

Higher-Order Dynamic Pattern Unification for Dependent Types and Records

Andreas Abel¹ and Brigitte Pientka²

¹ Institut für Informatik, Ludwig-Maximilians-Universität, München, Deutschland
andreas.abel@ifi.lmu.de

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 [7], Miller [9] 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 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 [12], Beluga [14], and Delphin [15] 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 [16]).

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 [11] 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. [10] and Pientka [13], 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. [2], 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 [16], the algorithm sketched in Dowek et al. [2] 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 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.

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 [16] to handle Σ -types. Previously, Elliot [4] described unification for Σ -types in a Huet-style unification algorithm. While it

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

is typically straightforward to incorporate η -expansions and lowering for meta-variables of Σ -type [18, 11], 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 [5] 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 [3] 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 Beluga language; 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.

Due to space restrictions, most proofs have been omitted; they can be found in an extended version of this article on the authors’ homepages.

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	$E ::= \bullet \mid EN \mid \pi E$
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 closed object $x, y. x(\text{succ } y)$. This 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. In general, the meta-variable u stands for a contextual object $\hat{\Psi}.M$ where $\hat{\Psi}$ describes the ordinary bound variables which may occur in M . This allows us to rename the free variables occurring in M if necessary. We use the following convention: If the meta-variable u is associated with the identity substitution, we simply write u instead of $u[\text{id}]$. A meta-variable u has the contextual type $A[\Psi]$ thereby characterizing an object of type A in the context Ψ . 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 . Identity substitutions id_{Φ} are defined recursively by $\text{id} = (\cdot)$ and $\text{id}_{\Phi, x:A} = (\text{id}_{\Phi}, x)$. The subscript Φ is dropped when unambiguous. If Φ is a sub-context of Ψ (in particular if $\Psi = \Phi$) then id_{Φ} is a well-formed substitution in Ψ , i.e., $\Psi \vdash \text{id}_{\Phi} : \Phi$ holds (see Fig. 1). We write $\hat{\Phi}$ for the list of variables $\text{dom}(\Phi)$ in order of declaration.

We write $E[M]$ 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 [1] simply writing $H M_1 \dots M_n$. Evaluation contexts are the proper generalization of spines to projections .

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 $\alpha\{\beta\}$ 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 [16] we indicate rigid occurrences by $\alpha\{\beta\}^{\text{rig}}$ and strongly rigid occurrences by $\alpha\{\beta\}^{\text{srig}}$.

We denote the set of free variables of α by $\text{FV}(\alpha)$ and the set of free meta-variables by $\text{FMV}(\alpha)$. A superscript $^{\text{rig}}$ 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 $A =_{\eta} C$ (rules omitted) compares A and C modulo η , i.e., modulo $R = \lambda x. R x$ ($x \notin \text{FV}(R)$) and $R = (\text{fst } R, \text{snd } R)$.

