
UNIFORM INTERPOLANTS IN \mathcal{EUF} : ALGORITHMS USING DAG-REPRESENTATIONS

SILVIO GHILARDI, ALESSANDRO GIANOLA, AND DEEPAK KAPUR

Università degli Studi di Milano (Italy)
e-mail address: silvio.ghilardi@unimi.it

Free University of Bozen-Bolzano (Italy)
e-mail address: gianola@inf.unibz.it

University of New Mexico (USA)
e-mail address: kapur@cs.unm.edu

ABSTRACT. The concept of uniform interpolant for a quantifier-free formula from a given formula with a list of symbols, while well-known in the logic literature, has been unknown to the formal methods and automated reasoning community for a long time. This concept is precisely defined. Two algorithms for computing quantifier-free uniform interpolants in the theory of equality over uninterpreted symbols (\mathcal{EUF}) endowed with a list of symbols to be eliminated are proposed. The first algorithm is non-deterministic and generates a uniform interpolant expressed as a disjunction of conjunctions of literals, whereas the second algorithm gives a compact representation of a uniform interpolant as a conjunction of Horn clauses. Both algorithms exploit efficient dedicated DAG representations of terms. Correctness and completeness proofs are supplied, using arguments combining rewrite techniques with model theory.

1. INTRODUCTION

The theory of equality over uninterpreted symbols, henceforth denoted by \mathcal{EUF} , is one of the simplest theories that have found numerous applications in computer science, compiler optimization, formal methods and logic. Starting with the works of Shostak [31] and Nelson and Oppen [28] in the early eighties, some of the first algorithms were proposed in the context of developing approaches for combining decision procedures for quantifier-free theories including freely constructed data structures and linear arithmetic over the rationals. \mathcal{EUF} was exploited for hardware verification of pipelined processors by Dill [4] and more widely subsequently in formal methods and verification using model checking frameworks. With the popularity of SMT solvers, where \mathcal{EUF} serves as a glue for combining solvers for different theories, numerous new graph-based algorithms have been proposed in the literature over the last two decades for checking unsatisfiability of a conjunction of (dis)equalities of terms built using function symbols and constants.

Key words and phrases: Uniform Interpolation, \mathcal{EUF} , DAG representation, term rewriting.

In [27], the use of interpolants for automatic invariant generation was proposed, leading to a plethora of research activities to develop algorithms for generating interpolants for specific theories as well as their combination. This new application is different from the role of interpolants for analyzing proof theories of various logics starting with the pioneering work of [12, 21, 30] (for a recent survey in the SMT area, see [3, 2]). Approaches like [27, 21, 30], however, assume access to a proof of $\alpha \rightarrow \beta$ for which an interpolant is being generated. Given that there can in general be many interpolants including infinitely many for some theories, little is known about what kind of interpolants are effective for different applications, even though some research has been reported on the strength and quality of interpolants [13, 32, 20].

In this paper, a different approach is taken, motivated by the insight connecting interpolating theories with those admitting quantifier-elimination, as advocated in [25]. Particularly, in the preliminaries, the concept of a uniform interpolant (UI) defined by a formula α , in the context of formal methods and verification, is proposed for \mathcal{EUF} , which is well-known not to admit quantifier elimination. We recall here uniform interpolants in general; we fix a logic or a theory T and a suitable fragment (propositional, first-order quantifier-free, etc.) of its language L . Given an L -formula $\alpha(\underline{x}, \underline{y})$ (here $\underline{x}, \underline{y}$ are the variables occurring in α), a *uniform interpolant* of α (w.r.t. \underline{y}) is an L -formula $\alpha'(\underline{x})$ where only the \underline{x} occur, and that satisfies the following two properties: (i) $\alpha(\underline{x}, \underline{y}) \vdash_T \alpha'(\underline{x})$; (ii) for any further L -formula $\beta(\underline{x}, \underline{z})$ such that $\alpha(\underline{x}, \underline{y}) \vdash_T \beta(\underline{x}, \underline{z})$, we have $\alpha'(\underline{x}) \vdash_T \beta(\underline{x}, \underline{z})$. Whenever uniform interpolants exist, one can compute an interpolant for an entailment like $\alpha(\underline{x}, \underline{y}) \vdash_T \beta(\underline{x}, \underline{z})$ in a way that is *independent* of β . A (quantifier-free) uniform interpolant for a formula α is in particular, for any formula β , an *ordinary* interpolant [12, 26] for the pair (α, β) such that $\alpha \rightarrow \beta$ (as well as a *reverse* interpolant [27] for an unsatisfiable pair $(\alpha, \neg\beta)$).¹ A uniform interpolant could be defined for theories irrespective of whether they admit quantifier elimination. For theories admitting quantifier elimination, a uniform interpolant can be obtained using quantifier elimination: indeed, this shows that a theory enjoying quantifier elimination admits uniform interpolants as well. A uniform interpolant, when it exists, is unique up to logical equivalence: this immediately follows from the definition, since any uniform interpolant implies all the other formulae that are implied by α (and, then implies any other uniform interpolant). The equivalent concept of cover is proposed in [19] (see also [7, 10, 9]).

Two algorithms with different characteristics for generating uniform interpolants from a formula in \mathcal{EUF} (with a list of symbols to be eliminated) are proposed in this paper. They share a common subpart based on concepts used in a ground congruence closure proposed in [22], which flattens the input and generates a canonical rewrite system on constants along with unique rules of the form $f(\dots)$, where f is an uninterpreted symbol and the arguments (\dots) are canonical forms of constants. Further, eliminated symbols are represented as a DAG (‘Directed Acyclic Graph’) to avoid any exponential blow-up. DAG representation of terms are commonly used in theorem provers and are essential to keep basic algorithms like unification polynomial (see e.g. the detailed analysis in textbooks like [1]). In this paper, we also introduce a new DAG-representation for terms, called ‘conditional DAG’-representation, see below. Both DAG- and conditional DAG-representations require an exponential blow-up

¹The third author recently learned from the first author that this concept has been used extensively in logic for decades [18, 29] to his surprise since he had the erroneous impression that he came up with the concept in 2012, which he presented in a series of talks [23, 24].

in order to get full string representations of terms (precise instructions for the related unravellings will be recalled/supplied in the paper). However a careful choice of adequate data structures inside the desired applications can avoid the need of such an unravelling and this is the reason for relying on DAG-compressed representations while designing symbolic algorithms.

Our first algorithm is non-deterministic where undecided equalities on constants are hypothesized to be true or false, generating a branch in each case, and recursively applying the algorithm. It could also be formulated as an algorithm similar in spirit to the use of equality interpolants in the Nelson and Oppen framework for combination, where different partitions on constants are tried, with each leading to a branch in the algorithm. New symbols are introduced along each branch to avoid exponential blow-up. Our second algorithm generalizes the concept of a DAG to conditional DAG in which subterms are replaced by new symbols under a conjunction of equality atoms, resulting in its compact and efficient representation. A fully or partially expanded form of a UI can be derived based on their use in applications. Because of their compact representation, UIs can be kept of polynomial size for a large class of formulas.

The former algorithm is tableaux-based, in the sense that it has a tree structure computed for the logical formulae involved and a finite collection of rules that specifies how to create branches. This algorithm produces the output in disjunctive normal form. The second algorithm is based on manipulation of Horn clauses and gives the output in (compressed) conjunctive normal form. We believe that the two algorithms are complementary to each other, especially from the point of view of applications. Model checkers typically synthesize safety invariants using conjunctions of clauses and in this sense they might better take profit from the second algorithm; however, model-checkers dually representing sets of backward reachable states as disjunctions of *cubes* (i.e., conjunctions of literals) would better adopt the first algorithm. Non-deterministic manipulations of cubes are also required to match certain PSPACE lower bounds, as in the case of SAS systems mentioned in [8]. On the other hand, regarding the overall complexity, it seems to be easier to avoid exponential blow-ups in concrete examples by adopting the second algorithm.

The termination, correctness and completeness of both the algorithms are proved using results in model theory about model completions; this relies on a basic result (Lemma 5.1 below) taken from [7].

Both our algorithms are simple, intuitive and easy to understand in contrast to other algorithms in the literature. In fact, the algorithm from [7] requires the full saturation of all the formulae deductively implied in a version of superposition calculus equipped with ad hoc settings (in that context, no compact representation of the involved formulae is considered), whereas the main merit of our second algorithm is to show that a very light form of completion is sufficient, thus simplifying the whole procedure and getting seemingly better complexity results.² The algorithm from [19] presents some issues/bugs that need to be fixed (see [7, 9] for details) and the correctness proof has never been published (the technical report mentioned in [19] is not available).

The paper is structured as follows: in the next paragraph we discuss about related work on the use UIs. In Section 2 we state the main problem, settle on some notation, discuss DAG representations and congruence closure. In Sections 3 and 4, we respectively give the

²Although we feel that some improvement is possible, the termination argument in [7] gives a double exponential bound, whereas we have a simple exponential bound for both algorithms (with optimal chances to keep the output polynomial in many concrete cases in the second algorithm).

two algorithms for computing uniform interpolants in \mathcal{EUF} (correctness and completeness of such algorithms are proved in Section 5). We conclude in Section 6. This paper extends a conference paper ([14]) in two respects: first, it improves the presentation and includes the full proofs, adding also further explanations; second, it contains additional material including detailed examples and some complexity considerations.

Related work on the use of UIs. The use of uniform interpolants in model-checking safety problems for infinite state systems was already mentioned in [19] and further exploited in a recent research line on the verification of data-aware processes [6, 5, 8, 15, 16]. Model checkers need to explore the space of all reachable states of a system; a precise exploration (either forward starting from a description of the initial states or backward starting from a description of unsafe states) requires quantifier elimination. The latter is not always available or might have prohibitive complexity; in addition, it is usually preferable to make over-approximations of reachable states both to avoid divergence and to speed up convergence. One well-established technique for computing over-approximations consists in extracting interpolants from spurious traces, see e.g. [27]. One possible advantage of uniform interpolants over ordinary interpolants is that they do not introduce over-approximations and so abstraction/refinements cycles are not needed in case they are employed (the precise reason for that goes through the connection between uniform interpolants, model completeness and existentially closed structures, see [8] for a full account). In this sense, computing uniform interpolants has the same advantages and disadvantages as computing quantifier eliminations, with two remarkable differences. The first difference is that uniform interpolants may be available also in theories not admitting quantifier elimination (\mathcal{EUF} being the typical example); the second difference is that computing uniform interpolants may be tractable when the language is suitably restricted e.g. to unary function symbols (this was already mentioned in [19], see also Remark 3.4 below). Restriction to unary function symbols is sufficient in database driven verification to encode primary and foreign keys [8]. It is also worth noticing that, precisely by using uniform interpolants for this restricted language, in [8] *new decidability results* have been achieved for interesting classes of infinite state systems. Notably, such results are also operationally mirrored in the MCMT [17] implementation since version 2.8.

2. PRELIMINARIES

We adopt the usual first-order syntactic notions, including signature, term, atom, (ground) formula; our signatures are always *finite* or *countable* and include equality. Without loss of generality, only functional signatures, i.e. signatures whose only predicate symbol is equality, are considered. A tuple $\langle x_1, \dots, x_n \rangle$ of variables is compactly represented 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} . This tuple is assumed to be formed by *distinct variables*, thus we underline that, when we write e.g. $\phi(\underline{x}, \underline{y})$, we mean that the tuples $\underline{x}, \underline{y}$ are made of distinct variables that are also disjoint from each other. 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 quantifier-free. Formulae with no free variables are called *sentences*.

