

SMT-based Verification of Data-Aware Processes: a Model-Theoretic Approach

Diego Calvanese¹, Silvio Ghilardi², Alessandro Gianola¹, Marco Montali¹ and Andrey Rivkin¹

¹*Free University of Bozen-Bolzano*

`{calvanese, gianola, montali, rivkin}@inf.unibz.it`

²*Università degli Studi di Milano*

`silvio.ghilardi@unimi.it`

Received November 2019

In recent times, Satisfiability-Modulo-Theories (SMT) techniques gained increasing attention and obtained remarkable success in model-checking infinite state systems. Still, we believe that whenever more expressivity is needed in order to specify the systems to be verified, more and more support is needed from mathematical logic and model-theory. This is the case of the applications considered in this paper: we study verification over a general model of relational, data-aware processes, to assess (parameterized) safety properties irrespectively of the initial database instance. Towards this goal, we take inspiration from array-based systems, and tackle safety algorithmically via backward reachability. To enable the adoption of this technique in our rich setting, we make use of the model-theoretic machinery of model completion, which surprisingly turns out to be an effective tool for verification of relational systems, and represents the main original contribution of this paper. In this way, we pursue a twofold purpose. On the one hand, we isolate three notable classes for which backward reachability terminates, in turn witnessing decidability. Two of such classes relate our approach to conditions singled out in the literature, whereas the third one is genuinely novel. On the other hand, we are able to exploit SMT technology in implementations, building on the well-known MCMT model checker for array-based systems, and extending it to make all our foundational results fully operational. All in all, the present contribution is deeply rooted in the long-standing tradition of the application of model theory in computer science. In particular, this paper applies these ideas in an original mathematical context and shows how these techniques can be used for the first time to empower algorithmic techniques for the verification of infinite-state systems based on arrays, so as to make such techniques applicable to the timely, challenging settings of data-aware processes.

1. Introduction

The main contribution of this paper comes from a rather surprising confluence of two well-established research traditions: *model-theoretic algebra* from mathematical logic and *satisfiability modulo theories* (SMT), an emerging technologically oriented area in computational logic. We believe that such seemingly very different scientific paradigms can

indeed cooperate in formal verification and we shall supply an evidence for this claim by developing an innovative application to the hot topic of the management of dynamic *data-aware processes*. In this introduction, we shall briefly explain how the above mentioned ingredients fit into the plan of our paper.

1.1. Model-Theoretic Algebra

Finding solutions to equations is a challenge at the heart of both mathematics and computer science. Model-theoretic algebra, originating with the ground-breaking work of (Robinson, 1951; Robinson, 1963), cast the problem of solving equations in a logical form, and used this setting to solve algebraic problems via model theory. The central notion is that of an *existentially closed model*, which we explain now. Call a quantifier-free formula with parameters in a model M *solvable* if there is an extension M' of M where the formula is satisfied. A model M is *existentially closed* if any solvable quantifier-free formula already has a solution in M itself. For example, the field of real numbers is not existentially closed, but the field of complex numbers is.

Although this definition is formally clear, it has a main drawback: it is not a first-order notion in general. However, in fortunate and important cases, the class of existentially closed models of T are exactly the models of another first-order theory T^* . In this case, the theory T^* can be characterized abstractly as the *model companion* of T . Model companions become *model completions* (cf. Definition 2.2) in the case of universal theories with the amalgamation property; in such model completions, quantifier elimination holds, unlike in the original theory T . The model companion/model completion of a theory identifies the class of those models where *all satisfiable existential statements can be satisfied*. For example, the theory of algebraically closed fields is the model companion of the theory of fields, and dense linear orders without endpoints give the model companion of linear orders.

In declarative approaches to model-checking, the runs of a system are identified with certain definable paths in the models of a theory T : we shall show that, without loss of generality, one may *restrict to paths within existentially closed models*, thus taking profit from the properties enjoyed by the model completion T^* (quantifier elimination being the key property to be carefully exploited).

Model completeness has other well-known applications in computer science. It has been applied to discover interesting connections between temporal logic and monadic second order logic (Ghilardi and van Gool, 2016; Ghilardi and van Gool, 2017). In automated reasoning, it has been used to design complete algorithms for constraint satisfiability in combined theories over non disjoint signatures (Ghilardi, 2004; Baader et al., 2006; Ghilardi et al., 2008b; Nicolini et al., 2010; Nicolini et al., 2009a; Nicolini et al., 2009b) and in theory extensions (Sofronie-Stokkermans, 2016; Sofronie-Stokkermans, 2018). Applications to combined interpolation (both for modal logics and for software verification theories) can be found in (Ghilardi and Gianola, 2017; Ghilardi and Gianola, 2018).

1.2. Satisfiability Modulo Theories

The SMT-LIB project <http://smtlib.cs.uiowa.edu/> (started in 2003) aims at bringing together people interested in developing powerful tools combining sophisticated techniques in SAT-solving with dedicated decision procedures involving specific theories used in applications (especially in software verification). SMT-tools are at the heart of declarative approaches to model checking, both in the bounded and in the unbounded case: they are employed in many advanced techniques, for instance in interpolation based (McMillan, 2006) and IC3 based (Hoder and Bjørner, 2012) techniques.

Specifically, our approach is grounded in *array-based systems*. Array-based systems are a declarative formalism originally introduced in (Ghilardi et al., 2008a; Ghilardi and Ranise, 2010a) to handle the verification of distributed systems, and afterwards successfully employed also to attack the static analysis of other types of systems (Alberti et al., 2017; Alberti et al., 2014a). Distributed systems are parameterized in their essence: the number N of interacting processes within a distributed system is unbounded, and the challenge is that of supplying certifications that are valid for all possible values of the parameter N . The overall state of the system is typically described by means of arrays indexed by process identifiers, and used to store the content of process variables like locations and clocks. These arrays are genuine *second order* variables. In addition, *first-order quantifiers* are used to represent sets of system states. Quantified formulae and second order function variables are at the heart of the model checking methodologies developed in (Ghilardi et al., 2008a; Ghilardi and Ranise, 2010a) and following papers.

It is worth noting that the term “array-based systems” is an umbrella term generically referring to transition systems specified with logical formulae having second-order variables (i.e., arrays). The precise formal notion depends on the application and is defined on the spot (in this paper we shall introduce a specific instance of the notion of an array-based system tailored to database-driven verification).

Model checkers for array-based systems handle safety problems by *backward reachability*: they iteratively *regress* bad states by computing their predecessors, the predecessors of the predecessors, etc., until a fixpoint is reached or the initial state(s) are intersected. This is done symbolically by manipulating logical formulae that describe sets of states. Depending on the specific features of the array-based system, to guarantee the regressability of such formulae the procedure may require to eliminate existentially quantified variables present in the formula.

There is an historical reason for choosing backward reachability: it was known since the seminal paper (Abdulla et al., 1996) that backward search decides safety problems for a large class of systems, called *well-structured transition systems*. What backward search in array-based systems is meant to achieve is precisely to reproduce the results of (Abdulla et al., 1996) in a declarative setting: in such a declarative setting, the abstract wqo underlying well-structured transitions systems is replaced by the standard model-theoretic notion of an *embedding* between finitely generated models (in many practical cases, in fact, such embeddability relation can be proved to be a wqo, using a suitable version of Dickson or of Higman lemma).

Backward search, once done in a declarative symbolic setting, requires discharging

proof obligations that can be reduced to satisfiability tests for quantified formulae, albeit of a restricted syntactic shape. This raises the question of how to handle such (first-order) quantifiers. In the original papers (Ghilardi et al., 2008a; Ghilardi and Ranise, 2010a) first order quantifiers were handled in satisfiability tests by *instantiation*, whereas in successive applications (Carioni et al., 2010; Alberti et al., 2014a) *quantifier elimination* was also used to handle quantifiers ranging over data structures (typically, real valued clocks). There quantifier elimination was made possible by the fact that theories axiomatizing such data structures were limited to light versions of arithmetics (mostly even strictly included in what is called ‘difference logic’ in SMT terminology), where quantifier elimination is indeed available and at least in the examples arising from benchmarks seems not to be as harmful as in the general arithmetic case. Suitable combinations of quantifier instantiations and quantifier eliminations are needed at the foundational level to design complete algorithms for the satisfiability tests that a model checker for array-based systems has to discharge during search: a specific form of such combination will be developed in this paper too. By means of such combinations, *satisfiability tests involving quantified formulae of special shape are reduced to satisfiability tests at quantifier-free level*, to be very efficiently discharged by existing SMT solvers (as confirmed by the extensive experiments, see Section 6 for a brief historical account).

1.3. Data-Aware Processes

To capture data-aware processes, we follow the traditional line of research focused on the formal representation of *artifact systems* (Hull, 2008; Vianu, 2009; Deutsch et al., 2009; Damaggio et al., 2012; Calvanese et al., 2013; Deutsch et al., 2018). Since their initial versions (e.g., in (Deutsch et al., 2009)), such systems are traditionally formalized using three components: (i) a *read-only database (DB)*, storing background information that does *not* change during the system evolution; (ii) an *artifact working memory*, storing data and lifecycle information about artifact(s) that *does* change during the system evolution; (iii) *actions* (also called *services*) that access the read-only database and the working memory, and determine how the working memory itself has to be updated.

Different variants of this framework have been considered towards decidability of verification, by carefully tuning the expressive power of the three components. For instance, for the working memory, radically different models are obtained depending on whether only a single artifact instance is evolved, or whether instead the co-evolution of multiple instances of possibly different artifacts is supported. In particular, early formal models for artifact systems merely considered a fixed set of so-called *artifact variables*, altogether instantiated into a single tuple of data. This, in turn, allows one to capture the evolution of a single artifact instance (Deutsch et al., 2009). We call artifact systems of this form *Simple Artifact System (SAS)*. Instead, more sophisticated types of artifact systems have been studied recently in (Deutsch et al., 2016; Li et al., 2017; Deutsch et al., 2019). Here, the working memory is not only equipped with artifact variables as in SAS, but also with so-called *artifact relations*, which supports storing arbitrarily many tuples, each accounting for a different artifact instance that can be separately evolved on its own. We call artifact systems of this form *Relational Artifact System (RAS)*.

Actions are usually specified using logical formulae relating the content of the read-only DB as well as the current configuration of the working memory to (possibly different) next configurations. An applicable action may be executed, nondeterministically transforming the current configuration of working memory in one of such next configurations.

1.4. Bringing all the Ingredients Together

RASs naturally fit the paradigm of array-based systems: the read-only database is axiomatized by a suitable universal first-order theory T and the artifact variables and relations are modeled by second-order variables. The identifiers of the tuples (i.e., the “entries”) of the artifact relations play the role of the identifiers of the processes in distributed systems: formally, in both cases, they are just sorted first order variables whose sort is the domain sort of a second order function variable.

The resulting framework, however, requires novel and non-trivial extensions of the array-based technology to make it operational. In fact, as we saw, quantifiers are handled in array-based systems both by quantifier instantiation and by quantifier elimination. Quantifier instantiation (ultimately referring to variants of the Herbrand Theorem) can be transposed to the new framework, whereas quantifier elimination becomes problematic. In fact, quantifier elimination should be applied to data variables, which do not simply range over data types (like integers, reals, or enumerated sets) as in standard array-based systems, but instead point to the content of a whole, full-fledged (read-only) relational database and there is no reason for the theory T axiomatizing it to enjoy quantifier elimination. Here model-theoretic algebra comes into the picture: we show that, without loss of generality, we can assume that system runs take place in existentially closed structures, so that we can exploit quantifier elimination, *provided T has a model completion*.

The question on whether T admits or not a model completion is related to the way we represent the read-only database. Model completions exist in case the read-only database is represented in the most simple way, as consisting on free n -ary relations, not subject to any kind of constraint. However, applications require the introduction of some minimal integrity constraints, like *primary* and *foreign* keys. A naif declarative modeling of such requirements (for instance via relations which are partially functional) would destroy amalgamation property, thus compromising the existence of model completions. Instead, we propose “functional view” of relations, where the read-only database and the artifact relations forming the working memory are represented with *sorted unary function symbols*. We introduce formally the above framework in Section 3 (Definitions 3.1 and 3.2): there we supply a detailed example (Example 3.1 and Figure 1) and we also show how to recover the traditional relational model (Subsection 3.1).¹

¹ We underline the fact that free n -ary relations (modeling plain relations without integrity constraints) can be added to our sorted unary functions framework without compromising our model-checking techniques to apply. We shall briefly mention such trivial extensions in the paper (see e.g. Definition 3.3).

1.5. Main Contributions

By exploiting the above explained machinery and its model-theoretic properties, we then provide a fourfold contribution.

Our first contribution is to define a general framework for SASs and RASs, in which artifacts are formalized in the spirit of *array-based systems*, one of the most sophisticated setting within the SMT tradition. In this setting, SASs are a particular class of RASs, where only artifact variables are allowed. RASs employ arrays to capture a very rich working memory that simultaneously accounts for artifact variables storing single data elements, and full-fledged artifact relations storing unboundedly many tuples. Each artifact relation is captured using a collection of arrays, so that a tuple in the relation can be retrieved by inspecting the content of the arrays with a given index. The elements stored therein may be fresh values injected into the RAS, or data elements extracted from the read-only DB, whose relations are subject to key and foreign key constraints.

To attack this modeling complexity within array-based systems, RASs encode the read-only DB using a functional, algebraic view, where relations and constraints are captured using multiple sorts and unary functions. To the best of our knowledge, this encoding has never been explored in the past, but is essential in our context. In fact, more direct attempts to encode the read-only DB into standard array-based systems would fail; for example, not using unary functions for relations with key dependencies would destroy amalgamability, which is a model-theoretic notion that is crucial towards decidability of verification (Bojańczyk et al., 2013).

Our resulting RAS model captures the essential features of (Li et al., 2017), which in turn is tightly related (though incomparable) to one of the most sophisticated formal models for artifact-centric systems of (Deutsch et al., 2016; Deutsch et al., 2019).

Our second contribution is the development of a new version of the backward reachability algorithm employed in traditional array-based systems (Ghilardi et al., 2008a; Ghilardi and Ranise, 2010a), making it able to assess safety of RASs (and consequently SASs) in a sound and complete way. The main technical difficulty, which makes the original algorithm not applicable anymore, is that transition formulae in RASs contain special existentially quantified “*data*” variables pointing to the read-only DB, which contains data elements possibly constrained by primary keys and foreign keys. Such data variables are central in our approach as they are needed:

- from the modeling point of view, to equip array-based systems with the ability of querying the read-only DB;
- again for modeling reasons, to express nondeterministic inputs from the external environment, such as users (a feature that is customary in business processes);
- to encode typical forms of updates employed in the artifact-centric literature (Deutsch et al., 2016; Li et al., 2017).

The presence of these quantified data variables constitutes a big leap from traditional array-based systems. In previous works on array-based systems, existentially quantified variables to be eliminated were just arithmetic variables, and the corresponding quantifier elimination procedures were consequently the standard ones studied in the context of arithmetics (such as Fourier-Motzkin and Presburger). Due to the peculiar nature of

data variables pointing to the read-only DB, and in particular of the constraints they must satisfy, such standard techniques do not carry over. Hence, genuinely novel research is needed in order to eliminate new existentially quantified data variables that are introduced during the computation of predecessors in the backward reachability procedure.

From a theoretical point of view, we solve this problem by introducing a dedicated machinery based on model completions. While in the case of arithmetic variables the corresponding theories admit themselves quantifier elimination, this is not the case anymore for our data variables. However, we show that quantifier elimination for data variables that is available in the model completion of their theory can actually be safely employed in the context of backward reachability, retaining soundness and completeness when checking safety of RASs. This requires to *significantly modify* the original procedure and the original proofs.

In the general case, backward reachability is not guaranteed to terminate when checking safety of SASs and RASs. As a third contribution, we consequently isolate three notable classes of RASs for which backward reachability is guaranteed to terminate, in turn witnessing decidability of safety. The first class restricts the working memory to variables only, i.e., focuses on SAS. The second class focuses on RAS operating under the restrictions imposed in (Li et al., 2017): it requires acyclicity of foreign keys and requires a sort of locality principle in the action guards, ensuring that different artifact tuples are not compared. Consequently, it reconstructs in our setting the essence of the decidability result exploited in (Li et al., 2017) if one restricts the verification logic used there to safety properties only. In addition, our second class supports full-fledged bulk updates, which greatly increase the expressive power of dynamic systems (Schmitz and Schnoebelen, 2013) and, in our setting, witness the incomparability of our results and those in (Li et al., 2017). The third class is genuinely novel, and while it further restricts foreign keys to form a tree-shaped structure, it does not impose any restriction on the shape of updates, consequently supporting not only bulk updates, but also comparisons between artifact tuples. To prove termination of backward reachability for this class, we resort to techniques based on well-quasi orders (in particular, a non-trivial application of Kruskal’s Tree Theorem (Kruskal, 1960)).

Our fourth and last contribution is to implement the new version of backward reachability required to handle the verification of RASs. We do so by extending the well-known MCMT model checker for array-based systems (Ghilardi and Ranise, 2010b). The resulting version of MCMT (version 2.8) provides a fully operational counterpart to all the foundational results presented in the paper. Even though implementation and experimental evaluation are not the central goal of this paper, we also note that our model checker correctly handles the examples produced to test VERIFAS (Li et al., 2017), as well as additional examples that go beyond the verification capabilities of VERIFAS, and report some interesting case here. The performance of MCMT to conduct verification of these examples is very encouraging, and indeed provides the first stepping stone towards effective, SMT-based verification techniques for artifact-centric systems.

All in all, the present contribution is deeply rooted in the long-standing tradition of the application of model theory in computer science, as witnessed by notable approaches like the one in (Ghilardi, 2004; Baader et al., 2006; Ghilardi et al., 2008b; Ghilardi and van

Gool, 2017; Nicolini et al., 2010; Nicolini et al., 2009a; Nicolini et al., 2009b; Sofronie-Stokkermans, 2008; Sofronie-Stokkermans, 2016; Sofronie-Stokkermans, 2018; Ghilardi and Gianola, 2017; Ghilardi and Gianola, 2018). In particular, this paper applies these ideas in a genuinely novel mathematical context and shows how these techniques can be used for the first time to empower algorithmic techniques for the verification of infinite-state systems based on arrays in the style of (Ghilardi et al., 2008a; Ghilardi and Ranise, 2010a; Alberti et al., 2014a; Alberti et al., 2014b; Alberti et al., 2017; Ghilardi and Ranise, 2010b; Conchon et al., 2012; Conchon et al., 2015; Conchon et al., 2018a; Delzanno, 2018; Conchon et al., 2018b; Cimatti et al., 2018), so as to make such techniques applicable to the timely, challenging settings of data-aware processes (Calvanese et al., 2019d). For an explicit linking between the use of model completeness in computer science and our application to verification, see in particular the survey (Calvanese et al., 2019b).

1.6. Plan of the paper and prerequisites

There are no specific prerequisites in model checking or in database theory to read this paper; we only assume some familiarity with some basic model-theory (like that supplied in a one-semester course); some elementary notions will be revised mainly to fix notations. All definitions we introduce concerning read-only databases and relational artifact systems are explained with the help of a running example; however, we did not give the straightforward (and rather boring) details showing how basic operations on artifact relations (like insertion/deletions, resetting, etc.) can be modeled in our systems: there is a preliminary but more elementary exposition of the content of this paper available from the network (Calvanese et al., 2018b) which covers these low level details and some more model-theoretic prerequisites.

The paper is organized as follows. After fixing preliminary notions in Section 2, in Section 3 we introduce our functional view of databases and discuss the related model-theoretic properties. In Section 4 we introduce a comprehensive framework for Artifact-Centric Systems: Simple Artifact Systems are studied in Subsection 4.1, whereas the most general framework of Relational Artifact Systems is investigated in Subsection 4.2. Termination and decidability results are supplied in Section 5. Finally, first experiments are discussed in Section 6 and Section 7 concludes.

2. Preliminaries

We adopt the usual first-order syntactic notions of signature, term, atom, (ground) formula, and so on; our signatures are multi-sorted and include equality for every sort. This implies that variables are sorted as well. For simplicity, most basic definitions in this Section will be supplied for single-sorted languages only (the adaptation to multi-sorted languages is straightforward). We compactly represent a tuple $\langle x_1, \dots, x_n \rangle$ of variables as \underline{x} . The notation $t(\underline{x}), \phi(\underline{x})$ means that the term t , the formula ϕ has free variables included in the tuple \underline{x} .

We assume that a function arity can be deduced from the context. Whenever we build terms and formulae, we always assume that they are well-typed, in the sense that the

sorts of variables, constants, and function sources/targets match. A formula is said to be *universal* (resp., *existential*) if it has the form $\forall \underline{x}(\phi(\underline{x}))$ (resp., $\exists \underline{x}(\phi(\underline{x}))$), where ϕ is a quantifier-free formula. Formulae with no free variables are called *sentences*.

From the semantic side, we use the standard notion of a Σ -structure \mathcal{M} and of truth of a formula in a Σ -structure under a free variables assignment.

A Σ -theory T is a set of Σ -sentences; a *model* of T is a Σ -structure \mathcal{M} where all sentences in T are true. We use the standard notation $T \models \phi$ (' ϕ is a logical consequence of T ') to say that ϕ is true in all models of T for every assignment to the variables occurring free in ϕ . We say that ϕ is *T -satisfiable* iff there is a model \mathcal{M} of T and an assignment to the variables occurring free in ϕ making ϕ true in \mathcal{M} . Thus, according to this definition, ϕ is *T -satisfiable* iff its *existential closure* is true in a model of T (notice that this convention might not be uniform in the literature). A Σ -theory T is *complete* iff for every Σ -sentence φ , either φ or $\neg\varphi$ is a logical consequence of T .

