

7.17 SCLT Clause Learning from Simple Models Modulo Theories

Let $\mathcal{T}^{\mathcal{B}}$ be first-order logic *background theory* over signature $\Sigma^{\mathcal{B}} = (\mathcal{S}^{\mathcal{B}}, \Omega^{\mathcal{B}}, \Pi^{\mathcal{B}})$ and term-generated $\Sigma^{\mathcal{B}}$ -algebras $\mathcal{C}^{\mathcal{B}}$: $\mathcal{T}^{\mathcal{B}} = (\Sigma^{\mathcal{B}}, \mathcal{C}^{\mathcal{B}})$. A constant $c \in \Omega^{\mathcal{B}}$ is called a *domain constant* if $c^{\mathcal{A}} \neq d^{\mathcal{A}}$ for all $\mathcal{A} \in \mathcal{C}^{\mathcal{B}}$ and for all $d \in \Omega^{\mathcal{B}}$ with $d \neq c$. Let $\Sigma^{\mathcal{F}} = (\mathcal{S}^{\mathcal{F}}, \Omega^{\mathcal{F}}, \Pi^{\mathcal{F}})$ be a *foreground signature* with respect to $\mathcal{T}^{\mathcal{B}}$ where $\mathcal{S}^{\mathcal{B}} \subseteq \mathcal{S}^{\mathcal{F}}$, $\Omega^{\mathcal{B}} \cap \Omega^{\mathcal{F}} = \emptyset$, and $\Pi^{\mathcal{B}} \cap \Pi^{\mathcal{F}} = \emptyset$.

Definition 7.17.1 (Hierarchic Specification). A *hierarchic specification* is a pair $\mathcal{H} = (\mathcal{T}^{\mathcal{B}}, \Sigma^{\mathcal{F}})$ with associated signature $\Sigma^{\mathcal{H}} = (\mathcal{S}^{\mathcal{F}}, \Omega^{\mathcal{B}} \cup \Omega^{\mathcal{F}}, \Pi^{\mathcal{B}} \cup \Pi^{\mathcal{F}})$. It generates *hierarchic* $\Sigma^{\mathcal{H}}$ -algebras. A $\Sigma^{\mathcal{H}}$ -algebra \mathcal{A} is called *hierarchic* with respect to its background theory $\mathcal{T}^{\mathcal{B}}$, if $\mathcal{A}^{\mathcal{H}}|_{\Sigma^{\mathcal{B}}} \in \mathcal{C}^{\mathcal{B}}$.

As usual, $\mathcal{A}^{\mathcal{H}}|_{\Sigma^{\mathcal{B}}}$ is obtained from a $\mathcal{A}^{\mathcal{H}}$ -algebra by removing all carrier sets $S^{\mathcal{A}}$ for all $S \in (\mathcal{S}^{\mathcal{F}} \setminus \mathcal{S}^{\mathcal{B}})$, all functions from $\Omega^{\mathcal{F}}$ and all predicates from $\Pi^{\mathcal{F}}$. We write $\models_{\mathcal{H}}$ for the entailment relation with respect to hierarchic algebras and formulas from $\Sigma^{\mathcal{H}}$ and $\models_{\mathcal{B}}$ for the entailment relation with respect to the $\mathcal{C}^{\mathcal{B}}$ algebras and formulas from $\Sigma^{\mathcal{B}}$.

Terms, atoms, literals build over $\Sigma^{\mathcal{B}}$ are called *pure background terms*, *pure background atoms*, and *pure background literals*, respectively. All terms, atoms, with a top-symbol from $\Omega^{\mathcal{B}}$ or $\Pi^{\mathcal{B}}$, respectively, are called *background terms*, *background atoms*, respectively. A background atom or its negation is a *background literal*. All terms, atoms, with a top-symbol from $\Omega^{\mathcal{F}}$ or $\Pi^{\mathcal{F}}$, respectively, are called *foreground terms*, *foreground atoms*, respectively. A foreground atom or its negation is a *foreground literal*. Given a set (sequence) of \mathcal{H} literals, the function `bgd` returns the set (sequence) of background literals and the function `fgd` the respective set (sequence) of foreground literals.

As a running example, I consider in detail the Bernays-Schoenfinkel clause fragment over linear arithmetic: BS(LRA). The background theory is linear rational arithmetic over the many-sorted signature $\Sigma^{\text{LRA}} = (\mathcal{S}^{\text{LRA}}, \Omega^{\text{LRA}}, \Pi^{\text{LRA}})$ with $\mathcal{S}^{\text{LRA}} = \{\text{LRA}\}$, $\Omega^{\text{LRA}} = \{0, 1, +, -\} \cup \mathbb{Q}$, $\Pi^{\text{LRA}} = \{\leq, <, \neq, =, >, \geq\}$ where LRA is the linear arithmetic sort, the function symbols consist of $0, 1, +, -$ plus the rational numbers and predicate symbols $\leq, <, =, \neq, >, \geq$. The linear arithmetic theory $\mathcal{T}^{\text{LRA}} = (\Sigma^{\text{LRA}}, \{\mathcal{A}^{\text{LRA}}\})$ consists of the linear arithmetic signature together with the standard model \mathcal{A}^{LRA} of linear arithmetic. This theory is then extended by the free (foreground) first-order signature $\Sigma^{\text{BS}} = (\{\text{LRA}\}, \Omega^{\text{BS}}, \Pi^{\text{BS}})$ where Ω^{BS} is a set of constants of sort LRA different from Ω^{LRA} constants, and Π^{BS} is a set of first-order predicates over the

sort LRA. We are interested in hierarchic algebras $\mathcal{A}^{\text{BS}(\text{LRA})}$ over the signature $\Sigma^{\text{BS}(\text{LRA})} = (\{\text{LRA}\}, \Omega^{\text{BS}} \cup \Omega^{\text{LRA}}, \Pi^{\text{BS}} \cup \Pi^{\text{LRA}})$ that are $\Sigma^{\text{BS}(\text{LRA})}$ algebras such that $\mathcal{A}^{\text{BS}(\text{LRA})}|_{\Sigma^{\text{LRA}}} = \mathcal{A}^{\text{LRA}}$.

Definition 7.17.2 (Simple Substitutions). A substitution σ is called *simple* if $x_S \sigma \in T_S(\Sigma^{\mathcal{B}}, \mathcal{X})$ for all $x_S \in \text{dom}(\sigma)$ and $S \in \mathcal{S}^{\mathcal{B}}$.

As usual, clauses are disjunctions of literals with implicitly universally quantified variables. We often write a $\Sigma^{\mathcal{H}}$ clause as a *constrained clause*, denoted $\Lambda \parallel C$ where Λ is a conjunction of background literals and C is a disjunction of foreground literals semantically denoting the clause $\neg\Lambda \vee C$. A *constrained closure* is denoted as $\Lambda \parallel C \cdot \sigma$ where σ is grounding for Λ and C . A constrained closure $\Lambda \parallel C \cdot \sigma$ denotes the ground constrained clause $\Lambda\sigma \parallel C\sigma$.

