
EFFICIENT FULL HIGHER-ORDER UNIFICATION

PETAR VUKMIROVIĆ, ALEXANDER BENTKAMP, AND VISA NUMMELIN

Vrije Universiteit Amsterdam, De Boelelaan 1111, 1081 HV Amsterdam, The Netherlands
e-mail address: p.vukmirovic@vu.nl

Vrije Universiteit Amsterdam, De Boelelaan 1111, 1081 HV Amsterdam, The Netherlands
e-mail address: a.bentkamp@vu.nl

Vrije Universiteit Amsterdam, De Boelelaan 1111, 1081 HV Amsterdam, The Netherlands
e-mail address: visa.nummelin@vu.nl

ABSTRACT. We developed a procedure to enumerate complete sets of higher-order unifiers based on work by Jensen and Pietrzykowski. Our procedure removes many redundant unifiers by carefully restricting the search space and tightly integrating decision procedures for fragments that admit a finite complete set of unifiers. We identify a new such fragment and describe a procedure for computing its unifiers. Our unification procedure, together with new higher-order term indexing data structures, is implemented in the Zipperposition theorem prover. Experimental evaluation shows a clear advantage over Jensen and Pietrzykowski's procedure.

1. INTRODUCTION

Unification is concerned with finding a substitution that makes two terms equal, for some notion of syntactic equality. Since the invention of Robinson's first-order unification algorithm [22], it has become an indispensable tool in theorem proving, logic programming, natural language processing, programming language compilation and other areas of computer science.

Many of these applications are based on higher-order formalisms and require higher-order unification. Due to its undecidability and explosiveness, the higher-order unification problem is considered one of the main obstacles on the road to efficient higher-order tools.

One of the reasons for higher-order unification's explosiveness lies in *flex-flex pairs*, which consist of two variable-headed terms, e.g., $F X \stackrel{?}{=} G a$, where F , G , and X are variables and a is a constant. Even this seemingly simple problem has infinitely many incomparable unifiers. One of the first methods designed to combat this explosion is Huet's preunification [12]. Huet noticed that some logical calculi would remain complete if flex-flex pairs are not eagerly solved but postponed as constraints. If only flex-flex constraints remain, we know that a unifier must exist and we do not need to solve them. Huet's preunification has been used in many reasoning tools including Isabelle [20], Leo-III [25], and Satallax [4]. However, recent

Key words and phrases: unification, higher-order logic, theorem proving, term rewriting, indexing data structures.

developments in higher-order theorem proving [2,3] require full unification—i.e., enumeration of unifiers even for flex-flex pairs, which is the focus of this article.

Jensen and Pietrzykowski’s (JP) procedure [13] is the best known procedure for this purpose (Section 2). Given two terms to unify, it first identifies a position where the terms disagree. Then, in parallel branches of the search tree, it applies suitable substitutions, involving a variable either at the position of disagreement or above, and repeats this process on the resulting terms until they are equal or trivially nonunifiable.

Building on the JP procedure, we designed a new procedure (Section 3) with the same completeness guarantees (Section 4). The new procedure addresses many of the issues that are detrimental to the performance of the JP procedure. First, the JP procedure does not terminate in many cases of obvious nonunifiability, e.g., for $X \stackrel{?}{=} f X$, where X is a non-functional variable and f is a function constant. This example also shows that the JP procedure does not generalize Robinson’s first-order procedure gracefully. To address this issue, our procedure detects whether a unification problem belongs to a fragment for which unification is decidable and finite complete sets of unifiers (CSUs) exist. We call algorithms that enumerate elements of the CSU for such fragments *oracles*. Noteworthy fragments with oracles are first-order terms, patterns [19], functions-as-constructors [16], and a new fragment we present in Section 5. The unification procedures of Isabelle and Leo-III check whether the unification problem belongs to a decidable fragment, but we take this idea a step further by checking this more efficiently and for every subproblem arising during unification.

Second, the JP procedure computes many redundant unifiers. Consider the example $F(G a) \stackrel{?}{=} F b$, where it produces, in addition to the desired unifiers $\{F \mapsto \lambda x. H\}$ and $\{G \mapsto \lambda x. b\}$, the redundant unifier $\{F \mapsto \lambda x. H, G \mapsto \lambda x. x\}$. The design of our procedure avoids computing many redundant unifiers, including this one. Additionally, as oracles usually return a small CSU, their integration reduces the number of redundant unifiers.

Third, the JP procedure applies more explosive rules than Huet’s preunification procedure to flex-rigid pairs. To gracefully generalize Huet’s procedure, we show that his rules for flex-rigid pairs suffice to enumerate CSUs if combined with appropriate rules for flex-flex pairs.

Fourth, the JP procedure repeatedly traverses the parts of the unification problem that have already been unified. Consider the problem $f^{100}(G a) \stackrel{?}{=} f^{100}(H b)$, where the exponents denote repeated application. It is easy to see that this problem can be reduced to $G a \stackrel{?}{=} H b$. However, the JP procedure will wastefully retrace the common context $f^{100}[\]$ after applying each new substitution. Since the JP procedure must apply substitutions to the variables occurring in the common context above the disagreement pair, it cannot be easily adapted to eagerly decompose unification pairs. By contrast, our procedure is designed to decompose the pairs eagerly, never traversing a common context twice.

Last, the JP procedure does not allow to apply substitutions and β -reduce lazily. The rules of simpler procedures (e.g., first-order [11] and pattern unification [19]) depend only on the heads of the unification pair. Thus, to determine the next step, implementations of these procedures need to substitute and β -reduce only until the heads of the current unification pair are not mapped by the substitution and are not λ -abstractions. Since the JP procedure is not based on the decomposition of unification pairs, it is unfit for optimizations of this kind. We designed our procedure to allow for this optimization.

To more efficiently find terms (in a large term set) that are unifiable with a given query term, we developed a higher-order extension of fingerprint indexing [23] (Section 6). We implemented our procedure, several oracles, and the fingerprint index in the Zipperposition prover (Section 7). Since a straightforward implementation of the JP procedure already

existed in Zipperposition, we used it as a baseline to evaluate the performance of our procedure (Section 8). The results show substantial performance improvements.

This invited article is an extended version of our FSCD-2020 paper [30]. Most notable extension is the Section 4, which gives the detailed proof of completeness of our new procedure. In addition, we give proofs for all the unproved statements from the paper, expand the examples and provide more detailed explanations.

2. BACKGROUND

Our setting is the simply typed λ -calculus. Types α, β, γ are either base types or functional types $\alpha \rightarrow \beta$. By convention, when we write $\alpha_1 \rightarrow \dots \rightarrow \alpha_n \rightarrow \beta$, we assume β to be a base type. Basic terms are free variables (denoted F, G, H, \dots), bound variables (x, y, z), and constants (f, g, h). Complex terms are applications of one term to another (st) or λ -abstractions ($\lambda x. s$). Following Nipkow [19], we use these syntactic conventions to distinguish free from bound variables. Bound variables with no enclosing binder, such as x in $\lambda y. x$, are called *loose bound variables*. We say that a term without loose bound variables is *closed* and a term without free variables is *ground*. Iterated λ -abstraction $\lambda x_1 \dots \lambda x_n. s$ is abbreviated as $\lambda \bar{x}_n. s$ and iterated application $(s t_1) \dots t_n$ as $s \bar{t}_n$, where $n \geq 0$. Similarly, we denote a sequence of terms t_1, \dots, t_n by \bar{t}_n , omitting its length $n \geq 0$ where it can be inferred or is irrelevant. Parameters and body for any term $\lambda \bar{x}. s$ are defined to be \bar{x} and s respectively, where s is not a λ -abstraction. The *size* of a term is inductively defined as $\text{size}(F) = 1$; $\text{size}(x) = 1$; $\text{size}(f) = 1$; $\text{size}(st) = \text{size}(s) + \text{size}(t)$; $\text{size}(\lambda x. s) = \text{size}(s) + 1$.

We assume the standard notions of α -, β -, η -conversions. A term is in *head normal form* (*hnf*) if it is of the form $\lambda \bar{x}. a \bar{t}$, where a is a free variable, bound variable, or a constant. In this case, a is called the *head* of the term. By convention, a and b denote heads. If a is a free variable, we call it a *flex* head; otherwise, we call it a *rigid* head. A term is called *flex* or *rigid* if its head is *flex* or *rigid*, respectively. By $s_{\downarrow h}$ we denote the term obtained from a term s by repeated β -reduction of the leftmost outermost redex until it is in hnf. Unless stated otherwise, we view terms syntactically, as opposed to $\alpha\beta\eta$ -equivalence classes. We write $s \leftrightarrow_{\alpha\beta\eta}^* t$ if s and t are $\alpha\beta\eta$ -equivalent. Substitutions $(\sigma, \varrho, \theta)$ are functions from free and bound variables to terms; σt denotes application of σ to t , which α -renames t to avoid variable capture. The composition $\varrho\sigma$ of substitutions is defined by $(\varrho\sigma)t = \varrho(\sigma t)$. A variable F is mapped by σ if $\sigma F \not\leftrightarrow_{\alpha\beta\eta}^* F$. We write $\varrho \subseteq \sigma$ if for all variables F mapped by ϱ , $\varrho F \leftrightarrow_{\alpha\beta\eta}^* \sigma F$. Given a substitution ϱ , which maps F to s , we write $\varrho \setminus \{F \mapsto s\}$ to denote a substitution that does not map F and otherwise coincides with ϱ . Given substitutions ϱ and σ , which map disjoint sets of variables, we write $\varrho \cup \sigma$ to denote $\varrho\sigma$.

Deviating from the standard notion of higher-order subterm, we define subterms on β -reduced terms as follows: a term t is a subterm of t at position ε . If s is a subterm of u_i at position p , then s is a subterm of $a \bar{u}_n$ at position $i.p$. If s is a subterm of t at position p , then s is a subterm of $\lambda x. t$ at position $1.p$. Our definition of subterm gracefully generalizes the corresponding first-order notion: \mathbf{a} is a subterm of $\mathbf{f a b}$ at position 1, but \mathbf{f} and $\mathbf{f a}$ are not subterms of $\mathbf{f a b}$. A context is a term with zero or more subterms replaced by a hole \square . We write $C[\bar{u}_n]$ for the term resulting from filling in the holes of a context C with the terms \bar{u}_n from left to right. The common context $\mathcal{C}(s, t)$ of two η -long β -reduced terms s and t of the same type is defined inductively as follows, assuming that $a \neq b$: $\mathcal{C}(\lambda x. s, \lambda y. t) = \lambda x. \mathcal{C}(s, \{y \mapsto x\}t)$; $\mathcal{C}(a \bar{s}_m, b \bar{t}_n) = \square$; $\mathcal{C}(a \bar{s}_m, a \bar{t}_m) = a \mathcal{C}(s_1, t_1) \dots \mathcal{C}(s_m, t_m)$.

A *unifier* for terms s and t is a substitution σ , such that $\sigma s \leftrightarrow_{\alpha\beta\eta}^* \sigma t$. Following JP [13], a *complete set of unifiers* (CSU) of terms s and t is defined as a set U of unifiers for s and t such that for every unifier ρ of s and t , there exists $\sigma \in U$ and substitution θ such that $\rho \subseteq \theta\sigma$. A *most general unifier* (MGU) is a one-element CSU. We use \subseteq instead of $=$ because a CSU element σ may introduce auxiliary variables not mapped by ρ .

3. THE UNIFICATION PROCEDURE

To unify two terms s and t , our procedure builds a tree as follows. The nodes of the tree have the form (E, σ) , where E is a multiset of unification constraints $\{(s_1 \stackrel{?}{=} t_1), \dots, (s_n \stackrel{?}{=} t_n)\}$ and σ is the substitution constructed up to that point. A unification constraint $s \stackrel{?}{=} t$ is an unordered pair of two terms of the same type. The root node of the tree is $(\{s \stackrel{?}{=} t\}, \text{id})$, where id is the identity substitution. The tree is then constructed applying the transitions listed below. The leaves of the tree are either failure nodes \perp or substitutions σ . Ignoring failure nodes, the set of all substitutions in the leaves forms a complete set of unifiers for s and t .

The transitions are parametrized by a mapping \mathcal{P} that assigns a set of substitutions to a unification pair; this mapping abstracts the concept of unification rules present in other unification procedures. Moreover, the transitions are parametrized by a selection function S mapping a multiset E of unification constraints to one of those constraints $S(E) \in E$, the *selected* constraint in E . The transitions, defined as follows, are only applied if the **grayed** constraint is selected.

Succeed: $(\emptyset, \sigma) \longrightarrow \sigma$

Normalize $_{\alpha\eta}$: $(\{\lambda\bar{x}_m. s \stackrel{?}{=} \lambda\bar{y}_n. t\} \uplus E, \sigma) \longrightarrow (\{\lambda\bar{x}_m. s \stackrel{?}{=} \lambda\bar{x}_m. t' x_{n+1} \dots x_m\} \uplus E, \sigma)$
 where $m \geq n$, $\bar{x}_m \neq \bar{y}_n$, and $t' = \{y_1 \mapsto x_1, \dots, y_n \mapsto x_n\}t$

Normalize $_{\beta}$: $(\{\lambda\bar{x}. s \stackrel{?}{=} \lambda\bar{x}. t\} \uplus E, \sigma) \longrightarrow (\{\lambda\bar{x}. s_{\downarrow h} \stackrel{?}{=} \lambda\bar{x}. t_{\downarrow h}\} \uplus E, \sigma)$
 where s or t is not in hnf

Dereference: $(\{\lambda\bar{x}. F \bar{s} \stackrel{?}{=} \lambda\bar{x}. t\} \uplus E, \sigma) \longrightarrow (\{\lambda\bar{x}. (\sigma F) \bar{s} \stackrel{?}{=} \lambda\bar{x}. t\} \uplus E, \sigma)$
 where none of the previous transitions apply and F is mapped by σ

Fail: $(\{\lambda\bar{x}. a \bar{s}_m \stackrel{?}{=} \lambda\bar{x}. b \bar{t}_n\} \uplus E, \sigma) \longrightarrow \perp$
 where none of the previous transitions apply, and a and b are different rigid heads

Delete: $(\{s \stackrel{?}{=} s\} \uplus E, \sigma) \longrightarrow (E, \sigma)$
 where none of the previous transitions apply

OracleSucc: $(\{s \stackrel{?}{=} t\} \uplus E, \sigma) \longrightarrow (E, \rho\sigma)$
 where none of the previous transitions apply, some oracle found a finite CSU U for $\sigma(s) \stackrel{?}{=} \sigma(t)$, and $\rho \in U$; if multiple oracles found a CSU, only one of them is considered

OracleFail: $(\{s \stackrel{?}{=} t\} \uplus E, \sigma) \longrightarrow \perp$
 where none of the previous transitions apply, and some oracle determined $\sigma(s) \stackrel{?}{=} \sigma(t)$ has no solutions

Decompose: $(\{\lambda\bar{x}. a \bar{s}_m \stackrel{?}{=} \lambda\bar{x}. a \bar{t}_m\} \uplus E, \sigma) \longrightarrow (\{s_1 \stackrel{?}{=} t_1, \dots, s_m \stackrel{?}{=} t_m\} \uplus E, \sigma)$
 where none of the transitions **Succeed** to **OracleFail** apply

Bind: $(\{s \stackrel{?}{=} t\} \uplus E, \sigma) \longrightarrow (\{s \stackrel{?}{=} t\} \uplus E, \rho\sigma)$
 where none of the transitions **Succeed** to **OracleFail** apply, and $\rho \in \mathcal{P}(s \stackrel{?}{=} t)$.

