

Cooperation of Background Reasoners in Theory Reasoning by Residue Sharing

Cesare Tinelli (tinelli@cs.uiowa.edu)

Department of Computer Science

The University of Iowa

Iowa City, IA, USA

Abstract. We propose a general way of combining background reasoners in theory reasoning. Using a restricted version of the Craig Interpolation Lemma, we show that background reasoner cooperation can be achieved as a form of constraint propagation, much in the spirit of existing combination methods for decision procedures. In this case, constraint information is propagated across reasoners by exchanging *residues* that are, in essence, disjunctions of ground literals over a common signature. As an application of our approach, we describe a multi-theory version of the semantic tableau calculus and we prove it sound and complete.

Keywords: Theory reasoning, multiple background reasoners

1. Introduction

Theory reasoning is a powerful deduction paradigm in which a general-purpose main reasoner is complemented by a *background* reasoner, a procedure specialized in (semi-)deciding formula satisfiability with respect to a fixed theory of interest, the *background theory*. In the resulting system, the inference steps of the main reasoner are typically subject to certain constraints over the background theory. The satisfiability of these constraints is not verified by the main reasoner itself but is instead delegated to the background reasoner. The main motivation for this delegation is that a background reasoners, for being domain-specific, is typically more efficient than the main one at processing constraints over the background theory. Alternatively, a background reasoner may be already available and ready to use. Finally, a certain theory may be decidable but impractical to express axiomatically.¹ In that case, a more viable option is to rely on an algorithmic representation of the theory, its decision procedure, and use it as a background reasoner.

Although the main idea of theory reasoning can be found in several early works (such as (Bibel, 1982; Plotkin, 1972) to name just a few), the first systematic treatment of it was given by Stickel in

¹ Presburger arithmetic for instance is decidable but its (first-order) axiomatization contains infinitely many instances of the induction schema.



(Stickel, 1985) which describes a theory version of the resolution calculus and the matings calculus. After that work, nearly all existing calculi for automatic reasoning have been extended to theory reasoning (see (Baumgartner et al., 1992) for a survey). Essentially all of them, however, consider the integration of just one background reasoner into the main one.

The reason for such a restriction, despite the clear desirability for modularity and scalability purposes of having several background reasoners at once, seems to be simply that no one really knew up to now how to achieve theory reasoning with multiple background reasoners in general. Note that, typically, it is not enough to integrate background reasoners *separately* into the main reasoner. To reason correctly with formulas spanning over several background theories, some sort of cooperation among the background reasoners is necessary. Finding a general way to achieve this cooperation in a sound and complete way is a non-trivial task.

We introduce one such way in this paper.² We show that the cooperation of background reasoners in theory reasoning is actually conceptually simple, and can be easily described in terms of partial theory reasoning in the sense of (Stickel, 1985). We appeal to a variant of a well-known interpolation result, the Craig Interpolation Lemma, to show that background reasoner cooperation can be achieved as a form of constraint propagation, much in the spirit of well-known combination methods for decision procedures (Nelson and Oppen, 1979). The main idea is to propagate information between reasoners by exchanging quantifier-free *residues* (see later) over a common signature.

The Craig Interpolation Lemma states that whenever two first-order theories \mathcal{T}_1 and \mathcal{T}_2 are jointly unsatisfiable they have an *interpolant*, a sentence φ made only of symbols shared by \mathcal{T}_1 and \mathcal{T}_2 , which is entailed by one theory and unsatisfiable with the other. Now, although the lemma is in principle enough for the type of reasoner cooperation we suggest, it is not useful in this general formulation because it provides no information on the syntactical form of interpolants—which could then be arbitrary formulas. This is unfortunate because most theory reasoning calculi effectively work only with certain types of formulas (quantifier-free, usually).

We provide a restricted version of the lemma which shows that, in essence, in the context of theory reasoning all needed interpolants are disjunctions of *ground literals*. Thanks to this result, background reasoner cooperation by constraint propagation becomes a viable option, as we will try to demonstrate. To do that we describe a multi-

² A preliminary version of this paper was presented at FTP 2000 (Tinelli, 2000).

theory extension of the semantic tableaux calculus, easily obtained from the corresponding single-theory version, that integrates and combines background reasoners by means of residue sharing. We show that the calculus is sound and complete under very general assumptions on the reasoners and their theories. We also show that for a large class of background theories the calculus remains complete even if residues are further restricted from disjunctions of literals to single literals.

After all this it will be clear that—like in all cases of partial theory reasoning—the real challenge lies in identifying specific situations in which the generation of residues can be implemented in a reasonably efficient and complete way. But since this is a research problem in its own right, we must leave its discussion to further work.

The current paper is organized as follows. Section 2 presents some formal preliminaries. Section 3 briefly describes the theory reasoning paradigm and explains how it can be extended to more than one background reasoner and theory. The same section also presents the specialized interpolation results that we will use to combine background reasoners. Section 4 describes a multi-theory version of free variable semantic tableaux in which background reasoners cooperate by sharing residues. The calculus is first proved sound in general, and then it is proved complete under the restriction that shared residues be just disjunctions of literals in the signature shared by the background theories. Section 5 presents a refinement of the interpolation results in Section 3 and the completeness result in Section 4 to the case of what we call Σ -convex theories. The class of Σ -convex theories is both large and significant for theory reasoning because it includes all universal Horn theories and also a number of prominent non-Horn background theories such as the theory of rational numbers under addition. Section 6 compares this work to the few existing results on the combination of background reasoners for theory reasoning. Section 7 concludes the paper with suggestions for further research. The proofs of the more technical lemmas in the paper can be found in the appendix.

2. Preliminaries

For convenience and generality, we will use first-order logic (FOL) with equality as our logical framework. In this logic, the equality symbol is a predefined logical constant, always interpreted as the identity relation. FOL without equality, the traditional logic of automated reasoning, can be obtained from FOL with equality by simply restricting the language to formulas without the predefined equality symbol.

α	α_1	α_2
$\phi \wedge \psi$	ϕ	ψ
$\neg(\phi \vee \psi)$	$\neg\phi$	$\neg\psi$
$\neg(\phi \rightarrow \psi)$	ϕ	$\neg\psi$
$\neg\neg\phi$	ϕ	$\neg\perp$

β	β_1	β_2
$\phi \vee \psi$	ϕ	ψ
$\neg(\phi \wedge \psi)$	$\neg\phi$	$\neg\psi$
$\phi \rightarrow \psi$	$\neg\phi$	ψ
$\phi \leftrightarrow \psi$	$\phi \wedge \psi$	$\neg\phi \wedge \neg\psi$
$\neg(\phi \leftrightarrow \psi)$	$\phi \wedge \neg\psi$	$\neg\phi \wedge \psi$

γ	$\gamma_1(y)$
$\forall x.\phi(x)$	$\phi\{x \mapsto y\}$
$\neg\exists x.\phi(x)$	$\neg\phi\{x \mapsto y\}$

δ	δ_t
$\neg\forall x.\phi(x)$	$\neg\phi\{x \mapsto t\}$
$\exists x.\phi(x)$	$\phi\{x \mapsto t\}$

Figure 1. Formula Types.

In this paper, a *signature* Σ consists of a set Σ^P of *relation symbols* and a set Σ^F of *function symbols*, each with an associated *arity*, an integer $n \geq 0$. A *constant* symbol is a function symbol of zero arity. Throughout the paper, we will fix a countably infinite set V of variables. Also, we will fix a constant symbol a and we will implicitly assume that every signature we consider includes a . As it will be clear later, this assumption leads to no loss of generality as far as the results presented here are concerned, but it will simplify some of the proofs in the paper.

For all signatures Σ , following (Fitting, 1996), we denote by Σ^{sko} the signature obtained by adding to Σ a countably infinite set of function symbols of arity n (not already in Σ), for all $n \geq 0$. Also, if X is any set disjoint from Σ , we denote by $\Sigma(X)$ the signature obtained by adding the elements of X as constant symbols to Σ .

We use the standard notions of formula, clause, literal, free and bound variable, substitution, structure (aka model) and so on. If φ is a formula, we denote by $\forall\varphi$ the universal closure of φ and by $\text{Var}(\varphi)$ the set of φ 's free variables, with $\text{Var}(\Phi)$ extending the notation to sets Φ of formulas in the obvious way. We write $\varphi(x)$ to indicate that x is a free variable of φ .

A *sentence* is a formula with no free variables. A *ground* formula is a formula with no variables. A *theory* is a set of sentences—we do not insist that the set be consistent. We will talk of Σ -*formula*, Σ -*structure*, Σ -*theory* and so on, whenever we want to specify that they have signature Σ . We denote by \perp the universally false formula and assume that it is a (ground) Σ -literal for every signature Σ .

If σ is a substitution we denote by $\text{Dom}(\sigma)$ the set of variables v such that $v\sigma \neq v$ and by $\text{Ran}(\sigma)$ the set $\text{Var}(\text{Dom}(\sigma)\sigma)$. We denote the

empty substitution by ε . We use the notation $\{x_1 \mapsto t_1, \dots, x_n \mapsto t_n\}$ to denote a substitution σ with $\text{Dom}(\sigma) = \{x_1, \dots, x_n\}$ and such that $x_i\sigma = t_i$ for all $i \in \{1, \dots, n\}$. As usual, we only consider idempotent substitutions. We extend the application of substitutions to first-order formulas as obvious but with the proviso that bound variables are renamed to fresh variables before the application of the substitution.

When needed, we will use Smullyan's uniform notation for first-order formulas (see, e.g., (Fitting, 1996)), which classifies them into formulas of type $\alpha, \beta, \gamma, \delta$ according to the tables in Figure 1.

We will also use the usual notions of satisfiability and entailment but extended to the case of formulas, as opposed to sentences, as done in the mathematical logic literature. Specifically, a set Φ of Σ -formulas is *satisfiable* in a Σ -structure \mathcal{A} , if there is a valuation θ of V into elements of \mathcal{A} that makes every formula in Φ true in \mathcal{A} . In that case, we say that θ *satisfies* Φ in \mathcal{A} . If every valuation θ of V into \mathcal{A} satisfies Φ in \mathcal{A} , we say that \mathcal{A} is a *model* of Φ .³ A Σ -formula φ is *satisfiable* in \mathcal{A} if $\{\varphi\}$ is satisfiable in \mathcal{A} . A set Φ of formulas (resp. a formula) is *satisfiable* if it is satisfiable in some structure \mathcal{A} , and it is *unsatisfiable* otherwise.

For all sets Φ, Ψ of formulas, Φ *entails* Ψ , in symbols $\Phi \models \Psi$, if for every structure \mathcal{A} in the signature of $\Phi \cup \Psi$ and valuation θ , θ satisfies Ψ in \mathcal{A} whenever it satisfies Φ in \mathcal{A} .⁴ The set Φ *entails* the formula φ , in symbols $\Phi \models \varphi$, if $\Phi \models \{\varphi\}$; equivalently, $\Phi \models \varphi$ if the set $\Phi \cup \{\neg\varphi\}$ is unsatisfiable. Notice that Φ is unsatisfiable if and only if $\Phi \models \perp$.

The reader unfamiliar with this notion of satisfiability/entailment for formulas should observe that in it free variables essentially behave as free constant symbols (they are *rigid*). This differs from the common practice in automated reasoning of treating free variables as implicitly universally quantified—perhaps as a consequence of the fact that clauses are written without their universal quantifier prefix. The distinction here is important and should be kept in mind because, for example, when we talk about the satisfiability of a clause we are quantifying its variables universally, whereas when we talk about the satisfiability of a quantifier-free formula we are not. This means for instance that if $\{p(x), \neg p(y)\}$ denotes a set of (unit) *clauses*, it is unsatisfiable; if it denotes instead a set of *literals* (no implicit universal quantifiers), it is satisfiable.

³ Note that if Φ contains only sentences, \mathcal{A} is a model of Φ iff Φ is satisfiable in \mathcal{A} .

⁴ Note that in $\Phi \models \Psi$, the set Ψ is essentially seen as the *conjunction* of its elements, not as the disjunction as often found in the literature.

In theory reasoning, satisfiability and entailment are also given with respect to a certain theory. The definitions below subsume the various, not always equivalent ones in the literature.