From the semantic side, the standard notion of Σ -structure \mathcal{M} is used: this is a pair formed of a set (the ‘support set’, indicated as $|\mathcal{M}|$) and of an interpretation function. The interpretation function maps n -ary function symbols to n -ary operations on $|\mathcal{M}|$ (in

particular, constants symbols are mapped to elements of $|\mathcal{M}|$). A free variables assignment \mathcal{I} on \mathcal{M} extends the interpretation function by mapping also variables to elements of $|\mathcal{M}|$; the notion of truth of a formula in a Σ -structure under a free variables assignment \mathcal{I} is the standard one.

It may be necessary to expand a signature Σ with a fresh name for every $a \in |\mathcal{M}|$: such expanded signature is called $\Sigma^{|\mathcal{M}|}$ and \mathcal{M} is by abuse seen as a $\Sigma^{|\mathcal{M}|}$ -structure itself by interpreting the name of $a \in |\mathcal{M}|$ as a (the name of a is directly indicated as a for simplicity).

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} .

A Σ -embedding [11] (or, simply, an embedding) between two Σ -structures \mathcal{M} and \mathcal{N} is a map $\mu : |\mathcal{M}| \rightarrow |\mathcal{N}|$ among the support sets $|\mathcal{M}|$ of \mathcal{M} and $|\mathcal{N}|$ of \mathcal{N} satisfying the condition $(\mathcal{M} \models \varphi \Rightarrow \mathcal{N} \models \varphi)$ for all $\Sigma^{|\mathcal{M}|}$ -literals φ (\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 $\mu(a)$). If $\mu : \mathcal{M} \rightarrow \mathcal{N}$ is an embedding that 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} .

Let \mathcal{M} be a Σ -structure. The *diagram* of \mathcal{M} , written $\Delta_\Sigma(\mathcal{M})$ (or just $\Delta(\mathcal{M})$), is the set of ground $\Sigma^{|\mathcal{M}|}$ -literals that are true in \mathcal{M} . An easy but important result, called *Robinson Diagram Lemma* [11], says that, given any Σ -structure \mathcal{N} , the embeddings $\mu : \mathcal{M} \rightarrow \mathcal{N}$ 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: the name of a is interpreted as $\mu(a)$. The typical use of the Robinson Diagram Lemma is the following: suppose we want to show that some structure \mathcal{M} can be embedded into a structure \mathcal{N} in such a way that some set of sentences Θ are true. Then, by the Lemma, this turns out to be equivalent to the fact that the set of sentences $\Delta(\mathcal{M}) \cup \Theta$ is consistent: thus, the Diagram Lemma can be used to transform an *embeddability* problem into a *consistency* problem (the latter is a problem of a logical nature, to be solved for instance by appealing to the compactness theorem for first-order logic).

2.1. Uniform Interpolants. Fix a theory T and an existential formula $\exists \underline{e} \phi(\underline{e}, \underline{z})$; call a *residue* of $\exists \underline{e} \phi(\underline{e}, \underline{z})$ any quantifier-free formula $\theta(\underline{z}, \underline{y})$ such that $T \models \exists \underline{e} \phi(\underline{e}, \underline{z}) \rightarrow \theta(\underline{z}, \underline{y})$ (equivalently, such that $T \models \phi(\underline{e}, \underline{z}) \rightarrow \theta(\underline{z}, \underline{y})$). The set of residues of $\exists \underline{e} \phi(\underline{e}, \underline{z})$ is denoted as $Res(\exists \underline{e} \phi(\underline{e}, \underline{z}))$. A quantifier-free formula $\psi(\underline{z})$ is said to be a *T -(quantifier-free) uniform interpolant*³ (or, simply, a *uniform interpolant*, abbreviated UI) of $\exists \underline{e} \phi(\underline{e}, \underline{z})$ iff $\psi(\underline{z}) \in Res(\exists \underline{e} \phi(\underline{e}, \underline{z}))$ and $\psi(\underline{z})$ implies (modulo T) all the formulae in $Res(\exists \underline{e} \phi(\underline{e}, \underline{z}))$. It is immediately seen that UIs are unique (modulo T -equivalence). A theory T has *uniform quantifier-free interpolation* iff every existential formula $\exists \underline{e} \phi(\underline{e}, \underline{z})$ has a UI.

Example 2.1. Consider the existential formula $\exists e (f(e, z_1) = z_2 \wedge f(e, z_3) = z_4)$: it can be shown that its \mathcal{EUF} -uniform interpolant is $z_1 = z_3 \rightarrow z_2 = z_4$.

Notably, if T has uniform quantifier-free interpolation, then it has ordinary quantifier-free interpolation, in the sense that if we have $T \models \phi(\underline{e}, \underline{z}) \rightarrow \phi'(\underline{z}, \underline{y})$ (for quantifier-free

³In some literature [19, 7] uniform interpolants are called *covers*.

formulae ϕ, ϕ' , then there is a quantifier-free formula $\theta(\underline{z})$ such that $T \models \phi(\underline{e}, \underline{z}) \rightarrow \theta(\underline{z})$ and $T \models \theta(\underline{z}) \rightarrow \phi'(\underline{z}, \underline{y})$. In fact, if T has uniform quantifier-free interpolation, then the interpolant θ is independent on ϕ' (the same $\theta(\underline{z})$ can be used as interpolant for all entailments $T \models \phi(\underline{e}, \underline{z}) \rightarrow \phi'(\underline{z}, \underline{y})$, varying ϕ'). Uniform quantifier-free interpolation has a direct connection to an important notion from classical model theory, namely model completeness (see [7] for more information).

2.2. Problem Statement. In this paper the problem of *computing UIs for the case in which T is pure identity theory in a functional signature Σ* is considered; this theory is called $\mathcal{EUF}(\Sigma)$ or just \mathcal{EUF} in the SMT-LIB2 terminology. Two different algorithms are proposed for that (while proving correctness and completeness of such algorithms, it is simultaneously shown that UIs exist in \mathcal{EUF}). The first algorithm computes a UI in disjunctive normal form format, whereas the second algorithm supplies a UI in conjunctive normal form format. Both algorithms use suitable DAG-compressed representation of formulae.

The following notation is used throughout the paper. Since it is easily seen that existential quantifiers commute with disjunctions, it is sufficient to compute UIs for *primitive* formulae, i.e. for formulae of the kind $\exists \underline{e} \phi(\underline{e}, \underline{z})$, where ϕ is a *constraint*, i.e. a conjunction of literals. We partition all the 0-ary symbols from the input as well as symbols newly introduced into disjoint sets. We use the following conventions:

- $\underline{e} = e_0, \dots, e_N$ (with N integer) are symbols to be eliminated, called *variables*,
- $\underline{z} = z_0, \dots, z_M$ (with M integer) are symbols *not* to be eliminated, called *parameters*,
- symbols a, b, \dots stand for both variables and parameters, and for (fresh) constants as well (usually introduced during skolemization).

In the following we will also use symbols \underline{y} for indicating variables that changed their status and do not need to be eliminated anymore: we use symbols a, b, \dots for them as well. Variables \underline{e} are usually skolemized during the manipulations of our algorithms and proofs below, in the sense that they have to be considered as fresh individual constants.

Remark 2.1. *UI computations eliminate symbols which are existentially quantified variables (or skolemized constants). Elimination of function symbols can be reduced to elimination of variables in the following way. Consider a formula $\exists f \phi(f, \underline{z})$, where ϕ is quantifier-free. Successively abstracting out functional terms, we get that $\exists f \phi(f, \underline{z})$ is equivalent to a formula of the kind $\exists \underline{e} \exists f (\bigwedge_i (f(\underline{t}_i) = e_i) \wedge \psi)$, where the \underline{e} are fresh variables (with $e_i \in \underline{e}$), \underline{t}_i are terms, f does not occur in $\underline{t}_i, e_i, \psi$ and ψ is quantifier-free. The latter is semantically equivalent to $\exists \underline{e} (\bigwedge_{i \neq j} (\underline{t}_i = \underline{t}_j \rightarrow e_i = e_j) \wedge \psi)$, where $\underline{t}_i = \underline{t}_j$ is the conjunction of the component-wise equalities of the tuples \underline{t}_i and \underline{t}_j .*

2.3. Flat Literals, DAGs and Congruence Closure. A *flat literal* is a literal of one of the following kinds

$$f(a_1, \dots, a_n) = b, \quad a_1 = a_2, \quad a_1 \neq a_2 \quad (2.1)$$

where a_1, \dots, a_n and b are (not necessarily distinct) variables or constants. A formula is flat iff all literals occurring in it are flat; flat terms are terms that may occur in a flat literal (i.e. terms like those appearing in (2.1)).

We call a *DAG-definition* (or simply a DAG) any formula $\delta(\underline{y}, \underline{z})$ of the following form (where $\underline{y} := y_1 \dots, y_n$):

$$\bigwedge_{i=1}^n (y_i = f_i(y_1, \dots, y_{i-1}, \underline{z})) .$$

Thus, $\delta(\underline{y}, \underline{z})$ provides an *explicit definition* of the \underline{y} in terms of the parameters \underline{z} .

Given a DAG δ , we can in fact associate to it the substitution σ_δ recursively defined by the mapping

$$(y_i)\sigma_\delta := f_i((y_1)\sigma_\delta, \dots, (y_{i-1})\sigma_\delta, \underline{z}).$$

DAGs are commonly used to represent formulae and substitutions in compressed form: in fact a formula like

$$\exists \underline{y} (\delta(\underline{y}, \underline{z}) \wedge \Phi(\underline{y}, \underline{z})) \tag{2.2}$$

is equivalent to $\Phi((\underline{y})\sigma_\delta, \underline{z})$, and is called *DAG-representation*. The formula $\Phi((\underline{y})\sigma_\delta, \underline{z})$ is said to be the *unravelling* of (2.2): notice that computing such an unravelling in uncompressed form by explicitly performing substitutions causes an exponential blow-up. This is why we shall systematically prefer DAG-representations (2.2) to their uncompressed forms.

As above stated, our main aim is to compute the UI of a primitive formula $\exists \underline{e} \phi(\underline{e}, \underline{z})$; using trivial logical manipulations (that have just linear complexity costs), it can be shown that, without loss of generality *the constraint $\phi(\underline{e}, \underline{z})$ can be assumed to be flat*. To do so, it is sufficient to perform a *preprocessing procedure* by applying well-known Congruence Closure Transformations: the reader is referred to [22] for a full account.

3. THE TABLEAUX ALGORITHM

The algorithm proposed in this section is tableaux-like. It manipulates formulae in the following *DAG-primitive* format

$$\exists \underline{y} (\delta(\underline{y}, \underline{z}) \wedge \Phi(\underline{y}, \underline{z}) \wedge \exists \underline{e} \Psi(\underline{e}, \underline{y}, \underline{z})) \tag{3.1}$$

where $\delta(\underline{y}, \underline{z})$ is a DAG and Φ, Ψ are flat constraints (notice that the \underline{e} do not occur in Φ). We call a formula of that format a *DAG-primitive formula*. To make reading easier, we shall omit in (3.1) the existential quantifiers, so as (3.1) will be written simply as

$$\delta(\underline{y}, \underline{z}) \wedge \Phi(\underline{y}, \underline{z}) \wedge \Psi(\underline{e}, \underline{y}, \underline{z}) . \tag{3.2}$$

We remark that Ψ can contain literals whose terms depend explicitly on \underline{e} , whereas Φ does *not* contain *any* occurrence of the \underline{e} variables. Initially the DAG δ and the constraint Φ are the empty conjunction. In the DAG-primitive formula (3.2), variables \underline{z} are called *parameter* variables, variables \underline{y} are called (*explicitly*) *defined* variables and variables \underline{e} are called (*truly*) *quantified* variables. Variables \underline{z} are never modified; in contrast, during the execution of the algorithm it could happen that some quantified variables may disappear or become defined variables (in the latter case they are renamed: a quantified variable e_i becoming defined is renamed as y_j , for a fresh y_j). Below, letters a, b, \dots range over $\underline{e} \cup \underline{y} \cup \underline{z}$.