A Σ -formula ϕ is a Σ -*constraint* (or just a constraint) iff it is a conjunction of literals. The constraint satisfiability problem for T is the following: we are given an existential formula $\exists \underline{y} \phi(\underline{x}, \underline{y})$ (with ϕ a constraint, but, for the purposes of this definition, we may equivalently take ϕ to be quantifier-free) and we are asking whether there exist a model \mathcal{M} of T and an assignment α to the free variables \underline{x} such that $\mathcal{M}, \alpha \models \exists \underline{y} \phi(\underline{x}, \underline{y})$.

A theory T has *quantifier elimination* iff for every formula $\phi(\underline{x})$ in the signature of T there is a quantifier-free formula $\phi'(\underline{x})$ such that $T \models \phi(\underline{x}) \leftrightarrow \phi'(\underline{x})$. It is well-known (and easily seen) that quantifier elimination holds in case we can eliminate quantifiers from *primitive* formulae, i.e. from formulae of the kind $\exists \underline{y} \phi(\underline{x}, \underline{y})$, where ϕ is a conjunction of literals (i.e. of atomic formulae and their negations). Since we are interested in effective computability, we assume that *whenever* we talk about quantifier elimination, an *effective procedure* for eliminating quantifiers is given.

Let Σ be a first-order signature. The signature obtained from Σ by adding to it a set \underline{a} of new constants (i.e., 0-ary function symbols) is denoted by $\Sigma^{\underline{a}}$. Analogously, given a Σ -structure \mathcal{A} , the signature Σ can be expanded to a new signature $\Sigma^{|\mathcal{A}|} := \Sigma \cup \{\bar{a} \mid a \in |\mathcal{A}|\}$ by adding a set of new constants \bar{a} (the *name* for a), one for each element $a \in |\mathcal{A}|$, with the convention that two distinct elements are denoted by different name constants (we use $|\mathcal{A}|$ to denote the support of the structure \mathcal{A}). \mathcal{A} can be expanded to a $\Sigma^{|\mathcal{A}|}$ -structure $\mathcal{A}' := (\mathcal{A}, a)_{a \in |\mathcal{A}|}$ just interpreting the additional constants over the corresponding elements. From now on, when the meaning is clear from the context, we will freely use the notation \mathcal{A} and \mathcal{A}' interchangeably: in particular, given a Σ -structure \mathcal{A} , a Σ -formula $\phi(\underline{x})$ and elements \underline{a} from $|\mathcal{A}|$, we will write, by abuse of notation, $\mathcal{A} \models \phi(\underline{a})$ instead of $\mathcal{A}' \models \phi(\underline{\bar{a}})$.

A Σ -*homomorphism* (or, simply, a homomorphism) between two Σ -structures \mathcal{A} and \mathcal{B} is any mapping $\mu : |\mathcal{A}| \rightarrow |\mathcal{B}|$ among the support sets $|\mathcal{A}|$ of \mathcal{A} and $|\mathcal{B}|$ of \mathcal{B} satisfying the condition

$$\mathcal{A} \models \varphi \quad \Rightarrow \quad \mathcal{B} \models \varphi \tag{1}$$

for all $\Sigma^{|\mathcal{A}|}$ -atoms φ (here - as above - \mathcal{A} is regarded as a $\Sigma^{|\mathcal{A}|}$ -structure, by interpreting each additional constant $a \in |\mathcal{A}|$ into itself and \mathcal{B} is regarded as a $\Sigma^{|\mathcal{A}|}$ -structure by interpreting each additional constant $a \in |\mathcal{A}|$ into $\mu(a)$). In case condition (1) holds for all $\Sigma^{|\mathcal{A}|}$ -literals, the homomorphism μ is said to be an *embedding* (denoted by $\mu : \mathcal{A} \hookrightarrow \mathcal{B}$)

and if it holds for all first order formulae, the embedding μ is said to be *elementary*. If $\mu : \mathcal{A} \rightarrow \mathcal{B}$ is an embedding which is just the identity inclusion $|\mathcal{A}| \subseteq |\mathcal{B}|$, we say that \mathcal{A} is a *substructure* of \mathcal{B} or that \mathcal{B} is an *extension* of \mathcal{A} . A Σ -structure \mathcal{A} is said to be *generated by* a set X included in its support $|\mathcal{A}|$ iff there are no proper substructures of \mathcal{A} including X .


Robinson Diagrams and Amalgamation

Let \mathcal{A} be a Σ -structure. The *diagram* of \mathcal{A} , denoted by $\Delta_\Sigma(\mathcal{A})$, is defined as the set of ground $\Sigma^{|\mathcal{A}|}$ -literals (i.e. atomic formulae and negations of atomic formulae) that are true in \mathcal{A} . For the sake of simplicity, once again by abuse of notation, we will freely say that $\Delta_\Sigma(\mathcal{A})$ is the set of $\Sigma^{|\mathcal{A}|}$ -literals which are true in \mathcal{A} .

An easy but nevertheless important basic result, called *Robinson Diagram Lemma* (Chang and Keisler, 1990), says that, given any Σ -structure \mathcal{B} , the embeddings $\mu : \mathcal{A} \rightarrow \mathcal{B}$ are in bijective correspondence with expansions of \mathcal{B} to $\Sigma^{|\mathcal{A}|}$ -structures which are models of $\Delta_\Sigma(\mathcal{A})$. The expansions and the embeddings are related in the obvious way: \bar{a} is interpreted as $\mu(a)$.

Amalgamation is a classical algebraic concept. We give the formal definition of this notion.

Definition 2.1 (Amalgamation). A theory T has the *amalgamation property* if for every couple of embeddings $\mu_1 : \mathcal{M}_0 \rightarrow \mathcal{M}_1$, $\mu_2 : \mathcal{M}_0 \rightarrow \mathcal{M}_2$ among models of T , there exists a model \mathcal{M} of T endowed with embeddings $\nu_1 : \mathcal{M}_1 \rightarrow \mathcal{M}$ and $\nu_2 : \mathcal{M}_2 \rightarrow \mathcal{M}$ such that $\nu_1 \circ \mu_1 = \nu_2 \circ \mu_2$



The triple $(\mathcal{M}, \mu_1, \mu_2)$ (or, by abuse, \mathcal{M} itself) is said to be a T -*amalgam* of $\mathcal{M}_1, \mathcal{M}_2$ over \mathcal{M}_0

Model Completions

We recall a standard notion in Model Theory, namely the notion of a *model completion* of a first order theory (Chang and Keisler, 1990) (we limit the definition to universal theories, because we shall use only this case):

Definition 2.2. Let T be a universal Σ -theory and let $T^* \supseteq T$ be a further Σ -theory; we say that T^* is a *model completion* of T iff: (i) every model of T can be embedded into a model of T^* ; (ii) for every model \mathcal{M} of T , we have that $T^* \cup \Delta_\Sigma(\mathcal{M})$ is a complete theory in the signature $\Sigma^{|\mathcal{M}|}$.

Since T is universal, condition (ii) is equivalent to the fact that T^* has *quantifier elimination*; on the other hand, a standard argument (based on diagrams and compactness) shows that condition (i) is the same as asking that T and T^* have the same universal consequences. Thus we have an equivalent definition (Ghilardi, 2004) (to be used in the following):

Proposition 2.1. Let T be a universal Σ -theory and let $T^* \supseteq T$ be a further Σ -theory; T^* is a model completion of T iff: (i) every Σ -constraint satisfiable in a model of T is also satisfiable in a model of T^* ; (ii) T^* has quantifier elimination.

As stated before, we assume that a model completion has an *effective procedure* for eliminating quantifiers. We recall also that the model completion T^* of a theory T is unique, if it exists (see (Chang and Keisler, 1990) for these results and for examples). It is well-known that a universal theory T which admits a model completion is also amalgamable (Chang and Keisler, 1990).

Example 2.1. The theory of undirected graphs admits a model completion (and, hence, is amalgamable); this is the theory T whose signature Σ contains only a binary relational symbol R , and whose axioms are specified as follows

$$T := \{\forall x \neg R(x, x), \forall x \forall y (R(x, y) \rightarrow R(y, x))\} .$$

Indeed, it is folklore that the model completion of T is the theory of the *Rado graph* (Rado, 1964): a Rado (also called *random*) graph is a countably infinite graph in which, given any non-empty sets $X = \{x_0, \dots, x_m\}$ and $Y = \{y_0, \dots, y_n\}$ of nodes with $X \cap Y = \emptyset$, there is a node z (with $z \notin X \cup Y$) such that there is an edge between z and all elements of X and there is no edge between z and any element of Y . This theory is first-order definable (Fagin, 1976).

Definable Extensions and λ -Notations

In the following, we specify transitions of artifact-centric systems using first-order formulae. To obtain a more compact representation, we make use of definable extensions as a means to introduce case-defined functions, abbreviating more complicated (still first-order) expressions. Let us fix a signature Σ and a Σ -theory T ; a T -partition is a finite set $\kappa_1(\underline{x}), \dots, \kappa_n(\underline{x})$ of quantifier-free formulae such that $T \models \forall \underline{x} \bigvee_{i=1}^n \kappa_i(\underline{x})$ and $T \models \bigwedge_{i \neq j} \forall \underline{x} \neg (\kappa_i(\underline{x}) \wedge \kappa_j(\underline{x}))$. Given such a T -partition $\kappa_1(\underline{x}), \dots, \kappa_n(\underline{x})$ together with Σ -terms $t_1(\underline{x}), \dots, t_n(\underline{x})$ (all of the same target sort), a *case-definable extension* is the Σ' -theory T' , where $\Sigma' = \Sigma \cup \{F\}$, with F a “fresh” function symbol (i.e., $F \notin \Sigma$), and $T' = T \cup \bigcup_{i=1}^n \{\forall \underline{x} (\kappa_i(\underline{x}) \rightarrow F(\underline{x}) = t_i(\underline{x}))\}$. Arity, source sorts, and target sort for F can be deduced from the context (considering that everything is well-typed).

Intuitively, F represents a case-defined function, which can be reformulated using nested if-then-else expressions and can be written as

$$F(\underline{x}) := \text{case of } \{\kappa_1(\underline{x}) : t_1; \dots; \kappa_n(\underline{x}) : t_n\}.$$

By abuse of notation, we shall identify T with any of its case-definable extensions T' . In

fact, it is easy to produce from a Σ' -formula ϕ' a Σ -formula ϕ that is equivalent to ϕ' in all models of T' : just remove (in the appropriate order) every occurrence $F(\underline{v})$ of the new symbol F in an atomic formula A , by replacing A with $\bigvee_{i=1}^n (\kappa_i(\underline{v}) \wedge A(t_i(\underline{v})))$.

We also exploit λ -abstractions (see, e.g., formula (13) below) for more “compact” representation of some complex expressions, and always use them in atoms like $b = \lambda y.F(y, \underline{z})$ as abbreviations of $\forall y. b(y) = F(y, \underline{z})$ (where, typically, F is a symbol introduced in a case-defined extension as above). Thus, also our formulae containing lambda abstractions, can be converted into plain first-order formulae.

3. Read-only Database Schemas

We now provide a formal definition of (read-only) DB-schemas by relying on an algebraic, functional characterization, and derive some key model-theoretic properties instrumental to the technical treatment.

Definition 3.1. A *DB schema* is a pair $\langle \Sigma, T \rangle$, where: (i) Σ is a *DB signature*, that is, a finite multi-sorted signature whose only symbols are equality, unary functions, and constants; (ii) T is a *DB theory*, that is, a set of universal Σ -sentences.

Next, we refer to a DB schema simply through its (DB) signature Σ and (DB) theory T . Given a DB signature Σ , we denote by Σ_{srt} the set of sorts and by Σ_{fun} the set of functions in Σ . Since Σ contains only unary function symbols and equality, all atomic Σ -formulae are of the form $t_1(v_1) = t_2(v_2)$, where t_1, t_2 are possibly complex terms, and v_1, v_2 are either variables or constants.

We associate to a DB signature Σ a characteristic (directed) graph $G(\Sigma)$ capturing the dependencies induced by functions over sorts. Specifically, $G(\Sigma)$ is an edge-labeled graph whose set of nodes is Σ_{srt} , and with a labeled edge $S \xrightarrow{f} S'$ for each $f : S \rightarrow S'$ in Σ_{fun} . We say that Σ is *acyclic* if $G(\Sigma)$ is so. The *leaves* of Σ are the nodes of $G(\Sigma)$ without outgoing edges. These terminal sorts are divided in two subsets, respectively representing *unary relations* and *value sorts*. Non-value sorts (i.e., unary relations and non-leaf sorts) are called *id sorts*, and are conceptually used to represent (identifiers of) different kinds of objects. Value sorts, instead, represent datatypes such as strings, numbers, clock values, etc. We denote the set of id sorts in Σ by Σ_{ids} , and that of value sorts by Σ_{val} , hence $\Sigma_{srt} = \Sigma_{ids} \uplus \Sigma_{val}$.

We now consider extensional data.

Definition 3.2. A *DB instance* of DB schema $\langle \Sigma, T \rangle$ is a Σ -structure \mathcal{M} that is a model of T and such that every id sort of Σ is interpreted in \mathcal{M} on a *finite* set.

Contrast this to arbitrary *models* of T , where no finiteness assumption is made. What may appear as not customary in Definition 3.2 is the fact that value sorts can be interpreted on infinite sets. This allows us, at once, to reconstruct the classical notion of DB instance as a finite model (since only finitely many values can be pointed from id sorts using functions), at the same time supplying a potentially infinite set of fresh values to be dynamically introduced in the working memory during the evolution of the artifact system. More details on this will be given in Section 3.1.

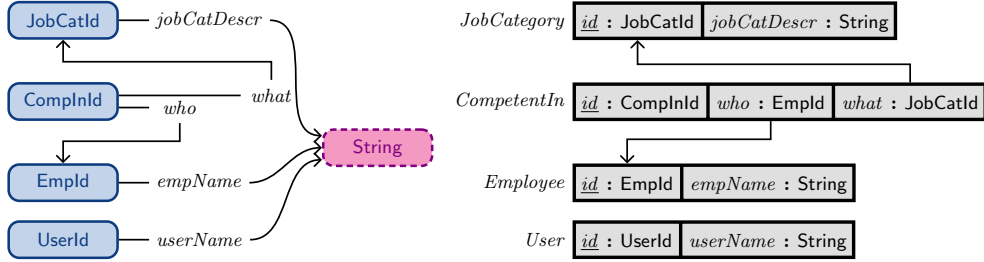


Figure 1. On the left: characteristic graph of the human resources DB signature from Example 3.1. On the right: relational view of the DB signature; each cell denotes an attribute with its type, underlined attributes denote primary keys, and directed edges capture foreign keys.

We respectively denote by $S^{\mathcal{M}}$, $f^{\mathcal{M}}$, and $c^{\mathcal{M}}$ the interpretation in \mathcal{M} of the sort S (this is a set), of the function symbol f (this is a set-theoretic function), and of the constant c (this is an element of the interpretation of the corresponding sort). Obviously, $f^{\mathcal{M}}$ and $c^{\mathcal{M}}$ must match the sorts in Σ . E.g., if f has source S and target U , then $f^{\mathcal{M}}$ has domain $S^{\mathcal{M}}$ and range $U^{\mathcal{M}}$.

Example 3.1. The human resource (HR) branch of a company stores the following information inside a relational database: (i) users registered to the company website, who are potentially interested in job positions offered by the company; (ii) the different, available job categories; (iii) employees belonging to HR, together with the job categories they are competent in (in turn indicating which job applicants they could interview). To formalize these different aspects, we make use of a DB signature Σ_{hr} consisting of: (i) four id sorts, used to respectively identify users, employees, job categories, and the competence relationship connecting employees to job categories; (ii) one value sort containing strings used to name users and employees, and describe job categories. In addition, Σ_{hr} contains five function symbols mapping: (i) user identifiers to their corresponding names; (ii) employee identifiers to their corresponding names; (iii) job category identifiers to their corresponding descriptions; (iv) competence identifiers to their corresponding employees and job categories. The characteristic graph of Σ_{hr} is shown in Figure 1 (left part). \triangleleft

We close the formalization of DB schemas by discussing DB theories. The role of a DB theory is to encode background axioms to express constraints on the different elements of the corresponding signature. We illustrate a typical background axiom, required to handle the possible presence of *undefined identifiers/values* in the different sorts. This, in turn, is essential to capture artifact systems whose working memory is initially undefined, in the style of (Deutsch et al., 2016; Li et al., 2017). To accommodate this, to specify an undefined value we add to every sort S of Σ a constant undef_S (written from now on, by abuse of notation, just as undef , used also to indicate a tuple). Then, for each function symbol f of Σ , we add the following axiom to the DB theory:

$$\forall x (x = \text{undef} \leftrightarrow f(x) = \text{undef}) \quad (2)$$

This axiom states that the application of f to the undefined value produces an undefined value, and it is the only situation for which f is undefined.

Remark 3.1. In the artifact-centric model in the style of (Deutsch et al., 2016; Li et al., 2017) that we intend to capture, the DB theory consists of Axioms (2) only. However, our technical results do not require this specific choice, and more general sufficient conditions will be discussed in Section 3.2.

Remark 3.2. If desired, we can freely extend DB schemas by adding arbitrary n -ary relation symbols to the signature Σ . For this purpose, we give the following definition.

Definition 3.3. A *DB extended-schema* is a pair $\langle \Sigma, T \rangle$, where: (i) Σ is a *DB extended-signature*, that is, a finite multi-sorted signature whose only symbols are equality, n -ary relations, unary functions, and constants; (ii) T is a *DB extended-theory*, that is, a set of universal Σ -sentences.

For simplicity, even if our implementation takes into account also the case of “free” relations, i.e. without key dependencies, we restrict our focus on DB schemas, which are sufficient to capture those constraints (as explained in the following subsection). The extension is straightforward and left to the reader. In fact, we can give in analogous the definitions of the characteristic graph $G(\Sigma)$ and of acyclicity for extended DB schemas. We notice that, in case Assumption 3.4 discussed below holds for DB extended-theories, all the results presented in Section 4 still hold even considering DB extended-schemas instead of DB schemas.

3.1. Relational View of DB Schemas

We now clarify how the algebraic, functional characterization of DB schema and instance can be actually reinterpreted in the classical, relational model. Definition 3.1 naturally corresponds to the definition of relational database schemas equipped with single-attribute *primary keys* and *foreign keys* (plus a reformulation of constraint (2)). To technically explain the correspondence, we adopt the *named perspective*, where each relation schema is defined by a signature containing a *relation name* and a set of *typed attribute names*. Let $\langle \Sigma, T \rangle$ be a DB schema. Each id sort $S \in \Sigma_{ids}$ corresponds to a dedicated relation R_S with the following attributes: (i) one identifier attribute id_S with type S ; (ii) one dedicated attribute a_f with type S' for every function symbol $f \in \Sigma_{fun}$ of the form $f : S \rightarrow S'$.

The fact that R_S is built starting from functions in Σ naturally induces different database dependencies in R_S . In particular, for each non-id attribute a_f of R_S , we get a *functional dependency* from id_S to a_f ; altogether, such dependencies in turn witness that id_S is the (*primary*) *key* of R_S . In addition, for each non-id attribute a_f of R_S whose corresponding function symbol f has id sort S' as image, we get an *inclusion dependency* from a_f to the id attribute $id_{S'}$ of $R_{S'}$; this captures that a_f is a *foreign key* referencing $R_{S'}$.

Example 3.2. The diagram on the right in Figure 1 graphically depicts the relational view corresponding to the DB signature of Example 3.1. \triangleleft

Given a DB instance \mathcal{M} of $\langle \Sigma, T \rangle$, its corresponding *relational instance* \mathcal{I} is the minimal set satisfying the following property: for every id sort $S \in \Sigma_{ids}$, let f_1, \dots, f_n be all functions in Σ with domain S ; then, for every identifier $o \in S^{\mathcal{M}}$, \mathcal{I} contains a *labeled fact* of the form $R_S(id_S : o^{\mathcal{M}}, a_{f_1} : f_1^{\mathcal{M}}(o^{\mathcal{M}}), \dots, a_{f_n} : f_n^{\mathcal{M}}(o^{\mathcal{M}}))$. With this interpretation, the *active domain* of \mathcal{I} is the set

$$\bigcup_{S \in \Sigma_{ids}} (S^{\mathcal{M}} \setminus \{\mathbf{undef}^{\mathcal{M}}\}) \cup \left\{ v \in \bigcup_{V \in \Sigma_{val}} V^{\mathcal{M}} \mid \begin{array}{l} v \neq \mathbf{undef}^{\mathcal{M}} \text{ and there exist } f \in \Sigma_{fun} \\ \text{and } o \in \text{dom}(f^{\mathcal{M}}) \text{ s.t. } f^{\mathcal{M}}(o) = v \end{array} \right\}$$

consisting of all (proper) identifiers assigned by \mathcal{M} to id sorts, as well as all values obtained in \mathcal{M} via the application of some function. Since such values are necessarily *finitely many*, one may wonder why in Definition 3.2 we allow for interpreting value sorts over infinite sets. The reason is that, in our framework, an evolving artifact system may use such infinite provision to inject and manipulate new values into the working memory. From the definition of active domain above, exploiting Axioms (2) we get that the membership of a tuple (x_0, \dots, x_n) to a generic $n + 1$ -ary relation R_S with key dependencies (corresponding to an id sort S) can be expressed in our setting by using just unary function symbols and equality:

$$R_S(x_0, \dots, x_n) \text{ iff } x_0 \neq \mathbf{undef} \wedge x_1 = f_1(x_0) \wedge \dots \wedge x_n = f_n(x_0) \quad (3)$$

Hence, the representation of negated atoms is the one that directly follows from negating (3):

$$\neg R_S(x_0, \dots, x_n) \text{ iff } x_0 = \mathbf{undef} \vee x_1 \neq f_1(x_0) \vee \dots \vee x_n \neq f_n(x_0) \quad (4)$$

This relational interpretation of DB schemas exactly reconstructs the requirements posed by (Deutsch et al., 2016; Li et al., 2017) on the schema of the *read-only* database: (i) each relation schema has a single-attribute primary key; (ii) attributes are typed; (iii) attributes may be foreign keys referencing other relation schemas; (iv) the primary keys of different relation schemas are pairwise disjoint.

