Full First-Order Sequent and Tableau Calculi With Preservation of Solutions and the Liberalized δ-Rule but Without Skolemization

Claus-Peter Wirth

Informatik 5, Universität Dortmund, D-44221, Germany
wirth@LS5.cs.uni-dortmund.de

Abstract. We present a combination of raising, explicit variable dependency representation, the liberalized δ-rule, and preservation of solutions for first-order deductive theorem proving. Our main motivation is to provide the foundation for our work on inductive theorem proving.

1 Introduction

1.1 Without Skolemization

We discuss how to analytically prove first-order theorems in contexts where Skolemization is not appropriate. It has at least three problematic aspects.

1. Skolemization enrichs the signature or introduces higher-order variables. Unless special care is taken, this may introduce objects into empty universes and change the notion of term-generatedness or Herbrand models. Above that, the Skolem functions occur in answers to goals or solutions of constraints which in general cannot be translated into the original signature. For a detailed discussion of these problems cf. Miller (1992).
2. Skolemization results in the following simplified quantification structure: "For all Skolem functions $\vec{u}$ there are solutions to the free existential variables $\vec{e}$ (i.e. the free variables of Fitting (1996)) such that the quantifier-free theorem $T(\vec{e},\vec{u})$ is valid. Short: $\forall\vec{u}\colon\ \exists\vec{e}\colon\ T(\vec{e},\vec{u})$." Since the state of a proof attempt is often represented as the conjunction of the branches of a tree (e.g. in sequent or (dual) tableau calculi), the free existential variables become "rigid" or "global", i.e. a solution for a free existential variable must solve all occurrences of this variable in the whole proof tree. This is because, for $B_0,\ldots,B_n$ denoting the branches of the proof tree,
$$\forall\vec{u}\colon\ \exists\vec{e}\colon\ (\ B_0\ \wedge\ \ldots\ \wedge\ B_n\ )$$
is logically strictly stronger than
$$\forall\vec{u}\colon\ (\ \exists\vec{e}\colon\ B_0\ \ \wedge\ \ \ldots\ \ \wedge\ \ \exists\vec{e}\colon\ B_n\ ).$$
Moreover, with this quantification structure it does not seem to be possible to do inductive theorem proving (*ITP*) by finding, for each assumed counterexample, another counterexample that is strictly smaller in some wellfounded ordering.[1] The reason for this is the following. When we have some counterexample $\vec{u}$ for $T(\vec{e},\vec{u})$ (i.e. there is no $\vec{e}$ such that $T(\vec{e},\vec{u})$ is valid) then for

[1] While this paradigm of ITP was already used by the Greeks, Pierre de Fermat

different $\vec{e}$ different branches B_i in the proof tree may cause the invalidity of the conjunction. If we have applied induction hypotheses in more than one branch, for different $\vec{e}$ we get different smaller counterexamples. What we would need, however, is one single smaller counterexample for all $\vec{e}$.

3. Skolemization increases the size of the formulas. (Note that in most calculi the only relevant part of Skolem terms is the top symbol and the set of occurring variables.)

The 1st and 2nd problematic aspects disappear when one uses *raising* (cf. Miller (1992)) instead of Skolemization. Raising is a dual of Skolemization and simplifies the quantification structure to something like: "There are raising functions $\vec{e}$ such that for all possible values of the free universal variables $\vec{u}$ (i.e. the nullary constants or "parameters") the quantifier-free theorem $T(\vec{e},\vec{u})$ is valid. Short: $\exists\vec{e}\colon\ \forall\vec{u}\colon\ T(\vec{e},\vec{u})$." Note that due to the two duality switches "unsatisfiability/validity" and "Skolemization/raising", in this paper raising will look much like Skolemization in refutational theorem proving. The inverted order of universal and existential quantification of raising (compared to Skolemization) is advantageous because now $\exists\vec{e}\colon\ \forall\vec{u}\colon\ (\ B_0 \wedge \ldots \wedge B_n\ )$ is indeed logically equivalent to $\exists\vec{e}\colon\ (\ \forall\vec{u}\colon B_0\ \wedge\ \ldots\ \wedge\ \forall\vec{u}\colon B_n\ )$. Furthermore, ITP works well: When, for some $\vec{e}$, we have some counterexample $\vec{u}$ for $T(\vec{e},\vec{u})$ (i.e. $T(\vec{e},\vec{u})$ is invalid) then one branch B_i in the proof tree must cause the invalidity of the conjunction. If this branch is closed, then it contains the application of an induction hypothesis that is invalid for this $\vec{e}$ and the $\vec{u}'$ resulting from the instantiation of the hypothesis. Thus, $\vec{u}'$ together with the induction hypothesis provides the strictly smaller counterexample we are searching for for this $\vec{e}$. The 3rd problematic aspect disappears when the dependency of variables is explicitly represented in a *variable-condition*, cf. Kohlhase (1995). This idea actually has a long history, cf. Prawitz (1960), Kanger (1963), Bibel (1987). Moreover, the use of variable-conditions admits the free existential variables to be first-order.

1.2 Sequent and Tableau Calculi

In Smullyan (1968), rules for analytic theorem proving are classified as α-, β-, γ-, and δ-rules independently from a concrete calculus. *α-rules* describe the simple and the *β-rules* the case-splitting propositional proof steps. *γ-rules* show existential properties, either by exhibiting a term witnessing to the existence or else by introducing a special kind of variable, called "dummy" in Prawitz (1960) and Kanger (1963), and "free variable" in footnote 11 of Prawitz (1960) and in Fitting (1996). We will call these variables *free existential variables*. With free existential variables we can delay the choice of a witnessing term until the state of the proof attempt gives us more information which choice is likely to result in a successful proof. It is the important addition of free existential variables that makes the major difference between the free variable calculi of Fitting (1996) and the calculi of Smullyan (1968). Since there use to be infinitely many

(1601-1665) rediscovered it under the name "descente infinie", and in our time it is sometimes called "implicit induction", cf. Wirth & Becker (1995), Wirth (1997).

possibly witnessing terms (and different branches may need different ones), the γ-rules (under assistance of the β-rules) often destroy the possibility to decide validity because they enable infinitely many γ-rule applications to the same formula. *δ-rules* show universal properties simply with the help of a new symbol, called a "parameter", about which nothing is known. Since the present free existential variables must not be instantiated with this new parameter, in the standard framework of Skolemization and unification the parameter is given the present free existential variables as arguments. In this paper, however, we will use nullary parameters, which we call *free universal variables*. These variables are not free in the sense that they may be chosen freely, but in the sense that they are not bound by any quantifier. Our free universal variables are similar to the parameters of Kanger (1963) because a free existential variable may not be instantiated with all of them. We will store the information on the dependency between free existential variables and free universal variables in *variable-conditions*.

1.3 Preservation of Solutions

Users even of pure Prolog are not so much interested in theorem proving as they are in answer computation. The theorem they want to prove usually contains some free existential variables that are instantiated during a proof attempt. When the proof attempt is successful, not only the input theorem is known to be valid but also the instance of the theorem with the substitution built-up during the proof. Since the knowledge of mere existence is much less useful than the knowledge of a term that witnesses to this existence (unless this term is a only free existential variable), theorem proving should—if possible—always provide these witnessing terms. Answer computation is no problem in Prolog's Horn logic, but also for the more difficult clausal logic, answer computation is possible: Cf. e.g. Baumgartner & al. (1997), where tableau calculi are used for answer computation in clausal logic. Answer computation becomes even harder when we consider full first-order logic instead of clausal logic. When δ-steps occur in a proof, the introduced free universal variables may provide no information on what kind of object they denote. Their excuse may be that they cannot do this in terms of computability or λ-terms. Nevertheless, they can provide this information in form of Hilbert's ε-terms, and the strong versions of our calculi will do so. When full first-order logic is considered, one should focus on *preservation of solutions* instead of computing answers. By this we mean at least the following property: "All solutions that transform a proof attempt for a proposition into a closed proof (i.e. the closing substitutions for the free existential variables) are also solutions of the original proposition." This is again closely related to ITP: Suppose that we finally have shown that for the reduced form $R(\vec{e}, \vec{u})$ (i.e. the state of the proof attempt) of the original theorem $T(\vec{e}, \vec{u})$ there is some solution $\vec{e}$ such that for each counterexample $\vec{u}$ of $R(\vec{e}, \vec{u})$ there is a counterexample $\vec{u}'$ for the original theorem that is strictly smaller than $\vec{u}$ in some wellfounded ordering. In this case we have proved $T(\vec{e}, \vec{u})$ only if the solution $\vec{e}$ for the reduced form $\forall \vec{u}\colon\ R(\vec{e}, \vec{u})$ is also a solution for the original theorem $\forall \vec{u}\colon\ T(\vec{e}, \vec{u})$.

