

Abbreviations, unification, and type inference

Didier Rémy

Draft version of May 6, 2020

Abstract

Modern ML-style type inference relies on an extension of first-order unification to efficiently deal with type generalization. However, ML-like type systems, and especially OCaml also allows type abbreviations, including recursive and deconstructive type abbreviations. We show how first-order unification can be extended to efficiently deal with type abbreviations and remain compatible with type inference.

Both deconstructive type abbreviations, which constrain their arguments to be of a certain shape, and projective type abbreviations, which just returns one of their argument, must be expanded eagerly, while purely constructive type abbreviations, which do not examine their parameters and construct a new type structure from their parameters, can be expanded lazily. General type abbreviations that mix both may be decomposed into a destructive part destructive part, which must be expanded eagerly and a constructive part which can still be expanded lazily.

Recursive type abbreviations are solved by encoding them into non-recursive type abbreviations that expand to recursive types.

Contents

1	Introduction	3
2	Type abbreviations in ML	3
3	ML-style type inference	5
3.1	Constraints generation	5
3.2	Constraints resolution	5
3.3	Unification in plain ML	5
4	Type abbreviations	5
4.1	Notations	5
4.2	Types	5
4.3	Ground types	6
4.4	Type abbreviations	6
4.5	Recursive definitions of type abbreviations	8
4.6	Deconstructive type abbreviations	10
5	Multi-equations	10
5.1	Constraints and multi-equations	10
5.2	Canonical forms for constraints	10
5.3	Multiequations and views	11
5.4	Unification without deconstructive abbreviations	11
5.5	Simplification rules	12
5.6	Invariants	13
6	Unification strategy	14
6.1	Invariants of composed rules	15
6.2	Algorithm	16
6.3	Termination	16
6.4	Efficiency of the strategy	18
6.5	Solved forms	18
7	Deconstructive type abbreviations	19
7.1	A purely deconstructive type abbreviation	19
7.2	Mixed type abbreviations	21
7.3	The general case	21
7.4	Further Encodings	24
8	Converting between types and canonical forms	24
8.1	Input types	24
8.2	Output types	25
8.3	Removing auxiliary type abbreviations	26
8.4	Example	27
8.5	Computing free variables	27

1 Introduction

Modern ML-style type inference relies on an extension of first-order unification to efficiently deal with type generalization. However, ML-like type systems, and especially OCaml also allows type abbreviations. In particular, OCaml allows deconstructive type abbreviations, which uses only subterms of their parameters in their expansion.

A semantically correct and trivial treatment of type abbreviations is to first expand while as a preliminary step. However, this defeats the original purpose of abbreviations, which is to improve not only writing of input types but also and especially reading of inferred types. That is, we wish to keep user-provided abbreviations as long as possible. Besides, type abbreviations may also improve efficiency of type inference by concisely representing some common patterns. This, however, requires to work with type abbreviations and expand them lazily, on demand while preserving sharing of expanded forms.

This is particularly crucial in OCaml for object types whose expanded forms can sometimes be quite large (more than an order of magnitude). Currently, OCaml already performs lazy expansion, but does not in an optimal way, and its approach has not been formalized.

We show how to efficiently extend first-order unification to deal with type abbreviations. Type abbreviations are kept in inferred types, but choice of which form—abbreviated or expanded—to show, or rather which strategy to use when printing types is left to the user.

Both deconstructive type abbreviations, which constrain their arguments to be of a certain shape, and projective type abbreviations, which just returns one of their argument, must be expanded eagerly, while purely constructive type abbreviations, which do not examine their parameters and construct a new type structure from their parameters, can be expanded lazily. General type abbreviations that mix both may be decomposed into a destructive part destructive part, which must be expanded eagerly and a constructive part which can still be expanded lazily.

Recursive type abbreviations are solved by encoding them into non-recursive type abbreviations that expand to recursive types.

2 Type abbreviations in ML

A basic type abbreviation could be of the form:

```
type var = string * int
```

This allows to write `var` instead of `string * int` just for conciseness, to avoid mistakes, or to be more readable. While abbreviations are introduced by the user in his programs,

```
let string_of_var (x : var) = (fst x) ^ string_of_int (snd x)
```

the user also expect the inferred type be printed with abbreviations:

```
unfreeze : var -> string
```

This requires that type inference keeps track of abbreviations, even when they must be expanded during type checking to check the compatibility with expanded forms. That is, the user expects to see types as abbreviations—where they were originally abbreviations and not otherwise. In particular, every pair `string * int` should be abbreviated as `var`. For example, when writing

```
let adult (name : string, age) = age > 18
```

the user would be surprised (and unhappy) if the inferred type were `var -> bool` instead of `string * int -> bool`.

That is, the user does not expect type inference to use abbreviation in the reserve way for contracting all types that could be printed as abbreviations (even if this could sometimes be useful in some very specific cases).

Odd usage of regular type abbreviations Type abbreviations may also be used in less usual ways to carry extra information along *phantom parameters*, i.e., parameters that do not appear in the expanded form, as in:

```
type 'a str = string
let one : int str = "1"
```

As these parameters are ignored, they are not decomposable and cannot be source of a clash:

```
let f : bool str = "false"
```

Indeed, one of type `int str` and `f` of type `bool str` are compatible types since they both expand to `string`.

Somewhat similarly, they may also be used as a marker to make a type look different while it is not—the type abbreviation marker will vanish when in conflict.

```
type 'a special = 'a
let one : int special = 1
let r = one + 2
```

Although, the type `int special` of `one` looks as a special kind of an `int`, it is nothing else but `int`.

Destructive type abbreviations More interestingly, we may define abbreviations whose type parameters are constrained:

```
type 'a elem = 'b constraint 'a = 'b list
```

That is, `'a elem` requires `'a` to be of the form `b list` and is then equal to `'b`. The above definition is intuitively equivalent to the following definition (illegal in OCaml):

```
type ('a list) elem = 'a
```

It allows to write:

```
let cons (h : 'a elem) (t : 'a) : 'a = h :: t
```

Notice that while type abbreviations have been available in OCaml for about two decades, they have never been formalized and typechecking them still raises issues in corner cases—typechecking the above expression erroneously loops in OCaml (version 4.09). Arguments of deconstructive type abbreviations may themselves be deconstructive, indeed:

```
let cons_elem (x : 'a elem elem) (t : 'a) : 'a = [x] :: t
```

Deconstructive abbreviations may still be constructive at the same time:

```
type 'a swap = 'c * 'b where 'a = 'b * 'c
let swap (x, y : 'a) : 'a swap = y, x
```

Examples with objects

Type abbreviations and modules

3 ML-style type inference

3.1 Constraints generation

3.2 Constraints resolution

3.3 Unification in plain ML

4 Type abbreviations

4.1 Notations

Given a relation \mathcal{R} , we write \mathcal{R}^* for the reflexive transitive closure of \mathcal{R} and \mathcal{R}^∞ for the saturation of \mathcal{R} , i.e., the relation defined by

$$x \mathcal{R}^\infty y \iff x \mathcal{R}^* y \wedge \forall z, \neg(y \mathcal{R} z)$$

We write $\mathcal{R}_1 ; \mathcal{R}_2$ for the composition of relations $\mathcal{R}_2 \circ \mathcal{R}_1$ in inverse order.

4.2 Types

We write α or β for type variables and F for type constructors, which may be primitive type constructors or type abbreviations. Primitive type constructors are such as `int`, `bool`, etc., or parameterized such as `list`, `*`, `→`, etc. Types may also be recursive. Recursive types are defined with the construction $\mu(\alpha) \tau$ and require τ to be productive, i.e., equal to a term of the form $F\bar{\sigma}$ where F is a primitive type constructor.

$$\tau, \sigma ::= \alpha \mid F\bar{\tau} \mid \mu(\alpha) \tau$$

We call primitive types those without type abbreviations. The equality on primitive types is that on iso-recursive types.

$$\begin{array}{c} \text{UNFOLDING} \\ \mu(\alpha) \tau = \tau[\alpha \leftarrow \mu(\alpha) \tau] \end{array} \quad \frac{\text{COINDUCTION} \quad \tau \text{ productive} \quad \tau_1 = \tau[\alpha \leftarrow \tau_1] \quad \tau_2 = \tau[\alpha \leftarrow \tau_2]}{\tau_1 = \tau_2}$$

Type equality also coincide with equality of their representation as infinite regular trees. It can be check by equality of automaton, or more efficiently by the unification algorithm.

Small types are types of depth at most 1, i.e., with type variables or a type constructor (which may be a type abbreviation) applied to a type variable.

We write \mathcal{V} for be the set of type variables We write $\text{ftv}(\tau)$ for the set of type variables occurring in τ . Type constructors come with a fix arity, so that all occurrences of F have the same number of arguments.

4.3 Ground types

Ground types are regular trees built out of type constructors and respecting their arities. That is, they do not contain type variables.

4.4 Type abbreviations

A type abbreviation is a constructor F that comes with a pair written $F\bar{\alpha} \triangleright \sigma_F$ where $\bar{\alpha}$ is a tuple of disjoint type variables such that $\text{ftv}(\sigma_F) \subseteq \bar{\alpha}$.

