true (depending on the use of negation).

**Example 6.3.** Getting back to the example above, its variant in  $\mu\mathcal{L}_P$  is

$$\begin{aligned} & \nu X. (\forall x. \text{LIVE}(x) \wedge \text{Stud}(x) \rightarrow \\ & \quad \mu Y. (\exists y. \text{LIVE}(y) \wedge \text{Grad}(x, y) \vee \langle - \rangle (\text{LIVE}(x) \wedge Y)) \wedge [-] X) \end{aligned}$$

which states that, along every path, it is always true, for each student  $x$ , that there exists an evolution in which  $x$  persists in the database until she eventually graduates (with some final mark  $y$ ). Formula

$$\begin{aligned} & \nu X. (\forall x. \text{LIVE}(x) \wedge \text{Stud}(x) \rightarrow \\ & \quad \mu Y. (\exists y. \text{LIVE}(y) \wedge \text{Grad}(x, y) \vee \langle - \rangle (\text{LIVE}(x) \rightarrow Y)) \wedge [-] X) \end{aligned}$$

instead states that, along every path, it is always true, for each student  $x$ , that there exists an evolution in which either  $x$  is not persisted, or eventually graduates (with final mark  $y$ ). ■

The bisimulation relation that captures  $\mu\mathcal{L}_P$  is as follows. Given two transition systems  $\Upsilon_1 = \langle \Delta_1, \mathcal{R}, \Sigma_1, s_{01}, db_1, \Rightarrow_1 \rangle$  and  $\Upsilon_2 = \langle \Delta_2, \mathcal{R}, \Sigma_2, s_{02}, db_2, \Rightarrow_2 \rangle$ , and the set  $H$  of partial bijections between  $\Delta_1$  and  $\Delta_2$ , a *persistence preserving bisimulation* between  $\Upsilon_1$  and  $\Upsilon_2$  is a relation  $\mathcal{B} \subseteq \Sigma_1 \times H \times \Sigma_2$  such that  $\langle s_1, h, s_2 \rangle \in \mathcal{B}$  implies that:

1.  $h$  is an isomorphism between  $db_1(s_1)$  and  $db_2(s_2)$ ;<sup>6</sup>
2. for each  $s'_1$ , if  $s_1 \Rightarrow_1 s'_1$  then there exists an  $s'_2$  with  $s_2 \Rightarrow_2 s'_2$  and a bijection  $h'$  that extends  $h|_{\text{ADOM}(db_1(s_1)) \cap \text{ADOM}(db_1(s'_1))}$ , such that  $\langle s'_1, h', s'_2 \rangle \in \mathcal{B}$ ;<sup>7</sup>

<sup>6</sup>Notice that this implies  $\text{DOM}(h) = \text{ADOM}(db_1(s_1))$  and  $\text{IM}(h) = \text{ADOM}(db_2(s_2))$ .

<sup>7</sup> Given a set  $D$ , we denote by  $f|_D$  the *restriction* of  $f$  to  $D$ , i.e.,  $\text{DOM}(f|_D) = \text{DOM}(f) \cap D$ , and  $f|_D(x) = f(x)$  for every  $x \in \text{DOM}(f) \cap D$ .

3. for each  $s'_2$ , if  $s_2 \Rightarrow_2 s'_2$  then there exists an  $s'_1$  with  $s_1 \Rightarrow_1 s'_1$  and a bijection  $h'$  that extends  $h|_{\text{ADOM}(db_1(s_1)) \cap \text{ADOM}(db_1(s'_1))}$ , such that  $\langle s'_1, h', s'_2 \rangle \in \mathcal{B}$ .

We say that a state  $s_1 \in \Sigma_1$  is *persistence preserving bisimilar* to  $s_2 \in \Sigma_2$  wrt a *partial bijection*  $h$ , written  $s_1 \sim_h s_2$ , if there exists a persistence preserving bisimulation  $\mathcal{B}$  between  $\Upsilon_1$  and  $\Upsilon_2$  such that  $\langle s_1, h, s_2 \rangle \in \mathcal{B}$ . A state  $s_1 \in \Sigma_1$  is *persistence preserving bisimilar* to  $s_2 \in \Sigma_2$ , written  $s_1 \sim s_2$ , if there exists a partial bijection  $h$  and a persistence preserving bisimulation  $\mathcal{B}$  between  $\Upsilon_1$  and  $\Upsilon_2$  such that  $\langle s_1, h, s_2 \rangle \in \mathcal{B}$ . A transition system  $\Upsilon_1$  is *persistence preserving bisimilar* to  $\Upsilon_2$ , written  $\Upsilon_1 \sim \Upsilon_2$ , if there exists a partial bijection  $h_0$  and a persistence preserving bisimulation  $\mathcal{B}$  between  $\Upsilon_1$  and  $\Upsilon_2$  such that  $\langle s_{01}, h_0, s_{02} \rangle \in \mathcal{B}$ . The next theorem shows the invariance of  $\mu\mathcal{L}_P$  under this notion of bisimulation.

**Theorem 6.2.** *Consider two transition systems  $\Upsilon_1$  and  $\Upsilon_2$  such that  $\Upsilon_1 \sim \Upsilon_2$ . Then for every  $\mu\mathcal{L}_P$  closed formula  $\Phi$ , we have:*

$$\Upsilon_1 \models \Phi \text{ if and only if } \Upsilon_2 \models \Phi.$$

## 7 Deterministic Services

Now we turn back to the semantics of DCDSs, and analyze them under the assumption that external services behave deterministically. This means that the evaluation of functions  $f \in \mathcal{F}$ , representing the service interfaces in the process layer, is independent from the moment in which the function is called: whenever an external service is called twice with the same parameters, it must return the same value. So, for example, if the function invocation  $f(a)$  returned  $b$  at a certain time, then in all successive moments the call  $f(a)$  will return  $b$  again. In particular, *stateless* services can be modeled with deterministic service calls.

Under this characterization of the services we can now define the transition system of a DCDS. We call such a transition system “concrete” transition system to avoid confusion with an “abstract” transition system that we are going to introduce for our verification technique.

### 7.1 Semantics

Let  $\mathcal{S} = \langle \mathcal{D}, \mathcal{P} \rangle$  be a DCDS with data layer  $\mathcal{D} = \langle \mathcal{C}, \mathcal{R}, \mathcal{E}, \mathcal{I}_0 \rangle$  and process layer  $\mathcal{P} = \langle \mathcal{F}, \mathcal{A}, \varrho \rangle$ .

First we focus on what is needed to characterize the states of the concrete transition system. One such state obviously needs to maintain the current instance of the data layer. This instance is a database made up of constants in  $\mathcal{C}$ , which conforms to the schema  $\mathcal{R}$  and satisfies the equality constraints in  $\mathcal{E}$ . Together with the current instance, however, we also need to remember all answers we had so far when calling the external services.

To meet the requirement that service calls behave deterministically, the states of the transition system keep track of all results of the service calls made so far, in the form of equalities between Skolem terms involving functions in  $\mathcal{F}$  and having as arguments constants and returned values in  $\mathcal{C}$ .<sup>8</sup> More precisely, we define the set of (Skolem terms representing) service calls as  $\text{SC} = \{f(v_1, \dots, v_n) \mid f/n \in \mathcal{F} \text{ and } \{v_1, \dots, v_n\} \subseteq \mathcal{C}\}$ , where  $f/n$  stands for a function  $f$  arity  $n$ . Then we introduce a *service call map*, which is a partial function  $\mathcal{M} : \text{SC} \rightarrow \mathcal{C}$ .

Now we are ready to formally define states of the concrete transition system. A *concrete state*, or simply state, is a pair  $\langle \mathcal{I}, \mathcal{M} \rangle$ , where  $\mathcal{I}$  is a relational instance of  $\mathcal{R}$  over  $\mathcal{C}$  satisfying each equality constraint in  $\mathcal{E}$ , and  $\mathcal{M}$  is a service call map. The *initial concrete state* is  $\langle \mathcal{I}_0, \emptyset \rangle$ .

Next we look at the result of executing an action in a state. For this it is convenient to denote the database instance  $\mathcal{M}(E)$  obtained by applying a service call map  $\mathcal{M}$  to a set  $E$  of facts including only constants in  $\mathcal{C}$  or terms in  $\text{DOM}(\mathcal{M})$ . Namely, we define  $\mathcal{M}(E)$  as

<sup>8</sup>Notice that, we have no knowledge of the specific functions adopted by the external services, and we simply assume that such functions return some value from  $\mathcal{C}$ . We are going to have different executions of the system corresponding to each way to assign values to the Skolem terms representing the service calls.

the application of  $\mathcal{M}$  to all the terms appearing in  $E$  where constants are preserved. Formally,  $\mathcal{M}(E) = \{R(c_1 \dots, c_n) \mid R(t_1, \dots, t_n) \in E \text{ and } c_i = t_i \text{ if } t_i \in \mathcal{C} \text{ and } \mathcal{M}(t_i) = c_i \text{ if } t_i \in \text{DOM}(\mathcal{M}) \text{ for } i \in \{1, \dots, n\}\}$ .

Let  $\alpha$  be an action in  $\mathcal{A}$  of the form  $\alpha(p_1, \dots, p_m) : \{e_1, \dots, e_m\}$  with  $e_i = q_i^+ \wedge Q_i^- \rightsquigarrow E_i$ . The parameters for  $\alpha$  are guarded by the condition-action rule  $Q \mapsto \alpha$  in  $\varrho$ . Let  $\sigma$  be a substitution for the input parameters  $p_1, \dots, p_m$  with values taken from  $\mathcal{C}$ . We say that  $\sigma$  is *legal* for  $\alpha$  in state  $\langle \mathcal{I}, \mathcal{M} \rangle$  if  $\langle p_1, \dots, p_m \rangle \sigma \in \text{ans}(Q, \mathcal{I})$ .

**Concrete action execution.** To capture what happens when  $\alpha$  is executed in a state using a substitution  $\sigma$  for its parameters, we introduce a transition relation  $\text{EXEC}_{\mathcal{S}}$  between states, called *concrete execution* of  $\alpha\sigma$ , such that  $\langle \langle \mathcal{I}, \mathcal{M} \rangle, \alpha\sigma, \langle \mathcal{I}', \mathcal{M}' \rangle \rangle \in \text{EXEC}_{\mathcal{S}}$  if the following holds:

1.  $\sigma$  is a legal parameter assignment for  $\alpha$  in state  $\langle \mathcal{I}, \mathcal{M} \rangle$ ,
2.  $\mathcal{M}' = \text{SERVICECALLS}(\mathcal{I}, \alpha\sigma, \mathcal{M})$ ,
3.  $\mathcal{I}' = \mathcal{M}'(\text{DO}(\mathcal{I}, \alpha\sigma))$ , and
4.  $\mathcal{I}'$  satisfies  $\mathcal{E}$ ,

where  $\text{DO}()$  and  $\text{SERVICECALLS}()$  are defined as follows.

$$\text{DO}(\mathcal{I}, \alpha\sigma) = \bigcup_{q_i^+ \wedge Q_i^- \rightsquigarrow E_i \in \text{EFFECT}(\alpha)} \bigcup_{\theta \in \text{ans}((q_i^+ \wedge Q_i^-)\sigma, \mathcal{I})} E_i\sigma\theta$$

applies the action  $\alpha$  to  $\mathcal{I}$ , using  $\sigma$  as the assignment for its parameters. The returned instance is the union of the results of applying the effects specifications  $\text{EFFECT}(\alpha)$ , where the result of each effect specification  $q_i^+ \wedge Q_i^- \rightsquigarrow E_i$  is, in turn, the set of facts  $E_i\sigma\theta$  obtained from  $E_i\sigma$  grounded on all the assignments  $\theta$  that satisfy the query  $q_i^+ \wedge Q_i^-$  over  $\mathcal{I}$ .

$$\begin{aligned} \text{SERVICECALLS}(\mathcal{I}, \alpha\sigma, \mathcal{M}) = \\ \mathcal{M} \cup \{t \mapsto \text{PICKVALUE}(\mathcal{C}) \mid t \text{ occurring in } \text{DO}(\mathcal{I}, \alpha\sigma) \\ \text{and not in } \text{DOM}(\mathcal{M})\} \end{aligned}$$

nondeterministically generates all possible values that can be returned by the service calls, guaranteeing that external services behave in a deterministic manner. More specifically, all the service calls already contained in  $\mathcal{M}$  are maintained, while new service calls are nondeterministically bound to an arbitrary value  $\text{PICKVALUE}(\mathcal{C})$  taken from  $\mathcal{C}$  (which will be the values assumed by such service calls in  $\mathcal{M}$  from now on in the execution).

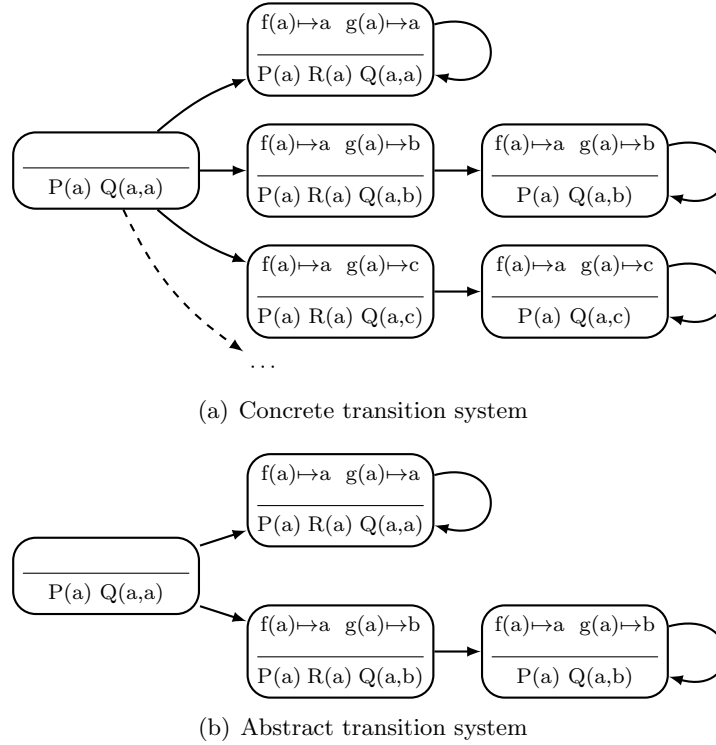
**Concrete transition system.** The *concrete transition system*  $\Upsilon_{\mathcal{S}}$  for  $\mathcal{S}$  is a possibly infinite-state transition system  $\langle \mathcal{C}, \mathcal{R}, \Sigma, s_0, db, \Rightarrow \rangle$  where  $s_0 = \langle \mathcal{I}_0, \emptyset \rangle$  and  $db$  is such that  $db(\langle \mathcal{I}, \mathcal{M} \rangle) = \mathcal{I}$ . Specifically, we define by simultaneous induction  $\Sigma$  and  $\Rightarrow$  as the smallest sets satisfying the following properties: (i)  $s_0 \in \Sigma$ ; (ii) if  $\langle \mathcal{I}, \mathcal{M} \rangle \in \Sigma$ , then for all substitutions  $\sigma$  for the input parameters of  $\alpha$  and for every  $\langle \mathcal{I}', \mathcal{M}' \rangle$  such that  $\langle \langle \mathcal{I}, \mathcal{M} \rangle, \alpha\sigma, \langle \mathcal{I}', \mathcal{M}' \rangle \rangle \in \text{EXEC}_{\mathcal{S}}$ , we have  $\langle \mathcal{I}', \mathcal{M}' \rangle \in \Sigma$  and  $\langle \mathcal{I}, \mathcal{M} \rangle \Rightarrow \langle \mathcal{I}', \mathcal{M}' \rangle$ .

Intuitively, to define the concrete transition system of the DCDS  $\mathcal{S}$  we start from the initial state  $s_0 = \langle \mathcal{I}_0, \emptyset \rangle$ , and for each rule  $Q \mapsto \alpha$  in  $\mathcal{P}$ , we evaluate  $Q$  over  $\mathcal{I}_0$ , and calculate all states  $s$  such that  $\langle s_0, \alpha\sigma, s \rangle \in \text{EXEC}_{\mathcal{S}}$ . Then we repeat the same steps considering each  $s$ , and so on. The computation of successor states can be done by picking all the possible combinations of resulting values for the newly introduced service calls, then checking if the successor obtained for a combination satisfies the equality constraints, filtering it away if this is not the case. It is worth noting that when new service calls are considered, the successors can be countably infinite.

**Example 7.1.** Let  $\mathcal{S} = \langle \mathcal{D}, \mathcal{P} \rangle$  be a DCDS with data layer  $\mathcal{D} = \langle \mathcal{C}, \mathcal{R}, \mathcal{E}, \mathcal{I}_0 \rangle$  and process layer  $\mathcal{P} = \langle \mathcal{F}, \mathcal{A}, \varrho \rangle$ , where  $\mathcal{F} = \{f/1, g/1\}$ ,  $\mathcal{R} = \{Q/2, P/1, R/1\}$ ,  $\mathcal{E} = \emptyset$ ,  $\mathcal{I}_0 = \{P(a), Q(a, a)\}$ ,  $\varrho = \{true \mapsto \alpha\}$ ,  $\mathcal{A} = \{\alpha\}$ , and

$$\alpha : \{Q(a, a) \wedge P(x) \rightsquigarrow \{R(x)\}, P(x) \rightsquigarrow \{P(x), Q(f(x), g(x))\}\}$$

The concrete transition system  $\Upsilon_S$  contains infinitely many successors connected to the initial state. These successors result from the assignment of each possible pair of values to  $f(a)$  and  $g(a)$  (see also Figure 3(a)). ■



**Figure 2** – Concrete and abstract transition systems of the DCDS with deterministic services described in Example 7.2; special relations that store the service calls results are represented using a  $call \mapsto value$  notation

**Example 7.2.** Consider a variation of the DCDS described in Example 7.1, where the data layer is equipped with an equality constraint, i.e.,  $\mathcal{E} = \{P(x) \wedge Q(y, z) \rightarrow x = y\}$ . The resulting concrete transition system has still infinitely many successors of the initial state, but the presence of the equality constraint requires to keep only those successors in which  $f(a)$  returns  $a$  (see also Figure 2(a)). ■

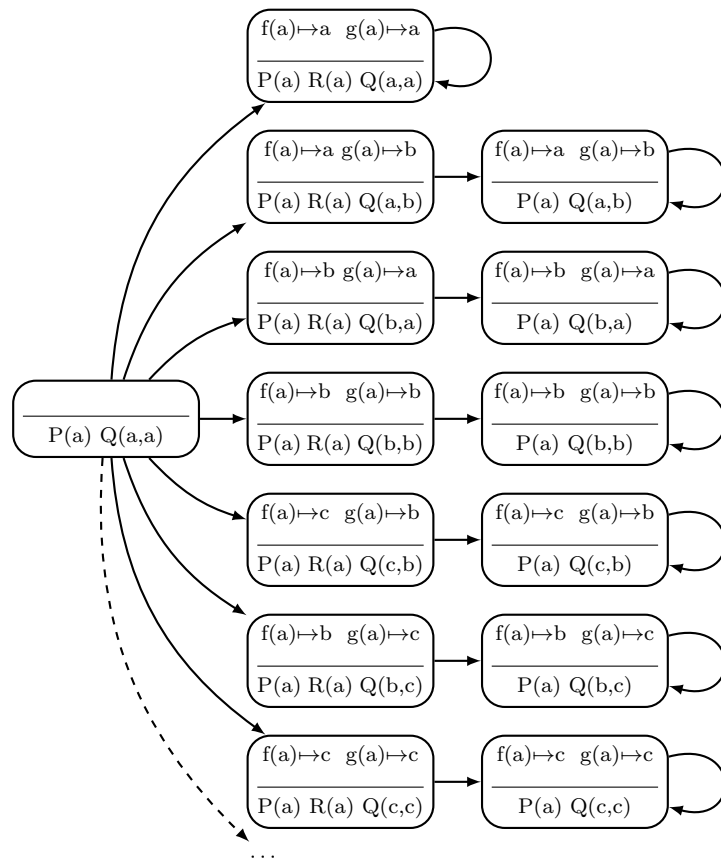
## 7.2 Run-Bounded Systems

We now study the verification of DCDSs with deterministic services. In particular, we are interested in the following problem: given a DCDS  $\mathcal{S}$  and a temporal property  $\Phi$ , check whether  $\Upsilon_S \models \Phi$ . Not surprisingly, given the expressive power of DCDS as a computation model, the verification problem is undecidable for all the  $\mu$ -calculus variants introduced in Section 6. In fact, we can show an even stronger undecidability result, for a very small fragment of propositional linear temporal logic (LTL) [58], namely the safety properties of the form  $Gp$  where  $p$  is propositional.

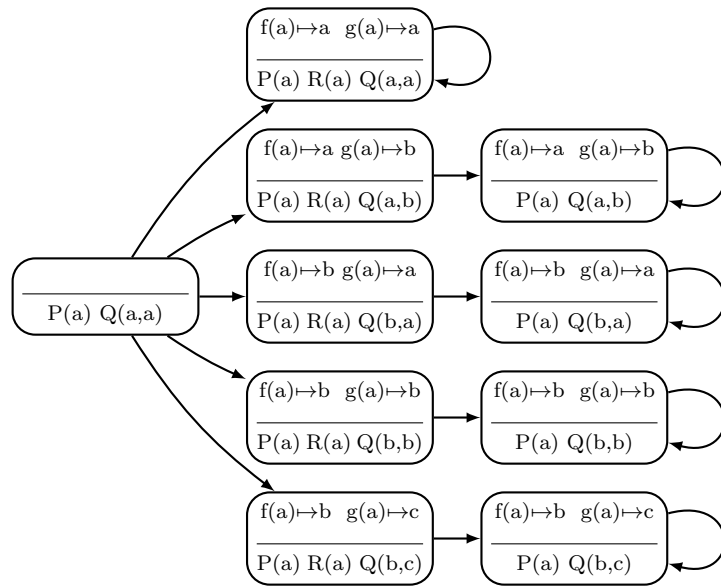
**Theorem 7.1.** *There exists a DCDS  $\mathcal{S}$  with deterministic services, and a propositional LTL safety property  $\Phi$ , such that checking  $\Upsilon_S \models \Phi$  is undecidable.*

In the following, we isolate a notable class of DCDS for which verification of  $\mu\mathcal{L}_A$  is not only decidable, but can also be reduced to standard model checking techniques.

Consider a transition system  $\Upsilon = \langle \Delta, \mathcal{R}, \Sigma, s_0, db, \Rightarrow \rangle$ . A run  $\tau$  in  $\Upsilon$  is a (finite or infinite) sequence of states  $s_0 s_1 s_2 \dots$  rooted at  $s_0$ , where  $s_i \Rightarrow s_{i+1}$ . We use  $\tau(i)$  to denote  $s_i$  and  $\tau[i]$



(a) Concrete transition system



(b) Abstract transition system

**Figure 3 – Concrete and abstract transition systems of the DCDS with deterministic services described in Example 7.1; special relations that store the service calls results are represented using a  $call \mapsto value$  notation**

to represent the finite prefix  $s_0 \cdots s_i$  of  $\tau$ . A run  $\tau = s_0 s_1 s_2 \cdots$  is (*data*) *bounded* if the number of values mentioned inside its databases is bounded, i.e., there exists a finite bound  $b$  such that  $|\bigcup_{s \text{ state of } \tau} \text{ADOM}(db(s))| < b$ . This is equivalent to saying that, for every finite prefix  $\tau[i]$  of  $\tau$ ,  $|\bigcup_{j \in \{0, \dots, i\}} \text{ADOM}(db(s_j))| < b$ . We say that  $\Upsilon$  is *run-bounded* if there exists a bound  $b$  such

that every run in  $\Upsilon$  is (data) bounded by  $b$ . A DCDS  $\mathcal{S}$  is run-bounded if its concrete transition system  $\Upsilon_{\mathcal{S}}$  is run-bounded.

Intuitively, a (data) unbounded run represents an execution of the DCDS in which infinitely many distinct values occur because infinitely many different service calls are issued. Since we model deterministic services whose number is finite, this can only happen if some service is repeatedly called with arguments that are the result of previous service calls. This means that the values of the run indirectly depend on arbitrarily many states in the past.

Notice that run boundedness does not impose any restriction about the branching of the transition system; in particular,  $\Upsilon_{\mathcal{S}}$  is typically infinite-branching because new service calls may return any possible value. We show that this restriction guarantees decidability for  $\mu\mathcal{L}_A$  verification of run-bounded DCDSs with deterministic services.

**Theorem 7.2.** *Verification of  $\mu\mathcal{L}_A$  properties on run-bounded DCDSs with deterministic services is decidable.*

We get this result by showing that for run-bounded DCDSs there always exists an abstract finite-state transition system that is history preserving bisimilar to the concrete one, and hence satisfies the same  $\mu\mathcal{L}_A$  formulae as the concrete transition system.

**Theorem 7.3.** *For every run-bounded DCDS  $\mathcal{S}$  with deterministic services, given its concrete transition system  $\Upsilon_{\mathcal{S}}$  there exists an (abstract) finite-state transition system  $\Theta_{\mathcal{S}}$  such that  $\Theta_{\mathcal{S}}$  is history preserving bisimilar to  $\Upsilon_{\mathcal{S}}$ , i.e.,  $\Theta_{\mathcal{S}} \approx \Upsilon_{\mathcal{S}}$ .*

Let  $\Sigma$  be the set of states of  $\Theta_{\mathcal{S}}$  and  $\text{ADOM}(\Theta_{\mathcal{S}}) = \bigcup_{s_i \in \Sigma} \text{ADOM}(db(s_i))$ . If  $\Theta_{\mathcal{S}}$  is finite-state, then there exists a bound  $b$  such that  $|\text{ADOM}(\Theta_{\mathcal{S}})| < b$ . Consequently, it is possible to transform a  $\mu\mathcal{L}_A$  property  $\Phi$  into an equivalent *finite* propositional  $\mu$ -calculus formula  $\text{PROP}(\Phi)$ , where  $\text{PROP}(\Phi)$  is inductively defined over the structure of  $\Phi$  as the identity, except for the following case:  $\text{PROP}(\exists x. \text{LIVE}(x) \wedge \Psi(x)) = \bigvee_{t_i \in \text{ADOM}(\mathcal{S})} \text{LIVE}(t_i) \wedge \text{PROP}(\Psi(t_i))$ . Clearly,  $\Theta_{\mathcal{S}} \models \Phi$  if and only if  $\Theta_{\mathcal{S}} \models \text{PROP}(\Phi)$ .

**Theorem 7.4.** *Verification of  $\mu\mathcal{L}_A$  properties for run-bounded DCDSs with deterministic services can be reduced to conventional model checking of propositional  $\mu$ -calculus over a finite transition system.*

By the above theorem, and recalling that model checking of propositional  $\mu$ -calculus formulae over finite transition systems is decidable [35], we get Theorem 7.2.

We conclude the section by observing that the approach presented above for  $\mu\mathcal{L}_A$  does not extend to full  $\mu\mathcal{L}$ .

**Theorem 7.5.** *There exists a DCDS  $\mathcal{S}$  for which it is impossible to find a faithful finite-state abstraction that satisfies the same  $\mu\mathcal{L}$  properties as  $\mathcal{S}$ .*

Theorem 7.5 is proved by exhibiting, for every  $n$ , a  $\mu\mathcal{L}$  property that requires the existence of at least  $n$  objects in the transition system. Even if this observation does not imply undecidability of model checking  $\mu\mathcal{L}$  properties over run-bounded DCDSs, it shows that there is no hope of reducing this problem to standard, finite-state model checking.

### 7.3 Weakly Acyclic DCDSs

The results presented in Section 7.2 rely on the hypothesis that the DCDS under study is run-bounded, which is a semantic restriction. A natural question is whether it is possible to check run-boundedness of a DCDS. We provide a negative answer to this question.

**Theorem 7.6.** *Checking run-boundedness of DCDSs with deterministic services is undecidable.*



To mitigate this issue, we investigate a sufficient syntactic condition that can be effectively tested over the process layer of the DCDS: if the condition is met, then the DCDS is guaranteed to be run-bounded, otherwise nothing can be said. To this end, we recast the approach of [7] in the more abstract and expressive framework presented here. In particular, we first introduce the “positive approximate” of a DCDS, which abstracts away some of its aspects. We do so for convenience, but we note that the definition of weak-acyclicity as well as our results can be stated directly over the original DCDS (in fact, we do so in condensed presentations of this work). Technically, given a DCDS  $\mathcal{S} = \langle \mathcal{D}, \mathcal{P} \rangle$  with data layer  $\mathcal{D} = \langle \mathcal{C}, \mathcal{R}, \mathcal{E}, \mathcal{I}_0 \rangle$  and process layer  $\mathcal{P} = \langle \mathcal{F}, \mathcal{A}, \varrho \rangle$ , its *positive approximate*  $\mathcal{S}^+$  is a DCDS  $\langle \mathcal{D}^+, \mathcal{P}^+ \rangle$ , where  $\mathcal{D}^+ = \langle \mathcal{C}, \mathcal{R}, \emptyset, \mathcal{I}_0 \rangle$  corresponds to  $\mathcal{D}$  without equality constraints, while  $\mathcal{P}^+ = \langle \mathcal{F}, \mathcal{A}^+, \varrho^+ \rangle$  is a process layer whose actions  $\mathcal{A}^+$  and process  $\varrho^+$  are obtained as follows:

- Each condition-action rule  $Q \mapsto \alpha$  in  $\varrho$  becomes  $\text{true} \mapsto \alpha^+$  in  $\varrho^+$ . Therefore,  $\varrho^+$  is a process that supports the execution of every action in  $\mathcal{A}^+$  at each step.
- Each action  $\alpha(p_1, \dots, p_n) : \{e_1, \dots, e_m\}$  in  $\mathcal{A}$  becomes  $\alpha^+ : \{e_1^+, \dots, e_m^+\}$  in  $\mathcal{A}^+$ , where each  $e_i = q_i^+ \wedge Q_i^- \rightsquigarrow E_i$  becomes in turn  $e_i^+ = q_i^+ \rightsquigarrow E_i$ . Intuitively, the positive approximate action is obtained from the original action by removing all the parameters from its signature, and by removing all “negative” components from the query used to instantiate its effect specifications; note that the variables of  $q_i^+$  that were parameters in  $\alpha$  are now free variables in  $\alpha^+$ .

The positive approximate fulfils the following key property.

**Lemma 7.7.** *Given a DCDS  $\mathcal{S}$ , if its positive approximate  $\mathcal{S}^+$  is run-bounded, then  $\mathcal{S}$  is run-bounded as well.*

To derive a sufficient condition for  $\mathcal{S}^+$  to be run-bounded, we can exploit a strict correspondence between the execution of an action in  $\mathcal{P}^+$  and a step in the chase of a set of tuple generating dependencies (TGDs) in data exchange [2, 36]. In particular, we resort to a well-known result in data exchange, namely chase termination for *weakly acyclic* TGDs [36].<sup>9</sup>

In our setting, the weak acyclicity of a process layer is a property over a dataflow graph constructed by analyzing the corresponding positive approximate process layer. A non-weakly acyclic DCDS contains a service that may be repeatedly called, every time using fresh values that are directly or indirectly obtained by manipulating previous results produced by the same service. This self-dependency can potentially lead to an infinite number of calls of the same service along an execution of the system, thus making it impossible to put a bound on the data used throughout the run (see also Example 7.3). Weak acyclicity rules out such self dependencies and is actually a sufficient condition for run-boundedness.

Given a DCDS  $\mathcal{S} = \langle \mathcal{D}, \mathcal{P} \rangle$  with positive approximate  $\mathcal{S}^+ = \langle \mathcal{D}^+, \mathcal{P}^+ \rangle$ , the *dependency graph* of  $\mathcal{P}^+$  is an edge-labeled directed graph  $\langle N, E \rangle$  where: (i)  $N \subseteq \mathcal{R} \times \mathbb{N}^+$  is a set of nodes such that  $\langle R, i \rangle \in N$  for every  $R/n \in \mathcal{R}$  and every  $i \in \{1, \dots, n\}$ ; (ii)  $E \subseteq N \times N \times \{\text{true}, \text{false}\}$  is a set of labeled edges where

- an ordinary edge  $\langle \langle R_1, j \rangle, \langle R_2, k \rangle, \text{false} \rangle \in E$  if there exists an action  $\alpha^+ \in \mathcal{A}^+$ , an effect  $q_i^+ \rightsquigarrow E_i \in \text{EFFECT}(\alpha^+)$  and a variable  $x$  such that  $R_1(\dots, t_{j-1}, x, t_{j+1}, \dots)$  occurs in  $q_i^+$  and  $R_2(\dots, t'_{k-1}, x, t'_{k+i}, \dots)$  occurs in  $E_i$ ;
- a special edge  $\langle \langle R_1, j \rangle, \langle R_2, k \rangle, \text{true} \rangle \in E$  if there exists an action  $\alpha^+ \in \mathcal{A}^+$ , an effect  $q_i^+ \rightsquigarrow E_i \in \text{EFFECT}(\alpha^+)$  and a variable  $x$  such that  $R_1(\dots, t_{j-1}, x, t_{j+1}, \dots)$  occurs in  $q_i^+$ ,  $R_2(\dots, t'_{k-1}, t, t'_{k+i}, \dots)$  occurs in  $E_i$ , and  $t = f(\dots, x, \dots)$ , with  $f \in \mathcal{F}$ .

$\mathcal{P}$  is *weakly acyclic* if the dependency graph of its approximate  $\mathcal{P}^+$  does not contain any cycle going through a special edge. We say that a DCDS is weakly acyclic if its process layer is weakly acyclic (e.g., see Figure 5(a)).

<sup>9</sup>Notice that using other variants of weak acyclicity is also possible [53].

Intuitively, ordinary edges represent the possible propagation (copy) of a value across states:  $\langle\langle R_1, j \rangle, \langle R_2, k \rangle, \text{false}\rangle \in E$  reflects the possibility that the value currently stored inside the  $j$ -th component of an  $R_1$  tuple will be moved to the  $k$ -th component of an  $R_2$  tuple in the next state. Contrariwise, special edges represent that a value can be taken as parameter of a service call, thus contributing to the creation of (possibly new) values across states:  $\langle\langle R_1, j \rangle, \langle R_2, k \rangle, \text{true}\rangle \in E$  means that the value currently stored inside the  $j$ -th component of an  $R_1$  tuple could be used as parameter for a service call, whose result is then stored inside the  $k$ -th component of an  $R_2$  tuple.

