

Decidability of the Monadic Shallow Linear First-Order Fragment with Straight Dismatching Constraints

Andreas Teucke^{1,2} and Christoph Weidenbach¹

¹ Max-Planck Institut für Informatik, Saarland Informatics Campus, 66123 Saarbrücken Germany

² Graduate School of Computer Science, Saarbrücken, Germany

Abstract. The monadic shallow linear Horn fragment is well-known to be decidable and has many application, e.g., in security protocol analysis, tree automata, or abstraction refinement. It was a long standing open problem how to extend the fragment to the non-Horn case, preserving decidability, that would, e.g., enable to express non-determinism in protocols. We prove decidability of the non-Horn monadic shallow linear fragment via ordered resolution further extended with dismatching constraints and discuss some applications of the new decidable fragment.

1 Introduction

Motivated by the automatic analysis of security protocols, the monadic shallow linear Horn (MSLH) fragment was shown to be decidable in [21]. In addition to the restriction to monadic Horn clauses, the main restriction of the fragment is positive literals of the form $S(f(x_1, \dots, x_n))$ or $S(x)$ where all x_i are different, i.e., all terms are shallow and linear. The fragment can be finitely saturated by superposition (ordered resolution) where negative literals with non-variable arguments are always selected. As a result, productive clauses with respect to the superposition model operator \mathcal{I}_N have the form $S_1(x_1), \dots, S_n(x_n) \rightarrow S(f(x_1, \dots, x_n))$. Therefore, the models of saturated MSLH clause sets can both be represented by tree automata [6] and shallow linear sort theories [8]. The models are typically infinite. The decidability result of MSLH clauses was rediscovered in the context of tree automata research [7] where in addition DEXPTIME-completeness of the MSLH fragment was shown. The fragment was further extended by disequality constraints [12,13] still motivated by security protocol analysis [14]. Although from a complexity point of view, the difference between Horn clause fragments and the respective non-Horn clause fragments is typically reflected by membership in the deterministic vs. the non-deterministic respective complexity fragment, for monadic shallow linear clauses so far there was no decidability result for the non-Horn case.

The results of this paper close this gap. We show the monadic shallow linear non-Horn (MSL) clause fragment to be decidable by superposition (ordered resolution). From a security protocol application point of view, non-Horn clauses

enable a natural representation of non-determinism. Our second extension to the fragment are unit clauses with disequations of the form $s \not\approx t$, where s and t are not unifiable. Due to the employed superposition calculus, such disequations do not influence saturation of an MSL clause set, but have an effect on potential models. They can rule out identification of syntactically different ground terms as it is, e.g., desired in the security protocol context for syntactically different messages or nonces. Our third extension to the fragment are straight mismatching constraints. These constraints are incomparable to the disequality constraints mentioned above [12,13]. They do not strictly increase the expressiveness of the MSL theory, but enable up to exponentially more compact saturations. For example, the constrained clause

$$(S(x), T(y) \rightarrow S(f(x, y)); y \neq f(x', f(a, y')))$$

over constants a, b describes the same set of ground clauses as the six unconstrained clauses

$$\begin{aligned} S(x), T(a) \rightarrow S(f(x, a)) \quad & S(x), T(b) \rightarrow S(f(x, b)) \quad \dots \\ & S(x), T(f(b, y')) \rightarrow S(f(x, f(b, y'))) \\ S(x), T(f(f(x'', y''), y')) \rightarrow & S(f(x, f(f(x'', y''), y'))). \end{aligned}$$

Furthermore, for a satisfiability equivalent transformation into MSL clauses, the nested terms in the positive literals would have to be factored out by the introduction of further predicates and clauses. E.g., the first clause is replaced by the two MSL clauses $S(x), T(a), R(y) \rightarrow S(f(x, y))$ and $R(a)$ where R is a fresh monadic predicate. The constrained clause belongs to the MSL(SDC) fragment. Altogether, the resulting MSL(SDC) fragment is shown to be decidable in Section 3.

The introduction of straight mismatching constraints (SDCs) enables an improved refinement step of our approximation refinement calculus [18]. Before, several clauses were needed to rule out a specific instance of a clause in an unsatisfiable core. For example, if due to a linearity approximation from clause $S(x), T(x) \rightarrow S(f(x, x))$ to $S(x), T(x), S(y), T(y) \rightarrow S(f(x, y))$ an instance $\{x \mapsto f(a, x'), y \mapsto f(b, y')\}$ is used in the proof, before [18] several clauses were needed to replace $S(x), T(x) \rightarrow S(f(x, x))$ in a refinement step in order to rule out this instance. With straight mismatching constraints the clause $S(x), T(x) \rightarrow S(f(x, x))$ is replaced by the two clauses $S(f(a, x), T(f(a, x)) \rightarrow S(f(f(a, x), f(a, x)))$ and $(S(x), T(x) \rightarrow S(f(x, x)); x \neq f(a, y))$. For the improved approximation refinement approach (FO-AR) presented in this paper, any refinement step results in just two clauses, see Section 4. The additional expressiveness of constraint clauses comes almost for free, because necessary computations, like, e.g., checking emptiness of SDCs, can all be done in polynomial time, see Section 2.

In addition to the extension of the known MSLH decidability result and the improved approximation refinement calculus FO-AR, we discuss in Section 5 the potential of the MSL(SDC) fragment in the context of FO-AR, Theorem 2, and its prototypical implementation in SPASS-AR (<http://www.mpi-inf.mpg.de/fileadmin/inf/rg1/spass-ar>). It turns out that for clause sets containing certain structures, FO-AR is superior to ordered resolution/superposition [1] and instance generating methods [10].

The paper ends with a discussion on challenges and future research directions, Section 6.

2 First-Order Clauses with Straight Dismatching Constraints: MSL(SDC)

We consider a standard first-order language where letters v, w, x, y, z denote variables, f, g, h functions, a, b, c constants, s, t terms, p, q, r positions and Greek letters $\sigma, \tau, \rho, \delta$ are used for substitutions. S, P, Q, R denote predicates, \approx denotes equality, A, B atoms, E, L literals, C, D clauses, N clause sets and \mathcal{V} sets of variables. \overline{L} is the complement of L . The signature $\Sigma = (\mathcal{F}, \mathcal{P})$ consists of two disjoint, non-empty, in general infinite sets of function and predicate symbols \mathcal{F} and \mathcal{P} , respectively. The set of all *terms* over variables \mathcal{V} is $\mathcal{T}(\mathcal{F}, \mathcal{V})$. If there are no variables, then terms, literals and clauses are called *ground*, respectively. A *substitution* σ is denoted by pairs $\{x \mapsto t\}$ and its update at x by $\sigma[x \mapsto t]$. A substitution σ is a *grounding* substitution for \mathcal{V} if $x\sigma$ is ground for every variable $x \in \mathcal{V}$.

The set of *free* variables of an atom A (term t) denoted by $\text{vars}(A)$ ($\text{vars}(t)$). A *position* is a sequence of positive integers, where ε denotes the empty position. As usual $t|_p = s$ denotes the subterm s of t at position p , which we also write as $t[s]_p$, and $t[p/s']$ then denotes the replacement of s with s' in t at position p . These notions are extended to literals and multiple positions.

A predicate with exactly one argument is called *monadic*. A term is *complex* if it is not a variable and *shallow* if it has at most depth one. It is called *linear* if there are no duplicate variable occurrences. A literal, where every argument term is shallow, is also called *shallow*. A variable and a constant are called *straight*. A term $f(s_1, \dots, s_n)$ is called *straight*, if s_1, \dots, s_n are different variables except for at most one straight term s_i .

A *clause* is a multiset of literals which we write as an implication $\Gamma \rightarrow \Delta$ where the atoms in the multiset Δ (the *succedent*) denote the positive literals and the atoms in the multiset Γ (the *antecedent*) the negative literals. We write \square for the empty clause. If Γ is empty we omit \rightarrow , e.g., we can write $P(x)$ as an alternative of $\rightarrow P(x)$. We abbreviate disjoint set union with sequencing, for example, we write $\Gamma, \Gamma' \rightarrow \Delta, L$ instead of $\Gamma \cup \Gamma' \rightarrow \Delta \cup \{L\}$. A clause $E, E, \Gamma \rightarrow \Delta$ is equivalent to $E, \Gamma \rightarrow \Delta$ and we call them equal *modulo duplicate literal elimination*. If every term in Δ is shallow, the clause is called *positive shallow*. If all atoms in Δ are linear and variable disjoint, the clause is called *positive linear*. A clause $\Gamma \rightarrow \Delta$ is called an *MSL* clause, if it is (i) positive shallow and linear, (ii) all occurring predicates are monadic, (iii) no equations occur in Δ , and (iv) no equations occur in Γ or $\Gamma = \{s \approx t\}$ and Δ is empty where s and t are not unifiable. *MSL* is the first-order clause fragment consisting of MSL clauses. Clauses $\Gamma, s \approx t \rightarrow \Delta$ where Γ, Δ are non-empty and s, t are not unifiable could be added to the MSL fragment without changing any of our results. Considering the superposition calculus, it will select $s \approx t$. Since the two terms are not unifiable, no inference will take place on such a clause and the

clause will not contribute to the model operator. In this sense such clauses do not increase the expressiveness of the fragment.

An *atom ordering* \prec is an irreflexive, well-founded, total ordering on ground atoms. It is lifted to literals by representing A and $\neg A$ as multisets $\{A\}$ and $\{A, A\}$, respectively. The multiset extension of the literal ordering induces an ordering on ground clauses. The clause ordering is compatible with the atom ordering; if the maximal atom in C is greater than the maximal atom in D then $D \prec C$. We use \prec simultaneously to denote an atom ordering and its multiset, literal, and clause extensions. For a ground clause set N and clause C , the set $N^{\prec C} = \{D \in N \mid D \prec C\}$ denotes the clauses of N smaller than C .

A *Herbrand interpretation* \mathcal{I} is a - possibly infinite - set of ground atoms. A ground atom A is called *true* in \mathcal{I} if $A \in \mathcal{I}$ and *false*, otherwise. \mathcal{I} is said to *satisfy* a ground clause $C = F \rightarrow \Delta$, denoted by $\mathcal{I} \models C$, if $\Delta \cap \mathcal{I} \neq \emptyset$ or $F \not\subseteq \mathcal{I}$. A non-ground clause C is satisfied by \mathcal{I} if $\mathcal{I} \models C\sigma$ for every grounding substitution σ . An interpretation \mathcal{I} is called a *model* of N , $\mathcal{I} \models N$, if $\mathcal{I} \models C$ for every $C \in N$. A model \mathcal{I} of N is considered *minimal* with respect to set inclusion, i.e., if there is no model \mathcal{I}' with $\mathcal{I}' \subset \mathcal{I}$ and $\mathcal{I}' \models N$. A set of clauses N is *satisfiable*, if there exists a model that satisfies N . Otherwise, the set is *unsatisfiable*.

A disequation $t \neq s$ is an *atomic straight dismatching constraint* if s and t are variable disjoint terms and s is straight. A straight dismatching constraint π is a conjunction of atomic straight dismatching constraints. Given a substitution σ , $\pi\sigma = \bigwedge_{i \in I} t_i\sigma \neq s_i$. $\text{lvar}(\pi) := \bigcup_{i \in I} \text{vars}(t_i)$ are the left-hand variables of π and the depth of π is the maximal term depth of the s_i . A *solution* of π is a grounding substitution δ such that for all $i \in I$, $t_i\delta$ is not an instance of s_i , i.e., there exists no σ such that $t_i\delta = s_i\sigma$. A dismatching constraint is solvable if it has a solution and unsolvable, otherwise. Whether a straight dismatching constraint is solvable, is decidable in linear-logarithmic time [19]. \top and \perp represent the true and false dismatching constraint, respectively.

We define constraint normalization $\pi \downarrow$ as the normal form of the following rewriting rules over straight dismatching constraints.

$$\begin{aligned} \pi \wedge f(t_1, \dots, t_n) \neq y &\Rightarrow \perp \\ \pi \wedge f(t_1, \dots, t_n) \neq f(y_1, \dots, y_n) &\Rightarrow \perp \\ \pi \wedge f(t_1, \dots, t_n) \neq f(s_1, \dots, s_n) &\Rightarrow \pi \wedge t_i \neq s_i \quad \text{if } s_i \text{ is complex} \\ \pi \wedge f(t_1, \dots, t_n) \neq g(s_1, \dots, s_m) &\Rightarrow \pi \\ \pi \wedge x \neq s \wedge x \neq s\sigma &\Rightarrow \pi \wedge x \neq s \end{aligned}$$

Note that $f(t_1, \dots, t_n) \neq f(s_1, \dots, s_n)$ normalizes to $t_i \neq s_i$ for some i , where s_i is the one straight complex argument of $f(s_1, \dots, s_n)$. Furthermore, the depth of $\pi \downarrow$ is less or equal to the depth of π and both have the same solutions.

A pair of a clause and a constraint $(C; \pi)$ is called a *constrained clause*. Given a substitution σ , $(C; \pi)\sigma = (C\sigma; \pi\sigma)$. $C\delta$ is called a ground clause of $(C; \pi)$ if δ is a solution of π . $\mathcal{G}((C; \pi))$ is the set of ground instances of $(C; \pi)$. If $\mathcal{G}((C; \pi)) \subseteq \mathcal{G}((C'; \pi'))$, then $(C; \pi)$ is an instance of $(C'; \pi')$. If $\mathcal{G}((C; \pi)) = \mathcal{G}((C'; \pi'))$, then $(C; \pi)$ and $(C'; \pi')$ are called variants. A Herbrand interpretation \mathcal{I} satisfies $(C; \pi)$, if $\mathcal{I} \models \mathcal{G}((C; \pi))$. A constrained clause $(C; \pi)$ is called *redundant* in N if for

every $D \in \mathcal{G}((C; \pi))$, there exist D_1, \dots, D_n in $\mathcal{G}(N)^{<D}$ such that $D_1, \dots, D_n \models D$. A constrained clause $(C'; \pi')$ is called a *condensation* of $(C; \pi)$ if $C' \subset C$ and there exists a substitution σ such that, $\pi\sigma = \pi'$, $\pi' \subseteq \pi$, and for all $L \in C$ there is an $L' \in C'$ with $L\sigma = L'$. A finite unsatisfiable subset of $\mathcal{G}(N)$ is called an unsatisfiable core of N .

An MSL clause with straight dismatching constraints is called an *MSL(SDC)* clause with MSL(SDC) being the respective first-order fragment. Note that any clause set N can be transformed into an equivalent constrained clause set by changing each $C \in N$ to $(C; \top)$.

3 Decidability of the MSL(SDC) fragment

In the following we will show that the satisfiability of the MSL(SDC) fragment is decidable. For this purpose we will define ordered resolution with selection on constrained clauses [19] and show that with an appropriate ordering and selection function, saturation of an MSL(SDC) clause set terminates.

