

Implementing Superposition in iProver (System Description)

André Duarte and Konstantin Korovin

The University of Manchester
{andre.duarte,konstantin.korovin}@manchester.ac.uk

Abstract iProver is a saturation theorem prover for first-order logic with equality, which is originally based on an instantiation calculus Inst-Gen. In this paper we describe an extension of iProver with the superposition calculus. We have developed a flexible simplification setup that subsumes and generalises common architectures such as DISCOUNT or Otter. This also includes the concept of “immediate simplification”, wherein newly derived clauses are more aggressively simplified among themselves, and the concept of “light normalisation”, wherein ground unit equations are stored in an interreduced TRS which is then used to simplify new clauses. We have also added support for associative-commutative theories (AC), namely by deletion of AC-joinable clauses, semantic detection of AC axioms, and preprocessing normalisation.

iProver¹ [11] is an automated theorem prover for first-order logic. It is a saturation prover, and is based primarily on the Inst-Gen calculus [8], but also implements resolution and supports running them in combination in an abstraction-refinement framework [10, 13]. In this work we detail how iProver was extended with support for the superposition calculus.

Currently, iProver deals with equality axiomatically, which can be inefficient for problems heavy on equality. At the same time, the superposition calculus is a set of complete inference rules specialised for first-order logic with equality. It complements the instantiation calculus since it is effective on problems where instantiation struggles, and vice-versa. We show that running the two calculi in combination yields better results than either plain instantiation or plain superposition.

Rules in a calculus can be classified as “generating”, if they derive new clauses, or “simplifying”, if some premise gets deleted. While generating rules are the ones necessary for completeness of a calculus, simplification rules are crucial for practical performance. Intuitively, simplification rules are beneficial to taming the growth of the search space as more clauses get generated. However, the computation of those simplifications itself takes time, so being too eager in applying them will also grind the prover to a halt. It is an open problem, what the optimal strategy to balance these conflicting requirements is, and although there is a huge amount of flexibility in how to perform simplifications (see e.g., [22]), most provers are rather restrictive about this. In iProver we developed a flexible simplification setup that subsumes and generalises most common architectures.

¹ Available at <http://www.cs.man.ac.uk/~korovink/iprover>.

Finally, we have also implemented specialised techniques to deal with associative-commutative (AC) theories. These are theories of great interest which arise in several domains [21], and are traditionally problematic for theorem provers to deal with, due to combinatorial explosion and the non-orientability of the AC axioms.

The paper is structured as follows: first, we give a quick overview of the architecture of iProver. Then, we describe the implementation of superposition in iProver, its modifications, the simplification architecture and given clause loop, and AC reasoning rules. For a more in-depth description of basic features of iProver, see [10].

1 Overview

iProver is based on the Inst-Gen calculus, which is based on the following idea. We approximate the first-order problem to a propositional problem, and submit it to a black box SAT solver. It either finds an inconsistency, which is also an inconsistency at the first-order level, or else it returns a model, which guides the instantiation of new clauses whose abstraction witnesses some inconsistency at the ground level. If no such instantiation exists, then the problem is satisfiable.

The SAT solver is also used to implement “global” simplification rules (in the sense that they involve reasoning with the clause set which is shared between different calculi), such as global propositional subsumption [10]. In a nutshell, we submit ground abstractions of clauses in S to a set S_{gr} . If the SAT solver finds that S_{gr} propositionally implies $D\gamma$, with D a strict subset of C and γ an injective substitution of variables to fresh constants, then we can replace $C\theta$, in S , by D .

$$\text{Glob. prop. subs.} \quad \frac{C\theta}{D}, \quad \text{where } D \subsetneq C \text{ and } C \in \tilde{S}, S \models S_{gr}, S_{gr} \models D\gamma \quad (1)$$

As mentioned before, iProver can also run other calculi. This is beneficial because (i) some problems are solved easily by one strategy and not by others, and (ii) clauses derived in e.g. resolution can be passed to the instantiation solver to participate in simplifications. For example, clauses derived in all calculi are submitted to a shared global propositional subsumption set, which is in turn used by all calculi to simplify its clauses.

Schematically, the high-level architecture of iProver is summarised in Figure 1. We can view it as a modular architecture where each calculus (Inst-Gen, resolution, superposition) runs its own saturation loop, and can (i) query external SAT and SMT solvers, and (ii) submit clauses to, and retrieve clauses from, the ‘Exchange’ module. The instantiation and resolution modules are discussed in [10]. Here we will focus on the superposition calculus.

