Proof-relevant unification

Dependent pattern matching with only the axioms of your type theory

Jesper Cockx and Dominique Devriese

Abstract

Dependently typed languages such as Agda, Coq and Idris use a syntactic first-order unification algorithm to check definitions by dependent pattern matching. However, standard unification algorithms implicitly rely on principles such as *uniqueness of identity proofs* and *injectivity of type constructors*. These principles are inadmissible in many type theories, particularly in the new and promising branch known as homotopy type theory (HoTT). As a result, programs and proofs in these new theories cannot make use of dependent pattern matching or other techniques relying on unification, and are as a result much harder to write, modify, and understand.

This paper proposes a proof-relevant framework for reasoning formally about unification in a dependently typed setting. In this framework, unification rules compute not just a unifier but also a corresponding soundness proof in the form of an *equivalence* between two sets of equations. By rephrasing the standard unification rules in a proof-relevant manner, they are guaranteed to preserve soundness of the theory. In addition, it enables us to safely add new rules that can exploit the dependencies between the types of equations, such as rules for eta-equality of record types and higher-dimensional unification rules for solving equations between equality proofs.

Using our framework, we implemented a complete overhaul of the unification algorithm used by Agda. As a result, we were able to replace previous ad hoc restrictions with formally verified unification rules, fixing a substantial number of bugs in the process. In the future, we may also want to integrate new principles with pattern matching, for example the higher inductive types introduced by HoTT. Our framework also provides a solid basis for such extensions to be built on.

1 Introduction

Unification is a generic method for solving symbolic equations algorithmically. It is a fundamental algorithm used in many areas in computer science, such as logic programming, type inference, term rewriting, automated theorem proving, and natural language processing. In particular, type checkers for languages with dependent pattern matching (Coquand, 1992) such as Agda (Norell, 2007), the Equations package for Coq (Sozeau, 2010), Idris (Brady, 2013), and Lean (de Moura, Kong, Avigad, van Doorn, and von Raumer, 2015) use first-order unification to determine whether a set of constructors covers all possible cases.¹

¹ They also use a different *higher-order* unification algorithm to solve constraints and derive the values of implicit arguments, but this is not the focus of this paper.

Although first-order unification is well understood in the untyped or simply typed setting, its interaction with dependent types has been mysterious so far. In particular, some standard unification rules are no longer valid in the presence of universes and indexed datatypes, two common features of dependently typed languages. Examples of how this can go wrong follow later in this introduction, but we start with an example where first-order unification works as intended.

Example 1. Consider the type of length-indexed vectors $\text{Vec } A \ n$ (i.e. the type of vectors containing n elements of type A) where $\text{nil} : \text{Vec } A \ \text{zero}$ is the empty vector and $\text{cons } n \ x \ xs$ is the vector with head $x : A$ and tail $xs : \text{Vec } A \ n$. We can define a safe tail function on vectors by dependent pattern matching as follows:

$$\begin{aligned} \text{tail} &: (n : \mathbb{N}) \rightarrow \text{Vec } A \ (\text{suc } n) \rightarrow \text{Vec } A \ n \\ \text{tail}.m \ (\text{cons } m \ x \ xs) &= xs \end{aligned} \tag{1}$$

The function tail needs only be defined in the case for $(\text{cons } n \ x \ xs)$: the case for nil is impossible because unification of zero (the length of nil) with $\text{suc } n$ reports an absurdity. This is all the better because there is no way to take the tail of an empty vector! In the cases where unification succeeds, it can also teach us something extra about the type of the right-hand side. For example, in the remaining case $\text{tail}.m \ (\text{cons } m \ x \ xs)$, unification of $\text{suc } m$ with $\text{suc } n$ tells us that n must be equal to m , as indicated by the so-called *inaccessible pattern* $.m$. This method of solving equations to either gain more information about the type of the right-hand side or to derive an absurdity is called *specialization by unification* (Goguen, McBride, and McKinna, 2006).

In a language that has dependent pattern matching as a primitive such as Agda (Norell, 2007), the particularities of the unification rules used become crucial for the language's notion of equality. Indeed, we can match on a proof of $u \equiv_A v$ with the constructor refl precisely when the unification algorithm is able to unify u with v . For example, if the unification algorithm is allowed to delete reflexive equations of the form $u = u$, then this allows us to prove uniqueness of identity proofs (UIP) by pattern matching (Coquand, 1992). So it is important to have a solid theoretical understanding of unification in order to study these languages.

When dependently typed terms themselves become the subject of unification, the unification algorithm can encounter *heterogeneous equations*: equations in which the left- and right-hand side have different types, that only become equal after previous equations have been solved. For example, consider the type $\Sigma_{A:\text{Set}} A$ with elements (A, a) packing a type A together with an element a of that type. By injectivity of the pair constructor $_, _$, an equation $(A, a) = (B, b)$ can be simplified to $A = B$ and $a = b$, but the type of the second equation is now heterogeneous since $a : A$ and $b : B$. Because traditional unification algorithms only look at the syntax of the terms they are trying to unify, they cause problems when applied to heterogeneous equalities.

Example 2. Consider the equation $(\text{Bool}, \text{true}) = (\text{Bool}, \text{false})$ of type $\Sigma_{A:\text{Set}} A$. By injectivity of the constructor $_, _$, we can simplify this equation to the two equations $\text{Bool} = \text{Bool}$ and $\text{true} = \text{false}$ and then derive an absurdity from the

second equation. However, this line of reasoning depends on the principle of *equality of second projections*, which is equivalent to UIP (Streicher, 1993). In a univalent theory such as HoTT (The Univalent Foundations Program, 2013) it is actually possible to prove that $(\text{Bool}, \text{true}) = (\text{Bool}, \text{false})$ of type $\Sigma_{A:\text{Set}} A$ under the equivalence between Bool and itself swapping true and false , refuting the use of injectivity above. So the naive injectivity rule above cannot be used in such a theory.

On the other hand, consider the exact same unification problem $(\text{Bool}, \text{true}) = (\text{Bool}, \text{false})$, but this time the type of the equation is a non-dependent product $\text{Set} \times \text{Bool}$ (defined as $\Sigma_{_:\text{Set}} \text{Bool}$). In this case it *is* possible to derive an absurdity, even in a univalent theory. However, a unification algorithm can never distinguish between these two equations unless it takes their types into account.

Example 3. Suppose we define two copies Bool_1 and Bool_2 of the boolean type with constructors $\text{true}_1, \text{false}_1$ and $\text{true}_2, \text{false}_2$ respectively, then it is unsound to apply the conflict rule on the (heterogeneous) equation $\text{true}_1 = \text{false}_2$. Doing so would allow us to prove that $\text{Bool}_1 \not\equiv_{\text{Set}} \text{Bool}_2$, again contradicting univalence.

Example 4. The problem is not limited to theories that do not support UIP, either. Problems can also occur when we use a naive injectivity rule for constructors of indexed datatypes. Let A be an arbitrary type and $\text{Singleton} : A \rightarrow \text{Set}$ be an indexed datatype with one constructor $\text{sing} : (x : A) \rightarrow \text{Singleton } x$ and consider the unification problem $(\text{Singleton } s, \text{sing } s) = (\text{Singleton } t, \text{sing } t)$. If we allow the injectivity rule to simplify $\text{sing } x = \text{sing } y$ to $x = y$ then this problem can be solved with solution $y \mapsto x$. However, this would allow us to prove injectivity of the type constructor Singleton . In general, injectivity of type constructors is an undesirable property because it is not only incompatible with the law of the excluded middle (Theorem 93), but also with univalence (Theorem 92) and with an impredicative universe of propositions (Miquel, 2010). In particular, if we let $A = \text{Set} \rightarrow \text{Set}$ in the above example then the injectivity of the Singleton type constructor allows us to refute the law of the excluded middle.

The unification algorithm used by Agda 2.4 (and older) for checking definitions by dependent pattern matching contains a number of restrictions to avoid bad unification steps like in the above examples. One of these restrictions is to not delete equations of the form $u = u$ if the theory doesn't support UIP (Cockx, Devriese, and Piessens, 2016a). The rule for simplifying equations of the form $c \ u_1 \ \dots \ u_n = c \ v_1 \ \dots \ v_n$ is also restricted in case c is a constructor of an indexed datatype. However, these ad hoc restrictions make the unification algorithm hard to prove correct, modify, or extend.

Contributions. In this paper, we give a *typed* and *proof-relevant* account of the first-order unification algorithm used for checking definitions by dependent pattern matching, in order to solve the problems with untyped unification in general and put unification in type theory back on a solid theoretical foundation. Concretely, we make the following contributions:

- We give a new representation of unification rules and most general unifiers in a dependently typed setting as *equivalences* between solution spaces represented by telescopic systems of equations, serving as evidence of their soundness (Section 3). This allows us to use type information to decide when a unification rule is applicable and avoid the use of any axioms such as UIP.
- We show how the standard first-order unification rules used by Goguen et al. (2006) can be implemented as equivalences (Section 4). We extend these rules to the case of indexed families of datatypes. These rules work on heterogeneous equations and can solve multiple equations at once, making them more general than the ones in our previous work (Cockx, Devriese, and Piessens, 2014). We also show how to extend the unification algorithm with rules for η -equality of record types.
- To guarantee good computational properties of the unifiers produced by our unification algorithm, we introduce the notion of a *strong unification rule* (Section 5). We show that all the unification rules used by our algorithm are strong ones, except for the (optional) deletion rule.
- We show how to make the injectivity rule for indexed datatypes more general by a technique called *higher-dimensional unification* (Section 6). In particular, this technique formalizes the concept of forced constructor arguments, a heuristic that allows unification to skip certain constructor arguments if they are determined by the type of the constructor (Corollary 67).
- We reimplement the unification algorithm used by Agda for pattern matching on indexed families of datatypes based on our framework for proof-relevant unification, eliminating the previous ad-hoc restrictions and fixing a number of bugs in the process (Section 8). This new unification algorithm has been released as part of Agda version 2.5.1².

This paper is based on the work of two conference papers (Cockx et al., 2016a; Cockx and Devriese, 2017), as well as the first author’s PhD thesis (Cockx, 2017). Compared to the conference papers, we made the following additions:

- We give a new definition of a strong unifier (Definition 53). Compared the previous definition (Cockx et al., 2016a), this definition is more natural to work with and allows us to prove that functions constructed by unification satisfy the expected computational behaviour (Lemma 56), while it is still satisfied by all the unification rules.
- We prove that lifting a strong unifier results again in a strong unifier (Lemma 73).
- We give a formal proof that (a suitably internalized notion of) a most general unifier is really equivalent to an equivalence (Lemma 17 and Lemma 18).
- The presentation of many examples, definitions, lemmas and theorems was improved compared to the conference version.

² Available from <http://wiki.portal.chalmers.se/agda/>.

Overview. First, we give an overview of the type theory we work in, including syntax and typing rules for telescopes and telescopic equality (Section 2). We start the paper proper with a general description of our framework for reasoning about unification in a dependently typed setting (Section 3). We phrase the basic unification rules in this framework and show how our algorithm can easily be extended by adding more unification rules (Section 4). We pay special attention to the computational behaviour of unification rules when viewed as terms in type theory (Section 5). To augment the power of the unification rules for indexed datatypes, we introduce a new technique called higher-dimensional unification (Section 6). We show that our unification rules are conservative over standard type theory by translating them to the standard datatype eliminators (Section 7). We also discuss the implementation of our unification algorithm in Agda (Section 8). We finish the paper with a discussion of related work (Section 9) and future work (Section 10). The appendix contains a proof of the incompatibility between injective type constructors on one hand and univalence and the law of the excluded middle on the other hand (Appendix A).

2 Preliminaries

We base ourselves on Martin-Löf’s Intuitionistic Theory of Types with dependent function types, inductive families, and universes (Martin-Löf, 1972; Martin-Löf, 1984). However, the results in this paper should be equally applicable in other type theories with inductive families such as the Unified Theory of Dependent Types (UTT) by Luo (1994) or the Calculus of Inductive Constructions (CIC) used by Coq. The main reason we do not use these more expressive calculi is because we do not need their additional features, and not using them makes our results more generally applicable.

2.1 Basic syntax and typing rules

We use a syntax closely resembling that of Agda (Figure 1). Types and terms share the same syntactic class. As a convention, types are indicated by capital letters $A, B, \dots$ and other terms by small letters $u, v, \dots$. Aside from the standard type-theoretic constructs, the syntax includes datatypes D , constructors c , and defined functions f .

The set of variables that occur freely in u is indicated by $FV(u)$. Simultaneous substitution $u[x_1 \mapsto v_1, \dots, x_n \mapsto v_n]$ is defined by simultaneously replacing all free occurrences of $x_1, \dots, x_n$ in u by $v_1, \dots, v_n$, avoiding variable capture by renaming bound variables when necessary.

The basic rules for context validity, typing, and definitional equality are given in Figure 2, Figure 3, and Figure 4 respectively. In addition to these rules, we assume one rule for each datatype D , constructor c , and defined function f , asserting that the symbol has its declared type in any valid context.

$\Gamma ::= ()$	(empty context)
$\Gamma(x : A)$	(context extension)
$A, B, u, v ::= x$	(variable)
$u v$	(application)
$\lambda x. u$	(lambda abstraction)
$(x : A) \rightarrow B$	(dependent function type)
$\mathbf{Set}_\ell$	(universe ℓ)
$\mathbf{D}$	(inductive datatype)
$\mathbf{c}$	(data constructor)
$\mathbf{f}$	(defined function)

Figure 1: The syntax for contexts, types, and terms.

$$\frac{}{() \text{ context}}$$

$$\frac{\Gamma \vdash A : \mathbf{Set}_\ell \quad x \notin FV(\Gamma)}{\Gamma(x : A) \text{ context}}$$

Figure 2: The typing rules for valid contexts.

2.2 Inductive families of datatypes.

Inductive families of datatypes are (dependent) types inductively defined by a number of constructors (Dybjer, 1991). Inductive families can also have *parameters* and *indices*. In the syntax (Figure 1), datatypes are written $\mathbf{D}$ and constructors $\mathbf{c}$.

$$\frac{\Gamma \text{ context} \quad x : A \in \Gamma}{\Gamma \vdash x : A}$$

$$\frac{\Gamma \vdash u : A_1 \quad \Gamma \vdash A_1 = A_2 : \mathbf{Set}_\ell}{\Gamma \vdash u : A_2}$$

$$\frac{\Gamma \text{ context}}{\Gamma \vdash \mathbf{Set}_\ell : \mathbf{Set}_{\ell+1}}$$

$$\frac{\Gamma \vdash A : \mathbf{Set}_\ell \quad \Gamma(x : A) \vdash B : \mathbf{Set}_{\ell'}}{\Gamma \vdash (x : A) \rightarrow B : \mathbf{Set}_{\max(\ell, \ell')}}}$$

$$\frac{\Gamma(x : A) \vdash u : B}{\Gamma \vdash \lambda x. u : (x : A) \rightarrow B}$$

$$\frac{\Gamma \vdash f : (x : A) \rightarrow B \quad \Gamma \vdash u : A}{\Gamma \vdash f u : B[x \mapsto u]}$$

Figure 3: The core typing rules, including dependent function types $(x : A) \rightarrow B$ and an infinite hierarchy of universes $\mathbf{Set}_0, \mathbf{Set}_1, \mathbf{Set}_2, \dots$

$$\begin{array}{c}
\frac{\Gamma(x : A) \vdash u : B \quad \Gamma \vdash v : A}{\Gamma \vdash (\lambda x. u) v = u[x \mapsto v] : B[x \mapsto v]} \\
\\
\frac{\Gamma \vdash u : A}{\Gamma \vdash u = u : A} \\
\\
\frac{\Gamma \vdash u_1 = u_2 : A}{\Gamma \vdash u_2 = u_1 : A} \\
\\
\frac{\Gamma \vdash u_1 = u_2 : A \quad \Gamma \vdash u_2 = u_3 : A}{\Gamma \vdash u_1 = u_3 : A} \\
\\
\frac{\Gamma \vdash u_1 = u_2 : A_1 \quad \Gamma \vdash A_1 = A_2 : \mathbf{Set}_\ell}{\Gamma \vdash u_1 = u_2 : A_2} \\
\\
\frac{\Gamma \vdash A_1 = A_2 : \mathbf{Set}_\ell \quad \Gamma(x : A_1) \vdash B_1 = B_2 : \mathbf{Set}_{\ell'}}{\Gamma \vdash (x : A_1) \rightarrow B_1 = (x : A_2) \rightarrow B_2 : \mathbf{Set}_{\max(\ell, \ell')}} \\
\\
\frac{\Gamma(x : A) \vdash u_1 = u_2 : B}{\Gamma \vdash \lambda x. u_1 = \lambda x. u_2 : (x : A) \rightarrow B} \\
\\
\frac{\Gamma \vdash f_1 = f_2 : (x : A) \rightarrow B \quad \Gamma \vdash u_1 = u_2 : A}{\Gamma \vdash f_1 u_1 = f_2 u_2 : B[x \mapsto u_1]}
\end{array}$$

Figure 4: The rules for definitional equality, including rules for β -equality, reflexivity, symmetry, transitivity, and congruence.

Example 5. $\mathbb{N}$ is defined as an inductive datatype with the constructors $\mathbf{zero} : \mathbb{N}$ and $\mathbf{suc} : \mathbb{N} \rightarrow \mathbb{N}$.

Example 6. $\mathbf{Vec} A n$ is an inductive family with one parameter $A : \mathbf{Set}$, one index $n : \mathbb{N}$, and two constructors $\mathbf{nil} : \mathbf{Vec} A \mathbf{zero}$ and $\mathbf{cons} : (n : \mathbb{N}) \rightarrow A \rightarrow \mathbf{Vec} A n \rightarrow \mathbf{Vec} A (\mathbf{suc} n)$.

In the original definition of indexed families by Dybjer (1991), parameters are required to occur uniformly everywhere in the definition of the datatype, while indices can vary from constructor to constructor. Agda is less restrictive and also allows parameters to occur non-uniformly in the types of recursive constructor arguments, but not in their return types. The work in this paper is valid for both variants of the definition.

The values of the parameters are not arguments to the constructor $\mathbf{c}$, not even implicitly. This is intentional: requiring constructors to remember their parameters is impractical from an implementation perspective, so we make sure they are never needed for the unification algorithm described in this paper.

2.3 The identity type

The propositional equality type $x \equiv_A y$ expresses the property that x and y are equal elements of type A (Martin-Löf, 1984). The basic way to prove a propositional equality is by using $\mathbf{refl} : x \equiv_A x$ (short for reflexivity); for example $\mathbf{refl}$ is a

proof of $\mathbf{zero} \equiv_{\mathbb{N}} \mathbf{zero}$. The identity type comes equipped with a number of useful reasoning principles besides $\mathbf{refl}$:

- $\mathbf{sym} : \{x\ y : A\} \rightarrow x \equiv_A y \rightarrow y \equiv_A x$ expresses the symmetry of propositional equality.³
- $\mathbf{trans} : \{x\ y\ z : A\} \rightarrow x \equiv_A y \rightarrow y \equiv_A z \rightarrow x \equiv_A z$ expresses the transitivity of propositional equality.
- $\mathbf{cong} : (f : A \rightarrow B) \rightarrow \{x\ y : A\} \rightarrow x \equiv_A y \rightarrow f\ x \equiv_A f\ y$ expresses congruence: applying a function to equal arguments gives equal results.
- $\mathbf{subst} : (P : A \rightarrow \mathbf{Set}) \rightarrow \{x\ y : A\} \rightarrow x \equiv_A y \rightarrow P\ x \rightarrow P\ y$ expresses substitution: if two types are equal up to some propositionally equal terms, then we can transport elements from one type to the other.

In particular, if we take $A = \mathbf{Set}$ and $P = \lambda X. X$ in the type of $\mathbf{subst}$, the result is a function of type $(X\ Y : \mathbf{Set}) \rightarrow (X \equiv_{\mathbf{Set}} Y) \rightarrow X \rightarrow Y$. This function is often called $\mathbf{coerce}$ because it allows us to coerce a term from one type to another if the types are propositionally equal.

The standard eliminator for the identity type $x \equiv_A y$ is called the $\mathbf{J}$ rule:

$$\mathbf{J} : \{x : A\} (P : (y : A) \rightarrow x \equiv_A y \rightarrow \mathbf{Set}) (p : P\ x\ \mathbf{refl}) (y : A) (e : x \equiv_A y) \rightarrow P\ y\ e \quad (2)$$

This rule is a generalization of $\mathbf{subst}$ where the type P is allowed to depend on the given equality proof. Using only $\mathbf{J}$, it is possible to define $\mathbf{sym}$, $\mathbf{trans}$, $\mathbf{cong}$, and $\mathbf{subst}$, and $\mathbf{coerce}$.

We will also rely on the concept of pointwise equality between two functions.

Definition 7 (Pointwise equality). Let $f, g : (x : A) \rightarrow B\ x$ be two functions. The pointwise equality type $f \doteq g$ is defined as $(x : A) \rightarrow f\ x \equiv_{B\ x} g\ x$. Similarly, if $\sigma, \tau : \Delta \rightarrow \Gamma$ are two telescope mappings then $\sigma \doteq \tau$ is defined as $(\bar{x} : \Delta) \rightarrow \sigma\ \bar{x} \equiv_{\Gamma} \tau\ \bar{x}$.

In case we have access to functional extensionality, pointwise equality becomes equivalent with regular propositional equality. However, we don't rely on functional extensionality for any of the work in this paper.

2.4 Equivalences

The central concept we use to represent unification rules in this paper is the notion of an *equivalence*. We use Definition 4.3.1 from The Univalent Foundations Program (2013):

Definition 8. A function $f : A \rightarrow B$ is an *equivalence* if we have two functions $g_1 : B \rightarrow A$ and $g_2 : B \rightarrow A$ and proofs that they are respectively a left and a right inverse of f (i.e. terms of type $(x : A) \rightarrow g_1\ (f\ x) \equiv_A x$ and $(y : B) \rightarrow f\ (g_2\ y) \equiv_B y$).

³ As in Agda, we use two function spaces $(x : A) \rightarrow B$ and $\{x : A\} \rightarrow B$ for explicit and implicit functions respectively. The values of implicit arguments can always be deduced from the types of the other arguments, so we omit them when applying the function. For example, we write $\mathbf{sym}\ e$ instead of $\mathbf{sym}\ x\ y\ e$. In case we do want to make these arguments explicit, we write them between curly brackets as well, for example $\mathbf{sym}\ \{x\}\ \{y\}\ e$.

$$\begin{array}{c}
\frac{\Gamma \text{ context}}{\Gamma \vdash () \text{ telescope}} \text{ (Tel-empty)} \\
\\
\frac{\Gamma \vdash A : \text{Set}_\ell \quad \Gamma(x : A) \vdash \Delta \text{ telescope}}{\Gamma \vdash (x : A) \Delta \text{ telescope}} \text{ (Tel-extend)} \\
\\
\frac{\Gamma \text{ context}}{\Gamma \vdash () : ()} \text{ (List-empty)} \\
\\
\frac{\Gamma \vdash u : A \quad \Gamma \vdash \bar{u} : \Delta[x \mapsto u]}{\Gamma \vdash u; \bar{u} : (x : A) \Delta} \text{ (List-extend)}
\end{array}$$

Figure 5: The typing rules for telescopes.

Two types A and B are *equivalent* if there exists an equivalence from A to B . The type of all equivalences $f : A \rightarrow B$ is written as $A \simeq B$.

If $f : A \simeq B$ is an equivalence, then we also write f for the function $f : A \rightarrow B$. We write $\text{linv } f$ for the function g_1 (for the **left inverse** of f), and $\text{isLinv } f$ for the proof of $(x : A) \rightarrow \text{linv } f (f x) \equiv_A x$. Similarly, $\text{rinv } f$ stands for the function g_2 and $\text{isRinv } f$ for the proof of $(y : B) \rightarrow f (\text{rinv } f y) \equiv_B y$. For many equivalences, $\text{linv } f$ and $\text{rinv } f$ are definitionally equal. In that case we write f^{-1} for their common value.

The notion of equivalence plays a central role in Voevodsky's univalence axiom. However, our work doesn't require any primitives on top of basic intuitionistic type theory (such as univalence). In fact, our work can be equally well understood without any knowledge of HoTT, and is still useful in a setting that assumes entirely different axioms (such as the law of the excluded middle from classical logic).

2.5 Telescopes

A telescope is a list of typed variable bindings where each type can depend on the previous variables (de Bruijn, 1991). For example, $(m : \mathbb{N})(p : m \equiv_{\mathbb{N}} \text{zero})$ is a telescope of length 2. Telescopes are much like contexts in the sense that they consist of a sequence of variable typings of the form $(x : A)$. However, they are used for different purposes so it is best to keep the two concepts separate. While contexts grow to the right, telescopes grow to the left. One way to think about a telescope is as the *tail* of a context: while a context must always be closed, a telescope can contain free variables from an ambient context, and the telescope can be added to that context to produce a new, extended context.

If there are multiple variables of the same type after each other, then we usually only write the type once. For example, $(x y z : A)$ stands for the telescope $(x : A)(y : A)(z : A)$.

Telescopes are used as the type of a list of terms. A list of terms is indicated by a bar above the letter, for example $\bar{t} = (\text{zero}; \text{refl}) : (m : \mathbb{N})(p : m \equiv_{\mathbb{N}} \text{zero})$. We also write $()$ for the empty list of terms. The typing rules for telescopes and lists of terms are given in Figure 5.