For the rest of this section we assume an atom ordering $<$ such that a literal $\neg Q(s)$ is not greater than a literal $P(t[s]_p)$, where $p \neq \varepsilon$. For example, a KBO where all symbols have weight one has this property.

Definition 1 (sel). *Given an MSL(SDC) clause $(C; \pi) = (S_1(t_1), \dots, S_n(t_n) \rightarrow P_1(s_1), \dots, P_m(s_m); \pi)$. The Superposition Selection function sel is defined by $S_i(t_i) \in \text{sel}(C)$ if (1) t_i is not a variable or (2) t_1, \dots, t_n are variables and $t_i \notin \text{vars}(s_1, \dots, s_m)$ or (3) $\{t_1, \dots, t_n\} \subseteq \text{vars}(s_1, \dots, s_m)$ and for some $1 \leq j \leq m$, $s_j = t_i$.*

The selection function sel (Definition 1) ensures that a clause $\Gamma \rightarrow \Delta$ can only be resolved on a positive literal if Γ contains only variables, which also appear in Δ at a non-top position. For example:

$$\begin{aligned} \text{sel}(P(f(x)), P(x), Q(z) \rightarrow Q(x), R(f(y))) &= \{P(f(x))\} \\ \text{sel}(P(x), Q(z) \rightarrow Q(x), R(f(y))) &= \{Q(z)\} \\ \text{sel}(P(x), Q(y) \rightarrow Q(x), R(f(y))) &= \{P(x)\} \\ \text{sel}(P(x), Q(y) \rightarrow Q(f(x)), R(f(y))) &= \emptyset. \end{aligned}$$

Note that given an MSL(SDC) clause $(C; \pi) = (S_1(t_1), \dots, S_n(t_n) \rightarrow P_1(s_1), \dots, P_m(s_m); \pi)$, if some $S_i(t_i)$ is maximal in C , then at least one literal is selected.

Definition 2. *A literal A is called [strictly] maximal in a constrained clause $(C \vee A; \pi)$ if and only if there exists a solution δ of π such that for all literals B in C , $B\delta \preceq A\delta$ [$B\delta < A\delta$].*

Definition 3 (SDC-Resolution).

$$\frac{(\Gamma_1 \rightarrow \Delta_1, A; \pi_1) \quad (\Gamma_2, B \rightarrow \Delta_2; \pi_2)}{((\Gamma_1, \Gamma_2 \rightarrow \Delta_1, \Delta_2)\sigma; (\pi_1 \wedge \pi_2)\sigma\downarrow)}, \text{ if}$$

1. $\sigma = \text{mgu}(A, B)$
2. $(\pi_1 \wedge \pi_2)\sigma\downarrow$ is solvable
3. $A\sigma$ is strictly maximal in $(\Gamma_1 \rightarrow \Delta_1, A; \pi_1)\sigma$ and $\text{sel}(\Gamma_1 \rightarrow \Delta_1, A) = \emptyset$
4. $B \in \text{sel}(\Gamma_2, B \rightarrow \Delta_2)$
5. $\text{sel}(\Gamma_2, B \rightarrow \Delta_2) = \emptyset$ and $\neg B\sigma$ maximal in $(\Gamma_2, B \rightarrow \Delta_2; \pi_2)\sigma$

Definition 4 (SDC-Factoring).

$$\frac{(\Gamma \rightarrow \Delta, A, B ; \pi)}{((\Gamma \rightarrow \Delta, A)\sigma; \pi\sigma\downarrow)} , \text{ if}$$

1. $\sigma = \text{mgu}(A, B)$
2. $\text{sel}(\Gamma \rightarrow \Delta, A, B) = \emptyset$
3. $A\sigma$ is maximal in $(\Gamma \rightarrow \Delta, A, B; \pi)\sigma$
4. $\pi\sigma\downarrow$ is solvable

Note that while the above rules do not operate on equations, we can actually allow unit clauses that consist of non-unifiable disequations, i.e., clauses $s \approx t \rightarrow$ where s and t are not unifiable. There are no potential superposition inferences on such clauses as long as there are no positive equations. So resolution and factoring suffice for completeness. Nevertheless, clauses such as $s \approx t \rightarrow$ affect the models of satisfiable problems. Constrained Resolution and Factoring are sound.

Lemma 1 (Soundness). *SDC-Resolution and SDC-Factoring are sound.*

Proof. Let $(\Gamma_1, \Gamma_2 \rightarrow \Delta_1, \Delta_2)\sigma\delta$ be a ground instance of $((\Gamma_1, \Gamma_2 \rightarrow \Delta_1, \Delta_2)\sigma; (\pi_1 \wedge \pi_2)\sigma)$. Then, δ is a solution of $(\pi_1 \wedge \pi_2)\sigma$ and $\sigma\delta$ is a solution of π_1 and π_2 . Hence, $(\Gamma_1 \rightarrow \Delta_1, A)\sigma\delta$ and $(\Gamma_2, B \rightarrow \Delta_2)\sigma\delta$ are ground instances of $(\Gamma_1 \rightarrow \Delta_1, A; \pi_1)$ and $(\Gamma_2, B \rightarrow \Delta_2; \pi_2)$, respectively. Because $A\sigma\delta = B\sigma\delta$, if $(\Gamma_1 \rightarrow \Delta_1, A)\sigma\delta$ and $(\Gamma_2, B \rightarrow \Delta_2)\sigma\delta$ are satisfied, then $(\Gamma_1, \Gamma_2 \rightarrow \Delta_1, \Delta_2)\sigma\delta$ is also satisfied. Therefore, SDC-Resolution is sound. Let $(\Gamma \rightarrow \Delta, A)\sigma\delta$ be a ground instance of $((\Gamma \rightarrow \Delta, A)\sigma; \pi\sigma)$. Then, δ is a solution of $\pi\sigma$ and $\sigma\delta$ is a solution of π . Hence, $(\Gamma \rightarrow \Delta, A, B)\sigma\delta$ is a ground instance of $(\Gamma \rightarrow \Delta, A, B; \pi)$. Because $A\sigma\delta = B\sigma\delta$, if $(\Gamma \rightarrow \Delta, A, B)\sigma\delta$ is satisfied, then $(\Gamma \rightarrow \Delta, A)\sigma\delta$ is also satisfied. Therefore, SDC-Factoring is sound. \square

Definition 5 (Saturation). *A constrained clause set N is called saturated up to redundancy, if for every inference between clauses in N the result $(R; \pi)$ is either redundant in N or $\mathcal{G}((R; \pi)) \subseteq \mathcal{G}(N)$.*

Note that our redundancy notion includes condensation and the condition $\mathcal{G}((R; \pi)) \subseteq \mathcal{G}(N)$ allows ignoring variants of clauses.

Lemma 2. *Let constrained clause $(C'; \pi')$ be a condensation of constrained clause $(C; \pi)$. Then, (i) $(C; \pi) \models (C'; \pi')$ and (ii) $(C; \pi)$ is redundant in $\{(C'; \pi')\}$.*

Proof. Let σ be a substitution such that $C' \subset C$, $\pi\sigma = \pi'$, $\pi' \subseteq \pi$, and for all $L \in C$ there is a $L' \in C'$ with $L\sigma = L'$.

(i) Let $C'\delta \in \mathcal{G}((C'; \pi'))$. Then $\sigma\delta$ is a solution of π and hence $C\sigma\delta \in \mathcal{G}((C; \pi))$. Let $\mathcal{I} \models C\sigma\delta$. Hence, there is a $L\sigma\delta \in \mathcal{I}$ for some $L \in C$ and thus $L'\delta \in \mathcal{I}$ for some $L' \in C'$ with $L\sigma = L'$. Therefore, $\mathcal{I} \models C'\delta$. Since \mathcal{I} and $C'\delta$ were arbitrary, $(C; \pi) \models (C'; \pi')$.

(ii) Let $C\delta \in \mathcal{G}((C; \pi))$. Because $\pi' \subseteq \pi$, δ is a solution of π' and hence, $C'\delta \in \mathcal{G}((C'; \pi'))$. Therefore, since $C'\delta \subset C\delta$, $C'\delta \in \mathcal{G}(\{(C'; \pi')\})^{<C\delta}$ and $C'\delta \models C\delta$. \square

Definition 6 (Partial Minimal Model Construction). *Given a constrained clause set N , an ordering \prec and the selection function sel , we construct an interpretation \mathcal{I}_N for N , called a partial model, inductively as follows:*

$$\begin{aligned}\mathcal{I}_C &:= \bigcup_{\substack{D \in \mathcal{G}(N) \\ D \prec C}} \delta_D, \text{ where } C \in \mathcal{G}(N) \\ \delta_D &:= \begin{cases} \{A\} & \text{if } D = \Gamma \rightarrow \Delta, A \\ & A \text{ strictly maximal, } \text{sel}(D) = \emptyset \text{ and } \mathcal{I}_D \not\models D \\ \emptyset & \text{otherwise} \end{cases} \\ \mathcal{I}_N &:= \bigcup_{C \in \mathcal{G}(N)} \delta_C\end{aligned}$$

Clauses D with $\delta_D \neq \emptyset$ are called *productive*.

Lemma 3 (Ordered SDC Resolution Completeness). *Let N be a constrained clause set saturated up to redundancy by ordered SDC-resolution with selection. Then N is unsatisfiable, if and only if $\square \in \mathcal{G}(N)$. If $\square \notin \mathcal{G}(N)$ then $\mathcal{I}_N \models N$.*

Proof. Assume N is unsatisfiable but $\square \notin \mathcal{G}(N)$. For the partial model \mathcal{I}_N , there exists a minimal false clause $C\sigma \in \mathcal{G}((C; \pi))$ for some $(C; \pi) \in N$.

$C\sigma$ is not productive, because otherwise $\mathcal{I}_N \models C\sigma$. Hence, either $\text{sel}(C) \neq \emptyset$ or no positive literal in $C\sigma$ is strictly maximal. Assume $C = \Gamma_2, B \rightarrow \Delta_2$ with $B \in \text{sel}(C)$ or $\neg B\sigma$ maximal. Then, $B\sigma \in \mathcal{I}_{C\sigma}$ and there exists a ground instance $(\Gamma_1 \rightarrow \Delta_1, A)\tau = D\tau \prec C\sigma$ of some clause $(D; \pi') \in N$, which produces $A\tau = B\sigma$. Therefore, there exists a $\rho = \text{mgu}(A, B)$ and ground substitution δ such that $C\sigma = C\rho\delta$, $D\tau = D\rho\delta$. Since $\rho\delta = \sigma$ is a solution of π and π' , δ is a solution of $(\pi \wedge \pi')\rho$. Under these conditions, SDC-Resolution can be applied to $(\Gamma_1 \rightarrow \Delta_1, A; \pi')$ and $(\Gamma_2, B \rightarrow \Delta_2; \pi)$. Their resolvent $(R; \pi_R) = ((\Gamma_1, \Gamma_2 \rightarrow \Delta_1, \Delta_2)\rho; (\pi \wedge \pi')\rho)$ is either redundant in N or $\mathcal{G}((R; \pi_R)) \subseteq \mathcal{G}(N)$. Its ground instance $R\delta$ is false in \mathcal{I}_N and $R\delta \prec C\sigma$. If $(R; \pi_R)$ is redundant in N , there exist C_1, \dots, C_n in $\mathcal{G}(N)^{\prec R\delta}$ with $C_1, \dots, C_n \models R\delta$. Because $C_i \prec R\delta \prec C\sigma$, $\mathcal{I}_N \models C_i$ and hence $\mathcal{I}_N \models R\delta$, which contradicts $\mathcal{I}_N \not\models R\delta$. Otherwise, if $\mathcal{G}((R; \pi_R)) \subseteq \mathcal{G}(N)$, then $R\delta \in \mathcal{G}(N)$, which contradicts $C\sigma$ being minimal false.

Now, assume $\text{sel}(C) = \emptyset$ and $C = \Gamma \rightarrow \Delta, B$ with $B\sigma$ maximal. Then, $C = \Gamma \rightarrow \Delta', A, B$ with $A\sigma = B\sigma$. Therefore, there exists a $\rho = \text{mgu}(A, B)$ and ground substitution δ such that $C\sigma = C\rho\delta$ and $\rho\delta$ is a solution of π . Hence, δ is a solution of $\pi\rho$. Under these conditions, SDC-Factoring can be applied to $(\Gamma \rightarrow \Delta', A, B; \pi)$. The result $(R; \pi_R) = ((\Gamma \rightarrow \Delta', A)\rho; \pi\rho)$ is either redundant in N or $\mathcal{G}((R; \pi_R)) \subseteq \mathcal{G}(N)$. Its ground instance $R\delta$ is false in \mathcal{I}_N and $R\delta \prec C\sigma$. If $(R; \pi_R)$ is redundant in N , there exist C_1, \dots, C_n in $\mathcal{G}(N)^{\prec R\delta}$ with $C_1, \dots, C_n \models R\delta$. Because $C_i \prec R\delta \prec C\sigma$, $\mathcal{I}_N \models C_i$ and hence $\mathcal{I}_N \models R\delta$, which contradicts $\mathcal{I}_N \not\models R\delta$. Otherwise, if $\mathcal{G}((R; \pi_R)) \subseteq \mathcal{G}(N)$, then $R\delta \in \mathcal{G}(N)$, which contradicts $C\sigma$ being minimal false.

Therefore, if $\square \notin \mathcal{G}(N)$, no minimal false clause exists and $\mathcal{I}_N \models N$. \square

Lemma 4. *Let N be a set of $MSL(SDC)$ clauses without variants or uncondensed clauses over a finite signature Σ . N is finite if there exists an integer d such that for every $(C; \pi) \in N$, $\text{depth}(\pi) \leq d$ and*

- (1) $C = S_1(x_1), \dots, S_n(x_n), S'_1(t), \dots, S'_m(t) \rightarrow \Delta$ or
(2) $C = S_1(x_1), \dots, S_n(x_n), S'_1(t), \dots, S'_m(t) \rightarrow S(t), \Delta$
with t shallow and linear, and $\text{vars}(t) \cap \text{vars}(\Delta) = \emptyset$.