DEFINITION 2.1. *Let \mathcal{T} be any theory. A set Φ of formulas is \mathcal{T} -satisfiable iff $\mathcal{T} \cup \Phi$ is satisfiable; otherwise, it is \mathcal{T} -unsatisfiable. The set Φ is literally \mathcal{T} -(un)satisfiable if the set $\{\varphi \in \Phi \mid \varphi \text{ is a literal}\}$ is \mathcal{T} -(un)satisfiable. The set Φ \mathcal{T} -entails a set Ψ of formulas, in symbols $\Phi \models_{\mathcal{T}} \Psi$, iff $\mathcal{T} \cup \Phi \models \Psi$.*

If φ is a formula, we will write $\Phi \models_{\mathcal{T}} \varphi$ whenever $\Phi \models_{\mathcal{T}} \{\varphi\}$. We will say that φ is \mathcal{T} -valid if $\emptyset \models_{\mathcal{T}} \varphi$. It is a simple exercise to show that, similarly to the entailment relation \models , the relation $\models_{\mathcal{T}}$ defined above is monotonic and transitive.

Traditionally, authors in theory reasoning define \mathcal{T} -satisfiability only for clauses (in terms of satisfiability of their ground instances in Herbrand models of \mathcal{T}) and then use another notion, \mathcal{T} -complementarity, for quantifier-free formulas. In essence, they say that a set S of quantifier-free formulas is \mathcal{T} -complementary whenever it is \mathcal{T} -unsatisfiable in the sense of Definition 2.1, and they say that it is \mathcal{T} -unsatisfiable whenever the set of the universal closures of the elements of S is \mathcal{T} -unsatisfiable again in the sense of Definition 2.1. We find it more convenient⁵ to adopt a single notion of \mathcal{T} -satisfiability, the one in Definition 2.1, and simply be careful in distinguishing (genuinely) free variables from implicitly universally quantified variables.

A formula is *universal* if it is in prefix normal form and its (possibly empty) quantifier prefix contains only universal quantifiers. A theory is universal if it is axiomatized by a set of universal sentences.⁶ Theory reasoning considers only universal theories as background theories because they are the only ones that can be “safely” built-in into a deduction calculus. The reason is that one of the main tools for proving properties of deduction calculi for automated reasoning, the Herbrand theorem, extends immediately to \mathcal{T} -satisfiability in a universal theory.

In this paper, we will not appeal to the Herbrand Theorem directly. Instead, we will use the satisfiability properties of a theory version of *Hintikka sets*. Hintikka sets (whether they are referred to as such or not) are often used to prove the completeness of calculi for first-order logic (see (Fitting, 1996; Socher-Ambrosius and Johann, 1997), among others).

⁵ And more in line with the established practice in Model Theory, which after all, provides all the semantical foundations for automated reasoning.

⁶ Some authors refer to universal theories as *open* or *quantifier-free theories*, again because of the common practice of writing their axioms without the quantifiers.

DEFINITION 2.2 (\mathcal{T} -Hintikka set). *Let \mathcal{T} be a theory of signature Ω and let Ω' be a signature including Ω . A set H of Ω' -sentences is a \mathcal{T} -Hintikka set iff the following holds:*

1. H is literally \mathcal{T} -satisfiable.
2. For all sentences $\alpha \in H$, $\alpha_1, \alpha_2 \in H$.
3. For all sentences $\beta \in H$, $\beta_1 \in H$ or $\beta_2 \in H$.
4. For all sentences $\gamma \in H$ and ground Ω' -terms t , $\gamma_1(y)\{y \mapsto t\} \in H$.
5. For all sentences $\delta \in H$, there is a ground Ω' -term t such that $\delta_t \in H$.

The usual definition of Hintikka set differs from the one above only in Point 1 where it requires instead that H contain neither \perp nor a complementary pair of (ground) literals. It should be clear that every \mathcal{T} -Hintikka set is a Hintikka set in the usual sense and, by compactness of first-order logic, every Hintikka set is a \mathcal{T} -Hintikka set when \mathcal{T} is the empty theory.

In FOL without equality, a *Herbrand structure* of some signature Ω is a structure whose domain coincides with the set of ground Ω -terms and that interprets every ground Ω -term as itself. In FOL with equality, the notion of Herbrand structure is generalized into that of *canonical structure*. A canonical structure (of signature Ω) is a structure each of whose elements is denoted by some ground Ω -term; equivalently, it is a structure generated by the empty set. Now, it is well-known that every Hintikka set has a canonical model (Fitting, 1996). That is also true for \mathcal{T} -Hintikka sets, provided that \mathcal{T} is universal. More precisely, the following holds.

LEMMA 2.3. *If \mathcal{T} is a satisfiable universal theory, then every \mathcal{T} -Hintikka set is satisfiable in a canonical model of \mathcal{T} .*

A proof of this result is provided in the appendix.

3. Theory Reasoning over Multiple Theories

Because of its generality, theory reasoning encompasses a vast array of seemingly different reasoning frameworks. Here we will focus on what (Baumgartner et al., 1992) calls *literal level theory reasoning*.⁷ The

⁷ The other forms of theory reasoning can be recast as essentially special cases of the literal level form.

basic operation in traditional refutation-based calculi is the detection of pairs of complementary literals; in other words, the detection of an unsatisfiable set Φ made of two quantifier-free formulas of a specific kind. Literal level theory reasoning generalizes this operation in two directions: the type and number of quantifier-free formulas in Φ , and the notion of (un)satisfiability, defined with respect to a certain *background theory* \mathcal{T} .

3.1. PARTIAL THEORY REASONING

In theory reasoning systems, the \mathcal{T} -satisfiability test is not performed by the main reasoner, the *foreground reasoner*, but is delegated instead to a specialized subsystem, the *background reasoner* for \mathcal{T} . At the ground level, we speak of *total theory reasoning* if the background reasoner gets a set Φ of formulas from the foreground one and simply confirms whether Φ is \mathcal{T} -unsatisfiable or not; we speak of *partial theory reasoning* if, whenever Φ is not \mathcal{T} -unsatisfiable, the background reasoner returns a *residue* for it, that is, a quantifier-free formula whose negation, if added to Φ , would make it \mathcal{T} -unsatisfiable. At the non-ground level, things are a bit more complicated because they involve the computation of substitutions that make Φ (partially) \mathcal{T} -unsatisfiable.

The precise general definition of residue varies in the literature, depending on the author and the partial theory reasoning calculus in question. But they are all instances of the one below.

DEFINITION 3.1 (Residue). *Let Φ be a set of quantifier-free formulas, called a key set. Let φ be a quantifier free formula and σ a substitution with $\text{Dom}(\sigma) \subseteq \text{Var}(\Phi)$ such that $\varphi\sigma = \varphi$. The pair (σ, φ) is a \mathcal{T} -residue of Φ iff the set $\Phi\sigma \cup \{\neg\varphi\}$ is \mathcal{T} -unsatisfiable or, equivalently, iff $\Phi\sigma \models_{\mathcal{T}} \varphi$.*

According to the definition above, the pair (σ, \perp) is a \mathcal{T} -residue of the key set Φ if and only if $\Phi\sigma$ is \mathcal{T} -unsatisfiable. More precisely then, we talk of total theory reasoning when the background reasoner computes only residues of the form (σ, \perp) , if any, and of partial theory reasoning otherwise.

In the following, we will simply say *residue* instead of \mathcal{T} -residue whenever \mathcal{T} is clear from context. Abusing the terminology, we will also call residue the second component of a residue (σ, φ) , especially when σ is the empty substitution.

3.2. COMBINING BACKGROUND REASONERS

Suppose we are interested in a background theory \mathcal{T} obtained as the union of $n > 1$ theories $\mathcal{T}_1, \dots, \mathcal{T}_n$. Also suppose that we do not have a background reasoner for \mathcal{T} but we do have one for each \mathcal{T}_i . From a practical standpoint, instead of implementing a reasoner for \mathcal{T} anew, it would be useful to integrate the reasoners for the various \mathcal{T}_i directly into a foreground reasoner and have them work together to detect the \mathcal{T} -unsatisfiability of formulas. The question then is how to make the reasoners cooperate in a sound and complete way.

In this section, we provide some interpolation results which suggest that background reasoners can cooperate by exchanging residues over a common quantifier-free language. In the next section, we will embed this kind of cooperation into a specific theory reasoning calculus and show that the resulting calculus is sound and complete. For simplicity, we will consider the case of just two background theories. From what follows, however, it should be clear that all the given results lift by iteration to the case of more than two theories.

We will impose no model-theoretic restrictions on the two theories other than universality. Also, we will make no assumptions on whether the theories share no, some or all predicate symbols. However, we will make the following assumption.

ASSUMPTION 1. *All the background theories to be combined will have exactly the same function symbols.*

For our purposes such an assumption is not as stringent as it sounds, at least in the case of refutation-based theory reasoning calculi.⁸ In fact, background reasoners used in such calculi must accept input formulas containing *Skolem* symbols, i.e., fresh function symbols produced by the Skolemization of existential variables. Technically then, all background theories in theory reasoning have a signature with arbitrarily many function symbols of arity n for every $n \geq 0$, even when the theories themselves are finitely axiomatized.⁹ More prosaically, if we make the very reasonable assumption that a reasoner for a background theory treats every *unknown* function symbol as a Skolem symbol, then we can always pass to it formulas with function symbols from some other

⁸ Actually, we are not aware of any theory reasoning calculus that is not refutation-based.

⁹ This is also the reason we assume without loss of generality that every signature we consider contains a designated constant symbol a (see Section 2). Also, observe that we said “arbitrarily many” function symbols and not “infinitely many”. The reason is that, in any given derivation in a refutation-based calculus, the number of Skolem symbols needed is always finite.

theory. This means that, given two background theories with reasoners of this sort, the function symbols of one theory can be always thought with no loss of generality as belonging to the signature of the other.

3.3. THE INTERPOLATION LEMMA

For the rest of this section we fix two signatures Σ_1, Σ_2 such that $\Sigma_1^F = \Sigma_2^F$, and two universal theories $\mathcal{T}_1, \mathcal{T}_2$ of respective signature Σ_1, Σ_2 . Also, let $\Sigma := \Sigma_1 \cap \Sigma_2$.

The main theoretical result of the paper is provided by the following restricted version of the Craig Interpolation Lemma, whose proof can be found in the appendix.

PROPOSITION 3.2 (Ground Interpolation Lemma). *If $\mathcal{T}_1 \cup \mathcal{T}_2$ is unsatisfiable, then*

$$\mathcal{T}_1 \models \varphi \text{ and } \mathcal{T}_2 \models \neg\varphi$$

for some ground Σ -formula φ .

We call the formula φ an *interpolant* of \mathcal{T}_1 and \mathcal{T}_2 . Note that, although this notion is not symmetric in \mathcal{T}_1 and \mathcal{T}_2 , φ is an interpolant of \mathcal{T}_1 and \mathcal{T}_2 iff $\neg\varphi$ is an interpolant of \mathcal{T}_2 and \mathcal{T}_1 .

For our purposes, the following corollary of Proposition 3.2 will be more useful.

PROPOSITION 3.3. *For $i = 1, 2$ let Φ_i be a set of Σ_i -literals. Then, the following are equivalent:*

1. $\Phi_1 \cup \Phi_2$ is $(\mathcal{T}_1 \cup \mathcal{T}_2)$ -unsatisfiable;
2. there is a finite set Ψ of disjunctions of Σ -literals with $\text{Var}(\Psi) \subseteq \text{Var}(\Phi_1 \cup \Phi_2)$ such that

$$\Phi_1 \models_{\mathcal{T}_1} \Psi \text{ and } \Phi_2 \cup \Psi \models_{\mathcal{T}_2} \perp .$$

Proof. (1 \Rightarrow 2) Let $X = \text{Var}(\Phi_1) \cup \text{Var}(\Phi_2)$. For $i = 1, 2$, let $\Omega_i := \Sigma_i(X)$ and consider Φ_i as a set of ground Ω_i -formulas. Then note that $\mathcal{T}_i \cup \Phi_i$ is a universal Ω_i -theory and $\Omega_1^F = \Omega_2^F$. By an application of Proposition 3.2 then, there is a ground Ω -formula φ such that $\mathcal{T}_1 \cup \Phi_1 \models \varphi$ and $\mathcal{T}_2 \cup \Phi_2 \models \neg\varphi$; equivalently, such that $\Phi_1 \models_{\mathcal{T}_1} \varphi$ and $\Phi_2 \cup \{\varphi\} \models_{\mathcal{T}_2} \perp$. The claim then follows by assuming, with no loss of generality, that φ is in conjunctive normal form and choosing Ψ to be the set of φ 's conjuncts.

