

REDUCING HIGHER-ORDER THEOREM PROVING TO A SEQUENCE OF SAT PROBLEMS

CHAD E. BROWN

ABSTRACT. We describe a complete theorem proving procedure for higher-order logic that uses SAT-solving to do much of the heavy lifting. The theoretical basis for the procedure is a complete, cut-free, ground refutation calculus that incorporates a restriction on instantiations. The refined nature of the calculus makes it conceivable that one can search in the ground calculus itself, obtaining a complete procedure without resorting to meta-variables and a higher-order lifting lemma. Once one commits to searching in a ground calculus, a natural next step is to consider ground formulas as propositional literals and the rules of the calculus as propositional clauses relating the literals. With this view in mind, we describe a theorem proving procedure that primarily generates relevant formulas along with their corresponding propositional clauses. The procedure terminates when the set of propositional clauses is unsatisfiable. We prove soundness and completeness of the procedure. The procedure has been implemented in a new higher-order theorem prover, Satallax, which makes use of the SAT-solver MiniSat. We also describe the implementation and give several examples. Finally, we include experimental results of Satallax on the higher-order part of the TPTP library.

1. INTRODUCTION

There are a number of distinct aspects of automated theorem proving. First, there is the usual combinatorial explosion already associated with search in the propositional case. Second, there is the problem of finding the correct instantiations for quantifiers. The instantiation problem appears in the first-order case. A third issue that appears in the higher-order case is how one builds in certain basic mathematical properties (e.g., extensionality and choice).

In this paper we give a complete theorem proving procedure for higher-order logic with extensionality and choice. The procedure separates the first issue from the second and third. We start from a

Date: May 2012, Technical Report version of an article submitted to Journal of Automated Reasoning.

complete ground calculus which already builds in extensionality and choice as well as certain restrictions on instantiations. Given a set of formulas to refute, the ground calculus can be used to suggest a sequence of relevant formulas which may be involved in a refutation. The procedure generates propositional clauses corresponding to the meaning of these relevant formulas. When the set of propositional clauses is unsatisfiable (in the propositional sense), then the original set of higher-order formulas is unsatisfiable (in the higher-order Henkin model sense). Conversely, when the original set of higher-order formulas is unsatisfiable, then an unsatisfiable set of propositional clauses will eventually be generated.

Such a procedure has been implemented in the higher-order theorem prover Satallax.¹ The first implementation of Satallax was in Steel Bank Common Lisp. The latest version of Satallax, Satallax 2.4, is implemented in Objective Caml. The SAT-solver MiniSat [10] (coded in C++) is used to determine propositional unsatisfiability.

This paper is a revised and expanded version of [7].

2. PRELIMINARIES

We begin with a brief presentation of Church’s simple type theory with a choice operator. For more details see a similar presentation in [3]. Simple types (σ, τ) are given inductively: $o|\iota|\sigma\sigma$. Types $\sigma\tau$ correspond to functions from σ to τ . Terms s, t are generated inductively $x|c|st|\lambda x.s$ where x ranges over variables and c ranges over the logical constants $\perp, \rightarrow, \forall_\sigma, =_\sigma, *$ and ε_σ . A name is either a variable or a logical constant. A decomposable name is either a variable or ε_σ for some σ . We use δ to range over decomposable names.

Each variable has a corresponding type σ , and for each type there is a countably infinite set of variables of this type. Likewise each logical constant has a corresponding type: $\perp : o, \rightarrow : ooo, \forall_\sigma : (\sigma o)o, =_\sigma : \sigma\sigma o, * : \iota$ and $\varepsilon_\sigma : (\sigma o)\sigma$. The constant ε_σ is a choice operator at type σ . The constant $*$ plays the role of a “default” element of the nonempty type ι . Types can be assigned to (some) terms in the usual way. From now on we restrict ourselves to typed terms and let Λ_σ be the set of terms of type σ . A *formula* is a term $s \in \Lambda_o$.

We adopt common notational conventions: stu means $(st)u$, $s =_\sigma t$ (or $s = t$) means $=_\sigma st$, $s \rightarrow t$ means $\rightarrow st$, $\neg s$ means $s \rightarrow \perp$, \top means $\neg \perp$, $s \neq_\sigma t$ (or $s \neq t$) means $\neg(s =_\sigma t)$, $\forall x.s$ means $\forall_\sigma \lambda x.s$ and $\varepsilon x.s$ means $\varepsilon_\sigma \lambda x.s$. Binders have as large a scope as is consistent with given

¹Satallax is available at satallax.com.

parenthesis. For example, in $\forall x.px \rightarrow qx$ the occurrence of x in qx is bound by the \forall . The set $\mathcal{V}t$ of *free variables of t* is defined as usual.

An *accessibility context* (\mathcal{C}) is a term with a hole $\llbracket \sigma$ of the form $\llbracket s_1 \cdots s_n, \neg(\llbracket s_1 \cdots s_n), (\llbracket s_1 \cdots s_n) \neq_\iota s$ or $s \neq_\iota (\llbracket s_1 \cdots s_n)$. We write $\mathcal{C}[s]$ for the term one obtains by putting s into the hole. A term s is *accessible* in a set A of formulas iff there is an accessibility context \mathcal{C} such that $\mathcal{C}[s] \in A$.

Let $[s]$ denote a $\beta\eta$ -normal form of s that makes a canonical choice of bound variables. That is, for any $s, t \in \Lambda_\sigma$, $[s] = [t]$ iff s and t are $\alpha\beta\eta$ -equivalent. (In the implementation, de Bruijn indices are used.) A term s is *normal* if $[s] = s$.

A *substitution* is a type preserving partial function from variables to terms. If θ is a substitution, x is a variable, and s is a term that has the same type as x , we write θ_s^x for the substitution that agrees everywhere with θ except $\theta_s^x x = s$. For each substitution θ let $\hat{\theta}$ be the usual extension of θ to all terms in a capture-avoiding manner.

A *frame* \mathcal{D} is a typed collection of nonempty sets such that $\mathcal{D}_o = \{0, 1\}$ and $\mathcal{D}_{\sigma\tau}$ is a set of total functions from \mathcal{D}_σ to \mathcal{D}_τ . An *assignment* \mathcal{I} into \mathcal{D} is a mapping from variables and logical constants of type σ into \mathcal{D}_σ . An assignment \mathcal{I} is *logical* if it interprets each logical constant to be an element satisfying the corresponding logical property. For example, if \mathcal{I} is logical, then $\mathcal{I}\perp = 0$. An assignment \mathcal{I} is an *interpretation* if it can be extended in the usual way to be a total function $\hat{\mathcal{I}}$ mapping each Λ_σ into \mathcal{D}_σ . A *Henkin model* $(\mathcal{D}, \mathcal{I})$ is a frame \mathcal{D} and a logical interpretation \mathcal{I} into \mathcal{D} . We say a formula s is *satisfied* by a Henkin model $(\mathcal{D}, \mathcal{I})$ if $\hat{\mathcal{I}}s = 1$. A set A of formulas is satisfied by a Henkin model if each formula in A is satisfied by the model.

Let A be a set of formulas. A term s is *discriminating* in A iff there is a term t such that $s \neq_\iota t \in A$ or $t \neq_\iota s \in A$. For each set A of formulas and each type σ we define a nonempty universe $\mathcal{U}_\sigma^A \subseteq \Lambda_\sigma$ as follows.

- Let $\mathcal{U}_o^A = \{\perp, \neg\perp\}$.
- Let \mathcal{U}_ι^A be the set of discriminating terms in A if there is some discriminating term in A .
- Let $\mathcal{U}_\iota^A = \{*\}$ if there are no discriminating terms in A .
- Let $\mathcal{U}_{\sigma\tau}^A = \{[s] \mid s \in \Lambda_{\sigma\tau}, \mathcal{V}s \subseteq \mathcal{V}A\}$.

When the set A is clear in context, we write \mathcal{U}_σ .

We call a finite set of normal formulas a *branch*. A cut-free tableau calculus for higher-order logic with extensionality is given in [8]. The calculus is complete with respect to Henkin models without choice. The details of the completeness proof indicated that one can restrict