Definition 3.1. *A term t (resp. a literal L) is \underline{e} -free when there is no occurrence of any of the variables \underline{e} in t (resp. in L). Two flat terms t, u of the kinds*

$$t := f(a_1, \dots, a_n) \quad u := f(b_1, \dots, b_n) \tag{3.3}$$

are said to be compatible iff for every $i = 1, \dots, n$, either a_i is identical to b_i or both a_i and b_i are \underline{e} -free. The difference set of two compatible terms as above is the set of disequalities $a_i \neq b_i$, where a_i is not equal to b_i .

3.1. The Tableaux Algorithm. Our algorithm applies the transformations below in a “don’t care” non-deterministic way. By saying this, we mean that the output is independent (up to logical equivalence) of the order of the application of the transformations, once the following priority is respected: the last transformation has lower priority with respect to the remaining transformations. The last transformation is also responsible for splitting the execution of the algorithm in several branches: each branch will produce a different disjunct in the output formula. Each state of the algorithm is a DAG-primitive formula like (3.2). We now provide the rules that constitute our ‘tableaux-like’ algorithm.

(1) : Simplification Rules:

(1.0): if an atom like $t = t$ belongs to Ψ , just remove it; if a literal like $t \neq t$ occurs somewhere, delete Ψ , replace Φ with \perp and stop;

(1.i): If t is not a variable and Ψ contains both $t = a$ and $t = b$, remove the former and replace it with $a = b$.

(1.ii): If Ψ contains $e_i = e_j$ with $i > j$, remove it and replace everywhere e_i by e_j .

(2) : DAG Update Rule: if Ψ contains $e_i = t(\underline{y}, \underline{z})$, remove it, rename everywhere e_i as y_j (for fresh y_j) and add $y_j = t(\underline{y}, \underline{z})$ to $\delta(\underline{y}, \underline{z})$. More formally:

$$\begin{aligned} & \delta(\underline{y}, \underline{z}) \wedge \Phi(\underline{y}, \underline{z}) \wedge (\Psi(\underline{e}, e_i, \underline{y}, \underline{z}) \wedge e_i = t(\underline{y}, \underline{z})) \\ & \quad \Downarrow \\ & (\delta(\underline{y}, \underline{z}) \wedge y_j = t(\underline{y}, \underline{z})) \wedge \Phi(\underline{y}, \underline{z}) \wedge \Psi(\underline{e}, y_j, \underline{y}, \underline{z}) \end{aligned}$$

(3) : e-Free Literal Rule: if Ψ contains a literal $L(\underline{y}, \underline{z})$, move it to $\Phi(\underline{y}, \underline{z})$. More formally:

$$\begin{aligned} & \delta(\underline{y}, \underline{z}) \wedge \Phi(\underline{y}, \underline{z}) \wedge (\Psi(\underline{e}, \underline{y}, \underline{z}) \wedge L(\underline{y}, \underline{z})) \\ & \quad \Downarrow \\ & \delta(\underline{y}, \underline{z}) \wedge (\Phi(\underline{y}, \underline{z}) \wedge L(\underline{y}, \underline{z})) \wedge \Psi(\underline{e}, \underline{y}, \underline{z}) \end{aligned}$$

(4) : Splitting Rule: If Ψ contains a pair of atoms $t = a$ and $u = b$, where t and u are compatible flat terms like in (3.3), and no disequality from the difference set of t, u belongs to Φ , then non-deterministically apply one of the following alternatives:

(4.0): remove from Ψ the atom $f(b_1, \dots, b_n) = b$, add to Ψ the atom $a = b$ and add to Φ all equalities $a_i = b_i$ such that $a_i \neq b_i$ is in the difference set of t, u ;

(4.1): add to Φ one of the disequalities from the difference set of t, u (notice that the difference set cannot be empty, otherwise Rule (1.i) applies).

When no more rule is applicable, delete $\Psi(\underline{e}, \underline{y}, \underline{z})$ from the resulting formula

$$\delta(\underline{y}, \underline{z}) \wedge \Phi(\underline{y}, \underline{z}) \wedge \Psi(\underline{e}, \underline{y}, \underline{z})$$

so as to obtain for any branch an output formula in DAG-representation of the kind

$$\exists \underline{y} (\delta(\underline{y}, \underline{z}) \wedge \Phi(\underline{y}, \underline{z})) .$$

We will see in Remark 3.4 that in the case of only unary functions, Rule (4) can be disregarded, and the algorithm becomes much simpler and computationally tractable.

Example 3.1. We give a simple example for the application of Splitting Rule. Let the pair of atoms $f(z_1, e) = z_2$ and $f(z_3, e) = z_4$ be in Ψ such that $z_1 \neq z_3$ is not in Φ . Since the difference set of $f(z_1, e)$ and $f(z_3, e)$ is $\{z_1 \neq z_3\}$ and $z_1 \neq z_3 \notin \Phi$, Splitting Rule applies. Applying the first alternative (Rule (4.0)), the atom $f(z_3, e) = z_4$ is removed from Ψ , and the atoms $z_1 = z_3$ and $z_2 = z_4$ are added to Ψ . Applying the second alternative (Rule (4.1)), the disequality $z_1 \neq z_3$ is added to Ψ .

Notice that in the above example the literals added to Ψ as a consequence of the Splitting Rule can then immediately be moved to Φ via Rule (3), as they are \underline{e} -free; this is always the case for the disequalities from difference sets in Rule (4.0) and Rule (4.1) (because they only involve \underline{e} -free terms, by definition of difference set), but not necessarily for the atom $a = b$ mentioned in Rule (4.0), because this atom might not be \underline{e} -free.

Remark 3.1. Splitting Rule (4) creates branches in an optimized way: indeed, the alternatives $a_i = b_i$ in (4.0) and $a_i \neq b_i$ in (4.1) are not generated for every pair of variables a, b , but the rule is applied only when the left members are compatible flat terms. Notice also that in case for every a_i in t and for every b_i in u the pairs a_i, b_i are all identical, Rule (1.i) applies instead.

The following remark will be useful to prove the correctness of our algorithm, since it gives a description of the kind of literals contained in a state triple that is *terminal* (i.e., when no rule applies).

Remark 3.2. Notice that if no transformation applies to (3.1), the set Ψ can only contain disequalities of the kind $e_i \neq a$, together with equalities of the kind $f(a_1, \dots, a_n) = a$. However, when it contains $f(a_1, \dots, a_n) = a$, one of the a_i must belong to \underline{e} (otherwise (2) or (3) applies). Moreover, if $f(a_1, \dots, a_n) = a$ and $f(b_1, \dots, b_n) = b$ are both in Ψ , then either they are not compatible or $a_i \neq b_i$ belongs to Φ for some i and for some variables a_i, b_i not in \underline{e} (otherwise (4) or (1.i) applies).

The following proposition states that, by applying the previous rules, termination is always guaranteed.

Proposition 3.1. *The non-deterministic procedure presented above always terminates.*

Proof. It is sufficient to show that every branch of the algorithm must terminate. In order to prove that, first observe that the total number of the variables involved never increases and it decreases if (1.ii) is applied (it might decrease also by the effect of (1.0)). Whenever such a number does not decrease, there is a bound on the number of disequalities that can occur in Ψ, Φ . Now transformation (4.1) decreases the number of disequalities that are actually *missing*; the other transformations do not increase this number. Finally, all transformations except (4.1) reduce the length of Ψ . \square

Remark 3.3. *The overall complexity of the above algorithm is exponential in time, because of the number of branches created by Splitting Rule (4). However, the number of rules applied in a single branch is quadratic in time in the dimension of the input: this fact can be proved by relying on the termination argument shown in Proposition 3.1. Indeed, every rule of the algorithm above except for Rule (4.1) reduces the length of Ψ , which has length $O(n)$ (where n is the dimension of the input). Let c_1 be a counter that decreases every time a rule (except for Rule (4.1)) is applied: hence, $c_1 := O(n)$. Moreover, whenever Rule (4.1) is applied, the length of Ψ remains the same, but the number of disequalities that are actually*

missing decreases: this number of missing disequalities is clearly $O(n^2)$. Let c_2 be a counter that decreases whenever Rule (4.1) is applied: hence, $c_2 := O(n^2)$. Now, consider a counter $c := c_1 + c_2$: this counter decreases every time a rule of the algorithm above is applied. Since $c = c_1 + c_2 = O(n) + O(n^2) = O(n + n^2) = O(n^2)$, we conclude that the number of rules applied in a single branch is quadratic, as wanted.

Remark 3.4. Notice that if function symbols are all unary, there is no need to apply Rule 4. Indeed, if all the function symbols are unary, compatible flat terms u and t can either have the form $u := f(e), t := f(e)$ (for some e in \underline{e}) or be \underline{e} -free. In the former case Rule 1 applies before, whereas in the latter case Rule 3 applies before (we recall that Rule 4 has the lowest priority). Hence, Rule 4 is never applied. Thanks to the argument shown in Remark 3.3, no branch is created and the number of rules applied is quadratic in time in the dimension of the input. Hence, for this restricted case computing UI is a tractable problem (this is consistent with [7, 9]). The case of unary functions has relevant applications in database driven verification [8, 6, 5, 15, 16] (where unary function symbols are used to encode primary and foreign keys).

Example 3.2. Let us compute the UI of the formula $\exists e_0 (g(z_4, e_0) = z_0 \wedge f(z_2, e_0) = g(z_3, e_0) \wedge h(f(z_1, e_0)) = z_0)$. Flattening gives the set of literals

$$g(z_4, e_0) = z_0 \wedge e_1 = f(z_2, e_0) \wedge e_1 = g(z_3, e_0) \wedge e_2 = f(z_1, e_0) \wedge h(e_2) = z_0 \quad (3.4)$$

where the newly introduced variables e_1, e_2 need to be eliminated too. Applying (4.0) removes $g(z_3, e_0) = e_1$ and introduces the new equalities $z_3 = z_4, e_1 = z_0$. This causes e_1 to be renamed as y_1 by (2). Applying again (4.0) removes $f(z_1, e_0) = e_2$ and adds the equalities $z_1 = z_2, e_2 = y_1$; moreover, e_2 is renamed as y_2 . To the literal $h(y_2) = z_0$ we can apply (3). The branch terminates with $y_1 = z_0 \wedge y_2 = y_1 \wedge z_1 = z_2 \wedge z_3 = z_4 \wedge h(y_2) = z_0 \wedge f(z_2, e_0) = y_1 \wedge g(z_4, e_0) = z_0$. This produces $z_1 = z_2 \wedge z_3 = z_4 \wedge h(z_0) = z_0$ as a first disjunct of the uniform interpolant.