We write $F_1 \prec F_2$ when F_1 is a base type constructor or F_1 is a type abbreviation $F_1\bar{\alpha} \triangleright \sigma_{F_1}$ and F_2 does not appear in σ_{F_1} .

We assume that \prec is a strict partial ordering, i.e., irreflexive and transitive. We call definitional-ordering any total ordering that extends \prec .

Type abbreviations define rewriting rules $F\bar{\alpha} \triangleright \sigma_F$, which can be applied in any context. This defines a rewriting system \rightarrow , which is terminating on finite terms, as well as on syntactic terms. That is, if a regular tree is defined syntactically with the $\mu(\alpha)\tau$ notation, then rewriting τ with \rightarrow eliminates all type abbreviations in τ . Hence, the fact that syntactic terms represent regular trees does not raise any problem regarding the expansion of type abbreviations, if we perform then on the finite syntactic representation and preserve sharing, i.e., do not perform unfolding.

We write $\nabla\tau$ the full expansion of τ , i.e., the type τ' such that $\tau \rightarrow^\infty \tau'$.

This also defines an equality $=_E$ on ground types obtained by expanding all type abbreviations and testing usual equality on the resulting canonical forms, that is:

$$\tau_1 =_E \tau_2 \iff \nabla\tau_1 = \nabla\tau_2$$

where simple equality is equality without equations but in the presence of recursive types, as usual.

We define the set of free type variables $\text{ftv}(\tau)$ that appear syntactically in τ , as usual, but also the set of type variables remaining after expansion $\text{ftv}^\nabla(\tau)$ equal to $\text{ftv}(\nabla\tau)$.

Phantom and decomposable parameters In particular, type abbreviations allow to define phantom type parameters: A position i in an abbreviation $F\bar{\alpha} \triangleright \sigma_F$ is phantom if $\alpha_i \notin \text{ftv}^\nabla(\sigma)$.

A phantom occurrence is an occurrence in a term whose path contains a phantom position. We can compute $\text{ftv}^\nabla(\tau)$ without computing $\nabla\tau$ just by not following phantom positions of type constructors. Therefore, we can also compute phantom positions of type constructors, in definitional ordering.

A position i in F is decomposable if $F\bar{\tau} = F\bar{\tau}'$ implies $\tau_i = \tau'_i$. A phantom occurrence is actually decomposable when it is not phantom. We write $\downarrow F$ the set of decomposable positions of F .

Head expansion Consider a type abbreviation $F\bar{\alpha} \triangleright \sigma_F$. The expansion $\nabla\sigma_F$ may be a variable or a type constructor.

If $\nabla\sigma_F$ is a type variable β , it must be some α_i of $\bar{\alpha}$: we write $\check{\vee}F = i$ and say that the type abbreviation F (or $F\bar{\alpha}$) is degenerate. Otherwise, $\nabla\sigma_F$ must be of the form $G\bar{\tau}$ where G a primitive constructor: we write $\check{\vee}F = G$ and say that F is productive. Notice that when $F\bar{\alpha}$ is degenerate, all its arguments but one are phantom. We extend $\check{\vee}$ to be the identity function on primitive type constructors.

We can compute $\check{\vee}F$ without expanding σ_F by searching in σ_F for the topmost occurrence of a symbol G that is productive and not in at a phantom path. Then, $\check{\vee}F$ is equal to G if G is primitive or $\check{\vee}G$ otherwise.

When F is productive, the topmost constructor of σ_F is a type constructor F_1 and we write $F \triangleright F_1$.

If $\check{\vee}F = G$, then $F \triangleright^\infty G$. We actually have a chain

$$F = F_0 \triangleright F_1 \triangleright \dots F_n = G$$

whose length n is called the rank of F and written $\text{rank } F$.

The rank can be inductively defined, without computing the chain: taking 0 for the rank of a primitive type constructor, we then have $\text{rank } F = 1 + \text{rank } G$ when $F \triangleright G$. The rank is only defined for productive type abbreviations.

Notice that the relation \triangleright can be represented by a forest where roots are exactly primitive type constructors, and all constructors in a given tree have the same head-constructor—the root of the tree. When F_1 and F_2 are in the same forest, we write $F_1 \vee F_2$ the smallest G such that $F_1 \triangleright^* G$ and $F_2 \triangleright^* G$. We say that a set of type constructors S is *connected* if for any pair of type constructors F and G in S , there is a path subset $F_i^{i \in 1..n}$ in S such that

$$F = F_0 \triangleright F_1 \triangleright \dots F_n = G$$

Compatibility Head expansion may also be used to check whether two type constructors F and G are compatible. Ideally, we wish $F \bowtie G$ to tell that two terms σ and σ' with top constructors F and G cannot be unified. Hence, the definition:

$$F \bowtie_\infty G \iff \exists \bar{\alpha} \bar{\beta} \cdot F\bar{\alpha} \doteq G\bar{\beta} \equiv \text{false}$$

This cannot in general be computed from information coming from F and G independently, and instead need to be computed for every pair (F, G) . In particular, tabulating the results so that \bowtie can be checked in constant time would require space proportional to the square of the number of symbols.

In fact, we need not an exact computation of \bowtie , just an approximation that contains the minimal relation \bowtie_0 defined as the restriction of \neq to primitive type constructors. This allows earlier detection of clashes—without further expansions—and might speed up when unification is more expected to fail than succeed.

That is, we need a relation \bowtie such that $\bowtie_0 \subseteq \bowtie \subseteq \bowtie_\infty$. A (good) possible choice is

$$F \bowtie_1 G \iff \forall F = F_0 \wedge \forall G = G_0 \wedge F_0 \bowtie_0 G_0$$

which can be computed efficiently, as it only requires the precomputation of \forall for every type constructor.

4.5 Recursive definitions of type abbreviations

We wish to also allow recursive type abbreviations $(F_i(\bar{\alpha}) = \sigma_i)^{i \in I}$ where F_j may appear in σ_i . We request however that all occurrences of F_j be applied to the same tuple of variables $\bar{\alpha}$. This ensures that recursion is regular. We may further assume that there is no j such that $F_j \prec F_i$ for all $i \neq j$. Otherwise, F_i could be defined prior to others.

Let $(\beta_i)^{i \in I}$ be a collection of fresh type variables. Let $\hat{\sigma}_i^{i \in I}$ be σ_i where each occurrence of $F\bar{\alpha}_i$ has been replaced by β_i .

We defined the collection of intermediate non recursive type abbreviations $(\hat{F}_i(\bar{\alpha}\bar{\beta}) = \hat{\sigma}_i)^{i \in I}$.

We then redefined $F_i(\bar{\alpha}) \triangleright \mu((\beta_j = \hat{F}_j(\bar{\alpha}, \bar{\beta}))^{j \in J}) \beta_i$ for each i in I , which are all independent, as they only depend on previously defined collection $(\hat{F}_i)^{i \in I}$.

We have used here a general form of recursive types $\mu(\alpha_i = \sigma_i) \sigma$ that generalizes $\mu(\alpha) \sigma$ —but does not add any difficulty: it stands for the type $\phi \sigma$ where ϕ is the principal solution of the equations $\alpha_i = \sigma_i$, that is if I is $1..n$ and ϕ_0 is the identity, then $\phi = \phi_n$ where $\phi_{i+i} = [\alpha_i \mapsto \mu(\alpha_i) \phi \sigma_i] \circ \phi$. However, the syntactic form $\mu((\beta_i = \sigma_i)^{i \in I}) \sigma$ is more compact and we may keep as a possible syntactic form.

If we only introduce \hat{F}_i 's symbols by expansion of F_i 's, then we establish the invariant that *for any occurrence of $\hat{F}_i(\bar{\sigma}, \bar{\tau})$, the term τ_j is always E-equal to $\hat{F}_j(\bar{\sigma}, \bar{\tau})$* . This invariant is preserved by all transformations that preserve E-equality. Besides, when this invariant holds, we also have $\hat{F}_i(\bar{\sigma}, \bar{\tau}) = F_i \bar{\sigma}$. That is, we can see $\hat{F}_i(\bar{\sigma}, \bar{\tau})$ as a version of $F_i \bar{\sigma}$ —instrumented to enable lazy expansion of recursive occurrences. In particular, we may always read back $F_i \bar{\sigma}$ form $\hat{F}_i(\bar{\sigma}, \bar{\tau})$, which will just loose expanded forms, and thus re-expand them again.

Example The most common case is a single recursively defined type abbreviation $F\alpha \triangleright H(\alpha, F\alpha)$. One may think of H as (or being replaced by) a very large

term, which you wish to allocate lazily.

Following, the instructions above, we define an auxiliary abbreviation $\hat{F}(\alpha, \beta) \triangleright H(\alpha, \beta)$ and redefine $F \triangleright \mu(\beta) \hat{F}(\alpha, \beta)$. Therefore, when $F\tau$ must be expanded in two steps, this will proceed in two steps: it first expands to the recursive terms $\mu(\beta) \hat{F}(\tau, \beta)$, which in turn can be expanded to $\mu(\beta) H(\tau, \beta)$, as expected. Notice in particular that this enforces sharing between this expansion of $F\tau$ of its recursive instances withing τ . Compare this with the naive expansion if we had not used the auxiliary form:

$$F\tau \rightarrow H(\alpha, F\tau) \rightarrow H(\alpha, H(\alpha, F\tau)) \rightarrow \dots$$

The intermediate abbreviation still allows to do this, but this would require “explicit” unfolding of $\mu(\beta) \hat{F}(\alpha, \beta)$, as this is not needed during unification.