2 Extension with superposition

The superposition inference system consists of the following rules [16]:²

$$\text{Superposition} \quad \frac{l \approx r \vee C \quad t[s] \approx u \vee D}{(t[s \mapsto r] \approx u \vee C \vee D)\theta} \quad (2)$$

² ‘ \approx ’ means ‘ \simeq ’ or ‘ $\not\approx$ ’; also rules are to be read modulo flipping the equalities.

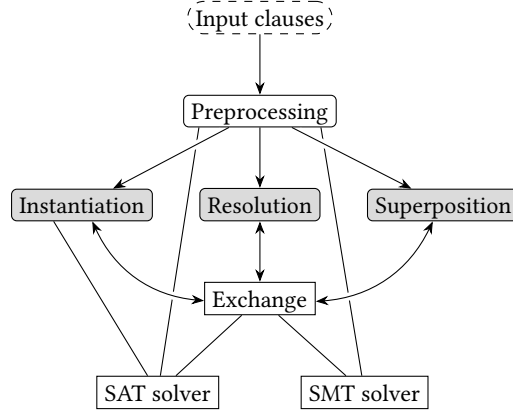


Figure 1: Architecture of iProver.

where $\theta = \text{mgu}(l, s)$, $l\theta \not\approx r\theta$, $t\theta \not\approx u\theta$, and s not a variable,

$$\text{Eq. Resolution} \quad \frac{l \not\approx r \vee C}{C\theta} \quad \text{where } \theta = \text{mgu}(l, r), \quad (3)$$

$$\text{Eq. Factoring} \quad \frac{l \approx r \vee l' \approx r' \vee C}{(l \approx r \vee r \not\approx r' \vee C)\theta} \quad \text{where } \theta = \text{mgu}(l, l'), \quad (4)$$

$l\theta \not\approx r\theta \text{ and } r\theta \not\approx r'\theta.$

We assume that \prec is a simplification ordering. Non-equality predicates $P(t)$ are encoded as $P(t) \approx \top$. The calculus can easily be generalised to the many-sorted case, which iProver uses even in untyped problems, since it can perform subtype inference during preprocessing. Superposition is sound and refutationally complete for first-order logic with equality (see [3, 16]) and implemented in a number of state-of-the-art theorem provers: Vampire [12], E [20] and SPASS [23]. Currently, iProver uses non-perfect discrimination trees to find unification candidates efficiently [9, 17]. For the literal selection, to ensure completeness, we must select either a negative literal, or all maximal literals. In iProver we use a variant of the Knuth-Bendix ordering which prioritises non-equational literals.

Simultaneous superposition In (2), by $t[s]$ and $t[s \mapsto r]$ we can mean resp. “a distinguished occurrence of s as a subterm of t ” and “replacing that subterm at that position by r ”. We call the variant *simultaneous superposition* where we mean instead “replacing all occurrences of s in t by r ”. This variant is still refutationally complete [4]. In cases where there are several occurrences of the same term, as in $f(\underbrace{s, s, \dots, s}_{n \text{ times}})$,

this avoids producing $2^n - 1$ intermediate clauses with $f(r, s, \dots, s)$, $f(s, r, \dots, s)$, $f(r, r, \dots, s)$, etc., instead producing only $f(r, r, r, \dots, r)$. In iProver, we implement this variant of superposition.

2.1 Simplifications

Apart from the generating inferences, necessary for completeness, we can add *simplification inferences*. In our implementation, we use the following rules: tautology deletion, syntactic equality resolution, subsumption, subset subsumption, subsumption resolution, demodulation and global subsumption [10, 12, 18, 23].

Light normalisation In addition, we introduce the following rule:

$$\text{Light Normalisation} \quad \frac{R \quad C[l]}{C[l \mapsto r]}, \quad \text{where } l \rightarrow r \in R \quad (5)$$

where R is a set of interreduced wrt. (5) oriented rewrite rules. It can be seen as a restricted case of the demodulation rule, but it is advantageous to formulate this separately because it may be implemented much more efficiently than demodulation, by simply looking up terms in a hashtable, rather than having to do matching with variable instantiations.

