

On deciding satisfiability by theorem proving with speculative inferences

Maria Paola Bonacina · Christopher A. Lynch ·
Leonardo de Moura

Received: 22 February 2010 / Accepted: 1 December 2010 / Published online: 22 December 2010

Abstract Applications in software verification often require determining the satisfiability of first-order formulæ with respect to background theories. During development, conjectures are usually false. Therefore, it is desirable to have a theorem prover that terminates on satisfiable instances. Satisfiability Modulo Theories (SMT) solvers have proven to be highly scalable, efficient and suitable for integrated theory reasoning. Inference systems with resolution and superposition are strong at reasoning with equalities, universally quantified variables, and Horn clauses. We describe a theorem-proving method that tightly integrates superposition-based inference system and SMT solver. The combination is refutationally complete if background theory symbols only occur in ground formulæ, and non-ground clauses are variable-inactive. Termination is enforced by introducing additional axioms as hypotheses. The system detects any unsoundness introduced by these speculative inferences and recovers from it.

Keywords Program checking · Theorem proving · Satisfiability modulo theories · Combination of theories

1 Introduction

Applications in software verification have benefited greatly from recent advances in automated reasoning. Reasoning about programs often requires determining the satisfiability of first-order formulæ with respect to some background theories. In numerous contexts in

The first author was supported in part by grant no. 2007-9E5KM8 of the Ministero dell'Istruzione Università e Ricerca.

Maria Paola Bonacina
Dipartimento di Informatica, Università degli Studi di Verona, Strada Le Grazie 15, I-37134 Verona, Italy
E-mail: mariapaola.bonacina@univr.it

Christopher A. Lynch
Department of Mathematics and Computer Science, Clarkson University, Potsdam, NY 13699-5815, USA
E-mail: clynch@clarkson.edu

Leonardo de Moura
Microsoft Research, One Microsoft Way, Redmond, WA 98052, USA
E-mail: leonardo@microsoft.com

software verification, quantifiers are also needed. For example, they are used for capturing frame conditions over loops, axiomatizing type systems, summarizing auxiliary invariants over heaps, and for supplying axioms of theories that are not already equipped with decision procedures for ground formulæ. Thus, many verification problems consist in determining the satisfiability of a set of formulæ $\mathcal{R} \uplus P$ modulo a background theory \mathcal{T} , where \mathcal{R} is a set of non-ground clauses without occurrences of \mathcal{T} -symbols, and P is a large ground formula, or set of ground clauses, that typically contains \mathcal{T} -symbols. The set of formulæ \mathcal{R} can be viewed as the axiomatization of an application specific theory. The background theory \mathcal{T} is a combination $\mathcal{T} = \bigcup_{i=1}^n \mathcal{T}_i$ of theories \mathcal{T}_i , $1 \leq i \leq n$, commonly used in hardware and software verification, such as linear arithmetic.

Satisfiability Modulo Theories (SMT) solvers have proven highly scalable, efficient and suitable for integrated theory reasoning. Most SMT solvers are restricted to ground formulæ, and integrate the Davis-Putnam-Logemann-Loveland procedure (DPLL) for propositional logic [26,25], with satellite solvers for ground satisfiability problems in the theories \mathcal{T}_i , $1 \leq i \leq n$, that are therefore built into the SMT solver. The resulting integration is called $\text{DPLL}(\mathcal{T})$, where $\mathcal{T} = \bigcup_{i=1}^n \mathcal{T}_i$. General treatments appeared in [55,61]. For quantifiers, there are situations where the needed instances of universally quantified variables can be computed without loss of completeness: for instance, for certain fragments of the theories of arrays [21,17] and pointers [52], or for local theory extensions [63,43]. Otherwise, techniques to guess how to instantiate variables, based on heuristics and user annotations, were investigated [34,27,39]. They are known as “triggering,” because the terms to be instantiated are called “triggers.” These techniques can be very efficient when they succeed, but they require expensive user guidance, and their incompleteness causes false positives in the program verification tools that use the SMT-solver.

In comparison with SMT solvers, generic inference systems based on superposition and resolution are refutationally complete for first-order logic with equality; they are strong at reasoning with equalities, universally quantified variables, and Horn clauses. Available surveys include [10,56,48]. A standard superposition-based inference system was proved to terminate and hence to be a satisfiability procedure for several theories of data structures, including arrays and recursive data structures, and their combinations [5,3,15,4].

The $\text{DPLL}(\Gamma)$ system of [28] integrates an SMT solver with a generic inference system Γ based on superposition and resolution. The DPLL engine works by building a candidate model: if the construction succeeds, it returns satisfiable; if it fails definitely, it returns unsatisfiable. The inference system Γ works by deducing clauses from clauses and removing redundant clauses: if it generates the empty clause \square , that represents contradiction, it returns unsatisfiable. The key to their integration is that the literals in the candidate model built by the DPLL engine can occur as premises of Γ -inferences. The resulting system aims at uniting the strengths of SMT solvers – propositional efficiency, fast theory solvers, tight integration – with those of superposition-based theorem provers, especially general reasoning about quantifiers without recurring to incomplete heuristics.

However, in general, $\text{DPLL}(\Gamma)$ is not refutationally complete, when both \mathcal{T} and \mathcal{R} are not empty, even when \mathcal{T} -symbols do not occur in \mathcal{R} . For example, assume $\mathcal{R} = \{x \simeq a \vee x \simeq b\}$ and $P = \emptyset$, and the background theory \mathcal{T} is arithmetic. The clause $x \simeq a \vee x \simeq b$ implies that any model has a domain with at most two elements, which is clearly incompatible with any model for arithmetic, that requires an infinite domain. In other words, $\text{DPLL}(\Gamma)$ does not have a way to detect unsatisfiability due to the lack of models with infinite domain. A first contribution of this article is a revised version of $\text{DPLL}(\Gamma)$, named $\text{DPLL}(\Gamma + \mathcal{T})$, with sufficient conditions to make it refutationally complete when \mathcal{T} is not empty.

DPLL($\Gamma + \mathcal{T}$) has to combine the built-in theories in $\mathcal{T} = \bigcup_{i=1}^n \mathcal{T}_i$ and the axiomatized theory \mathcal{R} . Combination of built-in theories is usually done by the *equality sharing* method of [53], later dubbed “Nelson-Open scheme” from the names of the authors. This method has three requirements. First, the theories cannot share function symbols. Second, each \mathcal{T}_i must be *stably infinite*. A theory \mathcal{T}_i is stably infinite if every \mathcal{T}_i -satisfiable ground formula has a \mathcal{T}_i -model with domain of infinite cardinality. Third, every \mathcal{T}_i -solver must be capable of generating all entailed disjunctions of equalities between shared constants. The third requirement is relaxed in *model-based theory combination* [29], which is a version of equality sharing, where it is sufficient that each \mathcal{T}_i -solver generates the equalities between shared constants that are true in the current candidate \mathcal{T}_i -model. Thus, this method requires that each \mathcal{T}_i -solver generates a candidate model. On the other hand, combination of axiomatized theories in a superposition-based engine requires that they do not share function symbols and are *variable-inactive*: under these hypotheses, if the superposition-based system terminates on satisfiability problems in each theory, it also terminates on satisfiability problems in their union [3,4]. A second contribution of this article is to explain how to apply known results on variable inactivity [18,4,17] to combine built-in theories by model-based theory combination and axiomatized theories in DPLL($\Gamma + \mathcal{T}$).

In software verification, during development time, several conjectures are false because of errors in the implementation or specification. Therefore, it is desirable to have a theorem prover that terminates on satisfiable instances. In general, this is not a realistic goal since pure first-order logic is not decidable, and, even worse, there is no sound and complete procedure for first-order logic formulæ of linear arithmetic with uninterpreted functions [41]. Axioms such as *transitivity*

$$\neg(x \sqsubseteq y) \vee \neg(y \sqsubseteq z) \vee x \sqsubseteq z$$

and *monotonicity*

$$\neg(x \sqsubseteq y) \vee f(x) \sqsubseteq f(y)$$

are problematic for any resolution-based Γ , since they tend to generate an unbounded number of clauses, even with a selection function that selects negative literals to prevent self-resolutions. Such axioms may arise in formalizations of type systems for programming languages. In that context, the symbol \sqsubseteq represents a subtype relationship, and the monadic function f represents a type constructor, such as `Array-of`.

As an example, assume that the axiomatization contains a monotonicity axiom

$$\neg(x \sqsubseteq y) \vee f(x) \sqsubseteq f(y).$$

Unrestricted resolution would resolve it with itself, that is, with a variant

$$\neg(x' \sqsubseteq y') \vee f(x') \sqsubseteq f(y')$$

to generate $\neg(x \sqsubseteq y) \vee f^2(x) \sqsubseteq f^2(y)$, hence the infinite series

$$\{\neg(x \sqsubseteq y) \vee f^i(x) \sqsubseteq f^i(y)\}_{i \geq 0},$$

where at each step the original axiom resolves with $\neg(x \sqsubseteq y) \vee f^{i-1}(x) \sqsubseteq f^{i-1}(y)$ to yield its successor. A selection function that selects negative literals prevents all these inferences, because $f(x') \sqsubseteq f(y')$ cannot resolve with $\neg(x \sqsubseteq y)$, if $\neg(x' \sqsubseteq y')$ is selected. However, even resolution with negative selection generates an infinite series

$$\{f^i(a) \sqsubseteq f^i(b)\}_{i \geq 0}$$

from monotonicity and each literal $a \sqsubseteq b$ in the input. In practice, it is seldom the case that we need to go beyond $f(a) \sqsubseteq f(b)$ or $f^2(a) \sqsubseteq f^2(b)$ to show satisfiability.

A third contribution of this article is a new version of $\text{DPLL}(\Gamma + \mathcal{T})$ with *speculative inferences*, a feature suggested by [49]. The idea is to allow the prover, or the experimenter, to guess additional axioms, that avoid infinitary behaviors such as that induced by the monotonicity axiom. If the additional axioms are a cause of unsoundness, by turning a satisfiable set into an unsatisfiable one, the prover detects it and recovers from it automatically. The resulting method yields decision procedures for several axiomatizations of type systems that are relevant to software verification.

This article is organized as follows. Section 2 provides the background. Section 3 shows how to apply previous results on superposition and variable-inactive theories in $\text{DPLL}(\Gamma + \mathcal{T})$; it states the requirements on the problem $\mathcal{R} \uplus P$ and the background theory \mathcal{T} , setting the stage for the sequel. Section 4 presents the new $\text{DPLL}(\Gamma + \mathcal{T})$ system with speculative inferences. Section 6 introduces the notion of *essentially finite theories* – a generalization of the *finite model property* – exhibits essentially finite theories, and shows how $\text{DPLL}(\Gamma + \mathcal{T})$ with speculative inferences yields a decision procedure for them and their combinations. Section 7 concludes with summary of the results, comparison with related work, and directions for further research. A short version of this article appeared in [20].

2 Background

We assume basic notions from logic used in theorem proving. Let Σ be a *signature* consisting of a set of *function* and *predicate* symbols, each with its *arity*, denoted by $\text{arity}(f)$, for symbol f . We call 0-arity function symbols *constant* symbols, and use a, b, c, d for constants, f, g, h for non-constant function symbols, p for a predicate symbol, x, y, z, u for variables, \simeq for equality, and \bowtie for either \simeq or $\not\simeq$: these three symbols are symmetric. Two signatures are *disjoint* if they share no function or predicate symbol other than \simeq . Terms, literals, clauses, sentences and formulæ are defined as usual. A clause (i.e., a disjunction of literals) is *positive* if all its literals are; it is *Horn* if it has at most one positive literal. We use t, s, l, r for terms, l, m for literals, C, D for clauses, \square for the empty clause, which denotes contradiction, and F, N, P, S for sets of clauses. $\text{Var}(l)$ is the set of variables occurring in l . The notation $l[t]$ means that t appears as subterm of l . A first-order Σ -theory is *presented*, or axiomatized, by a set of Σ -sentences. Two theories are *disjoint* if their signatures are. We reserve calligraphic letters, such as \mathcal{T} and \mathcal{R} , for presentations of theories. A *Horn presentation* is a set of non-negative Horn clauses. *Defined*, or *interpreted*, symbols are those symbols whose interpretation is restricted to the models of a theory, whereas *free*, or *uninterpreted*, symbols are those symbols whose interpretation is unrestricted.

A Σ -*structure* Φ consists of a non-empty universe, or domain, $|\Phi|$, and an interpretation for variables and symbols in Σ . We use v, w for elements of $|\Phi|$. For each $f \in \Sigma$, the interpretation of f is denoted by $\Phi(f)$. For a function symbol f with $\text{arity}(f) = n$, $\Phi(f)$ is an n -ary function on $|\Phi|$ with $\text{range}(\Phi(f)) = \{w \mid \exists v \in |\Phi|, \Phi(f)(v) = w\}$. For a predicate symbol p with $\text{arity}(p) = n$, $\Phi(p)$ is a subset of $|\Phi|^n$. The interpretation of a term t is denoted by $\Phi(t)$. If t is a variable or constant, $\Phi(t)$ is an element in $|\Phi|$. Otherwise, $\Phi(f(t_1, \dots, t_n)) = \Phi(f)(\Phi(t_1), \dots, \Phi(t_n))$. If S is a set of terms, $\Phi(S) = \{\Phi(t) \mid t \in S\}$. Satisfaction $\Phi \models C$ is defined as usual; if $\Phi \models C$, then Φ is a *model* of C . For refutational completeness in first-order theorem proving it is sufficient to consider *Herbrand interpretations*, where the domain is the Herbrand universe and constant and function symbols are interpreted as themselves.

