

Quantifier Elimination for Database Driven Verification

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

¹Faculty of Computer Science, Free University of Bozen-Bolzano (Italy)
{calvanese, gianola, montali, rivkin}@inf.unibz.it

²Dipartimento di Matematica, Università degli Studi di Milano (Italy)
silvio.ghilardi@unimi.it

Abstract. Running verification tasks in database driven systems requires solving quantifier elimination problems of a new kind. These quantifier elimination problems are related to the notion of a cover introduced in ESOP 2008 by Gulwani and Musuvathi.

In this paper, we show how covers are strictly related to model completions, a well-known topic in model theory. We also investigate the computation of covers within the Superposition Calculus, by adopting a constrained version of the calculus, equipped with appropriate settings and reduction strategies.

In addition, we show that cover computations are computationally tractable for the fragment of the language used in applications to database driven verification. This observation is confirmed by analyzing the preliminary results obtained using the MCMT tool on the verification of data-aware process benchmarks. These benchmarks can be found in the last version of the tool distribution.

1 Introduction

Declarative approaches to infinite state model checking [44] need to manipulate logical formulae in order to represent sets of reachable states. To prevent divergence, various abstraction strategies have been adopted, ranging from interpolation-based [35] to sophisticated search via counterexample elimination [27]. Precise computations of the set of reachable states require some form of quantifier elimination and hence are subject to two problems, namely that quantifier elimination might not be available at all and that, when available, it is computationally very expensive.

To cope with the first problem, [26] introduced the notion of a *cover* and proved that covers exist for equality with uninterpreted symbols (EUF) and its combination with linear arithmetic; also, it was shown that covers can be used instead of quantifier elimination and yield a precise computation of reachable states. Concerning the second problem, in [26] it was observed (as a side remark) that computing the cover of a conjunction of literals becomes tractable when only free unary function symbols occur in the signature. It can be shown (see Section 6 below) that the same observation applies when also free relational symbols occur.

In [10,12] we propose a new formalism for representing *read-only database schemata* towards the verification of integrated models of processes and data [9], in particular so-called *artifact systems* [47,16,32,7]; this formalism (briefly recalled in

Section 4.1 below) uses precisely signatures comprising unary function symbols and free n -ary relations. In [10,12] we apply model completeness techniques for verifying transition systems based on read-only databases, in a framework where such systems employ both individual and higher order variables.

In this paper we show (see Section 3 below) that covers are strictly related to *model completions* and to *uniform interpolation* [43], thus building a bridge between different research areas. In particular, we prove that computing covers for a theory is *equivalent* to eliminating quantifiers in its model completion. Model completeness has other well-known applications in computer science. It has been applied: (i) to reveal interesting connections between temporal logic and monadic second order logic [23,24]; (ii) in automated reasoning to design complete algorithms for constraint satisfiability in combined theories over non disjoint signatures [18,1,21,39,37,38] and theory extensions [45,46]; (iii) to obtain combined interpolation for modal logics and software verification theories [19,20].

In the last part of the paper (Section 5 below), we prove that covers for (EUF) can be computed through a constrained version of the *Superposition Calculus* [42] equipped with appropriate settings and reduction strategies; the related completeness proof requires a careful analysis of the constrained literals generated during the saturation process. Complexity bounds for the fragment used in database driven verification are investigated in Section 6; in Subsection 6.2 we give some details about our first implementation in our tool MCMT.

This paper is the extended version of [13].

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 $T \models \phi$ 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} .

We give now the definitions of constraint satisfiability problem and quantifier elimination for a theory 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 a constraint (equivalently, a quantifier-free formula) $\phi(\underline{x})$ and we are asked whether there exist a model \mathcal{M} of T and an assignment \mathcal{I} to the free variables \underline{x} such that $\mathcal{M}, \mathcal{I} \models \phi(\underline{x})$.

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 when we talk about quantifier elimination, an effective procedure for eliminating quantifiers is given.

We recall also some basic definitions and notions from logic and model theory. We focus on the definitions of diagram, embedding, substructure and amalgamation.

2.1 Substructures and embeddings

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{M} , the signature Σ can be expanded to a new signature $\Sigma^{|\mathcal{M}|} := \Sigma \cup \{\bar{a} \mid a \in |\mathcal{M}|\}$ by adding a set of new constants \bar{a} (the *name* for a), one for each element a in \mathcal{M} , with the convention that two distinct elements are denoted by different "name" constants. \mathcal{M} can be expanded to a $\Sigma^{|\mathcal{M}|}$ -structure $\overline{\mathcal{M}} := (\mathcal{M}, a)_{a \in |\mathcal{M}|}$ 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{M} and $\overline{\mathcal{M}}$ interchangeably: in particular, given a Σ -structure \mathcal{M} and a Σ -formula $\phi(\underline{x})$ with free variables that are all in \underline{x} , we will write, by abuse of notation, $\mathcal{M} \models \phi(\underline{a})$ instead of $\overline{\mathcal{M}} \models \phi(\underline{a})$.

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

$$\mathcal{M} \models \varphi \quad \Rightarrow \quad \mathcal{N} \models \varphi \tag{1}$$

for all $\Sigma^{|\mathcal{M}|}$ -atoms φ (here \mathcal{M} is regarded as a $\Sigma^{|\mathcal{M}|}$ -structure, by interpreting each additional constant $a \in |\mathcal{M}|$ into itself and \mathcal{N} is regarded as a $\Sigma^{|\mathcal{M}|}$ -structure by interpreting each additional constant $a \in |\mathcal{M}|$ into $\mu(a)$). In case condition (1) holds for all $\Sigma^{|\mathcal{M}|}$ -literals, the homomorphism μ is said to be an *embedding* and if it holds for all first order formulae, the embedding μ is said to be *elementary*. Notice the following facts:

- (a) since we have equality in the signature, an embedding is an injective function;
- (b) an embedding $\mu : \mathcal{M} \rightarrow \mathcal{N}$ must be an algebraic homomorphism, that is for every n -ary function symbol f and for every m_1, \dots, m_n in $|\mathcal{M}|$, we must have $f^{\mathcal{N}}(\mu(m_1), \dots, \mu(m_n)) = \mu(f^{\mathcal{M}}(m_1, \dots, m_n))$;
- (c) for an n -ary predicate symbol P we must have $(m_1, \dots, m_n) \in P^{\mathcal{M}}$ iff $(\mu(m_1), \dots, \mu(m_n)) \in P^{\mathcal{N}}$.

It is easily seen that an embedding $\mu : \mathcal{M} \longrightarrow \mathcal{N}$ can be equivalently defined as a map $\mu : |\mathcal{M}| \longrightarrow |\mathcal{N}|$ satisfying the conditions (a)-(b)-(c) above. If $\mu : \mathcal{M} \longrightarrow \mathcal{N}$ is an embedding which is just the identity inclusion $|\mathcal{M}| \subseteq |\mathcal{N}|$, we say that \mathcal{M} is a *substructure* of \mathcal{N} or that \mathcal{N} is an *extension* of \mathcal{M} . A Σ -structure \mathcal{M} is said to be *generated by* a set X included in its support $|\mathcal{M}|$ iff there are no proper substructures of \mathcal{M} including X .

The notion of substructure can be equivalently defined as follows: given a Σ -structure \mathcal{N} and a Σ -structure \mathcal{M} such that $|\mathcal{M}| \subseteq |\mathcal{N}|$, we say that \mathcal{M} is a Σ -*substructure* of \mathcal{N} if:

- for every function symbol f in Σ , the interpretation of f in \mathcal{M} (denoted using $f^{\mathcal{M}}$) is the restriction of the interpretation of f in \mathcal{N} to $|\mathcal{M}|$ (i.e. $f^{\mathcal{M}}(m) = f^{\mathcal{N}}(m)$ for every m in $|\mathcal{M}|$); this fact implies that a substructure \mathcal{M} must be a subset of \mathcal{N} which is closed under the application of $f^{\mathcal{N}}$.
- for every relation symbol P in Σ and every tuple $(m_1, \dots, m_n) \in |\mathcal{M}|^n$, $(m_1, \dots, m_n) \in P^{\mathcal{M}}$ iff $(m_1, \dots, m_n) \in P^{\mathcal{N}}$, which means that the relation $P^{\mathcal{M}}$ is the restriction of $P^{\mathcal{N}}$ to the support of \mathcal{M} .

We recall that a substructure *preserves* and *reflects* validity of ground formulae, in the following sense: given a Σ -substructure \mathcal{M}_1 of a Σ -structure \mathcal{M}_2 , a ground $\Sigma^{|\mathcal{M}_1|}$ -sentence θ is true in \mathcal{M}_1 iff θ is true in \mathcal{M}_2 .


2.2 Robinson Diagrams and Amalgamation

Let \mathcal{M} be a Σ -structure. The *diagram* of \mathcal{M} , denoted by $\Delta_{\Sigma}(\mathcal{M})$, is defined as the set of ground $\Sigma^{|\mathcal{M}|}$ -literals (i.e. atomic formulae and negations of atomic formulae) that are true in \mathcal{M} .

An easy but nevertheless important basic result, called *Robinson Diagram Lemma* [14], says that, given any Σ -structure \mathcal{N} , the embeddings $\mu : \mathcal{M} \longrightarrow \mathcal{N}$ are in bijective correspondence with expansions of \mathcal{N} to $\Sigma^{|\mathcal{M}|}$ -structures which are models of $\Delta_{\Sigma}(\mathcal{M})$. 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 1 (Amalgamation). *A theory T has the amalgamation property if for every couple of embeddings $\mu_1 : \mathcal{M}_0 \longrightarrow \mathcal{M}_1$, $\mu_2 : \mathcal{M}_0 \longrightarrow \mathcal{M}_2$ among models of T , there exists a model \mathcal{M} of T endowed with embeddings $\nu_1 : \mathcal{M}_1 \longrightarrow \mathcal{M}$ and $\nu_2 : \mathcal{M}_2 \longrightarrow \mathcal{M}$ 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 -*amalgama* of $\mathcal{M}_1, \mathcal{M}_2$ over \mathcal{M}_0

3 Covers, Uniform Interpolation and Model Completions

We report the notion of *cover* taken from [26]. Fix a theory T and an existential formula $\exists \underline{e} \phi(\underline{e}, \underline{y})$; call a *residue* of $\exists \underline{e} \phi(\underline{e}, \underline{y})$ any quantifier-free formula belonging to the set of quantifier-free formulae $Res(\exists \underline{e} \phi) = \{\theta(\underline{y}, \underline{z}) \mid T \models \phi(\underline{e}, \underline{y}) \rightarrow \theta(\underline{y}, \underline{z})\}$. A quantifier-free formula $\psi(\underline{y})$ is said to be a T -*cover* (or, simply, a *cover*) of $\exists \underline{e} \phi(\underline{e}, \underline{y})$ iff $\psi(\underline{y}) \in Res(\exists \underline{e} \phi)$ and $\psi(\underline{y})$ implies (modulo T) all the other formulae in $Res(\exists \underline{e} \phi)$. The following Lemma (to be widely used throughout the paper) supplies a semantic counterpart to the notion of a cover:

Lemma 1. *A formula $\psi(\underline{y})$ is a T -cover of $\exists \underline{e} \phi(\underline{e}, \underline{y})$ iff it satisfies the following two conditions: (i) $T \models \forall \underline{y} (\exists \underline{e} \phi(\underline{e}, \underline{y}) \rightarrow \psi(\underline{y}))$; (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 $\mathcal{N} \models \exists \underline{e} \phi(\underline{e}, \underline{a})$. \triangleleft*

Proof. Suppose that $\psi(\underline{y})$ satisfies conditions (i) and (ii) above. Condition (i) says that $\psi(\underline{y}) \in Res(\exists \underline{e} \phi)$, so ψ is a residue. In order to show that ψ is also a cover, we have to prove that $T \models \forall \underline{y}, \underline{z} (\psi(\underline{y}) \rightarrow \theta(\underline{y}, \underline{z}))$, for every $\theta(\underline{y}, \underline{z})$ that is a residue for $\exists \underline{e} \phi(\underline{e}, \underline{y})$. Given a model \mathcal{M} of T , take a pair of tuples $\underline{a}, \underline{b}$ of elements from $|\mathcal{M}|$ and suppose that $\mathcal{M} \models \psi(\underline{a})$. By condition (ii), there is a model \mathcal{N} of T such that \mathcal{M} embeds into \mathcal{N} and $\mathcal{N} \models \exists \underline{e} \phi(\underline{e}, \underline{a})$. Using the definition of $Res(\exists \underline{e} \phi)$, we have $\mathcal{N} \models \theta(\underline{a}, \underline{b})$, since $\theta(\underline{y}, \underline{z}) \in Res(\exists \underline{e} \phi)$. Since \mathcal{M} is a substructure of \mathcal{N} and θ is quantifier-free, $\mathcal{M} \models \theta(\underline{a}, \underline{b})$ as well, as required.