A cycle going through a special edge, forbidden by the weak acyclicity condition, represents that a service may be repeatedly called, every time using fresh values that are indirectly or directly obtained by manipulating previous results produced by the same service. This self-dependency can potentially lead to an infinite number of calls of the same service along an execution of the system, thus making it impossible to put a bound on the data used throughout the run.

**Example 7.3.** Let  $\mathcal{S} = \langle \mathcal{D}, \mathcal{P} \rangle$  be a DCDS with data layer  $\mathcal{D} = \langle \mathcal{C}, \mathcal{R}, \emptyset, \mathcal{I}_0 \rangle$  and process layer  $\mathcal{P} = \langle \mathcal{F}, \mathcal{A}, \varrho \rangle$ , where  $\mathcal{F} = \{f/1\}$ ,  $\mathcal{R} = \{R/1, Q/1\}$ ,  $\mathcal{I}_0 = \{R(a)\}$ ,  $\varrho = \{\text{true} \mapsto \alpha\}$  and  $\mathcal{A} = \{\alpha\}$ , where  $\alpha : \{R(x) \rightsquigarrow Q(f(x)), Q(x) \rightsquigarrow R(x)\}$ .

$\mathcal{S}$  is not weakly acyclic, due to the mutual dependency between  $R$  and  $Q$  that involves a call to service  $f$ . This can be easily seen from the dataflow graph (shown in Figure 5(b)), which contains a special edge from  $\langle R, 1 \rangle$  to  $\langle Q, 1 \rangle$ , and a normal edge from  $\langle Q, 1 \rangle$  to  $\langle R, 1 \rangle$ . Notice that, in this case, the positive approximate of  $\mathcal{S}$  coincides with  $\mathcal{S}$  itself. Starting from the initial state,  $\alpha$  calls  $f(a)$  and stores the result inside  $Q$ . A second execution of  $\alpha$  transfers the result of  $f(a)$  into  $R$ . When  $\alpha$  is executed for the third time,  $f$  is called again, but using as parameter the previously obtained result. Consequently,  $f$  may return a new, fresh result, because  $f(f(a))$  may be different from  $f(a)$ . This chain can be repeated forever, leading to possibly generating infinitely many distinct values along the run. The existence of a run in which  $a, f(a), f(f(a)), f(f(f(a))), \dots$ , are all distinct values, makes it impossible to obtain a finite-state abstraction for  $\mathcal{S}$  (see Figure 4(b)). ■

**Theorem 7.8.** *Every weakly acyclic DCDS with deterministic services is run-bounded.*

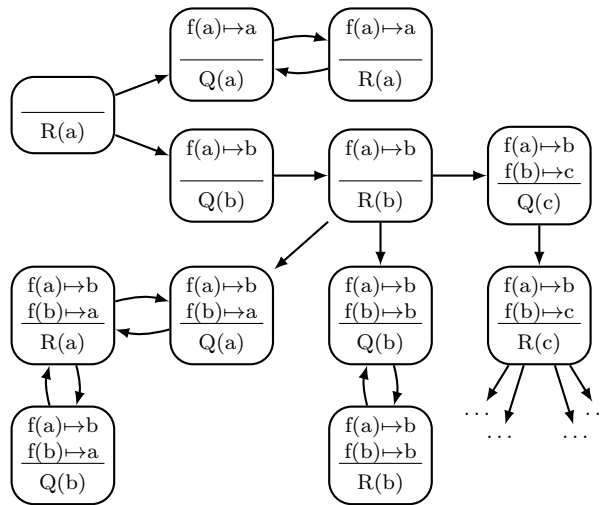
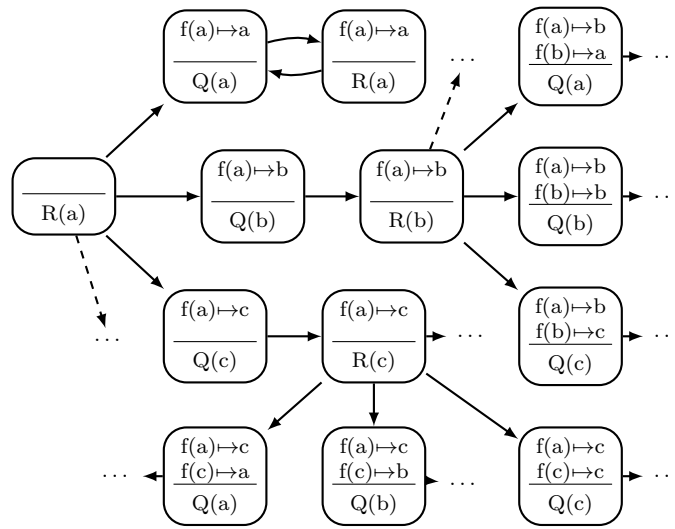
Checking weak acyclicity is polynomial in the size of the DCDS. Thus it gives us an effective way to verify DCDSs.

**Theorem 7.9.** *Verification of  $\mu\mathcal{L}_A$  properties for weakly acyclic DCDSs with deterministic services is decidable, and can be reduced to model checking of propositional  $\mu$ -calculus over a finite transition system.*



**Example 7.4.** Consider the DCDSs described in Example 7.1 and 7.2. They have the same dataflow graph, which is weakly acyclic (see Figure 5(a)). This guarantees that they are run-bounded and that it is possible to find a faithful finite-state abstraction from them. Two such abstractions are respectively shown in Figure 3(b) and 2(b). ■

## 8 Nondeterministic Services

We now consider DCDSs under the assumption that services behave nondeterministically, i.e., two calls of a service with the same arguments may return distinct results during the same run. This case captures both services that model a truly nondeterministic process (e.g., human operators, random processes), and services that model stateful servers. In the remainder of this section, whenever we refer to a DCDS, services are implicitly assumed nondeterministic.



**Figure 4 – Concrete and abstract transition systems of the run-unbounded DCDS with deterministic services  $\mathcal{S}$  described in Example 7.3; special relations that store the service calls results are represented using a *call  $\mapsto$  value* notation**



**Figure 5 – Examples of dataflow graphs for DCDSs with deterministic services; special edges are decorated with \***

## 8.1 Semantics

As in the case of deterministic services, we define the semantics of a DCDS  $\mathcal{S}$  in terms of a (possibly infinite) transition system  $\Upsilon_{\mathcal{S}}$ .

Let  $\mathcal{S} = \langle \mathcal{D}, \mathcal{P} \rangle$  be a DCDS with data layer  $\mathcal{D} = \langle \mathcal{C}, \mathcal{R}, \mathcal{E}, \mathcal{I}_0 \rangle$  and process layer  $\mathcal{P} = \langle \mathcal{F}, \mathcal{A}, \varrho \rangle$ . A *state* is simply a relational instance of  $\mathcal{R}$  over  $\mathcal{C}$  satisfying each constraint in  $\mathcal{E}$ . We denote the *initial state* with  $\mathcal{I}_0$ .

Next, we define the semantics of action application. Let  $\alpha$  be an action in  $\mathcal{A}$  of the form  $\alpha(p_1, \dots, p_m) : \{e_1, \dots, e_m\}$  with effects  $e_i = q_i^+ \wedge Q_i^- \rightsquigarrow E_i$ . The parameters for  $\alpha$  are guarded by the condition-action rule  $Q \mapsto \alpha \in \varrho$ . Let  $\sigma$  be a legal substitution for the input parameters  $p_1, \dots, p_m$  with values taken from  $\mathcal{C}$ .

We reuse the definition of  $\text{DO}(\mathcal{I}, \alpha\sigma)$  from Section 7.1, as the instance obtained by evaluating the effects of  $\alpha$  on instance  $\mathcal{I}$ . Recall that  $\text{DO}()$  generates an instance over values from the domain  $\mathcal{C}$  but also over Skolem terms, which model service calls. For any such instance  $\bar{\mathcal{I}}$ , we denote with  $\text{CALLS}(\bar{\mathcal{I}})$  the set of calls it contains. For a given set  $D \subseteq \mathcal{C}$ , we denote with  $\text{EVALS}_D(\mathcal{I}, \alpha, \sigma)$  the set of substitutions that replace all service calls in  $\text{DO}(\mathcal{I}, \alpha, \sigma)$  with values in  $D$ ,

$$\text{EVALS}_D(\mathcal{I}, \alpha, \sigma) = \{ \theta \mid \theta \text{ is a total function} \\ \theta : \text{CALLS}(\text{DO}(\mathcal{I}, \alpha, \sigma)) \rightarrow D \}.$$

Each substitution in  $\text{EVALS}_D(\mathcal{I}, \alpha, \sigma)$  models the simultaneous evaluation of all service calls, which replaces the calls with results selected nondeterministically from  $D$ . In the following, we refer to these substitutions as *evaluations*.

**Concrete action execution.** We introduce a transition relation  $\text{N-EXEC}_{\mathcal{S}}$  between states, called *concrete execution* of  $\alpha\sigma\theta$ , such that  $\langle \mathcal{I}, \alpha\sigma\theta, \mathcal{I}' \rangle \in \text{N-EXEC}_{\mathcal{S}}$  if the following holds:

1.  $\sigma$  is a legal parameter assignment for  $\alpha$  in state  $\mathcal{I}$ ,
2.  $\theta \in \text{EVALS}_{\mathcal{C}}(\mathcal{I}, \alpha, \sigma)$ ,
3.  $\mathcal{I}' = \text{DO}(\mathcal{I}, \alpha, \sigma)\theta$ , and
4.  $\mathcal{I}'$  satisfies the constraints  $\mathcal{E}$ .

Notice that, in contrast to the deterministic services case, the choice of evaluation  $\theta$  is not subject to the requirement that it evaluates a service call to the same result *across* concrete execution steps. However, notice that *within* a concrete execution step, all occurrences of the same service call evaluate to the same result (modeling the fact that a call with given arguments is invoked only once per transition, and the returned result is copied as needed).

**Concrete transition system.** The *concrete transition system*  $\Upsilon_{\mathcal{S}}$  for  $\mathcal{S}$  is a transition system whose states are labeled by databases. More precisely,  $\Upsilon_{\mathcal{S}} = \langle \mathcal{C}, \mathcal{R}, \Sigma, s_0, db, \Rightarrow \rangle$  where  $s_0 = \mathcal{I}_0$  and  $db$  is such that  $db(\mathcal{I}) = \mathcal{I}$ .  $\Sigma$  and  $\Rightarrow$  are defined by simultaneous induction as the smallest sets satisfying the following properties: (i)  $\mathcal{I}_0 \in \Sigma$ ; (ii) if  $\mathcal{I} \in \Sigma$ , then for all  $\alpha, \sigma, \theta$  and  $\mathcal{I}'$  such that  $\langle \mathcal{I}, \alpha\sigma\theta, \mathcal{I}' \rangle \in \text{N-EXEC}_{\mathcal{S}}$ , we have that  $\mathcal{I}' \in \Sigma$ , and  $\mathcal{I} \Rightarrow \mathcal{I}'$ .

## 8.2 State-Bounded Systems

We consider the verification problem for DCDS with nondeterministic services. As in the deterministic case, restrictions on both the processes and the properties are required, motivated by the following undecidability result.

**Theorem 8.1.** *There exists a DCDS  $\mathcal{S}$  with nondeterministic services, and a propositional LTL safety property  $\Phi$ , such that checking  $\Upsilon_{\mathcal{S}} \models \Phi$  is undecidable.*

**State-bounded DCDS.** Since we are interested in verifying more expressive temporal properties, we need to consider restricted classes of DCDS. We observe first that, with nondeterministic services, the run-boundedness restriction of Section 7.2 is very limiting on the form of the DCDS, as it boils down to imposing a bound on how many times each service may be called with the same arguments. Observe that this was not the case for deterministic services, where the unlimited same-argument calls are allowed, as they all return the same result. We propose a less restrictive alternative. We say that DCDS  $\mathcal{S}$  is *state-bounded* if there is a finite bound  $b$  such that for each state  $\mathcal{I}$  of  $\Upsilon_{\mathcal{S}}$ ,  $|\text{ADOM}(\mathcal{I})| < b$ . Notice that, in contrast to the notion of run-boundedness, state-boundedness does allow runs in which infinitely many distinct values occur because infinitely many service calls are issued. The unboundedly many call results are distributed *across* states of the run, but may not accumulate *within* a single state. The following result shows that we also need to restrict the logic, as the one used in the deterministic case is too expressive for decidability.

**Theorem 8.2.** *Verification of  $\mu\mathcal{L}_A$  properties on state-bounded DCDSs with nondeterministic services is undecidable.*

We therefore restrict the property class to the logic  $\mu\mathcal{L}_P \subset \mu\mathcal{L}_A$  presented in Section 6.2.

**Theorem 8.3.** *Verification of  $\mu\mathcal{L}_P$  properties by state-bounded DCDS with nondeterministic services is decidable.*

We relegate the proof of Theorem 8.3 to Appendix A.3, but provide the main ideas here.

**Abstract Transition System.** Given a DCDS  $\mathcal{S}$ , we show that if the concrete transition system  $\Upsilon_{\mathcal{S}}$  is state-bounded, then there is a finite-state abstract transition system  $\Theta_{\mathcal{S}}$  whose states and edges are subsets of those in  $\Upsilon_{\mathcal{S}}$ , such that  $\Theta_{\mathcal{S}}$  is persistence-preserving bisimilar to  $\Upsilon_{\mathcal{S}}$  (and hence satisfies the same  $\mu\mathcal{L}_P$  properties, by Theorem 6.2). Since  $\Theta_{\mathcal{S}}$  is finite-state, the verification of  $\mu\mathcal{L}_P$  properties on  $\Upsilon_{\mathcal{S}}$  reduces to finite-state model checking on  $\Theta_{\mathcal{S}}$ , and hence is decidable.

The existence of  $\Theta_{\mathcal{S}}$  follows from the key fact that if two states of  $\Upsilon_{\mathcal{S}}$  are isomorphic, then they are persistence-preserving bisimilar. This implies that one can construct a finitely-branching transition system  $\Theta_{\mathcal{S}}$  (i.e. with finite number of successors per state), such that  $\Theta_{\mathcal{S}}$  is persistence-preserving bisimilar to  $\Upsilon_{\mathcal{S}}$ , by dropping sibling states from  $\Upsilon_{\mathcal{S}}$  as follows: instead of listing among the successors of  $s$  one state for each possible instantiation of the service call results, just keep a *representative* state for each isomorphism type. Since the number of service calls made in each state is finite, the number of distinct isomorphism types is finite, so the finite branching follows. We call a transition system  $\Theta_{\mathcal{S}}$  obtained as above a *pruning* of  $\Upsilon_{\mathcal{S}}$ .

Notice that despite being finitely-branching, any pruning  $\Theta_{\mathcal{S}}$  can still have infinitely many states, as it may contain infinitely long simple runs<sup>10</sup>  $\tau$ , along which the service calls return in each state “fresh” values, i.e., values distinct from all values appearing in the predecessors of this state on  $\tau$ . This problem is solved by judiciously selecting which representatives to keep in  $\Theta_{\mathcal{S}}$  for the successors of a state  $s$ . Namely, whenever the representatives of a given isomorphism type  $\mathcal{T}$  include states generated exclusively by service calls that “recycle” values, select only such states (finitely many thereof, of course). By recycled values we mean values appearing on a path leading into  $s$ .

If  $\Upsilon_{\mathcal{S}}$  is state-bounded, then the number of service calls per state is bounded, and due to the construction’s preference for recycling, it follows that all simple runs in  $\Theta_{\mathcal{S}}$  must have finite length. Together with the finite branching, this implies finiteness of  $\Theta_{\mathcal{S}}$ .

Notice that proving the existence of  $\Theta_{\mathcal{S}}$  does not suffice for decidability, as the proof is non-constructive. We therefore provide an algorithm for constructing  $\Theta_{\mathcal{S}}$  (Algorithm RCYCL). One of the technical problems we need to overcome in developing the algorithm is that we evidently cannot start from the infinite-state concrete transition system, instead exploring a

<sup>10</sup>We call a run *simple* if no state appears more than once in the run.

portion thereof. This means that it is not obvious how to decide whether the successors of a state are generated by recycling service calls, since these calls may recycle from paths that RCYCL hasn't explored yet. Therefore, RCYCL may sometimes select non-recycling service calls even when a recycling alternative exists. However, we can prove that RCYCL constructs what we call an *eventually recycling pruning*, which in essence means it may fail to detect recycling service calls, but only a bounded number of times.

We formalize the above discussion in Appendix A.3, where we prove the following result:

**Theorem 8.4.** *If input DCDS  $\mathcal{S}$  is state-bounded, then every possible run of Algorithm RCYCL terminates, yielding a finite eventually recycling pruning  $\Theta_{\mathcal{S}}$  of  $\Upsilon_{\mathcal{S}}$ , with  $\Upsilon_{\mathcal{S}} \sim \Theta_{\mathcal{S}}$ .*

Theorem 8.4 and Theorem 6.2 directly imply Theorem 8.3.

Figures 7 and 6 illustrate two concrete transition systems, and possible recycling prunings for them.

### 8.3 GR-Acyclic DCDSs

As with run-boundedness in the deterministic services case, for nondeterministic services the state-boundedness restriction is a semantic property. We investigate whether it can be effectively checked.

**Theorem 8.5.** *Checking state-boundedness of DCDSs is undecidable.*

Consequently we propose a sufficient syntactic restriction.

Intuitively, for a run to have unbounded states, it must issue unboundedly many service calls. Since there are only a bounded number of effects in the process layer specification, there must exist some service-calling effect that “cyclically generates” fresh values (i.e., is invoked infinitely many times during the run). Notice that unbounded generation of fresh values is insufficient for state-unboundedness: these values must also accumulate in the states. But by definition of the DCDS semantics, a transition drops (“forgets”) all values that are not explicitly copied (“recalled”) into the successor. Therefore, to accumulate, a value must be “cyclically recalled” throughout the run (it must be copied infinitely many times from relation to relation).

GR-acyclicity is stated in terms of a dataflow graph constructed by analyzing the process layer. The graph identifies how service calls and value recalls can be sequenced. In essence, GR-acyclicity requires the absence of a “generate cycle” that feeds into a “recall cycle”.

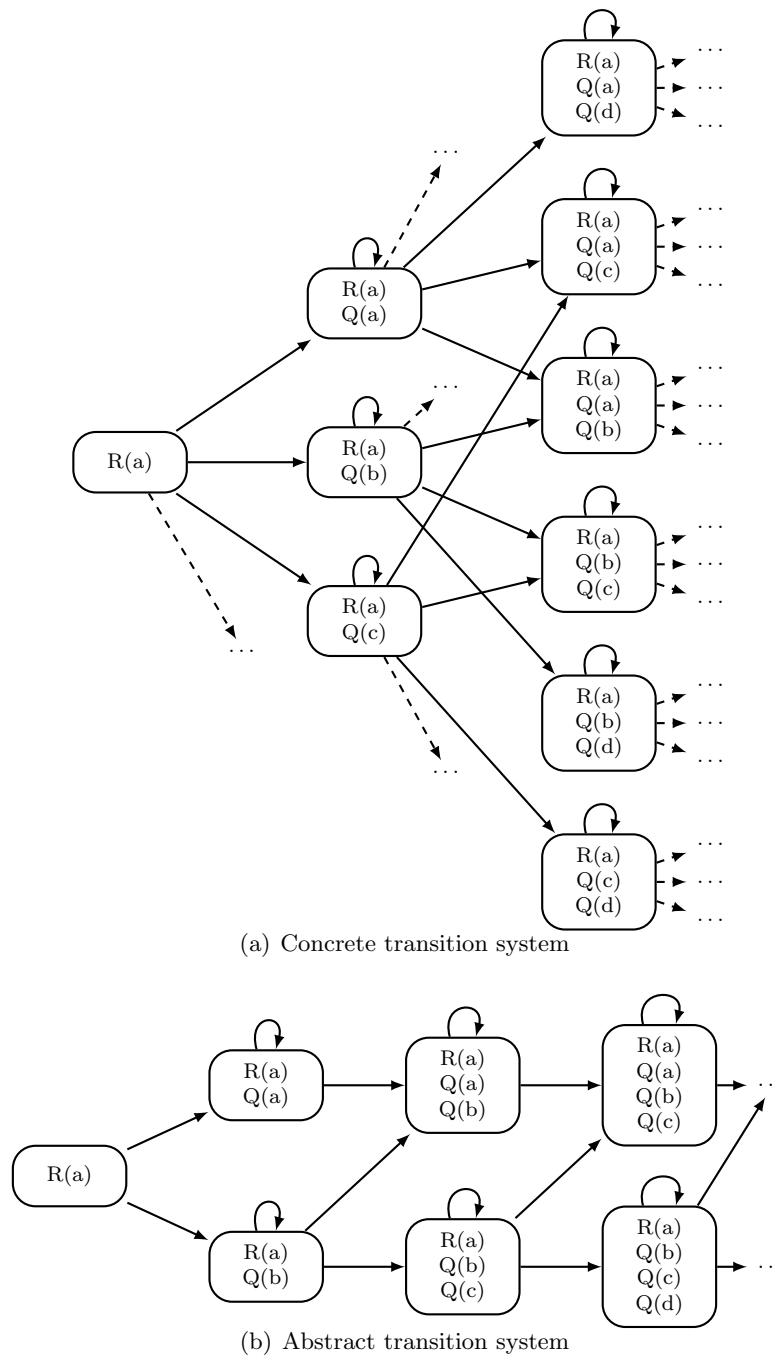
**GR-acyclicity.** Let  $\mathcal{A}$  be a set of actions, and  $\mathcal{A}^+$  its positive approximate (Section 7.3). We call *dataflow graph* of  $\mathcal{A}$  the directed edge-labeled graph  $\langle N, E \rangle$  whose set  $N$  of nodes is the set of relation names occurring in  $\mathcal{A}$ , and in which each edge in  $E$  is a 4-tuple  $(R_1, id, R_2, b)$ , where  $R_1$  and  $R_2$  are two nodes in  $N$ ,  $id$  is a (unique) edge identifier, and  $b$  is a boolean flag used to mark *special* edges. Formally,  $E$  is the minimal set satisfying the following condition: for each effect  $e$  of  $\mathcal{A}^+$ , each  $R(t_1, \dots, t_m)$  in the body of  $e$ , each  $Q(t'_1, \dots, t'_{m'})$  in the head of  $e$ , and each  $i \in \{1, \dots, m'\}$ :

- if  $t'_i$  is either an element of  $\text{ADOM}(\mathcal{I}_0)$  or a free variable, then  $(R, id, Q, \text{false}) \in E$ , where  $id$  is a fresh edge identifier.
- if  $t'_i$  is a service call, then  $(R, id, Q, \text{true}) \in E$ , where  $id$  is a fresh edge identifier.

We say that  $\mathcal{A}$  is *GR-acyclic* if there is no path  $\pi = \pi_1\pi_2\pi_3$  in the dataflow graph of  $\mathcal{A}$ , such that  $\pi_1, \pi_3$  are simple cycles and  $\pi_2$  is a path containing a special edge that is disjoint from the edges of  $\pi_1$ . We say that a process layer  $\mathcal{P} = \langle \mathcal{F}, \mathcal{A}, \rho \rangle$  is GR-acyclic, if  $\mathcal{A}$  is GR-acyclic. We call a DCDS GR-acyclic if its process layer is GR-acyclic.

Notice that GR-acyclicity is a purely syntactic notion. Moreover, it can be checked in PTIME since the dataflow graph has size polynomial in the size of the process layer specification.

**Theorem 8.6.** *Any GR-acyclic DCDS is state-bounded.*



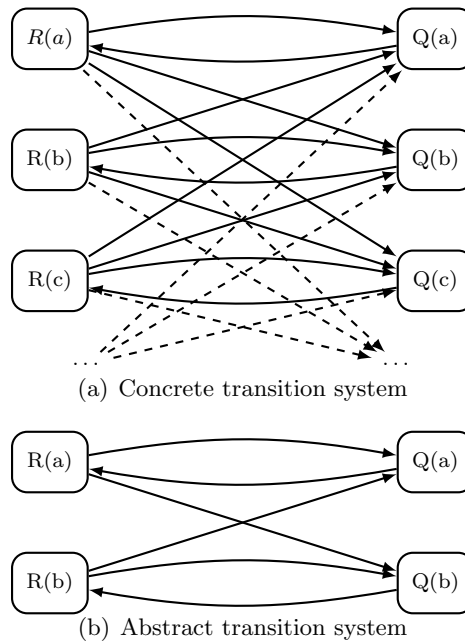
**Figure 6 – Concrete and abstract transition systems of the state-unbounded DCDS with nondeterministic services of Example 8.2**

We show the proof in Appendix A.3 but provide some intuition here, noting that the dataflow analysis is significantly more subtle than suggested above.

First, note that ordinary edges correspond to an effect copying a value from a relation of the current state to a relation of the successor state. Special edges correspond to feeding a value of the current state to a service call and storing the result in a relation of the successor state. Note that the cycles  $\pi_1$  and  $\pi_3$  allow both kinds of edges, reflecting the insight that the size of the state is affected in the same way regardless of whether a value is *copied* to the successor, or it is *replaced* with a service call result (see Example 8.2 and Example 8.3 for illustrations of state-unboundedness arising from each case).  $\pi_1, \pi_3$  are both “recall cycles”: the number of values moving around them does not decrease (this is of course a conservative statement; reality depends on the semantics of queries in the effects, which is abstracted away). Note that  $\pi_2$

contains a special edge  $E$ , which means that the values moving around  $\pi_1$  are cyclically fed into the service call  $f$  of  $E$ . The key insight here is that, even if the set of values moving around  $\pi_1$  does not change (no special edges in  $\pi_1$  replace them), and thus the service call  $f$  sees the same bounded set of distinct arguments over time, it can still generate an unbounded number of fresh values because  $f$  is nondeterministic.  $\pi_1\pi_2$  constitute the “generate cycle” we mention above. The generated values are stored in the recall cycle  $\pi_3$ , where they accumulate and force the size of the relations of  $\pi_3$  to grow unboundedly.

**Example 8.1.** Let us consider again the DCDS  $\mathcal{S}$  described in Example 7.3, this time considering  $f/1$  as a nondeterministic service. The resulting concrete transition system is shown in Figure 7(a). Even if  $\mathcal{S}$  is not run-bounded, it is state-bounded, because in every state its database consists of only one tuple. This is attested by the dataflow graph shown in Figure 7(a), and guarantees the existence of a faithful finite-state abstraction. One such finite-state abstraction is reported in Figure 7(b). ■




**Figure 7 – Concrete and abstract transition systems obtained when the DCDS described in Example 7.3 has nondeterministic services**

**Example 8.2.** Let  $\mathcal{S} = \langle \mathcal{D}, \mathcal{P} \rangle$  be a DCDS with data layer  $\mathcal{D} = \langle \mathcal{C}, \mathcal{R}, \emptyset, \mathcal{I}_0 \rangle$  and process layer  $\mathcal{P} = \langle \mathcal{F}, \mathcal{A}, \varrho \rangle$ , where  $\mathcal{F} = \{f/1\}$ ,  $\mathcal{R} = \{R/1, Q/1\}$ ,  $\mathcal{I}_0 = \{R(a)\}$ ,  $\varrho = \{true \mapsto \alpha\}$ ,  $\mathcal{A} = \{\alpha\}$  and  $\alpha : \{R(x) \rightsquigarrow R(x), R(x) \rightsquigarrow Q(f(x)), Q(x) \rightsquigarrow Q(x)\}$ .

$\mathcal{S}$  is not GR-acyclic, because each  $R$  tuple is continuously copied, and at the same time continuously issues a call to service  $f$  that is then stored into a  $Q$  tuple, which is continuously copied as well. This is attested by the dataflow graph of Figure 8(b)).

The overall effect caused by the iterated application of  $\alpha$  is that fresh values are continuously generated and accumulated, making  $\mathcal{S}$  state-unbounded. Consider for example the application



**Figure 8 – Examples of dependency graphs for DCDSs with nondeterministic services; special edges are decorated with \***



of action  $\alpha$  in state  $\mathcal{I}_0$ . It leads to an infinite number of successors, each one of the form  $\{R(a), Q(v)\}$  where  $v$  is the value returned by  $f(a)$ . Consider now a second application of  $\alpha$  in one of these states. It again leads to an infinite number of successors, due to the nondeterminism of  $f(a)$ . In particular, each successor has the form  $\{R(a), Q(v), Q(v')\}$ , where  $v'$  is the result of the second call  $f(a)$ . When  $v' \neq v$ , the number of tuples is increased from 2 to 3. By executing  $\alpha$  over and over again, for some successors the value returned by a new call  $f(a)$  will be distinct from all the ones already stored in  $Q$ . This causes an indefinite increment of the database size due to the continuous insertion of fresh  $Q$  tuples. Such behavior is clearly shown in the concrete transition system of  $\mathcal{S}$ , depicted in Figure 6(a). Figure 6(b) shows instead one possible corresponding abstraction; even if the abstraction approach ensures that the generated transition system is finite-branching, some of its runs pass through an infinite number of distinct, growing states. ■

**Example 8.3.** Let  $\mathcal{S} = \langle \mathcal{D}, \mathcal{P} \rangle$  be a DCDS with data layer  $\mathcal{D} = \langle \mathcal{C}, \mathcal{R}, \emptyset, \mathcal{I}_0 \rangle$  and process layer  $\mathcal{P} = \langle \mathcal{F}, \mathcal{A}, \varrho \rangle$ , where  $\mathcal{F} = \{f/1, g/1\}$ ,  $\mathcal{R} = \{R/1\}$ ,  $\mathcal{I}_0 = \{R(a)\}$ ,  $\varrho = \{true \mapsto \alpha\}$ ,  $\mathcal{A} = \{\alpha\}$  and  $\alpha : \{R(x) \rightsquigarrow \{R(f(x)), R(g(x))\}\}$ .

$\mathcal{S}$  is not GR-acyclic, because its dataflow graph, shown in Figure 8(c), contains a unique node  $R$  with two distinct special looping edges from  $R$  to  $R$  itself. Indeed, every time  $\alpha$  is executed, each  $R$  tuple contained in the current database may generate two  $R$  tuples in the next state, such that each such new tuple contains a value different from all the other ones. Therefore, even if the newly generated values are not accumulated, in the “worst” case the number of  $R$  tuples is doubled every time  $\alpha$  is executed. A sample run of the system could be the following. Starting from  $\mathcal{I}_0$ ,  $\alpha$  calls  $f(a)$  and  $g(a)$ , getting  $b$  and  $c$  as result and obtaining the state  $\{R(b), R(c)\}$ . A second execution of  $\alpha$  involves now 4 service calls ( $f(b), g(b), f(c), g(c)$ ), which may return 4 different new values, e.g. leading to state  $\{R(d), R(e), R(f), R(g)\}$ , and so on. ■

**GR<sup>+</sup>-Acyclicity.** This relaxation of GR-acyclicity is based on the insight that, for a cycle  $\Upsilon_{\mathcal{S}}$  in the dataflow graph to truly preserve the number of values moving in it,  $\Upsilon_{\mathcal{S}}$ 's edges must not all be simultaneously inactive. We say that an edge is *active* in a step of the run when some action corresponding to it executes. By the DCDS semantics, if all edges of  $\Upsilon_{\mathcal{S}}$  are simultaneously inactive, then none of the corresponding copy/call operations are executed and all relations involved in  $\Upsilon_{\mathcal{S}}$  forget their value in the next state.  $\Upsilon_{\mathcal{S}}$  is effectively flushed.

GR<sup>+</sup>-Acyclicity is a relaxation that does allow path  $\pi = \pi_1\pi_2\pi_3$  as in the definition of GR-Acyclicity, provided that  $\pi_2$  contains an edge  $e$  that cannot be active at the same time as any of the subsequent edges in  $\pi_2\pi_3$ .

Semantically this ensures that in order for the generate cycle  $\pi_1\pi_2$  to push fresh values toward recall cycle  $\pi_3$ , some action corresponding to  $e$  must execute, and in the meantime all actions maintaining the values in cycle  $\pi_3$  are disabled, thus flushing  $\pi_3$ .  $\pi_3$  thus receives an unbounded number of waves of fresh values from  $\pi_1\pi_2$ , but it forgets each wave before the next arrives.

Of course, the property of being active at the same time is semantic in nature, but we give a sufficient syntactic condition. Associate with every edge  $e$  in the dataflow graph the set *actions*( $e$ ) of actions it corresponds to (this set can be computed via simple inspection of the process layer). Then edges  $e_1, e_2$  are not simultaneously active if *actions*( $e_1$ )  $\cap$  *actions*( $e_2$ ) =  $\emptyset$ .

The DCDSs discussed in Example 8.2 and 8.3 are not GR<sup>+</sup>-acyclic. Indeed, they are not GR-acyclic, and all the edges contained in their dataflow graphs can be simultaneously active, because they all correspond to a single action.

We observe that GR-acyclicity is not related to weak acyclicity. In particular, a DCDS may be GR-acyclic but not weakly acyclic (see Example 8.1).

As with any sufficient syntactic condition for an undecidable semantic property, an infinite succession of refinements of GR-acyclicity is possible, each relaxing the condition to allow more DCDS classes. We propose a very powerful relaxation in Appendix A.3, GR<sup>+</sup>-acyclicity. Appendix 9 shows a full-fledged DCDS example that conforms to GR<sup>+</sup>-acyclicity, showing that it admits a practically relevant DCDS class.

Theorem 8.6 and Theorem 8.3 imply:

**Theorem 8.7.** *Verification of  $\mu\mathcal{L}_P$  properties for  $GR^+$ -acyclic DCDS with nondeterministic services is decidable.*

## 9 Example DCDS: Travel Reimbursement System

We model the process of reimbursing travel expenses in a university, and the corresponding audit system, in two different subsystems. In particular, the first subsystem, called the *request system* manages the submission of reimbursement requests by an employee, and preliminary inspection and approval of the request by an employee in the accounting department (we shall call her the *monitor*). A log of accepted requests will be submitted to the second subsystem, the *audit system*, in which requests can be accumulated, and they can be checked for accuracy by calling external web services (for instance to obtain the exchange rate from foreign currency to USD on a past date, or to check that the employee actually was on the declared flight).

**Request system.** To keep the example simple we model a travel reimbursement request as being associated to the name of the requester, and information related to the corresponding flight and hotel costs. After a request is submitted, a monitor will check the request and will decide to accept or reject the request. If a request is rejected, the employee needs to modify the information regarding hotel and flight, while employee name will not be changed while updating. After the update by the employee, the monitor will again check the request, and the reject-check loop continues until the monitor accepts the request. After a request is accepted a log of the request will be sent to the audit system, and the request system will be ready to process the next travel request.