An *inference system* Γ is a set of inference rules. *Ordering-based* inference systems use an ordering \succ on terms and literals to restrict *expansion inferences*, that expand the existing set by generating clauses, and to define *contraction inferences*, that contract the set by removing clauses. This ordering is assumed to be a *complete simplification ordering*: it is *stable* (if $s \succ t$ then $s\sigma \succ t\sigma$ for all substitutions σ), *monotone* (if $s \succ t$ then $l[s] \succ l[t]$ for all l), it has the *subterm property* ($l[t] \succ t$ for all t and $l \neq t$), hence it is *well-founded* (there is no infinite decreasing chain $t_1 \succ t_2 \succ \dots t_i \succ t_{i+1} \succ \dots$) [31], and it is *total* on ground terms and literals (if $s \neq t$, then either $s \succ t$ or $t \succ s$). The ordering is extended to equations and clauses by multiset extension, which preserves well-foundedness [33]. An *inference rule* with n premises has a *main premise* and $n - 1$ *side* premises. For an expansion rule, the main premise yields the conclusion in the context of the side premises; such a rule is *sound* if the conclusion is a logical consequence of the premises. For a contraction rule, the main premise is reduced to the conclusion or removed; such a rule is *sound* if the main premise is a logical consequence of side premises and conclusion, if present. Premises and conclusion of an inference γ are denoted by $P(\gamma)$ and $C(\gamma)$, respectively. We write $\gamma \in \Gamma$ to say that γ is an application of an inference rule in Γ .

Clauses deleted by contraction are *redundant*. Redundancy is defined based on well-founded orderings on clauses, whereby a ground clause is redundant in a set of clauses, if it is entailed by smaller ground clauses in the set, and a clause is redundant if all its ground instances are [6], or well-founded orderings on proofs, whereby a clause is redundant in a set of clauses, if it does not affect its minimal proofs [19]. An inference is redundant if it uses or generates a redundant clause. A clause or inference that is not redundant is *irredundant*. A set of clauses N is *saturated with respect to Γ* , or *Γ -saturated*, if all Γ -inferences in N are redundant. Given an input set of clauses S_0 , a *Γ -derivation* is a sequence $S_0 \vdash_{\Gamma} S_1 \vdash_{\Gamma} \dots S_i \vdash_{\Gamma} S_{i+1} \vdash_{\Gamma} \dots$ where at each step S_{i+1} is derived from S_i by a Γ -inference; its *limit* is the set of *persistent clauses* $S_{\infty} = \bigcup_{i \geq 0} \bigcap_{j \geq i} S_j$. A Γ -derivation is *fair*, if all expansion inferences from persistent irredundant premises are done eventually. If a Γ -derivation is fair, its limit is saturated. Since Γ is non-deterministic, there may be more than one Γ -derivation from a given S_0 . The combination of Γ with a *search plan*, that controls the choice of inferences, yields a deterministic *theorem-proving strategy*, or *proof procedure*, termed *Γ -strategy*, or *Γ -procedure*, whose derivation is unique given S_0 . A strategy or procedure is fair if all its derivations are. A recent abstract treatment of these notions, with references to their history, appeared in [12].

Let I be a mapping, called a *model functor*, that assigns to each set of ground clauses N , not containing \square , a Herbrand interpretation I_N , called the *candidate model*. A clause C is a counterexample for I_N if $I_N \not\models C$; C is a minimal counterexample if, in addition, there is no other counterexample D for I_N such that $C \succ D$. An inference system Γ has the *reduction property for counterexamples*, if for all sets N of clauses and counterexamples C for I_N in N , there is an inference in Γ from N with main premise C , side premises that are true in I_N , and conclusion D that is a smaller counterexample for I_N than C . This property is used in proofs of refutational completeness since at least [6] according to the following standard:

Theorem 1 *If N is a Γ -saturated set of ground clauses and Γ has the reduction property for counterexamples, then N is unsatisfiable if and only if it contains \square .*

Proof The “if” direction is trivial. To prove the “only if” direction, one proves its contrapositive: if N does not contain \square , then it is satisfiable. By way of contradiction, assume that it is not. Then, for every candidate model I_N there is a counterexample in N . Let C be a minimal counterexample for I_N in N . By the reduction property for counterexamples, there is a smaller counterexample D , conclusion of a Γ -inference. If $D \in N$, then C would not

have been minimal to begin with. If $D \notin N$, then N is not saturated, which contradicts the hypothesis. \square

We do not assume a specific inference system: Γ is a parameter for $\text{DPLL}(\Gamma + \mathcal{T})$. Since examples and proofs for specific theories mention concrete inference rules, and in order to make this article self-contained, typical expansion and contraction rules are collected in Figure 1: in contraction, what is above the double inference line is replaced by what is below, whereas in expansion, as usual, what is below the single inference line is added to what is above. In resolution, l and l' are the literals *resolved upon*. In paramodulation, $l[s']$ ($C \vee l[s']$) is the literal (clause) *paramodulated into*, and $s \simeq t$ ($D \vee s \simeq t$) is the literal (clause) *paramodulated from*. The same terminology applies to superposition.¹ Resolution, paramodulation and superposition originally appeared in [59], [57] and [45], respectively. Since then, they were the object of decades of research: contemporary versions of these rules appeared in [42, 60, 6, 19]; more references and history can be found in [10, 56, 12, 48, 11].

In addition to the ordering, expansion inference rules can be restricted by a *selection function*, that selects negative literals [6]. A clause can have all, some, or none of its negative literals selected, depending on the selection function. In *resolution with negative selection*, *paramodulation with negative selection* and *superposition with negative selection*, the ordering constraint on the negative literal resolved upon, on the literal paramodulated into (i.e., Condition (iii) in Figure 1) and on the literal superposed into (i.e., Condition (v) in Figure 1) is lifted, and replaced by two requirements: the negative literal resolved upon, the literal paramodulated into and the literal superposed into must be selected, and the other premise contains no selected literal. If some negative literal is selected for each clause containing one, one premise in each resolution, paramodulation or superposition inference will be a positive clause, yielding a *positive strategy*; if, in addition, all clauses are Horn, only positive unit clauses can be resolved upon or superposed or paramodulated from, yielding a *unit strategy* with *unit resolution* and *unit superposition* [32]. Contemporary theorem provers, including the implementation of $\text{DPLL}(\Gamma + \mathcal{T})$ on top of Z3 [30], use resolution with negative selection to implement *hyperresolution* [58]. In hyperresolution, the side premises, termed *satellites*, are positive clauses that resolve away all negative literals in the main premise, termed *nucleus*, generating a positive clause, in a single step with a simultaneous unifier of all pairs of literals resolved upon.² A selection function that selects some negative literal in each clause containing one induces resolution to simulate hyperresolution as a macro inference involving several steps of resolution. In this article, hyperresolution is realized via resolution with negative selection.

3 Variable inactivity in $\text{DPLL}(\Gamma + \mathcal{T})$

In this section we see how previous results from [4, 18, 17] can be imported into $\text{DPLL}(\Gamma + \mathcal{T})$ to combine a built-in theory \mathcal{T} and an axiomatized theory \mathcal{R} , where both \mathcal{T} and \mathcal{R} can be themselves unions of theories. In a purely rewrite-based approach there is no \mathcal{T} and all axioms are part of the input in \mathcal{R} . The ordering \succ of Γ is required to be *good* [4, 16], meaning that $t \succ c$ for all ground compound term t and constant c . A fair Γ -strategy

¹ We use superposition when the literal paramodulated into is equational, and paramodulation otherwise. Other articles reserve superposition to unit clauses or positive equations.

² This instance of the rule is called *positive hyperresolution*. The dual rule named *negative hyperresolution* operates in the same way with polarities exchanged. Since selection functions are defined to select negative literals, negative hyperresolution falls outside of this discussion.

Resolution	$\frac{C \vee \neg l \quad D \vee l'}{(C \vee D)\sigma}$	$\forall m \in C : \neg l \sigma \not\leq m \sigma, \forall m \in D : l' \sigma \not\leq m \sigma$
Factoring	$\frac{C \vee l \vee l'}{(C \vee l)\sigma}$	$\forall m \in C : l \sigma \not\leq m \sigma$
Paramodulation	$\frac{C \vee l[s'] \quad D \vee s \simeq t}{(C \vee D \vee l[t])\sigma}$	(i), (ii), (iii)
Superposition	$\frac{C \vee l[s'] \bowtie r \quad D \vee s \simeq t}{(C \vee D \vee l[t] \bowtie r)\sigma}$	(i), (ii), (iv), (v)
Reflection	$\frac{C \vee s' \not\leq s}{C\sigma}$	$\forall l \in C : (s' \not\leq s)\sigma \not\leq l \sigma$
Equational Factoring	$\frac{C \vee s \simeq t \vee s' \simeq t'}{(C \vee t \not\leq t' \vee s \simeq t')\sigma}$	(i), $\forall l \in \{s' \simeq t'\} \cup C : (s \simeq t)\sigma \not\leq l \sigma$

where σ is the most general unifier (mgu) of l and l' in resolution and factoring, and of s and s' in the other rules; s' is not a variable in paramodulation and superposition, and the following abbreviations hold:

- (i) is $s \sigma \not\leq t \sigma$,
- (ii) is $\forall m \in D : (s \simeq t)\sigma \not\leq m \sigma$,
- (iii) is $\forall m \in C : l[s']\sigma \not\leq m \sigma$,
- (iv) is $l[s']\sigma \not\leq r \sigma$, and
- (v) is $\forall m \in C : (l[s'] \bowtie r)\sigma \not\leq m \sigma$.

Strict Subsumption	$\frac{C \quad D}{C}$	$D \triangleright C$
Simplification	$\frac{C[l] \quad s \simeq t}{C[t\sigma] \quad s \simeq t}$	$l = s\sigma, \quad s\sigma \triangleright t\sigma, \quad C[l] \triangleright (s \simeq t)\sigma$
Deletion	$\frac{C \vee t \simeq t}{C}$	

where $D \triangleright C$ if $D \geq C$ and $C \not\geq D$; and $D \geq C$ if $C\sigma \subseteq D$ (as multisets) for some substitution σ . In practice, theorem provers also apply subsumption of variants (if $D \geq C$ and $C \geq D$, the oldest clause is retained) and tautology deletion (that removes clauses such as $C \vee l \vee \neg l$).

Fig. 1 Sample expansion and contraction rules: in expansion what is below the inference line is added to the clause set that contains what is above the inference line; in contraction what is above the double inference line is removed and what is below is added.

is shown to be an \mathcal{R} -satisfiability procedure, by showing that it is guaranteed to terminate on \mathcal{R} -satisfiability problems $\mathcal{R} \uplus S$, where S is a set of ground unit \mathcal{R} -clauses. Variable-inactivity was introduced in [3,4]:

Definition 1 A clause C is *variable-inactive* if no maximal literal in C is an equation $t \simeq x$ where $x \notin \text{Var}(t)$. A set of clauses is *variable-inactive* if all its clauses are.

Definition 2 A theory presentation \mathcal{R} is *variable-inactive* for an inference system Γ if the limit S_∞ of a fair Γ -derivation from $S_0 = \mathcal{R} \uplus S$ is variable-inactive, where S is a set of ground unit \mathcal{R} -clauses.

It was proved in [4] (cf. Theorem 4.1 and Corollary 3) that termination is *modular*:

Theorem 2 (Armando, Bonacina, Ranise and Schulz 2009) *Let $\mathcal{R}_1, \dots, \mathcal{R}_n$ be disjoint and variable-inactive for Γ , and let $\mathcal{R} = \bigcup_{i=1}^n \mathcal{R}_i$. If a fair Γ -strategy terminates on \mathcal{R}_i -satisfiability problems, for $1 \leq i \leq n$, then it terminates also on \mathcal{R} -satisfiability problems.*

A *cardinality constraint* is a clause containing only non-trivial (i.e., other than $x \simeq x$) positive equations between variables (e.g., $y \simeq x \vee y \simeq z$). Such a clause is clearly not variable-inactive. The following key lemma was proved in [18] (cf. Lemma 5.2):

Lemma 1 (Bonacina, Ghilardi, Nicolini, Ranise and Zucchelli 2006) *If S_0 is a finite satisfiable set of clauses, then S_0 admits no infinite models if and only if the limit S_∞ of a fair Γ -derivation from S_0 contains a cardinality constraint.*³

It follows that (cf. Theorem 4.5 in [4]):

Theorem 3 (Armando, Bonacina, Ranise and Schulz 2009) *If \mathcal{R} is variable-inactive for Γ , then it is stably-infinite.*

Thus, Γ reveals the lack of stable infiniteness by generating a cardinality constraint. The original versions of Theorem 2, Lemma 1 and Theorem 3 were proved in a context where equality was the only predicate and superposition the main expansion inference rule of Γ . It is trivial to extend them to the case where the signature of \mathcal{R} introduces predicate symbols other than equality, and Γ features also resolution and paramodulation. For instance, for Theorem 2, the essence of the proof is to show that there are only finitely many inferences across theories: disjointness of the signatures prevents not only superpositions, but also paramodulations from compound terms, and resolutions; variable inactivity prevents not only superpositions, but also paramodulations from variables; thus, the only inferences across theories are superpositions and paramodulations from shared constants, that are finitely many.