1.4 The Liberalized δ-Rule

We use '$\uplus$' for the union of disjoint classes and 'id' for the identity function. For a class R we define *domain*, *range*, and *restriction to* and *image* and *reverse-image of a class* A by $\mathrm{dom}(R) := \{\, a \mid \exists b\colon (a,b) \in R \,\}$; $\mathrm{ran}(R) := \{\, b \mid \exists a\colon (a,b) \in R \,\}$; ${}_A|R := \{\, (a,b) \in R \mid a \in A\}$; $\langle A\rangle R := \{\, b \mid \exists a \in A\colon (a,b) \in R\}$; $R\langle B\rangle := \{\, a \mid \exists b \in B\colon (a,b) \in R\}$. We define a *sequent* to be a list of formulas. The *conjugate* of a formula A (written: $\overline{A}$) is the formula B if A is of the form $\neg B$, and the formula $\neg A$ otherwise. In the tradition of Gentzen (1935) we assume the symbols for *free existential variables*, *free universal variables*, *bound variables* (i.e. variables for quantified use only), and the *constants* (i.e. the function and predicate symbols from the signature) to come from four disjoint sets $\mathrm{V}_\exists$, $\mathrm{V}_\forall$, $\mathrm{V}_{\mathrm{bound}}$, and Σ. We assume each of $\mathrm{V}_\exists$, $\mathrm{V}_\forall$, $\mathrm{V}_{\mathrm{bound}}$ to be infinite and set $\mathrm{V}_{\mathrm{free}} := \mathrm{V}_\exists \uplus \mathrm{V}_\forall$. Due to the possibility to rename bound variables w.l.o.g., we do not permit quantification on variables that occur already bound in a formula; i.e. e.g. $\forall x\colon A$ is only a formula in our sense if A does not contain a quantifier on x like $\forall x$ or $\exists x$. The simple effect is that our γ- and δ-rules can simply replace *all* occurrences of x. For a term, formula, sequent Γ etc., '$\mathcal{V}_\exists(\Gamma)$', '$\mathcal{V}_\forall(\Gamma)$', '$\mathcal{V}_{\mathrm{bound}}(\Gamma)$', '$\mathcal{V}_{\mathrm{free}}(\Gamma)$' denote the sets of variables from $\mathrm{V}_\exists$, $\mathrm{V}_\forall$, $\mathrm{V}_{\mathrm{bound}}$, $\mathrm{V}_{\mathrm{free}}$ occurring in Γ, resp.. For a substitution σ we denote with '$\Gamma\sigma$' the result of replacing in Γ each variable x in $\mathrm{dom}(\sigma)$ with $\sigma(x)$. We tacitly assume that each substitution σ satisfies $\mathcal{V}_{\mathrm{bound}}(\mathrm{dom}(\sigma) \cup \mathrm{ran}(\sigma)) = \emptyset$, such that no bound variables can be replaced and no additional variables become bound (i.e. captured) when applying σ.

A *variable-condition* R is a subset of $\mathrm{V}_\exists \times \mathrm{V}_\forall$. Roughly speaking, $(x^\exists, y^\forall) \in R$ says that $x^\exists$ is older than $y^\forall$, so that we must not instantiate the free existential variable $x^\exists$ with a term containing $y^\forall$. While the benefit of the introduction of free existential variables in γ-rules is to delay the choice of a witnessing term, it is sometimes unsound to instantiate such a free existential variable $x^\exists$ with a term containing a free universal variable $y^\forall$ that was introduced later than $x^\exists$:

Example 1.1 $\exists x\colon \forall y\colon (x = y)$ is not deductively valid. We can start a proof attempt via: γ-step: $\forall y\colon (x^\exists = y)$. δ-step: $(x^\exists = y^\forall)$. Now, if we were allowed to substitute the free existential variable $x^\exists$ with the free universal variable $y^\forall$, we would get the tautology $(y^\forall = y^\forall)$, i.e. we would have proved an invalid formula. In order to prevent this, the δ-step has to record $(x^\exists, y^\forall)$ in the variable-condition, which disallows the instantiation step.

In order to restrict the possible instantiations as little as possible, we should keep our variable-conditions as small as possible. Kanger (1963) and Bibel (1987) are quite generous in that they let their variable-conditions become quite big:

Example 1.2 $\exists x\colon (\mathsf{P}(x) \vee \forall y\colon \neg\mathsf{P}(y))$ can be proved the following way: γ-step: $(\mathsf{P}(x^\exists) \vee \forall y\colon \neg\mathsf{P}(y))$. α-step: $\mathsf{P}(x^\exists)$, $\forall y\colon \neg\mathsf{P}(y)$. δ-step: $\mathsf{P}(x^\exists)$, $\neg\mathsf{P}(y^\forall)$. Instantiation step: $\mathsf{P}(y^\forall)$, $\neg\mathsf{P}(y^\forall)$. The last step is not allowed in the above citations, so that another γ-step must be applied to the original formula in order to prove it. Our instantiation step, however, is perfectly sound: Since $x^\exists$ does not occur in $\forall y\colon \neg\mathsf{P}(y)$, the free variables $x^\exists$ and $y^\forall$ do not depend on each other and there is no reason to insist on $x^\exists$ being older

than $y^\forall$. Note that moving-in the existential quantifier transforms the original formula into the logically equivalent formula $\exists x\colon \mathsf{P}(x) \ \vee \ \forall y\colon \neg\mathsf{P}(y)$, which enables the δ-step introducing $y^\forall$ to come before the γ-step introducing $x^\exists$.

Keeping small the variable-conditions generated by the δ-rule results in non-elementary reduction of the size of smallest proofs. This "liberalization of the δ-rule" has a history ranging from Smullyan (1968) over Hähnle & Schmitt (1994) to Baaz & Fermüller (1995). While the liberalized δ-rule of Smullyan (1968) is already able to prove the formula of Ex. 1.2 with a single γ-step, it is much more restrictive than the more liberalized δ-rule of Baaz & Fermüller (1995). Note that liberalization of the δ-rule is not simple because it easily results in unsound calculi, cf. Kohlhase (1995) w.r.t. our Ex. 1.3 and Kohlhase (1998) w.r.t. our Ex. 4.12. The difficulty lies with instantiation steps that relate previously unrelated variables:

Example 1.3 $\exists x\colon \forall y\colon \mathsf{Q}(x,y) \ \vee \ \exists u\colon \forall v\colon \neg\mathsf{Q}(v,u)$ is not deductively valid (to wit, let Q be the identity relation on a non-trivial universe). Proof attempt: One α-, two γ-, and two liberalized δ-steps result in

$$\mathsf{Q}(x^\exists, y^\forall), \ \neg\mathsf{Q}(v^\forall, u^\exists) \qquad (*)$$

with variable-condition

$$R \ := \ \{(x^\exists, y^\forall), \ (u^\exists, v^\forall)\}. \qquad (\#)$$

(Note that the non-liberalized δ-rule would additionally have produced $(x^\exists, v^\forall)$ or $(u^\exists, y^\forall)$ or both, depending on the order of the proof steps.)
When we now instantiate $x^\exists$ with $v^\forall$, we relate the previously unrelated variables $u^\exists$ and $y^\forall$. Thus, our new goal $\mathsf{Q}(v^\forall, y^\forall), \ \neg\mathsf{Q}(v^\forall, u^\exists)$ must be equipped with the new variable-condition $\{(u^\exists, y^\forall)\}$. Otherwise we could instantiate $u^\exists$ with $y^\forall$, resulting in the tautology $\mathsf{Q}(v^\forall, y^\forall), \ \neg\mathsf{Q}(v^\forall, y^\forall)$.
Note that in the Skolemization framework, this new variable-condition is automatically generated by the occur-check of unification:
When we instantiate $x^\exists$ with $v^\forall(u^\exists)$ in $\mathsf{Q}(x^\exists, y^\forall(x^\exists)), \ \neg\mathsf{Q}(v^\forall(u^\exists), u^\exists)$ we get $\mathsf{Q}(v^\forall(u^\exists), y^\forall(v^\forall(u^\exists))), \ \neg\mathsf{Q}(v^\forall(u^\exists), u^\exists)$, which cannot be reduced to a tautology because $y^\forall(v^\forall(u^\exists))$ and $u^\exists$ cannot be unified.
When we instantiate the variables $x^\exists$ and $u^\exists$ in the sequence $(*)$ in parallel via

$$\sigma := \{x^\exists \mapsto v^\forall, \ u^\exists \mapsto y^\forall\}, \qquad (\$)$$

we have to check whether the newly imposed variable-conditions are consistent with the substitution itself. In particular, a cycle as given (for the R of $(\#)$) by $y^\forall \ \sigma^{-1} \ u^\exists \ R \ v^\forall \ \sigma^{-1} \ x^\exists \ R \ y^\forall$ must not exist.