Telescopes are useful for various other purposes: a telescope can be used

- ... as an extension to the context: $\Gamma\Delta$ is defined by $\Gamma() := \Gamma$ and $\Gamma((x : A)\Delta) := (\Gamma(x : A))\Delta$. In particular, if Δ is a valid telescope in the empty context then Δ can be used as the context $()\Delta$.
- ... as the names of the variables of a parallel substitution: $u[\Delta \mapsto \bar{v}]$ is defined by substituting the values $\bar{v}$ for the variables of Δ in u .
- ... as the argument types of an iterated function type: $\Delta \rightarrow B$ is defined by $() \rightarrow B := B$ and $(x : A)\Delta \rightarrow B := (x : A) \rightarrow (\Delta \rightarrow B)$.
- ... in the definition of an iterated lambda abstraction: $\lambda\Delta. u$ is defined by $\lambda(). u := u$ and $\lambda((x : A)\Delta). u := \lambda x. (\lambda\Delta. u)$.
- ... as a list of the variables in the telescope: $f \Delta$ is defined by $f() := f$ and $f((x : A)\Delta) := (f x) \Delta$.

These various interpretations of telescopes can be used together. For example, if $\Gamma \vdash f : \Delta \rightarrow B$ then we have $\Gamma\Delta \vdash f \Delta : B$.

If $\mathbf{D}$ is a datatype with indices Ξ , we write $\bar{\mathbf{D}}$ for the telescope $(\bar{u} : \Xi)(x : \mathbf{D} \bar{u})$, for example $\overline{\mathbf{Vec} A} = (n : \mathbb{N})(x : \mathbf{Vec} A n)$.

A function between telescopes is called a *telescope mapping*. A telescope mapping $f : \Delta \rightarrow \Delta'$ maps variables of type Δ to values of type Δ' . Telescope mappings generalize the concept of a (non-dependent) function to multiple inputs and multiple outputs. They could be encoded as normal functions by representing a telescope by an iterated Σ -type, but we find it useful to define them as a first-class concept.

Another way to view a telescope mapping $f : \Delta \rightarrow \Delta'$ is as a *typed* variant of a substitution. In particular, if we have a term $u : A$ with free variables coming from Δ' , then we can apply the substitution $[\Delta' \mapsto f \Delta]$ to it, replacing the variables from Δ' by the values given by $f \Delta$, to get a term u' with free variables coming from Δ .

Example 9. Suppose $\Delta = (k : \mathbb{N})$ and $\Delta' = (m n : \mathbb{N})$ and let $f k = (\mathbf{zero}; \mathbf{succ} k)$. We have $\Delta' \vdash m + n : \mathbb{N}$, so applying the substitution $[\Delta' \mapsto f \Delta] = [m \mapsto \mathbf{zero}; n \mapsto \mathbf{succ} k]$ gives us $\Delta \vdash \mathbf{zero} + \mathbf{succ} k : \mathbb{N}$.

When we use a telescope mapping $f : \Delta \rightarrow \Delta'$ as a substitution, the substitution goes in the ‘opposite’ direction: it takes terms with free variables Δ' to terms with free variables Δ . In this case, the type of f is often written as $\Delta \vdash f : \Delta'$. However, since in this paper we use telescope mappings mainly as functions rather than as substitutions, we stick to the notation $f : \Delta \rightarrow \Delta'$.

2.6 Heterogeneous and telescopic equality

The identity type $x \equiv_A y$ only allows equations between elements of the same type, so we still need a way to represent heterogeneous equations. For this purpose, McBride (2000) introduced a heterogeneous equality type $x_A \cong_B y$ where $x : A$ and $y : B$ can be of different types, but $x_A \cong_B y$ can only be proven if the types A and B are actually the same. Using this type, a unification problem can be represented by the (non-dependent) product of the individual equalities. By maintaining the

invariant that the leftmost equation is always homogeneous, the equations can be solved step by step, from left to right. However, using this heterogeneous equality type causes a number of problems:

- Turning a proof of heterogeneous equality between elements of the same type into a homogeneous one requires UIP. So in a theory without UIP (such as HoTT), heterogeneous equalities are worthless.
- Using heterogeneous equality causes information about dependencies between the equations to be lost. For example, if we have two equations $\text{Bool}_{\text{Set}} \cong_{\text{Set}} \text{Bool}$ and $\text{true}_{\text{Bool}} \cong_{\text{Bool}} \text{false}$, there is no way to see whether the type of the second equation depends on the first. Example 2 shows that both cases are possible, and that it is essential to know the difference!
- Finally, it is unsound to postpone an equation and continue with the next one when working with heterogeneous equality, since this allows us to prove things such as injectivity of certain type constructors (Example 4).

To avoid these problems and keep track of the dependencies between equations, we use the concept from HoTT of an equality “laying over” another one. There are multiple equivalent ways to define this type; for the sake of simplicity we use the following definition in terms of the regular homogeneous equality by substituting on the left:

Definition 10. Let $e : s \equiv_A t$ and $P : (x : A) \rightarrow \text{Set}$. We define the type $u \equiv_P^e v$ of equality proofs between $u : P s$ and $v : P t$ *laying over* e by

$$u \equiv_P^e v = (\text{subst } P e u) \equiv_{(P t)} v \quad (3)$$

In practice, the exact definition of $u \equiv_P^e v$ doesn’t have much impact, what’s important is that $(u \equiv_P^{\text{refl}} v) = (u \equiv_{P s} v)$ whenever $s = t$. An alternative definition would be to define $u \equiv_P^e v$ as a new datatype indexed over s, t, e, u , and v with a single constructor $\text{refl} : u \equiv_P^{\text{refl}} u$. Yet another alternative definition would be to define the type $u \equiv_P^e v$ by matching on e , giving $(u \equiv_P^{\text{refl}} v) = u \equiv_{P s} v$ in case e is refl . We prefer our definition to these more symmetric alternatives because it doesn’t require auxiliary datatypes or large eliminations.

We often write $u \equiv_{P e} v$ instead of $u \equiv_P^e v$. For example, if $e : m \equiv_{\mathbb{N}} n$ and $u : \text{Vec } A m$ and $v : \text{Vec } A n$ are two vectors, then we may form the type $u \equiv_{\text{Vec } A e} v$. This notation is inspired by cubical type theory (Cohen, Coquand, Huber, and Mörtberg, 2016), where a function $f : A \rightarrow B$ is automatically lifted to a function $x \equiv_A y \rightarrow f x \equiv_B f y$. In our setting it is merely a convenient abuse of notation.

Using this notion of an equality proof laying over another, we can define a version of cong that works for dependent functions:

$$\begin{aligned} \text{dcong} : (f : (x : A) \rightarrow B x) &\rightarrow \{x y : A\} \rightarrow (e : x \equiv_A y) \rightarrow f x \equiv_{B e} f y \\ \text{dcong } f \text{ refl} &= \text{refl} \end{aligned} \quad (4)$$

Telescopic equality is defined as follows:

Definition 11. Let Δ be a telescope and $\Gamma \vdash \bar{s}, \bar{t} : \Delta$. We define a new telescope $(\bar{e} : \bar{s} \equiv_{\Delta} \bar{t})$ called the *telescopic equality* between $\bar{s}$ and $\bar{t}$ inductively on the length of Δ by $() \equiv_{()} () = ()$ and

$$(e; \bar{e} : s; \bar{s} \equiv_{(x:A)\Delta} t; \bar{t}) = (e : s \equiv_A t) (\bar{e} : \bar{s} \equiv_{\Delta[x \mapsto e]} \bar{t}) \quad (5)$$

For each $\bar{t} : \Delta$, we define $\overline{\text{refl}} : \bar{t} \equiv_{\Delta} \bar{t}$ as $\text{refl}; \dots; \text{refl}$.

For example, $((e_1; e_2) : (m; u) \equiv_{(x:\mathbb{N})(y:\text{Vec } A \ x)} (n; v))$ stands for the telescope $(e_1 : m \equiv_{\mathbb{N}} n)(e_2 : u \equiv_{\text{Vec } A \ e_1} v)$.

Lemma 12. We have the telescopic equality eliminator

$$\bar{J} : (P : (\bar{s} : \Delta) \rightarrow \bar{r} \equiv \bar{s} \rightarrow \text{Set}_i) \rightarrow P \bar{r} \overline{\text{refl}} \rightarrow (\bar{s} : \Delta) \rightarrow (\bar{e} : \bar{r} \equiv \bar{s}) \rightarrow P \bar{s} \bar{e} \quad (6)$$

Construction. We define $\bar{J}$ by eliminating the equations $\bar{e}$ from left to right using $\bar{J}$:

$$\begin{aligned} \bar{J} P p () () &= p \\ \bar{J} P p (s; \bar{s}) (e; \bar{e}) &= \bar{J} (\lambda s; e. (\bar{s} : \Delta) (\bar{e} : \bar{r} \equiv \bar{s}) \rightarrow P (s; \bar{s}) (e; \bar{e})) \\ &\quad (\lambda \bar{s}; \bar{e}. \bar{J} (\lambda \bar{s}; \bar{e}. P (r; \bar{s}) (\text{refl}; \bar{e})) p \bar{e}) \\ &\quad e \bar{s} \bar{e} \end{aligned} \quad (7)$$

Each elimination of an equation $e_i : r_i \equiv s_i$ fills in refl for all occurrences of e_i , allowing the next equations to reduce and in particular ensuring that the following equation is of the correct form. $\square$

Using $\bar{J}$, we also define telescopic versions of subst , cong and dcong :

$$\begin{aligned} \overline{\text{subst}} &: (P : \Delta \rightarrow \text{Set}_\ell) \rightarrow \{\bar{u} \bar{v} : \Delta\} \rightarrow \bar{u} \equiv_{\Delta} \bar{v} \rightarrow P \bar{u} \rightarrow P \bar{v} \\ \overline{\text{cong}} &: (f : \Delta \rightarrow T) \rightarrow \{\bar{u} \bar{v} : \Delta\} \rightarrow \bar{u} \equiv_{\Delta} \bar{v} \rightarrow f \bar{u} \equiv_T f \bar{v} \\ \overline{\text{dcong}} &: (f : (\bar{x} : \Delta) \rightarrow T \bar{x}) \rightarrow \{\bar{u} \bar{v} : \Delta\} \rightarrow (\bar{e} : \bar{u} \equiv_{\Delta} \bar{v}) \rightarrow f \bar{u} \equiv_{T \bar{e}} f \bar{v} \end{aligned} \quad (8)$$

3 Unification in dependent type theory

In this section, we describe our new framework for unification of dependently typed data. First, we represent the input of the unification problem by a *telescopic equality* where each type in this telescope corresponds to one equation of the unification problem (Section 3.1). The output of the unification algorithm is then an equivalence between the original telescope of equations and a trivial one (Section 3.2). This equivalence contains not only the substitution computed by the unification algorithm, but also *evidence* that the output is sound. The unification algorithm itself works by successively applying unification rules, which are represented by equivalences between two telescopes of equations (Section 3.3). The aggregate equivalence produced by unification can be used for *specialization by unification*, an essential part of the translation of definitions by dependent pattern matching to eliminators (Section 3.4).

3.1 Unification problems as telescopes

To represent equations internally, we use the propositional equality type $x \equiv_A y$. For example, the equation $\text{succ } m = \text{succ } n$ is represented by the type $\text{succ } m \equiv_{\mathbb{N}} \text{succ } n$.

In general, a unification problem can consist of multiple equalities and the type of an equality may depend on the solution of the previous equalities. To keep track of these dependencies, we give a type to the list of equations in the form of a telescope. By the nature of a telescope, the type of each equation can depend on the previous equations.

Definition 13. A *unification problem* is a telescope of the form $\Gamma(\bar{e} : \bar{u} \equiv_{\Delta} \bar{v})$. The variables in Γ are called the *flexible variables*.

Example 14. A unification problem between two vectors can be represented by the telescope

$$(m\ n : \mathbb{N})(x\ y : A)(xs : \mathbf{Vec}\ A\ m)(ys : \mathbf{Vec}\ A\ n) \\ (e_1 : \mathbf{succ}\ m \equiv_{\mathbb{N}} \mathbf{succ}\ n)(e_2 : \mathbf{cons}\ m\ x\ xs \equiv_{\mathbf{Vec}\ A\ e_1} \mathbf{cons}\ n\ y\ ys) \quad (9)$$

3.2 Unifiers as equivalences

Traditionally, a unifier for a unification problem $\bar{u} = \bar{v}$ is defined as a substitution σ such that $\bar{u}\sigma$ and $\bar{v}\sigma$ are equal. So how do we translate this definition to type theory? Consider a unification problem of the form $\Gamma(\bar{e} : \bar{u} \equiv_{\Delta} \bar{v})$. We could represent a unifier as a telescope mapping $\sigma : \Gamma' \rightarrow \Gamma$ satisfying $\bar{u}[\Gamma \mapsto \sigma \Gamma] = \bar{v}[\Gamma \mapsto \sigma \Gamma]$, but then the soundness property is still external to the theory. Instead, we use the power of dependent types to express the fact that the equations are satisfied internally:

Definition 15. Let Γ and Δ be telescopes and $\bar{u}$ and $\bar{v}$ be lists of terms such that $\Gamma \vdash \bar{u}, \bar{v} : \Delta$. We define a *unifier* of $\bar{u}$ and $\bar{v}$ as a telescope mapping $\sigma : \Gamma' \rightarrow \Gamma(\bar{e} : \bar{u} \equiv_{\Delta} \bar{v})$ for some Γ' .

A unifier σ returns not only values of type Γ but also *evidence* that the equations are indeed satisfied by these values.

Example 16. The telescope mapping $\sigma : (k : \mathbb{N}) \rightarrow (k\ l : \mathbb{N})(e : \mathbf{succ}\ k \equiv_{\mathbb{N}} \mathbf{succ}\ l)$ defined by $\sigma = \lambda k.k; k; \mathbf{refl}$ is a unifier of $\mathbf{succ}\ k$ and $\mathbf{succ}\ l$. The evidence here is $\mathbf{refl} : \mathbf{succ}\ k = \mathbf{succ}\ k$, proving that $\mathbf{succ}\ k$ and $\mathbf{succ}\ l$ become equal under the substitution $[k \mapsto k; l \mapsto k]$.

Usually, the goal of a unification algorithm is not just to output any unifier but a *most general* one, i.e. a unifier $\sigma : \Gamma' \rightarrow \Gamma(\bar{e} : \bar{u} \equiv_{\Delta} \bar{v})$ such that any other unifier $\sigma' : \Gamma'' \rightarrow \Gamma(\bar{e} : \bar{u} \equiv_{\Delta} \bar{v})$ can be written as $\sigma \circ \nu$ for some $\nu : \Gamma'' \rightarrow \Gamma'$.

Again, we should think how to represent this concept internally. One way to do this is to translate the definition of most general unifier directly to a type. However, to do this we need to quantify over all possible telescopes Γ'' and unifiers σ' , making the definition more unwieldy than necessary. Can we find a better definition?

Lemma 17. Let Γ and Γ' be telescopes and $\sigma : \Gamma' \rightarrow \Gamma$. The following two statements are equivalent:

- For any telescope Γ'' and $\sigma' : \Gamma'' \rightarrow \Gamma$, there exists at least one $\nu : \Gamma'' \rightarrow \Gamma'$ such that $\sigma' \doteq \sigma \circ \nu$.
- There exists a $\tau_1 : \Gamma \rightarrow \Gamma'$ such that $\sigma \circ \tau_1 \doteq \mathbf{id}$.

Figure 6: Lemma 17 allows us to construct a right inverse τ_1 to σ from the existence of the telescope mapping ν , while Lemma 18 gives us a left inverse τ_2 from its uniqueness.

Proof. First suppose that we have a telescope mapping $\tau_1 : \Gamma \rightarrow \Gamma'$ such that $\sigma \circ \tau_1 \doteq \text{id}$ is the identity function on Γ . This allows us to define $\nu = \tau_1 \circ \sigma'$, which gives us $\sigma \circ \nu \doteq \sigma \circ \tau_1 \circ \sigma' \doteq \sigma'$, as we wanted.

For the other direction, we take $\Gamma'' = \Gamma$ and $\sigma' = \text{id}$. Then by assumption we have a $\tau_1 : \Gamma \rightarrow \Gamma'$ such that $\text{id} \doteq \sigma \circ \tau_1$. $\square$

It is often useful to require that the function ν is unique, for otherwise Γ' may contain ghost variables that are not actually used by σ . For example, for a unification problem with $\Gamma = (b : \text{Bool})$ and a single equation $b \equiv_{\text{Bool}} \text{true}$, we have the unifier $\sigma : () \rightarrow (b : \text{Bool})(e : b \equiv_{\text{Bool}} \text{true})$ that we would like to recognize as the most general one. However, if ν is not required to be unique, then there may be other most general unifiers with a non-equivalent choice of Γ' . For example, we could also have taken $\sigma' : (b' : \text{Bool}) \rightarrow (b : \text{Bool})(e : b \equiv_{\text{Bool}} \text{true})$ that ignores its argument b' .

Lemma 18. *Let Γ and Γ' be telescopes and let $\sigma : \Gamma' \rightarrow \Gamma$. The following two statements are equivalent:*

- *For any telescope Γ'' and $\sigma' : \Gamma'' \rightarrow \Gamma$, there exists at most one $\nu : \Gamma'' \rightarrow \Gamma'$ such that $\sigma' \doteq \sigma \circ \nu$.*
- *There exists a $\tau_2 : \Gamma \rightarrow \Gamma'$ such that $\tau_2 \circ \sigma \doteq \text{id}$.*

Proof. Suppose that we have a τ_2 such that $\tau_2 \circ \sigma \doteq \text{id}$. If ν and ν' are two telescope mappings such that $\sigma \circ \nu \doteq \sigma' \doteq \sigma \circ \nu'$ then we have $\nu \doteq \tau_2 \circ \sigma \circ \nu \doteq \tau_2 \circ \sigma \circ \nu' \doteq \nu'$, so ν is unique.

For the other direction, we assume that ν is unique. Let $\Gamma'' = \Gamma'$ and $\sigma' = \sigma$ and $\nu = \tau_1 \circ \sigma$ and $\nu' = \text{id}$. This gives us that $\sigma \circ \nu \doteq \sigma \doteq \sigma \circ \nu'$, so by uniqueness we have $\tau_1 \circ \sigma \doteq \text{id}$. Hence taking $\tau_2 = \tau_1$ gives us the desired telescope mapping τ_2 . $\square$

The proofs of Lemma 17 and Lemma 18 are illustrated in Figure 6. If we replace the telescope Γ by a unification problem $\Gamma(\bar{e} : \bar{u} \equiv_{\Delta} \bar{v})$, then Lemma 17 and Lemma 18 together give us that $\sigma : \Gamma(\bar{e} : \bar{u} \equiv_{\Delta} \bar{v}) \rightarrow \Gamma'$ is a most general unifier if and only if it is an *equivalence* between Γ' and $\Gamma(\bar{e} : \bar{u} \equiv_{\Delta} \bar{v})$. This brings us to the following definition of a most general unifier:

Definition 19. Let Γ and Δ be telescopes and $\bar{u}$ and $\bar{v}$ be lists of terms such that $\Gamma \vdash \bar{u}, \bar{v} : \Delta$. Then a *most general unifier* of $\bar{u}$ and $\bar{v}$ is an equivalence $f : \Gamma(\bar{e} : \bar{u} \equiv_{\Delta} \bar{v}) \simeq \Gamma'$ for some telescope Γ' .

The unifier $\sigma : \Gamma' \rightarrow \Gamma(\bar{e} : \bar{u} \equiv_{\Delta} \bar{v})$ corresponds to the inverse function f^{-1} . Intuitively, f allows us to recover the values of the variables in Γ' for any values of Γ that satisfy $\bar{u} \equiv_{\Delta} \bar{v}$.

The definition of a most general unifier doesn't prevent us from choosing $\Gamma' = \Gamma(\bar{e} : \bar{u} \equiv_{\Delta} \bar{v})$ and $f = \text{id}$. In fact, this is a valid (if trivial) most general unifier from a logical point of view.

In case unification succeeds negatively, we need evidence that the equations are indeed impossible. For this purpose, we use the empty type $\perp$:

Definition 20. Let Γ and Δ be telescopes and $\bar{u}$ and $\bar{v}$ be lists of terms such that $\Gamma \vdash \bar{u}, \bar{v} : \Delta$. A *disunifier* of $\bar{u}$ and $\bar{v}$ is an equivalence $f : \Gamma(\bar{e} : \bar{u} \equiv_{\Delta} \bar{v}) \simeq \perp$.

Any function $f : A \rightarrow \perp$ is automatically an equivalence $A \simeq \perp$, as the other components of the equivalence can be constructed by using the eliminator $\text{elim}_{\perp} : (A : \text{Set}_{\ell}) \rightarrow \perp \rightarrow A$. So the only interesting part of a disunifier is the function $f : A \rightarrow \perp$.

3.3 The unification algorithm

Now that we know how to represent the input and the output of the unification algorithm, we can start thinking about the unification algorithm itself. Since the end result of the unification process (the most general unifier) is an equivalence, it is natural to represent unification rules as equivalences as well. These unification rules can then be chained together by transitivity of $\simeq$ to produce the most general unifier f .

Definition 21. A *positive unification rule* is an equivalence of the form $r : \Gamma(\bar{e} : \bar{u} \equiv_{\Delta} \bar{v}) \simeq \Gamma'(\bar{e}' : \bar{u}' \equiv_{\Delta'} \bar{v}')$.

For example, the unification rule for injectivity of the `suc` constructor for $\mathbb{N}$ is

$$\text{injectivity}_{\text{suc}} : (e : \text{suc } m \equiv \text{suc } n) \simeq (e : m \equiv n) \quad (10)$$

Another important example of a positive unification rule is the solution rule used to solve equations where one side is a variable:

$$\text{solution} : (x : A)(e : x \equiv_A t) \simeq () \quad (11)$$

In addition to unification rules of this form, that transform one set of equations into another, there are also unification rules that refute absurd equations like `true` $\equiv_{\text{Bool}}$ `false`.

Definition 22. A *negative unification rule* is an equivalence of the form $r : \Gamma(\bar{e} : \bar{u} \equiv_{\Delta} \bar{v}) \simeq \perp$.

For example, the unification rule for conflict between `true` and `false` is

$$\text{conflict}_{\text{true}, \text{false}} : (\text{true} \equiv_{\text{Bool}} \text{false}) \simeq \perp \quad (12)$$

The unification algorithm tries to construct an equivalence $\Gamma(\bar{e} : \bar{u} \equiv_{\Delta} \bar{v}) \simeq \Gamma'$ by successively applying the unification rules to the unification problem, simplifying one or more equations in each step. This process continues until one of three possible situations occurs:

- If there are no more equations left, the algorithm succeeds positively. In this case, it returns an equivalence between the original problem $\Gamma(\bar{u} \equiv_{\Delta} \bar{v})$ and the reduced telescope Γ' .
- If a contradictory equation is encountered, the algorithm succeeds negatively. In this case, it returns an equivalence between the original problem $\Gamma(\bar{u} \equiv_{\Delta} \bar{v})$ and the empty type $\perp$.
- If there are no more applicable rules, the algorithm results in a failure.

We do not yet give an explicit strategy on which rule to apply in a specific situation. This leaves more freedom to the implementation to choose which rule to try first. When we discuss our implementation in Section 8, we give one concrete strategy.

Example 23. Consider the unification problem consisting of flexible variables $k : \mathbb{N}$ and a single equation between $\text{succ } k$ and $\text{succ } l$. First, we simplify the equation by applying the equivalence $\text{injectivity}_{\text{succ}} : (e : \text{succ } k \equiv_{\mathbb{N}} \text{succ } l) \simeq (e : k \equiv_{\mathbb{N}} l)$. Applying this rule leaves the two variables k and l unchanged. Next, we apply the solution rule, which tells us that $(l : \mathbb{N})(e : k \equiv_{\mathbb{N}} l) \simeq ()$. This leaves only the single variable $k : \mathbb{N}$. Since there are no more equations left in the telescope, unification ends in a positive success.

We write down the unification process as a series of telescopes (representing unification problems) with equivalences between them (representing the individual unification steps). In each step, we underline the variables and equations that are being solved or simplified.

$$\begin{aligned} (k \ l : \mathbb{N})(\underline{e : \text{succ } k \equiv_{\mathbb{N}} \text{succ } l}) \\ \simeq (k \ \underline{l : \mathbb{N}})(\underline{e : k \equiv_{\mathbb{N}} l}) \\ \simeq (k : \mathbb{N}) \end{aligned} \tag{13}$$

To get the substitution from $(k : \mathbb{N})$ to $(k \ l : \mathbb{N})(e : \text{succ } k \equiv_{\mathbb{N}} \text{succ } l)$ computed by the unification process, we only need to compose the functions embedded in the equivalences from the bottom to the top. The solution rule assigns l to be k and e to be refl , and the injectivity rule maps $e : k \equiv_{\mathbb{N}} l$ to $\text{cong succ } e : \text{succ } k \equiv_{\mathbb{N}} \text{succ } l$, so the aggregate unifier $f^{-1} : (k : \mathbb{N}) \rightarrow (k \ l : \mathbb{N})(\underline{e : \text{succ } k \equiv_{\mathbb{N}} \text{succ } l})$ here is $\lambda k. k; k; \text{refl}$ (since $\text{cong succ refl} = \text{refl}$).