(2 \Rightarrow 1) By the monotonicity of entailment (\models), it is immediate that $\Phi_1 \models_{\mathcal{T}_1 \cup \mathcal{T}_2} \Psi$ and $\Phi_2 \cup \Psi \models_{\mathcal{T}_1 \cup \mathcal{T}_2} \perp$. From the transitivity and the monotonicity of the relation $\models_{\mathcal{T}_1 \cup \mathcal{T}_2}$ it follows that $\Phi_1 \cup \Phi_2 \models_{\mathcal{T}_1 \cup \mathcal{T}_2} \perp$. \square

$\frac{\alpha}{\alpha_1}$	$\frac{\beta}{\beta_1 \mid \beta_2}$	$\frac{\gamma}{\gamma_1(y)}$ where y is a fresh free variable
α_2		
$\frac{\delta}{\delta_{f(\vec{x})}}$	where f is a fresh function symbol in $\Omega^{\text{sko}} \setminus \Omega$ and $\vec{x} = (x_1, \dots, x_n)$ with $\{x_1, \dots, x_n\} = \text{Var}(\delta)$	

Figure 2. Tableau Expansion Rules.

Note that the existence of an *interpolant set* Ψ for Φ_1 and Φ_2 above is already guaranteed by the Craig Interpolation Lemma. The contribution of Proposition 3.3 is to show that this set can be chosen so that all of its formulas are disjunctions of literals with no new variables.

Finally, we point out that Proposition 3.3 holds as stated in both flavors of first-order logic: the one with equality and the one without equality. The only difference is that, whereas in the first flavor the formulas of the interpolant set Ψ might contain equations, in the second flavor they will not.¹⁰ This entails in particular that, in FOL without equality, if the theories share no relation symbols at all, the only possible interpolant set of Φ_1 and Φ_2 in Proposition 3.3 is either $\{\perp\}$ or $\{\neg\perp\}$.

4. A Multi-Theory Tableau Calculus

The interpolation results of the previous section can be used to integrate multiple background reasoners into a theory reasoning calculus. In this section, we define a multi-reasoner extension of the partial theory version of free variable semantic tableaux.

We do this in two stages. First, we provide a generalized (partial) theory reasoning version of the semantic tableau calculus. Like every theory reasoning extensions of existing refutation calculi, this generalized calculus replaces the search for two complementary literals with the search for a \mathcal{T} -unsatisfiable key set, for some background theory \mathcal{T} . Then, we show that when \mathcal{T} is in fact the union of two theories \mathcal{T}_1 and \mathcal{T}_2 , it is enough to look only for \mathcal{T}_1 -unsatisfiable or \mathcal{T}_2 -unsatisfiable key sets. The main consequence of this fact is that a stand-alone background

¹⁰ Unless, of course, an equality predicate is explicitly axiomatized in one of the two theories, and the equality symbol used in the axiomatization also belongs to the signature of the other theory.

reasoner for \mathcal{T} is in principle not necessary if background reasoners for \mathcal{T}_1 and for \mathcal{T}_2 are already available.¹¹

Our treatment of the free variable semantic tableau calculus will follow closely the one given in (Fitting, 1996).

4.1. FREE VARIABLE SEMANTIC TABLEAUX

A tableau is a finite tree each of whose nodes is labeled by a formula. Since tableaux, for being trees, are directed acyclic graphs, we will represent every tableau S as the pair (V, E) where V is the set of S 's nodes and E is the set of S 's (directed) edges. In the following, to simplify the exposition, we blur the distinction between a node and the formula that labels it. Technically, this is incorrect because it is certainly possible for a tableau to have distinct nodes with the same label, but it will simplify our exposition. This imprecision should cause no problems if the distinction between a node and its label is kept in mind.

We denote an edge from a formula φ to a formula ψ in a tableau by the ordered pair $\langle \varphi, \psi \rangle$. A directed path from the root node of a tableau to one of its leaf nodes is called a *branch*. We denote by $leaf(B)$ the leaf node of a branch B . If B is a branch and σ a substitution, we denote by $B\sigma$ the branch obtained by replacing each node φ in B by the node $\varphi\sigma$ —similarly for set of branches or edges in a tableau. For notational convenience, we will often treat a branch B as the multiset of formulas in its nodes. A tableau branch is *closed* if it contains the node \perp and *open* otherwise. A tableau is *closed* if all of its branches are closed. In the following, the letters j and n will denote finite ordinal numbers whereas the letter κ will denote an ordinal smaller than or equal to the first infinite ordinal. For every κ then, we will denote a (possibly infinite) sequence a_0, a_1, a_2, \dots of κ elements by $(a_j)_{j < \kappa}$.

For the rest of this section we will fix a signature Ω and a *satisfiable* universal Ω -theory \mathcal{T} . Also, we will implicitly assume for all tableaux mentioned below that the signature of their formulas is (included in) Ω^{sko} .

DEFINITION 4.1. *Let $S = (V, E)$ and S' be two tableaux. We say that S' \mathcal{T} -derives from S iff S' is obtained from S in one of the following ways:*

1. *by applying one of the expansion rules in Fig. 2 to a formula φ in a branch B of S , i.e., by defining S' as follows for some $\varphi \in B$:*

¹¹ We must say “in principle” here because we are ignoring efficiency concerns.

$$S' := \begin{cases} (V \cup \{\alpha_1, \alpha_2\}, E \cup \{\langle \text{leaf}(B), \alpha_1 \rangle, \langle \alpha_1, \alpha_2 \rangle\}) & \text{if } \varphi \text{ has type } \alpha \\ (V \cup \{\beta_1, \beta_2\}, E \cup \{\langle \text{leaf}(B), \beta_1 \rangle, \langle \text{leaf}(B), \beta_2 \rangle\}) & \text{if } \varphi \text{ has type } \beta \\ (V \cup \{\gamma_1(y)\}, E \cup \{\langle \text{leaf}(B), \gamma_1(y) \rangle\}) & \text{if } \varphi \text{ has type } \gamma \\ (V \cup \{\delta_{f(\vec{x})}\}, E \cup \{\langle \text{leaf}(B), \delta_{f(\vec{x})} \rangle\}) & \text{if } \varphi \text{ has type } \delta \end{cases}$$

where y and $f(\vec{x})$ above are defined as in Figure 2, or

2. by adding a \mathcal{T} -residue to a branch B of S , i.e., by defining S' as follows:

$$S' := (V \cup \{\psi\}, E \cup \{\langle \text{leaf}(B), \psi \rangle\})\sigma$$

where (σ, ψ) is a \mathcal{T} -residue of some Ω -key set $\Phi \subseteq B$.

The definition above reduces to the usual one given for non-theory tableaux if \mathcal{T} is the empty theory. In that case, the key sets of interest consist of two literals with the same predicate symbol and opposite sign, and all residues have the form (σ, \perp) where σ is a (most general) syntactical unifier of the key set.

DEFINITION 4.2 (Derivation). *A (possibly infinite) sequence $(S_j)_{j < \kappa}$ of κ tableaux is a (tableau) \mathcal{T} -derivation iff for all $j > 0$, S_j derives from S_{j-1} .*

We say that a tableau branch B' *extends* a tableau branch B if there is a substitution σ such that $B\sigma$, as a path, is an initial segment of B' (possibly coinciding with B'). It is easy to verify that every branch of a tableau in a \mathcal{T} -derivation extends one, and only one, branch in each previous tableau in the derivation.

DEFINITION 4.3 (Proof). *Let φ be an Ω -sentence. A \mathcal{T} -derivation $(S_j)_{j < \kappa}$ is a (tableau) \mathcal{T} -derivation of φ iff S_0 is a tableau whose only node is $\neg\varphi$. The derivation $(S_j)_{j < \kappa}$ is a (tableau) \mathcal{T} -proof of φ iff there is an $n < \kappa$ such that S_n is closed.*

For convenience, we will simply say *derives*, *derivation*, *proof* and so on in place of \mathcal{T} -derives, \mathcal{T} -derivation, \mathcal{T} -proof whenever \mathcal{T} is clear from context or not important.

The tableau calculus induced by the above notions of derivation and proof is sound and complete in the sense that an Ω -formula φ is \mathcal{T} -valid iff there is a tableau proof for it. The soundness argument is very similar to that for non-theory tableaux. We provide a proof below, concentrating on the residue rule. As for completeness, we will actually show that a restricted version of the calculus is already complete. The most important restriction will concern the residue rule, as we will see.

Soundness

As usual, every tableau S can be seen as a (Ω^{sko} -)sentence: the universal closure of the disjunction of all the branches of S , where each branch is seen as the conjunction of the formulas in it. For simplicity, we will denote this sentence just by $\tilde{\forall} S$.

To prove the soundness of the calculus, it is enough to show that derivations preserve the \mathcal{T} -satisfiability of tableaux, when seen as sentences.

LEMMA 4.4. *Let S, S' be two tableaux such that S' \mathcal{T} -derives from S . If $\tilde{\forall} S$ is \mathcal{T} -satisfiable then $\tilde{\forall} S'$ is also \mathcal{T} -satisfiable.*

Proof. If S' derives from S by means of an expansion rule (cf. Definition 4.1(1)), bar the restriction to the Ω^{sko} -models of \mathcal{T} only, the claim is proved exactly as for (non-theory) free-variable semantic tableaux.¹² Suppose then that S' derives from S by an application of the residue rule. Let B, Φ, σ, ψ be as specified in Definition 4.1(2) and let φ_B be the conjunction of all the formulas in B .

Then, $\tilde{\forall} S$ is (logically equivalent to) a sentence of the form $\tilde{\forall}(\varphi \vee \varphi_B)$ and $\tilde{\forall} S'$ is (logically equivalent to) a sentence of the form $\tilde{\forall}(\varphi \sigma \vee (\varphi_B \sigma \wedge \psi))$. Now assume that $\tilde{\forall} S$ is \mathcal{T} satisfiable, and so there is a Ω^{sko} -model \mathcal{A} of \mathcal{T} such that $\varphi \vee \varphi_B$ is satisfied by every valuation into \mathcal{A} . From this it follows immediately that $\varphi \sigma \vee \varphi_B \sigma$ is also satisfied by every valuation into \mathcal{A} . Then, any such valuation θ satisfies either $\varphi \sigma$ or $\varphi_B \sigma$ in \mathcal{A} . Now, if θ satisfies $\varphi \sigma$ it clearly satisfies $\varphi \sigma \vee (\varphi_B \sigma \wedge \psi)$. If θ satisfies $\varphi_B \sigma$, it must satisfy ψ as well. The reason is that $\varphi_B \sigma \models_{\mathcal{T}} \psi$ by definition of residue and the inclusion $\Phi \sigma \subseteq B \sigma$. It follows that then θ satisfies $\varphi_B \sigma \wedge \psi$. Either way, θ satisfies $\varphi \sigma \vee (\varphi_B \sigma \wedge \psi)$ in \mathcal{A} . Since θ was arbitrary, we have that $\tilde{\forall}(\varphi \sigma \vee (\varphi_B \sigma \wedge \psi))$ is satisfiable in \mathcal{A} , which means that $\tilde{\forall} S'$ is \mathcal{T} -satisfiable. \square

The soundness of the calculus follows immediately from the lemma above.

PROPOSITION 4.5 (Soundness). *Every Ω -formula φ that has a tableau \mathcal{T} -proof is \mathcal{T} -valid.*

Completeness

We will now show that our tableaux calculus is complete even if subject to a number of restrictions on the possible derivations. Some of these

¹² See, e.g., (Fitting, 1996). There, a tableau S is called \forall -satisfiable if $\tilde{\forall} S$ is satisfiable in our sense.

restrictions correspond to the usual ones found in non-theory tableaux: each derivation is *strict*, in the sense that no occurrence of non- γ formula in a branch is used more than once to expand that branch; each derivation is constructed in a *fair* way, in the sense that every occurrence of a non-literal formula in a branch has a chance to be expanded later on in the derivation, and γ formulas have a chance to be expanded arbitrarily often.

The other restrictions are specific to our theory reasoning extension and hence concern the residue rule. We will see that it is enough to consider only key sets that are sets of literals, and residues that are disjunctions of literals. In addition, if the background theory \mathcal{T} is the union of two theories $\mathcal{T}_1, \mathcal{T}_2$ with respective signatures Σ_1, Σ_2 sharing all of their functions symbols, key sets can be chosen to contain only Σ_1 -literal or only Σ_2 -literals, and residues that contain predicate symbols not in $\Sigma_1 \cap \Sigma_2$ or variables not in the current branch can be ignored.