\mathcal{C}_\perp	\perp is not in A .
\mathcal{C}_\neg	If $\neg s$ is in A , then s is not in A .
\mathcal{C}_{\neq}	$s \neq_\iota s$ is not in A .
$\mathcal{C}_{\rightarrow}$	If $s \rightarrow t$ is in A , then $A \cup \{\neg s\}$ or $A \cup \{t\}$ is in Γ .
$\mathcal{C}_{\neg\rightarrow}$	If $\neg(s \rightarrow t)$ is in A , then $A \cup \{s, \neg t\}$ is in Γ .
\mathcal{C}_\forall	If $\forall_\sigma s$ is in A , then $A \cup \{[st]\}$ is in Γ for every $t \in \mathcal{U}_\sigma^A$.
$\mathcal{C}_{\neg\forall}$	If $\neg\forall_\sigma s$ is in A , then $A \cup \{\neg[sx]\}$ is in Γ for some variable x .
\mathcal{C}_{MAT}	If $\delta s_1 \dots s_n$ is in A and $\neg\delta t_1 \dots t_n$ is in A , then $n \geq 1$ and $A \cup \{s_i \neq t_i\}$ is in Γ for some $i \in \{1, \dots, n\}$.
\mathcal{C}_{DEC}	If $\delta s_1 \dots s_n \neq_\iota \delta t_1 \dots t_n$ is in A , then $n \geq 1$ and $A \cup \{s_i \neq t_i\}$ is in Γ for some $i \in \{1, \dots, n\}$.
\mathcal{C}_{CON}	If $s =_\iota t$ and $u \neq_\iota v$ are in A , then either $A \cup \{s \neq u, t \neq u\}$ or $A \cup \{s \neq v, t \neq v\}$ is in Γ .
\mathcal{C}_{BQ}	If $s =_o t$ is in A , then either $A \cup \{s, t\}$ or $A \cup \{\neg s, \neg t\}$ is in Γ .
\mathcal{C}_{BE}	If $s \neq_o t$ is in A , then either $A \cup \{s, \neg t\}$ or $A \cup \{\neg s, t\}$ is in Γ .
\mathcal{C}_{FQ}	If $s =_{\sigma\tau} t$ is in A , then $A \cup \{[\forall x.sx =_\tau tx]\}$ is in Γ for some $x \in \mathcal{V}_\sigma \setminus (\mathcal{V}s \cup \mathcal{V}t)$.
\mathcal{C}_{FE}	If $s \neq_{\sigma\tau} t$ is in A , then $A \cup \{\neg[\forall x.sx =_\tau tx]\}$ is in Γ for some $x \in \mathcal{V}_\sigma \setminus (\mathcal{V}s \cup \mathcal{V}t)$.
\mathcal{C}_ε	If $\varepsilon_\sigma s$ is accessible in A , then either $A \cup \{[s(\varepsilon s)]\}$ is in Γ or there is some $x \in \mathcal{V}_\sigma \setminus \mathcal{V}s$ such that $A \cup \{[\forall x.\neg(sx)]\}$ is in Γ .

FIGURE 1. Abstract consistency conditions (must hold for every $A \in \Gamma$)

instantiations for quantifiers on base types to terms occurring on one side of a disequation. This restriction is shown complete for the first-order case in [8]. The calculus is extended to include choice in [3] and the restriction on instantiations is proven complete in the higher-order case. The proof of completeness makes use of abstract consistency. A set Γ of branches is an *abstract consistency class* if it satisfies all the conditions in Figure 1. This definition differs slightly from the one in [3] because we are using \rightarrow instead of \neg and \forall . With obvious modifications to account for this difference, Theorem 2 in [3] implies that every $A \in \Gamma$ (where Γ is an abstract consistency class) is satisfiable by a Henkin model. We state this here as the *Model Existence Theorem*.

Theorem 2.1 (Model Existence Theorem). *Let Γ be an abstract consistency class. Each $A \in \Gamma$ is satisfiable by a Henkin model.*

3. MAPPING INTO SAT

We next describe a simple mapping from higher-order formulas into propositional literals and clauses. The essential idea is to abstract away the semantics of logical connectives. The general technique of using a propositional abstraction is standard and is used by SMT solvers (e.g., see [9]).

Let Atom be a countably infinite set of propositional *atoms*. For each atom a , let \bar{a} denote a distinct negated atom. A *literal* is an atom or a negated atom. Let Lit be the set of all literals. Let $\bar{\bar{a}}$ denote a . A *clause* is a finite set of literals, which we write as $l_1 \sqcup \dots \sqcup l_n$. A *propositional assignment* is a mapping Φ from Atom to $\{0, 1\}$. We extend any such Φ to literals by taking $\Phi(\bar{a}) = 1 - \Phi(a)$. We say an assignment Φ *satisfies a clause* \mathcal{C} if there is some literal $l \in \mathcal{C}$ such that $\Phi l = 1$. An assignment Φ *satisfies a set* \mathcal{S} of clauses if Φ satisfies \mathcal{C} for all $\mathcal{C} \in \mathcal{S}$.

Let $\lfloor \cdot \rfloor$ be a function mapping Λ_o into Lit such that $\lfloor \neg s \rfloor = \overline{\lfloor s \rfloor}$, $\lfloor s \rfloor = \lfloor \lfloor s \rfloor \rfloor$, and if $\lfloor s \rfloor = \lfloor t \rfloor$, then $\mathcal{I}s = \mathcal{I}t$ in every Henkin model $(\mathcal{D}, \mathcal{I})$.

Theorem 3.1. *In the implementation, $\lfloor s \rfloor = \lfloor t \rfloor$ whenever s and t are the same up to $\beta\eta$ and the removal of double negations. Under some flag settings, symmetric equations $u = v$ and $v = u$ are assigned the same literal.*

We say Φ is a *pseudo-model* of A if $\Phi \lfloor s \rfloor = 1$ for all $s \in A$. We say an assignment Φ is *Henkin consistent* if there is a Henkin model $(\mathcal{D}, \mathcal{I})$ such that $\Phi \lfloor s \rfloor = \hat{\mathcal{I}}s$ for all $s \in \Lambda_o$.

4. STATES AND SUCCESSORS

Theorem 4.1. *A quasi-state Σ is a 5-tuple $(\mathfrak{F}_p^\Sigma, \mathfrak{F}_a^\Sigma, \mathfrak{U}_p^\Sigma, \mathfrak{U}_a^\Sigma, \mathfrak{C}^\Sigma)$ where \mathfrak{F}_p^Σ and \mathfrak{F}_a^Σ are finite sets of normal formulas, \mathfrak{U}_p^Σ and \mathfrak{U}_a^Σ are finite sets of normal terms, and \mathfrak{C}^Σ is a finite set of clauses. We call formulas in \mathfrak{F}_p^Σ passive formulas, formulas in \mathfrak{F}_a^Σ active formulas, terms in \mathfrak{U}_p^Σ passive instantiations and terms in \mathfrak{U}_a^Σ active instantiations.*

Given a quasi-state Σ , we define the following notation:

$$\begin{aligned} \mathfrak{F}^\Sigma &:= \mathfrak{F}_p^\Sigma \cup \mathfrak{F}_a^\Sigma & \mathfrak{U}^\Sigma &:= \mathfrak{U}_p^\Sigma \cup \mathfrak{U}_a^\Sigma \\ \mathfrak{U}_{p,\sigma}^\Sigma &:= \mathfrak{U}_p^\Sigma \cap \Lambda_\sigma & \mathfrak{U}_{a,\sigma}^\Sigma &:= \mathfrak{U}_a^\Sigma \cap \Lambda_\sigma \end{aligned}$$

During the procedure, we will only consider quasi-states that satisfy certain invariants. Such a quasi-state will be called a *state*. Before giving the technical definition of a state, we consider a simple example.

In this example we will refer to the quasi-states as *states*, as they will always satisfy the relevant properties.

Each step of the search process will pass from one state to a successor state. The passive formulas and passive instantiations of a successor state will always include all the passive formulas and passive instantiations of the previous state. Likewise, all the clauses of the previous state will be clauses of the successor state. Often we obtain a successor state by moving an active formula (instantiation) to the set of passive formulas (instantiations). We will refer to this as *processing* the formula (instantiation).

This use of the terminology “active” and “passive” was introduced in [7]. If a formula or instantiation is active, it is waiting to be processed which will cause some action to be taken. An alternative would be to refer to active formulas and instantiations as “pending.” If a formula or instantiation is passive, then it has already been processed and will no longer directly cause any actions. Passive formulas and instantiations can be used as side information when processing active formulas and instantiations. An alternative would be to refer to passive formulas and instantiations as “usable.” To avoid confusion, we point out that in some theorem proving literature (e.g., the description of Inst-Gen in [12]) the adjectives “active” and “passive” are used in just the opposite way: a formula is “passive” if it is waiting to be processed and “active” after having been processed.

Theorem 4.1. *Let $p, q : o$ be variables. Suppose we wish to refute the branch with two formulas: p and $\forall q.p \rightarrow q$. We begin with a state Σ_0 with $\mathfrak{F}_p^{\Sigma_0} = \emptyset$, $\mathfrak{F}_a^{\Sigma_0} = \{p, \forall q.p \rightarrow q\}$, $\mathfrak{U}_p^{\Sigma_0} = \{\perp, \top\}$, $\mathfrak{U}_a^{\Sigma_0} = \emptyset$ and \mathfrak{C}^{Σ_0} contains exactly the two unit clauses $\lfloor p \rfloor$ and $\lfloor \forall q.p \rightarrow q \rfloor$. We will refute this branch in one step. In particular, we process the formula $\forall q.p \rightarrow q$ by moving it from being active to passive and by applying all the instantiations of type o in $\mathfrak{U}_p^{\Sigma_0}$. This results in a state Σ_1 in which $\mathfrak{F}_p^{\Sigma_1} = \{\forall q.p \rightarrow q\}$, $\mathfrak{F}_a^{\Sigma_1} = \{p, p \rightarrow \perp, p \rightarrow \top\}$, $\mathfrak{U}_p^{\Sigma_1} = \mathfrak{U}_p^{\Sigma_0}$, $\mathfrak{U}_a^{\Sigma_1} = \mathfrak{U}_a^{\Sigma_0}$ and \mathfrak{C}^{Σ_1} contains the two unit clauses from \mathfrak{C}^{Σ_0} as well as the two clauses $\overline{\lfloor \forall q.p \rightarrow q \rfloor} \sqcup \overline{\lfloor p \rightarrow \perp \rfloor}$ and $\overline{\lfloor \forall q.p \rightarrow q \rfloor} \sqcup \overline{\lfloor p \rightarrow \top \rfloor}$. Note that $\lfloor p \rightarrow \perp \rfloor$ is the same as $\overline{\lfloor p \rfloor}$. Clearly there is no propositional assignment satisfying the clauses in \mathfrak{C}^{Σ_1} . This completes the refutation. The two states can be displayed as in Figure 2.*