We make use of "[...]" for stating two definitions, lemmas, theorems etc. in one, where the parts between '[' and ']' are optional and are meant to be all included or all omitted. '$\mathbb{N}$' denotes the set of and '$\prec$' the ordering on natural numbers.

Validity is expected to be given with respect to some Σ-structure (Σ-algebra) $\mathcal{A}$, assigning a universe and an appropriate function to each symbol in Σ. For $\mathrm{X} \subseteq \mathrm{V}_{\mathrm{free}}$ we denote the set of total $\mathcal{A}$-valuations of X (i.e. functions mapping free variables to objects of the universe of $\mathcal{A}$) with $\mathrm{X} \to \mathcal{A}$ and the set of (possibly) partial $\mathcal{A}$-valuations of X with $\mathrm{X} \rightsquigarrow \mathcal{A}$. For $\pi \in \mathrm{X} \to \mathcal{A}$ we denote with '$\mathcal{A}\uplus\pi$' the extension of $\mathcal{A}$ to the variables of X which are then treated as

nullary constants. More precisely, we assume the existence of some evaluation function 'eval' such that $\mathrm{eval}(\mathcal{A}\uplus\pi)$ maps any term over $\Sigma\uplus\mathrm{X}$ into the universe of $\mathcal{A}$ such that for all $x \in \mathrm{X}$: $\mathrm{eval}(\mathcal{A}\uplus\pi)(x) = \pi(x)$. Moreover, $\mathrm{eval}(\mathcal{A}\uplus\pi)$ maps any formula B over $\Sigma\uplus\mathrm{X}$ to TRUE or FALSE, such that B is valid in $\mathcal{A}\uplus\pi$ iff $\mathrm{eval}(\mathcal{A}\uplus\pi)(B) = \mathsf{TRUE}$. We assume that the *Substitution-Lemma* holds in the sense that, for any substitution σ, Σ-structure $\mathcal{A}$, and valuation $\pi \in \mathrm{V}_{\mathrm{free}} \to \mathcal{A}$, validity of a formula B in $\mathcal{A} \uplus \big((\sigma \uplus {}_{\mathrm{V}_{\mathrm{free}}\setminus\mathrm{dom}(\sigma)}|\mathrm{id}) \circ \mathrm{eval}(\mathcal{A}\uplus\pi)\big)$ is logically equivalent to validity of $B\sigma$ in $\mathcal{A}\uplus\pi$. Finally, we assume that the value of the evaluation function on a term or formula B does not depend on the free variables that do not occur in B: $\mathrm{eval}(\mathcal{A}\uplus\pi)(B) = \mathrm{eval}(\mathcal{A} \uplus {}_{\mathcal{V}_{\mathrm{free}}(B)}|\pi)(B)$. Further properties of validity or evaluation are definitely not needed.

2 Two Versions of Variable-Conditions

We now describe two possible choices for the formal treatment of variable-conditions. The *weak* version works well with the non-liberalized δ-rule. The *strong* version is a little more difficult but can also be used for the liberalized versions. Several binary relations on free variables will be introduced. The overall idea is that when (x, y) occurs in such a relation this means something like "x is older than y" or "the value of y depends on or is described in terms of x".

Definition 2.1 (E_σ, U_σ, [Strong] Existential R-Substitution)
For a substitution σ with $\mathrm{dom}(\sigma) = \mathrm{V}_\exists$ we define the *existential relation* to be $E_\sigma := \{ (x', x) \mid x' \in \mathcal{V}_\exists(\sigma(x)) \wedge x \in \mathrm{V}_\exists \}$ and the *universal relation* to be $U_\sigma := \{ (y, x) \mid y \in \mathcal{V}_\forall(\sigma(x)) \wedge x \in \mathrm{V}_\exists \}$. Let R be a variable-condition. σ is an *existential R-substitution* if σ is a substitution with $\mathrm{dom}(\sigma) = \mathrm{V}_\exists$ for which $U_\sigma \circ R$ is irreflexive. σ is a *strong existential R-substitution* if σ is a substitution with $\mathrm{dom}(\sigma) = \mathrm{V}_\exists$ for which $(U_\sigma \circ R)^+$ is a wellfounded ordering.

Note that, regarding syntax, $(x^\exists, y^\forall) \in R$ is intended to mean that an existential R-substitution σ may not replace $x^\exists$ with a term in which $y^\forall$ occurs, i.e. $(y^\forall, x^\exists) \in U_\sigma$ must be disallowed, i.e. $U_\sigma \circ R$ must be irreflexive. Thus, the definition of a (weak) existential R-substitution is quite straightforward. The definition of a *strong* existential R-substitution requires an additional transitive closure because the strong version then admits a smaller R. To see this, take from Ex. 1.3 the variable-condition R of (#) and the σ of (\$). As explained there, σ must not be a strong existential R-valuation due to the cycle $y^\forall\ U_\sigma\ u^\exists\ R\ v^\forall\ U_\sigma\ x^\exists\ R\ y^\forall$ which just contradicts the irreflexivity of $U_\sigma \circ R \circ U_\sigma \circ R$. Note that in practice w.l.o.g. U_σ and R can always be chosen to be finite, so that irreflexivity of $(U_\sigma \circ R)^+$ is then equivalent to $(U_\sigma \circ R)^+$ being a wellfounded ordering.

After application of a [strong] existential R-substitution σ, in case of $(x^\exists, y^\forall) \in R$, we have to ensure that $x^\exists$ is not replaced with $y^\forall$ via a future application of another [strong] existential R-substitution that replaces a free existential variable $u^\exists$ occurring in $\sigma(x^\exists)$ with $y^\forall$. In this case, the new variable-condition has to contain $(u^\exists, y^\forall)$. This means that $E_\sigma \circ R$ must be a subset of the updated variable-condition. For the weak version this is already enough. For the strong version we have to add an arbitrary number of steps with $U_\sigma \circ R$ again.

Definition 2.2 ([Strong] σ-Update)
Let R be a variable-condition and σ be an [strong] existential R-substitution. The [*strong*] σ-*update of* R is $E_\sigma \circ R\ [\circ\ (U_\sigma \circ R)^*]$.

Example 2.3 In the proof attempt of Ex. 1.3 we applied the strong existential R-substitution $\sigma' := \{x^\exists \mapsto v^\forall\} \uplus {}_{\mathrm{V}_\exists \setminus \{x^\exists\}}|\mathrm{id}$ where $R = \{(x^\exists, y^\forall),\ (u^\exists, v^\forall)\}$. Note that $U_{\sigma'} = \{(v^\forall, x^\exists)\}$ and $E_{\sigma'} = {}_{\mathrm{V}_\exists \setminus \{x^\exists\}}|\mathrm{id}$. Thus: $E_{\sigma'} \circ R = \{(u^\exists, v^\forall)\}$; $E_{\sigma'} \circ R \circ U_{\sigma'} \circ R = \{(u^\exists, y^\forall)\}$; and $E_{\sigma'} \circ R \circ U_{\sigma'} \circ R \circ U_{\sigma'} \circ R = \emptyset$. The strong σ'-update of R is then the new variable-condition $\{(u^\exists, v^\forall),\ (u^\exists, y^\forall)\}$.

Let $\mathcal{A}$ be some Σ-structure. We now define a semantic counterpart of our existential R-substitutions, which we call "existential $(\mathcal{A}, R)$-valuation". Suppose that e maps each free existential variable not directly to an object of $\mathcal{A}$, but can additionally read the values of some free universal variables under an $\mathcal{A}$-valuation $\pi \in \mathrm{V}_\forall \to \mathcal{A}$, i.e. e gets some $\pi' \in \mathrm{V}_\forall \rightsquigarrow \mathcal{A}$ with $\pi' \subseteq \pi$ as a second argument; short: $e\colon \mathrm{V}_\exists \to ((\mathrm{V}_\forall \rightsquigarrow \mathcal{A}) \to \mathcal{A})$. Moreover, for each free existential variable x, we require the set of read free universal variables (i.e. $\mathrm{dom}(\pi')$) to be identical for all π; i.e. there has to be some "semantic relation" $S_e \subseteq \mathrm{V}_\forall \times \mathrm{V}_\exists$ such that for all $x \in \mathrm{V}_\exists$: $e(x)\colon (S_e\langle\{x\}\rangle \to \mathcal{A}) \to \mathcal{A}$. Note that, for each e, at most one semantic relation exists, namely $S_e := \{\ (y, x) \mid y \in \mathrm{dom}(\bigcup(\mathrm{dom}(e(x)))) \wedge x \in \mathrm{V}_\exists\ \}$.