In addition, we assume a well-founded, total, strict ordering \prec on ground literals, called an \mathcal{H} -order, such that background literals are smaller than foreground literals. This ordering is then lifted to constrained clauses and sets thereof by its respective multiset extension. We overload \prec for literals, constrained clauses, and sets of constrained clause if the meaning is clear from the context. We define \preceq as the reflexive closure of \prec and $N^{\preceq\Lambda \parallel C} := \{D \mid D \in N \text{ and } D \preceq \Lambda \parallel C\}$. For example, an instance of an LPO with according precedence can serve as \prec .

Definition 7.17.3 (Abstracted/Pure Clause). A clause $\Lambda \parallel C$ is *abstracted* if the arguments of $\mathcal{S}^{\mathcal{B}}$ sort of any predicate from $\Pi^{\mathcal{F}}$ in an atom in C are only variables. $\Lambda \parallel C$ is called *pure* if it does not contain symbols from $\Omega^{\mathcal{F}}$ ranging into a sort of $\mathcal{S}^{\mathcal{B}}$.

These two notions are extended to clause sets in the natural way. Any clause set can be transformed into an abstracted clause set.

Abstraction $N \uplus \{C \vee E[t]_p[s]_q\} \Rightarrow_{\text{ABSTR}} N \cup \{C \vee x_s \not\approx s \vee E[x_S]_q\}$
provided t, s are non-variable terms, $q \not\prec p$, $\text{sort}(s) = S$, and either $\text{top}(t) \in \Sigma^{\mathcal{F}}$ and $\text{top}(s) \in \Sigma^{\mathcal{B}}$ or $\text{top}(t) \in \Sigma^{\mathcal{B}}$ and $\text{top}(s) \in \Sigma^{\mathcal{F}}$

In case of BS(LRA) abstraction can only be applied to constants below a predicate.

Definition 7.17.4 (Clause Redundancy). A ground constrained clause $\Lambda \parallel C$ is *redundant* with respect to a set N of ground constrained clauses and an order \prec if $N^{\preceq\Lambda \parallel C} \models_{\mathcal{H}} \Lambda \parallel C$. A clause $\Lambda \parallel C$ is *redundant* with respect to a clause set N , an \mathcal{H} -order \prec , and a set of constants B if for all $\Lambda' \parallel C' \in \text{grd}((\mathcal{S}^{\mathcal{F}}, B, \Pi^{\mathcal{B}} \cup \Pi^{\mathcal{F}}), \Lambda \parallel C)$ the clause $\Lambda' \parallel C'$ is redundant with respect to $\cup_{D \in N} \text{grd}((\mathcal{S}^{\mathcal{F}}, B, \Pi^{\mathcal{B}} \cup \Pi^{\mathcal{F}}), D)$.

Assumption 7.17.5 (Considered Clause Sets). For the rest of this section I consider only pure, abstracted clause sets N . I assume that the background theory $\mathcal{T}^{\mathcal{B}}$ is term-generated, compact, contains an equality $=$, and that all constants of the background signature are domain constants. I further assume that the set $\Omega^{\mathcal{F}}$ contains infinitely many constants for each background sort.

Example 7.17.6 (Pure Clauses). With respect to BS(LRA) the unit clause $x \geq 5, 3x + 4y = z \parallel Q(x, y, z)$ is abstracted and pure while the clause $x \geq 5, 3x + 4y = a, z = a \parallel Q(x, y, z)$ is abstracted but not pure because of the foreground constant a of the LRA sort, and the clause $x \geq 5, 3x + 4y = 7 \parallel Q(x, y, 7)$ is not abstracted.

Note that for pure, abstracted clause sets, any unifier between two foreground literals is simple and its codomain consists of variables only.

In order for the SCL(T) calculus to be effective, decidability in $\mathcal{T}^{\mathcal{B}}$ is needed as well. For the calculus we implicitly use the following equivalence: A $\Sigma^{\mathcal{B}}$ sentence

$$\exists x_1, \dots, x_n \phi$$

where ϕ is quantifier free is true, i.e., $\models_{\mathcal{B}} \exists x_1, \dots, x_n \phi$ iff the ground formula

$$\phi\{x_1 \mapsto a_1, \dots, x_n \mapsto a_n\}$$

where the a_i are $\Omega^{\mathcal{F}}$ constants of the respective background sorts is \mathcal{H} satisfiable. Together with decidability in $\mathcal{T}^{\mathcal{B}}$ this guarantees decidability of the satisfiability of ground constraints from constrained clauses.

If not stated otherwise, satisfiability means satisfiability with respect to \mathcal{H} . The function $\text{adiff}(B)$ for some finite sequence of background sort constants denotes a constraint that implies different interpretations for the constants in B . In case the background theory enables a strict ordering $<$ as LRA does, then the ordering can be used for this purpose. For example, $\text{adiff}([a, b, c])$ is then the constraint $a < b < c$. In case the background theory does not enable a strict ordering, then inequations can express disjointness of the constants. For example, $\text{adiff}([a, b, c])$ is then constraint $a \neq b \wedge a \neq c \wedge b \neq c$. An ordering constraint has the advantage over an inequality constraint that it also breaks symmetries. Assuming all constants to be different will eventually enable a satisfiability test for foreground literals based on purely syntactic complementarity.

The inference rules of SCL(T) are represented by an abstract rewrite system. They operate on a problem state, a six-tuple $\Gamma = (M; N; U; B; k; D)$ where M is a sequence of annotated ground literals, the *trail*; N and U are the sets of *initial* and *learned* constrained clauses; B is a finite sequence of constants of background sorts for instantiation; k counts the number of decisions in M ; and D is a constrained closure that is either \top , $\Lambda \parallel \perp \cdot \sigma$, or $\Lambda \parallel C \cdot \sigma$. Foreground literals in M are either annotated with a number, a level; i.e., they have the form L^k meaning that L is the k -th guessed decision literal, or they are annotated with a constrained closure that propagated the literal to become true, i.e., they have the form $(L\sigma)^{(\Lambda \parallel C \vee L) \cdot \sigma}$. An annotated literal is called a decision literal if it is of the form L^k and a propagation literal or a propagated literal if it is of the form $L \cdot \sigma^{(\Lambda \parallel C \vee L) \cdot \sigma}$. A ground foreground literal L is of *level* i with respect to a problem state $(M; N; U; B; k; D)$ if L or $\text{comp}(L)$ occurs in M and the first decision literal left from L ($\text{comp}(L)$) in M , including L , is annotated with i . If there is no such decision literal then its level is zero. A ground constrained clause $\Lambda \parallel C$ is of *level* i with respect to a problem state $(M; N; U; B; k; D)$ if

i is the maximal level of a foreground literal in C ; the level of an empty clause $\Lambda \parallel \perp \cdot \sigma$ is 0. A ground literal L is *undefined* in M if neither L nor $\text{comp}(L)$ occur in M . The initial state for a first-order, pure, abstracted \mathcal{H} clause set N is $(\epsilon; N; \emptyset; B; 0; \top)$, where B is a finite sequence of foreground constants of background sorts. These constants cannot occur in N , because N is pure. The final state $(\epsilon; N; U; B; 0; \Lambda \parallel \perp)$ denotes unsatisfiability of N . Given a trail M and its foreground literals $\text{fgd}(M) = \{L_1, \dots, L_n\}$ an \mathcal{H} ordering \prec induced by M is any \mathcal{H} ordering where $L_i \prec L_j$ if L_i occurs left from L_j in M , or, L_i is defined in M and L_j is not.

