

Higher-Order Dynamic Pattern Unification for Dependent Types and Records

Andreas Abel¹ and Brigitte Pientka²

¹ Department of Computer Science and Engineering
Gothenburg University
andreas.abel@gu.se

² School of Computer Science, McGill University, Montreal, Canada
bpientka@cs.mcgill.ca

Abstract. While higher-order pattern unification for the λ^Π -calculus is decidable and unique unifiers exist, we face several challenges in practice: 1) the pattern fragment itself is too restrictive for many applications; this is typically addressed by solving sub-problems which satisfy the pattern restriction eagerly but delay solving sub-problems which are non-patterns until we have accumulated more information. This leads to a dynamic pattern unification algorithm. 2) Many systems implement $\lambda^{\Pi\Sigma}$ calculus and hence the known pattern unification algorithms for λ^Π are too restrictive.

In this paper, we present a constraint-based unification algorithm for $\lambda^{\Pi\Sigma}$ -calculus which solves a richer class of patterns than currently possible; in particular it takes into account type isomorphisms to translate unification problems containing Σ -types into problems only involving Π -types. We prove correctness of our algorithm and discuss its application.

1 Introduction

Higher-order unification is a key operation in logical frameworks, dependently-typed programming systems, or proof assistants supporting higher-order logic. It plays a central role in type inference and reconstruction algorithms, in the execution of programs in higher-order logic programming languages, and in reasoning about the totality of functions defined by pattern-matching clauses.

While full higher-order unification is undecidable [8], Miller [10] identified a decidable fragment of higher-order unification problems, called the *pattern* fragment. A pattern is a unification problem where all meta-variables (or logic variables) occurring in a term are applied to some distinct bound variables. For example, the problem $\lambda x y z. X x y = \lambda x y z. x (\text{succ } y)$ falls into the pattern fragment, because the meta-variable X is applied to distinct bound variables x and y ; the pattern condition allows us to solve the problem by a simple abstraction $X = \lambda x y. x (\text{succ } y)$. This is not possible for non-patterns; examples for non-pattern problems which have no unique most general unifier can be obtained by changing the left hand side of the previous problem to $\lambda x y z. X x x y$ (non-linearity),

$\lambda x y z. X (Y x) y$ (X applied to another meta-variable) or $\lambda x y z. X x (\text{succ } y)$ (X applied to non-variable term).

In practice we face several challenges: First, the pattern fragment is too restrictive for many applications. Systems such as Twelf [13], Beluga [15], and Delphin [16] solve eagerly sub-problems which fall into the pattern fragment and delay sub-problems outside the pattern fragment until more information has been gathered which in turn simplifies the delayed sub-problems. The meta-theory justifying the correctness of such a strategy is largely unexplored and complex (an exception is the work by Reed [17]).

Second, we often want to consider richer calculi beyond the λ^{Π} -calculus. In Beluga and Twelf for example we use Σ -types to group assumptions together. In Agda [12] we support Σ -types in form of records with associated η -equality in its general form. Yet, little work has been done on extending the pattern fragment to handle also Σ -types. The following terms may be seen as equivalent: (a) $\lambda y_1. \lambda y_2. X (y_1, y_2)$, (b) $\lambda y. X (\text{fst } y) (\text{snd } y)$ and (c) $\lambda y_1. \lambda y_2. X y_1 y_2$. Only the last term falls within the pattern fragment as originally described by Miller. However, the other two terms can be transformed such that they also fall into the pattern fragment: for term (a), we replace X with $\lambda y. X' (\text{fst } y) (\text{snd } y)$; for term (b), we unfold y which stands for a pair and replace y with (y_1, y_2) .

In this paper, we describe a higher-order unification algorithm for the $\lambda^{\Pi\Sigma}$ calculus; our algorithm handles lazily η -expansion and we translate terms into the pure pattern fragment where a meta-variable is applied to distinct bound variables. The key insight is to take into account type isomorphisms for Σ , the dependently typed pairs: $\Pi z. (\Sigma x. A. B). C$ is isomorphic to $\Pi x. A. \Pi y. B. [(x, y)/z]C$, and a function $f: \Pi x. A. \Sigma y. B. C$ can be translated into two functions $f_1 : \Pi x. A. B$ and $f_2 : \Pi x. A. [f_1 x/y]C$. These transformations allow us to handle a richer class of dependently-typed patterns than previously considered.

Following Nanevski et al. [11] and Pientka [14], our description takes advantage of modelling meta-variables as closures; instead of directly considering a meta-variable X at function type $\Pi \mathbf{x}. \mathbf{A}. B$ which is applied to \mathbf{x} , we describe them as contextual objects, i.e., objects of type B in a context $\mathbf{x}:\mathbf{A}$, which are associated with a delayed substitution for the local context $\mathbf{x}:\mathbf{A}$.³ This allows us to give a high-level description and analysis following Dowek et al. [3], but not resorting to explicit substitutions; more importantly, it provides a logical grounding for some of the techniques such as “pre-cooking” and handles a richer calculus including Σ -types. Our work also avoids some of the other shortcomings; as pointed out by Reed [17], the algorithm sketched in Dowek et al. [3] fails to terminate on some inputs. We give a clear specification of the pruning which eliminates bound variable dependencies for the dependently typed case and show correctness of the unification algorithms in three steps: 1) we show that it terminates, 2) we show that the transformations in our unification algorithm preserve types, and 3) that each transition neither destroys nor creates (additional) solutions.

³ We write $\mathbf{x}:\mathbf{A}$ for a vector $x_1:A_1, \dots, x_n:A_n$.

Our work is to our knowledge the first comprehensive description of constraint-based higher-order pattern unification for the $\lambda^{\Pi\Sigma}$ calculus. It builds on and extends prior work by Reed [17] to handle Σ -types. Previously, Elliot [5] described unification for Σ -types in a Huet-style unification algorithm. While it is typically straightforward to incorporate η -expansions and lowering for meta-variables of Σ -type [19,12], there is little work on extending the notion of Miller patterns to be able to handle meta-variables which are applied to projections of bound variables. Fettig and Löchner [6] describe a higher-order pattern unification algorithm with finite products in the simply typed lambda-calculus. Their approach does not directly exploit isomorphisms on types, but some of the ideas have a similar goal: for example abstractions $\lambda x. \text{fst } x$ is translated into $\lambda(x_1, x_2). \text{fst } (x_1, x_2)$ which in turn normalizes to $\lambda(x_1, x_2).x_1$ to eliminate projections. Duggan [4] also explores extended higher-order patterns for products in the simply-typed setting; he generalizes Miller’s pattern restriction for the simply-typed lambda-calculus by allowing repeated occurrences of variables to appear as arguments to meta-variables, provided such variables are prefixed by distinct sequences of projections.

Our work has been already tested in practice. Some of the ideas described in this paper are incorporated into the implementation of the dependently-typed language Agda and Beluge. Agda has dependently typed records, a generalization of Σ -types. In Beluga, Σ -types occur in a restricted form, i.e., only variable declarations in contexts can be of Σ -type and there is no nesting of Σ -types.

This is the extended version of a published conference paper [1].

2 $\lambda^{\Pi\Sigma}$ -calculus

In this paper, we are considering an extension of the $\lambda^{\Pi\Sigma}$ -calculus with meta-variables. Its grammar is mostly straightforward. We use x, y, z for bound variables to distinguish them from meta-variables u, v , and w .

Sorts	$s ::= \text{type} \mid \text{kind}$
Atomic types	$P, Q ::= \mathbf{a} \ M$
Types	$A, B, C, D ::= P \mid \Pi x:A.B \mid \Sigma x:A.B$
Kinds	$\kappa ::= \text{type} \mid \Pi x:A.\kappa$
(Rigid) heads	$H ::= \mathbf{a} \mid \mathbf{c} \mid x$
Projections	$\pi ::= \text{fst} \mid \text{snd}$
Evaluation contexts	$E ::= \bullet \mid E N \mid \pi E$
Neutral terms	$R ::= E[H] \mid E[u[\sigma]]$
Normal terms	$M, N ::= R \mid \lambda x.M \mid (M, N)$
Substitutions	$\sigma, \tau ::= \cdot \mid \sigma, M$
Variable substitutions	$\rho, \xi ::= \cdot \mid \rho, x$
Contexts	$\Psi, \Phi, \Gamma ::= \cdot \mid \Psi, x:A$
Meta substitutions	$\theta, \eta ::= \cdot \mid \theta, \hat{\Psi}.M/u$
Meta contexts	$\Delta ::= \cdot \mid \Delta, u:A[\Psi]$

Meta-variables are characterized as a closure $u[\sigma]$ which is the use of the meta-variable u under the suspended explicit substitution σ . The term $\lambda x y z. X x y$

with the meta-variable X which has type $\Pi x:A.\Pi y:B.C$ is represented in our calculus as $\lambda x y z. u[x, y]$ where u has type $C[x:A, y:B]$ and $[x, y]$ is a substitution with domain $x:A, y:B$ and the range x, y, z . Instead of an abstraction, we can directly replace u with a contextual object $x, y. x(\text{succ } y)$. In general, a contextual object $\hat{\Psi}.M$ is a term M whose free variables have to be contained in the variable list $\hat{\Psi}$. (In the example, $\hat{\Psi} = x, y$ and $M = x(\text{succ } y)$.) The use of contextual objects eliminates the need to craft a λ -prefix for the instantiation of meta-variables, avoids spurious reductions, and provides simple justifications for techniques such as *lowering* (see Section 3.2). In general, meta-variable u of contextual type $A[\hat{\Psi}]$ stands for a contextual object $\hat{\Psi}.M$ where $\hat{\Psi}$ is the domain of $\hat{\Psi}$, i. e., the list of variables declared in $\hat{\Psi}$ in the correct order. The use of the prefix $\hat{\Psi}$ allows us to rename the free variables occurring in M if necessary. Our grammar and our subsequent typing rules enforce that objects are β -normal.

A signature Σ is a collection of declarations, which take one of the forms: $\mathbf{a} : \kappa$ (type family declaration) or $\mathbf{c} : A$ (constructor declaration). Because variable substitutions ρ play a special role in the formulation of our unification algorithm, we recognize them as a subclass of general substitutions σ .

Weakening substitutions $\boxed{\text{wk}_\Phi}$, which are a special case of variable substitutions, are defined recursively by $\text{wk}. = (\cdot)$ and $\text{wk}_{\Phi, x:A} = (\text{wk}_\Phi, x)$. The subscript Φ is dropped when unambiguous. Incidentally, $\text{wk}_\Phi = \hat{\Phi}$, but conceptually, one is a substitution and one a list of binders, thus, we keep them separate. If and only if Φ is a sub-context of Ψ modulo η -equality, then wk_Φ is a well-formed substitution in Ψ , i. e., $\Psi \vdash \text{wk}_\Phi : \Phi$ holds (see Fig. 1).

We write $\boxed{E[M]}$ for plugging term M into the hole \bullet of evaluation context E . This will be useful when describing the unification algorithm, since we often need to have access to the head of a neutral term. In the λ^Π -calculus, this is often achieved using the spine notation [2] simply writing $H M_1 \dots M_n$. Evaluation contexts are the proper generalization of spines to projections (see also [19] for a similar generalization in the linear setting).

Occurrences and free variables. If α, β are syntactic entities such as evaluation contexts, terms, or substitutions, $\alpha, \beta ::= E \mid R \mid M \mid \sigma$, we write $\boxed{\alpha\{\beta\}}$ if β is a part of α . If we subsequently write $\alpha\{\beta'\}$ then we mean to replace the indicated occurrence of β by β' . We say an occurrence is *rigid* if it is not part of a delayed substitution σ of a meta-variable, otherwise it is termed *flexible*. For instance, in $\mathbf{c}(u[y_1])(x_1 x_2)(\lambda z. z x_3 v[y_2, w[y_3]])$ there are rigid occurrences of $x_{1..3}$ and flexible occurrences of $y_{1..3}$. The meta-variables u, v appear in a rigid and w in a flexible position. A rigid occurrence is *strong* if it is not in the evaluation context of a free variable. In our example, only x_2 does *not* occur strongly rigidly. Following Reed [17] we indicate rigid occurrences by $\boxed{\alpha\{\beta\}^{\text{rig}}}$ and strongly rigid occurrences by $\boxed{\alpha\{\beta\}^{\text{srig}}}$.

Flexible variable occurrences can vanish by instantiation of meta-variables, rigid ones not. In our example, y_1 will disappear by the substitution $y_1.c'/u$.

Rigid variables might disappear by instantiation of other free variables and subsequent normalization. For instance, x_2 disappears when we substitute $\lambda x_2.c'$ for x_1 . However, strongly rigid ones will only disappear when they are instantiated themselves, not through other variable instantiations.

We denote the set of free variables of α by $\boxed{\text{FV}(\alpha)}$ and the set of free meta variables by $\boxed{\text{FMV}(\alpha)}$. A superscript ^{rig} indicates to count only the rigid variables.

Typing. We rely on a bi-directional type system to guarantee that well-typed terms are in β -normal form (Fig. 1). The typing rules are divided into rules which check that an object has a given type (\Leftarrow judgments) and rules which synthesize a type for a given object (\Rightarrow judgments). We have record types $\Sigma x:A. B$ but no record kinds $\Sigma x:A. \kappa$. Our typing rules ensure that terms are in β -normal form, but they need not be η -long. The judgment $\boxed{A =_\eta C}$ (rules omitted) compares A and C modulo η , i.e., modulo $R = \lambda x. R.x$ ($x \notin \text{FV}(R)$) and $R = (\text{fst } R, \text{snd } R)$.

Hereditary substitution and meta substitution. For α a well-typed entity in context Ψ and $\Delta; \Phi \vdash \sigma : \Psi$ a well-formed substitution, we facilitate a simultaneous substitution operation $\boxed{[\sigma]_\Psi \alpha}$ that substitutes the terms in σ for the variables as listed by Ψ in α and produces a β -normal result. Such an operation exists for well-typed terms, since $\lambda^{\Pi\Sigma}$ is normalizing. A naive implementation just substitutes and then normalizes. A refined implementation, called *hereditary substitution* [20], proceeds by resolving created redexes on the fly through new substitutions. Both Agda and Beluga use that strategy, as well as other theoretical investigations of logical frameworks. Single hereditary substitution $\boxed{[N/x]_A \alpha}$ is conceived as a special case of simultaneous substitution. The type annotation A and the typing information in Ψ allow hereditary substitution to be defined by structural recursion; if no ambiguity arises, we may omit indices Ψ and A from substitutions.