DEFINITION 4.6. *Let $(S_j)_{j \leq n}$ be a derivation and B a branch of S_n . A node φ of B is reduced in B iff φ is a literal, or it is of type α, β or δ and there is a $j \in \{1, \dots, n\}$ such that S_j derives from S_{j-1} by the application of an expansion rule to φ in the branch of S_{j-1} extended by B .*

Note that according to the definition above, nodes of type γ are never reduced in any branch. We will use the following fact about reducible nodes in a tableaux branch of a derivation.

LEMMA 4.7. *Let $(S_j)_{j \leq n}$ be a derivation, B a branch of S_n and φ a reduced node of B . Then,*

- if φ is of type α , then both α_1 and α_2 are in B ;
- if φ is of type β , then either β_1 or β_2 is in B ;
- if φ is of type δ , then δ_t is in B for some Ω^{sko} -term t .

Proof. By a straightforward induction argument based on the definition of the expansion rules. \square

DEFINITION 4.8 (Strict Derivation). *A derivation is strict iff none of its tableaux derives from the previous one by the application of an expansion rule to a reduced node in a branch.*

DEFINITION 4.9 (Fair Derivation). *A derivation $(S_j)_{j < \kappa}$ is fair iff for every $j < \kappa$, every branch B of S_j , every formula φ in B and every $m > 0$, there is a tableau $S_{j'}$ with $j \leq j' < \kappa$ such that either φ is reduced in the branch B' of $S_{j'}$ extending B or—when φ is of type γ — φ has m (distinct) γ_1 instances in B' .*

It is easy to show that a derivation that is both strict and fair is infinite if, and only if, some tableau in the derivation contains a node of type γ .

The next lemma shows that every \mathcal{T} -valid sentence has a strict and fair derivation of a certain restricted form.

LEMMA 4.10. *Let φ be an Ω -sentence. Then, the following holds:*

1. *there is a strict and fair tableau derivation $(S_j)_{j < \kappa}$ of φ such that, for all $0 < j < \kappa$, S_j derives from S_{j-1} by means of a tableau expansion rule;*
2. *if φ \mathcal{T} -valid, then for each such derivation there is an $n < \kappa$ and a substitution σ such that every branch of $S_n\sigma$ is literally \mathcal{T} -unsatisfiable.*

Proof. Let \mathbf{D} be the set of all fair and strict derivations of φ as in Point 1 above. One can prove that \mathbf{D} is non-empty by using one of the usual *fair tableau construction rules* from the literature. Therefore, we prove just Point 2 here and refer the reader to (Fitting, 1996) for a proof of Point 1.¹³ We prove the claim in Point 2 by proving its contrapositive.

Assume that every derivation $(S_j)_{j < \kappa}$ in \mathbf{D} is such that

$$\text{for all } n < \kappa \text{ and } \sigma, S_n\sigma \text{ has a literally } \mathcal{T}\text{-satisfiable branch.} \quad (1)$$

We prove below that then $\neg\varphi$ is \mathcal{T} -satisfiable, which entails that φ is not \mathcal{T} -valid.

Given any derivation $(S_j)_{j < \kappa}$ in \mathbf{D} , let $S^* := (\bigcup_{j < \kappa} V_j, \bigcup_{j < \kappa} E_j)$, where for each $j < \kappa$, V_j is the set of nodes and E_j the set of edges of S_j . Note that S^* is itself a tree (albeit a possibly infinite one) and that it extends each S_j .

Now let X be the set of S^* 's free variables. Observe that (a) the first tableau in the derivation has no free variables, (b) a free variable can occur in a later tableau only as the result of the expansion of a γ formula, and (c) each such expansion introduces a fresh variable. This means that for each $x \in X$, if any, there is an $n > 0$ such that x *first occurs* in S_n —in the sense that it occurs in S_n but not in S_{n-1} . Moreover, x only occurs in S_n in a leaf $\gamma_1(x)$. Where B is the branch of S_n with leaf of the form $\gamma_1(x)$, let $d(x)$ be the number of free variables y other than x such that $\gamma_1(y)$ occurs in B . Then, let t_0, t_1, \dots be any

¹³ A fine but important point to notice here is that *every* derivation generated according to a fair tableau construction rule is fair and strict in our sense. In practice, this means that no search is necessary to produce a derivation that belongs to \mathbf{D} above.

enumeration of all the ground Ω^{sko} -terms and let σ be a substitution such that $x\sigma = t_{d(x)}$ for all $x \in X$. It should be clear from the above that σ is well-defined over X (although possibly non-injective).

To start with, we claim that at least one branch of $S^*\sigma$ is literally \mathcal{T} -satisfiable. In fact, assume by contradiction that every branch B^* of $S^*\sigma$ is literally \mathcal{T} -unsatisfiable. By compactness, for each B^* then there is a finite set Φ_{B^*} of literals included in B^* that is \mathcal{T} -unsatisfiable. It follows by construction of S^* that there is an $n < \kappa$ such that each set Φ_{B^*} is included in a branch of $S_n\sigma$. But that means that no branch of $S_n\sigma$ is literally \mathcal{T} -satisfiable, against (1) above.

Now, let $B^*\sigma$ be a literally \mathcal{T} -satisfiable branch of $S^*\sigma$ and assume that a formula α occurs in B^* . By construction of B^* , α occurs in some tableau S_j of the derivation in the branch B_j of S_j extended by B^* . Since the derivation is fair there is an n with $j \leq n < \kappa$ such that α is reduced in the branch of S_n extending B_j (and extended by B^*). By Lemma 4.7, α_1 and α_2 occur in that branch and so in B^* . In a similar way, we can show that if a formula β occurs in B^* then either β_1 or β_2 occurs in B^* , and if a formula δ occurs in B^* then some δ_t occurs in B^* for some Ω^{sko} -term t .

Finally, if a formula γ occurs in B^* , again by the fairness of the derivation, it is not difficult to show that B^* contains infinitely-many distinct variants of the formula $\gamma_1(x)$ —where x is replaced by a different variable. In addition, for each $i \geq 0$ there is a variant $\gamma_1(y_i)$ such that $d(y_i) = i$. By construction of the substitution σ , it follows that $\gamma_1(y_i)\{y_i \mapsto t_i\}$ occurs in $B^*\sigma$ for all $i \geq 0$, where t_0, t_1, \dots is the enumeration chosen earlier of all ground Ω^{sko} -terms.

All this shows that $B^*\sigma$ (which, notice, contains only Ω^{sko} -sentences) is a \mathcal{T} -Hintikka set, and so it is \mathcal{T} -satisfiable by Lemma 2.3. Now, $\neg\varphi$ belongs to $B^*\sigma$ because it is the root node of S^* and it equals $(\neg\varphi)\sigma$ for having no free variables. It follows that $\neg\varphi$ is \mathcal{T} -satisfiable, as claimed. \square

The completeness of our theory tableau calculus already follows from the lemma above. In fact, each branch B of S_n in the lemma, given that $B\sigma$ is literally unsatisfiable, must contain a set of literals that admits $\langle\sigma, \perp\rangle$ as a \mathcal{T} -residue. But then S_n can be turned into a closed tableau by finitely many applications of the residue rule. The first of these applications closes one branch of S_n by adding \perp to it and applying the substitution σ to the resulting tableau. Each of the following applications closes one of the remaining branches by simply adding \perp to them—the reason being that, after the substitution, all the branches admit $\langle\varepsilon, \perp\rangle$ as a residue.

Notice that this completeness argument is a direct generalization of the one given for the usual (non-theory) semantic tableaux calculus. There, each branch of S_n is closed by finding a substitution σ , computed by a simultaneous unification algorithm, that when applied to the tableau makes one pair of literals in each branch complementary. In effect then, the non-theory tableau calculus is an instance of the calculus described above in which the background theory \mathcal{T} is empty and the background reasoner used to compute \mathcal{T} -residues is just a procedure for (syntactic) unification.

In the general case, where \mathcal{T} is not necessarily empty, it is sufficient to have a background reasoner for \mathcal{T} that is able to enumerate all the \mathcal{T} -residues of signature Ω^{sko} for each given Ω^{sko} -key set. From that and Lemma 4.10, one can show that every \mathcal{T} -valid formula has a strict and fair proof. We show below that it is *not necessary* to have a background reasoner for \mathcal{T} if \mathcal{T} is the union of two theories each with its own background reasoner.

We will do that by assuming that $\mathcal{T} := \mathcal{T}_1 \cup \mathcal{T}_2$ where, for $i = 1, 2$, \mathcal{T}_i is a universal theory of signature Σ_i equipped with a background reasoner that can enumerate the \mathcal{T}_i -residues of signature Σ_i^{sko} for each given Σ_i^{sko} -key set, where $\Sigma := \Sigma_1 \cap \Sigma_2$. Since we assume that the background reasoner for \mathcal{T}_i accepts arbitrary function symbols in its input formulas¹⁴, we will also assume—with no loss of generality as pointed out in Section 3.2—that Σ_1 and Σ_2 share all their function symbols, and we will treat each \mathcal{T}_i as Σ_i^{sko} -theory.

For the main completeness result, we need the following lemma as well.

LEMMA 4.11. *Let $i \in \{1, 2\}$ and let S_m be a tableau all of whose branches contain a \mathcal{T}_i -unsatisfiable set of quantifier-free Σ_i^{sko} -formulas. Then, there is a strict tableau derivation $(S_j)_{m \leq j \leq n}$ for which the following holds:*

1. *for all j with $m < j \leq n$, S_j derives from S_{j-1} by means of a tableau expansion rule;*
2. *every branch of S_n contains a \mathcal{T}_i -unsatisfiable set of Σ_i^{sko} -literals.*

Proof. It is a simple, if tedious, exercise to show that there is a strict tableau derivation $(S_j)_{m \leq j \leq n}$ satisfying Point 1 above and such that each quantifier-free node of S_n is reduced in every branch of S_n in which it occurs. We prove by contradiction that S_n satisfies Point 2 above.

¹⁴ Those in $\Sigma_i^{\text{sko}} \setminus \Sigma_i$.

Suppose that S_n has a branch B_n such that the set of all Σ_i^{sko} -literals in B_n is \mathcal{T}_i -satisfiable. If B_m is the branch of S_m extended by B_n let Φ be a \mathcal{T}_i -unsatisfiable set of quantifier-free Σ_i^{sko} -formulas occurring in B_m . Then, consider Φ as a set of ground formulas in the expanded signature $\Sigma_i^{\text{sko}}(V)$. Using the fact that the elements of Φ are reduced nodes of B_n and that each of their subformulas occurring in B is also reduced, we can show by Lemma 4.7 that Φ is contained in a \mathcal{T}_i -Hintikka set of $\Sigma_i^{\text{sko}}(V)$ -sentences. But then, as a set of $\Sigma_i^{\text{sko}}(V)$ -sentences, Φ is \mathcal{T}_i -satisfiable by Lemma 2.3, which contradicts the assumption that, as a set of Σ_i^{sko} -formulas, Φ is \mathcal{T}_i -unsatisfiable. \square

Let us say that an application of the residue rule is *restricted* to \mathcal{T}_1 and \mathcal{T}_2 iff for $i = 1$ or $i = 2$,

1. the key set Φ of the branch B chosen by the rule consists of Σ_i^{sko} literals only,
2. the added residue (σ, ψ) is a \mathcal{T}_i -residue and such that ψ is a disjunction of Σ^{sko} -literals all of whose variables occur in $B\sigma$.

Our tableau calculus is complete even if every application of the residue rule is restricted to \mathcal{T}_1 and \mathcal{T}_2 . But before proving this claim, let us see with an example how a derivation with such a restriction would look like.