Before we continue with the general form of the unification rules in the next section, we first give three easy but useful manipulations on equivalences (and hence on unification rules) that allow us to postpone and reorder equations. These principles are used when we want to apply a unification rule, but the problem contains some additional variables or equations that aren't mentioned in the rule. For example, in the second step of Example 23 the solution rule didn't affect the variable k . In what follows, we often make use of these manipulations implicitly.

Lemma 24. *If we have an equivalence $f : \Delta \simeq \Delta'$ where Δ and Δ' possibly contain free variables from a telescope Γ , then we also have an equivalence $f_\Gamma : \Gamma\Delta \simeq \Gamma\Delta'$.*

Lemma 25. *If we have an equivalence $f : \Gamma \simeq \Gamma'$ and a telescope Δ possibly containing free variables from Γ , then we also have an equivalence $f^\Delta : \Gamma\Delta \simeq \Gamma'\Delta'$ where $\Delta' = \Delta[\Gamma \mapsto \text{linv } f \Gamma']$.*

Lemma 26. *If we have a telescope Γ , and Γ' is a reordering of the variable bindings in Γ that preserves the order of dependencies, then we have an equivalence $f : \Gamma \simeq \Gamma'$.*

Construction. The construction of the equivalence is in all three cases straightforward, relying on **J** to prove that the functions are mutual inverses. $\square$

3.4 Specialization by unification

When performing case analysis on a variable from an inductive family, the type-checker of a dependently typed language needs to determine which constructors can occur in a given position and how the variables need to be instantiated for the pattern to be well-typed. To do this, it applies unification to the indices of the datatype in question. If unification determines that there can be no such substitution, then we can skip the case for the corresponding constructor. This method of solving equations to either gain more information about the type of the right-hand side or to derive an absurdity is called *specialization by unification*.

Specialization by unification allows us to construct functions of the form

$$m : (\bar{x} : \Gamma)(\bar{e} : \bar{u} \equiv_\Delta \bar{v}) \rightarrow T \bar{x} \bar{e} \quad (14)$$

It can be seen as a generic method of constructing an *inversion principle* (McBride, 1998a). It is also a core component in the translation of definitions by pattern matching to eliminators (Goguen et al., 2006; Cockx, Devriese, and Piessens, 2016b).

Definition 27 (Specialization by unification). Consider a problem of the form $m : (\bar{x} : \Gamma)(\bar{e} : \bar{u} \equiv_\Delta \bar{v}) \rightarrow T \bar{x} \bar{e}$ and suppose that unification of $\bar{u}$ with $\bar{v}$ with Γ as flexible variables succeeds either positively or negatively, then we construct the function m :

- In case the unification succeeds positively with most general unifier $f : \Gamma(\bar{e} : \bar{u} \equiv_\Delta \bar{v}) \simeq \Gamma'$, we have

$$\text{isLinv } f : (\bar{x} : \Gamma)(\bar{e} : \bar{u} \rightarrow f^{-1}(f \bar{x} \bar{e}) \equiv_{\Gamma(\bar{e} : \bar{u} \bar{x}; \bar{e})} \bar{v}) \quad (15)$$

Then we define m by

$$m \bar{x} \bar{e} = \overline{\text{subst}} T (\text{isLinv } f \bar{x} \bar{e}) (m^s (f \bar{x} \bar{e})) \quad (16)$$

with the new subgoal of constructing $m^s : (\bar{x}' : \Gamma') \rightarrow T (\text{linv } f \bar{x}')$.

- In case the unification succeeds negatively with disunifier $f : \Gamma(\bar{e} : \bar{u} \equiv_\Delta \bar{v}) \simeq \perp$, then we define m by

$$m \bar{x} \bar{e} = \text{elim}_\perp (T \bar{x} \bar{e}) (f \bar{x} \bar{e}) \quad (17)$$

with no additional assumptions.

Example 28. We apply specialization by unification to construct two functions

$$\begin{aligned} m_{\mathbf{1z}} : (m : \mathbb{N})(k : \mathbb{N})(y : k \leq \mathbf{zero}) \rightarrow \\ (\mathbf{zero}; m; \mathbf{1z} \ m) \equiv_{(m : \mathbb{N})(x : m \leq n)} (k; \mathbf{zero}; y) \rightarrow \mathbf{zero} \equiv_{\mathbb{N}} k \end{aligned} \quad (18)$$

and

$$\begin{aligned} m_{\mathbf{1s}} : (m : \mathbb{N})(x : m \leq n)(k : \mathbb{N})(y : k \leq \mathbf{zero}) \rightarrow \\ (\mathbf{suc} \ m; \mathbf{suc} \ n; \mathbf{1s} \ m \ n \ x) \equiv_{(m : \mathbb{N})(x : m \leq n)} (k; \mathbf{zero}; y) \rightarrow \mathbf{zero} \equiv_{\mathbb{N}} k \end{aligned} \quad (19)$$

In case of $m_{\mathbf{1z}}$, unification of $\bar{u} = \mathbf{zero}; m; \mathbf{1z} \ m$ with $\bar{v} = k; \mathbf{zero}; y$ in the context $\Gamma = (m : \mathbb{N})(k : \mathbb{N})(y : k \leq \mathbf{zero})$ results in a positive success with most general unifier $f : \Gamma(\bar{u} \equiv_{(m : \mathbb{N})(x : m \leq n)} \bar{v}) \simeq ()$ with $f^{-1}() = \mathbf{zero}; \mathbf{zero}; \mathbf{1z} \ \mathbf{zero}; \mathbf{refl}$, where $\mathbf{refl} : \mathbf{zero}; \mathbf{zero}; \mathbf{1z} \ \mathbf{zero} \equiv_{(m : \mathbb{N})(x : m \leq n)} \mathbf{zero}; \mathbf{zero}; \mathbf{1z} \ \mathbf{zero}$. Specialization by unification gives us the function $m_{\mathbf{1z}}$ on the condition we can construct $m_{\mathbf{1z}}^s : \mathbf{zero} \equiv_{\mathbb{N}} \mathbf{zero}$, which we can do easily as $m_{\mathbf{1z}}^s = \mathbf{refl}$.

For $m_{\mathbf{1s}}$, unification of $\bar{u} = \mathbf{suc} \ m; \mathbf{suc} \ n; \mathbf{1s} \ m \ n \ x$ with $\bar{v} = k; \mathbf{zero}; y$ results in a negative success, so specialization by unification gives us the function $m_{\mathbf{1s}}$ without any additional assumptions.

4 Unification rules

In this section, we state the basic unification rules from McBride (1998b) in our framework. We first handle the unification rules for simple datatypes (Section 4.1) before moving on to the more challenging rules for indexed datatypes (Section 4.2). We also show how to extend the unification algorithm with rules for η -equality for record types (Section 4.3).

4.1 The basic unification rules

The first two rules are generic in the sense that they work for any type A .

Lemma 29. For any type A and term $t : A$, we have an equivalence

$$\mathbf{solution} : (x : A)(e : x \equiv_A t) \simeq () \quad (20)$$

satisfying $\mathbf{solution}^{-1}() = t; \mathbf{refl}$.

In this rule, the variable x should not occur freely in t .

Construction of $\mathbf{solution}$. The construction of the functions $\mathbf{solution} : (x : A)(x \equiv_A t) \rightarrow ()$ and $\mathbf{isRinv} \ \mathbf{solution} : () \rightarrow () \equiv_{()} ()$ is trivial since they both target an empty telescope. The function $\mathbf{solution}^{-1} : () \rightarrow (x : A)(e : x \equiv_A t)$ is defined by $\mathbf{solution}^{-1}() = t; \mathbf{refl}$. Finally, $\mathbf{isLinv} \ \mathbf{solution}$, of type $(x : A)(e : x \equiv_A t) \rightarrow t; \mathbf{refl} \equiv_{(x : A)(e : x \equiv_A t)} x; e$, is a direct application of the **J** rule. $\square$

Lemma 30. For any type A that satisfies UIP and any term $t : A$, we have an equivalence

$$\mathbf{deletion} : (e : t \equiv_A t) \simeq () \quad (21)$$

$$\frac{}{t_i \ s_1 \ \dots \ s_k \prec \mathbf{c} \ t_1 \ \dots \ t_n} \qquad \frac{s \prec t_i}{s \prec \mathbf{c} \ t_1 \ \dots \ t_n}$$

Figure 7: The structural order $\prec$ is used to check termination and to detect cycles during unification.

Construction of deletion. This follows directly from the uniqueness of identity proofs at type A . $\square$

The **injectivity**, **conflict**, and **cycle** rules are specific to an inductive datatype D . We present them here for a simple (non-indexed) datatype and for an indexed datatype in the next section.

Lemma 31. *Let $\mathbf{c} : \Delta_{\mathbf{c}} \rightarrow D$ be a constructor of the datatype $D : \mathbf{Set}_\ell$ and let $\bar{x}, \bar{x}' : \Delta_{\mathbf{c}}$. We have an equivalence*

$$\mathbf{injectivity}_{\mathbf{c}} : (\mathbf{c} \ \bar{x} \equiv_D \mathbf{c} \ \bar{x}') \simeq (\bar{x} \equiv_{\Delta_{\mathbf{c}}} \bar{x}') \quad (22)$$

such that $\mathbf{injectivity}_{\mathbf{c}}^{-1} \bar{e} = \overline{\mathbf{cong}} \ \mathbf{c} \ \bar{e}$.

Lemma 32. *Let $\mathbf{c}_1 : \Delta_1 \rightarrow D$ and $\mathbf{c}_2 : \Delta_2 \rightarrow D$ be two distinct constructors of the datatype $D : \mathbf{Set}_\ell$ and $\bar{x}_1 : \Delta_1$ and $\bar{x}_2 : \Delta_2$. We have an equivalence*

$$\mathbf{conflict}_{\mathbf{c}_1, \mathbf{c}_2} : (\mathbf{c}_1 \ \bar{x}_1 \equiv_D \mathbf{c}_2 \ \bar{x}_2) \simeq \perp \quad (23)$$

To state the **cycle** rule, we need the structural order $\prec$ defined in Figure 7. This definition is somewhat different from the one given by Goguen et al. (2006). It describes the same relation in case the left- and right-hand sides are elements of the datatype D , but it enforces that the left- and right-hand sides are actually elements of the datatype. This prevents odd structural orders that are allowed by the original definition. For example, if we have a datatype D with a constructor $\mathbf{c} : (A \rightarrow D) \rightarrow D$, then the definition by Goguen et al. allows us to derive for any $a : A$ that $a \prec \mathbf{c} \ a \prec \mathbf{c}$, even though a doesn't occur in $\mathbf{c}$.

Lemma 33. *Let $D : \mathbf{Set}_\ell$ be a datatype and let $x, t : D$ be such that $x \prec t$. We have an equivalence*

$$\mathbf{cycle}_{x, t} : (x \equiv_D t) \simeq \perp \quad (24)$$

Once again, the type of the equation should be exactly D . We postpone the proof of this lemma until Section 7.

Example 34. Consider the sum type $A \uplus B$ (where $A, B : \mathbf{Set}$ are arbitrary types) with two constructors **left** : $A \rightarrow A \uplus B$ and **right** : $B \rightarrow A \uplus B$. An expression of the form **left** x is never equal to **right** y , so any equality between those two terms is equivalent to $\perp$:

$$(x : A)(y : B)(\underline{e} : \mathbf{left} \ x \equiv_{A \uplus B} \mathbf{right} \ y) \simeq \perp \quad (25)$$

This is exactly the conflict rule between **left** and **right**.

The type of the equation on the left in Lemma 31 and Lemma 32 should be exactly D , in particular the constructor $\mathbf{c}$ must be fully applied.

Counterexample 35. An equation between the constructors **left** and **right** is not always absurd when they are not fully applied. Let $A = B = \perp$, then $(e : \text{left} \equiv \perp \rightarrow \perp \uplus \text{right})$ is not equivalent to $\perp$. This is because when viewed as functions of type $\perp \rightarrow \perp \uplus \perp$, the constructors **inj₁** and **inj₂** coincide on all possible inputs (i.e. none). The principle of functional extensionality then tells us that these two functions are equal. So if we would consider this equation to be absurd, we would prohibit ourselves from having a general rule for functional extensionality in our language, nevertheless a desirable property to have! Wrongly applying the conflict rule in this way led to the problem described by issue #1497 on the Agda bug tracker (Dijkstra, 2015).

4.2 Rules for indexed datatypes

The injectivity, conflict, and cycle rules defined in the previous section all work on regular datatypes, but unification only becomes really interesting once we consider indexed families of datatypes. Where the unification rules that we have seen so far only have a single equation on the left side, the rules for indexed datatypes have a telescope of equations: one equation for each index, and one final equation for the datatype itself.

Lemma 36. Let $c : \Delta \rightarrow \mathbf{D} \bar{u}$ be a constructor of the datatype $\mathbf{D} : \Xi \rightarrow \mathbf{Set}_\ell$. Then we have an equivalence

$$\text{injectivity}_c : (\bar{u}[\Delta \mapsto \bar{x}]; c \bar{x} \equiv_{\bar{\mathbf{D}}} \bar{u}[\Delta \mapsto \bar{x}']; c \bar{x}') \simeq (\bar{x} \equiv_{\Delta} \bar{x}') \quad (26)$$

where $\bar{x}, \bar{x}' : \Delta$ and $\text{injectivity}_c^{-1} \bar{e} = \overline{\text{dcong}}(\lambda \bar{x}. \bar{u}; c \bar{x}) \bar{e}$.

Lemma 37. Let $c_1 : \Delta_1 \rightarrow \mathbf{D} \bar{u}_1$ and $c_2 : \Delta_2 \rightarrow \mathbf{D} \bar{u}_2$ be two distinct constructors of the datatype $\mathbf{D} : \Xi \rightarrow \mathbf{Set}_\ell$. Then we have an equivalence

$$\text{conflict}_{c_1, c_2} : (\bar{u}_1[\Delta_1 \mapsto \bar{x}_1]; c_1 \bar{x}_1 \equiv_{\bar{\mathbf{D}}} \bar{u}_2[\Delta_2 \mapsto \bar{x}_2]; c_2 \bar{x}_2) \simeq \perp \quad (27)$$

where $\bar{x}_1 : \Delta_1$ and $\bar{x}_2 : \Delta_2$.

Lemma 38. Let $\mathbf{D} : \Xi \rightarrow \mathbf{Set}_\ell$ be a datatype and let $(\bar{u}; x), (\bar{v}; t) : \bar{\mathbf{D}}$ be such that $x \prec t$. Then we have an equivalence

$$\text{cycle}_{x, t} : (\bar{u}; x \equiv_{\bar{\mathbf{D}}} \bar{v}; t) \simeq \perp \quad (28)$$

Again, we postpone the proof of this lemma until Section 7. The unification rules are summarized in Figure 8. To these basic rules, we will add rules for η -equality for record types (Section 4.3, generalized rules for conflict and acyclicity (Section 6.1), and higher-dimensional unification (Section 6.2). We give a complete list of the unification rules with all these additions in Figure 9.

Example 39. Consider the indexed datatype $\text{Vec } A : \mathbb{N} \rightarrow \mathbf{Set}$ with the two constructors **nil** : $\text{Vec } A$ **zero** and **cons** : $(n : \mathbb{N}) \rightarrow A \rightarrow \text{Vec } A \ n \rightarrow \text{Vec } A \ (\text{succ } n)$. The injectivity rule for **cons** gives us the following equivalence:

$$\begin{aligned} & (\text{succ } m; \text{cons } m \ x \ xs \equiv_{\text{Vec } A} \text{succ } n; \text{cons } n \ y \ ys) \\ & \simeq (m; x; xs \equiv_{(n:\mathbb{N})(x:A)(xs:\text{Vec } A \ n)} n; y; ys) \end{aligned} \quad (29)$$

$$\begin{aligned}
\text{solution} &: (x : A)(e : x \equiv_A u) \simeq () & (\text{where } x \notin FV(u)) \\
\text{deletion} &: (e : u \equiv_A u) \simeq () \\
\text{injectivity}_{\mathbf{c}} &: (\bar{u}[\Delta \mapsto \bar{x}]; \mathbf{c} \bar{x} \equiv_{\bar{D}} \bar{u}[\Delta \mapsto \bar{x}']; \mathbf{c} \bar{x}') \simeq (\bar{x} \equiv_{\Delta} \bar{x}') \\
\text{conflict}_{\mathbf{c}_1, \mathbf{c}_2} &: (\bar{u}_1[\Delta_1 \mapsto \bar{x}_1]; \mathbf{c}_1 \bar{x}_1 \equiv_{\bar{D}} \bar{u}_2[\Delta_2 \mapsto \bar{x}_2]; \mathbf{c}_2 \bar{x}_2) \simeq \perp & (\text{where } \mathbf{c}_1 \neq \mathbf{c}_2) \\
\text{cycle}_{x,t} &: (\bar{u}; x \equiv_{\bar{D}} \bar{v}; t) \simeq \perp & (\text{where } x \prec t)
\end{aligned}$$

Figure 8: The basic unification rules can be formulated as equivalences.

This rule not only simplifies the equation between the two **cons** constructors, but simultaneously simplifies the equation between the indices **suc** m and **suc** n . Now let's see how this rule works in action:

$$\begin{aligned}
& (m : \mathbb{N})(x y : A)(xs : \mathbf{Vec} A m)(ys : \mathbf{Vec} A n) \\
& (e_1 : \mathbf{suc} m \equiv_{\mathbb{N}} \mathbf{suc} n)(e_2 : \mathbf{cons} m x xs \equiv_{\mathbf{Vec} A} e_1 \mathbf{cons} n y ys) \\
& \simeq (\underline{m} n : \mathbb{N})(x y : A)(xs : \mathbf{Vec} A m)(ys : \mathbf{Vec} A n) \\
& (e_1 : m \equiv_{\mathbb{N}} n)(e_2 : x \equiv_A y)(e_3 : xs \equiv_{\mathbf{Vec} A} e_1 ys) \\
& \simeq (n : \mathbb{N})(x : A)(xs : \mathbf{Vec} A n)
\end{aligned} \tag{30}$$

The first step is an application of the injectivity rule, while the next step consists of three applications of the solution rule.

To apply injectivity of **cons**, the type of the equation has to be of the form $\mathbf{Vec} A e$ where e refers to a previous equation. This implies that this rule cannot be applied directly to an equation of the form $\mathbf{cons} n x xs \equiv_{\mathbf{Vec} A} (\mathbf{suc} n) \mathbf{cons} n y ys$ where $xs : \mathbf{Vec} A n$ and $ys : \mathbf{Vec} A n$ have the same length ‘on the nose’. We show how to solve this deficiency in Section 6.

Example 40. In the previous example, it was not really necessary to simplify the equation between the indices together with the equation between the constructors, as we could also have applied **injectivity**_{**suc**} to the equation $\mathbf{suc} m \equiv_{\mathbb{N}} \mathbf{suc} n$. However, sometimes this simplification gives a real increase to the power of unification. For example, let $f : A \rightarrow B$ be a (possibly very complex) function, then in general there is no way to solve an equation of the form $f x \equiv_B f y$. Now consider the following datatype:

$$\begin{aligned}
& \text{data } \mathbf{Im} f : B \rightarrow \mathbf{Set} \text{ where} \\
& \quad \text{image} : (x : A) \rightarrow \mathbf{Im} f (f x)
\end{aligned} \tag{31}$$

The injectivity rule for **image** simultaneously solves the equations $e_1 : f x \equiv_B f y$ and $e_2 : \mathbf{image} x \equiv_{\mathbf{Im} f} e_1 \mathbf{image} y$:

$$\begin{aligned}
& (x y : A)(e_1 : f x \equiv_B f y)(e_2 : \mathbf{image} x \equiv_{\mathbf{Im} f} e_1 \mathbf{image} y) \\
& \simeq (x y : A)(e : x \equiv_A y) \\
& \simeq (x : A)
\end{aligned} \tag{32}$$

Having an injectivity rule that works in this way is useful when giving semantics to an embedded language (Danielsson, 2015).

Contrast this with the unification problem

$$(x\ y : A)(e_1 : \text{Im } f(f\ x) \equiv_{\text{Set}} \text{Im } f(f\ y))(e_2 : \text{image } x \equiv_{e_1} \text{image } y) \quad (33)$$

Here, it is not allowed to use injectivity on the second equation e_2 since its type is not a datatype but a variable e_1 . Like in Example 2, there is no way to distinguish between these two cases unless we keep track of the dependency of the type of e_2 on the equation e_1 . Wrongly applying injectivity in situations like this led to the problems described by Abel (2015a,c) on the Agda bug tracker.

Example 41. Let $D : \text{Bool} \rightarrow \text{Set}$ be an indexed datatype with two constructors $\text{tt} : D\ \text{true}$ and $\text{ff} : D\ \text{false}$. Then the conflict rule between tt and ff gives us the following equivalence:

$$(e_1 : \text{true} \equiv_{\text{Bool}} \text{false})(e_2 : \text{tt} \equiv_{D\ e_1} \text{ff}) \simeq \perp \quad (34)$$

On the other hand, the conflict rule cannot be applied if the first equation is between the *types* $D\ \text{true}$ and $D\ \text{false}$:

$$(e_1 : D\ \text{true} \equiv_{\text{Set}} D\ \text{false})(e_2 : \text{tt} \equiv_{e_1} \text{ff}) \not\simeq \perp \quad (35)$$

Allowing the conflict rule to apply in this case would mean that we can distinguish between $D\ \text{true}$ and $D\ \text{false}$, which means that the type constructor D is injective. In particular, this would be incompatible with univalence: there is an equivalence between $D\ \text{true}$ and $D\ \text{false}$ under which tt is identified with ff , so univalence allows us to prove that $D\ \text{true} \equiv_{\text{Set}} D\ \text{false}$. Note again that we need information about how the type of e_2 depends on e_1 to distinguish between these two cases. Wrongly applying conflict in situations like this led to the problems described by Danielsson (2010) and Vezzosi (2015) on the Agda bug tracker.

Example 42. This example is based on issue #1071 on the Agda bug tracker (Danielsson, 2014). Let $A : \text{Set}$, $F : \text{Set} \rightarrow \text{Set}$ and $P : \text{Set} \rightarrow \text{Set}_1$ be a datatype with one constructor $\text{c} : (A : \text{Set}) \rightarrow P\ (F\ A)$. Then we have:

$$\begin{aligned} & (f : F\ A)(R : \text{Set})(f' : F\ R) \\ & (e_1 : F\ A \equiv_{\text{Set}} F\ R)(e_2 : f \equiv_{e_1} f')(e_3 : \text{c } A \equiv_{P\ e_1} \text{c } R) \\ & \simeq (f : F\ A)(R : \text{Set})(f' : F\ R)(e'_3 : A \equiv_{\text{Set}} R)(e_2 : f \equiv_{F\ e'_3} f') \\ & \simeq (f : F\ A)(f' : F\ A)(e_2 : f \equiv_{F\ A} f') \\ & \simeq (f : F\ A) \end{aligned} \quad (36)$$

At each point during the unification process, there is only one valid way to proceed. At the first step, the second equation $f \equiv_{e_1} f'$ cannot be solved right away as the type is heterogeneous and the solution rule only applies to homogeneous equations. The first equation cannot be solved either as this would require injectivity of the functor F . The only possibility is to apply the injectivity of c to the third equation. At the second step, $f \equiv_{F\ e'_3} f'$ cannot be solved because the type $F\ e'_3$ is heterogeneous, so e'_3 has to be solved first instead.

4.3 Rules for record types

One of the big advantages of having a general notion of ‘unification rule’ and ‘most general unifier’ is that we have an easy way to check the soundness of new unification rules. Alternatively it can be used to assess the impact of adding a new unification rule to the algorithm. In this section, we extend our algorithm with two unification rules that deal with η -equality for record types.

A *record type* is a type for grouping values together. One of the properties that sets a record type apart from a regular datatype with a single constructor, are the additional laws for equality of records called η -laws (not to be confused with the η -law for functions).

Definition 43. A *record type* $\mathbf{R} : \mathbf{Set}_\ell$ is defined by a number of *fields* (also called *projections*):

$$\begin{aligned} \mathbf{f}_1 &: (r : \mathbf{R}) \rightarrow A_1 \\ \mathbf{f}_2 &: (r : \mathbf{R}) \rightarrow A_2 (\mathbf{f}_1 r) \\ &\vdots \\ \mathbf{f}_n &: (r : \mathbf{R}) \rightarrow A_n (\mathbf{f}_1 r) \dots (\mathbf{f}_{n-1} r) \end{aligned} \tag{37}$$

To construct an element of the record type from values $x_1 : A_1, \dots, x_n : A_n$ $x_1 \dots x_{n-1}$, we use the syntax $\mathbf{record}\{\mathbf{f}_1 = x_1; \dots; \mathbf{f}_n = x_n\}$. Applying one of the projections to a record constructed this way gives back the field:

$$\mathbf{f}_i (\mathbf{record}\{\mathbf{f}_1 = x_1; \dots; \mathbf{f}_n = x_n\}) = x_i \tag{38}$$

The type A_i of each field can depend on the values of the previous fields $\mathbf{f}_j r$ for $j < i$. For example, $\sum_{x:A} (B x)$ can be defined as a record with two projections $\mathbf{fst} : \sum_{x:A} (B x) \rightarrow A$ and $\mathbf{snd} : (p : \sum_{x:A} (B x)) \rightarrow B (\mathbf{fst} p)$. Then x, y is shorthand for $\mathbf{record}\{\mathbf{fst} = x; \mathbf{snd} = y\}$.