The other branches produce $z_1 = z_2 \wedge z_3 \neq z_4, z_1 \neq z_2 \wedge z_3 = z_4$ and $z_1 \neq z_2 \wedge z_3 \neq z_4$ as further disjuncts, so that the UI turns out to be equivalent (by trivial logical manipulations) to $z_1 = z_2 \wedge z_3 = z_4 \rightarrow h(z_0) = z_0$.

Example 3.3. Consider the following example, taken from [19]. We compute the cover of the primitive formula $\exists e (s_1 = f(z_3, e) \wedge s_2 = f(z_4, e) \wedge t = f(f(z_1, e), f(z_2, e)))$, where s_1, s_2, t are terms in \underline{z} .

Flattening gives the set of literals

$$f(z_3, e) = s_1 \wedge f(z_4, e) = s_2 \wedge f(z_1, e) = e_1 \wedge f(z_2, e) = e_2 \wedge f(e_1, e_2) = t \quad (3.5)$$

where the newly introduced variables e_1, e_2 need to be eliminated too. We use lists of integers to represent the nodes of the tree created by the tableaux-like algorithm (Figure 1).

Applying (4.0) to the first and the second literals of (3.5), the first branch is created: node [1] is generated, where $f(z_4, e) = s_2$ is removed and the new equalities $z_3 = z_4, s_1 = s_2$ are introduced. Then, applying (4.0) to the first and the third literals of (3.5), we get a new branch: node [1.1] is generated, where $f(z_1, e) = e_1$ is removed and the new equalities $z_3 = z_1, s_1 = e_1$ are introduced. As shown in Figure 2, this causes e_1 to be renamed as y_5 by (2). Applying again (4.0) to the first and the fourth literals of (3.5), new branch is created: in node [1.1.1], $f(z_2, e) = e_2$ is removed and the equalities $z_3 = z_2, e_2 = s_1$ are added; moreover, e_2 is renamed as y_6 by using (2). In addition, in node [1.1.2], we obtain that no literal is canceled and the inequality $z_3 \neq z_2$ from the difference set of $f(z_3, e)$ and $f(z_2, e)$ is added. To all the newly introduced literals in $\underline{y}, \underline{z}$ in all the branches, we can apply (3).

The branch of [1.1.1] terminates with $f(z_3, e) = s_1 \wedge f(y_5, y_6) = t \wedge z_3 = z_4 \wedge s_1 = s_2 \wedge z_1 = z_3 \wedge s_1 = y_5 \wedge z_3 = z_2 \wedge s_1 = y_6$. This produces Leaf 1: $f(y_5, y_6) = t \wedge z_3 = z_4 \wedge s_1 = s_2 \wedge z_1 = z_3 \wedge s_1 = y_5 \wedge z_3 = z_2 \wedge s_1 = y_6$ as a first disjunct of the uniform interpolant.

The branch of [1.1.2] terminates with $f(z_3, e) = s_1 \wedge f(y_5, e_2) = t \wedge f(z_2, e) = e_2 \wedge z_3 = y_4 \wedge s_1 = s_2 \wedge z_1 = z_3 \wedge s_1 = y_5 \wedge z_3 \neq z_2$. This produces Leaf 2: $z_3 = z_4 \wedge s_1 = s_2 \wedge z_1 = z_3 \wedge s_1 = y_5 \wedge z_3 \neq z_2$ as a second disjunct of the uniform interpolant.

We also analyze a portion of the branch starting with node [2]: as a consequence of (4.1), in node [2] the new inequality $z_3 \neq z_4$ is introduced. Then, applying (4.0) to the first and the third literals of (3.5), we get a new branch: as shown in Figure 3, node [2.1] is generated, where $f(z_1, e) = e_1$ is removed and the new equalities $z_3 = z_1$, $s_1 = e_1$ are introduced. The last equality causes e_1 to be renamed as y_5 by (2). Applying again (4.0) to the second and the fourth literals of (3.5), a new branch is created: in node [2.1.1], $f(z_2, e) = e_2$ is removed and the equalities $z_4 = z_2$, $s_2 = e_2$ are added; moreover, e_2 is renamed as y_6 by using again (2). To all the newly introduced literals in $\underline{y}, \underline{z}$ in all the branches, we can apply (3).

The branch of [2.1.1] terminates with $f(z_3, e) = s_1 \wedge f(z_4, e) = s_2 \wedge f(y_5, y_6) = t \wedge z_3 \neq z_4 \wedge z_3 = z_1 \wedge s_1 = y_5 \wedge z_4 = z_2 \wedge s_2 = y_6$. This produces Leaf 6: $f(y_5, y_6) = t \wedge z_3 \neq z_4 \wedge z_3 = z_1 \wedge s_1 = y_5 \wedge z_4 = z_2 \wedge s_2 = y_6$ as a first disjunct of the uniform interpolant.

The algorithm generates 16 branches, each of them produces one disjunct of the output formula, so that the UI turns out to be equivalent to

$$z_3 = z_4 \rightarrow s_1 = s_2 \wedge \bigwedge_{i,j \in \{3,4\}} (z_1 = z_i \wedge z_2 = z_j \rightarrow t = f(s_{i-2}, s_{j-2})) .$$

Notice that this is consistent with:

- (1) the formula in Leaf 1, where $z_3 = z_4$ implies $s_1 = s_2$, and where $z_1 = z_3$ and $z_2 = z_4$ implies $t = f(s_1, s_2)$;
- (2) the formula in Leaf 2, where $z_3 = z_4$ implies $s_1 = s_2$;
- (3) the formula in Leaf 6, where $z_1 = z_3$ and $z_2 = z_4$ implies $t = f(s_1, s_2)$.

4. THE CONDITIONAL ALGORITHM

This section discusses a new algorithm with the objective of generating a compact representation of the UI in \mathcal{EUF} : this representation avoids splitting and is based on conditions in Horn clauses generated from literals whose left sides have the same function symbol. A by-product of this approach is that the size of the output UI often can be kept polynomial. Further, the output of this algorithm generates the UI of $\exists \underline{e} \phi(\underline{e}, \underline{z})$ (where $\phi(\underline{e}, \underline{z})$ is a conjunction of literals and $\underline{e} = e_0, \dots, e_N$, $\underline{z} = z_0, \dots, z_M$, as usual) in conjunctive normal form as a conjunction of Horn clauses (we recall that a Horn clause is a disjunction of literals containing *at most one* positive literal). Toward this goal, a new data structure of a conditional DAG, a generalization of a DAG, is introduced so as to maximize sharing of sub-formulas.

Using the core preprocessing procedure explained in Subsection 2.3, it is assumed that ϕ is the conjunction $\bigwedge S_1$, where S_1 is a set of *flat* literals containing only literals of the following two kinds:

$$f(a_1, \dots, a_h) = a \tag{4.1}$$

$$a \neq b \tag{4.2}$$

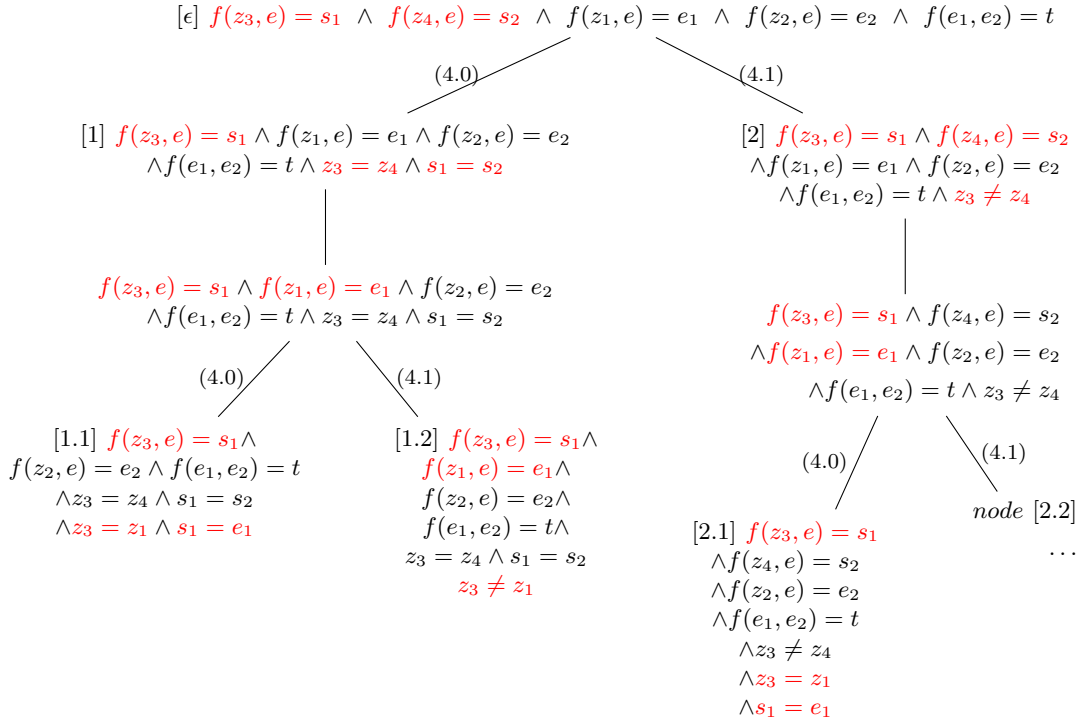


Figure 1: Portion of the tree created by the tableaux-like algorithm

(recall that we use letters a, b, \dots for elements of $\underline{e} \cup \underline{z}$). Since literals not involving variables to be eliminated or supplying an explicit definition of one of them can be moved directly to the output, we can assume that variables in \underline{e} must occur in (4.2) and in the left side of (4.1). We do not include equalities like $a = e$ because they can be eliminated by replacement.

4.1. The Conditional Algorithm. The algorithm requires two steps in order to get a set of clauses representing the output in a suitably compressed format.

Step 1. Out of every pair of literals $f(a_1, \dots, a_h) = a$ and $f(a'_1, \dots, a'_h) = a'$ of the kind (4.1) (where a is syntactically different from a') we produce the Horn clause

$$a_1 = a'_1, \dots, a_h = a'_h \rightarrow a = a' \quad (4.3)$$

which can be further simplified by deleting identities in the antecedent. Let us call S_2 the set of clauses obtained from S_1 by adding these new Horn clauses to it.