Definition 2.4 (S_e, [Strong] Existential $(\mathcal{A}, R)$-Valuation, ϵ)
Let R be a variable-condition, $\mathcal{A}$ a Σ-structure, and $e\colon \mathrm{V}_\exists \to ((\mathrm{V}_\forall \rightsquigarrow \mathcal{A}) \to \mathcal{A})$. The *semantic relation of* e is $S_e := \{\ (y, x) \mid y \in \mathrm{dom}(\bigcup(\mathrm{dom}(e(x)))) \wedge x \in \mathrm{V}_\exists\ \}$. e is an *existential* $(\mathcal{A}, R)$-*valuation* if $S_e \circ R$ is irreflexive and, for all $x \in \mathrm{V}_\exists$, $e(x)\colon (S_e\langle\{x\}\rangle \to \mathcal{A}) \to \mathcal{A}$. e is a *strong existential* $(\mathcal{A}, R)$-*valuation* if $(S_e \circ R)^+$ is a wellfounded ordering and, for all $x \in \mathrm{V}_\exists$, $e(x)\colon (S_e\langle\{x\}\rangle \to \mathcal{A}) \to \mathcal{A}$.
For applying [strong] existential $(\mathcal{A}, R)$-valuations in a uniform manner, we define the function

$$\epsilon\colon (\mathrm{V}_\exists \to ((\mathrm{V}_\forall \rightsquigarrow \mathcal{A}) \to \mathcal{A})) \to ((\mathrm{V}_\forall \to \mathcal{A}) \to (\mathrm{V}_\exists \to \mathcal{A}))$$

by $(e \in \mathrm{V}_\exists \to ((\mathrm{V}_\forall \rightsquigarrow \mathcal{A}) \to \mathcal{A}),\ \pi \in \mathrm{V}_\forall \to \mathcal{A},\ x \in \mathrm{V}_\exists)$

$$\epsilon(e)(\pi)(x) := e(x)({}_{S_e\langle\{x\}\rangle}|\pi).$$

3 The Weak Version

We define R-validity of a set of sequents with free variables, in terms of validity of a formula (where the free variables are treated as nullary constants).

Definition 3.1 (Validity) Let R be a variable-condition, $\mathcal{A}$ a Σ-structure, and G a set of sequents. G is R-*valid in* $\mathcal{A}$ if there is an existential $(\mathcal{A}, R)$-valuation e such that G is $(e, \mathcal{A})$-valid. G is $(e, \mathcal{A})$-*valid* if G is $(\pi, e, \mathcal{A})$-valid for all $\pi \in \mathrm{V}_\forall \to \mathcal{A}$. G is $(\pi, e, \mathcal{A})$-*valid* if G is valid in $\mathcal{A} \uplus \epsilon(e)(\pi) \uplus \pi$. G is *valid in* $\mathcal{A}$ if Γ is valid in $\mathcal{A}$ for all $\Gamma \in G$. A sequent Γ is *valid in* $\mathcal{A}$ if there is some formula listed in Γ that is valid in $\mathcal{A}$.
Validity in a class of Σ-structures is understood as validity in each of the Σ-structures of that class. If we omit the reference to a special Σ-structure we

mean validity (or reduction, cf. below) in some fixed class K of Σ-structures, e.g. the class of all Σ-structures (Σ-algebras) or the class of Herbrand Σ-structures (term-generated Σ-algebras), cf. Wirth (1997), Wirth & Gramlich (1994) for more interesting classes for establishing inductive validities.

Example 3.2 (Validity) For $x^\exists \in \mathrm{V}_\exists$, $y^\forall \in \mathrm{V}_\forall$, the sequent $x^\exists{=}y^\forall$ is $\emptyset$-valid in any $\mathcal{A}$ because we can choose $S_e := \mathrm{V}_\forall{\times}\mathrm{V}_\exists$ and $e(x^\exists)(\pi) := \pi(y^\forall)$ resulting in $\epsilon(e)(\pi)(x^\exists) = e(x^\exists)(_{S_e}\{\!\{x^\exists\}\!\}|\pi) = e(x^\exists)(_{\mathrm{V}_\forall}|\pi) = \pi(y^\forall)$. This means that $\emptyset$-validity of $x^\exists{=}y^\forall$ is the same as validity of $\forall y\colon\ \exists x\colon\ x{=}y$. Moreover, note that $\epsilon(e)(\pi)$ has access to the π-value of $y^\forall$ just as a raising function f for x in the raised (i.e. dually Skolemized) version $f(y^\forall){=}y^\forall$ of $\forall y\colon\ \exists x\colon\ x{=}y$. Contrary to this, for $R := \mathrm{V}_\exists{\times}\mathrm{V}_\forall$, the same formula $x^\exists{=}y^\forall$ is not R-valid in general because then the required irreflexivity of $S_e{\circ}R$ implies $S_e{=}\emptyset$, and $e(x^\exists)(_{S_e}\{\!\{x^\exists\}\!\}|\pi) = e(x^\exists)(_\emptyset|\pi) = e(x^\exists)(\emptyset)$ cannot depend on $\pi(y^\forall)$ anymore. This means that $(\mathrm{V}_\exists{\times}\mathrm{V}_\forall)$-validity of $x^\exists{=}y^\forall$ is the same as validity of $\exists x\colon\ \forall y\colon\ x{=}y$. Moreover, note that $\epsilon(e)(\pi)$ has no access to the π-value of $y^\forall$ just as a raising function c for x in the raised version $c{=}y^\forall$ of $\exists x\colon\ \forall y\colon\ x{=}y$.
For a more general example let $G = \{\ A_{i,0}\ldots A_{i,n_i-1} \mid\ i{\in}I\ \}$, where for $i \in I$ and $j \prec n_i$ the $A_{i,j}$ are formulas with free existential variables from $\vec{x}$ and free universal variables from $\vec{y}$. Then $(\mathrm{V}_\exists{\times}\mathrm{V}_\forall)$-validity of G means validity of

$$\exists\vec{x}\colon\ \forall\vec{y}\colon\ \forall i{\in}I\colon\ \exists j{\prec}n_i\colon\ A_{i,j}\,;$$

whereas $\emptyset$-validity of G means validity of

$$\forall\vec{y}\colon\ \exists\vec{x}\colon\ \forall i{\in}I\colon\ \exists j{\prec}n_i\colon\ A_{i,j}\,.$$

Besides the notion of validity we need the notion of reduction. Roughly speaking, a set G_0 of sequents reduces to a set G_1 of sequents if validity of G_1 implies validity of G_0. This, however, is too weak for our purposes here because we are not only interested in validity but also in preserving the solutions for the free existential variables: For ITP, answer computation, and constraint solving it becomes important that the solutions of G_1 are also solutions of G_0.

Definition 3.3 (Reduction)
G_0 R-reduces to G_1 in $\mathcal{A}$ if for all existential $(\mathcal{A},R)$-valuations e:

$$\text{if } G_1 \text{ is } (e,\mathcal{A})\text{-valid then } G_0 \text{ is } (e,\mathcal{A})\text{-valid, too.}$$

Now we are going to abstractly describe deductive sequent and tableau calculi. We will later show that the usual deductive first-order calculi are instances of our abstract calculi. The benefit of the abstract version is that every instance is automatically sound. Due to the small number of inference rules in deductive first-order calculi and the locality of soundness, this abstract version is not really necessary. For inductive calculi, however, due to a bigger number of inference rules that usually have to be improved now and then and the globality of soundness, such an abstract version is very helpful, cf. Wirth & Becker (1995), Wirth (1997).

Definition 3.4 (Proof Forest)
A *(deductive) proof forest in a sequent* (or else: *tableau*) *calculus* is a pair (F,R) where R is a variable-condition and F is a set of pairs (Γ,t), where Γ is a sequent and t is a tree whose nodes are labeled with sequents (or else: formulas).

Note that the tree t is intended to represent a proof attempt for Γ. The nodes of t are labeled with formulas in a tableau calculus and with sequents in a sequent calculus. While the sequents at the nodes of a tree in a sequent calculus stand for themselves, in a tableau calculus all the ancestors have to be included to make up a sequent and, moreover, the formulas at the labels are in negated form:

Definition 3.5 (Goals(), $\mathcal{AX}$, **Closedness)**
'Goals(T)' denotes the set of sequents labeling the leaves of the trees in the set T (or else: the set of sequents resulting from listing the conjugates of the formulas labeling a branch from a leaf to the root in a tree in T).
We assume $\mathcal{AX}$ to be some set of *axioms*. By this we mean that $\mathcal{AX}$ is $\mathrm{V}_\exists \times \mathrm{V}_\forall$-valid. (Cf. the last sentence in Def. 3.1.)
The tree t is *closed* if $\mathrm{Goals}(\{t\}) \subseteq \mathcal{AX}$.

The readers may ask themselves why we consider a forest instead of a single tree only. The possibility to have an empty forest provides a nicer starting point. Besides that, if we have trees (Γ, t), $(\Gamma', t') \in F$ we can apply Γ as a lemma in the tree t' of Γ', provided that the lemma application relation is acyclic. For deductive theorem proving the availability of lemma application is not really necessary. For ITP, however, lemma and induction hypothesis application of this form becomes necessary.

