

Comparing Instance Generation Methods for Automated Reasoning *

Swen Jacobs (sjacobs@mpi-inf.mpg.de)
and Uwe Waldmann (uwe@mpi-inf.mpg.de)
Max-Planck-Institut für Informatik, Saarbrücken, Germany

Abstract. The clause linking technique of Lee and Plaisted proves the unsatisfiability of a set of first-order clauses by generating a sufficiently large set of instances of these clauses that can be shown to be propositionally unsatisfiable. In recent years, this approach has been refined in several directions, leading to both tableau-based methods, such as the *Disconnection Tableau Calculus*, and saturation-based methods, such as *Primal Partial Instantiation* and *Resolution-Based Instance Generation*. We investigate the relationship between these calculi and answer the question to what extent refutation or consistency proofs in one calculus can be simulated in another one.

Keywords: Automated Reasoning, Theorem Proving, Disconnection Calculus, Primal Partial Instantiation, Resolution-Based Instance Generation

1. Introduction

In recent years, there has been a renewed interest in instantiation-based theorem proving for first-order logic. Much of the recent work in this field is based on the research of Plaisted and Lee [8]. They showed that the interleaving of production of instances with recombination of clauses, as done in resolution, leads to duplication of work in subsequent inference steps. As a means of avoiding this duplication, they proposed the *clause linking* approach. In clause linking, *links* between complementary unifiable literals are used to generate instances of the given clauses, based on the unifier of the linked literals. Unlike in the resolution calculus, the generated instances are not recombined, but added to the set of clauses as they are. As a consequence, in order to check satisfiability of the generated set, an additional SAT solving procedure is needed, which is usually in the spirit of the Davis-Putnam-Logemann-Loveland (DPLL) procedure [4].

Today, there are several methods which are based to some extent on clause linking and/or DPLL. They may be distinguished by the means

* This work was partly supported by the German Research Council (DFG) as part of the Transregional Collaborative Research Center “Automatic Verification and Analysis of Complex Systems” (SFB/TR 14 AVACS). See www.avacs.org for more information.



they use to detect unsatisfiability. Some calculi arrange instances in a tree or a tableau, integrating an implicit satisfiability check. This approach is used by the disconnection tableau calculus [3, 9, 12], as well as by Baumgartner’s FDPLL [1] and the model evolution calculus [2]. Other procedures separate instance generation and satisfiability test. These usually use a monotonically growing set of clause instances and call an external SAT solver on this set. Representatives of this approach are Plaisted’s (ordered semantic) hyper linking [8, 11], the partial instantiation methods [6] of Hooker et al., as well as resolution-based instance generation [5] by Ganzinger and Korovin. As they use interchangeable SAT solvers, the main difference between these methods lies in the guidance of instance generation.

The fact that tableau-based and saturation-based instance generation methods are somehow related can be considered as folklore knowledge; it has been mentioned repeatedly in the literature (e.g., [2, 6]). The precise relationship is, however, rather unclear. In this work, we will compare four instance generation methods that stay relatively close to the original clause linking approach, and can be seen as direct refinements of it:

The *Disconnection Calculus* [9, 12] (DCC) integrates the instance generation of the clause linking approach into a clausal tableau. The linking rule of DCC only allows inferences between literals which are on the same branch, strongly restricting the generation of clauses in the tableau. In this tableau structure, unsatisfiability of the given set of clauses can be decided by branch closure. Given a fair inference strategy, the calculus is refutationally complete.

Resolution-Based Instance Generation [5] (Inst-Gen) is a saturation-based method with a single resolution-like inference rule, representing the clause linking approach of generating instances. After a given set has been saturated under this inference rule, satisfiability of the original clause set is equivalent to propositional satisfiability and can be checked by any propositional decision procedure.

SInst-Gen is a refinement of Inst-Gen by semantic selection, based on a propositional model for the given set of clauses. Inconsistencies, which arise when extending this model to a model for the first-order clauses, are used to guide the inferences.

The *Primal Partial Instantiation* [6] method (PPI) is another saturation-based method with semantic selection similar to SInst-Gen.

After introducing the four calculi, we compare the saturation-based methods and show that SInst-Gen and PPI are to large parts equivalent. Then we compare SInst-Gen and PPI to DCC and we show to what extent refutation or consistency proofs in one calculus can be simulated in another one.

2. Introducing the Calculi

This section gives a short introduction to the methods we will compare. For a comprehensive description, we refer to Letz and Stenz [9, 12] for DCC, Ganzinger and Korovin [5] for (S)Inst-Gen, and Hooker et al. [6] for PPI.

2.1. LOGICAL PREREQUISITES

We use the usual symbols, notation and terminology of first-order logic with standard definitions. We consider all formulas to be in clausal normal form. This allows us to consider a formula as a set of clauses, thought to be connected by conjunction. In all of our methods, all variables in a clause are implicitly universally quantified in this clause. Thus, all clauses are considered to be variable-disjoint.

If F is some formula (clause, literal) and σ a substitution, then $F\sigma$ is an *instance* or *instantiation* of F . It is a *ground instance* if it is variable-free, otherwise it is a *partial instantiation*. It is a *proper instance*¹ if at least one variable is replaced by a non-variable term. A *variable renaming* is an injective substitution mapping variables to variables. Two formulas F and G are *variants* of each other if there is a variable renaming σ such that $F\sigma = G$. If F' is an instantiation of a formula F , then F is a *generalization* of F' . Given a set of formulas S , we say that $F \in S$ is a *most specific generalization* of F' with respect to S if F generalizes F' and there is no other formula $G \in S$ such that G generalizes F' , F generalizes G , and F is not a variant of G .

For any literal L , \bar{L} denotes its complement. By \perp we denote both a distinguished constant and the substitution mapping all variables to this constant. We say that two literals L, L' are \perp -*complementary* if $L\perp = \bar{L}'\perp$. A set of clauses S is called *propositionally unsatisfiable* if and only if $S\perp$ is unsatisfiable.

A clause is called *n-bounded* if all terms in the clause have depth less or equal to n . A set of clauses S is *n-satisfiable* if there is a Herbrand interpretation that makes all n -bounded ground instances of clauses in S true. By Herbrand's theorem, n -satisfiability for all n implies satisfiability.

2.2. THE DISCONNECTION TABLEAU CALCULUS

The *Disconnection Tableau Calculus* calculus has been developed by Billon [3] and Letz and Stenz [9, 12]. In order to define development

¹ This definition is due to Ganzinger and Korovin [5] and may deviate from definitions in other areas.

of a disconnection tableau, we need the notions of links, paths and tableaux:

A *literal occurrence* is a pair $\langle C, L \rangle$, where C is a clause and $L \in C$ a literal. If $\langle C, L \rangle$ and $\langle D, \bar{K} \rangle$ are two literal occurrences such that there is a most general unifier (mgu) σ of L and K , then the set $l = \{\langle C, L \rangle, \langle D, \bar{K} \rangle\}$ is called a *link* (between C and D). $C\sigma$ and $D\sigma$ are *linking instances* of C and D with respect to l .

A *path* P through a set of clauses S is a set of literal occurrences such that P contains exactly one literal occurrence $\langle C, L \rangle$ for every $C \in S$. A path P is \perp -*complementary* if it contains literal occurrences $\langle C, L \rangle$ and $\langle D, \bar{K} \rangle$ such that $L\perp = K\perp$, otherwise it is *open*.

A *disconnection tableau* (tableau, for short) is a possibly infinite downward tree with literal labels at all nodes except the root. Given a set of clauses S , a *tableau for S* is a tableau in which, for every tableau node N , the set of literals $C = L_1, \dots, L_m$ at the immediate successor nodes N_1, \dots, N_m of N is an instance of a clause in S . Every N_i is associated with the clause C and the literal occurrence $\langle C, L_i \rangle$.