A *light normalisation index* consists of (i) a hashtable that indexes rewrite rules in R by their left-hand sides for forward light normalisations and (ii) a map of all subterms in R for keeping R interreduced. When we derive a unit equality, we first normalise it wrt. R by recursively replacing each subterm by its normal form wrt. R . Then, if the simplified equality is orientable (wrt. \prec) we use it to normalise rules in R add it to R . If there is a conflict between two rules $t \rightarrow s$ and $t \rightarrow u$ where $t \succ s \succ u$, we keep the rules $t \rightarrow u$ and $s \rightarrow u$. If s and u are incomparable wrt. \succ we keep one of the rules in R . Since R is only used for simplifications this choice does not affect the completeness. In general, we can restrict which orientable equations we add to R (e.g. only ground ones, or small in size).

2.2 Given clause algorithm

In a standard given clause loop, the clause set is split into an active set, where inferences among the clauses have been performed, and a passive set, of clauses waiting to participate in inferences. Clauses are initially added to the passive, then in each iteration one given clause is picked from the passive set, added to the active set, and all inferences between given and active are performed. Newly derived clauses are pushed into the passive. The loop finishes when all clauses have been moved to the active set, meaning the initial set is satisfiable, or when a contradiction is derived.

Immediate simplification set Next, we introduce the idea of *immediate simplification*. The intuition is as follows. Clauses that are derived in each loop are “related” to each other. It may be beneficial to keep the set of immediate conclusions inter-simplified. Also, throughout the execution of the program the set of generated clauses in each loop remains small compared to the set of passive or active clauses. Therefore, we can get away with applying more expensive rules that we do not necessarily want to apply on the set of all clauses (e.g. only “light” simplifications between newly derived clauses

and passive clauses, but more expensive “full” simplifications among newly derived clauses). Finally, during this process, it is possible that the given clause itself becomes redundant (e.g. subsumed by one of its children). If this happens, we can add *only* the clauses responsible for making it redundant to the passive set, then remove the given clause from the active set, and throw away the rest of the iteration’s newly generated clauses, abort the rest of the iteration, and proceed to the next given clause. In some problems, a significant number of iterations may be aborted, which means that fewer new clauses are added to the passive queue, and that we avoid the work of computing those inferences. This can be seen as a variant of orphan elimination [18]. Even when the given clause is not eliminated it is often beneficial to extensively inter-simplify immediate descendants of the given clause.

Simplification setup How all these simplifications are performed can greatly impact the performance of the solver, so care is needed, and tuning this part of the solver can pay off significantly. There is a significant amount of choice in how to perform simplifications. We can choose which simplifications to perform, and at what times, and with respect to which clauses. Additionally, some of these simplifications require auxiliary data structures (here referred to generally as “indices”) to be done efficiently, and some indices support several simplification rules. Therefore we also need to choose which clauses to add to which indices at which stages.

For example, Otter-style loops [15] perform simplifications on clauses before adding them to the passive set. The problem with this is that the passive set is often orders of magnitude larger than the active set, therefore performance will degrade significantly as this set grows, and the system will spend most of its time performing simplifications on clauses that may not even end up being used. On the other hand, DISCOUNT-style loops [5] perform simplifications only with clauses that have been added to the active set. This has the benefit of reducing the time spent in simplifications, at the cost of potentially missing many valuable simplifications wrt. passive clauses. It is not clear where the “sweet spot” is, in terms of these setups.

It is possible, for example, to choose to apply only “cheap” simplifications to the full active + passive set (e.g. subset subsumption, and light normalisation), and use more expensive ones only on the small active set (e.g. full subsumption and demodulation). In Listing 1.1 we describe the ‘iProver-Sup’ given-clause saturation loop for superposition. A *simplification set* consists of a collection of indices, each of which supports one or more simplification rules. In our given clause loop we have four such sets: S_{passive} , to which we add the clauses added to the passive set, S_{active} , for the clauses in the active set, S_{immed} for newly derived clauses (this set is cleared at the end of every given-clause iteration, and the non-redundant clauses added to the passive queue), and S_{input} for preprocessing input clauses. Each set supports the following operations: *add*, which adds a clause to all indices in a set S , and *simplify*, which simplifies via some rules R wrt. a set S . These are called at several points in the loop (see Listing 1.1), and the user can configure which indices/rules are involved in each operation. When simplifying, some rules *forward simplify* the clause wrt. the existing set, and others *backward simplify* the clauses in the set wrt. the new clause.