Definition 3.6 (Invariant Condition) The *invariant condition on* (F, R) is that $\{\Gamma\}$ R-reduces to $\mathrm{Goals}(\{t\})$ for all $(\Gamma, t) \in F$.

Theorem 3.7 *Let the proof forest (F, R) satisfy the above invariant condition. Let $(\Gamma, t) \in F$. If t is closed, then Γ is R-valid.*

Theorem 3.8 *The above invariant condition is always satisfied when we start with an empty proof forest $(F, R) := (\emptyset, \emptyset)$ and then iterate only the following kinds of modifications of (F, R) (resulting in (F', R')):*

Hypothesizing: *Let R' be a variable-condition with $R \subseteq R'$. Let Γ be a sequent. Let t be the tree with a single node only, which is labeled with Γ (or else: with a single branch only, such that Γ is the list of the conjugates of the formulas labeling the branch from the leaf to the root). Then we may set $F' := F \cup \{(\Gamma, t)\}$.*

Expansion: *Let $(\Gamma, t) \in F$. Let R' be a variable-condition with $R \subseteq R'$. Let l be a leaf in t. Let Δ be the label of l (or else: result from listing the conjugates of the formulas labeling the branch from l to the root of t). Let G be a finite set of sequents. Now if $\{\Delta\}$ R'-reduces to G (or else: $\{ \Lambda\Delta \mid \Lambda \in G \}$), then we may set $F' := (F \setminus \{(\Gamma, t)\}) \cup \{(\Gamma, t')\}$ where t' results from t by adding to the former leaf l, exactly for each sequent Λ in G, a new child node labeled with Λ (or else: a new child branch such that Λ is the list of the conjugates of the formulas labeling the branch from the leaf to the new child node of l).*

Instantiation: *Let σ be an existential R-substitution. Let R' be the σ-update of R. Then we may set $F' := F\sigma$.*

While Hypothesizing and Instantiation steps are self-explanatory, Expansion steps are parameterized by a sequent Δ and a set of sequents G such that $\{\Delta\}$ R'-reduces to G. For tableau calculi, however, this set of sequents must actually have the form $\{\ \Lambda\Delta \mid \Lambda \in G\ \}$ because an Expansion step cannot remove formulas from ancestor nodes. This is because these formulas are also part of the goals associated with other leaves in the proof tree. Therefore, although tableau calculi may save repetition of formulas, sequent calculi have substantial advantages: Rewriting of formulas in place is always possible, and we can remove formulas that are redundant w.r.t. the other formulas in a sequent. But this is not our subject here. For the below examples of α-, β-, γ-, and δ-rules we will use the sequent calculi presentation because it is a little more explicit. When we write $\frac{\Delta}{\Pi_0 \ldots \Pi_{n-1}}\ R''$ we want to denote a sub-rule of the Expansion rule which is given by $G := \{\Pi_0, \ldots, \Pi_{n-1}\}$ and $R' := R \cup R''$. This means that for this rule really being a sub-rule of the Expansion rule we have to show that $\{\Delta\}$ R'-reduces to G. Moreover, note that in old times when trees grew upwards, Gerhard Gentzen would have written $\Pi_0 \ldots \Pi_{n-1}$ above the line and Δ below, such that passing the line meant implication. In our case, passing the line means reduction.

Let A and B be formulas, Γ and Π sequents, $x \in \mathrm{V}_{\mathrm{bound}}$, $x^{\exists} \in \mathrm{V}_{\exists} \setminus \mathcal{V}_{\exists}(A\Gamma\Pi)$, and $x^{\forall} \in \mathrm{V}_{\forall} \setminus \mathcal{V}_{\forall}(A\Gamma\Pi)$.

α-rules: $\dfrac{\Gamma\ \ (A \vee B)\ \ \Pi}{A\ \ B\ \ \Gamma\ \ \Pi}\ \emptyset$ $\qquad$ $\dfrac{\Gamma\ \ \neg(A \wedge B)\ \ \Pi}{\overline{A}\ \ \overline{B}\ \ \Gamma\ \ \Pi}\ \emptyset$ $\qquad$ $\dfrac{\Gamma\ \ \neg\neg A\ \ \Pi}{A\ \ \Gamma\ \ \Pi}\ \emptyset$

β-rules: $\dfrac{\Gamma\ \ (A \wedge B)\ \ \Pi}{A\ \ \Gamma\ \ \Pi \qquad\qquad B\ \ \Gamma\ \ \Pi}\ \emptyset$ $\qquad$ $\dfrac{\Gamma\ \ \neg(A \vee B)\ \ \Pi}{\overline{A}\ \ \Gamma\ \ \Pi \qquad\qquad \overline{B}\ \ \Gamma\ \ \Pi}\ \emptyset$

γ-rules: $\dfrac{\Gamma\ \ \exists x\colon A\ \ \Pi}{A\{x \mapsto x^{\exists}\}\ \ \Gamma\ \ \exists x\colon A\ \ \Pi}\ \emptyset$ $\qquad$ $\dfrac{\Gamma\ \ \neg\forall x\colon A\ \ \Pi}{\overline{A\{x \mapsto x^{\exists}\}}\ \ \Gamma\ \ \neg\forall x\colon A\ \ \Pi}\ \emptyset$

δ-rules: $\dfrac{\Gamma\ \ \forall x\colon A\ \ \Pi}{A\{x \mapsto x^{\forall}\}\ \ \Gamma\ \ \Pi}\ \mathcal{V}_{\exists}(A\Gamma\Pi) \times \{x^{\forall}\}$ $\qquad$ $\dfrac{\Gamma\ \ \neg\exists x\colon A\ \ \Pi}{\overline{A\{x \mapsto x^{\forall}\}}\ \ \Gamma\ \ \Pi}\ \mathcal{V}_{\exists}(A\Gamma\Pi) \times \{x^{\forall}\}$

Theorem 3.9 *The above examples of α-, β-, γ-, and δ-rules are all sub-rules of the Expansion rule of the sequent calculus of Th. 3.8.*

4 The Strong Version

The additional solutions (or existential substitutions) of the strong version (which admit additional proofs compared to the weak version) do not add much difficulty when one is interested in validity only, cf. e.g. Hähnle & Schmitt (1994). When also the preservation of solutions is required, however, the additional substitutions pose some problems because the new solutions may tear some free universal variables out of their contexts:

Example 4.1 In Ex. 1.2 a liberalized δ-step reduced $\mathsf{P}(x^{\exists}),\ \forall y\colon \neg\mathsf{P}(y)$ to $\mathsf{P}(x^{\exists}),\ \neg\mathsf{P}(y^{\forall})$ with empty variable-condition $R{:=}\emptyset$. The latter sequent is $(e, \mathcal{A})$-valid for the strong existential $(\mathcal{A}, R)$-valuation e given by $e(x^{\exists})(\pi) := \pi(y^{\forall})$. The former sequent, however, is not $(e, \mathcal{A})$-valid when $\mathsf{P}^{\mathcal{A}}(a)$ is true and $\mathsf{P}^{\mathcal{A}}(b)$ is false for some a, b from the universe of $\mathcal{A}$. To wit, take some π with $\pi(y^{\forall}) := b$.

How can we solve the problem exhibited in Ex. 4.1? I.e. how can we change the notion of reduction such that the liberalized δ-step becomes a reduction step?

Definition 4.2 (Choice-Condition and Compatibility)
C is a $(R,<)$-*choice-condition* if C is a (possibly) partial function from $\mathrm{V}_\forall$ to formulas, R is a variable-condition, $<$ is a wellfounded ordering on $\mathrm{V}_\forall$ with $(R \circ <) \subseteq R$, and, for all $y^\forall \in \mathrm{dom}(C)$: $z^\forall < y^\forall$ for all $z^\forall \in \mathcal{V}_\forall(C(y^\forall))\backslash\{y^\forall\}$ and $u^\exists \; R \; y^\forall$ for all $u^\exists \in \mathcal{V}_\exists(C(y^\forall))$.
Let C be a $(R,<)$-choice-condition, $\mathcal{A}$ a Σ-structure, and e a strong existential $(\mathcal{A},R)$-valuation. We say that π is $(e,\mathcal{A})$-*compatible with* C if $\pi \in \mathrm{V}_\forall \to \mathcal{A}$ and for each $y^\forall \in \mathrm{dom}(C)$: If $C(y^\forall)$ is $(\pi,e,\mathcal{A})$-valid, then $C(y^\forall)$ is $({}_{\mathrm{V}_\forall\backslash\{y^\forall\}}|\pi \uplus \eta, e, \mathcal{A})$-valid for all $\eta \in \{y^\forall\} \to \mathcal{A}$.