Theorem 4.2. *A quasi-state $\Sigma = (\mathfrak{F}_p^\Sigma, \mathfrak{F}_a^\Sigma, \mathfrak{U}_p^\Sigma, \mathfrak{U}_a^\Sigma, \mathfrak{C}^\Sigma)$ is a state if the conditions in Figures 3 and 4 hold and for every clause \mathcal{C} in \mathfrak{C}^Σ and every literal $l \in \mathcal{C}$, either $l = \lfloor s \rfloor$ for some $s \in \mathfrak{F}^\Sigma$ or $l = \overline{\lfloor s \rfloor}$ for some $s \in \mathfrak{F}_p^\Sigma$.*

	\mathfrak{F}_p	\mathfrak{F}_a	\mathfrak{U}_p	\mathfrak{U}_a	\mathfrak{C}
Σ_0		$p, \forall q.p \rightarrow q$	\perp, \top		$[p]$ $[\forall q.p \rightarrow q]$
Σ_1	$\forall q.p \rightarrow q$	$\forall q.p \rightarrow q$ $p \rightarrow \perp, p \rightarrow \top$			$[\forall q.p \rightarrow q] \sqcup [p]$ $[\forall q.p \rightarrow q] \sqcup [p \rightarrow \top]$

FIGURE 2. States from Example 4.1

- \mathcal{S}_\perp If \perp is in \mathfrak{F}_p , then $[\perp]$ is in \mathfrak{C} .
- \mathcal{S}_{\neq} If $s \neq_i s$ is in \mathfrak{F}_p , then $[s = s]$ is in \mathfrak{C} .
- $\mathcal{S}_{\rightarrow}$ If $s \rightarrow t$ is in \mathfrak{F}_p and t is not \perp , then $\{\neg s, t\} \subseteq \mathfrak{F}$ and $[s \rightarrow t] \sqcup [\neg s] \sqcup [t]$ is in \mathfrak{C} .
- $\mathcal{S}_{\rightarrow\rightarrow}$ If $\neg(s \rightarrow t)$ is in \mathfrak{F}_p , then $\{s, \neg t\} \subseteq \mathfrak{F}$, $[s \rightarrow t] \sqcup [s]$ and $[s \rightarrow t] \sqcup [\neg t]$ are in \mathfrak{C} .
- \mathcal{S}_\forall If $\forall_\sigma s$ is in \mathfrak{F}_p and $t \in \mathfrak{U}_{p,\sigma}$, then $[st] \in \mathfrak{F}$ and $[\forall_\sigma s] \sqcup [st]$ is in \mathfrak{C} .
- $\mathcal{S}_{\neg\forall}$ If $\neg\forall_\sigma s$ is in \mathfrak{F}_p , then there is some variable x of type σ such that $\neg[sx] \in \mathfrak{F}$ and $[\forall_\sigma s] \sqcup [sx]$ is in \mathfrak{C} .
- \mathcal{S}_{MAT} If $\delta s_1 \dots s_n$ and $\neg\delta t_1 \dots t_n$ are in \mathfrak{F}_p where $n \geq 1$, then $s_i \neq t_i$ is in \mathfrak{F} for each $i \in \{1, \dots, n\}$ and $[\delta s_1 \dots s_n] \sqcup [\delta t_1 \dots t_n] \sqcup [s_1 \neq t_1] \sqcup \dots \sqcup [s_n \neq t_n]$ is in \mathfrak{C} .
- \mathcal{S}_{DEC} If $\delta s_1 \dots s_n \neq_i \delta t_1 \dots t_n$ is in \mathfrak{F}_p where $n \geq 1$, then $s_i \neq t_i$ is in \mathfrak{F} for each $i \in \{1, \dots, n\}$ and $[\delta s_1 \dots s_n = \delta t_1 \dots t_n] \sqcup [s_1 \neq t_1] \sqcup \dots \sqcup [s_n \neq t_n]$ is in \mathfrak{C} .
- \mathcal{S}_{CON} If $s =_i t$ and $u \neq_i v$ are in \mathfrak{F}_p , then $\{s \neq u, t \neq u, s \neq v, t \neq v\} \subseteq \mathfrak{F}$ and the following four clauses are in \mathfrak{C} :
 $[s = t] \sqcup [u = v] \sqcup [s \neq u] \sqcup [s \neq v]$,
 $[s = t] \sqcup [u = v] \sqcup [s \neq u] \sqcup [t \neq v]$,
 $[s = t] \sqcup [u = v] \sqcup [t \neq u] \sqcup [s \neq v]$,
 $[s = t] \sqcup [u = v] \sqcup [t \neq u] \sqcup [t \neq v]$

FIGURE 3. Conditions on a quasi-state $\Sigma = (\mathfrak{F}_p, \mathfrak{F}_a, \mathfrak{U}_p, \mathfrak{U}_a, \mathfrak{C})$

- \mathcal{S}_{BQ} If $s =_o t$ is in \mathfrak{F}_p , then $\{s, t, \neg s, \neg t\} \subseteq \mathfrak{F}$ and $\overline{[s = t]} \sqcup [s] \sqcup [\neg t]$ and $\overline{[s = t]} \sqcup [\neg s] \sqcup [t]$ are in \mathfrak{C} .
- \mathcal{S}_{BE} If $s \neq_o t$ is in \mathfrak{F}_p , then $\{s, t, \neg s, \neg t\} \subseteq \mathfrak{F}$ and $[s = t] \sqcup [s] \sqcup [t]$ and $[s = t] \sqcup [\neg s] \sqcup [\neg t]$ are in \mathfrak{C} .
- \mathcal{S}_{FQ} If $s =_{\sigma\tau} t$ is in \mathfrak{F}_p , then there is some $x \in \mathcal{V}_\sigma \setminus (\mathcal{V}s \cup \mathcal{V}t)$ such that $[\forall x.sx =_\tau tx]$ is in \mathfrak{F} and $\overline{[s = t]} \sqcup [\forall x.sx = tx]$ is in \mathfrak{C} .
- \mathcal{S}_{FE} If $s \neq_{\sigma\tau} t$ is in \mathfrak{F}_p , then there is some $x \in \mathcal{V}_\sigma \setminus (\mathcal{V}s \cup \mathcal{V}t)$ such that $[\neg \forall x.sx =_\tau tx]$ is in \mathfrak{F} and $[s = t] \sqcup [\neg \forall x.sx = tx]$ is in \mathfrak{C} .
- \mathcal{S}_ε If $\varepsilon_\sigma s$ is accessible in \mathfrak{F}_p , then there is some $x \in \mathcal{V}_\sigma \setminus \mathcal{V}s$ such that $[s(\varepsilon s)]$ and $[\forall x.\neg(sx)]$ are in \mathfrak{F} and $[s(\varepsilon s)] \sqcup [\forall x.\neg(sx)]$ is in \mathfrak{C} .

FIGURE 4. More conditions on a quasi-state $\Sigma = (\mathfrak{F}_p, \mathfrak{F}_a, \mathfrak{U}_p, \mathfrak{U}_a, \mathfrak{C})$

We say a propositional assignment Φ satisfies a state Σ if Φ satisfies \mathfrak{C}^Σ . We say Σ is *propositionally satisfiable* if there is a Φ such that Φ satisfies Σ . Otherwise, we say Σ is *propositionally unsatisfiable*. Furthermore, we say Σ is *Henkin satisfiable* if there is a Henkin consistent propositional assignment satisfying \mathfrak{C}^Σ . Note that checking whether Σ is propositionally satisfiable is simply a SAT-problem.

A variable x is *fresh* for a state Σ if x is not free in any $s \in \mathfrak{F}^\Sigma \cup \mathfrak{U}^\Sigma$.

We now consider a second simple example.

Theorem 4.2. *Let $p : \iota o$ and $x : \iota$ be variables. Suppose we wish to prove the following basic property of the choice operator ε_ι : $\forall x.px \rightarrow p(\varepsilon_\iota p)$. The refutation will proceed in seven steps taking us from an initial state Σ_0 (corresponding to assuming the negation) to a state Σ_7 such that \mathfrak{C}^{Σ_7} is propositionally unsatisfiable. The states Σ_i for $i \in \{0, \dots, 7\}$ are indicated in Figure 5. In the first step we process $\neg \forall x.px \rightarrow p(\varepsilon p)$ by choosing a fresh variable $y : \iota$ and including the new formula $\neg(py \rightarrow p(\varepsilon p))$ and a clause relating the literals corresponding to the two formulas. The resulting state is Σ_1 . We obtain Σ_2 by processing $\neg(py \rightarrow p(\varepsilon p))$ and obtaining two new formulas py and $\neg p(\varepsilon p)$ and two new clauses. We obtain Σ_3 by processing py . In general, processing such a formula involves mating it with all passive formulas of the form $\neg pt$. Since there are no such passive formulas (in particular, $\neg p(\varepsilon p)$ is active), Σ_3 only differs from Σ_2 in that py has been made passive. We obtain Σ_4 by processing $\neg p(\varepsilon p)$. This involves mating it with the passive formula py to obtain the formula $y \neq \varepsilon p$ and adding a new clause. (The reader should note that the new clause in Σ_4 will not be used to show the final set of clauses is propositionally unsatisfiable.)*