Construction of a tableau starts from an *initial path* P_S through the set S of input clauses. The initial path may be chosen arbitrarily, but remains fixed through the construction of the tableau.

A *branch* of a tableau T is any maximal sequence $B = N_1, N_2, \dots$ of nodes in T such that N_1 is an immediate successor of the root node and any N_{i+1} is an immediate successor of N_i . With every branch B we associate a path P_B containing the literal occurrences associated with the nodes in B . The union $P_S \cup P_B$ of the initial path and the path of a branch B is called a *tableau path* of B .

To develop a tableau from the initial path and the empty tableau consisting of only the root node, we define the *linking rule*: Given an initial path P_S and a tableau branch B with literal occurrences $\langle C, L \rangle$ and $\langle D, \bar{K} \rangle$ in $P_S \cup P_B$, such that $l = \{\langle C, L \rangle, \langle D, \bar{K} \rangle\}$ is a link with mgu σ , the branch B is expanded with a linking instance with respect to l of one of the two clauses, say with $C\sigma$, and then, below the node labeled with $L\sigma$, the branch is expanded with the other linking instance $D\sigma$.

Example 1. (Linking Step) Figure 1 illustrates a tableau consisting of an initial path and the nodes produced by one linking step. Input clauses are inside the box, literals which are connected by lines inside this box give the initial path. The dashed lines marked with numbers denote links between selected literals. The “1” below the box means that link 1 is used and the clauses below are linking instances with respect to this link.

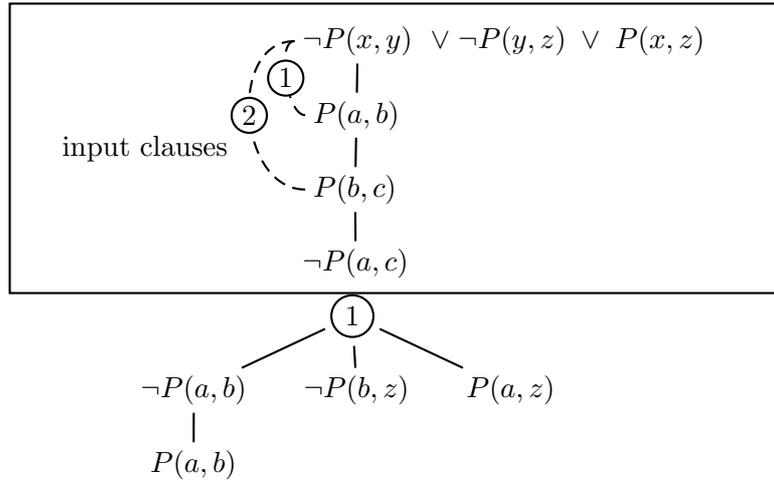


Figure 1. A single linking step

As all clauses are variable-disjoint, *variant-freeness* is required in order to restrict proof development to inferences which introduce “new” instances: A disconnection tableau T is *variant-free*, if no node N with clause C in T has an ancestor node N' with clause D in T such that C and D are variants of each other. In practice, variant-freeness is assured by two restrictions: First, a link that has already been used on the current branch may not be used again. Secondly, when a linking step is performed, variants of clauses which are already on the branch are not added to the tableau. This can result in linking steps where only one of the linking instances is added. By definition, variant-freeness does not extend to the initial path, i.e., variants of input clauses can be added to the tableau, and may be needed for completeness. This means that tableau paths can contain two variants of the same clause. In this case, the original clause on the initial path is redundant, i.e., need not be considered for linking steps on this branch, as we have shown in Jacobs [7].

Next, we define when tableau construction will terminate: A tableau branch B is *closed*, if P_B is \perp -complementary; if not, it is called *open*. A tableau is *closed* if it has no open branches. Note that this definition does not include the initial path P_S for branch closure. This means that \perp -complementary literals $L \in P_S$ and $K \in P_B$ do not close a branch. However, the link between L and K can be used to add a non-proper instance of the clause C containing L to the tableau. This results in a number of new branches, one of which is directly closed (the one which selects the instance of L). As a consequence, we can allow unit clauses

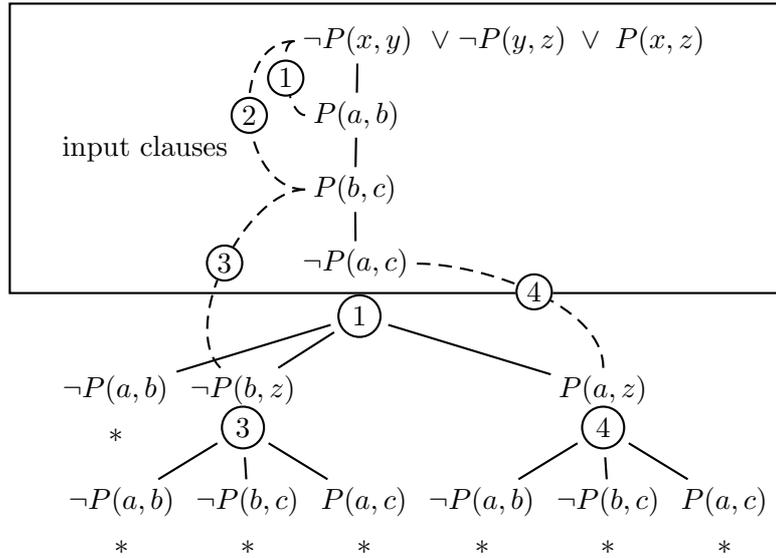


Figure 2. A closed disconnection tableau

on the initial path to close branches directly, as adding them to the tableau would result in only one new branch, which would be closed.

Example 2. (Closed Tableau) Figure 2 shows a closed disconnection tableau. All branches are closed by unit clauses on the initial path which have not been added to the tableau again. Closed branches are marked by a *.

A branch B in a (possibly infinite) tableau T is called *saturated* if B is open and there is no link on B that produces at least one linking instance that is not a variant of any clause on B . A tableau is *saturated* if either all its branches are closed or if it has a saturated branch. The propositional model represented by a saturated branch can be extended to a Herbrand model of the input clauses by assigning to every ground atom the truth value of its most specific generalization on the branch (see Letz and Stenz [9] for details).

For completeness of the calculus, we need a fairness condition. A derivation of DCC is *fair* if every infinite branch is saturated, i.e., if every linking step will be taken at some time or will become redundant because variants of its linking instances are produced by other linking steps.

2.3. RESOLUTION-BASED INSTANCE GENERATION

The *Inst-Gen* calculus is a saturation-based calculus due to Ganzinger and Korovin [5]. To extend a given set of clauses, it uses the following inference rule:

$$\text{Inst-Gen} \quad \frac{C \vee L \quad D \vee \overline{K}}{(C \vee L)\sigma \quad (D \vee \overline{K})\sigma}$$

where σ is the mgu of K and L and at least one of the literals $\overline{K}\sigma$ and $L\sigma$ is a proper instance of K or L , respectively.

For a set of clauses saturated under Inst-Gen, the satisfiability test can be reduced to the propositional case. As saturation may take infinitely long, however, satisfiability testing cannot be postponed until saturation is reached.

There is a formal notion of redundancy for Inst-Gen, which is, however, beyond the scope of this work. The only clauses we will consider as *redundant* are variants of clauses which are already present. An inference is redundant if it only produces such variants.

SInst-Gen is an refinement of Inst-Gen, which uses *semantic selection* in order to restrict the search space: Let S be a set of clauses such that $S \perp$ is satisfiable and let $I \perp$ be a model of $S \perp$. We define the *satisfiers* of a clause C to be the set $\text{sat} \perp(C) = \{L \in C \mid I \perp \models L \perp\}$. Now consider *selection functions* on clauses (modulo renaming), which select for every clause in S one of its satisfiers.