The *meta-substitution* operation, i.e., substitution of meta variables by contextual objects, is written as $\boxed{[\hat{\Psi}.M/u]N}$ and the simultaneous meta substitution is written as $\boxed{[\theta]N}$. Both operations restore β -normality. In the particular case when we apply $\hat{\Psi}.M/u$ to $u[\sigma]$, we first substitute $\hat{\Psi}.M$ for u in σ to obtain σ' . Subsequently, we continue to apply σ' to M hereditarily to obtain M' .

Detailed definitions of hereditary substitution and meta substitution can be found in the appendices A and B.

3 Constraint-based unification

We define the unification algorithm using rewrite rules which solve constraints incrementally. Constraints K and sets of constraints \mathcal{K} are defined as follows:

Neutral terms/types $\boxed{\Delta; \Psi \vdash R \Rightarrow A}$ (Δ and Ψ fixed)

$$\frac{\Sigma(\mathbf{a}) = \kappa}{\Delta; \Psi \vdash \mathbf{a} \Rightarrow \kappa} \quad \frac{\Sigma(\mathbf{c}) = A}{\Delta; \Psi \vdash \mathbf{c} \Rightarrow A} \quad \frac{\Psi(x) = A}{\Delta; \Psi \vdash x \Rightarrow A} \quad \frac{u:A[\Phi] \in \Delta \quad \Delta; \Psi \vdash \sigma \Leftarrow \Phi}{\Delta; \Psi \vdash u[\sigma] \Rightarrow [\sigma]_{\Phi} A}$$

$$\frac{\Delta; \Psi \vdash R \Rightarrow \Pi x:A.B \quad \Delta; \Psi \vdash M \Leftarrow A}{\Delta; \Psi \vdash RM \Rightarrow [M/x]_A B}$$

$$\frac{\Delta; \Psi \vdash R \Rightarrow \Sigma x:A.B}{\Delta; \Psi \vdash \text{fst } R \Rightarrow A} \quad \frac{\Delta; \Psi \vdash R \Rightarrow \Sigma x:A.B}{\Delta; \Psi \vdash \text{snd } R \Rightarrow [\text{fst } R/x]_A B}$$

Normal terms $\boxed{\Delta; \Psi \vdash M \Leftarrow A}$ (Δ fixed)

$$\frac{\Delta; \Psi \vdash R \Rightarrow A \quad A =_{\eta} C}{\Delta; \Psi \vdash R \Leftarrow C}$$

$$\frac{\Delta; \Psi, x:A \vdash M \Leftarrow B}{\Delta; \Psi \vdash \lambda x.M \Leftarrow \Pi x:A.B} \quad \frac{\Delta; \Psi \vdash M \Leftarrow A \quad \Delta; \Psi \vdash N \Leftarrow [M/x]_A B}{\Delta; \Psi \vdash (M, N) \Leftarrow \Sigma x:A.B}$$

Substitutions $\boxed{\Delta; \Psi \vdash \sigma \Leftarrow \Psi'}$ (Δ and Ψ fixed)

$$\frac{}{\Delta; \Psi \vdash \cdot \Leftarrow \cdot} \quad \frac{\Delta; \Psi \vdash \sigma \Leftarrow \Psi' \quad \Delta; \Psi \vdash M \Leftarrow [\sigma]_{\Psi'} A}{\Delta; \Psi \vdash \sigma, M \Leftarrow \Psi', x:A}$$

LF types and kinds $\boxed{\Delta; \Psi \vdash A \Leftarrow s}$ (Δ fixed)

$$\frac{\Delta; \Psi \vdash P \Rightarrow \text{type}}{\Delta; \Psi \vdash P \Leftarrow \text{type}} \quad \frac{\Delta; \Psi \vdash A \Leftarrow \text{type} \quad \Delta; \Psi, x:A \vdash B \Leftarrow \text{type}}{\Delta; \Psi \vdash \Sigma x:A.B \Leftarrow \text{type}}$$

$$\frac{}{\Delta; \Psi \vdash \text{type} \Leftarrow \text{kind}} \quad \frac{\Delta; \Psi \vdash A \Leftarrow \text{type} \quad \Delta; \Psi, x:A \vdash B \Leftarrow s}{\Delta; \Psi \vdash \Pi x:A.B \Leftarrow s}$$

LF typing contexts $\boxed{\Delta \vdash \Psi \text{ ctx}}$ (Δ fixed)

$$\frac{}{\Delta \vdash \cdot \text{ ctx}} \quad \frac{\Delta \vdash \Psi \text{ ctx} \quad \Delta; \Psi \vdash A \Leftarrow \text{type}}{\Delta \vdash \Psi, x:A \text{ ctx}}$$

Meta substitutions $\boxed{\Delta \vdash \theta \Leftarrow \Delta'}$ (Δ fixed)

$$\frac{\text{for all } u:A[\Phi] \in \Delta' \text{ and } \hat{\Phi}.M/u \in \theta : \quad \Delta; [\theta]\Phi \vdash M \Leftarrow [\theta]A}{\Delta \vdash \theta \Leftarrow \Delta'}$$

Meta contexts $\boxed{\vdash \Delta \text{ mctx}}$

$$\frac{\text{for all } u:A[\Psi] \in \Delta : \quad \Delta \vdash \Psi \text{ ctx} \quad \Delta; \Psi \vdash A \Leftarrow \text{type}}{\vdash \Delta \text{ mctx}}$$

Fig. 1. Typing rules for LF with meta-variables

Constraint	$K ::= \top \mid \perp$	Trivial constraint and inconsistency.
	$\mid \Psi \vdash M = N : C$	Unify term M with N .
	$\mid \Psi \mid R:A \vdash E = E'$	Unify evaluation context E with E' .
	$\mid \Psi \vdash u \leftarrow M : C$	Solution for u found.
C. sets	$\mathcal{K} ::= K \mid \mathcal{K} \wedge K$	(modulo laws of conjunction).

Our basic constraints are of the form $\Psi \vdash M = N : C$. The type annotation $\Psi \vdash C$ serves two purposes: First, we need to types to direct any hereditary substitutions we employ during constraint solving. Secondly, the type annotations in the context Ψ are necessary to eliminate Σ -types. For both purposes, simple types, i.e., the dependency-erasure of $\Psi \vdash C$ would suffice. However, we keep dependencies in this presentation to scale this work from $\lambda^{\text{II}\Sigma}$ to non-erasable dependent types such as featured by Agda.

A unification problem is described by $\boxed{\Delta \Vdash \mathcal{K}}$ where Δ contains the typings of all the meta variables in \mathcal{K} . A meta-variable u is *solved*, if there is a constraint $\Psi \vdash u \leftarrow M : C$ in \mathcal{K} ; otherwise we call u *active*. A solved metavariable does not appear in any other constraints nor in any type in Δ (nor in its solution M).

Intuitively, a set of constraints is well-formed if each constraint $\Psi \vdash M = N : C$ is well typed. Unfortunately, this is complicated by the fact that we may delay working on some sub-terms; to put it differently, we can work on subterms in an arbitrary order. Yet, the type of an equation may depend on the solvability of another postponed equation. Consider for example tuples. If (M_1, M_2) and (N_1, N_2) both have type $\Sigma x:A.B$, then M_1 and N_1 have type A . However, types may get out of sync when we consider M_2 and N_2 . M_2 has type $[M_1/x]B$ while N_2 has type $[N_1/x]B$, and we only know that their types agree, if we know that M_1 is equal to N_1 . Similar issues arise for function types and applications. Following Reed [17], we adopt here a weaker typing invariant, namely typing modulo constraints.

3.1 Typing modulo

For all typing judgments $\Delta; \Psi \vdash J$ defined previously, we define $\boxed{\Delta; \Psi \vdash_{\mathcal{K}} J}$ by the same rules as for $\Delta; \Psi \vdash J$ except replacing eta equality $=_{\eta}$ with $=_{\mathcal{K}}$. We write $\boxed{\alpha =_{\mathcal{K}} \beta}$ if for any ground meta-substitution θ that is a ground solution for \mathcal{K} , we have $\llbracket \theta \rrbracket \alpha =_{\eta} \llbracket \theta \rrbracket \beta$. To put it differently, if we can solve \mathcal{K} , we can establish that α is equal to β .

The following lemmas proven by Reed [17] hold also for the extension to Σ -types; we keep in mind that the judgment J stands for either a typing judgment or an equality judgment. We first prove that typing modulo is preserved under equality modulo.

Lemma 1 (Conversion modulo). *Let $\Delta_0 \Vdash \mathcal{K}$ and $\Delta =_{\mathcal{K}} \Delta'$ and $\Psi =_{\mathcal{K}} \Phi$ and $A =_{\mathcal{K}} B$.*

1. If $\Delta; \Psi \vdash_{\mathcal{K}} M \Leftarrow A$ then $\Delta'; \Phi \vdash_{\mathcal{K}} M \Leftarrow B$.
2. If $\Delta; \Psi \vdash_{\mathcal{K}} R \Rightarrow A$ then $\Delta'; \Phi \vdash_{\mathcal{K}} R \Rightarrow B$.
3. If $\Delta; \Psi \vdash_{\mathcal{K}} \sigma \Leftarrow \Psi'$ and $\Psi' =_{\mathcal{K}} \Phi'$ then $\Delta'; \Phi \vdash_{\mathcal{K}} \sigma \Leftarrow \Phi'$.

Proof. We generalize the statement to types and kinds and prove them by simultaneous induction on the typing derivation.

Lemma 2 (Substitution principle modulo). *Let $\Delta_0 \Vdash \mathcal{K}$. If $\Delta; \Psi \vdash_{\mathcal{K}} M \Leftarrow A$ and $\Delta; \Psi, x:B, \Psi' \vdash_{\mathcal{K}} J$ and $A =_{\mathcal{K}} B$ then $\Delta; \Psi, [M/x]_A \Psi' \vdash_{\mathcal{K}} [M/x]_A J$.*

Proof. This proof follows essentially the proof by Nanveski et al. [11] and we generalize the property to types, kinds, and contexts. Since we prove the substitution lemma modulo $A =_{\mathcal{K}} B$, we use Lemma 1 when we for example consider the variable case.

Lemma 3 (Meta-substitution principle modulo). *Let $\Delta_0 \Vdash \mathcal{K}$. If $\Delta_1 \vdash_{\mathcal{K}} \theta \Leftarrow \Delta$ and $\Delta; \Phi \vdash_{\mathcal{K}} J$ then $\mathcal{K}' = \llbracket \theta \rrbracket \mathcal{K}$ is defined and $\Delta_1; \llbracket \theta \rrbracket \Phi \vdash_{\mathcal{K}'} \llbracket \theta \rrbracket J$.*

Proof. Let $\cdot \vdash \eta \Leftarrow \Delta_0$ be a ground solution of $\Delta_0 \Vdash \mathcal{K}$. By assumption, $\llbracket \eta \rrbracket \Delta_1 \vdash \llbracket \eta \rrbracket \theta \Leftarrow \llbracket \eta \rrbracket \Delta$ and $\llbracket \eta \rrbracket \Delta; \llbracket \eta \rrbracket \Phi \vdash \llbracket \eta \rrbracket J$.

Induction on $\Delta; \Phi \vdash_{\mathcal{K}} J$. All cases are by inversion, appeal to the induction hypothesis (i.h.), reassembling the result and if necessary using Lemma 1 (see the case for meta-variables). We show the case where we transition between checking and synthesizing a type.

Case

$$\mathcal{D} = \frac{\Delta; \Phi \vdash_{\mathcal{K}} R \Rightarrow C_1 \quad C_1 =_{\mathcal{K}} C_2}{\Delta; \Phi \vdash_{\mathcal{K}} R \Leftarrow C_2}$$

$$\begin{array}{l} \Delta_1; \llbracket \theta \rrbracket \Phi \vdash_{\llbracket \theta \rrbracket \mathcal{K}} \llbracket \theta \rrbracket R \Rightarrow \llbracket \theta \rrbracket C_1 \\ \llbracket \theta \rrbracket C_1 =_{\llbracket \theta \rrbracket \mathcal{K}} \llbracket \theta \rrbracket C_2 \\ \Delta_1; \llbracket \theta \rrbracket \Phi \vdash_{\llbracket \theta \rrbracket \mathcal{K}} \llbracket \theta \rrbracket R \Leftarrow \llbracket \theta \rrbracket C_2. \end{array} \quad \begin{array}{l} \text{by i.h.} \\ \text{by i.h. } C_1 =_{\mathcal{K}} C_2 \end{array}$$

Case

$$\mathcal{D} = \frac{\Delta; \Psi \vdash_{\mathcal{K}} \sigma \Leftarrow \Psi' \quad u:B[\Psi'] \in \Delta}{\Delta; \Psi \vdash_{\mathcal{K}} u[\sigma] \Leftarrow [\sigma]_{\Psi'}(B)}$$

For $\hat{\Psi}'.M/u \in \theta$ and $u:B[\Psi'] \in \Delta$, $\Delta_1; \llbracket \theta \rrbracket \Psi' \vdash_{\mathcal{K}} M \Leftarrow \llbracket \theta \rrbracket B$.

$$\begin{array}{l} \Delta_1; \llbracket \theta \rrbracket \Psi \vdash_{\llbracket \theta \rrbracket \mathcal{K}} \llbracket \theta \rrbracket \sigma \Leftarrow \llbracket \theta \rrbracket \Psi' \\ \llbracket \theta \rrbracket (u[\sigma]) = \llbracket \llbracket \theta \rrbracket \sigma \rrbracket_{\Psi'}(M) \\ \Delta_1; \llbracket \theta \rrbracket \Psi \vdash_{\llbracket \theta \rrbracket \mathcal{K}} \llbracket \llbracket \theta \rrbracket \sigma \rrbracket M \Leftarrow \llbracket \llbracket \theta \rrbracket \sigma \rrbracket (\llbracket \theta \rrbracket B) \\ \Delta_1; \llbracket \theta \rrbracket \Psi \vdash_{\llbracket \theta \rrbracket \mathcal{K}} \llbracket \llbracket \theta \rrbracket \sigma \rrbracket M \Leftarrow \llbracket \llbracket \theta \rrbracket \sigma \rrbracket (\llbracket \theta \rrbracket B) \end{array} \quad \begin{array}{l} \text{by i.h.} \\ \text{by definition} \\ \text{by ord. subst. lemma} \\ \text{by definition of msubst. } \square \end{array}$$

Intuitively, a unification problem $\Delta \Vdash \mathcal{K}$ is well-formed if all constraints $(\Psi \vdash M = N : C) \in \mathcal{K}$ are well-typed modulo \mathcal{K} , i.e., $\Delta; \Psi \vdash_{\mathcal{K}} M \Leftarrow C$ and $\Delta; \Psi \vdash_{\mathcal{K}} N \Leftarrow C$. We will come back to this later when we prove correctness of our algorithm, but it is helpful to keep the typing invariant in mind when explaining the transitions in our algorithm.