	\mathfrak{F}_p	\mathfrak{F}_a	\mathfrak{U}_p	\mathfrak{U}_a	\mathfrak{C}
Σ_0		$\neg\forall x.px \rightarrow p(\varepsilon p)$			$\overline{[\forall x.px \rightarrow p(\varepsilon p)]}$
Σ_1	$\neg\forall x.px \rightarrow p(\varepsilon p)$	$\neg\forall x.px \rightarrow p(\varepsilon p)$ $\neg(py \rightarrow p(\varepsilon p))$			$[\forall x.px \rightarrow p(\varepsilon p)] \sqcup [py \rightarrow p(\varepsilon p)]$
Σ_2	$\neg(py \rightarrow p(\varepsilon p))$	$\neg(py \rightarrow p(\varepsilon p))$ $py, \neg p(\varepsilon p)$			$[py \rightarrow p(\varepsilon p)] \sqcup [py]$ $[py \rightarrow p(\varepsilon p)] \sqcup \overline{[p(\varepsilon p)]}$
Σ_3	py	py			
Σ_4	$\neg(p(\varepsilon p))$	$\neg p(\varepsilon p)$ $y \neq \varepsilon p$			$\overline{[py]} \sqcup [p(\varepsilon p)] \sqcup \overline{[y = \varepsilon p]}$
Σ_5	$y \neq \varepsilon p$	$y \neq \varepsilon p$ $\forall x.\neg px, p(\varepsilon p)$		$y, \varepsilon p$	$[p(\varepsilon p)] \sqcup [\forall x.\neg px]$
Σ_6	$\forall x.\neg px$	$\forall x.\neg px$, $p(\varepsilon p)$			
Σ_7		$\neg py$	y	y	$\overline{[\forall x.\neg px]} \sqcup \overline{[py]}$

FIGURE 5. States from Example 4.2

To obtain Σ_5 we process $y \neq \varepsilon p$. Since y and εp are discriminating terms in the set of passive formulas of Σ_5 , we add them to the set of active instantiations. Also, since εp is accessible in $\mathfrak{F}_p^{\Sigma_5}$, we include the formulas $\forall x.\neg px$ and $p(\varepsilon p)$ as well as a clause corresponding to the meaning of the choice operator ε . We obtain Σ_6 by processing $\forall x.\neg px$. In principle, this means instantiating with all passive instantiations of type ι , but we have no passive instantiations of this type. Finally, we obtain Σ_7 by processing the instantiation y . Since y has type ι , we will use it as an instantiation for the passive formula $\forall x.\neg px$. As a consequence, we add the formula $\neg py$ and a corresponding clause. At this point, the clauses are propositionally unsatisfiable and we are done.

Given a branch A , an *initial state* Σ for A is a state with $A \subseteq \mathfrak{F}^\Sigma$, and $\mathfrak{C}^\Sigma = \{[s] \mid s \in A\}$. (We require $A \subseteq \mathfrak{F}^\Sigma$ rather than $A \subseteq \mathfrak{F}_a^\Sigma$ to allow for the possibility that some formulas in A are passive rather than active in an initial state. In practice, this could result from some preprocessing of formulas in A .) To see that for any branch A there is an initial state, consider Σ with $\mathfrak{F}_p^\Sigma = \emptyset$, $\mathfrak{F}_a^\Sigma = A$, $\mathfrak{U}_p^\Sigma = \emptyset$, $\mathfrak{U}_a^\Sigma = \emptyset$ and $\mathfrak{C}^\Sigma = \{[s] \mid s \in A\}$.

Theorem 4.3. *We say a state Σ' is a successor of a state Σ (and write $\Sigma \rightarrow \Sigma'$) if $\mathfrak{F}_p^\Sigma \subseteq \mathfrak{F}_p^{\Sigma'}$, $\mathfrak{F}_a^\Sigma \subseteq \mathfrak{F}_a^{\Sigma'}$, $\mathfrak{U}_p^\Sigma \subseteq \mathfrak{U}_p^{\Sigma'}$, $\mathfrak{U}_a^\Sigma \subseteq \mathfrak{U}_a^{\Sigma'}$, $\mathfrak{C}^\Sigma \subseteq \mathfrak{C}^{\Sigma'}$, and if Σ is Henkin satisfiable, then Σ' is Henkin satisfiable.*

Note that the successor relation is reflexive and transitive. Also, soundness of the procedure is built into the definition of the successor relation.

Theorem 4.1 (Soundness). *Let A be a branch. If there is a propositionally unsatisfiable Σ' such that $\Sigma_A \rightarrow \Sigma'$, then A is unsatisfiable.*

Proof. Assume $(\mathcal{D}, \mathcal{I})$ is a Henkin model of A . Choose Φ such that $\Phi \upharpoonright s = \hat{\mathcal{I}}s$ for each $s \in A$. Clearly, Φ demonstrates that Σ_A is Henkin satisfiable. On the other hand, since Σ' is propositionally unsatisfiable, it is Henkin unsatisfiable. This contradicts the definition of $\Sigma_A \rightarrow \Sigma'$. \square

A strategy which chooses a successor state for each propositionally satisfiable state will yield a sound procedure. One such strategy is to interleave two kinds of actions: (1) process active formulas and instantiations while making the minimal number of additions of formulas and clauses consistent with the invariants in Figure 3 and (2) generate new active instantiations. To ensure soundness, when processing a formula $\neg \forall_\sigma s$ a procedure should choose a fresh variable x , add $\neg[sx]$ to \mathfrak{F}_a and add $[\forall_\sigma s] \sqcup [sx]$ to \mathfrak{C} .

If a strategy does not lead to a propositionally unsatisfiable state, then it will give a finite or infinite path of states. If the strategy is fair, this path will satisfy certain fairness properties. In this case, we can use the path to prove the original branch is satisfiable. That is, we can conclude that every fair strategy is complete.

Theorem 4.4. *Let $\alpha \in \omega \cup \{\omega\}$. An α -path (or, simply path) is an α -sequence $\bar{\Sigma} = (\Sigma_i)_{i < \alpha}$ of propositionally satisfiable states such that $\Sigma_i \rightarrow \Sigma_{i+1}$ for each i with $i + 1 < \alpha$. We say a type σ is a quantified type on the path if there exist $i < \alpha$ and s such that $\forall_\sigma s \in \mathfrak{F}^{\Sigma_i}$. Such a path is fair if the following conditions hold:*

- (1) *For all $i < \alpha$ and $s \in \mathfrak{F}_a^{\Sigma_i}$ there is some $j \in [i, \alpha)$ such that $s \in \mathfrak{F}_p^{\Sigma_j}$.*
- (2) *If σ is a quantified type, then for all $i < \alpha$, $A \subseteq \mathfrak{F}^{\Sigma_i}$ and $t \in \mathcal{U}_\sigma^A$ there is some $j \in [i, \alpha)$ such that $t \in \mathcal{U}_p^{\Sigma_j}$.*

Given a branch A_0 , we will start with an initial state Σ_0 for A_0 . Our theorem proving procedure will construct a sequence of successor states in such a way that, unless some state is propositionally unsatisfiable, the sequence will be a fair path. In order to prove completeness of this procedure, it is enough to prove that if there is a fair path starting from Σ_0 , then A_0 is satisfiable. This result will be given as Theorem 4.1 at the end of this section.

For the remainder of this section we assume a fixed α and fair α -path $\bar{\Sigma}$.

Theorem 4.5. *Let $i < \alpha$ be given. We say a branch A is i -supported if $A \subseteq \mathfrak{F}^{\Sigma_i}$ and there is a pseudo-model Φ of A satisfying Σ_i . We say a branch A is i -consistent if A is j -supported for all $j \in [i, \alpha)$.*

Theorem 4.1. *Let $i < \alpha$ and $j \in [i, \alpha)$ be given. If A is j -supported and $A \subseteq \mathfrak{F}^{\Sigma_i}$, then A is i -supported.*

Proof. This follows from $\mathfrak{C}^{\Sigma_i} \subseteq \mathfrak{C}^{\Sigma_j}$. □

Let Γ be the set of all branches A such that A is i -consistent for some $i < \alpha$. We will prove Γ is an abstract consistency class.

Theorem 4.2. *Let A be an j -consistent branch. Let A_1, \dots, A_n be branches such that $A \subseteq A_l \subseteq \mathfrak{F}^{\Sigma_j}$ for each $l \in \{1, \dots, n\}$. Either there is some $l \in \{1, \dots, n\}$ such that A_l is j -consistent or there is some $k \in [j, \alpha)$ such that A_l is not k -supported for each $l \in \{1, \dots, n\}$.*

Proof. Assume none of A_1, \dots, A_n is j -consistent. Let $k_1, \dots, k_n \in [j, \alpha)$ be such that A_l is not k_l -supported for each $l \in \{1, \dots, n\}$. Let k be the maximum of k_1, \dots, k_n . By Lemma 4.1 each A_l is not k -supported. □

Theorem 4.3. *Γ is an abstract consistency class.*

Proof. We verify a representative collection of cases.