Notice that the abbreviation $F\alpha \triangleright \mu(\beta) \hat{F}(\alpha, \beta)$ is a non recursive type abbreviation, as F does not occur in $\mu(\beta) \hat{F}(\alpha, \beta)$, but an abbreviation to a recursive type, which is perfectly fine.

For a single recursive type abbreviation, we could have directly defined $F\alpha \rightarrow \mu(\beta) H(\alpha, \beta)$, without the helper type abbreviation \hat{F} , as there is not recursively defined abbreviation to delay the expansion of. Hence, the benefit of helper type abbreviations really appears with multiple recursively defined type abbreviations.

Consider $F_1\alpha \rightarrow H_1(\alpha, F_2\alpha)$ and $F_2\alpha \rightarrow H_2(\alpha, F_1\alpha)$. The encoding becomes

$$\begin{aligned} \hat{F}_1(\alpha, \beta_1, \beta_2) &\rightarrow H_1(\alpha, \beta_2) & F_1\alpha &\rightarrow \mu(\beta_1 = \hat{F}_1(\alpha, \beta_1, \beta_2), \beta_2 = \hat{F}_2(\alpha, \beta_1, \beta_2)) \beta_1 \\ \hat{F}_2(\alpha, \beta_1, \beta_2) &\rightarrow H_2(\alpha, \beta_1) & F_2\alpha &\rightarrow \mu(\beta_1 = \hat{F}_1(\alpha, \beta_1, \beta_2), \beta_2 = \hat{F}_2(\alpha, \beta_1, \beta_2)) \beta_2 \end{aligned}$$

where

$$\begin{aligned} \mu(\beta_1 = \hat{F}_1(\alpha, \beta_1, \beta_2), \beta_2 = \hat{F}_2(\alpha, \beta_1, \beta_2)) \beta_1 &= \mu(\beta_1) \hat{F}_1(\alpha, \beta_1, \mu(\beta_2) \hat{F}_2(\alpha, \beta_1, \beta_2)) \\ \mu(\beta_1 = \hat{F}_1(\alpha, \beta_1, \beta_2), \beta_2 = \hat{F}_2(\alpha, \beta_1, \beta_2)) \beta_2 &= \mu(\beta_2) \hat{F}_2(\alpha, \mu(\beta_1) \hat{F}_1(\alpha, \beta_1, \beta_2), \beta_2) \end{aligned}$$

For example, head-expanding $F_1(\tau)$ leads to $\mu(\beta_1) \hat{F}_1(\alpha, \beta_1, \mu(\beta_2) \hat{F}_2(\alpha, \beta_1, \beta_2))$ which again head-expands into $\mu(\beta_1) H_1(\alpha, \beta_1, \mu(\beta_2) \hat{F}_2(\alpha, \beta_1, \beta_2))$ while leaving $\mu(\beta_2) \hat{F}_2(\alpha, \beta_1, \beta_2)$ unexpanded.

Relaxed, regular recursive definitions We requested all occurrences of F_j to be applied to the same tuple of variables, which ensure that recursions are regular. This condition could be relaxed, by allowing some of the parameters to be instantiated by closed types.

For instance $F\alpha = G_2(\alpha, F\alpha, FG_0)$ is regular, as it can be encoded as:

$$\hat{F}(\alpha, \beta) = G_2(\alpha, F\alpha, \beta) \quad F\alpha = \hat{F}(\alpha, \mu(\beta) \hat{F}(G_0, \beta, \beta))$$

By contrast $F\alpha = G_2(\alpha, F(G_1\alpha))$ is not regular.

Since those recursive schemas are rare and can still be encoded (manually), while the general case is involved and may be computationally costly, we do not include them in the formalization.

4.6 Deconstructive type abbreviations

Type abbreviation definitions may also deconstruct their argument, such as in $F(G(\alpha)) \triangleright \alpha$, or more generally, type abbreviations of the form $F(\bar{\sigma}) = \sigma_F$ where to be well-formed free type variable of σ_0 should also appear in a decomposable occurrence of $\bar{\sigma}$.

However, this generalization will be better explained once we understood the simpler cases of constructive type abbreviations, which can be categorized into degenerate or productive type abbreviations, while by contrast deconstructive type abbreviations mixes both concepts.

Hence, we will only consider them later in §7.

5 Multi-equations

5.1 Constraints and multi-equations

The grammar of constraints is

$$C ::= \text{true} \mid \text{false} \mid m \mid \exists\alpha \cdot C \mid C \wedge C$$

The existential quantification takes priority over conjunction of multiequations, i.e., $\exists\alpha \cdot C \wedge C$ stands for $\exists\alpha \cdot (C \wedge C)$. Conjunction of multiequations is commutative and associative.

Letter m stands for multiequations m . Multiequations are non multisets of terms. We write \doteq for the union of multiequations. That is, we write $\tau_1 \doteq \tau_2 \doteq \dots \tau_n$ for multiequations.

The semantics of a constraint C is given by the set $\llbracket C \rrbracket$ of its solutions, i.e., the ground substitutions, mapping variables to ground types, satisfy all the constraints simultaneously, or equivalently by the judgment $\phi \vdash C$.

$$\phi \vdash \text{true} \quad \frac{\phi \vdash C_1 \quad \phi \vdash C_2}{\phi \vdash C_1 \wedge C_2} \quad \frac{\phi, \alpha \mapsto \tau \vdash C_0}{\phi \vdash \exists\alpha \cdot C} \quad \frac{\forall\sigma_1, \sigma_2 \in m, \phi\sigma_1 =_E \phi\sigma_2}{\phi \vdash m}$$

where type equality $=_E$ is that of the equational theory.

A constraint C_1 implies a constraint C_2 and we write $C_1 \vdash C_2$ if any solution of C_1 is also a solution of C_2 , i.e., $\llbracket C_1 \rrbracket \subseteq \llbracket C_2 \rrbracket$. Two constraints C_1 and C_2 are equivalent and we write $C_1 \equiv C_2$ they imply one another. We also write $C_1 \Rightarrow C_2$ when we wish to orientate the equivalence for rewriting purposes.

5.2 Canonical forms for constraints

We may freely reorganize constraint by treating the conjunction of constraints \wedge as associative commutative with **true** as neutral element and **false** as absorbing element, which indeed preserves the semantics. Besides, we implicitly use the extrusion of existential bindings:

$$\text{EXAND} \quad \frac{\alpha \# C}{(\exists\alpha \cdot C_0) \wedge C \equiv \exists\alpha \cdot (C_0 \wedge C)}$$

The side condition $\alpha \# C_2$ means that $\alpha \notin \text{ftv}(C)$ where $\text{ftv}(C)$ is defined in the obvious way. Equivalences \equiv (or rewriting rules \Rightarrow) can be applied in any context:

$$\frac{\text{CONTEXT-EX} \quad C_1 \Rightarrow C_2}{\exists \alpha \cdot C_1 \Rightarrow \exists \alpha \cdot C_2} \qquad \frac{\text{CONTEXT-AND} \quad C_1 \Rightarrow C_2}{C_1 \wedge C \Rightarrow C_2 \wedge C}$$

This allows to put constraints in canonical form $\exists \bar{\alpha} \cdot \bigwedge^{i \in I} m_i$, i.e., a conjunction of multiequations preceded by existential quantification, or one of the trivial constraint **true** or **false**.

In the following, we implicitly maintain such canonical forms after every rewriting step. In particular, existential quantifiers are automatically extruded and rewriting may be aborted to return **false** as soon as a subconstraint is equivalent to **false**.

5.3 Multiequations and views

A multiequation is a multiset of terms. A substitution is a solution of a multiequation if it equates all terms of the multiequation, as indicated above in §5.1.

We introduce additional structure to multiequations by grouping terms into views. Hence, views are themselves multisets of terms. We write ν for views. We distinguish two kinds of views: a productive view $\langle \nu \rangle$ contains productive terms that whose head-constructors form a connected set; a degenerate view $[\nu]$ contains terms that all head expands to some variable. We write $\tau \doteq \nu$ for adding term τ to the view ν .

A multiequation becomes a multiset of both isolated terms (as before) and views. Views, i.e., the partitioning of the multiequation into views, can always be dropped, leaving the multiequation as just a plain multiset of terms. In particular, a solution of a multiequation with views is just a solution of the multiequation ignoring views.

We still write \doteq for union both within view and within multiequations.