EXAMPLE 4.12. Consider the universal theories \mathcal{T}_1 and \mathcal{T}_2 defined as follows and sharing the binary relation symbol R .

$$\mathcal{T}_1 := \left\{ \begin{array}{l} \forall u \forall v (P_1 u \wedge Q_1(v, v) \rightarrow R(u, v)), \\ \forall u \forall v (R(u, v) \rightarrow T_1(u, v)) \end{array} \right\}$$

$$\mathcal{T}_2 := \{ \forall v (P_2 v \wedge R(v, fv) \rightarrow R(fv, v)) \}$$

Now consider a tableau containing a branch B , represented below as the list of formulas in its nodes:

$$B = [\dots, P_1 x, P_2 x, Q_1(y, z), \neg T_1(y, x)]$$

It is not difficult to see that the substitution $\sigma := \{y \mapsto fx, z \mapsto fx\}$ is such that the subset $\{P_1 x, P_2 x, Q_1(fx, fx), \neg T_1(fx, x)\}$ of formulas in $B\sigma$ is $(\mathcal{T}_1 \cup \mathcal{T}_2)$ -unsatisfiable. The following describes a possible derivation that closes B with a sequence of applications of the residue rule restricted to \mathcal{T}_1 and \mathcal{T}_2 . At each step j below, \circ denotes list concatenation, $B^{(j)}$ is the current extension of B , $\Phi_i^{(j)}$ is a key set from $B^{(j)}$ in the signature of \mathcal{T}_i , and $(\sigma^{(j)}, \psi^{(j)})$ is a possible \mathcal{T}_i -residue of $\Phi_i^{(j)}$ in the shared signature.

$$\begin{aligned}
B^{(0)} &= B = [\dots, P_1x, P_2x, Q_1(y, z), \neg T_1(y, x)] \\
\Phi_1^{(0)} &= \{P_1x, Q_1(y, z)\} \\
\sigma^{(0)} &= \{z \mapsto y\} \\
\psi^{(0)} &= R(x, y) \\
\\
B^{(1)} &= (B^{(0)} \circ [\psi^{(0)}])\sigma^{(0)} \\
&= [\dots, P_1x, P_2x, Q_1(y, y), \neg T_1(y, x), R(x, y)] \\
\Phi_2^{(1)} &= \{P_2x, R(x, y)\} \\
\sigma^{(1)} &= \{y \mapsto fx\} \\
\psi^{(1)} &= R(fx, x) \\
\\
B^{(2)} &= (B^{(1)} \circ [\psi^{(1)}])\sigma^{(1)} \\
&= [\dots, P_1x, P_2x, Q_1(fx, fx), \neg T_1(fx, x), R(x, fx), R(fx, x)] \\
\Phi_1^{(2)} &= \{\neg T_1(fx, x), R(fx, x)\} \\
\sigma^{(2)} &= \{\} \\
\psi^{(2)} &= \perp \\
\\
B^{(3)} &= (B^{(2)} \circ [\psi^{(2)}])\sigma^{(2)} \\
&= [\dots, P_1x, P_2x, Q_1(fx, fx), \neg T_1(fx, x), R(x, fx), R(fx, x), \perp]
\end{aligned}$$

The example above was constructed for simplicity so that only residues containing a single literal would be needed to close the branch. In general, however, residues with a proper disjunctions of literals may be necessary in a proof, unless an additional condition is imposed on the theories. We will discuss this condition and its implications in the next section. First, we must prove our completeness claim.

PROPOSITION 4.13 (Completeness). *For every \mathcal{T} -valid sentence φ or signature $\Omega = \Sigma_1 \cup \Sigma_2$, there is a strict and fair tableau proof of φ in which every application of the residue rule is restricted to \mathcal{T}_1 and \mathcal{T}_2 .*

Proof. By Lemma 4.10, there is a strict and fair derivation $(S_j)_{j < \kappa}$ of φ , a substitution σ into Ω^{sko} -terms, and a $m < \kappa$. such that every branch of $S_m\sigma$ is literally \mathcal{T} -unsatisfiable. Let B one of the branches of S_m . First we show that there is a derivation starting with S_m that does not touch the other branches of S_m and extends B to a finite number of branches all of which contain a \mathcal{T}_2 -unsatisfiable set of quantifier-free Σ_2 -formulas.

Recall that $\Omega = \Sigma_1 \cup \Sigma_2$, $\Sigma_1^F = \Sigma_2^F$ and $\Sigma = \Sigma_1 \cap \Sigma_2$. Now, since $B\sigma$ is literally \mathcal{T} -unsatisfiable, for $i = 1, 2$, there must be a finite set Φ_i of Σ_i^{sko} -literals in B such that $\Phi_1\sigma \cup \Phi_2\sigma$ is \mathcal{T} -unsatisfiable. By Proposition 3.3, the following holds for some set $\Psi := \{\psi_1, \dots, \psi_l\}$ of disjunctions of Σ^{sko} -literals all of whose variables occur in $\Phi_1\sigma \cup \Phi_2\sigma$

and so in $B\sigma$.

$$\Phi_1\sigma \models_{\mathcal{T}_1} \Psi \quad (2)$$

$$\Phi_2\sigma \cup \Psi \models_{\mathcal{T}_2} \perp \quad (3)$$

By (2) $\Phi_1\sigma$ \mathcal{T}_1 -entails every formula in Ψ which means, say, that $\langle\sigma, \psi_1\rangle$ is a \mathcal{T}_1 -residue of Φ_1 and $\langle\varepsilon, \psi_k\rangle$ is a \mathcal{T}_1 -residue of $\Phi_1\sigma$ for all $k \in \{2, \dots, l\}$. From this it follows that there is a finite tableau derivation $(S_m, S_{m+1}, \dots, S_{m+l})$ such that (a) S_{m+1} is obtained, through an application of the residue rule, by first adding the formula ψ_1 to the branch B of S_m and then applying the substitution σ to the whole tableau, and (b) for all $k \in \{2, \dots, l\}$, S_{m+k} is obtained, also through an application of the residue rule, by adding the formula ψ_k to the branch of S_{m+k-1} that extends B .

The branch of S_{m+l} extending B then contains the set $\Phi_2\sigma \cup \Psi$ of quantifier-free Σ_2^{sko} -formulas, which is \mathcal{T}_2 -unsatisfiable by (3) above. Also, each of the other branches of S_{m+l} coincides with a branch of $S_m\sigma$ and so is literally \mathcal{T} -unsatisfiable. Using a similar argument for these branches we can then show that there is a finite (and strict) tableau derivation (S_m, \dots, S_{m+p}) such that every S_{m+k} (for $k \in \{1, \dots, p\}$) is obtained from S_{m+k-1} by an application of the residue rule restricted to \mathcal{T}_1 and \mathcal{T}_2 , and every branch of S_{m+p} contains a \mathcal{T}_2 -unsatisfiable set of quantifier-free Σ_2^{sko} -formulas.

Applying Lemma 4.11 to S_{m+p} and \mathcal{T}_2 , we can conclude that there is a strict derivation (S_m, \dots, S_{m+p+q}) such that every branch of S_{m+p+q} contains a \mathcal{T}_2 -unsatisfiable set of Σ_2^{sko} -literals. This is to say that each of these sets admits the \mathcal{T}_2 -residue $\langle\varepsilon, \perp\rangle$. But then, we can argue as before that there is a strict tableau derivation $(S_m, \dots, S_{m+p+q+r})$ such that $S_{m+p+q+r}$ is obtained by successive applications of the residue rule restricted to \mathcal{T}_2 , each adding \perp to an open branch of S_{m+p+q} until all of them are closed. The claim then follows by considering (any fair extension of) the derivation $(S_j)_{j \leq m+p+q+r}$. \square

In conclusion, we have proven the following about the calculus above.

THEOREM 4.14. *Let $\mathcal{T}_1, \mathcal{T}_2$ be two universal theories with respective signatures Σ_1, Σ_2 such that $\Sigma_1^{\text{F}} = \Sigma_2^{\text{F}}$ and $\mathcal{T} := \mathcal{T}_1 \cup \mathcal{T}_2$ is satisfiable. Let φ be a $(\Sigma_1 \cup \Sigma_2)$ -sentence. Then, φ is \mathcal{T} -valid iff there is a strict and fair tableau proof of φ in which every application of the residue rule is restricted to \mathcal{T}_1 and \mathcal{T}_2 .*

5. Refinements: The Convex Case

In this section we show that the Ground Interpolation Lemma of Section 3.3 can be further refined if the two theories \mathcal{T}_1 and \mathcal{T}_2 are also *convex* in a sense defined below. One consequence of this refinement is the possibility of strengthening the theory tableau calculus described in the previous section. As we will see, when two theories is convex, the completeness of the calculus is preserved even if one considers only *unit residues*, by which we mean residues of the form (σ, p) where p is a single literal. Such a restriction is significant because it considerably reduces the non-determinism of the residue rule, allowing more efficient implementations of the calculus.

The notion of theory convexity we adopt here is based on one due to Nelson and Oppen (Nelson and Oppen, 1979). It is also related to the notion of *independence of negative constraints* (see (Lassez and McAloon, 1990) for a general treatment) from the constraint programming literature.

DEFINITION 5.1 (Σ -Convex Theory). *Let Σ be a signature. A theory \mathcal{T} of signature Ω is Σ -convex iff for every set Φ of Ω -literals and every finite non-empty set Ψ of positive Σ -literals,*

$$\Phi \models_{\mathcal{T}} \bigvee_{p \in \Psi} p \quad \text{iff} \quad \Phi \models_{\mathcal{T}} p \quad \text{for some } p \in \Psi .$$

We prove in the appendix that universal Horn theories are Σ -convex for any Σ . There, we also provide references to examples of non-Horn universal theories that are Σ -convex for some signature Σ . This shows that an important portion of candidate background theories for theory reasoning are in fact convex, which justifies the particular relevance of the results presented in this section.

These results and their proofs will use Horn formulas. Following (Hodges, 1993a), we call a *basic Horn formula* a formula of the form

$$\neg p_1 \vee \cdots \vee \neg p_n \vee q$$

where $n \geq 0$ and each of p_1, \dots, p_n, q is a positive literal (possibly \perp).¹⁵

For the rest of this section, we will fix two signatures Σ_1, Σ_2 such that $\Sigma_1^F = \Sigma_2^F$ and two universal theories $\mathcal{T}_1, \mathcal{T}_2$ of respective signature Σ_1, Σ_2 such that $\mathcal{T}_1 \cup \mathcal{T}_2$ is satisfiable. Also, we assume that both theories is Σ -convex for $\Sigma := \Sigma_1 \cap \Sigma_2$.

¹⁵ Note that this definition rules out every disjunction $\neg p_1 \vee \cdots \vee \neg p_n$ of negative literals. However, we can treat such a disjunction as a basic Horn formula as well by identifying it with the logically equivalent formula $\neg p_1 \vee \cdots \vee \neg p_n \vee \perp$.

PROPOSITION 5.2 (Horn Ground Interpolation Lemma). *If $\mathcal{T}_1 \cup \mathcal{T}_2$ is unsatisfiable, then*

$$\mathcal{T}_1 \models \varphi \text{ and } \mathcal{T}_2 \models \neg\varphi$$

for some conjunction φ of ground basic Horn formulas of signature Σ .

Proof. By Proposition 3.2, there is a ground Σ -formula φ such that

$$\mathcal{T}_1 \models \psi \text{ and } \mathcal{T}_2 \models \neg\psi. \quad (4)$$

With no loss of generality we can assume that ψ has the conjunctive normal form $\psi_1 \wedge \cdots \wedge \psi_n$. For each $j \in \{1, \dots, n\}$, we can also assume that ψ_j has the form

$$\neg p_1^j \vee \cdots \vee \neg p_{m_j}^j \vee q_1^j \vee \cdots \vee q_{n_j}^j$$

where $m_j \geq 0$, $n_j \geq 1$, each p_k^j and each q_k^j is a positive ground literal, and q_1^j , say, is always \perp . Let $j \in \{1, \dots, n\}$. By (4) above, we have that $\mathcal{T}_1 \models \neg p_1^j \vee \cdots \vee \neg p_{m_j}^j \vee q_1^j \vee \cdots \vee q_{n_j}^j$. By the properties of logical entailment and the definition of $\models_{\mathcal{T}_1}$ it follows that

$$\{p_1^j, \dots, p_{m_j}^j\} \models_{\mathcal{T}_1} q_1^j \vee \cdots \vee q_{n_j}^j.$$

From the Σ -convexity of \mathcal{T}_1 we can conclude that there is a $k_j \in \{1, \dots, n_j\}$ such that $\{p_1^j, \dots, p_{m_j}^j\} \models_{\mathcal{T}_1} q_{k_j}^j$.¹⁶ In conclusion, for all $j \in \{1, \dots, n\}$ there is a $k_j \in \{1, \dots, n_j\}$ such that

$$\mathcal{T}_1 \models \neg p_1^j \vee \cdots \vee \neg p_{m_j}^j \vee q_{k_j}^j.$$

Let φ be the conjunctions of all the Horn formulas $\neg p_1^j \vee \cdots \vee \neg p_{m_j}^j \vee q_{k_j}^j$ above. It is immediate that $\mathcal{T}_1 \models \varphi$. To prove the proposition then it is enough to show that $\mathcal{T}_2 \cup \{\varphi\}$ is unsatisfiable.