Decomposition of functions

$$\begin{aligned} \Psi \vdash \lambda x.M = \lambda x.N : \Pi x:A. B & \mapsto_d \Psi, x:A \vdash M = N : B \\ \Psi \vdash \lambda x.M = R : \Pi x:A. B & \mapsto_d \Psi, x:A \vdash M = Rx : B \\ \Psi \vdash R = \lambda x.M : \Pi x:A. B & \mapsto_d \Psi, x:A \vdash Rx = M : B \end{aligned}$$

Decomposition of pairs

$$\begin{aligned} \Psi \vdash (M_1, M_2) = (N_1, N_2) : \Sigma x:A. B & \mapsto_d \Psi \vdash M_1 = N_1 : A \wedge \Psi \vdash M_2 = N_2 : [M_1/x]B \\ \Psi \vdash (M_1, M_2) = R : \Sigma x:A. B & \mapsto_d \Psi \vdash M_1 = \text{fst } R : A \wedge \Psi \vdash M_2 = \text{snd } R : [M_1/x]B \\ \Psi \vdash R = (M_1, M_2) : \Sigma x:A. B & \mapsto_d \Psi \vdash \text{fst } R = M_1 : A \wedge \Psi \vdash \text{snd } R = M_2 : [\text{fst } R/x]B \end{aligned}$$

Decomposition of neutrals

$$\begin{aligned} \Psi \vdash E[H] = E'[H] : C & \mapsto_d \Psi \mid H : A \vdash E = E' \text{ where } \Psi \vdash H \Rightarrow A \\ \Psi \vdash E[H] = E'[H'] : C & \mapsto_d \perp \text{ if } H \neq H' \end{aligned}$$

Decomposition of evaluation contexts

$$\begin{aligned} \Psi \mid R : A \vdash \bullet = \bullet & \mapsto_d \top \\ \Psi \mid R : \Pi x:A. B \vdash E[\bullet M] = E'[\bullet M'] & \mapsto_d \Psi \vdash M = M' : A \wedge \Psi \mid RM : [M/x]B \vdash E = E' \\ \Psi \mid R : \Sigma x:A. B \vdash E[\text{fst } \bullet] = E'[\text{fst } \bullet] & \mapsto_d \Psi \mid \text{fst } R : A \vdash E = E' \\ \Psi \mid R : \Sigma x:A. B \vdash E[\text{snd } \bullet] = E'[\text{snd } \bullet] & \mapsto_d \Psi \mid \text{snd } R : [\text{fst } R/x]B \vdash E = E' \\ \Psi \mid R : \Sigma x:A. B \vdash E[\pi \bullet] = E'[\pi' \bullet] & \mapsto_d \perp \text{ if } \pi \neq \pi' \end{aligned}$$

Orientation

$$\Psi \vdash M = u[\sigma] : C \text{ with } M \neq v[\dots] \mapsto_d \Psi \vdash u[\sigma] = M : C$$

 η -Contraction

$$\begin{aligned} \Psi \vdash u[\sigma\{\lambda x.Rx\}] = N : C & \mapsto_e \Psi \vdash u[\sigma\{R\}] = N : C \\ \Psi \vdash u[\sigma\{\text{fst } R, \text{snd } R\}] = N : C & \mapsto_e \Psi \vdash u[\sigma\{R\}] = N : C \end{aligned}$$

Eliminating projections

$$\begin{aligned} \Psi_1, x : \Pi \mathbf{y}:A. \Sigma z:B. C, \Psi_2 & \Psi_1, x_1 : \Pi \mathbf{y}:A. B, x_2 : \Pi \mathbf{y}:A. [(x_1 \mathbf{y})/z]C, \Psi_2 \\ \vdash u[\sigma\{\pi(xM)\}] = N : D & \mapsto_p \vdash u[\tau] = [\tau]N : [\tau]D \\ \text{where } \pi \in \{\text{fst}, \text{snd}\} & \text{ where } \tau = [\lambda \mathbf{y}. (x_1 \mathbf{y}, x_2 \mathbf{y})/x] \end{aligned}$$

Fig. 2. Local simplification $\boxed{K \mapsto_m \mathcal{K}}$.**3.2 A higher-order dynamic pattern unification algorithm for dependent types and records**

The higher-order dynamic pattern unification algorithm is presented as rewrite rules on the set of constraints \mathcal{K} in meta-variable context Δ . The *local simplification rules* (Figure 2) apply to a single constraint, decomposing it and molding it towards a pattern by η -contraction and projection elimination. Decomposition of neutral terms is defined using evaluation contexts to have direct access to the head.

Local simplification

$$\Delta \Vdash \mathcal{K} \wedge K \mapsto \Delta \Vdash \mathcal{K} \wedge \mathcal{K}' \quad \text{if } K \mapsto_m \mathcal{K}' \quad (m \in \{\mathbf{d}, \mathbf{e}, \mathbf{p}\})$$

Instantiation (notation)

$$\boxed{\Delta \Vdash \mathcal{K} + (\Phi \vdash u \leftarrow M : A)} = \llbracket \theta \rrbracket \Delta \Vdash \llbracket \theta \rrbracket \mathcal{K} \wedge \llbracket \theta \rrbracket \Phi \vdash u \leftarrow M : \llbracket \theta \rrbracket A$$

where $\theta = \hat{\Phi}.M/u$

Lowering

$$\begin{aligned} \Delta \Vdash \mathcal{K} &\mapsto \Delta, v:B[\Phi, x:A] \Vdash \mathcal{K} \\ u:(\Pi x:A.B)[\Phi] \in \Delta \text{ active} &+ \Phi \vdash u \leftarrow \lambda x.v : \Pi x:A. B \\ \Delta \Vdash \mathcal{K} &\mapsto \Delta, u_1:A[\Phi], u_2:([u_1/x]_A B)[\Phi] \Vdash \mathcal{K} \\ u:(\Sigma x:A.B)[\Phi] \in \Delta \text{ active} &+ \Phi \vdash u \leftarrow (u_1, u_2) : \Sigma x:A. B \end{aligned}$$

Flattening Σ -types

$$\begin{aligned} \Delta \Vdash \mathcal{K} \quad (u:A[\Phi] \in \Delta \text{ active}) &\mapsto \Delta, v:([\sigma^{-1}]A)[\Phi'] \Vdash \mathcal{K} + \Phi \vdash u \leftarrow v[\sigma] : A \\ \Phi = \Phi_1, x : \Pi \mathbf{y}:\mathbf{A}. \Sigma z:B.C, \Phi_2 &\quad \Phi' = \Phi_1, x_1 : \Pi \mathbf{y}:\mathbf{A}. B, x_2 : \Pi \mathbf{y}:\mathbf{A}. [x_1 \mathbf{y}/z]C, \Phi_2 \\ \sigma^{-1} = [\lambda \mathbf{y}. (x_1 \mathbf{y}, x_2 \mathbf{y})/x] &\quad \sigma = [\lambda \mathbf{y}. \text{fst}(x \mathbf{y})/x_1, \lambda \mathbf{y}. \text{snd}(x \mathbf{y})/x_2] \end{aligned}$$

Pruning

$$\begin{aligned} \Delta \Vdash \mathcal{K} &\mapsto \Delta' \Vdash \llbracket \eta \rrbracket \mathcal{K} \\ (\Psi \vdash u[\rho] = M : C) \in \mathcal{K} &\quad \text{if } \Delta \vdash \text{prune}_\rho M \Rightarrow \Delta'; \eta \text{ and } \eta \neq \text{id} \end{aligned}$$

Same meta-variable

$$\begin{aligned} \Delta \Vdash \mathcal{K} \wedge \Psi \vdash u[\rho] = u[\xi] : C &\mapsto \Delta, v:A[\Phi_0] \Vdash \mathcal{K} + \Phi \vdash u \leftarrow v[\text{wk}_{\Phi_0}] : A \\ u:A[\Phi] \in \Delta &\quad \text{if } \rho \cap \xi : \Phi \Rightarrow \Phi_0 \end{aligned}$$

Failing occurs check

$$\begin{aligned} \Delta \Vdash \mathcal{K} \wedge \Psi \vdash u[\rho] = M : C &\mapsto \perp \text{ if } \text{FV}^{\text{rig}}(M) \not\subseteq \rho \\ \Delta \Vdash \mathcal{K} \wedge \Psi \vdash u[\rho] = M : C &\mapsto \perp \text{ if } M = M'\{u[\xi]\}^{\text{srig}} \neq u[\xi] \end{aligned}$$

Solving (with successful occurs check)

$$\begin{aligned} \Delta \Vdash \mathcal{K} \wedge \Psi \vdash u[\rho] = M : C &\mapsto \Delta \Vdash \mathcal{K} + \Phi \vdash u \leftarrow M' : A \\ (u:A[\Phi]) \in \Delta; u \notin \text{FMV}(M) &\quad \text{if } M' = [\rho/\hat{\Phi}]^{-1}M \text{ exists and } \Delta; \Phi \vdash M' \Leftarrow A \end{aligned}$$

Fig. 3. Unification steps $\boxed{\Delta \Vdash \mathcal{K} \mapsto \Delta' \Vdash \mathcal{K}'}$.

Decomposition of pairs could maybe more concisely defined by

$$\begin{aligned} \Psi \vdash M = N : \Sigma x:A. B &\mapsto_{\mathbf{d}} \Psi \vdash \text{fst}@M = \text{fst}@N : A \\ \wedge \Psi \vdash \text{snd}@M = \text{snd}@N : [\text{fst}@M/x]B & \end{aligned}$$

where $\pi@M$ computes the β -normal form of πM . However, this would also apply to a constraint $\Psi \vdash R = R' : \Sigma x:A. B$ of neutral terms, just duplicating the work of decomposing the neutrals later. Similar reasoning justifies choice of function decomposition rules.

The other *unification steps* (Figure 3) work on a meta-variable and try to find an instantiation for it. We write $\boxed{\Delta \Vdash \mathcal{K} + \Phi \vdash u \leftarrow M : A}$ for instantiating the meta variable u with the term M both in the meta-context Δ and in the

constraints \mathcal{K} . This abbreviation is defined in Figure 3. Lowering rules transform a meta-variable of higher type to one of lower type. Flattening Σ -types concentrates on a meta-variable $u:A[\Phi]$ and eliminates Σ -types from the context Φ . The combination of the flattening Σ -types transition and the eliminating projections transition allow us to transform a unification problem into one which resembles our traditional pattern unification problem. The pruning transition is explained in detail in Section 3.4 and unifying a meta-variable with itself is discussed in Section 3.5.

To motivate our rules, let us consider some problems $\Psi \vdash u[\sigma] = M : C$ that fall outside of the Miller pattern fragment, meaning that σ is not a list of disjoint variables. We may omit types and/or context if appropriate.

η -contraction $u[\lambda x. y (\text{fst } x, \text{snd } x)] = M$.

Solved by contracting the l.h.s. to $u[y]$.

Eliminating projections $y : \Pi x:A. \Sigma z:B. C \vdash u[\lambda x. \text{fst } (y x)] = M$.

Applying substitution $\tau = [\lambda x. (y_1 x, y_2 x)/y]$ yields problem $y_1 : \Pi x:A. B$, $y_2 : \Pi x:A. [y_1 x/z]C \vdash u[\lambda x. y_1 x] = [\tau]M$ which is solved by η -contraction, provided $y_2 \notin \text{FV}([\tau]M)$.

Lowering $\Phi \vdash \text{fst } (u[y]) = \text{fst } y$ where $\Phi = (y : \Sigma x:A. B)$ and $u : (\Sigma x:A. B)[\Phi]$.

This equation determines only the first component of the tuple u . Thus, decomposition into $u[y] = y$, which also determines the second component, loses solutions. Instead we replace u by a pair (u_1, u_2) of meta-variables of lower type, $u_1 : A[\Phi]$ and $u_2 : ([u_1[y]/x]B)[\Phi]$, yielding $\Phi \vdash u_1[y] = \text{fst } y$.

Flattening Σ -types $\Psi \vdash u[\lambda x. (z_1 x, z_2 x)] = \mathbf{g} z_1 z_2$ where $\Psi(z_1) = \Pi x:A. B$ and $\Psi(z_2) = \Pi x:A. [z_1 x/y]C$ and $u : P[z : \Pi x:A. \Sigma y:B. C]$.

By splitting z into two functions z_1, z_2 in the context of u and replacing u by $v : P[z_1 : \Pi x:A. B, z_2 : \Pi x:A. [z_1 x/y]C]$ via meta substitution $z.v[\lambda x. \text{fst } z x, \lambda x. \text{snd } z x]/u$, we arrive at $\Psi \vdash v[\lambda x. z_1 x, \lambda x. z_2 x] = \mathbf{g} z_1 z_2$ and continue with η -contraction.

Solving in spite of non-linearity $u[x, x, z] = \text{suc } z$.

The non-linear occurrence of x on the l.h.s. can be ignored since x is not free on the r.h.s. We can solve this constraint by $u[x, y, z] = \text{suc } z$.