Proof. Let $(C; \pi) \in N$. $(C; \pi)$ can be separated into variable disjoint components $(\Gamma_1, \dots, \Gamma_n \rightarrow \Delta_1, \dots, \Delta_n; \pi_1 \wedge \dots \wedge \pi_n)$, where $|\Delta_i| \leq 1$ and $\text{lvar}(\pi_i) \subseteq \text{vars}(\Gamma_i \rightarrow \Delta_i)$. For each positive literal $P(s) \in \Delta$ there is a fragment

$$(A) \quad (S_1(x_1), \dots, S_k(x_k) \rightarrow P(s); \pi')$$

with $\{x_1, \dots, x_k\} \subseteq \text{vars}(s)$. If $m > 0$, there is another fragment

$$(B) \quad (S_1(x_1), \dots, S_k(x_k), S'_1(t), \dots, S'_m(t) \rightarrow; \pi')$$

or

$$(C) \quad (S_1(x_1), \dots, S_k(x_k), S'_1(t), \dots, S'_m(t) \rightarrow S(t); \pi')$$

with $\{x_1, \dots, x_k\} \subseteq \text{vars}(t)$, respectively. Lastly, for each variable $x \in \text{vars}(C)$ with $x \notin \text{vars}(t) \cup \text{vars}(\Delta)$ there is a fragment

$$(D) \quad (S_1(x), \dots, S_k(x) \rightarrow; \pi').$$

Since there are only finitely many terms s with $\text{depth}(s) \leq d$ modulo renaming, there are only finitely many atomic constraints $x \neq s$ for a given variable x different up to renaming s . Thus, a normal constraint can only contain finitely many combinations of subconstraints $\bigwedge_{i \in \mathcal{I}} x \neq s_i$ without some s_i being an instance of another s_j . Therefore, for a fixed set of variables x_1, \dots, x_k , there are only finitely many constraints $\pi = \bigwedge_{i \in \mathcal{I}} z_i \neq s_i$ with $\text{lvar}(\pi) \subseteq \{x_1, \dots, x_k\}$ up to variants.

Since the number of predicates, function symbols, and their ranks is finite, the number of possible shallow and linear atoms $S(t)$ different up to variants is finite. For a given shallow and linear t , there exist only finitely many clauses of the form $(S_1(t), \dots, S_n(t) \rightarrow S(t); \pi)$ or $(S_1(t), \dots, S_n(t) \rightarrow; \pi)$ with $\text{lvar}(\pi) \subseteq \text{vars}(t)$ modulo condensation and variants. For a fixed set of variables x_1, \dots, x_k , there exist only finitely many clauses of the form $(S_1(y_1), \dots, S_k(y_l) \rightarrow; \pi)$ with $\{y_1, \dots, y_l\} \cup \text{lvar}(\pi) \subseteq \{x_1, \dots, x_k\}$ modulo condensation and variants. Therefore, there are only finitely many distinct clauses of each form (A)-(D) without variants or condensations.

If in the clause $(C; \pi) = (\Gamma_1, \dots, \Gamma_n \rightarrow \Delta_1, \dots, \Delta_n; \pi_1 \wedge \dots \wedge \pi_n)$ for some $i \neq j$, $(\Gamma_i \rightarrow \Delta_i; \pi_i)$ is a variant of $(\Gamma_j \rightarrow \Delta_j; \pi_j)$, then $(C; \pi)$ has a condensation and is therefore not part of N . Hence, there can be only finitely many different $(C; \pi)$ without variants or condensations and thus N is finite. \square

Lemma 5 (Finite Saturation). *Let N be an $MSL(SDC)$ clause set. Then N can be finitely saturated up to redundancy by SDC -resolution with selection function sel .*

Proof. The general idea is that given the way sel is defined the clauses involved in constrained resolution and factoring can only fall into certain patterns. Any result of such inferences then is either strictly smaller than one of its parents by some terminating measure or falls into a set of clauses that is bounded by Lemma 4. Thus, there can be only finitely many inferences before N is saturated.

Let d be an upper bound on the depth of constraints found in N and Σ be the finite signature consisting of the function and predicate symbols occurring in N . Let $(\Gamma_1 \rightarrow \Delta_1, S(t); \pi_1)$ and $(\Gamma_2, S(t') \rightarrow \Delta_2; \pi_2)$ be clauses in N where sdc-resolution applies with $\sigma = \text{mgu}(S(t), S(t'))$ and resolvent $R = ((\Gamma_1, \Gamma_2 \rightarrow \Delta_1, \Delta_2)\sigma; (\pi_1 \wedge \pi_2)\sigma \downarrow)$.

Because no literal is selected by sel , $\Gamma_1 \rightarrow \Delta_1, S(t)$ can match only one of two patterns:

$$(A) \quad S_1(x_1), \dots, S_n(x_n) \rightarrow S(f(y_1, \dots, y_k)), \Delta$$

where $t = f(y_1, \dots, y_k)$ and $\{x_1, \dots, x_n\} \subseteq \{y_1, \dots, y_k\} \cup \text{vars}(\Delta)$.

$$(B) \quad S_1(x_1), \dots, S_n(x_n) \rightarrow S(y), \Delta$$

where $t = y$ and x_1, \dots, x_n are variables in $\text{vars}(\Delta)$, i.e., y occurs only once.

The literal $S(t')$ is selected by sel in $\Gamma_2, S(t') \rightarrow \Delta_2$, and therefore $\Gamma_2, S(t') \rightarrow \Delta_2$ can match only one of the following three patterns:

- (1) $S(f(t_1, \dots, t_k)), \Gamma' \rightarrow \Delta'$
- (2) $S(y'), \Gamma' \rightarrow \Delta'$ where Γ' has no function terms and $y \notin \text{vars}(\Delta')$.
- (3) $S(y'), \Gamma' \rightarrow S'(y'), \Delta'$ where Γ' has no function terms.

This means that the clausal part $(\Gamma_1, \Gamma_2 \rightarrow \Delta_1, \Delta_2)\sigma$ of R has one of six forms:

$$(A1) \quad S_1(x_1)\sigma, \dots, S_n(x_n)\sigma, \Gamma' \rightarrow \Delta, \Delta' \text{ with } \sigma = \{y_1 \mapsto t_1, \dots\}.$$

$\Delta\sigma = \Delta$ because $S(f(y_1, \dots, y_k))$ and Δ do not share variables.

$$(B1) \quad S_1(x_1), \dots, S_n(x_n), \Gamma' \rightarrow \Delta, \Delta'.$$

The substitution $\{y \mapsto f(t_1, \dots, t_k)\}$ is irrelevant since $S(y)$ is the only literal with variable y .

$$(A2) \quad S_1(x_1), \dots, S_n(x_n), \Gamma'\tau \rightarrow \Delta, \Delta' \text{ with } \tau = \{y' \mapsto f(y_1, \dots, y_k)\}.$$

$\Delta'\tau = \Delta'$ because $y' \notin \text{vars}(\Delta')$.

(B2) $S_1(x_1), \dots, S_n(x_n), \Gamma' \rightarrow \Delta, \Delta'$.

(A3) $S_1(x_1), \dots, S_n(x_n), \Gamma'\tau \rightarrow S'(f(y_1, \dots, y_k)), \Delta, \Delta'$ with $\tau = \{y \mapsto f(y_1, \dots, y_k)\}$.

$\Delta'\tau = \Delta'$ because $y' \notin \text{vars}(\Delta')$.

(B3) $S_1(x_1), \dots, S_n(x_n), \Gamma' \rightarrow S'(y'), \Delta, \Delta'$.

In the constraint $(\pi_1 \wedge \pi_2)\sigma \downarrow$ the maximal depth of the subconstraints is less or equal to the maximal depth of π_1 or π_2 . Hence, d is also an upper bound on the constraint of the resolvent. In each case, the resolvent is again an MSL(SDC) clause.

In the first and second case, the multiset of term depths of the negative literals in R is strictly smaller than for the right parent. In both, the Γ is the same between the right parent and the resolvent. Only the $f(t_1, \dots, t_k)$ term is replaced by $x_1\sigma, \dots, x_n\sigma$ and x_1, \dots, x_n respectively. In the first case, the depth of the $x_i\sigma$ is either zero if $x_i \notin \{y_1, \dots, y_k\}$ or at least one less than $f(t_1, \dots, t_k)$ since $x_i\sigma = t_i$. In the second case, the x_i have depth zero which is strictly smaller than the depth of $f(t_1, \dots, t_k)$. Since the multiset ordering on natural numbers is terminating, the first and second case can only be applied finitely many times by constrained resolution.

In the third to sixth case R is of the form $(S_1(x_1), \dots, S_l(x_l), S'_1(t), \dots, S'_m(t) \rightarrow \Delta; \pi)$ or $(S_1(x_1), \dots, S_l(x_l), S'_1(t), \dots, S'_m(t) \rightarrow S(t), \Delta; \pi)$ with $t = f(y_1, \dots, y_k)$. By Lemma 4, there are only finitely many such clauses after condensation and removal of variants. Therefore, these four cases can apply only finitely many times during saturation.

Let $(\Gamma \rightarrow \Delta, S(t), S(t'); \pi)$ be a clause in N where sdc-factoring applies with $\sigma = \text{mgu}(S(t), S(t'))$ and $R = ((\Gamma \rightarrow \Delta, S(t))\sigma; \pi\sigma \downarrow)$. Because in $\Gamma \rightarrow \Delta, S(t), S(t')$ no literal is selected, $\Gamma \rightarrow \Delta, S(t), S(t')$ and $(\Gamma \rightarrow \Delta, S(t))\sigma$ can only match one of three patterns.

(A) $S_1(x_1), \dots, S_n(x_n) \rightarrow S(f(y_1, \dots, y_k)), S(f(z_1, \dots, z_l)), \Delta$

where $t = f(y_1, \dots, y_k)$, $t' = f(z_1, \dots, z_l)$, and $\{x_1, \dots, x_n\} \subseteq \{y_1, \dots, y_k\} \cup \{z_1, \dots, z_l\} \cup \text{vars}(\Delta)$. The result is

$S_1(x_1)\sigma, \dots, S_n(x_n)\sigma \rightarrow S(f(y_1, \dots, y_k)), \Delta$ with $\sigma = \{z_1 \mapsto y_1, \dots\}$.

(B) $S_1(x_1), \dots, S_n(x_n) \rightarrow S(f(y_1, \dots, y_k)), S(z), \Delta$

where $t = f(y_1, \dots, y_k)$, $t' = z$ and $\{x_1, \dots, x_n\} \subseteq \{y_1, \dots, y_k\} \cup \text{vars}(\Delta)$, i.e., z occurs only once. The result is

$S_1(x_1), \dots, S_n(x_n) \rightarrow S(f(y_1, \dots, y_k)), \Delta$.

(C) $S_1(x_1), \dots, S_n(x_n) \rightarrow S(y), S(z), \Delta$

where $t = y$, $t' = z$ and $\{x_1, \dots, x_n\} \subseteq \text{vars}(\Delta)$, i.e., y and z occur only once. The result is

$$S_1(x_1), \dots, S_n(x_n) \rightarrow S(y), \Delta.$$

In the new constraint $\pi\sigma \downarrow$ the maximal depth of the subconstraints is less or equal to the maximal depth of π . Hence d is also an upper bound on the constraint of the resolvent. In each case, the resolvent is again an MSL(SDC) clause.

Furthermore, in each case the clause is of the form $(S_1(x_1), \dots, S_l(x_l) \rightarrow \Delta; \pi)$. By Lemma 4, there are only finitely many such clauses after condensation and removal of variants. Therefore, these three cases can apply only finitely many times during saturation. \square

Theorem 1 (MSL(SDC) Decidability). *Satisfiability of the MSL(SDC) first-order fragment is decidable.*

Proof. Follows from Lemma 5 and 3.

4 Approximation and Refinement

In the following, we show how decidability of the MSL(SDC) fragment can be used to improve the approximation refinement calculus presented in [18].

Our approach is based on a counter-example guided abstraction refinement (CEGAR) idea. The procedure loops through four steps: approximation, testing (un)satisfiability, lifting, and refinement. The approximation step transforms any first-order logic clause set into the decidable MSL(SDC) fragment while preserving unsatisfiability. The second step employs the decidability result for MSL(SDC), Section 3, to test satisfiability of the approximated clause set. If the approximation is satisfiable, the original problem is satisfiable as well and we are done. Otherwise, the third step, lifting, tests whether the proof of unsatisfiability found for the approximated clause set can be lifted to a proof of the original clause set. If so, the original clause set is unsatisfiable and we are again done. If not, we extract a cause for the lifting failure that always amounts to two different instantiations of the same variable in a clause from the original clause set. This is resolved by the fourth step, the refinement. The crucial clause in the original problem is replaced and instantiated in a satisfiability preserving way such that the different instantiations do not reoccur anymore in subsequent iterations of the loop.

As mentioned before, our motivation to use dismatching constraints is that for an unconstrained clause the refinement adds quadratically many new clauses to the clause set. In contrast, with constrained clauses the same can be accomplished with adding just a single new clause. This extension is rather simple as constraints are treated the same as the antecedent literals in the clause. Furthermore we present refinement as a separate transformation rule.

The second change compared to the previous version is the removal of the Horn approximation rule, where we have now shown in Section 3 that a restriction to Horn clauses is not required for decidability anymore. Instead, the linear and shallow approximations are extended to apply to non-Horn clauses instead.

The approximation consists of individual transformation rules $N \Rightarrow N'$ that are non-deterministically applied. They transform a clause that is not in the MSL(SDC) fragment in finite steps into MSL(SDC) clauses. Each specific property of MSL(SDC) clauses, i.e, monadic predicates, shallow and linear positive literals, is generated by a corresponding rule: the Monadic transformation encodes non-Monadic predicates as functions, the shallow transformation extracts non-shallow subterms by introducing fresh predicates and the linear transformation renames non-linear variable occurrences.

Starting from a constrained clause set N the transformation is parameterized by a single monadic projection predicate T , fresh to N and for each non-monadic predicate P a separate projection function f_P fresh to N . The clauses in N are called the original clauses while the clauses in N' are the approximated clauses. We assume all clauses in N to be variable disjoint.

Definition 7. *Given a predicate P , projection predicate T , and projection function f_P , define the injective function $\mu_P^T(P(\vec{t})) := T(f_P(\vec{t}))$ and $\mu_P^T(Q(\vec{s})) := Q(\vec{s})$ for $P \neq Q$. The function is extended to [constrained] clauses, clause sets and interpretations. Given a signature Σ with non-monadic predicates P_1, \dots, P_n , define $\mu_\Sigma^T(N) := \mu_{P_1}^T(\dots(\mu_{P_n}^T(N))\dots)$ and $\mu_\Sigma^T(\mathcal{I}) := \mu_{P_1}^T(\dots(\mu_{P_n}^T(\mathcal{I}))\dots)$.*

Monadic $N \Rightarrow_{\text{MO}} \mu_P^T(N)$

provided P is a non-monadic predicate in the signature of N .

Shallow
$$\begin{aligned} N \dot{\cup} \{(\Gamma \rightarrow E[s]_p, \Delta; \pi)\} &\Rightarrow_{\text{SH}} \\ N \cup \{(S(x), \Gamma_l \rightarrow E[p/x], \Delta_l; \pi); (\Gamma_r \rightarrow S(s), \Delta_r; \pi)\} \end{aligned}$$

provided s is complex, $|p| = 2$, x and S fresh, $\Gamma_l\{x \mapsto s\} \cup \Gamma_r = \Gamma$, $\Delta_l \cup \Delta_r = \Delta$, $\{Q(y) \in \Gamma \mid y \in \text{vars}(E[p/x], \Delta_l)\} \subseteq \Gamma_l$, $\{Q(y) \in \Gamma \mid y \in \text{vars}(s, \Delta_r)\} \subseteq \Gamma_r$.