It is useful to import results from the rewrite-based approach to $\text{DPLL}(\Gamma + \mathcal{T})$, applied to a problem $\mathcal{R} \uplus P$ modulo \mathcal{T} , because $\text{DPLL}(\Gamma + \mathcal{T})$ uses Γ as an \mathcal{R} -solver applied to \mathcal{R} -satisfiability problems $\mathcal{R} \uplus S$, where S is a set of ground unit \mathcal{R} -clauses. The initial set of ground clauses P typically contains also \mathcal{T} -symbols. However, P is subject to *purification*, which is a standard step in the Nelson-Oppen method. This transformation, also known as *separation* [37], separates occurrences of function symbols from different signatures occurring in ground terms, by introducing new constant symbols. For example, $f(g(a)) \simeq b$, where f and g belong to different signatures, becomes $f(c) \simeq b \wedge g(a) \simeq c$, where c is new. Since only constants are introduced, the set remains ground. Thus, P is transformed in two disjoint sets P_1 and P_2 , where P_1 contains only \mathcal{R} -symbols and P_2 only \mathcal{T} -symbols. Since a key feature of $\text{DPLL}(\Gamma + \mathcal{T})$ is that Γ deals only with *non-ground clauses* and *ground unit clauses*, it is indeed the case that Γ works on an \mathcal{R} -satisfiability problems $\mathcal{R} \uplus S$: initially, S will be the subset of unit clauses from P_1 .

$\text{DPLL}(\Gamma + \mathcal{T})$ needs to combine $\mathcal{T}_1, \dots, \mathcal{T}_n, \mathcal{R}$ in the Nelson-Oppen scheme, which requires that the theories are disjoint, stably infinite, and each solver generates all entailed disjunctions of equalities between shared constants. We assume that $\mathcal{T}_1, \dots, \mathcal{T}_n$ satisfy these requirements and that \mathcal{R} is disjoint from each of them. For stable infiniteness of \mathcal{R} , we require that \mathcal{R} is variable-inactive and apply Theorem 3. In practice, this condition is checked dynamically: in the implementation of $\text{DPLL}(\Gamma + \mathcal{T})$ on top of Z3, the superposition-based engine is equipped with a test that detects the generation of variable-inactive clauses, hence cardinality constraints, and discovers whether \mathcal{R} is not stably infinite. Such a test also excludes upfront a situation such as $\mathcal{R} = \{x \simeq a \vee x \simeq b\}$ of the example in Section 1. For the generation of disjunctions of equalities between shared constants by the \mathcal{R} -solver Γ ,

³ Lemma 5.2 in [18] requires that the superposition-based inference system is invariant with respect to renaming finitely many constants. Most inference systems satisfy a stronger requirement, namely they allow signature extensions, e.g., to introduce Skolem constants.

the fairness of the Γ -derivation ensures that every theorem is implied by some generated formulæ.⁴ An explicit proof that the superposition-based engine generates formulæ that entail all disjunctions of equalities between constants in the axiomatized theory was given in [17] (cf. Theorem 71). If contraction is also done systematically, only irredundant clauses generated by Γ are kept and passed to the DPLL(\mathcal{T}) core.

The following definition summarizes the problem requirements for the sequel:

Definition 3 A set of formulæ $S = \mathcal{R} \uplus P$ is *smooth* with respect to a background theory $\mathcal{T} = \bigcup_{i=1}^n \mathcal{T}_i$, or \mathcal{T} -*smooth* for short, if

- $\mathcal{T}_1, \dots, \mathcal{T}_n$ and \mathcal{R} are pairwise disjoint,
- $\mathcal{T}_1, \dots, \mathcal{T}_n$ are stably infinite,
- \mathcal{R} is variable-inactive, and
- P is a set of ground formulæ $P_1 \uplus P_2$, where P_1 contains only \mathcal{R} -symbols, and P_2 only \mathcal{T} -symbols.

Note that uninterpreted symbols are \mathcal{R} -symbols. In summary, variable inactivity is an ingredient for: (1) modularity of termination of Γ , when \mathcal{R} is a union of axiomatized theories [4]; (2) stable infiniteness of \mathcal{R} [18], hence combination of axiomatized theories and built-in theories; (3) refutational completeness of DPLL($\Gamma + \mathcal{T}$) when both \mathcal{T} and \mathcal{R} are not empty (cf. Theorem 4 in the next section).

4 A new DPLL($\Gamma + \mathcal{T}$) system with speculative inferences

DPLL($\Gamma + \mathcal{T}$) works on *hypothetical clauses* of the form $H \triangleright C$, where C is a clause, and the *hypothesis* H is a set of ground literals. The hypothesis is interpreted as a conjunction, and a hypothetical clause $(l_1 \wedge \dots \wedge l_n) \triangleright (l'_1 \vee \dots \vee l'_m)$ is interpreted as $\neg l_1 \vee \dots \vee \neg l_n \vee l'_1 \vee \dots \vee l'_m$. As we shall see, the literals in H come from the candidate model built by DPLL($\Gamma + \mathcal{T}$), and are the literals that C depends on, in the sense that they were used as premises to infer C by Γ -inferences. As it was done in [28] for DPLL(Γ), DPLL($\Gamma + \mathcal{T}$) is described as a *transition system* with two modes: *search mode* and *conflict resolution mode*.

In search mode, the state of the system has the form $M \parallel F$, where M is a sequence of *assigned literals*, and F a set of *hypothetical clauses*. Intuitively, M represents a partial assignment to ground literals, possibly with a justification, and therefore it represents a partial model, or a set of candidate models. An assigned literal can be either a *decided literal* or an *implied literal*. A decided literal represents a guess, and has no justification. An implied literal l_C is a literal l justified by a clause C : all other literals of C are false in M so that l needs to be true. No assigned literal occurs twice in M nor does it occur negated in M . If neither l nor $\neg l$ appears in M , then l is said to be *undefined*.

In conflict resolution mode, the state has the form $M \parallel F \parallel C$, where C is a ground clause whose literals are all false under M . Such a clause is in *conflict*. If C is $l_1 \vee \dots \vee l_n$, then $\neg C$ is the formula $\neg l_1 \wedge \dots \wedge \neg l_n$. We could state that C is in conflict by writing $M \models \neg C$. In DPLL($\Gamma + \mathcal{T}$), the DPLL engine accepts only propositional clauses, whereas the theory solvers accept ground first-order clauses and Γ accepts first-order clauses. To bridge this gap, an *abstraction function* maps first-order ground atoms to propositional atoms. Thus, it is customary to write $M \models_P \neg C$, read M “propositionally satisfies” $\neg C$, to say that M satisfies the propositional abstraction of $\neg C$.

⁴ Fairness guarantees even more: every theorem has a minimal proof in the limit; see [12] for details.

Definition 4 Given an input set of clauses $S = \mathcal{R} \uplus P$, a $DPLL(\Gamma + \mathcal{T})$ -derivation is a sequence of state transitions

$$\Delta_0 \Longrightarrow \Delta_1 \Longrightarrow \dots \Delta_i \Longrightarrow \Delta_{i+1} \Longrightarrow \dots$$

where $\forall i \geq 0$, Δ_i is of the form $M_i \parallel F_i$ or $M_i \parallel F_i \parallel C_i$, each transition is determined by a $DPLL(\Gamma + \mathcal{T})$ -rule, and $\Delta_0 = \parallel F_0$ for $F_0 = \{\emptyset \triangleright C \mid C \in S\}$.

In the sequel, we use C for $\emptyset \triangleright C$, $clauses(F)$ to denote the set $\{C \mid H \triangleright C \in F\}$, $ngclauses(F)$ for the subset of non-ground clauses of $clauses(F)$, $lits(M)$ to denote the set of assigned literals, $lits_{\mathcal{R}}(M)$ for the subset of assigned \mathcal{R} -literals and $clauses^*(M \parallel F)$ for $ngclauses(F) \cup lits_{\mathcal{R}}(M)$.

4.1 Speculative inferences in $DPLL(\Gamma + \mathcal{T})$

In theorem proving applied to mathematics, most conjectures are true. Thus, it is customary to sacrifice completeness for efficiency, and retain soundness, which is necessary to attribute unsatisfiability to the input set of clauses if a proof is found. A traditional example is *deletion by weight* [51], where clauses that are too “heavy” are deleted. In theorem proving applied to verification, most conjectures are false. Thus, it was suggested in [49] to sacrifice soundness for termination, and retain completeness, which is necessary to establish satisfiability if a proof is *not* found. Dually to deletion by weight, an unsound inference could suppress literals in clauses that are too heavy. We call *speculative* an inference that may turn out to be unsound.

We consider a single speculative inference rule: adding an arbitrary clause C . Such a step may be unsound because C may not be implied by the given set. This rule is simple, but can simulate different kinds of speculative inferences. Suppose we want to suppress the literals D in $C \vee D$, then we can simply add C , which subsumes $C \vee D$. Suppose a clause $C[t]$ contains a deep term t , and we want to replace it with a constant a . We can accomplish this by adding $t \simeq a$.

The idea is to extend $DPLL(\Gamma + \mathcal{T})$ with a *reversible* transition rule `SpeculativeIntro` for speculative inferences. Rather than merely adding a clause C , `SpeculativeIntro` introduces a hypothetical clause $\lceil C \rceil \triangleright C$ into F and it adds $\lceil C \rceil$ to M : $\lceil C \rceil$ is a new propositional variable used as a label for clause C . By adding $\lceil C \rceil$ to M , the system records the fact that it is *guessing* C . `SpeculativeIntro` is *reversible*, because the system uses $\lceil C \rceil$ to track the consequences of having added C . The hypothetical clause $\lceil C \rceil \triangleright C$ is semantically equivalent to $\neg \lceil C \rceil \vee C$. This clause does not change the satisfiability of the input formula because $\lceil C \rceil$ is a new propositional variable:

`SpeculativeIntro`

$$M \parallel F \quad \Longrightarrow \quad M \lceil C \rceil \parallel F, \lceil C \rceil \triangleright C \quad \mathbf{if} \quad \begin{cases} C \notin clauses(F), \\ \lceil C \rceil \text{ is new,} \\ \lceil C \rceil, \neg \lceil C \rceil \notin M. \end{cases}$$

Note that $\lceil C \rceil$ is added to M as a decided literal. The first condition says that we do not guess a clause that we already have. The second condition requires $\lceil C \rceil$ to be a new symbol with respect to the initial signature. The third condition prevents the system from adding C , if it was already done ($\lceil C \rceil \in M$), or if the addition was already discovered to be inconsistent with the current partial model M ($\neg \lceil C \rceil \in M$).

4.2 Model-based theory combination in $DPLL(\Gamma + \mathcal{T})$

In order to combine the theories in $\mathcal{T} = \bigcup_{i=1}^n \mathcal{T}_i$ and \mathcal{R} in the Nelson-Oppen scheme, every \mathcal{T}_i -solver, $1 \leq i \leq n$, needs to communicate to the other theories, including \mathcal{R} , the (disjunctions of) equalities between shared constants entailed by \mathcal{T}_i and \mathcal{P} . The next transition rule takes care of this requirement, according to model-based theory combination [29]. We assume that every \mathcal{T}_i -solver builds a specific candidate \mathcal{T}_i -model for M , that we denote by $\text{model}_i(M)$. For instance, solvers for linear arithmetic satisfy this requirement [35]. The idea is to inspect $\text{model}_i(M)$ and propagate all the equalities it implies, hedging that they are consistent with the other theories, including \mathcal{R} . Since these equalities are *guesses*, if one of them turns out to be inconsistent, backtracking will be used to fix $\text{model}_i(M)$. The rationale for this approach is practical: it is generally far less expensive to enumerate the equalities satisfied in a particular \mathcal{T}_i -model than those satisfied by all \mathcal{T}_i -models consistent with M ; in most experiments, the number of equalities that are really relevant turns out to be small.

PropagateEq

$$M \parallel F \quad \Longrightarrow \quad M \ t \simeq s \parallel F \quad \text{if} \quad \begin{cases} t \text{ and } s \text{ are ground,} \\ t, s \text{ occur in } F, \\ (t \simeq s) \text{ is undefined in } M, \\ \text{model}_i(M)(t) = \text{model}_i(M)(s), \end{cases}$$

for every theory \mathcal{T}_i , $1 \leq i \leq n$. Since the \mathcal{T}_i -solvers only deal with ground clauses, this rule treats only ground equalities, and therefore only ground terms that appear in F . The reason why it adds equalities between ground terms and not only between shared constants will be explained in relation to the Deduce rule.

4.3 The core transition rules of $DPLL(\Gamma + \mathcal{T})$

Figure 2 reports the basic and theory propagation rules of $DPLL(\Gamma + \mathcal{T})$ from [28].

The Decide rule is not concerned with literals in hypotheses, since such literals already come from M . The Deduce rule realizes the interface with Γ : assume γ is an inference of Γ with n premises, $\{H_1 \triangleright C_1, \dots, H_m \triangleright C_m\}$ is a set of hypothetical clauses in F , $\{l_{m+1}, \dots, l_n\}$ is a set of assigned literals in M , and $H(\gamma)$ denotes the set $H_1 \cup \dots \cup H_m \cup \{l_{m+1}, \dots, l_n\}$; if γ with premises $P(\gamma) = \{C_1, \dots, C_m, l_{m+1}, \dots, l_n\}$ yields $C(\gamma)$, the latter is added to F as $H(\gamma) \triangleright C(\gamma)$. The hypotheses of the clauses $H_i \triangleright C_i$ are hidden from the inference rules in Γ . Our Deduce rule differs from its predecessor in [28], named Deduce[#], in the range of allowed premises from F . Deduce[#] allowed Γ to use as premises $\text{lits}(M)$, and non-ground clauses and ground unit clauses from $\text{clauses}(F)$. Our Deduce allows Γ to use only $\text{clauses}^*(M \parallel F) = \text{ngclauses}(F) \cup \text{lits}_{\mathcal{R}}(M)$. This is a consequence of the addition of PropagateEq, which adds the relevant ground unit clauses directly to M , so that Γ finds them in $\text{lits}_{\mathcal{R}}(M)$. This is the reason why we let PropagateEq add equalities between ground terms and not only between shared constants.