We model the request system by a DCDS  $\mathcal{S}_R = \langle \mathcal{D}, \mathcal{P} \rangle$ , in which  $\mathcal{D} = \langle \mathcal{C}, \mathcal{R}, \mathcal{E}, \mathcal{I}_0 \rangle$  such that  $\mathcal{C}$  is a countably infinite set of *constants*,  $\mathcal{I}_0 = \{\text{Status}(\text{'readyForRequest'}, \text{true})\}$ , and  $\mathcal{R}$  is a database schema as follows:

- **Status** =  $\langle \text{status} \rangle$ , a unary relation that keeps the state of the request subsystem, and can take three different values: *'readyForRequest'*, *'readyToVerify'*, and *'readyToUpdate'*,
- **Travel** =  $\langle \text{eName} \rangle$ , holding the name of the employee;
- **Hotel** =  $\langle \text{hName}, \text{date}, \text{price}, \text{currency}, \text{priceInUSD} \rangle$ , holding the hotel cost information of the employee's travel, which might have been paid in some other currency than USD,
- **Flight** =  $\langle \text{date}, \text{fNum}, \text{price}, \text{currency}, \text{priceInUSD} \rangle$ , holding the flight cost information.

The process layer is defined as  $\mathcal{P} = \langle \mathcal{F}, \mathcal{A}, \rho \rangle$  where  $\mathcal{F}$  is a set of the following nondeterministic service calls, each modeling an input of an external value by the employee. Specifically,

- **INENAME()** : models the input of the employee name (filled in by the employee),
- **INHNAME()** : hotel name,
- **INHDATE()** : arrival date,
- **INHPRICE()** : sum paid to the hotel (possibly in foreign currency),
- **INHCURRENCY()** : currency exchange rate at that date,
- **INHPINUSD()** : amount paid to the hotel in USD,
- **INFDATE()** : flight date,
- **INFNUM()** : flight number,
- **INFPRICE()** : ticket price, possibly in foreign currency,
- **INFCURRENCY()** : currency exchange rate at date of ticket payment,
- **INFPUSD()** : ticket price in USD.

There is one additional service.

- `MAKEDECISION()` : a nondeterministic service modeling the decision of the human monitor. It returns *requestConfirmed* if the request is accepted, and returns *readyToUpdate* if the request needs to be updated by the employee.

The set  $\mathcal{A}$  of actions includes (among others):

*InitiateRequest* :

```

true  $\rightsquigarrow$  Status('readyToVerify')
true  $\rightsquigarrow$  Travel(INENAME())
true  $\rightsquigarrow$  Hotel(INHNAME(),
                INHDATE(),
                INHPRICE(),
                INHCURRENCY(),
                INHPIINUSD())
true  $\rightsquigarrow$  Flight(INFDATE(),
                    INFNUM(),
                    INFPRICE(),
                    INFCURRENCY(),
                    INFPUSD())

```

*VerifyRequest*:

```

true  $\rightsquigarrow$  Status(MAKEDECISION())
Travel( $n$ )  $\rightsquigarrow$  Travel( $n$ )
Hotel( $x_1, \dots, x_5$ )  $\rightsquigarrow$  Hotel( $x_1, \dots, x_5$ )
Flight( $x_1, \dots, x_5$ )  $\rightsquigarrow$  Flight( $x_1, \dots, x_5$ )

```

*UpdateRequest*:

```

true  $\rightsquigarrow$  Status('readyToVerify')
Travel( $n$ )  $\rightsquigarrow$  Travel( $n$ )
true  $\rightsquigarrow$  Hotel(INHNAME(),
                INHDATE(),
                INHPRICE(),
                INHCURRENCY(),
                INHPIINUSD())
true  $\rightsquigarrow$  Flight(INFDATE(),
                    INFNUM(),
                    INFPRICE(),
                    INFCURRENCY(),
                    INFPUSD())

```

*AcceptRequest*:


```

Status('requestConfirmed')  $\rightsquigarrow$  Status('readyForRequest')

```

When a request is initiated (modeled by the action *InitiateRequest*), (i) the system changes state “to waiting for verification”, (ii) a travel event is generated and the employee fills in his name, (iii) the employee fills in all hotel information, and (iv) the employee fills in all flight information.

Action *VerifyRequest* models the preliminary check by the monitor. Travel event, hotel and flight information are unchanged, but the system status is set by the non-deterministic service call `MAKEDECISION()`, which models the monitor’s decision for current active travel information.



**Figure 9 – Dataflow graph for request system**

If the monitor rejects, then she sets the next state to *‘readyToUpdate’*, which will trigger the action *UpdateRequest*, which in turn collects once again the hotel and flight information from the employee, and moves the status to *‘readyToVerify’*.

Finally, action *AcceptRequest* returns the system in the state *‘readyForRequest’*, in which it is ready to accept a new request.

Notice the use of the always true predicate *true*, with the evident meaning. A convenient way to model its meaning in the DCDS framework is to think of it as a nullary relation, initialized to contain the empty tuple, which is copied in perpetuity by each action (*true* never changes its interpretation). We omit the corresponding copy effects, treating them as built-in.

Notice how the condition-action rules in the set  $\rho$  below guard the actions by the current state of the system:

$$\begin{aligned}
 \text{Status} ('readyForRequest') &\mapsto \text{InitiateRequest} \\
 \text{Status} ('readyToVerify') &\mapsto \text{VerifyRequest} \\
 \text{Status} ('readyToUpdate') &\mapsto \text{UpdateRequest} \\
 \text{Status} ('requestConfirmed') &\mapsto \text{AcceptRequest}
 \end{aligned}$$

The dataflow graph corresponding to the request system is depicted in Figure 9, where special edges are starred. Notice that there can be multiple special edges between the same two nodes (these are distinguished by unique edge ids, which we omit in the figure to avoid clutter).

In particular, the group of special edges from the *true* node to the *Hotel* node corresponds to the action of employee filling in the hotel information, modeled by calls to such services as *INHNAME()*. Similarly for the special edges from *true* to *Flight*.

The special edge between *true* and *Travel* is due to the employee filling in his name into the created travel request. The special edge from *true* to *Status* reflects the monitor’s insertion of her decision (see the call to *MAKEDECISION()* in the first effect of action *VerifyRequest* in Example 9), while the normal edge corresponds to change of the status without calling a service (this happens in other actions). The self-loops on *Flight*, *Hotel*, and *Travel* are due to the remaining (copy) effects of *VerifyRequest*. The self-loop on node *true* is due to the modeling of this value by a singleton nullary relation containing the empty tuple, which keeps being copied in each action.

An inspection of this dataflow graph reveals that the request system is not GR-acyclic, since it contains several instances of two simple cycles connected by a path that includes a special edge. For instance, the path  $\pi$  comprised of the self-loops around *true* and *Travel*, and the special edge between them. However, the request system is GR<sup>+</sup>-acyclic. To illustrate this, notice that the path  $\pi$  is allowed by GR<sup>+</sup>-acyclicity because the special edge leading into the *Travel* loop is caused by action *InitiateRequest*, while all the subsequent edges in  $\pi$  are caused by other actions (in this case there is only one subsequent edge in  $\pi$ , namely the self-loop on *Travel*, caused by actions *VerifyRequest* and *UpdateRequest*).

We illustrate some  $\mu\mathcal{L}_P$  properties pertaining to the proper operation of the request system:

A property of interest is that once initiated, a request will eventually be decided by the monitor, and the decision can only be *‘readyToUpdate’* or *‘requestConfirmed’* (a liveness property). We show the property in the easier-to-read CTL syntactic sugar:

$$\begin{aligned}
 \mathbf{AG}(\forall n. \text{Travel}(n) \rightarrow \\
 \mathbf{A}(\text{Travel}(n)\mathbf{U}(\text{Status}('readyToUpdate') \vee \\
 \text{Status}('requestConfirmed')))
 \end{aligned}$$

The until operator  $\mathbf{U}$  (for this example, it is the strong flavor, in which  $\psi\mathbf{U}\phi$  means that  $\phi$  is guaranteed to eventually hold, and until it does  $\psi$  must hold in every step). We note that for a property to belong to  $\mu\mathcal{L}_P$ , it must require the bindings of quantified variables to be continuously live between the step when the quantification was evaluated and the step when the variable is used. This can be done by using `LIVE` or by using any relation, in our example `Travel`. The  $\mu\mathcal{L}_P$  version of the property is given below:

$$\begin{aligned} & \nu X. (\forall n. \text{Travel}(n) \rightarrow \\ & \quad \mu Y. (\text{Status}(\text{'readyToUpdate'}) \vee \text{Status}(\text{'requestConfirmed'}) \\ & \quad \vee [-](\text{Travel}(n) \wedge Y))) \wedge [-]X \end{aligned}$$

Another property of interest is that if the flight cost is not specified, then the request is not accepted (a safety property). We use the special constant  $\perp$  to denote a null value (this need not be treated specially in the semantics, any value of the domain can be reserved for this purpose):

$$\mathbf{G}\neg(\text{Status}(\text{'requestConfirmed'}) \wedge \exists x_1, \dots, x_4. \text{Flight}(x_1, x_2, \perp, x_3, x_4)).$$

The  $\mu\mathcal{L}_P$  version is given below:

$$\begin{aligned} & \nu X. (\neg(\text{Status}(\text{'requestConfirmed'}) \wedge \\ & \quad \exists x_1, \dots, x_4. \text{Flight}(x_1, x_2, \perp, x_3, x_4))) \wedge [-]X \end{aligned}$$

**Audit system.** After a request is verified by the monitor in the request system, it will be migrated to the audit system. The migration is performed by a logging subsystem which might perform such operations as: we extend each travel event with a freshly generated travel id, which guarantees uniqueness across the entire history of requests. We store these tuples in a database. We can model this migration using the DCDS formalism, but we omit the specification and focus directly on the audit system.

More specifically, we model the audit system by a DCDS  $\mathcal{S}_A = \langle \mathcal{D}_A, \mathcal{P}_A \rangle$ , in which  $\mathcal{D}_A = \langle \mathcal{C}, \mathcal{R}, \mathcal{E}, \mathcal{I}_0 \rangle$ .  $\mathcal{C}$  is a countably infinite set of *constants*.  $\mathcal{R}$  is a database schema as follows:

- **Status** =  $\langle \text{status} \rangle$  is a unary relation keeping the state of request subsystem, which can take two different values: *checkPrice*, and *checkTravel*, whose role is to sequence the actions of the audit system appropriately.
- **Travel** =  $\langle \text{id}, \text{eName}, \text{passed} \rangle$  extends the homonymous relation of the request system with two fields: *id* (the travel identifier), and *passed*, which will be set by the audit system to reflect whether both the hotel and the flight price checks succeed.
- **Hotel** =  $\langle \text{trld}, \text{hName}, \text{date}, \text{price}, \text{currency}, \text{priceInUSD}, \text{passed} \rangle$ , where *trld* is a foreign key to the travel id and *passed* is set by the audit system to reflect whether the claimed price and the calculated price match.
- **Flight** =  $\langle \text{trld}, \text{fNum}, \text{date}, \text{price}, \text{currency}, \text{priceInUSD}, \text{passed} \rangle$ , where *trld* and *passed* are analogous to the ones in the **Hotel** relation.

Finally,  $\mathcal{I}_0$  is the output of the logging subsystem to which we add the fact  $\text{Status}(\text{'checkPrice'})$ , to initialize the audit system status.

The process layer is defined as  $\mathcal{P} = \langle \mathcal{F}, \mathcal{A}, \varrho \rangle$  in which  $\mathcal{F}$  contains a deterministic service, where the call `CONVERTANDCHECK(price, currency, date, priceInUSD)` performs the official exchange rate acquisition and computation described above, returning true if and only if the claimed price and the computed one match.

$\mathcal{A} = \{ \text{CheckPrice}, \text{CheckTravel} \}$  includes the following actions.

*CheckPrice* :

$$\begin{aligned} & \text{true} \rightsquigarrow \text{Status}('checkTravel') \\ & \text{Travel}(i, n, v) \rightsquigarrow \text{Travel}(i, n, v) \\ & \text{Hotel}(x_1, x_2, date, price, currency, priceInUSD, x_7) \rightsquigarrow \\ & \quad \text{Hotel}(x_1, x_2, date, price, currency, priceInUSD, \\ & \quad \quad \text{CONVERTANDCHECK}(date, price, currency, priceInUSD)) \\ & \text{Flight}(x_1, x_2, date, price, currency, priceInUSD, x_7) \rightsquigarrow \\ & \quad \text{Flight}(x_1, x_2, date, price, currency, priceInUSD, \\ & \quad \quad \text{CONVERTANDCHECK}(date, price, currency, priceInUSD)) \end{aligned}$$

Notice that the first effect changes the audit system's state to enter the stage in which the the two checks (for hotel and flight) are combined. The second effect simply copies the request information. The third and fourth each check the claimed price (for hotel, respectively flight), performing the conversion described above.

The second action works on the result of the first (this is ensured by the appropriate status changes).

*CheckTravel*:

$$\begin{aligned} & \text{true} \rightsquigarrow \text{Status}('checkPrice') \\ & \text{Travel}(x_1, x_2, x_3) \wedge \\ & \text{Hotel}(x_1, y_1 \dots, y_5, p_h) \wedge \\ & \text{Flight}(x_1, z_1 \dots, z_7, p_f) \wedge \neg(p_h \wedge p_f) \rightsquigarrow \text{Travel}(x_1, x_2, \text{false}) \\ & \text{Travel}(x_1, x_2, x_3) \wedge \\ & \text{Hotel}(x_1, y_1 \dots, y_5, \text{true}) \wedge \\ & \text{Flight}(x_1, z_1 \dots, z_7, \text{true}) \rightsquigarrow \text{Travel}(x_1, x_2, \text{true}) \\ & \text{Hotel}(x_1, \dots, x_7) \rightsquigarrow \text{Hotel}(x_1, \dots, x_7) \\ & \text{Flight}(x_1, \dots, x_7) \rightsquigarrow \text{Flight}(x_1, \dots, x_7) \end{aligned}$$

Notice that the second and third effects set the `passed` field for the request, computed as the conjunction of the corresponding fields set by the price check on flight and hotel.

The process  $\varrho$  is defined as follows:


$$\begin{aligned} & \text{Status}('checkPrice') \mapsto \text{CheckPrice} \\ & \text{Status}('checkTravel') \mapsto \text{CheckTravel} \end{aligned}$$

The corresponding dependency graph is as shown in Figure 10. In this picture nodes correspond to the positions of the schema. To avoid clutter, we represent each relation by its first letter, and denote the position number with a subscript. For instance,  $T_1$  stands for the first (id) position of the relation `Travel`, and  $S$  stands for the only position of the relation `Status`. Moreover, the edges without label represent regular edges in the dependency graph, while the starred edges depict special edges. For instance, the edge  $F_5 \xrightarrow{*} F_7$  is introduced due to the fourth effect of action *CheckPrice*. It is starred because it reflects the service call of `CONVERTANDCHECK`, taking as argument the currency attribute of `Flight` (at position 5), and storing its result in a `Flight` tuple at position 7 (the `passed` attribute).

An inspection of the dependency graph reveals that the audit system is weakly acyclic, since there is no cycle including a special edge.

We illustrate a desirable property of the audit system: it guarantees that the request cannot pass the audit if one of the flight or hotel checks fail:

$$\begin{aligned} & \mathbf{A} \mathbf{G}(\exists i, n, v, x_2, \dots, x_6. \text{Travel}(i, n, v) \wedge \\ & \quad (\text{Hotel}(i, x_2, \dots, x_6, \text{false}) \vee \text{Flight}(i, x_2, \dots, x_6, \text{false})) \\ & \quad \rightarrow \mathbf{F} \text{Travel}(i, n, \text{false})) \end{aligned}$$



**Figure 10 – Weakly-acyclic dependency graph of the audit system**

The  $\mu\mathcal{L}_A$  version of the property is given below:

$$\begin{aligned} & \nu X. (\exists i, n, v, x_2, \dots, x_6. \text{Travel}(i, n, v) \wedge \\ & \quad (\text{Hotel}(i, x_2, \dots, x_6, \text{false}) \vee \text{Flight}(i, x_2, \dots, x_6, \text{false})) \\ & \quad \rightarrow \mu Y. (\text{Travel}(i, n, \text{false}) \vee \langle - \rangle Y)) \wedge [-] X \end{aligned}$$

Notice that, since the audit system uses deterministic services, if we wish to verify it in isolation from the other subsystems, we can verify an  $\mu\mathcal{L}_A$  property, which is what the above is (we are not enforcing the liveness of the variables  $i, v, n$  between the step at which the quantification was evaluated, and the eventual step when the passed attribute of `Travel` is set to false).

As we will show in Section 10, we can verify mixed semantics DCDS by reduction to non-deterministic services. If we wished to verify the above property over the collection of subsystems, we would have to express it as an  $\mu\mathcal{L}_P$  property. This is easily done using an until operator `U` (as illustrated above for the request system). Moreover, it is actually compatible with our expectation about the system's operation: while a request is being audited, we expect it to persist in the system.

## 10 Discussion

**Complexity.** Both in the case of weakly acyclic DCDSs with deterministic services and of  $\text{GR}^+$ -acyclic DCDSs with non-deterministic services, our construction generates a finite transition system whose number of states is exponential in the size of the DCDS. Let  $\Phi$  be a  $\mu\mathcal{L}_A$  or  $\mu\mathcal{L}_P$  formula of size  $\ell$  with  $k$  alternating nested fixpoints. Then, considering the complexity or propositional  $\mu$ -calculus model checking on finite transition systems [35], the complexity of verification of  $\Phi$  over a DCDS of size  $n$  is  $O(2^n \cdot n^\ell)^k$ , hence in EXPTIME.

**Comparison of the two semantics.** It is natural to ask how the expressivities of the two DCDS flavors compare. Interestingly, we can show that for unrestricted DCDSs, the two semantics are equivalent from the point of view of expressive power, i.e. any DCDS with deterministic services can be simulated by a DCDS with nondeterministic services, and conversely.

However, we show below that the two semantics are not equivalent with respect to decidability of verification.

Consider first the reduction from deterministic to non-deterministic services.

**Theorem 10.1.** *Let  $D$  be a DCDS with deterministic services and schema  $\mathcal{R}_D$ . Then one can rewrite  $D$  in linear time to a DCDS  $N$  with nondeterministic services and schema  $\mathcal{R}_N$ , such that (i)  $\mathcal{R}_N$  includes  $\mathcal{R}_D$  and (ii) the projection of  $\Upsilon_N$  to  $\mathcal{R}_D$  coincides with  $\Upsilon_D$ , and (iii) if  $D$  is run-bounded, then  $N$  is state-bounded.*

We turn next to the converse reduction.

**Theorem 10.2.** *Let  $N$  be a DCDS with nondeterministic services and schema  $\mathcal{R}_N$ . Then one can rewrite  $N$  in linear time to a DCDS  $D$  with deterministic services and schema  $\mathcal{R}_D$ , such that (i)  $\mathcal{R}_D$  includes  $\mathcal{R}_N$  and (ii) the projection of  $\Upsilon_D$  to  $\mathcal{R}_N$  coincides with  $\Upsilon_N$ .*

The above reductions show that for unrestricted DCDS, deterministic and nondeterministic services are equivalent with respect to expressive power. However, they are not equivalent with respect to decidability of verification. This is because state-boundedness of the DCDS with nondeterministic services does not imply run-boundedness of the rewritten DCDS with deterministic services. In fact, one can prove that there exists no reduction from state-bounded DCDS with nondeterministic services to run-bounded DCDS with deterministic services: recall that for properties from  $\mu\mathcal{L}_A - \mu\mathcal{L}_P$ , verification is decidable for the latter (by Theorem 7.2), and undecidable for the former (Theorem 8.2). In particular, the reduction we use to prove Theorem 10.2 yields a non-weakly-acyclic DCDS and is therefore not pertinent to verification decidability.

In contrast, for the converse reduction of Theorem 10.1, observe that whenever  $D$  is run-bounded,  $N$  is state-bounded. Therefore, if we restrict the property language to  $\mu\mathcal{L}_P$ , decidability of verification for run-bounded DCDSs with deterministic services follows as a corollary of the reduction. Recall however that decidability holds even for the larger logic  $\mu\mathcal{L}_A$  (by Theorem 7.2). Our proof of Theorem 7.2 exploits the reduction as well, though additional technical contribution is needed to handle  $\mu\mathcal{L}_A$ .

**Mixed semantics.** The reduction in Theorem 10.1 allows us to verify  $\mu\mathcal{L}_P$  properties for DCDSs with a *mix* of deterministic and nondeterministic services, by first rewriting to a DCDS with exclusively nondeterministic services (as long as the rewritten DCDS is GR-acyclic). We give an example of a DCDS with mixed service semantics in Section 9.

**Support for arbitrary integrity constraints.** We remark that, by exploiting the equality constraints, we can extend our decidability results to support integrity constraints on the database expressed as arbitrary FO sentences under the active domain semantics. First, note that the definition of DCDS semantics is independent of the type of constraints used, as it simply requires their satisfaction by each state of the concrete transition system. Now consider a DCDS  $\mathcal{S}$  with an FO integrity constraint  $IC$  defined on its schema. We can rewrite  $\mathcal{S}$  to enforce  $IC$  using equality constraints. To this end, we add a binary auxiliary relation  $aux$  to the schema, initialized in the initial state to contain the tuple  $\langle a, b \rangle$  of distinct constants. We add to each action an effect that simply copies  $aux$  between states, ensuring the persistence of fact  $aux(a, b)$  throughout the run. Finally, we add an equality constraint  $ec := \neg IC \wedge aux(x, y) \rightarrow x = y$ . Notice now that  $\mathcal{S}'$  will never execute an action that violates  $IC$ , because that would violate  $ec$ . Equality constraints also prove instrumental in modeling artifact systems, described next.

**Connection with the artifact model.** In terms of expressive capabilities, our DCDS model is equivalent to a business process model known in the literature as the *artifact model* (see Section II). While variations thereof abound, they are virtually all special cases of the following general model. In it, given a relational schema  $\mathcal{T}$ , an *artifact of type  $\mathcal{T}$*  (or  $\mathcal{T}$ -artifact) is a tuple of schema  $\mathcal{T}$ . The attributes of the tuple are known as *artifact variables*, and they



	DETERMINISTIC SERVICES			NONDETERMINISTIC SERVICES			
	$\mu\mathcal{L}$	$\mu\mathcal{L}_A$	$\mu\mathcal{L}_P$		$\mu\mathcal{L}$	$\mu\mathcal{L}_A$	$\mu\mathcal{L}_P$
<i>unrestricted</i>	$\mathbf{U} \leftarrow \mathbf{U} \leftarrow \mathbf{U}$	$\mathbf{U} \leftarrow \mathbf{U}$	$\mathbf{U} \leftarrow \mathbf{U}$	<sup>1</sup>	<i>unrestricted</i>	$\mathbf{U} \leftarrow \mathbf{U} \leftarrow \mathbf{U}$	$\mathbf{U} \leftarrow \mathbf{U}$
<i>bounded-run</i>	$\mathbf{?}$	$\mathbf{D} \rightarrow \mathbf{D}$	$\mathbf{D}$		<i>bounded-state</i>	$\mathbf{U} \leftarrow \mathbf{U}$	$\mathbf{D}$

<sup>1</sup> The result is even stronger: it holds for propositional LTL.

<sup>2</sup> Decidability cannot be established via a faithful finite-state abstraction.

<sup>3</sup> Decidability is obtained via reduction to finite-state model checking.

**Table 1 – Summary of our decidability and undecidability results for DCDSs**

must include an *id* attribute that uniquely identifies each artifact. An *artifact system* has a schema comprising a collection of types  $\{T_i\}_{i \in \{1, \dots, n\}}$ , and the schema  $\mathcal{R}_{DB}$  of an underlying relational database. The *instance* of an artifact system consists of a relation  $I_i$  for each type  $T_i$  and a database of schema  $\mathcal{R}_{DB}$ . The artifact system also has a collection of actions (usually called “services”, a term we avoid here to rule out confusion with the external services of the DCDS model). The execution of an action evolves the current instance into its successor. Each action has a *pre-condition* which is a FO sentence over the artifact schema, evaluated over the current artifact instance under the active domain semantics. The pre-condition must hold for an action to be eligible to execute. Actions are also equipped with a *post-condition* which is usually an  $\exists$ FO formula relating the current and the successor instances (if  $R$  is a relation in the schema, the post-condition’s  $R$ -atoms refer to the current instance, while  $R'$  atoms refer to the successor). By  $\exists$ FO we mean existential FO logic, in which only existential quantifiers are allowed, and they must appear in the scope of an even number of negations. Existentially quantified variables are not interpreted over the active domain, but over the possible infinite domain. They model external inputs from the environment the artifact system evolves in.

While we do not show a formal reduction between the two models, we sketch here how a DCDS process can simulate an artifact-based one. The DCDS can model the sets  $I_i$  of  $T_i$ -artifacts using an integrity constraint to enforce the uniqueness of the *id* attribute. The pre-conditions of artifact actions correspond to the conditions in the DCDS condition-action rules. Artifact post-conditions  $\psi$  can be simulated by DCDS effects, after rewriting  $\psi$  to Skolem normal form and introducing for each resulting Skolem term a nondeterministic service call. The fact that post-conditions can contain disjunction while effects are conjunctive and positive is no impediment: the additional expressivity needed can be transferred to the DCDS condition-action rules, if necessary modeling one artifact transition step with several DCDS transition steps.

**Summary of results for DCDSs.** We summarize our results in Table 1 (arrows denote implications between results). We note that exhibiting a finite faithful abstraction of a concrete transition system is more than a means towards showing decidability, being a desirable goal in its own right as the most promising avenue towards practical implementation. Notice that we list as open the verification of  $\mu\mathcal{L}$  properties on bounded-run DCDSs with deterministic services, but recall from Section 7.2 that in this case there exists no faithful finite-state abstract transition system.

We believe that DCDSs provide a natural and expressive model for business processes powered by an underlying database, and thus are an ideal vehicle for foundational research with potential to transfer to alternative models.

Note that the design space for FO extensions of propositional  $\mu$ -calculus is broad, and notoriously contains bounded-state settings for which satisfiability of even modest extensions of propositional LTL is highly undecidable (e.g. LTL with the freeze quantifier over infinite data words [33]). In light of this, our decidability results come as a pleasant surprise, and the two  $\mu\mathcal{L}$  variants studied here, paired with the respective DCDS classes, strike a fortuitous balance

between expressivity and verification feasibility.

## Part IV

## Semantic Artifacts

## 11 Knowledge Base Formalism

For expressing knowledge bases, we use  $DL-Lite_{\text{NU}}$ , a variant of the  $DL-Lite_{\mathcal{A}}$  language [59, 24] in which we drop the *unique name assumption* (UNA) [3]. The syntax of *concept* and *role expressions* in  $DL-Lite_{\text{NU}}$  is as follows:

$$\begin{array}{ll} B ::= N \mid \exists R & R ::= P \mid P^- \\ C ::= B \mid \neg B & V ::= R \mid \neg R \end{array}$$

where  $N$  denotes a *concept name*,  $B$  a *basic concept*,  $C$  an *arbitrary concept*,  $P$  a *role name*,  $P^-$  an *inverse role*,  $R$  a *basic role*, and  $V$  an *arbitrary role*. A  $DL-Lite_{\text{NU}}$  *knowledge base* (KB) is a pair  $(T, A)$ , where  $T$  is a TBox, i.e., a finite set of *TBox assertions*, and  $A$  is an ABox, i.e., a finite set of *ABox assertions*, of the form:

$$\begin{array}{l} \text{TBox assertions: } B \sqsubseteq C \mid R \sqsubseteq V \mid (\text{funct } R) \\ \text{ABox assertions: } N(t_1) \mid P(t_1, t_2) \mid t_1 = t_2 \end{array}$$

where  $t_1$  and  $t_2$  denote individuals. Note that  $(\text{funct } R)$  asserts functionality of a role, and since we don't have the UNA we allow for *equalities* between individuals in the ABox. As usual in  $DL-Lite_{\mathcal{A}}$ , an atomic role  $P$  that appears in a functionality assertion must not appear *positively* in the right hand side of a TBox assertion, i.e., it must not appear in  $R \sqsubseteq P$  or  $R \sqsubseteq P^-$ , for some  $R$ .

We adopt the standard FOL semantics of DLs based on FOL interpretations  $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$  such that  $c^{\mathcal{I}} \in \Delta^{\mathcal{I}}$ ,  $N^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$ , and  $P^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$ . The semantics of the construct, of TBox and ABox assertions, and the notions of *satisfaction* and of *model* are as usual. We also say that  $A$  is *consistent wrt*  $T$  if  $(T, A)$  is satisfiable, i.e., admits at least one model.

Next we introduce queries. As usual (cf. OWL 2), answers to queries are formed by terms denoting individuals explicitly mentioned in the ABox. The *domain of an ABox*  $A$ , denoted by  $\text{ADOM}(A)$ , is the (finite) set of terms appearing in  $A$ . A *union of conjunctive queries* (UCQ)  $q$  over a KB  $(T, A)$  is a FOL formula of the form  $\exists \vec{y}_1. \text{conj}_1(\vec{x}, \vec{y}_1) \vee \dots \vee \exists \vec{y}_n. \text{conj}_n(\vec{x}, \vec{y}_n)$ , with free variables  $\vec{x}$  and existentially quantified variables  $\vec{y}_1, \dots, \vec{y}_n$ . Each  $\text{conj}_i(\vec{x}, \vec{y}_i)$  in  $q$  is a conjunction of atoms of the form  $N(z)$ ,  $P(z, z')$ , where  $N$  and  $P$  respectively denote a concept and a role name occurring in  $T$ , and  $z, z'$  are constants in  $\text{ADOM}(A)$  or variables in  $\vec{x}$  or  $\vec{y}_i$ , for some  $i \in \{1, \dots, n\}$ . The (*certain*) *answers* to  $q$  over  $(T, A)$  is the set  $\text{ans}(q, T, A)$  of substitutions<sup>11</sup>  $\sigma$  of the free variables of  $q$  with constants in  $\text{ADOM}(A)$  such that  $q\sigma$  evaluates to true in every model of  $(T, A)$ . If  $q$  has no free variables, then it is called *boolean* and its certain answers are either true or false.

**Theorem 11.1** ([3]). *Computing  $\text{ans}(q, T, A)$  of an UCQs  $q$  over a KBs  $(T, A)$  is PTIME-complete in the size of  $T$  and  $A$ .*

We also consider an extension of UCQs, called ECQs, which are queries of the query language  $EQL-Lite(\text{UCQ})$  [25], that is, the FOL query language whose atoms are UCQs evaluated according to the certain answer semantics above. An *ECQ* over  $T$  and  $A$  is a possibly open formula of the form (where  $q$  is a UCQ):

$$Q ::= [q] \mid [x = y] \mid \neg Q \mid Q_1 \wedge Q_2 \mid \exists x. Q$$

<sup>11</sup>As customary, we can view each substitution simply as a tuple of constants, assuming some ordering of the free variables of  $q$ .

The *answer to  $Q$  over  $(\mathcal{T}, \mathcal{A})$* , is the set  $\text{ANS}(Q, \mathcal{T}, \mathcal{A})$  of tuples of constants in  $\text{ADOM}(\mathcal{A})$  defined as follows:

$$\begin{aligned} \text{ANS}([q], T, A) &= \text{ans}(q, T, A) \\ \text{ANS}([x = y], T, A) &= \{\langle x, y \rangle \in \text{ADOM}(A) \mid (T, A) \models x=y\} \\ \text{ANS}(\neg Q, T, A) &= \text{ADOM}(A) \setminus \text{ANS}(Q, T, A) \\ \text{ANS}(Q_1 \wedge Q_2, T, A) &= \text{ANS}(Q_1, T, A) \cap \text{ANS}(Q_2, T, A) \\ \text{ANS}(\exists x.Q, T, A) &= \bigcup_{t \in \text{ADOM}(A)} \text{ANS}(Q_{x/t}, T, A) \end{aligned}$$

where  $Q_{x/t}$  is obtained from  $Q$  by substituting each occurrence of  $x$  with  $t$ . Following the line of the proof in [25], but considering Theorem 11.1 for the basic step of evaluating an UCQ, we get:

**Theorem 11.2.** *Computing  $\text{ANS}(Q, T, A)$  of an ECQs  $Q$  over a KBs  $(T, A)$  is PTIME-complete in the size of  $T$  and  $A$ .*

We close by recalling that DL-Lite enjoys the *FO rewritability* property, which in our setting says that for every UCQ  $q$ ,  $\text{ans}(q, T, A) = \text{ans}(\text{rew}(q), \emptyset, A)$ , where  $\text{rew}(q)$  is a UCQ computed by the reformulation algorithm in [26]. Notice that, in this way, we have “compiled away” the TBox, though we still need to do logical implication w.r.t. ABox, which contains equality assertions. This result can be extended to ECQs as well [25].

## 12 Formalizing Knowledge and Action Bases

A *Knowledge and Action Base (KAB)* is a tuple  $\mathcal{K} = (\mathcal{T}, A_0, \Gamma, \Pi)$  where  $T$  and  $A_0$  form the *knowledge component* (or knowledge base), and  $\Gamma$  and  $\Pi$  form the *action component* (or action base). In practice,  $\mathcal{K}$  is a stateful device that stores the information of interest into a KB, formed by a fixed TBox  $\mathcal{T}$  and an initial ABox  $A_0$ , which can evolve by executing actions  $\Gamma$  according to the sequencing established by process  $\Pi$ . We describe such components in detail.

### 12.1 Knowledge Component of a KAB

The knowledge component of a KAB is constituted by the TBox and ABox of a *DL-Lite<sub>NU</sub>* knowledge base.

**TBox.**  $\mathcal{T}$  is a *DL-Lite<sub>NU</sub>* TBox, used to capture the intensional knowledge about the domain of interest. Such a TBox is fixed once and for all, and does not evolve during the execution of the KAB.

**ABox.**  $A_0$  is a *DL-Lite<sub>NU</sub>* ABox, which stores the extensional information of interest. Notice that  $A_0$  is the ABox of the *initial state* of the KAB, and as the KAB evolves due to the effect of actions, the ABox, which is indeed the state of the system, evolves accordingly to store up-to-date information. Through actions, we acquire new information from the external world, which results in new individuals. These individuals are denoted by (ground) *Skolem terms*. The presence of Skolem terms has an impact on the treatment of equality, since in principle we need to close equality w.r.t. congruence, i.e., if  $a = b$  holds, then also  $f(a) = f(b)$  must hold. Closure w.r.t. congruence generates an infinite number of logically implied equality assertions. However, we are going to keep such assertions implicit, computing them only when needed. Observe that, given two complex terms verifying their equality requires a PTIME computation.