The transition rules for $\text{SCL}(\mathbb{T})$ are

Propagate $(M; N; U; B; k; \top) \Rightarrow_{\text{SCL}(\mathbb{T})} (M, L\sigma^{(\Lambda \parallel C_0 \vee L)\delta \cdot \sigma}, \Lambda'\sigma; N; U; B; k; \top)$

provided $\Lambda \parallel C \in (N \cup U)$, σ is grounding for $\Lambda \parallel C$, $\text{adiff}(B) \wedge \text{bgd}(M) \wedge \Lambda\sigma$ is satisfiable, $C = C_0 \vee C_1 \vee L$, $C_1\sigma = L\sigma \vee \dots \vee L\sigma$, $C_0\sigma$ does not contain $L\sigma$, δ is the mgu of the literals in C_1 and L , $\Lambda'\sigma$ are the background literals from $\Lambda\sigma$ that are not yet on the trail, $\text{fgd}(M) \models \neg(C_0\sigma)$, $\text{codom}(\sigma) \subseteq B$, and $L\sigma$ is undefined in M

The rule Propagate applies exhaustive factoring to the propagated literal with respect to the grounding substitution σ and annotates the factored clause to the propagation. By writing $M, L\sigma^{(\Lambda \parallel C_0 \vee L)\delta \cdot \sigma}, \Lambda'\sigma$ we denote that all background literals from $\Lambda'\sigma$ are added to the trail.

Decide $(M; N; U; B; k; \top) \Rightarrow_{\text{SCL}(\mathbb{T})} (M, L\sigma^{k+1}, \Lambda\sigma; N; U; B; k+1; \top)$

provided $L\sigma$ is undefined in M , $|L\sigma| \in \text{atoms}(\text{grd}((\mathcal{S}, B, \Pi), N \cup U))$, $|K\sigma| \in \text{atoms}(\text{grd}((\mathcal{S}, B, \Pi), N \cup U))$ for all $K\sigma \in \Lambda\sigma$, σ is grounding for Λ , all background literals in $\Lambda\sigma$ are undefined in M , $\text{adiff}(B) \wedge \text{bgd}(M) \wedge \Lambda\sigma$ is satisfiable, and $\text{codom}(\sigma) \subseteq B$

Making sure that no duplicates of background literals occur on the trail by rules Propagate and Decide together with a fixed finite sequence B of constants and the restriction of Propagate and Decide to undefined literals guarantees that the number of potential trails of a run is finite. Requiring the constants from B to be different by the $\text{adiff}(B)$ constraint enables a purely syntactic consistency check for foreground literals.

Conflict $(M; N; U; B; k; \top) \Rightarrow_{\text{SCL}(\mathbb{T})} (M; N; U; B; k; \Lambda \parallel D \cdot \sigma)$

provided $\Lambda \parallel D \in (N \cup U)$, σ is grounding for $\Lambda \parallel D$, $\text{adiff}(B) \wedge \text{bgd}(M) \wedge \Lambda\sigma$ is satisfiable, $\text{fgd}(M) \models \neg(D\sigma)$, and $\text{codom}(\sigma) \subseteq B$

Resolve $(M, L\rho^{\Lambda \parallel C \vee L \cdot \rho}; N; U; B; k; (\Lambda' \parallel D \vee L') \cdot \sigma) \Rightarrow_{\text{SCL}(\mathbb{T})}$
 $(M, L\rho^{\Lambda \parallel C \vee L \cdot \rho}; N; U; B; k; (\Lambda \wedge \Lambda' \parallel D \vee C)\eta \cdot \sigma\rho)$

provided $L\rho = \text{comp}(L'\sigma)$, and $\eta = \text{mgu}(L, \text{comp}(L'))$

7.17. SCLT CLAUSE LEARNING FROM SIMPLE MODELS MODULO THEORIES 307

Note that Resolve does not remove the literal $L\rho$ from the trail. This is needed if the clause $D\sigma$ contains further literals complementary of $L\rho$ that have not been factorized.

Factorize $(M; N; U; B; k; (\Lambda \parallel D \vee L \vee L') \cdot \sigma) \Rightarrow_{\text{SCL(T)}} (M; N; U; B; k; (\Lambda \parallel D \vee L)\eta \cdot \sigma)$

provided $L\sigma = L'\sigma$, and $\eta = \text{mgu}(L, L')$

Note that Factorize is not limited with respect to the trail. It may apply to any two literals that become identical by application of the grounding substitution σ .

Skip $(M, L; N; U; B; k; \Lambda' \parallel D \cdot \sigma) \Rightarrow_{\text{SCL(T)}} (M; N; U; B; l; \Lambda' \parallel D \cdot \sigma)$

provided L is a foreground literal and $\text{comp}(L)$ does not occur in $D\sigma$, or L is a background literal; if L is a foreground decision literal then $l = k - 1$, otherwise $l = k$

Note that Skip can also skip decision literals. This is needed because we won't eventually require exhaustive propagation. While exhaustive propagation in CDCL is limited to the number of propositional variables, in the context of our logic, for example BS(LRA), it is exponential in the arity of foreground predicate symbols and can lead to an unfair exploration of the space of possible inferences, harming completeness, see Example 7.17.9.

Backtrack $(M, K^{i+1}, M'; N; U; B; k; (\Lambda \parallel D \vee L) \cdot \sigma) \Rightarrow_{\text{SCL(T)}} (M, L\sigma^{(\Lambda \parallel D \vee L)\sigma}, \Lambda'\sigma; N; U \cup \{\Lambda \parallel D \vee L\}; B; i; \top)$

provided $L\sigma$ is of level k , and $D\sigma$ is of level i , $\Lambda'\sigma$ are the background literals from $\Lambda\sigma$ that are not yet on the trail

The definition of Backtrack requires that if $L\sigma$ is the only literal of level k in $(D \vee L)\sigma$ then additional occurrences of $L\sigma$ in D have to be factorized first before Backtrack can be applied.

Grow $(M; N; U; B; k; \top) \Rightarrow_{\text{SCL(T)}} (\epsilon; N; U; B \cup B'; 0; \top)$

provided B' is a non-empty sequence of foreground constants of background sorts distinct from the constants in B

In case the adiff constraint is implemented by a strict ordering predicate on the basis of the sequence B , it can be useful to inject the new constants B' into $B \cup B'$ such that the ordering of the constants from B is not changed. This can help caching background theory results for testing trail satisfiability.