- \mathcal{C}_\perp Suppose $\perp \in A$ and A is i -consistent. By fairness there is some $j \in [i, \alpha)$ such that $\perp \in \mathfrak{F}_p^{\Sigma_j}$. By \mathcal{S}_\perp the unit clause $\overline{[\perp]}$ is in \mathfrak{C}^{Σ_j} . This contradicts A being j -supported.
- \mathcal{C}_\neg Suppose $\neg s$ and s are in A . Since no propositional assignment Φ can have $\Phi[\neg s] = 1$ and $\Phi[s] = 1$, A cannot be i -consistent for any i .
- \mathcal{C}_\rightarrow Suppose $s \rightarrow t$ is in an i -consistent branch A . If t is \perp , then $A \cup \{\neg s\}$ is the same as A and so $A \cup \{\neg s\}$ is i -consistent. Assume t is not \perp . Since A is i -consistent, we know $A \subseteq \mathfrak{F}^{\Sigma_i}$ and so $s \rightarrow t \in \mathfrak{F}^{\Sigma_i}$. By fairness there is some $j \in [i, \alpha)$ such that $s \rightarrow t \in \mathfrak{F}_p^{\Sigma_j}$. By \mathcal{S}_\rightarrow we know $\{\neg s, t\} \subseteq \mathfrak{F}^{\Sigma_j}$ and $\overline{[s \rightarrow t]} \sqcup \overline{[s]} \sqcup [t]$ is in \mathfrak{C}^{Σ_j} . Note that $A \cup \{\neg s\} \subseteq \mathfrak{F}^{\Sigma_k}$ and $A \cup \{t\} \subseteq \mathfrak{F}^{\Sigma_k}$ for every $k \in [j, \alpha)$. Assume neither $A \cup \{\neg s\}$ nor $A \cup \{t\}$ is j -consistent. By Lemma 4.2 there is some $k \in [j, \alpha)$ such that neither $A \cup \{\neg s\}$ nor $A \cup \{t\}$ is k -supported. Since A is i -consistent, A is k -supported and has some pseudo-model Φ satisfying Σ_k . Since $\overline{[s \rightarrow t]} \sqcup \overline{[s]} \sqcup [t]$ is in \mathfrak{C}^{Σ_k} and $\Phi[s \rightarrow t] = 1$, we must have $\Phi[s] = 0$ or $\Phi[t] = 1$. Thus Φ witnesses that either $A \cup \{\neg s\}$ or $A \cup \{t\}$ is k -supported, contradicting our choice of k . Hence either $A \cup \{\neg s\}$ or $A \cup \{t\}$ must be j -consistent.

- $\mathcal{C}_{\rightarrow}$ Suppose $\neg(s \rightarrow t)$ is in an i -consistent branch A . Since A is i -consistent, we know $\neg(s \rightarrow t) \in \mathfrak{F}^{\Sigma_i}$. By fairness there is some $j \in [i, \alpha)$ such that $\neg(s \rightarrow t) \in \mathfrak{F}_p^{\Sigma_j}$. By $\mathcal{S}_{\rightarrow}$ we know $\{s, \neg t\} \subseteq \mathfrak{F}^{\Sigma_j}$, and both $\lfloor s \rightarrow t \rfloor \sqcup \lfloor s \rfloor$ and $\lfloor s \rightarrow t \rfloor \sqcup \overline{\lfloor t \rfloor}$ are in \mathfrak{C}^{Σ_j} . We prove $A \cup \{s, \neg t\}$ is j -consistent. Let $k \in [j, \alpha)$ be given. Since A is i -consistent, it has some pseudo-model Φ satisfying Σ_k . Since $\Phi[\neg(s \rightarrow t)] = 1$, we must have $\Phi[s] = 1$ and $\Phi[\neg t] = 1$. Hence Φ is a pseudo-model of $A \cup \{s, \neg t\}$ and so $A \cup \{s, \neg t\}$ is k -supported. Therefore, $A \cup \{s, \neg t\}$ is j -consistent.
- \mathcal{C}_{\forall} Let A be an i -consistent branch such that $\forall_\sigma s \in A$ and $t \in \mathcal{U}_\sigma^A$. Note that $\forall_\sigma s \in A \subseteq \mathfrak{F}^{\Sigma_i}$ witnesses that σ is a quantified type on the path. By fairness there is some $j \in [i, \alpha)$ such that $\forall s \in \mathfrak{F}_p^{\Sigma_j}$ and $t \in \mathcal{U}_p^{\Sigma_j}$. By \mathcal{S}_{\forall} $\lfloor st \rfloor \in \mathfrak{F}^{\Sigma_j}$ and $\overline{\lfloor \forall_\sigma s \rfloor} \sqcup \lfloor st \rfloor$ is in \mathfrak{C}^{Σ_j} . We prove A is j -consistent. Let $k \in [j, \alpha)$ be given. Since A is i -consistent, it has some pseudo-model Φ satisfying Σ_k . Since $\Phi[\forall s] = 1$ and $\overline{\lfloor \forall_\sigma s \rfloor} \sqcup \lfloor st \rfloor$ is in \mathfrak{C}^{Σ_j} , we must have $\Phi[st] = 1$ and so $A \cup \{\lfloor st \rfloor\}$ is k -supported. (We know $\lfloor \lfloor st \rfloor \rfloor = \lfloor st \rfloor$ as a property of $\lfloor \cdot \rfloor$.)
- $\mathcal{C}_{\neg\forall}$ Let A be an i -consistent branch such that $\neg\forall_\sigma s \in A$. By fairness there is some $j \in [i, \alpha)$ such that $\neg\forall s \in \mathfrak{F}_p^{\Sigma_j}$. By $\mathcal{S}_{\neg\forall}$ there is some variable x such that $\neg[sx] \in \mathfrak{F}^{\Sigma_j}$ and $\lfloor \forall_\sigma s \rfloor \sqcup \overline{\lfloor sx \rfloor}$ is in \mathfrak{C}^{Σ_j} . Let $k \in [j, \alpha)$ be given. Let Φ be a pseudo-model of A satisfying Σ_k . Since $\Phi[\neg\forall s] = 1$ we must have $\Phi[\neg(sx)] = 1$ and so $A \cup \{\neg[sx]\}$ is k -supported.
- \mathcal{C}_{CON} Suppose $s =_i t$ and $u \neq_i v$ are in an i -consistent branch A . By fairness there is some $j \in [i, \alpha)$ such that $s =_i t$ and $u \neq_i v$ are $\mathfrak{F}_p^{\Sigma_j}$. By \mathcal{S}_{CON} we know the formulas $s \neq u$, $t \neq u$, $s \neq v$ and $t \neq v$ are in \mathfrak{F}^{Σ_j} and the following four clauses are in \mathfrak{C}^{Σ_j} :
- $$\begin{aligned} & \overline{\lfloor s = t \rfloor} \sqcup \lfloor u = v \rfloor \sqcup \lfloor s \neq u \rfloor \sqcup \lfloor s \neq v \rfloor, \\ & \overline{\lfloor s = t \rfloor} \sqcup \lfloor u = v \rfloor \sqcup \lfloor s \neq u \rfloor \sqcup \lfloor t \neq v \rfloor, \\ & \overline{\lfloor s = t \rfloor} \sqcup \lfloor u = v \rfloor \sqcup \lfloor t \neq u \rfloor \sqcup \lfloor s \neq v \rfloor, \\ & \overline{\lfloor s = t \rfloor} \sqcup \lfloor u = v \rfloor \sqcup \lfloor t \neq u \rfloor \sqcup \lfloor t \neq v \rfloor \end{aligned}$$
- Assume neither $A \cup \{s \neq u, t \neq u\}$ nor $A \cup \{s \neq v, t \neq v\}$ is j -consistent. We know by Lemma 4.2 that there is some $k \in [j, \alpha)$ such that neither $A \cup \{s \neq u, t \neq u\}$ nor $A \cup \{s \neq v, t \neq v\}$ is k -supported. Let Φ be a pseudo-model of A satisfying Σ_k . Note that $\Phi[s = t] = 1$ and $\Phi[u = v] = 0$. By examining the four clauses above, it is clear that we must either have $\Phi[s \neq u] = 1$ and $\Phi[t \neq u] = 1$ or have $\Phi[s \neq v] = 1$ and $\Phi[t \neq v] = 1$, a contradiction.

□

Theorem 4.1 (Model Existence). *Let A_0 be a branch and $\bar{\Sigma}$ be a fair α -path such that Σ_0 is an initial state for Σ_{A_0} . Then A_0 is satisfiable.*

Proof. By Theorem 2.1 it is enough to prove A_0 is 0-consistent. Let $j \in [0, \alpha)$ be given. Clearly $A_0 \subseteq \mathfrak{F}^{\Sigma_0} \subseteq \mathfrak{F}^{\Sigma_j}$. Let Φ satisfy Σ_j . For each $s \in A_0$, the unit clause $[s]$ is in \mathfrak{C}^{Σ_j} and so $\Phi[s] = 1$. □

5. IMPLEMENTATION

A procedure along the lines described above has been implemented in a theorem prover named Satallax. There are some minor differences from the abstract description. One difference is that double negations are eliminated during normalization in the implementation (e.g., the normal form of $p(\lambda x. \neg \neg x)$ is $p(\lambda x. x)$). Another difference is that there is no default constant $*$ of type ι . If there are no discriminating terms of type ι , then either a variable or the term $\varepsilon_\iota x. \perp$ is used as an instantiation of type ι . Also, there may be base types other than ι .

The first version of Satallax was written in Steel Bank Common Lisp. In this earlier version, MiniSat was restarted and sent all the clauses generated so far whenever propositional satisfiability was to be tested. The latest version of Satallax is implemented in Objective Caml. A foreign function interface allows Satallax to call MiniSat functions (coded in C++) in order to add new clauses to the current set of clauses and to test for satisfiability of the current set of clauses. This is a much more efficient way of using MiniSat.