Instance generation based on a selection function sel is defined as follows:

$$\text{SInst-Gen} \quad \frac{C \vee L \quad D \vee \overline{K}}{(C \vee L)\sigma \quad (D \vee \overline{K})\sigma}$$

where σ is the mgu of K and L and both \overline{K} and L are selected by sel .²

A selection function represents a model for the grounded set of clauses $S \perp$. When trying to extend it to a model of S , complementary unifiable literals represent inconsistencies in the extension. Every inference of SInst-Gen resolves such an inconsistency.

If selection functions are allowed to select more than one satisfier per clause, *hyper-inferences* can be enabled. If the literals L_1, \dots, L_n

² The second condition of the Inst-Gen rule is implied by the use of semantic selection.

are selected in clause C and there are clauses C_1, \dots, C_n with selected literals K_1, \dots, K_n such that σ is the mgu of $L_1 = \overline{K_1}, \dots, L_n = \overline{K_n}$, then a hyper-inference produces $C\sigma$ and all of the $C_i\sigma$ in one step.

In order to ensure that unsatisfiable input sets are saturated within finite time, the choice of inferences must be *fair*. Intuitively, fairness means that any inference which is available infinitely often must either be taken or become redundant by application of other inferences at some time [5].

Propositional satisfiability of the generated set of clauses S' is tested after every inference step by searching for a model of $S' \perp$. If this does not exist, unsatisfiability of the input set S has been proved. If it does, the model is used to determine the new selection function sel for the next step. This process is repeated until either no model can be found or no inferences are possible. In the latter case, the model of $S' \perp$ can be extended to a model of S without conflicts, i.e., satisfiability of S has been shown.

2.4. PRIMAL PARTIAL INSTANTIATION

The *Primal Partial Instantiation (PPI)* method of Hooker et al. [6] is a saturation-based calculus that is very close to SInst-Gen. Its inference rule differs from the SInst-Gen rule only in the side condition. To define this side condition, Hooker uses a notion of variants that is slightly different from the one used in (S)Inst-Gen and DCC: a formula F is a *PPI-variant* of a formula G if there is a (not necessarily injective) variable renaming σ such that $F\sigma = G\sigma$. For instance, $P(x, x)$ is a PPI-variant (but not a variant) of $P(y, z)$.

As in SInst-Gen, one literal per clause is selected at any time. In PPI, a selection is *admissible* if no complementary PPI-variants are selected. If no admissible selection is possible for the current set, PPI has shown unsatisfiability. The calculus also terminates if the current set of clauses is saturated with respect to the current selection, i.e., if no inferences are possible according to the following rule:

$$PPI \quad \frac{C \vee L \quad D \vee \overline{K}}{(C \vee L)\sigma \quad (D \vee \overline{K})\sigma}$$

where σ is the mgu of K and L , both \overline{K} and L are selected and at least one of $(C \vee L)\sigma$ and $(D \vee \overline{K})\sigma$ is not a generalization of a clause in the current set.

The last condition can be seen as a redundancy elimination criterion, but in this form, it makes the calculus unsound. Termination in case of saturation should indicate satisfiability of the given set, but with this

rule, satisfiability may be indicated when it is not the case. An example for this behaviour has already been given in [5]. We give a different one that carries more information on how the condition must be corrected (underlined literals are selected):

$$P(x, y, z) \vee \underline{Q(z)}, \quad (1)$$

$$\underline{P(x, x, a)} \vee Q(a), \quad (2)$$

$$\underline{\neg Q(a)}, \quad (3)$$

$$\underline{\neg P(a, b, a)} \quad (4)$$

Inconsistency of this set of clauses is apparent: with $[a/x, b/y, a/z]$ we get an unsatisfiable set of ground instances. Among the selected literals, only $Q(z)$ from clause (1) and $\neg Q(a)$ from clause (3) are complementarily unifiable. The resulting clauses are $P(x, y, a) \vee Q(a)$ and $\neg Q(a)$. The latter is already present in the set, while the former is a generalization of the second clause in the given set. Thus, PPI does not allow this inference, nor any other, on this set of clauses (with this selection).

In order to make the calculus sound, the last condition of the inference rule has to be weakened. It is not enough to replace “generalization” by “PPI-variant”, as the example above shows unsoundness even for this case. If “generalization” is replaced by “variant”, the calculus can be shown to be sound.³ We will consider this modification from now on.

What is still missing in order to have a refutationally complete calculus is a saturation strategy that restricts the order of inferences. To explain the scheme that PPI uses we need some additional definitions:

An inference is said to be *n-literal-bounded* if all terms in the unified literals have depth less or equal to n . A set of clauses is *n-literal-saturated* with respect to an admissible selection of literals if with this selection no n -literal-bounded inferences are possible.

Then, n -literal-saturation with respect to an admissible selection of literals implies n -satisfiability [6]. As n -satisfiability for all n implies satisfiability, we have a complete calculus if infinite derivations achieve n -literal-saturation for every n at some time.

The strategy of PPI reaches this goal in the following way: Starting with $n = 0$, the method only uses n -literal-bounded inferences until the set is n -literal-saturated for some literal selection. Only then, the depth limit is increased to $n + 1$, until $(n + 1)$ -literal-saturation is reached, and so on.

³ The original soundness proof from [6] works for the corrected inference rule, see Jacobs [7].

3. Comparing the Saturation-Based Methods

In this section, we want to compare Inst-Gen, SInst-Gen and PPI. Obviously, Inst-Gen is a proper generalization of both SInst-Gen and PPI. In the following, we show that PPI and SInst-Gen (without hyper-inferences) are equivalent up to the saturation strategy. After that, we compare the saturation strategies and give a notion of fairness which generalizes those of both calculi, while preserving completeness.

3.1. COMPARING SINST-GEN AND PPI

Considering selection of literals, one can see that two literals are complementary PPI-variants if and only if they are \perp -complementary, so selection is directly equivalent in both methods. It is also obvious that the revised inference rule of PPI is equivalent to the one of SInst-Gen coupled with redundancy of variants. Furthermore, equivalence of literal selection and inference rules implies equivalence of the respective termination requirements.

Thus, SInst-Gen and PPI are equivalent with respect to selection, inference rule and termination requirements. What remains to be compared is the saturation strategy, or the fairness criterion that restricts the possible strategies.

In SInst-Gen, a strategy must be fair, i.e., no inference may be enabled infinitely often without being used or becoming redundant. PPI seems to be more restricted, in that strategies must only use n -literal-bounded inferences until n -literal-saturation is reached, starting with $n = 0$.

It is obvious that there are strategies for SInst-Gen which are not allowed in PPI. For the other direction, one might assume that all of PPI's strategies are allowed in SInst-Gen. Example 3 shows however, that this is not the case.

Example 3. (Unfair run of PPI) The following is an infinite run of PPI on a satisfiable input set. We will indicate the growth of the term-depth counter n , which is in some sense responsible for this unfair behaviour of PPI. During this run, the selection function⁴ is changed frequently, in order to create the unfair behaviour. We start with $n = 0$ and the input set

$$\overline{\neg P(x)} \vee P(fx), \quad (1)$$

$$\overline{P(fa)}, \quad (2)$$

$$\overline{P(fb)} \vee P(c) \quad (3)$$

⁴ As we have shown PPI and SInst-Gen to be equivalent in large parts, we will use the notation of SInst-Gen also for PPI, where appropriate.