Linear 1
$$\begin{aligned} N \dot{\cup} \{(\Gamma \rightarrow \Delta, E'[x]_p, E[x]_q; \pi)\} &\Rightarrow_{\text{LI}} \\ N \cup \{(\Gamma\sigma, \Gamma \rightarrow \Delta, E'[x]_p, E[q/x']; \pi \wedge \pi\sigma)\} \end{aligned}$$

provided x' is fresh and $\sigma = \{x \mapsto x'\}$.

Linear 2
$$\begin{aligned} N \dot{\cup} \{(\Gamma \rightarrow \Delta, E[x]_{p,q}; \pi)\} &\Rightarrow_{\text{LI}} \\ N \cup \{(\Gamma\sigma, \Gamma \rightarrow \Delta, E[q/x']; \pi \wedge \pi\sigma)\} \end{aligned}$$

provided x' is fresh, $p \neq q$ and $\sigma = \{x \mapsto x'\}$.

Refinement $N \dot{\cup} \{(C, \pi)\} \Rightarrow_{\text{Ref}} N \cup \{(C; \pi \wedge x \neq t), (C; \pi)\{x \mapsto t\}\}$

provided $x \in \text{vars}(C)$, t straight and $\text{vars}(t) \cap \text{vars}((C, \pi)) = \emptyset$.

Note that variables are not renamed unless explicitly stated in the rule. This means that original clauses and their approximated counterparts share variable names. We use this to trace the origin of variables in the approximation.

The refinement transformation \Rightarrow_{Ref} is not needed to eventually generate MSL(SDC) clauses, but can be used to achieve a more fine-grained approximation of N , see below.

In the shallow transformation, Γ and Δ are separated into Γ_l , Γ_r , Δ_l , and Δ_r , respectively. The separation can be almost arbitrarily chosen as long as no atom from Γ , Δ is skipped. However, the goal is to minimize the set of shared variables, i.e., the variables of $(\Gamma \rightarrow E[s]_p, \Delta; \pi)$ that are inherited by both approximation clauses, $\text{vars}(\Gamma_r, s, \Delta_r) \cap \text{vars}(\Gamma_l, E[p/x], \Delta_l)$. If there are no shared variables, the shallow transformation is satisfiability equivalent. The conditions on Γ_l and Γ_r ensure that $S(x)$ atoms are not separated from the respective positive occurrence of x in subsequent shallow transformation applications.

Consider the clause $Q(f(x), y) \rightarrow P(g(f(x), y))$. The simple shallow transformation $S(x'), Q(f(x), y) \rightarrow P(g(x', y)); S(f(x))$ is not satisfiability equivalent – nor with any alternative partitioning of Γ . However, by replacing the occurrence of the extraction term $f(x)$ in $Q(f(x), y)$ with the fresh variable x' , the approximation $S(x'), Q(x', y) \rightarrow P(g(x', y)); S(f(x))$ is satisfiability equivalent. Therefore, we allow the extraction of s from the terms in Γ_l and require $\Gamma_l\{x \mapsto s\} \cup \Gamma_r = \Gamma$.

We consider Linear 1 and Linear 2 as two cases of the same linear transformation rule. Their only difference is whether the two occurrences of x are in the same literal or not. The duplication of literals and constraints in Γ and π is not needed if x does not occur in Γ or π .

Further, consider a linear transformation $N \cup \{(C; \pi)\} \Rightarrow_{\text{LI}} N \cup \{(C_a; \pi_a)\}$, where a fresh variable x' replaces an occurrence of a non-linear variable x in $(C; \pi)$. Then, $(C_a; \pi_a)\{x' \mapsto x\}$ is equal to $(C; \pi)$ modulo duplicate literal elimination. A similar property can be observed of a resolvent of $(C_l; \pi)$ and $(C_r; \pi)$ resulting from a shallow transformation $N \cup \{(C; \pi)\} \Rightarrow_{\text{SH}} N \cup \{(C_l; \pi), (C_r; \pi)\}$. Note that by construction, $(C_l; \pi)$ and $(C_r; \pi)$ are not necessarily variable disjoint. To simulate standard resolution, we need to rename at least the shared variables in one of them.

Definition 8 (\Rightarrow_{AP}). We define \Rightarrow_{AP} as the priority rewrite system [3] consisting of \Rightarrow_{Ref} , \Rightarrow_{MO} , \Rightarrow_{SH} and \Rightarrow_{LI} with priority $\Rightarrow_{\text{Ref}} > \Rightarrow_{\text{MO}} > \Rightarrow_{\text{SH}} > \Rightarrow_{\text{LI}}$, where \Rightarrow_{Ref} is only applied finitely many times.

Lemma 6 (\Rightarrow_{AP} is a Terminating Over-Approximation). (i) $\Rightarrow_{\text{AP}}^*$ terminates, (ii) if $N \Rightarrow_{\text{AP}} N'$ and N' is satisfiable, then N is also satisfiable.

Proof. (i) The transformations can be considered sequentially, because of the imposed rule priority. There are, by definition, only finitely many refinements at the beginning of an approximation $\Rightarrow_{\text{AP}}^*$. The monadic transformation strictly reduces the number of non-monadic atoms. The shallow transformation strictly reduces the multiset of term depths of the newly introduced clauses compared to

the removed parent clause. The linear transformation strictly reduces the number of duplicate variable occurrences in positive literals. Hence \Rightarrow_{AP} terminates.

(ii) Let $N \cup \{(C; \pi)\} \Rightarrow_{LI} N \cup \{(C_a; \pi_a)\}$ where an occurrence of a variable x in $(C; \pi)$ is replaced by a fresh x' . As $(C_a; \pi_a)\{x' \mapsto x\}$ is equal to $(C; \pi)$ modulo duplicate literal elimination, $\mathcal{I} \models (C; \pi)$ if $\mathcal{I} \models (C_a; \pi_a)$. Therefore, the linear transformation is an over-approximation.

Let $N \cup \{(C; \pi)\} \Rightarrow_{SH} N \cup \{(C_l; \pi_l), (C_r; \pi_r)\}$ and $(C_a; \pi_a)$ be the shallow ρ -resolvent. As $(C_a; \pi_a)\rho^{-1}$ equals $(C; \pi)$ modulo duplicate literal elimination, $\mathcal{I} \models (C; \pi)$ if $\mathcal{I} \models (C_l; \pi_l), (C_r; \pi_r)$. Therefore, the shallow transformation is an over-approximation.

Let $N \Rightarrow_{MO} \mu_P(N) = N'$. Then, $N = \mu_P^{-1}(N')$. Let \mathcal{I} be a model of N' and $(C; \pi) \in N$. Since $\mu_P((C; \pi)) \in N'$, $\mathcal{I} \models \mu_P((C; \pi))$ and thus, $\mu_P^{-1}(\mathcal{I}) \models (C; \pi)$. Hence, $\mu_P^{-1}(\mathcal{I})$ is a model of N . Therefore, the monadic transformation is an over-approximation. Actually, it is a satisfiability preserving transformation.

Let $N \cup \{(C; \pi)\} \Rightarrow_{Ref} N \cup \{(C; \pi \wedge x \neq t), (C; \pi)\{x \mapsto t\}\}$. Let $C\delta \in \mathcal{G}((C; \pi))$. If $x\delta$ is not an instance of t , then δ is a solution of $\pi \wedge x \neq t$ and $C\delta \in \mathcal{G}((C; \pi \wedge x \neq t))$. Otherwise, $\delta = \{x \mapsto t\}\delta'$ for some substitution δ' . Then, δ is a solution of $\pi\{x \mapsto t\}$ and thus, $C\delta = C\{x \mapsto t\}\delta' \in \mathcal{G}((C\{x \mapsto t\}; \pi\{x \mapsto t\}))$. Hence, $\mathcal{G}((C; \pi)) \subseteq \mathcal{G}((C; \pi \wedge x \neq t)) \cup \mathcal{G}((C; \pi)\{x \mapsto t\})$. Therefore, if \mathcal{I} is a model of $N \cup \{(C; \pi \wedge x \neq t), (C; \pi)\{x \mapsto t\}\}$, then \mathcal{I} is also a model of $N \cup \{(C; \pi)\}$. \square

Note that \Rightarrow_{Ref} and \Rightarrow_{MO} are also satisfiability preserving transformations.

Corollary 1. *If $N \Rightarrow_{AP}^* N'$ and N' is satisfied by a model \mathcal{I} , then $\mu_\Sigma^{-1}(\mathcal{I})$ is a model of N .*

Proof. Follows from Lemma 6 (ii)-(v). \square

On the basis of \Rightarrow_{AP} we can define an ancestor relation \Rightarrow_A that relates clauses, literal occurrences, and variables with respect to approximation. This relation is needed in order to figure out the exact clause, literal, variable for refinement.

Definition 9 (The Shallow Resolvent). *Let $N \cup \{(C; \pi)\} \Rightarrow_{SH} N \cup \{(C_l; \pi), (C_r; \pi)\}$ with $C = \Gamma \rightarrow E[s]_p, \Delta$, $C_l = S(x), \Gamma_l \rightarrow E[p/x], \Delta_l$ and $C_r = \Gamma_r \rightarrow S(s), \Delta_r$. Let x_1, \dots, x_n be the variables shared between C_l and C_r and $\rho = \{x_1 \mapsto x'_1, \dots, x_n \mapsto x'_n\}$ be a variable renaming with x'_1, \dots, x'_n fresh in C_l and C_r . We define $(\Gamma_l\{x \mapsto s\rho\}, \Gamma_r\rho \rightarrow E[p/s\rho], \Delta_l, \Delta_r\rho; \pi \wedge \pi\rho)$ as the shallow ρ -resolvent.*

Let $(C_a; \pi_a)$ be the shallow ρ -resolvent of $N \cup \{(C; \pi)\} \Rightarrow_{SH} N \cup \{(C_l; \pi), (C_r; \pi)\}$. Note that for any two ground instances $C_l\delta_l$ and $C_r\delta_r$, their resolvent is a ground instance of $(C_a; \pi_a)$. Furthermore, using the reverse substitution $\rho^{-1} = \{x'_1 \mapsto x_1, \dots, x'_n \mapsto x_n\}$, $(C_a; \pi_a)\rho^{-1} = (\Gamma_l\{x \mapsto s\}, \Gamma_r \rightarrow E[s]_p, \Delta_l, \Delta_r; \pi \wedge \pi)$ is equal to $(C; \pi)$ modulo duplicate literal elimination. This is because, $\Delta_l \cup \Delta_r = \Delta$ and $\Gamma_l\{x \mapsto s\} \cup \Gamma_r = \Gamma$ by definition of \Rightarrow_{SH} and $\pi \wedge \pi$ is equivalent to π .

Next, we establish parent relations that link original and approximated clauses, as well as their variables and literals. Together the parent, variable and literal

relations will allow us to not only trace any approximated clause back to their origin, but also predict what consequences changes to the original set will have on its approximations.

For the following definitions, we assume that clause and literal sets are lists and that μ_P^T and substitutions act as mappings. This means we can uniquely identify clauses and literals by their position in those lists. Further, for every shallow transformation $N \Rightarrow_{\text{SH}} N'$, we will also include the shallow resolvent in the parent relation as if it were a member of N' .

Definition 10 (Parent Clause). *For an approximation step $N \Rightarrow_{\text{AP}} N'$ and two clauses $(C; \pi) \in N$ and $(C'; \pi') \in N'$, we define $[(C; \pi), N] \Rightarrow_{\text{A}} [(C'; \pi'), N']$ expressing that $(C; \pi)$ in N is the parent clause of $(C'; \pi')$ in N' :*

If $N \Rightarrow_{\text{MO}} \mu_P^T(N)$, then

$$[(C; \pi), N] \Rightarrow_{\text{A}} [\mu_P^T((C; \pi)), \mu_P^T(N)] \text{ for all } (C; \pi) \in N.$$

If $N = N'' \cup \{(C; \pi)\} \Rightarrow_{\text{SH}} N'' \cup \{(C_l; \pi_l), (C_r; \pi_r)\} = N'$, then

$$[(D; \pi'), N] \Rightarrow_{\text{A}} [(D; \pi'), N'] \text{ for all } (D; \pi') \in N'' \text{ and}$$

$$[(C; \pi), N] \Rightarrow_{\text{A}} [(C_l; \pi_l), N'],$$

$$[(C; \pi), N] \Rightarrow_{\text{A}} [(C_r; \pi_r), N'] \text{ and}$$

$$[(C; \pi), N] \Rightarrow_{\text{A}} [(C_a; \pi_a), N'] \text{ for any shallow resolvent } (C_a; \pi_a).$$

If $N = N'' \cup \{(C; \pi)\} \Rightarrow_{\text{LI}} N'' \cup \{(C_a; \pi_a)\} = N'$, then

$$[(D; \pi'), N] \Rightarrow_{\text{A}} [(D; \pi'), N'] \text{ for all } (D; \pi') \in N'' \text{ and}$$

$$[(C; \pi), N] \Rightarrow_{\text{A}} [(C_a; \pi_a), N'].$$

If $N = N'' \cup \{(C; \pi)\} \Rightarrow_{\text{Ref}} N'' \cup \{(C; \pi \wedge x \neq t), (C; \pi)\{x \mapsto t\}\} = N'$, then

$$[(D; \pi'), N] \Rightarrow_{\text{A}} [(D; \pi'), N'] \text{ for all } (D; \pi') \in N'' ,$$

$$[(C; \pi), N] \Rightarrow_{\text{A}} [(C; \pi \wedge x \neq t), N'] \text{ and}$$

$$[(C; \pi), N] \Rightarrow_{\text{A}} [(C; \pi)\{x \mapsto t\}, N'].$$

Definition 11 (Parent Variable). *Let $N \Rightarrow_{\text{AP}} N'$ be an approximation step and $[(C; \pi), N] \Rightarrow_{\text{A}} [(C'; \pi'), N']$. For two variables x and y , we define $[x, (C; \pi), N] \Rightarrow_{\text{A}} [y, (C'; \pi'), N']$ expressing that $x \in \text{vars}(C)$ is the parent variable of $y \in \text{vars}(C')$:*

If $x \in \text{vars}((C; \pi)) \cap \text{vars}((C'; \pi'))$, then

$$[x, (C; \pi), N] \Rightarrow_{\text{A}} [x, (C'; \pi'), N'].$$

If $N \Rightarrow_{\text{SH}} N'$ and $(C'; \pi')$ is the shallow ρ -resolvent,

$$[x_i, (C; \pi), N] \Rightarrow_{\text{A}} [x_i \rho, (C'; \pi'), N'] \text{ for each } x_i \text{ in the domain of } \rho.$$

If $N \Rightarrow_{\text{LI}} N'$, $C = \Gamma \rightarrow \Delta[x]_{p,q}$ and $C' = \Gamma\{x \mapsto x'\}$, $\Gamma \rightarrow \Delta[q/x']$, then

$$[x, (C; \pi), N] \Rightarrow_{\text{A}} [x', (C'; \pi'), N'].$$

Note that if $N \Rightarrow_{\text{SH}} N'$ and x is the fresh extraction variable in $(C_l; \pi_l)$, then x has no parent variable. For literals, we actually further specify the relation on the positions within literals of a clause $(C; \pi)$ using pairs (L, r) of literals and positions. We write $(L, r) \in C$ to denote that (L, r) is a literal position in $(C; \pi)$ if $L \in C$ and $r \in \text{pos}(L)$. Note that a literal position (L, r) in $(C; \pi)$ corresponds to the term $L|_r$.