We stress that all such requirements are natively captured in our functional definition of a DB signature, and do not need to be formulated as axioms in the DB theory. The DB theory is used to express additional constraints, like that in Axiom (2). In the following subsection, we thoroughly discuss which properties must be respected by signatures and theories to guarantee that our verification machinery is well-behaved.

One may wonder why we have not directly adopted a relational view for DB schemas. This will become clear during the technical development. We anticipate the main, intuitive reasons. First, our functional view allows us to guarantee that our framework remains well-behaved even in the presence of key dependencies, since our DB theories *do* enjoy the crucial condition of Assumption 3.4 introduced below (i.e., that the DB theories admit a model completion), whereas relational structures with key constraints *do not*.

Second, our functional view makes the dependencies among different types explicit. In fact, our notion of characteristic graph, which is readily computed from a DB signature, exactly reconstructs the central notion of foreign key graph used in (Deutsch et al., 2016) towards the main decidability results.

3.2. Formal Properties of DB Schemas

The theory T from Definition 3.1 must satisfy few crucial requirements for our approach to work. In this section, we define such requirements and show that they are matched, e.g., when we are concerned with an acyclic signature Σ and with key dependencies (i.e., the setting presented in (Li et al., 2017)). Actually, acyclicity is a stronger requirement than needed, which, however, simplifies our exposition.

3.2.1. Finite Model Property We say that T has the *finite model property* (for constraint satisfiability) iff every constraint ϕ that is satisfiable in a model of T is satisfiable in a DB instance of T . It can be easily seen that this implies that ϕ is satisfiable also in a DB instance interpreting also value sorts into finite sets.

Observe that if Σ is acyclic, there are only finitely many terms involving a single variable x : in fact, there are as many terms as paths in $G(\Sigma)$ starting from the sort of x . If k_Σ is the maximum number of terms involving a single variable, then (since all function symbols are unary) there are at most $k_\Sigma \cdot n$ terms involving n variables.

Proposition 3.1. Let (Σ, T) be a DB schema (cf. Definition 3.1); T has the finite model property in case Σ is acyclic.

Proof. If $T := \emptyset$, then congruence closure ensures that the finite model property holds and decides constraint satisfiability in polynomial time (Bradley and Manna, 2007).

Otherwise, we reduce the argument to the Herbrand Theorem (recall that T is universal according to Definition 3.1). Indeed, suppose to have a finite set Φ of universal formulae and let $\phi(\underline{x})$ be the constraint we want to test for satisfiability. Replace the variables \underline{x} with free constants \underline{a} . Herbrand Theorem states that $\Phi \cup \{\phi(\underline{a})\}$ has a model iff the set of ground $\Sigma^{\underline{a}}$ -instances of $\Phi \cup \{\phi(\underline{a})\}$ has a model. These ground instances are finitely many by acyclicity, so we can apply congruence closure (as done in the case of the empty theory) to these ground instances. □

Remark 3.3. If T is finite, Proposition 3.1 ensures decidability of constraint satisfiability. In order to obtain a decision procedure, it is sufficient to instantiate the axioms of T and the axioms of equality (reflexivity, transitivity, symmetry, congruence) and to use a SAT-solver to decide constraint satisfiability. Alternatively, one can decide constraint satisfiability via congruence closure (Bradley and Manna, 2007) and avoid instantiating the equality axioms.

Remark 3.4. Acyclicity is a strong condition, often too strong. However, some condition must be imposed (otherwise we have undecidability, and then failure of finite model property, by reduction to word problem for finite presentations of monoids). In fact, the

empty theory and the theory axiomatized by Axioms (2) both have the finite model property even without acyclicity assumptions.

The finite model property implies decidability of the constraint satisfiability problem in case T is recursively axiomatized. Indeed, in this case it is possible to enumerate unsatisfiable constraints via a logical calculus and this enumeration can be interleaved with the enumeration of finite models, thus supplying a full decision procedure.

3.2.2. Model Completion of DB theories. A DB theory T does not necessarily have quantifier elimination; it is however often possible to strengthen T in a conservative way (with respect to constraint satisfiability) and get quantifier elimination. In order to do that, we study the model completion of T , when it exists, and we will show that model completion turns out to be quite effective to attack the verification of dynamic systems operating over relational databases.

The following Lemma gives a useful folklore technique for finding model completions:

Lemma 3.1. Suppose that for every primitive Σ -formula $\exists x \phi(x, \underline{y})$ it is possible to find a quantifier-free formula $\psi(\underline{y})$ such that

- (i) $T \models \forall x \forall \underline{y} (\phi(x, \underline{y}) \rightarrow \psi(\underline{y}))$;
- (ii) for every model \mathcal{M} of T , for every tuple of elements \underline{a} from the support of \mathcal{M} such that $\mathcal{M} \models \psi(\underline{a})$ it is possible to find another model \mathcal{N} of T such that \mathcal{M} embeds into \mathcal{N} and $\mathcal{N} \models \exists x \phi(x, \underline{a})$.

Then T has a model completion T^* axiomatized by the infinitely many sentences

$$\forall \underline{y} (\psi(\underline{y}) \rightarrow \exists x \phi(x, \underline{y})) . \quad (5)$$

Proof. From (i) and (5) we clearly get that T^* admits quantifier elimination: in fact, in order to prove that a theory enjoys quantifier elimination, it is sufficient to eliminate quantifiers from *primitive* formulae (then the quantifier elimination for all formulae can be easily shown by an induction over their complexity). This is exactly what is guaranteed by (i) and (5).

Let \mathcal{M} be a model of T . We show (by using a chain argument) that there exists a model \mathcal{M}' of T^* such that \mathcal{M} embeds into \mathcal{M}' . For every primitive formula $\exists x \phi(x, \underline{y})$, consider the pair $(\underline{a}, \exists x \phi(x, \underline{a}))$ such that $\mathcal{M} \models \psi(\underline{a})$ (where ψ is related to ϕ as in (i)). By Zermelo's Theorem, the set of all pairs $\{(\underline{a}, \exists x \phi(x, \underline{a}))\}$ can be well-ordered: let $\{(\underline{a}_i, \exists x \phi_i(x, \underline{a}_i))\}_{i \in I}$ be such a well-ordered set (where I is an ordinal). By transfinite induction on this well-order, we define $\mathcal{M}_0 := \mathcal{M}$ and, for each $i \in I$, \mathcal{M}_i as an extension of $\bigcup_{j < i} \mathcal{M}_j$ such that $\mathcal{M}_i \models \exists x \phi_i(x, \underline{a}_i)$, which exists for (ii) since $\bigcup_{j < i} \mathcal{M}_j \models \psi(\underline{a}_i)$ (remember that validity of ground formulae is preserved passing through substructures and superstructures, and $\mathcal{M}_0 \models \psi(\underline{a}_i)$).

Now we take the chain union $\mathcal{M}^1 := \bigcup_{i \in I} \mathcal{M}_i$: since T is universal, \mathcal{M}^1 is again a model of T , and it is possible to construct an analogous chain \mathcal{M}^2 as done above, starting from \mathcal{M}^1 instead of \mathcal{M} . Clearly, we get $\mathcal{M}_0 := \mathcal{M} \subseteq \mathcal{M}^1 \subseteq \mathcal{M}^2$ by construction. At this point, we iterate the same argument countably many times, so as to define a new chain

of models of T :

$$\mathcal{M}_0 := \mathcal{M} \subseteq \mathcal{M}^1 \subseteq \dots \subseteq \mathcal{M}^n \subseteq \dots$$

Defining $\mathcal{M}' := \bigcup_n \mathcal{M}^n$, we trivially get that \mathcal{M}' is a model of T such that $\mathcal{M} \subseteq \mathcal{M}'$ and satisfies all the sentences of type (5). The last fact can be shown using the following finiteness argument.

Fix ϕ, ψ as in (5). For every tuple $\underline{a}' \in \mathcal{M}'$ such that $\mathcal{M}' \models \psi(\underline{a}')$, by definition of \mathcal{M}' there exists a natural number k such that $\underline{a}' \in \mathcal{M}^k$: since $\psi(\underline{a}')$ is a ground formula, we get that also $\mathcal{M}^k \models \psi(\underline{a}')$. Therefore, we consider the step k of the countable chain: there, we have that the pair $(\underline{a}', \exists x \phi(x, \underline{a}'))$ appears in the enumeration given by the well-ordered set of pairs $\{(\underline{a}_i, \exists x \phi_i(x, \underline{a}_i))\}_{i \in I}$ such that $\mathcal{M}^k \models \psi_i(\underline{a}_i)$. Hence, by construction, we have that $\mathcal{M}_i^k \models \exists x \phi(x, \underline{a}')$ for some i . In conclusion, since the existential formulae are preserved passing to extensions, we obtain $\mathcal{M}' \models \exists x \phi(x, \underline{a}')$, as wanted. \square

Proposition 3.2. T has a model completion in case it is axiomatized by universal one-variable formulae and Σ is acyclic.

Proof. We freely take inspiration from an analogous result in (Wheeler, 1976). We preliminarily show that T is amalgamable. Then, for a suitable choice of ψ suggested by the acyclicity assumption, the amalgamation property will be used to prove the validity of the condition (ii) of Lemma 3.1: this fact (together with condition (i)) yields that T has a model completion which is axiomatized by the infinitely many sentences (5).

Let \mathcal{M}_1 and \mathcal{M}_2 two models of T with a submodel \mathcal{M}_0 of T in common (we suppose for simplicity that $|\mathcal{M}_1| \cap |\mathcal{M}_2| = |\mathcal{M}_0|$). We define a T -amalgam \mathcal{M} of $\mathcal{M}_1, \mathcal{M}_2$ over \mathcal{M}_0 as follows (we use in an essential way the fact that Σ contains only *unary* function symbols). Let the support of \mathcal{M} be the set-theoretic union of the supports of \mathcal{M}_1 and \mathcal{M}_2 , i.e. $|\mathcal{M}| := |\mathcal{M}_1| \cup |\mathcal{M}_2|$. \mathcal{M} has a natural Σ -structure inherited by the Σ -structures \mathcal{M}_1 and \mathcal{M}_2 : for every function symbol f in Σ , we define, for each $m_i \in |\mathcal{M}_i| (i = 1, 2)$, $f^{\mathcal{M}}(m_i) := f^{\mathcal{M}_i}(m_i)$, i.e. the interpretation of f in \mathcal{M} is the interpretation of f in \mathcal{M}_i for every element $m_i \in |\mathcal{M}_i|$. This is well-defined since, for every $a \in |\mathcal{M}_1| \cap |\mathcal{M}_2| = |\mathcal{M}_0|$, we have that $f^{\mathcal{M}}(a) := f^{\mathcal{M}_1}(a) = f^{\mathcal{M}_0}(a) = f^{\mathcal{M}_2}(a)$. It is clear that \mathcal{M}_1 and \mathcal{M}_2 are substructures of \mathcal{M} , and their inclusions agree on \mathcal{M}_0 .

We show that the Σ -structure \mathcal{M} , as defined above, is a model of T . By hypothesis, T is axiomatized by universal one-variable formulae: so, we can consider T as a theory formed by axioms ϕ which are universal closures of clauses with just one variable, i.e. $\phi := \forall x (A_1(x) \wedge \dots \wedge A_n(x) \rightarrow B_1(x) \vee \dots \vee B_m(x))$, where A_j and B_k ($j = 1, \dots, n$ and $k = 1, \dots, m$) are atoms.

We show that \mathcal{M} satisfies all such formulae ϕ . In order to do that, suppose that, for every $a \in |\mathcal{M}|$, $\mathcal{M} \models A_j(a)$ for all $j = 1, \dots, n$. If $a \in |\mathcal{M}_i|$, then $\mathcal{M} \models A_j(a)$ implies $\mathcal{M}_i \models A_j(a)$, since $A_j(a)$ is a ground formula. Since \mathcal{M}_i is model of T and so $\mathcal{M}_i \models \phi$, we get that $\mathcal{M}_i \models B_k(a)$ for some $k = 1, \dots, m$, which means that $\mathcal{M} \models B_k(a)$, since $B_k(a)$ is a ground formula. Thus, $\mathcal{M} \models \phi$ for every axiom ϕ of T , i.e. $\mathcal{M} \models T$ and, hence, \mathcal{M} is a T -amalgam of $\mathcal{M}_1, \mathcal{M}_2$ over \mathcal{M}_0 , as wanted.

Now, given a primitive formula $\exists x \phi(x, \underline{y})$, we find a suitable ψ such that the hypothesis of Lemma 3.1 holds. We define $\psi(\underline{y})$ as the conjunction of the set of all quantifier-free

$\chi(y)$ -formulae such that $\phi(x, y) \rightarrow \chi(y)$ is a logical consequences of T (they are finitely many - up to T -equivalence - because Σ is acyclic). By definition, clearly we have that (i) of Lemma 3.1 holds.

We show that also condition (ii) is satisfied. Let \mathcal{M} be a model of T such that $\mathcal{M} \models \psi(\underline{a})$ for some tuple of elements \underline{a} from the support of \mathcal{M} . Then, consider the Σ -substructure $\mathcal{M}[\underline{a}]$ of \mathcal{M} generated by the elements \underline{a} : this substructure is finite (since Σ is acyclic), it is a model of T and we trivially have that $\mathcal{M}[\underline{a}] \models \psi(\underline{a})$, since $\psi(\underline{a})$ is a ground formula. In order to prove that there exists an extension \mathcal{N}' of $\mathcal{M}[\underline{a}]$ such that $\mathcal{N}' \models \exists x \phi(x, \underline{a})$, it is sufficient to prove (by the Robinson Diagram Lemma) that the $\Sigma^{|\mathcal{M}[\underline{a}]| \cup \{e\}}$ -theory $\Delta(\mathcal{M}[\underline{a}]) \cup \{\phi(e, \underline{a})\}$ is T -consistent. For reduction to absurdity, suppose that the last theory is T -inconsistent. Then, there are finitely many literals $l_1(\underline{a}), \dots, l_m(\underline{a})$ from $\Delta(\mathcal{M}[\underline{a}])$ (remember that $\Delta(\mathcal{M}[\underline{a}])$ is a finite set of literals since $\mathcal{M}[\underline{a}]$ is a finite structure) such that $\phi(e, \underline{a}) \models_T \neg(l_1(\underline{a}) \wedge \dots \wedge l_m(\underline{a}))$. Therefore, defining $A(\underline{a}) := l_1(\underline{a}) \wedge \dots \wedge l_m(\underline{a})$, we get that $\phi(e, \underline{a}) \models_T \neg A(\underline{a})$, which implies that $\neg A(\underline{a})$ is one of the $\chi(y)$ -formulae appearing in $\psi(\underline{a})$. Since $\mathcal{M}[\underline{a}] \models \psi(\underline{a})$, we also have that $\mathcal{M}[\underline{a}] \models \neg A(\underline{a})$, which is a contradiction: in fact, by definition of diagram, $\mathcal{M}[\underline{a}] \models A(\underline{a})$ must hold. Hence, there exists an extension \mathcal{N}' of $\mathcal{M}[\underline{a}]$ such that $\mathcal{N}' \models \exists x \phi(x, \underline{a})$. Now, by amalgamation property, there exists a T -amalgam \mathcal{N} of \mathcal{M} and \mathcal{N}' over $\mathcal{M}[\underline{a}]$: clearly, \mathcal{N} is an extension of \mathcal{M} and, since $\mathcal{N}' \hookrightarrow \mathcal{N}$ and $\mathcal{N}' \models \exists x \phi(x, \underline{a})$, also $\mathcal{N} \models \exists x \phi(x, \underline{a})$ holds, as required. \square

Remark 3.5. The proof of Proposition 3.2 gives an algorithm for quantifier elimination in the model completion. The algorithm works as follows (see the formula (5)): to eliminate the quantifier $\exists x$ from $\exists x \phi(x, \underline{y})$ take the conjunction of the clauses $\chi(\underline{y})$ implied by $\phi(x, \underline{y})$. This algorithm is far from optimal from two points of view. First, contrary to what happens in linear arithmetics, the quantifier elimination needed to prove Proposition 3.2 has a much better behaviour (from the complexity point of view) if obtained via a suitable version of the Knuth-Bendix procedure (Baader and Nipkow, 1998) or of the Superposition Calculus (Nieuwenhuis and Rubio, 2001). Since these aspects concerning quantifier elimination are rather delicate, we started studying them in a dedicated paper (Calvanese et al., 2019c) and in its extended version (Calvanese et al., 2018a) (our MCMT implementation, however, already partially takes into account such development), where, by using a constrained version of Superposition we show that in the case of free unary functions and free relations the complexity has a quadratic bound even without assuming acyclicity.

Secondly, it is worth noting that the algorithm presented in Proposition 3.2 uses the acyclicity assumption, whereas such assumption, as just noticed, is in general not needed for Proposition 3.2 to hold: for instance, when $T := \emptyset$ or when T contains only Axiom (2), a model completion can be proved to exist, even if Σ is not acyclic, by using the constrained version of Superposition studied in (Calvanese et al., 2019c; Calvanese et al., 2018a).

Remark 3.6. Proposition 3.2 holds also for DB extended-schemas, in case the universal one-variable formulae do not involve the relation symbols (so, the relations are “free”): as explained in (Calvanese et al., 2018a), our implementation of the quantifier elimination

algorithm takes into account also this case. More generally, the model completion exists whenever we consider an acyclic DB extended-schema with a DB extended-theory T that enjoys the amalgamation property.

Hereafter, we make the following assumption:

Assumption 3.4. The DB theories we consider have decidable constraint satisfiability problem, finite model property, and admit a model completion.

This assumption is matched, for instance, in the following three cases: (i) when T is empty; (ii) when T is axiomatized by Axioms (2); (iii) when Σ is acyclic and T is axiomatized by universal one-variable formulae (such as Axioms (2)).

Hence, the artifact-centric model in the style of (Deutsch et al., 2016; Li et al., 2017) that we intend to capture *matches* Assumption 3.4.

Remark 3.7. Notice that the DB extended-schemas obtained by adding “free” relations to the DB schemas of (i), (ii), (iii) above match Assumption 3.4.

4. Artifact-Centric Systems

We are now in the position to define our formal models of SASs and RASs, and to study parameterized safety problems over SASs and RASs. Since RASs are formalized in the spirit of array-based systems, we start by recalling the intuition behind them.

In general terms, an array-based system is described using a multi-sorted theory that contains two types of sorts, one accounting for the indexes of arrays, and the other for the elements stored therein. The system variables changing over time are both individual first-order variables for data and second-order variables for arrays. The latter are referred to using second-order function variables, whose interpretation in a state is that of a total function mapping indexes to elements (so that applying the function to an index denotes the classical *read* operation for arrays). The definition of an array-based system with array and data variables a always requires: a formula $\iota(a)$ describing the *initial configuration* of the variables a , and a formula $\tau(a, a')$ describing a *transition* that transforms the content of the variables from a to a' . In such a setting, verifying whether the system can reach unsafe configurations described by a formula $v(a)$ amounts to check whether the formula $\iota(a_0) \wedge \tau(a_0, a_1) \wedge \dots \wedge \tau(a_{n-1}, a_n) \wedge v(a_n)$ is satisfiable for some n .

Next, we make these ideas formally precise by grounding array-based systems in the artifact-centric setting. We start considering the case where we only have individual variables for data and then we pass to the complete framework where we also have second order variables formalizing artifact relations (that is, relations which are mutable during system evolution).

4.1. Simple Artifact Systems

The SAS Formal Model. In this subsection we consider systems manipulating only individual variables and reading data from a given database instance. In order to introduce

verification problems in a symbolic setting, one first has to specify which formulae are used to represent sets of states, the system initializations, and system evolution. Given a DB schema $\langle \Sigma, T \rangle$ and a tuple $\underline{x} = x_1, \dots, x_n$ of variables, we introduce the following classes of Σ -formulae:

- a *state formula* is a quantifier-free Σ -formula $\phi(\underline{x})$;
- an *initial formula* is a conjunction of equalities of the form $\bigwedge_{i=1}^n x_i = c_i$, where each c_i is a constant (typically, c_i is an **undef** constant mentioned in Section 3 above);
- a *transition formula* $\hat{\tau}$ is an existential formula

$$\exists \underline{y} \left(G(\underline{x}, \underline{y}) \wedge \bigwedge_{i=1}^n x'_i = F_i(\underline{x}, \underline{y}) \right) \quad (6)$$

where \underline{x}' are renamed copies of \underline{x} , G is quantifier-free and F_1, \dots, F_n are case-defined functions. We call G the *guard* and F_i the *updates* of Formula (6).

Definition 4.1. A *Simple Artifact System* (SAS) has the form

$$\mathcal{S} = \langle \Sigma, T, \underline{x}, \iota(\underline{x}), \tau(\underline{x}, \underline{x}') \rangle$$

where: (i) $\langle \Sigma, T \rangle$ is a DB schema, (ii) $\underline{x} = x_1, \dots, x_n$ are variables (called *artifact variables*), (iii) ι is an initial formula, and (iv) τ is a disjunction of transition formulae of the type (6).

We notice that a formula $\hat{\tau}$ of the kind (6) is a *single* transition formula, where τ from Definition 4.1 is a disjunction of formulae of the kind (6); hence, such τ symbolically represents the union of all the possible transitions of the system.