Step 2. We saturate S_2 with respect to the following rewriting rule

$$\frac{\Gamma \rightarrow e_j = e_i \quad C}{\Gamma \rightarrow C[e_i]_p}$$

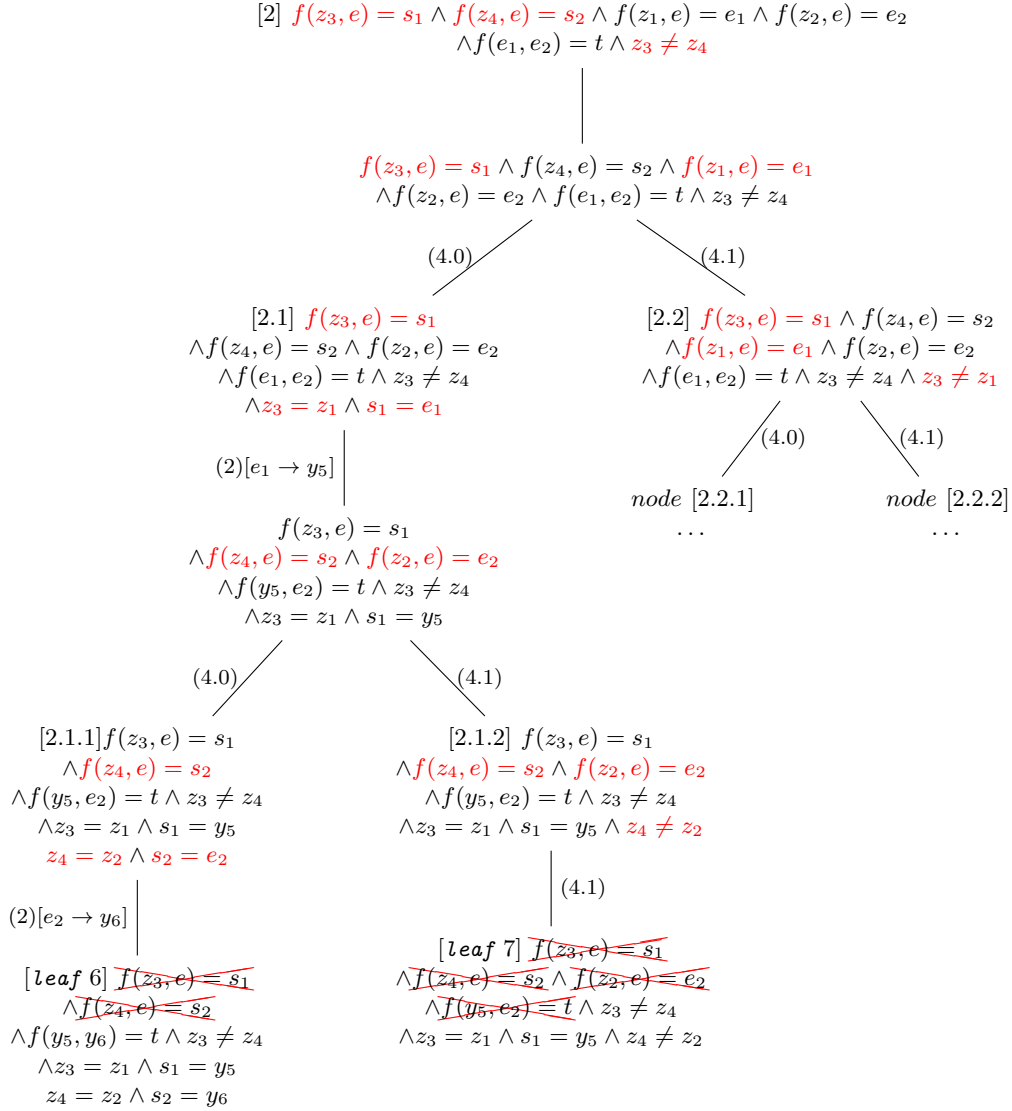


Figure 3: Portion of Branch 2

Step 1 produces the following set S_2 of Horn clauses

$$\begin{aligned} z_1 = z_2 \rightarrow e_1 = z_3, \quad z_4 = z_5 \rightarrow e_2 = z_6, \\ e_1 = z'_1 \rightarrow e_2 = z'_2, \quad e_2 = z''_1 \rightarrow e_1 = z''_2 \end{aligned}$$

Since there are no Horn clauses whose consequent is an equality of the kind $e_i = e_j$, Step 2 does not produce further clauses and we have $S_3 = S_2$.

4.2. Conditional DAGs. In order to be able to extract the output UI in an uncompressed format out of the above set of clauses S_3 , we must identify all the ‘implicit conditional definitions’ it contains. As for illustration, Example 4.1 contains, among the others, the following ‘implicit’ conditional definitions: the variable e_1 can be conditionally defined

with $e_1 := z_3$, when $z_1 = z_2$ holds (because $z_1 = z_2 \rightarrow e_1 = z_3$ is in S_3); moreover, when e_1 becomes conditionally defined, also e_2 becomes implicitly defined as $e_2 := z'_2$, under the condition that $e_1 = z'_1$ holds (because $e_1 = z'_1 \rightarrow e_2 = z'_2$ is in S_3 as well). All such conditional definitions need to be made explicit, and this is what we are going to formalize in the following.

Let \underline{w} be an ordered subset of the $\underline{e} = \{e_1, \dots, e_N\}$: that is, in order to specify \underline{w} we must take a subset of the \underline{e} and an ordering of this subset. Intuitively, these \underline{w} will play the role of *placeholders* inside a conditional definition.

If we let \underline{w} be w_1, \dots, w_s (where, say, w_i is some e_{k_i} with $k_i \in \{1, \dots, N\}$), we let L_i be the language restricted to \underline{z} and w_1, \dots, w_i (for $i \leq s$): in other words, an L_i -term or an L_i -clause may contain only terms built up from $\underline{z}, w_1, \dots, w_i$ by applying to them function symbols. In particular, L_s (also called $L_{\underline{w}}$) is the language restricted to $\underline{z} \cup \underline{w}$. We let L_0 be the language restricted to \underline{z} .

Given a set S of clauses and \underline{w} as above, a \underline{w} -conditional DAG δ (or simply a conditional DAG δ) built out of S is an s -tuple of Horn clauses from S

$$\Gamma_1 \rightarrow w_1 = t_1, \dots, \Gamma_s \rightarrow w_s = t_s \quad (4.4)$$

where Γ_i is a finite tuple of L_{i-1} -atoms and t_i is an L_{i-1} -term. Intuitively, a conditional DAG takes into consideration the given ordered subset \underline{w} of the \underline{e} symbols and on top of that it builds a sequence of conditional equations (i.e., Horn clauses), where step by step more and more \underline{e} symbols are employed and iteratively defined in terms of previously already defined \underline{e} symbols that are precedent in the order. Roughly speaking, each set of “dependencies” induced among the \underline{e} -symbols provides a suitable set of (conditional) definitions.

Conditional DAGs are used to define suitable formulae that are needed for the construction of the uniform interpolant. We now define such formulae. Given a \underline{w} -conditional DAG δ we can define the formulae ϕ_δ^i (for $i = 1, \dots, s+1$) as follows:

- ϕ_δ^{s+1} is the conjunction of all $L_{\underline{w}}$ -clauses belonging to S ;
- for $i \leq s$, the formula ϕ_δ^i is $\Gamma_i \rightarrow \forall w_i (w_i = t_i \rightarrow \phi_\delta^{i+1})$.

It can be seen that ϕ_δ^i is equivalent to a quantifier-free L_{i-1} formula,⁵ in particular ϕ_δ^1 (abbreviated as ϕ_δ) is equivalent to an L_0 -quantifier-free formula. The explicit computation of such quantifier-free formulae may however produce an exponential blow-up. The intuition behind these constructions is that in order to produce the correct uniform interpolant one needs to consider *all* the possible $L_{\underline{w}}$ -clauses in S and then iteratively ‘eliminate’ the \underline{e} symbols in them by exploiting the conditional definitions of the \underline{e} symbols (determined by the conditional DAG δ) in terms of previously defined \underline{e} symbols in the order of \underline{w} .

Example 4.2. *Let us analyze the conditional DAG δ that can be extracted out of the set S_3 of the Horn clauses mentioned in Example 4.1 (we disregard those δ such that ϕ_δ is the empty conjunction \top). The \underline{w}_1 -conditional DAG δ_1 with $\underline{w}_1 = e_1, e_2$ and conditional definitions*

$$z_1 = z_2 \rightarrow e_1 = z_3, \quad e_1 = z'_1 \rightarrow e_2 = z'_2$$

where e_2 depends upon e_1 , produces formula ϕ_{δ_1} , and similarly the \underline{w}_2 -conditional DAG δ_2 with $\underline{w}_2 = e_2, e_1$ and conditional definitions

$$z_4 = z_5 \rightarrow e_2 = z_6, \quad e_2 = z''_1 \rightarrow e_1 = z''_2.$$

⁵Since ϕ_δ^i is logically equivalent to $(\bigwedge \Gamma_i) \rightarrow \phi_\delta^{i+1}(t_i/w_i)$, it is immediate to see that it can be recursively turned, again up to equivalence, into a conjunction of Horn clauses.

where e_1 depends upon e_2 , produces formula ϕ_{δ_2} : notice that ϕ_{δ_1} and ϕ_{δ_2} are not logically equivalent.

Indeed, ϕ_{δ_1} is logically equivalent to

$$z_1 = z_2 \wedge z_3 = z'_1 \rightarrow \bigwedge S_3 \setminus \{e_0\}[z_3/e_1, z'_2/e_2] \quad (4.5)$$

where we used the notation $\bigwedge S_3 \setminus \{e_0\}[z_3/e_1, z'_2/e_2]$ to mean the result of the substitution of e_1 with z_3 and of e_2 with z'_2 in the conjunction of S_3 -clauses not involving e_0 . Notice that, intuitively, this formula is obtained by iteratively defining, step by step, bigger \underline{e} variables in terms of smaller ones: when $z_1 = z_2 \wedge z_3 = z'_1$ holds, e_1 is conditionally replaced by z_3 and e_2 by z'_2 .

Analogously, ϕ_{δ_2} is logically equivalent to

$$z_4 = z_5 \wedge z_6 = z''_1 \rightarrow \bigwedge S_3 \setminus \{e_0\}[z_6/e_2, z''_2/e_1] \quad (4.6)$$

(the explanation of the notation $\bigwedge S_3 \setminus \{e_0\}[z_6/e_2, z''_2/e_1]$ is the same as the explanation for the notation $\bigwedge S_3 \setminus \{e_0\}[z_3/e_1, z'_2/e_2]$ used above). A third possibility is to use the conditional definitions $z_1 = z_2 \rightarrow e_1 = z_3$ and $z_4 = z_5 \rightarrow e_2 = z_6$ with (equivalently) either \underline{w}_1 or \underline{w}_2 resulting in a conditional dag δ_3 with ϕ_{δ_3} logically equivalent to

$$z_1 = z_2 \wedge z_4 = z_5 \rightarrow \bigwedge S_3 \setminus \{e_0\}[z_3/e_1, z_6/e_2] . \quad (4.7)$$

The next lemma shows the relevant property of ϕ_δ :

Lemma 4.1. *For every set of clauses S and for every \underline{w} -conditional DAG δ built out of S , the formula $\bigwedge S \rightarrow \phi_\delta$ is logically valid.*

Proof. We prove that $\bigwedge S \rightarrow \phi_\delta^i$ is valid by induction on i . The base case is clear. For the case $i \leq s$, proceed, e.g., in natural deduction as follows: assume S, Γ_i and $\tilde{w}_i = t_i$ in order to prove $\phi_\delta^{i+1}(\tilde{w}_i/w_i)$. Since $\Gamma_i \rightarrow w_i = t_i \in S$, then by implication elimination you get $w_i = t_i$ and also $w_i = \tilde{w}_i$ by transitivity of equality. Now you get what you need from induction hypothesis and equality replacement. \square

Notice that it is not true that the conjunction of all possible ϕ_δ (varying δ and w) implies $\bigwedge S$: in fact, such a conjunction can be empty for instance in case S is just $\{e_1 = e_2\}$.

4.3. Extraction of UI's. We shall prove below that in order to get a UI of $\exists \underline{e} \phi(\underline{e}, \underline{a})$, one can take the conjunction of all possible ϕ_δ , varying δ among the conditional DAGs that can be built out of the set of clauses S_3 from Step 2 of the above algorithm. We highlight that in order to generate the correct output (i.e., the uniform interpolant), one needs to consider all the possible conditional DAGs built out of the set of clauses S_3 from Step 2, which implies to take into consideration all DAGs that can be defined considering all possible ordered subsets \underline{w} .