Note that $(e,\mathcal{A})$-compatibility of π with $\{(y^\forall, B)\}$ means that a different choice for the π-value of $y^\forall$ does not destroy the validity of the formula B in $\mathcal{A}\uplus\epsilon(e)(\pi)\uplus\pi$, or that $\pi(y^\forall)$ is chosen such that B becomes invalid if such a choice is possible, which is closely related to Hilbert's ε-operator ($y^\forall = \varepsilon y\colon(\neg B\{y^\forall \mapsto y\})$). Moreover, note that $\emptyset$ is a $(R,\emptyset)$-choice-condition for any variable-condition R.

Definition 4.3 (Extended Strong σ-Update) Let C be a $(R,<)$-choice-condition and σ a strong existential R-substitution. The *extended strong* σ*-update* $(C',R',<')$ *of* $(C,R,<)$ is given by $C' := \{\, (x, B\sigma) \mid (x,B) \in C \,\}$, R' is the strong σ-update of R, $<' \;:=\; < \circ (U_\sigma \circ R)^* \;\cup\; (U_\sigma \circ R)^+$.

We are now going to proceed like in the previous section, but using the strong versions instead of the weak ones. The tableau version is omitted because it differs from the sequent version just like before.

Definition 4.4 (Strong Validity) Let C be a $(R,<)$-choice-condition, $\mathcal{A}$ a Σ-structure, and G a set of sequents. G is C*-strongly* R*-valid in* $\mathcal{A}$ if there is a strong existential $(\mathcal{A},R)$-valuation e such that G is C-strongly $(e,\mathcal{A})$-valid. G is C*-strongly* $(e,\mathcal{A})$*-valid* if G is $(\pi,e,\mathcal{A})$-valid for each π that is $(e,\mathcal{A})$-compatible with C. The rest is given by Def. 3.1.

Example 4.5 (Strong Validity) Note that $\emptyset$-validity does not differ from $\emptyset$-strong $\emptyset$-validity and that $\mathrm{V}_\exists \times \mathrm{V}_\forall$-validity does not differ from $\emptyset$-strong $\mathrm{V}_\exists \times \mathrm{V}_\forall$-validity. This is because the notions of weak and strong existential valuations do not differ in these cases. Therefore, Ex. 3.2 is also an example for strong validity. Although $\emptyset$-strong R-validity always implies (weak) R-validity (because each strong existential $(\mathcal{A},R)$-valuation is a (weak) existential $(\mathcal{A},R)$-valuation), for R not being one of the extremes $\emptyset$ and $\mathrm{V}_\exists \times \mathrm{V}_\forall$, (weak) R-validity and $\emptyset$-strong R-validity differ from each other. E.g. the sequent $(*)$ in Ex. 1.3 is (weakly) R-valid but not $\emptyset$-strongly R-valid for the R of $(\#)$: For $S_e := \{(y^\forall, u^\exists),\ (v^\forall, x^\exists)\}$ we get $S_e \circ R = \{(y^\forall, v^\forall),\ (v^\forall, y^\forall)\}$, which is irreflexive. Since the sequent $(*)$ is $(e,\mathcal{A})$-valid for the (weak) $e$xistential $(\mathcal{A},R)$-valuation e given by $e(x^\exists)({}_{S_e\{\{x^\exists\}\}}|\pi) = \pi(v^\forall)$ and $e(u^\exists)({}_{S_e\{\{u^\exists\}\}}|\pi) = \pi(y^\forall)$, the sequent $(*)$ is (weakly) R-valid in $\mathcal{A}$. But $S_e \circ R \circ S_e \circ R$ is not irreflexive, so that this e is no *strong* existential $(\mathcal{A},R)$-valuation, which means that the sequent $(*)$ cannot be $\emptyset$-strongly R-valid in general.

Definition 4.6 (Strong Reduction)
Let C be a $(R,<)$-choice-condition, $\mathcal{A}$ a Σ-structure, and G_0, G_1 sets of sequents. G_0 *strongly* (R,C)*-reduces to* G_1 *in* $\mathcal{A}$ if for each strong existential $(\mathcal{A},R)$-valuation e and each π that is $(e,\mathcal{A})$-compatible with C:

if G_1 is $(\pi,e,\mathcal{A})$-valid, then G_0 is $(\pi,e,\mathcal{A})$-valid.

Definition 4.7 (Strong Proof Forest)
A *strong proof forest in a sequent calculus* is a quadruple $(F,C,R,<)$ where C is a $(R,<)$-choice-condition and F is a set of pairs (Γ,t), where Γ is a sequent and t is a tree whose nodes are labeled with sequents.

The notions of Goals(), $\mathcal{AX}$, and closedness of Def.3.5 are not changed. Note, however, that the $\mathrm{V}_\exists\times\mathrm{V}_\forall$-validity of $\mathcal{AX}$ immediately implies the $\emptyset$-strong $\mathrm{V}_\exists\times\mathrm{V}_\forall$-validity of $\mathcal{AX}$, i.e. the logically strongest kind of C-strong R-validity.

Definition 4.8 (Strong Invariant Condition)
The *strong invariant condition on* $(F,C,R,<)$ is that $\{\Gamma\}$ strongly (R,C)-reduces to Goals$(\{t\})$ for all $(\Gamma,t)\in F$.

Theorem 4.9 *Let the strong proof forest* $(F,C,R,<)$ *satisfy the above strong invariant condition. Let* $(\Gamma,t)\in F$ *and* t *be closed. Now:* Γ *is* C*-strongly* R*-valid and, for any injective* $\varsigma\in(\mathcal{V}_\forall(\Gamma)\cap\mathrm{dom}(C))\to(\mathrm{V}_\exists\backslash\mathcal{V}_\exists(\Gamma))$, $\Gamma\varsigma$ *is* $\emptyset$*-strongly* $_{\mathrm{V}_\exists\backslash\mathrm{ran}(\varsigma)}|R$*-valid and even* $\emptyset$*-strongly* R'*-valid for*
$R' := {}_{\mathrm{V}_\exists\backslash\mathrm{ran}(\varsigma)}|R \cup \bigcup_{y\in\mathrm{ran}(\varsigma)}\{y\}\times\langle\{\varsigma^{-1}(y)\}\rangle_< \cup \mathrm{V}_\exists\times\mathrm{dom}(C)$.

Theorem 4.10
The above strong invariant condition is always satisfied when we start with an empty strong proof forest $(F,C,R,<) := (\emptyset,\emptyset,\emptyset,\emptyset)$ *and then iterate only the following kinds of modifications of* $(F,C,R,<)$ *(resulting in* $(F',C',R',<')$*):*

Hypothesizing: *Let* $R' := R\cup R''$ *be a variable-condition with* $(R''\circ<)\subseteq R'$. *Set* $C' := C$ *and* $<' := <$. *Let* Γ *be a sequent. Let* t *be the tree with a single node only, which is labeled with* Γ. *Then we may set* $F' := F\cup\{(\Gamma,t)\}$.

Expansion: *Let* C' *be a* $(R',<')$*-choice-condition with* $C\subseteq C'$ *and* $R\subseteq R'$. *Let* $(\Gamma,t)\in F$. *Let* l *be a leaf in* t. *Let* Δ *be the label of* l. *Let* G *be a finite set of sequents. Now if* $\{\Delta\}$ *strongly* (R',C')*-reduces to* G, *then we may set* $F' := (F\backslash\{(\Gamma,t)\})\cup\{(\Gamma,t')\}$ *where* t' *results from* t *by adding to the former leaf* l, *exactly for each sequent* Λ *in* G, *a new child node labeled with* Λ.

Instantiation: *Let* σ *be a strong existential* R*-substitution. Let* $(C',R',<')$ *be the extended strong* σ*-update of* $(C,R,<)$. *Then we may set* $F' := F\sigma$.

When we write $\dfrac{\Delta}{\Pi_0\ldots\Pi_{n-1}}\ \begin{matrix}C''\\R''\\<''\end{matrix}$ we want to denote a sub-rule of the Expansion rule which is given by $G := \{\Pi_0,\ldots,\Pi_{n-1}\}$, $C' := C\cup C''$, $R' := R\cup R''$, and $<' := <\cup<''$. This means that for this rule really being a sub-rule of the Expansion rule we have to show that C' is a $(R',<')$-choice-condition and that $\{\Delta\}$ strongly (R',C')-reduces to G.