Suppose that $\psi(\underline{y})$ is a cover. The definition of residue implies condition (i). To show condition (ii) we have to prove that, given a model \mathcal{M} of T , for every tuple \underline{a} of elements from $|\mathcal{M}|$, if $\mathcal{M} \models \psi(\underline{a})$, then there exists a model \mathcal{N} of T such that \mathcal{M} embeds into \mathcal{N} and $\mathcal{N} \models \exists \underline{x} \phi(\underline{x}, \underline{a})$. By reduction to absurdity, suppose that this is not the case: this is equivalent (by using Robinson Diagram Lemma) to the fact that $\Delta(\mathcal{M}) \cup \{\phi(\underline{e}, \underline{a})\}$ is a T -inconsistent $\Sigma^{|\mathcal{M}| \cup \{\underline{e}\}}$ -theory. By compactness, there is a finite number of literals $\ell_1(\underline{a}, \underline{b}), \dots, \ell_m(\underline{a}, \underline{b})$ (for some tuple \underline{b} of elements from $|\mathcal{M}|$) such that $\mathcal{M} \models \ell_i$ (for all $i = 1, \dots, m$) and $T \models \phi(\underline{e}, \underline{a}) \rightarrow \neg(\ell_1(\underline{a}, \underline{b}) \wedge \dots \wedge \ell_m(\underline{a}, \underline{b}))$, which means that $T \models \phi(\underline{e}, \underline{y}) \rightarrow (\neg \ell_1(\underline{y}, \underline{z}) \vee \dots \vee \neg \ell_m(\underline{y}, \underline{z}))$, i.e. that $T \models \exists \underline{e} \phi(\underline{e}, \underline{y}) \rightarrow (\neg \ell_1(\underline{y}, \underline{z}) \vee \dots \vee \neg \ell_m(\underline{y}, \underline{z}))$. By definition of residue, clearly $(\neg \ell_1(\underline{y}, \underline{z}) \vee \dots \vee \neg \ell_m(\underline{y}, \underline{z})) \in Res(\exists \underline{x} \phi)$; then, since $\psi(\underline{y})$ is a cover, $T \models \psi(\underline{y}) \rightarrow (\neg \ell_1(\underline{y}, \underline{z}) \vee \dots \vee \neg \ell_m(\underline{y}, \underline{z}))$, which implies that $\mathcal{M} \models \neg \ell_j(\underline{a}, \underline{b})$ for some $j = 1, \dots, m$, which is a contradiction. Thus, $\psi(\underline{y})$ satisfies conditions (ii) too. \dashv

We say that a theory T has *uniform quantifier-free interpolation* iff every existential formula $\exists \underline{e} \phi(\underline{e}, \underline{y})$ (equivalently, every primitive formula $\exists \underline{e} \phi(\underline{e}, \underline{y})$) has a T -cover.

It is clear that if T has uniform quantifier-free interpolation, then it has ordinary quantifier-free interpolation [8], in the sense that if we have $T \models \phi(\underline{e}, \underline{y}) \rightarrow \phi'(\underline{y}, \underline{z})$ (for quantifier-free formulae ϕ, ϕ'), then there is a quantifier-free formula $\theta(\underline{y})$ such that $T \models \phi(\underline{e}, \underline{y}) \rightarrow \theta(\underline{y})$ and $T \models \theta(\underline{y}) \rightarrow \phi'(\underline{y}, \underline{z})$. In fact, if T has uniform quantifier-free interpolation, then the interpolant θ is independent on ϕ' (the same $\theta(\underline{y})$ can be used as interpolant for all entailments $T \models \phi(\underline{e}, \underline{y}) \rightarrow \phi'(\underline{y}, \underline{z})$, varying ϕ').

We say that a *universal* theory T has a *model completion* iff there is a stronger theory $T^* \supseteq T$ (still within the same signature Σ of T) such that (i) every Σ -constraint that is satisfiable in a model of T is satisfiable in a model of T^* ; (ii) T^* eliminates

quantifiers. Other equivalent definitions are possible [14]: for instance, (i) is equivalent to the fact that T and T^* prove the same quantifier-free formulae or again to the fact that every model of T can be embedded into a model of T^* . We recall that the model completion, if it exists, is unique and that its existence implies the amalgamation property for T [14]. The relationship between uniform interpolation in a propositional logic and model completion of the the equational theory of the variety algebraizing it was extensively studied in [25]. In the context of first order theories, we prove an even more direct connection:

Theorem 1. *Suppose that T is a universal theory. Then T has a model completion T^* iff T has uniform quantifier-free interpolation. If this happens, T^* is axiomatized by the infinitely many sentences*

$$\forall \underline{y} (\psi(\underline{y}) \rightarrow \exists \underline{e} \phi(\underline{e}, \underline{y})) \quad (2)$$

where $\exists \underline{e} \phi(\underline{e}, \underline{y})$ is a primitive formula and ψ is a cover of it. \triangleleft

Proof. Suppose first that there is a model completion T^* of T and let $\exists \underline{e} \phi(\underline{e}, \underline{y})$ be a primitive formula. Since T^* eliminates quantifiers, we have $T^* \models \exists \underline{e} \phi(\underline{e}, \underline{y}) \leftrightarrow \psi(\underline{y})$ for some quantifier-free formula $\psi(\underline{y})$. Since T and T^* prove the same quantifier-free formulae, we have that $\psi(\underline{y}) \in Res(\exists \underline{e} \phi)$. If $\theta(\underline{y}, \underline{z}) \in Res(\exists \underline{e} \phi)$, then we have $T \models \phi(\underline{e}, \underline{y}) \rightarrow \theta(\underline{y}, \underline{z})$; the same entailment holds in T^* too, where we have $T^* \models \psi(\underline{y}) \rightarrow \theta(\underline{y}, \underline{z})$. Since $\psi(\underline{y}) \rightarrow \theta(\underline{y}, \underline{z})$ is quantifier-free, we have also $T \models \psi(\underline{y}) \rightarrow \theta(\underline{y}, \underline{z})$, showing that ψ is a cover of $\exists \underline{e} \phi(\underline{e}, \underline{y})$. Thus T has uniform interpolation, because we found a cover for every primitive formula.

Suppose vice versa that T has uniform interpolation. Let T^* be the theory axiomatized by all the formulae (2) above. From (i) of Lemma 1 and (2) above, 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) of Lemma 1 and (2).

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 $\exists \underline{e} \phi(\underline{e}, \underline{y})$, consider the set $\{(\underline{a}, \exists \underline{e} \phi(\underline{e}, \underline{a}))\}$ such that $\mathcal{M} \models \psi(\underline{a})$ (where ψ is a cover of ϕ). By Zermelo's Theorem, the set $\{(\underline{a}, \exists \underline{e} \phi(\underline{e}, \underline{a}))\}$ can be well-ordered: let $\{(a_i, \exists \underline{e} \phi_i(\underline{e}, a_i))\}_{i \in I}$ 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 the extension of $\bigcup_{j < i} \mathcal{M}_j$ such that $\mathcal{M}_i \models \exists \underline{e} \phi_i(\underline{e}, a_i)$, which exists for (ii) of Lemma 1 since $\bigcup_{j < i} \mathcal{M}_j \models \psi_i(a_i)$ (remember that validity of ground formulae is preserved passing through substructures and superstructures, and $\mathcal{M}_0 \models \psi_i(a_i)$).

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 (2). The last fact can be shown using the following finiteness argument.

Fix ϕ, ψ as in (2). 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 $(\underline{a}', \psi(\underline{a}'))$ appears in the enumeration given by the well-ordered set of pairs $\{(\underline{a}_i, \exists \underline{e} \phi_i(\underline{e}, \underline{a}_i))\}_{i \in I}$ (for some ordinal I) such that $\mathcal{M}^k \models \psi_i(\underline{a})$. Hence, by construction and since $\psi(\underline{a}')$ is a ground formula, we have that there exists a $j \in I$ such that $\mathcal{M}_j^k \models \exists \underline{e} \phi(\underline{e}, \underline{a}')$. In conclusion, since the existential formulae are preserved passing to extensions, we obtain $\mathcal{M}' \models \exists \underline{e} \phi(\underline{e}, \underline{a}')$, as wanted. \dashv

4 Model-Checking Applications

In this section we supply old and new motivations for investigating covers and model completions in view of model-checking applications. We first report the considerations from [26,10,12] on symbolic model-checking via model completions (or, equivalently, via covers) in the basic case where system variables are represented as individual variables (for more advanced applications where system variables are both individual and higher order variables, see [10,12]). Similar ideas ('use quantifier elimination in the model completion even if T does not allow quantifier elimination') were used in [45] for interpolation and symbol elimination.

Definition 2. A (quantifier-free) transition system is a tuple

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

where: (i) Σ is a signature and T is a Σ -theory; (ii) $\underline{x} = x_1, \dots, x_n$ are individual variables; (iii) $\iota(\underline{x})$ is a quantifier-free formula; (iv) $\tau(\underline{x}, \underline{x}')$ is a quantifier-free formula (here the \underline{x}' are renamed copies of the \underline{x}). \triangleleft

A *safety* formula for a transition system \mathcal{S} is a further quantifier-free formula $v(\underline{x})$ describing undesired states of \mathcal{S} . We say that \mathcal{S} is *safe with respect to v* if the system has no finite run leading from ι to v , i.e. (formally) if there are no model \mathcal{M} of T and no $k \geq 0$ such that the formula

$$\iota(\underline{x}^0) \wedge \tau(\underline{x}^0, \underline{x}^1) \wedge \dots \wedge \tau(\underline{x}^{k-1}, \underline{x}^k) \wedge v(\underline{x}^k) \quad (3)$$

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

Suppose now that the theory T mentioned in Definition 2(i) is universal, has decidable constraint satisfiability problem and admits a model completion T^* . Algorithm 1 describes the *backward reachability algorithm* for handling the safety problem for \mathcal{S} (the dual algorithm working via forward search is described in equivalent terms in [26]). An integral part of the algorithm is to compute preimages. For that purpose, for any $\phi_1(\underline{x}, \underline{x}')$ and $\phi_2(\underline{x})$, we define $Pre(\phi_1, \phi_2)$ to be the formula $\exists \underline{x}' (\phi_1(\underline{x}, \underline{x}') \wedge \phi_2(\underline{x}'))$. The *preimage* of the set of states described by a state formula $\phi(\underline{x})$ is the set of states described by $Pre(\tau, \phi)$.