Problems are given to Satallax as a TPTP file in THF format [18]. We will refer to several problems from the TPTP problem library [15] (specifically, TPTP v5.3.0). Such a problem file may include axioms and optionally a conjecture. The conjecture, if given, is negated and treated as an axiom. Logical constants that occur in axioms are rewritten in favor of the basic logical constants \perp , \rightarrow , $=_\sigma$, \forall_σ and ε_σ . Also, all definitions are expanded and the terms are $\beta\eta$ -normalized. (De Bruijn indices are used to deal with α -convertibility.) If the normalized axiom s is of the particular form $\forall px. px \rightarrow p(ep)$ or $\forall p. (\neg \forall x. \neg px) \rightarrow p(ep)$ where e is a constant of type $(\sigma o)\sigma$ for some σ , then e is registered as a choice operator of type σ and the axiom s is omitted from the initial branch. Every other normalized axiom is an initial assumption. The choice rule can be applied with every name registered as a choice operator.

There are about a hundred flags that can be set in order to control the order in which the search space is explored. A collection of flag settings is called a *mode*. Currently, there are about 300 modes predefined in

Satallax. A particular mode can be chosen via a command line option. Otherwise, a default schedule of modes is used and each of the modes on the schedule is given a certain amount of time to try to refute the problem.

If the flag `SPLIT_GLOBAL_DISJUNCTIONS` is set to `TRUE`, then Satallax will decompose the topmost logical connectives including the topmost disjunctions. This is likely to result in a set of subgoals which can be solved independently. This is an especially good idea if, for example, the conjecture is a conjunction. It is, of course, a bad idea if there are many disjunctive axioms.

Once the initial branch is determined, the state is initialized to include a unit clause for each member and the set of active formulas is initialized to be the initial branch. The terms \perp and $\neg\perp$ are added as passive instantiations. Additionally, if the flag `INITIAL_SUBTERMS_AS_INSTANTIATIONS` is set to `TRUE`, then all subterms of the initial branch are added as passive instantiations. During the search, discriminating terms of type ι are added as active instantiations. If there is a quantifier at a function type $\sigma\tau$, a process of enumerating normal terms of type $\sigma\tau$ is started. Of course, this enumeration process is the least directed part of the search procedure.

The successor relation on states was defined very generally. In particular, it does not rule out adding more formulas, instantiations and clauses than the ones suggested by the invariants on states. These additions may be very useful, but they are not necessary for completeness. A simple example is that, if the flag

`INSTANTIATE_WITH_FUNC_DISEQN_SIDES` is set to `TRUE`, the terms s and t are added as active instantiations whenever an active formula $s \neq_{\sigma\tau} t$ is processed. In addition, if `HOUNIF1` is set to `TRUE`, then higher-order unification is used to suggest terms to add as active instantiations.

Theorem 5.1. *The conjecture of the problem SEU868⁵ from the TPTP states that one characterization of $C_{\alpha o}$ being a finite set implies another characterization.*

$$\begin{aligned} & (\forall w_{(\alpha o)o}.w(\lambda x.\perp) \wedge (\forall r_{\alpha o}x_{\alpha}.wr \rightarrow w(\lambda t.rt \vee t = x)) \rightarrow wC) \rightarrow \\ & \quad \forall P_{(\alpha o)o}.(\forall E_{\alpha o}.\neg(\exists t.Et) \rightarrow PE) \wedge \\ & (\forall Y_{\alpha o}z_{\alpha}Z_{\alpha o}.PY \wedge (\forall u_{\alpha}.Zu \equiv (Yu \vee u = z)) \rightarrow PZ) \rightarrow PC \end{aligned}$$

If `SPLIT_GLOBAL_DISJUNCTIONS` is true, the initial branch contains

$$\begin{aligned} & \forall w_{(\alpha o)o}.w(\lambda x.\perp) \wedge (\forall r_{\alpha o}x_{\alpha}.wr \rightarrow w(\lambda t.rt \vee t = x)) \rightarrow wC, \\ & \quad \forall E_{\alpha o}.(\forall t.\neg Et) \rightarrow PE, \\ & \forall Y_{\alpha o}z_{\alpha}Z_{\alpha o}.PY \wedge (\forall u_{\alpha}.Zu \equiv (Yu \vee u = z)) \rightarrow PZ, \\ & \quad \neg PC \end{aligned}$$

Several instantiations are used in the proof. Two subterms of this initial branch are P and $\lambda x.\perp$. *Satallax* uses P as an instantiation for w and $\lambda x.\perp$ as an instantiation for E if

`INITIAL_SUBTERMS_AS_INSTANTIATIONS` is true. This yields new formulas

$$P(\lambda x.\perp) \wedge (\forall r_{\alpha o}x_{\alpha}.Pr \rightarrow P(\lambda t.rt \vee t = x)) \rightarrow PC$$

and

$$(\forall t.\neg \perp) \rightarrow P(\lambda x.\perp).$$

Processing the former formula leads to the introduction of the variables r and x . Because of mating, the terms r , x and $\lambda z_{\alpha}.rz \vee z = x$ eventually appear as sides of disequations. This leads to the use of r as an instantiation for Y , x for z and finally $\lambda z_{\alpha}.rz \vee z = x$ for Z if `INSTANTIATE_WITH_FUNC_DISEQN_SIDES` is true. This yields the formula $Pr \wedge (\forall u.(ru \vee u = x) \equiv (ru \vee u = x)) \rightarrow P(\lambda z.rz \vee z = x)$. Once these appropriate instantiations have been made, the rest of the search is easy.

At each stage of the search there are a number of options for continuing the search. We call these options *commands*. There are commands for processing active formulas and active instantiation. Executing such a command creates a successor state in which the formula or instantiation is passive. Whenever a command is generated (e.g., when a new active formula is created), the command is put into a priority queue. Priorities of different commands are computed using the values of various flags.

As mentioned earlier, in order to have a complete procedure, we need to enumerate closed terms to use as instantiations for function types. There are various commands that correspond to making progress in such an enumeration process. There are two (independent) enumeration procedures implemented in *Satallax*. The first enumeration procedure uses commands to request normal terms of a certain type in a certain context. The execution of each command corresponds to making a further commitment about the term requested. Once a term is completely determined, it is added as an active instantiation. The second enumeration procedure is used if the flag `ENUM_ITER_DEEP` is true. This alternate procedure periodically generates all closed terms

using the current set of variables and logical constants up to a certain depth. For completeness, each time we enumerate again, we increase the depth. We next consider a simple example which demonstrates both enumeration procedures.

Theorem 5.2. *The conjecture in the TPTP problem SYO548^1 is*

$$\exists E_{(\iota)\iota} \forall P_{\iota\circ}. (\exists X_{\iota}. \neg PX) \rightarrow \neg(P(EP)).$$

This states that there is an operator which maps each predicate P over a type ι to an element of ι that does not satisfy P , if such an element exists. The idea is to obtain such an operator by applying ε_{ι} to the complement of P . That is, we should prove this by instantiating with $(\lambda p_{\iota\circ}. \varepsilon(\lambda x_{\iota}. \neg px))$ for E . Several modes can solve this problem in less than a second. We describe the search space generated by two modes: MODE81 (using the first enumeration procedure – with ENUM_ITER_DEEP false) and MODE261 (using the second enumeration procedure – with ENUM_ITER_DEEP true). Since $(\iota)\iota$ is a quantified type, both enumeration procedures will attempt to generate instantiation terms of this type.

With MODE81 Satallax executes a sequence of commands including the following. Request a term of type $(\iota)\iota$. After imitating ε_{ι} , request a term of the form $\lambda p_{\iota\circ}. \varepsilon_{\iota}[\]$ where the hole $[]$ is yet to be determined. To fill the hole, request a term of type $\iota\circ$ where $p_{\iota\circ}$ may occur free. After imitating negation, request a term of the form $\lambda x_{\iota}. \neg[\]$ where the hole $[]$ is yet to be determined. To fill the hole, request a term of type \circ where $p_{\iota\circ}$ and x_{ι} may occur free. After projecting p , request a term of the form $p[\]$ where $[]$ is yet to be determined. To fill this last hole, request a term of type ι in which $p_{\iota\circ}$ and x_{ι} may occur free. Projecting x fills this hole and the completed term $(\lambda p_{\iota\circ}. \varepsilon(\lambda x_{\iota}. \neg px))$ is added as an active instantiation. After this, the state quickly becomes propositionally unsatisfiable. No other potential instantiations of type $(\iota)\iota$ are generated before this useful one is generated. For this reason, by the time the state becomes propositionally unsatisfiable, only 15 clauses have been generated, 6 of which form an unsatisfiable core. Note that each request above is a separate command which is put into the priority queue and is hence interleaved with the other commands of the search procedure.

With MODE261 Satallax periodically enumerates terms of type $(\iota)\iota$ with sizes limited to a certain depth. The size of a term depends on the value of various integer flags. The first time Satallax enumerates terms with size bounded by 3 and each time the allowed size increases by 2. The flag settings for MODE261 mean that no terms are generated

until the allowed size is 103, at which point four terms are enumerated. On the next iteration when the allowed size is 105, 365 terms are enumerated to be used as instantiations. One of these terms is $(\lambda p_{\iota o}.\varepsilon(\lambda x_{\iota}.\perp = px))$, which leads to a successful proof. By the time the state becomes propositionally unsatisfiable, 6053 propositional clauses have been generated; 11 of these clauses form an unsatisfiable core.

The next example has only been proven using a mode in which `ENUM_ITER_DEEP` is true.

