

Depth-first Proof Search without Backtracking for Free-variable Clausal Tableaux

BERNHARD BECKERT

*Universität Karlsruhe, D-76128 Karlsruhe, Germany.
Email: beckert@ira.uka.de*

Abstract

We analyse the problem of constructing a deterministic proof procedure for free-variable clausal tableaux that performs depth-first proof search without backtracking; and we present a solution based on a fairness strategy. That strategy uses *weight orderings* and a notion of *tableau subsumption* to avoid proof cycles and it employs *reconstruction steps* to handle the destructiveness of free-variable calculi.

1. Introduction

In this paper, we analyse the problem of constructing a deterministic proof procedure for free-variable tableau calculi that performs depth-first proof search and is complete without backtracking. As an example, we present a solution for first-order clausal tableaux that is based on a fairness strategy. That strategy uses *weight orderings* and a notion of *tableau subsumption* to avoid proof cycles and it employs *reconstruction steps* to handle the destructiveness of clausal tableaux.

First-order clausal tableaux are *proof-confluent*, i.e., every tableau for an unsatisfiable clause set can be completed to a proof. They are, however, a *destructive* calculus because all occurrences of a (free) variable in a tableau have to be instantiated by the same term and, thus, a rule application can make another rule application impossible.

The proof search space can be visualised as a search tree where each possible choice of the next rule application to a tableaux T creates a node with as many successor nodes as T has different successor tableaux (Fig. 1). Since we use a proof-confluent calculus, all fully exhausted branches in the search tree (for an unsatisfiable input clause set) are either infinite or end in a node that is labelled with a proof, i.e., a closed tableau.

There are two main concepts for proof search: *breadth-first* and *depth-first*




Figure 1: A proof search tree.

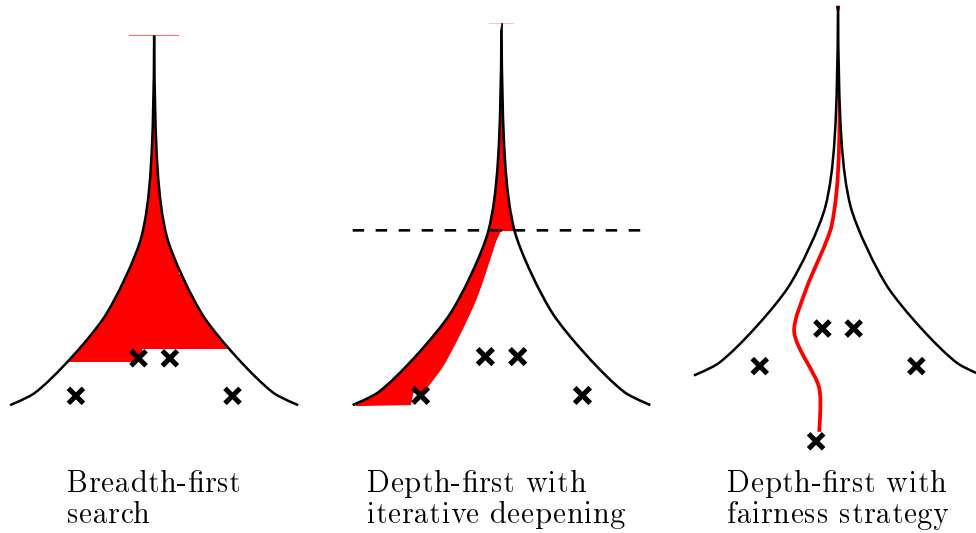


Figure 2: Comparison of proof search strategies.

search. Depth-first search requires that either there are no paths in the search tree that do not contain proofs or it is possible to avoid such paths using fairness strategies for the construction of tableaux.

As fairness strategies that allow depth-first search are difficult to construct for first-order clausal tableaux, most automated deduction systems use breadth-first search. It allows to find shorter proofs than depth-first search because all paths of the search tree are considered whereas, using depth-first search, paths in the search tree that contain short proofs may be missed; fairness strategies only guarantee that some proof is found but it may not be the shortest one. However, the length of found proofs is not of great importance in automated deduction (the only advantage of short proofs is that their construction requires less rule applications and that they are thus easier to find); and breadth-first search is “expensive” as compared to depth-first search because neighbouring paths in the search tree contain many similar or even identical tableaux that using breadth-first search all have to be considered.

For all (useful) iterative deepening strategies, i.e., all monotonic functions m from \mathbb{N} to sets of tableaux where $\bigcup_{i \in \mathbb{N}} m(i)$ includes *all* constructible tableaux

(i is the iteration level), the size of the partial search trees $m(i)$ grows *exponentially* in i . It is thus, even for small i , impossible to store all the tableaux in $m(i)$ in the memory of a machine. Therefore, most implementations use *depth-first iterative deepening* (DFID). The initial, partial search space consisting of all the tableaux in $M(i) = \bigcup_{j < i} m(j)$ for some $i \in \mathbb{N}$ is searched for proofs in a depth-first manner using backtracking, and if it turns out not to contain a proof, then i is increased (for example, the proof procedure described by Fitting [1996] is of this type). Then, however, the tableaux in $M(i)$ are not available for the construction of the tableaux in $M(i + 1)$; they have to be constructed again from scratch, which, however, merely causes polynomial overhead as compared to a breadth-first search at the “right” level i because $M(i + 1)$ is exponentially larger than $M(i)$. Although DFID search leads to acceptable performance of tableau-based automated theorem provers, it should be stressed that it is only a compromise used when no completeness preserving fairness strategy for depth-first search is available.

The advantage of depth-first proof search is that the information represented by the constructed tableaux increases at each proof step; no information is lost since there is no backtracking. In addition, considering similar tableaux or sequences of tableaux in different paths of the search tree is avoided.

Figure 2 shows how the different search strategies traverse the search space. The coloured part has to be searched before a proof is found. The form of the search space visualises its exponential growth.

In the case of *non-destructive* and proof-confluent tableau calculi—such as the ground version of first-order tableaux that does not use free variables—it is relatively easy to use depth-first proof search; it suffices to systematically add all possible conclusions until all branches of the constructed tableau are either fully expanded or closed. The situation is much more complicated in free-variable clausal tableau calculi, which are *destructive* (even if they are proof-confluent). Applying a substitution may destroy literals on a tableau that are needed for the proof, such that they have to be deduced again.

Up to now there was no useful solution to the problem of constructing a deterministic proof procedure for free-variable clause tableaux that performs depth-first search and is complete, i.e., that never fails to find a proof if there is one. Such procedures were only known for the special case where tableaux are expanded without instantiating variables and only a single substitution is finally applied that is known to allow to close all branches simultaneously. This solution is not really satisfactory as (1) free variables are only used in a very restricted way, and (2) checking whether there is a single substitution simultaneously closing all branches of a tableau is computationally expensive. Solving the problem for a similar calculus, Baumgartner et al. [1999] recently described a depth-first proof procedure for a *connection calculus*.

We propose in this paper a deterministic search strategy that is based upon:

- A *tableau subsumption* relation to detect “cycles” in the search (i.e., to make

sure that it is not possible to deduce the same literals or sub-tableaux again and again).