$$\begin{array}{c}
\text{Neutral Terms/Types} \quad \boxed{\Delta; \Psi \vdash R \Rightarrow A} \\
\frac{\Sigma(\mathbf{a}) = \kappa \quad \Sigma(\mathbf{c}) = A \quad \Psi(x) = A \quad u:A[\Phi] \in \Delta \quad \Delta; \Psi \vdash \sigma \Leftarrow \Phi}{\Delta; \Psi \vdash \mathbf{a} \Rightarrow \kappa \quad \Delta; \Psi \vdash \mathbf{c} \Rightarrow A \quad \Delta; \Psi \vdash x \Rightarrow A \quad \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 \quad \Delta; \Psi \vdash R \Rightarrow \Sigma x:A.B}{\Delta; \Psi \vdash \text{fst } R \Rightarrow A \quad \Delta; \Psi \vdash \text{snd } R \Rightarrow [\text{fst } R/x]_A B} \\
\text{Normal Terms} \quad \boxed{\Delta; \Psi \vdash M \Leftarrow A} \\
\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 \quad \Delta; \Psi \vdash M \Leftarrow A \quad \Delta; \Psi \vdash N \Leftarrow [M/x]_A B}{\Delta; \Psi \vdash \lambda x.M \Leftarrow \Pi x:A.B \quad \Delta; \Psi \vdash (M, N) \Leftarrow \Sigma x:A.B} \\
\text{Substitutions} \quad \boxed{\Delta; \Psi \vdash \sigma \Leftarrow \Psi'} \\
\frac{\Delta; \Psi \vdash \sigma \Leftarrow \Psi' \quad \Delta; \Psi \vdash M \Leftarrow [\sigma]_{\Psi'} A}{\Delta; \Psi \vdash \cdot \Leftarrow \cdot \quad \Delta; \Psi \vdash \sigma, M \Leftarrow \Psi', x:A} \\
\text{LF Types and Kinds} \quad \boxed{\Delta; \Psi \vdash A \Leftarrow s} \\
\frac{\Delta; \Psi \vdash P \Rightarrow \text{type} \quad \Delta; \Psi \vdash A \Leftarrow \text{type} \quad \Delta; \Psi, x:A \vdash B \Leftarrow \text{type}}{\Delta; \Psi \vdash P \Leftarrow \text{type} \quad \Delta; \Psi \vdash \Sigma x:A.B \Leftarrow \text{type}} \\
\frac{\Delta; \Psi \vdash A \Leftarrow \text{type} \quad \Delta; \Psi, x:A \vdash B \Leftarrow s}{\Delta; \Psi \vdash \text{type} \Leftarrow \text{kind} \quad \Delta; \Psi \vdash \Pi x:A.B \Leftarrow s} \\
\text{Meta-Substitutions} \quad \boxed{\Delta \vdash \theta \Leftarrow \Delta'} \\
\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'} \\
\text{Meta-Context} \quad \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}}
\end{array}$$

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

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 $[\sigma]_{\Psi}(\alpha)$ 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* [19], proceeds by resolving created redexes on the fly through new substitutions. Details can be found in Nanevski et al. [10], but do not concern us much here. Single substitution $[N/x]_A(\alpha)$ 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 is written as $\llbracket \hat{\Psi}.M/u \rrbracket(N)$ and the simultaneous meta-substitution is written as $\llbracket \theta \rrbracket(N)$. When we apply $\hat{\Psi}.M/u$ to $u[\sigma]$ we first substitute $\hat{\Psi}.M$ for u in the substitution σ to obtain σ' . Subsequently, we continue to apply σ' to M hereditarily to obtain M' .

The typing rules ensure that the type of the instantiation $\hat{\Psi}.M$ and the type of u agree, i.e. we can replace u which has type $A[\Psi]$ with a normal term M if M has type A in the context Ψ . Because of α -conversion, the variables that are substituted at different occurrences of u may be different, and we write $\hat{\Psi}.M$ where $\hat{\Psi}$ binds all the free variables in M . We can always appropriately rename the bound variable in $\hat{\Psi}$ such that they match the domain of the postponed substitution σ' . This complication can be eliminated in an implementation of the calculus based on de Bruijn indexes.

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:

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, it is necessary to ensure that all substitutions created and used in our transformations can be properly annotated and hence we can use the fact that their application will terminate and produce again normal forms. Second, 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 dependency in this presentation to scale this work from $\lambda^{\Pi\Sigma}$ to non-erasable dependent types such as in Agda.

A unification problem is described by $\Delta \Vdash \mathcal{K}$ 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 [16] 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 $\Delta; \Psi \vdash_{\mathcal{K}} J$ by the same rules as for $\Delta; \Psi \vdash J$ except replacing syntactic equality $=$ with $=_{\mathcal{K}}$. We write $\hat{\Psi}.M =_{\mathcal{K}} N$ if for any ground meta-substitution θ that is a ground solution for \mathcal{K} , we have $\llbracket \theta \rrbracket M =_{\eta} \llbracket \theta \rrbracket N$. To put it differently, if we can solve \mathcal{K} , we can establish that M is equal to N . Following Reed [16], substitution and meta-substitution preserve typing modulo.

