

Extensible records with scoped labels

Daan Leijen

Institute of Information and Computing Sciences, Utrecht University
P.O. Box 80.089, 3508 TB Utrecht, the Netherlands

daan@cs.uu.nl

Draft, Revision: 76 , July 23, 2005

Abstract

Records provide a safe and flexible way to construct data structures. We describe a natural approach to typing polymorphic and extensible records that is simple, easy to use in practice, and straightforward to implement. A novel aspect of this work is that records can contain duplicate labels, effectively introducing a form of scoping over the labels. Furthermore, it is a fully orthogonal extension to existing type systems and programming languages. In particular, we show how it can be used conveniently with standard Hindley-Milner, qualified types, and MLF.

1. Introduction

Tuples, or products, group data items together and are a fundamental concept to describe data structures. In ML and Haskell, we can construct a product of three integers as:

$$(7, 7, 1973)$$

Records are tuples where the individual components are labeled. Using curly braces to denote records, we can write the above product more descriptively as:

$$\{day = 7, month = 7, year = 1973\}$$

The record notation is arguably more readable than the plain product. It is also safer as we identify each component explicitly, preventing an accidental switch of the day and month for example.

Even though records are fundamental building blocks of data structures, most programming languages severely restrict their use: labels can not be reused at different types, records must be explicitly declared and are not extensible, etc. This is surprising given the large amount of research that has gone into type systems and compilation methods for records. We believe that the complexity of the proposed systems is one of the most important reasons that they are not yet part of mainstream programming languages. Most systems require non-trivial extensions to a type system that are hard to implement, and, perhaps even more important, that are difficult to explain to the user.

For all systems described in literature, it is assumed that records do not contain duplicate labels. In this article we take a novel view at records where duplicate labels are allowed and retained, effectively introducing a form of scoping over the labels. This leads to a simple and natural system for records that integrates seamlessly with many other type systems. In particular:

- The types are straightforward and basically what a naïve user would expect them to be. The system is easy to use in practice, as the user is not confronted with artificial type system constructs. Of course, all operations are checked and the type system statically prevents access to labels that are absent.

- The records support scoped labels since fields with duplicate labels are allowed and retained. As records are equivalent up to permutation of distinct labels, all basic operations are still well-defined. The concept of scoped labels is useful in its own right and can lead to new applications of records in practice.
- The system is straightforward to implement using a wide range of implementation techniques. For predicative type systems, we can guarantee constant-time field selection.
- The system works with just about any polymorphic type system with minimal effort. We only define a new notion of equality between (mono) types and present an extended unification algorithm. This is all completely independent of a particular set of type rules. We show how it can be used specifically with MLF [15], a higher-ranked, impredicative type system. Building on MLF, we can model a form of first-class modules with records.

The entire system is implemented in the experimental language Morrow [16]. The type system of Morrow is based on MLF, and all the examples in this article, including the first-class modules, are valid Morrow programs.

The work described here builds on numerous other proposals for records, in particular the work of Wand [35], Remy [29], and Gaster and Jones [7]. One can view our work as just a small variation of the previous systems. However, we believe that our design is an important variation, as it leads to a record system with much less complexity. This makes our design more suitable for integration with existing type systems and programming languages.

In the next section we introduce the basic record operations. We explain the type rules and discuss what effect scoped labels have on programs. In Section 4 we show how our system can be used with MLF to encode a form of first-class modules. In Section 5 we discuss how our systems can also support variant types. We formalize the type rules and inference in section 6 and 7. We conclude with an overview of implementation techniques and related work.

2. Record operations

Following Cardelli and Mitchell [2] we define three primitive operations on records: *selection*, *restriction*, and *extension*. Furthermore, we add the constant $\{\}$ as the empty record.

Extension. We can extend a record r with a label l and value e using the syntax $\{l = e \mid r\}$. For example:

$$origin = \{x = 0 \mid \{y = 0 \mid \{\}\}\}$$

To reduce the number of braces, we abbreviate a series of extensions using comma separated fields, and we leave the extension of

the empty record implicit. The above example can thus be written more conveniently as:

$$\text{origin} = \{x = 0, y = 0\}$$

The construction of the record is anonymous: we do not have to declare this record or its fields in advance. Furthermore, extension is polymorphic and not limited to records with a fixed type, but also applies to previously defined records, or records passed as an argument:

$$\begin{aligned} \text{origin3} &= \{z = 0 \mid \text{origin}\} \\ \text{named } s \ r &= \{\text{name} = s \mid r\} \end{aligned}$$

Selection. The selection operation $(r.l)$ selects the value of a label l from a record r . For example, we can define a function *distance* that calculates the distance of a point to the origin:

$$\text{distance } p = \text{sqrt} ((p.x * p.x) + (p.y * p.y))$$

In contrast to many programming languages, the *distance* function works for any record that contains an x and y field of a suitable numeric type. For example, we can use this function on records with a different set of fields:

$$\text{distance} (\text{named "2d" } \text{origin}) + \text{distance } \text{origin3}$$

Restriction. Finally, the restriction operation $(r - l)$ removes a label l from a record r . Using our primitive operations, we can now define the common *update* and *rename* operations:

$$\begin{aligned} \{l := x \mid r\} &= \{l = x \mid r - l\} && \text{-- update } l \\ \{l \leftarrow m \mid r\} &= \{l = r.m \mid r - m\} && \text{-- rename } m \text{ to } l \end{aligned}$$

Here is an example of using update to change the x and y components of a point:

$$\text{move } p \ dx \ dy = \{x := p.x + dx, y := p.y + dy \mid p\}$$

Note that *move* works on any record containing an x and y field, not just points. Effectively, we use parametric polymorphism to model a limited form of subtyping [2].