S_{passive}	SubsetSubsumption, PropSubs
S_{active}	Subsumption, LightNorm, FwDemod, BwDemod
S_{immed}	SubsetSubsumption, Subsumption, LightNorm, FwDemod, BwDemod
S_{input}	SubsetSubsumption, Subsumption, LightNorm, FwDemod, BwDemod
R_{passive}	TrivRules, ACJoinability, FwLightNormDemod, FwSubsumption
R_{active}	TrivRules, FwPropSubs, FwLightNormDemod, FwSubsumption, FwSubsumptionRes, BwDemod
R_{immed}	TrivRules, FwLightNormDemod, FwSubsumption, FwSubsumptionRes, BwDemod, BwSubsumption
R_{input}	TrivRules, FwLightNormDemod, FwSubsumption, FwSubsumptionRes, BwDemod, BwSubsumption, BwSubsumptionRes

Table 1: Default simplification options

The simplification setup is specified by the rules to apply at each stage (R_x), and the indices to which to add at each stage (S_x). These can be specified by the user via command-line options. S_x are lists of indices from {Subsumption, SubsetSubsumption, FwDemod, BwDemod, LightNorm, PropSubsSet}. R_x are lists of rules from {EqResSimp, TautologyElim, EqTautologyElim, TrivRules, FwPropSubs, FwSubsumption, FwSubsumptionStrict, FwSubsumptionRes, FwDemod, FwLightNorm, FwLightNormDemodLoop, ACJoinability, BwSubsumption, BwSubsumptionRes, BwDemod}. Their usage is documented in the command-line help. The default options are presented in Table 1.

Currently iProver uses non-perfect discrimination trees for implementing backward and forward demodulation [9, 17], feature vector indices for subsumption [19], tries for subset subsumption [9], and MiniSat [7] for global subsumption.

Generally, when during immediate simplification a parent clause of a newly derived clause is made redundant, we can remove all the children of that clause from the immediate set (and thus avoid adding them to the passive queues), except for the ones which caused it to be redundant. Currently, we restrict this feature to the given clause rather than to all the parent clauses, therefore, this simplifies to checking whether the given clause is made redundant in S_{immed} , and if so abort the loop, add only the clauses that made it redundant to the passive, and remove the given clause.

2.3 AC reasoning

If a problem contains associativity and commutativity axioms,

$$f(x, f(y, z)) = f(f(x, y), z), \quad f(x, y) = f(y, x), \quad (6)$$

then f is said to be AC.

AC axioms are particularly problematic in theorem proving, because they are non-orientable, which means they can generate permutations of arguments of AC functions. This leads to combinatorial explosion in the number of clauses. In particular, they will combine with each other to produce an exponential number of instances.

Listing 1.1: iProver-Sup given-clause loop algorithm

```

 $S_{\text{input}} = \emptyset$ 
for  $c$  in input_clauses:
    simplify( $c$  wrt  $S_{\text{input}}$  via  $R_{\text{input}}$ )
    add( $c$  to  $S_{\text{input}}$ )
add_to_passive_queue( $S_{\text{input}}$ )

 $S_{\text{active}} = S_{\text{passive}} = \emptyset$ 
loop:
     $S_{\text{immed}} = \emptyset$ 
    given = pop_from_passive_queue()
    simplify(given wrt  $S_{\text{active}} \cup S_{\text{passive}}$  via  $R_{\text{active}}$ )
    add_to_active_set(given)
    add(given to  $S_{\text{active}}$ )
    add(given to  $S_{\text{immed}}$ )
    for  $c$  in generating inferences between given and active:
        simplify( $c$  wrt  $S_{\text{immed}}$  via  $R_{\text{immed}}$ )
        if given was eliminated in  $S_{\text{immed}}$  by clauses  $U$ :
            add( $U$  to  $S_{\text{passive}}$ )
            remove(given from  $S_{\text{active}}$ )
            continue
        simplify( $c$  wrt  $S_{\text{active}} \cup S_{\text{passive}}$  via  $R_{\text{passive}}$ )
        add( $c$  to  $S_{\text{immed}}$ )
    add_to_passive_queue( $S_{\text{immed}}$ )
    add( $S_{\text{immed}}$  to  $S_{\text{passive}}$ )

```

AC problems are ubiquitous and appear in a variety of domains [21]. Although theoretical developments behind AC reasoning have a long history, AC support in most theorem provers is limited due to implementation complexity and is mainly restricted to unit equality problems. We extended some of the techniques to be applicable to the general clausal case, see Theorem 1 below, and implemented them in iProver.