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.

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 a neutral term is defined using evaluation contexts to have direct access to the head.

The other *unification steps* (Figure 3) work on a meta-variable and try to find an instantiation for it. We write $\Delta \Vdash \mathcal{K} + \Phi \vdash u \circ M : A$ for instantiating the eta-variable u with the term M in the meta-context Δ and in the constraints \mathcal{K} . This abbreviation is defined under the heading **Instantiation** 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.

Our algorithm can deal with a larger class of patterns where we typically require that meta-variables are associated with a linear substitution. 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.

Decomposition of functions

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

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

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

Orientation

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

η -Contraction

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

Eliminating projections

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

Fig. 2. Local simplification.

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

This equation determines only the first component of the tuple u . Thus, decomposition into $u[y] = y$, which also determines the second component,

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)

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

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

Lowering

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

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

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

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)

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

Fig. 3. Unification steps.

loses solutions. Instead we replace u by a pair (u_1, u_2) of meta-variables of *lower* type, yielding $u_1 : A[\Phi], u_2 : ([u_1/x]B)[\Phi] \Vdash u_1[y] = \text{fst } y$.

Flattening Σ -types $u : P[z : \Pi x:A. \Sigma y:B. C] \Vdash u[\lambda x. (z_1 x, z_2 x)] = M$

By splitting z into two functions z_1, z_2 we arrive at $u : P[z_1 : \Pi x:A. B, z_2 : \Pi x:A. [z_1 x/y]C] \Vdash u[\lambda x. z_1 x, \lambda x. z_2 x] = M$ 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 by $u[x, y, z] = \text{suc } z$.

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

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-variable. 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 [16] 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. [17, 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 $u[f] = \text{suc}(f(\text{suc zero}))$ 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/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.

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

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 ($\text{FMV}(\theta) = \emptyset$), we have $\text{FV}(\sigma) \supseteq \text{FV}([\theta]M)$. 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 fails either if pruned variables occur rigidly, like in $u[x] = c \ y \ v[x, y]$ (constraint unsolvable), or if the flexible occurrence is under another meta variable, like 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.

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. reffig:prune):

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

$\boxed{\text{prune_ctx}_\rho(\tau / \Psi_1) \Rightarrow \Psi_2}$ Prune $\tau : \Psi_1$, returning a sub-context Ψ_2 of Ψ_1 .

$$\frac{}{\text{prune_ctx}_\rho(\cdot / \cdot) \Rightarrow \cdot} \quad \frac{\text{prune_ctx}_\rho(\tau / \Psi_1) \Rightarrow \Psi_2 \quad \text{FV}^{\text{rig}}(M) \not\subseteq \rho}{\text{prune_ctx}_\rho(\tau, M / \Psi_1, x:A) \Rightarrow \Psi_2}$$

$$\frac{\text{prune_ctx}_\rho(\tau / \Psi_1) \Rightarrow \Psi_2 \quad \text{FV}(M) \subseteq \rho \quad \text{FV}(A) \subseteq \hat{\Psi}_2}{\text{prune_ctx}_\rho(\tau, M / \Psi_1, x:A) \Rightarrow \Psi_2, x:A}$$

$\boxed{\Delta \vdash \text{prune}_\rho M \Rightarrow \Delta'; \eta}$ Prune M , returning $\Delta' \vdash \eta \leftarrow \Delta$.

$$\frac{v:B[\Psi_1] \in \Delta \quad \text{prune_ctx}_\rho(\tau / \Psi_1) \Rightarrow \Psi_2 \quad \Psi_2 \neq \Psi_1 \quad \text{FV}(B) \subseteq \hat{\Psi}_2 \quad \eta = \hat{\Psi}_1.v'[\text{id}_{\Psi_2}]/v}{\Delta \vdash \text{prune}_\rho(v[\tau]) \Rightarrow \llbracket \eta \rrbracket(\Delta, v':B[\Psi_2]); \eta}$$