However, in case of non-linearity we need to make sure that the solution is well-typed. Consider $u : P z [\Phi]$ where $\Phi = (x:A, y : P x, z:A)$ and constraint

$$x:A, y : P x \vdash u[x, y, x] = y : P x.$$

The solution $x:A, y : P x \vdash u \leftarrow y : P z$ is ill-typed. While the non-linear variables do not appear in the term of the rhs, they appear in the type and make the constraint well-typed.

In our type theory, the above constraint is probably unsolvable (but we attempt no proof here). In richer type theories such as Agda one could imagine a solution for u which applies a type cast to y . Thus, we just leave such a constraint alone, instead of flagging unsolvability.

Pruning $u[x] = \text{suc}(v[x, y])$ and $v[x, \text{zero}] = f(x, \text{zero})$.

Since u depends only on x , necessarily v cannot depend on y . We can prune away the second parameter of v by setting $v[x, y] = v'[x]$. This turns the second constraint into the pattern $v'[x] = f(x, \text{zero})$, yielding the solution $u[x] = \text{suc}(f(x, \text{zero}))$.

Note that pruning is more difficult in case of nested meta variables. If instead $u[x] = \text{suc}(v[x, w[y]])$ then there are two cases: either v does not depend on its second argument or w is constant. Pruning as we describe it in this article cannot be applied to this case; Reed [17] proceeds here by replacing y by a placeholder “_”. Once w gets solved the placeholder might occur as argument to v , where it can be pruned. If the placeholder appears in a rigid position, the constraints have no solution.

Pruning and non-linearity $u[x, x] = v[x]$ and $u'[x, x] = v'[x, y]$.

Even though we cannot solve for u due to the non-linear x , pruning x from v could lose solutions. However, we can prune y from v' since only x can occur in $v'[x, y]$.

Failing occurs check $u[x] = \text{suc } y$.

Pruning y fails because it occurs rigidly. The constraint set has no solution.

Same meta-variable $u[x, y, x, z] = u[x, y, y, x]$.

Since variables x, y, z are placeholders for arbitrary *open* well-typed terms, of which infinitely many exists for every type, the above equation can only hold if u does not depend on its 3rd and 4th argument. Thus, we can solve by $u[x, y, z, x'] = v[x, y]$ where $[x, y]$ is the *intersection* of the two variable environments $[x, y, x, z]$ and $[x, y, y, x]$.

Recursive occurrence $u[x, y, x] = \text{suc } u[x, y, y]$.

Here, u has a *strong* rigid occurrence in its own definition. Even though not in the pattern fragment, this only has an infinite solution: consider the instance $u[z, z, z] = \text{suc } u[z, z, z]$. Consequently, the occurs check signals unsolvability. Reed [18, p. 105f] motivates why only *strong* rigid recursive occurrences force unsolvability. For instance, $f : \text{nat} \rightarrow \text{nat} \vdash u[f] = \text{suc}(f(u[\lambda x. \text{zero}]))$ has solution $u[f] = \text{suc}(f(\text{suc } \text{zero}))$ in spite of a rigid occurrence of u in its definition.

If u occurs flexibly in its own definition, like in $u[x] = v[u[x]]$, we cannot proceed until we know more of v . Using the other constraints, we might manage to prune v 's argument, arriving at $u[x] = v[]$, or find the solution of v directly; in these cases, we can revisit the constraint on u .

The examples suggest a *strategy* for implementation: Lowering can be integrated triggered by decomposition to resolve eliminations of a meta variable $E[u[\sigma]]$. After decomposition we have a set of $u[\sigma] = M$ problems. We try to turn the σ s into variable substitutions by applying η -contraction, and where this gets stuck, elimination of projections and Σ -flattening. Solution of constraints $u[\rho] = M$ can then be attempted by pruning, where a failing occurs check signals unsolvability.

3.3 Inverting substitutions

A most general solution for a constraint $u[\sigma] = M$ can only be hoped for if σ is a variable substitution. For instance $u[\text{true}] = \text{true}$ admits already two different solutions $u[x] = x$ and $u[x] = \text{true}$ that are pure λ -terms. In a language with computation such as Agda infinitely more solutions are possible, because $u[x]$ could be defined by cases on x and the value of $u[\text{false}]$ is completely undetermined.

But even constraints $u[\rho] = M$ can be ambiguous if the variable substitution ρ is not linear, i. e., no bijective variable renaming. For example, $u[x, x] = x$ has solutions $u[x, y] = x$ and $u[x, y] = y$. Other examples, like $u[x, x, z] = z$, which has unique solution $u[x, y, z] = z$, suggest that we can ignore non-linear variable occurrences as long as they do not occur on the r.h.s. Indeed, if we define a variable substitution ρ to be *invertible for term* M if there is exactly one M' such that $[\rho]M' = M$, then linearity is a sufficient, but not necessary condition. However, it is necessary that ρ must be linear if restricted to the free variables of (β -normal!) M . Yet instead of computing the free variables of M , checking that ρ is invertible, inverting ρ and applying the result to M , we can directly try to invert the effect of the substitution ρ on M .

For a variable substitution $\Psi \vdash \rho \Leftarrow \hat{\Phi}$ and a term or substitution $\alpha ::= M \mid R \mid \tau$ in context Ψ , we define the partial operation $[\rho/\hat{\Phi}]^{-1}\alpha$ by

$$\begin{aligned} [\rho/\hat{\Phi}]^{-1}x &= y && \text{if } x/y \in \rho/\hat{\Phi} \text{ and there is no } z \neq y \text{ with } x/z \in \rho/\hat{\Phi}, \\ &&& \text{undefined otherwise} \\ [\rho/\hat{\Phi}]^{-1}c &= c \\ [\rho/\hat{\Phi}]^{-1}(u[\tau]) &= u[\tau'] && \text{where } \tau' = [\rho/\hat{\Phi}]^{-1}\tau \end{aligned}$$

and homeomorphic in all other cases by

$$\begin{aligned} [\rho/\hat{\Phi}]^{-1}(RM) &= R' M' && \text{where } R' = [\rho/\hat{\Phi}]^{-1}R \text{ and } M' = [\rho/\hat{\Phi}]^{-1}M \\ [\rho/\hat{\Phi}]^{-1}(\pi R) &= \pi R' && \text{where } R' = [\rho/\hat{\Phi}]^{-1}R \\ [\rho/\hat{\Phi}]^{-1}(\lambda x. M) &= \lambda x. M' && \text{if } x \text{ not declared or free in } \rho \\ &&& \text{and } M' = [\rho, x/\hat{\Phi}, x]^{-1}M \\ [\rho/\hat{\Phi}]^{-1}(M, N) &= (M', N') && \text{where } M' = [\rho/\hat{\Phi}]^{-1}M \text{ and } N' = [\rho/\hat{\Phi}]^{-1}N \\ [\rho/\hat{\Phi}]^{-1}(\cdot) &= \cdot \\ [\rho/\hat{\Phi}]^{-1}(\tau, M) &= \tau', M' && \text{if } \tau' = [\rho/\hat{\Phi}]^{-1}\tau \text{ and } M' = [\rho/\hat{\Phi}]^{-1}M. \end{aligned}$$

We can show by induction on α , that inverse substitution $[\rho/\hat{\Phi}]^{-1}\alpha$ is correct and commutes with meta substitutions.

Lemma 4 (Inverse and meta-substitution commute). *Let ρ be a variable substitution, and $\alpha ::= M \mid R \mid \tau$. If $[\rho/\hat{\Phi}]^{-1}\alpha$ and $[\rho/\hat{\Phi}]^{-1}([\theta]\alpha)$ exist then $[\rho/\hat{\Phi}]^{-1}([\theta]\alpha) = [\theta]([\rho/\hat{\Phi}]^{-1}\alpha)$.*

Proof. By simultaneous induction on the structure of α . □

Lemma 5 (Soundness of inverse substitution). *If $[\rho/\hat{\Phi}]^{-1}\alpha$ exists then $[\rho]_{\hat{\Phi}}([\rho/\hat{\Phi}]^{-1}\alpha) = \alpha$.*

Proof. By simultaneous induction on the structure of α . □

Lemma 6 (Completeness of inverse substitution). *If $[\rho]_{\hat{\Phi}}\alpha = \alpha'$ and $\rho \upharpoonright \text{FV}(\alpha)$ is linear then $\alpha = [\rho/\hat{\Phi}]^{-1}\alpha'$ exists.*

Proof. By simultaneous induction on the structure of α . □

3.4 Pruning

If the constraint $u[\sigma] = M$ has a solution θ , then $[[\theta]]\sigma\theta(u) = [[\theta]]M$, and since θ is closed ($\text{FV}(\theta) = \emptyset$), we have $\text{FV}(\sigma) \supseteq \text{FV}([[\theta]])\sigma \supseteq \text{FV}([[\theta]])\sigma\theta(u) \supseteq \text{FV}([[\theta]])M$ (note that a hereditary substitution can remove free variables). Thus, if $\text{FV}(M) \not\subseteq \text{FV}(\sigma)$ we can try to find a most general meta-substitution η which *prunes* the free variables of M that are not in the range of σ , such that $\text{FV}([[\eta]])M \subseteq \text{FV}(\sigma)$. For instance, in case $u[x] = \text{succ } v[x, y]$, the meta-substitution $x, y. v'[x]/v$ does the job. However, pruning may fail for one of the following reasons:

1. Offending variables occur rigidly, like the y in $u[x] = \text{c } y \ v[x, y]$. This constraint is unsolvable.
2. The flexible occurrence of an offending variable is under another meta variable, like y in $u[x] = v[x, w[x, y]]$. Here, two minimal pruning substitutions $\eta_1 = x, y. v'[x]/v$ and $\eta_2 = x, y. w'[x]/w$ exist which are not instances of each other—applying pruning might lose solutions.
3. The offending flexible occurrence could be eliminated by the correct solution. For instance, consider the case $u : C[x:A]$ and $v : C[z : (A \rightarrow A \rightarrow A) \rightarrow A]$ and constraint

$$x:A, y:A \vdash u[x] = v[\lambda k. k x y] : C.$$

The offending variable y occurs flexibly in this constraint and rigidly in v 's substitution. If we pruned away v 's dependency on z , we would lose the partial solution $\theta(v) = z. z (\lambda x \lambda y. x)$ which would simplify the constraint to $u[x] = x$. The point here is that y , although it occurs rigidly in $\lambda k. k x y$, it is in an eliminateable position since the meta-substitution for v could place that term in a context that reduces y away. There are rigid occurrences that cannot be eliminated in such a way, we call these occurrences *bad*, see judgement $\boxed{\text{bad_occ}_y N}$ in Fig. 4.

We restrict pruning to situations $u[\rho] = M$ where ρ is a variable substitution. This is because we view pruning as a preparatory step to inverting ρ on M —which only makes sense for variable substitutions. Also, we do not consider partial pruning, as in pruning y from v in the situation $u[x] = v[x, y, w[x, y]]$, obtaining $u[x] = v'[x, w[x, y]]$. Such extensions to pruning are conceivable, but we have no data indicating that they strengthen unification significantly in practice. We employ the following judgments to define pruning (see Fig. 4):

$$\begin{aligned} \Delta \vdash \text{prune}_\rho M \Rightarrow \Delta'; \eta & \quad \text{prune } M \text{ such that } \text{FV}([[\eta]])M \subseteq \rho \\ \text{prune_ctx}_\rho(\tau / \Psi_1) \Rightarrow \Psi_2 & \quad \text{prune } \tau \text{ such that } \text{FV}^{\text{rig}}([\tau]\text{wk}_{\Psi_2}) \subseteq \rho. \end{aligned}$$

$\text{bad_occ}_x M$	Term M has a non-eliminateable occurrence of variable x .
$\frac{}{\text{bad_occ}_x E[x]}$	$\frac{\text{bad_occ}_x M}{\text{bad_occ}_x \lambda y. M} \quad x \neq y$
	$\frac{\text{bad_occ}_x M_1 \quad \text{bad_occ}_x M_2}{\text{bad_occ}_x (M_1, M_2)}$
$\text{prune_ctx}_\rho(\tau / \Psi_1) \Rightarrow \Psi_2$	Prune substitution $\tau : \Psi_1$, returning a sub-context Ψ_2 of Ψ_1 . (I.e., $\Psi_1 \vdash \text{wk}_{\Psi_2} : \Psi_2$.)
$\frac{}{\text{prune_ctx}_\rho(\cdot / \cdot) \Rightarrow \cdot}$	$\frac{\text{prune_ctx}_\rho(\tau / \Psi_1) \Rightarrow \Psi_2 \quad \text{bad_occ}_x M \text{ for some } x \in \rho}{\text{prune_ctx}_\rho(\tau, M / \Psi_1, y:A) \Rightarrow \Psi_2}$
$\frac{\text{prune_ctx}_\rho(\tau / \Psi_1) \Rightarrow \Psi_2 \quad \text{FV}(M) \subseteq \rho \quad A' = [\text{wk}_{\Psi_2/\hat{\Psi}_2}]^{-1} A \text{ exists}}{\text{prune_ctx}_\rho(\tau, M / \Psi_1, x:A) \Rightarrow \Psi_2, x:A'}$	
$\Delta \vdash \text{prune}_\rho M \Rightarrow \Delta'; \eta$	Prune term M , returning $\Delta' \vdash \eta \Leftarrow \Delta$.
$\frac{v:B[\Psi_1] \in \Delta \quad \text{prune_ctx}_\rho(\tau / \Psi_1) \Rightarrow \Psi_2 \quad \Psi_2 \neq \Psi_1 \quad B' = [\text{wk}_{\Psi_2/\hat{\Psi}_2}]^{-1} B \quad \eta = \hat{\Psi}_1.v'[\text{wk}_{\Psi_2}]/v}{\Delta \vdash \text{prune}_\rho(v[\tau]) \Rightarrow \llbracket \eta \rrbracket(\Delta, v':B'[\Psi_2]); \eta}$	
$\frac{v:B[\Psi_1] \in \Delta \quad \text{prune_ctx}_\rho(\tau / \Psi_1) \Rightarrow \Psi_1}{\Delta \vdash \text{prune}_\rho(v[\tau]) \Rightarrow \Delta; \text{id}_\Delta}$	$\frac{x \in \rho}{\Delta \vdash \text{prune}_\rho x \Rightarrow \Delta; \text{id}_\Delta}$
$\frac{}{\Delta \vdash \text{prune}_\rho c \Rightarrow \Delta; \text{id}_\Delta}$	$\frac{\Delta \vdash \text{prune}_\rho R \Rightarrow \Delta_1; \eta_1 \quad \Delta_1 \vdash \text{prune}_\rho(\llbracket \eta_1 \rrbracket M) \Rightarrow \Delta_2; \eta_2}{\Delta \vdash \text{prune}_\rho(RM) \Rightarrow \Delta_2; \llbracket \eta_2 \rrbracket \eta_1}$
$\frac{\Delta \vdash \text{prune}_\rho M \Rightarrow \Delta'; \eta}{\Delta \vdash \text{prune}_\rho(\pi M) \Rightarrow \Delta'; \eta}$	$\frac{\Delta \vdash \text{prune}_{\rho,x} M \Rightarrow \Delta'; \eta}{\Delta \vdash \text{prune}_\rho(\lambda x. M) \Rightarrow \Delta'; \eta}$
$\frac{\Delta \vdash \text{prune}_\rho M \Rightarrow \Delta_1; \eta_1 \quad \Delta_1 \vdash \text{prune}_\rho(\llbracket \eta_1 \rrbracket N) \Rightarrow \Delta_2; \eta_2}{\Delta \vdash \text{prune}_\rho(M, N) \Rightarrow \Delta_2; \llbracket \eta_2 \rrbracket \eta_1}$	

Fig. 4. Pruning.