The η -law states that for any $r : \mathbf{R}$, we have

$$r = \mathbf{record}\{\mathbf{f}_1 = \mathbf{f}_1 r; \dots; \mathbf{f}_n = \mathbf{f}_n r\} \tag{39}$$

We use the η -law to construct two unification rules. The first rule applies η to expand a variable of record type into its constituent fields, while the second rule performs a similar expansion on an equation between two elements of a record type.⁴

Lemma 44. Let $\mathbf{R} : \mathbf{Set}_\ell$ be a record type with fields given by (37). Then we have an equivalence:

$$\eta\mathbf{var}_{\mathbf{R}} : (r : \mathbf{R}) \simeq (f_1 : A_1) \dots (f_n : A_n f_1 \dots f_{n-1}) \tag{40}$$

Construction of $\eta\mathbf{var}_{\mathbf{R}}$. We define $\eta\mathbf{var}_{\mathbf{R}} r$ by $\mathbf{f}_1 r; \dots; \mathbf{f}_n r$, and $\eta\mathbf{var}_{\mathbf{R}}^{-1} f_1 \dots f_n$ by $\mathbf{record}\{\mathbf{f}_1 = f_1; \dots; \mathbf{f}_n = f_n\}$. The proofs of both $\mathbf{isLinV}$ and $\mathbf{isRinv}$ is $\mathbf{refl}$: in the former case this is type-correct because of the η -law (39), and in the latter case because of the computation rules for projections (38). $\square$

⁴ A cubical type theorist might say these are two instances of the same rule.

Example 45. This rule is especially useful for solving equations where one side is a projection applied to a variable. Consider the type $A \times B = \Sigma_{_} A.B$. Then we can solve the equation $\text{fst } p \equiv_{\mathbb{N}} \text{zero}$ as follows:

$$\begin{aligned} & (\underline{p} : \mathbb{N} \times \mathbb{N})(e : \text{fst } p \equiv_{\mathbb{N}} \text{zero}) \\ & \simeq (\underline{x} : \mathbb{N})(y : \mathbb{N})(\underline{e} : x \equiv_{\mathbb{N}} \text{zero}) \\ & \simeq (y : \mathbb{N}) \end{aligned} \quad (41)$$

Here the composite telescope mapping $\sigma : (y : \mathbb{N}) \rightarrow (p : \mathbb{N} \times \mathbb{N})(e : \text{fst } p \equiv_{\mathbb{N}} \text{zero})$ from bottom to top is $\lambda y. (\text{zero}, y); \text{refl}$.

Lemma 46. Let $\mathbf{R} : \text{Set}_{\ell}$ be a record type with fields given by (37). Then we have an equivalence:

$$\eta\text{eq}_{\mathbf{R}} : (e : r \equiv_{\mathbf{R}} s) \simeq (e_1 : \mathbf{f}_1 r \equiv_{A_1} \mathbf{f}_1 s) \dots (e_n : \mathbf{f}_n r \equiv_{A_n} e_1 \dots e_{n-1} \mathbf{f}_n s) \quad (42)$$

Construction of $\eta\text{eq}_{\mathbf{R}}$. To construct $\eta\text{eq}_{\mathbf{R}}$, we rely on $\eta\text{var}_{\mathbf{R}}$ and cong : we define $\eta\text{eq}_{\mathbf{R}} e = \text{cong } \eta\text{var}_{\mathbf{R}} e$ and $\eta\text{eq}_{\mathbf{R}}^{-1} \bar{e} = \overline{\text{cong}} \eta\text{var}_{\mathbf{R}}^{-1} \bar{e}$. The proofs of isLin and isRinv are straightforward applications of $\bar{\mathbf{J}}$. $\square$

Example 47. This rule is useful when one side of an equation is of the form $\text{record}\{\dots\}$. If $f : \mathbb{N} \rightarrow \mathbb{N} \times \mathbb{N}$, then we can solve the equation $x, y \equiv_{\mathbb{N} \times \mathbb{N}} f z$ as follows:

$$\begin{aligned} & (x y z : \mathbb{N})(\underline{e} : x, y \equiv_{\mathbb{N} \times \mathbb{N}} f z) \\ & \simeq (\underline{x} y z : \mathbb{N})(\underline{e}_1 : x \equiv_{\mathbb{N}} \text{fst } (f z))(\underline{e}_2 : y \equiv_{\mathbb{N}} \text{snd } (f z)) \\ & \simeq (\underline{y} z : \mathbb{N})(\underline{e}_2 : y \equiv_{\mathbb{N}} \text{snd } (f z)) \\ & \simeq (z : \mathbb{N}) \end{aligned} \quad (43)$$

5 Computational behaviour of unification rules

Until now, we have only been interested in an equivalence representing a most general unifier insofar that it has the correct type. But as a term in type theory, it also has a certain computational behaviour. This computational behaviour is important for the applications we have in mind, in particular the translation of pattern matching to eliminators (Cockx et al., 2016b).

In particular, an important step during the translation of a case split to the application of an eliminator is to generate auxiliary equations that are then solved by unification (Definition 27). However, in the end these equations are filled in with refl .

Example 48. Consider the definition of the function tail (Example 1):

$$\begin{aligned} & \text{tail} : (n : \mathbb{N}) \rightarrow \text{Vec } A (\text{succ } n) \rightarrow \text{Vec } A n \\ & \text{tail}.m (\text{cons } m x xs) = xs \end{aligned} \quad (44)$$

To translate this definition to eliminators, the first step is to generalize the problem to constructing $\text{tail}' : (k n : \mathbb{N}) \rightarrow \text{Vec } A k \rightarrow k \equiv_{\mathbb{N}} \text{succ } n \rightarrow \text{Vec } A n$ and then let $\text{tail } n xs = \text{tail}' (\text{succ } n) n xs \text{ refl}$. Note in particular that the final argument to tail' is refl !

The next step in the translation is to apply case analysis on xs , which instantiates k with **zero** in the case for **nil** and **suc** m in the case for **cons**. The equations **zero** $\equiv_{\mathbb{N}}$ **suc** n and **suc** $m \equiv_{\mathbb{N}}$ **suc** n are then solved by unification:

- In the first case, we get a negative unifier

$$f_1 : (n : \mathbb{N})(\text{zero} \equiv_{\mathbb{N}} \text{suc } n) \simeq \perp \quad (45)$$

and in particular $f \ n \ e : \perp$, so the case can be handled by applying **absurd** : $(A : \text{Set}_{\ell}) \rightarrow \perp \rightarrow A$.

- In the second case, we get a positive unifier

$$f_2 : (m \ n : \mathbb{N})(x : A)(xs : \text{Vec } A \ m)(\text{suc } m \equiv_{\mathbb{N}} \text{suc } n) \simeq (m : \mathbb{N})(x : A)(xs : \text{Vec } A \ m) \quad (46)$$

This unifier is used in the further translation of the right-hand side of **tail**.

When (the translated version of) **tail** is called with arguments m and **cons** $m \ x \ xs$, it will evaluate to **tail'** (**suc** m) m (**cons** $m \ x \ xs$) **refl** by definition. In particular, the proof of **suc** $m \equiv_{\mathbb{N}}$ **suc** n that is passed to the equivalence f_2 is **refl**.

So if we care about the computational behaviour of the output of this translation, we should worry about what happens when we apply a unification rule to **refl**, i.e. when the equations on one side of the equivalence hold in fact definitionally.

Intuitively, a unification rule $r : \Gamma(\bar{e} : \bar{u} \equiv_{\Delta} \bar{v}) \simeq \Gamma'(\bar{e}' : \bar{u}' \equiv_{\Delta'} \bar{v}')$ should satisfy the property that if the equations on the left hold definitionally then the ones on the right also hold definitionally and vice versa. Moreover, the proofs **isLinv** r and **isRinv** r should be trivial in those cases. In other words, the various components of the equivalence should satisfy the principle ‘**refl** in, **refl** out’. This leads us to the following definition of a strong unification rule:

Definition 49 (Strong unification rule). A positive unification rule $r : \Gamma(\bar{e} : \bar{u} \equiv_{\Delta} \bar{v}) \simeq \Gamma'(\bar{e}' : \bar{u}' \equiv_{\Delta'} \bar{v}')$ is a *strong unification rule* if for any Γ_0 and for any $\bar{s}$ and $\bar{s}'$ such that $\Gamma_0 \vdash \text{refl} : \bar{u}[\Gamma \mapsto \bar{s}] \equiv_{\Delta} \bar{v}[\Gamma \mapsto \bar{s}]$ and $\Gamma_0 \vdash \text{refl} : \bar{u}'[\Gamma' \mapsto \bar{s}'] \equiv_{\Delta'} \bar{v}'[\Gamma' \mapsto \bar{s}']$, it satisfies the following five properties:

1. $\Gamma_0 \vdash r \ \bar{s} \ \text{refl} = (\bar{t}'; \text{refl}) : \Gamma'(\bar{e}' : \bar{u}' \equiv_{\Delta'} \bar{v}')$ for some $\bar{t}' : \Gamma'$.
2. $\Gamma_0 \vdash \text{linv } r \ \bar{s}' \ \text{refl} = (\bar{t}_1; \text{refl}) : \Gamma(\bar{e} : \bar{u} \equiv_{\Delta} \bar{v})$ for some $\bar{t}_1 : \Gamma$.
3. $\Gamma_0 \vdash \text{rinv } r \ \bar{s}' \ \text{refl} = (\bar{t}_2; \text{refl}) : \Gamma(\bar{e} : \bar{u} \equiv_{\Delta} \bar{v})$ for some $\bar{t}_2 : \Gamma$.
4. $\Gamma_0 \vdash \text{isLinv } r \ \bar{s} \ \text{refl} = \text{refl} : \text{linv } r \ (r \ \bar{s} \ \text{refl}) \equiv_{\Gamma(\bar{e} : \bar{u} \equiv_{\Delta} \bar{v})} (\bar{s}; \text{refl})$.
5. $\Gamma_0 \vdash \text{isRinv } r \ \bar{s}' \ \text{refl} = \text{refl} : r \ (\text{rinv } r \ \bar{s}' \ \text{refl}) \equiv_{\Gamma'(\bar{e}' : \bar{u}' \equiv_{\Delta'} \bar{v}')} (\bar{s}'; \text{refl})$.

The trivial equivalence **id** : $\Gamma(\bar{e} : \bar{u} \equiv_{\Delta} \bar{v}) \simeq \Gamma(\bar{e} : \bar{u} \equiv_{\Delta} \bar{v})$ is clearly a strong unification rule, and we can compose strong unification rules:

Lemma 50. *If r and r' are strong unification rules, then $r \circ r'$ is one as well.*

Proof. This follows directly from the definition of a strong unification rule and the composition of two equivalences. $\square$

Most of the unification rules we have seen up until now are strong:

Lemma 51. **solution** and **injectivity**_c are strong unification rules.

Proof. This follows directly from the construction of these rules (see Lemma 29 for **solution**, and Lemma 81 for **injectivity_c**). $\square$

For the deletion rule (30), the computational behaviour depends on the type of the equation being eliminated and the construction of the proof of **K**. In a theory with a general **K** rule, the deletion rule is also a strong unification rule. However, if we construct **K** for a specific datatype such as **N** from the basic eliminator, then the resulting unification rule won't be strong as evaluation gets stuck in case the left- and right-hand side of the equation are not of the form **zero** or **suc** m .

Lemma 52. ηvar_R and ηeq_R are strong unification rules.

Proof. For ηvar_R , the first three properties are trivial as this rule does not involve equations, and the last two properties holds as well since $\text{isLinv } \eta\text{var}_R r$ and $\text{isRinv } \eta\text{var}_R r$ are equal to **refl** by definition.

For ηeq_R , the functions ηeq_R , ηeq_R^{-1} , $\text{isLinv } \eta\text{eq}_R$ and $\text{isRinv } \eta\text{eq}_R$ all map $\overline{\text{refl}}$ to $\overline{\text{refl}}$ by definition of $\overline{\text{cong}}$ and $\bar{J}$, so it is also trivially a strong rule. $\square$

In the special case of a most general unifier, the telescope Δ' on the right becomes trivial so we can give a simpler definition of strongness. In particular, the first property always holds so we may omit it, and for the second, third, and fifth property it is sufficient to require that they hold in the case that $\bar{s}'$ is a list of variables Γ' (since there are no equations in Δ' that should hold as a precondition). This leads us to the following definition of a strong unifier:

Definition 53 (Strong unifier). A most general unifier $f : \Gamma(\bar{e} : \bar{u} \equiv_{\Delta} \bar{v}) \simeq \Gamma'$ is *strong* if it satisfies the following four properties:

1. $\Gamma' \vdash \text{linv } f \Gamma' = (\bar{t}_1; \overline{\text{refl}}) : \Gamma(\bar{e} : \bar{u} \equiv_{\Delta} \bar{v})$ for some $\bar{t}_1$.
2. $\Gamma' \vdash \text{rinv } f \Gamma' = (\bar{t}_2; \overline{\text{refl}}) : \Gamma(\bar{e} : \bar{u} \equiv_{\Delta} \bar{v})$ for some $\bar{t}_2$.
3. For any Γ_0 and $\bar{s}$ such that $\Gamma_0 \vdash \overline{\text{refl}} : \bar{u}[\Gamma \mapsto \bar{s}] \equiv_{\Delta} \bar{v}[\Gamma \mapsto \bar{s}]$, we have $\Gamma_0 \vdash \text{isLinv } f \bar{s} \overline{\text{refl}} = \overline{\text{refl}} : \text{linv } f (f \bar{s} \overline{\text{refl}}) \equiv_{\Gamma(\bar{e} : \bar{u} \equiv_{\Delta} \bar{v})} (\bar{s}; \overline{\text{refl}})$.
4. $\Gamma' \vdash \text{isRinv } f \Gamma' = \overline{\text{refl}} : f (\text{rinv } f \Gamma') \equiv_{\Gamma'} \Gamma'$.

From the first two properties we deduce in particular that the equations $\bar{u} \equiv_{\Delta} \bar{v}$ are indeed satisfied definitionally under the substitution embedded in the most general unifier f .

Lemma 54. If $f = r_1 \circ r_2 \circ \dots \circ r_n : \Gamma(\bar{e} : \bar{u} \equiv_{\Delta} \bar{v}) \simeq \Gamma'$ is composed of strong unification rules $r_1, r_2, \dots, r_n$, then f is a strong unifier.

Proof. Since a strong unifier is a special case of a strong unification rule, this follows directly from Lemma 50. $\square$

Lemma 55. If $f : \Gamma(\bar{e} : \bar{u} \equiv_{\Delta} \bar{v}) \simeq \Gamma'$ is a strong unifier, then $\Gamma' \vdash \text{linv } f \Gamma' = \text{rinv } f \Gamma' : \Gamma$.

Proof. By the second property of a strong unifier, we have $\bar{s}_2$ such that

$$\Gamma' \vdash \text{rinv } f \Gamma' = (\bar{s}_2; \overline{\text{refl}}) : \Gamma(\bar{e} : \bar{u} \equiv_{\Delta} \bar{v}) \quad (47)$$

In particular, $(\bar{s}_2; \overline{\text{refl}})$ has type $\Gamma(\bar{e} : \bar{u} \equiv_{\Delta} \bar{v})$, so we know that

$$\Gamma' \vdash \overline{\text{refl}} : \bar{u}[\Gamma \mapsto \bar{s}_2] \equiv_{\Delta} \bar{v}[\Gamma \mapsto \bar{s}_2] \quad (48)$$

By the third property of a strong unifier, this implies that

$$\Gamma' \vdash \text{isLinv } f \bar{s}_2 \overline{\text{refl}} = \overline{\text{refl}} : (\bar{s}_2; \overline{\text{refl}}) \equiv_{\Gamma(\bar{e} : \bar{u} \equiv_{\Delta} \bar{v})} (\bar{s}_2; \overline{\text{refl}}) \quad (49)$$

Since the left- and right-hand side of a definitional equality always have definitionally equal types, it follows in particular that the type of $\text{isLinv } f \bar{s}_2 \overline{\text{refl}}$ must be definitionally equal to the type of $\overline{\text{refl}}$, i.e.

$$\Gamma' \vdash \text{linv } f (f \bar{s}_2 \overline{\text{refl}}) = (\bar{s}_2; \overline{\text{refl}}) : \Gamma(\bar{e} : \bar{u} \equiv_{\Delta} \bar{v}) \quad (50)$$

By similar reasoning, the fourth property gives us that the type of $\text{isRinv } f \Gamma'$ must be definitionally equal to that of $\overline{\text{refl}}$, so in particular

$$\Gamma' \vdash f (\text{rinv } f \Gamma') = \Gamma' : \Gamma' \quad (51)$$

But $\text{rinv } f \Gamma' = \bar{s}_2; \overline{\text{refl}}$, so we also have

$$\Gamma' \vdash f \bar{s}_2 \overline{\text{refl}} = \Gamma' : \Gamma' \quad (52)$$

Applying $\text{linv } f$ to both sides of this equations gives us that

$$\Gamma' \vdash \text{linv } f (f \bar{s}_2 \overline{\text{refl}}) = \text{linv } f \Gamma' : \Gamma(\bar{e} : \bar{u} \equiv_{\Delta} \bar{v}) \quad (53)$$

Putting this together with (50) gives us that

$$\Gamma' \vdash \text{linv } f \Gamma' = (\bar{s}_2; \overline{\text{refl}}) : \Gamma(\bar{e} : \bar{u} \equiv_{\Delta} \bar{v}) \quad (54)$$

Since $\text{rinv } f \Gamma' = (\bar{s}_2; \overline{\text{refl}})$, this gives us that $\Gamma' \vdash \text{linv } f \Gamma' = \text{rinv } f \Gamma' : \Gamma$, as we wanted to prove. $\square$

This lemma implies that we can write f^{-1} for both $\text{linv } f$ and $\text{rinv } f$ when f is a strong unifier.

When applying specialization by unification to construct a function $m : (\bar{x} : \Gamma) \rightarrow (\bar{e} : \bar{u} \equiv_{\Delta} \bar{v}) \rightarrow T \bar{x} \bar{e}$ from the subgoal $m^s : (\bar{x}' : \Gamma') \rightarrow T (\text{linv } f \bar{x}')$ we expect m to have ‘the same’ computational behaviour as m^s in case the equations $\bar{u} \equiv_{\Delta} \bar{v}$ are actually satisfied. This is the content of the following lemma.

Lemma 56. *If f is a strong unifier (Definition 53), then the function m constructed through specialization by unification (Definition 27) satisfies the definitional equality $m(f^{-1} \bar{x}') = m^s \bar{x}'$ for any $\bar{x}' : \Gamma'$.*

Proof. Remember that for a strong unifier f , $\text{linv } f$ and $\text{rinv } f$ are definitionally equal, so it is fine to write f^{-1} here. By the first property of a strong unifier, we have $f^{-1} \bar{x}' = \bar{s}; \overline{\text{refl}}$ for some $\bar{s} : \Gamma$. By the third property, this implies that $\text{isLinv } f (f^{-1} \bar{x}') = \overline{\text{refl}}$. By the fourth property, we also have $f(f^{-1} \bar{x}') = \bar{x}'$. So we have

$$\begin{aligned} m(f^{-1} \bar{x}') &= \overline{\text{subst}} T (\text{isLinv } f (f^{-1} \bar{x}')) (m^s (f(f^{-1} \bar{x}'))) \\ &= \overline{\text{subst}} T \overline{\text{refl}} (m^s \bar{x}') \\ &= m^s \bar{x}' \end{aligned} \quad (55)$$

$\square$

The fact that $m(\bar{t} \text{ refl})$ evaluates to $m^s(f \bar{t} \text{ refl})$ ensures that the computational behaviour corresponds to the clause written by the user in the translation of pattern matching to eliminators.

Discussion about the definition of a strong unifier. There are other possible definitions of a strong unifier. In particular, to guarantee the good computational properties of functions constructed through specialization by unification (Lemma 56), we only need properties 1, 3, and a weaker version of property 4. In our previous work (Cockx et al., 2016a) we used an even weaker definition of a strong unifier:

Definition 57 (DEPRECATED, version from Cockx et al. (2016a)). A most general unifier $f : \Gamma(\bar{e} : \bar{u} \equiv_{\Delta} \bar{v}) \simeq \Gamma'$ is *strong* if for any $\bar{x}' : \Gamma'$, it satisfies the following two properties:

- $f(f^{-1} \bar{x}') = \bar{x}'$
- $\text{isLinv } f(f^{-1} \bar{x}') = \text{refl}$

This definition ensures exactly the properties needed for Lemma 56 to hold. However, we require another property in order to preserve clauses as definitional equalities in the translation of dependent pattern matching to eliminators (Cockx, 2017), namely that $f^{-1} \bar{x}'$ must be definitionally equal to something of the form $\bar{s}; \text{refl}$. Additionally, in various places we relied implicitly on the fact that $\text{linv } f$ and $\text{rinv } f$ should be definitionally equal. These discoveries lead to our current definition of a strong unifier.

6 Higher-dimensional unification

When constructing the indexed versions of the injectivity, conflict, and cycle rules (Section 4.2), we required that the telescope of the equations on the left-hand side should be exactly $\bar{D} = (\bar{u} : \Xi)(x : \bar{D} \bar{u})$. This means these rules can only be applied to an equation where the type is *fully general*, i.e. a datatype applied to distinct equality proofs for its indices. This is convenient when the equations we start with are of this form because it allows us to simplify all equations at the same time.

The main question posed in this section is what we can do if we encounter an equation of the form $\text{c } \bar{u} \equiv_{\bar{D}} \text{c } \bar{v}$ but the indices $\bar{v}$ are not fully general.

Example 58. Suppose the unification algorithm is trying to solve an equation

$$(e : \text{cons } n \ x \ xs \equiv_{\text{Vec } A} (\text{suc } n) \ \text{cons } n \ y \ ys) \quad (56)$$

of type $\text{Vec } A (\text{suc } n)$ where n is a regular variable rather than an equality proof. In this case it is not possible to apply the $\text{injectivity}_{\text{cons}}$ rule directly.

There is no fundamental reason why unification should fail on this example. On the other hand, always applying the injectivity rule even when the indices are not fully general is unsound (Example 4). This is not just a theoretical problem either: see for example issues #1411 and #1775 on the Agda bug tracker (Abel, 2015b; Sicard-Ramírez, 2016).

In previous work, we tried different approaches to solve this problem that worked in some cases but were ultimately unsatisfactory. In Cockx et al. (2014) we restricted all unification rules to homogeneous equations and additionally imposed a *self-unifiability criterion* to the indices of the datatype when applying the injectivity rule. In practice, this meant that the injectivity rule could only be applied when the indices consisted of closed constructor forms only (e.g. $\text{succ}(\text{succ zero})$, but not $\text{succ } n$), a severe restriction to the applicability of the rule. In Cockx et al. (2016a) we used the general (heterogeneous) version of the injectivity rule and relied on *reverse unification* to generalize the indices. This method had some potential in theory, but turned out to be too difficult to implement in practice. Neither did we take into account the type of the constructor in question, so we were unable to include useful heuristics such as forced constructor arguments (Brady, McBride, and McKinna, 2003).

In this section, we describe a general technique for solving equations between constructors of indexed datatypes. First, we study why the problem is so difficult by looking at the analogous problem for the conflict and cycle rules, and make a first attempt at generalizing the injectivity rule (Section 6.1). We continue to show how to generalize the equality proofs in the indices in the general case by introducing new equations between equality proofs (Section 6.2). Borrowing terminology from homotopy type theory, we call them *higher-dimensional equations*. To solve these higher-dimensional equations, we show how to lift existing unification rules to higher dimensions (Section 6.3).

6.1 Generalizing unification rules

Before we try to tackle the problem of how to apply the injectivity rule on an equation when the indices are not fully general, we first consider the analogous problem for the conflict and cycle rules. The reason to take on these rules first is because they shed some light on why the problem is harder for the injectivity rule.

Lemma 59. *Consider a unification problem of the form $\bar{s}_1; \mathbf{c}_1 \bar{t}_1 \equiv_{\Phi(x:\mathbf{D} \bar{v})} \bar{s}_2; \mathbf{c}_2 \bar{t}_2$ where $\mathbf{D} : \Xi \rightarrow \mathbf{Set}_\ell$ is a datatype and $\mathbf{c}_1 : \Delta_1 \rightarrow \mathbf{D} \bar{u}_1$ and $\mathbf{c}_2 : \Delta_2 \rightarrow \mathbf{D} \bar{u}_2$ are two distinct constructors of $\mathbf{D}$. Then we have an equivalence*

$$\text{conflict}'_{\mathbf{c}_1, \mathbf{c}_2} : (\bar{s}_1; \mathbf{c}_1 \bar{t}_1 \equiv_{\Phi(x:\mathbf{D} \bar{v})} \bar{s}_2; \mathbf{c}_2 \bar{t}_2) \simeq \perp \quad (57)$$

The indices $\bar{v}$ are arbitrary, i.e. they do not have to be variables like in the standard conflict rule (Lemma 37). However, the type of the final equation still has to be the datatype $\mathbf{D}$ applied to these indices. In particular it cannot be a variable itself, or else we would run into the problem described in Example 41.

