University of Pennsylvania
ScholarlyCommons

Technical Reports (CIS)

Department of Computer & Information Science

April 1999

Union Types for Semistructured Data

Peter Buneman
University of Pennsylvania

Benjamin C. Pierce
University of Pennsylvania, bcpierce@cis.upenn.edu

Follow this and additional works at: https://repository.upenn.edu/cis_reports

Recommended Citation

Peter Buneman and Benjamin C. Pierce, "Union Types for Semistructured Data", . April 1999.

University of Pennsylvania Department of Computer and Information Science Technical Report No. MS-CIS-99-09.

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/cis_reports/173
For more information, please contact repository@pobox.upenn.edu.

Union Types for Semistructured Data

Abstract

Semistructured databases are treated as dynamically typed: they come equipped with no independent schema or type system to constrain the data. Query languages that are designed for semistructured data, even when used with structured data, typically ignore any type information that may be present. The consequences of this are what one would expect from using a dynamic type system with complex data: fewer guarantees on the correctness of applications. For example, a query that would cause a type error in a statically typed query language will return the empty set when applied to a semistructured representation of the same data. Much semistructured data originates in structured data. A semistructured representation is useful when one wants to add data that does not conform to the original type or when one wants to combine sources of different types. However, the deviations from the prescribed types are often minor, and we believe that a better strategy than throwing away all type information is to preserve as much of it as possible. We describe a system of untagged *union types* that can accommodate variations in structure while still allowing a degree of static type checking. A novelty of this system is that it involves non-trivial equivalences among types, arising from a law of distributivity for records and unions: a value may be introduced with one type (e.g., a record containing a union) and used at another type (a union of records). We describe programming and query language constructs for dealing with such types, prove the soundness of the type system, and develop algorithms for subtyping and typechecking.

Comments

University of Pennsylvania Department of Computer and Information Science Technical Report No. MS-CIS-99-09.

Union Types for Semistructured Data

MS-CIS-99-09

Peter Buneman, Benjamin Pierce

University of Pennsylvania
School of Engineering and Applied Science
Computer and Information Science Department
Philadelphia, PA 19104-6389

1999

Union Types for Semistructured Data

Peter Buneman Benjamin Pierce

University of Pennsylvania
Dept. of Computer & Information Science
200 South 33rd Street
Philadelphia, PA 19104-6389, USA
{peter,bcpierce}@cis.upenn.edu

Technical report MS-CIS-99-09

April 6, 1999

Abstract

Semistructured databases are treated as dynamically typed: they come equipped with no independent schema or type system to constrain the data. Query languages that are designed for semistructured data, even when used with structured data, typically ignore any type information that may be present. The consequences of this are what one would expect from using a dynamic type system with complex data: fewer guarantees on the correctness of applications. For example, a query that would cause a type error in a statically typed query language will silently return the empty set when applied to a semistructured representation of the same data.

Much semistructured data originates in structured data. A semistructured representation is useful when one wants to add data that does not conform to the original type or when one wants to combine sources of different types. However, the deviations from the prescribed types are often minor, and we believe that a better strategy than throwing away all type information is to preserve as much of it as possible. We describe a system of untagged *union types* that can accommodate variations in structure while still allowing a degree of static type checking.

A novelty of this system is that it involves non-trivial equivalences among types, arising from a law of distributivity for records and unions: a value may be introduced with one type (e.g., a record containing a union) and used at another type (a union of records). We describe programming and query language constructs for dealing with such types, prove the soundness of the type system, and develop algorithms for subtyping and typechecking.

1 Introduction

Although semistructured data has, by definition, no schema, there are many cases in which the data obviously possesses some structure, even if it has mild deviations from that structure. Moreover it typically has this structure because it is derived from sources that have structure. In the process of annotating data or combining data from different sources one needs to accommodate the irregularities that are introduced by these processes. Because there is no way of describing “mildly irregular” structure, current approaches start by ignoring the structure completely, treating the data as some dynamically typed object such as a labelled graph and then, perhaps, attempting to recover some structure by a variety of pattern matching and data mining techniques [NAM97, Ali99]. The purpose of this structure recovery is typically to provide optimization techniques for query evaluation or efficient storage structures, and it is partial. It is not intended as a technique for preserving the integrity of data or for any kind of static type-checking of applications.