Definition 7.17.7. The rules Propagate, Decide, Grow, and Conflict are called *conflict search* rules and the rules Resolve, Skip, Factorize, and Backtrack are called *conflict resolution* rules.

Recall that the goal of our calculus is to replace the ordering restrictions of the hierarchic superposition calculus with a guiding model assumption. All our inferences are hierarchic superposition inferences where the ordering restrictions are neglected.

Example 7.17.8 (Inconsistent Trail). Consider a clause set $N = \{R(x, y), x \leq y \parallel \neg R(x, y) \vee P(x), x \geq y \parallel \neg R(x, y) \vee \neg P(y)\}$; if we were to remove the $\text{adiff}(B)$ constraint from the side conditions of rule Propagate we would be able to obtain inconsistent trails. Starting with just $B = \{a, b\}$ as constants it is possible to propagate three times and obtain the trail $M = [R(a, b), P(a), a \leq b, \neg P(b), a \geq b]$, M is clearly inconsistent as $M \models P(a)$, $M \models \neg P(b)$ yet $a = b$.

Example 7.17.9 (Exhaustive Propagation). Consider a BS(LRA) clause set $N = \{x = 0 \parallel \text{Nat}(x), y = x + 1 \parallel \neg \text{Nat}(x) \vee \text{Nat}(y)\} \cup N'$ where N' is unsatisfiable and nothing can be propagated from N' . Let us further assume that N' is satisfiable with respect to any instantiation of variables with natural numbers. If propagation is not restricted, then the first two clauses will consume all constants in B . For example, if $B = [a, b, c]$ then the trail $[\text{Nat}(a), a = 0, \text{Nat}(b), b = a + 1, \text{Nat}(c), c = b + 1]$ will be derived. Now all constants are fixed to natural numbers. So there cannot be a refutation of N' anymore. An application of Grow will not solve the issue, because again the first two rules will fix all constants to natural numbers via exhaustive propagation.

Definition 7.17.10 (Well-formed States). A state $(M; N; U; B; k; D)$ is *well-formed* if the following conditions hold:

1. all constants appearing in $(M; N; U; B; k; D)$ are from B or occur in N .
2. $M \wedge \text{adiff}(B)$ is satisfiable
3. $N \models_{\mathcal{H}} U$,
4. Propagating clauses remain propagating and conflict clauses remain false:
 - (a) if $D = \Lambda \parallel C \cdot \sigma$ then $C\sigma$ is false in $\text{fgd}(M)$ and $\text{bgd}(M) \wedge \text{adiff}(B) \wedge \Lambda\sigma$ is satisfiable,
 - (b) if $M = M_1, L\sigma^{(\Lambda \parallel C \vee L) \cdot \sigma}, M_2$ then $C\sigma$ is false in $\text{fgd}(M_1)$, $L\sigma$ is undefined in M_1 , and $\text{bgd}(M_1) \wedge \text{adiff}(B) \wedge \Lambda\sigma$ is satisfiable.
5. All clauses in $N \cup U$ are pure. In particular, they don't contain any constants from B .

Lemma 7.17.11 (Rules preserve Well-Formed States). The rules of SCL(T) preserve well-formed states.

Proof. We prove each of the five properties by induction on the length of a derivation starting from the initial state $(\epsilon; N; \emptyset; B; k; \top)$. The induction step for the first two claims is

$$(M; N; U; B; k; D) \Rightarrow_{\text{SCL}(\text{T})} (M'; N'; U'; B'; k'; D').$$

1. In the initial state $(\epsilon; N; \emptyset; B; k; \top)$ constants can only appear in N and B , so it satisfies the claim. For the inductive step we do a case analysis on the rule application and prove $\text{con}((M'; N'; U'; B'; k'; D')) \subseteq \text{con}(N') \cup B'$ by case analysis on the rules of SCL(T). If we have applied Propagate or Decide then $N' = N$, $U' = U$, $B' = B$, $D' = D = \top$ and $M' = M, L\sigma$. So we only need to prove that $\text{con}(M') \subseteq \text{con}(N) \cup B$. Both rules require $|L\sigma| \in \text{atoms}(\text{grd}((\mathcal{S}, B, \Pi), N \cup U))$ satisfying the claim.

In case of the rules Grow, Skip, or Backtrack, then $\text{con}(N) \cup B \subseteq \text{con}(N') \cup B'$ and $\text{con}(M') \cup \text{con}(U') \cup \text{con}(D') \subseteq \text{con}(M) \cup \text{con}(U) \cup \text{con}(D)$.

If the rule Conflict was used then only the last component of the state changed $D' = \Lambda \parallel C \cdot \sigma$ with $\text{codom}(\sigma) \subseteq B$ and $\text{con}(\Lambda \parallel C) \subseteq \text{con}(N) \cup B$ by induction hypothesis.

If one of the rules Resolve or Factorize were used then as for Conflict the only component of the state that changed was the conflict clause and the constants in D' are a subset of the constants in M and D .

2. In the initial state $(\epsilon; N; \emptyset; B; k; \top)$ the condition $\text{adiff}(B)$ is satisfied as we assume constants in B to be distinct. For the inductive step we do a case analysis on the rule application. If the rule used was one of Conflict, Backtrack, Skip, Resolve, or Factorize, then $M = M', M''$ with M'' possibly empty and $B = B'$, so that $M \wedge \text{adiff}(B)$ satisfiable implies $M' \wedge \text{adiff}(B')$ to be satisfiable. If the rule Grow was used then we have $M' = \epsilon$ and that all constants in $B' = B \oplus B''$ are distinct, so that satisfiability of $\text{adiff}(B')$ is immediate. If the rule used was Propagate or Decide then we have $M' = M, L\sigma, \Lambda\sigma$ and $B' = B$, from the preconditions on the rules we also know that $L\sigma$ is undefined in M and that $\text{bgd}(M') \wedge \text{adiff}(B')$ is satisfiable.

3. By induction on the number of learned clauses. We prove that for each application of Backtrack

$$(M, K^{i+1}, M'; N; U; B; k; D \vee L \cdot \sigma) \Rightarrow_{\text{SCL(T)} \text{ Backtrack}} (M, L\sigma^{(\Lambda \parallel D \vee L) \cdot \sigma}, \Lambda'\sigma; N; U \cup \{\Lambda \parallel D \vee L\}; B; i; \top)$$

we have $N \cup U \models_{\mathcal{H}} D \vee L$. Following conflict resolution backward we can find a sequence of constrained closures $C_1 \cdot \sigma_1, \dots, C_n \cdot \sigma_n$, where $C_n \cdot \sigma_n = D \vee L \cdot \sigma$ such that $C_1 \in (N \cup U)$ is the most recent conflict clause, and C_{j+1} is either the result of a factorization on C_j or the result of a resolution inference between C_j and a clause in $(N \cup U)$. By induction on the length of conflict resolution and soundness of resolution and factoring we get $N \cup U \models_{\mathcal{H}} D \vee L$.