Example 4.1. We consider a SAS working over the DB schema of Example 3.1. It captures a global, single-instance artifact tracking the main, overall phases of a hiring process. The job hiring artifact employs a dedicated *pState* variable to store the current process state. Initially, hiring is disabled, which is captured by setting the *pState* variable to **undef**. A transition of the process from disabled to *enabled* may occur provided that the read-only HR DB contains at least one registered user (who, in turn, may decide to apply for a job). Technically, we introduce a dedicated artifact variable *uId* initialized to **undef**, and used to load the identifier of such a registered user, if (s)he exists. The enabling action is then captured by the following transition formula:

$$\exists y : \text{UserId} \left(\begin{array}{l} pState = \text{undef} \wedge y \neq \text{undef} \\ \wedge pState' = \text{enabled} \wedge uId' = y \end{array} \right)$$

The existential quantified variable $y : \text{UserId}$ is a data variable pointing to the read-only DB and is used to represent an external user input. Notice in particular how the existence of a user is checked using the typed variable y , checking that it is not **undef** and correspondingly assigning it to *uId*. ◁

Parameterized Safety via Backward Reachability for SAS. A *safety* formula for a SAS \mathcal{S} is a state formula $v(\underline{x})$ describing undesired states of \mathcal{S} . We say that \mathcal{S} is *safe with respect to* v if intuitively the system has no finite run leading from ι to v . Formally, there is no DB-instance \mathcal{M} of $\langle \Sigma, T \rangle$, no $k \geq 0$, and no assignment in \mathcal{M} to the variables

Algorithm 1: Backward reachability algorithm

```

Function BReach( $v$ )
1   $\phi \leftarrow v; B \leftarrow \perp;$ 
2  while  $\phi \wedge \neg B$  is  $T$ -satisfiable do
3    if  $\iota \wedge \phi$  is  $T$ -satisfiable then
4       $\perp$  return unsafe
5       $B \leftarrow \phi \vee B;$ 
6       $\phi \leftarrow \text{Pre}(\tau, \phi);$ 
7       $\phi \leftarrow \text{QE}(T^*, \phi);$ 
8  return safe;

```

$\underline{x}^0, \dots, \underline{x}^k$ such that the formula

$$\iota(\underline{x}^0) \wedge \tau(\underline{x}^0, \underline{x}^1) \wedge \dots \wedge \tau(\underline{x}^{k-1}, \underline{x}^k) \wedge v(\underline{x}^k) \quad (7)$$

is true in \mathcal{M} (here \underline{x}^i 's are renamed copies of \underline{x}). The *safety problem* for \mathcal{S} is the following: given a safety formula v decide whether \mathcal{S} is safe with respect to v .

Example 4.2. We provide an example of a safety formula for Example 4.1. Consider the *unsafe* configuration where the process is enabled but the identifier of the registered user loaded into uId is `undef`. Formally, this can be represented by the following state formula:

$$pState = \text{enabled} \wedge uId = \text{undef}$$

Notice that the following formula

$$\exists y (pState = \text{enabled} \wedge y \neq \text{undef} \wedge uId = y)$$

is *not* a safety formula, because of the existential quantified data variable y , but it is equivalent to

$$pState = \text{enabled} \wedge uId \neq \text{undef}$$

which is a safety formula. We will see in Lemma 4.2 that this equivalence (in some sense) holds in the general case of RASs (which SASs are a specific case of).

Algorithm 1 describes the *modified version of the backward reachability algorithm* (simply called *backward reachability* or *backward search* in the following) for handling the safety problem for \mathcal{S} (the original version of the backward reachability algorithm can be found in (Ghilardi et al., 2008a; Ghilardi and Ranise, 2010a)). An integral part of Algorithm 1 is to compute preimages. For that purpose, we define for any $\phi_1(\underline{z}, \underline{z}')$ and $\phi_2(\underline{z})$ (where \underline{z}' are renamed copies of \underline{z}), $\text{Pre}(\phi_1, \phi_2)$ as the formula $\exists \underline{z}' (\phi_1(\underline{z}, \underline{z}') \wedge \phi_2(\underline{z}'))$. The *preimage* of the set of states described by a state formula $\phi(\underline{x})$ is the set of states described by $\text{Pre}(\tau, \phi)$. Notice that, when $\tau = \bigvee \hat{\tau}$, then $\text{Pre}(\tau, \phi) = \bigvee \text{Pre}(\hat{\tau}, \phi)$. The modified algorithm presents a new subprocedure $\text{QE}(T^*, \phi)$ in Line 6 for computing quantifier elimination in the model completion T^* . The subprocedure $\text{QE}(T^*, \phi)$ applies the quantifier elimination algorithm of T^* to the existential formula ϕ . Algorithm 1 computes iterated preimages of v and applies to them quantifier elimination, until a fixpoint

is reached or until a set intersecting the initial states (i.e., satisfying ι) is found. *Inclusion* (Line 2) and *disjointness* (Line 3) tests can be discharged via proof obligations to be handled by SMT solvers. The fixpoint is reached when the test in Line 2 returns *unsat*, which means that the preimage of the set of the current states is included in the set of states reached by the backward search so far.

In the following, *partial correctness* means that, when the algorithm terminates, it gives a correct answer, whereas *effectiveness* means that all subprocedures in the algorithm can be effectively executed. We state now the main result of this subsection:

Theorem 4.2. Let $\langle \Sigma, T \rangle$ be a DB schema. Then, for every SAS $\mathcal{S} = \langle \Sigma, T, \underline{x}, \iota, \tau \rangle$ the following hold: (1) backward search is effective and partially correct for solving safety problems for \mathcal{S} ; (2) if Σ is acyclic, backward search terminates and decides safety problems for \mathcal{S} in PSPACE in the combined size of \underline{x} , ι , and τ .

Proof. Part (1). Recall formula (7)

$$\iota(\underline{x}^0) \wedge \tau(\underline{x}^0, \underline{x}^1) \wedge \cdots \wedge \tau(\underline{x}^{k-1}, \underline{x}^k) \wedge v(\underline{x}^k) .$$

By definition, \mathcal{S} is unsafe iff for some n , the formula (7) is satisfiable in a DB-instance of $\langle \Sigma, T \rangle$. Thanks to Assumption 3.4, T has the finite model property and consequently, as (7) is an existential Σ -formula, \mathcal{S} is unsafe iff for some n , formula (7) is satisfiable in a model of T ; furthermore, again by Assumption 3.4, \mathcal{S} is unsafe iff for some n , formula (7) is satisfiable in a model of T^* . Thus, we shall concentrate on satisfiability in models of T^* in order to prove the Theorem.

Let us call B_n (resp. ϕ_n), with $n \geq 0$, the status of the variable B (resp. ϕ) after n executions in Line 4 (resp. Line 6) of Algorithm 1 ($n = 0$ corresponds to the status of the variables in Line 1). Notice that we have

$$T^* \models \phi_{j+1} \leftrightarrow \text{Pre}(\tau, \phi_j) \tag{8}$$

for all j and that

$$T \models B_n \leftrightarrow \bigvee_{0 \leq j < n} \phi_j \tag{9}$$

is an invariant of the algorithm.

Since we are considering satisfiability in models of T^* , we can apply quantifier elimination and so the satisfiability of (7) is equivalent to the satisfiability of $\iota \wedge \phi_n$: this is a quantifier-free formula (because of line 6 of Algorithm 1), whose satisfiability (w.r.t. T or equivalently w.r.t. T^*)² is decidable by Assumption 1, so if Algorithm 1 terminates with an unsafe outcome, then \mathcal{S} is really unsafe.

Now consider the satisfiability test in Line 2. This is again a satisfiability test for a quantifier-free formula, thus it is decidable. In case of a safe outcome, we have that $T \models \phi_n \rightarrow B_n$; we claim that, if we continued executing the loop of Algorithm 1, we

² T -satisfiability and T^* -satisfiability are equivalent, by the definition of T^* , as far as existential (in particular, quantifier-free) formulae are concerned.

would nevertheless get that:

$$T^* \models B_m \leftrightarrow B_n \quad (10)$$

for all $m \geq n$. We justify Claim (10) below.

From $T \models \phi_n \rightarrow B_n$, taking into consideration that $T^* \supseteq T$ and that Formula (8) holds, we get $T^* \models \phi_{n+1} \rightarrow \text{Pre}(\tau, B_n)$. Since Pre commutes with disjunctions (i.e., $\text{Pre}(\tau, \bigvee_j \phi_j)$ is logically equivalent to $\bigvee_j \text{Pre}(\tau, \phi_j)$), we also have $T^* \models \text{Pre}(\tau, B_n) \leftrightarrow \bigvee_{1 \leq j \leq n} \phi_j$ by the Invariant (9) and by Formula (8) again. By using the entailment $T \models \phi_n \rightarrow B_n$ once more, we get $T^* \models \phi_{n+1} \rightarrow B_n$ and also that $T^* \models B_{n+1} \leftrightarrow B_n$, thus we finally obtain that $T^* \models \phi_{n+1} \rightarrow B_{n+1}$. Since $\phi_{n+1} \rightarrow B_{n+1}$ is quantifier-free, $T^* \models \phi_{n+1} \rightarrow B_{n+1}$ implies $T \models \phi_{n+1} \rightarrow B_{n+1}$. This argument can be repeated for all $m \geq n$, obtaining that $T^* \models B_m \leftrightarrow B_n$ for all $m \geq n$, i.e. Claim (10).

This would entail that $\iota \wedge \phi_m$ is always unsatisfiable (because of (9) and because $\iota \wedge \phi_j$ was unsatisfiable for all $j < n$), which is the same (as remarked above) as saying that all formulae (7) are unsatisfiable. Thus \mathcal{S} is safe.

Part (2). In case Σ is acyclic, there are only finitely many quantifier-free formulae (in which the finite set of variables \underline{x} occur), so it is evident that the algorithm must terminate: because of (9), the unsatisfiability test of Line 2 must eventually succeed, if the unsatisfiability test of Line 3 never does so.

Concerning complexity, we need to modify Algorithm 1 (we make it nondeterministic and use Savitch's Theorem saying that PSPACE = NPSpace).

Since Σ is acyclic, there are only finitely many terms involving a single variable, let this number be k_Σ (we consider T, Σ and hence k_Σ constant for our problems). Then, since all function symbols are unary, it is clear that we have at most $2^{O(n^2)}$ conjunctions of sets of literals involving at most n variables and that if the system is unsafe, unsafety can be detected with a run whose length is at most $2^{O(n^2)}$. Thus we introduce a counter to be incremented during the main loop (lines 2-6) of Algorithm 1. The fixpoint test in line 2 is removed and loop is executed only until the maximum length of an unsafe run is not exceeded (notice that an exponential counter requires polynomial space).

Inside the loop, line 4 is removed (we do not need anymore the variable B) and line 6 is modified as follows. We replace line 6 of the algorithm by

$$6'. \quad \phi \leftarrow \alpha(\underline{x});$$

where α is a non-deterministically chosen conjunction of literals implying $\text{QE}(T^*, \phi)$. Notice that to check the latter, there is no need to compute $\text{QE}(T^*, \phi)$: recalling the proof of Proposition 3.2 and Remark 3.5 it is sufficient to check that $T \models \alpha \rightarrow C$ holds for every clause $C(\underline{x})$ such that $T \models \phi \rightarrow C$.

The algorithm is now in PSPACE, because all the satisfiability tests we need are, as a consequence of the proof of Proposition 3.1, in NP: all such tests are reducible to T -satisfiability tests for quantifier-free Σ -formulae involving the variables \underline{x} and the additional (skolemized) quantified variables occurring in the transitions³. In fact, all

³ For the test in line 3, we just need replace in ϕ the \underline{x} by their values given by ι , conjoin the result with

these satisfiability tests are applied to formulae whose length is polynomial in the size of \underline{x} , of ι and of τ . \square

The proof of Theorem 4.2 shows that, whenever Σ is not acyclic, backward search is still a semi-decision procedure: if the system is unsafe, backward search always terminates and discovers it; if the system is safe, the procedure can diverge (but it is still correct).

Remark 4.1. We remark that Theorem 4.2 holds also for DB extended-schemas (so, even adding “free relations” to the DB signatures). Moreover, notice that it can be shown that every existential formula $\phi(\underline{x}, \underline{x}')$ can be turned into the form of Formula (6). Furthermore, we highlight that the proof of the decidability result of Theorem 4.2 requires that the considered background theory T : (i) admits a model completion; (ii) is *locally finite*, i.e., up to T -equivalence, there are only finitely many atoms involving a fixed finite number of variables (this condition is implied by acyclicity); (iii) is universal; and (iv) enjoys decidability of constraint satisfiability. Conditions (iii) and (iv) imply that one can decide whether a finite structure is a model of T . If (ii) and (iii) hold, it is well-known that (i) is equivalent to amalgamation (Wheeler, 1976), (Lipparini, 1982). Moreover, (ii) alone always holds for relational signatures and (iii) is equivalent to T being closed under substructures (this is a standard preservation theorem in model theory (Chang and Keisler, 1990)). It follows that *arbitrary relational signatures* (or *locally finite theories* in general, even allowing n -ary relation and n -ary function symbols) require only amalgamability and closure under substructures. Thanks to these observations, Theorem 4.2 is reminiscent of an analogous result in (Bojańczyk et al., 2013), i.e., Theorem 5, the crucial hypotheses of which are exactly amalgamability and closure under substructures, although the setting in that paper is different (there, key dependencies are not discussed, whereas we are interested only in DB (extended-)theories).

In our first-order setting, we can perform verification in a *purely symbolic* way, by using (semi-)decision procedures provided by SMT-solvers, even when local finiteness fails. As mentioned before, local finiteness is guaranteed in the relational context, but it does not hold anymore when *arithmetic operations* are introduced. Note that the theory of a single uninterpreted binary relation (i.e., the theory of directed graphs) has a model completion, whereas it can be easily seen that the theory of one binary relation which is a partial function *does not* (since it is *not* amalgamable). If primary key dependencies are formalized using partial functions, model completability is compromised. So, the second distinctive feature of our setting naturally follows from this observation: thanks to our *many-sorted functional representation* of DB schemas (with keys), the amalgamation property, required by Theorem 4.2, holds, witnessing that our framework remains well-behaved even in the presence of key dependencies.

all the ground instances of the axioms of T and finally decide satisfiability with congruence closure algorithm of a polynomial size ground conjunction of literals.

4.2. Relational Artifact Systems

The RAS Formal Model. Following the tradition of artifact-centric systems (Deutsch et al., 2009; Deutsch et al., 2016), a RAS consists of a read-only DB, a read-write working memory for artifacts, and a finite set of actions (also called services) that inspect the relational database and the working memory, and determine the new configuration of the working memory. In a RAS, the working memory consists of *individual* and *higher order* variables. These higher order variables (usually called *arrays*) are supposed to model evolving relations, so-called *artifact relations* in (Deutsch et al., 2016; Li et al., 2017). The idea is to treat artifact relations in a uniform way as we did for the read-only DB: we need extra sort symbols (recall that each sort symbol corresponds to a database relation symbol) and extra unary function symbols, the latter being treated as second-order variables.

Given a DB schema Σ , an *artifact extension* of Σ is a signature Σ_{ext} obtained from Σ by adding to it some extra sort symbols⁴. These new sorts (usually indicated with $E, E_1, E_2 \dots$) are called *artifact sorts* (or *artifact relations* by some abuse of terminology), whereas the old sorts from Σ are called *basic sorts*. In RAS, artifacts and basic sorts correspond, respectively, to the index and the elements sorts mentioned in the literature on array-based systems. Below, given $\langle \Sigma, T \rangle$ and an artifact extension Σ_{ext} of Σ , when we speak of a Σ_{ext} -model of T , a DB instance of $\langle \Sigma_{ext}, T \rangle$, or a Σ_{ext} -model of T^* , we mean a Σ_{ext} -structure \mathcal{M} whose reduct to Σ respectively is a model of T , a DB instance of $\langle \Sigma, T \rangle$, or a model of T^* .

An *artifact setting* over Σ_{ext} is a pair $(\underline{x}, \underline{a})$ given by a finite set \underline{x} of individual variables and a finite set \underline{a} of unary function variables: *the latter are required to have an artifact sort as source sort and a basic sort as target sort*. Variables in \underline{x} are called (as before) *artifact variables*, and variables in \underline{a} *artifact components*. Given a DB instance \mathcal{M} of Σ_{ext} , an *assignment* to an artifact setting $(\underline{x}, \underline{a})$ over Σ_{ext} is a map α assigning to every artifact variable $x_i \in \underline{x}$ of sort S_i an element $x_i^\alpha \in S_i^\mathcal{M}$ and to every artifact component $a_j : E_j \rightarrow U_j$ (with $a_j \in \underline{a}$) a set-theoretic function $a_j^\alpha : E_j^\mathcal{M} \rightarrow U_j^\mathcal{M}$.

We can view an assignment to an artifact setting $(\underline{x}, \underline{a})$ as a DB instance *extending* the DB instance \mathcal{M} as follows. Let all the artifact components in $(\underline{x}, \underline{a})$ having source E be $a_{i_1} : E \rightarrow S_1, \dots, a_{i_n} : E \rightarrow S_n$. Viewed as a relation in the artifact assignment (\mathcal{M}, α) , the artifact relation E “consists” of the set of tuples

$$\{ \langle e, a_{i_1}^\alpha(e), \dots, a_{i_n}^\alpha(e) \rangle \mid e \in E^\mathcal{M} \}$$

Thus each element of E is formed by an “entry” $e \in E^\mathcal{M}$ (uniquely identifying the tuple) and by “data” $\underline{a}_i^\alpha(e)$ taken from the read-only database \mathcal{M} . When the system evolves, the set $E^\mathcal{M}$ of entries remains fixed, whereas the components $\underline{a}_i^\alpha(e)$ may change: typically, we initially have $\underline{a}_i^\alpha(e) = \text{undef}$, but these values are changed when some defined values are inserted into the relation modeled by E ; the values are then repeatedly modified (and possibly also reset to undef , if the tuple is removed and e is re-set to point to undefined

⁴ By ‘signature’ we always mean ‘signature with equality’, so as, soon as new sorts are added, the corresponding equality predicates are added too.

values). In accordance with MCMT conventions, we denote the application of an artifact component a to a term (i.e., constant or variable) v also as $a[v]$ (standard notation for arrays), instead of $a(v)$.

To introduce Relational Artifact Systems, we discuss the kind of formulae we use. In such formulae, we use notations like $\phi(\underline{z}, \underline{b})$ to mean that ϕ is a formula whose free individual variables are among the \underline{z} and whose free unary function variables are among the \underline{b} .

Let $(\underline{x}, \underline{a})$ be an artifact setting over Σ_{ext} , where $\underline{x} = x_1, \dots, x_n$ are the artifact variables and $\underline{a} = a_1, \dots, a_m$ are the artifact components (their source and target sorts are left implicitly specified). We list the kind of formulae we shall use:

- An *initial formula* is a formula $\iota(\underline{x}, \underline{a})$ of the form

$$\left(\bigwedge_{i=1}^n x_i = c_i \right) \wedge \left(\bigwedge_{j=1}^m a_j = \lambda y. d_j \right) \quad (11)$$

where c_i, d_j are constants from Σ (typically, c_i and d_j are `undef`). Recall that $a_j = \lambda y. d_j$ abbreviates $\forall y a_j(y) = d_j$.

- A *state formula* has the form

$$\exists \underline{e} \phi(\underline{e}, \underline{x}, \underline{a}) \quad (12)$$

where ϕ is quantifier-free and the \underline{e} are individual variables of artifact sorts.

- A *transition formula* $\hat{\tau}$ has the form

$$\exists \underline{e} \left(\begin{array}{l} \gamma(\underline{e}, \underline{x}, \underline{a}) \wedge \bigwedge_i x'_i = F_i(\underline{e}, \underline{x}, \underline{a}) \\ \wedge \bigwedge_j a'_j = \lambda y. G_j(y, \underline{e}, \underline{x}, \underline{a}) \end{array} \right) \quad (13)$$

where the \underline{e} are individual variables (of *both* basic and artifact sorts), γ (the ‘guard’) is quantifier-free, $\underline{x}', \underline{a}'$ are renamed copies of $\underline{x}, \underline{a}$, and the F_i, G_j (the ‘updates’) are case-defined functions.

Note that transition formulae as above can express, e.g., (i) insertion (with/without duplicates) of a tuple in an artifact relation, (ii) removal of a tuple from an artifact relation, (iii) transfer of a tuple from an artifact relation to artifact variables (and vice-versa), and (iv) removal/modification of *all* the tuples satisfying a certain condition from an artifact relation. All the above operations can also be constrained; the formalization of the above operations in the formalism of our transition is straightforward (the reader can see all the details in the Appendix F from (Calvanese et al., 2018b)).

Definition 4.3. A *Relational Artifact System* (RAS) has the form

$$\mathcal{S} = \langle \Sigma, T, \Sigma_{ext}, \underline{x}, \underline{a}, \iota(\underline{x}, \underline{a}), \tau(\underline{x}, \underline{a}, \underline{x}', \underline{a}') \rangle$$

where: (i) $\langle \Sigma, T \rangle$ is a (read-only) DB schema, (ii) Σ_{ext} is an artifact extension of Σ , (iii) $(\underline{x}, \underline{a})$ is an artifact setting over Σ_{ext} , (iv) ι is an initial formula, and (v) τ is a disjunction of transition formulae of the type (13).

Notice that SASs are a particular class of RASs where the working memory consists *only* of artifact variables (without artifact relations).