With this selection, no 0-literal-bounded inference is possible, so n is set to 1. Then, we resolve the inference between (1) and (3), and afterwards change selection on the third clause (which we will mark with (3') with this selection):

$$\neg \underline{P(x)} \vee P(fx), \quad (1)$$

$$\underline{P(fa)}, \quad (2)$$

$$P(fb) \vee \underline{P(c)} \quad (3')$$

$$\neg P(fb) \vee \underline{P(ffb)} \quad (4)$$

Between (1) and (3') there is an inference which is 0-literal-bounded. This inference will reappear infinitely often, without being taken.

In addition to this one, there are inferences between (1) and (2), as well as between (1) and (4). As $n = 1$, we are allowed to take the 1-literal-bounded inference between (1) and (2), but not the 2-literal bounded one between (1) and (4). Thus, we resolve (1) and (2) and afterwards change selection on (3) again:

$$\neg \underline{P(x)} \vee P(fx), \quad (1)$$

$$\underline{P(fa)}, \quad (2)$$

$$\underline{P(fb)} \vee P(c) \quad (3)$$

$$\neg P(fb) \vee \underline{P(ffb)} \quad (4)$$

$$\neg P(fa) \vee \underline{P(ffa)} \quad (5)$$

Between (1) and (3) there is no inference, leaving two 2-literal-bounded inferences between (1) and (4) as well as between (1) and (5).

From here, we can carry on by induction: if the term-depth counter is at n and the first literal is selected in (3), the 0-literal-bounded inference is not enabled. The two possible inferences are both $n + 1$ -literal-bounded. We increase n by one, take one of the $(n + 1)$ -literal-bounded inferences and afterwards enable the 0-literal-bounded inference by changing selection to (3') again. We do not have to resolve it, however, as there is another $(n + 1)$ -literal-bounded inference. After resolving this one, we again switch selection to (3), disabling the 0-literal-bounded inference. The last two inferences have supplied us with clauses that now again allow two inferences, this time of depth $n + 2$.

In this way, the inference between (1) and (3') is available infinitely often, but is never resolved.

3.2. A NEW FAIRNESS CRITERION

As we have seen that the fairness criteria of SInst-Gen and PPI are incompatible, the question arises whether there is a fairness condition that generalizes both those of PPI and SInst-Gen, and preserves completeness. To preserve completeness, it is enough to ensure that if the given formula is not n -satisfiable for some n , then the derivation will terminate. A way to do this is to require n -literal-saturation to be reached for all n , but relaxing the restriction of PPI that until this point only n -literal-bounded inferences may be used. However, the following example shows that even in this case, SInst-Gen's fairness criterion is not fully included.

Example 4. The following derivation is fair with respect to the SInst-Gen definition, but never reaches 0-literal-saturation. We begin with the set

$$Q(x) \vee \underline{\neg P(y)} \vee P(fy), \quad (1)$$

$$\underline{\neg Q(b)}, \quad (2)$$

$$\underline{Q(x)} \vee \neg P(b), \quad (3)$$

$$Q(x) \vee \underline{P(a)}, \quad (4)$$

where we have 0-literal-bounded inferences between (2) and (3) as well as between (1) and (4). We can develop this derivation in such a way that we always have a 0-literal-bounded inference between (2) and some other clause containing $Q(x)$, but none of these inferences is available infinitely often.

To this end, we resolve the inference between (4) and (1) and afterwards change selection on clauses (3) and (4). This gives us

$$Q(x) \vee \underline{\neg P(y)} \vee P(fy), \quad (1)$$

$$\underline{\neg Q(b)}, \quad (2)$$

$$Q(x) \vee \underline{\neg P(b)}, \quad (3')$$

$$\underline{Q(x)} \vee P(a), \quad (4')$$

$$Q(x) \vee \neg P(a) \vee \underline{P(fa)} \quad (5)$$

Now there is no inference between (2) and (3'), but we have another 0-literal-bounded inference between (2) and (4').

From here on, we can again continue by induction: say we have a set of m clauses so far, selection in (1) and (2) is as above, the last literal is selected in clauses (3') to $(m - 2')$, $Q(x)$ is selected in $(m - 1)$. The last clause is

$$Q(x) \vee \neg P(f^i a) \vee \underline{P(f^{i+1} a)} \quad (m)$$

Then a 0-literal-bounded inference between (2) and $(m - 1)$ is available, as well as an inference between (1) and (m) . We resolve the latter and change selection in $(m - 1)$ to the last literal and in (m) to $Q(x)$. With

$$Q(x) \vee \neg P(f^{i+1}a) \vee \underline{P(f^{i+2}a)}, \quad (m + 1)$$

we have the premise for another induction step.

In this way, we always have a 0-literal-bounded inference available, without violating the fairness condition of SInst-Gen.

This example indicates that we need to relax the restrictions of PPI even more. The requirement of n -literal-saturation for all n is sufficient to preserve completeness, but not necessary. Ganzinger and Korovin use a different notion:

A set of clauses is called *n -clause-saturated* with respect to a selection function sel if sel does not allow any inference which produces non-redundant n -bounded clauses. Although weaker than n -literal-saturation, n -clause-saturation still implies n -satisfiability [5].

Definition 1. An infinite derivation consisting of pairs of clause sets and selection functions $\langle S_0, \text{sel}_0 \rangle, \langle S_1, \text{sel}_1 \rangle, \dots$ is *fair* if for every n there is an i such that S_i is n -clause-saturated with respect to sel_i .

Fairness of a derivation might be enforced by a relaxation of the PPI strategy, allowing only $(n + k)$ -clause-bounded inferences before n -clause-saturation is reached. That means that we can allow a flexible search depth for every n .

LEMMA 1. *A derivation is fair if and only if for every n there is a k such that until n -clause-saturation is reached only $(n + k)$ -bounded clauses are produced.*

Proof. If a derivation is fair, then for every n it reaches an n -clause-saturated set in finitely many steps. Thus, there must be an m such that all clauses produced until then are m -bounded. Let $k = m - n$.

For the other direction, let S be a finite set of clauses. Then there are only finitely many $(n + k)$ -bounded instances of clauses in S . As inferences that produce $(n + k)$ -bounded instances also have $(n + k)$ -bounded instances as premises, there are only finitely many $(n + k)$ -clause-bounded inferences between instances of S . As only these will be used until n -clause-saturation is reached, we will reach n -clause-saturation within finite time.

As n -literal-saturation implies n -clause-saturation, it is easy to see that the new fairness condition allows any possible behaviour of PPI.

What we need to show is that this definition also generalizes the fairness condition of SInst-Gen. The result stated in the following theorem is essentially included in the completeness proof of SInst-Gen [5], but we give a short proof of the property we need:

THEOREM 1. *If every inference that appears infinitely often in an SInst-Gen derivation is eventually taken, then the derivation is fair.*

Proof. As in the proof of Lemma 1, we can assume that for a given n there are only finitely many inferences that produce n -bounded clauses. If the derivation never reaches n -clause-saturation, at least one of these inferences must be enabled infinitely often, contradicting the assumption that every persistent inference is taken.

4. Comparing SInst-Gen and DCC: Refutation Proofs

In this section, we compare the saturation-based methods PPI and SInst-Gen to the tableau-based DCC. Based on the equivalence up to saturation strategies that we have established, we use SInst-Gen as a representative of the saturation-based methods. Only in case of important differences between SInst-Gen and PPI we will mention the latter explicitly.

In the previous section, the high degree of structural similarity between SInst-Gen and PPI allowed us to reduce the difference between these calculi to a difference between saturation strategies. When we want to compare saturation-based and tableau-based methods, the situation becomes much more complicated. We are no longer in the comfortable situation that the inference rules in these calculi are “essentially the same”, and consequently, the answer to the question what kind of comparison is appropriate for these calculi is much less obvious. In principle, several patterns of comparison are available: Assuming that our objects of investigation are abstract calculi, rather than concrete implementations, we could analyze for which inputs proofs are found, how long the proofs are, or which formulas are derived during the proof. Which of these choices are adequate for a comparison of DCC and SInst-Gen?