The transitions are designed so that only **OracleSucc**, **Decompose**, and **Bind** can introduce parallel branches in the constructed tree. **OracleSucc** can introduce branches using different

unifiers of the CSU, **Bind** can introduce branches using different substitutions in \mathcal{P} , and **Decompose** can be applied in parallel with **Bind**.

Our approach is to apply substitutions and $\alpha\beta\eta$ -normalize terms lazily. In particular, the transitions that modify the constructed substitution, **OracleSucc** and **Bind**, do not apply that substitution to the unification pairs directly. Instead, the transitions **Normalize $_{\alpha\eta}$** , **Normalize $_{\beta}$** , and **Dereference** partially normalize and partially apply the constructed substitution just enough to ensure that the heads are the ones we would get if the substitution was fully applied and the term was fully normalized. To support lazy dereferencing, **OracleSucc** and **Bind** must maintain the invariant that all substitutions are idempotent. The invariant is easily preserved if the substitution ϱ from the definition of **OracleSucc** and **Bind** is itself idempotent and no variable mapped by σ occurs in ϱF , for any variable F mapped by ϱ .

The **OracleSucc** and **OracleFail** transitions invoke oracles, such as pattern unification, to compute a CSU faster, produce fewer redundant unifiers, and discover nonunifiability earlier. In some cases, addition of oracles lets the procedure terminate more often.

In the literature, oracles are usually stated under the assumption that their input belongs to the appropriate fragment. To check whether a unification constraint is inside the fragment, we need to fully apply the substitution and β -normalize the constraint. To avoid these expensive operations and enable efficient oracle integration, oracles must be redesigned to lazily discover whether the terms belong to their fragment. Most oracles contain a decomposition operation which requires only a partial application of the substitution and only partial β -normalization. If one of the constraints resulting from decomposition is not in the fragment, the original problem is not in the fragment. This allows us to detect that the problem is not in the fragment without fully applying the substitution and β -normalizing.

The core of the procedure lies in the **Bind** step, parameterized by the mapping \mathcal{P} that determines which substitutions (called *bindings*) to create. The bindings are defined as follows:

JP-style projection for F : Let F be a free variable of type $\alpha_1 \rightarrow \dots \rightarrow \alpha_n \rightarrow \beta$, where some α_i is equal to β and $n > 0$. Then the JP-style projection binding is

$$F \mapsto \lambda \bar{x}_n. x_i$$

Huet-style projection for F : Let F be a free variable of type $\alpha_1 \rightarrow \dots \rightarrow \alpha_n \rightarrow \beta$, where some $\alpha_i = \gamma_1 \rightarrow \dots \rightarrow \gamma_m \rightarrow \beta$, $n > 0$ and $m \geq 0$. Huet-style projection is

$$F \mapsto \lambda \bar{x}_n. x_i (F_1 \bar{x}_n) \dots (F_m \bar{x}_n)$$

where the fresh free variables \bar{F}_m and bound variables \bar{x}_n are of appropriate types.

Imitation of g for F : Let F be a free variable of type $\alpha_1 \rightarrow \dots \rightarrow \alpha_n \rightarrow \beta$ and let g be a constant of type $\gamma_1 \rightarrow \dots \rightarrow \gamma_m \rightarrow \beta$ where $n, m \geq 0$. The imitation binding is

$$F \mapsto \lambda \bar{x}_n. g (F_1 \bar{x}_n) \dots (F_m \bar{x}_n)$$

where the fresh free variables \bar{F}_m and bound variables \bar{x}_n are of appropriate types.

Elimination for F : Let F be a free variable of type $\alpha_1 \rightarrow \dots \rightarrow \alpha_n \rightarrow \beta$, where $n > 0$. In addition, let $1 \leq j_1 < \dots < j_i \leq n$ and $i < n$. Elimination for the sequence $(j_k)_{k=1}^i$ is

$$F \mapsto \lambda \bar{x}_n. G x_{j_1} \dots x_{j_i}$$

where the fresh free variable G as well as all x_{j_k} are of appropriate type. We call fresh variables emerging from this binding in the role of *G elimination variables*.

Identification for F and G : Let F and G be different free variables. Furthermore, let the type of F be $\alpha_1 \rightarrow \dots \rightarrow \alpha_n \rightarrow \beta$ and the type of G be $\gamma_1 \rightarrow \dots \rightarrow \gamma_m \rightarrow \beta$, where $n, m \geq 0$. Then, the identification binding binds F and G with

$$F \mapsto \lambda \bar{x}_n. H \bar{x}_n (F_1 \bar{x}_n) \dots (F_m \bar{x}_n) \quad G \mapsto \lambda \bar{y}_m. H (G_1 \bar{y}_m) \dots (G_n \bar{y}_m) \bar{y}_m$$

where the fresh free variables H, \bar{F}_m, \bar{G}_n and bound variables \bar{x}_n, \bar{y}_m are of appropriate types. We call fresh variables emerging from this binding in the role of *H identification variables*.

Iteration for F : Let F be a free variable of the type $\alpha_1 \rightarrow \dots \rightarrow \alpha_n \rightarrow \beta_1$ and let some α_i be the type $\gamma_1 \rightarrow \dots \rightarrow \gamma_m \rightarrow \beta_2$, where $n > 0$ and $m \geq 0$. Iteration for F at i is

$$F \mapsto \lambda \bar{x}_n. H \bar{x}_n (\lambda \bar{y}. x_i (G_1 \bar{x}_n \bar{y}) \dots (G_m \bar{x}_n \bar{y}))$$

The free variables H and G_1, \dots, G_m are fresh, and \bar{y} is an arbitrary-length sequence of bound variables of arbitrary types. All new variables are of appropriate type. Due to indeterminacy of \bar{y} , this step is infinitely branching.

The following mapping $\mathcal{P}_c(\lambda \bar{x}. s \stackrel{?}{=} \lambda \bar{x}. t)$ is used as the parameter \mathcal{P} of the procedure:

- If the constraint is rigid-rigid, $\mathcal{P}_c(\lambda \bar{x}. s \stackrel{?}{=} \lambda \bar{x}. t) = \emptyset$.
- If the constraint is flex-rigid, let $\mathcal{P}_c(\lambda \bar{x}. F \bar{s} \stackrel{?}{=} \lambda \bar{x}. a \bar{t})$ be
 - an imitation of a for F , if a is a constant, and
 - all Huet-style projections for F , if F is not an identification variable.
- If the constraint is flex-flex and the heads are different, let $\mathcal{P}_c(\lambda \bar{x}. F \bar{s} \stackrel{?}{=} \lambda \bar{x}. G \bar{t})$ be
 - all identifications and iterations for both F and G , and
 - all JP-style projections for non-identification variables among F and G .
- If the constraint is flex-flex and the heads are identical, we distinguish two cases:
 - if the head is an elimination variable, $\mathcal{P}_c(\lambda \bar{x}. s \stackrel{?}{=} \lambda \bar{x}. t) = \emptyset$;
 - otherwise, let $\mathcal{P}_c(\lambda \bar{x}. F \bar{s} \stackrel{?}{=} \lambda \bar{x}. F \bar{t})$ be all iterations for F at arguments of functional type and all eliminations for F .

Comparison with the JP Procedure. In contrast to our procedure, the JP procedure constructs a tree with only one unification constraint per node and does not have a **Decompose** rule. Instead, at each node $(s \stackrel{?}{=} t, \sigma)$, the JP procedure computes the common context C of s and t , yielding term pairs $(s_1, t_1), \dots, (s_n, t_n)$, called *disagreement pairs*, such that $s = C[s_1, \dots, s_n]$ and $t = C[t_1, \dots, t_n]$. The procedure heuristically chooses one of the disagreement pairs (s_i, t_i) and applies a binding to the heads of s_i and t_i or to a free variable occurring above the disagreement pair in the common context C . Due to this application of bindings above the disagreement pair, lazy normalization and dereferencing cannot easily be integrated into the JP procedure.

The mapping \mathcal{P} uses many binding rules of the JP procedure, but our procedure explores the search space differently. In particular, the JP procedure allows iteration or elimination to be applied at a free variable in the common context of the unification constraint, even if bindings were already applied below that free variable. In contrast, our procedure forces the eliminations and iterations to be applied as soon as it observes a flex-flex pair with identical heads. After applying the **Decompose** transition, this flex-flex pair will be reduced to pairs representing the arguments of the identical heads. Therefore, unlike the JP algorithm, it will not apply bindings to the flex-flex pair after bindings have been applied to its arguments.

The JP procedure can be modified to solve the preunification problem by making it choose only flex-rigid disagreement pairs and terminate with a preunifier when no flex-rigid pair remains. However, such a procedure would be less efficient than Huet’s procedure because it would use the iteration binding instead of Huet-style projection to solve flex-rigid pairs. Our procedure applies Huet-style projections on flex-rigid pairs, which results in two important improvements over the JP procedure. First, our procedure terminates more often than the JP procedure because Huet-style projections cause only a finite branching, whereas iteration causes an infinite branching. Second, when our procedure is modified to solve the preunification problem by never selecting flex-flex pairs and stopping when only flex-flex pairs are left, it becomes an optimized variant of Huet’s procedure that supports oracles as well as lazy substitution and β -reduction.

The bindings of our procedure contain further optimizations that are absent in the JP procedure: The JP procedure applies eliminations for only one parameter at a time, yielding multiple paths to the same unifier. It applies imitations to flex-flex pairs, which we found to be unnecessary. Moreover, it does not keep track of which rules introduced which variables: iterations and eliminations are applied on elimination variables, and projections are applied on identification variables.

Examples. We present some illustrative derivations. The displayed branches of the constructed trees are not necessarily exhaustive. We abbreviate JP-style projection as **JP Proj**, imitation as **lmit**, identification as **ld**, **Decompose** as **Dc**, **Dereference** as **Dr**, **Normalize** $_{\beta}$ as **N** $_{\beta}$, and **Bind** of a binding x as **B**(x). Transitions of the JP procedure are denoted by \Longrightarrow . For the JP transitions we implicitly apply the generated bindings and fully normalize terms, which significantly shortens JP derivations.

Example 3.1. The JP procedure does not terminate on the problem $G \stackrel{?}{=} f G$:

$$(G \stackrel{?}{=} f G, \text{id}) \xrightarrow{\text{lmit}} (f G' \stackrel{?}{=} f^2 G', \sigma_1) \xrightarrow{\text{lmit}} (f^2 G'' \stackrel{?}{=} f^3 G'', \sigma_2) \xrightarrow{\text{lmit}} \dots$$

where $\sigma_1 = \{G \mapsto \lambda x. f G'\}$ and $\sigma_2 = \{G' \mapsto \lambda x. f G''\}\sigma_1$. By including any oracle that supports the first-order occurs check, such as the pattern oracle or the fixpoint oracle described in Section 7, our procedure gracefully generalizes first-order unification:

$$(\{G \stackrel{?}{=} f G\}, \text{id}) \xrightarrow{\text{OracleFail}} \perp$$

Example 3.2. The following derivation illustrates the advantage of the **Decompose** rule.

$$\begin{aligned} & (\{h^{100}(F a) \stackrel{?}{=} h^{100}(G b)\}, \text{id}) \xrightarrow{\text{Dc}^{100}} (\{F a \stackrel{?}{=} G b\}, \text{id}) \xrightarrow{\text{B}(\text{ld})} (\{F a \stackrel{?}{=} G b\}, \sigma_1) \\ & \xrightarrow{\text{Dr}+\text{N}_{\beta}} (\{H a (F' a) \stackrel{?}{=} H (G' b) b\}, \sigma_1) \xrightarrow{\text{Dc}} (\{a \stackrel{?}{=} G' b, F' a \stackrel{?}{=} b\}, \sigma_1) \\ & \xrightarrow{\text{B}(\text{lmit})} (\{a \stackrel{?}{=} G' b, F' a \stackrel{?}{=} b\}, \sigma_2) \xrightarrow{\text{Dr}+\text{N}_{\beta}} (\{a \stackrel{?}{=} a, F' a \stackrel{?}{=} b\}, \sigma_2) \xrightarrow{\text{Delete}} (\{F' a \stackrel{?}{=} b\}, \sigma_2) \\ & \xrightarrow{\text{B}(\text{lmit})} (\{F' a \stackrel{?}{=} b\}, \sigma_3) \xrightarrow{\text{Dr}+\text{N}_{\beta}} (\{b \stackrel{?}{=} b\}, \sigma_3) \xrightarrow{\text{Delete}} (\emptyset, \sigma_3) \xrightarrow{\text{Succeed}} \sigma_3 \end{aligned}$$

where $\sigma_1 = \{F \mapsto \lambda x. H x (F' x), G \mapsto \lambda y. H (G' y) y\}$; $\sigma_2 = \{G' \mapsto \lambda x. a\}\sigma_1$; and $\sigma_3 = \{F' \mapsto \lambda x. b\}\sigma_2$. The JP procedure produces the same intermediate substitutions σ_1 to σ_3 , but since it does not decompose the terms, it retraverses the common context $h^{100}[\]$ at every step to identify the contained disagreement pair:

$$\begin{aligned}
& (\mathbf{h}^{100}(F \mathbf{a}) \stackrel{?}{=} \mathbf{h}^{100}(G \mathbf{b}), \text{id}) \xrightarrow{\text{Id}} (\mathbf{h}^{100}(H \mathbf{a}(F' \mathbf{a})) \stackrel{?}{=} \mathbf{h}^{100}(H(G' \mathbf{b}) \mathbf{b}), \sigma_1) \\
& \xrightarrow{\text{Imit}} (\mathbf{h}^{100}(H \mathbf{a}(F' \mathbf{a})) \stackrel{?}{=} \mathbf{h}^{100}(H \mathbf{a} \mathbf{b}), \sigma_2) \xrightarrow{\text{Imit}} (\mathbf{h}^{100}(H \mathbf{a} \mathbf{b}) \stackrel{?}{=} \mathbf{h}^{100}(H \mathbf{a} \mathbf{b}), \sigma_3) \xrightarrow{\text{Succeed}} \sigma_3
\end{aligned}$$