Example 4.3. We transform the SAS of Example 4.1 into a RAS \mathcal{S}_{hr} containing a multi-instance artifact accounting for the evolution of *job applications*. Each job category may

receive multiple applications from registered users. Such applications are then evaluated, finally deciding which are accepted and which are rejected. The example is inspired by the job hiring process presented in (Silver, 2011) to show the intrinsic difficulties of capturing real-life processes with many-to-many interacting business entities using conventional process modeling notations (such as BPMN). An extended version of this example, capturing the co-evolution of multiple instances of two different artifacts, is presented in Appendix A.1 of (Calvanese et al., 2018b).

As for the read-only DB, \mathcal{S}_{hr} works over the DB schema of Example 3.1, extended with a further value sort `Score` used to score job applications. `Score` contains 102 values in the range $[-1, 100]$, where -1 denotes the non-eligibility of the application, and a score from 0 to 100 indicates the actual one assigned after evaluating the application. For the sake of readability, we use usual predicates $<$, $>$, and $=$ to compare variables of type `Score`. This is syntactic sugar and does not require to introduce *rigid* predicates (i.e., first-order logic relational symbols that have a pre-determined interpretation in a assigned quantified domain) in our framework: in the following, we do not need to introduce a new relational symbol $>$ that has to be interpreted in the model of natural numbers as the standard numerical relation $>$. Indeed, since in our case $>$ is applied to a fixed finite number of elements (i.e. the natural numbers in $[-1, 100]$), a literal of the form $s > 80$ can be substituted by a disjunction of a finite number of equalities (i.e. $s > 80$ iff $(s = 81$ or $s = 82$ or...or $s = 100)$).

As for the working memory, \mathcal{S}_{hr} consists of two artifacts: the single-instance *job hiring* artifact tracking the three main phases of the overall process (and described in Example 4.1), and a multi-instance artifact accounting for the evolution of *user applications*. To model applications, we take the DB signature Σ_{hr} of the read-only database of human resources, and enrich it with an artifact extension containing an artifact sort `appIndex` used to *index* (i.e., “internally” identify) job applications. The management of job applications is then modeled by an artifact setting with: (i) artifact components with domain `appIndex` capturing the artifact relation that stores the different job applications; (ii) additional individual variables as a temporary memory to manipulate the artifact relation. Specifically, each application consists of a job category, the identifier of the applicant user and that of an HR employee responsible for the application, the application score, and the final result (indicating whether the application is among the winners or the losers for the job offer). These information slots are encapsulated into dedicated artifact components, i.e., function variables with domain `appIndex` that collectively realize the application artifact relation:

$$\begin{aligned}
 appJobCat & : appIndex \longrightarrow JobCatId \\
 applicant & : appIndex \longrightarrow UserId \\
 appResp & : appIndex \longrightarrow EmplId \\
 appScore & : appIndex \longrightarrow Score \\
 appResult & : appIndex \longrightarrow String
 \end{aligned}$$

We now discuss the relevant transitions for inserting and evaluating job applications. When writing transition formulae, we make the following assumption: if an artifact variable/component is not mentioned at all, it is meant that it is updated identically;

otherwise, the relevant update function will specify how it is updated. Notice that, as mentioned also in the introduction, non-deterministic updates can be formalized using existentially quantified variables in the transition. The insertion of an application into the system can be executed when the hiring process is enabled (cf. Example 4.1), and consists of two consecutive steps. To indicate when a step can be applied, also ensuring that the insertion of an application is not interrupted by the insertion of another one, we manipulate a string artifact variable $aState$. The first step is executable when $aState$ is **undef**, and aims at loading the application data into dedicated artifact variables through the following simultaneous effects: (i) the identifier of the user who wants to submit the application, and that of the targeted job category, are selected and respectively stored into variables uId and jId ; (ii) the identifier of an HR employee who becomes responsible for the application is selected and stored into variable eId , with the requirement that such an employee must be competent in the job category targeted by the application; (iii) $aState$ evolves into state **received**. Formally:

$$\begin{array}{l} \exists u:\text{UserId}, j:\text{JobCatId}, e:\text{Empld}, c:\text{Complnd} \\ \left(\begin{array}{l} pState = \text{enabled} \wedge aState = \text{undef} \\ \wedge u \neq \text{undef} \wedge j \neq \text{undef} \wedge e \neq \text{undef} \wedge c \neq \text{undef} \\ \wedge \text{who}(c) = e \wedge \text{what}(c) = j \\ \wedge pState' = \text{enabled} \wedge aState' = \text{received} \\ \wedge uId' = u \wedge jId' = j \wedge eId' = e \wedge cId' = c \end{array} \right) \end{array}$$

The second step transfers the application data into the application artifact relation, using its corresponding function variables, at the same time resetting all application-related artifact variables to **undef** (including $aState$, so that new applications can be inserted). For the insertion, a “free” index (i.e., an index pointing to an undefined applicant) is picked. The newly inserted application gets a default score of -1 (thus initializing it to “not eligible”), while the final result is **undef**:

$$\begin{array}{l} \exists i:\text{appIndex} \\ \left(\begin{array}{l} pState = \text{enabled} \wedge aState = \text{received} \\ \wedge \text{applicant}[i] = \text{undef} \\ \wedge pState' = \text{enabled} \wedge aState' = \text{undef} \wedge cId' = \text{undef} \\ \wedge \text{appJobCat}' = \lambda j. (\text{if } j = i \text{ then } jId \text{ else } \text{appJobCat}[j]) \\ \wedge \text{applicant}' = \lambda j. (\text{if } j = i \text{ then } uId \text{ else } \text{applicant}[j]) \\ \wedge \text{appResp}' = \lambda j. (\text{if } j = i \text{ then } eId \text{ else } \text{appResp}[j]) \\ \wedge \text{appScore}' = \lambda j. (\text{if } j = i \text{ then } -1 \text{ else } \text{appScore}[j]) \\ \wedge \text{appResult}' = \lambda j. (\text{if } j = i \text{ then } \text{undef} \text{ else } \text{appResult}[j]) \\ \wedge jId' = \text{undef} \wedge uId' = \text{undef} \wedge eId' = \text{undef} \end{array} \right) \end{array}$$

Notice that such a transition does not prevent the possibility of inserting exactly the same application twice, at different indexes. If this is not wanted, the transition can be suitably changed so as to guarantee that no two identical applications can coexist in the same artifact relation (see the Appendix A.1 of (Calvanese et al., 2018b) for an example).

Each application currently considered as not eligible can be made eligible by assigning

a proper score to it:

$$\begin{aligned} & \exists i:\text{applIndex}, s:\text{Score} \\ & \left(\begin{aligned} & pState = \text{enabled} \wedge \text{appScore}[i] = -1 \wedge s \geq 0 \\ & \wedge pState' = \text{enabled} \wedge \text{appScore}'[i] = s \end{aligned} \right) \end{aligned}$$

Finally, application results are computed when the process moves to state `notified`. This is handled by the *bulk* transition:

$$\begin{aligned} & pState = \text{enabled} \wedge pState' = \text{notified} \\ & \wedge \text{appResult}' = \lambda j. \left(\begin{aligned} & \text{if } \text{appScore}[j] > 80 \text{ then } \text{winner} \\ & \text{else } \text{loser} \end{aligned} \right) \end{aligned}$$

which declares applications with a score above 80 as winning, and the others as losing. \triangleleft

Parameterized Safety via Backward Reachability for RAS. As for SAS, a *safety* formula for a RAS \mathcal{S} is a state formula $v(\underline{x})$. We say that \mathcal{S} is *safe with respect to* v if there is no DB-instance \mathcal{M} of $\langle \Sigma_{ext}, T \rangle$, no $k \geq 0$, and no assignment in \mathcal{M} to the variables $\underline{x}^0, \underline{a}^0 \dots, \underline{x}^k, \underline{a}^k$ such that the formula

$$\begin{aligned} & \iota(\underline{x}^0, \underline{a}^0) \wedge \tau(\underline{x}^0, \underline{a}^0, \underline{x}^1, \underline{a}^1) \\ & \wedge \dots \wedge \tau(\underline{x}^{k-1}, \underline{a}^{k-1}, \underline{x}^k, \underline{a}^k) \wedge v(\underline{x}^k, \underline{a}^k) \end{aligned} \quad (14)$$

is true in \mathcal{M} (here $\underline{x}^i, \underline{a}^i$ are renamed copies of $\underline{x}, \underline{a}$). The safety problem is defined as for SAS.

Example 4.4. We consider a safety property for the RAS from Example 4.3 that checks whether, after having received the evaluation notification, there are no applicants left without winner or loser status being assigned:

$$\begin{aligned} & \exists i:\text{applIndex} \\ & \left(\begin{aligned} & pState = \text{notified} \wedge \text{applicant}[i] \neq \text{undef} \\ & \wedge \text{appResult}[i] \neq \text{winner} \wedge \text{appResult}[i] \neq \text{loser} \end{aligned} \right) \end{aligned}$$

The job hiring RAS \mathcal{S}_{hr} turns out to be safe with respect to this property (cf. Section 6). \triangleleft

Interestingly, we can still run backward search for handling safety problems in RASs. In fact, Algorithm 1 presents the same structure. Notice that in this case the definition of $Pre(\tau, \phi)$ gives us $\exists \underline{x}' \exists \underline{a}' (\tau(\underline{x}, \underline{a}, \underline{x}', \underline{a}') \wedge \phi(\underline{x}', \underline{a}'))$. The subprocedure $QE(T^*, \phi)$ mentioned on Line 6 is extended so as to convert the preimage $Pre(\tau, \phi)$ of a state formula ϕ into a state formula (equivalent to it modulo the axioms of T^*), witnessing its *regressability*: this is possible since T^* eliminates from primitive formulae the existentially quantified variables over the basic sorts, whereas elimination of quantified variables over artifact sorts is not possible, because these variables occur as arguments of artifact components (see Lemma 4.1 and Lemma 4.2 below for details). In addition, the satisfiability tests from Lines 2-3 can still be discharged (in fact, we prove in Lemma 4.3 below that the entailment between state formulae can be decided via instantiation techniques). In the following, when we refer to Algorithm 1, we mean *Algorithm 1 adapted to RASs as explained above*.

In analogy to Statement (1) of Theorem 4.2, we obtain:

Theorem 4.4. Backward search (cf. Algorithm 1) is effective and partially correct for solving safety problems for RASs.

The proof of the above result occupies next paragraph.

Proof of Theorem 4.4.

When introducing our transition formulae in (6), (13) we made use of definable extensions and also of some function definitions via λ -abstraction. We already observed that such uses are due to notational convenience and do not really go beyond first-order logic. We are clarifying one more point now, before going into formal proofs. The lambda-abstraction definitions in (13) will make the proof of Lemma 4.1 below smooth. Recall that an expression like

$$b = \lambda y.F(y, \underline{z})$$

can be seen as a mere abbreviation of $\forall y \ b(y) = F(y, \underline{z})$. However, the use of such abbreviation makes clear that e.g. a formula like

$$\exists b \ (b = \lambda y.F(y, \underline{z}) \wedge \phi(\underline{z}, b))$$

is equivalent to

$$\phi(\underline{z}, \lambda y.F(y, \underline{z})/b) \ . \quad (15)$$

Since our $\phi(\underline{z}, b)$ is in fact a first-order formula, our b can occur in it only in terms like $b(t)$, so that in (15) all occurrences of λ can be eliminated by the so-called β -conversion: replace $\lambda y.F(y, \underline{z})(t)$ by $F(t, \underline{z})$. Thus, in the end, either we use definable extensions or definitions via lambda abstractions, *the formulae we manipulate can always be converted into plain first-order Σ - or Σ_{ext} -formulae.*

Let us call *extended state formulae* the formulae of the kind $\exists \underline{e} \ \phi(\underline{e}, \underline{x}, \underline{a})$, where ϕ is quantifier-free and the \underline{e} are individual variables of both artifact and basic sorts.

Lemma 4.1. The preimage of an extended state formula is logically equivalent to an extended state formula.

Proof. We manipulate the formula

$$\exists \underline{x}' \exists \underline{a}' \ (\tau(\underline{x}, \underline{a}, \underline{x}', \underline{a}') \wedge \exists \underline{e} \ \phi(\underline{e}, \underline{x}', \underline{a}')) \quad (16)$$

up to logical equivalence, where τ is given by⁵

$$\exists \underline{e}_0 \ (\gamma(\underline{e}_0, \underline{x}, \underline{a}) \wedge \underline{x}' = \underline{F}(\underline{e}_0, \underline{x}, \underline{a}) \wedge \underline{a}' = \lambda y.\underline{G}(y, \underline{e}_0, \underline{x}, \underline{a})) \quad (17)$$

(here we used plain equality for conjunctions of equalities, e.g. $\underline{x}' = \underline{F}(\underline{e}_0, \underline{x}, \underline{a})$ stands for $\bigwedge_i x'_i = F_i(\underline{e}, \underline{x}, \underline{a})$). Repeated substitutions show that (16) is equivalent to

$$\exists \underline{e} \exists \underline{e}_0 \ (\gamma(\underline{e}_0, \underline{x}, \underline{a}) \wedge \phi(\underline{e}, \underline{F}(\underline{e}_0, \underline{x}, \underline{a})/\underline{x}', \lambda y.\underline{G}(y, \underline{e}_0, \underline{x}, \underline{a})/\underline{a}')) \quad (18)$$

which is an extended state formula. \square

⁵ Actually, τ is a disjunction of such formulae, but it easily seen that disjunction can be accommodated by moving existential quantifiers back-and-forth through them.

Lemma 4.2. For every extended state formula there is a state formula equivalent to it in all Σ_{ext} -models of T^* .

Proof. Let $\exists \underline{e} \exists \underline{y} \phi(\underline{e}, \underline{y}, \underline{x}, \underline{a})$, be an extended state formula, where ϕ is quantifier-free, the \underline{e} are variables whose sort is an artifact sort and the \underline{y} are variables whose sort is a basic sort.

Now observe that, according to our definitions, the artifact components have an artifact sort as source sort and a basic sort as target sort; since equality is the only predicate, the literals in ϕ can be divided into equalities/inequalities between variables from \underline{e} and literals where the \underline{e} can only occur as arguments of an artifact component. Let $\underline{a}[\underline{e}]$ be the tuple of the terms among the terms of the kind $a_j[e_s]$ which are well-typed; using disjunctive normal forms, our extended state formula can be written as a disjunction of formulae of the kind

$$\exists \underline{e} \exists \underline{y} (\phi_1(\underline{e}) \wedge \phi_2(\underline{y}, \underline{x}, \underline{a}[\underline{e}]/\underline{z})) \quad (19)$$

where ϕ_1 is a conjunction of equalities/inequalities, $\phi_2(\underline{y}, \underline{x}, \underline{z})$ is a quantifier-free Σ -formula and $\phi_2(\underline{y}, \underline{x}, \underline{a}[\underline{e}]/\underline{z})$ is obtained from ϕ_2 by replacing the variables \underline{z} by the terms $\underline{a}[\underline{e}]$. Moving inside the existential quantifiers \underline{y} , we can rewrite (19) to

$$\exists \underline{e} (\phi_1(\underline{e}) \wedge \exists \underline{y} \phi_2(\underline{y}, \underline{x}, \underline{a}[\underline{e}]/\underline{z})) \quad (20)$$

Since T^* has quantifier elimination, we have that there is $\psi(\underline{x}, \underline{z})$ which is equivalent to $\exists \underline{y} \phi_2(\underline{y}, \underline{x}, \underline{z})$ in all models of T^* ; thus in all Σ_{ext} -models of T^* , the formula (20) is equivalent to

$$\exists \underline{e} (\phi_1(\underline{e}) \wedge \psi(\underline{x}, \underline{a}[\underline{e}]/\underline{z}))$$

which is a state formula. □

We underline that Lemmas 4.1 and 4.2 both give an explicit effective procedure for computing equivalent (extended) state formulae. Used one after the other, such procedures extends the procedure $QE(T^*, \phi)$ in Line 6 of Algorithm 1 to (non simple) artifact systems. Thanks to such procedure, the only formulae we need to test for satisfiability in lines 2 and 3 of the backward reachability algorithm are the $\exists \forall$ -formulae introduced below.

Let us call $\exists \forall$ -formulae the formulae of the kind

$$\exists \underline{e} \forall \underline{i} \phi(\underline{e}, \underline{i}, \underline{x}, \underline{a}) \quad (21)$$

where the variables $\underline{e}, \underline{i}$ are variables whose sort is an artifact sort and ϕ is quantifier-free. The crucial point for the following lemma to hold is that the *universally* quantified variables in $\exists \forall$ -formulae are all of artifact sorts:

Lemma 4.3. The satisfiability of a $\exists \forall$ -formula in a Σ_{ext} -model of T is decidable. Moreover, given a $\exists \forall$ -formula χ , the following three statements are equivalent:

- χ is satisfiable in a Σ_{ext} -model of T
- χ is satisfiable in a DB-instance of $\langle \Sigma_{ext}, T \rangle$
- χ is satisfiable in a Σ_{ext} -model of T^* .

Proof. First of all, notice that a $\exists\forall$ -formula (21) is equivalent to a disjunction of formulae of the kind

$$\exists \underline{e} (\text{AllDiff}(\underline{e}) \wedge \forall \underline{i} \phi(\underline{e}, \underline{i}, \underline{x}, \underline{a})) \quad (22)$$

where $\text{AllDiff}(\underline{e})$ says that any two variables of the same sort from the \underline{e} are distinct (to this aim, it is sufficient to guess a partition and to keep, via a substitution, only one element for each equivalence class).⁶ So we can freely assume that $\exists\forall$ -formulae are all of the kind (22).

Let us consider now the set of all (sort-matching) substitutions σ mapping the \underline{i} to the \underline{e} . The formula (22) is satisfiable (respectively: in a Σ_{ext} -model of T , in a DB-instance of $\langle \Sigma_{ext}, T \rangle$, in a Σ_{ext} -model of T^*) iff so it is the formula

$$\exists \underline{e} (\text{AllDiff}(\underline{e}) \wedge \bigwedge_{\sigma} \phi(\underline{e}, \underline{i}\sigma, \underline{x}, \underline{a})) \quad (23)$$

(here $\underline{i}\sigma$ means the componentwise application of σ to the \underline{i}): this is because, if (23) is satisfiable in \mathcal{M} , then we can take as \mathcal{M}' the same Σ_{ext} -structure as \mathcal{M} , but with the interpretation of the artifact sorts restricted only to the elements named by the \underline{e} and get in this way a Σ_{ext} -structure \mathcal{M}' satisfying (22) (notice that \mathcal{M}' is still a DB-instance of $\langle \Sigma_{ext}, T \rangle$ or a Σ_{ext} -model of T^* , if so was \mathcal{M}). Thus, we can freely concentrate on the satisfiability problem of formulae of the kind (23) only.

Now, by the way Σ_{ext} is built, the only atoms occurring in the subformula $\phi(\underline{e}, \underline{i}\sigma, \underline{x}, \underline{a})$ of (23) whose argument terms are terms of artifact sorts are of the kind $e_s = e_j$, so all such atoms can be replaced either by \top or by \perp (depending on whether we have $s = j$ or not). So we can assume that there are no such atoms in $\phi(\underline{e}, \underline{i}\sigma, \underline{x}, \underline{a})$ and as a result, the variables \underline{e} can only occur there as arguments of the \underline{a} .

Let now $\underline{a}[\underline{e}]$ be the tuple of the terms among the terms of the kind $a_j[e_s]$ which are well-typed. Since in (23) the \underline{e} can only occur as arguments of the artifact components, as observed above, the formula (23) is in fact of the kind

$$\exists \underline{e} (\text{AllDiff}(\underline{e}) \wedge \psi(\underline{x}, \underline{a}[\underline{e}]/\underline{z})) \quad (24)$$

where $\psi(\underline{x}, \underline{z})$ is a quantifier-free Σ -formula and $\psi(\underline{x}, \underline{a}[\underline{e}]/\underline{z})$ is obtained from ψ by replacing the variables \underline{z} by the terms $\underline{a}[\underline{e}]$ (notice that the \underline{z} are of basic sorts because the target sorts of the artifact components are basic sorts).

It is now evident that (24) is satisfiable (respectively: in a Σ_{ext} -model of T , in a DB-instance of $\langle \Sigma_{ext}, T \rangle$, in a Σ_{ext} -model of T^*) iff the formula

$$\psi(\underline{x}, \underline{z}) \quad (25)$$

is satisfiable (respectively: in a Σ -model of T , in a DB-instance of $\langle \Sigma, T \rangle$, in a Σ -model of T^*). In fact, if we are given a Σ -structure \mathcal{M} and an assignment satisfying (25), we can easily expand \mathcal{M} to a Σ_{ext} -structure by taking the e 's themselves as the elements of the interpretation of the artifact sorts; in the so-expanded Σ_{ext} -structure, we can interpret

⁶ In the MCMT implementation, state formulae are always maintained so that all existential variables occurring in them are differentiated and there is no need of this expensive computation step.

the artifact components \underline{a} by taking the $\underline{a}[\underline{e}]$ to be the elements assigned to the \underline{z} in the satisfying assignment for (25).

Thanks to Assumption 3.4, the satisfiability of (25) in a Σ -model of T , in a DB-instance of $\langle \Sigma, T \rangle$, or in a Σ -model of T^* are all equivalent and decidable. \square

The instantiation algorithm of Lemma 4.3 can be used to discharge the satisfiability tests in lines 2 and 3 of Algorithm 1 because the conjunction of a state formula and of the negation of a state formula is a $\exists\forall$ -formula (notice that ι is itself the negation of a state formula, according to (11)).