The subprocedure $\text{QE}(T^*, \phi)$ in Line 6 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 produce proof obligations that can be discharged thanks to the fact that T has decidable constraint satisfiability problem.

The proof of Proposition 1 consists just in the observation that, thanks to quantifier elimination in T^* , (3) is a quantifier-free formula and that a quantifier-free formula is satisfiable in a model of T iff so is it in a model of T^* :

Proposition 1. *Suppose that the universal Σ -theory T has decidable constraint satisfiability problem and admits a model completion T^* . For every transition system $\mathcal{S} = \langle \Sigma, T, \underline{x}, \iota, \tau \rangle$, the backward search algorithm is effective and partially correct for solving safety problems for \mathcal{S} .¹* \triangleleft

Despite its simplicity, Proposition 1 is a crucial fact. Notice that it implies decidability of the safety problems in some interesting cases: this happens, for instance, when in T there are only finitely many quantifier-free formulae in which \underline{x} occur, as in case T has a purely relational signature or, more generally, T is *locally finite*². Since a theory is universal iff it is closed under substructures [14] and since a universal locally finite theory has a model completion iff it has the amalgamation property [48,33], it follows that Proposition 1 can be used to cover the decidability result stated in Theorem 5 of [7] (once restricted to transition systems over a first-order definable class of Σ -structures).

4.1 Database Schemata

In this subsection, we provide a new application for the above explained model-checking techniques [10,12]. The application relates to the verification of integrated models of business processes and data [9], referred to as artifact systems [47], where the behavior of the process is influenced by data stored in a relational database (DB) with constraints. The data contained therein are read-only: they can be queried by the process and stored in a working memory, which in the context of this paper is constituted by a set of system variables. In this context, safety amounts to checking whether the system never reaches an undesired property, irrespectively of what is contained in the read-only DB.

We define next the two key notions of (read-only) DB schema and instance, by relying on an algebraic, functional characterization.

¹ *Partial correctness* means that, when the algorithm terminates, it gives a correct answer. *Effectiveness* means that all subprocedures in the algorithm can be effectively executed.

² We say that T is locally finite iff for every finite tuple of variables \underline{x} there are only finitely many non T -equivalent atoms $A(\underline{x})$ involving only the variables \underline{x} .