When data originates from some structured source, it is desirable to preserve that structure if at all possible. The typical cases in which one cannot require rigid conformance to a schema arise when one wants to annotate or modify the database with unanticipated structure or when one merges two databases with slight differences in structure. Rather than forgetting the original type and resorting to a completely dynamically type, we believe a more disciplined approach to maintaining type information is appropriate. We propose here a type system that can “degrade” gracefully if sources are added with variations in structure, while preserving the common structure of the sources where it exists.

The advantages of this approach include:

- The ability to check the correctness of programs and queries on semistructured data. Current semistructured query languages [AQM⁺96, BDHS96, DFF⁺] have no way of providing type errors – they typically return the empty answer on data whose type does not conform to the type assumed by the query.
- The ability to create data at one type and query it at another (equivalent) type. This is a natural consequence of using a flexible type system for semistructured data.
- New query language constructs that permit the efficient implementation of “case” expressions and increase the expressive power of a OQL-style query languages.

As an example, biological databases often have a structure that can be expressed naturally using a combination of tuples, records, and collection types. They are typically cast in special-purpose data formats, and there are groups of related databases, each expressed in some format that is a mild variation on some original format. These formats have an intended type, which could be expressed in a number of notations. For example a source (*source*₁) could have type

```
set[ id: Int,
     description: Str,
     bibl: set[ title: Str, authors: list[name: Str, address: Str], year: Int... ],
     ... ]
```

A second source (*source*₂) might yield a closely related structure:

```
set[ id: Int,
     description: Str,
     bibl: set[ title: Str, authors: list[fn: Str, ln: Str, address: Str], year: Int ... ],
     ... ]
```

This differs only in the way in which author names are represented. (This example is fictional, but not far removed from what happens in practice.)

The usual solution to this problem in conventional programming languages is to represent the union of the sources using some form of tagged union type:

```
set(( < tag1 : [ id: Int, ... ], tag2 : [ id: Int, ... ] )).
```

The difficulty with this solution is that a program such as

```
for each x in source1 do print(x.description)
```

(1)

that worked on *source*₁ must now be modified to

```
foreach x in source1 union source2 do
  case x of
    < tag1 = y1 > ⇒ print(y1.description)
  | < tag2 = y2 > ⇒ print(y2.description)
```

in order to work on the union of the sources, even though the two branches of the case statement contain identical code! This is also true for the (few) database query languages that deal with tagged union types [BLS⁺94].

Contrast this with a typical semi-structured query:

$$\begin{array}{l} \text{select } [\textit{description} = d, \textit{title} = t] \\ \text{where } [\textit{description} = d, \textit{bibl} = [\textit{Title} = t]] \leftarrow \textit{source}_1 \end{array} \quad (2)$$

This query works by pattern matching based on the (dynamically determined) structure of the data. Thus the *same* query works equally well against either of the two sources, and hence also against their union. The drawback of this approach, however, is that incorrect queries – for example, queries that use a field that does not exist in *either* source – yield the empty set rather than an error.

In this paper we define a system that combines the advantages of both approaches, based on a system of type-safe *untagged* union types. As a first example, consider the two forms of the *author* field in the types above. We may write the union of these types as:

$$[\textit{name} : \textit{Str}, \textit{address} : \textit{Str}] \vee [\textit{ln} : \textit{Str}, \textit{fn} : \textit{Str}, \textit{address} : \textit{Str}]$$

It is intuitively obvious that an *address* can always be extracted from a value of such a type. To achieve this formally, we begin by expressing a multi-field record type $[l_1 : T_1, l_2 : T_2, \dots]$ as a product of single-field record types: $[l_1 : T_1] \times [l_2 : T_2] \times \dots$. In this more basic form, the union type above is:

$$([\textit{name} : \textit{Str}] \times [\textit{address} : \textit{Str}]) \vee ([\textit{ln} : \textit{Str}] \times [\textit{fn} : \textit{Str}] \times [\textit{address} : \textit{Str}])$$

We now invoke a *distributivity* law that allows us to treat

$$[a : T_a] \vee ([b : T_b] \times [c : T_c]) \quad \text{and} \quad ([a : T_a] \times [b : T_b]) \vee ([a : T_a] \times [c : T_c])$$

as equivalent types. Using this, the union type above rewrites to:

$$([\textit{name} : \textit{Str}] \vee [\textit{fn} : \textit{Str} \times \textit{ln} : \textit{Str}]) \times [\textit{address} : \textit{Str}]$$

In this form, it is evident that the the selection of the *address* field is an allowable operation.