The second judgement is applied to subterms $v[\tau]$ of M to prune substitution τ with, say, domain Ψ_1 . We look at each term N in τ which substitutes for an $x:A$ of Ψ_1 . If N has a bad occurrence of a variable $y \notin \rho$, we discard the entry $x:A$ from the domain Ψ_1 , thus, effectively removing N from τ . If N has no occurrence of such an y we keep $x:A$. However, since we might have removed prior entries from Ψ_1 we need to ensure A is still well-formed, by validating that its free variables are bound in the pruned context. Pruning fails if N has a occurrence of a $y \notin \rho$ which is not bad, for instance flexible or rigid in an eliminateable

position. Examples:

- $\text{prune_ctx}_x(\text{c } x, \quad y \quad / \quad x':A, y':B) \Rightarrow x':A$ bad occurrence y
- $\text{prune_ctx}_y(\text{c } x, \quad u[y] \quad / \quad x':A, y':B) \Rightarrow y':B$ bad occurrence $\text{c } x$
- $\text{prune_ctx}_y(\lambda z. z x, y \quad / \quad x':A, y':B)$ fails occ. of x eliminateable
- $\text{prune_ctx}_y(u[x], \quad y \quad / \quad x':A, y':B)$ fails flexible occurrence $u[x]$

Pruning a term M with respect to ρ ensures that all rigid variables of M are in the range of ρ (see variable rule). Also, for each rigid occurrence of a meta-variable $v[\tau]$ in M we try to prune the substitution τ . If τ is already pruned, we leave v alone; otherwise, if the domain Ψ_1 of τ shrinks to Ψ_2 then we replace $v : B[\Psi_1]$ by a new meta-variable $v' : B[\Psi_2]$ with domain Ψ_2 . However, we need to ensure that the type B still makes sense in Ψ_2 ; otherwise, pruning fails. The last check is strengthening B from Ψ_1 to Ψ_2 and can be implemented as $[\text{wk}_{\Psi_2/\hat{\Psi}_2}]^{-1}B$ which exists whenever $\text{FV}(B) \subseteq \hat{\Psi}_2$.

Lemma 7 (Bad occurrences stay). *Let $\mathcal{D} :: \text{bad_occ}_x M$.*

1. *If $x \neq y$, then $\mathcal{D}' :: \text{bad_occ}_x [N/y]M$, and the derivation height of \mathcal{D}' is the same as the one of \mathcal{D} .*
2. *$x \in \text{FV}^{\text{rig}}(E[M])$.*
3. *If $y \in \text{FV}^{\text{rig}}(N)$ then $x \in \text{FV}^{\text{rig}}([M/y]N)$.*

Proof.

1. By induction on $\text{bad_occ}_x M$.
2. By induction on $\mathcal{D} :: \text{bad_occ}_x M$. In case $M = E'[x]$ we have trivially $x \in \text{FV}^{\text{rig}}(E[E'[x]])$. If $M = \lambda y. M'$ then either E is empty (trivial) or $E = E'[\bullet N]$. Then $E[M] = E'[[N/y]M']$, and by part 1 we get $\text{bad_occ}_x [N/y]M'$ with a smaller derivation height than \mathcal{D} . We conclude by induction hypothesis. If $M = (M_1, M_2)$ then either E is empty (trivial) or $E = E'[\text{fst } \bullet]$ or $E = E'[\text{snd } \bullet]$. In case of fst , $E[M] = E'[M_1]$ and we conclude by induction hypothesis. Case snd analogously.
3. By induction on β -normal form N . If N is a function $\lambda z. N'$ or a pair (N', N'') , proceed with N' . If $N = E[y]$ then apply part 2. Otherwise N is a neutral which does not have y as head variable, which implies that hereditarily substituting y will yield a neutral with the same spine form. If $N = \pi R$ then $y \in \text{FV}^{\text{rig}}(R)$ and $[M/y]N = \pi R'$ with $R' = [M/y]R$. By induction hypothesis, $x \in \text{FV}^{\text{rig}}(R') = \text{FV}^{\text{rig}}([M/y]N)$. If N is an application $R' N'$ then $[M/y]N = R'' N''$ with $R'' = [M/y]R'$ and $N'' = [M/y]N'$. Either $y \in \text{FV}^{\text{rig}}(R')$ or $y \in \text{FV}^{\text{rig}}(N')$ and we can conclude with the respective induction hypothesis, since $\text{FV}^{\text{rig}}(R'' N'') = \text{FV}^{\text{rig}}(R'') \cup \text{FV}^{\text{rig}}(N'')$. \square

Lemma 8 (Soundness and completeness of pruning).

1. *If $\Delta \vdash_{\mathcal{K}} \Psi_1 \text{ ctx}$ and $\text{prune_ctx}_\rho(\tau / \Psi_1) \Rightarrow \Psi_2$ then $\Delta \vdash_{\mathcal{K}} \Psi_2 \text{ ctx}$ and $\text{FV}([\tau]\text{wk}_{\Psi_2}) \subseteq \rho$. Additionally, if $x \in \Psi_1 \setminus \Psi_2$ then $\text{FV}^{\text{rig}}([\tau]x) \not\subseteq \rho$.*
2. *If $\Delta \vdash \text{prune}_\rho M \Rightarrow \Delta'$; η then $\Delta' \vdash_{\mathcal{K}} \eta \Leftarrow \Delta$ and $\text{FV}([\eta]M) \subseteq \rho$. Also, if θ solves $\Psi \vdash u[\rho] = M_0\{M\}^{\text{rig}} : C$ then there is some θ' such that $\theta = [[\theta']]$.*

Proof. Each by induction on the pruning derivation.

We detail 2., existence of θ' : Since θ is a solution of the constraint, $\llbracket \theta \rrbracket(u[\rho]) = \llbracket \llbracket \theta \rrbracket \rho \rrbracket(\theta(u)) = [\rho](\theta(u)) = \llbracket \theta \rrbracket M_0$, in particular $\text{FV}(\rho) = \rho \supseteq \text{FV}(\llbracket \theta \rrbracket M_0)$. This entails $\text{FV}(\llbracket \theta \rrbracket M) \subseteq \rho$.

Consider the interesting case $M = v[\tau]$ with $\text{prune_ctx}_\rho(\tau / \Psi_1) \Rightarrow \Psi_2$ and $\eta = \hat{\Psi}_1.v'[\text{wk}_{\Psi_2}]/v$. Let $N = \theta(v)$. We have $\text{FV}(\llbracket \llbracket \theta \rrbracket \tau \rrbracket N) \subseteq \rho$. If we can show $\text{FV}(N) \subseteq \hat{\Psi}_2$, then we can finish by setting $\theta' = \theta, \hat{\Psi}_2.N/v'$.

Assume now some $x \in \text{FV}(N)$ with $x \notin \hat{\Psi}_2$. By 1., $\text{FV}^{\text{rig}}([\tau]x) \not\subseteq \rho$, which entails that $\text{FV}(\llbracket \llbracket \theta \rrbracket \tau \rrbracket x) \not\subseteq \rho$. This is in contradiction to $\text{FV}(\llbracket \llbracket \theta \rrbracket \tau \rrbracket N) \subseteq \rho$. \square

In an implementation, we may combine pruning with inverse substitution and the occurs check. Since we already traverse the term M for pruning, we may also check whether $[\rho/\hat{\Phi}]^{-1}M$ exists and whether u occurs in M .

3.5 Unifying two identical existential variables

Any solution $\hat{\Phi}.N/u$ for a meta variable $u : A[\Phi]$ with constraint $u[\rho] = u[\xi]$ must fulfill $[\rho]N = [\xi]N$, which means that $[\rho]x = [\xi]x$ for all $x \in \text{FV}(N)$. This means that u can only depend on those of its variables in Φ that are mapped to the same term by ρ and ξ . Thus, we can substitute u by $\hat{\Phi}.v[\rho']$ where ρ' is the *intersection* of substitutions ρ and ξ . Similarly to context pruning, we obtain ρ' as $[\rho]\text{wk}_{\Phi'}$, which identical to $[\xi]\text{wk}_{\Phi'}$, where Φ' is a subcontext of Φ mentioning only the variables that have a common image under ρ and ξ . This process is given as judgement $\boxed{\rho \cap \xi : \Phi \Rightarrow \Phi'}$ with the following rules:

$$\frac{\cdot \cap \cdot \Rightarrow \cdot}{\frac{\rho \cap \xi : \Phi \Rightarrow \Phi'}{(\rho, y) \cap (\xi, y) : (\Phi, x:A) \Rightarrow (\Phi', x:A)} \quad \frac{\rho \cap \xi : \Phi \Rightarrow \Phi' \quad z \neq y}{(\rho, z) \cap (\xi, y) : (\Phi, x:A) \Rightarrow \Phi'}}$$

Lemma 9 (Soundness of intersection). *If $\Delta; \Psi \vdash_{\mathcal{K}} \rho, \xi \Leftarrow \Phi$ and $\rho \cap \xi : \Phi \Rightarrow \Phi'$, then $\Delta \vdash_{\mathcal{K}} \Phi' \text{ ctx}$ and $\Delta; \Phi \vdash_{\mathcal{K}} \text{wk}_{\Phi'} \Leftarrow \Phi'$ and $z \in \text{dom}(\Phi')$ iff $\rho(z) = \xi(z)$.*

Proof. By structural induction on the first derivation.

We consider the interesting case, where we actually retain a declaration $x:A$ because the two substitutions map x to the same variable y .

$$\frac{\rho \cap \xi : \Phi \Rightarrow \Phi'}{(\rho, y) \cap (\xi, y) : (\Phi, x:A) \Rightarrow \Phi', x:A}$$

The main challenge is to show that type A is well-formed even in the subcontext Φ' of Φ , which is the case if $\text{FV}(A) \subseteq \text{dom}(\Phi')$.

From the assumption $\Delta; \Psi \vdash_{\mathcal{K}} (\rho, y) \Leftarrow (\Phi, x:A)$ we get $\Delta; \Psi \vdash_{\mathcal{K}} \rho \Leftarrow \Phi$ and $\Psi(y) =_{\mathcal{K}} [\rho]A$ by inversion, and the same for ξ . This allows us to apply the induction hypothesis, which yields $\Delta \vdash_{\mathcal{K}} \Phi' \text{ ctx}$ and $\Delta; \Phi \vdash_{\mathcal{K}} \text{wk}_{\Phi'} \Leftarrow \Phi'$ and $\text{dom}(\Phi') = \{z \mid \rho(z) = \xi(z)\}$. Since $[\rho]A =_{\mathcal{K}} \Psi(y) =_{\mathcal{K}} [\xi]A$, we have for all

$z \in \text{FV}(A)$ that $\rho(z) =_{\mathcal{K}} \xi(z)$. Since this is an equation between variables, the constraints \mathcal{K} on meta-variables do not matter, and we even have $\rho(z) = \xi(z)$. Yet this means that $z \in \text{dom}(\Phi')$, thus, we obtain $\Delta; \Phi' \vdash_{\mathcal{K}} A \Leftarrow \text{type}$ by strengthening. This entails $\Delta \vdash_{\mathcal{K}} (\Phi', x:A) \text{ctx}$, hence, trivially, $\Delta; (\Phi, x:A) \vdash_{\mathcal{K}} \text{wk}_{\Phi', x:A} \Leftarrow (\Phi', x:A)$. Finally, $\text{dom}(\Phi', x:A) = \{z \mid \rho(z) = \xi(z)\}$ follows from this property for Φ' and the fact $\rho(x) = y = \xi(x)$. \square

Let us now reconsider the rule ‘‘Same meta-variable’’.

$$\begin{array}{c} \Delta \Vdash \mathcal{K} \wedge \Psi \vdash u[\rho] = u[\xi] : C \quad \mapsto \quad \Delta, v:A[\Phi'] \Vdash \mathcal{K} + \Phi \vdash u \Leftarrow v[\text{wk}_{\Phi'}] : A \\ u:A[\Phi] \in \Delta \quad \quad \quad \text{if } \rho \cap \xi : \Phi \Rightarrow \Phi' \end{array}$$

We have just shown that the resulting Φ' of computing the intersection of ρ and ξ , is indeed well-formed. We can also justify why in the unification rule itself the type A of two existential variables must be well-typed in the pruned context Φ' . Recall that by typing invariant, we know that $\Delta; \Psi \vdash_{\mathcal{K}} \rho \Leftarrow \Phi$ and $\Delta; \Psi \vdash_{\mathcal{K}} \xi \Leftarrow \Phi$ and $[\rho]A =_{\mathcal{K}} [\xi]A$. But this means that A can only depend on the variables mapped to the same term by ρ and ξ . Since Φ_0 is exactly the context which captures those shared variables, A must also be well-typed in Φ_0 modulo \mathcal{K} .