5.4 Unification without deconstructive abbreviations

We describe how to solve multisets of multiequations. We will initially establish and maintain the following invariants:

- A multiequation has a unique (possibly empty) degenerate view $[\nu]$ such that $[\nu] \doteq m \equiv m$. That is, the variable head expansions of terms in productive are already in m .
- A multiequation may have zero or several productive views. We maintain a *global* invariant all terms in a productive view are equivalent, i.e., for a whole system of constraint:

$$\langle \tau \doteq \tau_0 \doteq \nu \rangle \doteq m \wedge C \equiv \langle \tau \doteq \nu \rangle \doteq m \wedge C$$

This equivalence is not however used to simplify constraints (from left to right), as τ_0 is kept in the view both for efficiency reasons and for better presentation of the solutions of the multiequation.

Notice that variable terms are never put in views. Non variable terms are introduced as such outside of a view, awaiting for being dispatched in either a productive or degenerate view.

5.5 Simplification rules

$$\begin{array}{c} \text{STUTTER} \\ \alpha \doteq \alpha \doteq \nu \implies \alpha \doteq \nu \end{array}$$

$$\begin{array}{c} \text{MERGE} \\ \alpha \doteq [\nu_1] \doteq m_1 \wedge \alpha \doteq [\nu_2] \doteq m_2 \implies \alpha \doteq [\nu_1 \doteq \nu_2] \doteq m_1 \doteq m_2 \end{array}$$

$$\begin{array}{c} \text{DISPATCH-D} \\ \forall F = i \\ \hline \text{F}\bar{\alpha} \doteq [\nu] \doteq m \implies [\text{F}\bar{\alpha} \doteq \nu] \doteq \alpha_i \doteq m \end{array} \qquad \begin{array}{c} \text{DISPATCH-P} \\ \forall F = G \\ \hline \text{F}\bar{\tau} \doteq m \implies \langle \text{F}\bar{\tau} \rangle \doteq m \end{array}$$

$$\begin{array}{c} \text{TRIM} \\ \tau \notin \mathcal{V} \quad \alpha \# \tau_1, \tau_2, m \\ \hline \text{F}(\bar{\tau}_1, \tau, \bar{\tau}_2) \doteq m \implies \exists \alpha \cdot \text{F}(\bar{\tau}_1, \alpha, \bar{\tau}_2) \doteq m \wedge \alpha \doteq \tau \end{array}$$

$$\begin{array}{c} \text{TRIM-P} \\ \tau \notin \mathcal{V} \quad \alpha \# \tau_1, \tau_2, \nu, m \\ \hline \langle \text{F}(\bar{\tau}_1, \tau, \bar{\tau}_2) \doteq \nu \rangle \doteq m \implies \exists \alpha \cdot \langle \text{F}(\bar{\tau}_1, \alpha, \bar{\tau}_2) \doteq \nu \rangle \doteq m \wedge \alpha \doteq \tau \end{array}$$

$$\begin{array}{c} \text{DECOMPOSE} \\ \langle \text{F}\bar{\alpha} \doteq \nu \rangle \doteq \langle \text{F}\bar{\beta} \doteq \nu' \rangle \doteq m \implies \langle \text{F}\bar{\alpha} \doteq \text{F}\bar{\beta} \doteq \nu \doteq \nu' \rangle \doteq m \wedge (\alpha_i \doteq \beta_i)^{i \in \downarrow F} \end{array}$$

$$\begin{array}{c} \text{CLEAN} \\ \{ \text{F}\bar{\alpha} \doteq \text{F}\bar{\beta} \doteq \nu \} \doteq m \implies \{ \text{F}\bar{\alpha} \doteq \nu \} \doteq m \end{array}$$

$$\begin{array}{c} \text{CLASH} \\ \text{F} \bowtie \text{G} \\ \hline \langle \text{F}\bar{\alpha} \doteq \nu \rangle \doteq \langle \text{G}\bar{\beta} \doteq \nu' \rangle \doteq m \implies \text{false} \end{array}$$

$$\begin{array}{c} \text{EXPAND} \\ \text{F}\bar{\beta} \triangleright \text{G}\bar{\tau} \\ \hline \langle \text{F}\bar{\alpha} \doteq \nu \rangle \doteq m \implies \exists \bar{\beta} \cdot \langle \text{G}\bar{\tau} \doteq \text{F}\bar{\alpha} \doteq \nu \rangle \doteq m \wedge_{i \in \downarrow F} \alpha_i \doteq \beta_i \end{array}$$

Rule STUTTER removes duplicated variables. Rule MERGE merges two multiequations that share a variable term.

Rules DISPATCH-P and DISPATCH-D place non-variable terms in views, according to the degenerate *vs.* productive status of the head-symbol: for degenerate terms, we also add α_i in m , which amounts to fully expand F : this is needed to enforce further merging that could otherwise remain unnoticed; this is also

efficient as this kind of expansion does not allocate. By contrast, for productive abbreviations (or primitive type constructors) we just move them unexpanded in a productive view, letting expansion to be performed on demand.

Rule `TRIM` allows to turn large terms into small ones, i.e., of high at most one. This is done position by position, but could be performed more aggressively. `TRIM-P` is the same as `TRIM` but operates within a productive view. This is needed, as an expansion allocates a new (large) term directly inside a productive view.

Rule `DECOMPOSE` applies when two productive views have a common head constructor F : these views are merged and corresponding subterms are equated along decomposable positions $\downarrow F$. To avoid duplication of terms, decomposed terms should be small so that all subterms are variables.¹

Rule `CLEAN` eliminates duplicates within a view. Braces mean that it applies to both degenerate and productive views.

Rule `CLASH` detects obvious clashes based on \bowtie . A better (larger) relation \bowtie will detect clashes earlier, avoiding further expansions. However, the least relation \bowtie_0 is sufficient to detect all clashes, since in the absence of clashes expansion will be performed, as long as views do not share a common head-symbol, eventually ending with primitive symbols, on which \bowtie is complete.

Rule `EXPAND` applies when a term $\langle F\bar{\alpha} \rangle$ appears in a productive view of a multiequation and F is a type abbreviation $F\beta \triangleright G\tau$. We then add (a fresh copy of) $G\tau$ to the view containing $F\bar{\alpha}$ and add multiequations $\alpha_i \doteq \beta_i$ for all decomposable positions i in F .

5.6 Invariants

We formalize the invariants that are maintained by the unification rules. These subsume invariants informally described at the beginning of §5.4.

The invariants applies to a constraint $\exists \alpha \cdot C$ in canonical form, i.e., C is a conjunction of multiequations.

Invariant 1 *If C contains $[\tau \doteq \nu] = m$, then τ is of the form $F\bar{\sigma}$ and $\forall F = i$.*

Invariant 2 *If C is $[\nu] \doteq m \wedge C_0$, then $C \equiv m \wedge C_0$,*

Rule `DISPATCH-D`, which is the only one to add a term τ of the form $C\bar{\sigma}$ to the degenerate view, preserve both invariants: it only applies when $\forall C = i$ and it inserts the expansion α_i in m . Rule `MERGE`, which merges two degenerate views ν_1 and ν_2 preserves the invariant. Rule `CLEAN` relies on the invariant 2 to preserve the equivalence.

Invariant 3 *If C contains $\langle \nu \rangle \doteq m$, then m is nonempty and for all τ in m , τ is of the form $F\bar{\sigma}$ and $\forall \tau = G$.*

¹Notice, that only productive constructors are decomposed. Degenerate type constructors are decomposable along their non phantom position, but this will never be used. One could think that we miss an opportunity for efficient simplification, but this is not the case because degenerate type constructors must be eagerly expanded anyway and their expansion—a variable, is a quite efficient to compute and represent.

Invariant 4 *If C is $\langle \nu \doteq m \rangle \wedge C_0$ and ν' is a non-empty subset of ν then $C \equiv \langle \nu' \doteq m \rangle \wedge C_0$.*

Rule DISPATCH-P is the only one that creates a new productive (singleton) view $\langle \tau \rangle$ when τ is of the form $F\bar{\sigma}$ and $\forall F$ is a constructor G , so F is productive and the new view $\langle \tau \rangle$ satisfies the invariant 4. That view also satisfies 3 since it has no strict non-empty subview. Rule EXPAND increases a productive view. It introduces a new term $G\bar{\tau}$ in the productive view, which is productive, since it is the expansion of a productive term $F\bar{\beta}$. It therefore preserves the invariant 4; it also preserves the invariant 3. Rule DECOMPOSE merges two productive views of the same multiequation and preserves 3.

Invariant 5 *The set of toplevel symbols of any productive view ν in C forms a continuous subtree of the forest \triangleright . The root of this subtree, the type constructor of smallest rank.*

This invariant obviously holds for a singleton view $\langle \tau \rangle$ as introduced by rule DISPATCH-P. It is maintained by DECOMPOSE, since we join two continuous subtrees that share a node. Rule EXPAND extends the tree (towards the trunk) with an immediate successor, which preserves continuity.

Rules CLEAN and CLASH obviously preserve all five invariants.

Lemma 1 *The unification rules preserve the semantics of constraints.*

Proof: Rule MERGE and TRIM, DECOMPOSE, CLASS are standard.

Rule DISPATCH leaves the raw multiequation unchanged. Rule CLEAN relies on the invariant 2.