Before we give the proof of this lemma, we first want to show how a naive proof attempt fails. It goes as follows: to construct a function $(\bar{s}_1; \mathbf{c}_1 \bar{t}_1 \equiv_{\Phi(x:\mathbf{D} \bar{v})} \bar{s}_2; \mathbf{c}_2 \bar{t}_2) \rightarrow \perp$, it suffices (by the $\bar{\mathbf{J}}$ rule) to construct a function $\mathbf{c}_1 \bar{t}_1 \equiv_{\mathbf{D} \bar{v}} \mathbf{c}_2 \bar{t}_2 \rightarrow \perp$. This function is constructed by calling the indexed conflict rule (37) with refl for the proof of $\bar{u}_1[\Delta_1 \mapsto \bar{t}_1] \equiv_{\Xi} \bar{u}_2[\Delta_2 \mapsto \bar{t}_2]$. Since any function to $\perp$ is an equivalence, we are done.

Think a moment about what is wrong with this proof. It uses the $\bar{J}$ rule to eliminate the equations $\bar{s}_1 \equiv_{\Phi} \bar{s}_2$, but there is no guarantee that $\bar{s}_1$ or $\bar{s}_2$ are in fact variables. Moreover, their structure as a term may be important for satisfying the assumptions of the lemma, so simply generalizing the statement of the lemma is not possible. In other words, the error in this proof attempt stems from a confusion about the status of $\bar{s}_1$ and $\bar{s}_2$ as variables at the meta-level, while they can be arbitrary terms at the object level!

We work around this issue by using the following lemma:

Lemma 60. *Let $f : A \rightarrow B$, $P : B \rightarrow \mathbf{Set}$, $e : s \equiv_A t$, $u : P(f s)$ and $v : P(f t)$. Then the types $u \equiv_{P(f e)} v$ and $u \equiv_{P(\text{cong } f e)} v$ are equivalent.*

Proof. Notice the rather subtle difference between these two types: the first one expands to

$$\text{subst } (P \circ f) e u \equiv_{P t} v \quad (58)$$

while the second one expands to

$$\text{subst } P(\text{cong } f e) u \equiv_{P t} v \quad (59)$$

To prove that they are equivalent, it is sufficient to prove that $\text{subst } (P \circ f) e u \equiv_{P t} \text{subst } P(\text{cong } f e) u$. But this follows directly by eliminating e using J . $\square$

Construction of $\text{conflict}'_{c_1, c_2}$. We start by expanding the definition of telescopic equality: we have to derive an element of $\perp$ from

$$(\bar{e}_1 : \bar{s}_1 \equiv_{\Phi} \bar{s}_2)(e_2 : c_1 \bar{t}_1 \equiv_{\mathbf{D}[\Phi \mapsto \bar{e}_1]} c_2 \bar{t}_2) \quad (60)$$

By Lemma 60, the type of e_2 is equivalent to $c_1 \bar{t}_1 \equiv_{\mathbf{D}(\text{cong } (\lambda \Phi. \bar{v}) \bar{e})} c_2 \bar{t}_2$. So we call the conflict rule (37) with arguments $(\text{cong } (\lambda \Phi. \bar{v}) \bar{e}_1); e_2$ to get an element of type $\perp$. Since any function to $\perp$ is an equivalence, this finishes the proof. $\square$

Similarly, we can generalize the cycle rule:

Lemma 61. *Consider a unification problem of the form $\bar{s}_1; t_1 \equiv_{\Phi(x:\mathbf{D} \bar{v})} \bar{s}_2; t_2$ where $\mathbf{D} : \Xi \rightarrow \mathbf{Set}_{\ell}$ is a datatype and $t_1 \prec t_2$. Then we have an equivalence*

$$\text{cycle}'_{t_1, t_2} : (\bar{s}_1; t_1 \equiv_{\Phi(x:\mathbf{D} \bar{v})} \bar{s}_2; t_2) \simeq \perp \quad (61)$$

Construction of cycle'_{t_1, t_2} . Analogously to the construction of $\text{conflict}'_{c_1, c_2}$. $\square$

For injectivity, it is not as easy to generalize the rule to arbitrary indices like we just did for conflict and cycle. The problem here is harder because we also have to construct an inverse function and prove that it is indeed a left and right inverse, while this was trivial for the two negative rules. In the special case where the index telescope Ξ satisfies UIP, we can construct the generalization:

Lemma 62 (Generalized injectivity). *Consider a unification problem of the form $\bar{s}_1; c \bar{t}_1 \equiv_{\Phi(x:\mathbf{D} \bar{v})} \bar{s}_2; c \bar{t}_2$ where $\mathbf{D} : \Xi \rightarrow \mathbf{Set}_{\ell}$ is a datatype with (at least) one constructor $c : \Delta \rightarrow \mathbf{D} \bar{u}$, and assume Ξ satisfies UIP, i.e. we have $\text{deletion}_{\bar{x}} : (\bar{e} : \bar{x} \equiv_{\Xi} \bar{x}) \simeq ()$ for all $\bar{x} : \Xi$. Then we have an equivalence*

$$\text{injectivity}'_c : (\bar{s}_1; c \bar{t}_1 \equiv_{\Phi(x:\mathbf{D} \bar{v})} \bar{s}_2; c \bar{t}_2) \simeq (\bar{s}_1; \bar{t}_1 \equiv_{\Phi \Delta} \bar{s}_2; \bar{t}_2) \quad (62)$$

In case Φ is the empty telescope, this generalized injectivity rule is similar to the specialized injectivity rule from Cockx et al. (2014), but here we ask that the types of the indices Ξ satisfy UIP, instead of asking that the indices $\bar{u}$ are self-unifiable.

*Construction of **injectivity'**_c.* As for the previous lemma, we expand the definition of telescopic equality and apply Lemma 60 to get to

$$(\bar{e}_1 : \bar{s}_1 \equiv_{\Phi} \bar{s}_2)(e_2 : \bar{c} \bar{t}_1 \equiv_{\mathbf{D}} (\overline{\text{cong}}(\lambda\Phi.\bar{v}) \bar{e}_1) \bar{c} \bar{t}_2) \quad (63)$$

Since Ξ satisfies UIP, it follows that $(e'_1 : \bar{v}[\Phi \mapsto \bar{s}_1] \equiv_{\Xi} \bar{v}[\Phi \mapsto \bar{s}_2])$ is equivalent to $()$. So the previous telescope is equivalent to

$$\begin{aligned} &(\bar{e}_1 : \bar{s}_1 \equiv_{\Phi} \bar{s}_2)(\bar{e}'_1 : \bar{v}[\Phi \mapsto \bar{s}_1] \equiv_{\Xi} \bar{v}[\Phi \mapsto \bar{s}_2]) \\ &(e_2 : \bar{c} \bar{t}_1 \equiv_{\mathbf{D}} (\overline{\text{cong}}(\lambda\Phi.\bar{v}) \bar{e}_1) \bar{c} \bar{t}_2) \end{aligned} \quad (64)$$

Again by UIP, we have that the proofs $\overline{\text{cong}}(\lambda\Phi.\bar{v}) \bar{e}_1$ and $\bar{e}'_1$ of type $\bar{v}[\Phi \mapsto \bar{s}_1] \equiv_{\Xi} \bar{v}[\Phi \mapsto \bar{s}_2]$ are equal. This means the previous telescope is equivalent to

$$(\bar{e}_1 : \bar{s}_1 \equiv_{\Phi} \bar{s}_2)(\bar{e}'_1 : \bar{v}[\Phi \mapsto \bar{s}_1] \equiv_{\Xi} \bar{v}[\Phi \mapsto \bar{s}_2])(e_2 : \bar{c} \bar{t}_1 \equiv_{\mathbf{D}_{\bar{e}'_1}} \bar{c} \bar{t}_2) \quad (65)$$

Finally, we can apply the injectivity rule (36) to prove that the part of the telescope containing $\bar{e}'_1$ and e_2 is equivalent to $\bar{t}_1 \equiv_{\Delta} \bar{t}_2$, so the previous telescope is equivalent to

$$(\bar{e}_1 : \bar{s}_1 \equiv_{\Phi} \bar{s}_2)(\bar{e}_2 : \bar{t}_1 \equiv_{\Delta} \bar{t}_2) \quad (66)$$

which is what we wanted to prove. $\square$

Like for the deletion rule, this generalized injectivity rule usually won't be a strong rule because its computational behaviour depends on the construction of the proof UIP for the index types.

6.2 A generalized injectivity rule

The generalized injectivity rule from the previous section is unsatisfactory because it requires the index types of the datatype to satisfy UIP. This means we didn't actually solve the problem of depending on UIP yet, we only moved it to the indices. However, the proof taught us something about how to solve the problem in general: it introduced new equality proofs $\bar{e}'_1$ and used UIP to substitute these for the indices of $\mathbf{D}$, allowing us to apply the injectivity rule. In other words, it moved the problem from talking about equalities between *terms* to equalities between *equality proofs*.

In this section, we show how to apply this idea in a more general way to remove the dependency on UIP completely. We do this by applying—what else—unification to the equations between the indices of the datatype. Since the indices in the type of an equation can depend on the equality proofs of the previous equations, this means we have to solve not just equalities between terms but also equalities between other equality proofs, i.e. higher-dimensional equations.

At first sight, it would seem that an entirely new set of unification rules is needed to solve higher-dimensional equations (except for the solution rule, which can be used at any dimension). However, it is possible to reuse the existing unification

rules on higher-dimensional problems. For example, the `injectivitysuc` rule can be used not just to simplify equations of the form `suc x ≡N suc y` to `x ≡N y`, but also `cong suc e1 ≡suc x ≡N suc y cong suc e2` to `e1 ≡x ≡N y e2`.

In general, whenever the unification algorithm encounters a higher-dimensional unification problem $\bar{u} \equiv_{\bar{x} \equiv_{\Delta} \bar{y}} \bar{v}$ it lowers it by one dimension to the problem $\bar{u} \equiv_{\Delta} \bar{v}$ where the equation variables in $\bar{u}$ and $\bar{v}$ are treated as regular variables. If it manages to find a solution to this one-dimensional problem, it can then *lift* this solution to get a solution to the original problem. The technical result that makes this possible is Lemma 72 in the next section.

Let's first take a look of how this works on an example.

Example 63. Consider the unification problem:

$$\Gamma(e : \text{cons } n \ x \ xs \equiv_{\text{Vec } A} (\text{suc } n) \ \text{cons } n \ y \ ys) \quad (67)$$

where $\Gamma = (n : \mathbb{N})(x \ y : A)(xs \ ys : \text{Vec } A \ n)$. The `injectivitycons` rule cannot be applied, as the index `suc n` is not fully general (i.e. it is not an equation variable). Instead, we solve this unification problem in three steps: in the first step, we generalize over the indices in order to apply the injectivity rule, generating higher-dimensional equations in the process. In the second step, we bring down these equations by one dimension so we can solve them by applying known unification rules. Finally, we lift the one-dimensional unifier to the higher-dimensional problem.

Step 1: generalizing the indices. We generalize the problem by introducing an extra equation $e_1 : \text{suc } n \equiv_{\mathbb{N}} \text{suc } n$ to the telescope, together with a proof p that e_1 is equal to `refl`:

$$\begin{aligned} & \Gamma(e : \text{cons } n \ x \ xs \equiv_{\text{Vec } A} (\text{suc } n) \ \text{cons } n \ y \ ys) \\ & \simeq \Gamma(e_1 : \text{suc } n \equiv_{\mathbb{N}} \text{suc } n)(e_2 : \text{cons } n \ x \ xs \equiv_{\text{Vec } A} e_1 \ \text{cons } n \ y \ ys) \\ & \quad (p : e_1 \equiv_{\text{suc } n \equiv_{\mathbb{N}} \text{suc } n} \text{refl}) \end{aligned} \quad (68)$$

This is nothing but an application of the `solution` rule in the reverse direction, as applying `solution` to p would bring us back to the first equation.⁵

Since the index in the type of e_2 is now fully general, we are free to apply the `injectivitycons` rule:

$$\begin{aligned} & \Gamma(\underline{e}_1 : \text{suc } n \equiv_{\mathbb{N}} \text{suc } n)(\underline{e}_2 : \text{cons } n \ x \ xs \equiv_{\text{Vec } A} \underline{e}_1 \ \text{cons } n \ y \ ys) \\ & (p : \underline{e}_1 \equiv_{\text{suc } n \equiv_{\mathbb{N}} \text{suc } n} \text{refl}) \\ & \simeq \Gamma(e'_1 : n \equiv_{\mathbb{N}} n)(e'_2 : x \equiv_A y)(e'_3 : xs \equiv_{\text{Vec } A} e'_1 \ ys) \\ & \quad (p : \text{cong suc } e'_1 \equiv_{\text{suc } n \equiv_{\mathbb{N}} \text{suc } n} \text{refl}) \end{aligned} \quad (69)$$

Applying the injectivity rule to e_2 has instantiated the variable e_2 with `cong suc e'1`. This instantiation is determined by the computational behaviour of `injectivitycons`

⁵ This is the exact same technique as used by McBride (1998b): to do a case split on a variable $x : \text{Vec } A \ m$ where m is not fully general, he introduces a new variable $n : \mathbb{N}$ together with an equality $e : m \equiv_{\mathbb{N}} n$. This means that now $x : \text{Vec } A \ m$ where m are just variables, so it is possible to perform a case split on x . The only difference in our case is that we are working one dimension higher, i.e. we work with equations between elements of the datatype instead of elements of the datatype itself.

(Lemma 31). As you can see, p is a non-trivial equation between equality proofs, i.e. a *higher-dimensional equation*.

Step 2: lowering the dimension of equations. To solve the higher-dimensional equation p , we first consider a one-dimensional version of this problem:

$$(w'_1 : \mathbb{N})(w'_2 : A)(w'_3 : \mathbf{Vec} A w'_1)(p : \mathbf{succ} w'_1 \equiv_{\mathbb{N}} \mathbf{succ} n) \quad (70)$$

The equality proofs e'_1 , e'_2 and e'_3 from (69) have been replaced by regular variables w'_1 , w'_2 and w'_3 . To reflect this change, $\mathbf{cong} \mathbf{succ} e'_1 : \mathbf{succ} n \equiv_{\mathbb{N}} \mathbf{succ} n$ has been replaced by $\mathbf{succ} w'_1$ and $\mathbf{refl} : \mathbf{succ} n \equiv_{\mathbb{N}} \mathbf{succ} n$ by $\mathbf{succ} n$.

Now this is a problem we know how to solve: we apply $\mathbf{injectivity}_{\mathbf{succ}}$ and $\mathbf{solution}$ to find an equivalence f between this telescope and $(w'_2 : A)(w'_3 : \mathbf{Vec} A n)$.

This solves the one-dimensional problem.

Step 3: lifting unifiers to a higher dimension. How does this help us with the higher-dimensional problem? By Lemma 72 (Section 6.3), we can *lift* the equivalence f to get a new equivalence $f^\dagger$:

$$\begin{aligned} & (e'_1 : n \equiv_{\mathbb{N}} n)(e'_2 : x \equiv_A y)(e'_3 : xs \equiv_{\mathbf{Vec} A e'_1} ys) \\ & (p : \mathbf{cong} \mathbf{succ} e'_1 \equiv_{\mathbf{succ} n \equiv_{\mathbb{N}} \mathbf{succ} n} \mathbf{refl}) \\ & \simeq (e''_1 : x \equiv_A y)(e''_3 : xs \equiv_{\mathbf{Vec} A n} ys) \end{aligned} \quad (71)$$

This solves the higher-dimensional equation p , as well as the reflexive equation e'_1 , without relying on the fact that $\mathbb{N}$ satisfies uniqueness of identity proofs!

Finally, we apply the $\mathbf{solution}$ rule twice to solve the equations e''_2 and e''_3 . So putting everything together, we have found an equivalence between the original telescope (67) and $(n : \mathbb{N})(x : A)(xs : \mathbf{Vec} A n)$, solving the unification problem.

Now that we have seen how to solve the problem in an example, let's try to generalize the solution. The main result of this section is the following theorem.

Theorem 64. *Let $D : \Xi \rightarrow \mathbf{Set}_\ell$ be a datatype and $c : \Delta \rightarrow D \bar{u}$ be a constructor of D . Consider a unification problem of the form*

$$(\bar{e} : \bar{s}_1 ; c \bar{t}_1 \equiv_{\Phi(z:D \bar{v})} \bar{s}_2 ; c \bar{t}_2) \quad (72)$$

Suppose we have an equivalence $f : \Phi \Delta(\bar{p} : \bar{u} \equiv_{\Xi} \bar{v}) \simeq \Delta'$. Then we also have an equivalence $\mathbf{injectivity}''^f_c$ of type

$$(\bar{e} : \bar{s}_1 ; c \bar{t}_1 \equiv_{\Phi(z:D \bar{v})} \bar{s}_2 ; c \bar{t}_2) \simeq (\bar{e}' : f \bar{s}_1 \bar{t}_1 \overline{\mathbf{refl}} \equiv_{\Delta'} f \bar{s}_2 \bar{t}_2 \overline{\mathbf{refl}}) \quad (73)$$

Moreover, if f is a strong unification rule then so is this new equivalence.

The computational behaviour of the unifier f suddenly becomes relevant for the type of the resulting unification problem! In particular, we need the behaviour of f applied to $\mathbf{refl}$ to calculate the left- and right-hand sides of the new equations $\bar{e}'$.

Construction. We follow the same three steps as in Example 63, so if something is unclear it may help to take a look at the corresponding step in the example.

Step 1: generalizing the indices. First, we unfold the telescopic equality in (72) and apply Lemma 60 to get an equivalence with $(\bar{e}_1 : \bar{s}_1 \equiv_{\Phi} \bar{s}_2)(\bar{e}_2 : c \bar{t}_1 \equiv_{D \bar{v}_e} c \bar{t}_2)$

where $\bar{v}_e = \overline{\text{cong}}(\lambda\Phi. \bar{v}) \bar{e}_1$. The equality proofs $\bar{v}_e$ have type $\bar{u}_1 \equiv_{\Xi} \bar{u}_2$ where $\bar{u}_1$ and $\bar{u}_2$ stand for $\bar{u}[\Delta \mapsto \bar{t}_1]$ and $\bar{u}[\Delta \mapsto \bar{t}_2]$ respectively. To generalize $\bar{v}_e$, we introduce new variables $\bar{i} : \bar{u}_1 \equiv_{\Xi} \bar{u}_2$ together with equalities $\bar{p} : \bar{i} \equiv_{\bar{u}_1 \equiv_{\Xi} \bar{u}_2} \bar{v}_e$:

$$\begin{aligned} & (\bar{e}_1 : \bar{s}_1 \equiv_{\Phi} \bar{s}_2)(e_2 : \mathbf{c} \bar{t}_1 \equiv_{\mathbf{D} \bar{v}_e} \mathbf{c} \bar{t}_2) \\ & \simeq (\bar{e}_1 : \bar{s}_1 \equiv_{\Phi} \bar{s}_2)(\bar{i} : \bar{u}_1 \equiv_{\Xi} \bar{u}_2)(e_2 : \mathbf{c} \bar{t}_1 \equiv_{\mathbf{D} \bar{i}} \mathbf{c} \bar{t}_2) \\ & \simeq (\bar{p} : \bar{i} \equiv_{\bar{u}_1 \equiv_{\Xi} \bar{u}_2} \bar{v}_e) \end{aligned} \quad (74)$$

Since $\bar{i}$ consists of distinct equation variables, it's now possible to apply **injectivity_c** to the equation e_2 . This gives us an equivalence:

$$\begin{aligned} & (\bar{e}_1 : \bar{s}_1 \equiv_{\Phi} \bar{s}_2)(\bar{i} : \bar{u}_1 \equiv_{\Xi} \bar{u}_2)(e_2 : \mathbf{c} \bar{t}_1 \equiv_{\mathbf{D} \bar{i}} \mathbf{c} \bar{t}_2)(\bar{p} : \bar{i} \equiv_{\bar{u}_1 \equiv_{\Xi} \bar{u}_2} \bar{v}_e) \\ & \simeq (\bar{e}_1 : \bar{s}_1 \equiv_{\Phi} \bar{s}_2)(\bar{e}'_2 : \bar{t}_1 \equiv_{\Delta} \bar{t}_2)(\bar{p} : \bar{u}_e \equiv_{\bar{u}_1 \equiv_{\Xi} \bar{u}_2} \bar{v}_e) \end{aligned} \quad (75)$$

where $\bar{u}_e = \overline{\text{cong}}(\lambda\Delta. \bar{u}) \bar{e}'_2$.

Step 2: lowering the dimension of equations. Consider the one-dimensional version of this unification problem $\Phi\Delta(\bar{p} : \bar{u} \equiv_{\Xi} \bar{v})$, where the equality proofs $\bar{u}_e$ and $\bar{v}_e$ have been replaced by their lower-dimensional variants $\bar{u}$ and $\bar{v}$ respectively. Since this is a one-dimensional unification problem, we can apply the known unification rules from Figure 8 to solve it. By assumption of the theorem, unification succeeds positively with most general unifier f as a result.

Step 3: lifting unifiers to a higher dimension. Now we have to lift this solution back to the higher-dimensional problem. This lifting is explained in the next subsection. Lemma 72 gives us a lifted equivalence $f^\uparrow$:

$$\begin{aligned} & (\bar{e}_1 : \bar{s}_1 \equiv_{\Phi} \bar{s}_2)(\bar{e}'_2 : \bar{t}_1 \equiv_{\Delta} \bar{t}_2)(\bar{p} : \bar{u}_e \equiv_{\bar{u}_1 \equiv_{\Xi} \bar{u}_2} \bar{v}_e) \\ & \simeq (\bar{e}' : f \bar{s}_1 \bar{t}_1 \overline{\text{refl}} \equiv_{\Delta} f \bar{s}_2 \bar{t}_2 \overline{\text{refl}}) \end{aligned} \quad (76)$$

This is exactly what we need to solve the problem in (75).

Now we combine the equivalences in (74), (75) and (76) to get the final equivalence (73).

To see why this is a strong unification rule, note that it is the composition of four equivalences: Lemma 60, **solution**⁻¹, **injectivity_c** and $f^\uparrow$. By Lemma 50, it is sufficient to prove that these four equivalences are strong unification rules individually. The first two are strong by construction, and **injectivity_c** is a strong unification rule by Lemma 51. Finally, $f^\uparrow$ is strong too by Lemma 73. $\square$

This finishes the application of higher-dimensional unification to the equation e . We have solved the injectivity problem e , and there are no more higher-dimensional unification problems in the resulting equations $\bar{e}'_1$, so we can continue unification on the new problem as normal.

Having given the rule for higher-dimensional unification, the presentation of our full unification algorithm is finished. The rules are summarized in Figure 9, including η -rules for record types (Lemma 44 and Lemma 46), generalized conflict and cycle (Lemma 59 and Lemma 61) and higher-dimensional unification (Theorem 64).

$$\begin{aligned}
\text{solution} &: (x : A)(e : x \equiv_A u) \simeq () & (\text{where } x \notin FV(u)) \\
\text{deletion} &: (e : u \equiv_A u) \simeq () \\
\text{injectivity}''^f_{\mathbf{c}} &: (\bar{e} : \bar{s}_1; \mathbf{c} \bar{t}_1 \equiv_{\Phi(z:\mathbf{D} \bar{v})} \bar{s}_2; \mathbf{c} \bar{t}_2) \simeq (\bar{e}' : f \bar{s}_1 \bar{t}_1 \overline{\text{refl}} \equiv_{\Delta'} f \bar{s}_2 \bar{t}_2 \overline{\text{refl}}) \\
& \quad (\text{where } \mathbf{c} : \Delta \rightarrow \mathbf{D} \bar{u} \text{ and } f : \Phi\Delta(\bar{p} : \bar{u} \equiv_{\Xi} \bar{v}) \simeq \Delta') \\
\text{conflict}'_{\mathbf{c}_1, \mathbf{c}_2} &: (\bar{s}_1; \mathbf{c}_1 \bar{t}_1 \equiv_{\Phi(x:\mathbf{D} \bar{v})} \bar{s}_2; \mathbf{c}_2 \bar{t}_2) \simeq \perp & (\text{where } \mathbf{c}_1 \neq \mathbf{c}_2) \\
\text{cycle}'_{t_1, t_2} &: (\bar{s}_1; t_1 \equiv_{\Phi(x:\mathbf{D} \bar{v})} \bar{s}_2; t_2) \simeq \perp & (\text{where } t_1 \prec t_2) \\
\eta\text{var}_{\mathbf{R}} &: (r : \mathbf{R}) \simeq (f_1 : A_1) \dots (f_n : A_n f_1 \dots f_{n-1}) \\
\eta\text{eq}_{\mathbf{R}} &: (e : r \equiv_{\mathbf{R}} s) \simeq (e_1 : \mathbf{f}_1 r \equiv_{A_1} \mathbf{f}_1 s) \dots (e_n : \mathbf{f}_n r \equiv_{A_n} e_1 \dots e_{n-1} \mathbf{f}_n s)
\end{aligned}$$

Figure 9: The complete list of rules of the unification algorithm.

It is impossible for higher-dimensional unification to end in a negative success, as this would mean we are trying to solve an ill-typed equation. For example, we can never encounter a higher-dimensional conflict:

$$\overline{\text{cong}} \mathbf{c}_1 \bar{e}_1 \equiv_{???} \overline{\text{cong}} \mathbf{c}_2 \bar{u}' \bar{v}' \bar{e}_2 \quad (77)$$

because the left-hand side has a type of the form $\mathbf{c}_1 \bar{u} \equiv \mathbf{c}_1 \bar{v}$ while the right-hand side has type $\mathbf{c}_2 \bar{u}' \equiv \mathbf{c}_2 \bar{v}'$. Likewise, a higher-dimensional cycle would be:

$$e \equiv_{???} \overline{\text{cong}} \mathbf{c} e \quad (78)$$

where the left-hand side has some type $u \equiv v$ but the right-hand side has type $\mathbf{c} \bar{u}' \equiv \mathbf{c} \bar{v}'$ where u and v occur in $\bar{u}'$ and $\bar{v}'$ respectively.