A hypothetical clause $H \triangleright C$ is in *conflict* if every literal in C is complementary to an assigned literal. The Conflict rule converts a hypothetical conflict clause $H \triangleright C$ into a regular clause by negating its hypotheses, and puts the $DPLL(\Gamma + \mathcal{T})$ system in conflict resolution mode. The Explain rule unfolds literals from conflict clauses that were produced by unit propagation. Any clause derived by Explain can be added to F by the Learn rule, because it is a logical consequence of the original set of clauses. The Backjump rule drives the $DPLL(\Gamma + \mathcal{T})$ system back from conflict resolution mode to search mode, and it unassigns

Decide

$$M \parallel F \quad \Longrightarrow M \parallel I \parallel F \quad \text{if } \begin{cases} l \text{ is ground,} \\ l \text{ or } \neg l \text{ occurs in } F, \\ l \text{ is undefined in } M. \end{cases}$$

UnitPropagate

$$M \parallel F, H \triangleright (C \vee l) \quad \Longrightarrow M \parallel_{H \triangleright (C \vee l)} \parallel F, H \triangleright (C \vee l) \quad \text{if } \begin{cases} l \text{ is ground,} \\ M \models_P \neg C, \\ l \text{ is undefined in } M. \end{cases}$$

Deduce

$$M \parallel F \quad \Longrightarrow M \parallel F, H(\gamma) \triangleright C(\gamma) \quad \text{if } \begin{cases} \gamma \in \Gamma, \\ P(\gamma) \subseteq \text{clauses}^*(M \parallel F), \\ C(\gamma) \notin \text{clauses}(F). \end{cases}$$

Conflict

$$M \parallel F, H \triangleright C \quad \Longrightarrow M \parallel F, H \triangleright C \parallel \neg H \vee C \quad \text{if } M \models_P \neg C$$

Explain

$$M \parallel F \parallel C \vee \neg l \quad \Longrightarrow M \parallel F \parallel \neg H \vee D \vee C \quad \text{if } l_{H \triangleright (D \vee l)} \in M$$

Learn

$$M \parallel F \parallel C \quad \Longrightarrow M \parallel F, C \parallel C \quad \text{if } C \notin \text{clauses}(F)$$

Backjump

$$M \parallel l' \parallel M' \parallel F \parallel C \vee l \quad \Longrightarrow M \parallel_{C \vee l} \parallel F' \quad \text{if } \begin{cases} M \models_P \neg C, \\ l \text{ is undefined in } M, \\ F' = \left\{ \begin{array}{l} H \triangleright C \in F \\ H \cap \text{lits}(l' \parallel M') = \emptyset \end{array} \right\} \end{cases}$$

Unsat

$$M \parallel F \parallel \square \quad \Longrightarrow \text{unsat}$$

T-Propagate

$$M \parallel F \quad \Longrightarrow M \parallel_{(\neg l_1 \vee \dots \vee \neg l_n \vee l)} \parallel F \quad \text{if } \begin{cases} l \text{ is ground and occurs in } F, \\ l \text{ is undefined in } M, \\ l_1, \dots, l_n \in \text{lits}(M), \\ l_1, \dots, l_n \models_T l. \end{cases}$$

T-Conflict

$$M \parallel F \quad \Longrightarrow M \parallel F \parallel \neg l_1 \vee \dots \vee \neg l_n \quad \text{if } \begin{cases} l_1, \dots, l_n \in \text{lits}(M), \\ l_1, \dots, l_n \models_T \text{false}. \end{cases}$$

Fig. 2 Basic and theory propagation rules of $\text{DPLL}(\Gamma + \mathcal{T})$.

at least one decided literal, named l' in the rule definition in Figure 2. A typical choice is that l' be the least recently decided literal that satisfies the conditions of the rule. All hypothetical clauses $H \triangleright C$ which contain hypotheses that will be unassigned by the Backjump rule are deleted. Note that a learnt clause D may contain $\neg[C]$. In this case, the clause D is recording the context where guessing the clause C is unsound.

Figure 3 reproduces from [28] the *contraction transitions* that import the contraction rules of Γ in $\text{DPLL}(\Gamma + \mathcal{T})$: note that they apply only in search mode. These transitions and their explanation, that follows, refer to generic contraction rules schemas, and not to the concrete contraction rules of Figure 1, in order to show that this way of integrating contraction is general, and applies to the contraction rules of Figure 1 as well as to others. Any sound contraction inference taking a single premise (e.g., tautology deletion) can be easily incorporated into $\text{DPLL}(\Gamma + \mathcal{T})$. Given a hypothetical clause $H \triangleright C$, such a rule is just applied to C . Contraction rules with more than one premise need special treatment. We use $\gamma_d(C, C_2, \dots, C_m)$ to denote the application of a generic sound deletion rule

$$\frac{C, C_2, \dots, C_m}{C_2, \dots, C_m}$$

Delete	$M \parallel F, H \triangleright C \implies M \parallel F$	if $\begin{cases} \gamma_d(C, C_2, \dots, l_n), n \geq 2 \\ \text{level}(H) \geq \text{level}(H') \end{cases}$
Disable	$M \parallel F, H \triangleright C \implies M \parallel F, [H \triangleright C]_{\text{level}(H')}$	if $\begin{cases} \gamma_d(C, C_2, \dots, l_n), n \geq 2 \\ \text{level}(H) < \text{level}(H') \end{cases}$
Simplify	$M \parallel F, H \triangleright C \implies M \parallel F, (H \cup H') \triangleright C'$	if $\begin{cases} \gamma_s(C, C_2, \dots, l_n, C'), n \geq 2 \\ \text{level}(H) \geq \text{level}(H') \end{cases}$
Simplify-disable	$M \parallel F, H \triangleright C \implies M \parallel F, [H \triangleright C]_{\text{level}(H')}, (H \cup H') \triangleright C'$	if $\begin{cases} \gamma_s(C, C_2, \dots, l_n, C'), n \geq 2 \\ \text{level}(H) < \text{level}(H') \end{cases}$

Fig. 3 Contraction transitions of $\text{DPLL}(\Gamma + \mathcal{T})$: the notation in the conditions is explained in the text.

where a redundant main premise C is deleted, and $\gamma_s(C, C_2, \dots, C_m, C')$ to denote the application of a generic sound simplification rule

$$\frac{C, C_2, \dots, C_m}{C', C_2, \dots, C_m}$$

where a redundant main premise C is replaced by C' . $\text{DPLL}(\Gamma + \mathcal{T})$ assigns a *scope level* to each literal in M :

Definition 5 The *scope level* of a literal l , denoted $\text{level}(l)$, in $M \parallel M'$, is equal to the number of decided literals in $M \parallel l$. The *scope level* of a set of literals H is

$$\text{level}(H) = \begin{cases} \max\{\text{level}(l) \mid l \in H\} & \text{if } H \neq \emptyset, \\ 0 & \text{otherwise.} \end{cases}$$

A contraction inference γ from Γ is generalized to hypothetical clauses as follows: given main premise $H \triangleright C$, taken from $\text{ngclauses}(F)$, and side premises $H_2 \triangleright C_2, \dots, H_m \triangleright C_m, l_{m+1}, \dots, l_n$, taken from $\text{ngclauses}(F)$ and $\text{lits}_{\emptyset}(M)$, respectively, let $H' = H_2 \cup \dots \cup H_m \cup \{l_{m+1}, \dots, l_n\}$. Assume that γ has premises $C, C_2, \dots, C_m, l_{m+1}, \dots, l_n$. First, for a simplification γ_s , $H \triangleright C$ is replaced by $(H \cup H') \triangleright C'$. Second, for both γ_d and γ_s , $H \triangleright C$ is deleted only if $\text{level}(H) \geq \text{level}(H')$. Indeed, this condition prevents the situation where backjumping removes side premises (e.g., simplifiers or subsumers) before removing the main premise $H \triangleright C$ (i.e., the simplified or subsumed clause). Such a situation must be prevented, because otherwise the system would reach an unsound state, where $H \triangleright C$ was deleted, but the clauses that made it redundant and justified its deletion are no longer there. For this reason, if $\text{level}(H) < \text{level}(H')$, then $H \triangleright C$ is only *disabled*. In Figure 3 a disabled clause is surrounded by square brackets and bears as subscript the level of the set of side premises that disabled it. A disabled clause is not deleted, but it is not used as premise. When $\text{level}(H')$ is backjumped, all disabled clauses with subscript $\text{level}(H')$ will be re-enabled and will be available again as premises.

4.4 Refutational completeness of $\text{DPLL}(\Gamma + \mathcal{T})$

It was proved in [28] that $\text{DPLL}(\Gamma)$ is refutationally complete when \mathcal{T} is empty. We prove a stronger result for the case where both \mathcal{T} and \mathcal{R} are not empty. We start with definitions

that adapt to $DPLL(\Gamma + \mathcal{T})$ the classical notions of redundancy, fairness and saturation (cf. Section 2). We use Γ -based transitions for Deduce and the contraction transitions of Figure 3.

Definition 6 A Γ -based transition is *redundant* in state $M \parallel F$ if the corresponding Γ -inference is redundant in $clauses^*(M \parallel F)$.

Note that Γ -based transitions apply only in search mode.

Definition 7 A $DPLL(\Gamma + \mathcal{T})$ -derivation is *fair* if all applicable transitions are applied eventually, except SpeculativeIntro and redundant Γ -based transitions.

We recall that: (1) contraction rules are part of Γ ; (2) Γ , and therefore contraction rules, only sees $clauses^*(M \parallel F) = ngclauses(F) \cup lits_{\mathcal{R}}(M)$; (3) contraction inferences delete only clauses in $ngclauses(F)$. All other transitions do not use $ngclauses(F)$ and are therefore sheltered from contraction. Thus, the only transitions that are affected by contraction, and for which we need to stipulate that only irredundant inferences are considered, are Γ -based transitions.

Definition 8 A $DPLL(\Gamma + \mathcal{T})$ state is *saturated* if it is

- either *unsat*
- or a state $M \parallel F$ such that the only applicable transitions are SpeculativeIntro transitions or redundant Γ -based transitions.

Clearly, a fair derivation yields a saturated state eventually. In order to prove refutational completeness – whenever the input set S is unsatisfiable, $DPLL(\Gamma + \mathcal{T})$ reaches the *unsat* state – we prove as usual its contrapositive:

Theorem 4 *If the initial set of clauses $S = \mathcal{R} \uplus P$ is \mathcal{T} -smooth, and Γ has the reduction property for counterexamples, whenever $DPLL(\Gamma + \mathcal{T})$ reaches a saturated state $M \parallel F$, the input set S is satisfiable modulo \mathcal{T} .*

Proof We need to show that if $M \parallel F$ is saturated, then $clauses(F) \cup lits(M)$ is satisfiable. Satisfiability of S will follow, because the transition rules in $DPLL(\Gamma + \mathcal{T})$ are sound and therefore preserve satisfiability. Let N be $clauses(F) \cup lits(M)$. The set N has the form $\mathcal{R}' \uplus M_1 \uplus G_1 \uplus M_2 \uplus G_2$, where \mathcal{R}' contains non-ground clauses (i.e., $\mathcal{R}' = ngclauses(F)$), $M_1 \uplus G_1 \uplus M_2 \uplus G_2$ is ground, $M_1 \uplus M_2 = lits(M)$, $\mathcal{R}' \uplus G_1 \uplus G_2 = clauses(F)$, $G_1 \uplus M_1$ contains only \mathcal{R} -symbols (i.e., $M_1 = lits_{\mathcal{R}}(M)$), and $G_2 \uplus M_2$ contains only \mathcal{T} -symbols.

- We consider first $\mathcal{R}' \uplus G_1 \uplus M_1$. In a standard proof of completeness for Γ alone, we would have that $\mathcal{R}' \uplus G_1 \uplus M_1$ is Γ -saturated, because the Γ -derivation is fair. For $DPLL(\Gamma + \mathcal{T})$ we need to show that $\mathcal{R}' \uplus G_1 \uplus M_1$ is Γ -saturated, even if Γ -based transitions do not use G_1 . Since $M \parallel F$ is saturated, for every clause $C \in G_i$, for $i \in \{1, 2\}$, there is a literal l of C in M_i . Indeed, if this were not the case, the Decide rule could apply, violating the hypothesis that $M \parallel F$ is saturated. Thus, every clause $C \in G_1$ is subsumed by a literal in M_1 , therefore it is redundant in $M_1 \uplus G_1$, and every Γ -based transition that uses C is redundant. Then, $\mathcal{R}' \uplus M_1$ alone is Γ -saturated: if it were not, an irredundant Γ -based transition could apply, violating the hypothesis that $M \parallel F$ is saturated. It follows that $\mathcal{R}' \uplus G_1 \uplus M_1$ is Γ -saturated. Since $\mathcal{R}' \uplus G_1 \uplus M_1$ does not contain \square , and Γ has the reduction property for counterexamples, $\mathcal{R}' \uplus G_1 \uplus M_1$ is satisfiable by Theorem 1.