4. For the initial state the properties 4a and 4b obviously hold. For the induction step and an application of the rules Decide, Skip, and Grow there is nothing to show.

Consider a state $(M; N; U; B; k; \Delta \parallel D \cdot \delta)$ obtained by an application of Conflict. By the side conditions of Conflict $\text{adiff}(B) \wedge \text{bgd}(M) \wedge \Delta\delta$ is satisfiable and $\text{fgd}(M) \models \neg(D\delta)$ is shown for 4a. There is nothing to show for 4b.

Consider an application of rule Resolve

$$\begin{aligned} & (M, L\rho^{\Lambda\|(C\vee L)\cdot\rho}; N; U; B; k; \Lambda' \parallel (D \vee L') \cdot \sigma) \\ & \Rightarrow_{\text{SCL(T)}}^{\text{Resolve}} (M, L\rho^{\Lambda\|(C\vee L)\cdot\rho}; N; U; B; k; \Lambda \wedge \Lambda' \parallel (D \vee C)\eta \cdot \rho\sigma) \end{aligned}$$

by induction hypothesis $\text{bgd}(M) \wedge \text{adiff}(B) \wedge \Lambda\eta\rho\sigma$ is satisfiable and $C\eta\rho\sigma$ is false in $\text{fgd}(M)$, because $\Lambda\eta\rho\sigma = \Lambda\rho$ and $C\eta\rho\sigma = C\rho$ because η is the mgu and we always assume clauses to be variable disjoint. Using the same argument $\text{bgd}(M, L\sigma) \wedge \text{adiff}(B) \wedge \Lambda'\eta\delta\sigma$ is satisfiable and $(D \vee L')\eta\delta\sigma$ is false in $\text{fgd}(M, L\sigma)$. Therefore $(D \vee C)\eta \cdot \rho\sigma$ is false in $\text{fgd}(M, L\sigma)$ and $\text{bgd}(M, L\sigma) \wedge \text{adiff}(B) \wedge (\Lambda' \wedge \Lambda)\eta\delta\sigma$ is satisfiable, proving 4a. There is nothing to show for 4b.

For an application of the rule Factorize there is nothing to show for 4b and 4a obviously holds because the set of different ground literals in $(D \vee L \vee L')\sigma$ and $(D \vee L)\sigma$ is identical.

For an application of the rule Propagate there is nothing to show for 4a. For 4b consider the step

$$\begin{aligned} & (M; N; U; B; k; \top) \\ & \Rightarrow_{\text{SCL(T)}}^{\text{Propagate}} (M, L\sigma^{(\Lambda\|C_0\vee L)\delta\cdot\sigma}, \Lambda'\sigma; N; U; B; k; \top) \end{aligned}$$

where the side conditions of the rule imply the claim modulo the removal of duplicate literals $L\sigma$.

Finally, when applying Backtrack

$$\begin{aligned} & (M, K^{i+1}, M'; N; U; B; k; (\Lambda \parallel D \vee L) \cdot \sigma) \\ & \Rightarrow_{\text{SCL(T)}}^{\text{Backtrack}} (M, L\sigma^{(\Lambda\|D\vee L)\cdot\sigma}, \Lambda'\sigma; N; U \cup \{\Lambda \parallel D \vee L\}; B; i; \top) \end{aligned}$$

there is nothing to show for 4a. For 4b we know by induction hypothesis that $(D \vee L)\sigma$ is false in $\text{fgd}(M, K^{i+1}, M')$. The literal $L\sigma$ is of level k and $D\sigma$ of level i , $k > i$, hence $D\sigma$ is false in $\text{fgd}(M)$ and $L\sigma$ undefined in $\text{fgd}(M)$. Furthermore, $\text{bgd}(M, K^{i+1}, M') \wedge \text{adiff}(B) \wedge \Lambda\sigma$ is satisfiable by induction hypothesis, so $\text{bgd}(M) \wedge \text{adiff}(B) \wedge \Lambda\sigma$ is satisfiable as well.

4. For the initial state all clauses are pure by assumption. Conflict picks a clause from $N \cup U$ that is pure by induction hypothesis. Resolve and Factorize only apply unifiers between pure literals to the resulting clause, hence also only produce pure clauses from pure clauses. Finally, Backtrack adds the pure learned clause to $N \cup U$. \square

Definition 7.17.12 (Stuck State). A state $(M; N; U; B; k; D)$ is called *stuck* if $D \neq \Lambda \parallel \perp \cdot \sigma$ and none of the rules Propagate, Decide, Conflict, Resolve, Factorize, Skip, or Backtrack is applicable.

Proposition 7.17.13 (Form of Stuck States). If a run (without rule Grow) ends in a stuck state $(M; N; U; B; k; D)$ where Conflict was applied eagerly, then $D = \top$ and all ground foreground literals that can be build from the foreground literals in N by instantiation with constants from B are defined in M .

Proof. First we prove that stuck states never appear during conflict resolution. Consider a well-formed state $(M; N; U; B; k; \Delta \parallel D \cdot \delta)$, we prove by case analysis that either Skip, Resolve, Factorize or Backtrack can be applied. If $M = M', L\sigma$ and $L\sigma$ is either a background literal or a foreground literal such that $\text{comp}(L\sigma)$ is not contained in $D\delta$ then Skip can be applied. If $M = M', L\sigma^{\wedge \parallel C \cdot \sigma}$ with $D\delta = D' \vee \text{comp}(L\sigma)$ then Resolve can be applied. If $M = M', L\sigma^k, M''$ and D' contains multiple occurrences of $\text{comp}(L\sigma)$ then Factorize can be applied. In summary, we can reach a state with a unique literal $L\delta$ of level k in $D\delta$. Then Backtrack is applicable. Finally, if in some state $(M; N; U; B; k; \top)$ where Conflict is not applicable, some atom $|L| \in \text{atoms}(\text{grd}((\mathcal{S}, B, \Pi), N))$ is undefined, we can always apply Decide. \square

Lemma 7.17.14 (Stuck States Produce Ground Models). If a state $(M; N; U; B; k; \top)$ is stuck then $M \wedge \text{adiff}(B) \models \text{grd}((\mathcal{S}, B, \Pi), N \cup U)$.