AC preprocessing During preprocessing we can transform the input problem into any equisatisfiable form. We can normalise AC terms, by e.g. collecting nested AC subterms into a flat list, sorting wrt. some total extension of the term ordering, and making them right-associative.

Deletion of joinable equations A *rewrite system* is a set of rules $l \rightarrow r$, such that, if $l \rightarrow r$, then for any substitution σ , $l\sigma \rightarrow r\sigma$, and for any term u , $u[l] \rightarrow u[l \mapsto r]$. By abuse of notation we can also use unorientable equalities $l \leftrightarrow r$, in which case they stand for the set of its orientable instances, $\{l\sigma \rightarrow r\sigma \mid l\sigma \succ r\sigma\} \cup \{r\sigma \rightarrow l\sigma \mid r\sigma \succ l\sigma\}$. Two terms s and t are *joinable* wrt. a rewrite system R (written $s \downarrow_R t$) if $s \xrightarrow{*} c \xleftarrow{*} t$, where ‘ \rightarrow ’ denotes a rewrite step with a rule in R and ‘ $\xrightarrow{*}$ ’ its reflexive-transitive closure. Two terms are *ground joinable* ($s \Downarrow_R t$) if all its ground instances are joinable. Two terms are

strongly ground joinable (written $s \Downarrow_R t$) if, for all $s' = s\sigma, t' = t\sigma$ ground instances of s, t resp., with $s' \succeq t'$, either s' is t' or else $s' \xrightarrow{l \rightarrow r \in R} u' \Downarrow_R t'$ where either $l \prec s'$ or l is s' but not $u' \succ t'$ (see [1, 14]).

Theorem 1. *If $s \Downarrow_R t$, then $s \approx t \vee C$ is redundant wrt. R . If $s \Downarrow_R t$ then $s \not\approx t \vee C$ is redundant wrt. $R \cup \{C\}$.*

Theorem 1 was shown in the context of unit equality reasoning in [1]; we extended this theorem to general clauses and provided a different proof [6] (see Appendix).

This abstract theorem can be used for AC reasoning, provided we have a criterion to test $l \Downarrow r$. We use the following criterion [1]. Let R_{AC} be

$$f(x, y) \leftrightarrow f(y, x), \quad (7a)$$

$$f(f(x, y), z) \rightarrow f(x, f(y, z)), \quad (7b)$$

$$f(x, f(y, z)) \leftrightarrow f(y, f(x, z)), \quad (7c)$$

Unless $l \approx r$ is an instance of R_{AC} or can be simplified by an equation in R_{AC} , $l =_{AC} r$ implies $l \Downarrow_{R_{AC}} r$, which means we can use Theorem 1 to simplify/delete clauses wrt. R_{AC} . This is a cheap test for strong ground joinability to apply in practice, since in order to check whether $s =_{AC} t$ we can simply treat nested applications of f as a flat-list, and then sort wrt. some total order on terms (see above discussion on AC normalisation).

Semantic detection of axioms Some problems are AC even though the input does not contain the axioms explicitly. We say that a problem S is AC if $S \models AC$. The usual syntactic detection checks if $AC \in S$. But we wish also to detect AC problems even when this is not the case.

During preprocessing, we query an SMT solver to find out whether $S \models AC$. Since SMT solvers only accept ground problems, we need to use a sound approximation of the entailment relation. We do this using an injective substitution mapping variables to fresh constants similar as it is done for global subsumption [10]. This is a sound approximation, since $\phi(\bar{c}) \models \psi(\bar{c}) \Rightarrow \forall x \phi(\bar{x}) \models \forall x \psi(\bar{x})$. In order to make SMT reasoning more efficient we can further restrict reasoning to fast rules like unit propagation or place a limit on the number of backtracks. Apart from this, we also check if the AC axioms (6) get produced at some point during saturation, among binary symbols of sort $\alpha \times \alpha \rightarrow \alpha$.

3 Implementation and experimental results

We integrated the simultaneous superposition calculus, with the iProver-Sup saturation loop, into iProver and evaluated it over 15 168 first-order problems in TPTP-v7.2.0. The superposition loop can solve 7375 (49%), the instantiation loop (on the previous version of iProver) can solve 7884 (52%), and their combination can solve 8708 (57%). We can see that the combination with superposition and the iProver-Sup simplification setup improved the performance of iProver over the whole TPTP library.

Among problems that were solved by superposition, (excluding trivial problems solved by preprocessing), immediate simplification was used in 71.7% of problems and light normalisation was used in 64.5% of problems. AC axioms were detected in 1903 problems.