Theorem 5.3. *The TPTP problem `GRA027^1` is a problem in propositional type theory. It represents the fact that the Ramsey number $R(3,3)$ is greater than 4. The type `oo` has exactly four elements. The conjecture states that there is a symmetric relation $G : (oo)(oo)o$ with no 3-cliques and no 3-anticliques. `Satallax` can solve this problem with mode `MODE261`. Enumerating all terms of type `(oo)(oo)o` with size bounded by 3 leads to 13 possible instantiations. Either of the two terms $\lambda X : oo.\lambda Y : oo.X\perp = Y\perp$ or $\lambda X : oo.\lambda Y : oo.Y\perp = X\perp$ provide a solution.*

One of the most useful extensions implemented in `Satallax` is, under certain flag settings, to generate higher-order clauses with higher-order literals to be matched against formulas as the formulas are processed. Such higher-order clauses are only used when every existential variable in the clause has a strict occurrence in some literal. (A strict occurrence is essentially a pattern occurrence which is not below another existential variable [13].) The following example illustrates the use of higher-order clauses.

Theorem 5.4. *The problem `SEU506^2` of the TPTP is a simple set theory example. We assume a membership relation $\in : \iota o$ and an empty set $\emptyset : \iota$. We will write $s \in t$ for $\in st$ and $s \notin t$ for $\neg(\in st)$. The problem states that if we have*

$$\forall x_{\iota}.x \in \emptyset \rightarrow \forall p_{o}.p,$$

and

$$\forall XY_{\iota}.\left(\forall x_{\iota}.x \in X \rightarrow x \in Y\right) \rightarrow \left(\forall x_{\iota}.x \in Y \rightarrow x \in X\right) \rightarrow X = Y,$$

then we know

$$\forall A_{\iota}.\left(\forall x_{\iota}.x \notin A\right) \rightarrow A = \emptyset.$$

Over a hundred modes can solve this problem in less than 2 seconds, many of which do not use higher-order clauses. An example of a mode using higher-order clauses that solves the problem is `MODE205`, which solves it in less than 0.1 seconds. Since `SPLIT_GLOBAL_DISJUNCTIONS`

is true, the initial branch contains four formulas:

$$\begin{aligned} & A \neq \emptyset \\ & \forall x.x \notin A \\ \forall XY.\iota.(\forall x_\iota.x \in X \rightarrow x \in Y) \rightarrow (\forall x_\iota.x \in Y \rightarrow x \in X) \rightarrow X = Y \\ & \forall x_\iota.x \in \emptyset \rightarrow \forall p_o.p \end{aligned}$$

for a fixed $A : \iota$. As these formulas are processed the following higher-order clauses are created:

$$\begin{aligned} & ?x \notin A \\ ?X = ?Y \mid \neg(\forall x_\iota.x \in ?X \rightarrow x \in ?Y) \mid \neg(\forall x_\iota.x \in ?Y \rightarrow x \in ?X) \\ & ?p \mid ?z \notin \emptyset \end{aligned}$$

Here $?x$, $?X$, $?Y$, $?z$ and $?p$ are existential variables. Each time a passive formula is processed, Satallax attempts to match literals in higher-order clauses against the negation of the formula. Each time a ground instance of a higher-order clause is generated corresponding propositional clauses are created and new active formulas are added to the state. In this particular example, the literal $?X = ?Y$ matches the formula $A = \emptyset$ determining a ground instance of the second higher-order clause. This leads to the consideration of the two formulas $\neg(\forall x.x \in A \rightarrow x \in \emptyset)$ and $\neg(\forall x.x \in \emptyset \rightarrow x \in A)$. Two corresponding witnesses a and b are created and eventually the formulas $a \in A$, $a \notin \emptyset$, $b \in \emptyset$ and $b \notin A$. The sole literal in the unit clause $?x \notin A$ matches the negation of the formula $a \in A$. Likewise, the higher-order clause $?p \mid ?z \notin \emptyset$ contributes to the search since $?z \notin \emptyset$ matches $b \notin \emptyset$ and $?p$ matches an appropriate formula. (Many instantiations for $?p$ could be useful here. The one Satallax uses happens to be $b \in \emptyset \rightarrow b \in A$.)

6. RESULTS AND FURTHER EXAMPLES

TPTP v5.3.0 contains 2924 problems in THF0 format. Among these, 347 are known to be satisfiable. (Satallax 2.4 terminates on many of these problems, recognizing them as satisfiable.) For 2001 of the remaining 2577 problems (77.6%), there is some mode that Satallax 2.4 can use to prove the theorem (or show the assumptions are unsatisfiable) within 30 seconds. A strategy schedule running 61 modes for 5 minutes can solve each of the 2001 problems.

One reason for the success of Satallax is that it can solve some problems by brute force. SEV106^5 is a TPTP problem which codes a Ramsey-style theorem about graphs and cliques. In Example 5.3 we proved a lower bound for $R(3, 3)$ by providing an appropriate graph with 4 nodes. In this example we must prove an upper bound for $R(3, 3)$ by proving no such appropriate graph with 6 nodes exists. In

particular, we assume there are at least six distinct individuals and that there is a symmetric relation (i.e., an undirected graph) on individuals. There must be three distinct individuals all of whom are related or all of whom are unrelated. Since we are assuming there are six distinct individuals, we quickly have six corresponding discriminating terms. Satallax uses all six of these (blindly) as instantiations for the existential quantifiers, leading to 6^3 instantiations. Using mode MODE1 Satallax generates over 8000 propositional clauses which MiniSat can easily recognize as unsatisfiable. In most examples only a handful of the clauses are the cause of unsatisfiability. In this example a 284 clauses are used to show unsatisfiability.

Two higher-order examples from the TPTP that Satallax can solve are SYO378⁵ and SYO379⁵. These examples were created in TPS to illustrate the concept of quantificational depth, discussed at the end of [1]. Let $c : \iota$ be a variable and define $d_0 := \lambda x : \iota.x = c$, $d_1 := \lambda y : \iota.o.y = d_0 \wedge \exists x.yx$ and $d_2 := \lambda z : (\iota o)o.z = d_1 \wedge \exists y.zy$ (where $s \wedge t$ means $\neg(s \rightarrow \neg t)$ and $\exists x.s$ means $\neg\forall x.\neg s$). One of the examples is $\exists y.d_1y$ and the other is $\exists z.d_2z$. A high-level proof is simply to note that d_0c , d_1d_0 and d_2d_1 are all provable. However, if we expand all definitions, then these instantiations are no longer so easy to see. Fortunately, if the flag INSTANTIATE_WITH_FUNC_DISEQN_SIDES is set to TRUE, then d_0 and d_1 will appear as the side of a disequation and Satallax will include them as instantiations early. Verifying the instantiations work is not difficult. There are modes that can solve these problems within a second.

We also discuss two particularly interesting examples: SEV429¹ and SEV430¹. the TPTP. In both examples we use variables $f, g : \mu$ and $x, y : \iota$. SEV429¹ means every injective function f has a left inverse g and SEV430¹ means every surjective function f has a right inverse g .

$$(\forall x\forall y.fx = fy \rightarrow x = y) \rightarrow \exists g.\forall x.(g(fx)) = x \quad \text{SEV429}^1$$

$$(\forall y.\exists x.fx = y) \rightarrow \exists g.\forall x.(f(gx)) = x \quad \text{SEV430}^1$$

In both SEV429¹ and SEV430¹ Satallax must enumerate potential instantiations of type μ for g . Some of the instantiations (e.g., $\lambda x.x$, f and $\lambda x.f(fx)$) are unhelpful and only serve to make the search space large. In both cases the instantiation used in the refutation is $\lambda y.\varepsilon x.fx = y$. An equivalent instantiation, $\lambda y.\varepsilon x.y = fx$, is also generated. (While it seems likely that such an equivalent instantiation could be discarded without sacrificing completeness, there is no currently known meta-theoretic result to justify this intuition.)

Satallax can prove SEV429¹ using mode MODE214 in about 10 seconds. In the process it generates 15 candidates for g and 115609 propositional clauses. Only 10 of the clauses are needed. Satallax

can prove SEV430¹ using mode MODE207 in under 2 seconds. In the process it generates 16 higher-order instantiations (candidates for g) and 19806 propositional clauses. It turns out that only 6 of these clauses are required to determine propositional unsatisfiability.

Satallax can also prove SEV430¹ using mode MODE286 in a few milliseconds. The reason is the flags EXISTS_TO_CHOICE, HOUNIF1 and HOUNIFMATE1 are set to TRUE. (Each of these flags is new to version 2.4 of Satallax.) When the flag EXISTS_TO_CHOICE is set to TRUE, the problem is preprocessed by using the choice operator ε to (essentially) Skolemize the problem. In particular, each negative occurrence of a formula $\forall_o s$ is replaced by the logically equivalent formula $[s(\varepsilon x. \neg s x)]$ (where x is fresh). In SEV430¹ the problem is transformed to give two assumptions:

$$\forall y. f(\varepsilon x. f x = y) = y$$

and

$$\forall g. f(g(\varepsilon y. f(gy) \neq y)) \neq (\varepsilon y. f(gy) \neq y).$$

Note that the first assumption essentially contains the appropriate instantiation of g as a subterm. Since the flags HOUNIF1 and HOUNIFMATE1 are set to TRUE, Satallax will attempt to generate instantiations by creating existential variables, mating potentially complementary formulas and performing higher-order unification. In this particular case, existential variables ${}^?Y$ and ${}^?G$ will be created and a disagreement pair of the form

$$(f(\varepsilon x. f x = {}^?Y) = {}^?Y) = {}^? (f({}^?G(\varepsilon y. f({}^?Gy) \neq y)) = (\varepsilon y. f({}^?Gy) \neq y)).$$

Higher-order unification can solve for an appropriate instantiation for ${}^?G$ very quickly.