Finally, EXPAND is a combination of the addition of the immediate expansion of a term along \triangleright , which preserves equality followed by TRIM and DECOMPOSE, but it can also be seen directly:

- Assume that ϕ is a solution of $\langle F\bar{\alpha} \doteq \nu \rangle \doteq m$ and $F\beta \triangleright G\bar{\tau}$.
- Let ϕ' be $\phi[\beta_i \leftarrow \alpha_i]$ (notice that β_i 's are all distinct). Then $\phi'(F\bar{\beta}) = \phi'(G\bar{\tau})$, we also have $\phi'(F\bar{\alpha}) = \phi(F\bar{\beta})$ and therefore $\phi'(F\bar{\alpha}) = \phi'(G\bar{\tau})$.
- Therefore ϕ' is a solution of $\langle G\bar{\tau} \doteq F\bar{\alpha} \doteq \nu \rangle \doteq m \wedge_{i \in \downarrow F} \alpha_i \doteq \beta_i$, that is, ϕ is solution of $\exists \bar{\beta} \cdot \langle G\bar{\tau} \doteq F\bar{\alpha} \doteq \nu \rangle \doteq m \wedge_{i \in \downarrow F} \alpha_i \doteq \beta_i$.

The converse is obvious.

6 Unification strategy

The unification rules are sound and can be applied in any order. We describe a strategy to restrict the application of the rules that is both efficient and sufficient to obtain normal forms. We do so in two steps: we first defined a collection

$$\begin{aligned}
(\text{DISPATCH-ALL}) &\triangleq (\text{DISPATCH-P}) \cup (\text{DISPATCH-N}) \\
(\text{TRIM-DISPATCH}) &\triangleq (\text{TRIM-P})^\infty ; (\text{TRIM})^\infty ; (\text{DISPATCH-ALL})^\infty \\
(\text{DECOMPOSE-CLEAN}) &\triangleq (\text{CLASH}) \cup ((\text{DECOMPOSE}) ; (\text{CLEAN})^\infty) \\
(\text{MERGE-DECOMPOSE}) &\triangleq (\text{MERGE}) ; (\text{STUTTER})^\infty ; (\text{CLEAN})^\infty ; \\
&\quad (\text{DECOMPOSE-CLEAN})^\infty \\
(\text{EXPAND-DECOMPOSE}) &\triangleq (\text{EXPAND})^\dagger ; (\text{TRIM-DISPATCH})^\infty ; \\
&\quad (\text{DECOMPOSE-CLEAN})^\infty
\end{aligned}$$

Figure 1: Derived rules

of derived rules (Figure 1) that will be used instead of basic transformation rules. Derived rules ensure further invariants, following a transformation rule by further rules that restore some invariants broken by the first application. We then define an algorithm that applies derived rule in a deterministic order.

Derived rules are defined in Figure 1.

Rule DISPATCH-ALL is a helper rule to factor all forms of dispatch. Rule TRIM-DISPATCH implicitly uses EXAND to extrude existential quantifiers that TRIM introduced.

Rule MERGE-DECOMPOSE applies to a conjunction of multiequations, but each of the other rules applies to a single multiequation.

Rule EXPAND-DECOMPOSE uses Rule $(\text{EXPAND})^\dagger$ is a restriction of Rule EXPAND that requires:

- the expanded term to be the representative of the view of highest rank
- the rest of the multiequation m to contain at least another view ν'

6.1 Invariants of composed rules

Invariant 6 (Strategy) *The strategic rules maintain the additional invariants:*

- *All terms are small.*
- *All constructed terms are in a productive or degenerate view.*
- *The top constructors of two productive views are disjoint.*
- *All terms in a productive view have different toplevel symbols.*

The term of a productive view with the toplevel symbol of smallest rank is called the representative.

These invariants are preserved by each of the composed rules—often in an obvious way. These invariants can be ensured on the input constraint as follows:

- Input multiequations have only single terms views (one could also drop views that do not satisfy the invariants).
- Apply (TRIM-DISPATCH) to make all terms small and place them in the appropriate views.

6.2 Algorithm

The algorithm SOLVE is defined as the following composition of rules:

$$(\text{SOLVE}) \triangleq ((\text{MERGE-DECOMPOSE})^\infty ; (\text{EXPAND-DECOMPOSE}) ; (\text{MERGE-DECOMPOSE})^\infty)^\infty$$

In words, the algorithm makes the system merged and decomposed as far as possible; when no more merge, decomposition, nor clash can be performed, it does a single expansion; it iterates until no more (restricted) expansion is possible, i.e., all multiequations have at most one productive view. It is important to perform a single expansion because these may allocate new terms and may be expansive. It is also important to select the terms of lowest rank in the view of highest rank for expansion, as this is the one most likely to allow decomposition after merging. Otherwise, expansion could uselessly expand all type abbreviations away, which kills the purpose of this algorithm.

The algorithm SOLVE always terminates, with solved forms.

6.3 Termination

We show the termination of (SOLVE).

The weight $|F|$ of a type constructor F of rank n is the polynomial X^n . The weight $|\tau|$ of a term τ is defined as

$$|\alpha| = 0 \qquad |F\bar{\tau}| = |F| + 2 * (\sum_{i \in I} |\tau_i|)$$

In particular, the weight of a variable is null and the weight of a term is higher than the weight of its subtrees plus the weight of the toplevel type constructor.

The weight of a productive view is the weight of its representative, i.e., the term of the lowest rank.

The weight of a multiequation is the sum of the weights of its productive views plus twice the size of terms that are not in a view.

The weight of the constraint is the sum of the weight of its multiequations. The size of the constraint is $sY + mZ$ where s is its weight and m is the number of multiequations.

Rewriting rules reduce the size of C . More precisely, assuming that the initial constraint satisfies the invariant 6, i.e., in particular, where all terms are small:

- Rules STUTTER and CLEAN are auxiliary rules that cannot be applied forever, and do not increase the size of C , but may not decrease it either.
- Rule DISPATCH-ALL does not increase the weight of the multiequation as it moves a term whose size count into a degenerate view where its size does not count or a productive view where it counts less. (Currently, it does not increase the number of multiequations, but it could.) Hence, this rule can only be used after another rule that strictly decreases the weight of C .
- Rule TRIM-P and TRIM both reduce the weight of the multiequation they apply to, but also introduces new multiequations, which counts less.
- Rule TRIM-DISPATCH does not increase the size of the constraint. (It does not necessarily decreases it, as all terms may already be small.) The rule alone always terminates, as it does not create new type constructors while it decreases the weight of terms and the number of terms not in a view. Hence, this rule may be safely used after a rule that strictly decreases the weight of multiequations.
- Rule DECOMPOSE-CLEAN decreases the size of the constraint. Rule CLASH does, indeed! Rule DECOMPOSE followed by CLEAN decreases the weight of the system. Indeed, it replaces two views, with representatives $F_1\bar{\alpha}_1$ and $F_2\bar{\alpha}_2$, by a single view, with representative either $F_1\bar{\alpha}_1$ or $F_2\bar{\alpha}_2$. This is regardless of whether F is one of F_1 or F_2 . Simultaneously, the rule may introduce new multiequations composed of just variables, i.e., of null weight, hence the whole weight of the constraint decreases and so does its size.
- Rule Merge-Decompose decreases the size of the constraint: mainly, Rule MERGE reduces the number of multiequations without increasing anything else. Then, after cleaning steps, which leaves the size unchanged, it may decrease the size of the constraint (Rule DECOMPOSE-CLEAN) or leave it unchanged if this rule does not apply.
- Rule EXPAND-DECOMPOSE decreases the size of the constraint. Mainly, Rule Expand decreases the weight of the system by replacing in a view a representative of by another representative of lower rank, hence of smaller size, while increasing the number of multiequations. Notice that the original term is left in the view whose length increases, but it it will never be used as the representative anymore. Rule $(\text{TRIM-DISPATCH})^\infty$, does not increase the size of the constraint. Rule DECOMPOSE-CLEAN if applicable decreases the size of the constraint.
- Rule SOLVE uses three rules, MERGE-DECOMPOSE, EXPAND-DECOMPOSE, and MERGE-DECOMPOSE, each of which decreases the size of the constraint.

6.4 Efficiency of the strategy

By construction, the algorithm never performs an expansion (of a constructive type abbreviation) when a decomposition is possible.

Decompositions are in general preferable, because they do not allocate new terms, but on the opposite merges existing terms. Furthermore, when a multiequation is expanded, we also choose the view with highest rank hoping to be able to merge it (and decompose it) with a view of a lower rank. This choice ensures that the rank of the multiequation in its canonical form will be maximal, i.e., as high as the rank in any other canonical form.

However, this does not ensure a minimal number of expansions. The choice of the multiequation to expand is arbitrary among those that are candidate for expansion.