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
        $\perp$  return unsafe
4     $B \leftarrow \phi \vee B;$ 
5     $\phi \leftarrow \text{Pre}(\tau, \phi);$ 
6     $\phi \leftarrow \text{QE}(T^*, \phi);$ 
   return (safe,  $B$ );

```

Definition 3. 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 function symbols are all unary; (ii) T is a DB theory, that is, a set of universal Σ -sentences. \triangleleft

We now focus on extensional data conforming to a given DB schema.

Definition 4. A DB instance of DB schema $\langle \Sigma, T \rangle$ is a Σ -structure \mathcal{M} such that \mathcal{M} is a model of T .³ \triangleleft

One might be surprised by the fact that signatures in our DB schemata contain unary function symbols, beside relational symbols. As shown in [10,12], the algebraic, functional characterization of DB schema and instance can be actually reinterpreted in the classical, relational model so as to reconstruct the requirements posed in [32]. Definition 3 naturally corresponds to the definition of relational database schema equipped with single-attribute *primary keys* and *foreign keys*. To see this connection, 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 sort S from Σ 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 from Σ of the form $f : S \rightarrow S'$.

The fact that R_S is constructed starting from functions in Σ naturally induces corresponding functional dependencies within R_S , and inclusion dependencies from R_S to other relation schemas. 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 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 $R_{S'}$.

Given a DB instance \mathcal{M} of $\langle \Sigma, T \rangle$, its corresponding relational instance $\mathcal{R}[\mathcal{M}]$ is the minimal set satisfying the following property: for every id sort S from Σ , let f_1, \dots, f_n be all functions in Σ with domain S ; then, for every identifier $o \in S^{\mathcal{M}}$, $\mathcal{R}[\mathcal{M}]$ contains a *labeled fact* of the form $R_S(id_S : o^{\mathcal{M}}, a_{f_1} : f_1^{\mathcal{M}}(o), \dots, a_{f_n} : f_n^{\mathcal{M}}(o))$. In addition, $\mathcal{R}[\mathcal{M}]$ contains the tuples from $r^{\mathcal{M}}$, for every relational symbol r from Σ (these relational symbols represent plain relations, i.e. those not possessing a key).

We close our discussion by focusing on DB theories. Notice that EUF suffices to handle the sophisticated setting of database-driven systems from [12] (e.g., key dependencies). The role of a non-empty DB theory is to encode background axioms to express additional constraints. 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 [17,32]. To accommodate this, we add to every sort S of Σ a constant $undef_S$ (written by abuse of notation just $undef$ from now on), used to specify an undefined value. Then, for each function symbol f of Σ , we can impose additional constraints involving $undef$, for example by adding the following axioms to the DB theory:

$$\forall x (x = \text{undef} \leftrightarrow f(x) = \text{undef}) \quad (4)$$

³ One may restrict to models interpreting sorts as *finite* sets, as customary in database theory. Since the theories we are dealing with usually have finite model property for constraint satisfiability, assuming such restriction turns out to be irrelevant, as far as safety problems are concerned (see [10,12] for an accurate discussion).

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. A slightly different approach may handle *many* undefined values for each sort; the reader is referred to [10,12] for examples of concrete database instances formalized in our framework. We just point out that in most cases the kind of axioms that we need for our DB theories T are just *one-variable universal axioms* (like Axioms 4), so that they fit the hypotheses of Proposition 2 below.

We are interested in applying the algorithm of Proposition 1 to what we call *simple artifact systems*, i.e. transition systems $\mathcal{S} = \langle \Sigma, T, \underline{x}, \iota(\underline{x}), \tau(\underline{x}, \underline{x}') \rangle$, where $\langle \Sigma, T \rangle$ is a DB schema in the sense of Definition 3. To this aim, it is sufficient to identify a suitable class of DB theories having a model completion and whose constraint satisfiability problem is decidable. A first result in this sense is given below. We associate to a DB signature Σ the edge-labeled graph $G(\Sigma)$ whose nodes are the sorts in Σ , and such that $G(\Sigma)$ contains a labeled edge $S \xrightarrow{f} S'$ if and only if Σ contains a function symbol whose source sort is S and whose target sort is S' . We say that Σ is *acyclic* if $G(\Sigma)$ is so.

Proposition 2. *A DB theory T has decidable constraint satisfiability problem and admits a model completion in case it is axiomatized by finitely many universal one-variable formulae and Σ is acyclic.* \triangleleft

Proof. First, notice that, in case Σ is acyclic, T has the finite model property. In fact, if $T := \emptyset$, then congruence closure ensures that the finite model property holds and decides constraint satisfiability in time $O(n \log n)$. Otherwise, we reduce the argument to the Herbrand Theorem. 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 reduce to the case where T is empty. Hence, the constraint satisfiability problem for T is decidable.

To show the existence of a model completion, we freely take inspiration from an analogous result in [48,33]. 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 1: this fact (together with condition (i) of Lemma 1) yields that T has a model completion which is axiomatized by the infinitely many sentences (2).

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 and *n*-ary relation 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 $m_i \in |\mathcal{M}_i|$ ($i = 1, 2$), $f^{\mathcal{M}}(m_i) := f^{\mathcal{M}_1}(m_i)$, i.e. the interpretation of f in \mathcal{M} is the restriction of 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 $f^{\mathcal{M}}(a) := f^{\mathcal{M}_1}(a) = f^{\mathcal{M}_0}(a) = f^{\mathcal{M}_2}(a)$. For every *n*-ry relation symbol R , we let $R^{\mathcal{M}}$ be the union of $R^{\mathcal{M}_1} \cup R^{\mathcal{M}_2}$; notice that if the tuple \underline{a} belongs to $R^{\mathcal{M}}$, then we must have that $\underline{a} \in |\mathcal{M}_i|^n$ and that $\underline{a} \in R^{\mathcal{M}_i}$ for either $i = 1$ or $i = 2$ (or for both, but this happens

just in case $\underline{a} \in |\mathcal{M}_0|^n = |\mathcal{M}_1|^n \cap |\mathcal{M}_2|^n$, because \mathcal{M}_0 is a substructure of \mathcal{M}_1 and \mathcal{M}_2). 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 $\exists \underline{e} \phi(\underline{e}, \underline{y})$, we find a suitable $\psi(\underline{y})$ such that conditions (i) and (ii) of Lemma 1 hold. We define $\psi(\underline{y})$ as the conjunction of the set of all quantifier-free $\chi(\underline{y})$ -formulae such that $\phi(\underline{e}, \underline{y}) \rightarrow \chi(\underline{y})$ 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 1 holds.

We show that also condition (ii) of Lemma 1 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 $\mathcal{N}' \models \exists \underline{e} \phi(\underline{e}, \underline{a})$, it is sufficient to prove (by the Robinson Diagram Lemma) that the $\Sigma^{|\mathcal{M}[\underline{a}]| \cup \{\underline{e}\}}$ -theory $\Delta(\mathcal{M}[\underline{a}]) \cup \{\phi(\underline{e}, \underline{a})\}$ 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 $\phi(\underline{e}, \underline{a}) \models_T \neg(l_1(\underline{a}) \wedge \dots \wedge l_m(\underline{a}))$. Therefore, defining $A(\underline{a}) := l_1(\underline{a}) \wedge \dots \wedge l_m(\underline{a})$, we get that $\phi(\underline{e}, \underline{a}) \models_T \neg A(\underline{a})$, which implies that $\neg A(\underline{a})$ is one of the $\chi(\underline{y})$ -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 $\mathcal{N}' \models \exists \underline{e} \phi(\underline{e}, \underline{a})$. 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 $\mathcal{N}' \models \exists \underline{e} \phi(\underline{e}, \underline{a})$, also $\mathcal{N} \models \exists \underline{e} \phi(\underline{e}, \underline{a})$ holds, as required. \dashv

Since acyclicity of Σ yields local finiteness, we immediately get as a Corollary the decidability of safety problems for transitions systems based on DB schema satisfying the hypotheses of the above theorem.

5 Covers via Constrained Superposition

Of course, a model completion may not exist at all; Proposition 2 shows that it exists in case T is a DB theory axiomatized by universal one-variable formulae and Σ is acyclic.

The second hypothesis is unnecessarily restrictive and the algorithm for quantifier elimination suggested by the proof of Proposition 2 is highly impractical: for this reason we are trying a different approach. In this section, we drop the acyclicity hypothesis and examine the case where the theory T is empty and the signature Σ may contain function symbols of any arity. Covers in this context were shown to exist already in [26], using an algorithm that, very roughly speaking, determines all the conditional equations that can be derived concerning the nodes of the congruence closure graph. An algorithm for the generation of interpolants, still relying on congruence closure [29] and similar to the one presented in [26], is supplied in [30].

We follow a different plan and we want to produce covers (and show that they exist) using *saturation-based theorem proving*. The natural idea to proceed in this sense is to take the matrix $\phi(\underline{e}, \underline{y})$ of the primitive formula $\exists \underline{e} \phi(\underline{e}, \underline{y})$ we want to compute the cover of: this is a conjunction of literals, so we consider each variable as a free constant, we saturate the corresponding set of ground literals and finally we output the literals involving only the \underline{y} . For saturation, one can use any version of the superposition calculus [42]. This procedure however for our problem is not sufficient. As a trivial counterexample consider the primitive formula $\exists e (R(e, y_1) \wedge \neg R(e, y_2))$: the set of literals $\{R(e, y_1), \neg R(e, y_2)\}$ is saturated (recall that we view e, y_1, y_2 as constants), however the formula has a non-trivial cover $y_1 \neq y_2$ which is *not* produced by saturation. If we move to signatures with function symbols, the situation is even worse: the set of literals $\{f(e, y_1) = y'_1, f(e, y_2) = y'_2\}$ is saturated but the formula $\exists e (f(e, y_1) = y'_1 \wedge f(e, y_2) = y'_2)$ has the *conditional equality* $y_1 = y_2 \rightarrow y'_1 = y'_2$ as cover. *Disjunctions of disequations* might also arise: the cover of $\exists e h(e, y_1, y_2) \neq h(e, y'_1, y'_2)$ (as well as the cover of $\exists e f(f(e, y_1), y_2) \neq f(f(e, y'_1), y'_2)$, see Example 1 below) is $y_1 \neq y'_1 \vee y_2 \neq y'_2$.⁴

Notice that our problem is different from the problem of producing ordinary quantifier-free interpolants via saturation based theorem proving [31]: for ordinary Craig interpolants, we have as input *two* quantifier-free formulae $\phi(\underline{e}, \underline{y}), \phi'(\underline{y}, \underline{z})$ such that $\phi(\underline{e}, \underline{y}) \rightarrow \phi'(\underline{y}, \underline{z})$ is valid; here we have a *single* formula $\phi(\underline{e}, \underline{y})$ in input and we are asked to find an interpolant which is good *for all possible* $\phi'(\underline{y}, \underline{z})$ such that $\phi(\underline{e}, \underline{y}) \rightarrow \phi'(\underline{y}, \underline{z})$ is valid. Ordinary interpolants can be extracted from a refutation of $\phi(\underline{e}, \underline{y}) \wedge \neg \phi'(\underline{y}, \underline{z})$, here we are not given any refutation at all (and we are not even supposed to find one).

What we are going to show is that, nevertheless, saturation via superposition can be used to produce covers, if suitably adjusted. In this section we consider signatures with n -ary function symbols (for all $n \geq 1$). For simplicity, we omit n -ary relation symbols (you can easily handle them by rewriting $R(t_1, \dots, t_n)$ as $R(t_1, \dots, t_n) = true$, as customary in the paramodulation literature [42]).

We are going to compute the cover of a primitive formula $\exists \underline{e} \phi(\underline{e}, \underline{y})$ to be fixed for the remainder of this section. We call variables \underline{e} *existential* and variables \underline{y} *parameters*. By applying abstraction steps, we can assume that ϕ is *primitive flat*. i.e. that it is a conjunction of *\underline{e} -flat literals*, defined below. [By an abstraction step we mean replacing

⁴ This example points out a problem that needs to be fixed in the algorithm presented in [26]: that algorithm in fact outputs only equalities, conditional equalities and single disequalities, so it cannot correctly handle this example.

$\exists \underline{e} \phi$ with $\exists \underline{e} \exists e' (e' = u \wedge \phi')$, where e' is a fresh variable and ϕ' is obtained from ϕ by replacing some occurrences of a term $u(\underline{e}, y)$ by e'].

A term or a formula are said to be \underline{e} -free iff the existential variables do not occur in it. An \underline{e} -flat term is an \underline{e} -free term $t(\underline{y})$ or a variable from \underline{e} or again it is of the kind $f(u_1, \dots, u_n)$, where f is a function symbol and u_1, \dots, u_n are \underline{e} -free terms or variables from \underline{e} . An \underline{e} -flat literal is a literal of the form

$$t = a, \quad a \neq b$$