Type-equivalences like this distributivity rule allow us to introduce a value at one type and operate on it another type. Under this system both the program (1) and the query (2) above will type-check when extended to the union of the two sources. On the other hand, queries that reference a field that is not in either source will fail to type check.

Some care is needed in designing the operations for manipulating values of union types. Usually, the interrogation operation for records is field selection and the corresponding operation for unions is a case expression. However it is not enough simply to use these two operations. Consider the type $([a_1 : T_1] \vee [b_1 : U_1]) \times \dots \times ([a_n : T_n] \vee [b_n : U_n])$. The form of this type warrants neither selecting a field nor using a case expression. We can, if we want, use distributivity to rewrite it into a disjunct of products, but the size of this disjunct is exponential in n and so, presumably, would be the corresponding case expression. We propose, instead, an extended pattern matching syntax that allows us to operate on the type in its original, compact, form.

More sophisticated pattern matching operations may be useful additions even to existing semistructured query languages. Consider the problem of writing a query that produces a uniform output from a single source that contains two representations of names:

$$(\text{select } [\textit{description} = d, \textit{name} = n]$$

```

    where [ description = d, bibl = [ author = [ name = n ] ] ] ← source )
union
( select [ description = d, name = string-concat(f, l) ]
  where [ description = d, bibl = [ author = [ ln = l, fn = f ] ] ] ← source )

```

This is the only method known to the authors of expressing this query in current semistructured query languages. It suggests an inefficient execution model and may not have the intended semantics when, for example, the source is a list and one wants to preserve the order. Thus some enhancement to the syntax is desirable.

This paper develops a type system based on untagged union types along with operations to construct and deconstruct these types. In particular, we define a syntax of patterns that may be used both for an extended form of case expression and as an extension to existing query languages for semi-structured data. We should remark that we cannot capture all aspects of semistructured query languages. For example, we have nothing that corresponds to “regular path expressions” [BDHS96, AQM⁺96]. However, we believe that for most examples of “mildly” semistructured data – especially the forms that arise from the integration of typed data sources – a language such as proposed here will be adequate. Our main technical contribution is a proof of the decidability of subtyping for this type system (which is complicated by the non-trivial equivalences involving union and record types).

To our knowledge, untagged union types never been formalized in the context of database programming languages. *Tagged* union types have been suggested in several papers on data models [AH87, CM94] but have had minimal impact on the design of query languages. CPL [BLS⁺94], for example, can match on only one tag of a tagged union, and this is one of the few languages that makes use of union types. Pattern matching has been recently exploited in languages for semi-structured data and XML [BDHS96, DFF⁺]. In the programming languages and type theory communities, on the other hand, untagged union types have been studied extensively from a theoretical perspective [Pie91, BDCd95, Hay91, Dam94, DCdP96, etc.], but the interactions of unions with higher-order function types have been shown to lead to significant complexities; the present system, on the other hand, provides only a very limited form of function types (like most database query languages), and remains reasonably straightforward. .

Section 2 develops our language for programming with record and union types, including pattern matching primitives that can be used in both case expressions and query languages. Section 3 describes the system formally and demonstrates the decidability of subtyping and type equivalence. Proofs will be provided in the full paper. Section 4 offers concluding remarks.

2 Programming with Union Types

In this section we shall develop a syntax for the new programming constructs that are needed to deal with union types. The presentation is informal for the moment – more precise definitions appear in Section 3. We start with operations on records and extend these to work with unions of records; we then deal with operations on sets. Taken in conjunction with operations on records, these operations are enough to define a simple query language. We also look at operations on more general union types and give examples of a “typecase” operation.

2.1 Record formation

Just as we defined a record type $[l_1 : T_1, \dots, l_n : T_n]$ as the product $[l_1 : T_1] \times \dots \times [l_n : T_n]$ of elementary or “singleton” record types, we define a record value as the disjoint concatenation of singleton records. The operations for creating records are the empty record $[]$, the singleton record $[l = e]$ (where e is an expression), and the disjoint concatenation of records $e \# e$. We use $[l_1 = e_1, \dots, l_n = e_n]$ as shorthand for $[l_1 = e_1] \# \dots \# [l_n = e_n]$.

Unlike typed systems with tagged unions, in our system there is no basic formation operation directly associated with the union type. However we may want to introduce operators such as “relaxed union,” which takes two sets of type $\text{set}(t_1)$ and $\text{set}(t_2)$ and return a set of type $\text{set}(t_1 \vee t_2)$.