Note that intersection cannot easily be extended beyond variable substitutions. For instance, consider the following problem:

$$(\lambda f. f \circ f, \mathbf{g}, N) \cap (\lambda f. f, \mathbf{g} \circ \mathbf{g}, N) : (x : A^4, y : A^2, z : P(xy)) \Rightarrow ?$$

where $N : P(\mathbf{g} \circ \mathbf{g})$ and $A^4 := (A \rightarrow A) \rightarrow A \rightarrow A$ and $A^2 := A \rightarrow A$, and $f \circ g := \lambda n. f(gn)$ abbreviates function composition. We cannot remove x and y from the context only because the two substitutions differ at these positions. The resulting context $(z : P(xy))$ would be ill-typed. The difference to the variable case is that substitution is no longer injective, i.e., we can have $[\sigma]M = [\tau]M$ but $\sigma(x) \neq \tau(x)$ for some $x \in \text{FV}(M)$.

4 Correctness

Theorem 1 (Termination). *The algorithm terminates and results in one of the following states:*

- *A solved state where only assignments $\Psi \vdash u \Leftarrow M : A$ remain.*
- *A stuck state, i.e., no transition rule applies.*
- *Failure \perp .*

Proof. Let the size $|M|$ of a term be as usual the number of nodes and leaves in its tree representation, with the exception that we count λ -nodes twice. This modification has the effect that $|\lambda x.M| + |R| > |M| + |Rx|$, hence, an η -expanding decomposition step also decreases the sum of the sizes of the involved terms [7]. We define the size $|A[\Phi]|$ of a type A in context Φ by $|P[\Phi]| = 1 + \sum_{A \in \Phi} |A|$, $|(\Pi x:A. B)[\Phi]| = 1 + |B[\Phi, x:A]|$ and $|(\Sigma x:A. B)[\Phi]| = 1 + |A[\Phi]| + |B[\Phi]|$. The

size of a type can then be obtained as $|A| = |A|$ and the size of a context as $|\Phi| = \sum_{A \in \Phi} |A|$. The purpose of this measure is to give Σ -types a large weight that can “pay” for *flattening*.

Let the weight of a solved constraint be 0, whereas the weight $|K|$ for a constraint $\Psi \vdash M = M' : C$ be the ordinal $(|M| + |M'|)\omega + |\Psi|$ if a decomposition step can be applied, and simply $|\Psi|$ else. Similarly, let the weight of constraint $\Phi \mid R:A \vdash E = E'$ be $(|E| + |E'|)\omega + |\Psi|$. Finally, let the weight $|\Delta \Vdash \mathcal{K}|$ of a unification problem be the ordinal

$$\sum_{u:A[\Phi] \in \Delta \text{ active}} |A[\Phi]|\omega^2 + \sum_{K \in \mathcal{K}} |K|.$$

By inspection of the transition rules we can show that each unification step reduces the weight of the unification problem. \square

4.1 Solutions to unification

A solution to a set of equations \mathcal{K} is a meta-substitution θ for all the meta-variables in Δ s.t. $\Delta' \vdash \theta \Leftarrow \Delta$ and

1. for every $\Psi \vdash u \Leftarrow M : A$ in \mathcal{K} we have $\hat{\Psi}.M/u \in \theta$,
2. for all equations $\Psi \vdash M = N : A$ in \mathcal{K} , we have $\llbracket \theta \rrbracket M = \llbracket \theta \rrbracket N$.

A *ground* solution to a set of equations \mathcal{K} can be obtained from a solution to \mathcal{K} by applying a grounding meta-substitution θ' where $\cdot \vdash \theta' \Leftarrow \Delta'$ to the solution θ . We write $\theta \in \text{Sol}(\Delta \Vdash \mathcal{K})$ for a ground solution to the constraints \mathcal{K} .

Before we prove that transitions preserve solutions, we first prove that there always exists a meta-substitution relating the original meta-variable context Δ_0 to the meta-variable context Δ_1 we transition to. It is useful to state this property in isolation, although it is also folded into Theorem 2.

Lemma 10. *If $\Delta_0 \Vdash \mathcal{K}_0 \mapsto \Delta_1 \Vdash \mathcal{K}_1$ then there exists a meta-substitution θ s.t. $\Delta_1 \vdash_{\mathcal{K}_1} \theta \Leftarrow \Delta_0$.*

Proof. By case analysis on the unification steps. \square

We also observe that if we start in a state $\Delta_0 \Vdash K_0$ and transition to a state $\Delta_1 \Vdash K_1$ the meta-variable context strictly grows, i.e., $\text{dom}(\Delta_0) \subseteq \text{dom}(\Delta_1)$. We subsequently show that if we have a solution for $\Delta_0 \Vdash K_0$, then transitioning to a new state $\Delta_1 \Vdash K_1$ will not add any additional solutions nor will it destroy some solution we may already have. In other words, any additional constraints which may be added in $\Delta_1 \Vdash K_1$ are consistent with the already existing solution.

Theorem 2 (Transitions preserve solutions). *Let $\Delta_0 \Vdash \mathcal{K}_0 \mapsto \Delta_1 \Vdash \mathcal{K}_1$.*

1. *If $\theta_0 \in \text{Sol}(\Delta_0 \Vdash \mathcal{K}_0)$ then there exists a meta-substitution θ' s.t. $\Delta_1 \vdash_{\mathcal{K}_1} \theta' \Leftarrow \Delta_0$ and a solution $\theta_1 \in \text{Sol}(\Delta_1 \Vdash \mathcal{K}_1)$ such that $\llbracket \theta_1 \rrbracket \theta' = \theta_0$.*
2. *If $\theta_1 \in \text{Sol}(\Delta_1 \Vdash \mathcal{K}_1)$ then $\llbracket \theta_1 \rrbracket \text{wk}_{\Delta_0} \in \text{Sol}(\Delta_0 \Vdash \mathcal{K}_0)$.*

Proof. Proof by case analysis on the transitions. We only show the cases to prove that we are forward closed (statement 1). The second statement, that unification steps are backwards closed, is obvious: by if we transition from $\Delta_0 \Vdash \mathcal{K}_0$ to $\Delta_1 \Vdash \mathcal{K}_1$ then there exists a meta-substitution $\Delta_1 \vdash_{\mathcal{K}_1} \theta \Leftarrow \Delta_0$; hence, by composition $\llbracket \theta_1 \rrbracket \theta$ is ground and is a solution for $\Delta_0 \Vdash \mathcal{K}_0$.

In all the following cases, θ_0 is a solution for \mathcal{K}_0 .

Decomposition. Since Δ_0 does not change in any of the decomposition rules, the solution is almost trivially preserved; for the η -contraction rules, we simply observe that equality is always modulo η . For the eliminating projections transition, we use the substitution lemma and observe that meta-substitutions and ordinary substitutions commute.

Lowering. Let $u : (\Pi x:A.B)[\Phi] \in \Delta$ active and

$$\Delta \Vdash \mathcal{K} \quad \mapsto \quad (\Delta, v:B[\Phi, x:A] \Vdash \mathcal{K}) + (\Phi \vdash u \Leftarrow \lambda x.v : \Pi x:A. B)$$

Let $\hat{\Phi}.M := \theta_0(u)$ and observe $\cdot; \llbracket \theta_0 \rrbracket \Phi \vdash M \Leftarrow \llbracket \theta_0 \rrbracket (\Pi x:A.B)$. By inversion on typing, $M = \lambda x.N$ and $\cdot; \llbracket \theta_0 \rrbracket (\Phi, x:A) \vdash N \Leftarrow \llbracket \theta_0 \rrbracket B$. Hence, $(\hat{\Phi}, x).N$ is a solution for v . Choose $\theta' = \text{wk}_{\Delta_0}$ and $\theta_1 = \theta, (\hat{\Phi}, x).N/v$. Since v is a new meta variables, θ_1 is still a solution for the old state $\Delta \Vdash \mathcal{K}$.

The case for lowering Σ -types is similar.

Flattening. Using the substitution lemma, solutions are preserved.

Pruning. θ_0 is a solution for \mathcal{K}_0 and the active meta-variable $u:A[\Phi] \in \Delta_0$. Hence, $\hat{\Phi}.M/u \in \theta_0$. Moreover, if we have the constraint $\Psi \vdash u[\rho] = N : B$, we have $[\rho]M = \llbracket \theta_0 \rrbracket N$. By previous soundness lemma for pruning, $\Delta_p \vdash \theta_p \Leftarrow \Delta_0$ and there exists a θ' s.t. $\llbracket \theta' \rrbracket \theta_p = \theta$.

Same meta-variable. θ_0 is a solution for \mathcal{K}_0 and the active meta-variable $u:A[\Phi] \in \Delta_0$. Hence, $\hat{\Phi}.M/u \in \theta_0$ and $\cdot; \llbracket \theta_0 \rrbracket (\Phi) \vdash M \Leftarrow \llbracket \theta_0 \rrbracket A$. Moreover, $[\rho]M = [\xi]M$. Therefore, $\text{FV}([\rho]M) = \text{FV}([\xi]M)$ and $\hat{\Phi}_0$ contains exactly those meta-variables which are shared among ρ and ξ by definition of $\rho \cap \xi$; moreover, we must have $\llbracket \theta_0 \rrbracket ([\rho]A) = \llbracket \theta_0 \rrbracket ([\xi]A)$. Since meta-substitutions commute with variable substitutions, we have $[\rho](\llbracket \theta_0 \rrbracket A) = [\xi](\llbracket \theta_0 \rrbracket A)$, hence $\text{FV}(M) = \hat{\Phi}_0$ and $\text{FV}(\llbracket \theta_0 \rrbracket A) = \hat{\Phi}_0$ and $\cdot; \llbracket \theta_0 \rrbracket (\Phi_0) \vdash M \Leftarrow \llbracket \theta_0 \rrbracket A$; choosing id_{Δ_0} for θ' and for $\theta_1 = \theta, \hat{\Phi}_0.M/v$ solutions are preserved.

Solving. θ_0 is a solution for \mathcal{K}_0 and the active meta-variable $u:A[\Phi] \in \Delta_0$. Hence, $\hat{\Phi}.N/u \in \theta_0$. Therefore, we have $[\rho]N = \llbracket \theta \rrbracket M$. By completeness of inverse substitution, we know $N = [\rho/\Phi]^{-1}(\llbracket \theta \rrbracket M)$. By assumption we also know $[\rho/\Phi]^{-1}M = M'$ exists. Therefore, by lemma that inverse and meta-substitution commute, we have $N = \llbracket \theta \rrbracket ([\rho/\Phi]^{-1}M) = \llbracket \theta \rrbracket M'$. Therefore, the solution θ_0 is preserved. \square

4.2 Transitions preserve types

Our goal is to prove that if we start with a well-typed unification problem our transitions preserve the type, i.e., we can never reach an ill-typed state and hence, we cannot generate a solution which may contain an ill-typed term.

Lemma 11 (Equality modulo is preserved by transitions).

If $\Delta_0 \Vdash \mathcal{K}_0 \mapsto \Delta_1 \Vdash \mathcal{K}_1$ and $A =_{\mathcal{K}_0} B$, then $A =_{\mathcal{K}_1} B$.

Proof. Let θ be a solution for \mathcal{K}_0 ; by assumption, we have that $\llbracket \theta \rrbracket A = \llbracket \theta \rrbracket B$. By theorem 2, transitions preserve solutions, we know θ is also a solution for \mathcal{K}_1 , and therefore $A =_{\mathcal{K}_1} B$. \square

In the statement below it is again important to note that the meta-context strictly grows, i.e., $\Delta_0 \subseteq \Delta_1$ and that there always exists a meta-substitution θ which maps Δ_0 to Δ_1 . Moreover, since transitions preserve solutions, if we have a solution for \mathcal{K}_0 there exists a solution for \mathcal{K}_1 .

Lemma 12 (Transitions preserve typing). *Let $\Delta_0 \Vdash \mathcal{K}_0 \mapsto \Delta_1 \Vdash \mathcal{K}_1$ and $\Delta_1 \vdash_{\mathcal{K}_1} \theta \Leftarrow \Delta_0$.*

1. *If $\Delta_0; \Psi \vdash_{\mathcal{K}_0} M \Leftarrow A$ then $\Delta_1; \llbracket \theta \rrbracket \Psi \vdash_{\mathcal{K}_1} \llbracket \theta \rrbracket M \Leftarrow \llbracket \theta \rrbracket A$.*
2. *If $\Delta_0; \Psi \vdash_{\mathcal{K}_0} R \Rightarrow A$ then $\Delta_1; \llbracket \theta \rrbracket \Psi \vdash_{\mathcal{K}_1} \llbracket \theta \rrbracket R \Rightarrow A'$ and $\llbracket \theta \rrbracket A =_{\mathcal{K}_1} A'$.*

Proof. By induction on the derivation of $\Delta_0; \Psi \vdash_{\mathcal{K}_0} J$. Most cases are by inversion, appeal to induction hypothesis, and re-assembling the result. The most interesting case is transitioning between normal and neutral terms. Here we use the previous lemma on “Equality modulo preserved by transitions”. \square

Next, we define when a set of equations which constitute a unification problem are well-formed using the judgment $\Delta_0 \Vdash_{\mathcal{K}_0} \mathcal{K} \text{ wf}$, which states that each equation $\Psi \vdash M = N : A$ must be well-typed modulo the equations in \mathcal{K}_0 , i.e., $\Delta_0; \Psi \vdash_{\mathcal{K}_0} M \Leftarrow A$ and $\Delta_0; \Psi \vdash_{\mathcal{K}_0} N \Leftarrow A$. We simply write $\Delta_0 \Vdash \mathcal{K} \text{ wf}$ to mean $\Delta_0 \Vdash_{\mathcal{K}} \mathcal{K} \text{ wf}$.

Lemma 13 (Equations remain well-formed under meta-substitutions).

If $\Delta_0 \Vdash \mathcal{K} \text{ wf}$ and $\Delta_1 \vdash_{\llbracket \theta \rrbracket \mathcal{K}} \theta \Leftarrow \Delta_0$ then $\Delta_1 \Vdash \llbracket \theta \rrbracket \mathcal{K} \text{ wf}$.