where t is an \underline{e} -flat term and a, b are either \underline{e} -free terms or variables from \underline{e} .

We assume the reader is familiar with standard conventions used in rewriting and paramodulation literature: in particular $s|_p$ denotes the subterm of s in position p and $s[u]_p$ denotes the term obtained from s by replacing $s|_p$ with u . We use \equiv to indicate coincidence of syntactic expressions (as strings) to avoid confusion with equality symbol; when we write equalities like $s = t$ below, we may mean both $s = t$ or $t = s$ (an equality is seen as a multiset of two terms). For information on reduction ordering, see for instance [2].

We first replace variables $\underline{e} = e_1, \dots, e_n$ and $\underline{y} = y_1, \dots, y_m$ by free constants - we keep the names $e_1, \dots, e_n, y_1, \dots, y_m$ for these constants. Choose a reduction ordering $>$ total for ground terms such that \underline{e} -flat literals $t = a$ are always oriented from left to right in the following two cases: (i) t is not \underline{e} -free and a is \underline{e} -free; (ii) t is not \underline{e} -free, it is not equal to any of the \underline{e} and a is a variable from \underline{e} . To obtain such properties, one may for instance choose a suitable Knuth-Bendix ordering taking weights in some transfinite ordinal, see [34].

Given two \underline{e} -flat terms t, u , we indicate with $E(t, u)$ the following procedure:

- $E(t, u)$ fails if t is \underline{e} -free and u is not \underline{e} -free (or vice versa);
- $E(t, u)$ fails if $t \equiv e_i$ and (either $t \equiv f(t_1, \dots, t_k)$ or $u \equiv e_j$ for $i \neq j$);
- $E(t, u) = \emptyset$ if $t \equiv u$;
- $E(t, u) = \{t = u\}$ if t and u are different but both \underline{e} -free;
- $E(t, u)$ fails if none of t, u is \underline{e} -free, $t \equiv f(t_1, \dots, t_k)$ and $u \equiv g(u_1, \dots, u_l)$ for $f \neq g$;
- $E(t, u) = E(t_1, u_1) \cup \dots \cup E(t_k, u_k)$ if none of t, u is \underline{e} -free, $t \equiv f(t_1, \dots, t_k)$, $u \equiv f(u_1, \dots, u_k)$ and none of the $E(t_i, u_i)$ fails.

Notice that, whenever $E(t, u)$ succeeds, the formula $\bigwedge E(t, u) \rightarrow t = u$ is universally valid. The definition of $E(t, u)$ is motivated by the next lemma.

Lemma 2. *Let R be a convergent (i.e. terminating and confluent) ground rewriting system, whose rules consist of \underline{e} -free terms. Suppose that t and u are \underline{e} -flat terms with the same R -normal form. Then $E(t, u)$ does not fail and all pairs from $E(t, u)$ have the same R -normal form as well.* \triangleleft

Proof. This is due to the fact that if t is not \underline{e} -free, no R -rewriting is possible at root position because rules from R are \underline{e} -free. \dashv

In the following, we handle *constrained* ground flat literals of the form $L \parallel C$ where L is a ground flat literal and C is a conjunction of ground equalities among \underline{e} -free terms. The logical meaning of $L \parallel C$ is the Horn clause $\bigwedge C \rightarrow L$.

In the literature, various calculi with constrained clauses were considered, starting e.g. from the non-ground constrained versions of the Superposition Calculus of [4,41]. The calculus we propose here is inspired by such versions and it has close similarities

with a subcase of hierarchic superposition calculus [5], or rather to its "weak abstraction" variant from [6] (we thank an anonymous referee for pointing out this connection).

The rules of our *Constrained Superposition Calculus* follow; each rule applies provided the E subprocedure called by it does not fail. The symbol \perp indicates the empty clause. Further explanations and restrictions to the calculus are given in the Remarks below.

Superposition Right (Constrained)	$\frac{l = r \parallel C \quad s = t \parallel D}{s[r]_p = t \parallel C \cup D \cup E(s _p, l)}$	if $l > r$ and $s > t$
Superposition Left (Constrained)	$\frac{l = r \parallel C \quad s \neq t \parallel D}{s[r]_p \neq t \parallel C \cup D \cup E(s _p, l)}$	if $l > r$ and $s > t$
Reflexion (Constrained)	$\frac{t \neq u \parallel C}{\perp \parallel C \cup E(t, u)}$	
Demodulation (Constrained)	$\frac{L \parallel C, \quad l = r \parallel D}{L[r]_p \parallel C}$	if $l > r, L _p \equiv l$ and $C \supseteq D$

Remark 1. The first three rules are inference rules: they are non-deterministically selected for application, until no rule applies anymore. The selection strategy for the rule to be applied is not relevant for the correctness and completeness of the algorithm (some variant of a 'given clause algorithm' can be applied). An inference rule *is not applied in case one premise is \underline{e} -free* (we have no reason to apply inferences to \underline{e} -free premises, since we are not looking for a refutation). \triangleleft

Remark 2. The Demodulation rule is a simplification rule: its application not only adds the conclusion to the current set of constrained literals, but it also removes the first premise. It is easy to see (e.g., representing literals as multisets of terms and extending the total reduction ordering to multisets), that one cannot have an infinite sequence of consecutive applications of Demodulation rules. \triangleleft

Remark 3. The calculus takes $\{L \parallel \emptyset \mid L \text{ is a flat literal from the matrix of } \phi\}$ as the initial set of constrained literals. It terminates when a *saturated* set of constrained literals is reached. We say that S is saturated iff every constrained literal that can be produced by an inference rule, after being exhaustively simplified via Demodulation, is already in S (there are more sophisticated notions of 'saturation up to redundancy' in the literature, but we do not need them). When it reaches a saturated set S , the algorithm outputs the conjunction of the clauses $\bigwedge C \rightarrow L$, varying $L \parallel C$ among the \underline{e} -free constrained literals from S . \triangleleft

We need some rule application policy to ensure termination: without any such policy, a set like $\{e = y \parallel \emptyset, f(e) = e \parallel \emptyset\}$ may produce by Right Superposition the infinitely many literals (all oriented from right to left) $f(y) = e \parallel \emptyset, f(f(y)) = e \parallel \emptyset, f(f(f(y))) = e \parallel \emptyset$, etc. The next Remark explains the policy we follow.

Remark 4. First, we apply Demodulation *only in case the second premise is of the kind $e_j = t(y) \parallel D$, where t is \underline{e} -free*. Demodulation rule is applied with higher priority with respect to the inference rules. Inside all possible applications of Demodulation rule, we give priority to the applications where *both premises have the form $e_j = t(y) \parallel D$* (for the same e_j but with possibly different D 's - the D from the second premise being included in the D of the first). In case we have two constrained literals of the kind $e_j = t_1(y) \parallel D, e_j = t_2(y) \parallel D$ inside our current set of constrained literals (notice that the e_j 's and the D 's

here are the same), among the two possible applications of the Demodulation rule, we apply the rule that keeps the smallest t_i . Notice that in this way two different constrained literals cannot simplify each other. \triangleleft

We say that a constrained literal $L \parallel C$ belonging to a set of constrained literals S is *simplifiable in S* iff it is possible to apply (according to the above policy) a Demodulation rule removing it. A first effect of our policy is:

Lemma 3. *If a constrained literal $L \parallel C$ is simplifiable in S , then after applying to S any sequence of rules, it remains simplifiable until it gets removed. After being removed, if it is regenerated, it is still simplifiable and so it is eventually removed again.* \triangleleft

Proof. Suppose that $L \parallel C$ can be simplified by $e = t \parallel D$ and suppose that a rule is applied to the current set of constrained literals. Since there are simplifiable constrained literals, that rule cannot be an inference rule by the priority stated in Remark 4. For simplification rules, keep in mind again Remark 4. If $L \parallel C$ is simplified, it is removed; if none of $L \parallel C$ and $e = t \parallel D$ get simplified, the situation does not change; if $e = t \parallel D$ gets simplified, this can be done by some $e = t' \parallel D'$, but then $L \parallel C$ is still simplifiable - although in a different way - using $e = t' \parallel D'$ (we have that D' is included in D , which is in turn included in C). Similar observations apply if $L \parallel C$ is removed and re-generated. \dashv

Due to the above Lemma, if we show that a derivation (i.e. a sequence of rule applications) can produce terms only from a finite set, it is clear that when no new constrained literal is produced, saturation is reached. First notice that

Lemma 4. *Every constrained literal $L \parallel C$ produced during the run of the algorithm is \underline{e} -flat.* \triangleleft

Proof. The constrained literals from initialization are \underline{e} -flat. The Demodulation rule, applied according to Remark 4, produces an \underline{e} -flat literal out of an \underline{e} -flat literal. The same happens for the Superposition rules: in fact, since both the terms s and l from these rules are \underline{e} -flat, a Superposition may take place at root position or may rewrite some $l \equiv e_j$ with $r \equiv e_i$ or with $r \equiv t(y)$. \dashv

There are in principle infinitely many \underline{e} -flat terms that can be generated out of the \underline{e} -flat terms occurring in ϕ (see the above counterexample). We show however that only finitely many \underline{e} -flat terms can in fact occur during saturation and that one can determine in advance the finite set they are taken from.

To formalize this idea, let us introduce a hierarchy of \underline{e} -flat terms. Let D_0 be the \underline{e} -flat terms occurring in ϕ and let D_{k+1} be the set of \underline{e} -flat terms obtained by simultaneous rewriting of an \underline{e} -flat term from $\bigcup_{i \leq k} D_i$ via rewriting rules of the kind $e_j \rightarrow t_j(\underline{y})$ where the t_j are \underline{e} -flat \underline{e} -free terms from $\bigcup_{i \leq k} D_i$. The *degree* of an \underline{e} -flat term is the minimum k such that it belongs to set D_k (it is necessary to take the minimum because the same term can be obtained in different stages and via different rewritings).⁵

Lemma 5. *Let the \underline{e} -flat term t' be obtained by a rewriting $e_j \rightarrow u(\underline{y})$ from the \underline{e} -flat term t ; then, if t has degree $k > 1$ and u has degree at most $k - 1$, we have that t' has degree at most k .* \triangleleft

⁵ Notice that, in the above definition of degree, constraints (attached to the rewriting rules occurring in our calculus) are ignored.

Proof. This is clear, because at the k -stage one can directly produce t' instead of just t : in fact, all rewriting producing directly t' replace an occurrence of some e_i by an \underline{e} -free term, so they are all done in parallel positions. \dashv

Proposition 3. *The saturation of the initial set of \underline{e} -flat constrained literals always terminates after finitely many steps.* \triangleleft

Proof. We show that all \underline{e} -flat terms that may occur during saturation have at most degree n (where n is the cardinality of \underline{e}). This shows that the saturation must terminate, because only finitely many terms may occur in a derivation (see the above observations). Let the algorithm during saturation reach the status S ; we say that a constraint C allows the explicit definition of e_j in S iff S contains a constrained literal of the kind $e_j = t(\underline{y}) \parallel D$ with $D \subseteq C$. Now we show by mutual induction two facts concerning a constrained literal $L \parallel C \in S$:

- (1) if an \underline{e} -flat term u of degree k occurs in L , then C allows the explicit definition of k different e_j in S ;
- (2) if L is of the kind $e_i = t(\underline{y})$, for an \underline{e} -flat \underline{e} -free term t of degree k , then either $e_i = t \parallel C$ can be simplified in S or C allows the explicit definition of $k + 1$ different e_j in S (e_i itself is of course included among these e_j).

Notice that (1) is sufficient to exclude that any \underline{e} -flat term of degree bigger than n can occur in a constrained literal arising during the saturation process.

We prove (1) and (2) by induction on the length of the derivation leading to $L \parallel C \in S$. Notice that it is sufficient to check that (1) and (2) hold for the first time where $L \parallel C \in S$ because if C allows the explicit definition of a certain variable in S , it will continue to do so in any S' obtained from S by continuing the derivation (the definition may be changed by the Demodulation rule, but the fact that e_i is explicitly defined is forever). Also, by Lemma 3, a literal cannot become non simplifiable if it is simplifiable.

(1) and (2) are evident if S is the initial status. To show (1), suppose that u occurs for the first time in $L \parallel C$ as the effect of the application of a certain rule: we can freely assume that u does not occur in the literals from the premisses of the rule (otherwise induction trivially applies) and that u of degree k is obtained by rewriting in a non-root position some u' occurring in a constrained literal $L' \parallel D'$ via some $e_j \rightarrow t \parallel D$. This might be the effect of a Demodulation or Superposition in a non-root position (Superpositions in root position do not produce new terms). If u' has degree k , then by induction D' contains the required k explicit definitions, and we are done because D' is included in C . If u' has lower degree, then t must have degree at least $k - 1$ (otherwise u does not reach degree k by Lemma 5). Then by induction on (2), the constraint D (also included in C) has $(k - 1) + 1 = k$ explicit definitions (when a constraint $e_j \rightarrow t \parallel D$ is selected for Superposition or for making Demodulations in a non-root position, it is itself not simplifiable according to the procedure explained in Remark 4).