Suppose ad absurdum that $\mathcal{T}_2 \cup \{\varphi\}$ is satisfiable in some structure \mathcal{A} . In that case, $\mathcal{T}_2 \cup \{\neg p_1^j \vee \cdots \vee \neg p_{m_j}^j \vee q_{k_j}^j\}$ is satisfiable in \mathcal{A} for every $j \in \{1, \dots, n\}$. But then, $\mathcal{T}_2 \cup \{\psi_j\}$ with $\psi_j = \neg p_1^j \vee \cdots \vee \neg p_{m_j}^j \vee q_1^j \vee \cdots \vee q_{n_j}^j$ is clearly also satisfiable in \mathcal{A} for every $j \in \{1, \dots, n\}$. Noting that none of the elements of $\mathcal{T}_2 \cup \{\psi_j\}$ has free variables, it follows that $\mathcal{T}_2 \cup \{\psi_1 \wedge \cdots \wedge \psi_n\}$ is satisfiable in \mathcal{A} . But this contradicts the assumptions that $\psi = \psi_1 \wedge \cdots \wedge \psi_n$ and $\mathcal{T}_2 \models \neg\psi$. It follows that $\mathcal{T}_2 \cup \{\varphi\}$ is unsatisfiable. \square

¹⁶ If $\{p_1^j, \dots, p_{m_j}^j\}$ is \mathcal{T}_1 -unsatisfiable, one such k_j is 1, given the assumption that $q_1^j = \perp$.

As in Section 3.3, the following corollary of Proposition 3.2 is more useful for our purposes.

PROPOSITION 5.3. *For $i = 1, 2$ let Φ_i a set of Σ_i -literals. Then, the following are equivalent:*

1. $\Phi_1 \cup \Phi_2$ is $(\mathcal{T}_1 \cup \mathcal{T}_2)$ -unsatisfiable;
2. there is a finite set Ψ of basic Horn formulas of signature Σ with $\text{Var}(\Psi) \subseteq \text{Var}(\Phi_1 \cup \Phi_2)$ such that

$$\Phi_1 \models_{\mathcal{T}_1} \Psi \quad \text{and} \quad \Phi_2 \cup \Psi \models_{\mathcal{T}_2} \perp .$$

Proof. Similarly to Proposition 3.3. □

Now, we could use the result above in the proof of Proposition 4.13 to show that the tableau calculus in the previous section is complete even if residues are restricted to basic Horn formulas. In fact, we can do even better and restrict residues to just positive literals. The main reason for this is provided by the following lemma.

LEMMA 5.4. *For $i = 1, 2$ let Φ_i a set of Σ_i -literals such that $\Phi_1 \cup \Phi_2$ is $(\mathcal{T}_1 \cup \mathcal{T}_2)$ -unsatisfiable. Then there is a finite n , a sequence $(\Phi_j^i)_{j \leq n}$ of sets of Σ_i -literals with $\Phi_0^i = \Phi_i$ for $i = 1, 2$, and a sequence $(a_j)_{j \leq n}$ of positive Σ -literals with $a_n = \perp$ such that:*

1. for all $j < n$ and $i = 1, 2$, $\Phi_{j+1}^i = \Phi_j^i \cup \{a_j\}$;
2. for all $j \leq n$, $\text{Var}(a_j) \subseteq \text{Var}(\Phi_j^1 \cup \Phi_j^2)$ and either $\Phi_j^1 \models a_j$ or $\Phi_j^2 \models a_j$.

Proof. We construct recursively three sequences $(\Phi_j^1)_{j \leq n}$, $(\Phi_j^2)_{j \leq n}$ and $(a_j)_{j \leq n}$ that satisfy the statement of the lemma. To build $(a_j)_{j \leq n}$ we use an auxiliary sequence $(\Psi_j)_{j \leq n}$ in which each element Ψ_j is an interpolant set of Φ_j^1 and Φ_j^2 .

Let $\Phi_0^i := \Phi^i$ for $i = 1, 2$. By Proposition 5.3, there is a finite set

$$\Psi_0 := \left\{ \begin{array}{c} \neg p_{1,1} \vee \cdots \vee \neg p_{1,n_1} \vee q_1 \\ \vdots \\ \neg p_{m,1} \vee \cdots \vee \neg p_{m,n_m} \vee q_m \end{array} \right\}$$

of basic Horn formulas of signature Σ with $\text{Var}(\Psi_0) \subseteq \text{Var}(\Phi_0^1 \cup \Phi_0^2)$ such that

$$\Phi_0^1 \models_{\mathcal{T}_1} \Psi_0 \quad \text{and} \quad \Phi_0^2 \cup \Psi_0 \models_{\mathcal{T}_2} \perp .$$

We define a_0 according to the following (mutually exclusive and exhaustive) cases:

1. Ψ_0 is empty. Then $\Phi_0^2 \models_{\mathcal{T}_2} \perp$. We set $a_0 = \perp$.
2. $\Psi_0 \neq \emptyset$ and one of Ψ_0 's members is a single positive literal q_k —which happens if $n_k = 0$ for some k . Then $\Phi_0^1 \models_{\mathcal{T}_1} q_k$. We set $a_0 := q_k$.
3. $\Psi_0 \neq \emptyset$ but none of Ψ_0 's members is a single positive literal. Then $n_k > 0$ for all k . From the fact that $\Phi_0^2 \cup \Psi_0 \models_{\mathcal{T}_2} \perp$ it is easy to show by simple logical reasoning that $\Phi_0^2 \models_{\mathcal{T}_2} p_{1,1} \vee \dots \vee p_{m,1}$. If Φ_0^2 is \mathcal{T}_2 -unsatisfiable, we set $a_0 := \perp$. Otherwise, we know that, since \mathcal{T}_2 is convex, there is a k such that $\Phi_0^2 \models_{\mathcal{T}_2} p_{k,1}$. In that case, we set $a_0 := p_{k,1}$. In both cases we have that $\Phi_0^2 \models_{\mathcal{T}_2} a_0$.

It is immediate that $\mathcal{V}ar(a_0) \subseteq \mathcal{V}ar(\Phi_0^1 \cup \Phi_0^2)$ in all cases. Now, if $a_0 = \perp$ we stop; that is, $(a_j)_{j \leq n}$ is just $(a_j)_{j \leq 0}$. Otherwise, we consider the sets $\Phi_1^1 := \Phi_0^1 \cup \{a_0\}$ and $\Phi_1^2 := \Phi_0^2 \cup \{a_0\}$ and the set Ψ_1 obtained from Ψ_0 by removing (the disjunction with) the literal a_0 from all the formulas of Ψ_0 . It is not difficult to see that Ψ_1 is an interpolant set of Φ_1^1 and Φ_1^2 , i.e., $\Phi_1^1 \models_{\mathcal{T}_1} \Psi_1$ and $\Phi_1^2 \cup \Psi_1 \models_{\mathcal{T}_2} \perp$.

Applying the same process recursively to Φ_1^1 , Φ_1^2 and Ψ_1 we can generate the sequences $(\Phi_j^1)_j$, $(\Phi_j^2)_j$, $(\Psi_j)_j$ and $(a_j)_j$ where each Ψ_j is an interpolant set of Φ_j^1 , Φ_j^2 and $(\Phi_j^1)_j$ and $(\Phi_j^2)_j$, and $(a_j)_j$ have the desired properties. The finiteness of all the sequences is guaranteed by the fact that each Ψ_{j+1} is obtained from Ψ_j by removing all occurrences of a literal in Ψ_j . \square

Now let $\mathcal{T} := \mathcal{T}_1 \cup \mathcal{T}_2$. The lemma above is basically saying that a background reasoner for \mathcal{T}_1 and a one for \mathcal{T}_2 can together detect the \mathcal{T} -unsatisfiability of a set like $\Phi_1 \cup \Phi_2$ above by just exchanging entailed Σ -atoms back and forth. Using the lemma it is not difficult to modify the proof of Proposition 4.13 to show that the tableau calculus in the previous section satisfies the following completeness result.

PROPOSITION 5.5 (Convex Completeness). *For every \mathcal{T} -valid sentence φ of signature $\Sigma_1 \cup \Sigma_2$, there is a strict and fair tableau proof of φ in which every application of the residue rule is restricted to \mathcal{T}_1 and \mathcal{T}_2 and its residue has the form (σ, p) where p is a positive Σ -literal.*

6. Related Work

Our Ground Interpolation Lemma (Proposition 3.2) can also be seen as an instance of a general interpolation theorem for infinitary logic¹⁷ due to Malitz (Malitz, 1969). In our terms, the theorem states that any two theories \mathcal{T}_1 and \mathcal{T}_2 whose union is unsatisfiable, admit an interpolant ψ in their shared signature which is a universal sentence (of infinitary logic) whenever \mathcal{T}_2 is universal. Again, the contribution of Lemma 3.2 is to show that if \mathcal{T}_1 is also universal and \mathcal{T}_1 and \mathcal{T}_2 have the same function symbols and contain only finitary formulas, then ψ is a ground finitary formula.

The only research we are aware of that focuses on the cooperation of background reasoners in theory reasoning is that reported in (Käufel and Zabel, 1990; Tinelli and Harandi, 1998; Baalen and Roach, 1999; Petermann, 2000; Zarba, 2002). Except for (Petermann, 2000) and (Zarba, 2002), all of these works embed a well-known combination method by Nelson and Oppen (Nelson and Oppen, 1979) into a specific theory reasoning calculus: analytic tableaux in (Käufel and Zabel, 1990), a variant of the CLP scheme in (Tinelli and Harandi, 1998), and constrained resolution in (Baalen and Roach, 1999).

In essence, the approach in each of these papers is a specialization of the one presented here. One major difference is that the background reasoners return only residues of the form $\langle \varepsilon, \psi \rangle$, which basically means that they treat key sets as if they were ground. This is enough for completeness in both the CLP scheme (Jaffar and Maher, 1994) and in constrained resolution (Bürckert, 1994). It is also enough in (Käufel and Zabel, 1990) because in the used tableau calculus γ formulas are expanded into their ground instances—which makes the calculus very inefficient though. Another major difference with our approach is that the two background theories are *stably-infinite* (see, e.g., (Oppen, 1980)) and share at most the equality and the constant symbols, whereas in our case the theories, although possibly not stably-infinite, share all function symbols. The net effect of these differences—leading to more restricted but stronger computational results than ours—is that for each key set Φ it is enough to consider only the finitely many residues $\langle \varepsilon, \psi \rangle$ in which ψ is a disjunction of equations between certain subterms: the *alien* subterms in Φ (see (Tinelli and Harandi, 1996) for details).

In (Petermann, 2000), some special types of background theories are integrated into the theory connection calculus. A number of rather specific syntactical restrictions are imposed on the theories, including the disjointness of their sets of predicate symbols, none of which are

¹⁷ The extension of first-order logic that contains a conjunction and a disjunction symbol of countably infinite arity.

necessary in our approach. It is not clear, however, exactly how the approach in (Petermann, 2000) compares to ours.

A very recent work, (Zarba, 2002), extends semantic tableaux into a total multi-theory reasoning calculus by the addition of four new derivations rules. There are no restrictions of the background theories except that the satisfiability in them of ground literals must be decidable. The effect of the new rules is to purify mixed literals into conjunctions of pure ground literals (over a signature extended with Skolem constants), and to split branches on tautologies of the form $p \vee \neg p$, where p is *any* ground literal over the shared signature (including the Skolem constants). Since splitting on tautologies is clearly impractical in general, (Zarba, 2002) presents an alternative version of the calculus that replaces splitting by a residue rule inspired by our own calculus. The main advantage of this new calculus over ours is that the residue rule needs to add only disjunctions of *positive* literals.¹⁸ The main disadvantage of the calculus is that, for being based on semantic tableaux, it blindly expands γ formulas into their ground instances.

7. Conclusions

In this paper we have sought to demonstrate that, contrary to a common belief in the field, integrating multiple background reasoners in theory reasoning is conceptually straightforward. Thanks to a specialization of Craig's interpolations lemma, the needed cooperation between the background reasoners can be achieved as a simple form of constraint propagation over a common language.

Our main contribution was to show that, under some conditions on the background theories, the propagated constrains can be restricted to disjunctions of literals in the signature Σ shared by the theories. The requirements on the background theories, namely that they be universal and have all their function symbols in common, are very mild: the first is a given in all theory reasoning calculi; the second is typically easy to satisfy, as explained in Section 3.2. We have also shown that if the theories are Σ -convex as well, the propagated constrains can be further restricted to single Σ -literals.