- We consider next $G_2 \uplus M_2$: this set is satisfiable modulo \mathcal{T} , because if it were not, the T-Conflict rule would apply, and $M \parallel F$ would not be saturated.
- By the hypothesis that the initial set $S = \mathcal{R} \uplus P$ is \mathcal{T} -smooth (cf. Definition 3), \mathcal{R} is variable-inactive. By Definition 2, \mathcal{R}' , which is derived from \mathcal{R} and ground unit \mathcal{R} -clauses, is also variable-inactive, hence stably infinite by Theorem 3. Thus, $\mathcal{R}' \uplus G_1 \uplus M_1$ has a model with infinite domain. Again by the hypothesis that $\mathcal{R} \uplus P$ is \mathcal{T} -smooth, \mathcal{T} is a union of stably infinite theories. Thus, $G_2 \uplus M_2$ has a \mathcal{T} -model with infinite domain. Since all the requirements for a Nelson-Oppen combination are fulfilled, these two models can be combined in a \mathcal{T} -model of N by the completeness of equality sharing, establishing that $clauses(F) \cup lits(M)$ is satisfiable. \square

All inference systems considered in the sequel have the reduction property for counterexamples [6,56]. The proof of Theorem 4 shows that the integration of the components in $DPLL(\Gamma + \mathcal{T})$ is designed in such a modular way that its completeness descends from the completeness of its components.

5 Towards decision procedures: $DPLL(\Gamma + \mathcal{T})$ -strategies

The combination of the transition system $DPLL(\Gamma + \mathcal{T})$ with a *search plan*, which controls the application of transition rules, yields a $DPLL(\Gamma + \mathcal{T})$ -strategy, or $DPLL(\Gamma + \mathcal{T})$ -procedure. Similar to $DPLL(\mathcal{T})$, a search plan for $DPLL(\Gamma + \mathcal{T})$ is a depth-first search plan. A standard way to ensure fairness with a depth-first search plan is *iterative deepening*. This section describes first a $DPLL(\Gamma + \mathcal{T})$ -procedure with iterative deepening, and then a way to use it with `SpeculativeIntro` to get decision procedures for smooth sets.

Definition 9 For all states $M \parallel F$, for all $C \in clauses(F)$, for all implied literals $l_C \in lits(M)$, and for all decided literals $l \in lits(M)$, the *inference depth* is given by

- $infDepth(C) = \begin{cases} 0 & \text{if } C \in F_0, \\ n + 1 & \text{if } C = C(\gamma) \text{ and} \\ & n = \max \{infDepth(D) \mid D \in P(\gamma)\} \text{ in a Deduce step,} \end{cases}$
- $infDepth(l_C) = infDepth(C)$ and
- $infDepth(l) = \min \{infDepth(D) \mid D \in clauses(F), l \in D\}$.

Informally, the inference depth of a clause indicates the depth of the inference tree that produced it; the inference depth of an implied literal is the inference depth of the clause that implied it; and the inference depth of a decided literal is the minimum inference depth of a clause that includes it.

In order to have a $DPLL(\Gamma + \mathcal{T})$ -procedure with iterative deepening, both rules susceptible of yielding infinitely many steps need to be bounded:

Definition 10 $DPLL(\Gamma + \mathcal{T})$ is $\langle k_d, k_u \rangle$ -bounded, for $k_d, k_u > 0$, if Deduce is restricted to premises C with $infDepth(C) < k_d$, and `SpeculativeIntro` can be applied at most k_u times.

This notion leads to termination:

Theorem 5 $\langle k_d, k_u \rangle$ -bounded $DPLL(\Gamma + \mathcal{T})$ is guaranteed to terminate for all initial sets of clauses $S = \mathcal{R} \uplus P$.

Proof By Definition 10, there are only finitely many applications of Deduce and `SpeculativeIntro`. The other $DPLL(\mathcal{T})$ transition rules are known to terminate (e.g., [55], cf. Theorems 2.10 and 3.7).

Definition 11 $\text{DPLL}(\Gamma + \mathcal{T})$ is *stuck* at k_d in state $M \parallel F$ if the only applicable transitions are SpeculativeIntro transitions and Deduce transitions involving premises C with $\text{infDepth}(C) \geq k_d$.

A $\text{DPLL}(\Gamma + \mathcal{T})$ -procedure with *iterative deepening*, abbreviated *ID-DPLL*($\Gamma + \mathcal{T}$)-procedure, is a $\text{DPLL}(\Gamma + \mathcal{T})$ -procedure where $\text{DPLL}(\Gamma + \mathcal{T})$ is $\langle k_d, k_u \rangle$ -bounded, and k_d and k_u are increased whenever $\text{DPLL}(\Gamma + \mathcal{T})$ gets stuck. The following example shows how fairness is not obvious without iterative deepening:

Example 1 Let Γ be an inference system with resolution and let F_0 include the following clauses:

- (1) $\neg p(x, y) \vee p(f(x), f(y)) \vee p(g(x), g(y))$,
- (2) $p(a, b)$,
- (3) $g(x) \not\simeq x$,
- (4) $g(c) \simeq c \vee g(d) \simeq d$.

Initially, Γ sees clauses (1) and (3), because they are in $\text{ngclauses}(F)$, while $\text{lits}_{\mathcal{R}}(M)$ is empty. If Decide adds $p(a, b)$ to M , Γ sees also (2) and may generate

$$p(f(a), f(b)) \vee p(g(a), g(b))$$

from (1) and (2) by resolution. If Decide adds $p(f(a), f(b))$ to M , and Γ generates

$$p(f(f(a)), f(f(b))) \vee p(g(f(a)), g(f(b))),$$

this alternation of decision and resolution steps may yield an infinite unfair derivation that does not detect the unsatisfiability of F_0 . Iterative deepening prevents this kind of behavior: when the depth of the clauses generated by resolution reaches the bound, further such steps are forbidden and the system is forced to consider steps with clauses of lower depth. When Decide adds to M first $g(c) \simeq c$ and then $g(d) \simeq d$ and each yields \square by resolution with $g(x) \not\simeq x$, inconsistency is detected.

Let S be a smooth set, and let U denote a sequence of “speculative axioms,” in the signature of S , that are introduced by SpeculativeIntro. In order to get a decision procedure, one needs to show that for some sequence U , there exist bounds k_d and k_u , such that $\langle k_d, k_u \rangle$ -bounded $\text{DPLL}(\Gamma + \mathcal{T})$ is guaranteed to terminate in the *unsat* state, whenever S is unsatisfiable, and in a state $M \parallel F$ such that $\text{DPLL}(\Gamma + \mathcal{T})$ is not *stuck* at k_d , whenever S is satisfiable; note that this means that $M \parallel F$ is saturated. The second example illustrates this idea:

Example 2 Let \mathcal{R} be

$$\{\neg(x \sqsubseteq y) \vee \neg(y \sqsubseteq z) \vee x \sqsubseteq z, \neg(x \sqsubseteq y) \vee f(x) \sqsubseteq f(y)\},$$

and P be

$$\{a \sqsubseteq b, a \sqsubseteq f(c), \neg(a \sqsubseteq c)\}.$$

Assume Γ features resolution, superposition and simplification. If SpeculativeIntro adds

$$\lceil f(x) \simeq x \rceil \triangleright f(x) \simeq x,$$

the monotonicity axiom is rewritten to a tautology and $a \sqsubseteq f(c)$ is also rewritten. Note that $\lceil f(x) \simeq x \rceil$ is a decision literal, and $\text{level}(\lceil f(x) \simeq x \rceil) = 1$. Thus, the rewriting steps only disable the monotonicity axiom and $a \sqsubseteq f(c)$, whose scope level is 0, and add $\lceil f(x) \simeq$

$x] \triangleright a \sqsubseteq c$ to F . Resolution generates the conflict clause $[f(x) \simeq x] \triangleright \square$. In conflict resolution mode, the literal $\neg[f(x) \simeq x]$ is added to M , preventing $\text{DPLL}(\Gamma + \mathcal{T})$ from guessing $f(x) \simeq x$ again. Next, if SpeculativeIntro adds

$$[f(f(x)) \simeq x] \triangleright f(f(x)) \simeq x,$$

monotonicity and $a \sqsubseteq b$ produce only $f(a) \sqsubseteq f(b)$, while monotonicity and $a \sqsubseteq f(c)$ produce only $f(a) \sqsubseteq f(f(c))$, which is disabled and replaced by $[f(f(x)) = x] \triangleright f(a) \sqsubseteq c$. Then, $\text{DPLL}(\Gamma + \mathcal{T})$ reaches a saturated state, and satisfiability is detected.

The third example shows a case where PropagateEq plays the key rôle:

Example 3 Let Γ have hyperresolution, superposition and simplification, \mathcal{T} be the theory of linear integer arithmetic, \mathcal{R} be

$$\{\neg(x \sqsubseteq y) \vee \neg(y \sqsubseteq z) \vee x \sqsubseteq z\}$$

and P be

$$\{a \sqsubseteq b_1, b_2 \sqsubseteq c, \neg(a \sqsubseteq c), b_1 \leq b_2, b_1 > b_2 - 1\}.$$

UnitPropagate adds the literals of P to M . In the model $\text{model}_{LA}(M)$ maintained by the linear arithmetic solver, $\text{model}_{LA}(M)(b_1) = \text{model}_{LA}(M)(b_2)$. Thus, PropagateEq guesses the equation $b_1 \simeq b_2$. Say $b_2 \succ b_1$ in the ordering \succ of Γ : simplification rewrites $b_2 \sqsubseteq c$ to $b_1 \sqsubseteq c$. Hyperresolution derives $a \sqsubseteq c$ from $a \sqsubseteq b_1, b_1 \sqsubseteq c$ and the transitivity axiom, so that an inconsistency is detected. $\text{DPLL}(\Gamma + \mathcal{T})$ backtracks and adds $\neg(b_1 \simeq b_2)$ to M . T-Conflict detects the inconsistency between this literal and $\{b_1 \leq b_2, b_1 > b_2 - 1\}$. The conflict resolution rules are applied again and the empty clause is produced.

6 Decision procedures for axiomatizations of type systems

In this section we study specific theories of interest for software verification and we obtain decision procedures for them.

Definition 12 A structure Φ is *essentially finite* with respect to a function symbol f if $\text{range}(\Phi(f))$ is finite.

Essential finiteness is weaker than finiteness, because it admits an infinite domain provided $\text{range}(\Phi(f))$ is finite.

Theorem 6 *If Φ is an essentially finite structure with respect to a monadic function symbol f , then there exist $k_1, k_2 \geq 0, k_1 \neq k_2$, such that $\Phi \models f^{k_1}(x) \simeq f^{k_2}(x)$.*

Proof For all $v \in |\Phi|$, we call *f-chain* starting at v , the sequence:

$$v = \Phi(f)^0(v), \Phi(f)^1(v), \Phi(f)^2(v), \dots, \Phi(f)^i(v), \dots$$

Since $\Phi(f)$ has finite range, there exist q_1, q_2 , with $q_1 \neq q_2$, such that $\Phi(f)^{q_1}(v) = \Phi(f)^{q_2}(v)$. Say that $q_1 > q_2$. Then we call *size*, denoted $\text{sz}(\Phi, f, v)$, and *prefix*, denoted $\text{pr}(\Phi, f, v)$, of the *f-chain* starting at v , the smallest q_1 and q_2 , respectively, such that $\Phi(f)^{q_1}(v) = \Phi(f)^{q_2}(v)$ and $q_1 > q_2$. We term *lasso*, denoted $\text{ls}(\Phi, f, v)$, of the *f-chain* starting at v , the difference between size and prefix, that is, $\text{ls}(\Phi, f, v) = \text{sz}(\Phi, f, v) - \text{pr}(\Phi, f, v)$. We say that $\Phi(f)^n(v)$ is *in the lasso* of the *f-chain* starting at v , if $n \geq \text{pr}(\Phi, f, v)$. Clearly,

for all elements w in the lasso of the f -chain starting at v , $\Phi(f)^n(w) = w$, when $n = \text{ls}(\Phi, f, v)$. Also, for all multiples of the lasso, that is, for all $n = j \cdot \text{ls}(\Phi, f, v)$ for some integer $j > 0$, $\Phi(f)^n(w) = w$. Let $q = \max\{\text{pr}(\Phi, f, v) \mid v \in \text{range}(\Phi(f))\} + 1$ and $n = \text{lcm}\{\text{ls}(\Phi, f, v) \mid v \in \text{range}(\Phi(f))\}$, where lcm abbreviates least common multiple. We claim that $\Phi \models f^{q+n}(x) \simeq f^q(x)$, that is, $k_1 = q + n$ and $k_2 = q$. By way of contradiction, assume that for some $v \in |\Phi|$, $\Phi(f)^{q+n}(v) \neq \Phi(f)^q(v)$. Take the f -chain starting at v : $\Phi(f)^q(v)$ is in the lasso of this chain, because $q > \text{pr}(\Phi, f, v)$. Since n is a multiple of $\text{ls}(\Phi, f, v)$, we have $\Phi(f)^{q+n}(v) = \Phi(f)^n(\Phi(f)^q(v)) = \Phi(f)^q(v)$, a contradiction. \square

Example 4 Let Φ be a structure such that $|\Phi| = \{v_0, v_1, v_2, \dots, v_9, \dots\}$, and let $\Phi(f)$ be the function defined by the following mapping: $\{v_0 \mapsto v_1, v_1 \mapsto v_2, v_2 \mapsto v_3, v_3 \mapsto v_4, v_4 \mapsto v_2, v_5 \mapsto v_6, v_6 \mapsto v_7, v_7 \mapsto v_8, v_8 \mapsto v_5, * \mapsto v_9\}$, where $*$ stands for any other element. The f -chain starting at v_0 has $\text{pr}(\Phi, f, v_0) = 2$, $\text{sz}(\Phi, f, v_0) = 5$ and $\text{ls}(\Phi, f, v_0) = 3$. The f -chain starting at v_5 has $\text{pr}(\Phi, f, v_5) = 0$, $\text{sz}(\Phi, f, v_5) = 4$ and $\text{ls}(\Phi, f, v_5) = 4$. Then, $q = \max\{2, 0\} + 1 = 3$, $n = \text{lcm}\{3, 4\} = 12$, $k_1 = q + n = 15$ and $k_2 = q = 3$, and $\Phi \models f^{15}(x) \simeq f^3(x)$.