To show (2), we analyze the reasons why the non simplifiable constrained literal $e_i = t(\underline{y}) \parallel C$ is produced (let k be the degree of t). Suppose it is produced from $e_i = u' \parallel C$ via Demodulation with $e_j = u(\underline{y}) \parallel D$ (with $D \subseteq C$) in a non-root position; if u' has degree at least k , we apply induction for (1) to $e_i = u' \parallel C$: by such induction hypotheses, we get k explicit definitions in C and we can add to them the further explicit definition $e_i = t(\underline{y})$ (the explicit definitions from C cannot concern e_i because $e_i = t(\underline{y}) \parallel C$ is not simplifiable). Otherwise, u' has degree less than k and u has degree at least $k - 1$ by

Lemma 5 (recall that t has degree k): by induction, $e_j = u \parallel D$ is not simplifiable (it is used as the active part of a Demodulation in a non-root position, see Remark 4) and supplies k explicit definitions, inherited by $C \supseteq D$. Note that e_i cannot have a definition in D , otherwise $e_i = t(\underline{y}) \parallel C$ would be simplifiable, so with $e_i = t(\underline{y}) \parallel C$ we get the required $k + 1$ definitions.

The remaining case is when $e_i = t(\underline{y}) \parallel C$ is produced via Superposition Right. Such a Superposition might be at root or at a non-root position. We first analyse the case of a root position. This might be via $e_j = e_i \parallel C_1$ and $e_j = t(\underline{y}) \parallel C_2$ (with $e_j > e_i$ and $C = C_1 \cup C_2$ because $E(e_j, e_j) = \emptyset$), but in such a case one can easily apply induction. Otherwise, we have a different kind of Superposition at root position: $e_i = t(\underline{y}) \parallel C$ is obtained from $s = e_i \parallel C_1$ and $s' = t(\underline{y}) \parallel C_2$, with $C = C_1 \cup C_2 \cup E(s, s')$. In this case, by induction for (1), C_2 supplies k explicit definitions, to be inherited by C . Among such definitions, there cannot be an explicit definition of e_i otherwise $e_i = t(\underline{y}) \parallel C$ would be simplifiable, so again we get the required $k + 1$ definitions.

In case of a Superposition at a non root-position, we have that $e_i = t(\underline{y}) \parallel C$ is obtained from $u' = e_i \parallel C_1$ and $e_j = u(\underline{y}) \parallel C_2$, with $C = C_1 \cup C_2$; here t is obtained from u' by rewriting e_j to u . This case is handled similarly to the case where $e_i = t(\underline{y}) \parallel C$ is obtained via Demodulation rule. \dashv

Having established termination, we now prove that our calculus computes covers; to this aim, we rely on refutational completeness of unconstrained Superposition Calculus (thus, our technique resembles the technique used [5,6] in order to prove refutational completeness of hierarchic superposition, although it is not clear whether Theorem 2 below can be derived from the results concerning hierarchic superposition - we are not just proving refutational completeness and we need to build proper superstructures):

Theorem 2. *Suppose that the above algorithm, taking as input the primitive \underline{e} -flat formula $\exists \underline{e} \phi(\underline{e}, \underline{y})$, gives as output the quantifier-free formula $\psi(\underline{y})$. Then the latter is a cover of $\exists \underline{e} \phi(\underline{e}, \underline{y})$.* \triangleleft

Proof. Let S be the saturated set of constrained literals produced upon termination of the algorithm; let $S = S_1 \cup S_2$, where S_1 contains the constrained literals in which the \underline{e} do not occur and S_2 is its complement. Clearly $\exists \underline{e} \phi(\underline{e}, \underline{y})$ turns out to be logically equivalent to

$$\bigwedge_{L \parallel C \in S_1} (\bigwedge C \rightarrow L) \wedge \exists \underline{e} \bigwedge_{L \parallel C \in S_2} (\bigwedge C \rightarrow L)$$

so, as a consequence, in view of Lemma 1 it is sufficient to show that every model \mathcal{M} satisfying $\bigwedge_{L \parallel C \in S_1} (\bigwedge C \rightarrow L)$ via an assignment \mathcal{I} to the variables \underline{y} can be embedded into a model \mathcal{M}' such that for a suitable extension \mathcal{I}' of \mathcal{I} to the variables \underline{e} we have that $(\mathcal{M}', \mathcal{I}')$ satisfies also $\bigwedge_{L \parallel C \in S_2} (\bigwedge C \rightarrow L)$.

Fix \mathcal{M}, \mathcal{I} as above. The diagram $\Delta(\mathcal{M})$ of \mathcal{M} is obtained as follows. We take one free constant for each element of the support of \mathcal{M} (by Löwenheim-Skolem theorem you can keep \mathcal{M} at most countable, if you like) and we put in $\Delta(\mathcal{M})$ all the literals of the kind $f(c_1, \dots, c_k) = c_{k+1}$ and $c_1 \neq c_2$ which are true in \mathcal{M} (here the c_i are names for the elements of the support of \mathcal{M}). Let R be the set of ground equalities of the form $y_i = c_i$, where c_i is the name of $\mathcal{I}(y_i)$. Extend our reduction ordering in the natural way (so that $y_i = c_i$ and $f(c_1, \dots, c_k) = c_{k+1}$ are oriented from left to right). Consider now

the set of clauses

$$\Delta(\mathcal{M}) \cup R \cup \{\bigwedge C \rightarrow L \mid (L \parallel C) \in S\} \quad (5)$$

(below, we distinguish the positive and the negative literals of $\Delta(\mathcal{M})$ so that $\Delta(\mathcal{M}) = \Delta^+(\mathcal{M}) \cup \Delta^-(\mathcal{M})$). We want to saturate the above set in the standard Superposition Calculus. Clearly the rewriting rules in R , used as reduction rules, replace everywhere y_i by c_i inside the clauses of the kind $\bigwedge C \rightarrow L$. At this point, the negative literals from the equality constraints all disappear: if they are true in \mathcal{M} , they $\Delta^+(\mathcal{M})$ -normalize to trivial equalities $c_i = c_i$ (to be eliminated by standard reduction rules) and if they are false in \mathcal{M} they become part of clauses subsumed by true inequalities from $\Delta^-(\mathcal{M})$. Similarly all the e -free literals not coming from $\Delta(\mathcal{M}) \cup R$ get removed. Let \tilde{S} be the set of survived literals involving the e (they are not constrained anymore and they are $\Delta^+(\mathcal{M}) \cup R$ -normalized): we show that they cannot produce new clauses. Let in fact (π) be an inference from the Superposition Calculus [42] applying to them. Since no superposition with $\Delta(\mathcal{M}) \cup R$ is possible, this inference must involve only literals from \tilde{S} ; suppose it produces a literal \tilde{L} from the literals \tilde{L}_1, \tilde{L}_2 (coming via $\Delta^+(\mathcal{M}) \cup R$ -normalization from $L_1 \parallel C_1 \in S$ and $L_2 \parallel C_2 \in S$) as parent clauses. Then, by Lemma 2, our constrained inferences produce a constrained literal $L \parallel C$ such that the clause $\bigwedge C \rightarrow L$ normalizes to \tilde{L} via $\Delta^+(\mathcal{M}) \cup R$. Since S is saturated, the constrained literal $L \parallel C$, after simplification, belongs to S . Now simplifications via our Constrained Demodulation and $\Delta(\mathcal{M})^+ \cup R$ -normalization commute (they work at parallel positions, see Remark 4), so the inference (π) is redundant because \tilde{L} simplifies to a literal already in $\tilde{S} \cup \Delta(\mathcal{M})$.

Thus the set of clauses (5) saturates without producing the empty clause. By the completeness theorem of the Superposition Calculus [28,3,42] it has a model \mathcal{M}' . This \mathcal{M}' by construction fits our requests by Robinson Diagram Lemma. \dashv

Theorem 2 also proves the existence of the model completion of (EUF).

Example 1. We compute the cover of the primitive formula $\exists e f(f(e, y_1), y_2) \neq f(f(e, y'_1), y'_2)$. Flattening gives the set of literals

$$\{ f(e, y_1) = e_1, f(e_1, y_2) = e'_1, f(e, y'_1) = e_2, f(e_2, y'_2) = e'_2, e'_1 \neq e'_2 \} .$$

Superposition Right produces the constrained literal $e_1 = e_2 \parallel \{y_1 = y'_1\}$; supposing that we have $e_1 > e_2$, Superposition Right gives first $f(e_2, y_2) = e'_1 \parallel \{y_1 = y'_1\}$ and then also $e'_1 = e'_2 \parallel \{y_1 = y'_1, y_2 = y'_2\}$. Superposition Left and Reflexion now produce $\perp \parallel \{y_1 = y'_1, y_2 = y'_2\}$. Thus the clause $y_1 = y'_1 \wedge y_2 = y'_2 \rightarrow \perp$ will be part of the output (actually, this will be the only clause in the output). \triangleleft

Example 2. We add one more example, taken from [26]. We compute the cover of the primitive formula $\exists e (s_1 = f(y_3, e) \wedge s_2 = f(y_4, e) \wedge t = f(f(y_1, e), f(y_2, e)))$, where s_1, s_2, t are terms in \underline{y} . This example is taken from [26]. Flattening gives the set of literals

$$\{ f(y_3, e) = s_1, f(y_4, e) = s_2, f(y_1, e) = e_1, f(y_2, e) = e_2, f(e_1, e_2) = t \} .$$

Suppose that we have $e > e_1 > e_2 > t > s_1 > s_2 > y_1 > y_2 > y_3 > y_4$. Superposition Right between the 3rd and the 4th clauses produces the constrained 6th clause $e_1 = e_2 \parallel \{y_1 = y_2\}$. From now on, we denote the application of a Superposition Right to the i th and j th clauses with $R(i, j)$. We list a derivation performed by our calculus:

$$R(3, 4) \implies e_1 = e_2 \parallel \{y_1 = y_2\} \quad (6\text{th clause})$$

$$\begin{aligned}
R(1,2) &\implies s_1 = s_2 \parallel \{y_3 = y_4\} \quad (7\text{th clause}) \\
R(5,6) &\implies f(e_2, e_2) = t \parallel \{y_1 = y_2\} \quad (8\text{th clause}) \\
R(1,3) &\implies e_1 = s_1 \parallel \{y_1 = y_3\} \quad (9\text{th clause}) \\
R(1,4) &\implies e_2 = s_1 \parallel \{y_2 = y_3\} \quad (10\text{th clause}) \\
R(2,3) &\implies e_1 = s_2 \parallel \{y_1 = y_4\} \quad (11\text{th clause}) \\
R(2,4) &\implies e_2 = s_2 \parallel \{y_2 = y_4\} \quad (12\text{th clause}) \\
R(5,9) &\implies f(s_1, e_2) = t \parallel \{y_1 = y_3\} \quad (13\text{th clause}) \\
R(5,11) &\implies f(s_2, e_2) = t \parallel \{y_1 = y_4\} \quad (14\text{th clause}) \\
R(6,9) &\implies e_2 = s_1 \parallel \{y_1 = y_3, y_1 = y_2\} \quad (15\text{th clause}) \\
R(6,11) &\implies e_2 = s_2 \parallel \{y_1 = y_2, y_1 = y_4\} \quad (16\text{th clause}) \\
R(8,10) &\implies f(s_1, s_1) = t \parallel \{y_1 = y_3, y_2 = y_3\} \quad (17\text{th clause}) \\
R(8,12) &\implies f(s_2, s_2) = t \parallel \{y_1 = y_4, y_2 = y_4\} \quad (18\text{th clause}) \\
R(13,12) &\implies f(s_1, s_2) = t \parallel \{y_1 = y_3, y_2 = y_4\} \quad (19\text{th clause}) \\
R(14,10) &\implies f(s_2, s_1) = t \parallel \{y_1 = y_4, y_2 = y_3\} \quad (20\text{th clause}) \\
R(9,11) &\implies s_1 = s_2 \parallel \{y_1 = y_3, y_1 = y_4\} \quad (21\text{th clause})
\end{aligned}$$

The set of clauses above is saturated. The 7th, 17th, 18th, 19th and 20th clauses are exactly the output clauses of [26]. The non-simplified clauses that do not appear as output in [26] are redundant and they could be simplified by introducing a Subsumption rule as an additional simplification rule of our Constrained Superposition Calculus.

6 Complexity Analysis

In the special case where the signature Σ contains only unary function symbols, only empty constraints can be generated; in case Σ contains also relation symbols of arity $n > 1$, the only constrained clauses that can be generated have the form $\perp \parallel \{t_1 = t'_1, \dots, t_{n-1} = t'_{n-1}\}$. Also, it is not difficult to see that in a derivation at most one explicit definition $e_i = t(\underline{y}) \parallel \emptyset$ can occur for every e_i : as soon as this definition is produced, all occurrences of e_i are rewritten to t . This shows that Constrained Superposition computes covers in polynomial time for the empty theory, whenever the signature Σ matches the restrictions of Definition 3 for DB schemata. We give here a finer complexity analysis, in order to obtain a quadratic bound.

In this section, *we assume that our signature Σ contains only unary function and m -ary relation symbols*. As we already mentioned, the complexity of the saturation becomes polynomial in this restricted case. However, in order to attain an optimized quadratic complexity bound, we need to follow a *different strategy in applying the rules of our constrained superposition calculus* (this different strategy would not be correct for the general case). Thanks to this different strategy, we can make our procedure

close to the algorithm of [26]: in fact, such algorithm is correct for the case of unary functions and requires only a very minor adjustment for the case of unary functions and m -ary relations (the reason why only a very minor adjustment is sufficient is due to the fact that the empty clause is the only constrained clause that is generated during saturation, so that constraints do not become part of a recursive mechanism).

Since relations play a special role in the present restricted context, we prefer to treat them as such, i.e. not to rewrite $R(t_1, \dots, t_n)$ as $R(t_1, \dots, t_n) = true$; the consequence is that we need an additional Constrained Resolution Rule⁶. We preliminarily notice that when function symbols are all unary, the constraints remain all *empty* during the run of the saturation procedure, except for the case of the newly introduced Resolution Rule below. This fact follows from the observation that given two terms u_1 and u_2 , procedure $E(u_1, u_2)$ does not fail iff:

- (1) either u_1 and u_2 are both terms containing only variables from \underline{y} , or
- (2) u_1 and u_2 are terms that syntactically coincide.

In case (1), $E(u_1, u_2)$ is $\{u_1, u_2\}$ and in case (2), $E(u_1, u_2)$ is \emptyset . In case (1), Superposition Rules are not applicable. To show this, suppose that $u_1 \equiv s|_p$ and $u_2 \equiv l$; then, terms l and r use only variables from \underline{y} , and consequently cannot be fed into Superposition Rules, since Superposition Rules are only applied when variables from \underline{e} occur in both premises. Reflexion Rule does not apply too in case (1), because this rule (like any other rule) cannot be applied to an \underline{e} -free literal.

Thus, in the particular case of m -ary relations and unary functions, the rules of the calculus are the following:

Superposition	$\frac{l = r \quad L}{L[r]_p}$	if (i) $l > r$; (ii) if $L \equiv s = t$ or $L \equiv s \neq t$, then $s > t$ and $p \in Pos(s)$; (iii) $E(s _p, l)$ does not fail.
Resolution	$\frac{R(t_1, \dots, t_n) \quad \neg R(s_1, \dots, s_n)}{\perp \parallel \bigcup_i E(s_i, t_i)}$	if $E(s_i, t_i)$ does not fail for all $i = 1, \dots, n$
Reflexion	$\frac{t \neq u}{\perp}$	if $E(t, u)$ does not fail
Demodulation	$\frac{L \quad l = r}{L[r]_p}$	if $l > r$ and $L _p \equiv l$

We still restrict the use of our rules to the case where all premises are not \underline{e} -free literals; again Demodulation is applied only in the case where $l = r$ is of the kind $e_i = t(\underline{y})$. For the order of applications of the Rules, Lemma 6 below show that we can apply (restricted) Superpositions, Demodulations, Reflexions and Resolutions in this order and then stop.

An important preliminary observation to obtain such result is that *we do not need to apply Superposition Rules whose left premise $l = r$ is of the kind $e_i = t(\underline{y})$* : this is because constraints are always empty (unless the constrained clause is the empty

⁶ We extend the definition of an \underline{e} -flat literal so as to include also the literals of the kind $R(t_1, \dots, t_n)$ and $\neg R(t_1, \dots, t_n)$ where the terms t_i are either \underline{e} -free terms or variables from \underline{e} .

clause), so that a Superposition Rule with the left premise $e_i = t(\underline{y})$ can be replaced by a Demodulation Rule.⁷ If the left premise of Superposition is not of the kind $e_i = t(\underline{y})$, then since our literals are \underline{e} -flat, it can be either of the kind $e_i = e_j$ (with $e_i > e_j$) or of the kind $f(e_i) = t$. In the latter case t is either $e_k \in \underline{e}$ or it is an \underline{e} -free term; for Superposition Left (i.e. for Superposition applied to a negative literal), the left premise can only be $e_i = e_j$, because our literals are \underline{e} -flat and so negative literals L cannot have a position p such that $L|_p \equiv f(e_i)$.

Let S be a set of \underline{e} -flat literals with empty constraints; we say that S is *RS-closed* iff it is closed under *Restricted Superposition Rules*, i.e. under Superposition Rules whose left premise is not of the kind $e_i = t(\underline{y})$. In equivalent terms, as a consequence of the above discussion, S is RS-closed iff it satisfies the following two conditions:

- if $\{f(e_i) = t, f(e_i) = v\} \subseteq S$, then $t = v \in S$;
- if $\{e_i = e_j, L\} \subseteq S$ and $e_i > e_j$ and $L|_p \equiv e_i$, then $L[e_j]_p \in S$.

Since Restricted Superpositions do not introduce essentially new terms (newly introduced terms are just rewritings of variables with variables), it is clear that we can make a finite set S of \underline{e} -free literals RS-closed in finitely many steps. This can be naively done in time quadratic in the size of the formula. As an alternative, we can apply a *congruence closure algorithm* to S and produce a set of \underline{e} -free constraints S' which is RS-closed and logically equivalent to S : the latter can be done in $O(n \cdot \log(n))$ -time, as it is well-known from the literature [36,40,29].

Lemma 6. *Let S be a RS-closed set of empty-constrained \underline{e} -flat literals. Then, to saturate S it is sufficient to first exhaustively apply the Demodulation Rule, and then Reflexion and Resolution Rules.* ◁

Proof. Let \tilde{S} be the set obtained from S after having exhaustively applied Demodulation. Notice that the final effect of the reiterated application of Demodulation can be synthetically described by saying that literals in S are rewritten by using some explicit definitions

$$e_{i_1} = t_1(\underline{y}), \dots, e_{i_k} = t_k(\underline{y}) . \quad (6)$$

These definitions are either in S , or are generated through the Demodulations themselves (we can freely assume that Demodulations are done in appropriate order: first all occurrences of e_{i_1} are rewritten to t_1 , then all occurrences of e_{i_2} are rewritten to t_2 , etc.).⁸

Suppose now that a pair $L, l = r \in \tilde{S}$ can generate a new literal $L[r]_p$ by Superposition. We know from above that we can limit ourselves to Restricted Superposition, so l is either of the form e_j or of the form $f(e_j)$, where moreover e_j is not among the set $\{e_{i_1}, \dots, e_{i_k}\}$ from (6). The literals L and $l = r \in \tilde{S}$ happen to have been obtained from literals L' and $l = r'$ belonging to S by applying the rewriting rules (6) (notice that l cannot have been rewritten). Since such rewritings must have occurred in positions parallel to p and since S was closed under Restricted Superposition, we must have that

⁷ This is not true in the general case where constraints are not empty, because the Demodulation Rule does not merge incomparable constraints.

⁸ In addition, if we happen to have, say, two different explicit definitions of e_{i_1} as $e_{i_1} = t_1, e_{i_1} = t'_1$, we decide to use just one of them (and always the same one, until the other one is eventually removed by Demodulation).

S contained the literal $L'[r']_p$ that rewrites to $L[r]_p$ by the rewriting rules (6). This shows that $L[r]_p$ is already in S and proves the lemma, because Reflexion and Resolution can only produce the empty clause and no rule applies to the empty clause. \dashv

Thus the strategy of applying (in this order)

Restricted Superposition+Demodulation+Reflexion+Resolution

always saturates.

To produce an output in optimized format, it is convenient to get it in a dag-like form. This can be simulated via explicit acyclic definitions as follows. When we write $Def(\underline{e}, \underline{y})$ (where $\underline{e}, \underline{y}$ are tuples of distinct variables), we mean any flat formula of the kind (let $\underline{e} := e_1 \dots, e_n$)

$$\bigwedge_{i=1}^n e_i = t_i$$

where in the term t_i only the variables $e_1, \dots, e_{i-1}, \underline{y}$ can occur. We shall supply the output in the form

$$\exists \underline{e}' (Def(\underline{e}', \underline{y}) \wedge \psi(\underline{e}', \underline{y})) \quad (7)$$

where the \underline{e}' is a subset of the \underline{e} and ψ is quantifier-free. The *dag-format* (7) is not quantifier-free but can be converted to a quantifier-free formula by unravelling the acyclic definitions of the \underline{e}' .

Thus our procedure for computing a cover in dag-format of a primitive formula $\exists \underline{e} \phi(\underline{e}, \underline{y})$ (in case the function symbols of the signature Σ are all unary) runs by performing the following steps, one after the other. Let OUT be a quantifier-free formula (initially OUT is \top).

- (1) We preprocess ϕ in order to get out of it a RS-closed set S of empty-constrained \underline{e} -flat literals.
- (2) We *mark* the variables \underline{e} in the following way (initially, all variables are unmarked): we scan S and, as soon as we find an equality of the kind $e_i = t$ where all variables from \underline{e} occurring in t are marked, we mark e_i . This loop is repeated until no more variable gets marked.
- (3) If Reflexion is applicable, output \perp and exit.
- (4) Conjoin OUT with all literals where, besides the \underline{y} , only marked variables occur.
- (5) For every literal $R(t_1, \dots, e, \dots, t_m)$ that contains at least an unmarked e , scan S until a literal of the type $\neg R(t_1, \dots, e, \dots, t_m)$ is found: then, try to apply Resolution and if you succeed getting $\perp \parallel \{u_1 = u'_1, \dots, u_m = u'_m\}$ conjoin $\bigvee_j u_j \neq u'_j$ to OUT .
- (6) Prefix to OUT a string of existential quantifiers binding all marked variables and output the result.

One remark is in order: when running the subprocedures $E(s_i, t_i)$ required by the Resolution Rule in (5) above, *you must consider all marked variables as part of the \underline{y}* (thus, e.g. $R(e, t), \neg R(e, v)$ produces $\perp \parallel \{t = u\}$ if both t and u contain, besides the \underline{y} , only marked variables).

Proposition 4. *Let T be the theory (EUF) in a signature with unary functions and m -ary relation symbols. Consider a primitive formula $\exists \underline{e} \phi(\underline{e}, \underline{y})$; then, the above algorithm returns a cover of $\exists \underline{e} \phi(\underline{e}, \underline{y})$ in dag-format in time $O(n^2)$, where n is the size of $\exists \underline{e} \phi(\underline{e}, \underline{y})$. \triangleleft*

Proof. The preprocessing step (1) requires an abstraction phase for producing \underline{e} -flat literals and a second phase in order to get a RS-closed set: the first phase requires linear time, whereas the second one requires $O(n \cdot \log(n))$ time. All the remaining steps require linear time, except steps (2) and (5) that requires quadratic time. This is the dominating cost, thus the entire procedure requires $O(n^2)$ time. \dashv

Although we do not deeply investigate the problem here, we conjecture that it might be possible to further lower down the above complexity to $O(n \cdot \log(n))$.

6.1 An extension

We consider a useful extension of the above algorithm; let us assume that we have a nonempty theory whose axioms are (4), namely

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

for every function symbol f .

One side of the above axiom is equivalent to the ground literal $f(\text{undef}) = \text{undef}$ and as such it does not interfere with the completion process (we just add it to our constraints from the very beginning).

To accommodate the other side, we need to modify our Calculus. First, we add to the Constrained Superposition Calculus of Section 5 the following extra Rule

$$\text{Inference Rule } Ext(\text{undef}) \quad \frac{f(e_j) = u(y) \parallel D}{e_j = \text{undef} \parallel D \cup \{u(y) = \text{undef}\}} \\ \text{(Constrained)}$$

The Rule is sound because $u(y) = \text{undef} \wedge f(e_j) = u(y) \rightarrow e_j = \text{undef}$ follows from the axioms (4). For cover computation with our new axioms, we need a restricted version of Paramodulation Rule:

$$\text{Paramodulation} \quad \frac{e_j = r \parallel C \quad L \parallel D}{L[r]_p \parallel CUD} \quad (\text{if } e_j > r \text{ \& } L|_p \equiv e_j) \\ \text{(Constrained)}$$

Notice that we can have $e_j > r$ only in case r is either some existential variable e_i or it is an \underline{e} -free term $u(y)$. Paramodulation Rule (if it is not a Superposition) can only apply to a right member of an equality and such a right member must be e_j itself (because our literals are flat). Thus the rule cannot introduce new terms and consequently it does not compromise the termination argument of Proposition 3.

The proof of Theorem 2 can be easily adjusted as follows. We proceed as in the proof of Theorem 2, so as to obtain the set $\Delta(\mathcal{M}) \cup R \cup \tilde{\mathcal{S}}$ which is saturated in the standard (unconstrained) Superposition Calculus. Below, we refer to the general refutational completeness proof of the Superposition Calculus given in [42]. Since we only have unit literals here, in order to produce a model of $\Delta(\mathcal{M}) \cup R \cup \tilde{\mathcal{S}}$, we can just consider the convergent ground rewriting system \rightarrow consisting of the oriented equalities in $\Delta^+(\mathcal{M}) \cup R \cup \tilde{\mathcal{S}}$: the support of such model is formed by the \rightarrow -normal forms of our ground terms with the obvious interpretation for the function and constant sym-

bols. For simplicity, we assume that undef is in normal form.⁹ We need to check that whenever we have¹⁰ $f(t) \rightarrow^* \text{undef}$ then we have also $t \rightarrow^* \text{undef}$: we prove this by induction on the reduction ordering for our ground terms. Let t be a term such that $f(t) \rightarrow^* \text{undef}$: if t is \underline{e} -free then the claim is trivial (because the axioms (4) are supposed to hold in \mathcal{M}). Suppose also that induction hypothesis applies to all terms smaller than t . If t is not in normal form, then let \tilde{t} be its normal form; then we have $f(t) \rightarrow^+ f(\tilde{t}) \rightarrow^* \text{undef}$, by the fact that \rightarrow is convergent. By induction hypothesis, $\tilde{t} \rightarrow \text{undef}$, hence $t \rightarrow^+ \tilde{t} \rightarrow^* \text{undef}$, as desired. Finally, let us consider the case in which t is in normal form; since $f(t)$ is reducible in root position by some rule $l \rightarrow r$, our rules $l \rightarrow r$ are \underline{e} -flat and t is not \underline{e} -free, we have that $t \equiv e_j$ for some existential variable e_j . Then, we must have that S contains an equality of the kind $f(e_j) = u(\underline{y}) \parallel D$ or of the kind $f(e_j) = e_i \parallel D$ (the constraint D being true in \mathcal{M} under the given assignment to the \underline{y}). The latter case is reduced to the former, since $e_i \rightarrow^* \text{undef}$ (by the convergence of \rightarrow^*) and since S is closed under Paramodulation. In the former case, by the rule $Ext(\text{undef})$, we must have that S contains $e_j = \text{undef} \parallel D \cup \{u(\underline{y}) = \text{undef}\}$. Now, since $f(e_j) = u(\underline{y}) \parallel D$ belongs to S and D is true in \mathcal{M} , we have that the normal forms of $f(e_j)$ and of $u(\underline{y})$ are the same; since the normal form of $f(e_j)$ is undef , the normal form of $u(\underline{y})$ is undef too, which means that $u(\underline{y}) = \text{undef}$ is true in \mathcal{M} . But $e_j = \text{undef} \parallel D \cup \{u(\underline{y}) = \text{undef}\}$ belongs to S , hence $e_j = \text{undef}$ belongs to \tilde{S} , which implies $e_j \rightarrow^* \text{undef}$, as desired. \dashv