Expanding a view ν_1 in a multiequation m_1 may perhaps be avoided by another expansion that will induce merging of m_1 with another multiequation m_2 that will contain a view ν_2 that will be decomposable with ν_1 . For example, consider

$$\alpha_1 \doteq \langle \nu_1 \rangle \doteq \langle \nu_2 \doteq \tau \rangle \wedge \alpha_2 \doteq \langle \nu_2 \rangle \doteq \nu_1 \doteq \tau \wedge C$$

when the multiequation α_1 cannot be decomposed, hence is a candidate for expansion. Instead, an expansion in C might lead to a decomposition which in turn will require merging α_1 and α_2 , leading to a constraint of the form

$$\alpha_1 \doteq \alpha_2 \doteq \langle \nu_1 \doteq \tau \rangle \doteq \langle \nu_1' \rangle \langle \nu_1 \rangle \doteq \nu_1' \doteq \tau \wedge C'$$

Now, we first decompose the two views containing τ leading to

$$\alpha_1 \doteq \alpha_2 \doteq \langle \nu_1 \doteq \tau \doteq \nu_2 \rangle \doteq \langle \nu_1 \rangle \doteq \nu_2 \wedge C''$$

which in turn allows to decompose the two remaining views ending up with a multiequation with a single view—without doing any further expansion.

In summary, our strategy delays expansions until at least one is necessary, and chooses an expansion that is promising, but does not ensure optimality in the number of expansions performed.

6.5 Solved forms

The algorithm terminates with constraints in solved form, i.e., constraints C of the form $\exists \beta \cdot \wedge_{i \in I} \nu_i$ where

- Any variable occurring in C (free or bound) belongs to a unique multiequation.
- Every multiequation ν_i has at most one productive view, hence is of the form $\bar{\alpha}_i \doteq \langle \nu_i \rangle \doteq [\nu_i']$ where ν_i and ν_i' may be empty and otherwise satisfy the invariants 6 (in addition to invariants 1–4 of §5.6).

Indeed,

- This follows from the fact that MERGE does not apply.
- Since EXPAND does not apply, two views of the same multiequations cannot be such that one contains a type abbreviation.

Since DECOMPOSE does not apply, two views of the same multiequations cannot have the same head symbols.

Since CLASH does not apply, then there cannot be two such views in a multi-equations, i.e., there is be at most one productive view per multi-equation.

7 Deconstructive type abbreviations

We now extend the algorithm to allow deconstructive type abbreviations. We first consider a simple case in §7.1, relax it in §7.2, before we tackle the general case in §??

7.1 A purely deconstructive type abbreviation

The simplest for of deconstructive type abbreviation is $F(G\bar{\beta}) = \alpha_i$.

Notice that G may itself be a type abbreviation. Hence, we must assume that G is at least decomposable in position i , otherwise α_i is undetermined and the type definition does not make sense.

Such an abbreviation is non-productive: since $F(G\bar{\beta})$ expands to the type variable β_i , it should behave as β_i ; in particular $F(G\bar{\beta}) \doteq m$ should require the merge of this multiequation with with the one β_i belongs to. A simple way to ensure this is to expand $F(G\bar{\beta})$ eagerly: when a term $F\alpha$ appears in a multiequation m , we add β_i to m with an additional constraint $\alpha \doteq G\bar{\beta}$. We perform this when we dispatch $F\alpha$ into the degenerate view:

$$\text{DISPATCH-EXP-D} \quad \frac{F(G\bar{\beta}) \triangleright \beta_i}{F\alpha \doteq [\nu] \doteq m \implies \exists \bar{\beta} \cdot [F\alpha \doteq \nu] \doteq \beta_i \doteq m \wedge \alpha = G\bar{\beta}}$$

Rule DISPATCH-EXP-D is a mix of DISPATCH-N and EXPAND in the sense that, when we meet $F\alpha$, we need to allocate some structure $\alpha \doteq G\bar{\beta}$ to force α to have the requested shape and extract α_i from it.

Lemma 2 *Rule DISPATCH-C is an equivalence.*

Proof: Assume that ϕ is a solution of $F\alpha \doteq [\nu] \doteq m$. The definition of the type abbreviation of $F(G\bar{\beta}) \triangleright \beta_i$ implies that $\phi\alpha$ is of the form $F(G\bar{\tau})$ for some $\bar{\tau}$. Let ϕ' be $\phi[\beta_i \leftarrow \tau_i]$. By definition $F(G\bar{\tau})$ is equal to τ_i . Therefore, ϕ' is a solution of $[F\alpha \doteq \nu] \doteq \alpha_i \doteq m \wedge \alpha = G\bar{\beta}$. Hence, ϕ is a solution of $\exists \bar{\beta} \cdot [F\alpha \doteq \nu] \doteq \beta_i \doteq m \wedge \alpha = G\bar{\beta}$. The inverse direction is obvious.

Rule DISPATCH-EXP-D preserves invariant 2. To preserve invariant 1, we need to define \forall for type abbreviations such as F , replacing i by a sequence of projections to access the variable β_i from $F(\mathbf{G}\bar{\beta})$ in the definition of the type abbreviation². In our example, $\forall F$ is $1 \cdot i$. In fact, we may also allow two parameters β_i and β_j to be actually the same variable. Thus, $\forall F$ is rather a set of paths than a single path. Hence, we should write $1 \cdot i \in \forall F$ or $\pi \in \forall F$ when π stands for a path.

Adding rule DISPATCH-EXP-D to the definition of DISPATCH-ALL, we preserve properties of the unification algorithm. In particular, rule DISPATCH-ALL does not increase the weight of the constraint, although it may now increase the number of multiequations.

Therefore, the (SOLVE) still terminates, with the same solved forms (but more general degenerate views).

Solved forms Notice that a constraint such as $\alpha \doteq F\beta \wedge \beta \doteq F_0$ looks in solved form—but it is not! Indeed, $F\beta$ carries the implicit constraint that β must be of the form $\mathbf{G}\alpha_0$ which implies that the constraint $\mathbf{G}\alpha_0 \doteq F_0$ should also hold—but it is unsolvable (equivalent to false).

Fortunately, solved forms satisfy the invariant 6 which requires non variable terms to be either in degenerate views or productive views. Here, F must therefore have been moved to the degenerate view, which enforces its expansion and will produce a clash between F_0 and \mathbf{G} .

Decomposable paths If \mathbf{G} is unary, then F is decomposable and $\downarrow F = \{1\}$. Otherwise, F is not decomposable—unless all arguments of \mathbf{G} but one are phantom. For instance, assume F is unary and \mathbf{G} is binary and decomposable in both directions. Then $F(\mathbf{G}(\tau_0, \tau_1)) = F(\mathbf{G}(\tau_0, \tau_2))$ but $\mathbf{G}(\tau_0, \tau_1) \neq \mathbf{G}(\tau_0, \tau_2)$.

Thus, there are now positions that are neither decomposable nor phantom. This is not a problem however, since degenerate type systems will never appear in a productive view and need not be decomposed.

It may seem annoying that a position may be neither phantom nor decomposable. However, we should not consider positions but paths, which are a generalization of positions.

We redefine \downarrow to return a list of paths (starting from the root) instead of a list of positions. That is, it is not immediate parameters that are decomposable or phantom but paths rooted at F . Continuing with our example, we would have $\downarrow F$ is $\{1 \cdot 1\}$, i.e., the path $1 \cdot 1$ is decomposable, while the occurrence $1 \cdot 2$ is phantom.

Generalization We can easily generalize the previous example with type abbreviations of the form $F(\mathbf{G}_i \bar{\alpha}_i)^{i \in I} \triangleright \alpha_{ij}$.

²In the general case, $\mathbf{G}\bar{\beta}$ could even be replaced by a term σ where β_i also occurs deeper inside σ .

7.2 Mixed type abbreviations

Deconstructive type abbreviation pattern match on their argument to extract some subterm. They need not return just one subterm, but instead may also construct a new term by combining several subterms.

Consider for example $F(G(\alpha, \beta)) = G(\beta, \alpha)$ where G is a primitive type constructor (or more generally a productive one, decomposable on both directions). This abbreviation is both deconstructive and constructive. How should we treat it? Here are a few hints:

- F is decomposable: if $F\tau = F\tau'$ then τ and τ' must be of the form $G(\tau_1, \tau_2)$ and $G(\tau'_1, \tau'_2)$ and $F\tau = G(\tau_2, \tau_1)$ and $F\tau' = G(\tau'_2, \tau'_1)$.
- Hence, it looks like we could treat F as productive.
- Still, we must be aware that $\alpha \doteq F\beta \wedge \beta \doteq F\alpha$, which looks in canonical form is actually not. Indeed, there is an implicit constraint that β must be of the form $G(\tau_1, \tau_2)$ and $\alpha \doteq \tau_1$
- Hence, as with the previous example, we must eagerly expand F —but still put it in a productive views.

This leads to the following dispatching rule:

$$\frac{\text{DISPATCH-EXP-P} \quad F(G\bar{\beta}) \triangleright \sigma}{F\alpha \doteq m \implies \exists \bar{\beta} \cdot \langle \sigma \doteq F\alpha \rangle \doteq m \wedge \alpha \doteq G(\bar{\beta})}$$

This rule decreases the size of the constraint (it decreases its weight while increasing the number of multiequations, which counts less).

There (`SOLVE`) still terminates, with the same normal forms. By construction, we know that whenever $F\alpha$ appears in a productive view of a constraint C , then $C \vdash \exists \bar{\beta} \cdot \alpha \doteq G\bar{\beta}$ holds, since such constrains are injected into the system when F is initially placed into the productive view.