Proof. By assumption $\Delta_0 \Vdash \mathcal{K} \text{ wf}$. By definition, for every constraint $\Psi \vdash M = N : A \in \mathcal{K}$, we have $\Delta_0; \Psi \vdash_{\mathcal{K}} M \Leftarrow A$ and $\Delta_0; \Psi \vdash_{\mathcal{K}} N \Leftarrow A$. By meta-substitution principle modulo (lemma 3), we know $\Delta_1; \llbracket \theta \rrbracket \Psi \vdash_{\llbracket \theta \rrbracket \mathcal{K}} \llbracket \theta \rrbracket M \Leftarrow \llbracket \theta \rrbracket A$ and $\Delta_1; \llbracket \theta \rrbracket \Psi \vdash_{\llbracket \theta \rrbracket \mathcal{K}} \llbracket \theta \rrbracket N \Leftarrow \llbracket \theta \rrbracket A$, and hence $\Delta_1 \vdash_{\llbracket \theta \rrbracket \mathcal{K}} \llbracket \theta \rrbracket \mathcal{K} \text{ wf}$. \square

Lemma 14 (Well-formedness of equations is preserved by transitions).

If $\Delta_0 \Vdash \mathcal{K}_0 \mapsto \Delta_1 \Vdash \mathcal{K}_1$ and $\Delta_0 \Vdash_{\mathcal{K}_0} \mathcal{K} \text{ wf}$ then $\Delta_1 \vdash_{\mathcal{K}_1} \mathcal{K} \text{ wf}$.

Proof. By assumption $\Delta \Vdash_{\mathcal{K}_0} \mathcal{K} \text{ wf}$, we know that for each $\Psi \vdash M = N : A \in \mathcal{K}$, $\Delta; \Psi \vdash_{\mathcal{K}_0} M \Leftarrow A$ and $\Delta; \Psi \vdash_{\mathcal{K}_0} N \Leftarrow A$. By lemma 12, typing is preserved by transitions, we know that $\Delta; \Psi \vdash_{\mathcal{K}_1} M \Leftarrow A$ and $\Delta; \Psi \vdash_{\mathcal{K}_1} N \Leftarrow A$. Therefore $\Delta \Vdash_{\mathcal{K}_1} \mathcal{K} \text{ wf}$. \square

Theorem 3 (Unification preserves types).

If $\Delta_0 \Vdash \mathcal{K}_0 \text{ wf}$ and $\Delta_0 \Vdash \mathcal{K}_0 \mapsto \Delta_1 \Vdash \mathcal{K}_1$ then $\Delta_1 \Vdash \mathcal{K}_1 \text{ wf}$.

Proof. By case analysis on the transition rules and lemma 10.

Decomposition rules. We consider the decomposition rule for pairs. Let \mathcal{K}_0 be the set of equations which contains $\Psi \vdash (M_1, M_2) = (N_1, N_2) : \Sigma x:A.B$. By assumption we have $\Delta_0; \Psi \vdash_{\mathcal{K}_0} (M_1, M_2) \Leftarrow \Sigma x:A.B$. By inversion, we have $\Delta_0; \Psi \vdash_{\mathcal{K}_0} M_1 \Leftarrow A$ and $\Delta_0; \Psi \vdash_{\mathcal{K}_0} M_2 \Leftarrow [M_1/x]_A(B)$. By assumption we have $\Delta_0; \Psi \vdash_{\mathcal{K}_0} (N_1, N_2) \Leftarrow \Sigma x:A.B$. By inversion, we have $\Delta_0; \Psi \vdash_{\mathcal{K}_0} N_1 \Leftarrow A$ and $\Delta_0; \Psi \vdash_{\mathcal{K}_0} N_2 \Leftarrow [N_1/x]_A(B)$. Let $K_1 = \mathcal{K}_0 \wedge \Psi \vdash M_1 = N_1 : A$. Then $\Delta_0 \Vdash K_1$ wf. Moreover, $\Delta_0; \Psi \vdash_{K_1} N_2 \Leftarrow [M_2/x]_A(B)$. Hence, $\Psi \vdash M_2 = N_2 : [M_1/x]_A(B)$ is well-formed and $\Delta_0 \Vdash K_2$ wf where we replace the constraint $\Psi \vdash (M_1, M_2) : \Sigma x:A.B$ with $\Psi \vdash M_1 = N_1 : A \wedge \Psi \vdash M_2 = N_2 : [M_1/x]_A(B)$ in \mathcal{K} .

Next, we consider the decomposition rules for evaluation contexts. Let \mathcal{K}_0 be the set of equations containing $\Psi \mid R : \Pi x:A.B \vdash E[\bullet M] = E'[\bullet M']$. By assumption this constraint is well-typed, and hence $\Delta_0; \Psi \vdash_{\mathcal{K}_0} R \Rightarrow \Pi x:A.B$ and $\Delta_0; \Psi \vdash_{\mathcal{K}_0} R M \Rightarrow B_2$ and $\Delta_0; \Psi \vdash_{\mathcal{K}_0} R M' \Rightarrow B_1$ where $B_1 =_{\mathcal{K}_0} B_2$. In addition $\Delta_0; \Psi \vdash_{\mathcal{K}_0} M \Leftarrow A$ and $\Delta_0; \Psi \vdash_{\mathcal{K}_0} M' \Leftarrow A$.

Let $\mathcal{K}_1 = \mathcal{K}_0 \wedge \Psi \vdash M = M' : A$. Clearly, $\Delta_0 \Vdash \mathcal{K}_1$ wf. Moreover, since the evaluation $E[R M]$ and $E'[R M']$ are well-typed modulo \mathcal{K}_0 and the fact that $\Psi \vdash M = M' : A$, we have also that $E'[R M]$ is well-typed modulo \mathcal{K}_1 and $\Psi \mid R M : [M/x]B \vdash E = E'$ is well-typed in Δ_0 modulo \mathcal{K}_1 . Therefore, we have $\Delta_0 \Vdash \mathcal{K}_2$ wf where $\mathcal{K}_2 = \mathcal{K}_1 \wedge \Psi \mid R M : [M/x]B \vdash E = E'$.

Pruning rule. Let \mathcal{K}_0 be the set of equations containing $\Psi \vdash u[\rho] = M : A$. By assumption, we know that $\Delta_0; \Psi \vdash_{\mathcal{K}_0} u[\rho] \Leftarrow A$ and $\Delta_0; \Psi \vdash_{\mathcal{K}_0} M \Leftarrow A$.

By soundness of pruning (lemma 8), we know that $\Delta_1 \vdash_{\mathcal{K}_0} \eta \Leftarrow \Delta_0$. By lemma 13, we know that $\Delta_1 \Vdash \llbracket \eta \rrbracket (\mathcal{K}_0)$ wf.

Intersections. Let \mathcal{K}_0 be the set of equations containing $\Psi \vdash u[\rho] = u[\xi] : C$. By assumption, we know that $\Delta_0; \Psi \vdash_{\mathcal{K}_0} u[\rho] \Leftarrow C$ and $\Delta_0; \Psi \vdash_{\mathcal{K}_0} u[\xi] \Leftarrow C$. Let $u : A[\Phi] \in \Delta$. By inversion, we have $[\xi]A =_{\mathcal{K}_0} C =_{\mathcal{K}_0} [\rho]A$. This means $\text{FV}([\xi]A) = \text{FV}([\rho]A)$ and by definition of $\rho \cap \xi : \Phi \Rightarrow \Phi_0$, the context Φ_0 will contain exactly those variables shared in ξ and ρ . By soundness lemma 9, we have $\Delta_0 \vdash_{\mathcal{K}_0} \Phi_0$ ctx. Therefore, $\Delta_0; \Phi_0 \vdash_{\mathcal{K}_0} A \Leftarrow \text{type}$ and $(\Delta_0, v:A[\Phi_0])$ mctx. By typing rules, we have $(\Delta_0, v:A[\Phi_0]); \Phi \vdash_{\mathcal{K}_0} u \Leftarrow A$ and $(\Delta_0, v:A[\Phi_0]); \Phi \vdash_{\mathcal{K}_0} v[\text{wk}_{\Phi_0}] \Leftarrow A$. Hence, $\Phi \vdash u \Leftarrow v[\text{wk}_{\Phi_0}] : A$ is well-typed in Δ_0 modulo \mathcal{K}_0 . Hence, $\theta = \hat{\Phi}.v[\text{wk}_{\Phi_0}]/u$ is a well-formed meta-substitution. By lemma 3, we have $\llbracket \theta \rrbracket \Delta \Vdash (\llbracket \theta \rrbracket \mathcal{K}_0 \wedge \llbracket \theta \rrbracket \Phi \vdash u \Leftarrow M : \llbracket \theta \rrbracket A)$ wf.

Solving. Let \mathcal{K}_0 be the set of equations containing $\Psi \vdash u[\rho] = M : C$. By assumption $M' = [\rho/\hat{\Phi}]^{-1}M$ exists and $u:A[\Phi] \in \Delta$. Since $\Delta_0 \Vdash \mathcal{K}_0$ wf, we also have $\Delta_0; \Psi \vdash_{\mathcal{K}_0} M \Leftarrow C$ and $\Delta_0; \Psi \vdash_{\mathcal{K}_0} u[\rho] \Leftarrow C$. By inversion, we have $C =_{\mathcal{K}_0} [\rho]A$. By assumption we have $\Delta_0; \Phi \vdash_{\mathcal{K}_0} M' \Leftarrow A$. Hence, $\theta = \hat{\Phi}.M'/u$ is a well-formed meta-substitution and by lemma 3, we have $\llbracket \theta \rrbracket \Delta_0 \Vdash \llbracket \theta \rrbracket (\mathcal{K}_0 \wedge \Phi \vdash u \Leftarrow M' : A)$ wf. \square

5 Conclusion

We have presented a constraint-based unification algorithm which solves higher-order patterns dynamically and showed its correctness. There are several key aspects of our algorithm: First, we define pruning formally and show soundness in the dependently typed case. Our pruning operation differs from previous formulations in how it treats non-patterns which may occur in the term to be pruned: if it encounters a non-pattern term M where $FV(M) \subseteq \rho$, then pruning may succeed; otherwise it fails. This strategy avoids non-termination problems present in previous formulations [3], but is also less ambitious than the algorithm proposed by Reed [17]. We have extended higher-order pattern unification to handle Σ -types; this has been an open problem so far, yet it is of practical relevance:

1. In LF-based systems such as Beluga, Twelf or Delphin, a limited form of Σ -types arises due to context blocks: Σ -types are used to introduce several assumptions simultaneously. For Beluga, the second author has implemented the flattening of context blocks and it works well in type reconstruction.
2. In dependently typed languages such as Agda, Σ -types, or, more generally, record types, are commonly used but unification has original not been adapted to records. McBride [9, p.6] gives a practical example where the unification problem $T(\text{fst } \gamma)(\text{snd } \gamma) = T' \gamma$ appears. Using the techniques presented in this paper, the first author has extended Agda's unification algorithm to solve those kinds of problems.

Correctness of our unification constraint solver is proved using *typing modulo* [17]. This is possible since we have no constraints on the type level and we are dealing with terms whose normalization via hereditary substitutions can be defined by recursion on their type. Even in the presence of unsolvable constraints, which lead to ill-typed terms, normalization is terminating. This does not scale to Agda which has large eliminations and unification on the type level; there, ill-typed terms may lead to divergence of type reconstruction. A solution has been described by Norell [12]: unsolved constraints block normalization, thus guaranteeing termination of the unification algorithm. The idea has been implemented in Agda 2 and been extended to Σ -types and the unification rules described in this article.

Acknowledgments. We thank Jason Reed for his insightful work and his explanations given via email. Thanks to Conor McBride for in-depth comments on this text and his suggestions for improvement. We also acknowledge the anonymous referees who have given constructive critique on a previous version of this article.

References

1. Andreas Abel and Brigitte Pientka. Higher-order dynamic pattern unification for dependent types and records. In C.-H. Luke Ong, editor, *Typed Lambda Calculi*

- and Applications - 10th International Conference, TLCA 2011, Novi Sad, Serbia, June 1-3, 2011, Proceedings*, volume 6690 of *Lecture Notes in Computer Science*, pages 10–26. Springer, 2011.
2. Iliano Cervesato and Frank Pfenning. A linear spine calculus. *Journal of Logic and Computation*, 13(5):639–688, 2003.
 3. Gilles Dowek, Thérèse Hardin, Claude Kirchner, and Frank Pfenning. Unification via explicit substitutions: The case of higher-order patterns. In *Joint International Conference Logic Programming*, pages 259–273. MIT Press, 1996.
 4. Dominic Duggan. Unification with extended patterns. *Theoretical Computer Science*, 206(1-2):1–50, 1998.
 5. Conal Elliott. *Extensions and Applications of Higher-Order Unification*. PhD thesis, School of Computer Science, Carnegie Mellon University, 1990.
 6. Roland Fettig and Bernd Löchner. Unification of higher-order patterns in a simply typed lambda-calculus with finite products and terminal type. In *7th International Conference on Rewriting Techniques and Applications (RTA'96)*, LNCS 1103, pages 347–361. Springer, 1996.
 7. Healdene Goguen. Justifying algorithms for $\beta\eta$ conversion. In *8th Int. Conf. on Foundations of Software Science and Computational Structures (FoSSaCS'05)*, LNCS 3441, pages 410–424. Springer, 2005.
 8. W. D. Goldfarb. The undecidability of the second-order unification problem. *Theoretical Computer Science*, 13:225–230, 1981.
 9. Conor McBride. Outrageous but meaningful coincidences: Dependent type-safe syntax and evaluation. In *ACM SIGPLAN Workshop on Genetic Programming (WGP'10)*, pages 1–12. ACM, 2010.
 10. Dale Miller. Unification of simply typed lambda-terms as logic programming. In *Eighth International Logic Programming Conference*, pages 255–269. MIT Press, 1991.
 11. Aleksandar Nanevski, Frank Pfenning, and Brigitte Pientka. Contextual modal type theory. *ACM Transactions on Computational Logic*, 9(3):1–49, 2008.
 12. Ulf Norell. *Towards a Practical Programming Language Based on Dependent Type Theory*. PhD thesis, Department of Computer Science and Engineering, Chalmers University of Technology, Göteborg, Sweden, September 2007.
 13. Frank Pfenning and Carsten Schürmann. System description: Twelf — a meta-logical framework for deductive systems. In *16th International Conference on Automated Deduction (CADE-16)*, LNAI 1632, pages 202–206. Springer, 1999.
 14. Brigitte Pientka. *Tabled higher-order logic programming*. PhD thesis, Department of Computer Science, Carnegie Mellon University, 2003. CMU-CS-03-185.
 15. Brigitte Pientka and Joshua Dunfield. Beluga: a framework for programming and reasoning with deductive systems (System Description). In *5th International Joint Conference on Automated Reasoning (IJCAR'10)*, LNAI 6173, pages 15–21, 2010.
 16. Adam Poswolsky and Carsten Schürmann. System description: Delphin—a functional programming language for deductive systems. In *International Workshop on Logical Frameworks and Meta-Languages: Theory and Practice (LFMTP'08)*, ENTCS 228, pages 135–141. Elsevier, 2009.
 17. Jason Reed. Higher-order constraint simplification in dependent type theory. In *International Workshop on Logical Frameworks and Meta-Languages: Theory and Practice (LFMTP'09)*, 2009.
 18. Jason Reed. *A Hybrid Logical Framework*. PhD thesis, School of Computer Science, Carnegie Mellon University, 2009.