To identify classes of problems for which an ID-DPLL($\Gamma + \mathcal{T}$)-procedure is a decision procedure, we focus on theories \mathcal{R} that satisfy the following property:

Definition 13 A presentation \mathcal{R} is *essentially finite* if its signature contains a single monadic function symbol f , and for all sets P of ground \mathcal{R} -clauses, such that $\mathcal{R} \uplus P$ is satisfiable, $\mathcal{R} \uplus P$ has an essentially finite model Φ with respect to f .

We show that ID-DPLL($\Gamma + \mathcal{T}$) is a decision procedure for essentially finite theories if the number of literals in clauses is bounded:

Theorem 7 Let \mathcal{R} be an essentially finite presentation. Consider an ID-DPLL($\Gamma + \mathcal{T}$)-procedure where every SpeculativeIntro transition adds an equation $f^j(x) \simeq f^k(x)$ with $j > k$, for increasing values of j and k . If there exists an n such that no clause generated by DPLL($\Gamma + \mathcal{T}$) contains more than n literals, ID-DPLL($\Gamma + \mathcal{T}$) is a decision procedure for the satisfiability modulo \mathcal{T} of \mathcal{T} -smooth problems $\mathcal{R} \uplus P$.

Proof If $\mathcal{R} \uplus P$ is unsatisfiable, then, by refutational completeness, DPLL($\Gamma + \mathcal{T}$) will reach the state *unsat* when k_d becomes large enough. If $\mathcal{R} \uplus P$ is satisfiable, it has an essentially finite model Φ , because \mathcal{R} is essentially finite. Choose k_u large enough that the axiom $f^{k_1}(x) \simeq f^{k_2}(x)$ satisfied by Φ according to Theorem 6 is added by SpeculativeIntro. We need to prove that if k_d is large enough, DPLL($\Gamma + \mathcal{T}$) will not get stuck at k_d . To do that, we prove that only a finite number of clauses are generated for unbounded k_d for the chosen k_u . Say that $k_1 > k_2$: the axiom $f^{k_1}(x) \simeq f^{k_2}(x)$ is applied as a rewrite rule $f^{k_1}(x) \rightarrow f^{k_2}(x)$ to simplify⁵ all clauses that contain a term $f^k(t)$ with $k > k_1$. This guarantees that no such term will be kept and that the depth of terms in clauses is bounded. Since the number of literals in clauses is also bounded by the hypothesis that no clause can contain more than n literals, only a finite number of clauses can be derived for unbounded k_d . Thus, DPLL($\Gamma + \mathcal{T}$) will halt without getting stuck and will detect satisfiability. \square

From now on, unless otherwise stated, Γ is superposition with negative selection, hyperresolution, factoring and simplification.

⁵ Of course, this assumes that Γ features simplification.

Lemma 2 *If \mathcal{R} is Horn, the number of literals in clauses generated by $DPLL(\Gamma + \mathcal{T})$ from a \mathcal{T} -smooth $\mathcal{R} \uplus P$ is bounded.*

Proof In the Horn case, superposition is unit superposition, which does not increase the number of literals, and hyperresolution only generates positive unit clauses. \square

If \mathcal{R} is a set of non-equational clauses with no more than two literals each, and Γ is resolution, factoring and simplification (to apply $f^{k_1}(x) \rightarrow f^{k_2}(x)$), then all generated clauses contain at most two literals. To give further examples, we need the following:

Definition 14 A clause $C = \neg l_1 \vee \dots \vee \neg l_n \vee l_{n+1} \vee \dots \vee l_{n+q}$ is *ground-preserving* if

$$\bigcup_{j=n+1}^{n+q} \text{Var}(l_j) \subseteq \bigcup_{j=1}^n \text{Var}(l_j).$$

A set of clauses is *ground-preserving* if all its clauses are.

In a ground-preserving set the only positive clauses are ground.⁶

Lemma 3 *If \mathcal{R} is essentially finite and ground-preserving, and every SpeculativeIntro transition adds an equation $f^j(x) \simeq f^k(x)$ with $j > k$, for increasing values of j and k , $DPLL(\Gamma + \mathcal{T})$ generates finitely many clauses from a \mathcal{T} -smooth $\mathcal{R} \uplus P$.*

Proof Hyperresolution only generates positive ground clauses, because all variables get instantiated by resolving the negative literals with positive clauses. Superposition with negative selection superposes a ground positive clause into a ground-preserving clause, which generates either a ground clause, or a non-ground ground-preserving clause with no more variable positions than its non-ground parent. It follows that superposition creates no new non-ground term, and only finitely many non-ground ground-preserving clauses can be derived. Since term depth is limited by simplification by $f^{k_1}(x) \rightarrow f^{k_2}(x)$, only finitely many ground clauses can be generated. \square

Next, we consider some specific theories relevant to the axiomatization of type systems in programming languages. Given the axioms

$$\text{Reflexivity } x \sqsubseteq x \tag{1}$$

$$\text{Transitivity } \neg(x \sqsubseteq y) \vee \neg(y \sqsubseteq z) \vee x \sqsubseteq z \tag{2}$$

$$\text{Anti-Symmetry } \neg(x \sqsubseteq y) \vee \neg(y \sqsubseteq x) \vee x \simeq y \tag{3}$$

$$\text{Monotonicity } \neg(x \sqsubseteq y) \vee f(x) \sqsubseteq f(y) \tag{4}$$

$$\text{Tree-Property } \neg(z \sqsubseteq x) \vee \neg(z \sqsubseteq y) \vee x \sqsubseteq y \vee y \sqsubseteq x \tag{5}$$

$\{(1), (2), (3)\}$ presents a poset (partially ordered set), $MI = \{(1), (2), (3), (4)\}$ a type system with *multiple inheritance*, and $SI = MI \uplus \{(5)\}$ a type system with *single inheritance*, where \sqsubseteq is the subtype relationship and f is a type constructor. MI and SI are essentially finite, because they satisfy a stronger property:

⁶ Definition 14 is a weakening of that of positive variable dominated clause of [22] (cf. Definition 3.18), and it is dual to that of ground-preserving clause of [47], which required that negative literals do not contain variables that do not appear in positive ones. Our definition is for a positive strategy in the non-Horn case, hence forward reasoning, whereas that of [47] was for linear input proofs in the Horn case, hence backward reasoning.

Definition 15 \mathcal{R} has the *finite model property*, if for all sets P of ground \mathcal{R} -clauses, such that $\mathcal{R} \uplus P$ is satisfiable, $\mathcal{R} \uplus P$ has a model Φ with finite $|\Phi|$.

Theorem 8 *SI has the finite model property hence it is essentially finite.*

Proof Assume $SI \uplus P$ is satisfiable, and let Φ be a model for it. It is sufficient to show there is a finite model Φ' . Let T_P be the set of subterms of terms in P , and V_P be the set $\Phi(T_P)$. Since P is finite and ground, V_P is finite. Let $|\Phi'|$ be $V_P \cup \{v_m\}$, where v_m is an element not in V_P . Then, we define $\Phi'(\sqsubseteq)(v_1, v_2)$ as:

$$v_2 = v_m \text{ or } (v_1, v_2) \in \Phi(\sqsubseteq).$$

Intuitively, v_m is a new maximal element. $\langle |\Phi'|, \Phi'(\sqsubseteq) \rangle$ is a poset and $\Phi'(\sqsubseteq)$ satisfies the Tree-Property. Now, we define an auxiliary function $g: |\Phi'| \rightarrow |\Phi'|$ as:

$$g(v) = \begin{cases} \Phi(f)(v) & \text{if } f(t) \in T_P, \text{ and } \Phi(t) = v; \\ v_m & \text{otherwise.} \end{cases}$$

Let dom_f , the relevant domain of f , be the set $\{\Phi(t) \mid f(t) \in T_P\} \cup \{v_m\}$. With a small abuse of notation, we use $v \sqsubseteq w$ to denote $(v, w) \in \Phi'(\sqsubseteq)$. Then, we define $\Phi'(f)(v)$ as $g(w)$, where w is an element in $|\Phi'|$ such that $v \sqsubseteq w$, $w \in \text{dom}_f$, and for all $w', v \sqsubseteq w'$ and $w' \in \text{dom}_f$ imply $w \sqsubseteq w'$. This function is well defined because $v_m \in \text{dom}_f$, v_m is the maximal element of $|\Phi'|$, and $\Phi'(\sqsubseteq)$ satisfies the Tree-Property, which ensures uniqueness of the image. Moreover, $\Phi'(f)$ is monotonic with respect to $\Phi'(\sqsubseteq)$. \square

Definition 16 Let $\langle A, \sqsubseteq \rangle$ be a poset. The *Dedekind-MacNeille completion* [50] of $\langle A, \sqsubseteq \rangle$ is the unique complete lattice $\langle B, \preceq \rangle$ satisfying the following properties:

- There is an injection α from A to B such that: $v_1 \sqsubseteq v_2$ iff $\alpha(v_1) \preceq \alpha(v_2)$,
- Every subset of B has greatest and least lower bound, and
- B is finite if A is finite. Actually, B is a subset of 2^A .

Theorem 9 *MI has the finite model property hence it is essentially finite.*

Proof The construction used for SI does not work for MI, because without the Tree-Property the w in the definition of $\Phi'(f)(v)$ may not be unique for a given v . First, we define an auxiliary structure Φ_0 such that $|\Phi_0| = V_P$, $\Phi_0(\sqsubseteq) = \Phi(\sqsubseteq)|_{V_P}$, and $\Phi_0(f)$ is defined as:

$$\Phi_0(f)(v) = \begin{cases} \Phi(f)(v) & \text{if } f(t) \in T_P, \text{ and } \Phi(t) = v, \\ w & \text{otherwise,} \end{cases}$$

where w is some element of V_P . Note that $\langle V_P, \Phi_0(\sqsubseteq) \rangle$ is a poset. Let dom_f be the set $\{\Phi(t) \mid f(t) \in T_P\}$. Then, following [23] we use the Dedekind-MacNeille completion to complete $\langle V_P, \Phi_0(\sqsubseteq) \rangle$ into a complete lattice $\langle B, \preceq \rangle$. We use $\text{glb}(S)$ to denote the greatest lower bound of a subset S of B . Now, we define a finite model Φ' for $MI \uplus P$ with domain $|\Phi'| = B$, in the following way:

$$\begin{aligned} \Phi'(c) &= \alpha(\Phi_0(c)) \quad \text{for every constant } c \text{ in } T_P, \\ \Phi'(\sqsubseteq) &= \preceq, \\ \Phi'(f)(v) &= \text{glb}(\{\alpha(\Phi_0(f)(w)) \mid w \in V_P, w \in \text{dom}_f, v \preceq \alpha(w)\}). \end{aligned}$$

The function $\Phi'(f)$ is monotonic with respect to $\Phi'(\sqsubseteq)$. The structure Φ' satisfies P because for every term t in T_P , we have $\Phi'(t) = \alpha(\Phi(t))$. Moreover, the \sqsubseteq -literals in P are satisfied because the lattice $\langle B, \preceq \rangle$ is a Dedekind-MacNeille completion of Φ_0 which is a restriction of Φ . \square

Now we show that $ID\text{-}DPLL(\Gamma + \mathcal{T})$ is a decision procedure for MI and SI.

Theorem 10 *An $ID\text{-}DPLL(\Gamma + \mathcal{T})$ -procedure where every SpeculativeIntro transition adds an equation $f^j(x) \simeq f^k(x)$ with $j > k$, for increasing values of j and k , is a decision procedure for the satisfiability modulo \mathcal{T} of \mathcal{T} -smooth problems $MI \uplus P$.*

Proof It follows from Theorem 7 and Lemma 2, because MI is essentially finite and Horn. \square

Theorem 11 *An $ID\text{-}DPLL(\Gamma + \mathcal{T})$ -procedure where every SpeculativeIntro transition adds an equation $f^j(x) \simeq f^k(x)$ with $j > k$, for increasing values of j and k , is a decision procedure for the satisfiability modulo \mathcal{T} of \mathcal{T} -smooth problems $SI \uplus P$.*

Proof Since SI is essentially finite and ground-preserving, except for Reflexivity, it follows from Theorem 7 and Lemma 3, provided Reflexivity does not affect the result of Lemma 3. This is the case, since a hyperresolution involving Reflexivity generates either a tautology or a subsumed clause or a ground clause. \square

In Spec\# [7], the axiomatization of the type system also includes $TR = \{g(x) \not\approx null, h(g(x)) \simeq x\}$, where g represents the *type representative* of some type. The first axiom states that the representative is never the constant *null*, which means *null* has no pre-image, hence g is not surjective. The second axiom states that g has a left inverse, hence it is injective. It is well-known that a set with an injective but not surjective function is infinite (e.g., Lemma 1 in [24]), so that any model of TR is infinite.