2.1 Safe operations

The type system ensures statically that all record operations are safe. In particular, it ensures that record selection and restriction are only applied when the field is actually present. For example, the following expressions are both rejected by the type system:

$$\begin{aligned} \{x = 1\}.y \\ \text{distance } \{x = 1\} \end{aligned}$$

Our type system accepts the extension of a record with a field that is already present, and the following example is accepted:

$$\{x = 1 \mid \text{origin}\}$$

We call this *free* extension. Many type systems in the literature require that a record can only be extended with a label that absent, which we call *strict* extension. We believe that strict extension unnecessarily restricts the programs one can write. For example, the function *named* extends *any* record with a new *name* field. In a system with strict extension, we need to write two functions: one for records without the label, and one for records that already contain the label. In this particular example this is easy to do, but in general we might want to extend records locally with helper fields. Without free extension, the local extensions would artificially restrict the use of the function.

There are two possible semantics we can give to free extension. If a duplicate label is encountered we can choose to overwrite the previous field with the new field, or we can choose to retain the old field. All previous proposals that allow free extension [35, 29, 1] use the first approach. In those systems, extension is really a mixture of update and extension: if a field is absent, the record is extended. If the field is already present, the previous value is overwritten, after which it is no longer accessible.

We take a novel approach to free extension where the previous fields are always retained, both in the value and in the type. In our system, we clearly separate the concepts of update and extension. To keep selection and restriction well-defined, we need to explicitly define these operations to work on the first matching label in a record. Therefore, we can always unambiguously select a particular label:

$$\begin{aligned} \{x = 2, x = \text{True}\}.x &&& \text{-- select the first } x \text{ field} \\ (\{x = 2, x = \text{True}\} - x).x &&& \text{-- select the second } x \text{ field} \end{aligned}$$

Since previous fields are retained, our record system effectively introduces a form of scoping on labels. This is certainly useful in practice, where we can use scoped labels to model environments with access to previously defined values. For example, suppose we have an environment that includes the current text color:

$$\text{putText } \text{env } s = \text{putStr} (\text{ansiColor } \text{env}.color \ s)$$

We can define a combinator that temporarily changes the output color:

$$\text{warning } \text{env } f = f \ \{\text{color} = \text{red} \mid \text{env}\}$$

The function f passed to *warning* formats its output in a red color. However, it may want to format certain parts of its output in the color of the parent context. Using scoped labels, this is easily arranged: we can remove the first color field from the environment, thereby exposing the previous color field automatically (if present):

$$f \ \text{env} = \text{putText} (\text{env} - \text{color}) \ \text{"parent color"}$$

As we see in the next section, the type of the function f reflects that the environment is required to contain at least two *color* fields.

Another example of scoped labels occurs when encoding objects as records. Redefined members in a sub-class are simply extensions of the parent class. The scoped labels can now be used to access the overridden members in a parent class.

One can argue that free extension can lead to programming errors where one accidentally extends a record with a duplicate label. However, the type system can always issue a warning if a record with a fixed type contains duplicate labels, which could be attributed to a programmer mistake. This is comparable to a standard shadowed variable warning – and indeed, a warning is more appropriate here than a type error, since a program with duplicate labels can not go wrong!

3. The types of records

We write the type of a record as a sequence of labeled types. To closely reflect the syntax of record values, we enclose record types in curly braces $\{\}$ too:

$$\text{type } \text{Point} = \{x :: \text{Int}, y :: \text{Int}\}$$

As we will see during the formal development, it makes sense to talk about a sequence of labeled types as a separate concept. We call such sequence a *row*. Following Gaster and Jones [7], we consider an extensible row calculus where a row is either empty or an extension of a row. The empty row is written as $\langle \rangle$ and

the extension of a row r with a label l and type τ is written as $(l :: \tau \mid r)$. The full unabbreviated type of a *Point* is written with rows as:

$$\text{type Point} = \{(\{x :: \text{Int} \mid (\{y :: \text{Int} \mid (\{\})\})\})\}$$

Just like record extension, we abbreviate multiple extensions with a comma separated list of fields. Furthermore, we leave out the row brackets if they are directly enclosed by record braces.

3.1 Types of record operations

Using row types, we can now give the type signatures for the basic record operations:

$$\begin{aligned} (-.l) &:: \forall r\alpha. \{l :: \alpha \mid r\} \rightarrow \alpha \\ (- - l) &:: \forall r\alpha. \{l :: \alpha \mid r\} \rightarrow \{r\} \\ \{l = - \mid -\} &:: \forall r\alpha. \alpha \rightarrow \{r\} \rightarrow \{l :: \alpha \mid r\} \end{aligned}$$

Note that we assume a distfix notation where argument positions are written as “-”. Furthermore, we explicitly quantify all types in this paper, but practical systems can normally use implicit quantification. The selection operator $(-.l)$ takes a record that contains a field l of type α , and returns the value of type α . Similarly, the restriction operator $(- - l)$ returns the record without the l field. The type of extension is very natural: it takes a value α and any record $\{r\}$, and extends it with a new field $l :: \alpha$. Here is for example the inferred type for *origin*:

$$\begin{aligned} \text{origin} &:: \{x :: \text{Int}, y :: \text{Int}\} \\ \text{origin} &= \{x = 0, y = 0\} \end{aligned}$$

The type of selection naturally ensures that a label is present when it is selected. For example, $\text{origin}.x$ is well-typed, since the type of the record, $\{x :: \text{Int}, y :: \text{Int}\}$, is an instance of the type of the expected argument $\{x :: \alpha \mid r\}$ of the selector function $(-.x)$. Unfortunately, at this point, the type signatures are too strong: the valid expression $\text{origin}.y$ is still rejected as $\{x :: \text{Int}, y :: \text{Int}\}$ is just not an instance of $\{y :: \alpha \mid r\}$.