19. Anders Schack-Nielsen and Carsten Schürmann. Pattern unification for the lambda calculus with linear and affine types. In Karl Crary and Marino Miculan, editors, *5th International Workshop on Logical Frameworks and Meta-languages: Theory and Practice (LFMTP 2010), Edinburgh, Scotland, UK, July 14, 2010*, volume 34 of *Electronic Proceedings in Theoretical Computer Science*, pages 101–116, 2010.
20. Kevin Watkins, Iliano Cervesato, Frank Pfenning, and David Walker. A concurrent logical framework I: Judgements and properties. Technical report, School of Computer Science, Carnegie Mellon University, Pittsburgh, 2003.

A Hereditary substitution

Normal forms are maintained through the use of hereditary substitution, written as $[N/x]_A(B)$ to guarantee that when we substitute term N which has type A for the variable x in the type B , we obtain a type B' which is in normal form. Hereditary substitutions continue to substitute, if a redex is created; for example, when replacing naively x by $\lambda y.c y$ in the object $x z$, we would obtain $(\lambda y.c y) z$ which is not in normal form and hence not a valid term in our grammar. Hereditary substitutions continue to substitute z for y in $c y$ to obtain $c z$ as a final result.

Hereditary substitution can be defined recursively considering the term to which the substitution operation is applied and the type of the object which is being substituted. We define the hereditary substitution operations for normal object, neutral objects and substitutions. The hereditary substitution operations will be defined by nested induction, first on the structure of the type A and second on the structure of the objects N , R , and σ . In other words, we either go to a smaller type, in which case the objects themselves can become larger, or the type remains the same and the objects become smaller. We write $A \leq B$ and $A < B$ if A occurs in B (as a proper sub-expression in the latter case)⁴. Hereditary substitution is defined in Figure 5. For an in depth discussion, we refer the reader to Nanevski et al. [11].

If the original term is not well-typed, a hereditary substitution, though terminating, cannot always return a meaningful term. We formalize this as failure to return a result. However, on well-typed terms, hereditary substitution will always return well-typed terms. The definition for single hereditary substitutions can be easily extended to simultaneous substitutions substitution written as $[\sigma]_\Psi(M)$. We annotate the substitution with the sub-script Ψ for two reasons. First, σ itself does not carry its domain and hence we will look up the instantiation for a variable x in σ/Ψ . Second, we rely on the type of x in the context Ψ to guarantee that applying σ to an object terminates. Either we apply σ to sub-expressions or the type of the object we substitute will be smaller. Subsequently, we often omit the typing subscript at the substitution operation for better readability.

⁴ To ensure termination, it suffices to rely on type approximations of the dependent type; we leave this out from the discussion.

Normal terms / types		
$[M/x]_A(\Pi y:B_1.B_2)$	$= \Pi y:B'_1.B'_2$	where $B'_1 = [M/x]_A(B_1)$ and $B'_2 = [M/x]_A(B_2)$, $y \notin \text{FV}(M)$, and $y \neq x$
$[M/x]_A(\Sigma y:B_1.B_2)$	$= \Sigma y:B'_1.B'_2$	where $B'_1 = [M/x]_A(B_1)$ and $B'_2 = [M/x]_A(B_2)$, $y \notin \text{FV}(M)$, and $y \neq x$
$[M/x]_A(\text{type})$	$= \text{type}$	
$[M/x]_A(\lambda y.N)$	$= \lambda y.N'$	where $[M/x]_A(N) = N'$, $y \notin \text{FV}(M)$, and $y \neq x$
$[M/x]_A(N_1, N_2)$	$= (N'_1, N'_2)$	where $[M/x]_A(N_1) = N'_1$ and $[M/x]_A(N_2) = N'_2$
$[M/x]_A(R)$	$= M'$	if $[M/x]_A(R) = M' : A$
$[M/x]_A(R)$	$= R'$	if $[M/x]_A(R) = R'$
$[M/x]_A(N)$	fails	otherwise
Neutral terms		
$[M/x]_A(x)$	$= M : A$	
$[M/x]_A(y)$	$= y$	if $y \neq x$
$[M/x]_A(u[\sigma])$	$= u[\sigma']$	where $[M/x]_A(\sigma) = \sigma'$
$[M/x]_A(RN)$	$= R'N'$	where $[M/x]_A(R) = R'$ and $[M/x]_A(N) = N'$
$[M/x]_A(RN)$	$= M'' : B$	if $[M/x]_A(R) = \lambda y.M' : \Pi y:A_1.B$ where $\Pi x:A_1.B \leq A$ and $[M/x]_A(N) = N'$ and $[N'/y]_{A_1}(M') = M''$
$[M/x]_A(\pi R)$	$= \pi R'$	where $[M/x]_A(R) = R'$
$[M/x]_A(\text{fst } R)$	$= M_1 : B_1$	where $[M/x]_A(R) = (M_1, M_2) : \Sigma x:B_1.B_2$ where $\Sigma x:B_1.B_2 \leq A$
$[M/x]_A(\text{snd } R)$	$= M_2 : B_2$	where $[M/x]_A(R) = (M_1, M_2) : \Sigma x:B_1.B_2$ where $\Sigma x:B_1.B_2 \leq A$
$[M/x]_A(R)$	fails	otherwise
Substitution		
$[M/x]_A(\cdot)$	$= \cdot$	
$[M/x]_A(\sigma, N)$	$= (\sigma', N')$	where $[M/x]_A(\sigma) = \sigma'$ and $[M/x]_A(N) = N'$
$[M/x]_A(\sigma)$	fails	otherwise

Fig. 5. Hereditary substitutions for LF objects with contextual variables

B Meta substitution

The single meta-substitution operation is written as $\boxed{\llbracket \hat{\Psi}.M/u \rrbracket_{A[\hat{\Psi}]}(N)}$ and the simultaneous meta-substitution is written as $\boxed{\llbracket \theta \rrbracket_{\Delta}(N)}$. Subsequently, we define the application of the single meta-substitution to a given term and type, but the simultaneous meta-substitution definition can be easily derived from it.

As we annotate meta-substitutions with their type, we can appropriately annotate σ' with its domain Ψ to obtain M' . Without annotating meta-substitution with the type $C[\Psi]$, we would not be able to annotate the operation $[\sigma']M$ appropriately. Because M' may not be neutral, we may trigger a β -reduction and we return M' together with its (approximate) type C .

Applying the meta-substitution to an LF object will terminate for the same reasons as the ordinary substitution operation terminates; either we apply the substitution to a sub-expression or the objects we substitute are smaller. For an in-depth discussion, we refer the reader to Nanevski et al. [11].

Meta substitution on normal terms / types

$$\llbracket \hat{\Psi}.M/u \rrbracket_{C[\Psi]}(\Pi x:A.B) = \Pi x:A'.B' \text{ where } \llbracket \hat{\Psi}.M/u \rrbracket_{C[\Psi]}(A) = A' \\ \text{and } \llbracket \hat{\Psi}.M/u \rrbracket_{C[\Psi]}(B) = B'$$

$$\llbracket \hat{\Psi}.M/u \rrbracket_{C[\Psi]}(\Sigma x:A.B) = \Sigma x:A'.B' \text{ where } \llbracket \hat{\Psi}.M/u \rrbracket_{C[\Psi]}(A) = A' \\ \text{and } \llbracket \hat{\Psi}.M/u \rrbracket_{C[\Psi]}(B) = B'$$

$$\llbracket \hat{\Psi}.M/u \rrbracket_{C[\Psi]}(\text{type}) = \text{type}$$

$$\llbracket \hat{\Psi}.M/u \rrbracket_{C[\Psi]}(\lambda x.N) = \lambda x.N' \text{ where } \llbracket \hat{\Psi}.M/u \rrbracket_{C[\Psi]}(N) = N'$$

$$\llbracket \hat{\Psi}.M/u \rrbracket_{C[\Psi]}(N_1, N_2) = (N'_1, N'_2) \text{ where } \llbracket \hat{\Psi}.M/u \rrbracket_{C[\Psi]}(N_2) = N'_2 \\ \text{and } \llbracket \hat{\Psi}.M/u \rrbracket_{C[\Psi]}(N_1) = N'_1$$

$$\llbracket \hat{\Psi}.M/u \rrbracket_{C[\Psi]}(R) = N \text{ where } \llbracket \hat{\Psi}.M/u \rrbracket_{C[\Psi]}R = N : A$$

$$\llbracket \hat{\Psi}.M/u \rrbracket_{C[\Psi]}(R) = R' \text{ where } \llbracket \hat{\Psi}.M/u \rrbracket_{C[\Psi]}R = R'$$

$$\llbracket \hat{\Psi}.M/x \rrbracket_{C[\Psi]}(N) \text{ fails otherwise}$$

Meta substitution on neutral terms

$$\llbracket \hat{\Psi}.M/u \rrbracket_{A[\Psi]}(u[\sigma]) = M' : A \text{ where } \llbracket \hat{\Psi}.M/u \rrbracket_{C[\Psi]}(\sigma) = \sigma' \\ \text{and } [\sigma']_{\Psi}(M) = M'$$

$$\llbracket \hat{\Psi}.M/u \rrbracket_{A[\Psi]}(v[\sigma]) = v[\sigma'] \text{ where } \llbracket \hat{\Psi}.M/u \rrbracket_{C[\Psi]}(\sigma) = \sigma'$$

$$\llbracket \hat{\Psi}.M/u \rrbracket_{C[\Psi]}(RN) = R'N' \text{ where } \llbracket \hat{\Psi}.M/u \rrbracket_{C[\Psi]}(R) = R' \\ \text{and } \llbracket \hat{\Psi}.M/u \rrbracket_{C[\Psi]}(N) = N'$$

$$\llbracket \hat{\Psi}.M/u \rrbracket_{C[\Psi]}(RN) = M'' : B \text{ if } \llbracket \hat{\Psi}.M/u \rrbracket_{C[\Psi]}(R) = \lambda y.M' : \Pi x:A.B \\ \text{where } \Pi x:A.B \leq C, N' = \llbracket \hat{\Psi}.M/u \rrbracket_{C[\Psi]}(N) \\ \text{and } M'' = [N'/y]_A(M')$$

$$\llbracket \hat{\Psi}.M/u \rrbracket_{C[\Psi]}(\pi R) = \pi R' \text{ where } \llbracket \hat{\Psi}.M/u \rrbracket_{C[\Psi]}(R) = R'$$

$$\llbracket \hat{\Psi}.M/u \rrbracket_{C[\Psi]}(\text{fst } R) = N_1 : A \text{ if } \llbracket \hat{\Psi}.M/u \rrbracket_{C[\Psi]}(R) = (N_1, N_2) : \Sigma x:A.B \\ \text{where } \Sigma x:A.B \leq C$$

$$\llbracket \hat{\Psi}.M/u \rrbracket_{C[\Psi]}(\text{snd } R) = N_2 : B \text{ if } \llbracket \hat{\Psi}.M/u \rrbracket_{C[\Psi]}(R) = (N_1, N_2) : \Sigma x:A.B \\ \text{where } \Sigma x:A.B \leq C$$

$$\llbracket \hat{\Psi}.M/u \rrbracket_{C[\Psi]}(x) = x$$

$$\llbracket \hat{\Psi}.M/u \rrbracket_{C[\Psi]}(\mathbf{c}) = \mathbf{c}$$

$$\llbracket \hat{\Psi}.M/u \rrbracket_{C[\Psi]}(\mathbf{a}) = \mathbf{a}$$

$$\llbracket \hat{\Psi}.M/x \rrbracket_{C[\Psi]}(R) \text{ fails otherwise}$$

Meta substitution on substitutions

$$\llbracket \hat{\Psi}.M/u \rrbracket_{C[\Psi]}(\cdot) = \cdot$$

$$\llbracket \hat{\Psi}.M/u \rrbracket_{C[\Psi]}(\sigma, M) = \sigma', M' \text{ where } \llbracket \hat{\Psi}.M/u \rrbracket_{C[\Psi]}(\sigma) = \sigma' \\ \text{and } \llbracket \hat{\Psi}.M/u \rrbracket_{C[\Psi]}(M) = M'$$

$$\llbracket \hat{\Psi}.M/x \rrbracket_{C[\Psi]}(\sigma) \text{ fails otherwise}$$

Meta substitution on contexts

$$\llbracket \hat{\Psi}.M/u \rrbracket_{C[\Psi]}(\cdot) = \cdot$$

$$\llbracket \hat{\Psi}.M/u \rrbracket_{C[\Psi]}(\Psi, x:A) = \Psi', x:A' \text{ where } \llbracket \hat{\Psi}.M/u \rrbracket_{C[\Psi]}(\Psi) = \Psi' \\ \text{and } \llbracket \hat{\Psi}.M/u \rrbracket_{C[\Psi]}(A) = A'$$

$$\llbracket \hat{\Psi}.M/x \rrbracket_{C[\Psi]}(\Psi) \text{ fails otherwise}$$

Fig. 6. Meta substitution