Example 4.3. *If ϕ is the conjunction of the literals of Example 4.1, then the conjunction of (4.5), (4.6) and (4.7) is a UI of $\exists \underline{e} \phi$; in fact, no further non-trivial conditional dag δ can be extracted (if we take $\underline{w} = e_1$ or $\underline{w} = e_2$ or $\underline{w} = \emptyset$ to extract δ , then it happens that ϕ_δ is the empty conjunction \top).*

Example 4.4. *Let us turn to the literals (3.4) of Example 3.2. Step 1 produces out of them the conditional clauses*

$$z_3 = z_4 \rightarrow e_1 = z_0, \quad z_1 = z_2 \rightarrow e_2 = e_1 . \quad (4.8)$$

Step 2 produces by rewriting the further clauses $z_1 = z_2 \rightarrow f(z_1, e_0) = e_1$ and $z_1 = z_2 \rightarrow h(e_1) = z_0$. We can extract two conditional DAGs δ (using both the conditional definitions (4.8) or just the first one); in both cases ϕ_δ is $z_1 = z_2 \wedge z_3 = z_4 \rightarrow h(z_0) = z_0$, which is the UI.

As it should be evident from the two examples above, the conditional DAGs representation of the output considerably reduces computational complexity in many cases; this is a clear advantage of the present algorithm over the algorithm from Section 3 and over other approaches like, e.g. [7]. Still, the next example shows that in some cases the overall complexity remains exponential.

Example 4.5. Let \underline{e} be e_0, \dots, e_N and let \underline{z} be $\{z_0, z'_0\} \cup \{z_{i,j}, z'_{i,j} \mid 1 \leq i < j \leq N\}$. Let $\phi(\underline{e}, \underline{z})$ be the conjunction of the identities $f(e_0, e_1) = z_0$, $f(e_0, e_N) = z'_0$ and the set of identities $h_{ij}(e_0, z_{ij}) = e_i$, $h_{ij}(e_0, z'_{ij}) = e_j$, varying i, j such that $1 \leq i < j \leq N$. We now show that applying the conditional algorithm we get an UI which is exponentially long.

After applying Step 1 of the algorithm presented in Subsection 4.1, we get the Horn clauses $z_{ij} = z'_{ij} \rightarrow e_i = e_j$, as well as the clause $e_1 = e_N \rightarrow z_0 = z'_0$. If we now apply Step 2, we can never produce a conditional clause of the kind $\Gamma \rightarrow e_i = t$ with t being \underline{e} -free (because we can only rewrite some e_i into some e_j). Thus no sequence of clauses like (4.4) can be extracted from S_3 : notice in fact that the term t_1 from such a sequence must not contain the variables \underline{e} . In other words, the only \underline{w} -conditional DAG δ that can be extracted is based on the empty $\underline{w} \subseteq \underline{e}$ and is empty itself.

In order to extract the UI, we need to compute the formulae ϕ_δ from any \underline{w} -conditional DAG δ , which is only one in such a case. However, this unique δ produces a formula ϕ_δ that is quite big: it is the conjunction of the clauses from S_3 where the \underline{e} do not occur (S_3 contains in fact $\Gamma \rightarrow z_0 = z'_0$ for exponentially many \underline{e} -free Γ 's).

We conclude this example by commenting on the reason why ϕ_δ has an exponential size. In fact, for every minimal set of pairs $I \subseteq \{1, \dots, N\} \times \{1, \dots, N\}$ such that the equivalence relation generated by I contains the pair $(1, N)$, we have that S_3 contains the clause $\Gamma_I \rightarrow z_0 = z'_0$, where Γ_I is the set of equalities $z_{ij} = z'_{ij}$ varying $(i, j) \in I$.⁶

5. CORRECTNESS AND COMPLETENESS PROOFS

In this section we prove correctness and completeness of our two algorithms. To this aim, we need some preliminaries, both from model theory and from term rewriting.

Extensions and UI are related to each other by the following result we take from [7]:

Lemma 5.1 (Cover-by-Extensions). *Let T be a first order theory. A formula $\psi(\underline{y})$ is a UI in T 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 $\mathcal{N} \models \exists \underline{e} \phi(\underline{e}, \underline{a})$.

For term rewriting we refer to a textbook like [1]; we only recall the following classical result:

⁶You can easily find exponentially many such I , e.g. by selecting a subset X of $\{1, \dots, N\}$ containing both 1 and N and letting I be the set $\{(i, j) \mid i < j, i \in X, j \in X\}$.

Lemma 5.2. *Let R be a canonical ground rewrite system over a signature Σ . Then there is a Σ -structure \mathcal{M} such that for every pair of ground terms t, u we have that $\mathcal{M} \models t = u$ iff the R -normal form of t is the same as the R -normal form of u . Consequently R is consistent with a set of negative literals S iff for every $t \neq u \in S$ the R -normal forms of t and u are different.*

We are now ready to prove correctness and completeness of our algorithms. We first give the relevant intuitions for the proof technique, which is the same for both cases. By Lemma 5.1 characterizing uniform interpolants, what we need to show is that if a model \mathcal{M} satisfies the output formula of the algorithm, then it can be extended to a superstructure \mathcal{N} satisfying the input formula of the algorithm. By Robinson Diagram Lemma, this embeddability problem can be transformed into a consistency problem: in order to do so, we show that the Robinson Diagram $\Delta_\Sigma(\mathcal{M})$ is consistent with the input formula of the algorithm. For our purposes, it is convenient to see $\Delta_\Sigma(\mathcal{M})$ as a set of flat literals as follows: the positive part of $\Delta_\Sigma(\mathcal{M})$ contains the $\Sigma^{|\mathcal{M}|}$ -equalities $f(a_1, \dots, a_n) = b$ which are true in \mathcal{M} and the negative part of $\Delta_\Sigma(\mathcal{M})$ contains the $\Sigma^{|\mathcal{M}|}$ -inequalities $a \neq b$, varying a, b among the pairs of different elements of $|\mathcal{M}|$. The positive part of $\Delta_\Sigma(\mathcal{M})$ is a canonical rewriting system (equalities like $f(a_1, \dots, a_n) = a$ are obviously oriented from left-to-right) and every term occurring in $\Delta(\mathcal{M})$ is in normal form. If an algorithm works properly, it will be possible to see that the completion of the union of $\Delta_\Sigma(\mathcal{M})$ with the input constraint (or with a constraint equivalent to it) is trivial and does not produce inconsistencies. To sum up, the completeness proofs of both algorithms require the following technical ingredients:

- (1) Lemma 5.1, for transforming the problem of computing UI into an embeddability problem;
- (2) Robinson Diagram Lemma, for turning the previous problem into a consistency problem, which is more tractable;
- (3) the completion of the diagram joined with the input constraint, so as to get a canonical rewriting system.

Correctness and Completeness of the Tableaux Algorithm

In this subsection, we prove the correctness and completeness of the Tableaux Algorithm. We first summarize the structure of the proof by commenting on its main steps. As discussed above, the proof of Theorem 5.1 relies on Lemma 5.1: in order to prove that the output formula is a uniform interpolant, it is sufficient to show that the embeddability conditions stated in Lemma 5.1 hold. This is achieved, thanks to Robinson Diagram Lemma, by showing that the Robinson Diagram $\Delta_\Sigma(\mathcal{M})$, where \mathcal{M} satisfies the output formula, is consistent with the input formula, as manipulated up to logical equivalencies by the algorithm. In the case of the Tableaux Algorithm, after a normalization of a rewriting system suitably extending this Diagram, we show that all the obtained oriented equalities form a canonical rewriting system: this is an immediate consequence of Remark 3.2. We then conclude applying Lemma 5.2: this lemma allows us to exhibit a model of the canonical rewriting system, showing in turn the consistency of $\Delta_\Sigma(\mathcal{M})$ with the input formula.

Theorem 5.1. *Suppose that we apply the algorithm of Subsection 3.1 to the primitive formula $\exists \underline{e}(\phi(\underline{e}, \underline{z}))$ and that the algorithm terminates with its branches in the states*

$$\delta_1(\underline{y}_1, \underline{z}) \wedge \Phi_1(\underline{y}_1, \underline{z}) \wedge \Psi_1(\underline{e}_1, \underline{y}_1, \underline{z}), \dots, \delta_k(\underline{y}_k, \underline{z}) \wedge \Phi_k(\underline{y}_k, \underline{z}) \wedge \Psi_k(\underline{e}_k, \underline{y}_k, \underline{z})$$

then the UI of $\exists \underline{e}(\phi(\underline{e}, \underline{z}))$ in \mathcal{EUF} is the unravelling (see Subsection 2.3) of the formula

$$\bigvee_{i=1}^k \exists \underline{y}_i (\delta_i(\underline{y}_i, \underline{z}) \wedge \Phi_i(\underline{y}_i, \underline{z})) . \quad (5.1)$$

Proof. Since $\exists \underline{e}(\phi(\underline{e}, \underline{z}))$ is logically equivalent to $\bigvee_{i=1}^k \exists \underline{y}_i (\delta_i(\underline{y}_i, \underline{z}) \wedge \Phi_i(\underline{y}_i, \underline{z}) \wedge \exists \underline{e}_i \Psi_i(\underline{e}_i, \underline{y}_i, \underline{z}))$, it is sufficient to check that if a formula like (3.1) is terminal (i.e. no rule applies to it) then its UI is $\exists \underline{y} (\delta(\underline{y}, \underline{z}) \wedge \Phi(\underline{y}, \underline{z}))$. To this aim, we apply Lemma 5.1: we pick a model \mathcal{M} satisfying the formula $\delta(\underline{y}, \underline{z}) \wedge \Phi(\underline{y}, \underline{z})$ from the output via an assignment \mathcal{I} to the variables $\underline{y}, \underline{z}$ ⁷ and we show that \mathcal{M} can be embedded into a model \mathcal{M}' such that, for a suitable extensions \mathcal{I}' of \mathcal{I} to the variables \underline{e} , we have that $(\mathcal{M}', \mathcal{I}')$ satisfies also $\Psi(\underline{e}, \underline{y}, \underline{z})$ from the input. This embeddability problem can be transformed into a consistency problem as follows. In fact, what we need (by Robinson Diagram Lemma) is to find a model for the following set of literals

$$\Delta_{\Sigma}(\mathcal{M}) \cup \Psi \cup \{a = \tilde{a}\}_{a \in \underline{y} \cup \underline{z}} \quad (5.2)$$

where \tilde{a} is the value of a under the assignment \mathcal{I} (here all variables in (5.2) are seen as constants, so (5.2) is a set of ground literals). We can orient the equalities in (5.2) by letting function symbols having bigger precedence over constants and by letting a having bigger precedence over \tilde{a} . Normalizing (5.2) replaces a with \tilde{a} in Ψ : call $\tilde{\Psi}$ the resulting set of literals (we conventionally use \tilde{e}_i as an alias for e_i , for all $e_i \in \underline{e}$, to have a uniform notation). After this normalization, we show that all oriented equalities in $\Delta_{\Sigma}(\mathcal{M}) \cup \tilde{\Psi} \cup \{a = \tilde{a}\}_{a \in \underline{y} \cup \underline{z}}$ form a canonical rewriting system. This is due to Remark 3.2: in fact if $f(a_1, \dots, a_n)$ and $f(b_1, \dots, b_n)$ both occur in $\tilde{\Psi}$, it cannot happen that $f(\tilde{a}_1, \dots, \tilde{a}_n)$ and $f(\tilde{b}_1, \dots, \tilde{b}_n)$ are the same term because the \underline{e} are already in normal form and because $\mathcal{M} \models \Phi$ (and hence also $\mathcal{M} \models \tilde{\Phi}$): in particular, all disequalities between \underline{e} -free constants $a_i \neq b_i$ belonging to Φ are true in \mathcal{M} . In addition, if $f(a_1, \dots, a_n)$ occurs in $\tilde{\Psi}$, then one of the a_i belongs to \underline{e} , hence rules from $\tilde{\Psi}$ and $\Delta_{\Sigma}(\mathcal{M})$ cannot superpose. Since all oriented equalities in $\Delta_{\Sigma}(\mathcal{M}) \cup \tilde{\Psi} \cup \{a = \tilde{a}\}_{a \in \underline{y} \cup \underline{z}}$ form a canonical rewriting system, the inequalities in $\Delta_{\Sigma}(\mathcal{M}) \cup \tilde{\Psi}$ are in normal form and we can apply Lemma 5.2 to get the desired \mathcal{M}' : in fact, Lemma 5.2 provides a Σ -structure \mathcal{M}' that is a model of the canonical rewriting system $\Delta_{\Sigma}(\mathcal{M}) \cup \tilde{\Psi} \cup \{a = \tilde{a}\}_{a \in \underline{y} \cup \underline{z}}$, showing in turn the consistency of $\Delta_{\Sigma}(\mathcal{M})$ with the formula Ψ in input. \square