To accept the above selection, we need a new notion of equality between types where the rows are considered equal *up to permutation of distinct labels*. The new equality relation (\cong) is formalized in Figure 1. The first three rules are standard. Rule $(eq-trans)$ defines equality as a transitive relation. The last two rules define equality between rows. Rule $(eq-head)$ defines two rows as equal when their heads and tails are equal. The rule $(eq-swap)$ is the most interesting: it states that the first two fields of a row can be swapped if (and only if) their labels are different. Together with transitivity $(eq-trans)$ and row equality $(eq-head)$, this effectively allows us to swap a field repeatedly to the front of a record, but not past an equal label. With the new notion of equality, we can immediately derive that:

$$\{x :: \text{Int}, y :: \text{Int}\} \cong \{y :: \text{Int}, x :: \text{Int}\}$$

The expression $\text{origin}.y$ is now well-typed since the isomorphic type $\{y :: \text{Int}, x :: \text{Int}\}$ is an instance of $\{y :: \alpha \mid r\}$. The new notion of equality is the only addition needed to integrate our notion of records with a specific type system. Since no other concepts are introduced, the types of the primitive operations are basically what a naïve user would expect them to be. The same holds for the inferred types of derived operations such as update and rename:

$$\begin{aligned} \{l := - \mid -\} &:: \forall r\alpha\beta. \alpha \rightarrow \{l :: \beta \mid r\} \rightarrow \{l :: \alpha \mid r\} \\ \{l := x \mid r\} &= \{l = x \mid r - l\} \end{aligned}$$

$$\begin{aligned} \{l \leftarrow m \mid -\} &:: \forall r\alpha. \{m :: \alpha \mid r\} \rightarrow \{l :: \alpha \mid r\} \\ \{l \leftarrow m \mid r\} &= \{l = r.m \mid r - m\} \end{aligned}$$

$$\begin{aligned} (eq-var) & \quad \alpha \cong \alpha \\ (eq-const) & \quad c \cong c \\ (eq-app) & \quad \frac{\tau_1 \cong \tau'_1 \quad \tau_2 \cong \tau'_2}{\tau_1 \tau_2 \cong \tau'_1 \tau'_2} \\ (eq-trans) & \quad \frac{\tau_1 \cong \tau_2 \quad \tau_2 \cong \tau_3}{\tau_1 \cong \tau_3} \\ (eq-head) & \quad \frac{\tau \cong \tau' \quad r \cong s}{(l :: \tau \mid r) \cong (l :: \tau' \mid s)} \\ (eq-swap) & \quad \frac{l \neq l'}{(l :: \tau, l' :: \tau' \mid r) \cong (l' :: \tau', l :: \tau \mid r)} \end{aligned}$$

Figure 1. Equality between (mono) types

We see that the type of update is very natural: given a record with an l field of type β , we can assign it a new value of a possibly different type α .

3.2 Scoped labels

As remarked before, the type signature for record extension is free and does not reject duplicate labels. For example, both of the following expressions are well-typed:

$$\begin{aligned} \{x = 2, x = \text{True}\} &:: \{x :: \text{Int}, x :: \text{Bool}\} \\ \{x = \text{True}, x = 2\} &:: \{x :: \text{Bool}, x :: \text{Int}\} \end{aligned}$$

Note that the types of the two expressions are not equivalent though. Since rule $(eq-swap)$ only applies to distinct labels, selection and restriction are still well-defined operations. For example, the following expression selects the second field, as signified by the derived type:

$$(\{x = 2, x = \text{True}\} - x).x :: \text{Bool}$$

This example shows that it is essential to retain duplicate fields not only in the runtime value, but also in the static type of the record.

3.3 Related designs

In this section we briefly compare our records with closely related systems. The main point that we wish to demonstrate here is how other designs use types that are harder to understand, or lead to complicated implementations. A more complete overview of related work is deferred to Section 9.

Predicates. Gaster and Jones [7] presented an elegant type system for records and variants based on the theory of qualified types [10, 11]. Special *lacks* predicates are used to prevent records from containing duplicate labels. For example, extension is only valid on records that do not yet contain that label, and is thus qualified with a lacks predicate:

$$\{l = - \mid -\} :: \forall r\alpha. (r \setminus l) \Rightarrow \alpha \rightarrow \{r\} \rightarrow \{l :: \alpha \mid r\}$$

As we have argued before, strict extension unnecessarily restricts the programs we can write, and one can view the lacks predicate as an unnecessary part of the type. Also, each use of a label leads to a lacks predicate, which in turn can lead to large types that are hard to read. Furthermore, the system relies essentially on a type system that supports qualified types. Adding a fundamental extension like qualified types to an ML-like language just to support records could be hard to justify in practice.