7.3 The general case

Let us now tackle the general case of an abbreviation $F(\bar{\tau}) \triangleright \sigma_F$.

We need a few restrictions for this definition to be well formed:

- F should be new, i.e., not occur in $\bar{\tau}$ nor in σ_F .
- All variables should be bound, i.e., $\text{ftv}(\sigma_F) \subseteq \text{ftv}^\nabla(\bar{\tau})$
- $\nabla\sigma_F$ should exist, i.e., deconstructive type constructors that occur in σ_F should be applied to subterms of expected shapes.

There are several cases to consider:

- σ_F is a variable β ; then, F is degenerate, as in §7.1

$\frac{\text{DISPATCH-EXP-D}}{\frac{F(\bar{\tau}) \triangleright \sigma_F \quad i \cdot \pi \in \dot{\vee} F \quad \tau_i / \pi = \beta \quad \bar{\tau} \not\subseteq \mathcal{V}}{F\bar{\alpha} \doteq [\nu] \doteq m \implies \exists \text{ftv}(\bar{\tau}) \cdot [F\bar{\alpha} \doteq \nu] \doteq \beta \doteq m \wedge^{i \in I} \alpha_i \doteq \tau_i}}$
$\frac{\text{DISPATCH-EXP-P}}{\frac{F(\bar{\tau}) \triangleright \sigma_F \quad \dot{\vee} F = H \quad \bar{\tau} \not\subseteq \mathcal{V}}{F\bar{\alpha} \doteq m \implies \exists \text{ftv}(\bar{\tau}) \cdot \langle \sigma_F \doteq F\bar{\alpha} \rangle \doteq m \wedge^{i \in I} \alpha_i \doteq \tau_i}}$

Figure 2: Derived rules for deconstructive abbreviations

- σ_F is of the form $G\bar{\tau}$ and $G\bar{\tau} \rightarrow \beta$. This is quite similar to the previous case.
- σ_F is of the form $G\bar{\tau}$ and $\dot{\vee} G = H$. Then F is productive.
- σ_F is of the form $G\bar{\tau}$ and G is itself a deconstructive type abbreviation. This is in fact, a particular case of previous ones: we just look at $\dot{\vee} G$ which is the same as $\dot{\vee} F$ and tell us whether the type abbreviation is degenerate or productive.

In all cases, we compute $\dot{\vee} F$, directly by computing $\nabla \sigma_F$ directly, or indirectly by looking at σ_F . Then

- if $\nabla \sigma_F$ is a variable, then $\dot{\vee} F$ is a path π (of the form $i \cdot \pi$) and F is degenerate;
- otherwise, $\dot{\vee} F$ is a (primitive) type constructor H and F is productive

This leads to two new dispatching rules DISPATCH-EXP-D and DISPATCH-EXP-P defined on Figure 2. These rules are complementary and should not replace the rules DISPATCH-P and DISPATCH-D. Even though the DISPATCH-EXP versions of the rules are semantically correct, the (grayed) side condition $\bar{\tau} \not\subseteq \mathcal{V}$ requires that at least one of the τ_i 's is not a variable, so that the original rules do not apply.

XXX [Fix $\downarrow F$ for rule Expand Dispatch-Exp-P]

Degenerate type abbreviations For degenerate terms, we can actually see DISPATCH-D as DISPATCH-EXP-D followed by some cleanup of useless variables. Indeed, when $\bar{\tau}$ are all distinct type variables $\bar{\beta}$, the difference is just the introduction of auxiliary existential variables $\bar{\beta}$ with multiequations $\wedge^{i \in I} \alpha_i \doteq \beta_i$, which can then be removed, using the equivalence

$$\exists \beta \cdot C \wedge \alpha \doteq \beta \Rightarrow C \quad (\text{when } \beta \# C)$$

Productive type abbreviations: eager vs. lazy expansion The main difference between deconstructive rules and non-deconstructive ones is for productive terms: Rule DISPATCH-EXP-P performs an eager expansion while DISPATCH-P

does not perform expansion, which will then be done on demand, lazily. Hence, there is a discontinuity between the two rules.

We can in fact reduce this difference, by automatically splitting deconstructive type annotations as follows:

- Let $\bar{\beta}$ be $\text{ftv}(\sigma_F)$;
- Define $\hat{F}\bar{\beta} \triangleright \sigma_F$ where $\bar{\beta}$;
- Redefine $F\bar{\tau} \triangleright \hat{F}\bar{\beta}$.

With this encoding, the dispatch rule for F becomes:

$$\frac{\text{DISPATCH-EXP-P}' \quad F(\bar{\tau}) \triangleright \hat{F}\bar{\beta} \triangleright \sigma_F \quad \forall F = \forall \hat{F} = H}{F\bar{\alpha} \doteq m \implies \exists \text{ftv}(\bar{\tau}) \cdot \langle \hat{F}\bar{\beta} \doteq F\bar{\alpha} \rangle \doteq m \wedge^{i \in I} \alpha_i \doteq \beta_i}$$

The construction of σ_F is now delayed until $\hat{F}\bar{\beta}$ need to be expanded, as for non destructive type abbreviations.

This could be automated or used manually to delay the constructive part of deconstructive type abbreviations.

It can also be used to regain the continuity between the DISPATCH-EXP-P and DISPATCH-P, so that the latter can be seen as a particular case of the former (up to some details) when $\bar{\tau}$ is a tuple of disjoint variables: F becomes a pure alias for \hat{F} which could (intuitively) be removed. This would lead to the rule:

$$\frac{\text{DISPATCH-P}' \quad F\bar{\beta} \triangleright \sigma_F \quad \forall F = H}{F\bar{\alpha} \doteq m \implies \exists \bar{\beta} \cdot \langle F\bar{\alpha} \rangle \doteq m \wedge^{i \in I} \alpha_i \doteq \beta_i}$$

where the conclusion is equivalent and could be simplified to $\langle F\bar{\alpha} \rangle \doteq m$, which then coincides with the conclusion of Rule DISPATCH-P.

Soundness Rules DISPATCH-EXP-D and DISPATCH-EXP-D preserve the semantics of constraints.

Termination Both rules decrease the size of the constraint. Notice that this is different from the dispatching of non-deconstructive type abbreviations, which did not decrease the size of constraint. The deconstructive version allocate new terms, but this is coming and compensated by the expansion of the type abbreviation constructor, which as for regular expansions counts much more than the terms it allocates. Therefore, after including these rules in DISPATCH-ALL, Rule TRIM-DISPATCH will still terminate when used alone and will never increase the size of the constraint.

Therefore Rule SOLVE still terminates with the same solved forms.

7.4 Further Encodings

We have seen that the constructive part of deconstructive type abbreviations can be recovered by an helper type abbreviation.

Similarly, we could use helper type abbreviations for the deconstructive parts. Consider $F(\bar{\tau}) \triangleright \sigma_F$. Let $\bar{\alpha}$ be $\text{ftv}(V\tau)$ and $(F_i(\bar{\alpha}) \triangleright \tau)^{i \in I}$ a sequence of type abbreviations. We could then redefine $F(F_i\bar{\tau}) \triangleright \sigma_F$ with just one level of deconstruction.

This allows to limit the deconstruction of the arguments to their topmost constructor during the dispatch and treat the rest of decomposition lazily—or avoid it completely if the F_i 's are never mixed with some other type constructor.

Thanks to this encoding, we could also have limited the general case to type abbreviations of the form $F(G_i\bar{\alpha}_i)^{i \in I} \triangleright \sigma_F$ and solve others by encoding.

Alternative encodings According to the above schema, the deconstructive type abbreviation $F(G(G(\alpha_0, \alpha_2), \alpha_0)) \triangleright \alpha_0$ would be encoded as

$$\hat{F}(\alpha_0, \alpha_1, \alpha_2) \triangleright G(G(\alpha_0, \alpha_1), \alpha_2) \quad F(\hat{F}(\alpha_0, \alpha_1, \alpha_2)) \triangleright \alpha_0$$

where the deconstruction is done minimally delaying the allocation of the intermediate structure to \hat{F} .

Another possible encoding would be:

$$\hat{F}(G(\alpha_0, \alpha_1)) \triangleright \alpha_0 \quad F(G(\beta, \alpha_2)) \triangleright \hat{F}\beta$$

to deconstruct the argument step by step.

The difference is interesting:

- The former uses an auxiliary constructive type abbreviation \hat{F} as an argument to the deconstructive one F to destruct “more at once”. However, the expansion of \hat{F} will be performed lazily, only when $\hat{F}(\alpha_0, \alpha_1, \alpha_2)$ will be mixed with some other non variable type.

So this may behave better when the expansion of \hat{F} is a large term.

Still, when the encoding is automated, the abbreviations \hat{F} in the output result will have to be expanded away.

- The latter uses an auxiliary type abbreviation \hat{F} which as themselves Deconstructive as the result of the abbreviation F to deconstruct. The deconstruction is done “incrementally”, but not lazily, i.e., since \hat{F} must itself be destructed when dispatched, which happens at the very beginning before the constraints is really solved, so there is no real gain.