To see how higher-dimensional unification can be applied in a concrete situation, we now show a more limited but still very useful corollary of Theorem 64. In particular, this corollary formally justifies the notion of a *forced constructor argument* (Brady et al., 2003).

Definition 65 (Forced constructor argument). A variable x occurs *rigidly* in a term t if either $t = x$ or t is of the form $\mathbf{c} t_1 \dots t_n$ where x occurs rigidly in one of the t_i and t_i is not a forced argument of $\mathbf{c}$.

An argument of a constructor $\mathbf{c} : \Delta \rightarrow \mathbf{D} \bar{u}$ is *forced* if it occurs rigidly in one of the indices $\bar{u}$.

As an example, the argument $(n : \mathbb{N})$ is a forced argument of $\text{cons} : (n : \mathbb{N})(x : A)(xs : \text{Vec } A \ n) \rightarrow \text{Vec } A \ (\text{succ } n)$. The concept of a forced constructor argument was introduced by Brady et al. (2003) for efficiently compiling dependently typed programs. We use it here to describe when it is safe to apply the **injectivity** rule.

Definition 66 (Invertible constructor). A constructor $\mathbf{c} : \Delta \rightarrow \mathbf{D} \bar{u}$ is *invertible* if the indices $\bar{u}$ consist only of invertible constructors and variables bound in Δ , and no variable from Δ occurs more than once in a non-forced position in $\bar{u}$.

In particular, constructors of non-indexed datatypes such as $\mathbb{N}$ are always invertible. Likewise, the constructor $\text{cons} : (n : \mathbb{N})(x : A)(xs : \text{Vec } A \ n) \rightarrow \text{Vec } A \ (\text{succ } n)$ is

invertible. In contrast, $\text{refl} : u \equiv_A u$ is not an invertible constructor unless u consists itself completely of invertible constructors: $\text{refl} : \text{zero} \equiv_{\mathbb{N}} \text{zero}$ is invertible, but $\text{refl} : n \equiv_{\mathbb{N}} n$ for variable $n : \mathbb{N}$ is not.

Now we can justify why forced constructor arguments (of invertible constructors) can be skipped during unification.

Corollary 67. *Let $c : \Delta \rightarrow \mathbf{D} \bar{u}$ be an invertible constructor (Definition 66) of the datatype $\mathbf{D} : \Xi \rightarrow \mathbf{Set}_\ell$ and $\bar{t}_1, \bar{t}_2 : \Delta$. Then we have an equivalence*

$$(e : c \bar{t}_1 \equiv_{\mathbf{D} \bar{v}} c \bar{t}_2) \simeq (\bar{e}' : \bar{t}_1|_{\Delta'} \equiv_{\Delta'} \bar{t}_2|_{\Delta'}) \quad (79)$$

where $\bar{v} = \bar{u}[\Delta \mapsto \bar{t}_1] = \bar{u}[\Delta \mapsto \bar{t}_2]$, Δ' is the telescope of non-forced arguments of c with the forced arguments filled in with the corresponding values from $\bar{t}_1$, and $\bar{t}|_{\Delta'}$ is the sublist of $\bar{t}$ corresponding to the variables occurring in Δ' .

From the well-formedness of the type $c \bar{t}_1 \equiv_{\mathbf{D} \bar{v}} c \bar{t}_2$, it follows that the forced arguments of $\bar{t}_1$ and $\bar{t}_2$ are equal, so it doesn't matter from which side we take them.

Construction. By definition of an invertible constructor, applying unification to the unification problem $\Delta(\bar{p} : \bar{u} \equiv_{\Xi} \bar{v})$ ends in a positive success with result $f : \Delta(\bar{p} : \bar{u} \equiv_{\Xi} \bar{v}) \simeq \Delta'$ where the computational behaviour of f is to select the non-forced arguments of c from Δ . Applying Theorem 64 to f gives us the desired equivalence. $\square$

In particular, when applying **injectivity** to an equation $c \bar{s} = c \bar{t}$, this corollary tells us that it is safe to skip unification of the forced arguments of c . This allows us to avoid some situations where unification would otherwise require the **deletion** rule.

Example 68. Let $\text{Fin} : \mathbb{N} \rightarrow \mathbf{Set}$ be the following datatype:

$$\begin{aligned} \text{data } \text{Fin} : \mathbb{N} \rightarrow \mathbf{Set} \text{ where} \\ \text{fzero} & : (n : \mathbb{N}) \rightarrow \text{Fin} (\text{suc } n) \\ \text{fsuc} & : (n : \mathbb{N}) \rightarrow \text{Fin } n \rightarrow \text{Fin} (\text{suc } n) \end{aligned} \quad (80)$$

We apply Corollary 67 to solve the equation $\text{fsuc } n x \equiv_{\text{Fin} (\text{suc } n)} \text{fsuc } n y$. Since the first argument of the constructor fsuc is forced, the corresponding equation $n = n$ can be skipped. So we get an equivalence of type $\text{fsuc } n x \equiv_{\text{Fin} (\text{suc } n)} \text{fsuc } n y \simeq x \equiv_{\text{Fin } n} y$. In particular, we do not need to solve the (forced) equation $n \equiv_{\mathbb{N}} n$.

6.3 Lifting unifiers to higher dimensions

We have seen how to apply higher-dimensional unification to make the injectivity rule more generally applicable. In this section, we dive into the heart of the problem. Our core result that makes higher-dimensional unification work is Lemma 72, telling us exactly how to update the left- and right-hand sides of the equations when lifting a unifier.

Suppose we have a unifier that we want to lift to a higher dimension. As a first attempt, we try to apply the following theorem from The Univalent Foundations Program (2013):

Figure 10: The **Square** type represents the possible ways to fill a square defined by four equality proofs.

Theorem 69. *If a function $f : A \rightarrow B$ is an equivalence and $x, y : A$, then $\text{cong } f : x \equiv_A y \rightarrow f x \equiv_B f y$ is also an equivalence.*

Construction. This is Theorem 2.11.1 from The Univalent Foundations Program (2013). $\square$

Applying this theorem to a unifier $f : \Gamma(\bar{p} : \bar{a} \equiv_{\Delta} \bar{b}) \simeq \Gamma'$ results in an equivalence $\text{cong } f : (\bar{e} : \bar{u}; \bar{r} \equiv_{\Gamma(\bar{p} : \bar{a} \equiv_{\Delta} \bar{b})} \bar{v}; \bar{s}) \simeq (\bar{e}' : f \bar{u} \bar{r} \equiv_{\Gamma'} f \bar{v} \bar{s})$, or expanding the definition of telescopic equality:

$$\text{cong } f : (\bar{e} : \bar{u} \equiv_{\Gamma} \bar{v})(\bar{q} : \bar{r} \equiv_{\bar{a}_e \equiv_{\Delta_e} \bar{b}_e} \bar{s}) \simeq (\bar{e}' : f \bar{u} \bar{r} \equiv_{\Gamma'} f \bar{v} \bar{s}) \quad (81)$$

where $\bar{u}, \bar{v} : \Gamma$, $\bar{r} : \bar{a}_u \equiv_{\Delta_u} \bar{b}_u$, $\bar{s} : \bar{a}_v \equiv_{\Delta_v} \bar{b}_v$, and $\cdot_x$ is shorthand for $[\Gamma \mapsto \bar{x}]$. This is already *almost* what we need for higher-dimensional unification, but not quite.

To better visualize the problem, we make use of the concept of a *square*, also called a *2-path* by The Univalent Foundations Program (2013):

Definition 70 (Square). Let $A : \mathbf{Set}_\ell$, $w, x, y, z : A$, $t : w \equiv_A x$, $b : y \equiv_A z$, $l : w \equiv_A y$, and $r : x \equiv_A z$. The *square type* **Square** $t b l r$ is defined to be the dependent equality type $l \equiv_{t \equiv_A b} r$.

The type $l \equiv_{t \equiv_A b} r$ can be written a little more explicitly as $l \equiv_{\equiv_{\equiv_A}^{t; b}} r$, or even more explicitly as $\text{subst}(_ \equiv_A _)(t; b) l \equiv_{x \equiv_A z} r$. If we imagine a square with top side t , bottom side b , left side l , and right side r , then **Square** $t b l r$ can be thought of as the type of identity proofs that fill this square horizontally as visualized in Figure 10a.

There is a second way to construct a square type from four given points $w, x, y, z : A$ and equality proofs $t : w \equiv_A x$, $b : y \equiv_A z$, $l : w \equiv_A y$: we can ‘flip’ the square around its w - z axis, as illustrated by Figure 10b. To get to our desired result, we need to rely on the fact that both square types are in fact equivalent:

Lemma 71 (Flipping squares). *Let $A : \mathbf{Set}$, $w, x, y, z : A$, $t : w \equiv_A x$, $b : y \equiv_A z$, $l : w \equiv_A y$, and $r : x \equiv_A z$. Then we have an equivalence $\text{flip } t b l r : \text{Square } t b l r \simeq \text{Square } l r t b$.*

Proof. The proof of this lemma consists completely of repeated applications of **J**. We start by constructing the function $\text{flip } t b l r : \text{Square } t b l r \rightarrow \text{Square } l r t b$.

First, by **J** on t and b we can assume that $w = x$, $y = z$ and both t and b are **refl**, so we are left with the goal $l \equiv_{w \equiv_A y} r \rightarrow \text{refl} \equiv_{l \equiv_A r} \text{refl}$. The identity type in the function argument has become homogeneous, so we again apply **J**, giving us that $l = r$ and leaving us with the goal $\text{refl} \equiv_{l \equiv_A l} \text{refl}$. Finally, one more application of **J** on $l : w \equiv_A y$ leaves us with the goal $\text{refl} \equiv_{w \equiv_A w} \text{refl}$, which we solve with **refl**.

For the construction of the left and right inverse of **flip**, we just change the order of t, b, l and r in the construction of **flip**. For the proofs that they are in fact inverses, the same sequence of applications of **J** as used in the construction of **flip** suffices. $\square$

Now we prove the main lemma. When we applied this lemma in the last section, we only used it for $\bar{r} = \overline{\text{refl}}$ and $\bar{s} = \overline{\text{refl}}$, but the fully general version is not harder to prove so that's what we present here.

Lemma 72 (Lifting of unifiers). *Suppose we have a unifier $f : \Gamma(\bar{p} : \bar{a} \equiv_{\Delta} \bar{b}) \simeq \Gamma'$ and $\bar{u}, \bar{v} : \Gamma$, $\bar{r} : \bar{a}_u \equiv_{\Delta_u} \bar{b}_u$ and $\bar{s} : \bar{a}_v \equiv_{\Delta_v} \bar{b}_v$.⁶ Then we have a lifted unifier*

$$\begin{aligned} f^\uparrow : (\bar{e} : \bar{u} \equiv_{\Gamma} \bar{v})(\bar{p} : \overline{\text{cong}}(\lambda \Gamma. \bar{a}) \bar{e} \equiv_{\bar{r} \equiv_{\Delta_e} \bar{s}} \overline{\text{cong}}(\lambda \Gamma. \bar{b}) \bar{e}) \\ \simeq (\bar{e}' : f \bar{u} \bar{r} \equiv_{\Gamma'} f \bar{v} \bar{s}) \end{aligned} \quad (82)$$

Construction. By Theorem 69 we already have the equivalence in (81). By Lemma 60, the type of $\bar{q}$ is equivalent to $\bar{r} \equiv_{\overline{\text{cong}}(\lambda \Gamma. \bar{a}) \bar{e} \equiv_{\Delta_e} \overline{\text{cong}}(\lambda \Gamma. \bar{b}) \bar{e}} \bar{s}$. If we think of this type as a square type, then Lemma 71 gives us that this type is equivalent to $\overline{\text{cong}}(\lambda \Gamma. \bar{a}) \bar{e} \equiv_{\bar{r} \equiv_{\Delta_e} \bar{s}} \overline{\text{cong}}(\lambda \Gamma. \bar{b}) \bar{e}$. This is illustrated in Figure 11. Composing this equivalence with **cong** f gives us the desired equivalence $f^\uparrow$. $\square$

Lemma 73. *If $f : \Gamma(\bar{p} : \bar{a} \equiv_{\Delta} \bar{b}) \simeq \Gamma'$ is a strong unifier, then so is $f^\uparrow$.*

Proof. $f^\uparrow$ is constructed as a composition of the equivalences **cong** f (Theorem 69), Lemma 60, and **flip** (Lemma 71). By Lemma 50, we just have to verify that each of these equivalences is a strong unification rule. But this can be verified by looking at their construction (in the case of **cong** f also using the fact that f is a strong unifier). $\square$

7 Translation to eliminators

In this section, we translate the injectivity, conflict, and cycle rules from Section 4 to the ‘bare metal’ of type theory: datatype eliminators. These eliminators encode the basic induction principles associated to each datatype. By translating definitions by pattern matching to eliminators, we can be confident that they do not add any extra assumptions to the core theory.

For the rest of this section, let $\mathbf{D} : \Xi \rightarrow \mathbf{Set}_i$ be an inductive family (where Ξ is the telescope of the indices) with constructors $\mathbf{c}_1, \dots, \mathbf{c}_k$. Without loss of generality,

⁶ We again write $\cdot_x$ for $[\Gamma \mapsto \bar{x}]$

Figure 11: To construct the equivalence in Lemma 72, we apply Lemma 71 to transform the horizontal filling $\bar{p}$ into a vertical one $\bar{q}$.

we assume that the non-recursive constructor arguments come before the recursive ones, so $\mathbf{c}_i$ has type:

$$\mathbf{c}_i : \Delta_i \rightarrow (\Phi_{i,1} \rightarrow \mathbf{D} \bar{v}_{i,1}) \rightarrow \dots \rightarrow (\Phi_{i,n_i} \rightarrow \mathbf{D} \bar{v}_{i,n_i}) \rightarrow \mathbf{D} \bar{u}_i \quad (83)$$

We consider $\mathbf{D}$ to be already applied to its parameters, if it has any.

Definition 74. The standard datatype eliminator $\mathbf{elim}_{\mathbf{D}}$ for $\mathbf{D}$ has type

$$\begin{aligned} \mathbf{elim}_{\mathbf{D}} & : (P : \bar{\mathbf{D}} \rightarrow \mathbf{Set}_i)(m_1 : M_1) \dots (m_k : M_k) \\ & \rightarrow (\bar{x} : \bar{\mathbf{D}}) \rightarrow P \bar{x} \end{aligned} \quad (84)$$

where the methods $m_1, \dots, m_k$ have type

$$\begin{aligned} M_i & = (\bar{t} : \Delta_i)(x_1 : \Phi_{i,1} \rightarrow \mathbf{D} \bar{v}_{i,1}) \dots (x_{n_i} : \Phi_{i,n_i} \rightarrow \mathbf{D} \bar{v}_{i,n_i}) \\ & \rightarrow (h_1 : (\bar{s}_1 : \Phi_{i,1}) \rightarrow P \bar{v}_{i,1} (x_1 \bar{s}_1)) \rightarrow \dots \\ & \rightarrow (h_{n_i} : (\bar{s}_{n_i} : \Phi_{i,n_i}) \rightarrow P \bar{v}_{i,n_i} (x_{n_i} \bar{s}_{n_i})) \\ & \rightarrow P \bar{u}_i (\mathbf{c}_i \bar{t} x_1 \dots x_{n_i}) \end{aligned} \quad (85)$$

The evaluation behaviour of the standard datatype eliminator is given by the following rule for $i = 1, \dots, k$:

$$\begin{aligned} \mathbf{elim}_{\mathbf{D}} P m_1 \dots m_k \bar{u}_i (\mathbf{c}_i \bar{t} x_1 \dots x_{n_i}) & = \\ m_i \bar{t} x_1 \dots x_{n_i} & \\ (\lambda \bar{s}_1. \mathbf{elim}_{\mathbf{D}} P m_1 \dots m_k \bar{v}_{i,1} (x_1 \bar{s}_1)) & \\ \dots & \\ (\lambda \bar{s}_{n_i}. \mathbf{elim}_{\mathbf{D}} P m_1 \dots m_k \bar{v}_{i,n_i} (x_{n_i} \bar{s}_{n_i})) & \end{aligned} \quad (86)$$

Example 75. Following the pedagogy of McBride, Goguen, and McKinna (2006), we take binary trees as our example datatype. This type $\mathbf{Tree}$ is defined by the two

constructors **leaf** : **Tree** and **node** : **Tree** $\rightarrow$ **Tree** $\rightarrow$ **Tree**. The eliminator for **Tree** is

$$\begin{aligned} \text{elim}_{\text{Tree}} & : (P : \text{Tree} \rightarrow \text{Set}_i)(m_{\text{leaf}} : P \text{ leaf}) \\ & \rightarrow (m_{\text{node}} : (l r : \text{Tree}) \rightarrow P l \rightarrow P r \rightarrow P (\text{node } l r)) \\ & \rightarrow (x : \text{Tree}) \rightarrow P x \end{aligned} \quad (87)$$

The evaluation rules are

$$\text{elim}_{\text{Tree}} P m_{\text{leaf}} m_{\text{node}} \text{ leaf} = m_{\text{leaf}} \quad (88)$$

and

$$\begin{aligned} \text{elim}_{\text{Tree}} P m_{\text{leaf}} m_{\text{node}} (\text{node } l r) = \\ m_{\text{node}} l r (\text{elim}_{\text{Tree}} P m_{\text{leaf}} m_{\text{node}} l) (\text{elim}_{\text{Tree}} P m_{\text{leaf}} m_{\text{node}} r) \end{aligned} \quad (89)$$

Case analysis **case_D** is a weakened version of the standard eliminator without the inductive hypotheses.

Lemma 76. *We have a function **case_D** of type*

$$\begin{aligned} \text{case}_D & : (P : \bar{D} \rightarrow \text{Set}_i)(m_1 : M_1) \dots (m_k : M_k) \\ & \rightarrow (\bar{x} : \bar{D}) \rightarrow P \bar{x} \end{aligned} \quad (90)$$

where

$$\begin{aligned} M_i & : (\bar{t} : \Delta_i) \rightarrow (x_1 : \Phi_{i,1} \rightarrow D \bar{v}_{i,1}) \dots (x_{n_i} : \Phi_{i,n_i} \rightarrow D \bar{v}_{i,n_i}) \\ & \rightarrow P \bar{u}_i (\mathbf{c}_i \bar{t} x_1 \dots x_{n_i}) \end{aligned} \quad (91)$$

for $i = 1, \dots, k$.

Example 77. For the **Tree** type, we have:

$$\begin{aligned} \text{case}_{\text{Tree}} & : (P : \text{Tree} \rightarrow \text{Set}_i) \rightarrow P \text{ leaf} \\ & \rightarrow ((l r : \text{Tree}) \rightarrow P (\text{node } l r)) \rightarrow (x : \text{Tree}) \rightarrow P x \\ \text{case}_{\text{Tree}} P m_{\text{leaf}} m_{\text{node}} t & = \text{elim}_{\text{Tree}} P m_{\text{leaf}} (\lambda l r h_l h_r. m_{\text{node}} l r) t \end{aligned} \quad (92)$$

Example 78. For the type $m \leq n$, we have:

$$\begin{aligned} \text{case}_{\leq} & : (P : (m : \mathbb{N})(n : \mathbb{N})(x : m \leq n) \rightarrow \text{Set}_i) \\ & \rightarrow (m_{\text{lz}} : (m : \mathbb{N}) \rightarrow P \text{ zero } m (\text{lz } m)) \\ & \rightarrow (m_{\text{ls}} : (m : \mathbb{N})(n : \mathbb{N})(x : m \leq n) \rightarrow P (\text{suc } m) (\text{suc } n) (\text{ls } m n x)) \\ & \rightarrow (m : \mathbb{N})(n : \mathbb{N})(x : m \leq n) \rightarrow P m n x \end{aligned} \quad (93)$$

Construction of **case_D**.

$$\text{case}_D P m_1 \dots m_k = \text{elim}_D P (\lambda \bar{t} \bar{x} \bar{h}. m_1 \bar{t} \bar{x}) \dots (\lambda \bar{t} \bar{x} \bar{h}. m_k \bar{t} \bar{x}) \quad (94)$$

□

7.1 No confusion

Two of the unification rules, injectivity and conflict, are instances of a more general principle known as ‘no confusion’. In this section, we construct this principle internally as an equivalence **noConf_D**.

We first define an auxiliary type in order to give a general type to **noConf_D**.

Lemma 79. We have a type $\text{NoConfusion}_{\mathbf{D}} : \bar{\mathbf{D}} \rightarrow \bar{\mathbf{D}} \rightarrow \mathbf{Set}_d$ such that

$$\begin{aligned} \text{NoConfusion}_{\mathbf{D}}(\bar{u}; \mathbf{c}_i \bar{s})(\bar{v}; \mathbf{c}_i \bar{t}) &= \bar{s} \equiv_{\Delta_i} \bar{t} \\ \text{NoConfusion}_{\mathbf{D}}(\bar{u}; \mathbf{c}_i \bar{s})(\bar{v}; \mathbf{c}_j \bar{t}) &= \perp \quad (\text{when } i \neq j) \end{aligned} \quad (95)$$

On the diagonal (where we have two times the same constructor), $\text{NoConfusion}_{\mathbf{D}}$ only requires $\bar{s} \equiv_{\Delta_c} \bar{t}$. From this it follows that $\bar{u} \equiv_{\Xi} \bar{v}$ as well, since the indices are determined by the choice of constructor and its arguments.

Example 80. For the $\mathbf{Tree}$ datatype, $\text{NoConfusion}_{\mathbf{Tree}}$ is defined as follows:

$$\begin{aligned} \text{NoConfusion} : \mathbf{Tree} &\rightarrow \mathbf{Tree} \rightarrow \mathbf{Set} \\ \text{NoConfusion leaf leaf} &= \top \\ \text{NoConfusion leaf (node } l \ r) &= \perp \\ \text{NoConfusion (node } l \ r) \text{ leaf} &= \perp \\ \text{NoConfusion (node } l_1 \ r_1) \text{ (node } l_2 \ r_2) &= (l_1 \equiv_{\mathbf{Tree}} l_2) \times (r_1 \equiv_{\mathbf{Tree}} r_2) \end{aligned} \quad (96)$$

Construction of $\text{NoConfusion}_{\mathbf{D}}$. We apply $\text{case}_{\mathbf{D}}$ with the motive $\lambda _ . \bar{\mathbf{D}} \rightarrow \mathbf{Set}_i$. For each method $m_i \bar{x}$, we apply $\text{case}_{\mathbf{D}}$ again with motive $\lambda _ \rightarrow \mathbf{Set}$. This gives us k^2 methods $m_{i,j}$ to fill in, one for each pair of constructors. On the diagonal (where $i = j$) we define $m_{ii} = \lambda \bar{x}; \bar{x}' . \bar{x} \equiv_{\Delta_i} \bar{x}'$, and if $i \neq j$ we give $m_{i,j} = \lambda \bar{x}; \bar{x}' . \perp$. $\square$

Lemma 81. We have an equivalence

$$\text{noConf}_{\mathbf{D}} : (\bar{x} \ \bar{y} : \bar{\mathbf{D}}) \rightarrow (\bar{x} \equiv_{\bar{\mathbf{D}}} \bar{y}) \simeq \text{NoConfusion}_{\mathbf{D}} \ \bar{x} \ \bar{y} \quad (97)$$

Moreover, for any constructor $\mathbf{c} : \Delta \rightarrow \mathbf{D}$ and $\bar{s}, \bar{s}' : \Delta$, this equivalence satisfies $\text{noConf}_{\mathbf{D}}^{-1}(\bar{u}[\Delta \mapsto \bar{s}]; \mathbf{c} \ \bar{s})(\bar{u}[\Delta \mapsto \bar{s}']; \mathbf{c} \ \bar{s}') = \overline{\text{dcong}}(\lambda \bar{x} . \bar{u}; \mathbf{c} \ \bar{x})$.

Example 82. For $\mathbf{Tree}$, the function $\text{noConf}_{\mathbf{Tree}}$ gives for any two trees s and t that are equal a proof of $\text{NoConfusion}_{\mathbf{Tree}} \ s \ t$:

$$\text{noConf}_{\mathbf{Tree}} : (s \ t : \mathbf{Tree}) \rightarrow (s \equiv_{\mathbf{Tree}} t) \simeq \text{NoConfusion}_{\mathbf{Tree}} \ s \ t \quad (98)$$

If s and t are of the form $\text{node } l_1 \ r_1$ and $\text{node } l_2 \ r_2$ respectively, then this gives us the injectivity rule $(\text{node } l_1 \ r_1 \equiv_{\mathbf{Tree}} \text{node } l_2 \ r_2) \simeq (l_1 \equiv_{\mathbf{Tree}} l_2 \times r_1 \equiv_{\mathbf{Tree}} r_2)$. On the other hand, if s is of the form leaf and t is of the form $\text{node } l \ r$, then we get the conflict rule $(\text{leaf} \equiv_{\mathbf{Tree}} \text{node } l \ r) \simeq \perp$.