Let A and B be formulas, Γ and Π sequents, $x \in \mathrm{V}_{\mathrm{bound}}$, $x^{\exists} \in \mathrm{V}_{\exists} \backslash \mathcal{V}_{\exists}(A\Gamma\Pi)$, and $x^{\forall} \in \mathrm{V}_{\forall} \setminus (\mathcal{V}_{\forall}(A\Gamma\Pi) \cup \mathrm{dom}(<) \cup \mathrm{dom}(C))$. The α-, β-, and γ-rules differ from those of the weak version only in a triple instead of single $\emptyset$-annotation. Liberalized δ-rules:

$$\frac{\Gamma \quad \forall x\colon A \quad \Pi}{A\{x \mapsto x^{\forall}\} \quad \Gamma \quad \Pi} \begin{array}{l} \{(x^{\forall},\, A\{x \mapsto x^{\forall}\})\} \\ (\mathcal{V}_{\exists}(A) \cup R\langle \mathcal{V}_{\forall}(A)\rangle) \times \{x^{\forall}\} \\ \leq \langle \mathcal{V}_{\forall}(A)\rangle \times \{x^{\forall}\} \end{array} \qquad \frac{\Gamma \quad \neg\exists x\colon A \quad \Pi}{\overline{A\{x \mapsto x^{\forall}\}} \quad \Gamma \quad \Pi} \begin{array}{l} \{(x^{\forall},\, \overline{A\{x \mapsto x^{\forall}\}})\} \\ (\mathcal{V}_{\exists}(A) \cup R\langle \mathcal{V}_{\forall}(A)\rangle) \times \{x^{\forall}\} \\ \leq \langle \mathcal{V}_{\forall}(A)\rangle \times \{x^{\forall}\} \end{array}$$

Theorem 4.11 *The above examples of α-, β-, γ-, and liberalized δ-rules are all sub-rules of the Expansion rule of the sequent calculus of Th. 4.10.*

The following example shows that R'' of the above liberalized δ-rule must indeed contain $R\langle \mathcal{V}_{\forall}(A)\rangle \times \{x^{\forall}\}$.

Example 4.12 $\exists y\colon \forall x\colon (\neg \mathsf{Q}(x,y) \vee \forall z\colon \mathsf{Q}(x,z))$ is not deductively valid (to wit, let Q be the identity relation on a non-trivial universe). γ-step: $\forall x\colon (\neg \mathsf{Q}(x,y^{\exists}) \vee \forall z\colon \mathsf{Q}(x,z))$. Liberalized δ-step: $(\neg \mathsf{Q}(x^{\forall},y^{\exists}) \vee \forall z\colon \mathsf{Q}(x^{\forall},z))$ with choice-condition $(x^{\forall}, (\neg \mathsf{Q}(x^{\forall},y^{\exists}) \vee \forall z\colon \mathsf{Q}(x^{\forall},z)))$ and variable-condition $(y^{\exists}, x^{\forall})$. α-step: $\neg \mathsf{Q}(x^{\forall},y^{\exists})$, $\forall z\colon \mathsf{Q}(x^{\forall},z)$. Liberalized δ-step: $\neg \mathsf{Q}(x^{\forall},y^{\exists})$, $\mathsf{Q}(x^{\forall},z^{\forall})$ with additional choice-condition $(z^{\forall}, \mathsf{Q}(x^{\forall},z^{\forall}))$ and additional variable-condition $(y^{\exists}, z^{\forall})$. Note that the additional variable-condition arises although $y^{\exists}$ does not appear in $\mathsf{Q}(x^{\forall},z)$. The reason for the additional variable-condition is $y^{\exists}\ R\ x^{\forall} \in \mathcal{V}_{\forall}(\mathsf{Q}(x^{\forall},z))$. The variable-condition $(y^{\exists}, z^{\forall})$ is, however, essential for soundness, because without it we could complete the proof attempt by application of the strong existential $\{(y^{\exists}, x^{\forall})\}$-substitution $\sigma := \{y^{\exists} \mapsto z^{\forall}\} \uplus {}_{\mathrm{V}_{\exists} \backslash \{y^{\exists}\}}|\mathrm{id}$.

Another interesting point is that now that we have achieved our goal of liberalizing our δ-rule and strictly increasing our proving possibilities, we must not use our original non-liberalized δ-rule of section 3 anymore. This sounds quite strange on the first view, but is simply due to our changed notion of reduction: There is a fundamental difference related to the occurrence of the universal quantification on π between the notion of (weak) reduction

$$\ldots \big(\forall \pi\colon\ G_1\ (\pi, e, \mathcal{A})\text{-valid}\big) \Rightarrow \big(\forall \pi\colon\ G_0\ (\pi, e, \mathcal{A})\text{-valid}\big) \ldots$$

and the notion of strong reduction

$$\ldots \forall \pi\colon\ \big(G_1\ (\pi, e, \mathcal{A})\text{-valid} \Rightarrow G_0\ (\pi, e, \mathcal{A})\text{-valid}\big) \ldots .$$

This difference in the nature of reduction renders the weak version applicable in areas where the strong version is not. For this reason (and for the sake of stepwise presentation) we have included the weak version in this paper although the strong version will turn out to be superior in all aspects of the calculus of Th. 4.11 treated in this paper. For a more detailed discussion of this difference in the nature of reduction and its essentiality, cf. Wirth (1998).

Moreover, note that (as far as Th. 4.11 is concerned) the choice-conditions do not have any influence on our proofs and may be discarded. We could, however, use them for the following purposes:

1. We could use the choice-conditions in order to weaken our requirements for our set of axioms $\mathcal{AX}$: Instead of $\emptyset$-strong $\mathrm{V}_{\exists} \times \mathrm{V}_{\forall}$-validity of $\mathcal{AX}$ the weaker C-strong $\mathrm{V}_{\exists} \times \mathrm{V}_{\forall}$-validity of $\mathcal{AX}$ is sufficient for Th. 4.9.

2. If we add a functional behavior to a choice-condition C, i.e. if we require that for $(x^\forall, A) \in C$ the value for $x^\forall$ is not just an arbitrary one from the set of values that make A invalid, but a unique element of this set given by some choice-function, then we can use the choice-conditions for simulating the behavior of the δ^{+^+}-rule of Beckert & al. (1993) by using the same free universal variable for the same C-value and by later equating free universal variables whose C-values become equal during the proof.
3. Moreover, the choice-conditions may be used to get more interesting answers:

   Example 4.13 Starting with the empty proof tree and hypothesizing
   $$\forall x\colon \mathsf{Q}(x,x),\ \exists y\colon \big(\neg\mathsf{Q}(y,y) \wedge \neg\mathsf{P}(y)\big),\ \mathsf{P}(z^\exists)$$
   with the above rules we can produce a proof tree with the leaves
   $$\neg\mathsf{Q}(y^\exists, y^\exists),\ \mathsf{Q}(x^\forall, x^\forall),\ \exists y\colon \big(\neg\mathsf{Q}(y,y) \wedge \neg\mathsf{P}(y)\big),\ \mathsf{P}(z^\exists)$$
   and
   $$\neg\mathsf{P}(y^\exists),\ \mathsf{Q}(x^\forall, x^\forall),\ \exists y\colon \big(\neg\mathsf{Q}(y,y) \wedge \neg\mathsf{P}(y)\big),\ \mathsf{P}(z^\exists)$$
   and the $(\emptyset,\emptyset)$-choice-condition $\{(x^\forall, \mathsf{Q}(x^\forall, x^\forall))\}$.
   The strong existential $\emptyset$-substitution $\{y^\exists \mapsto x^\forall,\ z^\exists \mapsto x^\forall\} \uplus {}_{\mathrm{V}_\exists \setminus \{y^\exists, z^\exists\}}|\mathrm{id}$ closes the proof tree via an Instantiation step. The answer $x^\forall$ for our query variable $z^\exists$ is not very interesting unless we note that the choice-condition tells us to choose $x^\forall$ in such a way that $\mathsf{Q}(x^\forall, x^\forall)$ becomes false.
   The rules of our weak version of section 3 are not only unable to provide any information on free universal variables, but also unable to prove the hypothesized sequent, because they can only show
   $$\forall x\colon \mathsf{Q}(x,x),\ \exists y\colon \big(\neg\mathsf{Q}(y,y) \wedge \neg\mathsf{P}(y)\big),\ \exists z\colon \mathsf{P}(z)$$
   instead.

   Thus it is obvious that the calculus of Th. 4.11 is not only superior to the calculus of Th. 3.9 w.r.t. proving but also w.r.t. answer "computation".

Finally, note that (w.r.t. the calculus of Th. 4.11) the ordering $<$ is not needed at all when in the liberalized δ-steps we always choose a completely new free universal variable $x^\forall$ that does not occur elsewhere and when in the Hypothesizing steps we guarantee that $\mathrm{ran}(R'')$ contains only new free universal variables that have not occurred before. The former is reasonable anyhow, because the free universal variables introduced by previous liberalized δ-steps cannot be used because they are in $\mathrm{dom}(C)$ and the use of a free universal variable from the input hypothesis deteriorates the result of our proof by giving this free universal variable an existential meaning (because it puts it into $\mathrm{dom}(C)$) as explained in Th. 4.9. The latter does not seem to be restrictive for any reasonable application.