Both calculi are refutationally complete, and so they find (refutation) proofs for exactly the same input sets. (They may differ, however, in their ability to prove the consistency of satisfiable input sets. We will discuss this case later.)

DCC checks propositional satisfiability internally, whereas SInst-Gen uses an external program to solve this (NP-complete) subprob-

lem. Since the number of inference steps of the external SAT solver is unknown, a meaningful comparison of the length of proofs is impossible.

The only remaining choice is to compare the internal structure of the proofs, or more precisely, the sets of clauses that are generated during the proof.

Definition 2. A proof of method A *simulates* a given proof of method B if the instances generated by the simulating A -proof are a subset of the instances generated in the original B -proof.

Thus, if method A can simulate all proofs of method B , B can be seen as a special case (or refinement) of A . With respect to this definition, Inst-Gen is the most general of the instance generation methods we have introduced, as it can simulate DCC, PPI and SInst-Gen proofs. That is, all of these calculi are refinements of Inst-Gen. A method that can simulate any refutation proof of another method is more general, but usually not better in the sense that it finds more proofs with limited resources (time, space). On the contrary, a strictly more general method will usually have a larger search space.

The definition of simulation is also motivated by the following two lemmas:

LEMMA 2. *Let S be an unsatisfiable set of clauses, let T be a closed disconnection tableau for S . Then the set S' of all instances on the tableau is propositionally unsatisfiable.*

Proof. Suppose S' was not propositionally unsatisfiable. Then there must be an open path through $S' \perp$. Using the literal occurrences in this path, we can identify a branch in T which selects the same literals. This branch cannot be closed, contradicting the assumption that T is closed.

The lemma implies that an SInst-Gen proof will terminate as soon as all instances from a closed disconnection tableau have been generated. Thus, we know that, if generation of these instances is possible, the requirements for our notion of simulation will be fulfilled by SInst-Gen. The next lemma does the same for the other direction of simulation:

LEMMA 3. *Let S be an unsatisfiable set of clauses. If an SInst-Gen proof terminates after generating the set of instances $S' \supseteq S$, then there exists a closed tableau containing only instances from S' .*

Proof. We may simply add instances from S' to each branch of the tableau until it is closed. As there are no open paths through S' , this means that every branch containing all clauses from S' must be closed.

Note, however, that neither Lemma 2 nor Lemma 3 guarantee that the required clauses can actually be generated using the construction rules of the simulating calculus. Moreover, in the disconnection calculus, even if the instances can be generated somewhere in the tableau, this does not necessarily mean that they can be generated where they are needed.

Thus, we need to compare instance generation and the guidance of proofs in both methods in order to see if simulation is possible. Instance generation itself, represented by the inference rules, is identical in both methods: the main premise is a pair of selected literals L_1, L_2 in clauses C_1, C_2 such that $L_1\sigma = \overline{L_2}\sigma$ for some mgu σ . If those are present, the instances $C_1\sigma$ and $C_2\sigma$ are generated. In both methods, one literal per clause is selected and variants of existing clauses will not be added again. Open branches correspond essentially to selection functions; if a branch is closed in DCC, then in SInst-Gen a selection function for the set of clauses is not allowed to select the same literals. The other direction, however, does not hold, because of the special role of the initial path in DCC: It is chosen and fixed at the beginning of the proof, and literals on it must not be used to close branches. There is no equivalent notion in SInst-Gen. However, adding variants of clauses from the initial path to the tableau effectively allows a different selection on and closure by input clauses, albeit only if there is a link which allows generation of the variant. The fact that literals on the initial path do not close branches may also lead to the generation of non-proper instances of input clauses. This cannot happen in SInst-Gen, as it requires two \perp -complementary literals to be selected. The main difference between the two approaches, however, is that instances generated in DCC will only be available on the current branch, while in SInst-Gen all instances are available at all times, i.e., regardless of the current selection.

4.1. FROM SINST-GEN TO DCC

THEOREM 2. *There exist refuting SInst-Gen proofs that cannot be simulated by any DCC proof.*⁵

Proof. The following is an SInst-Gen proof for an unsatisfiable set of clauses. We claim that it cannot be simulated by any disconnection tableau, i.e., DCC cannot finish the proof with the same set of instances as SInst-Gen, or a subset thereof. The reason for this is that the needed instances cannot be generated on all branches without generating additional instances.

⁵ It is not known whether this result still holds if one considers some of the extensions of DCC from Stenz [12].

The proof which we consider starts with the set of input clauses

$$\underline{\neg P(x, y)} \vee \neg P(y, z) \vee P(x, z), \quad (1)$$

$$\underline{\neg P(x, y)} \vee P(fx, c), \quad (2)$$

$$\underline{P(a, b)}, \quad (3)$$

$$\underline{P(b, c)}, \quad (4)$$

$$\underline{\neg P(fa, c)}, \quad (5)$$

where a selection function is given by the underlined literals. SInst-Gen goes on to produce the following instances in the given order:

$$\neg P(a, b) \vee \underline{\neg P(b, z)} \vee P(a, z), \quad (6)$$

$$\neg P(a, b) \vee \neg P(b, c) \vee \underline{P(a, c)}, \quad (7)$$

$$\neg P(a, c) \vee P(fa, c) \quad (8)$$

One can easily see that addition of (8) makes the set propositionally unsatisfiable. SInst-Gen terminates, indicating unsatisfiability. The DCC derivation in Figure 3 tries to reproduce all steps of this SInst-Gen proof.

The initial path of the tableau selects the same literals as the initial selection function of the SInst-Gen proof does. One can confirm that links number 1, 5 and 6 are equivalent to the inference steps in SInst-Gen, as they produce the same instances. We see that there is one open branch. Our definition of simulation would allow us to close this branch by generating again any of the instances from the given proof. The available links on this branch are those with numbers 2 to 4 from the initial path, as well as two additional links to $P(a, z)$ which are not displayed in the figure. However, none of those links generate one of the needed instances. Therefore, simulation has failed with this strategy.

It might still be possible that there is a tableau simulating the given SInst-Gen proof which generates instances in a different order or uses a different initial path. We have shown that even in this case simulation of the given derivation is not possible. The proof is rather lengthy and can be found in Jacobs [7].

As the SInst-Gen proof given above obeys the term-depth restriction of Hooker's PPI method, it shows also that there are PPI proofs that cannot be simulated by any DCC proof:

COROLLARY 1. *There exist refuting PPI proofs that cannot be simulated by any DCC proof.*

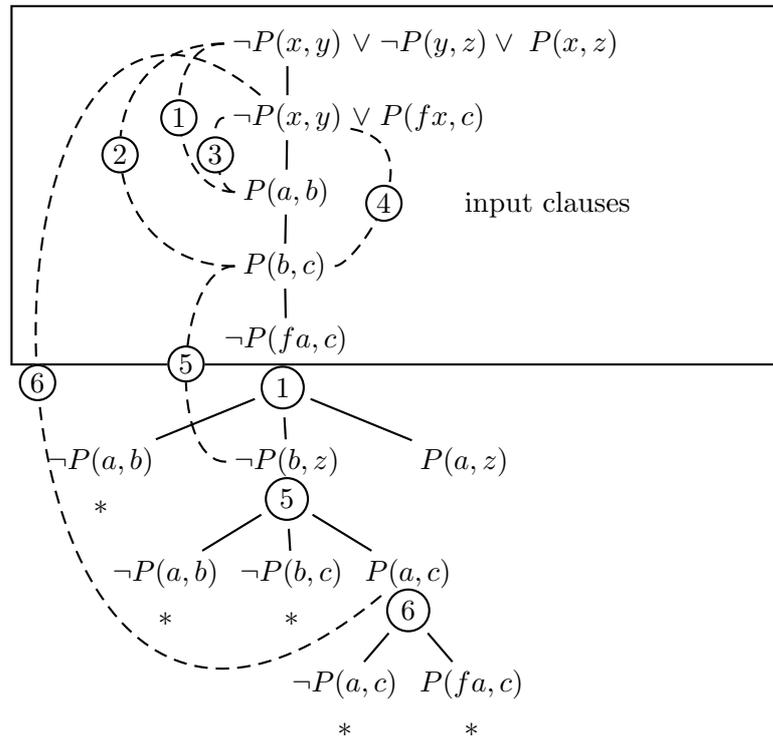


Figure 3. DCC needs additional instances

4.2. FROM DCC TO SINST-GEN

The fact that SInst-Gen proofs cannot always be simulated by DCC proofs is a consequence of the tree structure of DCC proofs. One might expect that in the other direction there is no such obstacle, but surprisingly, this is not the case.