Flags. Remy [29] developed a flexible system of extensible records where flags are used to denote the presence or absence of labels and rows. The type of free extension becomes:

$$\{l = - \mid -\} :: \forall r \varphi \alpha. \alpha \rightarrow \{l :: \varphi \mid r\} \rightarrow \{l :: \text{pre}(\alpha) \mid r\}$$

The type variable φ ranges over field types that can be either absent, written as *abs*, or present with type τ , written as $\text{pre}(\tau)$. By being polymorphic in the presence or absence of the field l , the type of extension encodes that a field that is already present is overwritten, while an absent field becomes present. As we can see here, such type would not be obvious to the naïve user; especially since absent labels in the value can be present in the type (with an absent flag *abs*). We also believe that the semantics of extension should not be a subtle mixture of update and extension. Indeed, a system based on flags can not give a type to our notion of proper free extension since it essentially views records as total functions from labels to values, where duplicate labels are certainly not allowed.

Row extension. Wand [35, 36] was the first to introduce row variables, and described a type system that is very close to our proposal. The types of the basic operations are given as type rules, but these rules have basically the same effect as the type signatures in our system. However, this system also uses an ambiguous interpretation of free extension as a mixture of update and extension. This gives rise to fundamental problems in the type system. Take for example the following expression:

$$\backslash r \rightarrow \text{if } \text{True} \text{ then } \{x = 1 \mid r\} \text{ else } \{x = 2 \mid \{\}\}$$

If extension overwrites previous fields, we can assign two different types to r , namely $\{x :: \text{Int}\}$ or the empty record $\{\}$. This is why Wand's original record system proved incomplete [35]. With scoped labels, the choice is unambiguous and r must have type $\{\}$. We show completeness (and soundness) of type inference formally in Section 7.

4. Modules

In this section, we give some examples how our record system can be used to encode (a form of) first-class modules and objects.

4.1 Polymorphic records

Since the type signatures for record operations are so general, we can conveniently package related functions together. Taking an example from Jones [14], we can write a type signature for complex numbers that is parameterized over the implementation type cpx :

$$\begin{aligned} \text{type } \text{Complex } cpx = \{ & \text{cart} :: \text{Float} \rightarrow \text{Float} \rightarrow cpx \\ & , \text{polar} :: \text{Float} \rightarrow \text{Float} \rightarrow cpx \\ & , \text{re} :: cpx \rightarrow \text{Float} \\ & , \text{im} :: cpx \rightarrow \text{Float} \\ & \dots \} \end{aligned}$$

Two obvious ways to implement complex numbers are using pairs of floating point numbers with either a cartesian or a polar representation:

$$\begin{aligned} \text{cartCpx} = \{ & \text{cart } x \ y = (x, y) \\ & , \text{polar } r \ a = (r * \cos a, r * \sin a) \\ & , \text{re } (x, y) = x \\ & \dots \} \end{aligned}$$

$$\begin{aligned} \text{polarCpx} = \{ & \text{cart } x \ y = (\text{sqrt } (x * x + y * y), \text{atan2 } y \ x) \\ & , \text{polar } r \ a = (r, a) \\ & , \text{re } (r, a) = r * \cos a \end{aligned}$$

...}

As convenient syntactic sugar, we abbreviate the binding of a functional value ($l = \backslash x_1 \dots x_n \rightarrow e$) as $(l \ x_1 \dots x_n = e)$. Of course, the types of both implementations are the same, namely $\text{Complex } (\text{Float}, \text{Float})$. The ability to treat the type of a complex number representation polymorphically, allows us to create new polymorphic records from any complex number. Take for example the following simple definition of an arithmetic package over elements of type α :

$$\begin{aligned} \text{type } \text{Arith } \alpha = \{ & \text{plus} :: \alpha \rightarrow \alpha \rightarrow \alpha \\ & , \text{zero} :: \alpha \} \end{aligned}$$

Relying on standard Hindley-Milner polymorphism, we can construct an arithmetic package from any complex number type:

$$\begin{aligned} \text{arithCpx} :: \forall \alpha. \text{Complex } \alpha \rightarrow \text{Arith } \alpha \\ \text{arithCpx } c = \{ & \text{plus } x \ y = c.\text{cart } (c.\text{re } x + c.\text{re } y) \\ & \quad (c.\text{im } x + c.\text{im } y) \\ & , \text{zero} = c.\text{cart } 0 \ 0 \} \end{aligned}$$

The last definition corresponds closely to a functor definition in SML. However, no special syntax or constructions are necessary and everything is done using standard record operations and normal functions. As noted by Jones [14] the use of the type variable α intuitively corresponds to a sharing specification in SML.

The above examples are relatively simple as the records do not contain polymorphic functions themselves. Things become more interesting when we define a signature for a *Prelude* package for example:

$$\begin{aligned} \text{type } \text{Prelude} = \{ & \text{id} :: \forall \alpha. \alpha \rightarrow \alpha \\ & , \text{const} :: \forall \alpha \beta. \alpha \rightarrow \beta \rightarrow \alpha \\ & , \text{zip} :: \forall \alpha \beta. [\alpha] \rightarrow [\beta] \rightarrow [(\alpha, \beta)] \\ & \dots \} \end{aligned}$$

In this signature, the types of *id* and *const* use impredicative higher-rank polymorphism since the quantifiers are nested inside the record structure, and type inference for impredicative rank- n polymorphism is a notoriously hard problem [15, 25, 21].