Construction of $\text{noConf}_{\mathbf{D}}$. First, we define the left-to-right function $\text{noConf}_{\mathbf{D}} \ \bar{a} \ \bar{b}$. To do this, we apply telescopic substitution $\overline{\text{subst}}$ with motive $\text{NoConfusion}_{\mathbf{D}} \ \bar{a}$. This reduces the problem to finding a function of type

$$(\bar{a} : \bar{\mathbf{D}}) \rightarrow \text{NoConfusion}_{\mathbf{D}} \ \bar{a} \ \bar{a} \quad (99)$$

But this can be done using $\text{case}_{\mathbf{D}}$ with motive $\lambda \bar{a} . \text{NoConfusion}_{\mathbf{D}} \ \bar{a} \ \bar{a}$, filling in $\overline{\text{refl}}$ for each method $m_i \bar{x}$.

For the inverse $\text{noConf}_{\mathbf{D}}^{-1} \ \bar{a} \ \bar{b}$, we need to do a little more work. First, we apply $\text{case}_{\mathbf{D}}$ twice as in the definition of $\text{NoConfusion}_{\mathbf{D}}$. Now we are left to give methods

$$m_{i,j} : \text{NoConfusion}_{\mathbf{D}}(\bar{u}_i; \mathbf{c}_i \bar{x})(\bar{u}'_j; \mathbf{c}_j \bar{x}') \rightarrow \bar{u}_i(\mathbf{c}_i \bar{x}) \equiv_{\bar{\mathbf{D}}} \bar{u}'_j(\mathbf{c}_j \bar{x}') \quad (100)$$

When $i \neq j$, this is easy: we get an element of type $\perp$ from $\text{NoConfusion}_{\mathbf{D}}$, from which we can conclude anything. On the diagonal (where $i = j$) we get a proof of

$\bar{x} \equiv_{\Delta_i} \bar{x}'$. Applying $\overline{\text{dcong}}$ to this equality gives us $\bar{u}_i; (\mathbf{c}_i \bar{x}) \equiv_{\bar{\mathbf{D}}} \bar{u}'_i; (\mathbf{c}_i \bar{x}')$, which is what we need.

Next, we prove that this is a left inverse by constructing a function of type

$$(\bar{a} \bar{b} : \bar{\mathbf{D}})(\bar{e} : \bar{a} \equiv_{\bar{\mathbf{D}}} \bar{b}) \rightarrow \text{noConf}_{\bar{\mathbf{D}}}^{-1} \bar{a} \bar{b} (\text{noConf}_{\bar{\mathbf{D}}} \bar{a} \bar{b} \bar{e}) \equiv_{\bar{a} \equiv_{\bar{\mathbf{D}}} \bar{b}} \bar{e} \quad (101)$$

By $\bar{\mathbf{J}}$, it is sufficient to give a function of type

$$(\bar{a} : \bar{\mathbf{D}}) \rightarrow \text{noConf}_{\bar{\mathbf{D}}}^{-1} \bar{a} \bar{a} (\text{noConf}_{\bar{\mathbf{D}}} \bar{a} \bar{a} \overline{\text{refl}}) \equiv_{\bar{a} \equiv_{\bar{\mathbf{D}}} \bar{a}} \overline{\text{refl}} \quad (102)$$

But this we can do by applying $\text{case}_{\bar{\mathbf{D}}}$ with methods $m_i \bar{x} = \overline{\text{refl}}$.

All that's left to do is to prove that it is a right inverse as well. To construct the proof isRinv that

$$(\bar{a} \bar{b} : \bar{\mathbf{D}})(e : \text{NoConfusion}_{\bar{\mathbf{D}}} \bar{a} \bar{b}) \rightarrow \text{noConf}_{\bar{\mathbf{D}}} \bar{a} \bar{b} (\text{noConf}_{\bar{\mathbf{D}}}^{-1} \bar{a} \bar{b} e) \equiv_{\text{NoConfusion}_{\bar{\mathbf{D}}} \bar{a} \bar{b}} e \quad (103)$$

we first apply case analysis on $\bar{a}$ and $\bar{b}$. In the cases where we have two distinct constructors $\mathbf{c}_i$ and $\mathbf{c}_k$, we have $e : \perp$ so we can conclude by $\text{elim}_{\perp}$. In the diagonal cases we have $e : \bar{s} \equiv_{\Delta_i} \bar{t}$. Eliminating these equations with $\bar{\mathbf{J}}$ leaves us with the goal $\overline{\text{refl}} \equiv_{\bar{s} \equiv_{\Delta_i} \bar{s}} \overline{\text{refl}}$, which we solve by giving $\overline{\text{refl}}$. $\square$

7.2 Acyclicity

The other property of datatypes we need is acyclicity: a term can never be structurally smaller than itself. This property is used for implementing the cycle detection rule of the unification algorithm. Internally, it is represented by the term $\text{noCycle}_{\bar{\mathbf{D}}}$. To express its type, we first define what it means for a term to (not) be structurally smaller than some other term.

We first define the auxiliary type $\text{Below}_{\bar{\mathbf{D}}}$: $\text{Below}_{\bar{\mathbf{D}}} P \bar{u} x$ is defined as a tuple type that is inhabited whenever $P \bar{v} y$ holds for all $y : \bar{\mathbf{D}} \bar{v}$ that are structurally smaller than $x : \bar{\mathbf{D}} \bar{u}$.

Lemma 83. *Let $P : \bar{\mathbf{D}} \rightarrow \text{Set}_i$. For any $x : \bar{\mathbf{D}} \bar{u}$, we have a type $\text{Below}_{\bar{\mathbf{D}}} P \bar{u} x$ such that for any $y \prec x$ we have a projection $\pi : \text{Below}_{\bar{\mathbf{D}}} P \bar{u} x \rightarrow P \bar{v} y$.*

Example 84. The type $\text{Below}_{\text{Tree}} P x$ expresses that the property $P : \text{Tree} \rightarrow \text{Set}$ holds for any subtree of $x : \text{Tree}$. In other words, we have

$$\begin{aligned} \text{Below}_{\text{Tree}} P \text{ leaf} &= \top \\ \text{Below}_{\text{Tree}} P (\text{node } l r) &= (\text{Below}_{\text{Tree}} P l \times P l) \times (\text{Below}_{\text{Tree}} P r \times P r) \end{aligned} \quad (104)$$

Construction of $\text{Below}_{\bar{\mathbf{D}}} P$. We apply the eliminator $\text{elim}_{\bar{\mathbf{D}}}$ to the motive $\Phi = \lambda _ . \text{Set}_i$. For the method m_i corresponding to the constructor $\mathbf{c}_i$ we give the following:

$$\begin{aligned} m_i &= \lambda \bar{t}; x_1; \dots; x_{n_i}; h_1; \dots; h_{n_i}. \\ &\quad ((\bar{s}_1 : \Phi_{i,1}) \rightarrow h_1 \bar{s}_1 \times P \bar{v}_{i,1} (x_1 \bar{s}_1)) \times \\ &\quad \dots \times ((\bar{s}_{n_i} : \Phi_{i,n_i}) \rightarrow h_{n_i} \bar{s}_{n_i} \times P \bar{v}_{i,n_i} (x_{n_i} \bar{s}_{n_i})) \end{aligned} \quad (105)$$

To construct the projection π , consider $x : \bar{\mathbf{D}} \bar{u}$ and any structurally smaller term $y : \bar{\mathbf{D}} \bar{v}$. If y is (an application of) a direct subterm of x , say $x = \mathbf{c} \bar{t} x_1 \dots x_n$ with

$y = x_i \bar{w}$, then we return the second component of the i th component of $\text{Below}_D P x$, i.e. we define

$$\pi H = \pi_2 (\pi_i H \bar{w}) : \text{Below}_D P \bar{u} x \rightarrow P \bar{v} y \quad (106)$$

Otherwise, y is a subterm of some direct subterm x_i of $x = \bar{c} \bar{t} x_1 \dots x_n$. In particular, by induction we have some $\pi' : \text{Below}_D P \bar{v}_i x_i \rightarrow P \bar{v} y$. This allows us to define π as follows:

$$\pi H = \pi' (\pi_1 (\pi_i H)) : \text{Below}_D P \bar{u} x \rightarrow P \bar{v} y \quad (107)$$

□

Lemma 85. We have a type $_ \not\prec_D _ : \bar{D} \rightarrow \bar{D} \rightarrow \text{Set}_d$ such that for any $x : D \bar{u}$ and $y : D \bar{v}$ with $x \prec y$, we have $x \not\prec_D y \rightarrow \perp$. We also define $\bar{a} \not\prec_D \bar{b} := \bar{a} \not\prec_D \bar{b} \times \bar{a} \not\equiv_D \bar{b}$.

If $x : D \bar{u}$ and $y : D \bar{v}$ then we often leave the indices implicit and write $x \not\prec_D y$ and $x \not\prec_D y$ instead of $\bar{u}; x \not\prec_D \bar{v}; y$ and $\bar{u}; x \not\prec_D \bar{v}; y$.

Example 86. The type $x \not\prec_{\text{Tree}} t$ expresses that x is *not* a subtree of t . In particular, we have the following equalities:

$$\begin{aligned} x \not\prec_{\text{Tree}} \text{leaf} &= \top \\ x \not\prec_{\text{Tree}} (\text{node } l \ r) &= ((x \not\prec_{\text{Tree}} l) \times (x \not\equiv_{\text{Tree}} l)) \times ((x \not\prec_{\text{Tree}} r) \times (x \not\equiv_{\text{Tree}} r)) \end{aligned} \quad (108)$$

Construction of $\not\prec_D$. We define $\not\prec_D$ in terms of Below_D :

$$\bar{a} \not\prec_D \bar{b} := \text{Below}_D (\lambda \bar{b}'. \bar{a} \not\equiv_D \bar{b}') \bar{b} \quad (109)$$

By definition of Below_D , we have a projection $\pi : \bar{u}; x \not\prec_D \bar{v}; y \rightarrow \bar{u}; x \not\equiv_D \bar{u}; x$ whenever $x \prec y$. Filling in $\overline{\text{refl}}$ for the proof of $\bar{u}; x \not\equiv_D \bar{u}; x$ gives us the desired proof of $x \not\prec_D y \rightarrow \perp$. □

Now we can state the property that no term can be structurally smaller than itself.

Lemma 87. We have a function $\text{noCycle}_D : (\bar{a} \bar{b} : \bar{D}) \rightarrow \bar{a} \equiv_D \bar{b} \rightarrow \bar{a} \not\prec_D \bar{b}$.

Example 88. $\text{noCycle}_{\text{Tree}}$ is the proof that no tree can ever be a subtree of itself, i.e. every well-typed tree is well-founded.

Construction of noCycle_D . Note that

$$\begin{aligned} x \not\prec_D \bar{c}_i \bar{t} x_1 \dots x_{n_i} &= ((\bar{s}_1 : \Phi_{i,1}) \rightarrow x \not\prec_D x_1 \bar{s}_1) \times \dots \\ &\quad \times ((\bar{s}_{n_i} : \Phi_{i,n_i}) \rightarrow x \not\prec_D x_{n_i} \bar{s}_{n_i}) \end{aligned} \quad (110)$$

by definition of Below_D and $\not\prec_D$. Now to construct noCycle_D , we start by eliminating the equation $\bar{a} \equiv_D \bar{b}$ using $\bar{J}$, which leaves us the goal $(\bar{a} : \bar{D}) \rightarrow \bar{a} \not\prec_D \bar{a}$. Next we apply case_D with motive $\lambda \bar{a}. \bar{a} \not\prec_D \bar{a}$, producing for each constructor $\bar{c}_i : \Delta_i \rightarrow (\Phi_{i,1} \rightarrow D \bar{v}_{i,1}) \rightarrow \dots \rightarrow (\Phi_{i,n_i} \rightarrow D \bar{v}_{i,n_i}) \rightarrow D \bar{u}_i$ the subgoal

$$\begin{aligned} (\bar{t} : \Delta_i)(x_1 : \Phi_{i,1} \rightarrow D \bar{v}_{i,1}) \dots (x_{n_i} : \Phi_{i,n_i} \rightarrow D \bar{v}_{i,n_i}) \rightarrow \\ (h_1 : (\bar{s}_1 : \Phi_{i,1}) \rightarrow x_1 \bar{s}_1 \not\prec_D x_1 \bar{s}_1) \dots \\ (h_{n_i} : (\bar{s}_{n_i} : \Phi_{i,n_i}) \rightarrow x_{n_i} \bar{s}_{n_i} \not\prec_D x_{n_i} \bar{s}_{n_i}) \rightarrow \\ \bar{c}_i \bar{t} x_1 \dots x_{n_i} \not\prec_D \bar{c}_i \bar{t} x_1 \dots x_{n_i} \end{aligned} \quad (111)$$

To continue, we define the auxiliary types $\mathbf{Step}_{i,j}$ for $i = 1, \dots, k$ and $j = 1, \dots, n_i$ as follows:

$$\begin{aligned} \mathbf{Step}_{i,j} : & (\bar{t} : \Delta_i)(x_1 : \Phi_{i,1} \rightarrow \mathbf{D} \bar{v}_{i,1}) \dots (x_{n_i} : \Phi_{i,n_i} \rightarrow \mathbf{D} \bar{v}_{i,n_i}) \rightarrow \\ & (\bar{s} : \Phi_{i,j})(\bar{a} : \bar{\mathbf{D}}) \rightarrow \mathbf{Set}_d \\ \mathbf{Step}_{i,j} \bar{t} x_1 \dots x_{n_i} \bar{s} (\bar{u}; b) = & (x_j \bar{s}) \not\prec_{\mathbf{D}} b \rightarrow (\mathbf{c}_i \bar{t} x_1 \dots x_{n_i}) \not\prec_{\mathbf{D}} b \end{aligned} \quad (112)$$

In what follows, we will construct

$$\mathbf{step}_{i,j} : (\bar{t} : \Delta_i)(x_1 : \Phi_{i,1} \rightarrow \mathbf{D} \bar{v}_{i,1}) \dots (x_{n_i} : \Phi_{i,n_i} \rightarrow \mathbf{D} \bar{v}_{i,n_i}) \rightarrow (\bar{s} : \Phi_{i,j})(\bar{a} : \bar{\mathbf{D}}) \rightarrow \mathbf{Step}_{i,j} \bar{t} x_1 \dots x_{n_i} \Phi_{i,j} \bar{a} \quad (113)$$

Once this is done, we solve the subgoal (111) by filling in

$$\begin{aligned} & \lambda \bar{t}; x_1; \dots; x_{n_i}; h_1; \dots; h_{n_i}. \\ & (\lambda \bar{s}_1. \mathbf{step}_{i,1} \bar{t} \bar{x} \bar{s}_1 (\bar{v}_{i,1}; (x_1 \bar{s}_1)) (h_1 \bar{s}_1)), \dots, \\ & (\lambda \bar{s}_{n_i}. \mathbf{step}_{i,n_i} \bar{t} \bar{x} \bar{s}_{n_i} (\bar{v}_{i,n_i}; (x_{n_i} \bar{s}_{n_i})) (h_{n_i} \bar{s}_{n_i})) \end{aligned} \quad (114)$$

So we only need to construct $\mathbf{step}_{i,j}$.

The construction of $\mathbf{step}_{i,j} \bar{t} x_1 \dots x_{n_i} \bar{s} : (\bar{a} : \bar{\mathbf{D}}) \rightarrow \mathbf{Step}_{i,j} \bar{t} x_1 \dots x_{n_i} \bar{s} \bar{a}$ proceeds by applying $\mathbf{elim}_{\mathbf{D}}$ with the motive $\mathbf{Step}_{i,j} \bar{t} x_1 \dots x_{n_i} \bar{s}$. The new subgoals are of the form

$$\begin{aligned} & (\bar{t}' : \Delta_p)(x'_1 : \Phi_{p,1} \rightarrow \mathbf{D} \bar{v}'_{p,1}) \dots (x'_{n_p} : \Phi_{p,n_p} \rightarrow \mathbf{D} \bar{v}'_{p,n_p}) \rightarrow \\ & (h'_1 : (\bar{s}'_1 : \Phi_{p,1}) \rightarrow \mathbf{Step}_{i,j} \bar{t} x_1 \dots x_{n_i} \bar{s} \bar{v}'_{p,1} (x'_1 \bar{s}'_1)) \dots \\ & (h'_{n_p} : (\bar{s}'_{n_p} : \Phi_{p,n_p}) \rightarrow \mathbf{Step}_{i,j} \bar{t} x_1 \dots x_{n_i} \bar{s} \bar{v}'_{p,n_p} (x'_{n_p} \bar{s}'_{n_p})) \rightarrow \\ & \mathbf{Step}_{i,j} \bar{t} x_1 \dots x_{n_i} \bar{s} \bar{u}'_p (\mathbf{c}_p \bar{t}' x'_1 \dots x'_{n_p}) \end{aligned} \quad (115)$$

We solve them by giving:

$$\lambda \bar{t}'; x'_1; \dots; x'_{n_p}; h'_1; \dots; h'_{n_p}; H. \alpha, \beta \quad (116)$$

where we still have to construct

$$\alpha : \bar{u}_i; (\mathbf{c}_i \bar{t} x_1 \dots x_{n_i}) \not\prec_{\mathbf{D}} \bar{u}'_p; (\mathbf{c}_p \bar{t}' x'_1 \dots x'_{n_p}) \quad (117)$$

and

$$\beta : \bar{u}_i; (\mathbf{c}_i \bar{t} x_1 \dots x_{n_i}) \not\equiv_{\bar{\mathbf{D}}} \bar{u}'_p; (\mathbf{c}_p \bar{t}' x'_1 \dots x'_{n_p}) \quad (118)$$

We have $H : x_j \bar{s} \not\prec_{\mathbf{D}} \mathbf{c}_p \Delta'_p x'_1 \dots x'_{n_p}$ or, by definition of $\not\prec_{\mathbf{D}}$, $H = (H_1, \dots, H_{n_p})$ where $H_q : (\bar{s}' : \Phi'_{p,q}) \rightarrow x_j \bar{s} \not\prec_{\mathbf{D}} x'_q \bar{s}'$. The construction of α reduces to the construction of components $\alpha_q : (\bar{s}' : \Phi'_{p,q}) \rightarrow \mathbf{c}_i \bar{t} x_1 \dots x_{n_i} \not\prec_{\mathbf{D}} x'_q \bar{s}'$. But these we can give as $\alpha_q = \lambda \bar{s}'. h'_q \bar{s}' (\pi_1 (H_p \bar{s}'))$.

For constructing β , we assume $\bar{u}_i; (\mathbf{c}_i \bar{t} x_1 \dots x_{n_i}) \equiv_{\bar{\mathbf{D}}} \bar{u}'_p; (\mathbf{c}_p \bar{t}' x'_1 \dots x'_{n_p})$ and derive an element of $\perp$. By $\mathbf{noConf}_{\mathbf{D}}$, it suffices to consider the case where $i = p$ and $\bar{t}; x_1; \dots; x_{n_i} = \bar{t}'; x'_1; \dots, x'_{n_p}$. But then we have $H_j \bar{s} : x_j \bar{s} \not\prec_{\mathbf{D}} x'_j \bar{s}$, hence $\pi_2 (H_j \bar{s}) \mathbf{refl} : \perp$. This finishes the construction of $\mathbf{noCycle}_{\mathbf{D}}$. $\square$

8 Implementation

Using our framework for proof-relevant unification described in this paper, we reimplemented the unification algorithm used by Agda for checking definitions by

dependent pattern matching. As a result, we were able to replace previous ad hoc restrictions with formally verified unification rules, fixing a number of bugs in the process. It also enabled us to add new unification rules dealing with η -equality for record types, as well as higher-dimensional unification for solving equations between constructors of indexed datatypes. Another advantage of our approach is that the implementation is now much cleaner than before, allowing it to be extended easily in the future. In this section, we take a look at our implementation from the point of view of an Agda user (Section 8.1) and an Agda developer (Section 8.2).

8.1 Impact on the Agda user

From the point of view of a user of Agda, unification happens behind the scenes while checking definitions by pattern matching, so a different algorithm doesn't impact the syntax of the language directly. Instead, the main criterion a user of Agda should judge the unification algorithm by is that it accepts the definitions that should be accepted, and rejects the definitions that should be rejected. The latter can be seen from the fact that our implementation directly resulted in a fix for issue #1408 (Example 41), dealing with an incompatibility between heterogeneous equations and the `-without-K` option (Vezzosi, 2015). Equally important, our implementation provides a much more principled solution to issues #292 (Danielsson, 2010, see also Example 41), #1071 (Danielsson, 2014), #1406 (Abel, 2015a, see also Example 40), #1411 (Abel, 2015b), and #1427 (Abel, 2015c, see also Example 40). All these issues are fixed without introducing special cases in the code and without limiting the power of the unification algorithm in any significant way, as can be seen from the fact that Agda's test suite and standard library are still typechecked correctly. This is in contrast to the previous ad-hoc fixes to some of these issues, which broke the unification algorithm in some cases, for example in issue #1435 (Danielsson, 2015).

The addition of the new unification rules for η -equality of record values also significantly improved the way Agda handles records. Before these unification rule were added to Agda, all variables of record type had to be fully eta-expanded before calling the unifier, for example in issue #473 (Danielsson, 2011). This caused a substantial overhead when dealing with deeply nested records, see issue #635 (Peebles, 2012). This also caused problems in combination with Agda's instance search mechanism, see for example issue #1613 (Abel, 2015d). In contrast, by using this unification rule we only eta-expand a variable when it is useful for the unification to proceed, thus eliminating this overhead.

We also implemented higher-dimensional unification (Section 6). This addition allows Agda to typecheck more definitions, such as the example given in issue #1775 (Sicard-Ramírez, 2016).

8.2 Impact on the Agda codebase

For the further development of Agda, it is important that the unification machinery is robust and easily extensible with further rules. For this reason, we separated it

```

data UnifyState = UState
  { varTel    :: Telescope
  , flexVars  :: FlexibleVars
  , eqTel     :: Telescope
  , eqLHS     :: [Term]
  , eqRHS     :: [Term]
  }

data UnifyStep
  = Deletion           { ... }
  | Solution           { ... }
  | Injectivity        { ... }
  | Conflict           { ... }
  | Cycle              { ... }
  | EtaExpandVar       { ... }
  | EtaExpandEquation  { ... }
  | LitConflict        { ... }
  | StripSizeSuc       { ... }
  | SkipIrrelevantEquation { ... }
  | TypeConInjectivity { ... }

```

Figure 12: The datatypes used for representing unification states and unification rules closely follow the theory. In addition to the unification rules presented in this paper, Agda also has unification rules for dealing with literals, sized types (Abel, 2010) and irrelevant equations (Abel, 2011), features not discussed in this paper. There is also a rule for injective type constructors that is only used when this is enabled explicitly by the user.

into two logical parts: a *unification strategy* and the *unification engine*. Both parts make use of the same data structures for representing the unification state and unification rules, as shown in Figure 12. The unification strategy takes a unification state as an argument and produces a lazy monadic list of unification rules to try (Figure 13), while the unification engine tries to apply these rules one by one until one succeeds (Figure 14).

A big difference between our implementation and Agda’s previous unification algorithm is that our version explicitly manipulates telescopes of free variables (`varTel`) and equations (`eqTel`) as well as explicit substitutions between these telescopes, while previously these had to be reconstructed after unification was finished. This change resulted in a significant simplification of the code for checking left-hand sides and coverage of definitions by pattern matching (the parts of Agda that use the unification algorithm).

An important choice when constructing a unification strategy is whether to start on the leftmost or the rightmost equation. It seems sensible to start on the left to avoid heterogeneous equations as much as possible, and this was also the preferred method for the old algorithm. However, our unification rules for indexed datatypes actually benefit from having unsolved equations in the telescope, so a unification strategy that starts from the right provides more opportunities to apply these rules.

```

type UnifyStrategy = UnifyState -> ListT TCM UnifyStep

skipIrrelevantStrategy basicUnifyStrategy
dataStrategy literalStrategy etaExpandVarStrategy
etaExpandEquationStrategy injectiveTypeConStrategy
simplifySizesStrategy checkEqualityStrategy
  :: Int -> UnifyStrategy

```

Figure 13: A unification strategy takes a unification state and produces a list of unification steps to try in order. For constructing unification strategies, we provide a number of basic strategies that can be combined in any order.

```

unifyStep :: UnifyState -> UnifyStep
          -> UnifyM (UnificationResult' UnifyState)

unify :: UnifyState -> UnifyStrategy
      -> UnifyM (UnificationResult' UnifyState)

```

Figure 14: The unification engine consists of an auxiliary function `unifyStep` that tries to apply one unification step, resulting in either a new state, an absurdity (e.g. for the conflict and cycle rules), or a failure, and the main function `unify` that tries all steps suggested by a given strategy, and continues until either the unification problem is solved (i.e. the equation telescope is empty) or there are no more rules left to try.

For this reason, our current implementation uses a right-to-left strategy, although plugging in a different strategy would be trivial.

Our implementation of higher-dimensional unification closely follows the steps in Section 6.2. In particular, when applying the injectivity rule to a unification problem of the form $(\bar{e}_1 : \bar{s}_1 \equiv_{\Phi} \bar{s}_2)(e_2 : \bar{c} \bar{t}_1 \equiv_{\bar{v}_e} \bar{t}_2)$ the unification algorithm constructs the new unification problem $\Phi\Delta(\bar{p} : \bar{u} \equiv_{\Xi} \bar{v})$ and recursively calls itself on this new problem.