THEOREM 3. *There exist refuting DCC proofs that cannot be simulated by any SInst-Gen proof.*

Proof. Figure 4 shows a closed disconnection tableau for an unsatisfiable set of input clauses. We claim that this tableau cannot be simulated by SInst-Gen.

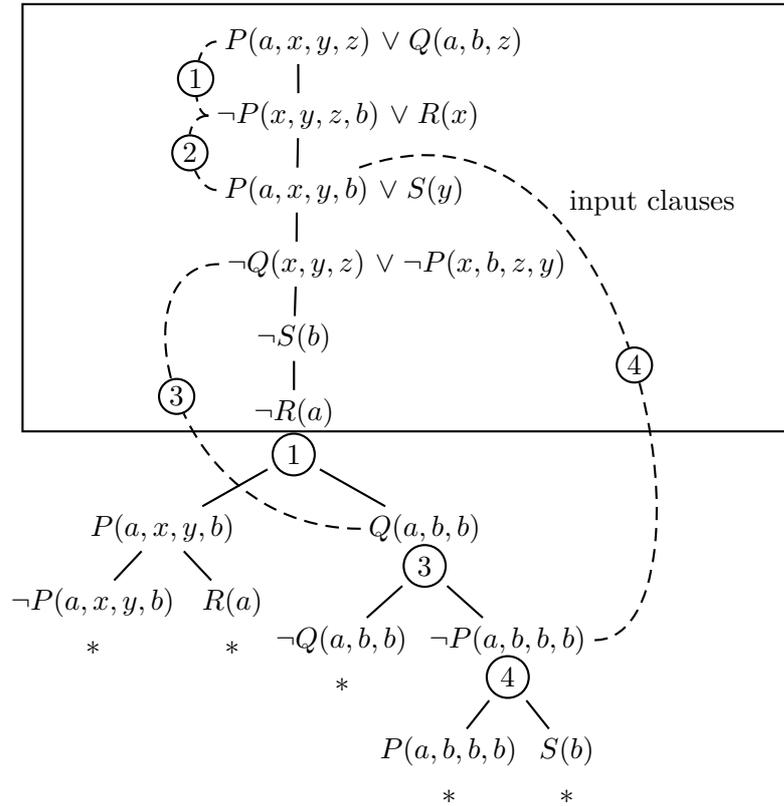


Figure 4. SInst-Gen needs additional instances

In order to verify this claim, let us consider all possible SInst-Gen proofs for the given set of clauses:

$$P(a, x, y, z) \vee Q(a, b, z), \quad (1)$$

$$\neg P(x, y, z, b) \vee R(x), \quad (2)$$

$$P(a, x, y, b) \vee S(y), \quad (3)$$

$$\neg Q(x, y, z) \vee \neg P(x, b, z, y), \quad (4)$$

$$\underline{\neg S(b)}, \quad (5)$$

$$\underline{\neg R(a)} \quad (6)$$

An underlined literal means that no other selection is possible in that clause. If there is no underlined literal, we consider all possible selections. There is an inconsistency between $P(a, x, y, z)$ in (1) and $\neg P(x, y, z, b)$ in (2), which is equivalent to link number 1 in the tableau (i.e., produces the same instances). Also, there are inconsistencies between $\neg P(x, y, z, b)$ in (2) and $P(a, x, y, b)$ in (3), and between $R(x)$

in (2) and the unit clause (6). Both of these generate only one of the instances produced by linking step 1, $\neg P(a, x, y, b) \vee R(a)$. There are four other possible inferences on this set of clauses, but all of them produce instances which are not on the given tableau.

Let us first consider a derivation which only produces

$$\underline{\neg P(a, x, y, b)} \vee R(a) \quad (7)$$

in the first step. Then, except the one mentioned first above, none of the possible inferences between the input clauses is usable for a simulation. Thus, we only need to consider inconsistencies referring to (7). Moreover, as $R(a)$ in (7) is complementary to the given unit clause $\neg R(a)$, we only need to consider inconsistencies between $\neg P(a, x, y, b)$ and literals of the given clauses. Only one new inconsistency is introduced by the new clause, related to $P(a, x, y, z)$ in (1). Furthermore, both of the inferences which are admissible for a simulating proof produce the same instance, which is

$$P(a, x, y, b) \vee \underline{Q(a, b, b)} \quad (8)$$

Thus, after either taking the first inference step mentioned above, or one of the other two possible inference steps and one of the two admissible subsequent inferences, our current set consists of the clauses (1)–(8). All admissible inferences between input clauses have been carried out and selection is fixed on all of the generated instances. Thus, we have only one possible inference, which is between $Q(a, b, b)$ in (8) and $\neg Q(x, y, z)$ in (4). This step is equivalent to linking step number 3 in the tableau, generating

$$\neg Q(a, b, b) \vee \underline{\neg P(a, b, b, b)} \quad (9)$$

We have to select $\neg P(a, b, b, b)$ in (9), which only gives us one new inference, connected to $P(a, x, y, z)$ in (1). Note that the inference equivalent to linking step number 4 is not possible, as $P(a, x, y, b)$ is \perp -complementary to $\neg P(a, x, y, b)$, which has to be selected in (7). The new inference generates $P(a, b, b, b) \vee Q(a, b, b)$, which is not on the given tableau. At this point, there is no inference which does not violate our simulation property, which means that simulation has failed.

This result also holds if hyper-inferences are allowed in SInst-Gen, as one can easily verify that the possible hyper-inferences either produce instances which violate simulation, or only produce the same instances as the standard inferences.

4.3. WEAK SIMULATION

We have shown that simulation of refutational SInst-Gen (or PPI) proofs by DCC, or vice versa, fails in general. We can get positive simulation results, however, if the definitions of the calculi are slightly changed and if the definition of simulation is modified in the following way:

Definition 3. A proof by method A *simulates* a given proof by method B *weakly*, if every instance generated by the simulating A -proof is a generalization of an instance generated in the original B -proof.

Recently, Letz and Stenz [10] introduced a more general version of the linking rule, which uses instead of standard unification a special form of a unifier: Given two literals L and L' , a substitution σ is called a \perp -unifier of L and L' if $L\sigma\perp = L'\sigma\perp$. σ is called a *most general* \perp -unifier of L and L' if it is more general than every \perp -unifier of L and L' . For instance, \perp -unification of $P(a, x, y)$ and $P(x', y', y')$ results in $P(a, x, y)$ and $P(a, y', y')$.