Theorem 4.4 *The backward search algorithm (cf. Algorithm 1) is effective and partially correct for solving the safety problem for RASs.*

Proof. Recall that \mathcal{S} is unsafe iff there is no DB-instance \mathcal{M} of $\langle \Sigma_{ext}, T \rangle$, no $k \geq 0$ and no assignment in \mathcal{M} to the variables $\underline{x}^0, \underline{a}^0 \dots, \underline{x}^k, \underline{a}^k$ such that the formula (14)

$$\iota(\underline{x}^0, \underline{a}^0) \wedge \tau(\underline{x}^0, \underline{a}^0, \underline{x}^1, \underline{a}^1) \wedge \dots \wedge \tau(\underline{x}^{k-1}, \underline{a}^{k-1}, \underline{x}^k, \underline{a}^k) \wedge v(\underline{x}^k, \underline{a}^k)$$

is true in \mathcal{M} . It is sufficient to show that this is equivalent to saying that there is no Σ_{ext} -model \mathcal{M} of T^* , no $k \geq 0$ and no assignment in \mathcal{M} to the variables $\underline{x}^0, \underline{a}^0 \dots, \underline{x}^k, \underline{a}^k$ such that (14) is true in \mathcal{M} (once this is shown, the proof goes in the same way as the proof of Theorem 4.2).

Now, the formula (14) is satisfiable in a Σ_{ext} -structure \mathcal{M} under a suitable assignment iff the formula

$$\begin{aligned} \iota(\underline{x}^0, \underline{a}^0) \wedge \exists \underline{a}^1 \exists \underline{x}^1 (\tau(\underline{x}^0, \underline{a}^0, \underline{x}^1, \underline{a}^1) \wedge \dots \\ \dots \wedge \exists \underline{a}^k \exists \underline{x}^k (\tau(\underline{x}^{k-1}, \underline{a}^{k-1}, \underline{x}^k, \underline{a}^k) \wedge v(\underline{x}^k, \underline{a}^k)) \dots) \end{aligned}$$

is satisfiable in \mathcal{M} under a suitable assignment; by Lemma 4.1, the latter is equivalent to a formula of the kind

$$\iota(\underline{x}, \underline{a}) \wedge \exists \underline{e} \exists \underline{y} \phi(\underline{e}, \underline{y}, \underline{x}, \underline{a}) \tag{26}$$

where $\exists \underline{e} \exists \underline{y} \phi(\underline{e}, \underline{y}, \underline{x}, \underline{a})$ is an extended state formula (thus ϕ is quantifier-free, the \underline{e} are variables of artifact sorts and the \underline{y} are variables of basic sorts - we renamed $\underline{x}^0, \underline{a}^0$ as $\underline{x}, \underline{a}$). However the satisfiability of (26) is the same as the satisfiability of $\exists \underline{e} (\iota(\underline{x}, \underline{a}) \wedge \phi(\underline{e}, \underline{y}, \underline{x}, \underline{a}))$; the latter, in view of (11), is a $\exists\forall$ -formula and so Lemma 4.3 applies and shows that its satisfiability in a DB-instance of $\langle \Sigma_{ext}, T \rangle$ is the same as its satisfiability in a Σ_{ext} -model of T^* . \square

To sum up, in this subsection we remarked that for Algorithm 1, to be effective, we need decision procedures for discharging the satisfiability tests in Lines 2-3. Thanks to the subprocedure $\text{QE}(T^*, \phi)$, the only formulae we need to test in these lines have a specific form (i.e. they are $\exists\forall$ -formulae). In fact, by our hypotheses in Assumption 3.4, we can freely assume that all the runs we are interested in take place inside models of T^* (where we can eliminate quantifiers binding variables of basic sorts). Then, in two first technical lemmas (Lemmas 4.1 and 4.2) we show that the preimage of a state formula is an extended state formula and that such an extended state formula can be converted back (modulo T^*) into a state formula; finally, in a third technical lemma (Lemma 4.3), we

show that entailments between state formulae (more generally, satisfiability of formulae of the kind $\exists\forall$) can be decided via finite instantiation techniques. These observations make both safety and fixpoint tests effective and constitute the skeleton of the proof of Theorem 4.4.

Remark 4.2. Notice that the role of quantifier elimination (Line 6 of Algorithm 1) is twofold: (i) It allows to discharge the fixpoint test of Line 2 (see Lemma 4.3). (ii) It ensures termination in significant cases, namely those where (*strongly*) *local formulae*, introduced in the next section, are involved.

5. Termination Results for RASs

Theorem 4.4 gives a semi-decision procedure for unsafety: if the system is unsafe, the procedure discovers it, but if the system is safe, the procedure (still correct) may not terminate. Termination is much more difficult to achieve for RASs, since acyclicity of Σ seems not to be sufficient to guarantee it. We present two termination results for RASs, both obtained via the use of well quasi-orders. The strategy for proving termination consists of isolating sufficient conditions that imply that the embeddability relation between DB instances is a well-quasi-ordering. Since there is no guarantee that this fact holds in general, RASs are *not* well-structured transition systems.

5.1. Termination with Local Updates

Consider an acyclic signature Σ , a theory T (satisfying our Assumption 3.4), and an artifact setting $(\underline{x}, \underline{a})$ over an artifact extension Σ_{ext} of Σ . We call a state formula *local* if it is a disjunction of the formulae

$$\exists e_1 \cdots \exists e_k (\delta(e_1, \dots, e_k) \wedge \bigwedge_{i=1}^k \phi_i(e_i, \underline{x}, \underline{a})), \quad (27)$$

and *strongly local* if it is a disjunction of the formulae

$$\exists e_1 \cdots \exists e_k (\delta(e_1, \dots, e_k) \wedge \psi(\underline{x}) \wedge \bigwedge_{i=1}^k \phi_i(e_i, \underline{a})). \quad (28)$$

In (27) and (28), δ is a conjunction of variable equalities and disequalities, ϕ_i , ψ are quantifier-free, and e_1, \dots, e_k are individual variables varying over artifact sorts. The key expressivity limitation of local state formulae is that they cannot compare entries belonging to different tuples of artifact relations: in fact, each ϕ_i in (27) and (28) can contain only the existentially quantified variable e_i .

A transition formula $\hat{\tau}$ is *local* (resp., *strongly local*) if whenever a formula ϕ is local (resp., strongly local), so is $Pre(\hat{\tau}, \phi)$ (modulo the axioms of T^*).

Below in Theorem 5.4 we show that (for acyclic Σ) Algorithm 1 terminates when applied to a local safety formula in a RAS whose τ is a disjunction of local transition formulae. Note that Theorem 5.4 can be used to reconstruct (restricted to safety problems) the essence of the decidability results of (Li et al., 2017). Specifically, it can be shown by a direct computation that transitions in (Li et al., 2017) are strongly local which, in turn, can be shown using quantifier elimination (see Appendix F in (Calvanese et al.,

2018b) for all the details, where we also show how to represent transitions from (Li et al., 2017) by the means of existentially quantified data variables). Interestingly, Theorem 5.4 can be applied to more cases not covered in (Li et al., 2017). For example, one can provide transitions enforcing *updates over unboundedly many* tuples (bulk updates) that are strongly local (cf. Appendix F in (Calvanese et al., 2018b)).

Example 5.1. By inspecting the transitions of our running example, one can see that all of them are strongly local. Consequently, it is decidable to check safety of local state formulae. For example, we show that the first transition of Example 4.3 is strongly local (the computations for all the other transitions are analogous, and all the details about the format of transitions that are (strongly) local can be found in Appendix F of (Calvanese et al., 2018b)).

The first transition of Example 4.3 represents the first step of the insertion of an application into the system. This step is executable when the artifact variable $aState$ is `undef`, aims at loading the application data (user ID, job ID and employee ID) into dedicated artifact variables (uId , jId , eId , respectively) and evolves $aState$ into state `received`. Formally, we have:

$$\begin{array}{l} \exists u:\text{UserId}, j:\text{JobCatId}, e:\text{Empld}, c:\text{Complnd} \\ \left(\begin{array}{l} pState = \text{enabled} \wedge aState = \text{undef} \\ \wedge u \neq \text{undef} \wedge j \neq \text{undef} \wedge e \neq \text{undef} \wedge c \neq \text{undef} \\ \wedge who(c) = e \wedge what(c) = j \\ \wedge pState' = \text{enabled} \wedge aState' = \text{received} \\ \wedge uId' = u \wedge jId' = j \wedge eId' = e \wedge cid' = c \end{array} \right) \end{array} \quad (29)$$

For simplicity, we can rewrite Formula 29 into the following equivalent but more succinct formula:

$$\exists \underline{d} \left(\begin{array}{l} \pi(\underline{x}_1, \underline{x}_2) \wedge \psi(\underline{d}) \wedge d_1 = \text{enabled} \wedge d_2 = \text{received} \\ \wedge (\underline{x}'_1 := \underline{x}_1 \wedge \underline{x}'_2 := \underline{d} \wedge \underline{a}' := \underline{a}) \end{array} \right) \quad (30)$$

where $\underline{d} := \langle d_1, d_2, u, j, e, c \rangle$, \underline{x}_1 are the artifact variables of the system that are *not* updated, \underline{x}_2 are the artifact variables of the system that are updated, $\pi(\underline{x}_1, \underline{x}_2)$ and $\psi(\underline{d})$ are quantifier-free Σ -formulae and \underline{a} are the artifact components of the systems.

We show that the preimage along (30) of a strongly local formula is strongly local.

Given a strongly local state formula ϕ , we can easily suppose that ϕ has the following format:

$$\phi := \psi'(\underline{x}) \wedge \exists \underline{i} (\text{AllDiff}(\underline{i}) \wedge \Theta(\underline{a}))$$

where \underline{x} are all the artifact variables of the system, \underline{i} are variables of artifact sorts and Θ is a formula involving all the artifact components \underline{a} .

We compute the preimage $Pre(30, \phi)$:

$$\exists \underline{d} \left(\begin{array}{l} \pi(\underline{x}_1, \underline{x}_2) \wedge \psi(\underline{d}) \wedge d_1 = \text{enabled} \wedge d_2 = \text{received} \\ \wedge (\underline{x}'_1 := \underline{x}_1 \wedge \underline{x}'_2 := \underline{d} \wedge \underline{a}' := \underline{a}) \\ \wedge \psi'(\underline{x}') \wedge \exists \underline{i} (\text{AllDiff}(\underline{i}) \wedge \Theta(\underline{a}')) \end{array} \right) \quad (31)$$

which can be rewritten as follows:

$$\exists \underline{d} \left(\begin{array}{l} \pi(\underline{x}_1, \underline{x}_2) \wedge \psi(\underline{d}) \wedge d_1 = \text{enabled} \wedge d_2 = \text{received} \\ \wedge \psi'(\underline{x}_1, \underline{d}) \wedge \exists i (\text{AllDiff}(i) \wedge \Theta(\underline{a})) \end{array} \right) \quad (32)$$

Now, we can move the existential quantifier $\exists \underline{d}$ in front of $\chi(\underline{d}, \underline{x}_1) := (\psi(\underline{d}) \wedge d_1 = \text{enabled} \wedge d_2 = \text{received} \wedge \psi'(\underline{x}_1, \underline{d}))$. We eliminate the quantifiers (applying the quantifier elimination procedure for T^*) from the subformula $\exists \underline{d}(\chi(\underline{d}, \underline{x}_1))$ obtaining a formula of the kind $\theta(\underline{x}_1)$.

The final result is

$$\pi(\underline{x}_1, \underline{x}_2) \wedge \theta(\underline{x}_1) \wedge \exists i (\text{AllDiff}(i) \wedge \Theta(\underline{a})) \quad (33)$$

which is a strongly local formula.

The interested reader can find additional details about applications of (strongly) local RASs to data-aware business processes in (Calvanese et al., 2019a). Specifically, this paper contains a running example (verified against several properties) that can be represented using a RAS that is strongly-local.

In addition, Theorem 5.4 covers also problems coming from a different source, like coverability problems for broadcast protocols (Esparza et al., 1999; Delzanno et al., 1999): these problems can be encoded using local formulae over the trivial one-sorted signature containing just one basic sort, finitely many constants and one artifact sort with one artifact component. We remark that coverability for broadcast protocols can be decided with a non-primitive recursive lower bound (Schmitz and Schnoebelen, 2013); this proves that our framework is quite expressive (the problems in (Li et al., 2017) have for instance an EXPSpace upper bound). Recalling that (Li et al., 2017) handles verification of LTL-FO, thus going beyond safety problems, this shows that the two settings are incomparable. Finally, notice that Theorem 5.4 implies also the decidability of the safety problem for SASs, in case of Σ acyclic.

Before stating and proving Theorem 5.4, we need to recall some basic facts about well-quasi-orders. Recall that a *well-quasi-order* (wqo) is a set W endowed with a reflexive-transitive relation \leq having the following property: for every infinite succession

$$w_0, w_1, \dots, w_i, \dots$$

of elements from W there are i, j such that $i < j$ and $w_i \leq w_j$. The fundamental result about wqo's is the following theorem, which is a recursive version of Higman's lemma (Higman, 1952) and is a special case of the well-known Kruskal's Tree Theorem (Kruskal, 1960):

Theorem 5.1. If (W, \leq) is a wqo, then so is the partial order of the finite lists over W , ordered by componentwise subword comparison (i.e. $w \leq w'$ iff there is a subword w_0 of w' of the same length as w , such that the i -th entry of w is less or equal to—in the sense of (W, \leq) —the i -th entry of w_0 , for all $i = 0, \dots, |w|$).

Various wqo's can be recognized by applying the above theorem; in particular, the theorem implies that the cartesian product of wqo's is a wqo. As an application, notice that \mathbb{N} is a wqo, hence the following corollary (known as Dickson's Lemma) follows:

Corollary 5.2. The cartesian product of k -copies of \mathbb{N} (and also of $\mathbb{N} \cup \{\infty\}$), with componentwise ordering, is a wqo.

Let $\tilde{\Sigma}$ be $\Sigma_{ext} \cup \{\underline{a}, \underline{x}\}$, that is, Σ_{ext} expanded with function symbols \underline{a} and constants \underline{x} (thus, a $\tilde{\Sigma}$ -structure is a Σ_{ext} -structure endowed with an assignment to \underline{x} and \underline{a} , which were variables and now are treated as symbols of $\tilde{\Sigma}$). For the following, we need the following definition:

Definition 5.3. A $\tilde{\Sigma}$ -structure \mathcal{M} is called *cyclic*⁷ if it is generated by a *single* element $e \in E^{\mathcal{M}}$ (called *generator* of \mathcal{M}), where E is an artifact sort (i.e. e belongs to the interpretation of an artifact sort E).

The previous definition intuitively means that *all* the elements of the cyclic structures are obtained from the generator by applying the function symbols of $\tilde{\Sigma}$ to the generator.

Since Σ is acyclic, so is $\tilde{\Sigma}$, and then one can show that there are only finitely many cyclic $\tilde{\Sigma}$ -structures $\mathcal{C}_1, \dots, \mathcal{C}_N$ up to isomorphism. With a $\tilde{\Sigma}$ -structure \mathcal{M} we associate the tuple of numbers $k_1(\mathcal{M}), \dots, k_N(\mathcal{M}) \in \mathbb{N} \cup \{\infty\}$ counting the numbers of elements generating (as singletons) cyclic substructures isomorphic to $\mathcal{C}_1, \dots, \mathcal{C}_N$, respectively.

Now, we show that, if the tuple associated with \mathcal{M} is component-wise bigger than the one associated with \mathcal{N} , then \mathcal{M} satisfies all the local formulae satisfied by \mathcal{N} .

Lemma 5.1. Let \mathcal{M}, \mathcal{N} be $\tilde{\Sigma}$ -structures. If the inequalities

$$k_1(\mathcal{M}) \leq k_1(\mathcal{N}), \dots, k_N(\mathcal{M}) \leq k_N(\mathcal{N})$$

hold, then all local formulae true in \mathcal{M} are also true in \mathcal{N} .

Proof. Notice that local formulae (viewed in $\tilde{\Sigma}$) are sentences, because they do not have free variable occurrences - the $\underline{a}, \underline{x}$ are now constant function symbols and individual constants, respectively. The proof of the lemma is fairly obvious: notice that, once we assigned some $\alpha(e_i)$ in \mathcal{M} to the variable e_i , the truth of a formula like $\phi(e_i, \underline{x}, \underline{a})$ under such an assignment depends only on the $\tilde{\Sigma}$ -substructure generated by $\alpha(e_i)$, because ϕ is quantifier-free and e_i is the only $\tilde{\Sigma}$ -variable occurring in it. In fact, if a local state formula $\exists e_1 \dots \exists e_k \left(\delta(e_1, \dots, e_k) \wedge \bigwedge_{i=1}^k \phi_i(e_i, \underline{x}, \underline{a}) \right)$ is true in \mathcal{M} , then there exist elements $\bar{e}_1, \dots, \bar{e}_k$ (in the interpretation of some artifact sorts), each of which makes ϕ_i true. Hence, ϕ_i is also true in the corresponding cyclic structure generated by \bar{e}_i . Since $k_1(\mathcal{M}) \leq k_1(\mathcal{N}), \dots, k_N(\mathcal{M}) \leq k_N(\mathcal{N})$ hold, then also in \mathcal{N} there are at least as many elements in the interpretation of artifact sorts as there are in \mathcal{M} that validate all the ϕ_i . Thus, we get that the formula $\exists e_1 \dots \exists e_k \left(\delta(e_1, \dots, e_k) \wedge \bigwedge_{i=1}^k \phi_i(e_i, \underline{x}, \underline{a}) \right)$ is true also in \mathcal{N} , as wanted. \square

Now we are ready to prove our first termination and decidability result.

Theorem 5.4. If Σ is acyclic, backward search (cf. Algorithm 1) terminates when applied to solve the safety problem, with respect to a (*strongly*) local safety formula $v(\underline{x}, \underline{a})$,

⁷ This is unrelated to cyclicity of Σ defined in Section 3, and comes from universal algebra terminology.

for a RAS $\langle \Sigma, T, \Sigma_{ext}, \underline{x}, \underline{a}, \iota(\underline{x}, \underline{a}), \tau(\underline{x}, \underline{a}, \underline{x}', \underline{a}') \rangle$, where τ is a disjunction of (*strongly*) *local* transition formulae.

Proof. Suppose the algorithm does not terminate. Then the fixpoint test of Line 2 fails infinitely often. Recalling that the T -equivalence of B_n and of $\bigvee_{0 \leq k < n} \phi_k$ is an invariant of the algorithm (here ϕ_n, B_n are the status of the variables ϕ, B after n execution of the main loop), this means that there are models

$$\mathcal{M}_0, \mathcal{M}_1, \dots, \mathcal{M}_j, \dots$$

such that for all j , we have that $\mathcal{M}_j \models \phi_j$ and $\mathcal{M}_j \not\models \phi_i$ (all $i < j$). But the ϕ_j are all local formulae, so considering the tuple of cardinals $k_1(\mathcal{M}_j), \dots, k_N(\mathcal{M}_j)$ and Lemma 5.1, we get a contradiction, in view of Dickson's Lemma. This is because, by Dickson's Lemma, $(\mathbb{N} \cup \{\infty\})^N$ is a wqo, so there exist i, j such that $i < j$ and $k_1(\mathcal{M}_i) \leq k_1(\mathcal{M}_j), \dots, k_N(\mathcal{M}_i) \leq k_N(\mathcal{M}_j)$. Using Lemma 5.1, we get that ϕ_i , which is local and true in \mathcal{M}_i , is also true in \mathcal{M}_j , which is a contradiction. \square

5.2. Termination with Tree-like Signatures

Σ is *tree-like* if it is acyclic and all non-leaf nodes have outdegree 1. An artifact setting over Σ is tree-like if $\tilde{\Sigma} := \Sigma_{ext} \cup \{\underline{a}, \underline{x}\}$ is tree-like. In tree-like artifact settings, artifact relations have a single “data” component, and basic relations are unary or binary.

Proving termination for RAS with a tree-like artifact setting is more complex, but follows a similar schema as in the case of local transition formulae.

If (W, \leq) is a partial order, we consider the set $M(W)$ of finite multisets of W as a partial order in the following way:⁸ say that $M \leq N$ holds iff there is an injection $p : M \rightarrow N$ such that $m \leq p(m)$ holds for all $m \in M$ (in other words, p associates with every occurrence of an element m of M an occurrence $p(m)$ of an element of N such that $p(m) \geq m$ - this is moreover done injectively, i.e. in such a way that different occurrences are associated to different occurrences).

Corollary 5.5. If (W, \leq) is a wqo, then so is $(M(W), \leq)$ as defined above.

Proof. This is due to the fact that one can convert a multiset M to a list $L(M)$ so that if $L(M) \leq L(N)$ holds, then also $M \leq N$ holds (such a conversion L can be obtained by ordering the occurrences of elements in M in any arbitrarily chosen way). \square

We assume that the graph $G(\tilde{\Sigma})$ associated to $\tilde{\Sigma}$ is a tree (the generalization to the case where such a graph is a forest is trivial). This means in particular that each sort is the domain of at most one function symbol and that there just one sort which is not the domain of any function symbol (let us call it the *root sort* of $\tilde{\Sigma}$ and let us denote it with S_r).

By induction on the height of a sort S (defined as the length of the longest path from

⁸ This is not the canonical ordering used for multisets, as introduced, e.g., in (Baader and Nipkow, 1998).

S to a leaf) in the above graph, we define a wqo $w(S)$ (in the definition we use the fact the cartesian product of wqo's is a wqo and Corollary 5.5). Let S_1, \dots, S_n be the sons of S in the tree; put

$$w(S) := M(w(S_1)) \times \dots \times M(w(S_n)) \quad (34)$$

(thus, if S is a leaf, $w(S)$ is the trivial one-element wqo - its only element is the empty tuple).