- *Weight orderings* that assign each literal a “weight” in such a way that there are only finitely many different literals (up to variable renaming) of a certain weight; thus, if literals with lesser weight are deduced first, then sooner or later each possible conclusion is added to all branches containing its premiss.
- *Reconstruction steps* to handle the destructiveness of free-variable clausal tableaux. Immediately after a rule application that destroys literals, the construction steps that are needed to recreate the destroyed sub-tableaux are executed.

The main difficulty is to define a tableau subsumption relation that on the one hand is restrictive enough to avoid cycles in the proof construction and on the other hand is not too restrictive such that completeness is preserved.

Our fairness strategy considers the whole tableau tree (and not only a single branch) both for the subsumption check and for choosing a conclusion of minimal weight; a procedure based on this strategy may extend any branch of a tableau at any time. Note that this does not imply a large memory consumption; at least it is not worse than that of proof strategies where a “current” branch is extended until it is closed before other branches are considered and where DFID-based breadth-first search is used to ensure completeness, as in that case all closed branches have to be stored for backtracking.

As said above, no *useful* deterministic proof procedures for free-variable clausal tableaux were known up to now. There is trivially a (not useful) deterministic proof procedure for all proof-confluent calculi, namely a procedure performing a *breadth-first* search in the background. “Useful” means that the computational complexity of deciding what the next rule application should be in each situation has to be reasonably low. In addition, the number of construction steps that are necessary to find a proof has to be reasonably small as compared to the number of necessary steps when a breadth-first search strategy is used.

If the fairness strategy we present in the following sections is used, then the complexity of deciding what the next expansion step should be is in the worst case quadratic in the size of the tableau to be expanded and its possible successor tableaux. In the average case the complexity is much lower as only those parts of a tableau have to be considered that are affected by one of the possible tableau rule applications. The number of steps required to construct a proof is smaller than (or equal to) the worst-case number of construction steps using DFID.

The structure of the paper is as follows: In Section 2, we describe the calculus of clausal tableaux. After introducing our notion of tableau subsumption in Section 3 and that of weight orderings in Section 4, our method for constructing deterministic proof procedures for free-variable clausal tableaux is presented in Section 5.

In [Beckert, 1998] the same techniques for constructing deterministic proof

procedures are applied to a more general class of tableau calculi that are in a certain sense well-behaved, and of which free-variable clausal tableaux are an instance.

2. First-order Clausal Tableaux

A *first-order signature* $\Sigma = \langle P_\Sigma, F_\Sigma, \alpha_\Sigma \rangle$ consists of a non-empty set P_Σ of predicate symbols, a set F_Σ of function symbols, and a function $\alpha_\Sigma : P_\Sigma \cup F_\Sigma \rightarrow \mathbb{N}$ assigning an arity to the predicate and function symbols; functions of arity 0 are called constants. Moreover, there is an infinite set Var of *object variables*.

The *logical operators* of clause logic are the connectives \neg (negation) and \vee (disjunction) and the universal quantifier \forall . In addition, we use the logical constants \top (true) and \perp (false).

The notions of *free* and *bound variable*, *term*, *atom*, *literal*, and *substitution* are defined as usual. We use x, y, z etc. to denote quantified variables and X, Y, Z etc. to denote free variables. The logical constants \top and \perp are considered to be literals (but not atoms). The complement of a literal L is denoted with \overline{L} .

DEFINITION 2.1: *A variable renaming is a substitution that replaces all occurring variables by distinct new variables (i.e., variables that are new w.r.t. the context where the variable renaming is applied).*

DEFINITION 2.2: *A clause C over a signature Σ is a first-order formula of the form*

$$(\forall x_1) \cdots (\forall x_n)(L_1 \vee \cdots \vee L_r)$$

where the L_i are literals over Σ and x_1, \dots, x_n are all the variables occurring in L_1, \dots, L_r .

A new instance of a clause C is a formula $(L_1 \vee \cdots \vee L_r)\sigma$ where σ is a variable renaming.

DEFINITION 2.3: *A clausal tableau for a set S of clauses is built by a sequence of applications of the following construction rules. Each rule has a premiss (a set of literals) and a conclusion (consisting of a set of literals and a substitution).*

Initialisation: The tree consisting of a single node labelled with \top is a tableau for S .

Expansion: If T is a tableau for S , B is a branch of T , and $L_1 \vee \cdots \vee L_r$ is a new instance of a clause in S , then a tableau T' is a tableau for S if obtained by extending B with r nodes L_1, \dots, L_r . (In this case, the premiss is empty and the conclusion is $\langle \{L_1, \dots, L_r\}, id \rangle$.)

Closure: If T is a tableau for S , B is a branch of T , L, L' are literals on B , and $L, \overline{L'}$ are unifiable with most general unifier σ , then T' is a tableau for S if obtained by appending \perp to B and applying σ to each node of T . (In this case, the premiss is $\{L, L'\}$ and the conclusion is $\langle \{\perp\}, \sigma \rangle$.)

Note, that a branch is closed by adding the special literal \perp ; therefore, branch closure can be considered to be a special kind of branch expansion.

DEFINITION 2.4: *A tableau T is closed if all its branches are closed, i.e., contain \perp .*

We use a slightly non-standard definition of the notion of successor tableau.

DEFINITION 2.5: *A tableau T' is a successor tableau of a tableau T if it is constructed from T by one or more “identical” rule applications, i.e., there are*

1. *different branches B_1, \dots, B_n ($n \geq 1$) of T ,*
2. *premisses Π_i on the B_i such that Π_i, Π_j are identical up to variable renaming,*
3. *a (single) conclusion $\langle C, \sigma \rangle$ such that $\Pi_i \sigma = \Pi_j \sigma$ ($1 \leq i, j \leq n$),*

and T' is constructed from T by extending each of the branches B_i with the literals in C and applying the substitution σ to T .

DEFINITION 2.6: *Let S be a clause set.*

A tableau sequence for S is a sequence $(T_i)_{i \geq 1}$ of tableaux for S such that T_{i+1} is a successor tableau of T_i ($i \geq 1$).

A tableau proof for (the unsatisfiability of) S is a finite tableau sequence T_1, \dots, T_n for S such that T_n is closed.

Clausal tableaux as defined above are a complete and proof-confluent calculus.

THEOREM 2.1: *Let S be an unsatisfiable clause set.*

Completeness: There is a tableau proof for the unsatisfiability of S .

Proof confluence: Every finite tableau sequence T_1, \dots, T_k for S can be extended to a tableau proof $T_1, \dots, T_k, T_{k+1}, \dots, T_n$ for S .

Our technique for constructing a backtracking-free proof procedure (as described in the following sections) is compatible with many search space restrictions that preserve proof confluence. Examples are selection functions [Hähnle and Pape, 1997] and connection tableaux with restart, where (except for restart steps) a clause used for expansion must have a link into the branch being expanded.

3. Tableau Subsumption Relation

Assume that a sequence T_1, \dots, T_n of tableaux has already been constructed. A rule application to T_n is forbidden if the successor tableau T_{n+1} is *subsumed* by one of the predecessor tableaux T_j —in particular, if T_{n+1} is subsumed by T_n . In that case, the sequence T_j, \dots, T_{n+1} constitutes a cycle in the proof search because T_{n+1} does not contain any information that is not already in T_j .