$$\frac{v:B[\Psi_1] \in \Delta \quad \text{prune_ctx}_\rho(\tau / \Psi_1) \Rightarrow \Psi_1}{\Delta \vdash \text{prune}_\rho(v[\tau]) \Rightarrow \Delta; \text{id}_\Delta} \quad \frac{x \in \rho}{\Delta \vdash \text{prune}_\rho x \Rightarrow \Delta; \text{id}_\Delta}$$

$$\frac{\Delta \vdash \text{prune}_\rho R \Rightarrow \Delta_1; \eta_1 \quad \Delta_1 \vdash \text{prune}_\rho(\llbracket \eta_1 \rrbracket M) \Rightarrow \Delta_2; \eta_2}{\Delta \vdash \text{prune}_\rho(RM) \Rightarrow \Delta_2; \llbracket \eta_2 \rrbracket \eta_1}$$

$$\frac{\Delta \vdash \text{prune}_\rho M \Rightarrow \Delta'; \eta}{\Delta \vdash \text{prune}_\rho(\pi M) \Rightarrow \Delta'; \eta} \quad \frac{\Delta \vdash \text{prune}_{\rho,x} M \Rightarrow \Delta'; \eta}{\Delta \vdash \text{prune}_\rho(\lambda x. M) \Rightarrow \Delta'; \eta}$$

$$\frac{\Delta \vdash \text{prune}_\rho M \Rightarrow \Delta_1; \eta_1 \quad \Delta_1 \vdash \text{prune}_\rho(\llbracket \eta_1 \rrbracket N) \Rightarrow \Delta_2; \eta_2}{\Delta \vdash \text{prune}_\rho(M, N) \Rightarrow \Delta_2; \llbracket \eta_2 \rrbracket \eta_1}$$

Fig. 4. Pruning.

When pruning substitution τ with domain Ψ_1 we look at each term M in τ which substitutes for an $x:A$ of Ψ_1 . If M has a rigid occurrence of a variable $y \notin \rho$, we discard the entry $x:A$ from the domain Ψ_1 , thus, effectively removing M from τ . If M 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. Finally, if M has a flexible occurrence of a $y \notin \rho$, pruning fails. Examples:

1. $\text{prune_ctx}_x(\mathbf{c} x, y / x':A, y':B) \Rightarrow x':A$
2. $\text{prune_ctx}_y(\mathbf{c} x, y / x':A, y':B) \Rightarrow y':B$
3. $\text{prune_ctx}_y(\mathbf{c} x, u[y] / x':A, y':B) \Rightarrow y':B$
4. $\text{prune_ctx}_y(u[x], y / x':A, y':B)$ fails

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.

Lemma 1 (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 $\text{FV}([\tau]\text{id}_{\Psi_2}) \subseteq \rho$. 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 $\theta = [[\theta']]\eta$.

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 $[\rho]\text{id}_{\Phi'}$, which is identical to $[\xi]\text{id}_{\Phi'}$, 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:

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

Lemma 2 (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 $\Delta; \Phi \vdash_{\mathcal{K}} \text{id}_{\Phi'} \Leftarrow \Phi'$ and $[\rho]\text{id}_{\Phi'} = [\xi]\text{id}_{\Phi'}$.*

Although we have restricted intersection to variable substitutions, it could be extended to meta-ground substitutions, i.e., substitutions that do not contain any meta-variables.

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 [6].

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

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 $\llbracket \theta_1 \rrbracket \text{id}_{\Delta_0} \in \text{Sol}(\Delta_0 \Vdash \mathcal{K}_0)$.*

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. 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 4 (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 $A =_{\mathcal{K}_0} B$, then $A =_{\mathcal{K}_1} B$.*
2. *If $\Delta_0; \Psi \vdash_{\mathcal{K}_0} M \Leftarrow A$ then $\Delta_1; \llbracket \theta \rrbracket \Psi \vdash_{\mathcal{K}_1} \llbracket \theta \rrbracket M \Leftarrow \llbracket \theta \rrbracket A$.*
3. *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 $\llbracket \theta \rrbracket A =_{\mathcal{K}_1} A'$.*

Next, we define when a set of equations which constitute a unification problem are well-formed using the judgment $\Delta_0 \Vdash_{\mathcal{K}_0} \mathcal{K}$ 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 $\Delta_0 \Vdash \mathcal{K}$ wf to mean $\Delta_0 \Vdash_{\mathcal{K}} \mathcal{K}$ wf.