4.2 Higher-ranked impredicative records

When we move to more complicated type systems, our framework of records proves its value, since it only relies on a new notion of equality between (mono) types and no extra type rules are introduced. This means that it becomes relatively easy to add our system to just about any polymorphic type system. In particular, it integrates seamlessly with MLF, an elegant impredicative higher-ranked type inference system by Le Botlan and Remy [15]. We have a full implementation of this system in the experimental Morrow compiler [16] and all the examples in this article are valid Morrow programs.

The combination of MLF with anonymous polymorphic and extensible records (and variants) leads to a powerful system where fields can have full polymorphic type signatures. For example, we can now implement our previous *Prelude* example as a first-class structure:

$$\begin{aligned} \text{prelude} :: \text{Prelude} \\ \text{prelude} = \{ & \text{id } x = x \\ & , \text{const } x \ y = x \\ & \dots \} \end{aligned}$$

Since Morrow uses the MLF type system, the (higher-ranked) type for *prelude* is automatically inferred and the type annotation is not necessary. Neither is it necessary to declare the *Prelude* type; we

can just construct an anonymous record. Indeed, we can use the *id* member polymorphically without any type declaration:

```
identity :: ∀α. α → α
identity x = prelude.id prelude.id x
```

In particular, as imposed by MLF, type annotations are only necessary when function arguments of an unknown type are used polymorphically. For example, we can not leave out the type signature if we use an unknown implementation of a prelude passed as an argument:

```
twice :: Prelude → (Int, Bool)
twice p = (p.id 1, p.id True)
```

In the above example, the *id* field of the argument *p* is used polymorphically at two different types, and thus the type signature is required to reveal that the *id* field of a *Prelude* is indeed a polymorphic function.

Apart from abstract types, these examples are very close to the goals of the XHM system, described by Jones [14] as an approach to treat modules as first-class citizens. We believe that our notion of records in combination with higher-order polymorphic MLF is therefore a significant step towards a realistic implementation of polymorphic and extensible first-class modules.

5. Variants

The design that we described lends itself well to also include *variants*. Where records model labeled tuples, variants model a labeled choice among values. It is beyond the scope of this paper to describe the full implications of programming with variants, but we will describe how our basic design can be easily extended with support for variants. We denote variant types by enclosing them in angled brackets ($\langle \rangle$). Here is an example of a variant that models an event:

```
type Event = ⟨key :: Char, mouse :: Point⟩
```

Note that whereas a record includes all fields, the variant describes a choice: an *Event* is either *key* event with *Char* value, or a *mouse* event with a *Point* value. As we can see however, the types of records and variants can both share the concept of a row as a sequence of labeled sequence types. Indeed, the unabbreviated type of an *Event* is:

```
type Event = ⟨(key :: Char | mouse :: Point | ())⟩
```

Since the rules for type equality in Figure 1 directly work on rows, we do not need to change anything to our basic type system. Just like records, there are three primitive operations on variants: *injection*, *embedding*, and *decomposition*.

```
⟨l = _⟩ :: ∀αr. α → ⟨l :: α | r⟩ -- injection
⟨l | _⟩ :: ∀αr. r → ⟨l :: α | r⟩ -- embedding
⟨l ∈ -? _ : _⟩ :: ∀αβr. ⟨l :: α | r⟩ → (α → β) → (⟨r⟩ → β) → β -- decomposition
```

The injection operation creates a variant value tagged with a label *l*. For example, we can construct a key event for the tab key as follows:

```
tab :: ∀r. ⟨key :: Char | r⟩
tab = ⟨key = '\t'⟩
```

The type signature of *tab* is polymorphic in the row *r* and can thus be used in any context where a variant with a *key* field is expected.

The embedding operation embeds a value in a variant type that also allows for a label *l*:

```
⟨mouse | tab⟩ :: ∀αr. ⟨mouse :: α, key :: Char, r⟩
```

In general, this operation is not so useful since injection already constructs polymorphic variants. However, we will see that in combination with duplicate labels this becomes an essential operation for variants.

The decomposition operations allows us to check if a variant is tagged with a particular label *l* and act accordingly. Here is an example of a function to show an *Event*:

```
showEvent :: Event → String
showEvent e
= (key ∈ e ? (\c → showChar c)
  : (\e' → (mouse ∈ e' ? (\p → showPoint p)
    : error "unmatched variant")))
```

If the variant is tagged with *key*, the function $(\backslash c \rightarrow \text{showChar } c)$ is applied to its value. If it doesn't match, we continue with a check for the *mouse* label. Note that the passed argument *e'* is the same value as *e* but has a more specific type, namely $\langle \text{mouse} :: \text{Point} \rangle$. If the variant doesn't match either label, an exception is raised. Of course, programming languages can add extra syntax to make it more convenient to match on variants. For example, in Morrow we can write the above match as:

```
showEvent :: Event → String
showEvent e = case e { key c = showChar c
                      , mouse p = showPoint p }
```

A curious property of our system is that variants can now contain duplicate tags. However, all operations are still well-defined and can not "go wrong". For example, we can match on the second tag of a variant by doing two decompositions in a row:

```
f :: ∀αβ. ⟨l :: α, l :: β⟩ → Int
f v = case v { l x = 1 -- first l tag
              , l y = 2 -- second l tag }
```

Even though this may look too liberal, we believe this is not the case because it will take conscious effort to actually construct a variant type with a scoped label. Indeed, the only operation that can do this is embedding. Since it does not serve any other purpose, it can be viewed as a programmers annotation that a scoped label is really required.