We define a tableau T_j to subsume a tableau T_{n+1} iff each branch of T_j subsumes one of the branches of T_{n+1} . Intuitively, the tableau T_{n+1} is in that case redundant because, if closed sub-tableaux can be constructed below all branches of T_{n+1} , it is possible to construct closed sub-tableaux of the same complexity below all branches of T_j as each of them subsumes a branch of T_{n+1} .

When does a tableau *branch* subsume another branch? A first approximate answer to that question is: A branch B subsumes a branch B' if B contains a variant of each literal occurring on B' . That, however, is an over-simplification; three additional aspects have to be taken into concern.

First additional aspect. For a branch B to subsume a branch B' , it is in general *not* sufficient if the branch B contains *one* variant of each literal L occurring in B' , namely in case B' contains two variants of L that are both needed to close the branch. However, since the premiss for a single rule application contains at most two literals, it is sufficient if B contains a variant of each set of (at most) two literals occurring on B' . This implies that at most two variable-disjoint variants of each literal on B' are needed on B (where however, as described below, literals may have to be considered to be effectively different although they are variants of each other on first sight).

EXAMPLE 3.1: *If the literals $\neg p(X)$, $p(f(X))$, $\neg p(X')$, $p(f(X'))$ occur on B' whereas the branch B only contains $\neg p(X)$ and $p(f(X))$ (and B and B' are otherwise identical), then B contains a variant of each literal on B' . Nevertheless, the transition from B to B' is definitely not a cycle in proof search because—contrary to B —the branch B' can be closed.*

Second additional aspect. The second important aspect is that, in order to decide whether some literal L on a branch B can be considered to be a variant of some literal L' on a branch B' , not only the literals L and L' have to be taken into concern but also *associated* literals (on B and B' as well as on other branches) that have free variables in common with L and L' .

DEFINITION 3.1: *Literals L_1 and L_2 are associated if there is a variable occurring in both L_1 and L_2 . The set of all literals in a tableau T that are associated with a literal L , excluding L itself, is denoted with $\text{Assoc}(T, L)$. Accordingly, if Φ is a set of literals, then $\text{Assoc}(T, \Phi)$ is the literal set $(\bigcup_{L \in \Phi} \text{Assoc}(T, L)) \setminus \Phi$.*

Associated literals play a role because the ordering of tableau rule applications used by a deterministic proof procedure as described in Section 5 has to take all literals into account that are generated by an application. So, if $L(X)$ is a premiss for a certain tableau rule application that leads to the instantiation of X with a term t and there is a literal $L'(X)$ on the tableau, then that application will generate the new literal $L'(t)$ (besides $L(t)$). Since the choice of which rule is applied next to a tableau depends the new literals that possible rule applications

generate, the form of $L'(t)$ —and thus the form of the associated literal $L'(X)$ —can affect the choice of the next rule application (in particular if $L'(X)$ has a higher weight than $L(X)$, see Section 4).

Third additional aspect. As said above, a tableau T subsumes a tableau T' if for each branch B in T there is a branch B' in T' such that B subsumes B' . That includes the possibility that the same branch B' of T' is assigned to two different branches B_1 and B_2 of T . In that case there is for each set Φ' (of at most two literals on B') a literal set Φ_1 on B_1 and a literal set Φ_2 on B_2 that are variants of Φ' . The basic idea behind the definition of our subsumption relation implies that every possible rule application on branch B' with the premiss Φ' can as well be applied—simultaneously—on the branches subsuming B' with the premisses Φ_1 resp. Φ_2 . That, however, requires the two variable renamings constructing Φ' from Φ_1 resp. Φ_2 to be compatible. The same holds if B' is assigned to more than two branches in T .

Formal definition of the subsumption relation. We now formally define our tableau subsumption relation. It is transitive and reflexive.



DEFINITION 3.2: *Let T and T' be tableaux that do not have any variables in common. The tableau T subsumes the tableau T' , denoted by $T' \subseteq T$, if*

- i. it is possible to assign to each branch B of T a branch B' of T'*
- ii. and then, for each pair B, B' respectively, to each set Φ' of at most two literals on B' , a set Φ of literals B and a variable renaming π can be assigned such that:*

- 1. The following holds for Φ , Φ' and π :*
 - (a) $\Phi\pi = \Phi'$;*
 - (b) for each of the literals L in $\text{Assoc}(T, \Phi)$ there is (at least) one literal L' in $\text{Assoc}(T', \Phi')$ such that $L\pi$ and L' are identical up to the renaming of variables not occurring in $\Phi\pi$ resp. Φ' .*
- 2. If a branch B' of T' is assigned to different branches B_1, \dots, B_s of T ($s \geq 2$), then, for all Φ' on B' , the variable renamings π_1, \dots, π_s assigned to Φ' in connection with B_1, \dots, B_s are compatible in the following way: there is a substitution π such that the restriction of π to the variables occurring in $\Phi_i \cup \text{Assoc}(T, \Phi_i)$ is identical to π_i ($1 \leq i \leq s$).*

Now, let T and T' be tableaux that have variables in common; and let ρ be a variable renaming such that T and $T'\rho$ do not have any variables in common. Then, T subsumes T' iff T subsumes $T'\rho$.

If a tableau T subsumes a tableau T' , then each branch B of T is assigned a branch B' of T' . In that case, we say that B subsumes B' .

**Figure 3:** The tableaux from Example 3.3.**Figure 4:** The tableaux from Example 3.4.



Using the tableau subsumption relation for restricting the search space preserves completeness of clausal tableaux: Given a partial proof T_1, \dots, T_i it is forbidden to derive a successor tableau T_{i+1} from T_i that is subsumed by any of the tableau T_1, \dots, T_i . On the other hand, this restriction is strong enough to ensure that every sequence of tableaux built accordingly, i.e., every tableau sequence not containing a tableau that is subsumed by one of its predecessors, has the following property: If the sequence is infinite, then it contains infinitely many different literals or, equivalently, if the sequence only contains finitely many different literals (up to the renaming of variables) then it is finite.

To check whether a tableau T subsumes one of its successor tableaux T' and, thus, whether the rule application deriving T' from T is allowed, it is sufficient to only consider those parts of the tableaux that are affected, i.e., the expanded branch and the formulae on the tableaux that are associated with it. The check does not involve unifiability tests because free variables may only be renamed but not instantiated with terms.

EXAMPLE 3.2: Let $\Phi = \{p(X)\}$ and let $\Phi' = \{p(X')\}$; further let $\text{Assoc}(T, \Phi)$ consist of $q(X, Y_1)$ and $q(X, Y_2)$.

Then, Condition 1b in Definition 3.2 is, for example, satisfied if $q(X', Y')$ occurs in $\text{Assoc}(T', \Phi')$. But it is neither satisfied if $\text{Assoc}(T', \Phi') = \emptyset$ nor if $\text{Assoc}(T', \Phi') = \{q(Y', X')\}$ (because to make $q(X', Y_1)$ and $q(Y', X')$ identical would require to rename the variable X' that occurs in Φ').