8 Converting between types and canonical forms

8.1 Input types

Types constraints are usually introduced as a collection of simple equations of the form

$$(\sigma_1^j \doteq \sigma_2^j)^{j \in J}$$

Types may be large types, which can always be turned into small types with TRIM, which can be applied before hand or on demand. Types, may also contain $\mu((\alpha_i = \tau_i)^{i \in I}) \tau$ types, which can be eliminated using the following rule:

$$\begin{array}{c} \text{TRIM-REC} \\ \mu(\alpha_i = \tau_i) \tau = m \implies \exists \alpha_i^{i \in I} . \tau = m \wedge^{i \in I} (\alpha_i = \tau_i) \end{array}$$

(If we count the $\mu(\alpha) \tau$ form as a constructor in the weight of term, then TRIM-REC reduces the size of the system.)

8.2 Output types

A principal solution of a constraint C is a substitution ϕ of domain $\text{ftv}(C)$ that maps type variables to syntactic types such that solutions of C are exactly those of the form $\phi' \circ \phi$ when ϕ' range over ground substitutions—up to expansion of type abbreviations.

When C in canonical form $\exists \bar{\beta} . \wedge^{i \in 1..n} m_i$, a principal solution for C can be computed inductively as follows:

- Take ϕ_0 be the identity.
- Then, for i in $1..n$:
 - let $\bar{\alpha}_i$ be $\nu_i \cap \mathcal{V}$;
 - if m_i contains a productive view ν_i (which is then unique as m is canonical), let $\hat{\tau}_i$ be any element of ν_i and let $\hat{\phi}_i$ be $(\bar{\alpha}_i \mapsto \mu(\hat{\alpha}_i) \phi_{i-1} \hat{\tau}_i)$, where $\mu(\bar{\beta}) \sigma$ stands for $\mu(\hat{\beta}) (\sigma[\bar{\beta} \leftarrow \hat{\beta}])$ where β is an arbitrary element of $\bar{\beta}$;
 - otherwise, let $\hat{\alpha}_i$ be any element of $\bar{\alpha}_i$ and ϕ_i be $(\bar{\alpha}_i \mapsto \hat{\alpha}_i)$;
 - let ϕ_i be $\hat{\phi}_i \circ \phi_{i-1}$.
- Return ϕ_n restricted outside of $\bar{\beta}$.

Lemma 3 *The solution ϕ is a principal solution.*

Notice that types $\hat{\tau}_i$ may be chosen arbitrarily in the productive view for each productive multiequation m_i : all choices are semantically equivalent.

Still, some choices may be less arbitrary than others:

- (repr) Pick the representative of ν_i for $\hat{\tau}_i$. This is the most canonical choice: the representative is unique; its is the least unfolded presentation that sufficed to make all terms of the view agreed during constraint resolutions.
- (all) Return all terms of the view: there head symbol form a tree for \succ : starting from the leaf, the next term may be obtain by a single expansion at the root.

(leaves) Return all the leaves of the view, when ordered by \succ , e.g., terms with toplevel symbol of minimal ranks. (This would mean extending the syntax of terms to allow printing terms with synonyms, which the current syntax does not allow.)

The interest of the (repr) strategy is to be canonical. However, it tends to lose type abbreviations while the (leaves) strategy will always preserve them. The (leaves) strategy may return multiple views, which then forces a choice again, unless there is a single leave.

Degenerate type abbreviations Notice that all of these strategies ignore degenerate type abbreviations, which will then never be used in output types. One reason is that the expansion of a degenerate term is smaller than the term itself. Another reason is that showing a degenerate term does not in general prevent from also showing a productive term when the multi-equation is productive.

8.3 Removing auxiliary type abbreviations

We have shown some encodings that use auxiliary type abbreviations. These should be eliminated in output types.

Recursive type abbreviations Recursive type abbreviations uses non-recursive helper type abbreviations $F_i(\bar{\alpha}) \triangleright \mu((\beta_j = \hat{F}_j(\bar{\alpha}, \bar{\beta}))^{j \in J} \beta_j)$, but introduced in such a way that we can always read back $F_i(\bar{\tau})$ from $\hat{F}_i(\bar{\tau}, \bar{\sigma})$ at anytime (an invariant that we preserve during unification). This amounts to performing a contraction. We can indeed, perform all such contractions on canonical forms and eliminate all such helper type abbreviations before outputting terms.

Alternatively, helper type abbreviations can be retained in principal solutions and only eliminated on the fly when show printing or converting terms.

Other auxiliary type abbreviations We have also proposed type abbreviations for the productive part of deconstructive type abbreviations. To eliminate those, we need to expand them. This seems to be removing the purpose of introducing them in the first place, but this is not quite the case: some of these might have been introduced to solve intermediate constraints on existential type variables and will not appear in the output solution, while we only need to expand those appearing in the output types.

Scoped type annotations Scoped type annotations are similar, they must be removed in the output, and we can do so by forcing their expansion in the constraint at anytime. In particular, we can deal with scoped type annotations during resolution of the constraint if needed.

8.4 Example

Here is an example that loops in OCaml (version 4.09)

```
type 'a f = 'b constraint 'a = 'b option
let _ = fun (x : ('a f) option) (y : 'a) -> (x = y)
```

That is, we have an abbreviation:

$$F(G\alpha) \triangleright \alpha$$

where G , which stands for the option type, is a primitive type. Typechecking the program produces the constraint C equal to $\alpha \doteq H(F\alpha)$.

This will be reduced as follows:

$$\begin{array}{lcl}
& & C \\
\text{SPLIT} & \Rightarrow & \exists \beta \cdot \alpha \doteq H\beta \wedge \beta \doteq F\alpha \\
\text{DISPATCH-P} & \Rightarrow & \exists \beta \cdot \alpha \doteq \langle H\beta \rangle \wedge \beta \doteq F\alpha \\
\text{DISPATCH-EXP-D} & \Rightarrow & \exists \beta \cdot \alpha \doteq \langle H\beta \rangle \wedge (\exists \beta_1 \cdot \beta \doteq [F\alpha] \doteq \beta_1 \wedge \alpha \doteq H\beta_1) \\
\text{MERGE} & \Rightarrow & \exists \beta \beta_1 \cdot \alpha \doteq \langle H\beta \rangle \doteq H\beta_1 \wedge \beta \doteq \beta_1 \doteq [F\alpha] \\
\text{DISPATCH-P} & \Rightarrow & \exists \beta \beta_1 \cdot \alpha \doteq \langle H\beta \rangle \doteq \langle H\beta_1 \rangle \wedge \beta \doteq \beta_1 \doteq [F\alpha] \\
\text{DECOMPOSE} & \Rightarrow & \exists \beta \beta_1 \cdot \alpha \doteq \langle H\beta \doteq H\beta_1 \rangle \wedge \beta \doteq \beta_1 \doteq [F\alpha] \wedge \beta \doteq \beta_1 \\
\text{MERGE} & \Rightarrow & \exists \beta \beta_1 \cdot \alpha \doteq \langle H\beta \doteq H\beta_1 \rangle \wedge \beta \doteq \beta_1 \doteq \beta_1 \doteq [F\alpha] \\
\text{CLEAN-P} & \Rightarrow & \exists \beta \beta_1 \cdot \alpha \doteq \langle H\beta \rangle \wedge \beta \doteq \beta_1 \doteq \beta_1 \doteq [F\alpha] \\
\text{STUTTER} & \Rightarrow & \exists \beta \beta_1 \cdot \alpha \doteq \langle H\beta \rangle \wedge \beta \doteq \beta_1 \doteq [F\alpha]
\end{array}$$

The solution ϕ , is defined as follows:

$$\begin{aligned}
\phi_0 &= id \\
\phi_1 &= (\alpha \mapsto \mu(\alpha) H\beta) \circ \phi_0 = (\alpha \mapsto H\beta) \\
\phi_2 &= (\beta, \beta_1 \mapsto \beta) \circ \phi_1 = (\beta_1 \mapsto \beta, \alpha \mapsto H\beta) \\
\phi &= \phi_2 \setminus \{\beta, \beta_1\} = (\alpha \mapsto H\beta)
\end{aligned}$$

That is ϕ is equal to $\alpha \mapsto H\beta$. As noticed earlier, F does not appear in the output type since it is degenerate. We may still verify that ϕ solves the input constraint, indeed:

$$\phi\alpha = H\beta =_E H(F(H\beta)) = \phi(H(F\alpha))$$

8.5 Computing free variables

In section 4.2, we defined the set $\text{ftv}(\tau)$ of syntactic free variables, computed as usual. We also defined the set $\text{ftv}^\nabla(\tau)$ of semantic free variables as $\text{ftv}(\nabla\tau)$, which can still be computed on τ without fully expanding τ .

For this purpose, we only need to expand (the) deconstructive (part of) type abbreviations—if only to check that the constraints are satisfied. Then, we can

Acknowledgments

This core of this work, including the treatment constructive type abbreviations, is entirely based on the master internship of Morel [2019].

References

Carine Morel. Type inference and modular elaboration with constraints for ML extended with type abbreviations. Master's thesis, Université Paris Diderot-Paris 7, September 2019. URL <https://hal.inria.fr/hal-02361707>.