Definition 12 (Parent literal position). *Let $N \Rightarrow_{\text{AP}} N'$ be an approximation step and $[(C; \pi), N] \Rightarrow_{\text{A}} [(C'; \pi'), N']$. For two literal positions (L, r) and (L', r') , we define $[r, L, (C; \pi), N] \Rightarrow_{\text{A}} [r', L', (C'; \pi'), N']$ expressing that (L, r)*

in $(C; \pi)$ is the parent literal position of (L', r') in $(C'; \pi')$:
 If $(C; \pi) = (C'; \pi')$, then
 $[r, L, (C; \pi), N] \Rightarrow_A [r, L, (C'; \pi'), N']$ for all $(L, r) \in C$.
 If $N \Rightarrow_{\text{Ref}} N'$ and $(C', \pi') = (C; \pi \wedge x \neq t)$, then
 $[r, L, (C; \pi), N] \Rightarrow_A [r, L, (C'; \pi'), N']$ for all $(L, r) \in C$.
 If $N \Rightarrow_{\text{Ref}} N'$ and $(C', \pi') = (C; \pi)\{x \mapsto t\}$, then
 $[r, L, (C; \pi), N] \Rightarrow_A [r, L\{x \mapsto t\}, (C'; \pi'), N']$ for all $(L, r) \in C$.
 If $N \Rightarrow_{\text{MQ}} \mu_P^T(N) = N'$, then
 $[\varepsilon, P(\vec{t}), (C; \pi), N] \Rightarrow_A [\varepsilon, T(f_p(\vec{t})), (C'; \pi'), N']$ for all $P(\vec{t}) \in C$ and
 $[r, P(\vec{t}), (C; \pi), N] \Rightarrow_A [1.r, T(f_p(\vec{t})), (C'; \pi'), N']$ for all $(P(\vec{t}), r) \in C$.
 If $N \Rightarrow_{\text{SH}} N'$, $C = \Gamma \rightarrow E[s]_p, \Delta$ and $C' = S(x), \Gamma_l \rightarrow E[p/x], \Delta_l$, then
 $[r, E[s]_p, (C; \pi), N] \Rightarrow_A [r, E[p/x], (C'; \pi'), N']$ for all $r \in \text{pos}(E[p/x])$,
 $[p, E[s]_p, (C; \pi), N] \Rightarrow_A [r, S(x), (C'; \pi'), N']$ for all $r \in \text{pos}(S(x))$,
 $[r, L\{x \mapsto s\}, (C; \pi), N] \Rightarrow_A [r, L, (C'; \pi'), N']$ for all $(L, r) \in \Gamma_l$,
 $[r, L, (C; \pi), N] \Rightarrow_A [r, L, (C'; \pi'), N']$ for all $(L, r) \in \Delta_l$.
 If $N \Rightarrow_{\text{SH}} N'$, $C = \Gamma \rightarrow E[s]_p, \Delta$ and $C' = \Gamma_r \rightarrow S(s), \Delta_r$, then
 $[p, E[s]_p, (C; \pi), N] \Rightarrow_A [\varepsilon, S(s), (C'; \pi'), N']$,
 $[pr, E[s]_p, (C; \pi), N] \Rightarrow_A [1.r, S(s), (C'; \pi'), N']$ for all $r \in \text{pos}(s)$, and
 $[r, L, (C; \pi), N] \Rightarrow_A [r, L, (C'; \pi'), N']$ for all $(L, r) \in \Gamma_r \cup \Delta_r$.
 If $N \Rightarrow_{\text{SH}} N'$, $C = \Gamma \rightarrow E[s]_p, \Delta$ and (C', π') is the shallow ρ -resolvent, then
 $[r, E[s]_p, (C; \pi), N] \Rightarrow_A [r, E[p/s\rho], (C'; \pi'), N']$ for all $r \in \text{pos}(E[p/s\rho])$,
 $[r, L\{x \mapsto s\}, (C; \pi), N] \Rightarrow_A [r, L\{x \mapsto s\rho\}, (C'; \pi'), N']$ for all $(L, r) \in \Gamma_l$,
 $[r, L, (C; \pi), N] \Rightarrow_A [r, L\rho, (C'; \pi'), N']$ for all $(L, r) \in \Gamma_r \cup \Delta_r$, and
 $[r, L, (C; \pi), N] \Rightarrow_A [r, L, (C'; \pi'), N']$ for all $(L, r) \in \Delta_l$.
 If $N \Rightarrow_{\text{LI}} N'$, $C = \Gamma \rightarrow \Delta, E'[x]_p, E[x]_q$ and $C' = \Gamma\{x \mapsto x'\}, \Gamma \rightarrow \Delta, E'[x]_p, E[q/x']$,
 $[r, E'[x]_p, (C; \pi), N] \Rightarrow_A [r, E'[x]_p, (C'; \pi'), N']$ for all $r \in \text{pos}(E'[x]_p)$,
 $[r, E[x]_q, (C; \pi), N] \Rightarrow_A [r, E[q/x'], (C'; \pi'), N']$ for all $r \in \text{pos}(E[q/x'])$,
 $[r, L, (C; \pi), N] \Rightarrow_A [r, L\{x \mapsto x'\}, (C'; \pi'), N']$ for all $(L, r) \in \Gamma$,
 $[r, L, (C; \pi), N] \Rightarrow_A [r, L, (C'; \pi'), N']$ for all $(L, r) \in \Gamma$, and
 $[r, L, (C; \pi), N] \Rightarrow_A [r, L, (C'; \pi'), N']$ for all $(L, r) \in \Delta$.
 If $N \Rightarrow_{\text{LI}} N'$, $C = \Gamma \rightarrow \Delta, E[x]_{p,q}$ and $C' = \Gamma\{x \mapsto x'\}, \Gamma \rightarrow \Delta, E[q/x']$, then
 $[r, E[x]_{p,q}, (C; \pi), N] \Rightarrow_A [r, E[q/x'], (C'; \pi'), N']$ for all $r \in \text{pos}(E[q/x'])$,
 $[r, L, (C; \pi), N] \Rightarrow_A [r, L\{x \mapsto x'\}, (C'; \pi'), N']$ for all $(L, r) \in \Gamma$,
 $[r, L, (C; \pi), N] \Rightarrow_A [r, L, (C'; \pi'), N']$ for all $(L, r) \in \Gamma$, and
 $[r, L, (C; \pi), N] \Rightarrow_A [r, L, (C'; \pi'), N']$ for all $(L, r) \in \Delta$.

The transitive closures of each parent relation are called ancestor relations.

The over-approximation of a clause set N can introduce resolution refutations that have no corresponding equivalent in N which we consider a lifting failure. Compared to our previous calculus [18], the lifting process is identical with the exception that there is no case for the removed Horn transformation. We only update the definition of conflicting cores to consider constrained clauses.

Definition 13 (Conflicting Core). A finite set of unconstrained clauses and a solvable constraint $(N^\perp; \pi)$ are a conflicting core if $N^\perp \delta$ is unsatisfiable for




Fig. 1. Visual representation of the parent literal position relation (Definition 12)

all solutions δ of π over $\text{vars}(N^\perp) \cup \text{lvar}(\pi)$. A conflicting core $(N^\perp; \pi)$ is a conflicting core of the constrained clause set N if for every $C \in N^\perp$ there is a clause $(C'; \pi') \in N$ such that $(C; \pi)$ is an instance of $(C'; \pi')$ modulo duplicate literal elimination. The clause $(C'; \pi')$ is then called the instantiated clause of $(C; \pi)$ in $(N^\perp; \pi)$. We call $(N^\perp; \pi)$ complete if for every clause $C \in N^\perp$ and literal $L \in C$, there exists a clause $D \in N^\perp$ with $\bar{L} \in D$.

A conflicting core is a generalization of a ground unsatisfiability core that allows global variables to act as parameters. This enables more efficient lifting and refinement compared to a simple ground unsatisfiable core. We show some examples at the end of this section.

We discuss the potential lifting failures and the corresponding refinements only for the linear and shallow case because lifting the satisfiability equivalent monadic and refinement transformations always succeeds. To reiterate from our previous work: in the linear case, there exists a clause in the conflicting core that is not an instance of the original clauses. In the shallow case, there exists a pair of clauses whose resolvent is not an instance of the original clauses. We combine these two cases by introducing the notion of a lift-conflict.

Definition 14 (Conflict). Let $N \cup \{(C, \pi)\} \Rightarrow_{\text{LI}} N \cup \{(C_a, \pi_a)\}$ and N^\perp be a complete ground conflicting core of $N \cup \{(C_a, \pi_a)\}$. We call a conflict clause $C_c \in N^\perp$ with the instantiated clause (C_a, π_a) a lift-conflict if C_c is not an instance of (C, π) modulo duplicate literal elimination. Then, C_c is an instance of (C_a, π_a) , which we call the conflict clause of C_c .

Let $N \cup \{(C, \pi)\} \Rightarrow_{\text{SH}} N \cup \{(C_l, \pi_l), (C_r, \pi_r)\}$, $(C_a; \pi_a)$ be the shallow resolvent and N^\perp be a complete ground conflicting core of $N \cup \{(C_l, \pi_l), (C_r, \pi_r)\}$. We call the resolvent C_c of $C_l \delta_l \in N^\perp$ and $C_r \delta_r \in N^\perp$ a lift-conflict if C_c is not an instance of (C, π) modulo duplicate literal elimination. Then, C_c is an instance of $(C_a; \pi_a)$, which we call the conflict clause of C_c .

The goal of refinement is to instantiate the original parent clause in such a way that is both satisfiability equivalent and prevents the lift-conflict after approximation. Solving the refined approximation will then either necessarily produce a complete saturation or a new refutation proof, because its conflicting core has to be different. For this purpose, we use the refinement transformation to segment the original parent clause $(C; \pi)$ into two parts $(C; \pi \wedge x \neq t)$ and $(C; \pi)\{x \mapsto t\}$.

For example, consider N and its linear transformation N' .

$$\begin{array}{ccc} \rightarrow P(x, x) & \Rightarrow_{\text{LI}} & \rightarrow P(x, x') \\ P(a, b) \rightarrow & \Rightarrow_{\text{AP}}^0 & P(a, b) \rightarrow \end{array}$$

The ground conflicting core of N' is

$$\begin{array}{c} \rightarrow P(a, b) \\ P(a, b) \rightarrow \end{array}$$

Because $P(a, b)$ is not an instance of $P(x, x)$, lifting fails. $P(a, b)$ is the lift-conflict. Specifically, $\{x \mapsto a\}$ and $\{x \mapsto b\}$ are conflicting substitutions for the parent variable x . We pick $\{x \mapsto a\}$ to segment $P(x, x)$ into $(P(x, x); x \neq a)$ and $P(x, x)\{x \mapsto a\}$. Now, any descendant of $(P(x, x); x \neq a)$ cannot have a at the position of the first x , and any descendant of $P(x, x)\{x \mapsto a\}$ must have an a at the position of the second x . Thus, $P(a, b)$ is excluded in both cases and no longer appears as a lift-conflict.

To show that the lift-conflict will not reappear in the general case, we use that the conflict clause and its ancestors have strong ties between their term structures and constraints.

Definition 15 (Constrained Term Skeleton). *The constrained term skeleton of a term t under constraint π , $\text{skt}(t, \pi)$, is defined as the normal form of the following transformation:*

$$(t[x]_{p,q}; \pi) \Rightarrow_{\text{skt}} (t[q/x']; \pi \wedge \pi\{x \mapsto x'\}), \text{ where } p \neq q \text{ and } x' \text{ is fresh.}$$

The constrained term skeleton of a term t is essentially a linear version of t where the restrictions on each variable position imposed by π are preserved. For (t, π) and a solution δ of π , $t\delta$ is called a ground instance of (t, π) .

Lemma 7. *Let $N_0 \Rightarrow_{\text{AP}}^* N_k$, $(C_k; \pi_k)$ in N with the ancestor clause $(C_0; \pi_0) \in N_0$ and N_k^\perp be a complete ground conflicting core of N_k . Let δ be a solution of π_k such that $C_k\delta$ is in N_k^\perp . If (L', q') is a literal position in $(C_k; \pi_k)$ with the ancestor (L, q) in (C_0, π_0) , then (i) $L'\delta|_{q'}$ is an instance of $\text{skt}(L|_q, \pi_0)$, (ii) $q = q'$ if L and L' have the same predicate, and (iii) if $L'|_{q'} = x$ and there exists an ancestor variable y of x in (C_0, π_0) , then $L|_q = y$.*

Proof. By induction on the length of the approximation $N_0 \Rightarrow_{\text{AP}}^* N_k$.

The base case $N_k = N_0$, is trivial. Let $N_0 = N \cup \{(C; \pi)\} \Rightarrow_{\text{SH}} N \cup \{(C_l; \pi_l), (C_r; \pi_r)\} = N_k$, $(C_k; \pi_k)$ be the shallow ρ -resolvent and $C_k\delta$ be the resolvent of two instances of $(C_l; \pi_l)$ and $(C_r; \pi_r)$ in N_k^\perp . Then, $(C_k; \pi_k)\rho^{-1}$ is equal to $(C; \pi)$ modulo duplicate literal elimination. Thus, by definition $(L, q) = (L', q')\rho^{-1}$. Therefore, (i) $L'\delta|_{q'}$ is an instance of $\text{skt}(L|_q, \pi_0)$, (ii) $q = q'$ if L and

L' have the same predicate, and (iii) if $L'|_{q'} = x$ and there exists an ancestor variable y of x in (C_0, π_0) , then $L|_q = y$.