References

- [1] J. Avenhaus, T. Hillenbrand and B. Löchner. “On using ground joinable equations in equational theorem proving”. In: *J. Symb. Comput.* 36.1-2 (2003), pp. 217–233. DOI: [10.1016/S0747-7171\(03\)00024-5](https://doi.org/10.1016/S0747-7171(03)00024-5) (cit. on p. 8).
- [2] L. Bachmair and H. Ganzinger. “Resolution Theorem Proving”. In: *Handbook of Automated Reasoning (in 2 volumes)*. 2001, pp. 19–99. DOI: [10.1016/b978-044450813-3/50004-7](https://doi.org/10.1016/b978-044450813-3/50004-7) (cit. on p. 11).
- [3] L. Bachmair and H. Ganzinger. “Rewrite-Based Equational Theorem Proving with Selection and Simplification”. In: *J. Log. Comput.* 4.3 (1994), pp. 217–247. DOI: [10.1093/logcom/4.3.217](https://doi.org/10.1093/logcom/4.3.217) (cit. on p. 3).
- [4] D. Benanav. “Simultaneous Paramodulation”. In: *10th International Conference on Automated Deduction, Kaiserslautern, FRG, July 24-27, 1990, Proceedings*. 1990, pp. 442–455 (cit. on p. 3).
- [5] J. Denzinger, M. Kronenburg and S. Schulz. “DISCOUNT — A Distributed and Learning Equational Prover”. In: *Journal of Automated Reasoning* 18.2 (Apr. 1997), pp. 189–198. DOI: [10.1023/A:1005879229581](https://doi.org/10.1023/A:1005879229581) (cit. on p. 5).
- [6] A. Duarte and K. Korovin. “AC Reasoning Revisited”. to appear. 2020 (cit. on p. 8).
- [7] N. Eén and N. Sörensson. “An Extensible SAT-solver”. In: *Theory and Applications of Satisfiability Testing, 6th International Conference, SAT 2003. Santa Margherita Ligure, Italy, May 5-8, 2003 Selected Revised Papers*. 2003, pp. 502–518. DOI: [10.1007/978-3-540-24605-3_37](https://doi.org/10.1007/978-3-540-24605-3_37) (cit. on p. 6).
- [8] H. Ganzinger and K. Korovin. “New Directions in Instantiation-Based Theorem Proving”. In: *Proc. 18th IEEE Symposium on Logic in Computer Science, (LICS'03)*. IEEE Computer Society Press, 2003, pp. 55–64 (cit. on p. 1).
- [9] P. Graf. *Term indexing*. Lecture notes in computer science ; Lecture notes in artificial intelligence 1053. Berlin ; New York: Springer, 1995. 284 pp. ISBN: 978-3-540-61040-3 (cit. on pp. 3, 6).
- [10] K. Korovin. “Inst-Gen — A Modular Approach to Instantiation-Based Automated Reasoning”. In: *Programming Logics*. Ed. by A. Voronkov and C. Weidenbach. Vol. 7797. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 239–270 (cit. on pp. 1, 2, 4, 8).
- [11] K. Korovin. “iProver — An Instantiation-Based Theorem Prover for First-Order Logic (System Description)”. In: *Proceedings of the 4th International Joint Conference on Automated Reasoning, (IJCAR 2008)*. Ed. by A. Armando, P. Baumgartner and G. Dowek. Vol. 5195. Lecture Notes in Computer Science. Springer, 2008, pp. 292–298 (cit. on p. 1).