Example 3.3. The search space restrictions also allow us to prune some redundant unifiers. Consider the problem $F(G \mathbf{a}) \stackrel{?}{=} F \mathbf{b}$, where \mathbf{a} and \mathbf{b} are of base type. Our procedure produces only one failing branch and the following two successful branches:

$$\begin{aligned}
& (\{F(G \mathbf{a}) \stackrel{?}{=} F \mathbf{b}\}, \text{id}) \xrightarrow{\text{Dc}} (\{G \mathbf{a} \stackrel{?}{=} \mathbf{b}\}, \text{id}) \xrightarrow{\text{B(Imit)}} (\{G \mathbf{a} \stackrel{?}{=} \mathbf{b}\}, \{G \mapsto \lambda x. \mathbf{b}\}) \\
& \xrightarrow{\text{Dr+N}_\beta} (\{\mathbf{b} \stackrel{?}{=} \mathbf{b}\}, \{G \mapsto \lambda x. \mathbf{b}\}) \xrightarrow{\text{Delete}} (\emptyset, \{G \mapsto \lambda x. \mathbf{b}\}) \xrightarrow{\text{Succeed}} \{G \mapsto \lambda x. \mathbf{b}\} \\
& (\{F(G \mathbf{a}) \stackrel{?}{=} F \mathbf{b}\}, \text{id}) \xrightarrow{\text{B(Elim)}} (\{F(G \mathbf{a}) \stackrel{?}{=} F \mathbf{b}\}, \{F \mapsto \lambda x. F'\}) \\
& \xrightarrow{\text{Dr+N}_\beta} (\{F' \stackrel{?}{=} F'\}, \{F \mapsto \lambda x. F'\}) \xrightarrow{\text{Delete}} (\emptyset, \{F \mapsto \lambda x. F'\}) \xrightarrow{\text{Succeed}} \{F \mapsto \lambda x. F'\}
\end{aligned}$$

The JP procedure additionally produces the following redundant unifier:

$$\begin{aligned}
& (F(G \mathbf{a}) \stackrel{?}{=} F \mathbf{b}, \text{id}) \xrightarrow{\text{JP Proj}} (F \mathbf{a} = F \mathbf{b}, \{G \mapsto \lambda x. x\}) \\
& \xrightarrow{\text{Elim}} (F' = F', \{G \mapsto \lambda x. x, F \mapsto \lambda x. F'\}) \xrightarrow{\text{Succeed}} \{G \mapsto \lambda x. x, F \mapsto \lambda x. F'\}
\end{aligned}$$

Moreover, the JP procedure does not terminate because an infinite number of iterations is applicable at the root. Our procedure terminates in this case since we only apply iteration binding for non base-type arguments, which F does not have.

Pragmatic Variant. We structured our procedure so that most of the unification machinery is contained in the **Bind** step. Modifying \mathcal{P} , we can sacrifice completeness and obtain a pragmatic variant of the procedure that often performs better in practice. Our preliminary experiments showed that using mapping \mathcal{P}_p defined as follows is a reasonable compromise between completeness and performance:

- If the constraint is rigid-rigid, $\mathcal{P}_p(\lambda \bar{x}. s \stackrel{?}{=} \lambda \bar{x}. t) = \emptyset$.
- If the constraint is flex-rigid, let $\mathcal{P}_p(\lambda \bar{x}. F \bar{s} \stackrel{?}{=} \lambda \bar{x}. a \bar{t})$ be
 - an imitation of a for F , if a is a constant, and
 - all Huet-style projections for F if F is not an identification variable.
- If the constraint is flex-flex and the heads are different, let $\mathcal{P}_p(\lambda \bar{x}. F \bar{s} \stackrel{?}{=} \lambda \bar{x}. G \bar{t})$ be
 - an identification binding for F and G , and
 - all Huet-style projections for F if F is not an identification variable
- If the constraint is flex-flex and the heads are identical, we distinguish two cases:
 - if the head is an elimination variable, $\mathcal{P}_p(\lambda \bar{x}. F \bar{s} \stackrel{?}{=} \lambda \bar{x}. F \bar{t}) = \emptyset$;
 - otherwise, let $\mathcal{P}_p(\lambda \bar{x}. F \bar{s} \stackrel{?}{=} \lambda \bar{x}. F \bar{t})$ be the set of all eliminations bindings for F .

The pragmatic variant of our procedure removes all iteration bindings to enforce finite branching. Moreover, it imposes limits on the number of bindings applied, counting the applications of bindings locally, per constraint. It is useful to distinguish the Huet-style projection cases where α_i is a base type (called *simple projection*), which always reduces the problem size, and the cases where α_i is a functional type (called *functional projection*). We limit the number applications of the following bindings: functional projections, eliminations, imitations and identifications. In addition, a limit on the total number of applied bindings

can be set. An elimination binding that removes k arguments counts as k elimination steps. Due to these limits, the pragmatic variant terminates.

To fail as soon as any of the limits is reached, the pragmatic variant employs an additional oracle. If this oracle determines that the limits are reached and the constraint is of the form $\lambda\bar{x}. F \bar{s}_m \stackrel{?}{=} \lambda\bar{x}. G \bar{t}_n$, it returns a *trivial unifier* – a substitution $\{F \mapsto \lambda\bar{x}_m. H, G \mapsto \lambda\bar{x}_n. H\}$, where H is a fresh variable; if the limits are reached and the constraint is flex-rigid, the oracle fails; if the limits are not reached, it reports that terms are outside its fragment. The trivial unifier prevents the procedure from failing on easily unifiable flex-flex pairs.

Careful tuning of each limit optimizes the procedure for a specific class of problems. For problems originating from proof assistants, shallow unification depth usually suffices. However, hard hand-crafted problems often need deeper unification.

4. PROOF OF COMPLETENESS

Like the JP procedure, our procedure misses no unifiers:

Theorem 4.1. *The procedure described in Section 3 parametrized by \mathcal{P}_c is complete, meaning that the substitutions on the leaves of the constructed tree form a CSU. In other words, for any unifier ϱ of a multiset of constraints E there exists a derivation $(E, \text{id}) \longrightarrow^* \sigma$ and a substitution θ such that $\varrho \subseteq \theta\sigma$.*

Taking a high-level view, this theorem is proved by incrementally defining states (E_j, σ_j) and *remainder substitutions* ϱ_j starting with $(E_0, \sigma_0) = (E, \text{id})$ and $\varrho_0 = \varrho$. Intuitively, ϱ_j is what remains to be added to σ_j to reach a unifier more general than ϱ_0 . States are defined so that the shape of the selected constraint from E_j and the remainder substitution guide the choice of applicable transition rule of our procedure. We employ a measure based on values of E_j and ϱ_j that decreases with each application of the rules. Therefore, eventually, we will reach the target substitution σ .

In the remaining of this section, we view terms as $\alpha\beta\eta$ -equivalence classes, with the η -long β -normal form as their canonical representative. Moreover, we consider all substitutions to be fully applied. These assumptions are justified because all bindings depend only on the head of terms and hence replacing the lazy transitions **Normalize** $_{\alpha\eta}$, **Normalize** $_{\beta}$, and **Dereference** by eager counterparts only affects the efficiency but not the overall behavior of our procedure.

We now give the detailed completeness proof of Theorem 4.1. Our proof is an adaptation of the proof given by Jensen and Pietrzykowski [13]. Definitions and lemmas are reused, but are combined together differently to suit our procedure. We start by listing all reused definitions and lemmas from the original JP proof. The “JP” labels in their statements refer to the corresponding lemmas and definitions from the original proof.

Definition 4.2 (JP D1.6). Given two terms t and s and their common context C , we can write t as $C[\bar{t}]$ and s as $C[\bar{s}]$ for some \bar{t} and \bar{s} . The pairs (s_j, t_j) are called *disagreement pairs*.

Definition 4.3 (JP D3.1). Given two terms t and s , let $\lambda\bar{x}. t'$ and $\lambda\bar{y}. s'$ be respective α -equivalent terms such that their parameters \bar{x} and \bar{y} are disjoint. Then the disagreement pairs of t' and s' are called *opponent pairs* in t and s .

Lemma 4.4 (JP L3.3 (1)). *Let ρ be a substitution and X, Y be free variables such that $\rho(X \bar{s}) = \rho(Y \bar{t})$ for some term tuples \bar{s} and \bar{t} . Then for every opponent pair u, v in ρX and ρY (Definition 4.3), the head of u or v is a parameter of ρX or ρY .*

In contrast to first order, a higher order context $C[]$ need not to be injective as a mapping from terms to terms. As a result, eagerly decomposing $C[s] \stackrel{?}{=} C[t]$ into $s \stackrel{?}{=} t$ is not always complete. The concept of ω -simplicity captures an important subset of injective contexts $C[]$. Non- ω -simplicity on the other hand is the main trigger of **Iteration**—the most explosive binding of our procedure.

Definition 4.5 (JP D3.2). An occurrence of a parameter x of term t in the body of t is ω -simple if both

- (1) the arguments of x are distinct and are exactly (the η -long forms of) all of the variables bound in the body of t , and
- (2) this occurrence of x is not in an argument of any parameter of t .

This definition is slightly too restrictive for our purposes. It is unfortunate that condition 1 requires x to be applied to *all* instead of just some of the bound variables. The JP proof would probably work with such a relaxation, and the definition would be equivalent to injectivity with respect to the parameter for any values of other parameters. However, to reuse the JP lemmas, we stick to the original notion of ω -simplicity and introduce the following relaxation:

Definition 4.6. An occurrence of a parameter x of term t in the body of t is *base-simple* if it is ω -simple or both

- (1) x is of base type, and
- (2) this occurrence of x is not in an argument of any parameter of t .

Lemma 4.7. *Let s have parameters \bar{x} and a subterm $x_j \bar{v}$ where this occurrence of x_j is base-simple. Then for any sequence \bar{t} of (at least j) terms, the body of t_j is a subterm of $s \bar{t}$ (after normalization) at the position of $x_j \bar{v}$ up to renaming of the parameters of t_j . To compare positions of s and $s \bar{t}$, ignore the parameter count mismatch.*

Proof. Consider the process of β -normalizing $s \bar{t}$. After substituting terms \bar{t} into the body of s , a further reduction can only take place when some t_k is an abstraction that gets arguments in s . The arguments \bar{v} to the x_j are distinct variables bound in the body of s . This follows easily from either case of the definition of base-simplicity. So t_j is applied to the unmodified \bar{v} after substituting terms \bar{t} into the body of s . Base-simplicity also implies that $t_j \bar{v}$ does not occur in an argument to another t_k . Hence only the reduction of $t_j \bar{v}$ itself affects this subterm. The variables \bar{v} match the parameter count of t_j because we consider the η -long form of t_j ; so $t_j \bar{v}$ reduces to the body of t_j (modulo renaming). The position is obviously that of $x_j \bar{v}$. \square

Lemma 4.8 (JP C3.4 strengthened). *Let ρ be a substitution and X a free variable. If $\rho(\lambda \bar{x}. X \bar{s}) = \rho(\lambda \bar{x}. X \bar{t})$ and some occurrence of the i^{th} parameter of ρX is base-simple, then $\rho s_i = \rho t_i$.*

Proof. By Lemma 4.7, ρs_i occurs in $\rho X(\rho \bar{s})$ at certain position that depends only on ρX . Similarly ρt_i occurs in $\rho X(\rho \bar{t}) = \rho X(\rho \bar{s})$ at the same position, and hence $\rho s_i = \rho t_i$. \square

We define more properties to determine which binding to apply to a given constraint. Roughly speaking, the simple comparison form will trigger identification bindings, projectivity will trigger Huet-style projections, and simple projectivity will trigger JP-style projections.

Definition 4.9 (JP D3.4). We say that s and t are in *simple comparison form* if all ω -simple heads of opponent pairs in s and t are distinct, and each opponent pair has an ω -simple head.

Definition 4.10 (JP D3.5). A term t is called *projective* if the head of t is a parameter of t . If the whole body is just the parameter, then t is called *simply projective*.

A central part of the proof is to find a suitable measure for the remaining problem size. Showing that the measure is strictly decreasing and well-founded guarantees that the procedure finds a suitable substitution in finitely many steps. We reuse the measure for remainder substitutions from JP [13], but embed it into a lexicographic measure to handle the decomposition steps and oracles of our procedure.

Definition 4.11 (JP D3.7). The *free weight* of a term t is the total number of occurrences of free variables and constants in t . The *bound weight* of t is the total number of occurrences (excluding occurrences λx) of bound variables in t , but with the particular exemption: if a prefix variable u has one or more ω -simple occurrences in the body, then one such occurrence and its arguments are not counted. It does not matter which occurrence is not counted because in η -long form the bound weight of the arguments of an ω -simple variable is the same for all occurrences of that variable.

Definition 4.12 (JP D3.8). For multisets E of unification constraints and substitutions ϱ , our measure on pairs (E, ϱ) is the lexicographic comparison of

A: the sum of the sizes of the terms in E

B: the sum over the free weight of ϱF , for all variables F mapped by ϱ

C: the sum over the bound weight of ϱF , for all variables F mapped by ϱ

D: the sum over the number of parameters of ϱF , for all variables F mapped by ϱ

We denote the quadruple containing these numbers as $\text{ord}(E, \varrho)$. We denote the triple containing only the last three components of $\text{ord}(E, \varrho)$ as $\text{ord } \varrho$. We write $<$ for the lexicographic comparison of these tuples.

The next six lemmas correspond to the bindings of our procedure and sufficient conditions for the binding to bring us closer to a given solution. This is expressed as a decrease of the ord measure of the remainder. In each of these lemmas, let u be a term with a variable head a and v a term with an arbitrary head b . Let ϱ be a unifier of u and v . The conclusion, let us call it **C**, is always the same: there exists a binding δ applicable to the problem $u \stackrel{?}{=} v$, and there exists a substitution ϱ' such that $\varrho \subseteq \varrho' \delta$ and $\text{ord } \varrho' < \text{ord } \varrho$. For most of these lemmas, we refer to JP [13] for proofs. However, our bindings have more preconditions, yielding additional orthogonal hypotheses in our lemmas.