Correctness and Completeness of the Conditional Algorithm

In this subsection we provide the full proof of correctness and completeness of the Conditional Algorithm. First of all, we briefly present the main ideas behind this proof. As in the case of the Tableaux algorithm, exploiting Lemma 5.1, we need to show that the embeddability conditions stated in Lemma 5.1 hold. We do so by using Robinson Diagram Lemma: we prove that the Robinson Diagram $\Delta(\mathcal{M})$, where \mathcal{M} satisfies the output formula, is consistent with the input formula. However, in the case of the Conditional Algorithm, the proof of this fact is more involved: indeed, we use the ground Knuth-Bendix completion in order to prove that no inconsistent literal can be produced, and this requires a careful analysis of the equalities that can be generated during the completion. Specifically, a particular attention is needed for equalities involving only symbols from a certain subset

⁷Actually the values of the assignment \mathcal{I} to the \underline{z} uniquely determines the values of \mathcal{I} to the \underline{y} .

$\underline{e} \setminus \underline{w}$ (called \underline{u} in the proof of the theorem): this analysis is carried out in Lemma 5.4 below. The fact that no inconsistency can be produced concludes the proof of the theorem.

In order to prove Lemma 5.4 (used in the proof of Theorem 5.2), we need to show the following preliminary lemma:

Lemma 5.3. *If the clauses $\Gamma \rightarrow f(a_1, \dots, a_h) = b$ and $\Gamma' \rightarrow f(a'_1, \dots, a'_h) = b'$ both belong to the set of clauses S_3 obtained after Step 2 in Subsection 4.1 and b is not the same term as b' , then S_3 contains also a clause subsuming the clause*

$$\Gamma, \Gamma', a_1 = a'_1, \dots, a_h = a'_h \rightarrow b = b'$$

Proof. By induction on the number K of applications of the rewriting rule of Step 2 needed to derive $\Gamma \rightarrow f(a_1, \dots, a_h) = b$ and $\Gamma' \rightarrow f(a'_1, \dots, a'_h) = b'$. If K is 0, the claim is clear by the instruction of Step 1. Suppose that $K > 0$ and let $\Gamma' \rightarrow f(a'_1, \dots, a'_h) = b'$ be obtained from $\Gamma_1 \rightarrow e_i = e_j$ by rewriting e_j to e_i from some clause C . We need to distinguish cases depending on the position p of the rewriting. All cases being treated in the same way, suppose for instance that p is in the antecedent,⁸ so that C is $\Gamma_2 \rightarrow f(a'_1, \dots, a'_h) = b'$ and that $\Gamma' \rightarrow f(a'_1, \dots, a'_h) = b'$ is $\Gamma_1, \Gamma_2[e_i]_p \rightarrow f(a'_1, \dots, a'_h) = b'$. Then by induction hypothesis S_3 contains a clause subsuming

$$\Gamma, \Gamma_2, a_1 = a'_1, \dots, a_h = a'_h \rightarrow b = b'$$

and rewriting with $\Gamma_1 \rightarrow e_i = e_j$ produces

$$\Gamma, \Gamma_1, \Gamma_2[e_i]_p, a_1 = a'_1, \dots, a_h = a'_h \rightarrow b = b'$$

as required. □

We now state and prove the theorem of correctness and completeness of the Conditional Algorithm.

Theorem 5.2. *Let S_3 be obtained from $\exists \underline{e} \phi(\underline{e}, \underline{z})$ as in Steps 1-2 of Subsection 4.1. Then the conjunction \mathcal{C} of all possible ϕ_δ (varying δ among the conditional DAGs that can be built out of S_3) is a UI of $\exists \underline{e} \phi(\underline{e}, \underline{z})$ in \mathcal{EUF} .*

Proof. We use Lemma 5.1 in order to show that the output \mathcal{C} is the UI of the input formula $\exists \underline{e} \phi(\underline{e}, \underline{z})$. Condition (i) of that Lemma is ensured by Lemma 4.1 above because $\bigwedge S_3$ is logically equivalent to ϕ . So let us take a model \mathcal{M} and elements $\tilde{\underline{a}}$ from its support such that we have $\mathcal{M} \models \bigwedge_\delta \phi_\delta$ under the assignment of the $\tilde{\underline{a}}$ to the parameters \underline{z} . We need to expand it to a superstructure \mathcal{N} in such a way that we have $\mathcal{N} \models \bigwedge S_1$, under some assignment to $\underline{z}, \underline{e}$ extending the assignment $\underline{z} \mapsto \tilde{\underline{a}}$ (recall that $\bigwedge S_1$ is logically equivalent to ϕ too). From now on, we consider the assignment $\underline{z} \mapsto \tilde{\underline{a}}$ fixed, so that when we write $\mathcal{M} \models C$ for a clause $C(\underline{z})$ we mean that $\mathcal{M} \models C$ holds under the assignment $\underline{z} \mapsto \tilde{\underline{a}}$.

Now, we can transform the embeddability problem of finding the aforementioned superstructure \mathcal{N} into a consistency problem as follows. First of all, notice that every \underline{w} -conditional DAG δ extracted from S_3 (let it be given by the clauses (4.4)) is naturally

⁸There is a case, where p is in the consequent, that is treated in a slightly different way than all the other ones. If, using $\Gamma_1 \rightarrow e_i = e_j$, we rewrite a clause of the form $\Gamma'' \rightarrow f(a'_1, \dots, a'_n) = e_j$ into $\Gamma_1, \Gamma'' \rightarrow f(a'_1, \dots, a'_n) = e_i$ (so that b' is the same as e_i) and if b is e_j , then, instead of applying induction, we can directly take $\Gamma_1 \rightarrow e_i = e_j$ as the clause we are looking for. If b is not e_j , induction applies as in all the other cases.

equipped with a substitution σ_δ which is given in DAG form by $w_1 \mapsto t_1, \dots, w_s \mapsto t_s$. We say that δ is *realized* in \mathcal{M} iff we have that

$$\mathcal{M} \models \bigwedge \Gamma_1 \sigma_\delta, \dots, \mathcal{M} \models \bigwedge \Gamma_s \sigma_\delta$$

Let δ be a \underline{w} -conditional DAG which is realized in \mathcal{M} and let it be maximal with this property (a \underline{w} -conditional DAG δ is said to be bigger than a \underline{w}' -conditional DAG δ' iff \underline{w} includes \underline{w}' - the inclusion is as sets, the order is disregarded). Since $\mathcal{M} \models \phi_\delta$ and δ is realized in \mathcal{M} , it is clear that all $L_{\underline{w}}$ -clauses from S_3 (hence also all $L_{\underline{w}}$ -literals from S_1) are true in \mathcal{M} . Let $\underline{u} = u_1, \dots, u_k$ be the variables from $\underline{e} \setminus \underline{w}$ and let $S_{\underline{u}}$ be the literals from S_1 which are not $L_{\underline{w}}$ -literals. What we need (by Robinson Diagram Lemma) is to find a model for the following set of literals

$$\Delta_\Sigma(\mathcal{M}) \cup S_{\underline{u}} \cup \{w_i = \tilde{b}_i \mid i = 1, \dots, s\} \cup \{z_i = \tilde{a}_i \mid i = 1, \dots, M\} \quad (5.3)$$

where \tilde{b}_i is the value of $w_i \sigma_\delta$ under the assignment $\underline{z} \mapsto \tilde{\underline{a}}$. This is the consistency problem we solve in the remaining part of the proof: in order to do so, we obtain by completion a suitable canonical rewriting system that does not introduce inconsistencies.

We orient the functional equalities in (5.3) from left to right and the equalities $w_i = \tilde{b}_i$ also from left to right. The \underline{e} are ordered as $e_1 > \dots > e_N$ and are bigger than the constants naming the elements of $|\mathcal{M}|$; function symbols are bigger than constant symbols. We show that the ground Knuth-Bendix completion of (5.3) cannot produce any inconsistent literal of the kind $t \neq t$ (this completes the proof of the Theorem).

First notice that the rules $\{w_i = \tilde{b}_i \mid i = 1, \dots, s\} \cup \{z_i = \tilde{a}_i \mid i = 1, \dots, M\}$ simply eliminates the \underline{w} and the \underline{z} from $S_{\underline{u}}$ (they become inactive after such normalization steps). Let $\tilde{S}_{\underline{u}}$ be the set of equalities resulting after this elimination. It turns out that $\tilde{S}_{\underline{u}}$ can only contain equalities of the kinds

$$f(a_1, \dots, a_h) = a \quad (5.4)$$

$$u_j \neq a \quad (5.5)$$

where a_1, \dots, a_h, a can be either among the \underline{u} or constants naming elements of $|\mathcal{M}|$. However some of the \underline{u} must be among a_1, \dots, a_h for each equality of the kind (5.4) because atoms not containing the \underline{u} are removed by $\Delta_\Sigma(\mathcal{M})$ and atoms like $u_i = t$ (where t does not contain any of the \underline{u}) cannot be there because δ is maximal. During completion, in addition to these kinds of atoms, only atoms of the kind

$$u_i = u_j \quad (5.6)$$

can possibly be produced. This is a consequence of the next Lemma. Below we say that a tuple of atoms Γ is *realized in \mathcal{M}* iff the Γ are $L_{\underline{w}}$ -atoms and $\mathcal{M} \models \bigwedge \Gamma \sigma_\delta$; similarly we say that a literal Θ is *conditionally realized in \mathcal{M}* if there exists Γ realized in \mathcal{M} with $\Gamma \rightarrow \Theta \in S_3$ (if Θ is a negative literal, $\Gamma \rightarrow \Theta$ stands for $\Gamma, \neg\Theta \rightarrow$).