- [12] L. Kovács and A. Voronkov. “First-Order Theorem Proving and Vampire”. In: *Proceedings of the 25th International Conference on Computer Aided Verification*. Ed. by N. Sharygina and H. Veith. Lecture Notes in Artificial Intelligence 8044. Springer-Verlag, 2013, pp. 1–35 (cit. on pp. 3, 4).
- [13] J. C. López-Hernández and K. Korovin. “An Abstraction-Refinement Framework for Reasoning with Large Theories”. In: *Automated Reasoning - 9th International Joint Conference, IJCAR 2018, Proceedings*. Ed. by D. Galmiche, S. Schulz and R. Sebastiani. Vol. 10900. Lecture Notes in Computer Science. Springer, 2018, pp. 663–679. DOI: [10.1007/978-3-319-94205-6_43](https://doi.org/10.1007/978-3-319-94205-6_43) (cit. on p. 1).
- [14] U. Martin and T. Nipkow. “Ordered Rewriting and Confluence”. In: *10th International Conference on Automated Deduction, Kaiserslautern, FRG, July 24-27, 1990, Proceedings*. Ed. by M. E. Stickel. Vol. 449. Lecture Notes in Computer Science. Springer, 1990, pp. 366–380. DOI: [10.1007/3-540-52885-7_100](https://doi.org/10.1007/3-540-52885-7_100) (cit. on p. 8).
- [15] W. McCune. “OTTER 3.3 Reference Manual”. In: *CoRR* cs.SC/0310056 (2003). arXiv: [cs.SC/0310056](https://arxiv.org/abs/cs.SC/0310056) (cit. on p. 5).
- [16] R. Nieuwenhuis and A. Rubio. “Paramodulation-Based Theorem Proving”. In: *Handbook of Automated Reasoning (in 2 volumes)*. Ed. by J. A. Robinson and A. Voronkov. Elsevier and MIT Press, 2001, pp. 371–443. ISBN: 0-444-50813-9 (cit. on pp. 2, 3).
- [17] J. A. Robinson and A. Voronkov, eds. *Handbook of Automated Reasoning (in 2 volumes)*. Elsevier and MIT Press, 2001. ISBN: 0-444-50813-9 (cit. on pp. 3, 6).
- [18] S. Schulz. “E – A Brainiac Theorem Prover”. In: *Journal of AI Communications* 15.2/3 (2002), pp. 111–126 (cit. on pp. 4, 5).
- [19] S. Schulz. “Simple and Efficient Clause Subsumption with Feature Vector Indexing”. In: *Automated Reasoning and Mathematics - Essays in Memory of William W. McCune*. Ed. by M. P. Bonacina and M. E. Stickel. Vol. 7788. Lecture Notes in Computer Science. Springer, 2013, pp. 45–67. DOI: [10.1007/978-3-642-36675-8_3](https://doi.org/10.1007/978-3-642-36675-8_3) (cit. on p. 6).
- [20] S. Schulz. “System Description: E 1.8”. In: *Proc. of the 19th LPAR, Stellenbosch*. Ed. by K. McMillan, A. Middeldorp and A. Voronkov. Vol. 8312. LNCS. Springer, 2013 (cit. on p. 3).
- [21] G. Sutcliffe. “The TPTP Problem Library and Associated Infrastructure. From CNF to TH0, TPTP v6.4.0”. In: *Journal of Automated Reasoning* 59.4 (2017), pp. 483–502 (cit. on pp. 2, 7).
- [22] U. Waldmann, S. Turret, S. Robillard and J. Blanchette. “A Comprehensive Framework for Saturation Theorem Proving”. In: *IJCAR 2020*. Ed. by N. Peltier and V. Sofronie-Stokkermans. Vol. this volume. Lecture Notes in Computer Science. Springer, 2020 (cit. on p. 1).
- [23] C. Weidenbach, R. A. Schmidt, T. Hillenbrand, R. Rusev and D. Topic. “System Description: SpassVersion 3.0”. In: *Automated Deduction - CADE-21, 21st International Conference on Automated Deduction, Bremen, Germany, July 17-20, 2007, Proceedings*. Ed. by F. Pfenning. Vol. 4603. Lecture Notes in Computer Science. Springer, 2007, pp. 514–520 (cit. on pp. 3, 4).

Appendix A Proof of Theorem 1

Preliminaries

Consider a function signature Σ and a countable set of variables V . We assume the signature Σ contains an infinite number of fresh constant symbols. The set of all terms over a set of variables X , will be denoted $\mathcal{T}(\Sigma, X)$, the set of all ground terms $\mathcal{T}(\Sigma)$. An expression is either a term, a literal, a clause, or a list of thereof.

A *substitution* is a mapping from variables to terms which is the identity on all but finitely many variables. A *ground substitution* is a substitution whose codomain are ground terms. A substitution σ is *grounding* for an expression e if $e\sigma$ is ground. The set of all grounding substitutions for e will be denoted by $\text{GSubs}(e)$. The set of ground instances of an expression e is defined as $\text{GInsts}(e) = \{e\sigma \mid \sigma \in \text{GSubs}(e)\}$.

We write $s[t]$ if t is a *subterm* of s . We write $s[t \mapsto t']$ to denote the term s with that distinguished occurrence of t replaced by t' .