EXAMPLE 3.3: The tableau T_1 in Figure 3 subsumes each of the tableaux T'_1, T'_2, T'_3 . The tableau T_2 subsumes only T'_1 .

**Figure 5:** The tableaux from Example 3.5.**Figure 6:** The tableaux from Example 3.6.

EXAMPLE 3.4: *Neither of the two tableaux in Figure 4 subsumes the other one. The tableau T_1 does not subsume the tableau T_2 because the (single) branch of T_2 contains an additional literal; and, although a variant of each literal set on T_1 occurs on T_2 , the tableau T_2 does not subsume T_1 as for $r(X')$, which is an element of $\text{Assoc}(T_2, q(X'))$, there is no corresponding element in $\text{Assoc}(T_1, q(X))$ and, thus, Condition 1b in Definition 3.2 is not satisfied.*

EXAMPLE 3.5: *The tableau T_1 in Figure 5 subsumes the tableau T_2 . But T_2 does not subsume T_1 because the literals $p(X')$ and $q(X')$ in T_2 are associated, whereas the corresponding literals in T_1 are not associated.*

Indeed, a transition from T_2 to T_1 does not constitute a cycle in proof search because the tableau T_1 can be closed whereas T_2 cannot be closed.

EXAMPLE 3.6: *The tableau T_1 in Figure 6 does not subsume the tableau T' .*

That would (only) be possible if the branches B_1, B_2 of T both subsumed the single branch of T' . Both B_1 and B_2 contain a variant $p(X, Y)$ resp. $p(Y, X)$ of the (single) literal $p(U, V)$ on the branch of T' . But the required variable renamings $\{X \mapsto U, Y \mapsto V\}$ and $\{X \mapsto V, Y \mapsto U\}$ are not compatible, which violates Condition 2 in the definition of the subsumption relation (Def. 3.2). This problem does not occur with the tableau T_2 in Figure 6. It subsumes T' because the two required variable renamings $\{X_1 \mapsto U, Y_1 \mapsto V\}$ and $\{X_2 \mapsto U, Y_2 \mapsto V\}$ are compatible.

EXAMPLE 3.7: *Consider the tableau T_1 shown on the left in Figure 7. The rule application that derives the conclusion $\langle \{\perp\}, \text{id} \rangle$ from the premiss $\{\neg p, p\}$ can be used to close both of its branches.*




Figure 7: The tableaux from Example 3.7.




Figure 8: The tableaux from Example 3.8.

Closing the tableau requires two consecutive applications. However, the intermediate tableau T_2 that results from closing the left branch (in the middle of Figure 7) is subsumed by T_1 because both branches of T_1 subsume the right (not yet expanded) branch of T_2 . Thus, this first rule application is not allowed. The tableau T_3 , however, that results from closing both branches (shown in the right in Figure 7) is neither subsumed by T_2 nor by T_1 . Indeed, since both rule applications use the same premiss and conclusion, T_3 is by definition a successor tableau of T_1 (without considering the intermediate step), and deriving T_3 from T_1 is an allowed rule application.

EXAMPLE 3.8: An important type of tableau construction steps that generate a tableau T' subsumed by its predecessor T and that are, therefore, forbidden, is the following: Assume that a branch B_1 of T is extended using a conclusion $\langle C, \sigma \rangle$, and some branch $B'_2\sigma$ in the new tableau T' is subsumed by all branches B of T affected by the rule application, i.e., the branch B_1 (which is extended) and all other branches containing variables that are instantiated by σ . This is in particular the case if $B'_2\sigma$ is “contained” in an initial sub-branch R_0 of B_2 that ends above the first occurrence of any free variable in the domain of σ . As an example consider the tableau T shown in Figure 8 on the left, and assume that its branch B_1 is closed using the premiss consisting of the two literals $\neg p(a)$ and $p(X)$ to derive the conclusion $\langle \{\perp\}, \{X \mapsto a\} \rangle$. The right branch $B'_2\sigma$ of the resulting tableau T' (shown in Figure 8 on the right) whose nodes are labelled with

the literals $\neg p(a)$ and twice $q(a)$ is “contained” in the sub-branch R_0 of T whose nodes are labelled with $\neg p(a)$ and $q(a)$; and R_0 ends above the first occurrence of X in T which is the only variable instantiated by σ . Intuitively, the application is useless because any closed sub-tableau that can be constructed below $B'_2\sigma$ can be constructed as well below both B_1 and B_2 .

A forbidden rule application as described above is irregular according to the definition of regularity that is usually given in the literature (e.g. [Beckert and Hähnle, 1998]) since the branch $B'_2\sigma$ contains the same branch extension multiply. In general, however, irregular expansion steps are not necessarily forbidden according to our subsumption condition—nor are all forbidden expansion steps irregular.

The proof of Theorem 5.1, which is the main theorem of this paper, makes use of the following lemma. It states an important property of tableau sequences that satisfy the subsumption condition, i.e., that do not contain a tableau that subsumes one of its successors.

LEMMA 3.1: *Let Γ be a finite set of literals; and let $(T_i)_{i \geq 1}$ be a sequence of tableaux such that*

1. *no tableau T_i subsumes any of the tableaux T_j for $i < j$,*
2. *all literals occurring in $(T_i)_{i \geq 1}$ are variants of the literals in Γ .*

Then the sequence $(T_i)_{i \geq 1}$ is finite.

Proof: Let the equivalence relation \sim on tableaux be defined by: $T \sim T'$ iff T and T' subsume each other, i.e., $T \subseteq T'$ and $T' \subseteq T$.

We show that there is only a finite number of equivalence classes of tableaux (w.r.t. \sim) that consist of variants of literals from Γ .

Since Γ is finite, there are—up to variable renaming—only finitely many pairs $\langle \Phi, \Psi \rangle$ of literal sets such that:

- All elements of Φ and Ψ are variants of the literals in Γ ;
- Φ has at most two elements;
- every literal in Ψ has at least one variable in common with a literal in Φ ;
- there are no different literals ψ, ψ' in Ψ that are identical up to the renaming of variables not occurring in Φ .

Let $r(\Phi)$ be the number of different, incompatible renamings of the variables in a set Φ of at most two literals (where the notion of compatible variable renamings is defined as in Condition 2 in Definition 3.2); $r(\Phi)$ only depends on the number of different variables in Φ .

Now, let $\widehat{\Gamma}$ be a set of pairs of literal sets as described above, such that (a) $\widehat{\Gamma}$ contains $r(\Phi)$ different representatives of each class of pairs that are identical up to variable renaming to some $\langle \Phi, \Psi \rangle$ and (b) the elements of $\widehat{\Gamma}$ are all pairwise variable disjoint. Since $\widehat{\Gamma}$ is finite, the power set \mathcal{P} of the power set of $\widehat{\Gamma}$ is finite.

Each tableau T is assigned an element of \mathcal{P} as follows: To each set Φ of at most two literals occurring on some branch of T , the pair $\langle \Phi, \text{Assoc}(T, \Phi) \rangle$ is assigned. To each branch B of T , a subset of $\widehat{\Gamma}$ is assigned that contains variants of the pairs assigned to sets Φ of at most two literals occurring on B , where at most $r(\Phi)$ of the pairs that are identical up to variable renaming are included. To the tableau T , an element of \mathcal{P} is assigned consisting of variants of the sets assigned to the branches of T .