¹⁸ (Zarba, 2002) puts a lot of emphasis on the fact that it does not require, like we do, that the background theories have all their function symbols in common. However, (Zarba, 2002) requires the background reasoners to accept inputs with Skolem constants, which makes its results no more general than ours. In fact, it is easy to show that a reasoner deciding the satisfiability of ground formulas with Skolem constants can be always augmented to decide the satisfiability of quantifier-free formulas with Skolem function symbols of non-zero arity, which is enough for our purposes, as explained in Section 3.2.

For concreteness, we have proved our claims here in the context of a specific theory reasoning calculus. We have described a multi-theory version of the semantic tableau calculus in which the cooperation among the background reasoners is achieved by the sort of constraint propagation mentioned above, and we have proved the calculus sound and complete.

We stress that our combination results are not limited to the theory calculus considered here. For instance, although not shown in this paper, we have been able to extend theory resolution (Stickel, 1985; Baumgartner, 1992) in a similar way and produce corresponding soundness and completeness results. We conjecture that analogous multi-theory extensions can be obtained for all the major literal-level theory reasoning calculi.

7.1. FURTHER RESEARCH

Further research is obviously needed to assess the practical utility of the combination results presented here. The two major practical issues for actual theory reasoning systems are—on the foreground reasoner side—how to choose key sets effectively and—on the background reasoner side—how to generate residues efficiently.¹⁹ These same issues remain crucial in our approach as well. We did show that, under the right conditions, it is enough to consider only certain types of key sets and residues. But even within these restrictions the number of possible choices is still large enough to make actual applications impractical without further optimizations. As we mentioned in the introduction, the research challenge is now to identify specific combinations of theories and more or less general implementation techniques like those described in (Beckert and Pape, 1996) for which our cooperation approach is feasible.

Focusing on the specialized results from Section 5, one potentially interesting application could come in conjunction with Baumgartner’s results on *linearizing completion* (Baumgartner, 1996), a technique for producing background reasoners automatically for certain universal Horn theories. Linearizing completion takes a finite set \mathcal{T} of Horn clauses and either diverges or produces a finite set $\mathcal{I}(\mathcal{T})$ of unit-resulting inference rules. This set constitutes an inference system that in turn can be automatically “compiled” into a specialized reasoner $R_{\mathcal{T}}$ for \mathcal{T} . The reasoner $R_{\mathcal{T}}$ is refutationally complete with respect to \mathcal{T} -unsatisfiability. More interestingly for us though, $R_{\mathcal{T}}$ can be used as a rather efficient background reasoner for partial theory reasoning calculi: one that preserves the completeness of the overall calculus while

¹⁹ A noteworthy approach partially addressing these issues and based on incremental methods is described in (Beckert and Pape, 1996).

only needing key sets up to a certain cardinality, and producing “few” \mathcal{T} -residues, all of them unit (see (Baumgartner, 1996) for more details).

The major limitation of linearizing completion is that it often diverges on the input set \mathcal{T} of clauses. In some cases, however, it is possible to partition \mathcal{T} into two sets \mathcal{T}_1 and \mathcal{T}_2 such that linearizing completion on each of them separately succeeds (Baumgartner, 2001). Until now this fact was not extremely useful because, clearly, neither of the two reasoners produced this way would be a background reasoner for \mathcal{T} when taken separately. The two reasoners, however, can now be combined with our approach to reason over \mathcal{T} in cooperation. Although this idea looks promising, the cases mentioned in (Baumgartner, 2001) involve very simple theories. More work needs to be done to find more interesting cases of theories that can be partitioned into subtheories on which linearizing completion converges.

Acknowledgements

I would like to thank Melvin Fitting for pointing out Malitz’s interpolation theorem to me. Special thanks go Peter Baumgartner for some illuminating discussions on several aspects of theory reasoning and for his numerous and insightful comments on a draft of this paper which have helped improve it. Finally, I thank Calogero Zarba for pointing out an erroneous result in a previous version of the paper.

Appendix

A. Proofs

In the following, we will use the standard model-theoretic notions of embedding, isomorphism, substructure, generators, reducts and so on. The reader is referred to (Hodges, 1993b), among others, for their definition. The results given here are expressed and hold in first-order logic with equality. However, all of them can be shown to hold as stated in first-order logic without equality as well.

Where \mathcal{A} is a structure of signature Ω , we denote by $diag(\mathcal{A})$ the set of all ground Ω -literals that are true in \mathcal{A} ; if Σ is a subsignature of Ω , we denote by \mathcal{A}^Σ the reduct of \mathcal{A} to Σ ; if X is a subset of the universe of \mathcal{A} , we denote by $\langle X \rangle_{\mathcal{A}}$ the substructure of \mathcal{A} generated by X .

We will appeal to the following three basic results from model theory. The first is an elementary fact. For a proof of the other two, again see (Hodges, 1993b), among others.

LEMMA A.1. *Let \mathcal{A} be an Ω -structure, Σ a subsignature of Ω and X a subset of \mathcal{A} 's universe. If $\Sigma^F = \Omega^F$, then $\langle X \rangle_{\mathcal{A}}^{\Sigma} = \langle X \rangle_{\mathcal{A}^{\Sigma}}$.*

LEMMA A.2 (Robinson's Diagram Lemma). *Let \mathcal{A} be a Σ -structure generated by the empty set and \mathcal{B} a structure whose signature includes Σ . Then, \mathcal{A} is embeddable in \mathcal{B}^{Σ} whenever \mathcal{B} models the set $\text{diag}(\mathcal{A})$.*

LEMMA A.3. *The set of models of a universal Σ -theory \mathcal{T} is closed under substructures. That is, every substructure of a (Σ) -model of \mathcal{T} is a model of \mathcal{T} .*

We will also appeal to the notion of fusion from (Tinelli and Ringeissen, 2002) and some of its properties, proved in (Tinelli and Ringeissen, 2002).

DEFINITION A.4 (Fusion). *For $i = 1, 2$ let \mathcal{A}_i be a structure of signature Σ_i . A $(\Sigma_1 \cup \Sigma_2)$ -structure \mathcal{F} is a fusion of \mathcal{A}_1 and \mathcal{A}_2 if \mathcal{F}^{Σ_i} is isomorphic to \mathcal{A}_i for $i = 1, 2$.*

Fusions of structures do not always exist. The following proposition establishes a necessary and sufficient condition for their existence.

PROPOSITION A.5. *Let \mathcal{A} and \mathcal{B} be two structure and let Σ be the intersection of their signatures. Then, \mathcal{A} and \mathcal{B} admit a fusion exactly when \mathcal{A}^{Σ} is isomorphic to \mathcal{B}^{Σ} .*

Fusions of structures are related to unions of theories as follows.

PROPOSITION A.6. *For $i = 1, 2$ let \mathcal{T}_i be a theory of signature Σ_i . A $(\Sigma_1 \cup \Sigma_2)$ -structure is a model of $\mathcal{T}_1 \cup \mathcal{T}_2$ iff it is the fusion of a model of \mathcal{T}_1 and a model of \mathcal{T}_2 .*

A.1. \mathcal{T} -HINTIKKA SETS

We start by proving Lemma 2.3, stating that for each satisfiable universal theory \mathcal{T} , every \mathcal{T} -Hintikka set is satisfiable in a canonical model of \mathcal{T} .

LEMMA 2.3. *If \mathcal{T} is a satisfiable universal theory, then every \mathcal{T} -Hintikka set is satisfiable in a canonical model of \mathcal{T} .*

Proof. Assume that \mathcal{T} has signature Ω . Where Ω' is a signature including Ω (and having at least a constant symbol), let H be a \mathcal{T} -Hintikka set of signature Ω' and let Φ be the (possibly infinite) set of all the literals

in H .²⁰ By Definition 2.2(1), the universal theory $\Phi \cup \mathcal{T}$ is satisfiable. Let \mathcal{A} be any Ω' -model of $\Phi \cup \mathcal{T}$ and let \mathcal{B} be the substructure of \mathcal{A} generated by the empty set. By Lemma A.3, \mathcal{B} as well is a model of $\Phi \cup \mathcal{T}$. Since it is generated by the empty set, we know that, in addition, \mathcal{B} is (isomorphic to) a canonical model of $\Phi \cup \mathcal{T}$. We prove by structural induction that every sentence of H holds in \mathcal{B} .

(Base case) Every literal of H holds in \mathcal{B} by construction, for being an element of Φ .

(Induction step) We consider just the β and γ sentences of H . For α or δ sentences the argument is similar. If a sentence β is in H then $\beta_i \in H$ for $i = 1$ or $i = 2$. By induction hypothesis, β_i holds in \mathcal{B} . But then, by definition, β is also holds in \mathcal{B} . If a sentence γ is in H then $\gamma_1(t) \in H$ for all ground Ω' -terms t . By induction hypothesis, each $\gamma(t)$ holds in \mathcal{B} . Since every element of \mathcal{B} is denoted by some ground term t , given that \mathcal{B} is a canonical model, it follows that the formula $\gamma(x)$ in the free variable x is satisfied in \mathcal{B} by every interpretation of x . But by the semantics of universal quantification, this means that γ holds in \mathcal{B} . \square

A.2. GROUND INTERPOLATION LEMMA

In this subsection we prove Proposition 3.2, the Ground Interpolation Lemma from Section 3.3. As in that section, we fix two signatures Σ_1, Σ_2 such that $\Sigma_1^F = \Sigma_2^F$ and two universal theories $\mathcal{T}_1, \mathcal{T}_2$ of respective signature Σ_1, Σ_2 . Also, let $\Sigma := \Sigma_1 \cap \Sigma_2$.

The proof of the lemma will be facilitated by the following intermediate result. In its proof, we use \bar{l} to denote the complement of a literal l .

LEMMA A.7. *Let Ψ be the set of all disjunctions ψ of ground Σ -literals such that $\mathcal{T}_1 \models \psi$. If every finite subset of Ψ is \mathcal{T}_2 -satisfiable, then the theory $\mathcal{T}_1 \cup \mathcal{T}_2$ is satisfiable.*

Proof. Assume that every finite subset of Ψ is \mathcal{T}_2 -satisfiable. Then, every finite subset of $\Psi \cup \mathcal{T}_2$ is satisfiable. By the compactness of first-order logic, this entails that the whole $\Psi \cup \mathcal{T}_2$ is satisfiable. Let \mathcal{A}_2 be a Σ_2 -model of $\Psi \cup \mathcal{T}_2$ and assume with no loss of generality that \mathcal{A}_2 is generated by the empty set.²¹ Since $\Sigma^F = \Sigma_1^F = \Sigma_2^F$ by assumption, the Σ -reduct \mathcal{A}_2^Σ as well is generated by the empty set.

²⁰ Note, that since H is a set of sentences, each literal in it must be ground.

²¹ Otherwise, one can consider in its place the substructure of \mathcal{A}_2 generated by the empty set. This substructure exists because Σ_2 contains at least a constant symbol; moreover, it is a model of $\Psi \cup \mathcal{T}_2$ by Lemma A.3 because $\Psi \cup \mathcal{T}_2$ is universal.

We start by showing by contradiction that the Σ_1 -theory $\mathcal{T}_1 \cup \text{diag}(\mathcal{A}_2^\Sigma)$ is satisfiable.

If $\mathcal{T}_1 \cup \text{diag}(\mathcal{A}_2^\Sigma)$ is not satisfiable, then by compactness again we can show that there is a finite subset $\{l_1, \dots, l_n\}$ of $\text{diag}(\mathcal{A}_2^\Sigma)$ that is \mathcal{T}_1 -unsatisfiable. This implies that the formula $\psi := \bar{l}_1 \vee \dots \vee \bar{l}_n$ is entailed by \mathcal{T}_1 . For being a disjunction of ground Σ -literals, ψ must then be an element of Ψ . Now, since \mathcal{A}_2 models Ψ and ψ is a Σ -formula, we then have that $\bar{l}_1 \vee \dots \vee \bar{l}_n$ is true in \mathcal{A}_2^Σ . But this is impossible because every l_j is in $\text{diag}(\mathcal{A}_2^\Sigma)$, the set of ground Σ -literals true in \mathcal{A}_2^Σ , and so every \bar{l}_j is false in \mathcal{A}_2^Σ .

Now let \mathcal{A}_1 be a Σ_1 -model of $\mathcal{T}_1 \cup \text{diag}(\mathcal{A}_2^\Sigma)$, and assume, again with no loss of generality, that \mathcal{A}_1 is generated by the empty set. By Lemma A.2, since \mathcal{A}_1 models $\text{diag}(\mathcal{A}_2^\Sigma)$, \mathcal{A}_2^Σ is embeddable into \mathcal{A}_1^Σ . Recalling that \mathcal{A}_2^Σ is generated by the empty set, this means that \mathcal{A}_2^Σ is isomorphic to $\langle \emptyset \rangle_{\mathcal{A}_1^\Sigma}$, the substructure of \mathcal{A}_1^Σ generated by the empty set. By Lemma A.1, since $\Sigma^F = \Sigma_1^F$, this substructure coincides with $\langle \emptyset \rangle_{\mathcal{A}_1}^\Sigma$, that is, with \mathcal{A}_1^Σ .