It has already been stated [5] that the degree of instantiation of an SInst-Gen inference can be chosen flexibly, as long as at least one variable is instantiated properly. As for properly instantiated variables there is no difference between standard and \perp -unification, we can safely assume that we can use \perp -unification also for SInst-Gen.

THEOREM 4. *Let S be an unsatisfiable set of clauses, S' a set of instances of clauses from S such that $(S \cup S')\perp$ is unsatisfiable. If M is a finite subset of $S \cup S'$ such that $M\perp$ is unsatisfiable, then SInst-Gen with \perp -unification can prove unsatisfiability of S by only generating generalizations of clauses from M .*

Proof. As the set of all generalizations of clauses in M is finite (up to renaming), it is not necessary to consider infinite derivations using such clauses. So the only way how the construction of an SInst-Gen proof using generalizations of clauses in M can fail is that at some point of the proof, SInst-Gen has generated a set of clauses M_1 such that $M_1\perp$ is satisfiable and every possible inference on M_1 results in generation of an instance which is not a generalization of any clause in M . As $M_1\perp$ is satisfiable, we can choose a selection function sel on M_1 . Every clause $C\sigma \in M$ has at least one most specific generalization C with respect to M_1 . Suppose we choose one of these most specific generalizations for every $C\sigma \in M$ and select the literal $L\sigma \in C\sigma$ if L is selected by sel in C . As $M\perp$ is unsatisfiable, we must have selected at least one pair of \perp -complementary literals, say $L_1\sigma_1 \in C_1\sigma_1$ and $L_2\sigma_2 \in C_2\sigma_2$.

Thus, in M_1 , sel selects $L_1 \in C_1$ and $L_2 \in C_2$. As $L_1\sigma_1$ and $L_2\sigma_2$ are \perp -complementary, we can state that L_1 and L_2 are complementary unifiable, say by τ . As they are not \perp -complementary, there is an SInst-Gen inference with \perp -unification between them. The substitution used in this inference is a most general \perp -unifier, therefore the clauses produced by this inference will also be generalizations of clauses from M . Note that this would not hold for SInst-Gen without \perp -unification. The inference generates at least one proper instance with respect to the premises, say $C_1\tau$ is a proper instance of C_1 . There cannot be a variant of $C_1\tau$ in M_1 , as C_1 was chosen to be a most specific generalization of $C_1\sigma$. Thus, we have produced a new generalization of a clause from M , contradicting our assumption that no such inference is possible.

COROLLARY 2. *For every refuting DCC proof (with or without \perp -unification) there exists a weakly simulating SInst-Gen proof (with \perp -unification).*

Proof. For a DCC proof that shows the unsatisfiability of a set of clauses S , let S' be the finite set of all clauses on the DCC tableau. Since $(S \cup S')\perp$ is unsatisfiable, we can apply the previous theorem.

THEOREM 5. *Let S be an unsatisfiable set of clauses, S' a set of instances of clauses from S such that $(S \cup S')\perp$ is unsatisfiable. If M is a finite subset of $S \cup S'$ such that $M\perp$ is unsatisfiable, then DCC with \perp -unification can prove unsatisfiability of S by only generating generalizations of clauses from M .*

Proof. Again, the fact that the set of all generalizations of clauses in M is finite ensures that unfair derivations need not be considered. Suppose a DCC tableau for S has an open branch B that cannot be extended without generating instances that are not generalizations of clauses in M . Every clause in $C\sigma \in M$ has at least one most specific generalization C with respect to the clauses on the tableau path $P_S \cup P_B$. We select in each $C\sigma \in M$ the literal corresponding to the literal of C on the tableau path. As $M\perp$ is unsatisfiable, \perp -complementary literals $L\sigma \in C\sigma$ and $K\tau \in D\tau$ are selected. Thus the literals L and K of the most specific generalizations C and D are on the tableau path. The DCC inference with \perp -unification from C and D uses a most general \perp -unifier, so the instances produced by this inference are again generalizations of $C\sigma$ and $D\tau$; moreover, at least one of them is a proper instance of a premise. This instance cannot be a variant of a clause on the tableau path, since C and D were chosen as most specific generalizations of $C\sigma$ and $D\tau$. Therefore, the tableau can be extended with a generalization of a clause in M , contradicting our assumption.

COROLLARY 3. *For every refuting SInst-Gen or PPI proof (with or without \perp -unification) there exists a weakly simulating DCC proof (with \perp -unification).*

5. Comparing SInst-Gen and DCC: Consistency Proofs

5.1. FROM SINST-GEN TO DCC

The case of consistency proofs differs in several aspects from the case of refuting proofs. First, it is clear that no theorem proving method is guaranteed to terminate for satisfiable sets of clauses. Second, in order to declare a set of clauses unsatisfiable, SInst-Gen must check *all* propositional interpretations for the given set of clauses and show that none of them is a model. In contrast to this, termination on a satisfiable set of clauses only depends on *one* interpretation which can be extended to a first-order model of the input clauses. Essentially, the same holds for DCC. Simulation can therefore be based on the final set of clauses and selection function. In the following, we will show that this fact enables us to simulate any SInst-Gen proof which terminates on satisfiable input by DCC.

THEOREM 6. *Let S be a satisfiable set of input clauses, $S \cup S'$ a finite set of clauses saturated under SInst-Gen with selection function sel . Then the given consistency proof can be simulated by DCC.*

Proof. We prove our claim by induction on the number of proof steps of the simulating proof, where a proof step consists of both the selection of literals and the generation of instances. Simulation is not based on the steps of the given proof, but only on the final set of clauses $S \cup S'$ and the selection function sel . We will show that every step of the simulating proof produces only instances from $S \cup S'$, while the tableau path we follow is always equivalent to sel for the clauses which are on this path.

First, we choose the initial path P_S of our simulating proof to select the same literals as sel on S . Then, every link on P_S can only produce instances from $S \cup S'$, as otherwise there would be an inconsistency of sel . Thus, we may carry out an arbitrary link from those which are available on P_S , resulting in a set of clauses which is a subset of $S \cup S'$.

Now, suppose an arbitrary number of steps has been carried out by our simulating proof, always following a tableau path which selects the same literals as sel for clauses on the path and only generating instances from $S \cup S'$. We can extend the path from the last step such that it selects the same literals as sel on the new instances. As sel does not

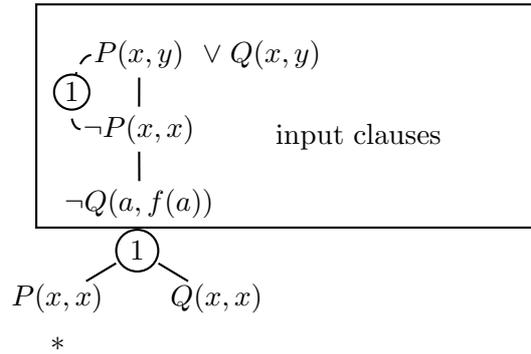


Figure 5. Saturated disconnection tableau

select \perp -complementary literals, the path must also be open. By the same argument as above, we may again carry out any possible linking step, producing only instances from $S \cup S'$.

In this way, we carry out all possible linking steps. As we choose our tableau path such that only instances from $S \cup S'$ are produced, the process must terminate after a finite number of steps. As termination of the process means that we have saturated the current path, we can state that we have reached our goal to simulate the given proof.

It may happen that a linking step adds two instances, but sel is such that we cannot select a tableau path which is equivalent to the selection function *and* considers both of these clauses. This happens if both of the selected literals are not linked literals of the linking step that produces them. In this case, our tableau path can only consider one of the produced clauses. However, as we want to saturate the branch we are following, missing clauses and links are not a problem, but a benefit in this case.

Note that this result also implies that there is no satisfiable set of clauses on which SInst-Gen can terminate, but DCC cannot.