Theorem 12 *An $ID\text{-}DPLL(\Gamma + \mathcal{T})$ -procedure where every SpeculativeIntro transition adds an equation $f^j(x) \simeq f^k(x)$ with $j > k$, for increasing values of j and k , is a decision procedure for the satisfiability modulo \mathcal{T} of \mathcal{T} -smooth problems $MI \uplus TR \uplus P$ and $SI \uplus TR \uplus P$.*

Proof Superposition applied to an axiom in TR and a ground equation generates a ground equation smaller than its ground parent in the \succ ordering (e.g., under a precedence $g \succ h \succ null$), so that Γ terminates on TR-satisfiability problems. Since MI (or SI) and TR are disjoint and variable-inactive, Γ terminates also on satisfiability problems in $MI \uplus TR$ (or $SI \uplus TR$) by Theorem 2. Thus, the combination with TR does not change that only finitely many clauses can be generated. The claim follows from Theorem 10 and this observation for problems $MI \uplus TR \uplus P$, and from Theorem 11 and this observation for problems $SI \uplus TR \uplus P$. \square

7 Discussion

The $DPLL(\Gamma + \mathcal{T})$ system integrates $DPLL(\mathcal{T})$ with a first-order engine Γ , to combine the strengths of $DPLL$ and efficient solvers for special theories, such as linear arithmetic, with those of superposition and resolution. $DPLL(\mathcal{T})$ -based SMT-solvers and general theorem provers grew independently for several years. The increasing recognition that their features are complementary, and necessary to solve frontier problems in fields such as program verification, is leading to study their interaction.

The rewrite-based approach to satisfiability procedures developed in [5, 3, 18, 15, 4] was concerned with using first-order engines as decision procedures for satisfiability problems. In [14, 16] it was generalized from satisfiability problems, given by sets of ground unit clauses, to decision problems given by sets of ground clauses. The first-order engine alone

was used as decision procedure, with no integration with an SMT-solver. The two-stage method of [13, 17] lets a first-order engine and an SMT-solver cooperate, including allowing both a union \mathcal{R} of variable-inactive axiomatized theories, and a union \mathcal{T} of Nelson-Oppen built-in theories: the first-order engine was applied as a pre-processor to compile \mathcal{R} and reduce it to a theory that DPLL(\mathcal{T}) alone could handle. Thus, the two reasoners were applied in sequence. The systems in [46, 1] explored embedding a \mathcal{T} -solver for linear arithmetic into a superposition-based theorem prover, while the study of *iterated schemata* in [2] offers another perspective on the integration of propositional and arithmetical reasoning.

In DPLL($\Gamma + \mathcal{T}$) the first-order engine Γ is tightly integrated within DPLL(\mathcal{T}), resulting in one single system. This is a main difference with respect to the two-stage approach of [13, 17]. DPLL($\Gamma + \mathcal{T}$) and systems such as those in [46, 1] can be considered symmetric: in DPLL($\Gamma + \mathcal{T}$) the superposition-based engine Γ is a satellite of DPLL(\mathcal{T}); in [46, 1] the \mathcal{T} -solver is a satellite of the superposition-based engine. A first version of DPLL($\Gamma + \mathcal{T}$) appeared in [28]. It was called DPLL(Γ), because it was known to be refutationally complete only in the case where the background theory \mathcal{T} is empty. A first contribution of this article was to advance the DPLL($\Gamma + \mathcal{T}$) approach by giving conditions under which it is refutationally complete when both \mathcal{R} and \mathcal{T} are not empty.

Combination of theories is of paramount importance to reason about software. In previous work, it was known how to combine built-in theories, according to the equality sharing method pioneered by Nelson and Oppen [53], and studied since then by many authors (e.g., [37, 54, 65] for some of the most recent extensions). This style of combination requires one to embed in the prover a decision procedure for each theory of interest. While decision procedures are available for several theories, it might not be the case for each and every group of axioms that may appear in program checking problems. This is a main reason why we need an axiomatized theory \mathcal{R} and Γ to reason about it. A second main contribution of this article was to show how to let combination of built-in theories à la Nelson-Oppen, and union of axiomatized theories under variable inactivity, coexist and work together in DPLL($\Gamma + \mathcal{T}$).

We presented a new DPLL($\Gamma + \mathcal{T}$) system that combines DPLL($\Gamma + \mathcal{T}$) with speculative inferences. The purpose is to enforce termination by introducing additional axioms as hypotheses. This idea was inspired by the “unsound theorem proving” concept of [49]. The additional axioms may cause unsoundness, by making unsatisfiable what was a satisfiable set. We provided a mechanism for the prover to detect any unsoundness introduced by the added axioms and recover from it. This mechanism is based on the backtracking scheme that is native of a DPLL search: an inconsistency due to a speculative inference is an “*unnatural failure*” that the prover treats like a “*natural failure*” (a proper inconsistency) by backtracking. Furthermore, it keeps memory of the failure to avoid repeating it. An idea of speculative inferences may be implicit in bottom-up model generation approaches [9]. In those contexts the speculation consists of trying in turn each case in a case analysis. In our method, the speculative inferences assert additional clauses, on top of the native case analysis of DPLL on the input clauses. Considering each of the two horns of a case analysis can be seen either as an inference step or as a search step on existing data. Our speculative inferences are more like guessing additional features that a model may satisfy.

DPLL($\Gamma + \mathcal{T}$) equipped with an iterative deepening search plan forms an *ID-DPLL*($\Gamma + \mathcal{T}$)-*procedure*. We showed that ID-DPLL($\Gamma + \mathcal{T}$) with speculative inferences is a decision procedure for theories that axiomatize type systems relevant to program checking. Their crucial feature is that they are *essentially finite*: they have one unary function symbol whose range is finite. However, we gave examples where ID-DPLL($\Gamma + \mathcal{T}$) is a decision procedure also when more function symbols are involved via combination of theories. Another way to approach the axiomatizations in Section 6 is *locality*, proposed for Horn theories [40], and

then extended beyond the Horn case and developed in [8,63,64,43,44]. In a local theory, validity of a conjecture can be decided by considering only finitely many of its ground instances. We emphasize that $\text{DPLL}(\Gamma + \mathcal{T})$ yields a decision procedure for the axiomatizations in Section 6 united with an arbitrary built-in \mathcal{T} , provided the problem $\mathcal{T} \uplus \mathcal{R} \uplus P$ is \mathcal{T} -smooth. In applications, there is no guarantee that all relevant instances of $\mathcal{T} \uplus \mathcal{R}$ will be local. Thus, speculative inferences and locality can be considered complementary.

There are several directions for future work. One is to extend the approach to more presentations, including cases where the signature of \mathcal{R} features also non-monadic function symbols. For example, consider the axiom $y \sqsubseteq x \wedge u \sqsubseteq z \Rightarrow \text{map}(x, u) \sqsubseteq \text{map}(y, z)$, where \sqsubseteq is a subtype relation, and $\text{map}(x, u)$ represents the type of maps from type x to type u . If $y \sqsubseteq x$, a value of type y can be used whenever a value of type x is used. For maps $f \in \text{map}(x, u)$ and $g \in \text{map}(y, z)$ this is the case if $y \sqsubseteq x$, which means f can take whatever g takes, and $u \sqsubseteq z$, which means whatever f yields is within what g yields. Such an axiom with a dyadic function symbol may be useful for an axiomatization of maps. Another open issue is the duplication of reasoning on ground unit equational clauses in $\text{DPLL}(\Gamma + \mathcal{T})$, due to the fact that they are seen by both Γ and the congruence closure (CC) algorithm within $\text{DPLL}(\mathcal{T})$. Using the CC algorithm to compute the completion of the set of ground equations [38, 62], and pass the resulting canonical system to Γ , would not solve the problem, because this solution is not *incremental*, as the addition of a single ground equation requires recomputing the canonical system. It would be ideal to automate the choice of clauses to be added by *SpeculativeIntro*. However, this manual component of $\text{DPLL}(\Gamma + \mathcal{T})$ is at a higher level of abstraction than triggering, and it is certainly not heavier.

Another topic for future investigation is how to improve the capability of the system to discover unsatisfiability due to the lack of finite models. By detecting the generation of a cardinality constraint by Γ , $\text{DPLL}(\Gamma + \mathcal{T})$ can discover that an axiomatized theory \mathcal{R} is not variable-inactive and not stably infinite. In other words, it can discover the lack of infinite models. On the other hand, it does not have a general way to discover the lack of finite models, or that there are only models with the “wrong” cardinality: for example, \mathcal{R} only features models with a certain finite cardinality, when \mathcal{T} requires a different finite cardinality.

The class of formulæ that can be decided by $\text{DPLL}(\Gamma + \mathcal{T})$ includes axiomatizations of type systems, used in tools such as ESC/Java [36] and Spec# [7], which represents significant evidence of the relevance of this work to applications.

Acknowledgements Part of this research initiated during a visit of the first author with the Software Reliability Group of Microsoft Research in Redmond. We thank the anonymous reviewers whose suggestions allowed us to improve an earlier version of this article.

References

1. Ernst Althaus, Evgeny Kruglov, and Christoph Weidenbach. Superposition modulo linear arithmetic SUP(LA). In Silvio Ghilardi and Roberto Sebastiani, editors, *Proceedings of the Seventh Symposium on Frontiers of Combining Systems (FroCoS)*, volume 5749 of *Lecture Notes in Artificial Intelligence*, pages 84–99. Springer, 2009.
2. Vincent Aravantinos, Ricardo Caferra, and Nicolas Peltier. A decidable class of nested iterated schemata. In Jürgen Giesl and Reiner Hähnle, editors, *Proceedings of the Fifth International Joint Conference on Automated Reasoning (IJCAR)*, volume 6173 of *Lecture Notes in Artificial Intelligence*, pages 293–308. Springer, 2010.
3. Alessandro Armando, Maria Paola Bonacina, Silvio Ranise, and Stephan Schulz. On a rewriting approach to satisfiability procedures: extension, combination of theories and an experimental appraisal.

- In Bernhard Gramlich, editor, *Proceedings of the Fifth Workshop on Frontiers of Combining Systems (FroCoS)*, volume 3717 of *Lecture Notes in Artificial Intelligence*, pages 65–80. Springer, 2005.
4. Alessandro Armando, Maria Paola Bonacina, Silvio Ranise, and Stephan Schulz. New results on rewrite-based satisfiability procedures. *ACM Transactions on Computational Logic*, 10(1):129–179, 2009.
 5. Alessandro Armando, Silvio Ranise, and Michaël Rusinowitch. A rewriting approach to satisfiability procedures. *Information and Computation*, 183(2):140–164, 2003.
 6. Leo Bachmair and Harald Ganzinger. Rewrite-based equational theorem proving with selection and simplification. *Journal of Logic and Computation*, 4(3):217–247, 1994.
 7. Michael Barnett, K. Rustan M. Leino, and Wolfram Schulte. The Spec# programming system: An overview. In Gilles Barthe, Lilian Burdy, Marieke Huisman, Jean-Louis Lanet, and Traian Muntean, editors, *Proceedings of the Workshop on Construction and Analysis of Safe, Secure, and Interoperable Smart Devices (CASSIS 2004)*, volume 3362 of *Lecture Notes in Computer Science*, pages 49–69. Springer, 2005.
 8. David A. Basin and Harald Ganzinger. Automated complexity analysis based on ordered resolution. *Journal of the ACM*, 48(1):70–109, 2001.
 9. Peter Baumgartner and Renate A. Schmidt. Blocking and other enhancements for bottom-up model generation methods. In Ulrich Furbach and Natarajan Shankar, editors, *Proceedings of the Third International Joint Conference on Automated Reasoning (IJCAR)*, volume 4130 of *Lecture Notes in Artificial Intelligence*, pages 125–139. Springer, 2006.
 10. Maria Paola Bonacina. A taxonomy of theorem-proving strategies. In Michael J. Wooldridge and Manuela Veloso, editors, *Artificial Intelligence Today – Recent Trends and Developments*, volume 1600 of *Lecture Notes in Artificial Intelligence*, pages 43–84. Springer, 1999.
 11. Maria Paola Bonacina. On theorem proving for program checking – historical perspective and recent developments. In Maribel Fernandez, editor, *Proceedings of the Twelfth International Symposium on Principles and Practice of Declarative Programming (PPDP)*, pages 1–11. ACM Press, July 2010.
 12. Maria Paola Bonacina and Nachum Dershowitz. Abstract canonical inference. *ACM Transactions on Computational Logic*, 8(1):180–208, 2007.
 13. Maria Paola Bonacina and Mnacho Echenim. \mathcal{T} -decision by decomposition. In Frank Pfenning, editor, *Proceedings of the Twenty-first Conference on Automated Deduction (CADE)*, volume 4603 of *Lecture Notes in Artificial Intelligence*, pages 199–214. Springer, 2007.
 14. Maria Paola Bonacina and Mnacho Echenim. Rewrite-based decision procedures. In Myla Archer, Thierry Boy de la Tour, and Cesar Munoz, editors, *Proceedings of the Sixth Workshop on Strategies in Automated Deduction (STRATEGIES), Federated Logic Conference 2006*, volume 174(11) of *Electronic Notes in Theoretical Computer Science*, pages 27–45. Elsevier, 2007.
 15. Maria Paola Bonacina and Mnacho Echenim. Rewrite-based satisfiability procedures for recursive data structures. In Byron Cook and Roberto Sebastiani, editors, *Proceedings of the Fourth Workshop on Pragmatics of Decision Procedures in Automated Reasoning (PDPAR), Federated Logic Conference 2006*, volume 174(8) of *Electronic Notes in Theoretical Computer Science*, pages 55–70. Elsevier, 2007.
 16. Maria Paola Bonacina and Mnacho Echenim. On variable-inactivity and polynomial T-satisfiability procedures. *Journal of Logic and Computation*, 18(1):77–96, 2008.
 17. Maria Paola Bonacina and Mnacho Echenim. Theory decision by decomposition. *Journal of Symbolic Computation*, 45(2):229–260, 2010.
 18. Maria Paola Bonacina, Silvio Ghilardi, Enrica Nicolini, Silvio Ranise, and Daniele Zucchelli. Decidability and undecidability results for Nelson-Oppen and rewrite-based decision procedures. In Ulrich Furbach and Natarajan Shankar, editors, *Proceedings of the Third International Joint Conference on Automated Reasoning (IJCAR)*, volume 4130 of *Lecture Notes in Artificial Intelligence*, pages 513–527. Springer, 2006.
 19. Maria Paola Bonacina and Jieh Hsiang. Towards a foundation of completion procedures as semidecision procedures. *Theoretical Computer Science*, 146:199–242, 1995.
 20. Maria Paola Bonacina, Christopher A. Lynch, and Leonardo de Moura. On deciding satisfiability by $DPLL(\Gamma + \mathcal{T})$ and unsound theorem proving. In Renate Schmidt, editor, *Proceedings of the Twenty-second Conference on Automated Deduction (CADE)*, volume 5663 of *Lecture Notes in Artificial Intelligence*, pages 35–50. Springer, 2009.
 21. Aaron R. Bradley, Zohar Manna, and Henny B. Sipma. What’s decidable about arrays? In E. Allen Emerson and Kedar S. Namjoshi, editors, *Proceedings of the Seventh Conference on Verification, Model Checking, and Abstract Interpretation (VMCAI)*, volume 3855 of *Lecture Notes in Computer Science*, pages 427–442. Springer, 2006.
 22. Ricardo Caferra, Alexander Leitsch, and Nicolas Peltier. *Automated Model Building*. Kluwer Academic Publishers, Amsterdam, The Netherlands, 2004.
 23. Domenico Cantone and Calogero G. Zarba. A decision procedure for monotone functions over bounded and complete lattices. In Harrie de Swart, editor, *Proc. TARSKI II*, volume 4342 of *Lecture Notes in Artificial Intelligence*, pages 318–333. Springer, 2006.