A relation ' \rightarrow ' over terms is a *rewrite relation* if (i) $t \rightarrow s \Rightarrow t\sigma \rightarrow s\sigma$ and (ii) $t \rightarrow s \Rightarrow u[t] \rightarrow u[t \mapsto s]$. A *transitive and reflexive closure* of \rightarrow will be denoted by \rightarrow^* .

Two terms are *joinable* ($s \downarrow t$) if $s \xrightarrow{*} u \xleftarrow{*} t$. They are *ground joinable* ($s \Downarrow t$) if for all $\sigma \in \text{GSubs}(s, t)$ we have $s\sigma \downarrow t\sigma$. Obviously joinability implies ground joinability.

Two literals are *strongly ground joinable* (written $s \Downarrow_{\triangleright} t$) if for all ground instances s', t' of s, t respectively, either $s' = t'$, or else (letting s' denote wlog. the greater of s', t'), $s' \xrightarrow{l \approx r} u' \downarrow t'$ with either: $l \prec s'$, or $l = s'$ but not $u' \succ t'$.

A *simplification ordering* is a rewrite relation which is a well-founded and total ordering on ground terms. Let \succ be a simplification ordering on terms. We assume that \succ is extended to an ordering on literals and clauses using the multi-set extensions of \succ . In particular we represent a literal $l \approx r$ as $\{l, r\}$, $l \not\approx r$ as $\{l, l, r, r\}$ and a clause as a multi-set of its literals.

Redundancy

The standard redundancy criterion is defined as follows [2]. A ground clause C is *redundant* with respect to a set of ground clauses S if there exists a subset of clauses S' of S such that

- (i) $S' \models C$, and
- (ii) $\forall D \in S'. D \prec C$.

This notion can be extended to non-ground clauses as follows. A possibly non-ground clause C is redundant wrt. a set of clauses S if every $C' \in \text{GInsts}(C)$ is redundant wrt. to $\text{GInsts}(S)$.

Proof of Theorem 1

We prove the [Theorem 1](#) using the standard notion of redundancy. We say that a clause C is redundant wrt R if it is either subsumed by an equation in R or C follows from smaller equations in R .

Theorem A.1. *If $s \Downarrow_{\triangleright R} t$, then $s \approx t \vee C$ is redundant wrt. R . If $s \Downarrow_R t$, then $s \not\approx t \vee C$ is redundant wrt. $R \cup \{C\}$.*

Proof. Assume that $s \Downarrow_{\triangleright R} t$, let us show that $s \approx t \vee C$ is redundant. For this it is sufficient to show that all clauses in $\text{GInsts}(s \approx t \vee C)$ follow from smaller rules in $\text{GInsts}(R)$. Consider the ground instance $s' \approx t' \vee C'$ in $\text{GInsts}(s \approx t \vee C)$. If $s' = t'$ then this instance is a tautology and follows from $\emptyset \subseteq R$, therefore redundant. Now assume that $s' \succ t'$ (the case when $t' \succ s'$ is similar). Since $s \Downarrow_{\triangleright R} t$ we have that $s' \xrightarrow{l \rightarrow r \in R} u' \Downarrow_R t'$ where either i) $l \prec s'$ or ii) l is s' but not $u' \succ t'$. Let R' be the ground instances of rules in R used in the joinability proof. It is clear that $R' \models s' \approx t'$ and therefore $R' \models s' \approx t' \vee C'$. Let us show that all rules in R' are smaller in \succ than $s' \approx t'$. In the case i) we have $l \prec s'$ therefore $l \approx r \prec s' \approx t'$. Likewise $t', u' \prec s'$ and therefore all rules used in rewrites of u' and t' have left-hand-side smaller than s' and therefore are smaller than $s' \approx t'$. The case ii) is similar: we have $s' = l$ and $u' = r$ and either $u' = t'$, so $l \approx r = s' \approx t'$ and in this case the clause is subsumed by $l \approx r$, or $u' \prec t'$ and in this case all rewrite rules used to reduce u' and t' are smaller than $s' \approx t'$.

To prove the second part of the theorem we can use the first part to note that for any ground instance $s' \not\approx t' \vee C'$ of $s \not\approx t \vee C$, $s' \approx t'$ follows from a smaller ground instances of rules from R , denote them by R' . Therefore $s' \not\approx t' \vee C'$ follows from $R' \cup \{C'\}$. As $C \prec s' \not\approx t' \vee C'$, the statement is proved.