5.2. FROM DCC TO SINST-GEN

We have shown that every SInst-Gen consistency proof can be simulated by DCC. The reverse, however, does not hold: Figure 5 shows a saturated disconnection tableau for a satisfiable set of input clauses. Saturation of the tableau is achieved by generating a single non-proper instance of an input clause.

An SInst-Gen proof for the given set of clauses would have to select $Q(x, y)$ in the first input clause, as \perp -complementary literals must not be selected. Thus, we have an inconsistency between this literal and

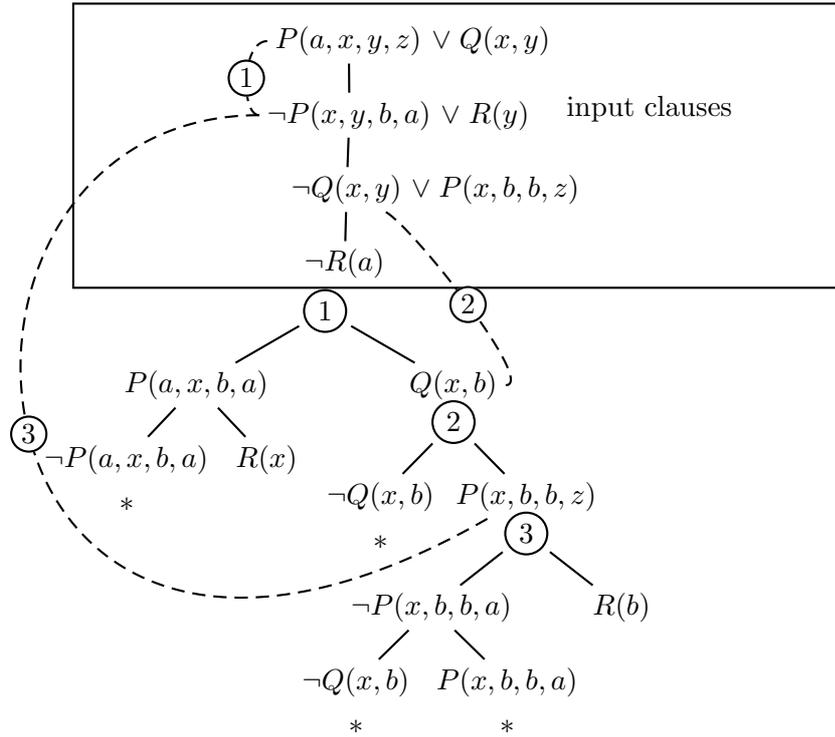


Figure 6. Simulation still fails in satisfiable case

$\neg Q(a, f(a))$, which produces an instance we do not have on the tableau. Without producing this instance however, satisfiability of the input set cannot be detected by SInst-Gen.

If \perp -unification is used for both DCC and SInst-Gen, examples like this one are ruled out, as only proper instances or variants of input clauses will be added to the tableau. However, even in this case simulation is in general not possible. Simulation of the tableau in Figure 6 by SInst-Gen is not possible, as we have shown in detail in Jacobs [7]. In this tableau, the rightmost branch is saturated. What makes simulation impossible for SInst-Gen are the clauses in the tableau which are not on this branch. They do not influence saturation of the tableau branch, as the linking rule only considers literal occurrences on the current branch and the initial path. In SInst-Gen however, these additional instances give rise to inconsistencies which have to be resolved before saturation can be reached.

As for refutation proofs, this result still holds when hyper-inferences are allowed in the simulating SInst-Gen proof.

6. Conclusions

We have compared the four instance generation methods DCC, Inst-Gen, SInst-Gen, and PPI. Inst-Gen, which does not make any attempt to restrict the search space, is obviously the most general of these calculi: any DCC, SInst-Gen, or PPI proof can be simulated by an Inst-Gen proof.

PPI is essentially a special case of SInst-Gen, except for its term-depth-based saturation strategy, which ensures completeness of the calculus but is not subsumed by the fairness criterion of SInst-Gen. We have given a new fairness criterion that preserves completeness and allows all strategies that are allowed in PPI or SInst-Gen.

For DCC and SInst-Gen, we have demonstrated that, for refutation proofs, simulation between (the basic versions of) the methods is in general not possible. This implies in particular that neither of the methods can be considered as a special case of the other. In case of consistency proofs, we have shown that SInst-Gen proofs that terminate on satisfiable input can always be simulated by DCC, while this does not hold for the other direction. All of these results still hold when SInst-Gen is allowed to use hyper-inferences.

We face a very different situation when we consider weak simulation, so that we can not only use clauses from the given proof but also generalizations thereof. We have shown that DCC and SInst-Gen with \perp -unification can weakly simulate each other; in fact we conjecture that DCC and SInst-Gen with \perp -unification can weakly simulate any instance-based calculus.

For DCC and SInst-Gen there are various refinements which are beyond the scope of this work. It is not clear how our results would translate to the refined calculi, e.g. DCC with lemma generation and subsumption or SInst-Gen with redundancy elimination. Additionally, for both methods extensions to equality reasoning are available. A comparison of the different approaches to these refinements and extensions might give more useful insights on the relation between DCC and SInst-Gen.

References

1. Baumgartner, P.: 2000, ‘FDPLL – A First-Order Davis-Putnam-Logeman-Loveland Procedure’. In: D. McAllester (ed.): *CADE-17 – The 17th International Conference on Automated Deduction*. pp. 200–219.
2. Baumgartner, P. and C. Tinelli: 2003, ‘The Model Evolution Calculus’. In: F. Baader (ed.): *CADE-19 – The 19th International Conference on Automated Deduction*. pp. 350–364.
3. Billon, J.-P.: 1996, ‘The Disconnection Method: a Confluent Integration of Unification in the Analytic Framework’. In: P. Miglioli, U. Moscato, D. Mundici, and M. Ornaghi (eds.): *Automated Reasoning with Analytic Tableaux and Related Methods: International Conference*. pp. 110–126.
4. Davis, M., G. Logemann, and D. Loveland: 1962, ‘A Machine Program for Theorem Proving’. *Communications of the ACM* **5**, 201–215.
5. Ganzinger, H. and K. Korovin: 2003, ‘New Directions in Instantiation-Based Theorem Proving’. In: *Proc. 18th IEEE Symposium on Logic in Computer Science*. pp. 55–64.
6. Hooker, J. N., G. Rago, V. Chandru, and A. Rivastava: 2002, ‘Partial Instantiation Methods for Inference in First-Order Logic’. *Journal of Automated Reasoning* **28**, 371–396.
7. Jacobs, S.: 2004, ‘Instance Generation Methods for Automated Reasoning’. Diploma thesis, Universität des Saarlandes. Available at <http://www.mpi-sb.mpg.de/~sjacobs/publications>.
8. Lee, S.-J. and D. A. Plaisted: 1992, ‘Eliminating Duplication with the Hyper-Linking Strategy’. *Journal of Automated Reasoning* **13**, 25–42.
9. Letz, R. and G. Stenz: 2001, ‘Proof and Model Generation with Disconnection Tableaux’. In: R. Nieuwenhuis and A. Voronkov (eds.): *Logic for Programming, Artificial Intelligence, and Reasoning: 8th International Conference*. pp. 142–156.
10. Letz, R. and G. Stenz: 2004, ‘Generalised Handling of Variables in Disconnection Tableaux’. In: D. Basin and M. Rusinowitch (eds.): *Automated Reasoning: Second International Joint Conference*.
11. Plaisted, D. A. and Y. Zhu: 2000, ‘Ordered Semantic Hyper-Linking’. *Journal of Automated Reasoning* **25**, 167–217.
12. Stenz, G.: 2002, ‘The Disconnection Calculus’. Ph.D. thesis, TU München.