Let now \mathcal{M} be a finite $\tilde{\Sigma}$ -structure; we indicate with $S^{\mathcal{M}}$ the interpretation in \mathcal{M} of the sort S (it is a finite set). For $a \in S^{\mathcal{M}}$, we define $M_{\mathcal{M}}(a) \in w(S)$, again by induction on the height of S . Suppose that S_1, \dots, S_n are the sons of S and that the arc from S_i to S is labeled by the function symbol f_i ; then we put

$$M_{\mathcal{M}}(a) := \langle \{M_{\mathcal{M}}(b_1) \mid b_1 \in S_1^{\mathcal{M}} \text{ and } f_1^{\mathcal{M}}(b_1) = a\}, \dots \\ \dots, \{M_{\mathcal{M}}(b_n) \mid b_n \in S_n^{\mathcal{M}} \text{ and } f_n^{\mathcal{M}}(b_n) = a\} \rangle$$

where $f_i^{\mathcal{M}}$ ($i = 1, \dots, n$) is the interpretation of the symbol f_i in \mathcal{M} .

Moreover, for every sort S , we let

$$M_{\mathcal{M}}(S) := \{M_{\mathcal{M}}(a) \mid a \in S^{\mathcal{M}}\} \quad . \quad (35)$$

Finally, we define

$$M(\mathcal{M}) := M_{\mathcal{M}}(S_r) \quad . \quad (36)$$

For termination, the relevant lemma is the following:

Lemma 5.2. Suppose that $\tilde{\Sigma}$ is tree-like and does not contain constant symbols; given two finite $\tilde{\Sigma}$ -structures \mathcal{M} and \mathcal{N} , we have that if $M(\mathcal{M}) \leq M(\mathcal{N})$, then \mathcal{M} embeds into \mathcal{N} . As a consequence, the finite $\tilde{\Sigma}$ -structures are a wqo with respect to the embeddability quasi-order.

Proof. Again, we make an induction on the height of S , proving the claim for the subsignature of $\tilde{\Sigma}$ having S as a root (let us call this the S -subsignature).

Let \mathcal{M} be a model over the S -subsignature. For every $a \in S^{\mathcal{M}}$, and for every $f_i : S_i \rightarrow S$, if we restrict \mathcal{M} to the elements in the f_i -fibers of a , we get a model $\mathcal{M}_{f_i, a}$ for the S_i -subsignature (an element $c \in \tilde{S}^{\mathcal{M}}$ is in the f_i -fiber of a if, taking the term t corresponding to the composition of the functions symbols going from \tilde{S} to S_i , we have that $f_i^{\mathcal{M}}(t^{\mathcal{M}}(c)) = a$). In addition, if $M_{\mathcal{M}}(a) = (M_1, \dots, M_n)$, then $M_i = M(\mathcal{M}_{f_i, a})$ by definition. Finally, observe that the restriction of \mathcal{M} to the S_i -subsignature is the disjoint union of the f_i -fibers models $\mathcal{M}_{f_i, a}$, varying $a \in S^{\mathcal{M}}$.

Suppose now that \mathcal{M}, \mathcal{N} are models over the S -subsignature such that $M(\mathcal{M}) \leq M(\mathcal{N})$; this means that we can find an injective map μ mapping $S^{\mathcal{M}}$ into $S^{\mathcal{N}}$ so that $M_{\mathcal{M}}(a) \leq M_{\mathcal{N}}(\mu(a))$. If $M_{\mathcal{M}}(a) = (M_1, \dots, M_n)$ and $M_{\mathcal{N}}(\mu(a)) = (N_1, \dots, N_n)$, we then have that $M_i \leq N_i$ for every $i = 1, \dots, n$. Considering that, as noticed above, $M_i = M(\mathcal{M}_{f_i, a})$ and $N_i = M(\mathcal{N}_{f_i, \mu(a)})$, by induction hypothesis, we have embeddings $\nu_{i, a}$ for the f_i -fibers models of a and $\mu(a)$ (for every $a \in S^{\mathcal{M}}$ and $i = 1, \dots, n$). Glueing these embeddings to the disjoint union (varying i, a) and adding them μ as S -component, we get the desired embedding of \mathcal{M} into \mathcal{N} . \square

Theorem 5.6. Backward search (cf. Algorithm 1) terminates when applied to a safety problem in a RAS with a tree-like artifact setting.

Proof. For simplicity, we start giving the argument for the case where we do not have constants and artifact variables. Similarly to the proof of Theorem 5.4, suppose the algorithm does not terminate. Then the fixpoint test of Line 2 fails infinitely often. Recalling that the T -equivalence of B_n and of $\bigvee_{0 \leq k < n} \phi_k$ is an invariant of the algorithm (here ϕ_n, B_n are the status of the variables ϕ, B after n execution of the main loop), this means that there are models

$$\mathcal{M}_0, \mathcal{M}_1, \dots, \mathcal{M}_j, \dots$$

such that for all j , we have that $\mathcal{M}_j \models \phi_j$ and $\mathcal{M}_j \not\models \phi_i$ (all $i < j$). The models can be taken to be all finite, by Lemma 4.3. But the ϕ_j are all existential sentences in $\tilde{\Sigma}$, so this is incompatible to the fact that, by Lemma 5.2, there are $i < j$ with \mathcal{M}_i embeddable into \mathcal{M}_j .

Concerning the general case, it is sufficient to consider the following observation that shows how to extend the proof to the case where we have constants and artifact variables. Recall that in $\tilde{\Sigma}$ the artifact variables are seen as constants, so we need to consider only the case of constants. Let $\tilde{\Sigma}^+$ be $\tilde{\Sigma}$ where each constant symbol c of sort S is replaced by a new sort S_c and a new function symbol $f_c : S_c \rightarrow S$. Now every model \mathcal{M} of $\tilde{\Sigma}$ can be transformed into a model \mathcal{M}^+ of $\tilde{\Sigma}^+$ by interpreting S_c as a singleton set $\{*\}$ and f_c as the map sending $*$ to $c^{\mathcal{M}}$. This transformation has the following property: $\tilde{\Sigma}$ -embeddings of \mathcal{M} into \mathcal{N} are in bijective correspondence with $\tilde{\Sigma}^+$ -embeddings of \mathcal{M}^+ into \mathcal{N}^+ . Since $\tilde{\Sigma}^+$ is still tree-like and does not have constant symbols, this shows that Theorem 5.6 holds for $\tilde{\Sigma}$ too. \square

While tree-like RAS restrict artifact relations to be unary, their transitions are not subject to any locality restriction. This allows for expressing rich forms of updates, including general bulk updates (which allow us to capture non-primitive recursive verification problems⁹) and transitions comparing at once different tuples in artifact relations. The flight management process presented in the following example shows these advanced features, with a tree-like RAS whose safety verification is indeed decidable. Finally, notice that tree-like RASs are incomparable: (i) with the “tree” classes of (Bojańczyk et al., 2013), since the former use artifact relations, whereas the latter only individual variables; (ii) with the decidability class of (Li et al., 2017), since tree-like RASs express transitions able to compare at once values stored in different tuples in artifact relations.

Example 5.2.

We consider a simple RAS that falls in the scope of the tree-like decidability result. Specifically, this example has a tree-like artifact setting (see Figure 2), thus assuring that, when solving the safety problem for it, the backward search algorithm is guaranteed to terminate. Note, however, that the termination result adopted here is the one of Theorem 5.6 due to the non-locality of certain transitions, as explained in detail below.

⁹ Notice that the artifact setting described above to capture coverability problems for broadcast protocols is both local and tree-like.




Figure 2. A characteristic graph of the flight management process, where blue and yellow boxes respectively represent basic and artifact sorts.

The flight management process represents a simplified version of a flight management system adopted by an airline. To prepare a flight, the company picks a corresponding destination (that meets the aviation safety compliance indications) and consequently reports on a number of passengers that are going to attend the flight. Then, an airport dispatcher may pick a manned flight and put it in the airports flight plan. In case the flight destination becomes unsafe (e.g., it was struck by a hurricane or the hosting airport had been seized by terrorists), the dispatcher uses the system to inform the airline about this condition. In turn, the airline notifies all the passengers of the affected destination about the contingency, and temporary cancels their flights.

To formalize these different aspects, we make use of a DB signature Σ_{fm} that consists of: (i) two id sorts, used to identify flights and cities; (ii) one function symbol $destination : \text{FlightId} \rightarrow \text{CityId}$ mapping flight identifiers to their corresponding destinations (i.e., city identifiers). Note that, in a classical relational model (cf. Section 3.1), our signature would contain two relations: one binary R_{FlightId} that defines flights and their destinations, and another unary R_{CityId} identifying cities, that are referenced by R_{FlightId} using $destination$.

We assume that the read-only flight management database contains data about at least one flight and one city. To start the process, one needs at least one city to meet the aviation safety compliances. It is assumed that, initially, all the cities are unsafe. An airport dispatcher, at once, may change the safety status only of one city.

We model this action by performing two consequent actions. First, we select the city identifier and store it in the designated artifact variable $safeCityId$:

$$\exists c:\text{CityId} (c \neq \text{undef} \wedge safeCityId = \text{undef} \wedge safeCityId' = c)$$

Then, we place the extracted city identifier into a unary artifact relation $safeCity : \text{CityIndex} \rightarrow \text{CityId}$, that is used to represent safe cities and where CityIndex is its artifact sort.

$$\exists i:\text{CityIndex} \left(\begin{array}{l} safeCity[i] = \text{undef} \wedge safeCityId \neq \text{undef} \wedge safeCityId' = \text{undef} \\ \wedge safeCity' = \lambda j. \left(\begin{array}{l} \text{if } j = i \text{ then } safeCityId \\ \text{else if } safeCity[j] = safeCityId \text{ then } \text{undef} \\ \text{else } safeCity[j] \end{array} \right) \end{array} \right)$$

Note that the two previous transitions can be rewritten as a unique one, hence showing a more compact way of specifying RAS transitions. This, in turn, can augment the performance of the verifier while working with large-scale cases. The unified transition actually looks as follows:

$$\left(\begin{array}{l} \exists c:\text{CityId}, \exists i:\text{CityIndex} \\ c \neq \text{undef} \wedge \text{safeCity}[i] = \text{undef} \\ \wedge \text{safeCity}' = \lambda j. \left(\begin{array}{l} \text{if } j = i \text{ then } c \\ \text{else if } \text{safeCity}[j] = c \text{ then } \text{undef} \\ \text{else } \text{safeCity}[j] \end{array} \right) \end{array} \right)$$

Then, to register passengers with booked tickets on a flight, the airline needs to make sure that a corresponding flight destination is actually safe. To perform the passenger registration, the airline selects a flight identifier that is assigned to the route and uses it to populate entries in an unary artifact relation $\text{regdPassenger} : \text{PassengerIndex} \rightarrow \text{FlightId}$. Note that there may be more than one passenger taking the flight, and therefore, more than one entry in regdPassenger with the same flight identifier.

$$\left(\begin{array}{l} \exists i:\text{CityIndex}, f:\text{FlightId}, p:\text{PassengerIndex} \\ f \neq \text{undef} \wedge \text{destination}(f) = \text{safeCity}[i] \wedge \text{regdPassenger}[p] = \text{undef} \\ \wedge \text{regdPassenger}' = \lambda j. \left(\begin{array}{l} \text{if } j = p \text{ then } f \\ \text{else } \text{regdPassenger}[j] \end{array} \right) \end{array} \right)$$

We also assume that the airline owns aircraft of one type that can contain no more than k passengers. In case there were more than k passengers registered on the flight, the airline receives a notification about its overbooking and temporarily suspends all passenger registrations associated to this flight. This is modelled by checking whether there are at least $k + 1$ entries in regdPassenger . If so, the flight identifier is added to a unary artifact relation $\text{overbooked} : \text{FlightIndex} \rightarrow \text{FlightId}$ and all the passenger registrations in regdPassenger that reference this flight identifier are nullified by updating unboundedly many entries in the corresponding artifact relation:¹⁰

$$\left(\begin{array}{l} \exists p_1:\text{PassengerIndex}, \dots, p_{k+1}:\text{PassengerIndex}, m:\text{FlightIndex} \\ \left(\bigwedge_{i,i' \in \{1, \dots, k+1\}, i \neq i'} (p_i \neq p_{i'} \wedge \text{regdPassenger}[p_i] \neq \text{undef} \wedge \text{regdPassenger}[p_i] = \text{regdPassenger}[p_{i'}]) \right) \\ \wedge \text{overbooked}[m] = \text{undef} \\ \wedge \text{regdPassenger}' = \lambda j. \left(\begin{array}{l} \text{if } \text{regdPassenger}[j] = \text{regdPassenger}[p_1] \text{ then } \text{undef} \\ \text{else } \text{regdPassenger}[j] \end{array} \right) \\ \wedge \text{overbooked}'[m] = \text{regdPassenger}[p_1] \end{array} \right)$$

Notice that this transition is not local, since its guard contains literals of the form $\text{regdPassenger}[p_i] = \text{regdPassenger}[p_{i'}]$ (with $p_i \neq p_{i'}$), which involve more than one element of one artifact sort.

In case of any contingency, the airport dispatcher may change the city status from *safe* to *unsafe*. To do it, we first select one of the safe cities, make it unsafe (i.e., remove it

¹⁰ For simplicity of presentation, we simply remove such data from the artifact relation. In a real setting, this information would actually be transferred to a dedicated, historical table, so as to reconstruct the status of past, overbooked flights.

from *safeCity* relation) and store its identifier in the artifact variable *unsafeCityId*:

$$\exists i:\text{CityIndex} \left(\begin{array}{l} \text{unsafeCityId} = \text{undef} \wedge \text{safeCity}[i] \neq \text{undef} \wedge \\ \wedge \text{unsafeCityId}' = \text{safeCity}[i] \wedge \text{safeCity}'[i] = \text{undef} \end{array} \right)$$

Then, we use the remembered city identifier to cancel all the passenger registrations for flights that use this city as their destination:

$$\left(\begin{array}{l} \text{unsafeCityId} \neq \text{undef} \wedge \text{unsafeCityId}' = \text{undef} \wedge \\ \wedge \text{regdPassenger}' = \lambda j. \left(\begin{array}{l} \text{if } \text{destination}(\text{regdPassenger}[j]) = \text{unsafeCityId} \text{ then } \text{undef} \\ \text{else } \text{regdPassenger}[j] \end{array} \right) \end{array} \right)$$

Similarly to the previous case, this transition performs the intended action by updating unboundedly many entries in the artifact relation.

Also in this case, we can shrink the last two transitions into a single transition:

$$\exists i:\text{CityIndex} \left(\begin{array}{l} \text{safeCity}[i] \neq \text{undef} \wedge \\ \wedge \text{regdPassenger}' = \lambda j. \left(\begin{array}{l} \text{if } \text{destination}(\text{regdPassenger}[j]) = \text{safeCity}[i] \text{ then } \text{undef} \\ \text{else } \text{regdPassenger}[j] \end{array} \right) \end{array} \right)$$

However, as in the previous case, the transition turns out to be not local. Specifically, it is due to the literal $\text{destination}(\text{regdPassenger}[j]) = \text{safeCity}[i]$ that involves more than one element of (different) artifact sorts.

6. First experiments

We implemented a prototype of our backward reachability algorithm for artifact systems on top of the MCMT model checker, extending it with the features required to formalize and verify RASs. MCMT manages verification in the infinite-state case by exploiting as its model-theoretic framework the declarative formalism of array-based systems. Since their first introduction in (Ghilardi et al., 2008a; Ghilardi and Ranise, 2010a), array-based systems have been provided with various implementations of the standard backward reachability algorithms (including more sophisticated variants and heuristics). Starting from its first version (Ghilardi and Ranise, 2010b), MCMT was successfully applied to cache coherence and mutual exclusions protocols (Ghilardi and Ranise, 2010a), timed (Carioni et al., 2010) and fault-tolerant (Alberti et al., 2012b; Alberti et al., 2010) distributed systems, and then to imperative programs (Alberti et al., 2014b; Alberti et al., 2017); interesting case studies concerned waiting time bounds synthesis in parameterized timed networks (Bruttomesso et al., 2012) and internet protocols (Bruschi et al., 2017). Further related tools include SAFARI (Alberti et al., 2012a) and ASASP (Alberti et al., 2011); finally, (Conchon et al., 2012; Conchon et al., 2015; Conchon et al., 2013; Conchon et al., 2018a) implement the array-based setting on a parallel architecture with further powerful extensions.

The work principle of MCMT is rather simple: the tool generates the proof obligations arising from the safety and fixpoint tests in backward search (Lines 2-3 of Algorithm 1) and passes them to the background SMT-solver (currently it is YICES (Dutertre and De Moura, 2006)). In practice, the situation is more complicated because SMT-solvers are quite efficient in handling satisfiability problems in combined theories at quantifier-free level, but may encounter difficulties with quantifiers. For this reason, MCMT implements

	Example	#(AC)	#(AV)	#(T)
E1	JobHiring	9	18	15
E2	Acquisition-following-RFQ	6	13	28
E3	Book-Writing-and-Publishing	4	14	13
E4	Customer-Quotation-Request	9	11	21
E5	Patient-Treatment-Collaboration	6	17	34
E6	Property-and-Casualty-Insurance-Claim-Processing	2	7	15
E7	Amazon-Fulfillment	2	28	38
E8	Incident-Management-as-Collaboration	3	20	19

Table 1. Summary of the experimental examples

modules for *quantifier elimination* and *quantifier instantiation*. A specific module for the quantifier elimination problems mentioned in Line 6 of Algorithm 1 has been added to Version 2.8 of MCMT.

We base our experimental evaluation on the already existing benchmark provided in (Li et al., 2017), that samples 32 real-world BPMN workflows published at the official BPM website (<http://www.bpmn.org/>). Specifically, we select seven examples of varying complexity (see Table 1) and provide their faithful encoding in the array-based specification using MCMT Version 2.8. Moreover, we enrich our experimental set with an extended version of the running example of this paper (see Appendix A.1 of (Calvanese et al., 2018b)). Each example has been checked against at least one safe and one unsafe conditions. Since MCMT performs safety verification parameterized on the read-only DB, the result is *independent* on the size of specific DB instances. Moreover, MCMT can in principle handle unbounded DB schemas and unboundedly many DB constants: we have ascertained that the size of the DB schema does not affect the performances as much as the number of transitions involved in the verified RAS. We leave for future work a systematic experimental evaluation of those preliminary observations.

Experiments were performed on a machine with Ubuntu 16.04, 2.6 GHz Intel Core i7 and 16 GB RAM. The benchmark set is available as part of the last distribution 2.8 of MCMT

<http://users.mat.unimi.it/users/ghilardi/mcmt/>

(see the subdirectory `/examples/dbdriven` of the distribution); the user manual, also included in the distribution, contains a dedicated section (pages 36–39) giving essential information on how to use the capabilities of the new version of MCMT (by activating the “`db_driven`” mode), how to encode RASs in MCMT specifications and how to produce user-defined examples in the database driven framework. All the verified examples include transitions with quantified “data” variables, and rely on and the algebraic framework of DB theories introduced in the paper. Consequently, the experiments were carried out on the new version of MCMT and could not be verified with the previous versions.

In Table 1, the columns **#(AV)**, **#(AC)** and **#(T)** represent, respectively, the number of artifact variables, artifact components and transitions used in the example specification; in Table 2, the column **Time** is the MCMT execution time. The most critical measures in Table 2 are **#(N)**, **depth** and **#(SMT-calls)** that respectively define the number of nodes and the depth of the tree used for the backward reachability procedure

Example	Property	Result	Time	#(N)	depth	#(SMT-calls)
E1	E1P1	SAFE	0.06	3	3	1238
	E1P2	UNSAFE	0.36	46	10	2371
	E1P3	UNSAFE	0.50	62	11	2867
	E1P4	UNSAFE	0.35	42	10	2237
E2	E2P1	SAFE	0.72	50	9	3156
	E2P2	UNSAFE	0.88	87	10	4238
	E2P3	UNSAFE	1.01	92	9	4811
	E2P4	UNSAFE	0.83	80	9	4254
E3	E3P1	SAFE	0.05	1	1	700
	E3P2	UNSAFE	0.06	14	3	899
E4	E4P1	SAFE	0.12	14	6	1460
	E4P2	UNSAFE	0.13	18	8	1525
E5	E5P1	SAFE	4.11	57	9	5618
	E5P2	UNSAFE	0.17	13	3	2806
E6	E6P1	SAFE	0.04	7	4	512
	E6P2	UNSAFE	0.08	28	10	902
E7	E7P1	SAFE	1.00	43	7	5281
	E7P2	UNSAFE	0.20	7	4	3412
E8	E8P1	SAFE	0.70	77	11	3720
	E8P2	UNSAFE	0.15	25	7	1652

Table 2. *Experimental results for safety properties*

adopted by MCMT, and the number of the SMT-solver calls. Indeed, MCMT computes the iterated preimages of the formula describing the unsafe states along the various transitions. Such computation produces a tree, whose nodes are labelled by formulae describing sets of states that can reach an unsafe state and whose arcs are labelled by a transition.

To stress test our encoding, we came up with a few formulae describing unsafe configurations (sets of “bad” states), that is, the configurations that the system should not incur throughout its execution. **Property** references encodings of examples endowed with specific (un)safety properties done in MCMT, whereas **Result** shows their verification outcome that can be of the two following types: **SAFE** and **UNSAFE**. The MCMT tool returns **SAFE**, if the undesirable property it was asked to verify represents a configuration that the system cannot reach. At the same time, the result is **UNSAFE** if there exists a path of the system execution that reaches “bad” states.

To conclude, we would like to point out that seemingly high number of SMT solver calls in **#(SMT-calls)** against relatively small execution time demonstrates that MCMT could be considered as a promising tool supporting the presented line of research. This is due to the following two reasons. On the one hand, the SMT technology underlying solvers like YICES (Dutertre and De Moura, 2006) is quite mature and impressively well-performing. On the other hand, the backward reachability algorithm generates proof obligations which are relatively easy to be analyzed as (un)satisfiable by the solver.