All in all, when interested in proving only, the (compared to the weak version) additional choice-condition and ordering of the strong version do not produce any overhead because they can simply be omitted. This is interesting because choice-conditions or Hilbert's ε-expressions are sometimes considered to make proofs quite complicated. When interested in answer "computation", however, they could turn out to be useful.

W.r.t. the calculus of Th. 4.11 we thus may conclude that the strong version is generally better than the weak version and the only overhead seems to be that we have to compute transitive closures when checking whether a substitution σ

is really a strong existential R-substitution and when computing the strong σ-update of R. But we actually do not have to compute the transitive closure at all, because the only essential thing is the circularity-check which can be done on a bipartite graph generating the transitive closures. This checking is in the worst case linear in $|R| + \sum_\sigma \big(|U_\sigma| + |E_\sigma| \big)$ and is expected to perform at least as well as an optimally integrated version (i.e. one without conversion of term-representation) of the linear unification algorithm of Paterson & Wegman (1978) in the standard framework of Skolemization and unification. Note, however, that the checking for strong existential R-substitutions can also be implemented with any other unification algorithm.

Not really computing the transitive closure enables another refinement that allows us to go even beyond the fascinating *strong Skolemization* of Nonnengart (1996), whose basic idea can be translated into our framework in the following simplified way. Instead of proving $\forall x\colon (A \vee B)$ it may be advantageous to prove the stronger $\forall x\colon A \vee \forall x\colon B$, because after applications of α- and liberalized δ-rules to $\forall x\colon A \vee \forall x\colon B$, resulting in $A\{x \mapsto x^{\forall}_A\}$, $B\{x \mapsto x^{\forall}_B\}$, the variable-conditions introduced for $x^{\forall}_A$ and $x^{\forall}_B$ may be smaller than the variable-condition introduced for $y^{\forall}$ after applying these rules to $\forall x\colon (A \vee B)$, resulting in $A\{x \mapsto y^{\forall}\}$, $B\{x \mapsto y^{\forall}\}$, i.e. $R\langle\{x^{\forall}_A\}\rangle$ and $R\langle\{x^{\forall}_B\}\rangle$ may be *proper* subsets of $R\langle\{y^{\forall}\}\rangle$. Therefore the proof of $\forall x\colon A \vee \forall x\colon B$ may be simpler than the proof of $\forall x\colon (A \vee B)$. The nice aspect of Nonnengart (1996) is that the proofs of $\forall x\colon A$ and $\forall x\colon (A \vee B)$ can be done in parallel without extra costs, such that the bigger variable-condition becomes active only if we decide that the smaller variable-condition is not enough to prove $\forall x\colon A$ and we had better prove the weaker $\forall x\colon (A \vee B)$. The disadvantage of the strong Skolemization approach, however, is that we have to decide whether to prove either $\forall x\colon A$ or else $\forall x\colon B$ in parallel to $\forall x\colon (A \vee B)$. In terms of Hilbert's ε-operator, this asymmetry can be understood from the argumentation of Nonnengart (1996), which, for some new variable $z \in \mathrm{V}_{\mathrm{bound}}$ and t denoting the term $\varepsilon z\colon (\neg A\{x \mapsto z\} \wedge (A \vee x{=}z))$, employs the logical equivalence of $\forall x\colon (A \vee B)$ with $\forall x\colon A \vee \forall x\colon (B\{x \mapsto t\})$ and of $\forall x\colon A$ with $\exists x\colon (A\{x \mapsto t\})$. Now, if we do not really compute the transitive closures in our strong version, we can prove $A\{x \mapsto x^{\forall}_A\}$, $B\{x \mapsto x^{\forall}_B\}$ in parallel and may later decide to prove the stronger $A\{x \mapsto y^{\forall}\}$, $B\{x \mapsto y^{\forall}\}$ instead, simply by merging the nodes for $x^{\forall}_A$ and $x^{\forall}_B$ and substituting $x^{\forall}_A$ and $x^{\forall}_B$ with $y^{\forall}$.

5 Conclusion

All in all, we have presented a combination of raising, explicit variable dependency representation, the liberalized δ-rule, and preservation of solutions for first-order deductive theorem proving. Our main motivation was to provide the foundation for our work on inductive theorem proving (cf. Wirth (1999)) where the preservation of solutions is indispensable.

To our knowledge we have presented on the one hand the first sound combination of explicit variable dependency representation and the liberalized δ-rule and on the other hand the first framework for preservation of solutions in full first-order logic.

References

Matthias Baaz, Christian G. Fermüller (1995). *Non-elementary Speedups between Different Versions of Tableaux.* 4th TABLEAUX 1995, LNAI 918, pp. 217-230, Springer.

Peter Baumgartner, Ulrich Furbach, Frieder Stolzenburg (1997). *Computing Answers with Model Elimination.* Artificial Intelligence **90**, pp. 135-176.

Bernhard Beckert, Reiner Hähnle, Peter H. Schmitt (1993). *The Even More Liberalized δ-Rule in Free Variable Semantic Tableaux.* Kurt Gödel Colloquium, LNCS 713, pp. 108-119, Springer.

Wolfgang Bibel (1987). *Automated Theorem Proving.* 2nd rev.ed., Vieweg, Braunschweig.

Wolfgang Bibel, Peter H. Schmitt (eds.) (1998). *Automated Deduction — A Basis for Applications.* Kluwer Acad. Publ..

Melvin Fitting (1996). *First-Order Logic and Automated Theorem Proving.* 2nd extd. ed., Springer.

Gerhard Gentzen (1935). *Untersuchungen über das logische Schließen.* Mathematische Zeitschrift **39**, pp. 176-210, 405-431.

Reiner Hähnle, Peter H. Schmitt (1994). *The Liberalized δ-Rule in Free Variable Semantic Tableaux.* J. Automated Reasoning **13**, pp. 211-221, Kluwer Acad. Publ..

S. Kanger (1963). *A Simplified Proof Method for Elementary Logic.* In: Siekmann & Wrightson (1983), Vol. 1, pp. 364-371.

Michaël Kohlhase (1995). *Higher-Order Tableaux.* 4th TABLEAUX 1995, LNAI 918, pp. 294-309, Springer. Revised version is: Kohlhase (1998).

Michaël Kohlhase (1998). *Higher-Order Automated Theorem Proving.* In: Bibel & Schmitt (1998), Vol. 1, pp. 431-462.

Dale Miller (1992). *Unification under a Mixed Prefix.* J. Symbolic Computation **14**, pp. 321-358.

Andreas Nonnengart (1996). *Strong Skolemization.* MPI–I–96–2–010, Max-Planck-Institut für Informatik, Saarbrücken.

Michael S. Paterson, Mark N. Wegman (1978). *Linear Unification.* J. Computer and System Sci. **16**, pp. 158-167, Academic Press.

Dag Prawitz (1960). *An Improved Proof Procedure.* In: Siekmann & Wrightson (1983), Vol. 1, pp. 159-199.

Jörg Siekmann, G. Wrightson (eds.) (1983). *Automation of Reasoning.* Springer.

Raymond M. Smullyan (1968). *First-Order Logic.* Springer.

Claus-Peter Wirth (1997). *Positive/Negative-Conditional Equations: A Constructor-Based Framework for Specification and Inductive Theorem Proving.* Dissertation (Ph.D. thesis), Verlag Dr. Kovač, Hamburg.

Claus-Peter Wirth (1998). *Full First-Order Sequent and Tableau Calculi With Preservation of Solutions and the Liberalized δ-Rule but Without Skolemization.* Report 698/1998, FB Informatik, Univ. Dortmund. `http://LS5.cs.uni-dortmund.de/~wirth/publications/gr698/all.ps.gz`.

Claus-Peter Wirth (1999). *Full First-Order Free Variable Sequents and Tableaux in Implicit Induction.* 8th TABLEAUX 1999, LNAI, Springer, to appear.

Claus-Peter Wirth, Klaus Becker (1995). *Abstract Notions and Inference Systems for Proofs by Mathematical Induction.* 4th CTRS 1994, LNCS 968, pp. 353-373, Springer.

Claus-Peter Wirth, Bernhard Gramlich (1994). *On Notions of Inductive Validity for First-Order Equational Clauses.* 12th CADE 1994, LNAI 814, pp. 162-176, Springer.

Acknowledgements: I would like to thank Ulrich Furbach and his whole group for all they taught me on tableau calculi. Furthermore, I would like to thank Michaël Kohlhase for his encouragement to drop Skolemization, and Peter Padawitz for the possibility to finish this work. Finally, I am indebted to an anonymous referee for his careful reading of an earlier version of this paper and his most useful remarks.

Editorial Remark: Wirth (1998) is a long version of this paper including all proofs.