A useful application of scoped labels arises when modelling open-ended data types like user-defined exceptions. Suppose we have a function *throw* that raises an exception as a tagged value:

```
throw :: ∀αr. ⟨r⟩ → α
```

Each library can throw its own kinds of exception variants and it is not unlikely that the names for exceptions might clash between different libraries. The embedding operator can be used to prioritize identically named exceptions of different libraries, where the global exception handler matches on duplicate labels.

6. Type rules

In this section, we formalize the concept of rows and define the structure of types. First, we have to make some basic assumptions about the structure of types. This structure is needed as not all types are well-formed. For example, a row type can not extend an integer and a *Maybe* type needs a parameter:

```
(l = Maybe | Int)
```


If we look closely at the rules in Figure 3 there are only two possible substitutions as the outcome of a row rewrite. When a label l can be found, the substitution will be empty as only the rules (*row-swap*) and (*row-head*) apply. If a label is not present, the rule (*row-var*) applies and a singleton substitution $[\alpha \mapsto (\lambda :: \gamma \mid \beta)]$ is returned, where α is the tail of s . Therefore, the side-condition $\text{tail}(r) \notin \text{dom}(\theta_1)$ prevents us from unifying rows with a common tail but a distinct prefix¹. Here is an example where unification would not terminate without the side condition:

$$\backslash r \rightarrow \text{if } \text{True} \text{ then } \{x = 2 \mid r\} \text{ else } \{y = 2 \mid r\}$$

During type inference, the rows in both **if** branches are unified:

$$(\lambda :: \text{Int} \mid \alpha) \sim (\lambda :: \text{Int} \mid \alpha) : \theta_3 \circ \theta_2 \circ \theta_1$$

Which implies that $(\lambda :: \text{Int} \mid \alpha)$ is rewritten as:

$$(\lambda :: \text{Int} \mid \alpha) \simeq (\lambda :: \gamma \mid y :: \text{Int} \mid \beta) : \theta_1$$

Where $\theta_1 = [\alpha \mapsto (\lambda :: \gamma \mid \beta)]$. After unification of γ and Int , the unification of the row tails is now similar to the initial situation and thus loops forever:

$$\begin{aligned} \theta_2(\theta_1 \alpha) &\sim \theta_2(\theta_1((\lambda :: \text{Int} \mid \beta))) : \theta_3 \\ &= \\ (\lambda :: \text{Int} \mid \beta) &\sim (\lambda :: \text{Int} \mid \beta) : \theta_3 \end{aligned}$$

However, with the side condition in place, no such thing will happen since $\text{tail}(r) = \alpha \in \{\alpha\} = \text{dom}(\theta_1)$. Not all record systems described in literature correctly ensure termination of record unification for this class of programs. For example, the unification rules of TREX fail to terminate for this particular example [7].

The reader might be worried that the side condition endangers the soundness or completeness of unification, but such is not the case, as asserted by the following theorems.

Theorem 1. *Unification is sound. If two types unify they are equal under the resulting substitution: $\tau \sim \tau' : \theta \Rightarrow \theta\tau \cong \theta\tau'$.*

Proof. Proved by straightforward induction over the cases of the unification algorithm. A full proof can be found in [17].

Theorem 2. *Unification is complete. If two types are equal under some unifier, unification will succeed and find a most general unifier: $\theta\tau \cong \theta\tau' \Rightarrow \tau \sim \tau' : \theta_1 \wedge \theta_1 \sqsubseteq \theta$.*

Proof. Standard proof of completeness over the structure of types. A constructive proof is given in a separate technical report [17].

As most type inference algorithms reduce to a set of unification constraints, soundness and completeness results carry over directly with the above results for row unification. In particular, the proofs for Hindley-Milner, qualified types [11], and MLF[15] are easily adapted to hold in the presence of row unification.

8. Implementing records

Providing an efficient implementation for extensible and polymorphic records is not entirely straightforward. In this section we discuss several implementation techniques and show in particular how standard compilation techniques can be used to provide constant-time access for label selection.

Association lists. A naïve implementation of records uses a simple association list of label-value pairs. Selection is implemented as

¹ In practice, this side condition can also be implemented by passing $\text{tail}(r)$ to the (\simeq) function and checking in (*row-var*) that $\alpha \neq \text{tail}(r)$

a linear search over this list, where the type system ensures that such label is always found. Extension is a constant time operation that adds an element to the head of the list. However, in practice the most common operation by far is label selection, and a linear algorithm would be slow for larger records. This holds especially when using records as modules that can easily contain hundreds of labels.

Labeled vectors. A more efficient representation for label selection is a vector of label-value pairs where the fields are sorted on the label according to some order on the labels. Extension is now a linear operation as the full vector must be copied. This can be made somewhat more efficient by processing multiple extensions together in a single operation. Label selection can now use a binary search to select a particular label and becomes an $O(\log(n))$ operation. When labels can be compared efficiently, for example by using a Garrigue’s hashing scheme [4], the binary search over the vector can be implemented very efficiently. It is also possible to improve the search time for small and medium sized records by using a partially evaluated header [32], but at the price of a potentially more expensive extension operation.