### 12.2 Action Component of a KAB

The action component of a KAB is constituted by a finite set actions that modify the current ABox, and by a possibly nondeterministic program that uses the KAB ABoxes to store its computation results, and the actions as atomic instructions.

**Actions.**  $\Gamma$  is a finite set actions. An *action*  $\gamma \in \Gamma$  modifies the current ABox  $A$  by adding or deleting assertions, thus generating a new ABox  $A'$ .  $\gamma$  is constituted by a signature and an effect specification. The *action signature* is constituted by a name and a list of individual *input parameters*. Such parameters need to be instantiated with individuals for the execution of the action. Given a substitution  $\theta$  for the input parameters, we denote by  $\gamma\theta$  the instantiated action with the *actual* parameters coming from  $\theta$ .<sup>12</sup> The *effect specification* consists of a set  $\{e_1, \dots, e_n\}$  of effects, which take place simultaneously. An *effect*  $e_i$  has the form

$$[q_i^+] \wedge Q_i^- \rightsquigarrow A'_i,$$

where

- $q_i^+$  is an UCQ;
- $Q_i^-$  is an arbitrary ECQ whose free variables occur all among the free variables of  $q_i^+$ ;
- $A'_i$  is a set of facts (over the alphabet of  $\mathcal{T}$ ) which include as terms: individuals in  $A_0$ , free variables of  $q_i^+$ , and Skolem terms  $f(\vec{x})$  having as arguments free variables  $\vec{x}$  of  $q_i^+$ .

Given the current ABox  $A$  of  $\mathcal{K}$  and a substitution  $\theta$  for the parameters of the action  $\gamma$ , the new state  $A'$  resulting from firing the action  $\gamma$  with parameters  $\theta$  on the state  $A$ , is computed as follows:

- (i) each effect  $e_i \in \gamma$  extracts from  $A$  the set  $\text{ANS}([q_i^+] \wedge Q_i^-)\theta, \mathcal{T}, A$  of tuples of terms in  $\text{ADOM}(A)$ , and for each such tuple  $\sigma$  asserts a set  $A'_i\theta\sigma$  of facts obtained from  $A'_i\theta$  by applying the substitution  $\sigma$  for the free variables of  $q_i^+$ . For each Skolem term  $f(\vec{x})\theta$  appearing in  $A'_i\theta$ , a new ground term is introduced having the form  $f(\vec{x})\theta\sigma$ . These terms represent new “constants” denoting “unknown” individuals. We denote by  $e_i\theta(A)$  the overall set of facts, i.e.,

$$e_i\theta(A) = \bigcup_{\sigma \in \text{ANS}([q_i^+] \wedge Q_i^-)\theta, \mathcal{T}, A} A'_i\theta\sigma.$$

- (ii) Moreover, let  $\text{EQ}(A) = \{t_1 = t_2 \mid \langle t_1, t_2 \rangle \in \text{ANS}([x_1 = x_2], \mathcal{T}, A)\}$ . Observe that, due to the semantics of queries, the terms in  $\text{EQ}(A)$  must appear explicitly in  $\text{ADOM}(A)$ , that is, the possibly infinite number of equalities due to congruence do not appear in  $\text{EQ}(A)$ , though they are logically implied.

The overall *effect* of the action  $\gamma$  with parameter substitution  $\theta$  over  $A$  is the new ABox  $A' = \text{DO}(\mathcal{T}, A, \gamma\theta)$  where

$$\text{DO}(\mathcal{T}, A, \gamma\theta) = \text{EQ}(A) \cup \bigcup_{1 \leq i \leq n} e_i\theta(A).$$

Let us make some observations on such actions. The effects of an action are a form of update of the previous state<sup>13</sup>, and not of belief revision [42]. That is, we never learn new facts on the state in which an action is executed, but only on the state resulting from the action execution. Skolem terms introduced by actions effects can be though of witnesses of new information coming from an external user/environment when executing the action. Their presence makes the domain of the ABoxes obtained by executing actions continuously changing. For simplicity, we do not make any persistence (or frame) assumption in our formalization (except for equality) [60]. In principle at every move we substitute the whole old state, i.e., ABox, with a new one. On the other hand, it should be clear that we can easily write effect specifications that *copy* big chunks of the old state into the new one. For example,  $P(x, y) \rightsquigarrow P(x, y)$  copies the entire

<sup>12</sup>We disregard a specific treatment of the output to the user, and assume instead that she can freely pose queries to the ontology, extracting implicit or explicit information from the states through which the KAB evolves.

<sup>13</sup>Our approach sidesteps the semantical and computational difficulties of description logic knowledge base update [50]. Adopting such forms of update in our setting is an interesting research issue.




Figure 11 – The TBox of the KAB for Example 11

set of assertions involving the role  $P$ . We do have a persistence assumption on equalities, we implicitly copy all equalities holding in the current state to the new one. This implies that, as the system evolves, we acquire new information on equalities between terms, but never lose equality information already acquired.

**Process.** The process component of a KAB is a possibly nondeterministic program that uses the KAB ABoxes to store its (intermediate and final) computation results, and the actions in  $\Gamma$  as atomic instructions. The ABoxes can be arbitrarily queried through the KAB TBox  $\mathcal{T}$ , while they can be updated only through actions in  $\Gamma$ . There are many ways to specify processes. We adopt a rule-based specification.

A *process* is a finite set  $\Pi$  of condition/action rules. A *condition/action rule*  $\pi \in \Pi$  is an expression of the form  $Q \mapsto \gamma$ , where  $\gamma$  is an action in  $\Gamma$  and  $Q$  is an ECQ over  $T$ , whose free variables are exactly the parameters of  $\gamma$ . The rule expresses that, for each tuple  $\theta$  for which condition  $Q$  holds, the action  $\gamma$  with actual parameters  $\theta$  can be executed. Processes do not force the execution of actions but constrain them: the user of the process will be able to choose any action that the rules forming the process allow. Moreover, our processes inherit entirely their states from the KAB knowledge component (TBox and ABox), see e.g., [31]. Other choices are also possible: the process could maintain its own state besides the one of the KAB. As long as such an additional state is finite, or embeddable into the KAB itself, the results here would easily extend to such a case.

**Example 12.1.** Let us consider a KAB  $\mathcal{K} = (T, A_0, \Gamma, \Pi)$  describing a super-heroes comics world, where we have cities in which characters live. Figure 11 shows a UML representation of the TBox  $T$  (see [24] for the correspondence between *DL-Lite<sub>A</sub>* and UML). Characters can be superheroes or (super)villains, who fight each other. As in the most classic paradigm, superheroes help the endeavors of law enforcement fighting villains threatening the city they live in. In fact, as a villain reveals himself for perpetrating his nefarious purposes against the city’s peace, he consequently becomes a declared enemy of all superheroes living in that city. Each character can live in one city at the time. A common trait of almost all superheroes is a secret identity: a superhero is said to be the alter ego of some character, which is his identity in common life. Hence, the ABox assertion  $\text{alterEgo}(s, p)$  means that the superhero  $s$  is the alter ego of character  $p$ . Villains always try to unmask superheroes, i.e., find their secret identity, in order to exploit such a knowledge to defeat them. Notice the subtle difference here: we use an  $\text{alterEgo}(s, p)$  assertion to model the fact that  $s$  is the alter ego of  $p$ , whereas only by asserting  $s = p$  we can capture the knowledge that  $s$  and  $p$  actually semantically denote the same individual.  $\Gamma$  includes the following actions:

$$\text{BecomeSH}(p, c) : \{ [\text{Character}(p) \wedge \exists v. \text{Villain}(v) \wedge \text{livesIn}(v, c)] \\ \rightsquigarrow \{ \text{Superhero}(f(p)), \text{alterEgo}(f(p), p) \}, \\ \text{CopyAll} \}$$

states that if there exists at least one villain living in the city  $c$ , a new superhero  $f(p)$  is created, with the purpose of protecting  $c$ . Such a superhero has  $p$  as alter ego. *CopyAll* is a shortcut for copying everything in the new state.

$$\text{Unmask}(s, p) : \{ [\text{alterEgo}(p, s)] \rightsquigarrow \{ s = p \}, \text{CopyAll} \}$$

states that superhero  $s$ , who is the alter ego of  $p$ , gets unmasked by asserting the equality between  $s$  and  $p$  (it is now known that  $s = p$ ).

$$\text{Fight}(v, s) : \{ \exists p. [\text{Villain}(v) \wedge \text{Character}(p) \wedge \text{alterEgo}(s, p)] \wedge [s = p] \rightsquigarrow \{\text{defeats}(v, s)\}, \\ \text{CopyAll} \}$$

states that when villain  $v$  fights superhero  $s$ , he defeats  $s$  if  $s$  has been unmasked, i.e., it is known that  $s$  is equal to his alter ego.

$$\text{Challenge}(v, s) : \{ [\text{Villain}(v) \wedge \text{Superhero}(s) \wedge \exists p. \text{alterEgo}(s, p) \wedge \text{livesIn}(p, sc)] \wedge \neg[\text{defeats}(v, s)] \\ \rightsquigarrow \{\text{livesIn}(v, sc), \text{enemy}(v, s)\}, \\ \text{CopyAll} \}$$

states that when villain  $v$  challenges superhero  $s$  and has not defeated him, next he lives in the same city as  $s$  and is enemy of  $s$ .

$$\text{ThreatCity}(v, c) : \{ [\text{Villain}(v) \wedge \text{Superhero}(s) \wedge \exists p. \text{alterEgo}(s, p) \wedge \text{livesIn}(p, c)] \\ \rightsquigarrow \{\text{enemy}(v, s) \wedge \text{livesIn}(v, c)\}, \\ \text{CopyAllExceptEnemy}, \\ [\text{Villain}(v) \wedge \text{enemy}(v', s')] \wedge \neg[v = v'] \rightsquigarrow \text{enemy}(v', s') \}$$

states that when villain  $v$  threatens city  $c$ , then he becomes an enemy of all and only superheroes that live in  $c$ . *CopyAllExceptEnemy* has the intended meaning.

Consider an initial ABox

$$A_0 = \{ \text{Superhero}(\text{batman}), \text{Villain}(\text{joker}), \text{alterEgo}(\text{batman}, \text{bruce}), \\ \text{livesIn}(\text{bruce}, \text{gotham}), \text{livesIn}(\text{batman}, \text{gotham}), \text{livesIn}(\text{joker}, \text{city1}) \}.$$

In this state, `bruce` and `batman` live in the same city, and `batman` is the alterego of `bruce` but it is not known whether they denote the same individual. Executing `Challenge(joker, batman)` in  $A_0$  generates a new ABox with added facts `enemy(joker, batman)`, `livesIn(joker, gotham)`, and `gotham = city1` is implied by the functionality on `livesIn`.

A process  $\Pi$  might include the following rules:

$$\begin{aligned} & [\text{Character}(p)] \wedge \neg[\text{Superhero}(p)] \wedge [\text{livesIn}(p, c)] \mapsto \text{BecomeSH}(p, c), \\ & [\text{Superhero}(s) \wedge \text{Character}(c)] \mapsto \text{Unmask}(s, c), \\ & [\text{enemy}(v, s)] \mapsto \text{Fight}(v, s), \\ & [\text{Villain}(v) \wedge \text{Superhero}(s)] \mapsto \text{Challenge}(v, s), \\ & [\text{Villain}(v) \wedge \text{City}(c)] \wedge \neg\exists v'([\text{Villain}(v') \wedge \text{livesIn}(v', c)] \wedge \neg[v = v']) \mapsto \text{ThreatCity}(v, c) \end{aligned}$$

For instance, the first one states that a character can become a superhero if it is not already one. ■

## 12.3 KAB Semantics

The semantics of KABs is given in terms of possibly infinite transition systems that represents the possible evolutions of the KAB over time as actions are executed according to the process. Notice that such transition systems must be equipped with semantically rich states, since a full ontology is associated to them.

**Transition systems.** A transition system  $\Upsilon$  is a tuple of the form  $(\mathbb{U}, \mathcal{T}, \Sigma, s_0, \text{abox}, \Rightarrow)$ , where:  $\mathbb{U}$  is a countably infinite set of terms denoting individuals;  $\mathcal{T}$  is a TBox;  $\Sigma$  is a set of states;  $s_0 \in \Sigma$  is the initial state; *abox* is a function that, given a state  $s \in \Sigma$ , returns an ABox associated to  $s$ , which has as individuals elements of  $\mathbb{U}$  and conforms to  $\mathcal{T}$ ;  $\Rightarrow \subseteq \Sigma \times \Sigma$  is a transition relation between pairs of states. For convenience, we define the active domain of the whole transition system as  $\text{ADOM}(\Upsilon) = \bigcup_{s \in \Sigma} \text{ADOM}(\text{abox}(s))$ .

**Transition system generated by a KAB.** Given a KAB  $\mathcal{K} = (\mathcal{T}, A_0, \Gamma, \Pi)$ , we define its (*generated*) *transition system*  $\Upsilon = (\mathbb{U}, \mathcal{T}, \Sigma, s_0, \text{abox}, \Rightarrow)$  as follows:

- $\mathbb{U}$  is formed by all constants and all Skolem terms inductively formed starting from  $\text{ADOM}(A_0)$  by applying the Skolem functions occurring in the actions in  $\Gamma$ ;
- $\mathcal{T}$  is the TBox of the KAB;
- $\text{abox}$  is the identity function (i.e., each state is simply an ABox);
- $s_0 = A_0$  is the initial state;
- $\Sigma$  and  $\Rightarrow$  are defined by mutual induction as the smallest sets satisfying the following property: if  $s \in \Sigma$  then for each rule  $Q \mapsto \gamma$ , evaluate  $Q$ , and for each tuple  $\theta$  returned, if  $\text{DO}(T, \text{abox}(s), \gamma\theta)$  is consistent w.r.t.  $\mathcal{T}$ , then  $s' = \text{DO}(T, \text{abox}(s), \gamma\theta)$  and  $s \Rightarrow s'$ .

Notice that  $\Upsilon$  is an infinite tree in general and, in fact, it is enough to perform infinitely a single action to obtain an infinite tree. Hence the classical results on model checking [30], which are developed for finite transition systems, cannot be applied directly.

## 13 Verification

To specify dynamic properties over a semantic artifacts, we use a first-order variant of  $\mu$ -calculus [61, 57]. (Temporal)  $\mu$ -calculus is virtually the most powerful temporal logic used for model checking of finite-state transition systems, and is able to express both linear time logics such as LTL and PSL, and branching time logics such as CTL and CTL\* [30]. The main characteristic of  $\mu$ -calculus is its 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 properties defined by induction or co-induction. This is the reason why virtually all logics used in verification can be considered as fragments of  $\mu$ -calculus. Technically,  $\mu$ -calculus separates local properties, asserted on the current state or on states that are immediate successors of the current one, and properties talking about states that are arbitrarily far away from the current one [61]. The latter are expressed through the use of fixpoints.

In this work, we use a first-order variant of temporal  $\mu$ -calculus where: on the one hand we allow local properties to be expressed as EQL queries, and on the other hand we allow arbitrary first-order quantification across states. Given the nature of EQL queries used for formulating local properties, first-order quantification must range over terms denoting individuals.

Formally, we introduce the logic  $\mu\mathcal{L}_A$  defined as follows:

$$\Phi ::= Q \mid \neg\Phi \mid \Phi_1 \wedge \Phi_2 \mid \exists x.\Phi \mid \langle - \rangle\Phi \mid Z \mid \mu Z.\Phi$$

where  $Q$  is a possibly open EQL query, and  $Z$  is a second order predicate variable (of arity 0). We make use of the following abbreviations:  $\forall x.\Phi = \neg(\exists x.\neg\Phi)$ ,  $\Phi_1 \vee \Phi_2 = \neg(\neg\Phi_1 \wedge \neg\Phi_2)$ ,  $[-]\Phi = \neg\langle - \rangle\neg\Phi$ , and  $\nu Z.\Phi = \neg\mu Z.\neg\Phi[Z/\neg Z]$ . The formulae  $\mu Z.\Phi$  and  $\nu Z.\Phi$  respectively denote the least and greatest fixpoint of the formula  $\Phi$  (seen as the predicate transformer  $\lambda Z.\Phi$ ). As usual in  $\mu$ -calculus, formulae of the form  $\mu Z.\Phi$  (and  $\nu Z.\Phi$ ) must obey to the *syntactic monotonicity* of  $\Phi$  wrt  $Z$ , which states that every occurrence of the variable  $Z$  in  $\Phi$  must be within the scope of an even number of negation symbols. This ensures that the least fixpoint  $\mu Z.\Phi$  (as well as the greatest fixpoint  $\nu Z.\Phi$ ) always exists.

The semantics of  $\mu\mathcal{L}_A$  formulae is defined over possibly infinite transition systems  $\langle \mathbb{U}, \mathcal{T}, \Sigma, s_0, \text{abox}, \Rightarrow \rangle$ . Since  $\mu\mathcal{L}_A$  also contains formulae with both individual and predicate free variables, given a transition system  $\Upsilon$ , we introduce an individual variable valuation  $v$ , i.e., a mapping from individual variables  $x$  to  $\mathbb{U}$ , and a predicate variable valuation  $V$ , i.e., a mapping from the predicate variables  $Z$  to a subset of  $\Sigma$ . With these three notions in place, we assign meaning to formulae by associating to  $\Upsilon$ ,  $v$ , and  $V$  an *extension function*  $(\cdot)_{v,V}^\Upsilon$ , which



$$\begin{aligned}
(Q)_{v,V}^{\Upsilon} &= \{s \in \Sigma \mid \text{ANS}(Qv, T, \text{abox}(s)) = \text{true}\} \\
(\neg\Phi)_{v,V}^{\Upsilon} &= \Sigma \setminus (\Phi)_{v,V}^{\Upsilon} \\
(\Phi_1 \wedge \Phi_2)_{v,V}^{\Upsilon} &= (\Phi_1)_{v,V}^{\Upsilon} \cap (\Phi_2)_{v,V}^{\Upsilon} \\
(\exists x.\Phi)_{v,V}^{\Upsilon} &= \{s \in \Sigma \mid \exists t.t \in \text{ADOM}(\text{abox}(s)) \text{ and } s \in (\Phi)_{v[x/t],V}^{\Upsilon}\} \\
(\langle - \rangle\Phi)_{v,V}^{\Upsilon} &= \{s \in \Sigma \mid \exists s'.s \Rightarrow s' \text{ and } s' \in (\Phi)_{v,V}^{\Upsilon}\} \\
(Z)_{v,V}^{\Upsilon} &= V(Z) \\
(\mu Z.\Phi)_{v,V}^{\Upsilon} &= \bigcap \{S \subseteq \Sigma \mid (\Phi)_{v,V[Z/S]}^{\Upsilon} \subseteq S\}
\end{aligned}$$

**Figure 12** – Semantics of  $\mu\mathcal{L}_A$ .

maps formulae to subsets of  $\Sigma$ . Formally, the extension function  $(\cdot)_{v,V}^{\Upsilon}$  is defined inductively as shown in Figure 12, where  $Qv$  stands for the query obtained from  $Q$  by substituting its free variables according to  $v$ .<sup>14</sup> Intuitively,  $(\cdot)_{v,V}^{\Upsilon}$  assigns to such constructs the following meaning:

- The boolean connectives have the expected meaning.
- The quantification of individuals is done over the objects of the “current” ABox.
- The extension of  $\langle - \rangle\Phi$  consists of the states  $s$  such that for *some* state  $s'$  with  $s \Rightarrow s'$ , we have that  $\Phi$  holds in  $s'$ , while the extension of  $[-]\Phi$  consists of the states  $s$  such that for *all* states  $s'$  with  $s \Rightarrow s'$ ,  $\Phi$  holds in  $s'$ .
- The extension of  $\mu X.\Phi$  is the *smallest subset*  $\mathcal{S}_\mu$  of  $\Sigma$  such that, assigning to  $Z$  the extension  $\mathcal{S}_\mu$ , the resulting extension of  $\Phi$  is contained in  $\mathcal{S}_\mu$ . That is, the extension of  $\mu X.\Phi$  is the *least fixpoint* of the operator  $(\Phi)_{v,V[Z/\mathcal{S}]}$ , where  $V[Z/\mathcal{S}]$  denotes the predicate valuation obtained from  $V$  by forcing the valuation of  $Z$  to be  $\mathcal{S}$ .
- Similarly, the extension of  $\nu X.\Phi$  is the *greatest subset*  $\mathcal{S}_\nu$  of  $\Sigma$  such that, assigning to  $X$  the extension  $\mathcal{S}_\nu$ , the resulting extension of  $\Phi$  contains  $\mathcal{S}_\nu$ . That is, the extension of  $\nu X.\Phi$  is the *greatest fixpoint* of the operator  $(\Phi)_{v,V[X/\mathcal{S}]}$ . Formally,  $(\nu Z.\Phi)_{v,V}^{\Upsilon} = \bigcup \{S \subseteq \Sigma \mid S \subseteq (\Phi)_{v,V[Z/S]}^{\Upsilon}\}$ .

**Example 13.1.** An example of  $\mu\mathcal{L}_A$  formula is:

$$\nu X.(\forall x.[\text{Superhero}(x)] \rightarrow \mu Y.([\text{alterEgo}(x, x)] \vee \langle - \rangle Y)) \wedge [-]X$$

It states that, along every path, it is always true, for each superhero, that there exists an evolution that eventually leads to unmask him. ■

When  $\Phi$  is a closed formula,  $(\Phi)_{v,V}^{\Upsilon}$  does not depend on  $v$  nor  $V$ , and we denote the extension of  $\Phi$  simply by  $(\Phi)^{\Upsilon}$ . A closed formula  $\Phi$  holds in a state  $s \in \Sigma$  if  $s \in (\Phi)^{\Upsilon}$ . In this case, we write  $\Upsilon, s \models \Phi$ . A closed formula  $\Phi$  holds in  $\Upsilon$ , denoted by  $\Upsilon \models \Phi$ , if  $\Upsilon, s_0 \models \Phi$ . We call *model checking* verifying whether  $\Upsilon \models \Phi$  holds.

Two transition systems are *behaviourally equivalent* if they satisfy exactly the same  $\mu\mathcal{L}_A$  formulas. To formally capture such an equivalence, we make use of the notion of *bisimulation* [54], suitably extended to deal with query answering over ontologies.

For transition systems  $\Upsilon_1 = \langle \mathbb{U}, \mathcal{T}_1, \Sigma_1, s_{01}, \text{abox}_1, \Rightarrow_1 \rangle$  and  $\Upsilon_2 = \langle \mathbb{U}, \mathcal{T}_2, \Sigma_2, s_{02}, \text{abox}_2, \Rightarrow_2 \rangle$ , a *bisimulation* between  $\Upsilon_1$  and  $\Upsilon_2$  is a relation  $\mathcal{B} \subseteq \Sigma_1 \times \Sigma_2$  such that:  $(s_1, s_2) \in \mathcal{B}$  implies that:

- $(T_1, \text{abox}(s_1))$  and  $(T_2, \text{abox}(s_2))$  are logically equivalent, i.e., for each ABox assertion  $\alpha_1 \in \text{abox}(s_1)$  we have that  $(T_2, \text{abox}(s_2)) \models \alpha_1$ , and for each ABox assertion  $\alpha_2 \in \text{abox}(s_2)$  we have that  $(T_1, \text{abox}(s_1)) \models \alpha_2$ ;

<sup>14</sup>Notice that it is built in the semantics of EQL queries that if  $v$  substitutes some free variable with an element of  $\mathbb{U}$  not occurring in  $\text{abox}(s)$ , then  $\text{ANS}(Qv, T, \text{abox}(s)) = \text{false}$  (c.f. Sec. 11); However, this does not happen for the KAB generated transition system by construction, due to the preservation of equality between domain individuals, which is reflexive.

- if  $s_1 \Rightarrow_1 s'_1$  then there exists  $s'_2$  such that  $s_2 \Rightarrow_2 s'_2$  and  $(s'_1, s'_2) \in \mathcal{B}$ ;
- if  $s_2 \Rightarrow_2 s'_2$  then there exists  $s'_1$  such that  $s_1 \Rightarrow_1 s'_1$  and  $(s'_1, s'_2) \in \mathcal{B}$ .

We say that two states  $s_1$  and  $s_2$  are *bisimilar*, if there exists a bisimulation  $\mathcal{B}$  such that  $(s_1, s_2) \in \mathcal{B}$ . Two transition systems  $\Upsilon_1$  with initial state  $s_{01}$  and  $\Upsilon_2$  with initial state  $s_{02}$  are *bisimilar* if  $(s_{01}, s_{02}) \in \mathcal{B}$ .

The following theorem states that the formula evaluation in  $\mu\mathcal{L}_A$  is indeed invariant wrt bisimulation, so we can equivalently check any bisimilar transition systems.

**Theorem 13.1.** *Let  $\Upsilon_1$  and  $\Upsilon_2$  be two bisimilar transition systems. Then, for two states  $s_1$  of  $\Upsilon_1$  and  $s_2$  of  $\Upsilon_2$  (including the initial ones) that are bisimilar, for all closed  $\mu\mathcal{L}_A$  formulas  $\Phi$ , we have that  $s_1 \in (\Phi)^{\Upsilon_1}$  iff  $s_2 \in (\Phi)^{\Upsilon_2}$ .*

*Proof.* The proof is analogous to the standard proof of bisimulation invariance of  $\mu$ -calculus [61], though taking into account our bisimulation, which guarantees that ECQs are evaluated identically over bisimilar states.  $\square$

We observe that in this invariance result, we make no use of the restriction on existentially quantified individuals belonging to the “current” active domain. However, we enforce this restriction to be able to talk only about individuals that are explicitly mentioned in the ABoxes of the transition system (i.e., in the active domain of the transition system) and not those that are implicitly present because of the congruence of equalities.

Making use of such a notion of bisimulation, we can redefine the transition system generated by KAB  $\mathcal{K} = (\mathcal{T}, A_0, \Gamma, \Pi)$  while maintaining bisimilarity, by modifying the definition of  $\Upsilon = \langle \mathbb{U}, \mathcal{T}, \Sigma, s_0, \text{abox}, \Rightarrow \rangle$  given in Section 12.3 as follows:

- We modify  $\text{DO}()$  so that no Skolem term  $t'$  is introduced in the generated ABox  $A'$  if in the current ABox<sup>15</sup>  $A$  there already is a term  $t$  such that  $(T, A) \models t = t'$ .
- If the ABox  $A' = \text{DO}(T, \text{abox}(s), \gamma\theta)$  obtained from the current state  $s$  is logically equivalent to the ABox  $\text{abox}(s'')$ , for some already generate state  $s''$ , we do not generate a new state, but simply add  $s \Rightarrow s''$  to  $\Upsilon$ .

## 13.1 Weakly Acyclic KABs

Verification of KABs is undecidable in general. Indeed, we have:

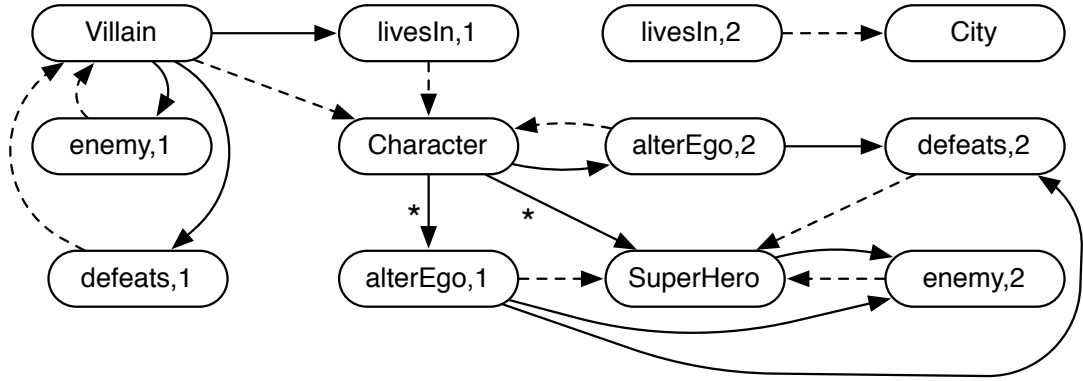
**Theorem 13.2.** *Verification of CTL reachability-like formulas of the form  $\mu Z.(N(a) \vee \langle - \rangle Z)$  (with  $N$  an atomic concept and  $a$  an individual occurring in  $A_0$ ) on KABs with empty TBoxes and actions that make use only of UCQs (no negation nor equality) is undecidable.*

*Proof.* The proof is by reduction to answering boolean UCQs in a relational database under a set of tuple-generating dependencies (TGDs), which is undecidable [36].

Given a boolean UCQ  $q$ , we consider the KAB  $\mathcal{K}_q = (\emptyset, A_0, \{\gamma\}, \{\text{true} \mapsto \gamma\})$ , where  $\text{Flag}(\text{flag}) \in A_0$  is a concept not mentioned in  $q$ ,  $\gamma$  has no parameter and contains:

- a set of effects of the form  $[q_i^+] \rightsquigarrow A_i$ , where  $A_i$  does not contain any equality assertion nor the  $\text{Flag}$  concept;
- a *CopyAll* effect, used to copy all concept and role assertions;
- a further additional effect specification of the form  $[q \wedge \text{Flag}(\text{flag})] \rightsquigarrow \text{Up}(\text{flag})$ , where  $\text{Up}(\text{flag})$  denotes a concept assertion over the known individual  $\text{flag}$ .

<sup>15</sup>Note that all terms that are present in the current ABox are preserved in the new ABox, together with equalities between terms.



**Figure 13** – Weakly acyclic dependency graph for Example 11

We now observe that  $\gamma$  can be seen as a set of tuple generating dependencies (TGDs) [43]. Indeed, it suffices to consider one TGD for each disjunct in the UCQ on the left-hand side of each effect. The problem of answering a boolean UCQ  $q$  in a relational database under a set of TGDs, which is undecidable [12], can therefore be reduced to the verification problem  $\Upsilon \models \mu Z.(\text{Up}(flag) \vee \langle - \rangle Z)$ .  $\square$

Hence it is of interest to isolate interesting special cases in which verification is decidable. Next, we introduce a notable class of KABs that enjoys such a property. In particular, we show that a suitable syntactic restriction, which resembles the notion of *weak acyclicity* in data exchange [36]<sup>16</sup>, guarantees boundedness of ABoxes generated by the execution of the KAB, which in turn guarantees decidability of verification.

To do so, we introduce the edge-labeled directed *dependency graph* of a KAB  $\mathcal{K} = (T, A_0, \Gamma, \Pi)$ , defined as follows. *Nodes*, called *positions*, are obtained from the TBox  $T$ . In particular, there is a node for every concept name  $N$  in  $T$ , and two nodes for every role name  $P$  in  $T$ , the first corresponding to the domain and the second to the range of  $P$ . *Edges* are drawn by considering every effect specification  $[q^+] \wedge Q^- \rightsquigarrow A'$  of each action contained in  $\Gamma$ , tracing how values are copied or contribute to generate new values as the system progresses. In particular, let  $p$  be a position corresponding to a concept/role component in the rewriting  $rew(q^+)$  of  $q^+$  with variable  $x$ . For every position  $p'$  in  $A'$  with the same variable  $x$ , we include a *normal edge*  $p \rightarrow p'$ . For every position  $p''$  in  $A'$  with a Skolem term  $f(\vec{t})$  such that  $x \in \vec{t}$ , we include a *special edge*  $p \xrightarrow{*} p''$ . We say that  $\mathcal{K}$  is *weakly-acyclic* if its dependency graph has no cycle going through a special edge.

**Theorem 13.3.** *Verification of  $\mu\mathcal{L}_A$  properties for weakly acyclic KABs is decidable.*

*Proof.* We can define a simplified KAB, such that the size of its ABoxes bounds the size of those of the original KAB, and then relate the execution of such simplified KAB to the chase of a set of TGDs. If the original KAB is weakly acyclic, so is such a set of TGDs, hence we can apply the bound in [36]. Since all ABoxes are bounded, this implies that we can generate a *finite-state* transition system which is bisimilar to  $\Upsilon$ , and do verification there.  $\square$

**Example 13.2.** The KAB of Example 11 is indeed weakly acyclic. Its dependency graph, shown in Figure 13, does not contain any cycle going through special edges. For readability, self-loops are not shown in the Figure (but are present for all nodes), and dashed edges are used to compactly represent the contributions given by the rewriting of the queries. E.g., the dashed edge from Villain to Character denotes that for every outgoing edge from Character, there exists an outgoing edge from Villain with the same type and target. Hence, w.r.t. weak acyclicity dashed edges can be simply replaced by normal edges.  $\blacksquare$

<sup>16</sup>We use the original definition of weak acyclicity. However, our results depend only on the ability of finding a finite bound for the chase. So, other variants of weak acyclicity, such as [52, 53], can also be adopted.

## 14 Discussion

In this part of the report, we have studied verification of knowledge and action bases, which are dynamic systems constituted by a knowledge base expressed in description logic, and by an action specification that changes the knowledge base over time. We obtain an interesting decidability result by relying on the notion of weak acyclicity, based on a connection with the theory of chase of TGDs in relational databases. Our technique for decidability also gives hints on how to get complexity results. However, a complexity characterization requires further study.

We observe that the restriction imposed by weak acyclicity (or variants) is not too severe, and that in many real cases KABs are indeed weakly acyclic or can be transformed into weakly acyclic ones at cost of redesign. Indeed, if a KAB is not weakly acyclic, it indefinitely generates new values from the old ones, which then depend on a chain of unboundedly many previous values. In other words, current values depend on an unbounded number of old values that are arbitrarily far in the past. If this is not the case, then the KAB can in principle be rewritten into a weakly acyclic one. While such unbounded systems exist in theory, e.g., Turing machines, higher level processes, as those used in business process management or service-oriented modeling, do not typically require such an unboundedness in practice. How to systematically transform systems into weakly acyclic ones remains an open issue.

We close by observing that while we fully used the knowledge base for querying the KAB state, we use it in a limited way when *updating* the state, namely only to guarantee consistency. Description logic knowledge base update has its own semantical and computational difficulties [50, 28, 45], which the approach presented here sidesteps. However, we tackle this issue in the ACSI Deliverable D2.5.