A thorough comparison with VERIFAS (Li et al., 2017) is at the moment rather problematic, for various reasons, due for instance to the different specification languages and to the different set of benchmarks covered. In fact, the two systems tackle incomparable verification problems: on the one hand, we deal with safety problems, whereas VERIFAS handles more general LTL-FO properties. On the other hand, we tackle features not avail-

able in VERIFAS, like bulk updates and comparisons between artifact tuples. We leave this for future, more experimentally oriented, work: the comparison might be interesting because the two tools apply quite different technologies (VERIFAS is based on VASS encoding, whereas MCMT follows a purely declarative paradigm). We just point out that Table 2 shows the very encouraging results: in fact, MCMT seems to effectively handle the benchmark with a similar performance to that shown in other, well-established settings, with verification times below 1s in most cases.

7. Conclusions and Future Work

We have laid the foundations of SMT-based verification for artifact systems, focusing on safety problems and relying on array-based systems as underlying inspiring model. We have shown how to overcome the main technical difficulty arising from this approach, namely reconstructing quantifier elimination techniques in the rich setting of artifact systems, using the model-theoretic machinery of model completions. On top of this framework, we have identified three classes of systems for which safety is decidable, which impose different combinations of restrictions on the form of actions and the shape of DB constraints. The presented techniques have been implemented on top of the well-established MCMT model checker, making our approach fully operational. Notably, the machinery presented in this paper has been already employed to formalize and verify a data-aware extension of the de-facto process modeling standard BPMN (Calvanese et al., 2019a), starting a line of research that aims at transferring our technical results into practical, end user- oriented settings.

It is important to stress that the artifact systems we study here are radically different from other formal models integrating dynamics with data, such as Data Petri Nets (Lazic et al., 2008), ν -PNs (Rosa-Velardo and de Frutos-Escrig, 2011) and multiset rewriting systems with data and constraints (Delzanno, 2002). Let us consider Data Petri Nets as a representative example of this class of approaches. In Data Petri Nets, one can generate tokens that carry fresh values not already present in the current marking. The requirement that a value is fresh can be encoded in the model (Rosa-Velardo and de Frutos-Escrig, 2011). Such values can only be mutually related using the comparison predicates over the underlying domain. In our setting, instead, the working memory of an artifact system may contain data elements arbitrarily taken from value sorts, or extracted from the (active domain of a) read-only DB. When loading a data element from a value sort, this may or not be present in the active domain, and thus it may be possibly fresh (but notice that freshness cannot be enforced in our model). When loading data elements from the read-only DB, it is crucial to consider that they are mutually related via constraints present therein. These constraints are primary keys, foreign keys, and additional axioms present in the DB theory. The read-only DB it is fixed within a run, but model checking of safety properties is studied parametrically with respect to all possible read-only DBs over a given schema. During model checking, we are examining sets of reachable states described by logical formulae, whose validity depends on properties that might happen to be true in the read-only DB, depending on the constraints present therein. To handle data elements coming from the read-only DB and their corresponding constraints, we

therefore need a specialized machinery that is different from the one typically used to tame the infinity brought by freshness. In fact, it is not enough to embed the read-only DB in a larger model that admits fresh values to obtain quantifier elimination, which is essential in our model checking algorithm. Quantifier elimination becomes available only when such a “larger model” possesses suitable model-theoretic properties, which we have studied in the paper. In particular, we have argued that such properties are captured by the well-known model-theoretic notion of existentially closed structure and its intimately related notion of model completion. Notably, resorting to model completion can be seen as the “most natural” way to obtain quantifier elimination, as it is the “closest” theory to the original one that at once admits quantifier elimination and preserves satisfiability of existential formulae. This is precisely what we intensively exploit in our model checking algorithm.

From the *foundational* point of view, it is an open, non-trivial research question to see whether our framework and Data Petri nets (or similar approaches) can be inter-reduced to each other when we restrict our attention to the three decidable fragments studied in the paper. We are also interested in using the present contribution as the starting point for a full line of research dedicated to SMT-based techniques for the effective verification of data-aware processes, considering richer forms of verification going beyond safety, and richer classes of artifact systems incorporating concrete data types and arithmetic operations. We also intend to investigate more integrity constraints used in database theory: this study should extend decidability and model-completeness results beyond the cases covered in this paper. In addition, it would be interesting to study whether the decidable classes considered here are tight, or whether interesting variations can be found for which decidability is preserved, possibly guaranteeing termination of the our backward reachability procedure.

From the *algorithmic* point of view, we intend to develop more sophisticated techniques for the quantifier elimination module required in our backward reachability procedure. A first stepping stone in this direction, relying on a constrained version of the Superposition Calculus (SC) (Nieuwenhuis and Rubio, 2001), can be found in (Calvanese et al., 2019c): indeed, thanks to our constrained version of the SC, suitably combined with congruence closure, we show that it is possible to obtain a quadratic bound for the complexity of the quantifier elimination procedure in the case of interest for our applications (Calvanese et al., 2019c; Calvanese et al., 2018a).

As for *experiments*, we aim at building on the encouraging results reported here towards an extensive experimental evaluation of our approach, using the benchmark of (Li et al., 2017) and the concrete specification language in (Calvanese et al., 2019a) as a starting point. A natural next step is then to study how the computation of over-approximations, abstractions and invariants (a capability that MCMT already supports but that should be adapted to the “db_driven” mode) and well-established techniques for SMT-based model checking like CEGAR (McMillan, 2006; Alberti et al., 2014a) and IC3 (Bradley, 2011; Hoder and Bjørner, 2012; Bradley, 2012) can be used to speed up the verification of artifact systems.

Financial Support. This research has been partially supported by the UNIBZ CRC

project *REKAP: Reasoning and Enactment for Knowledge-Aware Processes* and by the CHIST-ERA project *PACMEL: Process-aware Analytics Support based on Conceptual Models for Event Logs*.

References

- Abdulla, P. A., Cerans, K., Jonsson, B., and Tsay, Y.-K. (1996). General decidability theorems for infinite-state systems. In *Proceedings of the 11th Annual IEEE Symposium on Logic in Computer Science (LICS)*, pages 313–321.
- Alberti, F., Armando, A., and Ranise, S. (2011). ASASP: Automated symbolic analysis of security policies. In *Proceedings of the 23rd International Conference on Automated Deduction (CADE)*, volume 6803 of *LNCS (LNAI)*, pages 26–33. Springer.
- Alberti, F., Bruttomesso, R., Ghilardi, S., Ranise, S., and Sharygina, N. (2012a). SAFARI: SMT-based abstraction for arrays with interpolants. In *Proceedings of the 24th International Conference on Computer Aided Verification (CAV)*, volume 7358 of *LNCS*, pages 679–685. Springer.
- Alberti, F., Bruttomesso, R., Ghilardi, S., Ranise, S., and Sharygina, N. (2014a). An extension of lazy abstraction with interpolation for programs with arrays. *Formal Methods in System Design*, 45(1):63–109.
- Alberti, F., Ghilardi, S., Pagani, E., Ranise, S., and Rossi, G. P. (2010). Brief announcement: Automated support for the design and validation of fault tolerant parameterized systems - A case study. In *Proceeding of 24th International Symposium on Distributed Computing DISC*, volume 6343 of *LNCS*, pages 392–394. Springer.
- Alberti, F., Ghilardi, S., Pagani, E., Ranise, S., and Rossi, G. P. (2012b). Universal guards, relativization of quantifiers, and failure models in Model Checking Modulo Theories. *Journal on Satisfiability, Boolean Modeling and Computation*, 8(1/2):29–61.
- Alberti, F., Ghilardi, S., and Sharygina, N. (2014b). Booster: An acceleration-based verification framework for array programs. In *Proceedings of the 12th International Symposium on Automated Technology for Verification and Analysis (ATVA)*, volume 8837 of *LNCS*, pages 18–23. Springer.
- Alberti, F., Ghilardi, S., and Sharygina, N. (2017). A framework for the verification of parameterized infinite-state systems. *Fundamenta Informaticae*, 150(1):1–24.
- Baader, F., Ghilardi, S., and Tinelli, C. (2006). A new combination procedure for the word problem that generalizes fusion decidability results in modal logics. *Information and Computation*, pages 1413–1452.
- Baader, F. and Nipkow, T. (1998). *Term Rewriting and All That*. Cambridge University Press.
- Bojańczyk, M., Segoufin, L., and Toruńczyk, S. (2013). Verification of database-driven systems via amalgamation. In *Proceedings of the 32nd ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems (PODS)*, pages 63–74.
- Bradley, A. R. (2011). SAT-based model checking without unrolling. In *Proceedings of the 12th International Conference on Verification, Model Checking, and Abstract Interpretation VMCAI*, volume 6538 of *LNCS*, pages 70–87. Springer.
- Bradley, A. R. (2012). IC3 and beyond: Incremental, inductive verification. In *Proceedings of the 24th International Conference on Computer Aided Verification (CAV)*, volume 7358 of *LNCS*, page 4. Springer.
- Bradley, A. R. and Manna, Z. (2007). *The Calculus of Computation - Decision Procedures with Applications to Verification*. Springer.

- Bruschi, D., Di Pasquale, A., Ghilardi, S., Lanzi, A., and Pagani, E. (2017). Formal verification of ARP (address resolution protocol) through SMT-based model checking - A case study. In *Proceedings of the 13th International Conference on Integrated Formal Methods (IFM)*, volume 10510 of *LNCS*, pages 391–406. Springer.
- Bruttomesso, R., Carioni, A., Ghilardi, S., and Ranise, S. (2012). Automated analysis of parametric timing-based mutual exclusion algorithms. In *Proceedings of the 4th International Symposium on NASA Formal Methods (NFM)*, volume 7226 of *LNCS*, pages 279–294. Springer.
- Calvanese, D., De Giacomo, G., and Montali, M. (2013). Foundations of data-aware process analysis: A database theory perspective. In *Proceedings of the 32nd ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems (PODS)*, pages 1–12.
- Calvanese, D., Ghilardi, S., Gianola, A., Montali, M., and Rivkin, A. (2018a). Quantifier elimination for database driven verification. Technical Report arXiv:1806.09686, arXiv.org.
- Calvanese, D., Ghilardi, S., Gianola, A., Montali, M., and Rivkin, A. (2018b). Verification of data-aware processes via array-based systems (extended version). Technical Report arXiv:1806.11459, arXiv.org.
- Calvanese, D., Ghilardi, S., Gianola, A., Montali, M., and Rivkin, A. (2019a). Formal modeling and SMT-based parameterized verification of data-aware BPMN. In *Proceeding of the 17th International Conference on Business Process Management (BPM)*, volume 11675 of *LNCS*, pages 157–175. Springer.
- Calvanese, D., Ghilardi, S., Gianola, A., Montali, M., and Rivkin, A. (2019b). From model completeness to verification of data aware processes. In *Description Logic, Theory Combination, and All That*, volume 11560 of *LNCS*, pages 212–239. Springer.
- Calvanese, D., Ghilardi, S., Gianola, A., Montali, M., and Rivkin, A. (2019c). Model completeness, covers and superposition. In *Proceedings of 27th International Conference on Automated Deduction (CADE)*, volume 11716 of *LNCS (LNAI)*, pages 142–160. Springer.
- Calvanese, D., Ghilardi, S., Gianola, A., Montali, M., and Rivkin, A. (2019d). Verification of data-aware processes: Challenges and opportunities for automated reasoning. In *Proceedings of the 2nd International Workshop on Automated Reasoning: Challenges, Applications, Directions, Exemplary Achievements (ARCADE)*, volume 311. EPTCS.
- Carioni, A., Ghilardi, S., and Ranise, S. (2010). MCMT in the land of parametrized timed automata. In *Proceedings of the 6th International Verification Workshop (VERIFY)*, pages 47–64.
- Chang, C.-C. and Keisler, J. H. (1990). *Model Theory*. North-Holland Publishing Co.
- Cimatti, A., Stojic, I., and Tonetta, S. (2018). Formal specification and verification of dynamic parametrized architectures. In *Proceedings of the 22nd International Symposium on Formal Methods (FM)*, volume 10951 of *LNCS*, pages 625–644. Springer.
- Conchon, S., A.Goel, Krstic, S., Mebsout, A., and Zaïdi, F. (2013). Invariants for finite instances and beyond. In *Proceedings of the International Conference on Formal Methods in Computer-Aided Design (FMCAD)*, pages 61–68.
- Conchon, S., Declerck, D., and Zaidi, F. (2018a). Cubicle-W: Parameterized model checking on weak memory. In *Proceedings of the 9th International Joint Conference on Automated Reasoning (IJCAR)*, volume 10900 of *LNCS (LNAI)*, pages 152–160. Springer.
- Conchon, S., Delzanno, G., and Ferrando, A. (2018b). Declarative parameterized verification of topology-sensitive distributed protocols. In *Proceedings of the 6th International Conference on Networked Systems (NETYS)*, volume 11028 of *LNCS*, pages 209–224. Springer.
- Conchon, S., Goel, A., Krstic, S., Mebsout, A., and Zaïdi, F. (2012). Cubicle: A parallel SMT-based model checker for parameterized systems - Tool paper. In *Proceedings of the 24th*

- International Conference on Computer Aided Verification (CAV)*, volume 7358 of *LNCS*, pages 718–724. Springer.
- Conchon, S., Mebsout, A., and Zaïdi, F. (2015). Certificates for parameterized model checking. In *Proceeding of the 20th International Symposium on Formal Methods (FM)*, volume 9109 of *LNCS*, pages 126–142. Springer.
- Damaggio, E., Deutsch, A., and Vianu, V. (2012). Artifact systems with data dependencies and arithmetic. *ACM Transactions on Database Systems*, 37(3):22:1–22:36.
- Delzanno, G. (2002). An overview of MSR(C): A CLP-based framework for the symbolic verification of parameterized concurrent systems. *Electronic Notes in Theoretical Computer Science*, 76:65–82.
- Delzanno, G. (2018). Parameterized verification of publish/subscribe protocols via infinite-state model checking. In *Proceedings of the 33rd Italian Conference on Computational Logic (CILC)*, pages 97–111.
- Delzanno, G., Esparza, J., and Podelski, A. (1999). Constraint-based analysis of broadcast protocols. In *Proceeding of 13th International Workshop on Computer Science Logic (CSL)*, volume 1683 of *LNCS*, pages 50–66. Springer.
- Deutsch, A., Hull, R., Li, Y., and Vianu, V. (2018). Automatic verification of database-centric systems. *SIGLOG News*, 5(2):37–56.
- Deutsch, A., Hull, R., Patrizi, F., and Vianu, V. (2009). Automatic verification of data-centric business processes. In *Proceedings of the 12th International Conference on Database Theory (ICDT)*, pages 252–267.
- Deutsch, A., Li, Y., and Vianu, V. (2016). Verification of hierarchical artifact systems. In *Proceedings of the 35th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems (PODS)*, pages 179–194.
- Deutsch, A., Li, Y., and Vianu, V. (2019). Verification of hierarchical artifact systems. *ACM Transactions on Database Systems*, 44(3):12:1–12:68.
- Dutertre, B. and De Moura, L. (2006). The YICES SMT solver. Technical report, SRI International.
- Esparza, J., Finkel, A., and Mayr, R. (1999). On the verification of broadcast protocols. In *Proceedings of the 14th Annual IEEE Symposium on Logic in Computer Science (LICS)*, pages 352–359.
- Fagin, R. (1976). Probabilities on finite models. *Journal of Symbolic Logic*, 41(1):50–58.
- Ghilardi, S. (2004). Model theoretic methods in combined constraint satisfiability. *Journal of Automated Reasoning*, 33(3-4):221–249.
- Ghilardi, S. and Gianola, A. (2017). Interpolation, amalgamation and combination (the non-disjoint signatures case). In *Proceedings of the 11th International Symposium on Frontiers of Combining Systems (FroCoS)*, volume 10483 of *LNCS (LNAI)*, pages 316–332. Springer.
- Ghilardi, S. and Gianola, A. (2018). Modularity results for interpolation, amalgamation and superamalgamation. *Annals of Pure and Applied Logic*, 169(8):731–754.
- Ghilardi, S., Nicolini, E., Ranise, S., and Zucchelli, D. (2008a). Towards SMT model checking of array-based systems. In *Proceedings of the 4th International Joint Conference on Automated Reasoning (IJCAR)*, volume 5195 of *LNCS (LNAI)*, pages 67–82. Springer.
- Ghilardi, S., Nicolini, E., and Zucchelli, D. (2008b). A comprehensive framework for combined decision procedures. *ACM Transactions on Computational Logic*, 9(2):8:1–8:54.
- Ghilardi, S. and Ranise, S. (2010a). Backward reachability of array-based systems by SMT solving: Termination and invariant synthesis. *Logical Methods in Computer Science*, 6(4).
- Ghilardi, S. and Ranise, S. (2010b). MCMT: A model checker modulo theories. In *Proceedings*

- of the 5th International Joint Conference on Automated Reasoning (IJCAR), volume 6173 of LNCS (LNAI), pages 22–29. Springer.
- Ghilardi, S. and van Gool, S. J. (2016). Monadic second order logic as the model companion of temporal logic. In *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 417–426.
- Ghilardi, S. and van Gool, S. J. (2017). A model-theoretic characterization of monadic second order logic on infinite words. *Journal of Symbolic Logic*, 82(1):62–76.
- Higman, G. (1952). Ordering by divisibility in abstract algebras. *Proceedings of the London Mathematical Society*, 3(2):326–336.
- Hoder, K. and Bjørner, N. (2012). Generalized property directed reachability. In *Proceedings of the 15th International Conference on Theory and Applications of Satisfiability Testing (SAT)*, volume 7317 of LNCS, pages 157–171. Springer.
- Hull, R. (2008). Artifact-centric business process models: Brief survey of research results and challenges. In *Proceedings of the OTM Confederated International Conferences*, volume 5332 of LNCS, pages 1152–1163. Springer.
- Kruskal, J. B. (1960). Well-quasi-ordering, the Tree Theorem, and Vazsonyi’s conjecture. *Transactions of the American Mathematical Society*, 95:210–225.
- Lazic, R., Newcomb, T. C., Ouaknine, J., Roscoe, A. W., and Worrell, J. (2008). Nets with tokens which carry data. *Fundamenta Informaticae*, 88(3):251–274.
- Li, Y., Deutsch, A., and Vianu, V. (2017). VERIFAS: A practical verifier for artifact systems. *Proceedings of the VLDB Endowment*, 11(3):283–296.
- Lipparini, P. (1982). Locally finite theories with model companion. In *Atti della Accademia Nazionale dei Lincei. Classe di Scienze Fisiche, Matematiche e Naturali. Rendiconti, Serie 8*, volume 72. Accademia Nazionale dei Lincei.
- McMillan, K. L. (2006). Lazy Abstraction with Interpolants. In *Proceedings of the 18th International Conference on Computer Aided Verification (CAV)*, volume 4144 of LNCS, pages 123–136. Springer.
- Nicolini, E., Ringeissen, C., and Rusinowitch, M. (2009a). Data structures with arithmetic constraints: a non-disjoint combination. In *Proceedings of the 7th International Symposium on Frontiers of Combining Systems (FroCoS)*, volume 5749 of LNCS (LNAI), pages 319–334. Springer.
- Nicolini, E., Ringeissen, C., and Rusinowitch, M. (2009b). Satisfiability procedures for combination of theories sharing integer offsets. In *Proceedings of the 15th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, volume 5505 of LNCS, pages 428–442. Springer.
- Nicolini, E., Ringeissen, C., and Rusinowitch, M. (2010). Combining satisfiability procedures for unions of theories with a shared counting operator. *Fundamenta Informaticae*, 105(1-2):163–187.
- Nieuwenhuis, R. and Rubio, A. (2001). Paramodulation-based theorem proving. In *Handbook of Automated Reasoning (in 2 volumes)*, pages 371–443. MIT Press.
- Rado, R. (1964). Universal graphs and universal functions. *Acta Arithmetica*, 9:331–340.
- Robinson, A. (1951). *On the Metamathematics of Algebra*. North-Holland.
- Robinson, A. (1963). *Introduction to model theory and to the metamathematics of algebra*. Studies in logic and the foundations of mathematics. North-Holland.
- Rosa-Velardo, F. and de Frutos-Escrig, D. (2011). Decidability and complexity of Petri nets with unordered data. *Theoretical Computer Science*, 412(34):4439–4451.
- Schmitz, S. and Schnoebelen, P. (2013). The power of well-structured systems. In *Proceedings*

- of the 24th International Conference on Concurrency Theory (CONCUR), volume 8052 of LNCS, pages 5–24. Springer.
- Silver, B. (2011). *BPMN Method and Style*. Cody-Cassidy, 2nd edition.
- Sofronie-Stokkermans, V. (2008). Interpolation in local theory extensions. *Logical Methods in Computer Science*, 4(4).
- Sofronie-Stokkermans, V. (2016). On interpolation and symbol elimination in theory extensions. In *Proceedings of the 8th International Joint Conference on Automated Reasoning (IJCAR)*, volume 9706 of LNCS (LNAI), pages 273–289. Springer.
- Sofronie-Stokkermans, V. (2018). On interpolation and symbol elimination in theory extensions. *Logical Methods in Computer Science*, 14(3).
- Vianu, V. (2009). Automatic verification of database-driven systems: a new frontier. In *Proceedings of the 12th International Conference on Database Theory (ICDT)*, pages 1–13.
- Wheeler, W. H. (1976). Model-companions and definability in existentially complete structures. *Israel Journal of Mathematics*, 25(3-4):305–330.