Labeled vectors + constant folding. Label selection can be divided into two separate operations: looking up the index of the field in the record (*lookup*), and selecting the value using that index (*select*). When the labels of the record are known it is possible to partially evaluate the *lookup* operation using standard compiler techniques. For example, the expression $\{l = \text{expr}\}.l$ could be translated with *lookup* and *select* as:

$$\begin{aligned} \text{let } r &= \{l = \text{expr}\} \\ & \quad i = \text{lookup } r \ l \\ \text{in } \text{select } r \ i \end{aligned}$$

Since the type of r is known, the compiler can statically evaluate *lookup* $r \ l$ and replace it by 0, avoiding a binary search at runtime:

$$\text{let } r = \{l = \text{expr}\} \text{ in } \text{select } r \ 0$$

This optimization by itself guarantees constant-time label selection for all records with a fixed set of labels. In contrast, if the record type is open, the *lookup* operation can not be evaluated statically. Here is an example:

$$\begin{aligned} \text{dist} &:: \forall r. \{\lambda :: \text{Int}, y :: \text{Int} \mid r\} \rightarrow \text{Int} \\ \text{dist } r &= r.x + r.y \end{aligned}$$

However, at the call site, the type of the row variable r is generally known. A compiler can take advantage of this by floating all *lookup* operations upwards and split the function *dist* into two functions, a wrapper and a worker:

$$\begin{aligned} \text{dist } r &= \text{dist_work } (\text{lookup } r \ x) (\text{lookup } r \ y) \ r \\ \text{dist_work } i \ j \ r &= \text{select } r \ i + \text{select } r \ j \end{aligned}$$

If the wrapper function *dist* is inlined at a call site there is a good chance that the type of r is known. In that situation, the offsets for both labels are passed as a runtime parameter giving constant-time access even for record selection on open records. This technique is basically a variant of the worker-wrapper transformation in GHC [24]. An aggressive compiler would also push the *lookup* operation through extensions in order to float properly. For example:

$$\backslash r \rightarrow \{x = 2 \mid r\}.y$$

The compiler would first translate this to:

Wand [35, 36] was the first to use row variables to capture the subtype relationship between records using standard parametric polymorphism. As we have seen in Section 3.3 not all programs have a principal type in this system. The work of Wand is later refined by Berthomieu and Sagazan [1] where a polynomial unification algorithm is presented. Remy [29] extended the work of Wand with a flexible system of extensible records with principle types. Flags are used to denote the presence or absence of labels and rows.

Ohori [23] was the first to present an efficient compilation method for polymorphic records with constant time label selection, but only for non-extensible rows. Subsequently, Gaster and Jones [7, 6] presented an elegant type system for records and variants based on the theory of qualified types [10, 11]. They use strict extension with special *lacks* predicates to prevent duplicate labels. The predicates correspond to runtime label offsets, and standard evidence translation gives a straightforward compilation scheme with constant time label selection.

The primitive operations in our system are based on extension, but many calculi use record concatenation as the primitive operation [8, 30, 31, 36, 34, 27, 19]. Pottier [28] describes a general framework for row based type systems which has an effective type inference algorithm for record concatenation. Unfortunately, the types of those systems are generally hard to use in practice. Here is for example the type of free extension in Pottier's system:

$$\{l = - \mid -\} :: (\{l\} : \partial(\text{pre}(\alpha)) \leq r_2, (\mathcal{L} / \{l\}) : r_1 \leq r_2) \\ \Rightarrow \alpha \rightarrow r_1 \rightarrow r_2$$

The first constraint is a singleton filter to indicate that the result contains field l , while the second constraint is a co-singleton filter that indicates that all fields other than l have the same status as in the argument.

10. Conclusion

We believe that polymorphic and extensible records are a flexible and fundamental concept to program with data structures. The complexity of type systems for such records used to prevent widespread adoption in mainstream languages. We presented polymorphic and extensible records based on scoped labels that are safe, convenient, easy to check, and straightforward to compile – every programming language should have them!

Acknowledgments

We would like to thank François Pottier for his feedback and for pointing out the similarities between our system and the label selective calculus.