## Part V

# References

- [1] S. Abiteboul, P. Bourhis, A. Galland, and B. Marinoiu. The AXML artifact model. In *Proc. of the 16th Int. Symp. on Temporal Representation and Reasoning (TIME 2009)*, pages 11–17, 2009.
- [2] S. Abiteboul, R. Hull, and V. Vianu. *Foundations of Databases*. Addison Wesley, 1995.
- [3] A. Artale, D. Calvanese, R. Kontchakov, and M. Zakharyashev. The *DL-Lite* family and relations. *J. of Artificial Intelligence Research*, 36:1–69, 2009.
- [4] F. Baader, D. Calvanese, D. McGuinness, D. Nardi, and P. F. Patel-Schneider, editors. *The Description Logic Handbook: Theory, Implementation and Applications*. Cambridge University Press, 2003.
- [5] F. Baader, S. Ghilardi, and C. Lutz. LTL over description logic axioms. *ACM Trans. on Computational Logic*, 13(3), 2012.
- [6] B. Bagheri-Hariri, D. Calvanese, G. De Giacomo, and R. De Masellis. Verification of conjunctive-query based semantic artifacts. In *Proc. of the 24th Int. Workshop on Description Logic (DL 2011)*, volume 745 of *CEUR Electronic Workshop Proceedings*, <http://ceur-ws.org/>, 2011.
- [7] B. Bagheri-Hariri, D. Calvanese, G. De Giacomo, R. De Masellis, and P. Felli. Foundations of relational artifacts verification. In *Proc. of the 9th Int. Conference on Business Process Management (BPM 2011)*, volume 6896 of *Lecture Notes in Computer Science*, pages 379–395. Springer, 2011.
- [8] B. Bagheri Hariri, D. Calvanese, G. De Giacomo, R. De Masellis, and M. Montali. Verification of OWL 2 QL based semantic artifacts. Submitted for publication, 2012.
- [9] B. Bagheri Hariri, D. Calvanese, G. De Giacomo, A. Deutsch, and M. Montali. Verification of relational data-centric dynamic systems with external services. CoRR Technical Report arXiv:1203.0024, arXiv.org e-Print archive, 2012. Available at <http://arxiv.org/abs/1203.0024>. Submitted for publication.
- [10] B. Bagheri Hariri, D. Calvanese, M. Montali, G. De Giacomo, R. De Masellis, and P. Felli. Verification of description logic Knowledge and Action Bases. In *Proc. of the 20th Eur. Conf. on Artificial Intelligence (ECAI 2012)*, 2012. To appear.
- [11] C. Baier and J.-P. Katoen. *Principles of Model Checking*. MIT Press, 2008.
- [12] C. Beeri and M. Y. Vardi. The implication problem for data dependencies. In *Proc. of the 8th Coll. on Automata, Languages and Programming (ICALP'81)*, volume 115 of *Lecture Notes in Computer Science*, pages 73–85. Springer, 1981.
- [13] F. Belardinelli, A. Lomuscio, and F. Patrizi. Verification of deployed artifact systems via data abstraction. In *Proc. of the 9th Int. Joint Conf. on Service Oriented Computing (IC-SOC 2011)*, volume 7084 of *Lecture Notes in Computer Science*, pages 142–156. Springer, 2011.
- [14] K. Bhattacharya et al. A model-driven approach to industrializing discovery processes in pharmaceutical research. *IBM Systems Journal*, 44(1), 2005.

- [15] K. Bhattacharya, C. Gerede, R. Hull, R. Liu, and J. Su. Towards formal analysis of artifact-centric business process models. In *Proc. of the 5th Int. Conference on Business Process Management (BPM 2007)*, volume 4714 of *Lecture Notes in Computer Science*, pages 288–234. Springer, 2007.
- [16] M. Bojanczyk, A. Muscholl, T. Schwentick, L. Segoufin, and C. David. Two-variable logic on words with data. In *Proc. of the 21st IEEE Symp. on Logic in Computer Science (LICS 2006)*, pages 7–16, 2006.
- [17] A. Bouajjani, P. Habermehl, Y. Jurski, and Y. Sighireanu. Rewriting systems with data. In *Proc. of the 16th Int. Symp. on Fundamentals of Computation Theory (FCT 2007)*, pages 1–22, 2007.
- [18] A. Bouajjani, P. Habermehl, and R. Mayr. Automatic verification of recursive procedures with one integer parameter. *Theoretical Computer Science*, 295, 2003.
- [19] P. Bouyer. A logical characterization of data languages. *Information Processing Lett.*, 84(2), 2002.
- [20] P. Bouyer, A. Petit, and D. Thérien. An algebraic approach to data languages and timed languages. *Information and Computation*, 182(2), 2003.
- [21] J. Bradfield and C. Stirling. Modal mu-calculi. In *Handbook of Modal Logic*, volume 3, pages 721–756. Elsevier, 2007.
- [22] O. Burkart, D. Caucal, F. Moller, and B. Steffen. Verification of infinite structures. In *Handbook of Process Algebra*. Elsevier Science, 2001.
- [23] D. Calvanese, G. De Giacomo, R. Hull, and J. Su. Artifact-centric workflow dominance. In *Proc. of the 7th Int. Joint Conf. on Service Oriented Computing (ICSOC 2009)*, volume 5900 of *Lecture Notes in Computer Science*, pages 130–143. Springer, 2009.
- [24] D. Calvanese, G. De Giacomo, D. L., M. Lenzerini, A. Poggi, M. Rodríguez-Muro, and R. Rosati. Ontologies and databases: The *DL-Lite* approach. In S. Tessaris and E. Franconi, editors, *Semantic Technologies for Informations Systems – 5th Int. Reasoning Web Summer School (RW 2009)*, volume 5689 of *Lecture Notes in Computer Science*, pages 255–356. Springer, 2009.
- [25] D. Calvanese, G. De Giacomo, D. Lembo, M. Lenzerini, and R. Rosati. EQL-Lite: Effective first-order query processing in description logics. In *Proc. of the 20th Int. Joint Conf. on Artificial Intelligence (IJCAI 2007)*, pages 274–279, 2007.
- [26] D. Calvanese, G. De Giacomo, D. Lembo, M. Lenzerini, and R. Rosati. Tractable reasoning and efficient query answering in description logics: The *DL-Lite* family. *J. of Automated Reasoning*, 39(3):385–429, 2007.
- [27] D. Calvanese, G. De Giacomo, M. Lenzerini, and R. Rosati. Actions and programs over description logic knowledge bases: A functional approach. In G. Lakemeyer and S. A. McIlraith, editors, *Knowing, Reasoning, and Acting: Essays in Honour of Hector Levesque*. College Publications, 2011.
- [28] D. Calvanese, E. Kharlamov, W. Nutt, and D. Zheleznyakov. Evolution of *DL-Lite* knowledge bases. In *Proc. of the 9th Int. Semantic Web Conf. (ISWC 2010)*, volume 6496 of *Lecture Notes in Computer Science*, pages 112–128. Springer, 2010.
- [29] P. Cangialosi, G. De Giacomo, R. De Masellis, and R. Rosati. Conjunctive artifact-centric services. In *Proc. of the 8th Int. Joint Conf. on Service Oriented Computing (ICSOC 2010)*, volume 6470 of *Lecture Notes in Computer Science*, pages 318–333. Springer, 2010.

- [30] E. M. Clarke, O. Grumberg, and D. A. Peled. *Model checking*. The MIT Press, 1999.
- [31] D. Cohn and R. Hull. Business artifacts: A data-centric approach to modeling business operations and processes. *Bull. of the IEEE Computer Society Technical Committee on Data Engineering*, 32(3):3–9, 2009.
- [32] E. Damaggio, A. Deutsch, and V. Vianu. Artifact systems with data dependencies and arithmetic. In *Proc. of the 14th Int. Conf. on Database Theory (ICDT 2011)*, pages 66–77, 2011.
- [33] S. Demri and R. Lazić. LTL with the freeze quantifier and register automata. *ACM Trans. on Computational Logic*, 10(3):1–30, 2009.
- [34] A. Deutsch, R. Hull, F. Patrizi, and V. Vianu. Automatic verification of data-centric business processes. In *Proc. of the 12th Int. Conf. on Database Theory (ICDT 2009)*, pages 252–267, 2009.
- [35] E. A. Emerson. Model checking and the mu-calculus. In *Descriptive Complexity and Finite Models*, 1996.
- [36] R. Fagin, P. G. Kolaitis, R. J. Miller, and L. Popa. Data exchange: Semantics and query answering. *Theoretical Computer Science*, 336(1):89–124, 2005.
- [37] C. Gerede and J. Su. Specification and verification of artifact behaviors in business process models. In *Proc. of the 5th Int. Conf. on Service Oriented Computing (ICSOC 2007)*, 2007.
- [38] C. E. Gerede, K. Bhattacharya, and J. Su. Static analysis of business artifact-centric operational models. In *IEEE Int. Conf. on Service-Oriented Computing and Applications*, 2007.
- [39] D. G. Giuseppe, Y. Lesperance, and F. Patrizi. Bounded situation calculus action theories and decidable verification. In *Proc. of the 13th Int. Conf. on the Principles of Knowledge Representation and Reasoning (KR 2012)*, 2012.
- [40] R. Hull. Artifact-centric business process models: Brief survey of research results and challenges. In *Proc. of the 7th Int. Conf. on Ontologies, DataBases, and Applications of Semantics (ODBASE 2008)*, volume 5332 of *Lecture Notes in Computer Science*, pages 1152–1163. Springer, 2008.
- [41] M. Jurdzinski and R. Lazić. Alternation-free modal mu-calculus for data trees. In *Proc. of the 22nd IEEE Symp. on Logic in Computer Science (LICS 2007)*, 2007.
- [42] H. Katsuno and A. Mendelzon. On the difference between updating a knowledge base and revising it. In *Proc. of the 2nd Int. Conf. on the Principles of Knowledge Representation and Reasoning (KR'91)*, pages 387–394, 1991.
- [43] P. G. Kolaitis. Schema mappings, data exchange, and metadata management. In *Proc. of the 24rd ACM SIGACT SIGMOD SIGART Symp. on Principles of Database Systems (PODS 2005)*, pages 61–75, 2005.
- [44] R. Lazić, T. Newcomb, J. Ouaknine, A. Roscoe, and J. Worrell. Nets with tokens which carry data. In *ICATPN*, 2007.
- [45] M. Lenzerini and D. F. Savo. On the evolution of the instance level of *DL-Lite* knowledge bases. In *Proc. of the 24th Int. Workshop on Description Logic (DL 2011)*, volume 745 of *CEUR Electronic Workshop Proceedings*, <http://ceur-ws.org/>, 2011.
- [46] H. J. Levesque. Foundations of a functional approach to knowledge representation. *Artificial Intelligence*, 23:155–212, 1984.

- [47] L. Limonad, P. De Leenheer, M. Linehan, R. Hull, and R. Vaculin. Ontology of dynamic entities. In *Proc. of the 31st Int. Conf. on Conceptual Modeling (ER 2012)*, 2012. To appear.
- [48] F. Lin and R. Reiter. State constraints revisited. *J. of Logic Programming*, 4(5):655–678, 1994.
- [49] H. Liu, C. Lutz, M. Milicic, and F. Wolter. Reasoning about actions using description logics with general TBoxes. In *Proc. of the 10th Eur. Conference on Logics in Artificial Intelligence (JELIA 2006)*, volume 4160 of *Lecture Notes in Computer Science*. Springer, 2006.
- [50] H. Liu, C. Lutz, M. Milicic, and F. Wolter. Updating description logic ABoxes. In *Proc. of the 10th Int. Conf. on the Principles of Knowledge Representation and Reasoning (KR 2006)*, pages 46–56, 2006.
- [51] D. C. Luckham, D. M. R. Park, and M. Paterson. On formalised computer programs. *J. of Computer and System Sciences*, 4(3), 1970.
- [52] B. Marnette and F. Geerts. Static analysis of schema-mappings ensuring oblivious termination. In *Proc. of the 13th Int. Conf. on Database Theory (ICDT 2010)*, pages 183–195, 2010.
- [53] M. Meier, M. Schmidt, F. Wei, and G. Lausen. Semantic query optimization in the presence of types. In *Proc. of the 19th ACM SIGACT SIGMOD SIGART Symp. on Principles of Database Systems (PODS 2000)*, pages 111–122, 2010.
- [54] R. Milner. An algebraic definition of simulation between programs. In *Proc. of the 2nd Int. Joint Conf. on Artificial Intelligence (IJCAI’71)*, pages 481–489, 1971.
- [55] F. Neven, T. Schwentick, and V. Vianu. Finite state machines for strings over infinite alphabets. *ACM Trans. on Computational Logic*, 5(3), 2004.
- [56] A. Nigam and N. S. Caswell. Business artifacts: An approach to operational specification. *IBM Systems Journal*, 42(3):428–445, 2003.
- [57] D. M. R. Park. Finiteness is Mu-ineffable. *Theoretical Computer Science*, 3(2):173–181, 1976.
- [58] A. Pnueli. The temporal logic of programs. In *Proc. of the 18th Annual Symp. on the Foundations of Computer Science (FOCS’77)*, pages 46–57, 1977.
- [59] A. Poggi, D. Lembo, D. Calvanese, G. De Giacomo, M. Lenzerini, and R. Rosati. Linking data to ontologies. *J. on Data Semantics*, X:133–173, 2008.
- [60] R. Reiter. *Knowledge in Action: Logical Foundations for Specifying and Implementing Dynamical Systems*. The MIT Press, 2001.
- [61] C. Stirling. *Modal and Temporal Properties of Processes*. Springer, 2001.
- [62] W. M. P. van der Aalst, P. Barthelmess, C. A. Ellis, and J. Wainer. Proclets: A framework for lightweight interacting workflow processes. *Int. J. of Cooperative Information Systems*, 10(4), 2001.
- [63] H. van Ditmarsch, W. van der Hoek, and B. Kooi. *Dynamic epistemic logic*. Springer, 2007.
- [64] V. Vianu. Automatic verification of database-driven systems: a new frontier. In *Proc. of the 12th Int. Conf. on Database Theory (ICDT 2009)*, pages 1–13, 2009.



- [65] F. Wolter and M. Zakharyashev. Temporalizing description logic. In D. Gabbay and M. de Rijke, editors, *Frontiers of Combining Systems*, pages 379–402. Studies Press/Wiley, 1999.
- [66] X. Zhao, J. Su, H. Yang, and Z. Qiu. Enforcing constraints on life cycles of business artifacts. In *TASE*, 2009.

## Part VI

## Appendix

## A Proofs for DCDSs

A.1 Bisimulation Invariance for  $\mu$ -calculi

We give here the bisimulation invariance results for  $\mu\mathcal{L}_A$  and  $\mu\mathcal{L}_P$ .

## History Preserving Mu-Calculus

We prove *history preserving bisimulation invariance* for  $\mu\mathcal{L}_A$ . We adopt a two-step approach. We first prove the result for the logic  $\mathcal{L}_A$ , obtained from  $\mu\mathcal{L}_A$  dropping the predicate variables and the fixpoint constructs. Such a logic corresponds to a first-order variant of the Hennessy Milner logic; note that the semantics of this logic is completely independent from the second-order valuation. We then extend the result to the whole  $\mu\mathcal{L}_A$  by dealing with fixpoints.

**Lemma A.1.** *Consider two transition systems  $\Upsilon_1 = \langle \Delta_1, \mathcal{R}, \Sigma_1, s_{01}, db_1, \Rightarrow_1 \rangle$  and  $\Upsilon_2 = \langle \Delta_2, \mathcal{R}, \Sigma_2, s_{02}, db_2, \Rightarrow_2 \rangle$ , a partial bijection  $h$  between  $\Delta_1$  and  $\Delta_2$  and two states  $s_1 \in \Sigma_1$  and  $s_2 \in \Sigma_2$  such that  $s_1 \approx_h s_2$ . Then for every (open) formula  $\Phi$  of  $\mathcal{L}_A$ , and every valuations  $v_1$  and  $v_2$  that assign to each of its free variables a value  $d_1 \in \text{ADOM}(db_1(s_1))$  and  $d_2 \in \text{ADOM}(db_2(s_2))$ , such that  $d_2 = h(d_1)$ , we have that*

$$\Upsilon_1, s_1 \models \Phi v_1 \quad \text{if and only if} \quad \Upsilon_2, s_2 \models \Phi v_2.$$

*Proof.* We proceed by induction on the structure of  $\Phi$ .

**Local first-order queries (base case)** Consider  $\Phi = Q$ , where  $Q$  is an (open) FO query. Since  $h$  induces an isomorphism between  $db(s_1)$  and  $db(s_2)$ , for every valuations  $v_1$  and  $v_2$  that assign to each free variable of  $Q$  a value  $d_1 \in \text{ADOM}(db_1(s_1))$  and  $d_2 \in \text{ADOM}(db_2(s_2))$ , such that  $d_2 = h(d_1)$ , we have that  $\text{ans}(Qv_1, db(s_1)) \equiv \text{ans}(Qv_2, db(s_2))$ .

**Negation** By induction hypothesis, for every (open) formula  $\Phi$  and every valuations  $v_1$  and  $v_2$  that assign to each of its free variables a value  $d_1 \in \text{ADOM}(db_1(s_1))$  and  $d_2 \in \text{ADOM}(db_2(s_2))$ , such that  $d_2 = h(d_1)$ , we have that  $\Upsilon_1, s_1 \models \Phi v_1$  if and only if  $\Upsilon_2, s_2 \models \Phi v_2$ . By definition,  $\Upsilon_1, s_1 \models (\neg\Phi)v_1$  if and only if  $\Upsilon_1, s_1 \not\models \Phi v_1$ , and, by induction hypothesis,  $\Upsilon_1, s_1 \not\models \Phi v_1$  if and only if  $\Upsilon_2, s_2 \not\models \Phi v_2$ , which corresponds to  $\Upsilon_2, s_2 \models (\neg\Phi)v_2$ .

**Conjunction** By induction hypothesis, for every (open) formula  $\Phi$  and every valuations  $v_1$  and  $v_2$  that assign to each of its free variables a value  $d_1 \in \text{ADOM}(db_1(s_1))$  and  $d_2 \in \text{ADOM}(db_2(s_2))$ , such that  $d_2 = h(d_1)$ , we have that  $\Upsilon_1, s_1 \models \Phi_i v_1$  if and only if  $\Upsilon_2, s_2 \models \Phi_i v_2$ , with  $i \in \{1, 2\}$ . Hence,  $\Upsilon_1, s_1 \models \Phi_1 v_1$  and  $\Upsilon_1, s_1 \models \Phi_2 v_1$  if and only if  $\Upsilon_2, s_2 \models \Phi_1 v_2$  and  $\Upsilon_2, s_2 \models \Phi_2 v_2$ . By definition, we therefore have  $\Upsilon_1, s_1 \models (\Phi_1 \wedge \Phi_2)v_1$  if and only if  $\Upsilon_2, s_2 \models (\Phi_1 \wedge \Phi_2)v_2$ .

**Modal operator** Consider two states  $s_1 \in \Sigma_1$  and  $s_2 \in \Sigma_2$  such that  $s_1 \approx_h s_2$ . By definition, given a valuation  $v_1$  that assigns to each free variable of  $\Phi$  a value  $d_1 \in \text{ADOM}(db_1(s_1))$ , we have that  $\Upsilon_1, s_1 \models (\langle - \rangle \Phi)v_1$  if there exists a transition  $s_1 \Rightarrow_1 s'_1$  such that  $\Upsilon_1, s'_1 \models \Phi v_1$ . Since  $s_1 \approx_h s_2$ , there exists a transition  $s_2 \Rightarrow_2 s'_2$  such that  $s'_1 \approx_{h'} s'_2$ , where  $h'$  extends  $h$ . By induction hypothesis, for every valuation  $v_2$  that assigns to each free variable  $x$  of  $\Phi$  a value  $d_2 \in \text{ADOM}(db_2(s'_2))$ , such that  $d_2 = h'(d_1)$  with  $x/d_1 \in v_1$ , we have that  $\Upsilon_1, s'_1 \models \Phi v_1$  if and only if  $\Upsilon_2, s'_2 \models \Phi v_2$ . Since  $h'$  is an extension of  $h$ , and  $v_1$  assigns to each free variable of  $\Phi$  a value  $d_1 \in \text{ADOM}(db_1(s_1)) \subseteq \text{DOM}(h)$ , we observe that for every pair of assignments  $x/d_1 \in v_1$  and  $x/d_2 \in v_2$ , it holds that  $d_2 = h'(d_1) = h(d_1)$ . Furthermore, since

$h$  induces an isomorphism between  $db_1(s_1)$  and  $db_2(s_2)$ , for each assignment  $x/d_2 \in v_2$ , we have that  $d_2 \in \text{ADOM}(db_2(s_2))$ . Considering that  $s_2 \Rightarrow_2 s'_2$ , by definition we therefore get  $\Upsilon_2, s_2 \models (\langle - \rangle \Phi)v_2$ .

The other direction can be proven in a symmetric way.

**Quantification** Consider two states  $s_1 \in \Sigma_1$  and  $s_2 \in \Sigma_2$  such that  $s_1 \approx_h s_2$ . By definition, given a formula  $\Phi$  and a valuation  $v'_1$  that assigns to each free variable of  $\Phi$  a value  $d_1 \in \text{DOM}(h)$ , we have that  $\Upsilon_1, s_1 \models (\exists x.\text{LIVE}(x) \wedge \Phi)v'_1$  if and only if there exists  $d \in \text{ADOM}(db_1(s_1))$  such that  $\Upsilon_1, s_1 \models \Phi v_1$ , where  $v_1 = v'_1[x/d]$ . By induction hypothesis, for every valuation  $v_2$  that assigns to each free variable  $y$  of  $\Phi$  a value  $d_2 \in \text{ADOM}(db_2(s_2))$ , such that  $d_2 = h(d_1)$  with  $y/d_1 \in v_1$ , we have that  $\Upsilon_1, s_1 \models \Phi v_1$  if and only if  $\Upsilon_2, s_2 \models \Phi v_2$ . More specifically, the structure of  $v_2$  is  $v_2 = v'_2[x/d']$ , where  $d' = h(d) \in \text{ADOM}(db_2(s_2))$  because  $h$  induces an isomorphism between  $db_1(s_1)$  and  $db_2(s_2)$ . Hence, we get  $\Upsilon_2, s_2 \models (\exists x.\text{LIVE}(x) \wedge \Phi)v'_2$ .

The other direction can be proven in a symmetric way.  $\square$

*Proof of Theorem 6.1.* We prove the theorem in two steps. First, we show that Lemma A.1 can be extended to the infinitary version of  $\mathcal{L}_A$  that supports arbitrary countable disjunction. Then, we recall that fixpoints can be translated into this infinitary logic, thus guaranteeing invariance for the whole  $\mu\mathcal{L}_A$  logic.

Let  $\Psi$  be a countable ordered set of open  $\mathcal{L}_A$  formulae. Given a transition system  $\Upsilon = \langle \Delta, \mathcal{R}, \Sigma, s_0, db, \Rightarrow \rangle$ , the semantics of  $\bigvee \Psi$  is  $(\bigvee \Psi)_v^\Upsilon = \bigcup_{\psi \in \Psi} (\psi)_v^\Upsilon$ . Therefore, given a state  $s$  of  $\Upsilon$  and a variable valuation  $v$  that assigns to each free variable of  $\Psi$  a value  $d \in \text{ADOM}(db(s))$ , we have  $\Upsilon, s \models \bigvee \Psi v$  if and only if  $\Upsilon, s \models \psi v$  for some  $\psi \in \Psi$ . Arbitrary countable conjunction is obtained for free because of negation.

We show that the invariance result proven in Lemma A.1 trivially extends to this arbitrary countable disjunction. Lemma A.1 guarantees that invariance is preserved for any finite disjunction. Formally, let  $\{\Phi_1, \dots, \Phi_n\}$  be a finite set of open  $\mathcal{L}_A$  formulae. Consider two states  $s_1 \in \Sigma_1$  and  $s_2 \in \Sigma_2$  such that  $s_1 \approx_h s_2$ . Then, for every valuations  $v_1$  and  $v_2$  that assign to each free variable of  $\{\Phi_1, \dots, \Phi_n\}$  a value  $d_1 \in \text{ADOM}(db_1(s_1))$  and  $d_2 \in \text{ADOM}(db_2(s_2))$ , such that  $d_2 = h(d_1)$ , we have that  $\Upsilon_1, s_1 \models (\bigvee_{i \in \{1, \dots, n\}} \Phi_i)v_1$  if and only if  $\Upsilon_2, s_2 \models (\bigvee_{i \in \{1, \dots, n\}} \Phi_i)v_2$ .

Now consider two valuations  $v'_1$  and  $v'_2$  that assign to each free variable of  $\bigvee \Psi$  a value  $d_1 \in \text{ADOM}(db_1(s_1))$  and  $d_2 \in \text{ADOM}(db_2(s_2))$ , such that  $d_2 = h(d_1)$ . By definition,  $\Upsilon_1, s_1 \models (\bigvee \Psi)v'_1$  if and only if there exists  $\psi_k \in \Psi$  such that  $\Upsilon_1, s_1 \models \psi_k v'_1$ . The proof of invariance for the infinitary  $\mathcal{L}_A$  logic is then obtained by observing that  $\Upsilon_1, s_1 \models (\bigvee \Psi)v'_1$  if and only if  $\Upsilon_1, s_1 \models (\bigvee_{i \in \{1, \dots, k\}} \psi_i)v'_1$  if and only if  $\Upsilon_2, s_2 \models (\bigvee_{i \in \{1, \dots, k\}} \psi_i)v'_2$  if and only if  $\Upsilon_2, s_2 \models (\bigvee \Psi)v'_2$ .

In order to extend the result to the whole  $\mu\mathcal{L}_A$ , we resort to the well-known result stating that fixpoints of the  $\mu$ -calculus can be translated into the infinitary Hennessy Milner logic by iterating over *approximants*, where the approximant of index  $\alpha$  is denoted by  $\mu^\alpha Z.\Phi$  ( $\nu^\alpha Z.\Phi$ ). This is a standard result that also holds for  $\mu\mathcal{L}_A$ . In particular, approximants are built as follows:

$$\begin{array}{ll} \mu^0 Z.\Phi = \text{false} & \nu^0 Z.\Phi = \text{true} \\ \mu^{\beta+1} Z.\Phi = \Phi[Z/\mu^\beta Z.\Phi] & \nu^{\beta+1} Z.\Phi = \Phi[Z/\nu^\beta Z.\Phi] \\ \mu^\lambda Z.\Phi = \bigvee_{\beta < \lambda} \mu^\beta Z.\Phi & \nu^\lambda Z.\Phi = \bigwedge_{\beta < \lambda} \nu^\beta Z.\Phi \end{array}$$

where  $\lambda$  is a limit ordinal, and where fixpoints and their approximants are connected by the following properties: given a transition system  $\Upsilon$  and a state  $s$  of  $\Upsilon$

- $s \in (\mu Z.\Phi)_{v,V}^\Upsilon$  if and only if there exists an ordinal  $\alpha$  such that  $s \in (\mu^\alpha Z.\Phi)_{v,V}^\Upsilon$  and, for every  $\beta < \alpha$ , it holds that  $s \notin (\mu^\beta Z.\Phi)_{v,V}^\Upsilon$ ;
- $s \notin (\nu Z.\Phi)_{v,V}^\Upsilon$  if and only if there exists an ordinal  $\alpha$  such that  $s \notin (\nu^\alpha Z.\Phi)_{v,V}^\Upsilon$  and, for every  $\beta < \alpha$ , it holds that  $s \in (\nu^\beta Z.\Phi)_{v,V}^\Upsilon$ .  $\square$

## Persistence Preserving Mu-Calculus

We prove *persistence preserving bisimulation invariance* for  $\mu\mathcal{L}_P$ . To prove the invariance result, we adopt a two-step approach. We first prove the result for the logic  $\mathcal{L}_P$ , obtained from  $\mu\mathcal{L}_P$  dropping the predicate variables and the fixpoint constructs. Such a logic corresponds to a first-order variant of the Hennessy Milner logic; note that the semantics of this logic is completely independent from the second-order valuation. We then extend the result to the whole  $\mu\mathcal{L}_P$  by dealing with fixpoints.

**Lemma A.2.** *Consider two transition systems  $\Upsilon_1 = \langle \Delta_1, \mathcal{R}, \Sigma_1, s_{01}, db_1, \Rightarrow_1 \rangle$  and  $\Upsilon_2 = \langle \Delta_2, \mathcal{R}, \Sigma_2, s_{02}, db_2, \Rightarrow_2 \rangle$ , a partial bijection  $h$  between  $\Delta_1$  and  $\Delta_2$  and two states  $s_1 \in \Sigma_1$  and  $s_2 \in \Sigma_2$  such that  $s_1 \sim_h s_2$ . Then for every (open) formula  $\Phi$  of  $\mathcal{L}_P$ , and every valuations  $v_1$  and  $v_2$  that assign to each of its free variables a value  $d_1 \in \text{ADOM}(db_1(s_1))$  and  $d_2 \in \text{ADOM}(db_2(s_2))$ , such that  $d_2 = h(d_1)$ , we have that*

$$\Upsilon_1, s_1 \models \Phi v_1 \text{ if and only if } \Upsilon_2, s_2 \models \Phi v_2.$$

*Proof.* We proceed by induction on the structure of  $\Phi$ . In particular, we discuss the two base cases of  $\langle - \rangle(\text{LIVE}(x) \wedge \Phi)$  and  $[-](\text{LIVE}(x) \wedge \Phi')$  with one variable. For convenience, we rewrite the latter case to  $\langle - \rangle(\text{LIVE}(x) \rightarrow \Phi)$ , where  $\Phi = -\Phi'$ . The other cases are derived, or proven in the same way as done for Lemma A.1.

**Modal operator (conjunction)** Consider two states  $s_1 \in \Sigma_1$  and  $s_2 \in \Sigma_2$  such that  $s_1 \sim_h s_2$ .

Let  $x$  be the only free variable of  $\Phi$ , and  $x/d$  a valuation such that  $d \in \text{ADOM}(db_1(s_1))$ . Then, by definition we have that  $\Upsilon_1, s_1 \models (\langle - \rangle(\text{LIVE}(x) \wedge \Phi)[x/d])$  if there exists a transition  $s_1 \Rightarrow_1 s'_1$  such that  $d \in \text{ADOM}(db_1(s'_1))$  and  $\Upsilon_1, s'_1 \models \Phi[x/d]$ . Since  $s_1 \sim_h s_2$ , there exists a transition  $s_2 \Rightarrow_2 s'_2$  such that  $s'_1 \sim_{h'} s'_2$ , where  $h'$  is compatible with  $h$ . By induction hypothesis and by considering that  $h'$  is an isomorphism between  $db_1(s'_1)$  and  $db_2(s'_2)$ , we have that  $\Upsilon_1, s'_1 \models \Phi[x/d]$  if and only if  $h'(d) \in \text{ADOM}(db_2(s'_2))$  and  $\Upsilon_2, s'_2 \models \Phi[x/h'(d)]$ . Now we observe that  $d \in \text{ADOM}(db_1(s_1)) \cap \text{ADOM}(db_1(s'_1))$  and  $h'$  is an extension of  $h|_{\text{ADOM}(db_1(s_1)) \cap \text{ADOM}(db_1(s'_1))}$ . This implies that  $h'(d) = h(d) \in \text{ADOM}(db_2(s_2))$ , because  $h$  is an isomorphism between  $db_1(s_1)$  and  $db_2(s_2)$ . Considering that  $s_2 \Rightarrow_2 s'_2$ , by definition we therefore get  $\Upsilon_2, s_2 \models (\langle - \rangle(\text{LIVE}(x) \wedge \Phi))[x/h(d)]$ .

The other direction can be proven in a symmetric way.

**Modal operator (implication)** Consider two states  $s_1 \in \Sigma_1$  and  $s_2 \in \Sigma_2$  such that  $s_1 \sim_h s_2$ .

Let  $x$  be the only free variable of  $\Phi$ , and  $x/d$  a valuation such that  $d \in \text{ADOM}(db_1(s_1))$ . Then, by definition we have that  $\Upsilon_1, s_1 \models (\langle - \rangle(\text{LIVE}(x) \rightarrow \Phi)[x/d])$  if there exists a transition  $s_1 \Rightarrow_1 s'_1$  such that  $d \notin \text{ADOM}(db_1(s'_1))$  or  $\Upsilon_1, s'_1 \models \Phi[x/d]$ . Since  $s_1 \sim_h s_2$ , there exists a transition  $s_2 \Rightarrow_2 s'_2$  such that  $s'_1 \sim_{h'} s'_2$ , where  $h'$  is an extension of  $h|_{\text{ADOM}(db_1(s_1)) \cap \text{ADOM}(db_1(s'_1))}$ . Now we discuss the two cases in which  $d \notin \text{ADOM}(db_1(s'_1))$  and  $d \in \text{ADOM}(db_1(s'_1))$ .

- Assume that  $d \notin \text{ADOM}(db_1(s'_1))$ . Since  $s_1 \sim_h s_2$ , we have that  $h(d) \in \text{ADOM}(db_2(s_2))$ . Now, towards contradiction, let us assume that  $h(d) \in \text{ADOM}(db_2(s'_2))$ . Hence, we have  $h(d) \in \text{ADOM}(db_2(s_2)) \cap \text{ADOM}(db_2(s'_2))$ . Observe that  $h'$  is an extension of  $h|_{\text{ADOM}(db_1(s_1)) \cap \text{ADOM}(db_1(s'_1))}$ , which is equivalent to state that  $h'^{-1}$  is an extension of  $h^{-1}|_{\text{ADOM}(db_2(s_2)) \cap \text{ADOM}(db_2(s'_2))}$ . This implies that  $h^{-1}(d) = h'^{-1}(d) = d$ . Since  $h'$  is an isomorphism between  $db_1(s'_1)$  and  $db_2(s'_2)$ , then  $d \in \text{ADOM}(db_1(s'_1))$ , and this contradicts the hypothesis.
- Assume that  $d \in \text{ADOM}(db_1(s'_1))$ . Then we can proceed following the line of reasoning used for the case of  $\langle - \rangle(\text{LIVE}(x) \wedge \Phi)$ .

The other direction can be proven in a symmetric way. □

*Proof of Theorem 6.2.* The proof is analogous to that of Theorem 6.1, but now using Lemma A.2. □

## A.2 Deterministic Services

### Run-Bounded Systems

*Proof of Theorem 7.1.* The proof is by reduction from the halting problem. Given a deterministic Turing Machine  $TM$ , we define DCDS  $\mathcal{S}$  with deterministic services and propositional safety property  $\Phi$ , such that  $TM$  halts if and only if  $\Upsilon_{\mathcal{S}} \models \Phi$ .