Now, let $N_0 \Rightarrow_{AP} N_1 \Rightarrow_{AP}^* N_k$. Since (L', p) has an ancestor literal position in (C_0, π_0) , the ancestor clause of $(C_k; \pi_k)$ in $N_1, (C_1, \pi_1)$, contains the ancestor literal position (L_1, q_1) , which has (L, q) as its parent literal position. By the induction hypothesis on $N_1 \Rightarrow_{AP}^* N_k$, (i) $L'\delta|_{q'}$ is an instance of $\text{skt}(L_1|_{q_1}, \pi_1)$, (ii) $q_1 = q'$ if L_1 and L' have the same predicate, and (iii) if $L'|_{q'} = x$ and there is an ancestor variable y_1 of x in (C_1, π_1) , then $L_1|_{q_1} = y_1$.

Let $N_0 = N \cup \{(C; \pi)\} \Rightarrow_{\text{Ref}} N \cup \{(C; \pi \wedge x \neq t), (C; \pi)\{x \mapsto t\}\} = N_1$. If (C_1, π_1) is neither $(C; \pi \wedge x \neq t)$ nor $(C; \pi)\{x \mapsto t\}$, then trivially $(C_0, \pi_0) = (C_1, \pi_1)$. Otherwise, $(C_1, \pi_1) = (C; \pi \wedge x \neq t)$ or $(C_1, \pi_1) = (C; \pi)\{x \mapsto t\}$. Then $(L_1, q_1) = (L, q)$ or $(L_1, q_1) = (L, q)\{x \mapsto t\}$. In either case, (i) $L'\delta|_{q'}$ is an instance of $\text{skt}(L|_q, \pi_0)$, (ii) $q = q'$ if L and L' have the same predicate, and (iii) if $L'|_{q'} = x$ and there exists an ancestor variable y of x in (C_0, π_0) , then $L|_q = y$.

Let $N_0 \Rightarrow_{MO} \mu_P(N) = N_1$. If P is not the predicate of L , then trivially $(L, q) = (L_1, q_1)$. If P is the predicate of L , then $(L, q) = (P(t_1, \dots, t_n), q)$ and $(L_1, q_1) = (T(f_p(t_1, \dots, t_n)), 1.q)$. Thus, (i) $L'\delta|_{q'}$ is an instance of $\text{skt}(L|_q, \pi_0) = \text{skt}(T(f_p(t_1, \dots, t_n))|_{1.q}, \pi_0)$. (ii) The predicate of L' is not P by definition. (iii) Let $L'|_{q'} = x$ and y be the ancestor variable of x in (C_0, π_0) . Then, y is also the ancestor variable of x in (C_1, π_1) and $L_1|_{q_1} = y$. Therefore, $L|_q = P(t_1, \dots, t_n)|_q = T(f_p(t_1, \dots, t_n))|_{1.q} = L_1|_{q_1} = y$.

Let $N_0 = N \cup \{(C; \pi)\} \Rightarrow_{LI} N \cup \{(C_a; \pi_a)\} = N_1$ where an occurrence of a variable x is replaced by a fresh x' . If $(C_1, \pi_1) \neq (C_a; \pi_a)$, then trivially $(C_0, \pi_0) = (C_1, \pi_1)$. Otherwise, $(C_1, \pi_1) = (C_a; \pi_a)$, $(C_0, \pi_0) = (C; \pi)$. By definition, $(L, q) = (L_1\{x' \mapsto x\}, q_1)$ and $\pi_0 = \pi_1\{x' \mapsto x\}$. Thus, $\text{skt}(L|_q, \pi_0) = \text{skt}(L_1|_{q_1}, \pi_1)$. Therefore, $L'\delta|_{q'}$ is an instance of $\text{skt}(L|_q, \pi_0)$. Since L and L_1 have the same predicate and $q = q_1$, $q = q'$ if L and L' have the same predicate. Let $L'|_{q'} = z$ and y be the ancestor variable of z in (C_1, π_1) . If $y \neq x'$, then y is the ancestor variable of z in (C_0, π_0) and $L|_q = L_1\{x' \mapsto x\}|_{q_1} = y_1$. Otherwise, x is the ancestor variable of z in (C_0, π_0) and $L|_q = L_1\{x' \mapsto x\}|_{q_1} = x$.

Let $N_0 = N \cup \{(C; \pi)\} \Rightarrow_{SH} N \cup \{(C_l; \pi_l), (C_r; \pi_r)\} = N_1$ where a term s is extracted from a positive literal $Q(s'[s]_p)$ via introduction of fresh predicate S and variable x . If (C_1, π_1) is neither $(C_l; \pi_l)$ nor $(C_r; \pi_r)$, then trivially $(C_0, \pi_0) = (C_1, \pi_1)$.

If $(C_1, \pi_1) = (C_l; \pi_l)$ and $L_1 = S(x)$, then $(C_0, \pi_0) = (C; \pi)$, $q_1 = 1$, $(L', q') = (S(x), 1)$ and $(Q(s'[s]_p), 1.p)$ is the parent literal position of $(S(x), 1)$. Let $L'\delta = S(t)$. Because N_k^\perp is complete and ground, there is a clause $C'_k\delta' \in N_k^\perp$ that contains the positive literal $S(t)$. The ancestor of $(C'_k, \pi'_k) \in N_k$ in N_1 is $(C_r; \pi_r)$ because it is the only clause in N_1 with a positive S -literal. Then, by the inductive hypothesis, $(S(s), 1)$ in $(C_r; \pi_r)$ is the ancestor literal position of $(S(x), 1)$ in (C'_k, π'_k) . Thus, t is an instance of $\text{skt}(S(s)|_1, \pi_r) = \text{skt}(s, \pi_r)$. Therefore, $t = L'\delta|_{q'}$ is an instance of $\text{skt}(Q(s'[s]_p)|_{1.p}, \pi) = \text{skt}(s, \pi_r)$. Further, Q and S are not the same predicate because S is fresh. Since x has no parent variable, $L'|_{q'} = x$ has no ancestor variable in (C_0, π_0) .

If $(C_1, \pi_1) = (C_l; \pi_l)$ and $L_1 = Q(s'[p/x])$, then $(C_0, \pi_0) = (C; \pi)$ and $(Q(s'[s]_p), q_1)$ in $(C; \pi)$ is the parent literal position of (L_1, q_1) in (C_1, π_1) and ancestor literal position of (L', q') in (C_k, π_k) . If q_1 is not a position at or above p , the subterm at p is irrelevant and thus $\text{skt}(Q(s'[s]_p)|_{q_1}, \pi) = \text{skt}(Q(s'[p/x])|_{q_1}, \pi_l)$. Otherwise, let r be a position such that $q_1 r = 1.p$. Since $|p| = 2$, no following shallow transformation step extracts a subterm of $s'[p/x]$ containing x . Thus by definition of \Rightarrow_{AP} , $L' = Q(t'[x]_p)$ and C_k also contains the negative literal $S(x)$. Let $S(x)\delta = S(t)$. Analogously to the previous case, t is an instance of $\text{skt}(s, \pi_r)$. Combined with $L'\delta|_{q'}$ being an instance of $\text{skt}(L_1|_{q_1}, \pi_1) = \text{skt}(Q(s'[p/x])|_{q_1}, \pi_l)$ and $L'\delta|_{1.p} = t$, $L'\delta|_{q'}$ is an instance of $\text{skt}(Q(s'[s]_p)|_q, \pi)$. Since L and L_1 have the same predicate and $q = q_1$, $q = q'$ if L and L' have the same predicate. Let $L'|_{q'} = z$ and y in (C_1, π_1) be the ancestor variable of z in (C_k, π_k) . Since x has no parent, $y \neq x$ and y in (C_0, π_0) is the ancestor variable of z . Therefore, $Q(s'[s]_p)|_{q_1} = y$ because $Q(s'[p/x])|_{q_1} = y$.

If $(C_1, \pi_1) = (C_r; \pi_r)$ and $L_1 = S(s)$, let $q_1 = 1.q'_1$. Then, $(C_0, \pi_0) = (C; \pi)$ and $(L, q) = (Q(s'[s]_p), 1.pq'_1)$ in (C_0, π_0) is the parent literal position of (L_1, q_1) in (C_1, π_1) . Thus, $L'\delta|_{q'}$ is an instance of $\text{skt}((Q(s'[s]_p)|_{1.pq'_1}, \pi) = \text{skt}(s|_{q'_1}, \pi) = \text{skt}(L_1|_{q_1}, \pi_r)$. Because S is fresh, Q is not the predicate of L' . Let $L'|_{q'} = z$ and y in (C_1, π_1) be the ancestor variable of z in (C_k, π_k) . Then, y in (C_0, π_0) is the ancestor variable of z and $Q(s'[s]_p)|_q = s|_{q'_1} = y$ because $s|_{q'_1} = L_1|_{q_1} = y$.

Otherwise, (L_1, q_1) in (C_0, π_0) is the parent literal position of (L_1, q_1) in (C_1, π_1) , by definition. Then, $\text{skt}(L_1, \pi) = \text{skt}(L_1, \pi_l)$ or $\text{skt}(L_1, \pi) = \text{skt}(L_1, \pi_r)$, respectively. \square

Next, we define the notion of descendants and descendant relations to connect lift-conflicts in ground conflicting cores with their corresponding ancestor clauses. The goal, hereby, is that if a ground clause D is not a descendant of a clause in N , then it can never appear in a conflicting core of an approximation of N .

Definition 16 (Descendants). Let $N \Rightarrow_{\text{AP}}^* N'$, $[(C; \pi), N] \Rightarrow_{\text{A}}^* [(C'; \pi'), N']$ and D be a ground instance of $(C'; \pi')$. Then, we call D a descendant of $(C; \pi)$ and define the $[(C; \pi), N] \Rightarrow_{\text{A}}^* [(C'; \pi'), N']$ -descendant relation \Rightarrow_D that maps literals in D to literal positions in $(C; \pi)$ using the following rule:

$$L'\delta \Rightarrow_D (L, r) \text{ if } L'\delta \in D \text{ and } [r, L, (C; \pi), N] \Rightarrow_{\text{A}}^* [\varepsilon, L', (C'; \pi'), N']$$

For the descendant relations it is of importance to note that while there are potentially infinite ways that a lift-conflict C_c can be a descendant of an original clause $(C; \pi)$, there are only finitely many distinct descendant relations over C_c and $(C; \pi)$. This means, if a refinement transformation can prevent one distinct descendant relation without generating new distinct descendant relations (Lemma 8), a finite number of refinement steps can remove the lift-conflict C_c from the descendants of $(C; \pi)$ (Lemma 9). Thereby, preventing any conflicting cores containing C_c from being found again.

A clause $(C; \pi)$ can have two descendants that are the same except for the names of the S -predicates introduced by shallow transformations. Because the

used approximation $N \Rightarrow_{\text{AP}}^* N'$ is arbitrary and therefore also the choice of fresh S -predicates, if D is a descendant of $(C; \pi)$, then any clause D' equal to D up to a renaming of S -predicates is also a descendant of $(C; \pi)$. On the other hand, the actual important information about an S -predicate is which term it extracts. Two descendants of $(C; \pi)$ might be identical but their S -predicate extract different terms in $(C; \pi)$. For example, $P(a) \rightarrow S(f(a))$ is a descendant of $P(x), P(y) \rightarrow Q(f(x), g(f(x)))$ but might extract either occurrence of $f(x)$. These cases are distinguished by their respective descendant relations. In the example, we have either $S(f(a)) \Rightarrow_D (Q(f(x), g(f(x))), 1)$ or $S(f(a)) \Rightarrow_D (Q(f(x), g(f(x))), 2.1)$.

Lemma 8. *Let $N_0 = N \cup \{(C; \pi)\} \Rightarrow_{\text{Ref}} N \cup \{(C; \pi \wedge x \neq t), (C; \pi)\{x \mapsto t\}\} = N_1$ be a refinement transformation and D a ground clause. If there is a $[(C; \pi \wedge x \neq t), N_1] \Rightarrow_{\text{A}}^* [(C'; \pi'), N_2]$ - or $[(C; \pi)\{x \mapsto t\}, N_1] \Rightarrow_{\text{A}}^* [(C'; \pi'), N_2]$ -descendant relation \Rightarrow_D^1 , then there is an equal $[(C; \pi), N_0] \Rightarrow_{\text{A}}^* [(C'; \pi'), N_2]$ -descendant relation \Rightarrow_D^0 .*

Proof. Let L_D be a literal of D and $L' \Rightarrow_D^1 (L, r)$. If D is a descendant of $(C; \pi \wedge x \neq t)$, then $[r, L, (C; \pi \wedge x \neq t), N_1] \Rightarrow_{\text{A}}^* [\varepsilon, L', (C'; \pi'), N_2]$. Because $[r, L, (C; \pi), N_0] \Rightarrow_{\text{A}} [r, L, (C; \pi \wedge x \neq t), N_1]$, $L' \Rightarrow_D^0 (L, r)$. If D is a descendant of $(C; \pi)\{x \mapsto t\}$, the proof is analogous. \square

Lemma 9 (Refinement). *Let $N \Rightarrow_{\text{AP}} N'$ and N^\perp be a complete ground conflicting core of N' . If $C_c \in N^\perp$ is a lift-conflict, then there exists a finite refinement $N \Rightarrow_{\text{Ref}}^* N_R$ such that for any approximation $N_R \Rightarrow_{\text{AP}}^* N'_R$ and ground conflicting core N_R^\perp of N'_R , C_c is not a lift-conflict in N_R^\perp modulo duplicate literal elimination.*

Proof. Let (C_a, π_a) be the conflict clause of C_c and $(C; \pi) \in N$ be the parent clause of (C_a, π_a) . C_c is a descendant of $(C; \pi)$ with the corresponding $[(C; \pi), N] \Rightarrow_{\text{A}}^* [(C_a; \pi_a), N']$ -descendant relation $\Rightarrow_{C_c}^0$. We apply induction on the number of distinct $[(C; \pi), N] \Rightarrow_{\text{A}}^* [(C'; \pi'), N'']$ -descendant relations \Rightarrow_{C_c} for arbitrary approximations $N \Rightarrow_{\text{AP}}^* N''$.

Since only the shallow and linear transformations can produce lift-conflicts, the clause $(C; \pi)$ is replaced by either a linearized clause $(C'; \pi')$ or two shallow clauses $(C_l; \pi)$ and $(C_r; \pi)$. Then, the conflict clause $(C_a; \pi_a)$ of C_c is either the linearized $(C'; \pi')$ or the resolvent of $(C_l; \pi)$ and $(C_r; \pi)$. In either case, $C_c = C_a \delta$ for some solution δ of π_a . Furthermore, there exists a substitution $\tau = \{x'_1 \mapsto x_1, \dots, x'_n \mapsto x_n\}$ such that $(C; \pi)$ and $(C_a; \pi_a)\tau$ are equal modulo duplicate literal elimination. That is, $\tau = \{x' \mapsto x\}$ for a linear transformation and $\tau = \rho^{-1}$ for shallow transformation (Definition 9).