Lemma 5.4. *Suppose that a literal Λ is produced during the completion of $\tilde{S}_{\underline{u}}$. Then it must be of the kinds (5.4), (5.5), (5.6). Moreover there exists a literal Λ' such that (i) Λ' is conditionally realized in \mathcal{M} ; (ii) Λ is obtained from Λ' by rewriting $\underline{z}, \underline{w}$ respectively to $\tilde{\underline{a}}, \tilde{\underline{b}}$.*

Proof. By straightforward case analysis; we analyze the most interesting case given by the superposition of two rules of the kind (5.4). Suppose that

$$f(a_1, \dots, a_h) = a \text{ and } f(a_1, \dots, a_h) = b$$

produce the equality $a = b$. Then by induction hypothesis, there are in S_3 two clauses like

$$\Gamma' \rightarrow f(a'_1, \dots, a'_h) = a', \quad \Gamma'' \rightarrow f(a''_1, \dots, a''_h) = a''$$

with Γ', Γ'' realized in \mathcal{M} , with a'_1, \dots, a'_h, a' rewritable (using $\underline{z}, \underline{w} \mapsto \tilde{a}, \tilde{b}$) to a_1, \dots, a_h, a , respectively, and with a''_1, \dots, a''_h, a'' also rewritable (using $\underline{z}, \underline{w} \mapsto \tilde{a}, \tilde{b}$) to a_1, \dots, a_h, b , respectively. By Lemma 5.3, S_3 contains a clause subsuming

$$\Gamma', \Gamma'', a'_1 = a''_1, \dots, a'_h = a''_h \rightarrow a' = a'' \quad (5.7)$$

which is as required because $\Gamma', \Gamma'', a'_1 = a''_1, \dots, a'_h = a''_h$ is realized in \mathcal{M} and $a' = a''$ rewrites (using $\underline{z}, \underline{w} \mapsto \tilde{a}, \tilde{b}$) to $a = b$. It remains to check that $a = b$ is of the kind (5.6). If both a', a'' taken from the consequent of (5.7) belong to $\underline{z} \cup \underline{w}$, then since the antecedent of (5.7) is realized in \mathcal{M} and (5.7) belongs to S_3 , a and b must be the same element from $|\mathcal{M}|$, so that $a = b$ is a trivial identity (which does not enter into the completion). It cannot be that only one between a' and a'' belongs to $\underline{z} \cup \underline{w}$ (the other one being from \underline{u}) because δ is maximal among conditional DAGs realized by \mathcal{M} and thus it cannot be properly enlarged by adding to it the additional conditional definition which would be supplied by (5.7). Thus it must be the case that both a', a'' are from \underline{u} , which implies that they cannot be rewritten (using $\underline{z}, \underline{w} \mapsto \tilde{a}, \tilde{b}$), so that a' is a , a'' is b and $a = b$ is of the kind (5.6). \square

Proof of Theorem 5.2 (continued). Once $\tilde{S}_{\underline{u}}$ (standing alone) is completed, only literals of the kinds (5.4), (5.5), (5.6) are produced. No completion inference is possible between literals of the kinds (5.4), (5.5), (5.6) on one side and literals from $\Delta_{\Sigma}(\mathcal{M}) \cup \{w_i = \tilde{b}_i \mid i = 1, \dots, s\} \cup \{z_i = \tilde{a}_i \mid i = 1, \dots, M\}$ on the other side; hence the completion of $\tilde{S}_{\underline{u}}$ alone, once joined to $\Delta_{\Sigma}(\mathcal{M}) \cup \{w_i = \tilde{b}_i \mid i = 1, \dots, s\} \cup \{z_i = \tilde{a}_i \mid i = 1, \dots, M\}$ yields a completion of (5.3). The only possible inconsistencies that can arise are given by literals of the kind $u_i \neq u_i$. Suppose that indeed one such a literal $u_i \neq u_i$ is produced during the completion of $\tilde{S}_{\underline{u}}$. Applying the above lemma, there should be in S_3 a clause like $\Gamma, u_i = u_i \rightarrow$ (i.e. after simplification, a clause like $\Gamma \rightarrow$) with Γ being realized in \mathcal{M} . The last means that $\mathcal{M} \models \bigwedge \Gamma \sigma_{\delta}$. This cannot be, because $\Gamma \rightarrow$ is a $L_{\underline{w}}$ -clause from S_3 : in fact, we have that $\mathcal{M} \models \phi_{\delta}$ and that δ is realized in \mathcal{M} , which imply that $\mathcal{M} \models C \sigma_{\delta}$ holds for every $L_{\underline{w}}$ -clause C from S_3 by the definition of ϕ_{δ} . In particular, we should have $\mathcal{M} \models \neg \bigwedge \Gamma \sigma_{\delta}$, taking $\Gamma \rightarrow$ as C . \square

6. CONCLUSIONS

Two different algorithms for computing uniform interpolants (UIs) from a formula in \mathcal{EUF} with a list of symbols to be eliminated are presented. They share a common subpart as well as they are different in their overall objectives. The first algorithm generates a UI expressed as a disjunction of conjunctions of literals, whereas the second algorithm gives a compact representation of a UI as a conjunction of Horn clauses. The output of both algorithms needs to be expanded if a fully (or partially) unravelled uniform interpolant is needed for an application. This restriction/feature is similar in spirit to syntactic unification where also efficient unification algorithms never produce output in fully expanded form to avoid an exponential blow-up.

For generating a compact representation of the UI, both algorithms make use of DAG representations of terms by introducing new symbols to stand for subterms arising in the

full expansion of the UI. Moreover, the second algorithm uses a conditional DAG, a new data structure introduced in the paper, to represent subterms under conditions.

The complexity of the algorithms is also analyzed. It is shown that the first algorithm generates exponentially many branches with each branch of at most quadratic length; the UIs produced by the second algorithm have polynomial size in all the hand-made examples we tried (but the worst case size is still exponential as witnessed by *ad hoc* examples like Example 4.5). A fully expanded UI can easily be of exponential size. An implementation of both the algorithms, along with a comparative study are planned as future work. In parallel with the implementation, a characterization of classes of formulae for which computation of UIs requires polynomial time in our algorithms (especially in the second one) needs further investigation.

Acknowledgments. The third author has been partially supported by the National Science Foundation award CCF -1908804.

REFERENCES

- [1] F. Baader and T. Nipkow. *Term Rewriting and All That*. Cambridge University Press, United Kingdom, 1998.
- [2] M. P. Bonacina and M. Johansson. Interpolation systems for ground proofs in automated deduction: a survey. *J. Autom. Reasoning*, 54(4):353–390, 2015.
- [3] M. P. Bonacina and M. Johansson. On interpolation in automated theorem proving. *J. Autom. Reasoning*, 54(1):69–97, 2015.
- [4] J. R. Burch and D. L. Dill. Automatic verification of pipelined microprocessor control. In David L. Dill, editor, *Proc. of CAV*, volume 818 of *LNCS*, pages 68–80. Springer, 1994.
- [5] 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*, volume 11675 of *LNCS*, pages 157–175. Springer, 2019.
- [6] 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*, volume 11560 of *LNCS*, pages 212–239. Springer, 2019.
- [7] D. Calvanese, S. Ghilardi, A. Gianola, M. Montali, and A. Rivkin. Model completeness, covers and superposition. In *Proc. of CADE*, volume 11716 of *LNCS (LNAI)*, pages 142–160. Springer, 2019.
- [8] D. Calvanese, S. Ghilardi, A. Gianola, M. Montali, and A. Rivkin. SMT-based verification of data-aware processes: a model-theoretic approach. *Math. Struct. Comput. Sci.*, 30(3):271–313, 2020.
- [9] D. Calvanese, S. Ghilardi, A. Gianola, M. Montali, and A. Rivkin. Model completeness, uniform interpolants and superposition calculus. *J. Autom. Reasoning*, 65(7):941–969, 2021.
- [10] D. Calvanese, S. Ghilardi, A. Gianola, M. Montali, and A. Rivkin. Combined Covers and Beth Definability. In *Proc. of IJCAR*, volume 12166 of *LNCS (LNAI)*, pages 181–200. Springer, 2020.
- [11] C.-C. Chang and J. H. Keisler. *Model Theory*. North-Holland Publishing Co., Amsterdam-London, third edition, 1990.
- [12] W. Craig. Three uses of the Herbrand-Gentzen theorem in relating model theory and proof theory. *J. Symbolic Logic*, 22:269–285, 1957.
- [13] V. D’Silva, D. Kroening, M. Purandare, and G. Weissenbacher. Interpolant strength. In *Proc. of VMCAI 2010*, volume 5944 of *LNCS*, pages 129–145. Springer, 2010.
- [14] S. Ghilardi, A. Gianola, and D. Kapur. Computing uniform interpolants for \mathcal{EUF} via (conditional) DAG-based compact representations. In *Proc. of CILC*, volume 2710 of *CEUR Workshop Proceedings*, pages 67–81. CEUR-WS.org, 2020.
- [15] S. Ghilardi, A. Gianola, M. Montali, and A. Rivkin. Petri nets with parameterised data - modelling and verification. In *Proc. of BPM*, volume 12168 of *LNCS*, pages 55–74. Springer, 2020.
- [16] S. Ghilardi, A. Gianola, M. Montali, and A. Rivkin. Delta-BPMN: A Concrete Language and Verifier for Data-Aware BPMN. In *Proc. of BPM*, volume 12875 of *LNCS*, pages 179–196. Springer, 2021.

- [17] S. Ghilardi and S. Ranise. MCMT: A model checker modulo theories. In *Proc. of IJCAR*, pages 22–29, 2010.
- [18] S. Ghilardi and M. Zawadowski. *Sheaves, games, and model completions*, volume 14 of *Trends in Logic—Studia Logica Library*. Kluwer Academic Publishers, Dordrecht, 2002.
- [19] S. Gulwani and M. Musuvathi. Cover algorithms and their combination. In *Proc. of ESOP, Held as Part of ETAPS*, pages 193–207, 2008.
- [20] K. Hoder, L. Kovács, and A. Voronkov. Playing in the grey area of proofs. In *Proc. of POPL 2012*, pages 259–272. ACM, 2012.
- [21] G. Huang. Constructing Craig interpolation formulas. In *Computing and Combinatorics COCOON*, pages 181–190. LNCS, 959, 1995.
- [22] D. Kapur. Shostak’s congruence closure as completion. In *Proc. of RTA*, pages 23–37, 1997.
- [23] 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.
- [24] D. Kapur. Conditional congruence closure over uninterpreted and interpreted symbols. *J. Systems Science & Complexity*, 32(1):317–355, 2019.
- [25] D. Kapur, R. Majumdar, and C. G. Zarba. Interpolation for data structures. In *Proc. of SIGSOFT FSE*, pages 105–116, 2006.
- [26] R. C. Lyndon. An interpolation theorem in the predicate calculus. *Pacific J. Math.*, 9(1):129–142, 1959.
- [27] K. L. McMillan. Lazy abstraction with interpolants. In *Proc. of CAV*, pages 123–136, 2006.
- [28] G. Nelson and D. C. Oppen. Simplification by cooperating decision procedures. *ACM Trans. Program. Lang. Syst.*, 1(2):245–257, 1979.
- [29] A. M. Pitts. On an interpretation of second order quantification in first order intuitionistic propositional logic. *J. Symb. Log.*, 57(1):33–52, 1992.
- [30] P. Pudlák. Lower bounds for resolution and cutting plane proofs and monotone computations. *J. Symb. Log.*, 62(3):981–998, 1997.
- [31] R. E. Shostak. Deciding combinations of theories. *J. ACM*, 31(1):1–12, 1984.
- [32] G. Weissenbacher. Interpolant strength revisited. In *Proc. of SAT 2012*, volume 7317 of *LNCS*, pages 312–326. Springer, 2012.