Intuitively, every run of  $\Upsilon_{\mathcal{S}}$  simulates a run of  $TM$ . Each state  $s$  of  $\Upsilon_{\mathcal{S}}$  models a configuration of  $TM$ . A transition in  $\Upsilon_{\mathcal{S}}$  models a transition in  $TM$ . We give the construction next.

**The DCDS.** To model a configuration of  $TM$  in a relation of the DCDS state, we model the visited tape segment as a graph whose nodes are cell identifiers, and whose edges form a linear path. The edge relation is called *right*, with the intended meaning that  $right(x, y)$  declares cell  $y$  to be the right neighbor of cell  $x$  on the tape. We also introduce a relation *sym*, with  $sym(c, s)$  intended to model that cell  $c$  holds symbol  $s$ . Unary relation *head* models the head position:  $head(c)$  means that the head points to cell  $c$ . Finally, unary relation *state* keeps the state of  $TM$ , and a boolean predicate *halted* is meant to detect that  $TM$  has halted. In summary, the data layer  $\mathcal{D} = \langle \mathcal{C}, \mathcal{R}, \mathcal{E}, \mathcal{I}_0 \rangle$  of  $\mathcal{S}$  contains the schema  $\mathcal{R} = \{right/2, sym/2, head/1, halted/0\}$ . We detail  $\mathcal{E}$  and  $\mathcal{I}_0$  after sketching the process layer.

There is a single action  $\alpha$ , in charge of simulating the transitions of  $TM$ . It has no parameters, and its guard is always true:  $true \mapsto \alpha$ .  $\alpha$  contains the following effects.

$e_{copy}$  simply copies the part of the tape that stays unchanged in the transition because the head doesn't point to it:

$$\begin{aligned} e_{copy} : & \quad right(X, Y) \wedge right(Y, Z) \wedge sym(X, SX) \wedge sym(Y, SY) \wedge sym(Z, SZ) \wedge \\ & \quad \neg(head(X) \wedge head(Y) \wedge head(Z)) \\ & \quad \rightsquigarrow \{ right(X, Y), right(Y, Z), sym(X, SX), sym(Y, SY), sym(Z, SZ) \} \end{aligned}$$

In addition, we add effects for each entry of  $TM$ 's transition relation  $\delta$ .

For instance, if  $(p, b, \rightarrow) \in \delta(s, a)$  (i.e.,  $\delta$  prescribes that in state  $s$ , if the head points to a cell containing symbol  $a$ ,  $TM$  changes state to  $p$ , the cell's symbol is overwritten with  $b$ , and the head moves to the right), we introduce two effects. One for the case when the tape needs no extension to the right,

$$\begin{aligned} e_{s,a,p,b,\rightarrow}^{noext} : & \quad right(X, Y) \wedge sym(X, a) \wedge sym(Y, SY) \wedge SY \neq \omega \wedge head(X) \wedge state(s) \\ & \quad \rightsquigarrow \{ right(X, Y), sym(X, b), sym(Y, SY), head(Y), state(p) \}, \end{aligned}$$

and one when it does:

$$\begin{aligned} e_{s,a,p,b,\rightarrow}^{ext} : & \quad right(X, Y) \wedge sym(X, a) \wedge sym(Y, \omega) \wedge head(X) \wedge state(s) \\ & \quad \rightsquigarrow \{ right(X, Y), right(Y, newCell(Y)), sym(X, b), sym(Y, \perp), \\ & \quad \quad sym(newCell(Y), \omega), head(Y), state(p) \}. \end{aligned}$$

To distinguish among constants and variables in the above effect specifications, we use capital letters for the latter and lower-case letters for the former. Notice that the extension is performed by calling service *newCell*, which is meant to return a fresh cell id (we show below how to ensure this). Also notice the use of special symbol  $\omega$ , which is reserved for labeling the end of the tape segment. Finally, special symbol  $\perp$  is by convention used to initialize the tape prior to starting the run.

We are not quite done, as we still need to ensure that *right* induces a linear order on the collection of cell identifiers generated during the run. Notice that this cannot be achieved exclusively by declaring FO constraints in *right*, as linear orders are not FO-axiomatizable. The solution must exploit the interplay between constraints on *right* and the way  $\Upsilon_{\mathcal{S}}$  transitions.

Observe that, by definition of the effects that extend *right* (e.g.  $e_{s,a,p,b,\rightarrow}^{ext}$  above), at each step the current right end of the tape segment obtains at most one new successor. However, if the call of service *newCell* returns a cell id that already appears in the tape segment, then there

can be some cell with several predecessors according to *right*. We rule out this case by declaring the second component of *right* to be a key. It follows that *right* must be either (i) a linear path (possibly starting from a source node that has a self-loop), or (ii) it must contain a simple cycle involving more than one cell id. The simple cycle is created at the step when *newCell* returns the id of the leftmost cell.

We wish to force case (i). To rule out case (ii), we proceed as follows: we initialize *right* to contain a source node 0 that can never be a cell id because it cannot be returned by *newCell* without violating the key constraint on *right*. To this end, we initialize  $\mathcal{I}_0$  to

- $right^{\mathcal{I}_0} = \{(0, 0), (0, 1), (1, 2)\}$ ,
- $sym^{\mathcal{I}_0} = \{(1, \$), (2, \omega)\}$ ,
- $head^{\mathcal{I}_0} = \{2\}$ ,
- $state^{\mathcal{I}_0} = \{s_0\}$ ,
- $halted^{\mathcal{I}_0} = \{\}$ ,

where  $s_0$  is the initial state of *TM*.

Notice that, if we disregard cell 0,  $\mathcal{I}_0$  contains the representation of an empty tape (symbol \$ labels the left end, symbol  $\omega$  the right end). Also notice that  $right^{\mathcal{I}_0}$  has type (i). An easy induction shows that every run prefix must also construct a *right* relation of type (i), since any attempt to extend *right* with an edge back to one of its existing nodes violates the key constraint.

Because symbol \$ denotes the left end of the tape, it also follows easily from the behavior of *TM* that during the run, the head will never reach the special cell 0, so *head* can only take values from the suffix of *right* starting at cell 1, which is a true linear path.

Now assume without loss of generality that the *TM* is normalized to enter a particular sink state *h* when it halts. We add effect  $e_h$ , which detects the halting state and sets the boolean predicate *halted*.

Observe that for the cases when the head stays in place or moves left, no tape extension is required, so each such entry in the transition relation corresponds to a single effect.

**The property.** We define the propositional safety property  $\Phi$  as

$$\Phi : \mathbf{G}\text{-halted}.$$

It is easy to see that the runs of  $\Upsilon_{\mathcal{S}}$  correspond one-to-one to the runs of *TM*. Since  $\Phi$  is a linear-time property, this run correspondence suffices to guarantee that  $\Upsilon_{\mathcal{S}} \models \Phi$  if and only if *TM* does not halt.  $\square$

*Proof of Theorem 7.2.* The proof is directly obtained from Theorem 7.4, noticing that model checking of propositional  $\mu$ -calculus formulae over finite transition systems is decidable [35].  $\square$

*Proof of Theorem 7.3.* In view of proving this result, we first introduce a key lemma. We say that a transition system is *adom-inflationary* if the active domain of every state is included in its successor's active domain. We say that a DCDS  $\mathcal{S}$  is *adom-inflationary* if  $\Upsilon_{\mathcal{S}}$  is *adom-inflationary*. We can show that, for *adom-inflationary* transition systems, persistence-preserving bisimilarity coincides with history-preserving bisimilarity.

**Lemma A.3.** *Consider two adom-inflationary DCDSs with non-deterministic services,  $\mathcal{S}_1, \mathcal{S}_2$ . Then  $\Upsilon_{\mathcal{S}_1} \sim \Upsilon_{\mathcal{S}_2}$  if and only if  $\Upsilon_{\mathcal{S}_1} \approx \Upsilon_{\mathcal{S}_2}$ .*

*Proof.* A comparison of the two notions of bisimilarity reveals that the difference is in the local condition, as follows.

Notice first that what both bisimilarity notions have in common is that they mention bisimilar states  $s_1$  and  $s_2$  and witness isomorphism  $h$ , and their successors  $s_1 \Rightarrow s'_1, s_2 \Rightarrow s'_2$  such that  $s'_1$  and  $s'_2$  are bisimilar as witnessed by isomorphism  $h'$ . The key difference lies in how  $h'$  and  $h$

are related. In the history-preserving flavor,  $h'$  must extend  $h$ , while in the persistence-preserving flavor  $h'$  must only extend  $h \upharpoonright_{\text{ADOM}(s_1) \cap \text{ADOM}(s'_1)}$ .

Clearly, history-preserving bisimilarity implies persistence-preserving bisimilarity. However, notice that if the transition systems are adom-inflationary, then the converse also holds. Indeed, assume  $s_1 \sim_h s_2$ . By definition,  $h'$  extends  $h \upharpoonright_{\text{ADOM}(s_1) \cap \text{ADOM}(s'_1)}$ . But because of adom-inflation,  $\text{ADOM}(s_1) \subseteq \text{ADOM}(s'_1)$  and hence  $\text{ADOM}(s_1) \cap \text{ADOM}(s'_1) = \text{ADOM}(s_1)$ , yielding  $h \upharpoonright_{\text{ADOM}(s_1) \cap \text{ADOM}(s'_1)} = h$ . Hence,  $h'$  extends  $h$ , which is the condition for history-preserving bisimilarity.  $\square$

We proceed now with the proof of Theorem 7.3, by exploiting the reduction postulated by Theorem 10.1.

Starting from run-bounded DCDS  $D$  with deterministic services, the reduction gives us state-bounded DCDS  $N$  with non-deterministic services. Moreover, the two transition systems have the same domains, and the projection of  $\Upsilon_N$  on the schema of  $D$  coincides with  $\Upsilon_D$  (Theorem 10.1(ii)). In more detail, denoting the schema of  $D$  with  $\mathcal{R}_D$  and the schema of  $N$  with  $\mathcal{R}_N$ , there is a bijection  $\beta$  between the states of  $\Upsilon_D$  and the states of  $\Upsilon_N$ , such that  $s = \beta(s) \upharpoonright_{\mathcal{R}_D}$ . Clearly, this implies that  $\Upsilon_D$  and  $\Upsilon_N$  satisfy the same  $\mu\mathcal{L}$  formulae.

However, a weaker statement suffices for our purpose. By definition of history-preserving bisimilarity, Theorem 10.1(ii) implies that

$$(1) \quad \Upsilon_D \approx \Upsilon_N.$$

We recall that on the way to proving Theorem 8.3, it is shown in Theorem 8.4 that since  $N$  is state-bounded, we can construct using algorithm RCYCL a finite-state abstract transition system  $F$  such that  $\Upsilon_N \sim \Upsilon_F$  ( $F$  is an eventually recycling pruning of  $\Upsilon_N$ ).

An inspection of the reduction in Theorem 10.1 reveals that  $\Upsilon_N$  is adom-inflationary. But since  $\Upsilon_N \sim \Upsilon_F$ , it follows that  $\Upsilon_F$  is adom-inflationary as well (by the local condition of persistence-preserving bisimilarity). Thus Lemma A.3 applies, yielding

$$(2) \quad \Upsilon_N \approx \Upsilon_F.$$

By (1) and (2), and by transitivity of  $\approx$ , we obtain that  $\Upsilon_D \approx \Upsilon_F$ .  $\square$

*Proof of Theorem 7.4.* Theorem 7.3 implies that, given a DCDS  $\mathcal{S}$ , there exists a finite-state transition system  $\Theta_{\mathcal{S}} = \langle \mathbb{U}, \mathcal{R}, \Sigma_a, s_0^a, db_a, \Rightarrow_a \rangle$  that is history preserving bisimilar to the concrete transition system  $\Upsilon_{\mathcal{S}} = \langle \mathbb{U}, \mathcal{R}, \Sigma, s_0, db, \Rightarrow \rangle$ . Thus, it is possible to use  $\Theta_{\mathcal{S}}$  in place of  $\Upsilon_{\mathcal{S}}$  for verification. In particular, given a  $\mu\mathcal{L}_A$  property  $\Phi$ , the verification problem is reduced to  $\Theta_{\mathcal{S}} \models \Phi$ . Let  $\text{ADOM}(\Theta_{\mathcal{S}}) = \bigcup_{s_i \in \Sigma} \text{ADOM}(db(s_i))$ . If  $\Theta_{\mathcal{S}}$  is finite-state, then there exists a bound  $b$  such that  $|\text{ADOM}(\Theta_{\mathcal{S}})| < b$ . Consequently, it is possible to transform  $\Phi$  into an equivalent *finite* propositional  $\mu$ -calculus formula  $\text{PROP}(\Phi)$  as follows:

$$\begin{aligned} \text{PROP}(Q) &= Q \\ \text{PROP}(\neg\Psi) &= \neg\text{PROP}(\Psi) \\ \text{PROP}(\Psi_1 \wedge \Psi_2) &= \text{PROP}(\Psi_1) \wedge \text{PROP}(\Psi_2) \\ \text{PROP}(\langle - \rangle \Psi) &= \langle - \rangle \text{PROP}(\Psi) \\ \text{PROP}(Z) &= Z \\ \text{PROP}(\mu Z. \Psi) &= \mu Z. \text{PROP}(\Psi) \\ \text{PROP}(\exists x. \text{LIVE}(x) \wedge \Psi(x)) &= \bigvee_{t_i \in \text{ADOM}(\mathcal{S})} \text{LIVE}(t_i) \wedge \text{PROP}(\Psi(t_i)) \end{aligned}$$

Clearly,  $\Theta_{\mathcal{S}} \models \Phi$  if and only if  $\Theta_{\mathcal{S}} \models \text{PROP}(\Phi)$ . The proof is then obtained by observing that verification of  $\mu$ -calculus formulae over finite transition systems is decidable [35].  $\square$

*Proof of Theorem 7.5.* The theorem is proved by exhibiting, for every  $n$ , a  $\mu\mathcal{L}$  property that requires the existence of at least  $n$  objects in the transition system.

Let  $\mathcal{S} = \langle \mathcal{D}, \mathcal{P} \rangle$  be a DCDS with data layer  $\mathcal{D} = \langle \mathcal{C}, \mathcal{R}, \emptyset, \mathcal{I}_0 \rangle$  and process layer  $\mathcal{P} = \langle \mathcal{F}, \mathcal{A}, \varrho \rangle$ , where  $\mathcal{F} = \{f/1\}$ ,  $\mathcal{R} = \{R/1, Q/1\}$ ,  $\mathcal{I}_0 = \{R(a)\}$ ,  $\varrho = \{R(x) \mapsto \alpha(x)\}$  and  $\mathcal{A} = \{\alpha(p)\}$ , where  $\alpha(p) : \{\text{true} \rightsquigarrow \{Q(f(p))\}\}$ . The concrete transition system  $\Upsilon_{\mathcal{S}}$  has the following shape:

- The initial state is  $s_0 = \langle \{R(a)\}, \emptyset \rangle$ ;
- $s_0$  is connected to infinitely many successor states, each one storing into  $Q$  a distinct value  $d$  resulting from the service call  $f(a)$ ; each such state has then the form  $s_d = \langle \{Q(d)\}, \{f(a) \mapsto d\} \rangle$ ;
- each  $s_d$  has no outgoing edge, because there is no applicable action in  $s_d$ .

$\mathcal{S}$  is clearly run-bounded, in particular by a bound  $b = 3$ .

Let us now consider the following  $\mu\mathcal{L}$  property without fixpoints:

$$\Phi_n = \exists x_1, \dots, x_n. \bigwedge_{i \neq j} x_i \neq x_j \wedge \bigwedge_{i \in \{1, \dots, n\}} \langle - \rangle Q(x_i)$$

The property states that there are  $n$  distinct values, each of which is stored into relation  $Q$  in one of the successors of the initial state. It is easy to see that  $\Upsilon_{\mathcal{S}} \models \Phi_n$  for every  $n$ . On the other hand, for every finite state abstraction  $\Theta_{\mathcal{S}}$  with  $k$  successors of the initial state, we have that  $\Theta_{\mathcal{S}} \not\models \Phi_{k+1}$ .  $\square$

## Weakly Acyclic DCDSs

*Proof of Theorem 7.6.* The proof is by reduction from the halting problem. We reuse without change the reduction in the proof of Theorem 7.1. This reduction yields for any Turing Machine  $TM$  a DCDS with deterministic services  $\mathcal{S}$ , such that  $\mathcal{S}$  simulates  $TM$ 's computation. That is, the runs of  $TM$  correspond one-to-one to the runs of  $\Upsilon_{\mathcal{S}}$ . It follows immediately that  $TM$  halts if and only if  $\mathcal{S}$  is run-bounded.  $\square$

*Proof of Lemma 7.7.* Let  $\mathcal{S} = \langle \mathcal{D}, \mathcal{P} \rangle$  be a DCDS with data layer  $\mathcal{D} = \langle \mathcal{C}, \mathcal{R}, \mathcal{E}, \mathcal{I}_0 \rangle$  and process layer  $\mathcal{P} = \langle \mathcal{F}, \mathcal{A}, \varrho \rangle$ . Consider now  $\Upsilon_{\mathcal{S}} = \langle \mathcal{C}, \mathcal{R}, \Sigma, s_0, db, \Rightarrow \rangle$  and  $\Upsilon_{\mathcal{S}^+} = \langle \mathcal{C}, \mathcal{R}, \Sigma^+, s_0, db, \Rightarrow^+ \rangle$ . Since  $\mathcal{S}^+$  is weakly acyclic by hypothesis, to prove that run boundedness of  $\Upsilon_{\mathcal{S}^+}$  implies run boundedness of  $\Upsilon_{\mathcal{S}}$ , we show the following stronger result: for every run  $\tau$  in  $\Upsilon_{\mathcal{S}}$ , there exists a run  $\tau^+$  in  $\Upsilon_{\mathcal{S}^+}$  such that, for all pairs of states  $\tau(i) = \langle \mathcal{I}_i, \mathcal{M}_i \rangle$  and  $\tau^+(i) = \langle \mathcal{I}_i^+, \mathcal{M}_i^+ \rangle$ , we have

1.  $\mathcal{M}_i^+$  extends  $\mathcal{M}_i$ ;
2.  $\mathcal{I}_i \subseteq \mathcal{I}_i^+$ ;
3. for the mappings mentioned in  $\mathcal{M}_i^+$  but not in  $\mathcal{M}_i$ ,  $\mathcal{M}_i^+$  “agrees” with the maps contained in the suffix of  $\tau[i]$ , i.e.,

$$\mathcal{M}_i^+|_{C_i} = \left( \bigcup_{j>i} \mathcal{M}_j \right)|_{C_i}$$

where  $C_i = \text{DOM}(\mathcal{M}_i^+) \cap \bigcup_{j>i} \text{DOM}(\mathcal{M}_j)$ .

We prove this by induction on the length of  $\tau$ :

**(base case)** The initial state of both runs is  $\tau(0) = \tau^+(0) = \langle \mathcal{I}_0, \emptyset \rangle$ , and therefore all the three conditions are trivially satisfied.

**(inductive step)** Consider a pair of corresponding states  $\tau(i)$  and  $\tau^+(i)$ , with  $i > 0$ . By definition,  $\tau(i) \Rightarrow \tau(i+1)$  means that there exists an action  $\alpha \in \mathcal{A}$  and a substitution  $\sigma$  for the parameters of  $\alpha$  such that  $\langle \tau(i), \alpha\sigma, \tau(i+1) \rangle \in \text{EXEC}_{\mathcal{S}}$ . We first observe that  $\alpha^+$  can be executed in  $\tau^+(i)$ , since  $\mathcal{P}^+$  does not impose any restriction on the executability of



actions. Let  $Next^+ = \{s^+ \in \Sigma^+ \mid \langle \tau^+(i), \alpha, s^+ \rangle \in EXEC_{\mathcal{S}^+}\}$  be the set of successor states of  $\tau^+(i)$  that are obtained from the application of  $\alpha^+$ .

We now show that there exists  $\underline{s} \in Next^+$  that satisfies the three claims above. The proof is then obtained by simply imposing  $\tau^+(i+1) = \underline{s}$ .

1. By definition,  $\text{DOM}(\mathcal{M}_{i+1}) = \text{DOM}(\mathcal{M}_i) \cup \text{CALLS}(\text{DO}(\mathcal{I}_i, \alpha\sigma))$ , and, for every  $s_k = \langle \mathcal{M}_k^+, \mathcal{I}_k^+ \rangle \in Next^+$ , we have  $\text{DOM}(\mathcal{M}_k^+) = \text{DOM}(\mathcal{M}_i^+) \cup \text{CALLS}(\text{DO}(\mathcal{I}_i^+, \alpha^+\sigma))$ . Consider each effect specification  $q_j^+ \wedge Q_j^- \rightsquigarrow E_j \in \text{EFFECT}(\alpha)$ . By definition of  $q_j^+$  and  $Q_j^-$ ,  $\theta \in \text{ans}((q_j^+ \wedge Q_j^-)\sigma, \mathcal{I}_i)$  implies  $\theta \in \text{ans}(q_j^+\sigma, \mathcal{I}_i)$ , which in turn implies  $\theta \in \text{ans}(q_j^+\sigma, \mathcal{I}_i^+)$ , because  $\mathcal{I}_i \subseteq \mathcal{I}_i^+$  by induction hypothesis. Consequently, we have  $\text{DO}(\mathcal{I}_i, \alpha\sigma) \subseteq \text{DO}(\mathcal{I}_i^+, \alpha^+\sigma)$ , and hence  $\text{CALLS}(\text{DO}(\mathcal{I}_i, \alpha\sigma)) \subseteq \text{CALLS}(\text{DO}(\mathcal{I}_i^+, \alpha^+\sigma))$ . Since  $\text{DOM}(\mathcal{M}_i) \subseteq \text{DOM}(\mathcal{M}_i^+)$  by induction hypothesis, then we obtain  $\text{DOM}(\mathcal{M}_{i+1}) \subseteq \text{DOM}(\mathcal{M}_k^+)$ . Since  $\mathcal{S}^+$  has no equality constraint, the states in  $Next^+$  cover every possible result obtained by calling the service call in  $\mathcal{M}_k^+ \setminus \mathcal{M}_i^+$ , including those states for which  $\mathcal{M}_k^+$  is an extension of  $\mathcal{M}_{i+1}$ . We use  $\underline{Next}^+$  to denote such states.
2. By definition, for each state  $s_k = \langle \mathcal{M}_k^+, \mathcal{I}_k^+ \rangle \in \underline{Next}^+$ , we have that  $\mathcal{M}_k^+$  extends  $\mathcal{M}_{i+1}$ . Therefore, since  $\text{DO}(\mathcal{I}_i, \alpha\sigma) \subseteq \text{DO}(\mathcal{I}_i^+, \alpha^+\sigma)$ , we have  $\mathcal{I}_{i+1} = \mathcal{M}_{i+1}(\text{DO}(\mathcal{I}_i, \alpha\sigma)) \subseteq \mathcal{I}_k^+ = \mathcal{M}_k^+(\text{DO}(\mathcal{I}_i^+, \alpha^+\sigma))$ .
3. Since  $\mathcal{S}^+$  has no equality constraints, we observe that the states in  $\underline{Next}^+$  cover all possible values for the service calls that are not mentioned in  $\mathcal{M}_{i+1}$ . Therefore, there must exist at least one state  $\underline{s} \in \underline{Next}^+$  that satisfies the third claim. In other words, by imposing  $\tau^+(i+1) = \underline{s}$ , we have

$$\mathcal{M}_{i+1}^+|_{C_{i+1}} = \left( \bigcup_{j>i+1} \mathcal{M}_j \right)|_{C_{i+1}} \quad \square$$

*Proof of Theorem 7.8.* Let  $\mathcal{S} = \langle \mathcal{D}, \mathcal{P} \rangle$  be a DCDS with data layer  $\mathcal{D} = \langle \mathcal{C}, \mathcal{R}, \mathcal{E}, \mathcal{I}_0 \rangle$  and process layer  $\mathcal{P} = \langle \mathcal{F}, \mathcal{A}, \rho \rangle$ . We consider the positive approximate  $\mathcal{S}^+$ , showing that if the dependency graph  $G = \langle N, E \rangle$  of  $\mathcal{S}$  (which corresponds by definition to the one of  $\mathcal{S}^+$ ) is weakly acyclic, then  $\mathcal{S}^+$  is run-bounded. The complete proof is then directly obtain by appealing to Lemma 7.7, which states that if  $\mathcal{S}^+$  is run-bounded, then  $\mathcal{S}^+$  is run-bounded as well.

To prove that weak acyclicity of  $\mathcal{S}$  implies that  $\mathcal{S}^+$  is run-bounded, we exploit the connection with the chase of a set of tuple generating dependencies (TGDs) in data exchange. In particular, we resort to the proof given in [36], Theorem 3.9. For every node  $p \in N$ , we consider an incoming path to be any (finite or infinite) path ending in  $p$ . For simplicity, we say that a value appears in position  $p = \langle R_k, j \rangle \in N$  if it appears in the  $j$ -th component of an  $R_k$  tuple. We define the rank of  $p$ , denoted  $\text{rank}(p)$ , as the maximum number of special edges on any such incoming path. Since  $\mathcal{S}^+$  is weakly acyclic by hypothesis,  $G$  does not contain cycles going through special edges, and therefore  $\text{rank}(p)$  is finite. Let  $r$  be the maximum among  $\text{rank}(p_i)$  over all nodes. We observe that  $r \leq |N|$ ; indeed no path can lead to the same node twice using special edges, otherwise  $G$  would contain a cycle going through special edges, thus breaking the weak acyclicity hypothesis. Notice also that  $|N|$  is a constant value, because it is obtained from  $\mathcal{R}$ , which is fixed. We now partition the nodes in  $N$  according to their rank, obtaining a set of sets  $\{N_0, N_1, \dots, N_r\}$ , where  $N_i$  is the set of all nodes with rank  $i$ . The proof is then a natural consequence of the following claim:

**Claim.** Consider a trace  $\tau$  in  $\Upsilon_{\mathcal{S}^+}$ . For every  $i \in \{1, \dots, r\}$ , the total number of distinct values occurring in the databases of  $\tau$  inside position  $p \in N_i$  is bounded by a polynomial  $P_i(|\text{ADOM}(\mathcal{I}_0)|)$ .

We prove the claim by induction on  $i$ :

**(Base case)** Consider  $p \in N_0$ . By definition,  $p$  has no incoming path containing special edges. Therefore, no new values are stored in  $p$  along the run:  $p$  can just store values that are

part of the initial database  $\mathcal{I}_0$ . This holds for all nodes in  $N_0$ , and hence we can fix  $P_0(|\text{ADOM}(\mathcal{I}_0)|) = |\text{ADOM}(\mathcal{I}_0)|$ .

**(Inductive step)** Consider  $p \in N_i$ , with  $i \in \{1, \dots, r\}$ . The first kind of values that may be stored inside  $p$  are those values that were stored inside the component itself in  $\mathcal{I}_0$ . The number of such values is at most  $|\text{ADOM}(\mathcal{I}_0)|$ . In addition, a value may be stored in  $p$  for two reasons: either it is copied from some other position  $p' \in N_j$  with  $i \neq j$ , or it is generated by means of a service call.

We first determine how many fresh values can be generated by service calls. The possibility of generating and storing a new value in  $p$  as a result of an action is reflected by the presence of special edges. By definition, any special edge entering  $p$  must start from a node  $p' \in N_0 \cup \dots \cup N_{i-1}$ . By induction hypothesis, the number of distinct values that can exist in  $p'$  is bounded by  $H(|\text{ADOM}(\mathcal{I}_0)|) = \sum_{j \in \{0, \dots, i-1\}} P_j(|\text{ADOM}(\mathcal{I}_0)|)$ . Let  $b_a$  be the maximum number of special edges that enter a position, over all positions in the schema;  $b_a$  bounds the arity taken by service calls in  $\mathcal{F}$ . Then for every choice of  $b_a$  values in  $N_0 \cup \dots \cup N_{i-1}$  (one for each special edge that can enter a position) and for every action in  $\mathcal{A}^+$ , the number of new values generated at position  $p$  is bounded by  $t_f \cdot H(n)^{b_a}$ , where  $t_f$  is the total number of facts mentioned in the effects of actions that belong to  $\mathcal{A}^+$ . Notice that this number does not depend on the data in  $\mathcal{I}_0$ . By considering all positions in  $N_i$ , the total number of values that can be generated is then bounded by  $G(|\text{ADOM}(\mathcal{I}_0)|) = |N_i| \cdot t_f \cdot H(|\text{ADOM}(\mathcal{I}_0)|)^{b_a}$ . Obviously,  $G(\cdot)$  is a polynomial, because  $t_f$  and  $b_a$  are values extracted from the schema  $\mathcal{R}$  of the DCDS, which is fixed.

We count next the number of distinct values that can be copied to positions of  $N_i$  from positions of  $N_j$ , with  $j \neq i$ . A copy is represented in the graph as a normal edge going from a node in  $N_j$  to a node in  $N_i$ , with  $j \neq i$ . We observe first that such normal edges can start only from nodes in  $N_0 \cup \dots \cup N_{i-1}$ , that is, they cannot start from nodes in  $N_j$  with  $j > i$ . We prove this by contradiction. Assume that there exists  $\langle p', p, \text{false} \rangle \in E$ , such that  $p \in N_i$  and  $p' \in N_j$  with  $j > i$ . In this case, the rank of  $p$  would be  $j > i$ , which contradicts the fact that  $p \in N_i$ . As a consequence, the number of distinct values that can be copied to positions in  $N_i$  is bounded by the total number of values in  $N_0 \cup \dots \cup N_{i-1}$ , which corresponds to  $H(|\text{ADOM}(\mathcal{I}_0)|)$  from our previous consideration. Putting it all together, we define  $P_i(|\text{ADOM}(\mathcal{I}_0)|) = |\text{ADOM}(\mathcal{I}_0)| + G(|\text{ADOM}(\mathcal{I}_0)|) + H(|\text{ADOM}(\mathcal{I}_0)|)$ .  $P_i(\cdot)$  is a polynomial, and therefore the claim is proven.

In the above claim,  $i$  is bounded by the maximum rank  $r$ , which is a constant. Hence, there exists a fixed polynomial  $P(\cdot)$  such that the number of distinct values that can exist in the active domains of the run  $\tau$  is bounded by  $P(|\text{ADOM}(\mathcal{I}_0)|)$ . Technically, given  $\Upsilon_{\mathcal{S}^+} = \langle \mathcal{C}, \mathcal{R}, \Sigma, s_0, db, \Rightarrow \rangle$ , we have:

$$\left| \bigcup_{s \text{ state of } \tau} db(s) \right| < P(|\text{ADOM}(\mathcal{I}_0)|)$$

which attests that  $\tau$  is (data) bounded, and consequently that  $\mathcal{S}$  is run-bounded.  $\square$

## A.3 Nondeterministic Services

### State-bounded Systems

*Proof of Theorem 8.1.* We reuse the proof of Theorem 7.1. Recall that the reduction in this proof constructs for every Turing Machine  $TM$  a DCDS *with deterministic services*  $\mathcal{S}$  that simulates the computation of  $TM$ . It also constructs a propositional safety property  $\Phi$  such that  $\Upsilon_{\mathcal{S}} \models \Phi$  if and only if  $TM$  halts.

What we need here is a reduction to a DCDS with *nondeterministic services*. However, we recall from the proof of Theorem 7.1 that the only service in the process layer, service *newCell*, is guaranteed to be called only with distinct arguments across distinct transitions, and

so its behavior is unaffected by the choice of deterministic versus nondeterministic semantics. Therefore, the reduction applies unchanged to DCDS with nondeterministic services.  $\square$

*Proof of Theorem 8.2.* We prove a stronger result, namely for *linear-time*  $\mu\mathcal{L}_A$  sentences. Such sentences can be written using LTL syntax.

We reduce from the problem of satisfiability of LTL with freeze quantifier over infinite data words, known to be highly undecidable ( $\Sigma_1^1$ -hard) [33].

**Infinite data words [33].** Let  $\Sigma$  be a finite alphabet of labels and  $D$  an infinite set of data values. An infinite *data word*  $w = \{w_i\}$  is an infinite sequence over  $\Sigma \times D$ , i.e., each  $w_i$  is of the form  $(a_i, d_i)$  with  $a_i \in \Sigma$  and  $d_i \in D$ .

**LTL with freeze quantifier ( $\text{LTL}^\downarrow$ ).** This logic operates over infinite data words, seen as runs. It extends propositional LTL with a finite number of *registers*, which can record the data value at the current step of the run (position in the data word), and recall it at subsequent steps. The operation of recording the data value at the current position into register  $i$  is denoted with  $\downarrow_i$ .  $\uparrow_i$  denotes the boolean comparison of the data value at the current position with the value stored in register  $i$ .

As an example, consider the  $\text{LTL}^\downarrow$  sentence

$$\varphi_{ex} = \downarrow_1 X(G(a \implies \neg \uparrow_1))$$

over alphabet  $\{a, b\}$ , which states that the data value assigned to each label  $a$  at positions greater than one is different from the data value at the first position of the data word. Notice that the data value at the first position is recorded in register 1 by operation  $\downarrow_1$ , and it is compared to subsequent data values by  $\uparrow_1$ .

**The DCDS construction.** Given a finite alphabet  $\Sigma = \{\sigma_i\}_{i \in \{1, \dots, n\}}$ , we build a DCDS  $\mathcal{S} = \langle \mathcal{D}_\Sigma, \mathcal{P}_\Sigma \rangle$  with nondeterministic services, such that each run of  $\Upsilon_{\mathcal{S}}$  represents an infinite data word over  $\Sigma$ . In particular, each state in the run holds the label and data value for a single position in the data word. Moreover, given an  $\text{LTL}^\downarrow$  sentence  $\varphi$  over  $\Sigma$ , we construct a  $\mu\mathcal{L}_A$  formula  $\Phi$ , such  $\Upsilon_{\mathcal{S}} \models \Phi$  if and only if  $\varphi$  is *unsatisfiable*.

The idea is to model the registers with existentially quantified variables, which  $\mu\mathcal{L}_A$  allows us to introduce at any given point in the run and use subsequently, even if in between their binding does not persist in the run.