Proof. By contradiction. Note that $M \wedge \text{adiff}(B)$ is satisfiable, Lemma 7.17.11.2. Consider any clause $(\Lambda \parallel C)\sigma \in \text{grd}((\mathcal{S}, B, \Pi), N \cup U)$. It can only be not true in $M \wedge \text{adiff}(B)$ if $\text{fgd}(M) \models \neg(C\sigma)$ and $\text{bgd}(M) \wedge \text{adiff}(B) \wedge \Lambda\sigma$ is satisfiable. But then Conflict would be applicable, a contradiction. \square

Example 7.17.15 (SCL(T) Model Extraction). In some cases it is possible to extract an overall model from the ground trail of a stuck state of an SCL(T) derivation. Consider $B = [a, b, c]$ and a satisfiable BS(LRA) constrained clause set $N = \{x \geq 1 \parallel P(x), x < 0 \parallel P(x), 0 \leq x \wedge x < 1 \parallel \neg P(x), 2x \geq 1 \parallel P(x) \vee Q(x)\}$. Starting from state $(\epsilon; N; \emptyset; B; 0; \top)$ and applying Propagate fairly a regular run can derive the following trail

$$M = P(a)^{x \geq 1 \parallel P(x) \cdot \{x \mapsto a\}}, a \geq 1, P(b)^{x < 0 \parallel P(x) \cdot \{x \mapsto b\}}, b < 0, \\ \neg P(c)^{0 \leq x \wedge x < 1 \parallel \neg P(x) \cdot \{x \mapsto c\}}, 0 \leq c, c < 1, Q(c)^{2x \geq 1 \parallel P \vee Q(x) \cdot \{x \mapsto c\}}, 2c \geq 1$$

The state $(M; N; \emptyset; B; 0; \top)$ is stuck and $M \models_{\mathcal{H}} \text{grd}((\mathcal{S}, B, \Pi), N)$. Moreover from M we can generate an interpretation $\mathcal{A}^{\text{BS(LRA)}}$ of N by generalizing the foreground constants used for instantiation and interpreting the predicates P and Q as formulas over $\Sigma^{\mathcal{B}}$, $P^{\mathcal{A}} = \{q \in \mathbb{Q} \mid q < 0 \vee q \geq 1\}$ and $Q^{\mathcal{A}} = \{q \in \mathbb{Q} \mid 2q \geq 1 \wedge q < 1\}$.

Lemma 7.17.16 (Soundness). If a derivation reaches the state $(M; N; U; B; k; \Lambda \parallel \perp \cdot \sigma)$, then N is unsatisfiable.

Proof. All learned clauses are consequences of $N \cup U$, Lemma 7.17.11.3. Furthermore $\text{bgd}(M) \wedge \text{adiff}(B) \wedge \Lambda\sigma$ is satisfiable, Lemma 7.17.11.4a. \square

Definition 7.17.17 (Reasonable Run). A sequence of SCL(T) rule applications is called a *reasonable run* if the rule Decide is only applied if there exists no application of the rule Propagate that would generate a conflict.

Definition 7.17.18 (Regular Run). A sequence of SCL(T) rule applications is called a *regular run* if it is a reasonable run the rule Conflict has precedence over all other rules, and Resolve resolves away at least the rightmost foreground literal from the trail.

Example 7.17.19 (SCL(T) Refutation). Given a set of foreground constants $B = [a, b, c]$ and a BS(LRA) constrained clause set $N = \{C_1: x = 0 \parallel P(x), C_2: y = x + 1 \parallel \neg P(x) \vee P(y), C_3: z = 2 \parallel \neg P(z)\}$ the following is a regular derivation

$$\begin{aligned}
& (\epsilon; N; \emptyset; B; 0; \top) \\
\Rightarrow_{\text{SCL(T)} \text{ Propagate}} & (P(a)^{C_1 \cdot \{x \mapsto a\}}, a = 0; N; \emptyset; B; 0; \top) \\
\Rightarrow_{\text{SCL(T)} \text{ Propagate}} & (\dots, P(b)^{C_2 \cdot \{x \mapsto a, y \mapsto b\}}, b = a + 1; N; \emptyset; B; 0; \top) \\
\Rightarrow_{\text{SCL(T)} \text{ Propagate}} & (\dots, P(c)^{C_2 \cdot \{x \mapsto b, y \mapsto c\}}, c = b + 1; N; \emptyset; B; 0; \top) \\
\Rightarrow_{\text{SCL(T)} \text{ Conflict}} & (\dots, P(c)^{C_2 \cdot \{x \mapsto b, y \mapsto c\}}, c = b + 1; N; \emptyset; B; 0; z = 2 \parallel \neg P(z) \cdot \{z \mapsto c\}) \\
\Rightarrow_{\text{SCL(T)} \text{ Resolve}} & (\dots, P(c)^{C_2 \cdot \{x \mapsto b, y \mapsto c\}}, c = b + 1; N; \emptyset; B; 0; \\
& z = x + 1 \wedge z = 2 \parallel \neg P(x) \cdot \{z \mapsto c, x \mapsto b\}) \\
\Rightarrow_{\text{SCL(T)} \text{ Skip}} & (\dots, P(b)^{C_2 \cdot \{x \mapsto a, y \mapsto b\}}, b = a + 1; N; \emptyset; B; 0; \\
& z = x + 1 \wedge z = 2 \parallel \neg P(x) \cdot \{z \mapsto c, x \mapsto b\}) \\
\Rightarrow_{\text{SCL(T)} \text{ Resolve}} & (\dots, P(b)^{C_2 \cdot \{x \mapsto a, y \mapsto b\}}, b = a + 1; N; \emptyset; B; 0; \\
& z = x + 1 \wedge z = 2 \wedge x = x_1 + 1 \parallel \neg P(x_1) \cdot \{z \mapsto c, x \mapsto b, x_1 \mapsto a\}) \\
\Rightarrow_{\text{SCL(T)} \text{ Skip}} & (P(a)^{C_1 \cdot \{x \mapsto a\}}, a = 0; N; \emptyset; B; 0; \\
& z = x + 1 \wedge z = 2 \wedge x = x_1 + 1 \parallel \neg P(x_1) \cdot \{z \mapsto c, x \mapsto b, x_1 \mapsto a\}) \\
\Rightarrow_{\text{SCL(T)} \text{ Resolve}} & (P(a)^{C_1 \cdot \{x \mapsto a\}}, a = 0; N; \emptyset; B; 0; \\
& z = x + 1 \wedge z = 2 \wedge x = x_1 + 1 \wedge x_1 = 0 \parallel \perp \cdot \{z \mapsto c, x \mapsto b, x_1 \mapsto a\})
\end{aligned}$$

N is proven unsatisfiable as we reach a state in the form $(M; N; U; B; k; \Lambda \parallel \perp \cdot \sigma)$.