Satallax 1.4 (the last version of Satallax coded in Common Lisp) competed in the higher-order division of CASC in 2010 [16] and was able to prove 120 out of 200 problems, coming in second to LEO-II 1.2 [4] which proved 125 out of 200 problems.

Satallax 2.1 competed in the higher-order division of CASC in 2011 [17] and came in first by proving 246 out of 300 problems. LEO-II 1.2.8 came in second proving 208 of 300 problems.

Satallax 2.4 will compete in the higher-order division of CASC in 2012.²

²If this paper is accepted, results should be available by the time the final version is due.

7. RELATED WORK

Smullyan introduced the notion of abstract consistency in 1963 [14]. One of his applications of abstract consistency is to justify reducing first-order unsatisfiability of a set M to propositional unsatisfiability of an extended set $R \cup M$. The procedure described in this paper and implemented in Satallax was developed without Smullyan’s application in mind. Nevertheless, one can consider the procedure to be both an elaboration of Smullyan’s idea as well as an extension to the higher-order case.

A different instantiation-based method Inst-Gen is described in [11]. Inst-Gen generates ground instances of first-order clauses and searches by interacting with a SAT-solver. This method is implemented in the first-order prover iProver [11, 12]. Note that iProver is also coded in Objective Caml and uses MiniSat via a foreign function interface. Two differences between the Inst-Gen method and the method in this paper should be noted. First, Inst-Gen assumes the problem is in clausal normal form. We do not make this assumption. As is well known, a substitution into a higher-order clause may lead to the need for further clause normalization. Second, Inst-Gen assumes an appropriate ordering on closures (clauses with substitutions). This ordering leads to important restrictions on inferences that can significantly improve the performance of Inst-Gen. We do not make use of any such ordering. In fact, a straightforward attempt to find such an ordering for the higher-order case is doomed to failure. This can be briefly indicated by an example. Suppose we define a closure to be a pair $C \cdot \theta$ of an atomic formula C and a substitution θ . The basic condition of a closure ordering \succ (see [11]) is that $C \cdot \sigma \succ D \cdot \tau$ whenever $C\sigma = D\tau$ and $C\theta = D$ for some “proper instantiator” θ . In the higher-order case, we would consider equality of normal forms instead of strict syntactic equality. Consider two atomic formulas $C := p(\lambda xy.fxy)$ and $D := p(\lambda yx.fxy)$ where p , f , x and y are variables of appropriate types. Consider the substitution $\theta p := \lambda fxy.p(\lambda yx.fxy)$. Clearly $C\theta$ is β -equivalent to D and $D\theta$ is β -equivalent to C . An appropriate ordering (assuming θ would be considered a “proper instantiator”) would need to have $C \cdot \emptyset \succ D \cdot \emptyset \succ C \cdot \emptyset$ where \emptyset plays the identity substitution.

Regarding higher-order theorem provers, two well-known examples are TPS [2] and LEO-II [4]. Automated search in TPS is based on expansion proofs while search in LEO-II is based on a resolution calculus. Both TPS and LEO-II make essential use of existential variables which are partially instantiated during search. LEO-II was the first higher-order prover to take a cooperative approach. LEO-II makes

calls to a first-order theorem prover to determine if the current set of higher-order clauses maps to an unsatisfiable set of first-order clauses.

8. CONCLUSION

We have given an abstract description of a search procedure for higher-order theorem proving. The key idea is to start with a notion of abstract consistency which integrates a restriction on instantiations. We gave a notion of a state which consists of finite sets of formulas, instantiations and propositional clauses. The invariants in the definition of a state correspond to the abstract consistency conditions. We have given a successor relation on states. Any fair strategy for choosing successors (until the set of propositional clauses is unsatisfiable) will give a complete theorem prover.

We have also described the implementation of this procedure as a higher-order theorem prover Satallax. Satallax is still new and there is a lot of room for improvement and further research. For example, it would improve Satallax significantly if one integrated procedures for solving for set variables as described in [5, 6]. Also, at the moment Satallax is very weak on problems that involve equational reasoning. Ideas for integrating efficient equational reasoning with the larger search procedure are needed.

Acknowledgements Thanks to Andreas Teucke whose ongoing work on translating Satallax refutations to Coq proof terms has led to the consideration of some of the examples in this paper. Thanks also to Gert Smolka for his support and stimulating conversations.

REFERENCES

- [1] Andrews, P.B., Bishop, M., Brown, C.E.: System description: TPS: A theorem proving system for type theory. In: D. McAllester (ed.) Proceedings of the 17th International Conference on Automated Deduction, *Lecture Notes in Artificial Intelligence*, vol. 1831, pp. 164–169. Springer-Verlag, Pittsburgh, PA, USA (2000)
- [2] Andrews, P.B., Brown, C.E.: TPS: A hybrid automatic-interactive system for developing proofs. *Journal of Applied Logic* **4**(4), 367–395 (2006)
- [3] Backes, J., Brown, C.E.: Analytic Tableaux for Higher-Order Logic with Choice. In: R.H. Jürgen Giesl (ed.) Automated Reasoning: 5th International Joint Conference, IJCAR 2010, Edinburgh, UK, July 16-19, 2010, Proceedings, *LNCS/LNAI*, vol. 6173, pp. 76–90. Springer (2010)
- [4] Benzmüller, C., Paulson, L., Theiss, F., Fietzke, A.: LEO-II — A cooperative automatic theorem prover for classical higher-order logic. In: Fourth International Joint Conference on Automated Reasoning (IJCAR’08), *LNCS (LNAI)*, vol. 5195, pp. 162–170. Springer (2008). URL www.ags.uni-sb.de/~chris/papers/C25.pdf

- [5] Bledsoe, W.W., Feng, G.: Set-Var. *Journal of Automated Reasoning* **11**, 293–314 (1993)
- [6] Brown, C.E.: Solving for Set Variables in Higher-Order Theorem Proving. In: A. Voronkov (ed.) *Proceedings of the 18th International Conference on Automated Deduction, Lecture Notes in Artificial Intelligence*, vol. 2392, pp. 408–422. Springer-Verlag, Copenhagen, Denmark (2002)
- [7] Brown, C.E.: Reducing Higher-Order Theorem Proving to a Sequence of SAT Problems. In: N. Bjørner, V. Sofronie-Stockkermans (eds.) *CADE - the 23rd International Conference on Automated Deduction, LNCS/LNAI 6803*, pp. 147 – 161. Springer (2011)
- [8] Brown, C.E., Smolka, G.: Analytic tableaux for simple type theory and its first-order fragment. *Logical Methods in Computer Science* **6**(2) (2010)
- [9] De Moura, L., Bjørner, N.: Satisfiability modulo theories: introduction and applications. *Communications of the ACM* **54**(9), 69–77 (2011). DOI 10.1145/1995376.1995394. URL <http://doi.acm.org/10.1145/1995376.1995394>
- [10] Eén, N., Sörensson, N.: An Extensible SAT-solver. In: E. Giunchiglia, A. Tacchella (eds.) *Theory and Applications of Satisfiability Testing, Lecture Notes in Computer Science*, vol. 2919, pp. 333–336. Springer Berlin / Heidelberg (2004)
- [11] Korovin, K.: iProver – an instantiation-based theorem prover for first-order logic (system description). In: A. Armando, P. Baumgartner, G. Dowek (eds.) *Proceedings of the 4th International Joint Conference on Automated Reasoning, (IJCAR 2008), Lecture Notes in Computer Science*, vol. 5195, pp. 292–298. Springer (2008)
- [12] Korovin, K., Sticksel, C.: iprover-eq: An instantiation-based theorem prover with equality. In: R.H. Jürgen Giesl (ed.) *Automated Reasoning: 5th International Joint Conference, IJCAR 2010, Edinburgh, UK, July 16-19, 2010, Proceedings, LNCS/LNAI*, vol. 6173, pp. 196–201. Springer (2010)
- [13] Pfenning, F., Schürmann, C.: Algorithms for equality and unification in the presence of notational definitions. In: T. Altenkirch, W. Naraschewski, B. Reus (eds.) *TYPES 1998, Lecture Notes in Computer Science*, vol. 1657, pp. 179–193. Springer (1999)
- [14] Smullyan, R.M.: A unifying principle in quantification theory. *Proceedings of the National Academy of Sciences, U.S.A.* **49**, 828–832 (1963)
- [15] Sutcliffe, G.: The TPTP Problem Library and Associated Infrastructure: The FOF and CNF Parts, v3.5.0. *Journal of Automated Reasoning* **43**(4), 337–362 (2009)
- [16] Sutcliffe, G.: The 5th IJCAR Automated Theorem Proving System Competition - CASC-J5. *AI Communications* **24**(1), 75–89 (2011)
- [17] Sutcliffe, G.: The CADE-23 Automated Theorem Proving System Competition - CASC-23. *AI Communications* **25**(1), 49–63 (2012)
- [18] Sutcliffe, G., Benzmüller, C.: Automated Reasoning in Higher-Order Logic using the TPTP THF Infrastructure. *Journal of Formalized Reasoning* **3**(1), 1–27 (2010)