References

- [1] B. Berthomieu and C. de Sagazan. A calculus of tagged types, with applications to process languages. In *TAPSOFT Workshop on Types for Program Analysis*, May 1995.
- [2] L. Cardelli and J. Mitchell. Operations on records. *Journal of Mathematical Structures in Computer Science*, 1(1):3–48, Mar. 1991.
- [3] J. P. Furuse and J. Garrigue. A label-selective lambda-calculus with optional arguments and its compilation method. RIMS Preprint 1041, Kyoto University, Oct. 1995.
- [4] J. Garrigue. Typing deep pattern-matching in presence of polymorphic variants. In *Proceedings of the JSSST Workshop on Programming and Programming Languages*, Mar. 2004.
- [5] J. Garrigue and H. Ait-Kaci. The typed polymorphic label selective calculus. In *21th ACM Symp. on Principles of Programming Languages (POPL'94)*, pages 35–47, Portland, OR, Jan. 1994.
- [6] B. R. Gaster. *Records, Variants, and Qualified Types*. PhD thesis, Dept. of Computer Science, University of Nottingham, July 1998.
- [7] B. R. Gaster and M. P. Jones. A polymorphic type system for extensible records and variants. Technical Report NOTTCS-TR-96-3, Dept. of Computer Science, University of Nottingham, 1996.
- [8] R. Harper and B. C. Pierce. A record calculus based on symmetric concatenation. In *18th ACM Symp. on Principles of Programming Languages (POPL'91)*, pages 131–142, Jan. 1991.
- [9] J. Hindley. The principal type scheme of an object in combinatory logic. *Transactions of the American Mathematical Society*, 146:29–60, Dec. 1969.
- [10] M. P. Jones. A theory of qualified types. In *4th. European Symposium on Programming (ESOP'92)*, volume 582 of *Lecture Notes in Computer Science*, pages 287–306. Springer-Verlag, Feb. 1992.
- [11] M. P. Jones. *Qualified types in Theory and Practice*. Distinguished Dissertations in Computer Science. Cambridge University Press, 1994.
- [12] M. P. Jones. From Hindley-Milner types to first-class structures. In *Proceedings of the Haskell Workshop*, June 1995. Yale University Research Report YALEU/DCS/RR-1075.
- [13] M. P. Jones. A system of constructor classes: overloading and implicit higher-order polymorphism. *Journal of Functional Programming*, 5(1):1–35, Jan. 1995.
- [14] M. P. Jones. Using parameterized signatures to express modular structure. In *23th ACM Symp. on Principles of Programming Languages (POPL'96)*, pages 68–78. ACM Press, 1996.
- [15] D. Le Botlan and D. Rémy. MLF: Raising ML to the power of System-F. In *Proceedings of the International Conference on Functional Programming (ICFP 2003)*, Uppsala, Sweden, pages 27–38. ACM Press, aug 2003.
- [16] D. Leijen. Morrow: a row-oriented programming language. <http://www.cs.uu.nl/~daan/morrow.html>, Jan. 2005.
- [17] D. Leijen. Unqualified records and variants: proofs. Technical Report UU-CS-2005-00, Dept. of Computer Science, Universiteit Utrecht, 2005.
- [18] D. Leijen and E. Meijer. Domain specific embedded compilers. In *2nd USENIX Conference on Domain Specific Languages (DSL'99)*, pages 109–122, Austin, Texas, Oct. 1999. Also appeared in *ACM SIGPLAN Notices* 35, 1, (Jan. 2000).
- [19] H. Makhholm and J. B. Wells. Type inference and principal typings for symmetric record concatenation and mixin modules. Technical Report HW-MACS-TR-0030, School of Mathematical and Computer Sciences, Heriot-Watt University, Mar. 2005.
- [20] R. Milner. A theory of type polymorphism in programming. *Journal of Computer and System Sciences*, 17:248–375, Aug. 1978.
- [21] M. Odersky and K. Läufer. Putting type annotations to work. In *23th ACM Symp. on Principles of Programming Languages (POPL'96)*, pages 54–67, Jan. 1996.
- [22] M. Odersky, P. Wadler, and M. Wehr. A second look at overloading. In *FPCA '95: Proceedings of the seventh international conference on Functional Programming languages and Computer Architecture*, pages 135–146, 1995.
- [23] A. Ohori. A polymorphic record calculus and its compilation. *ACM Transactions on Programming Languages and Systems*, 17(6):844–895, 1995.
- [24] S. Peyton Jones and A. Santos. A transformation-based optimiser for Haskell. *Science of Computer Programming*, 32(1–3):3–47, Sept. 1998.
- [25] S. Peyton-Jones and M. Shields. Practical type inference for arbitrary-rank types. Submitted to the *Journal of Functional Programming (JFP)*, 2004.
- [26] B. C. Pierce and D. N. Turner. Simple type theoretic foundations for object-oriented programming. *Journal of Functional Programming*, 4(2):207–247, Apr. 1994.
- [27] F. Pottier. A 3-part type inference engine. In *European Symposium on Programming (ESOP)*, volume 1782 of *Lecture Notes in Computer Science*, pages 320–335. Springer Verlag, Mar. 2000.
- [28] F. Pottier. A constraint-based presentation and generalization of rows. In *Eighteenth Annual IEEE Symposium on Logic In Computer Science (LICS'03)*, pages 331–340, Ottawa, Canada, June 2003.

- [29] D. Rémy. Type inference for records in a natural extension of ML. In C. A. Gunter and J. C. Mitchell, editors, *Theoretical Aspects Of Object-Oriented Programming. Types, Semantics and Language Design*. MIT Press, 1993.
- [30] D. Rémy. Typing record concatenation for free. In C. A. Gunter and J. C. Mitchell, editors, *Theoretical Aspects Of Object-Oriented Programming. Types, Semantics and Language Design*. MIT Press, 1993.
- [31] D. Rémy. A case study of typechecking with constrained types: Typing record concatenation. Workshop on Advances in Types for Computer Science, Aug. 1995.
- [32] D. Rémy. Efficient representation of extensible records. INRIA Rocquencourt, 2001.
- [33] J. A. Robinson. A machine-oriented logic based on the resolution principle. *Journal of the ACM*, 12(1):23–41, Jan. 1965.
- [34] M. Sulzmann. Designing record systems. Technical Report YALEU/DCS/RR-1128, Dept. of Computer Science, Yale University, Apr. 1997.
- [35] M. Wand. Complete type inference for simple objects. In *Proceedings of the 2nd. IEEE Symposium on Logic in Computer Science*, pages 37–44, 1987. Corrigendum in LICS'88, page 132.
- [36] M. Wand. Type inference for record concatenation and multiple inheritance. *Information and Computation*, 93:1–15, 1991.