6.2 Remarks on MCMT implementation

As evident from Subsection 4.1, our main motivation for investigating covers originated from the verification of data-aware processes. Such applications require database (DB) signatures to contain only unary function symbols (besides relations of every arity). We observed that computing covers of primitive formulae in such signatures requires only polynomial time. In addition, if relation symbols are at most binary, *the cover of a primitive formula is a conjunction of literals*: this is crucial in applications, because model checkers like MCMT [22] and CUBICLE [15] represent sets of reachable states as primitive formulae. This makes cover computations a quite attractive technique in database-driven model checking.

Our cover algorithm for DB signatures has been implemented in the model checker MCMT. The implementation is however still partial, nevertheless the tool is able to compute covers for the *EUF*-fragment with unary function symbol, unary relations and binary relations. The optimized procedure of Section 6 has not yet been implemented, instead MCMT uses a customary Knuth-Bendix completion (in fact, for the above mentioned fragments constraints are always trivial and our constrained Superposition Calculus essentially boils down to Knuth-Bendix completion for ground literals in *EUF*).

Axioms (4) are also covered in the following way. We assume that constraints of which we want to compute the cover always contains either the literal $e_j = \text{undef}$ or the literal $e_j \neq \text{undef}$ for every existential variable e_j . Whenever a constraint contains the

⁹ To be pedantic, according to the definition of $\Delta^+(\mathcal{M})$, there should be an equality $\text{undef} = c_0$ in $\Delta^+(\mathcal{M})$ so that c_0 is the normal form of undef .

¹⁰ We use \rightarrow^* for the reflexive-transitive closure of \rightarrow and \rightarrow^+ for the transitive closure of \rightarrow .

literal $e_j \neq \text{undef}$, the completion procedure adds the literal $u(y_i) \neq \text{undef}$ whenever it had produced a literal of the kind $f(e_j) = u(y_i)$.¹¹

We wonder whether we are justified in assuming that all onstraints of which we want to compute the cover always contains either the literal $e_j = \text{undef}$ or the literal $e_j \neq \text{undef}$ for every existential variable e_j . The answer is the affirmative: according to the backward search algorithm implemented in array-based systems tools, the variable e_j to be eliminated always comes from the guard of a transition and we can assume that such a guard contains the literal $e_j \neq \text{undef}$ (if we need a transition with $e_j = \text{undef}$ - for an existentially quantified variable e_j - it is possible to write trivially this condition without using a quantified variable). The MCMT User Manual (available from the distribution) contains precise instructions on how to write specifications following the above prescriptions.

A first experimental evaluation (based on the existing benchmark provided in [32], which samples 32 real-world BPMN workflows taken from the BPMN official website <http://www.bpmn.org/>) is described in [10]. 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 giving essential information on how to encode relational artifact systems (comprising *both first order and second order variables*) in MCMT specifications and how to produce user-defined examples in the database driven framework. The first experiments were very encouraging: the tool was able to solve in few seconds all the proposed benchmarks and the cover computations generated automatically during model-checking search were discharged instantaneously, see [10] for more information about our experiments.

7 Conclusions and Future Work

The above experimental setup motivates new research to extend Proposition 2 to further theories axiomatizing integrity constraints used in DB applications. Combined cover algorithms (along the perspectives in [26]) could be crucial also in this setting. Practical algorithms for the computation of covers in the theories falling under the hypotheses of Proposition 2 need to be designed: as a little first example, in Subsection 6.1 above we showed how to handle Axiom (4) by light modifications to our techniques. Symbol elimination of function and predicate variables should also be combined with cover computations.

We consider the present work, together with [10,12], as the starting point for a full line of research dedicated to SMT-based techniques for the effective verification of data-aware processes, addressing richer forms of verification beyond safety (such as liveness, fairness, or full LTL-FO) and richer classes of artifact systems, (e.g., with concrete data types and arithmetics), while identifying novel decidable classes (e.g., by restricting the structure of the DB and of transition and state formulae) beyond the ones presented in [10,12]. Concerning implementation, we plan to further develop our tool to incorporate in it the plethora of optimizations and sophisticated search strategies available

¹¹ This is sound because $e \neq \text{undef}$ implies $f(e) \neq \text{undef}$ according to (4), so $u(y_i) \neq \text{undef}$ follows from $f(e_j) = u(y_i)$ and $e \neq \text{undef}$.

in infinite-state SMT-based model checking. Finally, in [11] we tackle more conventional process modeling notations, concerning in particular data-aware extensions of the de-facto standard BPMN¹²: we plan to provide a full-automated translator from the data-aware BPMN model presented in [11] to the artifact systems setting of [12].

References

1. F. Baader, S. Ghilardi, and C. Tinelli. A new combination procedure for the word problem that generalizes fusion decidability results in modal logics. *Inform. and Comput.*, pages 1413–1452, 2006.
2. F. Baader and T. Nipkow. *Term Rewriting and All That*. Cambridge University Press, United Kingdom, 1998.
3. L. Bachmair and H. Ganzinger. Rewrite-based equational theorem proving with selection and simplification. *J. Log. Comput.*, 4(3):217–247, 1994.
4. L. Bachmair, H. Ganzinger, C. Lynch, and W. Snyder. Basic paramodulation. *Inform. and Comput.*, 121(2):172–192, 1995.
5. L. Bachmair, H. Ganzinger, and U. Waldmann. Refutational theorem proving for hierarchic first-order theories. *Appl. Algebra Eng. Commun. Comput.*, 5:193–212, 1994.
6. P. Baumgartner and U. Waldmann. Hierarchic superposition with weak abstraction. In *Proc. of CADE*, pages 39–57, 2013.
7. M. Bojańczyk, L. Segoufin, and S. Toruńczyk. Verification of database-driven systems via amalgamation. In *Proc. of PODS*, pages 63–74, 2013.
8. R. Bruttomesso, S. Ghilardi, and S. Ranise. Quantifier-free interpolation in combinations of equality interpolating theories. *ACM Trans. Comput. Log.*, 15(1):5:1–5:34, 2014.
9. D. Calvanese, G. De Giacomo, and M. Montali. Foundations of data aware process analysis: A database theory perspective. In *Proc. of PODS*, 2013.
10. D. Calvanese, S. Ghilardi, A. Gianola, M. Montali, and A. Rivkin. Verification of data-aware processes via array-based systems (extended version). Technical Report arXiv:1806.11459, arXiv.org, 2018.
11. D. Calvanese, S. Ghilardi, A. Gianola, M. Montali, and A. Rivkin. Formal modeling and SMT-based parameterized verification of data-aware BPMN. In *Proc. of BPM*, 2019.
12. D. Calvanese, S. Ghilardi, A. Gianola, M. Montali, and A. Rivkin. From model completeness to verification of data aware processes. In *Description Logic, Theory Combination, and All That*. Springer, 2019.
13. D. Calvanese, S. Ghilardi, A. Gianola, M. Montali, and A. Rivkin. Model completeness, covers and superposition. In *Proc. of CADE*, 2019.
14. C.-C. Chang and J. H. Keisler. *Model Theory*. North-Holland Publishing Co., Amsterdam-London, third edition, 1990.
15. S. Conchon, A. Goel, S. Krstic, A. Mebsout, and F. Zaïdi. Cubicle: A parallel smt-based model checker for parameterized systems - tool paper. In *Proc. of CAV*, pages 718–724, 2012.
16. A. Deutsch, R. Hull, F. Patrizi, and V. Vianu. Automatic verification of data-centric business processes. In *Proc. of ICDT*, pages 252–267, 2009.
17. A. Deutsch, Y. Li, and V. Vianu. Verification of hierarchical artifact systems. In *Proc. of PODS*, pages 179–194. ACM Press, 2016.
18. S. Ghilardi. Model theoretic methods in combined constraint satisfiability. *J. Autom. Reasoning*, 33(3-4):221–249, 2004.

¹² <http://www.bpmn.org/>

19. S. Ghilardi and A. Gianola. Interpolation, amalgamation and combination (the non-disjoint signatures case). In *Proc. FroCoS*, pages 316–332. Springer LNAI, 2017.
20. S. Ghilardi and A. Gianola. Modularity results for interpolation, amalgamation and superamalgamation. *Ann. Pure Appl. Logic*, 2018.
21. S. Ghilardi, E. Nicolini, and D. Zucchelli. A comprehensive framework for combined decision procedures. *ACM Trans. Comput. Log.*, pages 1–54, 2008.
22. S. Ghilardi and S. Ranise. MCMT: A model checker modulo theories. In *Proc. of IJCAR*, pages 22–29, 2010.
23. S. Ghilardi and S. J. van Gool. Monadic second order logic as the model companion of temporal logic. In *Proc. LICS*, pages 417–426, 2016.
24. S. Ghilardi and S. J. van Gool. A model-theoretic characterization of monadic second order logic on infinite words. *J. Symb. Log.*, 82(1):62–76, 2017.
25. S. Ghilardi and M. Zawadowski. *Sheaves, games, and model completions*, volume 14 of *Trends in Logic—Studia Logica Library*. Kluwer Academic Publishers, Dordrecht, 2002. A categorical approach to nonclassical propositional logics.
26. S. Gulwani and M. Musuvathi. Cover algorithms and their combination. In *Proc. of ESOP, Held as Part of ETAPS*, pages 193–207, 2008.
27. K. Hoder and N. Bjørner. Generalized property directed reachability. In *Proc. of SAT*, pages 157–171, 2012.
28. J. Hsiang and M. Rusinowitch. Proving refutational completeness of theorem-proving strategies: The transfinite semantic tree method. *J. ACM*, 38(3):559–587, 1991.
29. D. Kapur. Shostak’s congruence closure as completion. In *Rewriting Techniques and Applications, 8th International Conference, RTA-97, Sitges, Spain, June 2-5, 1997, Proceedings*, pages 23–37, 1997.
30. D. Kapur. Nonlinear polynomials, interpolants and invariant generation for system analysis. In *Proc. of the 2nd International Workshop on Satisfiability Checking and Symbolic Computation co-located with ISSAC*, 2017.
31. L. Kovács and A. Voronkov. Interpolation and symbol elimination. In *Proc. of CADE*, pages 199–213, 2009.
32. Y. Li, A. Deutsch, and V. Vianu. VERIFAS: A practical verifier for artifact systems. *PVLDB*, 11(3):283–296, 2017.
33. P. Lipparini. 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, 1982.
34. M. Ludwig and U. Waldmann. An extension of the Knuth-Bendix ordering with lpo-like properties. In *Proc. of LPAR*, pages 348–362, 2007.
35. K. L. McMillan. Lazy abstraction with interpolants. In *Proc. of CAV*, pages 123–136, 2006.
36. G. Nelson and D. C. Oppen. Fast decision procedures based on congruence closure. *J. ACM*, 27(2):356–364, 1980.
37. E. Nicolini, C. Ringeissen, and M. Rusinowitch. Data structures with arithmetic constraints: a non-disjoint combination. In *Proc. FroCoS*, pages 319–334. Springer LNCS, 2009.
38. E. Nicolini, C. Ringeissen, and M. Rusinowitch. Satisfiability procedures for combination of theories sharing integer offsets. In *Proc. TACAS*. Springer LNCS, 2009.
39. E. Nicolini, C. Ringeissen, and M. Rusinowitch. Combining satisfiability procedures for unions of theories with a shared counting operator. *Fund. Inform.*, pages 163–187, 2010.
40. R. Nieuwenhuis and A. Oliveras. Fast congruence closure and extensions. *Inf. Comput.*, 205(4):557–580, 2007.
41. R. Nieuwenhuis and A. Rubio. Theorem proving with ordering and equality constrained clauses. *J. Symb. Comput.*, 19(4):321–351, 1995.
42. R. Nieuwenhuis and A. Rubio. Paramodulation-based theorem proving. In *Handbook of Automated Reasoning (in 2 volumes)*, pages 371–443. MIT Press, 2001.

43. A. M. Pitts. On an interpretation of second order quantification in first order intuitionistic propositional logic. *J. Symb. Log.*, 57(1):33–52, 1992.
44. T. Rybina and A. Voronkov. A logical reconstruction of reachability. In *Perspectives of Systems Informatics, 5th International Andrei Ershov Memorial Conference, PSI 2003, Revised Papers*, pages 222–237, 2003.
45. V. Sofronie-Stokkermans. On interpolation and symbol elimination in theory extensions. In *Proc. IJCAR*, Lecture Notes in Computer Science. Springer, 2016.
46. V. Sofronie-Stokkermans. On interpolation and symbol elimination in theory extensions. *Log. Methods Comput. Sci.*, 14(3), 2018.
47. V. Vianu. Automatic verification of database-driven systems: a new frontier. In *Proc. of ICDT*, pages 1–13, 2009.
48. W. H. Wheeler. Model-companions and definability in existentially complete structures. *Israel J. Math.*, 25(3-4):305–330, 1976.