Lemma 4.13 (JP L3.9). *If $a = b$ is not an elimination variable and ϱa discards any of its parameters, then **C** by elimination. Moreover, for the elimination variable G introduced by this elimination, $\varrho' G$ discards none of its parameters and has the body of ϱa .*

Proof. Let $\varrho a = \lambda \bar{x}. t$ and let $(x_{j_k})_{k=1}^i$ be the subsequence of \bar{x} consisting of those variables which occur in the body t . It is a strict subsequence, since ϱa is assumed to discard some parameter. Since the equal heads $a = b$ of the constraint $u \stackrel{?}{=} v$ are not elimination variables, elimination for $(j_k)_{k=1}^i$ can be applied. Let $\delta = \{a \mapsto \lambda \bar{x}. G x_{j_1} \dots x_{j_i}\}$ be the corresponding binding. Define ϱ' to be like ϱ except

$$\varrho' a = a \quad \text{and} \quad \varrho' G = \lambda x_{j_1} \dots x_{j_i}. t.$$

Obviously $\varrho'G$ is a closed term and $\varrho \subseteq \varrho' \delta$ holds. Moreover $\text{ord } \varrho' < \text{ord } \varrho$, because free and bound weights stay the same (ϱa and $\varrho'G$ have the same body t) whereas the number of parameters strictly decreases. The definition of $(j_k)_{k=1}^i$ implies that $\varrho'G$ discards none of its parameters. \square

Lemma 4.14 (JP L3.10). *Assume that there exists a parameter x of ϱa such that x has a non- ω -simple (Definition 4.5) occurrence in ϱa , which is not below another parameter, or such that x has at least two ω -simple occurrences in ϱa . Moreover, if $a = b$, to make iteration applicable, a must not be an elimination variable, and x must be of functional type. Then \mathbf{C} is achieved by iteration.*

Lemma 4.15 (JP L3.11). *Assume that a and b are different free variables. If ϱa is simply projective (Definition 4.10) and a is not an identification variable, then \mathbf{C} by JP-style projection.*

Lemma 4.16 (JP L3.12). *If ϱa is not projective and b is rigid, then \mathbf{C} by imitation.*

Lemma 4.17 (JP L3.13). *Let $a \neq b$. Assume that $\varrho a \neq a$ and $\varrho b \neq b$ are in simple comparison form (Definition 4.9) and neither is projective. Then \mathbf{C} by identification. Moreover, $\varrho'H$ is not projective, where H is the identification variable introduced by this application of the identification binding.*

Proof. This is JP's Lemma 3.13, plus the claim that $\varrho'H$ is not projective. Inspecting the proof of that lemma, it is obvious that $\varrho'H$ cannot be projective because ϱa and ϱb are not projective. \square

Lemma 4.18. *Assume that ϱa is projective (Definition 4.10), a is not an identification variable, and b is rigid. Then \mathbf{C} by Huet-style projection.*

Proof. Since ϱa is projective, we have $\varrho a = \lambda \bar{x}_n. x_k \bar{t}_m$ for some k and some terms \bar{t}_m . If ϱa is also simply projective, then x_k must be non-functional and since Huet-style projection and JP-style projection coincide in that case, Lemma 4.15 applies. Hence, in the following we may assume that ϱa is not simply projective, i.e., that $m > 0$.

Let δ be the Huet-style projection binding:

$$\delta = \{a \mapsto \lambda \bar{x}_n. x_i (F_1 \bar{x}_n) \dots (F_m \bar{x}_n)\}$$

for fresh variables F_1, \dots, F_m . This binding is applicable because b is rigid. Let ϱ' be the same as ϱ except that we set $\varrho'a = a$ and for each $1 \leq j \leq m$ we set

$$\varrho'F_j = \lambda \bar{x}_n. t_j$$

It remains to show that $\text{ord } \varrho' < \text{ord } \varrho$. The free weight of ϱa is the same as the sum of the free weights of $\varrho'F_j$ for $1 \leq j \leq m$. Thus, the free weight is the same for ϱ and ϱ' . The bound weight of ϱa however is exactly 1 larger than the sum of the bound weights of $\varrho'F_j$ for $1 \leq j \leq m$ because of the additional occurrence of x_k in ϱa . The exemption for ω -simple occurrences in the definition of the bound weight cannot be triggered by this occurrence of x_k because $m > 0$ and thus x_k is not ω -simple. It follows that $\text{ord } \varrho' < \text{ord } \varrho$. \square

We are now ready to prove the completeness theorem (Theorem 4.1).

Proof. Let $E_0 = E$ and $\sigma_0 = \text{id}$. Let $\varrho_0 = \tau \varrho$ for some renaming τ , such that the free variables of $\varrho_0 E$ and E are disjoint. Then ϱ_0 unifies E_0 because ϱ unifies E by assumption. Moreover, $\varrho_0 = \varrho_0 \sigma_0$. We proceed to inductively define E_j , σ_j and ϱ_j until we reach some j

such that $E_j = \emptyset$. To guarantee well-foundedness, we ensure that the measure $\text{ord}(\varrho_j, E_j)$ decreases with each step. We maintain the following invariants for all j :

- $(E_j, \sigma_j) \longrightarrow (E_{j+1}, \sigma_{j+1})$;
- $\text{ord}(\varrho_j, E_j) > \text{ord}(\varrho_{j+1}, E_{j+1})$;
- ϱ_j unifies E_j ;
- $\varrho_0 \subseteq \varrho_j \sigma_j$;
- the free variables in $\varrho_j E_j$ and E_j are disjoint;
- for every identification variable X , $\varrho_j X$ is not projective; and
- for every elimination variable X , each parameter of $\varrho_j X$ has occurrences in $\varrho_j X$, all of which are base-simple.

If $E_j \neq \emptyset$, let $u \stackrel{?}{=} v$ be the selected constraint $S(E_j)$ in E_j .

First assume that an oracle is able to find a CSU for the constraint $u \stackrel{?}{=} v$. Since ϱ_j unifies u and v , by the definition of a CSU, the CSU discovered by the oracle contains a unifier δ of u and v such that there exists ϱ_{j+1} satisfying $\varrho_j \subseteq \varrho_{j+1} \delta$. Thus, an `OracleSucc` transition is applicable and yields the node $(E_{j+1}, \sigma_{j+1}) = (\delta(E_j \setminus \{u \stackrel{?}{=} v\}), \delta \sigma_j)$. The strict containment $\varrho_{j+1} E_{j+1} \subset \varrho_{j+1} \delta E_j = \varrho_j E_j$ implies $\text{ord}(E_{j+1}, \varrho_{j+1}) < \text{ord}(E_j, \varrho_j)$. It also shows that the constraints $\varrho_{j+1} E_{j+1}$ are unified when $\varrho_j E_j$ are. By direct computation $\varrho_0 \subseteq \varrho_j \sigma_j \subseteq \varrho_{j+1} \delta \sigma_j = \varrho_{j+1} \sigma_{j+1}$. Variable disjointness follows by picking ϱ_{j+1} in the right way. New identification or elimination variables are not introduced; so their properties are preserved. Hence all invariants are preserved.

Otherwise we proceed by a case distinction on the form of $u \stackrel{?}{=} v$. Typically, one of the Lemmas 4.13–4.18 is going to apply. Any one of them gives substitutions ϱ' and δ with properties that let us define $E_{j+1} = \delta E_j$, $\sigma_{j+1} = \delta \sigma_j$ and $\varrho_{j+1} = \varrho'$. The problem size always strictly decreases, because these lemmas imply $\varrho_{j+1} E_{j+1} = \varrho_{j+1} \delta E_j = \varrho_j E_j$ and $\text{ord} \varrho_{j+1} = \text{ord} \varrho' < \text{ord} \varrho_j$. Regarding the other invariants, the former equation guarantees that ϱ_{j+1} unifies E_{j+1} , and $\varrho_0 \subseteq \varrho_j \sigma_j \subseteq \varrho_{j+1} \delta \sigma_j = \varrho_{j+1} \sigma_{j+1}$ holds by computation. The conditions on variables must be checked separately when new ones are introduced. Let a be the head of $u = \lambda \bar{x}. a \bar{u}$ and b be the head of $v = \lambda \bar{x}. b \bar{v}$. Consider the following cases:

u and v have the same head symbol $a = b$:

- (1) Suppose that $\varrho_j a$ has a parameter with non-base-simple occurrence. By one of the induction invariants, a is not an elimination variable. Among all non-base-simple occurrences of parameters in $\varrho_j a$, choose the leftmost one, which we call x . This occurrence of x cannot be below another parameter, because having x occur in one of its arguments would make that other parameter non-base-simple, contradicting the occurrence of x being leftmost. Thus x is neither base-simple nor below another parameter; so x is of functional type. Moreover, non-base-simplicity implies non- ω -simplicity. Hence, we can apply Lemma 4.14 (iteration).
- (2) Otherwise suppose that $\varrho_j a$ discards some of its parameters. By one of the induction invariants, $\varrho_j a$ is not an elimination variable. Hence Lemma 4.13 (elimination) applies. The newly introduced elimination variable G satisfies the required invariants, because Lemma 4.13 guarantees that $\varrho_{j+1} G$ uses its parameters and shares the body with $\varrho_j a$ which by assumption of this case contains only base-simple occurrences.
- (3) Otherwise every parameter of $\varrho_j a$ has occurrences and all of them are base-simple. We are going to show that `Decompose` is a valid transition and decreases $\varrho_j E_j$. By Lemma 4.8 we conclude from $\varrho_j u = \varrho_j v$ that $\varrho_j u_i = \varrho_j v_i$ for every i . Hence the new constraints $E_{j+1} = E_j \setminus \{u \stackrel{?}{=} v\} \cup \{u_i \stackrel{?}{=} v_i \mid \text{for all } i\}$ after `Decompose` are unified by ϱ_j .

This allows us to define $\varrho_{j+1} = \varrho_j$ and $\sigma_{j+1} = \sigma_j$. To check that $\varrho_{j+1}E_{j+1} = \varrho_jE_{j+1}$ is smaller than ϱ_jE_j it suffices to check that constraints $\varrho_ju_i \stackrel{?}{=} \varrho_jv_i$ together are smaller than $\varrho_ju \stackrel{?}{=} \varrho_jv$. Since all parameters of ϱ_ja have base-simple occurrences, ϱ_ju_i is a subterm of $\varrho_ju = \lambda\bar{x}.\varrho_ja(\varrho_j\bar{u})$ by Lemma 4.7. Similarly for ϱ_jv . It follows that $\varrho_{j+1}E_{j+1}$ is smaller than ϱ_jE_j . Since $\varrho_{j+1} = \varrho_j$ and $\sigma_{j+1} = \sigma_j$, the other invariants are obviously preserved.

u and v is a flex-flex pair with different heads:

- (5) First, suppose that ϱ_ja or ϱ_jb is simply projective (Definition 4.10). By the induction hypothesis, the simply projective head cannot be an identification variable. Thus Lemma 4.15 (JP-style projection) applies.
- (6) Otherwise suppose that ϱ_ja is projective but not simply. Then the head of ϱ_ja is its some parameter x_k . But this occurrence cannot be ω -simple because it has arguments which cannot be bound above the head x_k . Thus Lemma 4.14 (iteration) applies. If ϱ_jb is projective but not simply, the same argument applies.
- (7) Otherwise suppose that ϱ_ja, ϱ_jb are in simple comparison form (Definition 4.9). By one of the the induction invariants, the free variables in ϱ_jE_j and E_j are disjoint. Thus $\varrho_ja \neq a$ and $\varrho_jb \neq b$. Then Lemma 4.17 (identification) applies.
- (8) Otherwise ϱ_ja, ϱ_jb are not in simple comparison form. By Lemma 4.4 and by the definition of simple comparison form, there is some opponent pair $x_k\bar{r}, b$ in ϱ_ja and ϱ_jb (after possibly swapping u and v) where either the occurrence of x_k is not ω -simple (Definition 4.5) or else x_k has another ω -simple occurrence in the body of ϱ_ja . Then Lemma 4.14 (iteration) applies.

u and v is a flex-rigid pair: Without loss of generality, assume that a is flex and b is rigid.

- (9) Suppose first that ϱ_ja is projective. By one of the induction invariants, a cannot be an identification variable. Thus Lemma 4.18 (Huet-style projection) applies.
- (10) Otherwise ϱ_ja is not projective. The head of ϱ_ja must be b because b is rigid, and ϱ_j unifies u and v . Since ϱ_ja is not projective, that means that b is not a bound variable. Therefore, b must be a constant. Then Lemma 4.16 (imitation) applies.

We have now constructed a run $(E_0, \sigma_0) \longrightarrow (E_1, \sigma_1) \longrightarrow (E_2, \sigma_2) \longrightarrow \dots$ of the procedure. This run cannot be infinite because because the measure $\text{ord}(E_j, \varrho_j)$ strictly decreases as j increases. Hence, at some point we reach a j such that $E_j = \emptyset$ and $\varrho_0 \subseteq \varrho_j \sigma_j$. Therefore, $(E, \text{id}) \longrightarrow^* (\emptyset, \sigma_j) \longrightarrow \sigma_j$ and $\varrho \subseteq \tau^{-1} \varrho_j \sigma_j$, completing the proof. \square

5. A NEW DECIDABLE FRAGMENT

We discovered a new fragment that admits a finite CSU and a simple oracle. The oracle is based on work by Prehofer and the PT procedure [21], a modification of Huet's procedure. PT transforms an initial multiset of constraints E_0 by applying bindings ϱ . If there is a sequence $E_0 \Longrightarrow^{\varrho_1} \dots \Longrightarrow^{\varrho_n} E_n$ such that E_n has only flex-flex constraints, we say that PT produces a preunifier $\sigma = \varrho_n \dots \varrho_1$ with constraints E_n . A sequence fails if $E_n = \perp$. As in the previous section, we consider all terms to be $\alpha\beta\eta$ -equivalence classes with the η -long β -reduced form as their canonical representative. Unlike previously, in this section we view unification constraints $s \stackrel{?}{=} t$ as ordered pairs.