2.2 Case expressions

Records are decomposed through the use of case expressions. These allow us to take alternative actions based on the structure of values. We shall also be able to use components of the syntax of case expressions in the development of matching constructs for query languages. The idea in developing a relatively complex syntax for the body of case expressions is that the structure of the body can be made to match the expected structure of the type of the value on which it is operating. There should be no need to “flatten” the type into disjunctive normal form and write a much larger case expression at that type.

We start with a simple example:

$$\text{case } e \text{ of } \begin{array}{l} [fn = f:Str, ln = l:Str] \Rightarrow \text{string-concat}(f, l) \\ \quad | \quad [name = n:Str] \Rightarrow n \end{array}$$

This matches the result r of evaluating e to one of two record types. If r is a record with fn and ln fields, the variables f and l are bound and the right-hand side of the first clause is evaluated. If the first pattern does not match, the second clause is tried. This case expression will work provided r has type $[fn: Str, ln: Str] \vee [name: Str]$.

We should note that pattern matching introduces identifiers such as l, f, n in this example, and we shall make a short-sighted assumption that identifiers are introduced when they are associated with a type $(x : T)$. This ignores the possibility of type inference. See [BLS⁺94] for a more sophisticated syntax for introducing identifiers in patterns.

Field selection is given by a one-clause case expression: $\text{case } e \text{ of } [l = x:T] \Rightarrow x$.

We shall also allow case expressions to dispatch on the run-time type of an argument:

$$\text{case } e \text{ of } \begin{array}{l} x:Int \Rightarrow x \\ \quad | \quad y:\text{set}(Int) \Rightarrow \text{sum}(y) \end{array}$$

This will typecheck when $e : Int \vee \text{set}(Int)$

The clauses of a case expression have the form $p \Rightarrow e$, where p is a *pattern* that introduces (binds) identifiers which may occur free in the expression e . Thus each clause defines a *function*. Two or more functions can be combined by writing $p_1 \Rightarrow e_1 \mid p_2 \Rightarrow e_2 \mid \dots$ to form another function. The effect of the case expression $\text{case } e \text{ of } f$ is to apply this function to the result of evaluating e .

Now suppose we want to extract information from a value of type

$$([name: Str] \vee [ln: Str, fn: Str]) \times [age: Int] \tag{2}$$

The *age* field may be extracted using field selection using a one-clause case expression as described above. However information from the left-hand component cannot be extracted by extending this case expression. What we need is need something that will turn a multi-clause function back into a pattern that binds a new identifier. We propose the syntax $x : T \text{ as } f$, in which f is a multi-clause function. In the evaluation of $x \text{ as } f$, f is applied to the appropriate structure and x is bound to the result.

$$\begin{array}{l} \text{case } e \text{ of} \\ \quad x \text{ as } ([fn = f:Str, ln = l:Str] \Rightarrow \text{string-concat}(f, l) \mid [name = n:Str] \Rightarrow n) \\ \quad \# [name = a:Int] \\ \quad \Rightarrow [name = x, age = a + 1] \end{array}$$

could be applied to an expression e of type (2) above. Note the use of $\#$ to combine two patterns so that they match on a product type. This symbol is used to concatenate patterns in the same way that it is used to concatenate record values.

2.3 Sets

We shall follow the approach to collection types given in [BNTW95]. It is known that both relational and complex-object languages can be expressed using this formalism. The operations for forming sets are: $\{\}$ (empty set), $\{e\}$ (singleton set) and $e \text{ union } e$ (set union). For “iterating” over a set we use the form

collect e where $p \leftarrow e'$.

Here, e and e' are both expressions of set type, and p is a pattern as described above. The meaning of this is (informally) $\bigcup \{\sigma(e) \mid \sigma(p) \in e'\}$, in which σ is a substitution that binds the variables of p to match an element of e' .

These operations, taken in conjunction with the record operations described above and an equality operation, may be used as the basis of practical query languages. Conditionals and booleans may be added, but they can also be simulated with case expressions and some appropriately chosen constants.

2.4 Examples

We conclude this section with some remarks on high-level query languages. A typical form of a query that makes use of pattern matching is:

```
select e
where  $p_1 \leftarrow e_1$ ,
       $p_2 \leftarrow e_2$ ,
      ...
      condition
```

Here the p_i are patterns and the expressions $e_1, \dots, e_i$ have set types. Variables introduced in pattern p_i may be used in expression e_j and (as constants) in pattern p_j where $j > i$. They may also be used in the expression e and the *condition*, which is simply a boolean expression. This query form can be implemented using the operations described in the previous section.