Lemma 5 (Equations remain well-formed under meta-substitutions). *If $\Delta_0 \Vdash \mathcal{K}$ wf and $\Delta_1 \vdash_{\llbracket \theta \rrbracket \mathcal{K}} \theta \Leftarrow \Delta_0$ then $\Delta_1 \Vdash \llbracket \theta \rrbracket \mathcal{K}$ wf.*

Lemma 6 (Well-formedness of equations is preserved by transitions). *If $\Delta_0 \Vdash \mathcal{K}_0 \mapsto \Delta_1 \Vdash \mathcal{K}_1$ and $\Delta_0 \Vdash_{\mathcal{K}_0} \mathcal{K}$ wf then $\Delta_1 \vdash_{\mathcal{K}_1} \mathcal{K}$ wf.*

Theorem 3 (Unification preserves types). *If $\Delta_0 \Vdash \mathcal{K}_0$ wf and $\Delta_0 \Vdash \mathcal{K}_0 \mapsto \Delta_1 \Vdash \mathcal{K}_1$ then $\Delta_1 \Vdash \mathcal{K}_1$ wf.*

5 Conclusion

We have presented a constraint-based unification algorithm which solves higher-order patterns dynamically and showed its correctness. We have extended higher-order pattern unification to handle Σ -types; this has been an open problem so far, and 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 not been adapted to records. McBride [8, p.6] gives a practical example where the unification problem $T(\text{fst } \gamma) (\text{snd } \gamma) = T' \gamma$ appears, which is not solved by Agda at this point. With the techniques described in this article we will be able to solve such problems and make systems such as Agda more robust.

Correctness of our unification constraint solver is proved using *typing modulo* [16]. 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 [11]: unsolved constraints block normalization, thus guaranteeing termination of the unification algorithm. The idea has been implemented in Agda 2 and can be 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 this article.

References

1. Iliano Cervesato and Frank Pfenning. A linear spine calculus. *Journal of Logic and Computation*, 13(5):639–688, 2003.
2. 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.
3. Dominic Duggan. Unification with extended patterns. *Theoretical Computer Science*, 206(1-2):1–50, 1998.
4. Conal Elliott. *Extensions and Applications of Higher-Order Unification*. PhD thesis, School of Computer Science, Carnegie Mellon University, 1990.
5. 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.
6. 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.
7. W. D. Goldfarb. The undecidability of the second-order unification problem. *Theoretical Computer Science*, 13:225–230, 1981.
8. 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.
9. Dale Miller. Unification of simply typed lambda-terms as logic programming. In *Eighth International Logic Programming Conference*, pages 255–269. MIT Press, 1991.
10. Aleksandar Nanevski, Frank Pfenning, and Brigitte Pientka. Contextual modal type theory. *ACM Transactions on Computational Logic*, 9(3):1–49, 2008.
11. Ulf Norell. *Towards a practical programming language based on dependent type theory*. PhD thesis, Department of Computer Science and Engineering, Chalmers University of Technology, 2007. Technical Report 33D.

12. 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.
13. Brigitte Pientka. *Tabled higher-order logic programming*. PhD thesis, Department of Computer Science, Carnegie Mellon University, 2003. CMU-CS-03-185.
14. 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.
15. 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.
16. 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.
17. Jason Reed. *A Hybrid Logical Framework*. PhD thesis, School of Computer Science, Carnegie Mellon University, 2009.
18. Anders Schack-Nielsen and Carsten Schürmann. Pattern unification for the lambda calculus with linear and affine types. In *International Workshop on Logical Frameworks and Meta-Languages: Theory and Practice (LFMTP'10)*, EPTCS 34, pages 101–116, 2010.
19. 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.