The definitions of \subseteq resp. \sim imply that two tableaux belong to the same equivalence class w.r.t. \sim if they can be assigned the same element of \mathcal{P} . Thus, there are at most as many equivalence classes as there are elements in \mathcal{P} —which is a finite set.

Now we can complete the proof. Let $n \in \mathbb{N}$ be the number of equivalence classes of tableaux w.r.t. \sim . It is an upper bound for the length of the sequence $(T_i)_{i \geq 1}$, because if there were more than n tableaux in the sequence $(T_i)_{i \geq 1}$, then it had to contain two tableaux T_i and T_j ($i < j$) belonging to the same equivalence class. That, however, implies $T_j \subseteq T_i$ which contradicts the lemma's assumptions. \square

4. Weight Orderings

Weight orderings are the second important concept (besides the concept of tableau subsumption) on which our fairness strategy is based. The properties an ordering on literals for ensuring fairness must have are: (1) It is a well-ordering on the set of literals (up to renaming of free-variables), i.e., it is well-founded and there are only finitely many literals that are incomparable to a given literal—such that for each literal L all but a *finite* number of literals (up to variable renaming) are greater than L w.r.t. the ordering. (2) Proper instances of a literal L have a higher weight than L . (3) Literals that are identical up to variable renaming have the same weight. Intuitively, these are typical properties of orderings on literals that are defined by assigning a “weight” to the symbols of a signature (which is why we call them *weight* orderings).

DEFINITION 4.1: *A weight function w assigns to each literal over Σ a natural number (its weight), such that the following conditions are satisfied:*

1. *Let Φ be a set of literals. If*
 - (a) *Φ does not contain different literals that are identical up to variable renaming, and*
 - (b) *there is an $n \in \mathbb{N}$ such that $w(\phi) \leq n$ for all $\phi \in \Phi$,**then Φ is finite.*
2. *If X is a variable occurring in some literal ϕ and t is a term that is not a variable, then*

$$w(\phi) < w(\phi\{X \mapsto t\}) .$$

3. If literals ϕ, ϕ' are identical up to the renaming of variables, then

$$w(\phi) = w(\phi') .$$

Let w be a weight function. Then the weight ordering \leq_w (that is induced by w) is for all literals ϕ, ψ over Σ defined by:

$$\phi \leq_w \psi \text{ iff } w(\phi) \leq w(\psi) .$$

A weight ordering is extended to *sets* of literals by comparing the *maximal* weight of the literals they contain. This extension is a well-ordering as well, provided the literal sets are only allowed to contain a certain number of variants of each literal.

EXAMPLE 4.1: Assume that every function symbol $f \in F_\Sigma$ is assigned a weight $w_f \in \mathbb{N}$ and every predicate symbol $p \in P_\Sigma$ is assigned a weight $w_p \in \mathbb{N}$, such that only finitely many symbols are assigned the same weight.

For all literals ϕ , let $w(\phi)$ be the sum of the weights of all function and predicate symbols in ϕ , where the weight of a symbol occurring multiply is added multiply. Then, w is a weight function (Def. 4.1).

5. Deterministic Proof Procedures for Clausal Tableaux

In this section, we define a (class of) complete deterministic proof procedure(s) for clausal tableaux; this proof procedure can be used to perform depth-first search for proofs without backtracking. It is constructed using the notions of subsumption and weight orderings as described in Sections 3 and 4.

To ensure that a deterministic proof procedure is complete, i.e., a proof is found if there is one, we demand that the constructed sequence of tableaux satisfies the following two conditions: (1) The creation of a tableau that is subsumed by one of its predecessors is forbidden. (2) At each step, from all possible rule applications not violating Condition 1, an application is chosen that creates a successor tableau in which the maximal weight of literals is as small as possible (i.e., successor tableaux are compared according to the maximal weight of the literals they contain). If several rule applications satisfy these conditions, arbitrary heuristics may be employed to choose one of them; e.g., rule applications creating less new sub-branches may be preferred.

Note that conclusions are not necessarily added to a tableau branch in the order defined by the maximal weight of their literals because a literal L can only be added if the necessary premiss Π is present on the branch; and the weight of the literals in Π may be higher than that of L . Also, at which point a certain conclusion is added, depends on its literal with the highest weight. Thus, literals with a lower weight that can only be added as part of a conclusion containing other literals of higher weight are added to the tableau later.

To comply with the condition that *all* rule applications adding literals of

less weight have to be executed before literals of higher weight are added to a tableau, it may be necessary to expand branches that are already closed. That is not always redundant, because closed branches still contain useful information and can influence other branches by the substitutions that are applied when they are expanded (the first substitution that is applied to close a branch is not necessarily the “right one” that allows to complete the proof). If a closed branch has no free variables in common with other branches, it needs not be further expanded.

Unfortunately, the restriction of the search space as described above is difficult to implement; it requires to compare a tableau T_{n+1} with *all* its predecessors T_1, \dots, T_n and not only with the tableau T_n from which it is derived. Such a subsumption check is prohibitively expensive w.r.t. both space and time. Moreover, if a subsumption is encountered, i.e., if T_{n+1} is subsumed by one of the predecessor tableaux T_j , then other successor tableaux of T_j (besides T_{j+1}) must be considered, which in a certain sense amounts to backtracking. The reason for this is the following: A tableau T_{n+1} that is subsumed by a tableau T_j does not have to be considered for proof search as all the proofs that may be constructed from T_{n+1} can be constructed from T_j . Now, if $j = n$, then we can just exclude the successor tableau T_{n+1} and be sure that if there is a proof derivable from T_{n+1} then it is derivable from T_n without considering T_{n+1} . If, however, $j < n$, then the closed tableau that is known to be derivable from T_{n+1} and thus from T_j may not involve T_{j+1}, \dots, T_n but require to proceed with an alternative successor tableau T'_{j+1} different from T_{j+1} . This situation is visualised in the upper part of Figure 9.

The following example demonstrates that cycles in the proof search are not all of length one or two but can indeed be of arbitrary length (the cycle in the example is of length four).

EXAMPLE 5.1: *Figure 10 shows a tableau sequence $(T_i)_{i \geq 1}$ for the clause set*

$$S = \{ \neg p(a), \neg r(b), (\forall x)(p(x) \vee q(x)), (\forall y)(r(y) \vee s(y)) \} .$$

For $i \geq 8$, the tableau T_i subsumes the tableau T_{i+4} (and none of the tableaux T_{i+1}, \dots, T_{i+3}). Thus, the sequence contains a cycle in proof search of length four.