The following rules, however, are stated modulo orientation. The PT transition rules, adapted for our presentation style, are as follows:

Deletion: $\{s \stackrel{?}{=} s\} \uplus E \xRightarrow{\text{id}} E$

Decomposition: $\{\lambda\bar{x}. a \bar{s}_m \stackrel{?}{=} \lambda\bar{x}. a \bar{t}_m\} \uplus E \xRightarrow{\text{id}} \{s_1 \stackrel{?}{=} t_1, \dots, s_m \stackrel{?}{=} t_m\} \uplus E$

where a is rigid

Failure: $\{\lambda\bar{x}. a \bar{s} \stackrel{?}{=} \lambda\bar{x}. b \bar{t}\} \uplus E \xRightarrow{\text{id}} \perp$

where a and b are different rigid heads

Solution: $\{\lambda\bar{x}. F \bar{x} \stackrel{?}{=} \lambda\bar{x}. t\} \uplus E \xRightarrow{\varrho} \varrho(E)$

where F does not occur in t , t does not have a flex head, and $\varrho = \{F \mapsto \lambda\bar{x}. t\}$

Imitation: $\{\lambda\bar{x}. F \bar{s}_m \stackrel{?}{=} \lambda\bar{x}. f \bar{t}_n\} \uplus E \xRightarrow{\varrho} \varrho(\{G_1 \bar{s}_m \stackrel{?}{=} t_1, \dots, G_n \bar{s}_m \stackrel{?}{=} t_n\} \uplus E)$

where $\varrho = \{F \mapsto \lambda\bar{x}_m. f(G_1 \bar{x}_m) \dots (G_n \bar{x}_m)\}$, \bar{G}_n are fresh variables of appropriate types

Projection: $\{\lambda\bar{x}. F \bar{s}_m \stackrel{?}{=} \lambda\bar{x}. a \bar{t}\} \uplus E \xRightarrow{\varrho} \varrho(\{s_i(G_1 \bar{s}_m) \dots (G_j \bar{s}_m) \stackrel{?}{=} a \bar{t}\} \uplus E)$

where $\varrho = \{F \mapsto \lambda\bar{x}_m. x_i(G_1 \bar{x}_m) \dots (G_j \bar{x}_m)\}$, \bar{G}_j are fresh variables of appropriate types

The `grayed` constraints are required to be selected by a given selection function S . We call S *admissible* if it prioritizes selection of constraints applicable for **Failure** and **Decomposition**, and of descendant constraints of **Projection** transitions with $j = 0$ (i.e., for x_i of base type), in that order of priority. In the remainder of this section we consider only admissible selection functions, an assumption that Prehofer also makes implicitly in his thesis. Additionally, whenever we compare multisets, we use the multiset ordering defined by Dershowitz and Manna [8].

The following Lemma states that PT is *complete for preunification*:

Lemma 5.1. *Let ϱ be a unifier of a multiset of constraints E_0 . Then PT produces a preunifier σ with constraints E_n , and there exists a unifier θ of E_n such that $\varrho \subseteq \theta\sigma$.*

Proof. This lemma is a refinement of Lemma 4.1.7 from Prehofer's PhD thesis [21], and this proof closely follows the proof of that lemma.

We prove the lemma by induction using a well-founded measure on (ϱ, E_0) . The ordering is the lexicographic comparison of the following properties:

- A:** sum of the abstraction-free sizes of the terms ϱF for each variable F mapped by ϱ
- B:** multiset of the sizes of constraints in E_0

Here, the *abstraction-free size* is inductively defined by $\text{afsize}(F) = 1$; $\text{afsize}(x) = 1$; $\text{afsize}(f) = 1$; $\text{afsize}(st) = \text{afsize}(s) + \text{afsize}(t)$; $\text{afsize}(\lambda x. s) = \text{afsize}(s)$.

If E_0 consists only of flex-flex constraints, then we can take an empty transition sequence (i.e., $\sigma = \text{id}$) and $\theta = \varrho$. Otherwise, there exists a constraint that is selected by an admissible selection function. We show that for each such constraint, there is going to be a PT transition bringing us closer to the desired preunifier.

Let $E_0 = \{s \stackrel{?}{=} t\} \uplus E'_0$. Since $s \stackrel{?}{=} t$ is selected, at least one of s and t must have a rigid head. We distinguish several cases based on the form of $s \stackrel{?}{=} t$ (modulo the constraint order):

- $s \stackrel{?}{=} s$: in this case, **Deletion** applies. This transition does not alter ϱ , but removes an equation obtaining E_1 which is smaller than E_0 and still unifiable by ϱ . By induction hypothesis, the preunifier is reachable.
- $\lambda\bar{x}. a \bar{s}_m \stackrel{?}{=} \lambda\bar{x}. b \bar{t}_n$, where a and b are rigid: since ϱ is a unifier, then $a = b$ and $n = m$. Therefore, **Decomposition** applies. Similarly to the above case, we conclude that the preunifier is reachable.
- $\lambda\bar{x}. F \bar{x} \stackrel{?}{=} \lambda\bar{x}. t$ where t has a rigid head and F does not occur in t : **Solution** applies. Huet showed [12, proof of L5.1] that in this case $\varrho = \varrho\sigma_1$, where $\sigma_1 = \{F \mapsto \lambda\bar{x}. t\}$. Let $\varrho_1 = \varrho \setminus \{F \mapsto \varrho(F)\}$. Then $\varrho = \varrho_1\sigma_1$. Since ϱ unifies E'_0 , ϱ_1 unifies $E_1 = \sigma_1(E'_0)$. Clearly, ϱ_1 is smaller than ϱ . Now we can apply induction hypothesis on ϱ_1 and E_1 , from which we obtain a sequence $E_1 \xRightarrow{\sigma_2} \dots \xRightarrow{\sigma_n} E_n$, θ which unifies E_n and $\sigma' = \sigma_n \dots \sigma_2$,

such that $\varrho_1 \subseteq \theta\sigma'$. Finally, it is clear that sequence $E_0 \Longrightarrow^{\sigma_1} E_1 \Longrightarrow^{\sigma_2} \dots \Longrightarrow^{\sigma_n} E_n$, θ , and $\sigma = \sigma'\sigma_1$ are as wanted in the lemma statement.

- $\lambda\bar{x}. F \bar{s}_m \stackrel{?}{=} \lambda\bar{x}. a \bar{t}$ where a is rigid: depending on the value of ϱ , we can either take an **Imitation** or **Projection** step. In either case, we show that the measure reduces. Let $\varrho(F) = \lambda\bar{x}_m. b \bar{u}_n$ where b is either a constant or a bound variable. If b is a constant, we choose the **Imitation** step; otherwise we choose the **Projection** step. In either case, $\sigma_1 = \{F \mapsto \lambda\bar{x}_m. b (G_1 \bar{x}_m) \dots (G_n \bar{x}_m)\}$. Then it is easy to check that $\varrho_1 = \varrho \setminus \{F \mapsto \lambda\bar{x}_m. b \bar{u}_n\} \cup \{G_1 \mapsto \lambda\bar{x}_m. u_1, \dots, G_n \mapsto \lambda\bar{x}_m. u_n\}$ unifies E_1 created as the result of imitation or projection. Clearly, ϱ_1 has smaller abstraction-free size than ϱ . Therefore, by the induction hypothesis we obtain the transitions $E_1 \Longrightarrow^{\sigma_2} \dots \Longrightarrow^{\sigma_n} E_n$ and substitutions θ and $\sigma' = \sigma_n \dots \sigma_2$, where $\varrho_1 \subseteq \theta\sigma'$ and θ is a unifier of E_n . Our goal is to prove $\varrho \subseteq \theta\sigma'\sigma_1$. The only variable that ϱ maps and on which ϱ and ϱ_1 differ is F . Since ϱ_1 and $\theta\sigma'$ agree on G_1, \dots, G_n , we conclude that ϱ and $\theta\sigma'\sigma_1$ agree on F . Therefore, by taking $\sigma = \sigma'\sigma_1$, and prepending E_0 to the sequence from the induction hypothesis, we show that the preunifier is reachable. \square

Prehofer showed that PT terminates for some classes of constraints. We call a term *linear* if no free variable has repeated occurrences in it. We call a term *solid* if the arguments of its free variables are either bound variables or ground base-type terms. We call it *strictly solid*¹ if the arguments of its free variables are either bound variables or second-order ground base-type terms. For example, if G , \mathbf{a} , and x are of base type, and F , H , \mathbf{g} , and y are binary, the terms $F G \mathbf{a}$, and $H (\lambda x. x) \mathbf{a}$ are not solid; $\lambda x. F x x$ is strictly solid; $F \mathbf{a} (\mathbf{g} (\lambda y. y \mathbf{a} \mathbf{a})) \mathbf{a}$ is solid, but not strictly. Prehofer's thesis states that PT terminates on $\{s \stackrel{?}{=} t\}$ if s is linear, s shares no free variables with t , s is strictly solid, and t is second-order.

We extend this result in Theorem 5.7 along two axes: we create an oracle for the full unification problem (rather than the preunification problem), and we lift the second-order constraints by requiring s and t to be solid.

As a first step, we consider the matching problem. A *matching problem* is a unification problem $\{s \stackrel{?}{=} t\}$ where t is ground. In what follows, we establish some useful properties of matching problems in which both s and t are solid (*solid matching problems*).

Lemma 5.2. *PT terminates on a solid matching problem $\{s \stackrel{?}{=} t\}$.*

Proof. We show termination by designing a measure on a matching problem E that decreases with each application of a PT transition (possibly followed by **Decomposition** or **Failure** steps). Our measure function compares the following properties lexicographically:

A: multiset of the sizes of right-hand sides of the constraints in E

B: number of free variables in E

Clearly, when applying PT transitions, the problem stays a matching problem because the applied bindings do not introduce free variables on the right-hand sides. It is also easy to check that each applied binding keeps the terms in the solid fragment. Namely, bindings for both **Imitation** and **Projection** transitions are patterns, which means that, after applying the binding, fresh free variables are applied only to bound variables or ground base-type terms. The **Solution** transition effectively replaces the variable with a ground term, which is obviously solid. We show each transition either trivially terminates or reduces the measure:

Deletion: A decreases.

Decomposition: A decreases.

¹In Prehofer's thesis, strictly solid terms are called weakly second-order.

Failure: Trivial – represents a terminal node.

Solution: This transition applies on constraints of the form $F \stackrel{?}{=} t$. This reduces A, since the constraint $F \stackrel{?}{=} t$ is removed and F cannot appear on the right-hand side.

Imitation: The rule is applicable only on $\lambda\bar{x}. F \bar{s}_m \stackrel{?}{=} \lambda\bar{x}. f \bar{t}_n$. The Imitation transition replaces the constraint with $\{H_i \bar{s}_m \stackrel{?}{=} \bar{t}_i \mid 1 \leq i \leq n\}$. This reduces A, as F does not appear on any right-hand side.

Projection: The rule is applicable only on $\lambda\bar{x}. F \bar{s}_m \stackrel{?}{=} \lambda\bar{x}. a \bar{t}_n$. If a is a bound variable, and we project F to argument s_i different than a , Failure applies and PT trivially terminates. If we project F to a , we apply Decomposition, which reduces A. If a is a constant, for Failure not to apply, we have to project to a base-type ground term s_i . This does not increase A, since no variables appear on right-hand sides, but removes the variable F from E , reducing B by one. \square

We say σ is a grounding substitution if for every variable F mapped by σ , σF is ground.

Lemma 5.3. *All unifiers produced by PT for the solid matching problem $\{s \stackrel{?}{=} t\}$ are grounding substitutions.*

Proof. Closely following the proof of Lemma 5.2.5 in Prehofer’s PhD thesis [21], we prove our claim by induction on the length of the PT transition sequence that leads to the unifier. We know this sequence is finite by Lemma 5.2. The base case of induction, for the empty sequence, is trivial. The induction step is made using one of the following transitions:

Deletion: Trivial.

Decomposition: Trivial.

Failure: This rule is not relevant, since it will not lead to the unifier.

Solution: This rule applies the substitution $\{F \mapsto t\}$, which is grounding since t is ground.

Imitation: This transition applies on constraints of the form

$$\lambda\bar{x}. F \bar{s}_k \stackrel{?}{=} \lambda\bar{x}. f \bar{t}_l$$

The binding for Imitation is $\varrho = \{F \mapsto \lambda\bar{x}_k. f (G_1 \bar{s}_k) \dots (G_l \bar{s}_k)\}$, and reduces the problem to $\{G_i \bar{s}_k \stackrel{?}{=} t_i \mid 1 \leq i \leq l\}$. Since the right-hand sides are ground, any unifier σ produced by PT must map all of the variables G_1, \dots, G_l . By induction hypothesis, $\sigma(G_i)$ is ground. Therefore, $\sigma\varrho$ must map F to a ground term.

Projection: This transition applies on constraints of the form

$$\lambda\bar{x}. F \bar{s}_k \stackrel{?}{=} \lambda\bar{x}. t$$

The binding for Projection is $\varrho = \{F \mapsto \lambda\bar{x}_k. x_i (G_1 \bar{x}_k) \dots (G_j \bar{x}_k)\}$, and reduces the problem to $\{s_i (G_1 \bar{s}_k) \dots (G_j \bar{s}_k) \stackrel{?}{=} a \bar{t}_l\}$. Since the right-hand side is ground, any unifier σ produced by PT must map all of the variables G_1, \dots, G_j . By induction hypothesis, $\sigma(G_i)$ is ground. Therefore, $\sigma\varrho$ must map F to a ground term. \square

Lemma 5.4. *If s and t are solid, s is linear and shares no free variables with t , PT terminates for the preunification problem $\{s \stackrel{?}{=} t\}$, and all remaining flex-flex constraints are solid.*

Proof. This lemma is a modification of Lemma 5.2.1 and Theorem 5.2.6 from Prehofer’s PhD thesis [21]. Correspondingly, the following proof closely follows the proofs for these two lemmas. We also adopt the notion of an *isolated* variable from Prehofer’s thesis: a variable is isolated in a multiset E of constraints if it appears exactly once in E .

First, similarly to the proof of previous lemma we conclude each transition maintains the condition that the terms remain solid. Second, we have to show that variables on left-hand

sides remain isolated. For **Imitation** and **Projection** rules, the preservation of this invariant is obvious. Later, we also show that **Solution** preserves it. Third, since s and t share no variables, and s is linear, no rule can introduce a variable from the right-hand side to the left-hand side.

To prove termination of PT, we devise a measure that decreases with each application of a PT transition. The measure lexicographically compares the following properties:

- A:** number of occurrences of constant symbols and bound variables that are not below free variables on left-hand sides
- B:** number of free variables on right-hand sides
- C:** multiset of the sizes of right-hand sides

We show that each transition either trivially terminates or reduces the measure:

Deletion: A does not increase and at least one of A or B reduces.

Decomposition: A reduces.

Failure: Trivial.

Solution: This rule applies in two cases:

- $F \stackrel{?}{=} t$: since F is isolated, A is unchanged, B is not increased, and C reduces. All variables on left-hand sides remain isolated since they are unaffected by this substitution.
- $t \stackrel{?}{=} F$: since F does not occur on any left-hand side and the head of t must be a constant or bound variable (otherwise the rule would not be applicable), A reduces. Even though t might contain some free variables and F can have multiple occurrences on right-hand sides, all free variables on left-hand sides remain isolated. Namely, all free variables in t occur exactly once in the multiset, and since $t \stackrel{?}{=} F$ is removed, all of them either disappear or end up on right-hand sides.