In conclusion, we have shown that the structures \mathcal{A}_1 and \mathcal{A}_2 have isomorphic reducts over their shared signature Σ . Therefore, they admit a fusion \mathcal{F} by Proposition A.5. Since by construction \mathcal{A}_1 is model of \mathcal{T}_1 and \mathcal{A}_2 is model of \mathcal{T}_2 , we can conclude by Proposition A.6 that \mathcal{F} is a model of $\mathcal{T}_1 \cup \mathcal{T}_2$, which makes $\mathcal{T}_1 \cup \mathcal{T}_2$ satisfiable. \square

PROPOSITION 3.2. (Ground Interpolation Lemma) *If $\mathcal{T}_1 \cup \mathcal{T}_2$ is unsatisfiable, then $\mathcal{T}_1 \models \varphi$ and $\mathcal{T}_2 \models \neg\varphi$ for some ground Σ -formula φ .*

Proof. Assume that $\mathcal{T}_1 \cup \mathcal{T}_2$ is unsatisfiable and let Ψ be the set of all disjunctions ψ of ground Σ -literals such that $\mathcal{T}_1 \models \psi$. By the contrapositive of Lemma A.7, we know that there is a finite subset $\{\psi_1, \dots, \psi_n\}$ of Ψ that is \mathcal{T}_2 -unsatisfiable. Let $\varphi := \psi_1 \wedge \dots \wedge \psi_n$. By construction, φ is a ground Σ -formula such that $\mathcal{T}_1 \models \varphi$ and $\mathcal{T}_2 \models \neg\varphi$. \square

A.3. Σ -CONVEX THEORIES

In this subsection we show that every Horn theory, and in particular every universal Horn theory, is Σ -convex for any Σ . Then, we point to some examples of non-Horn Σ -convex theories. We start by defining (universal) Horn theories, again following (Hodges, 1993a).

Recall that a basic Horn formula is a formula of the form $\neg p_1 \vee \dots \vee \neg p_n \vee q$ where $n \geq 0$ and each of p_1, \dots, p_n, q is a positive literal (possibly \perp). A *Horn formula* is a formula of the form $Q.(\varphi_1 \wedge \dots \wedge \varphi_n)$

where Q is an arbitrary quantifier prefix, $n > 0$ and each φ_j is a basic Horn formula. A *Horn sentence* is Horn formula with no free variables and a *universal Horn sentence* is a Horn sentence whose quantifier prefix contains only universal quantifiers. A *universal Horn theory* is a set of universal Horn sentences.²²

The convexity of Horn theories is an almost immediate consequence of a well-known result by McKinsey, one of whose formulations is the following (see (Hodges, 1993a)).

LEMMA A.8 (McKinsey’s Lemma). *Let \mathcal{T} be a satisfiable Horn theory and let Ψ be a set of positive ground literals. If every model of \mathcal{T} is a model of at least one element of Ψ , then there is a $p \in \Psi$ such that $\mathcal{T} \models p$.*

PROPOSITION A.9. *Every Horn theory is Σ -convex for any signature Σ .*

Proof. Let Σ be any signature and \mathcal{T} a Horn theory of signature Ω . Let Φ be a set of Ω -literals and Ψ a finite, non-empty set of positive Σ -literals such that $\Phi \models_{\mathcal{T}} \bigvee_{p \in \Psi} p$. We show that $\Phi \models_{\mathcal{T}} p$ for some $p \in \Psi$.

If Φ is \mathcal{T} -unsatisfiable, the claim is trivially true for any $p \in \Psi$. Therefore assume that Φ is \mathcal{T} -satisfiable and consider Φ and Ψ as sets of ground literals in the signature $\Omega(X)$ and $\Sigma(X)$, respectively, where $X := \text{Var}(\Phi \cup \Psi)$. Then observe that $\mathcal{T}' := \mathcal{T} \cup \Phi$ is a satisfiable Horn theory of signature $\Omega(X)$ and that $\mathcal{T}' \models \bigvee_{p \in \Psi} p$. The claim then follows immediately from Lemma A.8. \square

To conclude, in (Nelson and Oppen, 1979) an Ω -theory \mathcal{T} is called convex iff whenever a conjunction of Ω -literals \mathcal{T} -entails a disjunction of equalities between variables, it \mathcal{T} -entails one of the equalities in the disjunction. That paper also provides some example of convex, universal theories two of which—a theory of rational numbers under addition and a theory of lists—are not (axiomatizable by) Horn theories.

It is easy to see that, in FOL with equality, convex theories in the sense of (Nelson and Oppen, 1979) are Σ -convex theories according to Definition 5.1, where Σ is the empty signature. There is, however, an even stronger connection between the two definitions.

PROPOSITION A.10. *Every convex theory of signature Ω is Σ -convex with $\Sigma := \Omega^F$.*

²² Note that a universal Horn sentence is a universal sentence as defined in Section 2. Similarly, a universal Horn theory is a universal theory.

Proof. Let \mathcal{T} be the convex Ω -theory and $\Sigma := \Omega^F$. Let Φ be any set of Ω -literals and Ψ any finite non-empty set of positive Σ -literals such that $\Phi \models_{\mathcal{T}} \bigvee_{p \in \Psi} p$. We prove that $\Phi \models_{\mathcal{T}} p$ for some $p \in \Psi$.

Since Σ contains only function symbols, Ψ must be a set of equalities. So let $\Psi := \{s_1 \approx t_1, \dots, s_n \approx t_n\}$ where \approx denotes the equality symbol. By compactness, we can assume with no loss of generality that Φ is finite. Now, from the assumption that $\Phi \models_{\mathcal{T}} \bigvee_{j=1, \dots, n} s_j \approx t_j$ we can deduce by elementary logical reasoning that

$$\bigwedge_{p \in \Phi} p \wedge \left(\bigwedge_j x_j \approx s_j \right) \wedge \left(\bigwedge_j y_j \approx t_j \right) \models_{\mathcal{T}} \bigvee_j x_j \approx y_j$$

where all the x_j and y_j are fresh variables. Since the conjunction above is a conjunction of Ω -literals, by the convexity of \mathcal{T} there is a $j \in \{1, \dots, n\}$ such that

$$\bigwedge_{p \in \Phi} p \wedge \left(\bigwedge_j x_j \approx s_j \right) \wedge \left(\bigwedge_j y_j \approx t_j \right) \models_{\mathcal{T}} x_j \approx y_j$$

But this entails that $\Phi \models_{\mathcal{T}} s_j \approx t_j$ for some $j \in \{1, \dots, n\}$. \square

References

- Baalen, J. V. and S. Roach: 1999, ‘Using Decision Procedures to Accelerate Domain-Specific Deductive Synthesis Systems’. In: P. Flener (ed.): *Proceedings of the 8th International Workshop on Logic Programming Synthesis and Transformation (LOPSTR’98), Manchester, UK*, Vol. 1559 of *Lecture Notes in Computer Science*. pp. 61–70, Springer-Verlag.
- Baumgartner, P.: 1992, ‘An Ordered Theory Resolution Calculus’. In: A. Voronkov (ed.): *Proceedings of the 1st International Conference on Logic Programming and Automated Reasoning*, Vol. 624 of *Lecture Notes in Artificial Intelligence*. pp. 119–130, Springer Verlag.
- Baumgartner, P.: 1996, ‘Linear and Unit-Resulting Refutations for Horn Theories’. *Journal of Automated Reasoning* **16**(3), 241–319.
- Baumgartner, P.: 2001, ‘Personal communication’.
- Baumgartner, P., U. Furbach, and U. Petermann: 1992, ‘A Unified Approach to Theory Reasoning’. Research Report 15–92, Universität Koblenz-Landau, Koblenz, Germany. Fachberichte Informatik.
- Beckert, B. and C. Pape: 1996, ‘Incremental Theory Reasoning Methods for Semantic Tableaux’. In: P. Miglioli, U. Moscato, D. Mundici, and M. Ornaghi (eds.): *Proceedings of the 5th Workshop on Theorem Proving with Analytic Tableaux and Related Methods, TABLEUX’92 Palermo, (Italy)*, Vol. 1071 of *Lecture Notes in Computer Science*. pp. 93–109, Springer.
- Bibel, W.: 1982, *Automated Theorem Proving*. Braunschweig, Germany: Friedr. Vieweg & Sohn.
- Bürckert, H.-J.: 1994, ‘A Resolution Principle for Constraint Logics’. *Artificial Intelligence* **66**, 235–271.

- Fitting, M.: 1996, *First-Order Logic and Automated Theorem Proving*, Graduate Texts in Computer Science. Berlin: Springer, 2nd edition.
- Hodges, W.: 1993a, 'Logical features of Horn clauses'. In: D. Gabbay, C. Hogger, and J. Robinson (eds.): *Handbook of Logic in Artificial Intelligence and Logic Programming*, Vol. 1. Oxford University Press, pp. 449–503.
- Hodges, W.: 1993b, *Model Theory*, Vol. 42 of *Encyclopedia of mathematics and its applications*. Cambridge University Press.
- Jaffar, J. and M. Maher: 1994, 'Constraint Logic Programming: A Survey'. *Journal of Logic Programming* **19/20**, 503–581.
- Käuffl, T. and N. Zabel: 1990, 'Cooperation of decision procedures in a tableau-based theorem prover'. *Revue d'Intelligence Artificielle* **4**(3), 99–126.
- Lassez, J.-L. and K. McAloon: 1990, 'A Constraint Sequent Calculus'. In: *Proceedings of the 5th IEEE Symposium on Logic in Computer Science, LICS'90*. Philadelphia, pp. 52–61, IEEE Computer Society Press.
- Malitz, J.: 1969, 'Universal Classes in Infinitary Languages'. *Duke Mathematical Journal* **36**(3), 621–630.
- Nelson, G. and D. C. Oppen: 1979, 'Simplification by Cooperating Decision Procedures'. *ACM Trans. on Programming Languages and Systems* **1**(2), 245–257.
- Oppen, D. C.: 1980, 'Complexity, Convexity and Combinations of Theories'. *Theoretical Computer Science* **12**, 291–302.
- Petermann, U.: 2000, 'Connection Calculus Theorem Proving with Multiple Built-in Theories'. *Journal of Symbolic Computation* **29**(2), 373–392.
- Plotkin, G.: 1972, 'Building in Equational Theories'. In: B. Meltzer and D. Michie (eds.): *Machine Intelligence*, Vol. 7. Edinburgh, Scotland, pp. 73–90, Edinburgh University Press.
- Socher-Ambrosius, R. and P. Johann: 1997, *Deduction Systems*. Berlin: Springer-Verlag.
- Stickel, M. E.: 1985, 'Automated Deduction by Theory Resolution'. *Journal of Automated Reasoning* **1**(4), 333–355.
- Tinelli, C.: 2000, 'Cooperation of Background Reasoners in Theory Reasoning by Residue Sharing'. In: *International Workshop on First Order Theorem Proving, FTP'2000, St Andrews (Scotland)*.
- Tinelli, C. and M. T. Harandi: 1996, 'A new correctness proof of the Nelson–Oppen combination procedure'. In: F. Baader and K. Schulz (eds.): *Frontiers of Combining Systems: Proceedings of the 1st International Workshop (Munich, Germany)*. pp. 103–120, Kluwer Academic Publishers.
- Tinelli, C. and M. T. Harandi: 1998, 'Constraint Logic Programming over Unions of Constraint Theories'. *The Journal of Functional and Logic Programming* **1998**(6).
- Tinelli, C. and C. Ringeissen: 2002, 'Unions of Non-Disjoint Theories and Combinations of Satisfiability Procedures'. *Theoretical Computer Science*. (to appear).
- Zarba, C.: 2002, 'A Tableau Calculus for Combining Non-Disjoint Theories'. In: C. Fermüller and U. Egly (eds.): *Proceedings of the International Conference on Automated Reasoning with Analytic Tableaux and Related Methods, TABLEUX 2002 (Copenhagen, Scotland)*. Springer. (to appear).

Address for Offprints: Department of Computer Science
The University of Iowa

14 MacLean Hall
52242 Iowa City, IA
USA