The sequence is constructed as follows: After the tableau T_8 has been derived, the following rule applications are repeatedly executed:

1. *A conclusion of the form $\langle \{r(Y), s(Y)\}, id \rangle$ is derived and used to expand the rightmost branch (e.g., to derive tableau T_9 from tableau T_8).*
2. *The third branch from the right is closed using a premiss consisting of literals $\neg p(a)$ and $p(X)$ and a conclusion $\langle \{\perp\}, \{X \mapsto a\} \rangle$ (e.g., to derive T_{10} from T_9).*
3. *A conclusion of the form $\langle \{p(X), q(X)\}, id \rangle$ is derived and used to expand the rightmost branch (e.g., to derive tableau T_{11} from tableau T_{10}).*

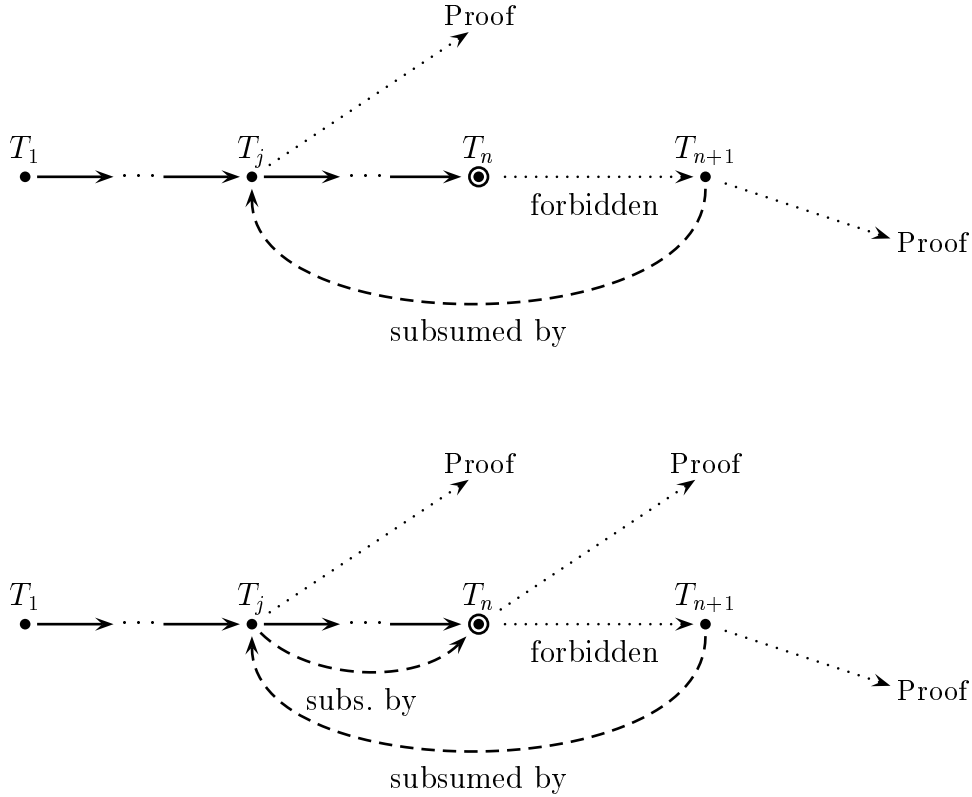



Figure 9: Proof search with a calculus that is *destructive* (top) resp. *non-destructive* (bottom).

4. The third branch from the right is closed using a premiss consisting of literals $\neg r(b)$ and $r(Y)$ and a conclusion $\langle \{\perp\}, \{Y \mapsto b\} \rangle$ (e.g., to derive T_{12} from T_{11}).

All problems stem from the fact that a tableau T_j is not necessarily subsumed by its successor tableau T_{j+1} as the clausal tableau calculus is destructive and literals occurring in T_j may not occur in T_{j+1} any more. However, if we make the calculus *weakly non-destructive* in the sense that a tableau is always subsumed by all its successor tableaux, then we have the situation shown in the lower part Figure 9. Now, the tableau T_j is subsumed by the tableau T_n ensuring that every proof that can be constructed from T_{n+1} can as well be constructed from T_n —without deriving T_{n+1} as an intermediate result. In a certain sense, a (weakly) non-destructive calculus is proof-confluent w.r.t. the restricted search space (where no tableaux subsumed by a predecessor are allowed).

To make clausal tableaux weakly non-destructive, i.e., to make sure that a tableau T_{i+1} always subsumes its predecessor tableau T_i , we impose the following additional restriction on the proof construction: Immediately after a tableau construction step destroying literals, the construction steps needed to recreate the destroyed literals must be executed. In the worst case, a new copy of the sub-tableau that was affected by the variable instantiation is created and appended



$$S = \{ \neg p(a), \neg r(b), (\forall x)(p(x) \vee q(x)), (\forall y)(r(y) \vee s(y)) \}$$

Figure 10: A sequence of tableaux containing a proof search cycle; see Example 5.1.

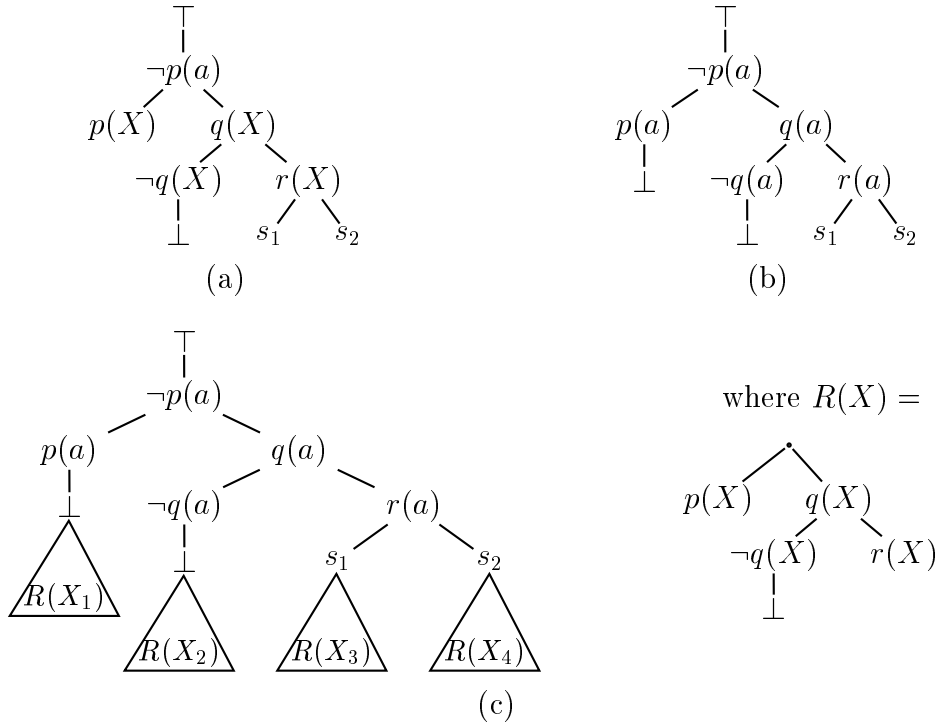


Figure 11: A tableau reconstruction step (Example 5.2).

to all sub-branches that have been affected. The result is a tableau T_{i+1}^+ that subsumes both T_i and T_{i+1} and all the tableaux occurring as intermediate results during the reconstruction.