Imitation: We distinguish the following two forms of the selected constraint:

- $\lambda\bar{x}. F \bar{v}_n \stackrel{?}{=} \lambda\bar{x}. f \bar{u}_m$: We replace this constraint by $\{H_i \bar{v}_n \stackrel{?}{=} u_i \mid 1 \leq i \leq m\}$ and apply the **Imitation** binding $F \mapsto \lambda\bar{x}_n. f (H_1 \bar{x}_n) \dots (H_m \bar{x}_n)$. As F is isolated, A and B do not increase. Since F is isolated, it does not occur on any right-hand side, and hence C decreases.
- $\lambda\bar{x}. f \bar{u}_m \stackrel{?}{=} \lambda\bar{x}. F \bar{v}_n$: applying the **Imitation** transition as above will reduce A since F cannot appear on any left-hand side.

Projection: Similarly to the previous transition rule, we have two cases:

- $\lambda\bar{x}. F \bar{v}_n \stackrel{?}{=} \lambda\bar{x}. a \bar{u}_m$: if a is a bound variable, then for **Failure** not to be applicable afterwards, we have to project F onto argument v_i equal to a . Then we proceed like for **Imitation**. If a is not a bound variable, then we have to project to some base-type term v_j (otherwise **Failure** would be applicable afterwards). This reduces the problem to $\lambda\bar{x}. v_j \stackrel{?}{=} \lambda\bar{x}. a \bar{u}_m$. This is a solid matching problem, whose solutions computed by PT are grounding substitutions (see Lemma 5.3). Applying one of those solutions will eliminate all the free variables in $\lambda\bar{x}. a \bar{u}_m$. Since PT is parametrized by an admissible selection function, we know that there are no constraints descending from a simple projection in E since the constraint $\lambda\bar{x}. F \bar{v}_n \stackrel{?}{=} \lambda\bar{x}. a \bar{u}_m$ was chosen, which, due to solidity restrictions, cannot be such a descendant. Therefore, we know that PT will transform the descendants of the matching problem $\lambda\bar{x}. v_j \stackrel{?}{=} \lambda\bar{x}. a \bar{u}_m$ until either **Failure** is observed (making PT trivially terminating) or until no descendant exists and the grounding matcher is computed (see Lemmas 5.3 and 5.2). This results in removal of the original constraint $\lambda\bar{x}. F \bar{v}_n \stackrel{?}{=} \lambda\bar{x}. a \bar{u}_m$ and application of the computed grounding matcher which will either remove all the free variables in the right-hand side of the

constraint (not increasing A and reducing B) or not increasing A and B, reduce C if no free variables occur in the right-hand side.

- $\lambda\bar{x}. a \bar{v}_n \stackrel{?}{=} \lambda\bar{x}. F \bar{u}_m$: if a is a bound variable, projecting F onto argument u_i will either enable application of **Decomposition** as the next step reducing A, or it will result in **Failure**, trivially terminating. If a is a constant, then projecting F onto some u_j will either yield **Failure** or enable **Decomposition**, reducing A. \square

Enumerating a CSU for a solid flex-flex pair may seem as hard as for any other flex-flex pair; however, the following two lemmas show that solid pairs admit an MGU:

Lemma 5.5. *The unification problem $\{\lambda\bar{z}. F \bar{s}_m \stackrel{?}{=} \lambda\bar{z}. F' \bar{s}'_{m'}\}$, where both terms are solid, has an MGU of the form $\sigma = \{F \mapsto \lambda\bar{x}_m. G x_{j_1} \dots x_{j_r}\}$ where G is a fresh variable, and $1 \leq j_1 < \dots < j_r \leq m$ are exactly those indices j_i for which $s_{j_i} = s'_{j_i}$.*

Proof. Let ρ be a unifier for the given unification problem. Let $\lambda\bar{x}. u = \rho(F)$. Take an arbitrary subterm of u whose head is a bound variable x_i . If x_i is of function type, it corresponds to either s_i or s'_i which, due to solidity restrictions, has to be a bound variable. Furthermore, since ρ is a unifier, s_i and s'_i have to be syntactically equal. Similarly, if x_i is of base type, it corresponds to two ground terms s_i and s'_i which have to be syntactically equal. We conclude that ρ can use variables from \bar{x}_m only if they correspond to syntactically equal terms. Therefore, there is a substitution θ such that $\rho \subseteq \theta\sigma$. Due to arbitrary choice of ρ , we conclude that σ is an MGU. \square

Lemma 5.6. *Let $\{\lambda\bar{x}. F \bar{s}_m \stackrel{?}{=} \lambda\bar{x}. F' \bar{s}'_{m'}\}$ be a solid unification problem where $F \neq F'$. By Lemma 5.2, there exists a finite CSU $\{\sigma_i^1, \dots, \sigma_i^{k_i}\}$ of the problem $\{s_i \stackrel{?}{=} H_i \bar{s}'_{m'}\}$, where H_i is a fresh free variable. Let $\lambda\bar{y}_{m'}. s_i^j = \lambda\bar{y}_{m'}. \sigma_i^j(H_i) \bar{y}_{m'}$. Similarly, also by Lemma 5.2, there exists a finite CSU $\{\tilde{\sigma}_i^1, \dots, \tilde{\sigma}_i^{l_i}\}$ of the problem $\{s'_i \stackrel{?}{=} \tilde{H}_i \bar{s}_m\}$, where \tilde{H}_i is a fresh free variable. Let $\lambda\bar{x}_m. s'^j_i = \lambda\bar{x}_m. \tilde{\sigma}_i^j(\tilde{H}_i) \bar{x}_m$. Let Z be a fresh free variable. An MGU σ for the given problem is*

$$\begin{aligned} F &\mapsto \lambda\bar{x}_m. Z \underbrace{x_1 \dots x_1}_{k_1 \text{ times}} \dots \underbrace{x_m \dots x_m}_{k_m \text{ times}} s_1^1 \dots s_1^{l_1} \dots s_{m'}^1 \dots s_{m'}^{l_{m'}} \\ F' &\mapsto \lambda\bar{y}_{m'}. Z s_1^1 \dots s_1^{k_1} \dots s_m^1 \dots s_m^{k_m} \underbrace{y_1 \dots y_1}_{l_1 \text{ times}} \dots \underbrace{y_{m'} \dots y_{m'}}_{l_{m'} \text{ times}} \end{aligned}$$

Proof. Let ρ be an arbitrary unifier for the problem $\lambda\bar{x}. F \bar{s}_m \stackrel{?}{=} \lambda\bar{x}. F' \bar{s}'_{m'}$. We prove σ is an MGU by showing that there exists a substitution θ such that $\rho \subseteq \theta\sigma$. Without loss of generality, we can focus only on the values of ρF and $\rho F'$; F and F' are the only free variables that appear in the original problem and for all other variables mapped by ρ , θ can be chosen to coincide with ρ .

Let $\lambda\bar{x}_m. u = \rho F$ and $\lambda\bar{y}_{m'}. u' = \rho F'$, where the bound variables \bar{x}_m and $\bar{y}_{m'}$ have been α -renamed apart. We also assume that the names of variables bound inside u and u' are α -renamed so that they are different from \bar{x}_m and $\bar{y}_{m'}$. Finally, bound variables from the definition of σ have been α -renamed to match \bar{x}_m and $\bar{y}_{m'}$. We define θ to be the substitution

$$\theta = \{Z \mapsto \lambda z_1^1 \dots z_1^{k_1} \dots z_m^1 \dots z_m^{k_m} w_1^1 \dots w_1^{l_1} \dots w_{m'}^1 \dots w_{m'}^{l_{m'}}. \mathbf{diff}(u, u')\}$$

where $\mathbf{diff}(v, v')$ is defined recursively as

$$\mathbf{diff}(\lambda x. v, \lambda y. v') = \lambda x. \mathbf{diff}(v, \{y \mapsto x\}v') \quad (5.1)$$

$$\mathbf{diff}(a \bar{v}_n, a \bar{v}'_n) = a \mathbf{diff}(v_1, v'_1) \dots \mathbf{diff}(v_n, v'_n) \quad (5.2)$$

$$\mathbf{diff}(x_i, v') = z_i^k, \text{ if } v' = s_i^k \quad (5.3)$$

$$\mathbf{diff}(v, y_i) = w_i^l, \text{ if } v = s_i^l \quad (5.4)$$

$$\mathbf{diff}(x_i \bar{v}_n, y_j \bar{v}'_n) = z_i^k \mathbf{diff}(v_1, v'_1) \dots \mathbf{diff}(v_n, v'_n), \text{ if } y_j = s_i^k \quad (5.5)$$

From \mathbf{diff} 's definition it is clear that there are terms v, v' for which it is undefined. However, we will show that for each u and u' which are bodies of bindings from a unifier ϱ , \mathbf{diff} is defined and has the desired property. In equations 5.3, 5.4, and 5.5, if there are multiple numbers k or l that fulfill the condition, choose an arbitrary one. We need to show that $\varrho \subseteq \theta\sigma$, i.e., $\varrho F = \theta\sigma F$ and $\varrho F' = \theta\sigma F'$. By the definitions of u, u', θ and σ and β -reduction, this is equivalent to

$$\begin{aligned} \lambda \bar{x}_m. u &= \lambda \bar{x}_m. \{z_i^k \mapsto x_i, w_i^l \mapsto s_i^l \text{ for all } k, l, i\} \mathbf{diff}(u, u') \\ \lambda \bar{y}'_m. u' &= \lambda \bar{y}'_m. \{z_i^k \mapsto s_i^k, w_i^l \mapsto y_i \text{ for all } k, l, i\} \mathbf{diff}(u, u') \end{aligned}$$

We will show by induction that for any $\lambda \bar{x}_m. v, \lambda \bar{y}'_m. v'$ such that

$$\{x_1 \mapsto s_1, \dots, x_m \mapsto s_m\}v = \{y_1 \mapsto s'_1, \dots, y_{m'} \mapsto s'_{m'}\}v' \quad (\star)$$

we have

$$\begin{aligned} v &= \{z_i^k \mapsto x_i, w_i^l \mapsto s_i^l \text{ for all } k, l, i\} \mathbf{diff}(v, v') \\ v' &= \{z_i^k \mapsto s_i^k, w_i^l \mapsto y_i \text{ for all } k, l, i\} \mathbf{diff}(v, v') \end{aligned} \quad (\dagger)$$

The equation (\star) holds for $v = u$ and $v' = u'$ because ϱ is a unifier of $\lambda \bar{x}. F \bar{s}_m \stackrel{?}{=} \lambda \bar{x}. F' \bar{s}'_{m'}$. Therefore, once we have shown that (\star) implies (\dagger) , we know that (\dagger) holds for $v = u$ and $v' = u'$ and we are done.

We prove that (\star) implies (\dagger) by induction on the size of v and v' . We consider the following cases:

$v = \lambda x. v_1$: For (\star) to hold v and v' must be of the same type. Therefore, the λ -prefixes of their η -long representatives must have the same length and we can apply equation 5.1.

By the induction hypothesis, (\dagger) holds.

$v = x_i$: In this case, $\{x_1 \mapsto s_1, \dots, x_m \mapsto s_m\}v = s_i$. Since (\star) holds, v' must be an instance of a unifier from the CSU of $s_i = H_i \bar{s}'_{m'}$. However, since s_i and all terms in $\bar{s}'_{m'}$ are ground, $\lambda \bar{y}'_m. v' = \sigma_i^k(H_i)$, for some k . Then, $\mathbf{diff}(x_i, v') = z_i^k$, and it is easy to check that (\dagger) holds.

$v = x_i \bar{v}_n, n > 0$: In this case, x_i is mapped to s_i which, due to solidity restrictions, has to be a functional bound variable. Since (\star) holds, we conclude that the head of $\{y_1 \mapsto s'_1, \dots, y_{m'} \mapsto s'_{m'}\}v'$ must be s'_j , such that $s'_j = s_i$; this also means that $v' = y_j \bar{v}'_n$. Therefore, it is easy to check that some $\tau = \{H_i \mapsto \lambda \bar{y}'_m. y_j\}$ is a matcher for the problem $s_i = H_i \bar{s}'_{m'}$. For some $k, \sigma_i^k = \tau$, i.e., $\mathbf{diff}(v, v') = z_i^k \mathbf{diff}(v_1, v'_1) \dots \mathbf{diff}(v_n, v'_n)$. By induction hypothesis, we get that (\dagger) holds.

$v = a \bar{v}_n$: In the remaining cases a is either a free variable, a loose bound variable different than x_1, \dots, x_m , or a constant. If a is a free variable or a loose bound variable different than x_1, \dots, x_m , then $v' = a \bar{v}'_n$, since (\star) holds, all of $\bar{s}'_{m'}$ are ground, and bound variables different than x_1, \dots, x_m and $y_1, \dots, y_{m'}$ are renamed to match by equation 5.1. By the

induction hypothesis and by equation 5.2, we obtain (\dagger) . If a is a constant, we consider two cases: either $v' = a \overline{v'_n}$ which allows us to apply the induction hypothesis and obtain (\dagger) as above, or $v' = y_j \overline{v'_m}$. Since a is a constant, y_j cannot be a functional bound variable, since then it would be mapped to s_j , which due to solidity restrictions also has to be a functional bound variable and (\star) would not hold. Therefore $v' = y_j$. In this case, we proceed as in the case $v = x_i$ with the roles of v and v' swapped. \square

Theorem 5.7. *Let s and t be solid terms that share no free variables, and let s be linear. Then the unification problem $\{s \stackrel{?}{=} t\}$ has a finite CSU.*

Proof. By Lemma 5.4, PT terminates on $\{s \stackrel{?}{=} t\}$ with a finite set of preunifiers σ , each associated with a multiset E of solid flex-flex pairs.

An MGU δ_E of the remaining multiset E of solid flex-flex constraints can be found as follows. Choose one constraint $(u \stackrel{?}{=} v) \in E$ and determine an MGU ρ for it using Lemma 5.5 or 5.6. Then the set $\rho(E \setminus \{u \stackrel{?}{=} v\})$ also contains only solid flex-flex constraints, and we iterate this process by choosing a constraint from $\rho(E \setminus \{u \stackrel{?}{=} v\})$ next until there are no constraints left, eventually yielding an MGU ρ' of $\rho(E \setminus \{u \stackrel{?}{=} v\})$. Finally, we obtain the MGU $\delta_E = \rho' \rho$ of E .

Let $U = \{\delta_E \sigma \mid \text{PT produces preunifier } \sigma \text{ with constraints } E\}$. By termination of PT, U is finite. We show that U is a CSU. Let ρ be an arbitrary unifier for $\{s \stackrel{?}{=} t\}$. By Lemma 5.1, PT produces a preunifier σ with flex-flex constraints E such that there is a unifier θ of E and $\rho \subseteq \theta \sigma$. Since δ_E is an MGU of E , there exists a substitution θ' such that $\theta \subseteq \theta' \delta_E$. If we make sure that the fresh variables introduced by σ and δ_E are distinct, it follows that $\rho \subseteq \theta' \delta_E \sigma$. Therefore, U is a CSU. \square

The proof of Theorem 5.7 provides an effective way to calculate a CSU using PT and the results of Lemmas 5.5 and 5.6.