Example 7.17.20 (SCL(T) Clause learning). Given an initial constant set $B = [a]$ of fresh foreground constants and a BS(LRA) constrained clause set $N = \{C_1: x \geq y \parallel \neg P(x, y) \vee Q(z), C_2: z = u + v \parallel \neg P(u, v) \vee \neg Q(z), \}$ the following is an example of a regular run

$$\begin{aligned}
& (\epsilon; N; \emptyset; B; 0; \top) \\
\Rightarrow_{\text{SCL(T)} \text{ Decide}} & (P(a, b)^1; N; \emptyset; B; 1; \top) \\
\Rightarrow_{\text{SCL(T)} \text{ Propagate}} & (P(a, a)^1, Q(a)^{C_1 \cdot \{x \mapsto a, y \mapsto a, z \mapsto a\}}, a \geq a; N; \emptyset; B; 1; \top) \\
\Rightarrow_{\text{SCL(T)} \text{ Conflict}} & (P(a, a)^1, Q(a)^{C_1 \cdot \{u \mapsto a, v \mapsto a, z \mapsto a\}}, a \geq a; N; \emptyset; B; 1; \\
& C_2 \cdot \{x \mapsto a, y \mapsto a, z \mapsto a\}) \\
\Rightarrow_{\text{SCL(T)} \text{ Resolve}} & (P(a, a)^1, Q(a)^{C_1 \cdot \{x \mapsto a, y \mapsto a, z \mapsto a\}}, a \geq a; N; \emptyset; B; 1; x \geq y \wedge z = u + v \parallel \\
& \neg P(x, y) \vee \neg P(u, v) \cdot \{x \mapsto a, y \mapsto a, z \mapsto a, u \mapsto a, v \mapsto a\}) \\
\Rightarrow_{\text{SCL(T)} \text{ Skip}^*} & (P(a, a)^1; N; \emptyset; B; 1; x \geq y \wedge z = u + v \parallel \\
& \neg P(x, y) \vee \neg P(u, v) \cdot \{x \mapsto a, y \mapsto a, z \mapsto a, u \mapsto a, v \mapsto a\}) \\
\Rightarrow_{\text{SCL(T)} \text{ Factorize}} & (P(a, a)^1; N; \emptyset; B; 1; x \geq y \wedge z = x + y \parallel \neg P(x, y) \cdot \{x \mapsto a, y \mapsto a, z \mapsto a\}) \\
\Rightarrow_{\text{SCL(T)} \text{ Backtrack}} & (\neg P(a, a)^{(x \geq y \wedge z = x + y \parallel \neg P(x, y)) \cdot \{x \mapsto a, y \mapsto a\}}, a \geq a, a = a + a; N; \\
& \{x \geq y \wedge z = x + y \parallel \neg P(x, y)\}; B; 1; \top)
\end{aligned}$$

In this example the learned clauses $x \geq y \wedge z = x + y \parallel \neg P(x, y)$; note how there are two distinct variables in the learned clause even if we had to use a single constant for instantiations in conflict search.

Proposition 7.17.21. Let N be a set of constrained clauses. Then any application of Decide in an SCL(T) regular run from starting state $(\epsilon; N; \emptyset; B; 0; \top)$ does not create a conflict.

Proof. Assume the contrary: then Propagate would have been applicable before Decide, contradicting with the definition of a regular and hence reasonable run. \square

Corollary 7.17.22. Let N be a set of constrained clauses. Then any conflict in an SCL(T) regular run from starting state $(\epsilon; N; \emptyset; B; 0; \top)$ admits a regular conflict resolution.

Proof. We need to prove that it is possible to apply Resolve during conflict resolution. By Proposition 7.17.21 the rightmost foreground literal on the trail is a propagation literal and by regularity we know that this literal appears in the conflict clause. So a conflict resolution can start by skipping over the background literals and then resolving once with the rightmost foreground literal. \square

Lemma 7.17.23 (Non-Redundant Clause Learning). Let N be a set of constrained clauses. Then clauses learned in an SCL(T) regular run from starting state $(\epsilon; N; \emptyset; B; 0; \top)$ are not redundant.

Proof. Consider the following fragment of a derivation learning a clause:

$$\begin{aligned} &\Rightarrow_{\text{SCL(T)}^{\text{Conflict}}} (M''; N; U; B; k; \Lambda_0 \parallel C_0 \cdot \sigma_0) \\ &\Rightarrow_{\text{SCL(T)}^{\{\text{Skip}, \text{Factorize}, \text{Resolve}\}^*}} (M, K^{i+1}, M'; N; U; B; k; \Lambda_n \parallel C_n \cdot \sigma_n) \\ &\Rightarrow_{\text{SCL(T)}^{\text{Backtrack}}} (M, L\sigma^{(\Lambda_n \parallel D \vee L)} \cdot \sigma, \Lambda'_n \sigma; N; U \cup \{\Lambda_n \parallel D \vee L\}; B; i; \top). \end{aligned}$$

where $C_n = D \vee L$ and $\sigma = \sigma_n$. Let \prec be any \mathcal{H} order induced by M . We prove that $\Lambda_n \sigma \parallel C_n \sigma$ is not redundant with respect to \prec , B , and $(N \cup U)$. By soundness of hierarchic resolution $(N \cup U) \models \Lambda_n \parallel C_n$ and $\Lambda_n \sigma$ is satisfiable with $M \wedge \text{adiff}(B)$, and $C_n \sigma$ is false under both M and M, K^{i+1}, M' , Lemma 7.17.11. For a proof by contradiction, assume there is a $N' \subseteq \text{grd}((\mathcal{S}, B, \Pi), N \cup U) \preceq^{\Lambda_n \sigma \parallel C_n \sigma}$ such that $N' \models_{\mathcal{H}} \Lambda_n \sigma \parallel C_n \sigma$. As $\Lambda_n \sigma \parallel C_n \sigma$ is false under M , there is a ground constrained clause $\Lambda' \parallel C' \in N'$ with $\Lambda' \parallel C' \preceq \Lambda_n \sigma \parallel C_n \sigma$, and all literals from C' are defined in M and false by the definition of \prec . Furthermore, we can assume that $\text{adiff}(B) \wedge \text{bgd}(M) \wedge \Lambda'$ is satisfiable or $C_n \sigma$ would be a tautology, because $\text{adiff}(B) \wedge \text{bgd}(M) \wedge \Lambda_n \sigma$ is satisfiable.

The clause $\Lambda_0 \sigma_0 \parallel C_0 \sigma_0$ has at least one literal of level k and due to a regular run, Definition 7.17.18, the rightmost trail literal is resolved away in $\Lambda_n \sigma \parallel C_n \sigma$, Corollary 7.17.22. Therefore, the rightmost foreground literal does not appear in $\Lambda' \parallel C'$, so by regularity $\Lambda' \parallel C'$ would have created a conflict at a previous state. \square

Of course, in a regular run the ordering of foreground literals on the trail will change, i.e., the ordering underlying Lemma 7.17.23 will change as well. Thus the non-redundancy property of Lemma 7.17.23 reflects the situation at the time of creation of the learned clause. A non-redundancy property holding for an overall run must be invariant against changes on the ordering. However, the ordering underlying Lemma 7.17.23 also entails a fixed subset ordering that is invariant against changes on the overall ordering. This means that our dynamic ordering entails non-redundancy criteria based on subset relations including forward redundancy. From an implementation perspective, this means that learned clauses need not to be tested for forward redundancy. Current resolution, or superposition based provers spent a reasonable portion of their time in testing forward redundancy of newly generated clauses. In addition, also tests for backward reduction can be restricted knowing that learned clauses are not redundant.