Assume $C_c = C_a \tau \sigma$ for some grounding substitution σ , where $\tau \sigma$ is a solution of π_a . Thus, σ is a solution of $\pi_a \tau$, which is equivalent to π . Then, C_c is equal to $C \sigma$ modulo duplicate literal elimination an instance of $(C; \pi)$, which contradicts with C_c being a lift-conflict. Hence, $C_c = C_a \delta$ is not an instance of $C_a \tau$ and thus, $x_i \delta \neq x'_i \delta$ for some x_i in the domain of τ .

Because $x_i\delta$ and $x'_i\delta$ are ground, there is a position p where $x_i\delta|_p$ and $x'_i\delta|_p$ have different function symbols. We construct the straight term t using the path from the root to p on $x_i\delta$ with variables that are fresh in (C, π) . Then, we can use x_i and t to segment $(C; \pi)$ into $(C; \pi \wedge x_i \neq t)$ and $(C; \pi)\{x_i \mapsto t\}$ for the refinement $N \Rightarrow_{\text{Ref}} N_R$. Note, that $x_i\delta$ is a ground instance of t , while $x'_i\delta$ is not.

Let (L'_1, r'_1) and (L'_2, r'_2) in (C_a, π_a) be literal positions of the variables x_i and x'_i in C_a , and (L_1, r_1) and (L_2, r_2) in (C, π) be the parent literal positions of (L'_1, r'_1) and (L'_2, r'_2) , respectively. Because $(C_a, \pi_a)\tau$ is equal to $(C; \pi)$ modulo duplicate literal elimination, $L_1|_{r_1} = L_2|_{r_2} = x_i$. Let $N \Rightarrow_{\text{Ref}} N_1$ be the refinement where $(C; \pi)$ is segmented into $(C; \pi \wedge x_i \neq t)$ and $(C; \pi)\{x_i \mapsto t\}$.

By Lemma 8, all $[(C; \pi \wedge x_i \neq t), N_1] \Rightarrow_A^* [(C'_a; \pi'_a), N_2]$ - or $[(C; \pi)\{x_i \mapsto t\}, N_1] \Rightarrow_A^* [(C'_a; \pi'_a), N_2]$ -descendant relations correspond to an equal $[(C; \pi), N] \Rightarrow_A^* [(C'_a; \pi'_a), N_2]$ -descendant relation. Assume there is a $[(C; \pi \wedge x_i \neq t), N_1] \Rightarrow_A^* [(C'_a; \pi'_a), N_2]$ -descendant relation $\Rightarrow_{C_c}^1$ that is not distinct from $\Rightarrow_{C_c}^0$. Because $L'_1\delta \Rightarrow_{C_c}^0 (L_1, r)$ for some literal position (L_1, r) in $(C; \pi)$, which is the parent literal position of (L_1, r) in $(C; \pi \wedge x_i \neq t)$, $L'_1\delta \Rightarrow_{C_c}^1 (L_1, r)$. However, this contradicts Lemma 7 because $x_i\delta$ is not an instance of $\text{skt}(L_1|_{r_1}, \pi \wedge x_i \neq t) = \text{skt}(x_i, \pi \wedge x_i \neq t)$. The case that there is a $[(C; \pi)\{x_i \mapsto t\}, N_1] \Rightarrow_A^* [(C'_a; \pi'_a), N_2]$ -descendant relation that is not distinct from $\Rightarrow_{C_c}^0$ is analogous using the argument that $x'_i\delta$ is not an instance of $\text{skt}(L_2\{x_i \mapsto t\}|_{r_2}, \pi) = \text{skt}(t, \pi)$. Hence, there are strictly less distinct descendant relations over C_c and $(C; \pi \wedge x_i \neq t)$ or $(C; \pi)\{x_i \mapsto t\}$ than there are distinct descendant relations over C_c and (C, π) .

If there are no descendant relations, then C_c can no longer appear as a lift conflict. Otherwise, by the inductive hypothesis, there exists a finite refinement $N \Rightarrow_{\text{Ref}} N_1 \Rightarrow_{\text{Ref}}^* N_R$ such that for any approximation $N_R \Rightarrow_{\text{AP}} N'_R$ and ground conflicting core N_R^\perp of N'_R , C_c is not a lift-conflict in N_R^\perp modulo duplicate literal elimination. \square

Theorem 2 (Soundness and Completeness of FO-AR). *Let N be an unsatisfiable clause set and N' its MSL(SDC) approximation: (i) if N is unsatisfiable then there exists a conflicting core of N' that can be lifted to a refutation in N , (ii) if N' is satisfiable, then N is satisfiable too.*

Proof. (Idea) By Lemma 6 and Lemma 9, where the latter can be used to show that a core of N' that cannot be lifted also excludes the respective instance for unsatisfiability of N .

Let (C_a, π_a) be the conflict clause of C_c and $(C; \pi) \in N$ be the parent clause of (C_a, π_a) . C_c is a descendant of $(C; \pi)$ with the corresponding $[(C; \pi), N] \Rightarrow_A^* [(C_a; \pi_a), N']$ -descendant relation $\Rightarrow_{C_c}^0$. We apply induction on the number of distinct $[(C; \pi), N] \Rightarrow_A^* [(C'; \pi'), N'']$ -descendant relations \Rightarrow_{C_c} for arbitrary approximations $N \Rightarrow_{\text{AP}}^* N''$.

Since only the shallow and linear transformations can produce lift-conflicts, the clause $(C; \pi)$ is replaced by either a linearized clause $(C'; \pi')$ or two shallow clauses $(C_l; \pi)$ and $(C_r; \pi)$. Then, the conflict clause $(C_a; \pi_a)$ of C_c is either the linearized $(C'; \pi')$ or the resolvent of $(C_l; \pi)$ and $(C_r; \pi)$. In either case,

$C_c = C_a\delta$ for some solution δ of π_a . Furthermore, there exists a substitution $\tau = \{x'_1 \mapsto x_1, \dots, x'_n \mapsto x_n\}$ such that $(C; \pi)$ and $(C_a; \pi_a)\tau$ are equal modulo duplicate literal elimination. That is, $\tau = \{x' \mapsto x\}$ for a linear transformation and $\tau = \rho^{-1}$ for shallow transformation (Definition 9).

Assume $C_c = C_a\tau\sigma$ for some grounding substitution σ , where $\tau\sigma$ is a solution of π_a . Thus, σ is a solution of $\pi_a\tau$, which is equivalent to π . Then, C_c is equal to $C\sigma$ modulo duplicate literal elimination an instance of $(C; \pi)$, which contradicts with C_c being a lift-conflict. Hence, $C_c = C_a\delta$ is not an instance of $C_a\tau$ and thus, $x_i\delta \neq x'_i\delta$ for some x_i in the domain of τ .

Because $x_i\delta$ and $x'_i\delta$ are ground, there is a position p where $x_i\delta|_p$ and $x'_i\delta|_p$ have different function symbols. We construct the straight term t using the path from the root to p on $x_i\delta$ with variables that are fresh in (C, π) . Then, we can use x_i and t to segment $(C; \pi)$ into $(C; \pi \wedge x_i \neq t)$ and $(C; \pi)\{x_i \mapsto t\}$ for the refinement $N \Rightarrow_{\text{Ref}} N_R$. Note, that $x_i\delta$ is a ground instance of t , while $x'_i\delta$ is not.

Let (L'_1, r'_1) and (L'_2, r'_2) in (C_a, π_a) be literal positions of the variables x_i and x'_i in C_a , and (L_1, r_1) and (L_2, r_2) in (C, π) be the parent literal positions of (L'_1, r'_1) and (L'_2, r'_2) , respectively. Because $(C_a, \pi_a)\tau$ is equal to $(C; \pi)$ modulo duplicate literal elimination, $L_1|_{r_1} = L_2|_{r_2} = x_i$. Let $N \Rightarrow_{\text{Ref}} N_1$ be the refinement where $(C; \pi)$ is segmented into $(C; \pi \wedge x_i \neq t)$ and $(C; \pi)\{x_i \mapsto t\}$.

By Lemma 8, all $[(C; \pi \wedge x_i \neq t), N_1] \Rightarrow_{\text{A}}^* [(C'_a; \pi'_a), N_2]$ - or $[(C; \pi)\{x_i \mapsto t\}, N_1] \Rightarrow_{\text{A}}^* [(C'_a; \pi'_a), N_2]$ -descendant relations correspond to an equal $[(C; \pi), N] \Rightarrow_{\text{A}}^* [(C'_a; \pi'_a), N_2]$ -descendant relation. Assume there is a $[(C; \pi \wedge x_i \neq t), N_1] \Rightarrow_{\text{A}}^* [(C'_a; \pi'_a), N_2]$ -descendant relation $\Rightarrow_{C_c}^1$ that is not distinct from $\Rightarrow_{C_c}^0$. Because $L'_1\delta \Rightarrow_{C_c}^0 (L_1, r)$ for some literal position (L_1, r) in $(C; \pi)$, which is the parent literal position of (L_1, r) in $(C; \pi \wedge x_i \neq t)$, $L'_1\delta \Rightarrow_{C_c}^1 (L_1, r)$. However, this contradicts Lemma 7 because $x_i\delta$ is not an instance of $\text{skt}(L_1|_{r_1}, \pi \wedge x_i \neq t) = \text{skt}(x_i, \pi \wedge x_i \neq t)$. The case that there is a $[(C; \pi)\{x_i \mapsto t\}, N_1] \Rightarrow_{\text{A}}^* [(C'_a; \pi'_a), N_2]$ -descendant relation that is not distinct from $\Rightarrow_{C_c}^0$ is analogous using the argument that $x'_i\delta$ is not an instance of $\text{skt}(L_2\{x_i \mapsto t\}|_{r_2}, \pi) = \text{skt}(t, \pi)$. Hence, there are strictly less distinct descendant relations over C_c and $(C; \pi \wedge x \neq t)$ or $(C; \pi)\{x \mapsto t\}$ than there are distinct descendant relations over C_c and (C, π) .

If there are no descendant relations, then C_c can no longer appear as a lift conflict. Otherwise, by the inductive hypothesis, there exists a finite refinement $N \Rightarrow_{\text{Ref}} N_1 \Rightarrow_{\text{Ref}}^* N_R$ such that for any approximation $N_R \Rightarrow_{\text{AP}} N'_R$ and ground conflicting core N_R^\perp of N'_R , C_c is not a lift-conflict in N_R^\perp modulo duplicate literal elimination. \square

Actually, Lemma 9 can be used to define a fair strategy on refutations in N' in order to receive also a dynamically complete FO-AR calculus, following the ideas presented in [18].

In Lemma 9, we segment the conflict clause's immediate parent clause. If the lifting later successfully passes this point, the refinement is lost and will be possibly repeated. Instead, we can refine any ancestor of the conflict clause as long as it contains the ancestor of the variable used in the refinement. By Lemma 7-(iii), such an ancestor will contain the ancestor variable at the same

positions. If we refine the ancestor in the original clause set, the refinement is permanent because lifting the refinement steps always succeeds. Only variables introduced by shallow transformation cannot be traced to the original clause set. However, these shallow variables are already linear and the partitioning in the shallow transformation can be chosen such that they are not shared variables. Assume a shallow, shared variable y , that is used to extract the term t , in the shallow transformation of $\Gamma \rightarrow E[s]_p, \Delta$ into $S(x), \Gamma_l \rightarrow E[p/x], \Delta_l$ and $\Gamma_r \rightarrow S(s), \Delta_r$. Since $\Delta_l \dot{\cup} \Delta_r = \Delta$ is a partitioning, y can only appear in either $E[p/x], \Delta_l$ or $S(s), \Delta_r$. If $y \in \text{vars}(E[p/x], \Delta_l)$ we instantiate Γ_r with $\{y \mapsto t\}$ and Γ_l , otherwise. Now, y is no longer a shared variable.

The refinement Lemmas only guarantee a refinement for a given ground conflicting core. In practice, however, conflicting cores contain free variables. We can always generate a ground conflicting core by instantiating the free variables with ground terms. However, if we only exclude a single ground case via refinement, next time the new conflicting core will likely have overlaps with the previous one. Instead, we can often remove all ground instances of a given conflict clause at once.

The simplest case is when unifying the conflict clause with the original clause fails because their instantiations differ at some equivalent positions. For example, consider $N = \{P(x, x); P(f(x, a), f(y, b)) \rightarrow\}$. N is satisfiable but the linear transformation is unsatisfiable with conflict clause $P(f(x, a), f(y, b))$ which is not unifiable with $P(x, x)$, because the two terms $f(x, a)$ and $f(y, b)$ have different constants at the second argument. A refinement of $P(x, x)$ is

$$\begin{aligned} & (P(x, x); x \neq f(v, a)) \\ & (P(f(x, a), f(x, a)); \top) \end{aligned}$$

$P(f(x, a), f(y, b))$ shares no ground instances with the approximations of the refined clauses.

Next, assume that again unification fails due to structural difference, but this time the differences lie at different positions. For example, consider $N = \{P(x, x); P(f(a, b), f(x, x)) \rightarrow\}$. N is satisfiable but the linear transformation of N is unsatisfiable with conflict clause $P(f(a, b), f(x, x))$ which is not unifiable with $P(x, x)$ because in $f(a, b)$ the first and second argument are different but the same in $f(x, x)$. A refinement of $P(x, x)$ is

$$\begin{aligned} & (P(x, x); x \neq f(a, v)) \\ & (P(f(a, x), f(a, x)); x \neq a) \\ & (P(f(a, a), f(a, a)); \top) \end{aligned}$$

$P(f(a, b), f(x, x))$ shares no ground instances with the approximations of the refined clauses.

It is also possible that the conflict clause and original clause are unifiable by themselves, but the resulting constraint has no solutions. For example, consider $N = \{P(x, x); (P(x, y) \rightarrow; x \neq a \wedge x \neq b \wedge y \neq c \wedge y \neq d)\}$ with signature $\Sigma = \{a, b, c, d\}$. N is satisfiable but the linear transformation of N is unsatisfiable with conflict clause $(\rightarrow P(x, y); x \neq a \wedge x \neq b \wedge y \neq c \wedge y \neq d)$. While $P(x, x)$ and $P(x, y)$ are unifiable, the resulting constraint $x \neq a \wedge x \neq b \wedge x \neq c \wedge x \neq d$

has no solutions. A refinement of $P(x, x)$ is

$$\begin{aligned} & (P(x, x); x \neq a \wedge x \neq b) \\ & (P(a, a); \top) \\ & (P(b, b); \top) \end{aligned}$$

$(P(x, y); x \neq a \wedge x \neq b \wedge y \neq c \wedge y \neq d)$ shares no ground instances with the approximations of the refined clauses.