Example 5.8. Let $\{F(\mathbf{f} \mathbf{a}) \stackrel{?}{=} \mathbf{g} \mathbf{a} (G \mathbf{a})\}$ be the unification problem to solve. Projecting F onto the first argument will lead to a nonunifiable problem, so we perform imitation of \mathbf{g} building a binding $\sigma_1 = \{F \mapsto \lambda x. \mathbf{g} (F_1 x) (F_2 x)\}$. This yields the problem $\{F_1(\mathbf{f} \mathbf{a}) \stackrel{?}{=} \mathbf{a}, F_2(\mathbf{f} \mathbf{a}) \stackrel{?}{=} G \mathbf{a}\}$. Again, we can only imitate \mathbf{a} for F_1 – building a new binding $\sigma_2 = \{F_1 \mapsto \lambda x. \mathbf{a}\}$. Finally, this yields the problem $\{F_2(\mathbf{f} \mathbf{a}) \stackrel{?}{=} G \mathbf{a}\}$. According to Lemma 5.6, we find CSUs for the problems $J_1 \mathbf{a} = \mathbf{f} \mathbf{a}$ and $I_1(\mathbf{f} \mathbf{a}) \stackrel{?}{=} \mathbf{a}$ using PT. The latter problem has a singleton CSU $\{I_1 \mapsto \lambda x. \mathbf{a}\}$, whereas the former has a CSU containing $\{J_1 \mapsto \lambda x. \mathbf{f} x\}$ and $\{J_1 \mapsto \lambda x. \mathbf{f} \mathbf{a}\}$. Combining these solutions, we obtain an MGU $\sigma_3 = \{F_2 \mapsto \lambda x. H x x \mathbf{a}, G \mapsto \lambda x. H(\mathbf{f} \mathbf{a})(\mathbf{f} x) x\}$ for $F_2(\mathbf{f} \mathbf{a}) \stackrel{?}{=} G \mathbf{a}$. Finally, we get the MGU $\sigma = \sigma_3 \sigma_2 \sigma_1 = \{F \mapsto \lambda x. \mathbf{g} \mathbf{a} (H x x \mathbf{a}), G \mapsto \lambda x. H(\mathbf{f} \mathbf{a})(\mathbf{f} x) x\}$ of the original problem.

Small examples that violate conditions of Theorem 5.7 and admit only infinite CSUs can be found easily. The problem $\{\lambda x. F(\mathbf{f} x) \stackrel{?}{=} \lambda x. \mathbf{f} (F x)\}$ violates variable distinctness and is a well-known example of a problem with only infinite CSUs. Similarly, $\lambda x. \mathbf{g} (F(\mathbf{f} x)) F \stackrel{?}{=} \lambda x. \mathbf{g} (\mathbf{f} (G x)) G$, which violates linearity, reduces to the previous problem. Only ground arguments to free variables are allowed because $\{F X \stackrel{?}{=} G \mathbf{a}\}$ has only infinite CSUs. Finally, it is crucial that functional arguments to free variables are only bound variables: the problem $\{\lambda y. X(\lambda x. x) y \stackrel{?}{=} \lambda y. y\}$ has only infinite CSUs.

6. AN EXTENSION OF FINGERPRINT INDEXING

A fundamental building block for almost all automated reasoning tools is the operation of retrieving term pairs that satisfy certain conditions, e.g., unifiable terms, instances or generalizations. Indexing data structures are used to implement this operation efficiently. If the data structure retrieves precisely the terms that satisfy the condition, it is called *perfect*.

Higher-order indexing has received little attention compared to its first-order counterpart. However, recent research in higher-order theorem proving increased the interest in higher-order indexing [3, 17]. A *fingerprint index* [23, 31] is an imperfect index based on the idea that the skeleton of the term consisting of all non-variable positions is not affected by substitutions. Therefore, we can easily determine that the terms are not unifiable (or matchable) if they disagree on a fixed set of sample positions.

More formally, when we sample an untyped first-order term t on a sample position p , the *generic fingerprinting function* gfpf distinguishes four possibilities:

$$\text{gfpf}(t, p) = \begin{cases} \mathbf{f} & \text{if } t|_p \text{ has a symbol head } \mathbf{f} \\ \mathbf{A} & \text{if } t|_p \text{ is a variable} \\ \mathbf{B} & \text{if } t|_q \text{ is a variable for some proper prefix } q \text{ of } p \\ \mathbf{N} & \text{otherwise} \end{cases}$$

We define the *fingerprinting function* $\text{fp}(t) = (\text{gfpf}(t, p_1), \dots, \text{gfpf}(t, p_n))$, based on a fixed tuple of positions \bar{p}_n . Determining whether two terms are compatible for a given retrieval operation reduces to checking their fingerprints' componentwise compatibility. The following matrices determine the compatibility for retrieval operations:

	f_1	f_2	A	B	N
f_1		\mathbf{X}			\mathbf{X}
A					\mathbf{X}
B					
N	\mathbf{X}	\mathbf{X}	\mathbf{X}		

	f_1	f_2	A	B	N
f_1		\mathbf{X}	\mathbf{X}	\mathbf{X}	\mathbf{X}
A				\mathbf{X}	\mathbf{X}
B					
N	\mathbf{X}	\mathbf{X}	\mathbf{X}	\mathbf{X}	

The left matrix determines unification compatibility, while the right matrix determines compatibility for matching term s (rows) onto term t (columns). Symbols f_1 and f_2 stand for arbitrary distinct constants. Incompatible features are marked with \mathbf{X} . For example, given a tuple of term positions $(1, 1.1.1, 2)$, and terms $\mathbf{f}(\mathbf{g}(X), \mathbf{b})$ and $\mathbf{f}(\mathbf{f}(\mathbf{a}, \mathbf{a}), \mathbf{b})$, their fingerprints are $(\mathbf{g}, \mathbf{B}, \mathbf{b})$ and $(\mathbf{f}, \mathbf{N}, \mathbf{b})$, respectively. Since the first fingerprint component is incompatible, terms are not unifiable.

Fingerprints for the terms in the index are stored in a trie data structure. This allows us to efficiently filter out terms that are not compatible with a given retrieval condition. For the remaining terms, a unification or matching procedure must be invoked to determine whether they satisfy the condition or not.

The fundamental idea of first-order fingerprint indexing carries over to higher-order terms – application of a substitution does not change the rigid skeleton of a term. However, to extend fingerprint indexing to higher-order terms, we must address the issues of $\alpha\beta\eta$ -normalization and figure how to cope with λ -abstractions and bound variables. To that end, we define a function $\lfloor t \rfloor$, defined on β -reduced η -long terms in De Bruijn [5] notation:

$$\lfloor F \bar{s} \rfloor = F \quad \lfloor x_i \bar{s}_n \rfloor = \mathbf{db}_i^\alpha(\lfloor s_1 \rfloor, \dots, \lfloor s_n \rfloor) \quad \lfloor \mathbf{f} \bar{s}_n \rfloor = \mathbf{f}(\lfloor s_1 \rfloor, \dots, \lfloor s_n \rfloor) \quad \lfloor \lambda \bar{x}. s \rfloor = \lfloor s \rfloor$$

We let x_i be a bound variable of type α with De Bruijn index i , and \mathbf{db}_i^α be a fresh constant corresponding to this variable. All constants \mathbf{db}_i^α must be fresh. Effectively, $\lfloor _ \rfloor$ transforms

a η -long β -reduced higher-order term to an untyped first-order term. Let $t_{\downarrow\beta\eta}$ be the η -long β -reduced form of t ; the higher-order generic fingerprinting function gfpf_{ho} , which relies on conversion $\langle t \rangle_{\text{db}}$ from named to De Bruijn representation, is defined as

$$\text{gfpf}_{\text{ho}}(t, p) = \text{gfpf}(\lfloor \langle t_{\downarrow\beta\eta} \rangle_{\text{db}} \rfloor, p)$$

If we define $\text{fp}_{\text{ho}}(t) = \text{fp}(\lfloor \langle t_{\downarrow\beta\eta} \rangle_{\text{db}} \rfloor)$, we can support fingerprint indexing for higher-order terms with no changes to the compatibility matrices. For example, consider the terms $s = (\lambda xy. x y) \mathbf{g}$ and $t = \mathbf{f}$, where \mathbf{g} has the type $\alpha \rightarrow \beta$ and \mathbf{f} has the type $\alpha \rightarrow \alpha \rightarrow \beta$. For the tuple of positions $(1, 1.1.1, 2)$ we get

$$\begin{aligned} \text{fp}_{\text{ho}}(s) &= \text{fp}(\lfloor \langle s_{\downarrow\beta\eta} \rangle_{\text{db}} \rfloor) = \text{fp}(\mathbf{g}(\text{db}_0^\alpha)) = (\text{db}_0^\alpha, \mathbf{N}, \mathbf{N}) \\ \text{fp}_{\text{ho}}(t) &= \text{fp}(\lfloor \langle t_{\downarrow\beta\eta} \rangle_{\text{db}} \rfloor) = \text{fp}(\mathbf{f}(\text{db}_1^\alpha, \text{db}_0^\alpha)) = (\text{db}_1^\alpha, \mathbf{N}, \text{db}_0^\alpha) \end{aligned}$$

Since the first and third fingerprint component are incompatible, the terms are not unifiable.

Other first-order indexing techniques such as feature vector indexing and substitution trees can probably be extended to higher-order terms using the method described here as well.

7. IMPLEMENTATION

Zipperposition [6, 7] is an open-source² theorem prover written in OCaml. It is a versatile testbed for prototyping extensions to superposition-based theorem provers. It was initially designed as a prover for polymorphic first-order logic and then extended to higher-order logic. A recent addition is a complete mode for Boolean-free higher-order logic [2], which depends on a unification procedure that can enumerate a CSU. We implemented our procedure in Zipperposition.

We used OCaml’s functors to create a modular implementation. The core of our procedure is implemented in a module which is parametrized by another module providing oracles and implementing the `Bind` step. In this way we can obtain the complete or pragmatic procedure and seamlessly integrate oracles while reusing as much common code as possible.

To enumerate all elements of a possibly infinite CSU, we rely on lazy lists whose elements are subsingletons of unifiers (either one-element sets containing a unifier or empty sets). The search space must be explored in a *fair* manner, meaning that no branch of the constructed tree is indefinitely postponed.

Each `Bind` step will give rise to new unification problems E_1, E_2, \dots to be solved. Solutions to each of those problems are lazy lists p_1, p_2, \dots containing subsingletons of unifiers. To avoid postponing some unifier indefinitely, we use the dovetailing technique: we first take one subsingleton from p_1 , then one from each of p_1 and p_2 . We continue with one subsingleton from each of p_1, p_2 and p_3 , and so on. Empty lazy lists are ignored in the traversal. To ensure we do not remain stuck waiting for a unifier from a particular lazy list, the procedure will periodically return an empty set, indicating that the next lazy list should be probed.

The implemented selection function for our procedure prioritizes selection of rigid-rigid over flex-rigid pairs, and flex-rigid over flex-flex pairs. However, since the constructed substitution σ is not applied eagerly, heads can appear to be flex, even if they become rigid after dereferencing and normalization. To mitigate this issue to some degree, we dereference the heads with σ , but do not normalize, and use the resulting heads for prioritization.

²<https://github.com/sneeuwballen/zipperposition>

We implemented oracles for the pattern, solid, and fixpoint fragment. Fixpoint unification [12] is concerned with problems of the form $\{F \stackrel{?}{=} t\}$. If F does not occur in t , $\{F \mapsto t\}$ is an MGU for the problem. If there is a position p in t such that $t|_p = F\bar{u}_m$ and for each prefix $q \neq p$ of p , $t|_q$ has a rigid head and either $m = 0$ or t is not a λ -abstraction, then we can conclude that $F \stackrel{?}{=} t$ has no solutions. Otherwise, the fixpoint oracle is not applicable.

8. EVALUATION

We evaluated the implementation of our unification procedure in Zipperposition, assessing the complete variant and the pragmatic variant, the latter with several different combinations of limits for number of bindings. As part of the implementation of the complete mode for Boolean-free higher-order logic in Zipperposition [2], Bentkamp implemented a straightforward version of JP procedure. This version is faithful to the original description, with a check as to whether a (sub)problem can be solved using a first-order oracle as the only optimization. Our evaluations were performed on StarExec Miami [26] servers with Intel Xeon E5-2620 v4 CPUs clocked at 2.10 GHz with 60s CPU limit.

Contrary to first-order unification, there is no widely available corpus of benchmarks dedicated solely to evaluating performance of higher-order unification algorithms. Thus, we used all 2606 monomorphic higher-order theorems from the TPTP library [28] and 832 monomorphic higher-order Sledgehammer (SH) generated problems [27] as our benchmarks³. Many TPTP problems require synthesis of complicated unifiers, whereas Sledgehammer problems are only mildly higher-order – many of them are solved with first-order unifiers.

We used the naive implementation of the JP procedure (**jp**) as a baseline to evaluate the performance of our procedure. We compare it with the complete variant of our procedure (**cv**) and pragmatic variants (**pv**) with several different configurations of limits for applied bindings. All other Zipperposition parameters have been fixed to the values of a variant of a well-performing configuration we used for the CASC-27 theorem proving competition [29]. The cv configuration and all of the pv configurations use only pattern unification as an underlying oracle. To test the effect of oracle choice, we evaluated the complete variant in 8 combinations: with no oracles (**n**), with only fixpoint (**f**), pattern (**p**), or solid (**s**) oracle, and with their combinations: **fp**, **fs**, **ps**, **fps**.

Figure 1 compares different variants of the procedure with the naive JP implementation. Each pv configuration is denoted by pv_{bcde}^a where a is the limit on the total number of applied bindings, and b , c , d , and e are the limits of functional projections, eliminations, imitations, and identifications, respectively. Figure 2 summarizes the effects of using different oracles.

The configuration of our procedure with no oracles outperforms the JP procedure with the first-order oracle. This suggests that the design of the procedure, in particular lazy normalization and lazy application of the substitution, already reduces the effects of the JP procedure’s main bottlenecks. Raw evaluation data shows that on TPTP benchmarks, complete and pragmatic configurations differ in the set of problems they solve – cv solves 19 problems not solved by pv_{2222}^4 , whereas pv_{2222}^4 solves 34 problems cv does not solve. Similarly, comparing the pragmatic configurations with each other, pv_{3333}^6 and pv_{2222}^4 each solve 13 problems that the other one does not. The overall higher success rate of pv_{1020}^2 compared to pv_{1222}^2 suggests that solving flex-flex pairs by trivial unifiers often suffices for superposition-based theorem proving.