Lemma 7.17.24 (Termination of SCL(T)). Let N be a set of constrained clauses and B be a finite set of background constants. Then any regular run with start state $(\epsilon; N; \emptyset; B; 0; \top)$ that uses Grow only finitely often terminates.

Proof. Since Grow can only be used a finite number of times we consider as a start state the state after the final application of Grow and prove termination of runs that never use Grow. We do so by giving an explicit termination measure on the SCL(T) states. Given a state $(M; N; U; B; k; D)$ we define a termination measure μ as $\mu(M; N; U; B; k; D) = (u, s, m, r, d) \in \mathbb{N}^5$ with a lexicographical combination of $>$ where

- $l = |\text{atoms}(\text{grd}((\mathcal{S}, B, \Pi), N \cup U))|$, $u = 3^l - |\text{grd}((\mathcal{S}, B, \Pi), U)|$, and $m = |M|$,
- in the case $D = \top$:
 - * $s = 1 + l - m$, $d = 0$, and $r = 0$,
- otherwise if $D = \Delta \parallel D' \cdot \delta$:
 - * $s = 0$,
 - * if $M = M', L$ with L foreground literal then r is the number of copies of L in $D'\delta$
 - * if the rightmost literal of M is a background literal or if M is empty then $r = 0$
 - * d is the number of literals in D'

The number of ground atoms $l = |\text{atoms}(\text{grd}((\mathcal{S}, B, \Pi), N \cup U))|$ is an upper bound to the length of the trail because the trail is consistent and no literal can appear more than once on the trail. Similarly, every learned clause has at least one non-redundant ground instance so $|\text{grd}((\mathcal{S}, B, \Pi), U)|$ increases whenever SCL(T) learns a new clause and 3^l is an upper bound to the ground instances of all learned clauses in a regular run. This means that Backtrack strictly decreases

u , Decide, Propagate, and Conflict strictly decrease s without modifying u , Skip strictly decreases m without modifying u or s , Resolve strictly decreases r without modifying u , s , or m , and finally Factorize strictly decreases d possibly decreases r and does not modify u , s , or m . \square

Theorem 7.17.25 (Hierarchic Herbrand Theorem). Let N be a finite set of clauses. N is unsatisfiable iff there exists a finite set $N' = \{\Lambda_1 \parallel C_1, \dots, \Lambda_n \parallel C_n\}$ of variable renamed copies of clauses from N and a finite set B of fresh constants and a substitution σ , grounding for N' where $\text{codom}(\sigma) = B$ such that $\bigwedge_i \Lambda_i \sigma$ is \mathcal{T}^B satisfiable and $\bigwedge_i C_i \sigma$ is first-order unsatisfiable over $\Sigma^{\mathcal{F}}$.

Proof. Recall that N is a pure, abstracted clause set and that \mathcal{T}^B is term-generated, compact background theory that contains an equality $=$, and that all constants of the background signature are domain constants. Then by completeness of hierarchic superposition [?], N is unsatisfiable iff there exists a refutation by hierarchic superposition. Let $N' = \{\Lambda_1 \parallel C_1, \dots, \Lambda_n \parallel C_n\}$ be a finite set renamed copies of clauses from N such that there is a refutation by hierarchic superposition such that each clause in N' and each derived clause is used exactly once. This set exists because the refutation is finite and any hierarchic superposition refutation can be transformed into a refutation where every clause is used exactly once. Now let δ be the overall unifier of this refutation. This unifier exists, because all clauses in N' have disjoint variables and all clauses in the refutation are used exactly once. Now we consider a finite set of constants B and a substitution σ , $\text{codom}(\sigma) = B$, σ grounding for N' , and for all $x, y \in \text{dom}(\delta)$ we have $x\sigma = y\sigma$ iff $x\delta = y\delta$. Now there is also a refutation for $N'\sigma$ by hierarchic superposition where the clauses are inferred exactly in the way they were inferred for N' . It remains to be shown that $\bigwedge_i \Lambda_i \sigma$ is \mathcal{T}^B satisfiable and $\bigwedge_i C_i \sigma$ is $\mathcal{A}^{\mathcal{H}}$ unsatisfiable. The hierarchic superposition refutation terminates with the clause $\bigwedge_i \Lambda_i \sigma \parallel \perp$ where $\bigwedge_i \Lambda_i \sigma$ is satisfiable. Furthermore, the refutation derives \perp from $\{C_1\sigma, \dots, C_n\sigma\}$ via superposition, proving the theorem. \square

Finally, we show that an unsatisfiable clause set can be refuted by SCL(T) with any regular run if we start with a sufficiently large sequence of constants B and apply Decide in a fair way. In addition, we need a Restart rule to recover from a stuck state.

Restart $(M; N; U; B; k; \top) \Rightarrow_{\text{SCL(T)}} (\epsilon; N; U; B; 0; \top)$

Of course, an unrestricted use of rule Restart immediately leads to non-termination.

Theorem 7.17.26 (Refutational Completeness of SCL(T)). Let N be an unsatisfiable clause set. Then any regular SCL(T) run will derive the empty clause provided (i) Rule Grow and Decide are operated in a fair way, such that all possible trail prefixes of all considered sets B during the run are eventually explored, and (ii) Restart is only applied to stuck states.

Proof. If N is unsatisfiable then by Theorem 7.17.25 there exists a finite set $N' = \{\Lambda_1 \parallel C_1, \dots, \Lambda_n \parallel C_n\}$ of variable renamed copies of clauses from N and a finite set B of fresh constants and a substitution σ , grounding for N' where $\text{codom}(\sigma) = B$ such that $\bigwedge_i \Lambda_i \sigma$ is \mathcal{T}^B satisfiable and $\bigwedge_i C_i \sigma$ is first-order unsatisfiable over $\Sigma^{\mathcal{F}}$. If the SCL(T) rules are applied in a fair way, then they will in particular produce trails solely consisting of literals from $N' \sigma$. For these trails all theory literals are satisfiable, because $\bigwedge_i \Lambda_i \sigma$ is \mathcal{T}^B satisfiable. Furthermore, the states corresponding to these trails cannot end in a stuck state, because this contradicts the unsatisfiability of $\bigwedge_i C_i \sigma$. Instead, they all end in a conflict with some clause in $N' \sigma$. In addition, there are only finitely many such trails, because the number of literals in $N' \sigma$ is finite. Now let $\mu((M; N; U; B; k; \top))$ be the multiset of the levels of all states with trails from $N' \sigma$ until a conflict occurs. Each time a state with a trail from $N' \sigma$ results in a conflict, SCL(T) learns a non-redundant clause that propagates at a strictly smaller level, Lemma 7.17.23. Thus $\mu((M; N; U; B; k; \top))$ strictly decreases after each Backtrack step after a conflict on a trail with atoms from $N' \sigma$. The clause learnt at level zero is the empty clause. \square

Condition (i) of the above theorem is quite abstract. It can, e.g., be made effective by applying rule Grow only after all possible trail prefixes with respect to the current set B have been explored and to make sure that Decide does not produce the same stuck state twice.