Lastly, we should mention that there are cases where the refinement process does not terminate. For example, consider the clause set $N = \{P(x, x); P(y, g(y)) \rightarrow\}$. N is satisfiable but the linear transformation of N is unsatisfiable with conflict clause $P(y, g(y))$, which is not unifiable with $P(x, x)$. A refinement of $P(x, x)$ based on the ground instance $P(a, g(a))$ is

$$\begin{aligned} & (P(x, x); x \neq g(v)) \\ & (P(g(x), g(x)); \top) \end{aligned}$$

While $P(y, g(y))$ is not an instance of the refined approximation, it shares ground instances with $P(g(x), g(x'))$. The new conflict clause is $P(g(y), g(g(y)))$ and the refinement will continue to enumerate all $P(g^i(x), g^i(x))$ instances of $P(x, x)$ without ever reaching a satisfiable approximation. Satisfiability of first-order clause sets is undecidable, so termination cannot be expected by any calculus, in general.

5 Experiments

In the following we discuss several first-order clause classes for which FO-AR implemented in SPASS-AR immediately decides satisfiability but superposition and instantiation-based methods fail. We argue both according to the respective calculi and state-of-the-art implementations, in particular SPASS 3.9 [22], Vampire 4.1 [11,20], for ordered-resolution/superposition, iProver 2.5 [9] an implementation of Inst-Gen [10], and Darwin v1.4.5 [4] an implementation of the model evolution calculus [5]. All experiments were run on a 64-Bit Linux computer (Xeon(R) E5-2680, 2.70GHz, 256GB main memory). For Vampire and Darwin we chose the CASC-sat and CASC settings, respectively. For iProver we set the schedule to “sat” and SPASS, SPASS-AR were used in default mode. Please note that Vampire and iProver are portfolio solvers including implementations of several different calculi including superposition (ordered resolution), instance generation, and finite model finding. SPASS, SPASS-AR, and Darwin only implement superposition, FO-AR, and model evolution, respectively.

For the first example

$$P(x, y) \rightarrow P(x, z), P(z, y); \quad P(a, a)$$

and second example,

$$Q(x, x); \quad Q(v, w), P(x, y) \rightarrow P(x, v), P(w, y); \quad P(a, a)$$

the superposition calculus produces independently of the selection strategy and ordering an infinite number of clauses of form

$$\rightarrow P(a, z_1), P(z_1, z_2), \dots, P(z_n, a).$$

Using linear approximation, however, FO-AR replaces $P(x, y) \rightarrow P(x, z), P(z, y)$ and $\rightarrow Q(x, x)$ with $P(x, y) \rightarrow P(x, z), P(z', y)$ and $\rightarrow Q(x, x')$, respectively. Consequently, ordered resolution derives $\rightarrow P(a, z_1), P(z_2, a)$ which subsumes any further inferences $\rightarrow P(a, z_1), P(z_2, z_3), P(z_4, a)$. Hence, saturation of the approximation terminates immediately. Both examples belong to the Bernays-Schönfinkel fragment, so model evolution (Darwin) and Inst-Gen (iProver) can decide them as well. Note that the concrete behavior of superposition is not limited to the above examples but potentially occurs whenever there are variable chains in clauses.

On the third problem

$$P(x, y) \rightarrow P(g(x), z); \quad P(a, a)$$

superposition derives all clauses of the form $\rightarrow P(g(\dots g(a) \dots), z)$. With a shallow approximation of $P(x, y) \rightarrow P(g(x), z)$ into $S(v) \rightarrow P(v, z)$ and $P(x, y) \rightarrow S(g(x))$, FO-AR (SPASS-AR) terminates after deriving $\rightarrow S(g(a))$ and $S(x) \rightarrow S(g(x))$. Again, model evolution (Darwin) and Inst-Gen (iProver) can also solve this example.

The next example

$$P(a); \quad P(f(a)) \rightarrow; \quad P(f(f(x))) \rightarrow P(x); \quad P(x) \rightarrow P(f(f(x)))$$

is already saturated under superposition. For FO-AR, the clause $P(x) \rightarrow P(f(f(x)))$ is replaced by $S(x) \rightarrow P(f(x))$ and $P(x) \rightarrow S(f(x))$. Then ordered resolution terminates after inferring $S(a) \rightarrow$ and $S(f(x)) \rightarrow P(x)$.

The Inst-Gen and model evolution calculi, however, fail. In either, a satisfying model is represented by a finite set of literals, i.e., a model of the propositional approximation for Inst-Gen and the trail of literals in case of model evolution. Therefore, there necessarily exists a literal $P(f^n(x))$ or $\neg P(f^n(x))$ with a maximal n in these models. This contradicts the actual model where either $P(f^n(a))$ or $P(f^n(f(a)))$ is true. However, iProver can solve this problem using its built-in ordered resolution solver whereas Darwin does not terminate on this problem.

Lastly consider an example of the form

$$f(x) \approx x \rightarrow; \quad f(f(x)) \approx x \rightarrow; \quad \dots; \quad f^n(x) \approx x \rightarrow$$

which is trivially satisfiable, e.g., saturated by superposition, but any model has at least $n+1$ domain elements. Therefore, adding these clauses to any satisfiable clause set containing f forces calculi that explicitly consider finite models to consider at least $n+1$ elements. The performance of final model finders [15] typically degrades in the number of different domain elements to be considered.

Combining each of these examples into one problem is then solvable by neither superposition, Inst-Gen, or model evolution and not practically solvable with increasing n via testing finite models. For example, we tested

$$\begin{aligned} &P(x, y) \rightarrow P(x, z), P(z, y); \quad P(a, a); \quad P(f(a), y) \rightarrow; \\ &P(f(f(x)), y) \rightarrow P(x, y); \quad P(x, y) \rightarrow P(f(f(x)), y); \\ &f(x) \approx x \rightarrow; \dots, f^n(x) \approx x \rightarrow; \end{aligned}$$

for $n = 20$ against SPASS, Vampire, iProver, and Darwin for more than one hour each without success. Only SPASS-AR solved it in less than one second.

For iProver we added an artificial positive equation $b \approx c$. For otherwise, iProver throws away all disequations while preprocessing. This is a satisfiability

ity preserving operation, however, the afterwards found (finite) models are not models of the above clause set due to the collapsing of ground terms.

6 Conclusion

The previous section showed FO-AR is superior to superposition, instantiation-based methods on certain classes of clause sets. Of course, there are also classes of clause sets where superposition and instantiation-based methods are superior to FO-AR, e.g., for unsatisfiable clause sets where the structure of the clause set forces FO-AR to enumerate failing ground instances due to the approximation in a bottom-up way.

Our prototypical implementation SPASS-AR cannot compete with systems such as iProver or Vampire on the respective CASC categories of the TPTP [17]. This is already due to the fact that they are all meanwhile portfolio solvers. For example, iProver contains an implementation of ordered resolution and Vampire an implementation of Inst-Gen. Our results, Section 5, however, show that these systems may benefit from FO-AR by adding it to their portfolio.

The DEXPTIME-completeness result for MSLH strongly suggest that both the MSLH and also our MSL(SDC) fragment have the finite model property. However, we are not aware of any proof. If MSL(DSC) has the finite model property, the finite model finding approaches are complete on MSL(SDC). The models generated by FO-AR and superposition are typically infinite. It remains an open problem, even for fragments enjoying the finite model property, e.g., the first-order monadic fragment, to design a calculus that combines explicit finite model finding with a structural representation of infinite models. For classes that have no finite models this problem seems to become even more difficult. To the best of our knowledge, SPASS is currently the only prover that can show satisfiability of the clauses $R(x, x) \rightarrow; R(x, y), R(y, z) \rightarrow R(x, z); R(x, g(x))$ due to an implementation of chaining [2,16]. Apart from the superposition calculus, it is unknown to us how the specific inferences for transitivity can be combined with any of the other discussed calculi, including the abstraction refinement calculus introduced in this paper.

Finally, there are not many results on calculi that operate with respect to models containing positive equations. Even for fragments that are decidable with equality, such as the Bernays-Schoenfinkel-Ramsey fragment or the monadic fragment with equality, there seem currently no convincing suggestions compared to the great amount of techniques for these fragments without equality. Adding positive equations to MSL(SDC) while keeping decidability is, to the best of our current knowledge, only possible for at most linear, shallow equations $f(x_1, \dots, x_n) \approx h(y_1, \dots, y_n)$ [8]. However, approximation into such equations from an equational theory with nested term occurrences typically results in an almost trivial equational theory. So this does not seem to be a very promising research direction.

References

1. Leo Bachmair and Harald Ganzinger. Rewrite-based equational theorem proving with selection and simplification. *Journal of Logic and Computation*, 4(3):217–247, 1994. Revised version of Max-Planck-Institut für Informatik technical report, MPI-I-91-208, 1991.
2. Leo Bachmair and Harald Ganzinger. Ordered chaining calculi for first-order theories of transitive relations. *Journal of the ACM*, 45(6):1007–1049, 1998.
3. Jos C. M. Baeten, Jan A. Bergstra, Jan Willem Klop, and W. P. Weijland. Term-rewriting systems with rule priorities. *Theor. Comput. Sci.*, 67(2&3):283–301, 1989.
4. Peter Baumgartner, Alexander Fuchs, and Cesare Tinelli. Implementing the model evolution calculus. *International Journal on Artificial Intelligence Tools*, 15(1):21–52, 2006.
5. Peter Baumgartner and Cesare Tinelli. The model evolution calculus. In Franz Baader, editor, *Automated Deduction - CADE-19, 19th International Conference on Automated Deduction Miami Beach, FL, USA, July 28 - August 2, 2003, Proceedings*, volume 2741 of *Lecture Notes in Computer Science*, pages 350–364. Springer, 2003.
6. H. Comon, M. Dauchet, R. Gilleron, C. Löding, F. Jacquemard, D. Lugiez, S. Tison, and M. Tommasi. Tree automata techniques and applications. Available on: <http://www.grappa.univ-lille3.fr/tata>, 2007. release October, 12th 2007.
7. Jean Goubault-Larrecq. Deciding \mathcal{H}_1 by resolution. *Information Processing Letters*, 95(3):401 – 408, 2005.
8. Florent Jacquemard, Christoph Meyer, and Christoph Weidenbach. Unification in extensions of shallow equational theories. In Tobias Nipkow, editor, *Rewriting Techniques and Applications, 9th International Conference, RTA-98*, volume 1379 of *LNCS*, pages 76–90. Springer, 1998.
9. Konstantin Korovin. iprover - an instantiation-based theorem prover for first-order logic (system description). In Alessandro Armando, Peter Baumgartner, and Gilles Dowek, editors, *Automated Reasoning, 4th International Joint Conference, IJCAR 2008, Sydney, Australia, August 12-15, 2008, Proceedings*, volume 5195 of *Lecture Notes in Computer Science*, pages 292–298. Springer, 2008.
10. Konstantin Korovin. Inst-Gen - A modular approach to instantiation-based automated reasoning. In *Programming Logics - Essays in Memory of Harald Ganzinger*, pages 239–270, 2013.
11. Laura Kovács and Andrei Voronkov. First-order theorem proving and vampire. In Natasha Sharygina and Helmut Veith, editors, *Computer Aided Verification - 25th International Conference, CAV 2013, Saint Petersburg, Russia, July 13-19, 2013. Proceedings*, volume 8044 of *Lecture Notes in Computer Science*, pages 1–35. Springer, 2013.
12. Helmut Seidl and Andreas Reuß. Extending \mathcal{H}_1 -Clauses with Disequalities. *Information Processing Letters*, 111(20):1007–1013, 2011.
13. Helmut Seidl and Andreas Reuß. *Foundations of Software Science and Computational Structures: 15th International Conference, FOSSACS 2012, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2012, Tallinn, Estonia, March 24 – April 1, 2012. Proceedings*, chapter Extending \mathcal{H}_1 -Clauses with Path Disequalities, pages 165–179. Springer Berlin Heidelberg, Berlin, Heidelberg, 2012.
14. Helmut Seidl and Kumar Neeraj Verma. Cryptographic protocol verification using tractable classes of horn clauses. In *Program Analysis and Compilation, Theory and Practice*, pages 97–119. Springer, Juni 2007. Lecture Notes in Computer Science.

15. John K. Slaney and Timothy Surendonk. Combining finite model generation with theorem proving: Problems and prospects. In Franz Baader and Klaus U. Schulz, editors, *Frontiers of Combining Systems, First International Workshop FroCoS 1996, Munich, Germany, March 26-29, 1996, Proceedings*, volume 3 of *Applied Logic Series*, pages 141–155. Kluwer Academic Publishers, 1996.
16. Martin Suda, Christoph Weidenbach, and Patrick Wischniewski. On the saturation of yago. In *Automated Reasoning, 5th International Joint Conference, IJCAR 2010*, volume 6173 of *LNAI*, pages 441–456, Edinburgh, United Kingdom, 2010. Springer.
17. G. Sutcliffe. The TPTP Problem Library and Associated Infrastructure: The FOF and CNF Parts, v3.5.0. *Journal of Automated Reasoning*, 43(4):337–362, 2009.
18. Andreas Teucke and Christoph Weidenbach. First-order logic theorem proving and model building via approximation and instantiation. In Carsten Lutz and Silvio Ranise, editors, *Frontiers of Combining Systems: 10th International Symposium, FroCoS 2015, Wroclaw, Poland, September 21-24, 2015, Proceedings*, pages 85–100, Cham, 2015. Springer International Publishing.
19. Andreas Teucke and Christoph Weidenbach. Ordered resolution with straight dis-matching constraints. In Pascal Fontaine, Stephan Schulz, and Josef Urban, editors, *Proceedings of the 5th Workshop on Practical Aspects of Automated Reasoning co-located with International Joint Conference on Automated Reasoning (IJCAR 2016), Coimbra, Portugal, July 2nd, 2016.*, volume 1635 of *CEUR Workshop Proceedings*, pages 95–109. CEUR-WS.org, 2016.
20. Andrei Voronkov. AVATAR: the architecture for first-order theorem provers. In Armin Biere and Roderick Bloem, editors, *Computer Aided Verification - 26th International Conference, CAV 2014, Held as Part of the Vienna Summer of Logic, VSL 2014, Vienna, Austria, July 18-22, 2014. Proceedings*, volume 8559 of *Lecture Notes in Computer Science*, pages 696–710. Springer, 2014.
21. Christoph Weidenbach. Towards an automatic analysis of security protocols in first-order logic. In Harald Ganzinger, editor, *16th International Conference on Automated Deduction, CADE-16*, volume 1632 of *LNAI*, pages 314–328. Springer, 1999.
22. Christoph Weidenbach, Dilyana Dimova, Arnaud Fietzke, Rohit Kumar, Martin Suda, and Patrick Wischniewski. SPASS version 3.5. In Renate A. Schmidt, editor, *Automated Deduction - CADE-22, 22nd International Conference on Automated Deduction, Montreal, Canada, August 2-7, 2009. Proceedings*, volume 5663 of *Lecture Notes in Computer Science*, pages 140–145. Springer, 2009.