More precisely, we define the data layer  $\mathcal{D}_\Sigma$  of  $\mathcal{S}$  as  $\mathcal{D}_\Sigma = \langle \mathcal{C}, \mathcal{R}, \emptyset, \mathcal{I}_0 \rangle$ , where  $\mathcal{C} = \Sigma \cup \{0\}$ ,  $\mathcal{R} = \{\text{LABEL}/1, \text{DATUM}/1\}$ , and  $\mathcal{I}_0 = \emptyset$ . Intuitively, LABEL stores the label and DATUM the data value. We then define the process layer  $\mathcal{P}_\Sigma$  of  $\mathcal{S}$  as  $\mathcal{P}_\Sigma = \langle \mathcal{F}, \mathcal{A}_\Sigma, \varrho_\Sigma \rangle$ , where:

- $\mathcal{F} = \{f/0\}$ .
- For each  $l \in \{1, \dots, n\}$ ,  $\varrho_\Sigma$  contains an action  $\alpha_l$  with no parameters and no guard ( $\text{true} \mapsto \alpha_l$ ).
- Each  $\alpha_l \in \mathcal{A}_\Sigma$  contains a single effect  $e_l$ , which creates the position of a data word corresponding to label  $\sigma_l \in \Sigma$ :

$$e_l : \text{true} \rightsquigarrow \text{LABEL}(\sigma_l) \wedge \text{DATUM}(f())$$

The service call  $f()$  is used to get an arbitrary data value from the domain during the action execution. It is nondeterministic, and will therefore return possibly distinct values across the run.

Since actions are always executable, at each step of a run all of them qualify, and one is nondeterministically chosen. In this way, the collection of all runs corresponds to all possible infinite data words. Observe that  $\mathcal{S}$  is state-bounded, as each state contains just one label and one data value.

**The property.** We now define the property. For simplicity of presentation, we show it using an LTL-based syntax (branching is irrelevant here), though it is clearly expressible in  $\mu\mathcal{L}_A$ .

We obtain  $\varphi'$  from  $\varphi$  by:

1. replacing each freeze quantifier  $\downarrow_n$  with  $\exists x_n.\text{DATUM}(x_n)$ , and
2. replacing each occurrence of  $\uparrow_n$  with  $\text{DATUM}(x_n)$ , and
3. replacing each proposition  $\sigma \in \Sigma$  with  $\text{LABEL}(\sigma)$ .

Now let  $\Phi := \neg\varphi'$ .

We illustrate the rewrite on property  $\varphi_{ex}$  above, obtaining

$$\varphi'_{ex} := \exists x_1.\text{DATUM}(x_1) \wedge \mathbf{X} \mathbf{G}(\text{LABEL}(a) \implies \neg\text{DATUM}(x_1)).$$

It is easy to see that  $\varphi$  is unsatisfiable over infinite data words using alphabet  $\Sigma$  if and only if  $\Upsilon_S \models \Phi$ .

As a result,  $\mu\mathcal{L}_A$  verification by state-bounded DCDSs with nondeterministic services is undecidable, which concludes the proof of Theorem 8.2.  $\square$

## Abstract Transition System

*Proof of Theorem 8.3.* We formalize the discussion from Section 8.2. Since DCDSs with nondeterministic services are modeled by means of transition systems whose states are constituted by database instances, with a slight abuse of notation we will directly use the state to refer to its database instance.

**Equality commitments.** Consider a set  $D$  comprised of constants and of Skolem terms built by applying a Skolem function to constant arguments. An *equality commitment*  $\mathcal{H}$  on  $D$  is a partition of  $D$ , i.e. a set of disjoint subsets of  $D$ , called *cells*, such that the union of the cells in  $\mathcal{H}$  is  $D$ . Moreover, each cell contains at most one constant (but arbitrarily many Skolem terms). For any  $e \in D$ ,  $[e]_{\mathcal{H}}$  denotes the cell  $e$  belongs to. The intention of the partition is to model equality and non-equality commitments on the members of  $D$  as follows: for every  $e_1, e_2 \in D$ ,  $e_1 = e_2$  if and only if  $[e_1]_{\mathcal{H}} = [e_2]_{\mathcal{H}}$ .

**Service call evaluations that respect equality commitments.** It is convenient to view the concrete transition system  $\Upsilon_S$  in the following equivalent formulation, which emphasizes equality commitments on the service calls: successor states are built by picking an equality commitment  $\mathcal{H}$ , and then picking a service call evaluation that respects  $\mathcal{H}$ . More specifically,

- for each state  $\mathcal{I}$ ,
- for each action  $\alpha$ ,
- for each parameter choice  $\sigma$ , and
- for each equality commitment  $\mathcal{H}$  involving the service calls in  $\text{CALLS}(\text{DO}(\mathcal{I}, \alpha, \sigma))$  and the values in  $\text{ADOM}(\mathcal{I}) \cup \text{ADOM}(\mathcal{I}_0)$ ,

$\Upsilon_S$  contains possibly infinitely many successor states  $\mathcal{I}_{next}$ , each obtained from  $\text{DO}(\mathcal{I}, \alpha, \sigma)$  by picking a service call evaluation that respects  $\mathcal{H}$ . We say that evaluation  $\theta$  *respects*  $\mathcal{H}$  if for every two terms  $t_1, t_2 \in \text{CALLS}(\text{DO}(\mathcal{I}, \alpha, \sigma)) \cup \text{ADOM}(\mathcal{I}) \cup \text{ADOM}(\mathcal{I}_0)$ , we have  $[t_1]_{\mathcal{H}} = [t_2]_{\mathcal{H}}$  if and only if  $t_1\theta = t_2\theta$ .

Given  $\mathcal{I}, \alpha, \sigma$  and  $\mathcal{H}$ , we denote the set of all legal evaluations with

$$\text{EVALS}^{\mathcal{H}}(\mathcal{I}, \alpha, \sigma) := \{\theta \mid \theta \in \text{EVALS}_{\mathcal{C}}(\mathcal{I}, \alpha, \sigma), \theta \text{ respects } \mathcal{H}, \\ \text{DO}(\mathcal{I}, \alpha, \sigma)\theta \models \mathcal{E}\}.$$

Notice that we consider legal only those evaluations that respect the equality commitment  $\mathcal{H}$  and that, conforming to the semantics of the concrete transition system, generate successors which satisfy the constraints  $\mathcal{E}$ . Finally, notice that  $\mathcal{H}$  determines an isomorphism type, as all successors of  $\mathcal{I}$  generated by the evaluations in  $\text{EVALS}^{\mathcal{H}}(\mathcal{I}, \alpha, \sigma)$  are isomorphic to each other.

**Prunings.** We observe that for each state  $\mathcal{I}$  of the concrete transition system  $\Upsilon_{\mathcal{S}}$ , the number of possible choices of  $\alpha, \sigma$  and  $\mathcal{H}$  are finite. The sole reason for infinite branching in  $\Upsilon_{\mathcal{S}}$  are the infinitely many distinct evaluations that respect  $\mathcal{H}$ , whenever  $\mathcal{H}$  states that at least one service call result is distinct from  $\text{ADOM}(\mathcal{I}) \cup \text{ADOM}(\mathcal{I}_0)$ : in that case, the service call can be substituted with any value in  $\mathcal{C} \setminus (\text{ADOM}(\mathcal{I}) \cup \text{ADOM}(\mathcal{I}_0))$ .

In contrast, we obtain a finitely-branching transition system if instead of keeping the successors generated by *all* evaluations in  $\text{EVALS}^{\mathcal{H}}(\mathcal{I}, \alpha, \sigma, \mathcal{H})$ , we keep the successors generated by a *finite subset* of these evaluations (if  $\text{EVALS}^{\mathcal{H}}(\mathcal{I}, \alpha, \sigma)$  is non-empty, we pick a non-empty subset, to ensure that if  $\mathcal{H}$  is represented among the successors of  $\mathcal{I}$  in  $\Upsilon_{\mathcal{S}}$ , it is also represented among the successors of  $\mathcal{I}$  in  $\Theta_{\mathcal{S}}$ ). We call any transition system obtained in this way a *pruning* of  $\Upsilon_{\mathcal{S}}$ , and we denote with  $\text{PRUNINGS}(\Upsilon_{\mathcal{S}})$  the set of all such prunings. By construction, every pruning of  $\Upsilon_{\mathcal{S}}$  is finitely branching.

Formally, let  $\mathcal{S}$  be a DCDS and  $\Upsilon_{\mathcal{S}}$  its concrete transition system, with states  $\Sigma_C$  and initial state  $\mathcal{I}_0$ . A *pruning* of  $\Upsilon_{\mathcal{S}}$  is the restriction of  $\Upsilon_{\mathcal{S}}$  to a subset of states  $\Sigma_P \subseteq \Sigma_C$ , where  $\Sigma_P$  satisfies the following properties:

- (i)  $\mathcal{I}_0 \in \Sigma_P$ , and
- (ii) for each  $\mathcal{I} \in \Sigma_C$  and each equality commitment  $\mathcal{H}$ , if  $\mathcal{H}$  is represented by some successor of  $\mathcal{I}$  in  $\Upsilon_{\mathcal{S}}$ , it is also represented by a successor of  $\mathcal{I}$  in  $\Theta_{\mathcal{S}}$ . We say that  $\mathcal{H}$  is *represented* by successor  $\mathcal{I}'$  of  $\mathcal{I}$  if there exist  $\alpha, \sigma$  and  $\theta \in \text{EVALS}^{\mathcal{H}}(\mathcal{I}, \alpha, \sigma)$  such that  $\langle \mathcal{I}, \alpha\sigma\theta, \mathcal{I}' \rangle \in \text{N-EXEC}_{\mathcal{S}}$ .
- (iii) for each  $\mathcal{I} \in \Sigma_C$ , the number of successors of  $\mathcal{I}$  that are also in  $\Sigma_P$  is finite.

Clearly, a concrete transition system  $\Upsilon_{\mathcal{S}}$  admits (potentially infinitely) many prunings, but we show next that they all are persistence-preserving bisimilar to  $\Upsilon_{\mathcal{S}}$  (and therefore to each other, due to transitivity of the  $\sim$  relation):

**Lemma A.4.** *For every concrete transition system  $\Upsilon_{\mathcal{S}}$  and pruning  $\Theta_{\mathcal{S}} \in \text{PRUNINGS}(\Upsilon_{\mathcal{S}})$ , we have that  $\Theta_{\mathcal{S}} \sim \Upsilon_{\mathcal{S}}$ .*

The result follows from the fact that state isomorphism implies persistence-preserving bisimilarity. In the following, we denote with  $s \mapsto_h s'$  the fact that  $h$  is an isomorphism from state  $s$  to  $s'$ .

**Lemma A.5.** *Consider a concrete transition system  $\Upsilon_{\mathcal{S}}$  with initial state  $s_0$  and one of its prunings  $\Theta_{\mathcal{S}}$ . Let  $s_C$  be a state of  $\Upsilon_{\mathcal{S}}$  and  $s_P$  a state of  $\Theta_{\mathcal{S}}$ . If there exists function  $h$  such that  $h$  fixes  $\text{ADOM}(s_0)$  and  $s_P \mapsto_h s_C$ , then  $s_P \sim_h s_C$ .*

*Proof.* Let  $\Upsilon_{\mathcal{S}} = \langle \mathcal{C}, \mathcal{R}, \Sigma_C, s_0, db, \Rightarrow_C \rangle$  and  $\Theta_{\mathcal{S}} = \langle \mathcal{C}, \mathcal{R}, \Sigma_P, s_0, db, \Rightarrow_P \rangle$ . The proof follows from the following claim:

- Claim 1.** Given  $s_C \in \Sigma_C$  and  $s_P \in \Sigma_P$ , if  $s_P \mapsto_h s_C$  and  $h$  is the identity on  $\text{ADOM}(\mathcal{I}_0)$ , then for each  $s'_C$  such that  $s_C \Rightarrow_C s'_C$  there exist  $s'_P$  and  $h'$  such that
- (i)  $s_P \Rightarrow_P s'_P$ ;
  - (ii)  $h'$  is an extension of  $h|_{\text{ADOM}(s_P) \cap \text{ADOM}(s'_P)}$ ;
  - (iii)  $h'$  is the identity on  $\text{ADOM}(\mathcal{I}_0)$ ;
  - (iv)  $s'_P \mapsto_{h'} s'_C$ .

Indeed, this claim allows us to exhibit the bisimilarity relation

$$R = \{(x, i, y) \mid x \in \Sigma_P, y \in \Sigma_C, x \mapsto_i y\}.$$

$R$  is a bisimilarity relation because it satisfies the forth condition in the definition of persistence-preserving bisimilarity by Claim 1. It trivially satisfies the back condition because  $P$  is constructed by picking a subset of the states of  $\Upsilon_S$ . Since by construction  $(s_P, h, s_C) \in R$ , we have  $s_P \sim_h s_C$ .

To prove Claim 1, we observe that the successor  $s'_C$  of  $s_C$  is generated by a particular choice of the action  $\alpha$  (with condition-action rule  $Q \mapsto \alpha$ ), the parameter instantiation  $\sigma_C$  (such that  $s_C \models Q\sigma_C$ ), the equality commitment  $\mathcal{H}_C$  on  $\text{CALLS}(\text{DO}(s_C, \alpha, \sigma_C)) \cup \text{ADOM}(s_C) \cup \text{ADOM}(\mathcal{I}_0)$ , and the service call evaluation  $\theta_C \in \text{EVALS}^{\mathcal{H}_C}(s_C, \alpha, \sigma_C)$ :  $s'_C = \text{DO}(s_C, \alpha, \sigma_C)\theta_C$ . We show how to construct  $\sigma_P, \mathcal{H}_P$  and  $\theta_P \in \text{EVALS}^{\mathcal{H}_P}(s_P, \alpha, \sigma_P)$  such that  $s_P \models Q\sigma_P$  and  $s'_P = \text{DO}(s_P, \alpha, \sigma_P)\theta_P$  satisfies the claim.

We let  $\sigma_P = h^{-1}(\sigma_C)$ , observing that since  $Q$  is a first-order query, it is preserved under isomorphism, so  $s_C \models Q\sigma_C$  implies  $s_P \models Q\sigma_P$ . Thus,  $\sigma_P$  is a legal parameter instantiation.

To construct  $\mathcal{H}_P, \theta_P$ , we first show that  $\bar{s}_C = \text{DO}(s_C, \alpha, \sigma_C)$  and  $\bar{s}_P = \text{DO}(s_P, \alpha, \sigma_P)$  are isomorphic, as witnessed by the function  $\bar{h} : \text{ADOM}(\bar{s}_P) \rightarrow \text{ADOM}(\bar{s}_C)$  defined as follows:

$$\begin{aligned} \bar{h} := & \{c \mapsto h(c) \mid c \in \text{ADOM}(s_P) \cup \text{ADOM}(\mathcal{I}_0)\} \\ & \cup \{f(m_P, \dots, m_n) \mapsto f(h(m_P), \dots, h(m_n)) \mid \\ & \quad f(m_P, \dots, m_n) \in \text{CALLS}(\bar{s}_P)\}. \end{aligned}$$

From the definition of  $\bar{h}$  and the fact that the service calls are generated by queries preserved under isomorphism, it follows immediately that  $\bar{s}_P \mapsto_{\bar{h}} \bar{s}_C$ . It is easy to see that  $\bar{h}$  is also an isomorphism between  $\text{CALLS}(\text{DO}(s_C, \alpha, \sigma_C)) \cup \text{ADOM}(s_C) \cup \text{ADOM}(\mathcal{I}_0)$  and  $\text{CALLS}(\text{DO}(s_P, \alpha, \sigma_P)) \cup \text{ADOM}(s_P) \cup \text{ADOM}(\mathcal{I}_0)$ , (i.e.  $\bar{h}$  preserves the structure of Skolem terms), and therefore between the sets of corresponding equality commitments.

We therefore pick  $\mathcal{H}_P = \bar{h}^{-1}(\mathcal{H}_C)$ . By construction of  $P$ , each equality type is represented among a state's successors in  $P$ , i.e. there exists  $\theta_P$  that respects  $\mathcal{H}_P$ , and there exists  $s'_P \in \Sigma_P$ , such that  $s'_P = \bar{s}_P\theta_P$ .

The existence of legal choices for  $\alpha, \sigma_P, \mathcal{H}_P$  and  $\theta_P$  proves item (i) of Claim 1, namely that  $s_P \Rightarrow_P s'_P$ .

To prove the remaining items, we exhibit  $h'$  defined as follows:

$$h'(t) := \bar{h}(\bar{t})\theta_C,$$

for some choice of  $\bar{t}$  such that  $\bar{t}\theta_P = t$ .

To see why  $h'$  is well-defined, observe that, by construction of the successor states in  $\Upsilon_S$ , for each  $t \in \text{ADOM}(s'_P)$  there must exist  $\bar{t} \in \text{ADOM}(\bar{s}_P)$  such that  $\theta_P$  evaluates  $\bar{t}$  to  $t$  ( $\bar{t}\theta_P = t$ ). Moreover, observe that if there are distinct  $\bar{t}, \bar{u} \in \text{ADOM}(\bar{s}_P)$  such that  $t = \bar{t}\theta_P = \bar{u}\theta_P$ , it does not matter which one we pick in the definition of  $h'$ , since  $\bar{h}(\bar{t})\theta_C = \bar{h}(\bar{u})\theta_C$ . This is because  $\theta_P$  respects  $\mathcal{H}_P$ , and therefore  $[\bar{t}]_{\mathcal{H}_P} = [\bar{u}]_{\mathcal{H}_P}$ . Since  $\mathcal{H}_P \mapsto_{\bar{h}} \mathcal{H}_C$ , it follows that  $[\bar{h}(\bar{t})]_{\mathcal{H}_C} = [\bar{h}(\bar{u})]_{\mathcal{H}_C}$ , and since  $\theta_C$  respects  $\mathcal{H}_C$ , we have that  $\bar{h}(\bar{t})\theta_C = \bar{h}(\bar{u})\theta_C$ .

Items (ii), (iii), and (iv) of Claim 1 follow by similar reasoning from the fact that service call evaluations respect the equality commitments, which are isomorphic.  $\square$

*Proof of Lemma A.4.* This is a corollary of Lemma A.5.

Indeed, by definition,  $\Theta_S \sim \Upsilon_S$  holds if and only if the initial state  $s_0^P$  of  $\Theta_S$  is bisimilar to the initial state  $s_0^C$  of  $\Upsilon_S$ , i.e. there exists isomorphism  $h$  such that  $s_0^P \sim_h s_0^C$ .

By definition, a concrete transition system shares the initial state with all its prunings, so  $s_0^P = s_0^C$ . The identity mapping  $id$  witnesses isomorphism:  $s_0^P \mapsto_{id} s_0^C$ . By Lemma A.5, we have  $s_0^P \sim_{id} s_0^C$ .  $\square$

**Eventually Recycling Prunings.** While all prunings of a concrete transition system are finitely-branching, they are not guaranteed to be finite. The reason is that they don't necessarily rule out infinitely long simple runs  $\tau$ , along which the service calls return in each state  $\mathcal{I}$  “fresh” values, i.e. values distinct from all values appearing in  $\mathcal{I}$  and its predecessors on  $\tau$ . Towards addressing this problem, we focus on prunings in which the evaluations are not chosen arbitrarily.

Given a finite run  $\tau$  ending in state  $\mathcal{I}$  of  $\Upsilon_{\mathcal{S}}$ , an action  $\alpha$ , a parameter choice  $\sigma$  and an equality commitment  $\mathcal{H}$  on  $\text{CALLS}(\text{DO}(\mathcal{I}, \alpha, \sigma))$ , we say that evaluation  $\theta \in \text{EVALS}^{\mathcal{H}}(\mathcal{I}, \alpha, \sigma)$  *recycles from*  $\tau$  if each value in the range of  $\theta$  occurs in  $\tau$ . We say that pruning  $\Theta_{\mathcal{S}}$  is *eventually recycling* if every (finite or infinite) path  $\tau$  in  $\Theta_{\mathcal{S}}$  contains only finitely many states generated by non-recycling evaluations. Formally, if  $\tau = s_0 s_1 \cdots$  and the service call evaluation used in  $s_i \Rightarrow_C s_{i+1}$  is denoted as  $\theta_i$ , then there are only finitely many indexes  $j$  such that  $\theta_j$  does not recycle from  $\tau[j]$ .

**Lemma A.6.** *Let  $\Upsilon_{\mathcal{S}}$  be a concrete transition system.*

- (i) *All eventually recycling prunings of  $\Upsilon_{\mathcal{S}}$  are finite.*
- (ii) *If  $\Upsilon_{\mathcal{S}}$  is state-bounded, then it has at least one eventually recycling pruning.*

*Proof.* (i): All eventually recycling prunings are finite.

Let  $\Theta_{\mathcal{S}}$  be an eventually recycling pruning of concrete transition system  $\Upsilon_{\mathcal{S}}$ . By virtue of being a pruning,  $\Theta_{\mathcal{S}}$  is finitely branching. We show next that every simple path in  $\Theta_{\mathcal{S}}$  has finite length, which together with finite branching implies finiteness by König's Lemma.

Towards a contradiction, assume that there exists infinite simple run  $\tau$  in  $\Theta_{\mathcal{S}}$ . Since  $\Theta_{\mathcal{S}}$  is eventually recycling, there is a finite prefix of  $\tau$  such that all values occurring in  $\tau$  occur also in this prefix. Therefore,  $\tau$  contains only finitely many distinct values, and hence only finitely many distinct states (databases of given schema over these values). If  $\tau$  has infinite length, then a pigeonhole argument contradicts the assumption that  $\tau$  is simple.

(ii): If  $\Upsilon_{\mathcal{S}}$  is state-bounded, then it has an eventually recycling pruning.

Let  $\Theta_{\mathcal{S}}$  be a pruning obtained from  $\Upsilon_{\mathcal{S}}$  by picking the finite subset of evaluations  $SE \subseteq \text{EVALS}^{\mathcal{H}}(s, \alpha, \sigma)$  as follows: if there is a run  $\tau$  in  $\Theta_{\mathcal{S}}$  from  $s_0$  to  $s$  such that  $\text{EVALS}^{\mathcal{H}}(s, \alpha, \sigma)$  includes at least one evaluation that recycles from  $\tau$ , then  $SE$  contains exclusively recycling evaluations (i.e. for each evaluation  $\theta \in SE$ , there is a run  $\tau$  from  $s_0$  to  $s$  in  $\Theta_{\mathcal{S}}$  such that  $\theta$  recycles from  $\tau$ ). Otherwise,  $SE$  is an arbitrary finite subset of  $\text{EVALS}^{\mathcal{H}}(s, \alpha, \sigma)$ .

We prove that pruning  $\Theta_{\mathcal{S}}$  is eventually recycling. By definition, if  $\Upsilon_{\mathcal{S}}$  is state-bounded then  $|\text{ADOM}(s)| \leq b$  for each state  $s$ , where  $b$  is the size bound on the state. Assume towards a contradiction that  $\Theta_{\mathcal{S}}$  contains a run  $\tau = s_0 s_1 s_2 \cdots$  that includes infinitely many states generated by evaluations that do not recycle from  $\tau$ . It follows that there must exist a finite  $k \geq 0$  such that  $|\bigcup_{i=0}^k \text{ADOM}(s_i)| > 3b$  and such that  $\text{ADOM}(s_{k+1})$  contains at least one fresh value, i.e.  $\text{ADOM}(s_{k+1}) - \bigcup_{i=0}^k \text{ADOM}(s_i) \neq \emptyset$ . Let  $\theta_{k+1} \in \text{EVALS}^{\mathcal{H}}(s_k, \alpha, \sigma)$  be the service call evaluation that generates  $s_{k+1}$ . Clearly  $\theta_{k+1}$  does not recycle from  $\tau$ , since it contains at least one fresh value in its range. However observe that, since the  $k$ -length prefix of  $\tau$  contains at least  $3b$  distinct values, this prefix contains at least  $b$  values that are distinct from the values in  $\text{ADOM}(\mathcal{I}_0) \cup \text{ADOM}(s_k)$  (since by state-boundedness,  $|\text{ADOM}(\mathcal{I}_0) \cup \text{ADOM}(s_k)| \leq 2b$ ). Call the set of these values  $\mathcal{V}$ .

Also by state-boundedness,  $\theta_{k+1}$  introduces at most  $b$  fresh values. Any one of the values in  $\mathcal{V}$  can be used instead of the fresh values introduced by  $\theta_{k+1}$ , to obtain another evaluation  $\theta_{k+1}^r$  that respects  $\mathcal{H}$ . Hence  $\theta_r$  witnesses an evaluation in  $\text{EVALS}^{\mathcal{H}}(s_k, \alpha, \sigma)$  that *does* recycle from  $\tau$ . But this contradicts the definition of  $\Theta_{\mathcal{S}}$ , which mandates that  $\theta_{k+1}$  be dropped in favor of  $\theta_{k+1}^r$ .  $\square$

This result implies that if  $\Upsilon_{\mathcal{S}}$  is state-bounded, then there exists a finite-state abstract transition system  $\Theta_{\mathcal{S}}$  that is persistence-preserving bisimilar to  $\Upsilon_{\mathcal{S}}$ . Indeed, any eventually recycling pruning of  $\Upsilon_{\mathcal{S}}$  can play the role of  $\Theta_{\mathcal{S}}$  (it is finite by Lemma A.6(i), it is bisimilar to  $\Upsilon_{\mathcal{S}}$  by Lemma A.4, and one is guaranteed to exist by Lemma A.6(ii)).

**Construction of Eventually Recycling Pruning.** The existence result in Lemma A.6 is non-constructive and therefore does not yet yield decidability of verification even if the concrete transition system  $\Upsilon_{\mathcal{S}}$  is state-bounded. We next present Algorithm RCYCL, which is guaranteed to construct an eventually recycling pruning when its input DCDS is state-bounded, but which may diverge otherwise.

**Algorithm RCYCL**

**Input:**  $\mathcal{S} = \langle \mathcal{D}, \mathcal{P} \rangle$ , a DCDS with data layer  $\mathcal{D} = \langle \mathcal{C}, \mathcal{R}, \mathcal{E}, \mathcal{I}_0 \rangle$   
and process layer  $\mathcal{P} = \langle \mathcal{F}, \mathcal{A}, \varrho \rangle$ .

$\Sigma := \{\mathcal{I}_0\}$ ,  $\Longrightarrow := \emptyset$ ,  $UsedValues := \text{ADOM}(\mathcal{I}_0)$ ,  $Visited := \emptyset$

**repeat**

pick state  $\mathcal{I} \in \Sigma$ , action  $\alpha$  and legal parameters  $\sigma$   
such that  $(\mathcal{I}, \alpha, \sigma) \notin Visited$

$RecyclableValues := UsedValues - (\text{ADOM}(\mathcal{I}_0) \cup \text{ADOM}(\mathcal{I}))$

pick set  $\mathcal{V}$  of  $n$  service call results such that:

$|\mathcal{V}| = n = |\text{CALLS}(\text{DO}(\mathcal{I}, \alpha, \sigma))|$  and

**if**  $|RecyclableValues| \geq n$

**then**  $\mathcal{V} \subseteq RecyclableValues$  % recycled values

**else**  $\mathcal{V} \subseteq \mathcal{C} - UsedValues$  % fresh values

$F := \text{ADOM}(\mathcal{I}_0) \cup \text{ADOM}(\mathcal{I}) \cup \mathcal{V}$

**for each**  $\theta \in \text{EVALS}_F(\mathcal{I}, \alpha, \sigma)$  such that  $\mathcal{I}_{next} \models \mathcal{E}$

where  $\mathcal{I}_{next} := \text{DO}(\mathcal{I}, \alpha, \sigma)\theta$  **do**

$\Sigma := \Sigma \cup \{\mathcal{I}_{next}\}$

$\Longrightarrow := \Longrightarrow \cup \{(\mathcal{I}, \mathcal{I}_{next})\}$

$UsedValues := UsedValues \cup \text{ADOM}(\mathcal{I}_{next})$

$Visited := Visited \cup \{(\mathcal{I}, \alpha, \sigma)\}$

**end**

**until**  $\Sigma$  and  $\Longrightarrow$  no longer change.

**return**  $\langle \mathcal{C}, \mathcal{R}, \Sigma, \mathcal{I}_0, \Longrightarrow \rangle$

Observe that algorithm RCYCL performs several nondeterministic choices in each iteration. The particular choices (and their order) do not matter, by Theorem 8.4.

*Proof of Theorem 8.4 (Sketch).* First, we show that algorithm RCYCL builds a pruning. Items (i) and (iii) in the definition of pruning are trivially satisfied in every run of RCYCL. Item (ii) follows from the following claim:

**Claim:** for any choice of  $\mathcal{V}$  such that  $|\mathcal{V}| \geq |\text{CALLS}(\text{DO}(\mathcal{I}, \alpha, \sigma))|$ , the set of equality commitments represented by the successors of  $\mathcal{I}$  generated by the evaluations in  $\text{EVALS}_F(\mathcal{I}, \alpha, \sigma)$  coincides with the set of commitments represented by the successors of  $\mathcal{I}$  in  $\Upsilon_{\mathcal{S}}$ .

Next, we show that if  $\mathcal{S}$  is state-bounded, every run of RCYCL terminates. Indeed, state-boundedness guarantees that in each iteration, only at most  $b$  service call values are needed, where  $b$  is the state size bound. But after running “sufficiently” long, RCYCL variable  $UsedValues$  accumulates at least  $3b$  distinct values. At each subsequent step of the algorithm, there will therefore exist at least  $b$  values distinct from the active domains of  $\mathcal{I}_0$  and  $\mathcal{I}$ , so the pick of  $\mathcal{V}$  will always recycle values (observe that RCYCL only picks evaluations from set  $\Theta_{\mathcal{S}}$ ).  $UsedValues$  will no longer change, and therefore  $\Sigma$  and  $\Longrightarrow$  must eventually saturate (a key reason for this is the bookkeeping of variable  $Visited$ , which avoids repeating the nondeterministic pick for any combination of state, action and parameter instantiation  $(\mathcal{I}, \alpha, \sigma)$ ).



Finally, since RCYCL terminates, then it outputs a finite-state pruning, which is trivially eventually recycling.  $\square$

Theorem 8.4 and Theorem 6.2 directly imply Theorem 8.3.  $\square$

## GR-Acyclic DCDSs

*Proof of Theorem 8.5.* For the proof, we reduce from the undecidable problem of checking if the run of a deterministic Turing Machine is confined to a bounded-length segment of the tape (we say that the TM is tape-bounded). This in turn is undecidable by reduction from the halting problem: Given deterministic TM  $T$ , build TM  $T'$  such that  $T'$  is tape-bounded if and only if  $T$  halts.  $T'$  simulates  $T$  but also records on the tape the historical configurations of  $T$ . At each step,  $T'$  checks if the most recent configuration of  $T$  was seen in the history. If so,  $T'$  stops simulating  $T$  and enters a loop in which it keeps extending the right end of its tape. It is easy to see that  $T'$  is tape-bounded if and only if  $T$  halts.

We reuse without change the reduction exhibited in the proof of Theorem 7.1. Recall that the reduction constructs for every Turing Machine  $TM$  a DCDS with deterministic services  $\mathcal{S}$  that simulates the computation of  $TM$ . We recall from the proof that the only service in the process layer, service *newCell*, is guaranteed to be called only with distinct arguments across distinct transitions, and so its behavior is unaffected by the choice of deterministic versus nondeterministic semantics. We also note that the state of the DCDS has size linear in the length of the tape segment visited by  $TM$ , so tape-boundedness reduces to state-boundedness.  $\square$

*Proof of Theorem 8.6 (Sketch).* We prove the result by counting the maximum number of different values in a state of the transition system.

Since this task is undecidable (by Theorem 8.5), we necessarily have to approximate this value. The approximation is performed by analysing a different, much more abstract transition system we call *dataflow* transition system (to distinguish from the abstract system that is bisimilar to the concrete system).

The dataflow system is a DCDS obtained as follows from the dataflow graph and  $\mathcal{I}_0$ : For each node of the dataflow graph, there is a unary relation in the dataflow system, and for each normal (special) edge in the dataflow graph, there is a normal (special) transition in the dataflow system between the corresponding relations. The schema of the dataflow system is a set of relation names with arity one, in correspondence to the nodes of the dataflow graph. A state of the dataflow system is an instantiation of its schema using values from the domain  $\mathcal{C}$ .

For each term  $t$  appearing in a relation in the initial state of the concrete system, there is a term  $t$  in the corresponding relation of the initial state of the dataflow system. Being in one state of the dataflow system, the next state is constructed as follows:

- for each normal transition from a relation  $A$  to a relation  $B$ , for each term  $t$  in the relation  $A$  of the current state, there is a term  $t$  in the relation  $B$  of the next state.
- for each special edge from a relation  $A$  to a relation  $B$ , for each term  $t$  in the node  $A$  of the current state, there is a fresh term  $t'$  in node  $B$  of the next state.

It is easy to see the following claim:

**Claim 1.** For any run  $\tau$  of length  $m \geq 0$  in the concrete system, there is a run  $\tau^d$  of length  $m$  in the dataflow system, such that the size of the active domain of state  $\tau(i)$  is at most the size of the active domain of state  $\tau^d(i)$ .

As a result, any state bound for the dataflow system also bounds the state of the concrete system. We compute such a bound next.

Consider the dataflow graph of  $\mathcal{A}$ . GR-acyclicity forces cycles with special edges to not be connected to any other cycles in the dataflow graph. More specifically, each connected component of the dataflow graph must have one of the following types:

- A: A simple cycle  $C$  (possibly with special edges), possibly connected with several directed acyclic graphs (DAG)s, such that the component contains no additional cycle beyond  $C$ .
- B: Several cycles  $C_1, \dots, C_m$  containing only normal edges, each  $C_i$  possibly connected to several DAGs, such that the component contains no other cycle beyond the  $C_i$ 's, and there is no path with special edges connecting two cycles  $C_i, C_j$ .
- C: A DAG, possibly containing normal and special edges.

Denote with

- d: the longest path of the dataflow graph after deleting the cycles,
- b: the maximum number of special edges going out of a node of the dataflow graph plus one, and
- n: the number of nodes of the dataflow graph.

It is easy to see that in each transition of the dataflow system, for each term in the current state, there can be at most  $n \cdot b$  distinct terms in the next state.

First, consider the components of type A. Call the DAGs  $D$  connected to the unique cycle  $C$  via edges from  $D$  to  $C$ , *input DAGs*. Call *output DAGs* the DAGs connected via edges from  $C$  to  $D$ . It is easy to see that after  $d$  transition steps, in any run of the dataflow system there is no term in any relation of an input DAG (all have been forgotten), and at most

$$m := |\text{ADOM}(\mathcal{I}_0)| + n \cdot b \cdot |\text{ADOM}(\mathcal{I}_0)| + \dots + n^d \cdot b^d \cdot |\text{ADOM}(\mathcal{I}_0)|$$

distinct terms may co-exist within the relations of the cycle. Moreover, after  $d$  steps, the total number of distinct terms in the cycle will no longer increase in any run suffix starting from step  $d + 1$ . Consider now how the  $m$  terms can be copied into the output DAGs. It is again easy to see that there can be at most  $n^d \cdot b^d \cdot m$  distinct terms in any relation of an outgoing DAG. As a result, in any step at most  $n^{d+1} \cdot b^d \cdot m$  distinct terms may co-exist within a type A component.