24. Koen Claessen and Ann Lillieström. Automated inference of finite unsatisfiability. In Renate Schmidt, editor, *Proceedings of the Twenty-second Conference on Automated Deduction (CADE)*, volume 5663 of *Lecture Notes in Artificial Intelligence*, pages 388–403. Springer, 2009.
25. Martin Davis, George Logemann, and Donald Loveland. A machine program for theorem-proving. *Communications of the ACM*, 5(7):394–397, 1962.
26. Martin Davis and Hilary Putnam. A computing procedure for quantification theory. *Journal of the ACM*, 7:201–215, 1960.
27. Leonardo de Moura and Nikolaj Bjørner. Efficient E-matching for SMT-solvers. In Frank Pfenning, editor, *Proceedings of the Twenty-first Conference on Automated Deduction (CADE)*, volume 4603 of *Lecture Notes in Artificial Intelligence*, pages 183–198. Springer, 2007.
28. Leonardo de Moura and Nikolaj Bjørner. Engineering DPLL(T) + saturation. In Alessandro Armando, Peter Baumgartner, and Gilles Dowek, editors, *Proceedings of the Fourth International Joint Conference on Automated Reasoning (IJCAR)*, volume 5195 of *Lecture Notes in Artificial Intelligence*, pages 475–490. Springer, 2008.
29. Leonardo de Moura and Nikolaj Bjørner. Model-based theory combination. In Sava Krstić and Albert Oliveras, editors, *Proceedings of the Fifth Workshop on Satisfiability Modulo Theories (SMT), Conference on Automated Verification 2007*, volume 198(2) of *Electronic Notes in Theoretical Computer Science*, pages 37–49. Elsevier, 2008.
30. Leonardo de Moura and Nikolaj Bjørner. Z3: an efficient SMT solver. In C. R. Ramakrishnan and Jakob Rehof, editors, *Proceedings of the Fourteenth Conference on Tools and algorithms for the construction and analysis of systems (TACAS)*, volume 4963 of *Lecture Notes in Computer Science*, pages 337–340. Springer, 2008.
31. Nachum Dershowitz. Orderings for term-rewriting systems. *Theoretical Computer Science*, 17(3):279–301, 1982.
32. Nachum Dershowitz. A maximal-literal unit strategy for Horn clauses. In Stephan Kaplan and Mitsuhiro Okada, editors, *Proceedings of the Second Workshop on Conditional and Typed Term Rewriting Systems (CTRS 1990)*, volume 516 of *Lecture Notes in Computer Science*, pages 14–25. Springer, 1991.
33. Nachum Dershowitz and Zohar Manna. Proving termination with multiset orderings. *Communications of the ACM*, 22(8):465–476, 1979.
34. David L. Detlefs, Greg Nelson, and James B. Saxe. Simplify: a theorem prover for program checking. *Journal of the ACM*, 52(3):365–473, 2005.
35. Bruno Dutertre and Leonardo de Moura. A fast linear-arithmetic solver for DPLL(T). In Tom Ball and R. B. Jones, editors, *Proceedings of the Eighteenth Conference on Automated Verification (CAV)*, volume 4144 of *Lecture Notes in Computer Science*, pages 81–94. Springer, 2006.
36. Cormac Flanagan, K. Rustan M. Leino, Mark Lillibridge, Greg Nelson, James B. Saxe, and Raymie Stata. Extended static checking for Java. In Laurie J. Hendren, editor, *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, pages 234–245, 2002.
37. Pascal Fontaine. Combinations of theories for decidable fragments of first-order logic. In Silvio Ghilardi and Roberto Sebastiani, editors, *Proceedings of the Seventh Symposium on Frontiers of Combining Systems (FroCoS)*, volume 5749 of *Lecture Notes in Artificial Intelligence*, pages 263–278. Springer, 2009.
38. Jean Gallier, Paliath Narendran, David A. Plaisted, Stan Raatz, and Wayne Snyder. Finding canonical rewriting systems equivalent to a finite set of ground equations in polynomial time. *Journal of the ACM*, 40(1):1–16, 1993.
39. Yeting Ge, Clark Barrett, and Cesare Tinelli. Solving quantified verification conditions using satisfiability modulo theories. In Frank Pfenning, editor, *Proceedings of the Twenty-first Conference on Automated Deduction (CADE)*, volume 4603 of *Lecture Notes in Artificial Intelligence*, pages 167–182. Springer, 2007.
40. Robert Givan and David A. McAllester. Polynomial-time computation via local inference relations. *ACM Transactions on Computational Logic*, 3(4):521–541, 2002.
41. Joseph Y. Halpern. Presburger arithmetic with unary predicates is π_1^1 complete. *Journal of Symbolic Logic*, 56:637–642, 1991.
42. Jieh Hsiang and Michaël Rusinowitch. Proving refutational completeness of theorem proving strategies: the transfinite semantic tree method. *Journal of the ACM*, 38(3):559–587, 1991.
43. Carsten Ihlemann, Swen Jacobs, and Viorica Sofronie-Stokkermans. On local reasoning in verification. In C. R. Ramakrishnan and Jakob Rehof, editors, *Proceedings of the Fourteenth Conference on Tools and algorithms for the construction and analysis of systems (TACAS)*, volume 4963 of *Lecture Notes in Computer Science*, pages 265–281. Springer, 2008.
44. Swen Jacobs. Incremental instance generation in local reasoning. In Franz Baader, Silvio Ghilardi, Miki Hermann, Ulrike Sattler, and Viorica Sofronie-Stokkermans, editors, *Notes of the First Workshop on Complexity, Expressibility and Decidability (CEDAR), International Joint Conference on Automated Reasoning 2008*, pages 47–62, 2008.

45. Donald E. Knuth and Peter B. Bendix. Simple word problems in universal algebras. In J. Leech, editor, *Proceedings of the Conference on Computational Problems in Abstract Algebras*, pages 263–298. Pergamon Press, 1970.
46. Konstantin Korovin and Andrei Voronkov. Integrating linear arithmetic into superposition calculus. In Jacques Duparc and Thomas A. Henzinger, editors, *Proceedings of the Sixteenth EACSL Annual Conference on Computer Science Logic (CSL)*, volume 4646 of *Lecture Notes in Computer Science*, pages 223–237. Springer, 2007.
47. Emmanuel Kounalis and Michaël Rusinowitch. On word problems in Horn theories. *Journal of Symbolic Computation*, 11(1–2):113–128, 1991.
48. Vladimir Lifschitz, Leora Morgenstern, and David A. Plaisted. Knowledge representation and classical logic. In Frank van Harmelen, Vladimir Lifschitz, and Bruce Porter, editors, *Handbook of Knowledge Representation*, volume 1, pages 3–88. Elsevier, 2008.
49. Christopher A. Lynch. Unsound theorem proving. In Jan Marcinkowski and Andrzej Tarlecki, editors, *Proceedings of the Thirteenth EACSL Annual Conference on Computer Science Logic (CSL)*, volume 3210 of *Lecture Notes in Computer Science*, pages 473–487. Springer, 2004.
50. Holbrook Mann MacNeille. Partially ordered sets. In *Transactions of the American Mathematical Society*, volume 42, pages 416–460, 1937.
51. William W. McCune. Otter 3.3 reference manual. Technical Report ANL/MCS-TM-263, MCS Division, Argonne National Laboratory, Argonne, IL, USA, 2003.
52. Scott McPeak and George C. Necula. Data structure specifications via local equality axioms. In Kousha Etessami and Sriram. K. Rajamani, editors, *Proceedings of the Seventeenth Conference on Automated Verification (CAV)*, volume 3576 of *Lecture Notes in Computer Science*, pages 476–490. Springer, 2005.
53. Greg Nelson and Derek C. Oppen. Simplification by cooperating decision procedures. *ACM Transactions on Programming Languages and Systems*, 1(2):245–257, 1979.
54. Enrica Nicolini, Christophe Ringeissen, and Michaël Rusinowitch. Data structures with arithmetic constraints: a non-disjoint combination. In Silvio Ghilardi and Roberto Sebastiani, editors, *Proceedings of the Seventh Symposium on Frontiers of Combining Systems (FroCoS)*, volume 5749 of *Lecture Notes in Artificial Intelligence*, pages 319–334. Springer, 2009.
55. Robert Nieuwenhuis, Albert Oliveras, and Cesare Tinelli. Solving SAT and SAT modulo theories: from an abstract Davis-Putnam-Logemann-Loveland procedure to DPLL(T). *Journal of the ACM*, 53(6):937–977, 2006.
56. Robert Nieuwenhuis and Albert Rubio. Paramodulation-based theorem proving. In Alan Robinson and Andrei Voronkov, editors, *Handbook of Automated Reasoning*, volume 1, pages 371–443. Elsevier, 2001.
57. G. Robinson and Larry Wos. Paramodulation and theorem-proving in first-order theories with equality. In D. Michie and R. Meltzer, editors, *Machine Intelligence*, volume IV, pages 135–150. Edinburgh University Press, 1969.
58. J. Alan Robinson. Automatic deduction with hyper-resolution. *International Journal of Computer Mathematics*, 1:227–234, 1965.
59. J. Alan Robinson. A machine oriented logic based on the resolution principle. *Journal of the ACM*, 12(1):23–41, 1965.
60. Michaël Rusinowitch. Theorem-proving with resolution and superposition. *Journal of Symbolic Computation*, 11:21–50, 1991.
61. Roberto Sebastiani. Lazy satisfiability modulo theories. *Journal of Satisfiability, Boolean Modeling and Computation*, 3:141–224, 2007.
62. Wayne Snyder. A fast algorithm for generating reduced ground rewriting systems from a set of ground equations. *Journal of Symbolic Computation*, 15(4):415–450, 1993.
63. Viorica Sofronie-Stokkermans. Hierarchic reasoning in local theory extensions. In Robert Nieuwenhuis, editor, *Proceedings of the Twentieth Conference on Automated Deduction (CADE)*, volume 3632 of *Lecture Notes in Artificial Intelligence*, pages 219–234. Springer, 2005.
64. Viorica Sofronie-Stokkermans and Carsten Ihlemann. Automated reasoning in some local extensions of ordered structures. *Journal of Multiple-Valued Logics and Soft Computing*, 13(4–6):397–414, 2007.
65. Thomas Wies, Ruzica Piskac, and Viktor Kuncak. Combining theories with shared set operations. In Silvio Ghilardi and Roberto Sebastiani, editors, *Proceedings of the Seventh Symposium on Frontiers of Combining Systems (FroCoS)*, volume 5749 of *Lecture Notes in Artificial Intelligence*, pages 366–382. Springer, 2009.