EXAMPLE 5.2: Consider the clause set

$$S = \{ (\forall x)(p(x) \vee q(x)), (\forall x)(\neg q(x) \vee r(x)), s_1 \vee s_2 \} .$$

A tableau T_i for S is shown in Figure 11 (a). The left branch of T_i is closed using the conclusion $\{\{\perp\}, \{X \mapsto a\}\}$. The result is the tableau T_{i+1} (Figure 11 (b)), in which all literals containing the free variable X have been destroyed. They are reconstructed by appending a copy of the sub-tableau $R(X)$ that consists of all literals in T_i in which X occurs to all the branches in T_{i+1} from which literals are missing; the resulting tableau T_{i+1}^+ (Figure 11 (c)) subsumes both T_i and T_{i+1} .

If a deterministic proof procedure executes a reconstruction step after each tableau rule application, then a sequence T_1^+, T_2^+, \dots of tableaux is constructed where T_{i+1}^+ is derived from T_i^+ by executing a construction step (that does not lead to a tableau subsumed by its predecessor) and then reconstructing the destroyed literals. To ensure that such a sequence meets all conditions, it is sufficient to test whether the immediate successor tableau T_{i+1}^+ of T_i^+ is subsumed by T_i^+ . The earlier predecessors do not have to be considered as they are all subsumed by T_i^+ . Theorem 5.1 below states completeness of such a proof procedure; it is the main theorem of this paper.

THEOREM 5.1: *If a clause set S is unsatisfiable, then every sequence $(T_i^+)_{i \geq 1}$ of tableaux for S that is constructed as described below contains a closed tableau.*

The tableau T_1^+ is an initial tableau for S . And for all $i > 1$ the following holds:

1. T_{i+1} is a successor tableau of T_i^+ (Def. 2.5) such that
 - (a) T_i^+ does not subsume T_{i+1} and
 - (b) there is no successor tableau T'_{i+1} of T_i^+ that satisfies Condition (a) and has a smaller maximal literals weight than T_{i+1} (w.r.t. an arbitrary but fixed weight ordering).
2. Let $\langle C_i, \tau_i \rangle$ be the conclusion (derived from some premiss on T_i^+) that is used to construct T_{i+1} ; and let R_i be the minimal sub-tableau of T_i^+ that contains all occurrences of the variables instantiated by τ_i . The tableau T_{i+1}^+ is constructed from T_{i+1} by (repeatedly) executing all rule applications that are necessary to generate R_i ; R_i is appended to all branches that go through the sub-tableau of T_{i+1} corresponding to R_i (which results from applying τ_i to R_i).

Proof: Part 1. The definition of the tableau construction rules immediately implies that, if a conclusion $\langle C, \sigma \rangle$ is derivable from a premiss Π , then the conclusion $\langle C, id \rangle$ is derivable from the premiss $\Pi\sigma$ for all substitutions σ . Therefore, the reconstruction step (i.e., the rule applications that are necessary to construct T_i^+ from T_i) does not require the instantiation of free variables. Consequently, the reconstruction is non-destructive, and we have

$$T_i \subseteq T_i^+$$

for all $i \geq 1$.

By construction of T_{i+1}^+ , all literals that are destroyed by applying the substitution τ_i to T_i^+ are appended by the reconstruction step to all branches from which they are missing in T_{i+1} . Therefore $T_i^+ \subseteq T_{i+1}^+$ and, thus,

$$T_i^+ \subseteq T_j^+ \quad \text{for all } i \leq j.$$

Part 2. We show by contradiction that the sequence $(T_i^+)_{i \geq 1}$ does not contain a tableau that subsumes one of its successors. Assume the contrary, i.e., there are tableaux $T_i^+, T_j^+, i < j$, such that $T_j^+ \subseteq T_i^+$. Using the results of Part 1, that implies

$$T_{i+1} \subseteq T_{i+1}^+ \subseteq T_j^+ \subseteq T_i^+ .$$

By construction of the sequence $(T_i^+)_{i \geq 1}$, however, the tableau T_{i+1} is *not* subsumed by T_i^+ . Thus, the assumption is wrong, and $(T_i^+)_{i \geq 1}$ does indeed not contain a tableau that subsumes one of its successors.

Part 3. Let $w_{\max} \in \mathbb{N}$ be an arbitrary weight. We prove finiteness of the initial sub-sequence of $(T_i^+)_{i \geq 1}$ that only contains literals ϕ of weight $w(\phi) \leq w_{\max}$.

Let Γ be a set of (variants of) representatives of each equivalence class of literals in $(T_i^+)_{i \geq 1}$ that are identical up to variable renaming and whose weight is not bigger than w_{\max} . That is, if ϕ is a literal in $(T_i^+)_{i \geq 1}$ with $w(\phi) \leq w_{\max}$, then there is a variant of ϕ in Γ . Let the literals in Γ be chosen in such a way that they are all pairwise variable-disjoint.

The definition of weight orderings implies that the set Γ must be finite. Thus, by Lemma 3.1, the initial sub-sequence of $(T_i^+)_{i \geq 1}$ that only contains literals of weight $w(\phi) \leq w_{\max}$ must be finite.

Part 4. Since S is unsatisfiable and the (unrestricted) clausal tableau calculus is complete, there is a tableau proof T'_1, \dots, T'_m for S (it possibly violates the subsumption condition and is not constructed as described in the theorem). Let w_{\max} be the maximum of the weights of all the literals occurring in this tableau sequence; and let T_n^+ be the last tableau in the initial sub-sequence of $(T_i^+)_{i \geq 1}$ that only contains literals of weight not bigger than w_{\max} (as has been shown in Part 3, this sub-sequence is finite).

We prove by induction on j that there are substitutions σ_j for $1 \leq j \leq m$ such that

$$T'_j \subseteq T_n^+ \sigma_j \text{ .}$$

Base case $j = 1$. Let $\sigma_1 = id$. The tableaux T'_1 and T_1^+ are both initial tableaux for S and do not contain any free variables. Thus, $T'_1 = T_1^+ \subseteq T_n^+ = T_n^+ \sigma_1$ holds trivially.

Induction step $j \rightarrow j + 1$. Let Π' be the premiss and $\langle C, \rho \rangle$ be the conclusion used to derive T'_{j+1} from T'_j . And let B'_j be the branch in T'_j that is extended. Since ρ is a most general unifier of literals in T'_j , it instantiates only variables occurring in T'_j .

If no branch in $T_n^+ \sigma_j$ subsumes the branch B'_j , then every branch of $T_n^+ \sigma_j$ subsumes one of the branches of T'_j that are not expanded. In that case we have $T'_{j+1} \subseteq T_n^+ \sigma_j \rho$ and thus $T'_{j+1} \subseteq T_n^+ \sigma_{j+1}$ with $\sigma_{j+1} = (\rho \circ \sigma_j)$.

Otherwise, there are branches B^1, \dots, B^s in $T_n^+ \sigma_j$ that subsume B'_j . Thus, there are sets Π^l of literals on B^l and a substitution π such that $\Pi^l \pi = \Pi'$. According to Condition 2 in the definition of the subsumption relation (Def. 3.2), the substitution π does not depend on l . Therefore, B can be extended using the conclusion $\langle C, \rho \circ \pi \rangle$. Now, let \widehat{T} be the tableau, that is derived from $T_n^+ \sigma_j$ by extending each of the branches B^1, \dots, B^s using the conclusion $\langle C, \rho \circ \pi \rangle$. The successor tableau \widehat{T} of $T_n^+ \sigma_j$ subsumes T'_{j+1} , i.e., $T'_{j+1} \subseteq \widehat{T}$.