Second, consider the components of type B. A similar argument yields at most  $n^{d+1} \cdot b^d \cdot m$  different terms that may co-exist within a type B component.

Third, it is easy to see that there can be at most  $n^d \cdot b^d \cdot |\text{ADOM}(\mathcal{I}_0)|$  different terms within a type C component.

All in all, at most  $|\text{ADOM}(\mathcal{I}_0)| \cdot n^{2d+1} \cdot b^{2d}$  distinct terms may co-exist in a state of the concrete transition system.  $\square$

*Proof of Theorem 10.1 (Sketch).* The technical problem here is to force the results of nondeterministic service calls to conform to historic evaluations.

Let  $D$  be a deterministic DCDS. We rewrite  $D$  to obtain a new DCDS  $N$  whose semantics under nondeterministic services coincides with that of  $D$  under deterministic services. For each term  $f(a_1, \dots, a_n)$  appearing in some effect of  $D$   $e := q^+ \wedge Q^- \rightsquigarrow E$ , we rewrite  $D$  as follows. We extend the schema with a new  $n + 1$ -ary relation  $R_f$ . Intuitively,  $R_f(a_1, \dots, a_n, r)$  states that the call  $f(a_1, \dots, a_n)$  evaluates to  $r$ . We extend the effect to record this fact, replacing  $e$  with  $e' := q^+ \wedge Q^- \rightsquigarrow E \wedge R_f(a_1, \dots, a_n, f(a_1, \dots, a_n))$ . To ensure that  $R_f$  records all past calls of  $f$ , we add to each action an effect that simply copies  $R_f$ . We also add the functional dependency  $a_1, \dots, a_n \rightarrow r$  on  $R_f$ . Notice that any attempt to record a service call with a result distinct from a past invocation violates the functional dependency and the transition does not occur. It is easy to see that, if we project the states of  $\Upsilon_N$  on the schema of  $D$ , we obtain  $\Upsilon_D$ .  $\square$

*Proof of Theorem 10.2 (Sketch).* The challenge here lies in forcing a deterministic service  $f_d$  to return possibly distinct results for same-argument calls of the nondeterministic service  $f_n$  it corresponds to.

The trick is to call  $f_d$  with one additional argument, which plays the role of a timestamp, where each state in the run has its own unique timestamp. This way same-argument calls of  $f_n$  at distinct steps in the run correspond to distinct-argument calls of  $f_d$ , which therefore simulates the desired nondeterministic behavior.

An additional trick is used to get the run to generate a sequence of unique values to act as timestamps. We use a deterministic service  $new(x)$  to generate a new timestamp, we record the successor relation over timestamps in binary relation  $succ$ , and the most recent timestamp in unary relation  $now$ . We add to each action

- the effect

$$now(x) \rightsquigarrow now(new(x)) \wedge succ(x, new(x))$$

which extends the successor relation by one timestamp and sets the new timestamp as most recent; and

- the effect

$$succ(x, y) \rightsquigarrow succ(x, y)$$

which accumulates the historical  $succ$  entries.

We are not quite done, as we still need to ensure that  $succ$  induces a linear order on the collection of timestamps generated during the run. For this purpose we employ the same trick as the proof of Theorem 7.1. We describe it below for the sake of proof self-containment.

Observe that, by definition of the effect extending  $succ$ , at each step, the generated timestamp has at most one successor.

However, if the call  $new(x)$  returns a previously seen timestamp, then there can be some timestamp with several predecessors in  $succ$ . We rule out this case by declaring the second component of  $succ$  to be a key. It follows that  $succ$  must be either (i) a linear path over timestamps (possibly starting from a source node that has a self-loop), or (ii) must contain a simple cycle involving more than one timestamp. The simple cycle is created when  $new$  returns the minimal element of the  $succ$  relation.

We wish to force case (i). To rule out case (ii), we proceed as follows: we initialize  $succ$  to contain a source node 0 that can never be a timestamp because it cannot be returned by  $new$  without violating the key constraint on  $succ$ . To this end, we initialize  $succ$  in  $\mathcal{I}_0$  to  $succ^{\mathcal{I}_0} = \{(0, 0), (0, 1)\}$  and  $now^{\mathcal{I}_0} = \{1\}$ . Notice that  $succ^{\mathcal{I}_0}$  has type (i). An easy induction shows that every run prefix must also construct a  $succ$  relation of type (i), since any attempt to extend  $succ$  with an edge back to one of its existing nodes violates the key constraint. It also follows easily that during the run,  $now$  takes values from the linear path starting at 1, and never includes 0.  $\square$

## B Proofs for Semantic Artifacts

In this section, we provide the proof of Theorem 13.3. Given a KAB  $\mathcal{K} = (\mathcal{T}, A_0, \Gamma, \Pi)$ , in what follows, we consider each ABox  $A$  appearing in some state of the transition system generated by  $\mathcal{K}$  as closed under the equalities present in  $A$  itself or implied by the TBox  $\mathcal{T}$ , i.e., we assume that  $A$  is replaced by<sup>17</sup>

$$A \cup \text{EQ}_{\mathcal{T}}(A).$$

We start by observing that we can perform a sequence of transformations of  $\Gamma$  that preserve the semantics of  $\mathcal{K}$  and that are applied to the effect specifications of each action appearing in  $\Gamma$ .

1. We replace each ECQ  $q^+ \wedge Q^-$  that constitutes the left-hand side of an effect specification by  $rew(q^+) \wedge rew(Q^-)$  [25], exploiting the results in [26, 3], which guarantee that, for every ECQ  $q$ ,

$$\text{ANS}(q, T, A) = \text{ANS}(rew(q), \emptyset, A).$$

2. We replace each effect specification  $[q^+] \wedge Q^- \rightsquigarrow A'$ , resulting from Step 1, by a set of effect specifications  $[q_i^+] \wedge Q^- \rightsquigarrow A'$ , one for each CQ  $q_i$  in the UCQ  $q^+$ .

<sup>17</sup>For readability, we make here the TBox  $\mathcal{T}$  in  $\text{EQ}_{\mathcal{T}}()$  explicit.

3. For each effect specification  $[q_i^+] \wedge Q^- \rightsquigarrow A'$ , we re-express  $q_i^+$  so as to make equalities used to join terms explicit and so as to remove constants from  $q_i^+$ . Specifically, we replace the effect specification by

$$[q_i^{++}] \wedge q^- \wedge Q^- \rightsquigarrow A',$$

where

- $q_i^{++}$  is the CQ without repeated variables obtained from  $q_i^+$  by (i) replacing for each variable  $x$  occurring in  $q_i^+$ , the  $j$ -th occurrence of  $x$  except for the first one, by  $x^{[j]}$ ; and (ii) replacing each constant  $c$  with a new variable  $x_c$ ;
- $q^- = \bigwedge [x = x^{[j]}] \wedge \bigwedge [x_c = c]$ , where (i) the first conjunction contains one equality  $[x = x^{[j]}]$  for each variable  $x$  in  $q_i^+$  and for each variable  $x^{[j]}$  introduced in the step above, and (ii) the second conjunction contains one equality for each constant  $c$  in  $q_i^+$ .

For example, in Step 3, a query  $[q_i^+] = [N(x) \wedge P_1(x, y) \wedge P_2(c, x)]$  is replaced by  $[q_i^{++}] \wedge q^-$ , where  $q_i^{++} = N(x) \wedge P_1(x^{[2]}, y) \wedge P_2(x_c, x^{[3]})$  and  $q^- = [x = x^{[2]}] \wedge [x = x^{[3]}] \wedge [x_c = c]$ .

As for the correctness of Step 3, notice that it is immediate that  $[q_i^+]$  is equivalent to  $[q_i^+ \wedge \bigwedge (x = x^{[j]})]$ . The equivalence between the latter and  $[q_i^+] \wedge q^-$  is a consequence of the construction in [3], which shows that query entailment in the presence of equalities can be reduced to query evaluation by saturating equalities with respect to transitivity, reflexivity, symmetry, and functionality.

Given a KAB  $\mathcal{K}$ , we denote with  $\mathcal{K}_{\text{SMP}}$  its simplified version obtained by applying the above transformations to all effect specifications in  $\mathcal{K}$ . Hence, all the effect specifications in  $\mathcal{K}_{\text{SMP}}$  have the form specified at Step 3 above.

## B.1 Simplified Transition System

We introduce a simplified but equivalent representation of  $\text{DO}()$ , denoted by  $\text{DO}_{\text{SMP}}()$ . In particular, given an action  $\gamma = \{e_1, \dots, e_n\}$  with parameters substitution  $\theta$  and an ABox  $A$ , we have

$$\text{DO}_{\text{SMP}}(\mathcal{T}, A, \gamma\theta) = \bigcup_{e \in \gamma} \text{APPLY}_{\text{SMP}}(\mathcal{T}, e, \theta, A),$$

where, for an effect specification

$$e : [q^{++}] \wedge q^- \wedge Q^- \rightsquigarrow A',$$

we have

$$\text{APPLY}_{\text{SMP}}(\mathcal{T}, e, \theta, A) = \bigcup_{\sigma \in \text{ANS}([q^{++}]\theta, \emptyset, A^-) \cap \text{ANS}((q^- \wedge Q^-)\theta, \emptyset, A)} A'\theta\sigma \cup \text{EQ}_{\mathcal{T}}(A).$$

Here,  $A^- = A \setminus \text{EQ}_{\mathcal{T}}(A)$ , i.e.,  $A^-$  is the set of ABox assertions in  $A$  of the form  $N(t_1)$  and  $P(t_1, t_2)$ .

The following lemma shows that  $\text{DO}()$  and  $\text{DO}_{\text{SMP}}()$  are logically equivalent.

**Lemma B.1.**  *$(\mathcal{T}, \text{DO}(\emptyset, A, \gamma\theta))$  and  $(\mathcal{T}, \text{DO}_{\text{SMP}}(\emptyset, A, \gamma\theta))$  are logically equivalent.*

*Proof.* We can rewrite the definition of  $\text{DO}()$  given in Section 12 as follows:

$$\text{DO}(\mathcal{T}, A, \gamma\theta) = \bigcup_{e \in \gamma} \text{APPLY}(\mathcal{T}, e, \theta, A),$$

where

$$\text{APPLY}(\mathcal{T}, e, \theta, A) = \bigcup_{\sigma \in \text{ANS}([q^{++}]\theta, \emptyset, A) \cap \text{ANS}((q^- \wedge Q^-)\theta, \emptyset, A)} A'\theta\sigma \cup \text{EQ}_{\mathcal{T}}(A).$$

In order to prove the lemma it is enough to show that for each ABox assertion  $\alpha_1 \in \text{DO}(\emptyset, A, \gamma\theta)$  we have that  $(\mathcal{T}, \text{DO}_{\text{SMP}}(\emptyset, A, \gamma\theta)) \models \alpha_1$ , and for each ABox assertion  $\alpha_2 \in \text{DO}_{\text{SMP}}(\emptyset, A, \gamma\theta)$  we have that  $(\mathcal{T}, \text{DO}(\emptyset, A, \gamma\theta)) \models \alpha_2$ .

We prove the following stronger result:

(i) for each ABox assertion  $\alpha \in \text{APPLY}_{\text{SMP}}(\mathcal{T}, e, \theta, A)$ , we have that  $(\mathcal{T}, \text{APPLY}(\mathcal{T}, e, \theta, A)) \models \alpha$ , and

(ii) for each ABox assertion  $\alpha \in \text{APPLY}(\mathcal{T}, e, \theta, A)$ , we have that  $(\mathcal{T}, \text{APPLY}_{\text{SMP}}(\mathcal{T}, e, \theta, A)) \models \alpha$ .

Using the monotonicity of  $q^{++}$  and the fact that  $A^- \subseteq A$ , we have that  $\bigcup_{\sigma \in (\text{ANS}([q^{++}]\theta, \emptyset, A^-) \cap \text{ANS}((q^- \wedge Q^-)\theta, \emptyset, A))} A'\theta\sigma$  is contained in  $\bigcup_{\sigma \in (\text{ANS}([q^{++}]\theta, \emptyset, A) \cap \text{ANS}((q^- \wedge Q^-)\theta, \emptyset, A))} A'\theta\sigma$ , which results in (i).

For the proof of the other direction, (ii), consider an ABox assertion  $\alpha \in \text{APPLY}(\mathcal{T}, e, \theta, A)$  and  $e : [q^{++}] \wedge Q^- \wedge q^- \rightsquigarrow A'$ . Using the definition of  $\text{APPLY}()$ , we know that there exists an assignment  $\sigma$  to the free variables of  $q^{++}$  such that  $\sigma \in (\text{ANS}([q^{++}]\theta, \emptyset, A) \cap \text{ANS}((q^- \wedge Q^-)\theta, \emptyset, A))$  and  $\alpha \in A'\theta\sigma$ . Let  $\{x_1, \dots, x_n\}$  be all free variables in  $q^{++}\theta$ , and  $\sigma = \{x_1 \rightarrow t_1, \dots, x_n \rightarrow t_n, \}$ . For each variable  $x_i$ , let  $N(x_i)$  be the (unique) atom in  $q^{++}\theta$  in which  $x_i$  occurs (similar considerations hold when  $x_i$  occurs in a role atom). Then, either  $N(t_i) \in A^-$ , or for some  $t'_i$ ,  $N(t'_i) \in A^-$  and  $(t_i = t'_i) \in \text{EQ}_{\mathcal{T}}(A)$ . In the former case, we use  $t''_i$  to denote  $t_i$ , while in the latter case we use  $t''_i$  to denote  $t'_i$ . Then, let  $\sigma' = \{x_1 \rightarrow t''_1, \dots, x_n \rightarrow t''_n, \}$ . By construction, we have that  $\sigma' \in \text{ANS}([q^{++}]\theta, \emptyset, A^-)$ , and since  $\sigma \in \text{ANS}((q^- \wedge Q^-)\theta, \emptyset, A)$ , and  $(t''_i = t_i) \in \text{EQ}_{\mathcal{T}}(A)$  for each  $i \in \{1, \dots, n\}$ , we also have that  $\sigma' \in \text{ANS}([q^{++}]\theta, \emptyset, A^-) \cap \text{ANS}((q^- \wedge Q^-)\theta, \emptyset, A)$ . Since (i)  $\alpha \in A'\theta\sigma$ , (ii)  $\sigma$  and  $\sigma'$  are identical modulo  $\text{EQ}_{\mathcal{T}}(A)$ , and (iii)  $\text{EQ}_{\mathcal{T}}(A) \in \text{APPLY}_{\text{SMP}}(\mathcal{T}, e, \theta, A)$ , we can infer that  $(\mathcal{T}, \text{APPLY}_{\text{SMP}}(\mathcal{T}, e, \theta, A)) \models \alpha$ .  $\square$

**Simplified transition system.** Given a KAB  $\mathcal{K}$  and its simplified version  $\mathcal{K}_{\text{SMP}}$ , we call the transition system generated in the same way as  $\Upsilon$ , but using  $\text{DO}_{\text{SMP}}()$  on  $\mathcal{K}_{\text{SMP}}$  instead of  $\text{DO}()$  on  $\mathcal{K}$ , the *simplified transition system* generated by  $\mathcal{K}_{\text{SMP}}$ , and denote it with  $\Upsilon_{\mathcal{K}, \text{SMP}}$ .

**Lemma B.2.** *Given a KAB  $\mathcal{K}$ , the transition systems  $\Upsilon$  and  $\Upsilon_{\mathcal{K}, \text{SMP}}$  are bisimilar.*

*Proof.* Let  $\Upsilon = (\mathbb{U}, \mathcal{T}, \Sigma, s_0, \text{abox}, \Rightarrow)$  and  $\Upsilon_{\mathcal{K}, \text{SMP}} = (\mathbb{U}, \mathcal{T}, \Sigma_{\text{SMP}}, s_0, \text{abox}_{\text{SMP}}, \Rightarrow_{\text{SMP}})$ . We define the relation  $\mathcal{B} \subseteq \Sigma \times \Sigma_{\text{SMP}}$  as follows:  $(s_1, s_2) \in \mathcal{B}$  iff  $(\mathcal{T}, \text{abox}(s_1))$  and  $(\mathcal{T}, \text{abox}_{\text{SMP}}(s_2))$  are logically equivalent. We now show that  $\mathcal{B}$  is a bisimulation. We start by making the following observation. If  $(\mathcal{T}, A_1)$  and  $(\mathcal{T}, A_2)$  are logically equivalent, then  $(\mathcal{T}, \text{DO}(\mathcal{T}, A_1, \gamma\theta))$  and  $(\mathcal{T}, \text{DO}_{\text{SMP}}(\mathcal{T}, A_2, \gamma\theta))$  are logically equivalent as well. This is an easy consequence of Lemma B.1 and the observation that logical equivalence is transitive.

To show that  $\mathcal{B}$  is a bisimulation, we prove that  $\mathcal{B}$  is closed under the definition of bisimulation itself. Indeed, if  $(s_1, s_2) \in \mathcal{B}$ , then:

1.  $(\mathcal{T}, \text{abox}(s_1))$  and  $(\mathcal{T}, \text{abox}(s_2))$  are logically equivalent by definition.
2. If  $s_1 \Rightarrow s'_1$  then there exists an action  $\gamma$  and a substitution  $\theta$  such that  $s'_1 = \text{DO}(\mathcal{T}, \text{abox}(s_1), \gamma\theta)$  (notice that  $\text{abox}(s_1) = s_1$ ) and  $s'_1$  is consistent wrt  $\mathcal{T}$ . Now let us consider  $s'_2 = \text{DO}_{\text{SMP}}(\mathcal{T}, \text{abox}(s_2), \gamma\theta)$ . Since  $(\mathcal{T}, s_1)$  is logically equivalent to  $(\mathcal{T}, s_2)$ , then  $(\mathcal{T}, s'_1)$  is logically equivalent to  $(\mathcal{T}, s'_2)$ . Therefore,  $s'_2$  is consistent wrt  $\mathcal{T}$ , and hence  $s_2 \Rightarrow_{\text{SMP}} s'_2$ , and  $(s'_1, s'_2) \in \mathcal{B}$ .
3. Similarly, if  $s_2 \Rightarrow_{\text{SMP}} s'_2$  then there exists an action  $\gamma$  and a substitution  $\theta$  such that  $s'_2 = \text{DO}_{\text{SMP}}(\mathcal{T}, \text{abox}(s_2), \gamma\theta)$  and  $s'_2$  is consistent wrt  $\mathcal{T}$ . Now let us consider  $s'_1 = \text{DO}(\mathcal{T}, \text{abox}(s_1), \gamma\theta)$ . Since  $(\mathcal{T}, s_2)$  is logically equivalent to  $(\mathcal{T}, s_1)$ , then  $(\mathcal{T}, s'_2)$  is logically equivalent to  $(\mathcal{T}, s'_1)$ . Therefore,  $s'_1$  is consistent wrt  $\mathcal{T}$ , and hence  $s_1 \Rightarrow s'_1$ , and  $(s'_1, s'_2) \in \mathcal{B}$ .  $\square$

## B.2 Positive Dominant

Given a KAB  $\mathcal{K}$  and its simplified form  $\mathcal{K}_{\text{SMP}} = (\mathcal{T}, A_0, \Gamma, \Pi)$ , we define the *positive dominant* of  $\mathcal{K}$  as the KAB

$$\mathcal{K}^+ = (\emptyset, A_0^-, \{\gamma^+\}, \{\text{true} \mapsto \gamma^+\}).$$

(Recall that  $A_0^- = A_0 \setminus \text{EQ}_{\mathcal{T}}(A_0)$ .) The only action  $\gamma^+$  is constituted by *CopyAll* and one effect of the form

$$[q_i^{++}] \rightsquigarrow A''_i$$

for each effect  $[q_i^{++}] \wedge q_i^- \wedge Q_i^- \rightsquigarrow A_i'$  in every action of  $\Gamma$ , where  $A_i'$  is  $A_i'$  without the equality assertions. Moreover, the parameters of the actions in  $\Gamma$  become free variables in  $\gamma^+$ .

Notice that  $\gamma^+$  is applicable at every step because the process always allows it. The resulting state is always consistent, since  $\mathcal{K}^+$  has an empty TBox. Moreover, no equality assertion is ever generated. Hence, when constructing  $\Upsilon$ , there is no difference between using  $\text{DO}()$  and  $\text{DO}_{\text{SMP}}()$ . Observe that the resulting transition system  $\Upsilon$  is constituted by a single run, which incrementally accumulates all the facts that can be derived by the iterated application of  $\gamma^+$  over such increasing ABox. This behavior closely resembles the chase of TGDs in data exchange, where an application of  $\gamma^+$  corresponds to a “parallel” chase step.

We first relate the dependency graph of  $\mathcal{K}^+$  to the one of  $\mathcal{K}$ , showing that it preserves weak acyclicity.

**Lemma B.3.** *Given a KAB  $\mathcal{K}$ , if  $\mathcal{K}$  is weakly acyclic then also its positive dominant  $\mathcal{K}^+$  is weakly acyclic.*

*Proof.* Consider each effect specification  $[q^+] \wedge Q^- \rightsquigarrow A'$  belonging to an action in  $\mathcal{K}$ . The contribution of such an effect specification to the dependency graph  $\mathcal{G}$  of  $\mathcal{K}$  is limited to each CQ  $q_i$  in the UCQ  $\text{rew}(q^+)$ , and to the set of concept and role assertions of  $A_i'$ , denoted  $A_i''$ . We observe that each such  $q_i$  corresponds to a query  $q_i^{++}$  in  $\mathcal{K}_{\text{SMP}}$ , where each variable of  $q_i$  occurs exactly once.

For every free variable  $x$  of  $q_i$  that also appears in  $A_i''$ , and for every occurrence of  $x$  in  $q_i$  itself, an edge is included in  $\mathcal{G}$ . In the dependency graph  $\mathcal{G}_{\text{SMP}}$  of  $\mathcal{K}_{\text{SMP}}$ , only one of such edges appears, corresponding to the single occurrence of the variable  $x$  in  $q_i^{++}$ . Therefore,  $\mathcal{G}_{\text{SMP}}$  is a subgraph of  $\mathcal{G}$ , and hence weak acyclicity of  $\mathcal{G}$  implies weak acyclicity of  $\mathcal{G}_{\text{SMP}}$ .

Finally, the claim follows from the fact that, by construction, the dependency graph  $\mathcal{G}^+$  of  $\mathcal{K}^+$  is equal to  $\mathcal{G}_{\text{SMP}}$ . In fact, both  $q_i^{++}$  and its connection with  $A_i''$  are preserved by  $\mathcal{K}^+$ .  $\square$

We can now prove the following fundamental result related to weak acyclicity.

**Lemma B.4.** *Given a KAB  $\mathcal{K} = (\mathcal{T}, A_0, \Gamma, \Pi)$ , if  $\mathcal{K}$  is weakly acyclic, then there exists a polynomial  $P(\cdot)$  such that  $|\text{ADOM}(\Upsilon)| < P(|\text{ADOM}(A_0)|)$ , where  $\Upsilon$  is the transition system of the positive dominant  $\mathcal{K}^+$  of  $\mathcal{K}$ .*

*Proof.* We observe that there exists a strict connection between the execution of  $\mathcal{K}^+$  and the chase of a set of TGDs in data exchange. Therefore, the proof closely resembles the one of [36] – Theorem 3.9, where it is shown that for weakly acyclic TGDs, every chase sequence is bounded.

Let  $\Upsilon = (\mathbb{U}, \emptyset, \Sigma, A_0^-, \text{abox}, \Rightarrow)$ , let  $\mathcal{G} = (N, E)$  be the dependency graph of  $\mathcal{K}^+$ , and let  $n = |\text{ADOM}(A_0^-)|$ . For every node  $p \in N$ , we consider an incoming path to be any (finite or infinite) path ending in  $p$ . Let  $\text{rank}(p)$  be the maximum number of special edges on any such incoming path. Since  $\mathcal{K}$  is weakly acyclic by hypothesis,  $\mathcal{G}$  does not contain cycles going through special edges, and therefore  $\text{rank}(p)$  is finite. Let  $r$  be the maximum among  $\text{rank}(p_i)$  over all nodes. We observe that  $r \leq |N|$ ; indeed no path can lead to the same node twice using special edges, otherwise  $\mathcal{G}$  would contain a cycle going through special edges, thus breaking the weak acyclicity hypothesis. Notice also that  $|N|$  is a constant value, because it is obtained from  $\mathcal{T}$ , which is fixed. The next observation is that we can partition the nodes in  $N$  according to their rank, obtaining a set of sets  $\{N_0, N_1, \dots, N_r\}$ , where  $N_i$  is the set of all nodes with rank  $i$ .

Let us now consider a state  $s$  obtained from  $A_0^-$  by applying the only action  $\gamma^+$  contained in  $\mathcal{K}^+$  an arbitrary number of times. We now prove, by induction on  $i$ , the following claim: for every  $i$  there exists a polynomial  $P_i$  such that the total number of distinct values  $c$  that occur in  $s$  at positions in  $N_i$  is at most  $P_i(n)$ .

**(Base case)** Consider  $p \in N_0$ . By definition,  $p$  has no incoming path containing special edges. Therefore, no new values are stored in  $p$  along the run  $A_0^- \Rightarrow \dots \Rightarrow s$ :  $p$  can just store values that are part of the initial ABox  $A_0^-$ . This holds for all nodes in  $N_0$ , and hence we can fix  $P_0(n) = n$ .

**(Inductive step)** Consider  $p \in N_i$ , with  $i \in \{1, \dots, r\}$ . The first kind of values that may be stored inside  $p$  are those values that were stored inside  $p$  itself in  $A_0^-$ . The number of such values

is at most  $n$ . In addition, a value may be stored in  $p$  for two reasons: either it is copied from some other position  $p' \in N_j$  with  $i \neq j$ , or it is generated as a possibly new Skolem term, built when applying effects that contain a Skolem function in their head.

We first determine the number of fresh individuals that can be generated from Skolem terms. The possibility of generating and storing a new value in  $p$  as a result of an action is reflected by the presence of special edges. By definition, any special edge entering  $p$  must start from a node  $p' \in N_0 \cup \dots \cup N_{i-1}$ . By induction hypothesis, the number of distinct values that can exist in  $p'$  is bounded by  $H(n) = \sum_{j \in \{0, \dots, i-1\}} P_j(n)$ . Let  $b_a$  be the maximum number of special edges that enter a position, over all positions in the TBox;  $b_a$  bounds the arity taken by each Skolem term contained in  $\gamma$ . Then for every choice of  $b_a$  values in  $N_0 \cup \dots \cup N_{i-1}$  (one for each special edge that can enter a position), the number of new values generated at position  $p$  is bounded by  $t_f \cdot H(n)^{b_a}$ , where  $t_f$  is the total number of facts contained in all effects of  $\gamma^+$ . Note that this number does not depend on the data in  $A_0^-$ . By considering all positions in  $N_i$ , the total number of values that can be generated is then bounded by  $G(n) = |N_i| \cdot t_f \cdot H(n)^{b_a}$ . Clearly,  $G(\cdot)$  is a polynomial, because  $t_f$  and  $b_a$  are determined by  $\gamma^+$ .

We count next the number of distinct values that can be copied to positions of  $N_i$  from positions of  $N_j$ , with  $j \neq i$ . A copy is represented in the graph as a normal edge going from a node in  $N_j$  to a node in  $N_i$ , with  $j \neq i$ . We observe first that such normal edges can start only from nodes in  $N_0 \cup \dots \cup N_{i-1}$ , that is, they cannot start from nodes in  $N_j$  with  $j > i$ . We prove this by contradiction. Assume that there exists  $p' \rightarrow p \in E$ , such that  $p \in N_i$  and  $p' \in N_j$  with  $j > i$ . In this case, the rank of  $p$  would be  $j > i$ , which contradicts the fact that  $p \in N_i$ . As a consequence, the number of distinct values that can be copied to positions in  $N_i$  is bounded by the total number of values in  $N_0 \cup \dots \cup N_{i-1}$ , which corresponds to  $H(n)$  from our previous consideration.

Putting it all together, we define  $P_i(n) = n + G(n) + H(n)$ . Since  $P_i(\cdot)$  is a polynomial, the claim is proven.

Notice that, in the above claim,  $i$  is bounded by  $r$ , which is a constant. Hence, there exists a fixed polynomial  $P(\cdot)$  such that the number of distinct values that can exist in every state  $s \in \Sigma$  is bounded by  $P(n)$ .  $\mathcal{K}^+$  is inflationary, because when  $\gamma^+$  is applied it copies all concept and role assertions from the current to the next state. Since  $\Upsilon$  contains only a single run,  $P(n)$  is a bound for  $\text{ADOM}(\Upsilon)$  as well.  $\square$

**Lemma B.5.** *Let  $\mathcal{K}$  be a KAB. For its simplified version  $\mathcal{K}_{\text{SMP}}$  and its positive dominant  $\mathcal{K}^+$ , the following holds:*

$$\text{ADOM}(\Upsilon) \subseteq \text{ADOM}(\Upsilon)$$

*Proof.* Let  $\mathcal{K} = (\mathcal{T}, \mathcal{A}_0, \Gamma, \Pi)$  and  $\mathcal{K}^+ = (\emptyset, \mathcal{A}_0^-, \{\gamma^+\},)$

We first observe that, for every ABox  $\mathcal{A}$  in  $\Upsilon$ ,  $\text{ADOM}(\mathcal{A}) = \text{ADOM}(\mathcal{A}^-)$ . This can be shown by induction on the construction of  $\Upsilon$  using  $\text{APPLY}_{\text{SMP}}()$ , which instantiates the facts in the new ABox  $\mathcal{A}'$  by using individuals in the answer of queries over  $\mathcal{A}^-$  suitably filtered.

We show by induction on the construction of  $\Upsilon(\mathbb{U}, \mathcal{T}, \Sigma_1, A_0, \text{abox}, \Rightarrow_1)$  and  $\Upsilon = (\mathbb{U}, \emptyset, \Sigma_2, A_0^-, \text{abox}, \Rightarrow_2)$ , that for each state  $\mathcal{A}_1 \in \Sigma_1$  we have that there exists a state  $\mathcal{A}_2 \in \Sigma_2$  such that  $\mathcal{A}_1^- \subseteq \mathcal{A}_2$ . The base case holds for the initial states  $A_0$  and  $A_0^-$ . For the inductive case, we have to show that, given  $\mathcal{A}_1 \in \Sigma_1$  and  $\mathcal{A}_2 \in \Sigma_2$  with  $\mathcal{A}_1^- \subseteq \mathcal{A}_2$ , and for each  $\mathcal{A}'_1 \in \Sigma_1$  with  $\mathcal{A}_1 \Rightarrow_1 \mathcal{A}'_1$ , the unique state  $\mathcal{A}'_2 \in \Sigma_2$  with  $\mathcal{A}_2 \Rightarrow_2 \mathcal{A}'_2$  is such that  $\mathcal{A}'_1^- \subseteq \mathcal{A}'_2$ . To show this, note that  $\mathcal{A}_1 \Rightarrow_1 \mathcal{A}'_1$  if there exists an action  $\gamma$  of  $\mathcal{K}_{\text{SMP}}$  and a substitution  $\theta$  for the parameters of  $\gamma$  such that  $\mathcal{A}'_1 = \text{DO}_{\text{SMP}}(\mathcal{T}, \mathcal{A}_1, \gamma\theta)$ . Similarly, taking into account that  $\gamma^+$  has no parameters and is always executable in  $\Upsilon$ , we have that  $\mathcal{A}'_2 = \text{DO}_{\text{SMP}}(\mathcal{T}, \mathcal{A}_2, \gamma^+)$ . By construction of  $\mathcal{K}^+$ , for each effect  $e_1 \in \gamma$  of the form

$$e_1 : [q^{++}] \wedge q^- \wedge Q^- \rightsquigarrow F'_{e_1},$$

there is an effect  $e_2 \in \gamma^+$  of the form

$$e_2 : [q^{++}] \rightsquigarrow F'_{e_2},$$

where  $F'_{e_2}$  is obtained from  $F'_{e_1}$  by removing all equality assertions. By induction hypothesis, we have that  $\mathcal{A}_1^- \subseteq \mathcal{A}_2$ . Hence, by observing that  $\text{ANS}([q^{++}]\theta, \emptyset, \mathcal{A}_1^-) \cap \text{ANS}((q^= \wedge Q^-)\theta, \emptyset, \mathcal{A}_1) \subseteq \text{ANS}([q^{++}], \emptyset, \mathcal{A}_2)$ , we obtain that  $A'_{e_1} \setminus \text{EQ}_{\mathcal{T}}(A'_{e_1}) \subseteq A'_{e_2}$ , where  $A'_{e_1} = \text{APPLY}_{\text{SMP}}(\mathcal{T}, e_1, \theta, \mathcal{A}_1)$  and  $A'_{e_2} = \text{APPLY}(\emptyset, e_2, \emptyset, \mathcal{A}_2)$ . Hence, we get the claim that  $\mathcal{A}_1^- \subseteq \mathcal{A}'_2$ .  $\square$

We now combine the above results in order to study the impact of weak acyclicity in terms of the verification of KABs. In particular, we exploit the notion of positive dominant to show that verification of  $\mu\mathcal{L}_A$  properties for weakly acyclic KABs is decidable.

Given a weakly acyclic KAB  $\mathcal{K}$ , by Theorem 13.1 and Lemma B.2 we know that the simplified and original transition systems  $\Upsilon$  and  $\Upsilon$  of  $\mathcal{K}$  satisfy the same  $\mu\mathcal{L}_A$  properties. Moreover, by Lemma B.3, we know that also  $\mathcal{K}^+$  is weakly acyclic, and therefore Lemma B.4 guarantees that  $\text{ADOM}(\Upsilon)$  is polynomially bounded in the size of its initial state. By Lemma B.5 this property holds also for  $\Upsilon$ . Considering that the number of atomic concept and role symbols in  $\mathcal{K}_{\text{SMP}}$  is finite this in turn implies that  $\Upsilon$  is finite state. Hence, verification of  $\mu\mathcal{L}_A$  properties over  $\Upsilon$  reduces to standard finite-state model checking, which is known to be decidable [30].  $\square$