³An archive with raw results, all used problems, and scripts for running each configuration is available at http://matryoshka.gforge.inria.fr/pubs/hounif_data.zip

	jp	cv	PV_{6666}^{12}	PV_{3333}^6	PV_{2222}^4	PV_{1222}^2	PV_{1121}^2	PV_{1020}^2
TPTP	1551	1717	1722	1732	1732	1715	1712	1719
SH	242	260	253	255	255	254	259	257

Figure 1: Proved problems, per configuration

	n	f	p	s	fp	fs	ps	fps
TPTP	1658	1717	1717	1720	1719	1724	1720	1723
SH	245	255	260	259	255	254	258	254

Figure 2: Proved problems, per used oracle

Counterintuitively, in some cases, using oracles can hurt the performance of Zipperposition. Using oracles typically results in generating smaller CSUs, whose elements are more general substitutions than the ones we obtain without oracles. These more general substitutions usually contain more applied variables, which Zipperposition’s heuristics avoid due to their explosive nature. This can make Zipperposition postpone necessary inferences for too long. Configuration n benefits from this effect and therefore solves 18 TPTP problems that no other configuration in Figure 2 solves. The same effect also gives configurations with only one oracle an advantage over configurations with multiple oracles on some problems.

The evaluation sheds some light on how often solid unification problems appear in practice. The raw data show that configuration s solves 5 TPTP problems that neither f nor p solve. Configuration f solves 8 TPTP problems that neither s nor p solve, while p solves 9 TPTP problems that two other configurations do not. This suggests that the solid oracle is slightly less beneficial than the fixpoint or pattern oracles, but still presents a useful addition the set of available oracles.

A subset of TPTP benchmarks, concerning operations on Church numerals, is designed to test the efficiency of higher-order unification. Our procedure performs exceptionally well on these problems – it solves all of them, usually faster than other competitive higher-order provers. There are 11 benchmarks in NUM category of TPTP that contain conjectures about Church numerals: $NUM020^1$, $NUM021^1$, $NUM415^1$, $NUM416^1$, $NUM417^1$, $NUM418^1$, $NUM419^1$, $NUM798^1$, $NUM799^1$, $NUM800^1$, and $NUM801^1$. We evaluated those problems using the same CPU nodes and the same time limits as above. In addition to Zipperposition, we used all higher-order provers that took part in the 2019 edition of CASC [29] (in the THF category) for this evaluation: CVC4 1.7 [1], Leo-III 1.4 [25], Satallax 3.4 [4], Vampire 4.4 THF [14]. Figure 3 shows the CPU time needed to solve a problem or “–” if the prover timed out.

9. DISCUSSION AND RELATED WORK

The problem addressed in this paper is that of finding a complete and efficient higher-order unification procedure. Three main lines of research dominated the research field of higher-order unification over the last forty years.

The first line of research went in the direction of finding procedures that enumerate CSUs. The most prominent procedure designed for this purpose is the JP procedure [13]. Snyder and Gallier [24] also provide such a procedure, but instead of solving flex-flex pairs systematically, their procedure blindly guesses the head of the necessary binding by considering all constants

	CVC4	Leo-III	Satallax	Vampire	cv
NUM020 [^] 1	–	0.46	–	–	0.03
NUM021 [^] 1	–	–	–	–	4.10
NUM415 [^] 1	45.80	0.34	0.21	0.42	0.03
NUM416 [^] 1	47.37	0.92	0.21	0.41	0.07
NUM417 [^] 1	–	49.73	0.30	0.40	0.45
NUM418 [^] 1	–	0.40	1.29	0.38	0.03
NUM419 [^] 1	–	0.42	23.33	0.37	0.03
NUM798 [^] 1	46.29	0.35	4.01	0.38	0.03
NUM799 [^] 1	–	5.05	–	–	0.03
NUM800 [^] 1	–	–	–	0.37	3.15
NUM801 [^] 1	–	0.73	38.77	–	0.50

Figure 3: Time needed to prove a problem, in seconds.

in the signature and fresh variables of all possible types. Another approach, based on higher-order combinators, is given by Dougherty [9]. This approach blindly creates (partially applied) S-, K-, and I-combinator bindings for applied variables, which results in returning many redundant unifiers, as well as in nonterminating behavior even for simple problems such as $X a = a$.

The second line of research is concerned with enumerating preunifiers. The most prominent procedure in this line of research is Huet’s [12]. The Snyder–Gallier procedure restricted not to solve flex-flex pairs is a version of the PT procedure presented in Section 5. It improves Huet’s procedure by featuring a Solution rule.

The third line of research gives up the expressiveness of the full λ -calculus and focuses on decidable fragments. Patterns [19] are arguably the most important such fragment in practice, with implementations in Isabelle [20], Leo-III [25], Satallax [4], λ Prolog [18], and other systems. Functions-as-constructors [16] unification subsumes pattern unification but is significantly more complex to implement. Prehofer [21] lists many other decidable fragments, not only for unification but also preunification and unifier existence problems. Most of these algorithms are given for second-order terms with various constraints on their variables. Finally, one of the first decidability results is Farmer’s discovery [10] that higher-order unification of terms with unary function symbols is decidable.

Our procedure draws inspiration from and contributes to all three lines of research. Accordingly, its advantages over previously known procedures can be laid out along those three lines. First, our procedure mitigates many issues of the JP procedure. Second, it can be modified not to solve flex-flex pairs, and become a version of Huet’s procedure with important built-in optimizations. Third, it can integrate any oracle for problems with finite CSUs, including the one we discovered.

The implementation of our procedure in Zipperposition was one of the reasons this prover evolved from proof-of-concept prover for higher-order logic to competitive higher-order prover. In the 2019 edition of CASC, Zipperposition featured a less efficient version of our procedure. Nevertheless, it finished neck-and-neck with state-of-the-art prover Leo-III and has beaten all the newcomers in higher-order competition category.

10. CONCLUSION

We presented a procedure for enumerating a complete set of higher-order unifiers that is designed for efficiency. Due to a design that restricts the search space and a tight integration of oracles, it reduces the number of redundant unifiers returned and gives up early in cases of nonunifiability. In addition, we presented a new fragment of higher-order terms that admits finite CSUs. Our evaluation shows a clear improvement over previously known procedures.

In future work, we will focus on designing intelligent heuristics that automatically adjust unification parameters according to the type of the problem. For example, we should usually choose shallow unification for mostly first-order problems and deeper unification for hard higher-order problems. We plan to investigate other heuristic choices, such as the order of bindings and the way in which search space is traversed (breadth- or depth-first). We are also interested in further improving the termination behavior of the procedure, without sacrificing completeness. Finally, following the work of Libal [15] and Zaionc [32], we would like to consider the use of regular grammars to finitely present infinite CSUs. For example, the grammar $G ::= \lambda x. x \mid \lambda x. f(Gx)$ represents all elements of the CSU for the problem $\lambda x. G(fx) \stackrel{?}{=} \lambda x. f(Gx)$.

Acknowledgment. We are grateful to the maintainers of StarExec for letting us use their service. We thank Ahmed Bhayat, Jasmin Blanchette, Daniel El Ouraoui, Mathias Fleury, Pascal Fontaine, Predrag Janičić, Robert Lewis, Femke van Raamsdonk, Hans-Jörg Schurr, Sophie Turrett, Dmitriy Traytel, and the anonymous reviewers for suggesting many improvements to this text. Vukmirović and Bentkamp’s research has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement No. 713999, Matryoshka). Nummelin has received funding from the Netherlands Organization for Scientific Research (NWO) under the Vidi program (project No. 016.Vidi.189.037, Lean Forward).

REFERENCES

- [1] Clark W. Barrett, Christopher L. Conway, Morgan Deters, Liana Hadarean, Dejan Jovanovic, Tim King, Andrew Reynolds, and Cesare Tinelli. CVC4. In Ganesh Gopalakrishnan and Shaz Qadeer, editors, *CAV 2011*, volume 6806 of *LNCS*, pages 171–177. Springer, 2011.
- [2] Alexander Bentkamp, Jasmin Blanchette, Sophie Turrett, Petar Vukmirović, and Uwe Waldmann. Superposition with lambdas. In Pascal Fontaine, editor, *CADE-27*, volume 11716 of *LNCS*, pages 55–73. Springer, 2019.
- [3] Ahmed Bhayat and Giles Reger. Restricted combinatory unification. In Pascal Fontaine, editor, *CADE-27*, volume 11716 of *LNCS*, pages 74–93. Springer, 2019.
- [4] Chad E. Brown. Satallax: An automatic higher-order prover. In Bernhard Gramlich, Dale Miller, and Uli Sattler, editors, *IJCAR 2012*, volume 7364 of *LNCS*, pages 111–117. Springer, 2012.
- [5] Nicolaas G. De Bruijn. Lambda calculus notation with nameless dummies, a tool for automatic formula manipulation, with application to the Church-Rosser theorem. *J. Symb. Log.*, 40(3):470–470, 1975.
- [6] Simon Cruanes. *Extending Superposition with Integer Arithmetic, Structural Induction, and Beyond*. PhD thesis, École polytechnique, 2015.
- [7] Simon Cruanes. Superposition with structural induction. In Clare Dixon and Marcelo Finger, editors, *FroCoS 2017*, volume 10483 of *LNCS*, pages 172–188. Springer, 2017.
- [8] Nachum Dershowitz and Zohar Manna. Proving termination with multiset orderings. *Commun. ACM*, 22(8):465–476, 1979.
- [9] Daniel J. Dougherty. Higher-order unification via combinators. *Theor. Comput. Sci.*, 114(2):273–298, 1993.
- [10] William M. Farmer. A unification algorithm for second-order monadic terms. *Ann. Pure Appl. Logic*, 39(2):131–174, 1988.

- [11] Krystof Hoder and Andrei Voronkov. Comparing unification algorithms in first-order theorem proving. In Bärbel Mertsching, Marcus Hund, and Muhammad Zaheer Aziz, editors, *KI 2009*, volume 5803 of *LNCS*, pages 435–443. Springer, 2009.
- [12] Gérard P. Huet. A unification algorithm for typed lambda-calculus. *Theor. Comput. Sci.*, 1(1):27–57, 1975.
- [13] Don C. Jensen and Tomasz Pietrzykowski. Mechanizing omega-order type theory through unification. *Theor. Comput. Sci.*, 3(2):123–171, 1976.
- [14] Laura Kovács and Andrei Voronkov. First-order theorem proving and vampire. In Natasha Sharygina and Helmut Veith, editors, *CAV 2013*, volume 8044 of *LNCS*, pages 1–35. Springer, 2013.
- [15] Tomer Libal. Regular patterns in second-order unification. In Amy P. Felty and Aart Middeldorp, editors, *CADE-25*, volume 9195 of *LNCS*, pages 557–571. Springer, 2015.
- [16] Tomer Libal and Dale Miller. Functions-as-constructors higher-order unification. In Delia Kesner and Brigitte Pientka, editors, *FSCD 2016*, volume 52 of *LIPICs*, pages 26:1–26:17. Schloss Dagstuhl, 2016.
- [17] Tomer Libal and Alexander Steen. Towards a substitution tree based index for higher-order resolution theorem provers. In Pascal Fontaine, Stephan Schulz, and Josef Urban, editors, *IJCAR 2016*, volume 1635 of *CEUR-WS*, pages 82–94. CEUR-WS, 2016.
- [18] Dale Miller and Gopalan Nadathur. *Programming with Higher-Order Logic*. Cambridge University Press, 2012.
- [19] Tobias Nipkow. Functional unification of higher-order patterns. In E. Best, editor, *LICS '93*, pages 64–74. IEEE Computer Society, 1993.
- [20] Tobias Nipkow, Lawrence C. Paulson, and Markus Wenzel. *Isabelle/HOL - A Proof Assistant for Higher-Order Logic*, volume 2283 of *LNCS*. Springer, 2002.
- [21] Christian Prehofer. *Solving higher order equations: from logic to programming*. PhD thesis, Technical University Munich, Germany, 1995.
- [22] John Alan Robinson. A machine-oriented logic based on the resolution principle. *J. ACM*, 12(1):23–41, 1965.
- [23] Stephan Schulz. Fingerprint indexing for paramodulation and rewriting. In Bernhard Gramlich, Dale Miller, and Uli Sattler, editors, *IJCAR 2012*, volume 7364 of *LNCS*, pages 477–483. Springer, 2012.
- [24] Wayne Snyder and Jean H. Gallier. Higher-order unification revisited: Complete sets of transformations. *J. Symb. Comput.*, 8(1/2):101–140, 1989.
- [25] Alexander Steen and Christoph Benzmüller. The higher-order prover Leo-III. In Didier Galmiche, Stephan Schulz, and Roberto Sebastiani, editors, *IJCAR 2018*, volume 10900 of *LNCS*, pages 108–116. Springer, 2018.
- [26] Aaron Stump, Geoff Sutcliffe, and Cesare Tinelli. Starexec: A cross-community infrastructure for logic solving. In Stéphane Demri, Deepak Kapur, and Christoph Weidenbach, editors, *IJCAR 2014*, volume 8562 of *LNCS*, pages 367–373. Springer, 2014.
- [27] Nik Sultana, Jasmin Christian Blanchette, and Lawrence C. Paulson. LEO-II and Satallax on the Sledgehammer test bench. *J. Applied Logic*, 11(1):91–102, 2013.
- [28] Geoff Sutcliffe. The TPTP problem library and associated infrastructure - from CNF to TH0, TPTP v6.4.0. *J. Autom. Reasoning*, 59(4):483–502, 2017.
- [29] Geoff Sutcliffe. The CADE-27 automated theorem proving system competition - CASC-27. *AI Commun.*, 32(5-6):373–389, 2019.
- [30] Petar Vukmirovic, Alexander Bentkamp, and Visa Nummelin. Efficient full higher-order unification. In *FSCD*, volume 167 of *LIPICs*, pages 5:1–5:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.
- [31] Petar Vukmirovic, Jasmin Christian Blanchette, Simon Cruanes, and Stephan Schulz. Extending a brainiac prover to lambda-free higher-order logic. In Tomás Vojnar and Lijun Zhang, editors, *TACAS 2019*, volume 11427 of *LNCS*, pages 192–210. Springer, 2019.
- [32] Marek Zaionc. The set of unifiers in typed lambda-calculus as regular expression. In Jean-Pierre Jouannaud, editor, *RTA-85*, volume 202 of *LNCS*, pages 430–440. Springer, 1985.