One noteworthy fact about the implementation is how the left- and right-hand sides $f \bar{s}_1 \bar{t}_1 \text{ refl}$ and $f \bar{s}_2 \bar{t}_2 \text{ refl}$ of the new unification problem in (73) are computed. The implementation doesn't have an explicit representation of the function f , so it's not possible to calculate them directly. Instead, the recursive call produces a substitution ρ of type $\Delta' \rightarrow \Phi\Delta$. This allows us to calculate $f^{-1} : \Delta' \rightarrow (\bar{x} : \Phi)(\bar{y} : \Delta)(\bar{p} : \bar{u} \equiv_{\Xi} \bar{v})$ as $\lambda \bar{x}'. \bar{x}' \rho; \text{refl}$, but doesn't give us a direct way to compute f . To go in the opposite direction, we note that ρ is a *pattern* with free variables Δ' . So we can match the values from $\Phi\Delta$ against this pattern. The proofs of $\bar{u} \equiv_{\Xi} \bar{v}$ (assumed to be `refl` in our implementation, since all the unifiers we compute are strong unifiers) ensure that this matching cannot fail, so this allows us to recover the values of the variables in Δ' , thus computing the function $f : \Phi\Delta(\bar{p} : \bar{u} \equiv_{\Xi} \bar{v}) \rightarrow \Delta'$.

9 Related work

Unification is a large area of research that we cannot hope to cover here in full. We refer the interested reader to Jouannaud and Kirchner (1990) and Baader and Snyder (2001) for a general overview of the subject. Most extensions to unification that are studied, such as higher-order unification and E-unification, are orthogonal to the work in this paper, although it would be interesting to see how they fit within our framework.

Type checkers of dependently typed languages typically have some facility for meta-variables that are solved by higher order pattern unification (Reed, 2009; Abel and Pientka, 2011). This is not directly related to the work in this paper as the requirements on the unification algorithm are different. For example, these unification algorithms suppose all rigid symbols (including type constructors) to be ‘injective’ for the purpose of unification. Some algorithms even consider defined functions to be rigid (Ziliani and Sozeau, 2015) or make use of user-provided hints to choose one solution over the other (Asperti, Ricciotti, Coen, and Tassi, 2009), thereby giving up on finding most general unifiers in favour of finding solutions more often. In this case, the only problem is that the solution to the metavariable may not be what the user intended. In contrast, our algorithm produces evidence of unification internal to the theory we’re working in, and it is actually important that the unifier found by the algorithm is indeed the most general one (otherwise we might lose e.g. coverage of functions by pattern matching). Still, it would be interesting to further investigate the similarities and differences between these two unification algorithms.

Goguen (1989) takes a categorical view on unification, representing most general unifiers as *equalizers* in a category of types and substitutions. It shouldn’t be surprising that many of the category-theoretic notions are analogous to the type-theoretic ones presented in this paper. For example, giving an explicit type to the domain of substitutions helps to avoid problems with non-uniqueness in the definition of a most general unifier in other presentations. Compared to the category-theoretical presentation of unification, our work adds support for indexed datatypes, and it also differs in the fact that type theory allows an internal representation of equations as (telescopic) equality types.

The idea to represent unification problems at the object level by using the identity type stems from McBride (1998b). In McBride’s paper, the types of equations are limited to simple (non-dependent) types, and the injectivity rule is likewise limited to simple datatypes. Later, he solves this by introducing a heterogeneous identity type (McBride, 2002). However, UIP is needed to turn heterogeneous equalities back into homogeneous ones. Additionally, postponing equations is not supported, as heterogeneous equations can only be turned into homogeneous ones if the types are equal. In our previous work, we solved the problem of requiring UIP, but the unification rules still only worked on the first equation in a telescope (Cockx et al., 2014). As a consequence, we had to limit the injectivity, conflict, and cycle rules to work only in homogeneous situations, while here we can use them in their fully general form.

Our approach to unification is closely related to the notion of inversion of an inductive hypothesis (Cornes and Terrasse, 1995; Monin, 2010). The usual approach to inversion works by crafting a *diagonalizer* that is used as the motive for an eliminator. Unification can also be as an alternative method for proving inversion lemmas (McBride, 1998b). One advantage of the diagonalizer approach is that it moves most of the work to the type level, potentially improving performance of the resulting function. The process of constructing diagonalizers has recently also been automated (Braibant, 2013). However, it requires that the indices of the inductive hypothesis we are inverting can be written as a pattern, which is not always the case (e.g. they may be non-linear), so the approach based on unification seems to be more general. It would be interesting to try to implement an inversion tactic based on the unification algorithm in this paper to compare the power of the two approaches.

The idea to view equality proofs themselves as the subjects of unification is inspired by cubical type theory, where equality proofs are terms viewed ‘one level up’ (Cohen et al., 2016). In fact, if we were working in a cubical type theory, there would be no difference between regular unification and higher-dimensional unification, so the work in this paper could be seen as ‘backporting’ some of the power of cubical type theory back to the (currently) better-understood world of standard intuitionistic type theory.

Compared to our reverse unification rules from Cockx et al. (2016a), higher-dimensional unification takes information into account from the types of the constructors as well as the types of the equation. This difference is similar to the inversion of an inductive hypothesis by using a diagonalizer (Cornes and Terrasse, 1995) versus using unification for the problem (McBride, 1998b).

10 Discussion and future work

In this paper, we present a proof-relevant unification algorithm for use in dependent type theory, where unification rules are represented as terms internal to the type theory. Thus the type system itself enforces the soundness of these unification rules. Moreover, this lets us extend the unification algorithm with new principles in a safe and modular way. For example, we showed how to add two new unification rules for η -equality of record types. As another example, higher-dimensional unification augments the power of the injectivity rule by allowing us to skip unification of forced arguments, yet would be impossible to even formulate for an untyped unification algorithm.

Having an elegant theoretical framework for unification also helped us a lot when implementing it in practice. As a result, the implementation of our algorithm for Agda has become cleaner, more robust, and more easily extensible. We hope this will also be the case for implementers of other dependently typed languages, as it has already been for the Lean theorem prover and the Equations package for Coq.

Other applications of proof-relevant unification. In this work, we focus on one application of proof-relevant unification, namely specialization by unification

and its role in the compilation of dependent pattern matching. However, we believe firmly that it could also be applied elsewhere, for example for metaprogramming or tactic systems.

More unification rules. It would be interesting to further explore the correspondence between unification rules and new features of type theory. For example, it seems that E-unification (unification modulo a set of equations) could correspond to new unification rules for higher inductive types from HoTT. As another example, higher-order (pattern) unification could correspond to functional extensionality as a unification rule. And since the univalence axiom is itself an equivalence, maybe it could be seen as a unification rule as well?

Custom unification rules. We can put the power of unification in the hands of the user by allowing them to define custom unification rules in the form of *hints* (Asperti et al., 2009). For example, if the user provides a proof of $(f\ x \equiv_B f\ y) \simeq (x \equiv_A y)$ for some function $f : A \rightarrow B$, then this could be used as an injectivity rule for f by the unifier. One obstacle is that these rules might not be strong unification rules, so we either have to give up on some computational properties of unifiers or implement a check of strongness of a given unification rule.

Unification with higher inductive types. HoTT introduces the concept of *higher inductive types*, which can have non-trivial identity proofs between their constructors. This implies that in general they do not satisfy the injectivity, disjointness, or acyclicity properties. So to adapt our unification algorithm to a context with higher inductive types, we should start by limiting the unification algorithm further, for example by cutting out the “no confusion” and “cycle” properties for types to which they do not apply.

As a second step, these principles can be replaced by type-specific solvers that exploit any extra structure that may be available.

Example 89. The interval $\mathbf{I}$ is a higher inductive type with two point constructors $0 : \mathbf{I}$ and $1 : \mathbf{I}$ and one path constructor $\text{line} : 0 \equiv_{\mathbf{I}} 1$. We have the following equivalence:

$$\text{contract} : (e : 0 \equiv_{\mathbf{I}} 1) \simeq () \quad (119)$$

By definition we have $\text{contract}^{-1}() = \text{line}$, so if we use this equivalence as a unification rule, we won’t get a strong unification rule as a result. Maybe it is possible to weaken this requirement a bit by not requiring refl as such, but merely *some* canonical form. But this means that we also need computation rules for functions applied to higher constructors, which is still an open problem. So for now, we have to settle for a weaker kind of unification rules that do not have the proper definitional behaviour, but still produce an equivalence of the correct type.

More generally, the “no confusion” principle is similar to the encode/decode technique used by Licata and Shulman (2013) and McKinna and Forsberg (2015) to calculate the fundamental group of the circle. In particular, they also construct an equivalence between an equality/path type and a type of *codes* taking the role

of our `NoConfusion` type. So it may be possible to construct a new unification rule for the circle type based on this equivalence. However, be aware that these custom unification rules can introduce additional variables, for example the rule for the circle introduces a variable of type $\mathbb{Z}$!

Unification in cubical type theory. Our unification algorithm is developed for an Agda-like theory based on standard MLTT. In such a theory, principles such as functional extensionality or univalence can be postulated but they do not get any computational behaviour. On the other hand, a new and promising theory called *cubical* type theory gives a constructive interpretation to the univalence axiom, and hence also functional extensionality (Bezem, Coquand, and Huber, 2014; Cohen et al., 2016). In the future we would like to adapt the work in this paper to this setting, so it would become usable in languages based on cubical type theory as well.

One obstacle for this adaptation is the fact that the representation of datatypes in our theory (and also that of Agda, Coq, Idris, ...) is computationally incompatible with functional extensionality. We give an example to illustrate the problem.⁷

Example 90. Let `Favourite` : $(\mathbb{N} \rightarrow \mathbb{N}) \rightarrow \mathbf{Set}$ be a datatype with one constructor `favourite` : `Favourite` $(\lambda x. 0 + x)$. We can give a proof p of $(x : \mathbb{N}) \rightarrow 0 + x \equiv x + 0$, so we have `funext` p : $\lambda x. 0 + x \equiv \lambda x. x + 0$ and thence

$$\mathbf{subst} \text{ Favourite } (\mathbf{funext} \ p) \ \mathbf{favourite} : \text{Favourite } (\lambda x. x + 0) \quad (120)$$

However, there is no closed canonical form of type `Favourite` $(\lambda x. x + 0)$ so this term doesn't reduce to a canonical form. This cannot be fixed by taking the constructor itself to be the canonical form (i.e. by letting `favourite` : `Favourite` $(\lambda x. x + 0)$), as this would require the typechecker to check whether two functions are extensionally equal, which is undecidable in general.

This incompatibility could be solved by disallowing indexed datatypes and instead having each constructor carry explicit proofs of the constraints it imposes on the former indices. For example, `favourite` would have the internal type $(e : f \equiv (\lambda x. 0 + x)) \rightarrow \text{Favourite } f$. The surface-level constructor is then represented as `favourite refl`, while `subst Favourite (funext p) favourite` computes to `favourite (funext p)`. With this representation of datatypes, the work done in this paper is just as necessary as before, since we still need unification to solve the (telescopic) equations embedded in the constructors, as well as equations between these embedded equality proofs.

Example 91. We illustrate this by working out Example 63 again for a version of the `Vec` datatype with embedded equality proofs instead of indices. Suppose `Vec` A n is defined with constructors `nil` : $n \equiv_{\mathbb{N}} \mathbf{zero} \rightarrow \text{Vec } A \ n$ and `cons` : $(m : \mathbb{N})(x : A)(xs : \text{Vec } A \ m) \rightarrow n \equiv_{\mathbb{N}} \mathbf{suc} \ m \rightarrow \text{Vec } A \ n$ and consider the unification problem:

$$(e : \mathbf{cons} \ n \ x \ xs \ \mathbf{refl} \equiv_{\text{Vec } A} (\mathbf{suc} \ n) \ \mathbf{cons} \ n \ y \ ys \ \mathbf{refl}) \quad (121)$$

⁷ Thanks to Conor McBride for pointing out the problem and giving this example.

Since this version of the `Vec` datatype doesn't have an index, we can apply the `injectivitycons` rule to simplify this equation to:

$$\begin{aligned} & (e_1 : n \equiv_{\mathbb{N}} n)(e_2 : x \equiv_A y)(e_3 : xs \equiv_{\text{Vec}} A\ e_1\ ys) \\ & (e_4 : \text{refl} \equiv_{\text{Suc}} n \equiv_{\mathbb{N}} \text{Suc}\ e_1\ \text{refl}) \end{aligned} \quad (122)$$

Now e_4 is an equation between equality proofs, much like the one we obtained in (69), except that the equality $(p : \text{Suc}\ e_1 \equiv_{\text{Suc}} n \equiv_{\mathbb{N}} \text{Suc}\ n\ \text{Suc}\ n)$ is replaced with an equality $(e_4 : \text{refl} \equiv_{\text{Suc}} n \equiv_{\mathbb{N}} \text{Suc}\ e_1\ \text{refl})$. Lemma 71 shows that these two types are in fact equivalent. So higher-dimensional unification problems also occur in languages without indexed datatypes, and hence that a general way to solve this kind of equations is equally useful in these languages.

Acknowledgements

We want to thank the anonymous reviewers for their hard work and insightful comments. Dominique Devriese holds a Postdoctoral fellowship from the Research Foundation Flanders (FWO).

Bibliography

- Andreas Abel. MiniAgda: Integrating sized and dependent types. In *Workshop on Partiality and Recursion in Interactive Theorem Provers*, PAR, 2010.
- Andreas Abel. Irrelevance in type theory with a heterogeneous equality judgement. In *Foundations of Software Science and Computational Structures*. 2011.
- Andreas Abel. Injectivity of type constructors is partially back. Agda refutes excluded middle, 2015a. URL <https://github.com/agda/agda/issues/1406>. (on the Agda bug tracker).
- Andreas Abel. Order of patterns matters for checking left hand sides, 2015b. URL <https://github.com/agda/agda/issues/1411>. (on the Agda bug tracker).
- Andreas Abel. Circumvention of forcing analysis brings back easy proof of Fin injectivity, 2015c. URL <https://github.com/agda/agda/issues/1427>. (on the Agda bug tracker).
- Andreas Abel. Eta-expanded implicit patterns are not used for instance search, 2015d. URL <https://github.com/agda/agda/issues/1613>. (on the Agda bug tracker).
- Andreas Abel and Brigitte Pientka. Higher-order dynamic pattern unification for dependent types and records. In *International Conference on Typed Lambda Calculi and Applications*, TLCA, pages 10–26. Springer, 2011.
- Andrea Asperti, Wilmer Ricciotti, Claudio Sacerdoti Coen, and Enrico Tassi. Hints in unification. In *Theorem Proving in Higher Order Logics*, 2009.
- Franz Baader and Wayne Snyder. Unification theory. *Handbook of automated reasoning*, 2001.
- Marc Bezem, Thierry Coquand, and Simon Huber. A model of type theory in cubical sets. In *Types for Proofs and Programs*, TYPES, 2014.
- Edwin Brady. Idris, a general purpose dependently typed programming language: Design and implementation. *Journal of Functional Programming*, 23(5), 2013.

- Edwin Brady, Conor McBride, and James McKinna. Inductive families need not store their indices. In *Types for Proofs and Programs*, TYPES, 2003.
- Thomas Braibant. A new Coq tactic for inversion, 2013. URL <http://gallium.inria.fr/blog/a-new-Coq-tactic-for-inversion>.
- Jesper Cockx. *Dependent pattern matching and proof-relevant unification*. PhD thesis, KU Leuven, 2017.
- Jesper Cockx and Dominique Devriese. Lifting proof-relevant unification to higher dimensions. In *6th Conference on Certified Programs and Proofs*, CPP. ACM, 2017.
- Jesper Cockx, Dominique Devriese, and Frank Piessens. Pattern matching without K. In *19th International Conference on Functional Programming*, ICFP. ACM, 2014.
- Jesper Cockx, Dominique Devriese, and Frank Piessens. Unifiers as equivalences: proof-relevant unification of dependently typed data. In *21th International Conference on Functional Programming*, ICFP. ACM, 2016a.
- Jesper Cockx, Dominique Devriese, and Frank Piessens. Eliminating dependent pattern matching without K. *Journal of Functional Programming*, 26, 2016b.
- Cyril Cohen, Thierry Coquand, Simon Huber, and Anders Mörtberg. Cubical type theory: a constructive interpretation of the univalence axiom. *CoRR*, 2016. URL <http://arxiv.org/abs/1611.02108>.
- Thierry Coquand. Pattern matching with dependent types. In *Types for Proofs and Programs*, TYPES, 1992.
- Cristina Cornes and Delphine Terrasse. Automating inversion of inductive predicates in Coq. In *Types for Proofs and Programs*, TYPES. 1995.
- Nils Anders Danielsson. Heterogenous equality is crippled by the Bool /= Fin 2 fix, 2010. URL <https://github.com/agda/agda/issues/292>. (on the Agda bug tracker).
- Nils Anders Danielsson. The unification machinery does not respect η -equality, 2011. URL <https://github.com/agda/agda/issues/473>. (on the Agda bug tracker).
- Nils Anders Danielsson. Regression in unifier, possibly related to modules and/or heterogeneous constraints, 2014. URL <https://github.com/agda/agda/issues/1071>. (on the Agda bug tracker).
- Nils Anders Danielsson. Dependent pattern matching is broken, 2015. URL <https://github.com/agda/agda/issues/1435>. (on the Agda bug tracker).
- Nicolaas Govert de Bruijn. Telescopic mappings in typed lambda calculus. *Information and Computation*, 91(2), 1991.
- Leonardo de Moura, Soonho Kong, Jeremy Avigad, Floris van Doorn, and Jakob von Raumer. The Lean theorem prover (system description). In *25th International Conference on Automated Deduction*, CADE, 2015.
- Gabe Dijkstra. Disunifying non-fully applied constructors is inconsistent with function extensionality, 2015. URL <https://github.com/agda/agda/issues/1497>. (on the Agda bug tracker).
- Peter Dybjer. Inductive sets and families in Martin-Löf's type theory and their set-theoretic semantics. In *Proceedings of the first workshop on Logical frameworks*, 1991.

- Healfdene Goguen, Conor McBride, and James McKinna. Eliminating dependent pattern matching. In *Algebra, Meaning, and Computation*. 2006.
- Joseph A. Goguen. What is unification? – A categorical view of substitution, equation and solution. In *Resolution of Equations in Algebraic Structures, Volume 1: Algebraic Techniques*, 1989.
- Chung Kil Hur. Agda with the excluded middle is inconsistent?, 2010. URL <https://lists.chalmers.se/pipermail/agda/2010/001522.html>. On the Agda mailing list.
- Jean-Pierre Jouannaud and Claude Kirchner. *Solving equations in abstract algebras: A rule-based survey of unification*. 1990.
- Daniel R. Licata and Michael Shulman. Calculating the fundamental group of the circle in homotopy type theory. In *28th Symposium on Logic in Computer Science, LICS*, 2013.
- Zhaohui Luo. *Computation and reasoning: a type theory for computer science*, volume 11 of *International Series of Monographs on Computer Science*. 1994.
- Per Martin-Löf. An intuitionistic theory of types. In *Twenty-five years of constructive type theory (Venice, 1995)*, pages 127–172. Oxford University Press, 1972.
- Per Martin-Löf. *Intuitionistic type theory*. Number 1 in Studies in Proof Theory. 1984.
- Conor McBride. Towards dependent pattern matching in LEGO. Unpublished, 1998a.
- Conor McBride. Inverting inductively defined relations in LEGO. In *Types for Proofs and Programs, TYPES*, 1998b.
- Conor McBride. *Dependently typed functional programs and their proofs*. PhD thesis, University of Edinburgh, 2000.
- Conor McBride. Elimination with a motive. In *Types for Proofs and Programs, TYPES*, 2002.
- Conor McBride, Healfdene Goguen, and James McKinna. A few constructions on constructors. In *Types for Proofs and Programs, TYPES*, 2006.
- James McKinna and Fredrik Nordvall Forsberg. The encode-decode method, relationally. In *Types for Proofs and Programs, TYPES*, 2015.
- Alexandre Miquel. Re: Agda with the excluded middle is inconsistent?, 2010. URL <https://lists.chalmers.se/pipermail/agda/2010/001543.html>. Proof posted by Chung-Kil Hur on the Agda mailing list.
- Jean-François Monin. Proof trick: Small inversions. In *Second Coq Workshop*, 2010.
- Ulf Norell. *Towards a practical programming language based on dependent type theory*. PhD thesis, Chalmers University of Technology, 2007.
- Daniel Peebles. Case splitting emits hidden record patterns that should remain implicit, 2012. URL <https://github.com/agda/agda/issues/635>. (on the Agda bug tracker).
- Jason Reed. Higher-order constraint simplification in dependent type theory. In *4th International Workshop on Logical Frameworks and Meta-Languages: Theory and Practice*, pages 49–56. ACM, 2009.

- Andrés Sicard-Ramírez. The `--without-K` option generates unsolved metas, 2016. URL <https://github.com/agda/agda/issues/1775>. (on the Agda bug tracker).
- Matthieu Sozeau. Equations: A dependent pattern-matching compiler. In *Interactive theorem proving*, ITP, 2010.
- Thomas Streicher. Investigations into intensional type theory, 1993. Habilitation thesis, Ludwig Maximilian University of Munich.
- The Univalent Foundations Program. *Homotopy Type Theory: Univalent Foundations of Mathematics*. <http://homotopytypetheory.org/book>, Institute for Advanced Study, 2013.
- Andrea Vezzosi. Heterogeneous equality incompatible with univalence even `--without-k`, 2015. URL <https://github.com/agda/agda/issues/1408>. (on the Agda bug tracker).
- Beta Ziliani and Matthieu Sozeau. A unification algorithm for Coq featuring universe polymorphism and overloading. In *International Conference on Functional Programming*, ICFP, 2015.

A Incompatibility of injective type constructors with univalence and excluded middle

The proofs of these two theorems are not essential for the understanding of the work in this paper. However, to our knowledge there is no easy reference for them and we think they are interesting enough to mention here.

Theorem 92. *MLTT extended with univalence and injective type constructors is inconsistent.*

Proof. Let $D : \mathbf{Set} \rightarrow \mathbf{Set}$ be an inductive family with no constructors. Then $D \top \simeq \perp \simeq D \perp$, so $D \top \equiv_{\mathbf{Set}} D \perp$ by univalence. But if D is injective, this means that $\top \equiv_{\mathbf{Set}} \perp$, which is clearly a contradiction. $\square$

Theorem 93. *MLTT extended with the excluded middle and injective type constructors is inconsistent.*

Proof. This proof is based on the proof given by Hur (2010).

Assume the datatype $D : (\mathbf{Set} \rightarrow \mathbf{Set}) \rightarrow \mathbf{Set}$ is injective (it doesn't matter what the constructors of D are). We define a right inverse E of D as follows: if A is equal to $D F$ for some $F : \mathbf{Set} \rightarrow \mathbf{Set}$, then $E A$ is defined to be that F , otherwise it is $\lambda _ . \perp$. Formally, E is defined by case analysis on `excluded-middle` applied to $(\mathbf{Image} D A)$, where $\mathbf{Image} D$ is a datatype with a single constructor `image` : $(F : \mathbf{Set} \rightarrow \mathbf{Set}) \rightarrow \mathbf{Image} (D F)$.

We have

$$E (D F) \equiv_{\mathbf{Set} \rightarrow \mathbf{Set}} F \tag{A 1}$$

for any F : because $D F$ is certainly in the image of D , $E (D F)$ must be equal to G for *some* G with $D G \equiv_{\mathbf{Set}} D F$, but then this G must be equal to F by injectivity of D .

Now we construct by diagonalization a $\mathbf{C} : \mathbf{Set} \rightarrow \mathbf{Set}$ that is not in the image of $\mathbf{E}$, thus leading to a contradiction. $\mathbf{C} A$ is defined by case analysis on `excluded-middle` applied to $(\mathbf{E} A A \equiv_{\mathbf{Set}} \perp)$: if $\mathbf{E} A A$ is equal to $\perp$, then $\mathbf{C} A = \top$, otherwise $\mathbf{C} A = \perp$.

To come to the contradiction, consider the term $\mathbf{B} = \mathbf{E} (\mathbf{D} \mathbf{C}) (\mathbf{D} \mathbf{C})$. By (A 1), we have $\mathbf{B} \equiv_{\mathbf{Set}} \mathbf{C} (\mathbf{D} \mathbf{C})$. Is $\mathbf{B}$ equal to $\perp$ or not? By the excluded middle, there are two cases:

$\mathbf{B} \equiv_{\mathbf{Set}} \perp$: then we have $\mathbf{B} \equiv_{\mathbf{Set}} \mathbf{C} (\mathbf{D} \mathbf{C}) \equiv_{\mathbf{Set}} \top$ by definition of $\mathbf{C}$, but this is a contradiction with $\mathbf{B} \equiv_{\mathbf{Set}} \perp$.

$(\mathbf{B} \equiv_{\mathbf{Set}} \perp) \rightarrow \perp$: then we have $\mathbf{B} \equiv_{\mathbf{Set}} \mathbf{C} (\mathbf{D} \mathbf{C}) \equiv_{\mathbf{Set}} \perp$ again by definition of $\mathbf{C}$, but this is a contradiction with $(\mathbf{B} \equiv_{\mathbf{Set}} \perp) \rightarrow \perp$.

We have constructed an element of $\perp$ in the empty context, so we conclude that MLTT extended with the excluded middle and injective type constructors is inconsistent. $\square$