Let \widehat{K} be the set of all literals that are added by the rule applications used to derive \widehat{T} from $T_n^+ \sigma_j$, which are the literals in C and in $(Assoc(T_n^+, \Pi))\pi\rho$. Because $T'_{j+1} \subseteq \widehat{T}$, by the definition of \subseteq and by the construction of \widehat{T} , all literals in \widehat{K} already occur in T'_{j+1} . Therefore, by definition of w_{\max} , the weight of all literals in \widehat{K} is not bigger than w_{\max} .

Using rule applications corresponding to those used to derive \widehat{T} from $T_n^+ \sigma_j$, it is possible to derive a Tableau \widetilde{T} from T_n^+ such that $\widetilde{T}\sigma_{j+1} = \widehat{T}$ for some

substitution σ_{j+1} . Let \tilde{K} be the set of all literals that are added by the rule applications used to derive \tilde{T} from T_n^+ . Then there is a substitution $\tilde{\sigma}$ such that $\tilde{K}\tilde{\sigma} = \hat{K}$. Therefore, the weight of the literals in \tilde{K} is not bigger than w_{\max} . Also, by definition, the weight of the literals in T_n^+ is not bigger than w_{\max} . Thus, the weight of all literals in \tilde{T} is not bigger than w_{\max} .

Since, by construction, T_n^+ is the last tableau in the sequence $(T_i)_{i \geq n}$ that does not contain literals of weight bigger than w_{\max} , every rule application that leads to a tableau with this property must be forbidden because it leads to a tableau subsumed by T_n^+ (because otherwise it would have to be derived instead of the tableau T_{n+1} , which contains a literal of weight bigger than w_{\max}). Thus, we have $\tilde{T} \subseteq T_n^+$ and, therefore, $\tilde{T}\sigma_{j+1} \subseteq T_n^+\sigma_{j+1}$. Consequently,

$$T'_{j+1} \subseteq \hat{T} = \tilde{T}\sigma_{j+1} \subseteq T_n^+\sigma_{j+1} .$$

Part 5. We now can conclude the proof of the theorem. As a consequence of what is proven in Part 4, the tableau $T_n^+\sigma_m$ subsumes the tableau T'_m . Since T'_m is closed, each branch of T'_m contains the literal \perp . Thus, each branch of $T_n^+\sigma_m$ contains \perp ; and, since \perp does not contain variables, each branch of T_n^+ contains \perp . Therefore, T_n^+ is closed. \square

EXAMPLE 5.3: *As an example for the proof construction as described in this section, Figure 12 shows a closed tableau for the clause set consisting of the clauses $\neg p(a)$, $\neg p(b)$, $\neg q(b)$, $(\forall x)(p(x) \vee q(x))$. The proof construction starts with adding the unit literals to the initial tableau; the result is the tableau T_1 . At this point only one rule application is possible, which results in the tableau T_2 . Then there are several possibilities to proceed; the left branch of T_2 can be closed instantiating X_1 with either a or b and the right branch can be closed instantiating X_1 with b . We assume that according to the weight ordering, $p(a) \leq_w p(b)$ and $q(a) \leq_w q(b)$. Consequently, the “bad” instantiation $\{X_1 \mapsto a\}$ is preferred and the tableau T_3 is constructed, because the maximal weight of its literals is less than that of the literals in the alternative tableaux. Since the variable X_1 is instantiated, a reconstruction step is required; the result of that step is the tableau T_4 . Now there are again several possibilities. If the weight of literals were the only criterion, then a tableau T'_5 would have to be derived from T_4 that is identical to T_5 except that X_3 is instantiated with a instead of b (i.e., repeating the useless instantiation). However, deriving T'_5 from T_4 is not allowed as T'_5 is subsumed by T_4 (it is easy to check that each branch of T_4 subsumes one of the branches of T'_5). Therefore, the tableau T_5 is derived instead of T'_5 ; and the variable X_3 is instantiated with b instead of a . Again, a reconstruction step is required, which results in the tableau T_6 . From T_6 the closed tableau T_7 can easily be constructed.*

A proof procedure as described in Theorem 5.1 constructs a tableau sequence T_1^+, T_2^+, \dots such that no tableau is subsumed by any of its predecessors and all tableaux are subsumed by their successors. Such a procedure simulates (in a




Figure 12: The closed tableau described in Example 5.3.

certain sense) a depth-first iterative deepening search (as described in the introduction). The weight of the literals that can occur in the tableaux increases stepwise. If some (unrestricted) tableau proof exists that does not contain literals of weight bigger than w_{\max} , then there is a closed tableau T_n^+ that is the last in the constructed sequence not containing literals of weight bigger than some $w_{\max}^+ \in \mathbb{N}$. It subsumes all tableaux that can be constructed from literals L of weight $w(L) \leq w_{\max}$. The big advantage of this simulated DFID over classical DFID search based on backtracking is that the tableau T_n^+ is a very compact representation of the search space. All the information that is contained in tableaux whose literals are of weight less than w_{\max} is present in the single structure T_n^+ ; and all the tableaux in the search space that are identical or in some way symmetrical to each other are represented by only one sub-tableau of T_n^+ . Since no backtracking occurs, no information that has been derived is ever lost. There may be parts of the tableau T_n^+ that represent redundant information and are therefore useless (i.e., non-closed sub-tableaux that should not have been created); but these can be easily be detected and removed (e.g., using the *pruning* technique [Beckert and Hähnle, 1998]).

References

- Peter Baumgartner, Norbert Eisinger, and Ulrich Furbach. A confluent connection calculus. In H. Ganzinger, editor, *Proceedings, Conference on Automated Deduction (CADE), Trento, Italy*, LNCS 1632, pages 329–343. Springer, 1999.
- Bernhard Beckert. *Integration und Uniformierung von Methoden des tableau-basierten Theorembeweisens*. PhD thesis, Universität Karlsruhe, Fakultät für Informatik, July 1998. Available at www.ubka.uni-karlsruhe.de/eva/.
- Bernhard Beckert and Reiner Hähnle. Analytic tableaux. In Wolfgang Bibel and Peter H. Schmitt, editors, *Automated Deduction — A Basis for Applications*, volume I: Foundations. Kluwer, Dordrecht, 1998.
- Melvin Fitting. *First-Order Logic and Automated Theorem Proving*. Springer, New York, second edition, 1996.
- Reiner Hähnle and Christian Pape. Ordered tableaux: Extensions and applications. In *Proceedings, International Conference on Theorem Proving with Analytic Tableaux and Related Methods, Pont-à-Mousson, France*, LNCS 1227, pages 173–187. Springer, 1997.

Acknowledgement

I thank Reiner Hähnle and Peter H. Schmitt for fruitful discussions, and the anonymous referees for useful comments on earlier versions of this paper.