As an example, here is a query based on the example types in the introduction. We make use of the syntax of patterns as developed for case expressions, but here we are using them to match on elements of one or more input sets.

```
select [description = d, authName = a, address = s]
where [description = d:Str, bibl = b:BT] ← source1 union source2,
      [authors = aa:AT, year = y:Int] ← b,
      a as ([fn = f:Str, ln = l:Str] ⇒ string-concat(f, l) | [name = n:Str] ⇒ n) ← aa,
      y > 1991
```

Note that we have assumed a “relaxed” union to combine the two sources. In the interests of consistency with the formal development, we have also inserted all types for identifiers, so *AT* and *BT* are names for the appropriate fragments of the expected source type. In many cases such types can be inferred.

Here are two examples that show the use of patterns in matching on types rather than record structures. Examples of this kind are commonly used to illustrate the need for semistructured data.

```

select x
where x as (s : set(Num) ⇒ average(s) | r : Num ⇒ r) ← source

select s
where s as (n : Str ⇒ n | [fn = f:Str, ln = l:Str] ⇒ string-concat(f, l)) ← source'

```

In the first case we have a set *source* that may contain both numbers and sets of numbers. In the second case we have a set that may contain both base types and record types. Both of these can be statically type-checked. If, for example, in the first query, *s* has type `set(Str)`, the query would not type-check.

To demonstrate the proposed syntax for the use of functions in patterns, here is one last (slightly contrived) example. We want to calculate the mass of a solid object that is either rectangular or a sphere. Each measure of length can be either integer or real. The type is

```

[ density: Real ]
×
( [ intRadius: Int ] ∨ [ realRadius: Real ]
  ∨
  ( ([ intHeight: Int ] ∨ [ realHeight: Real ])
    ×
    ([ intWidth: Int ] ∨ [ realWidth: Real ])
    ×
    ([ intDepth: Int ] ∨ [ realDepth: Real ]) ) )

```

The following case expression makes use of matching based on both unions and products of record structures. Note that the structure of the expression follows that of the type. It would be possible to write an equivalent case expression for the disjunctive normal form for the type and avoid the use of the form *x asf*, but such an expression would be much larger than the one given here.

```

case e of
  [ density = d:Real ]
  #
  v as
    ( r as ([ intRadius = ir:Int ] ⇒ float(ir) | [ realRadius = rr:Real ] ⇒ rr)
      ⇒ r**3)
    |
    ( h as ([ intHeight = ih:Int ] ⇒ float(ih) | [ realHeight = rh:Real ] ⇒ rh)
      #
      w as ([ intWidth = iw:Int ] ⇒ float(iw) | [ realWidth = rw:Real ] ⇒ rw)
      #
      d as ([ intDepth = id:Int ] ⇒ float(id) | [ realDepth = rd:Real ] ⇒ rd)
        ⇒ h * w * d)
    ⇒ d * v

```

3 Formal Development

With the foregoing intuitions and examples in mind, we now proceed to the formal definition of our language, its type system, and its operational semantics. Along the way, we establish fundamental properties such as run-time safety and the decidability of subtyping and type-checking.

3.1 Types

We develop a type system that is based on conventional complex object types, those that are constructed from the base types with record (tuple) and set constructors. As described in the introduction, the record constructors are $[]$, the empty record type, $[l : t]$, the singleton record type, and $R \times R$, the *disjoint* concatenation of two records types. (By disjoint we mean that the two record types have no field names in common.) Thus a conventional record type $[l_1 : T_1, \dots, l_n : T_n]$ is shorthand for $[l_1 : Y_1] \times \dots \times [l_n : T_n]$. To this we add an untagged union type $T \vee T$. We also assume a single base type B and a set type $\text{set}(T)$. Other collection types such as lists and multisets would behave similarly,

The syntax of types is described by the following grammar:

$T ::= B$	base type
$[]$	empty record type
$[l : T]$	labeling (single-field record type)
$T_1 \times T_2$	record type concatenation
$T_1 \vee T_2$	union type
$\text{set}(T)$	set type

3.2 Kinding

We have already noted that certain operations on types are restricted. For example, we cannot take the product of two record types with a common field name. This in turn means that any operation on records whose typing rules make improper use of a type constructor is also illegal. In order to control the formation of types we introduce a system of *kinds*. This consists of the kind of all types, Type , and a subkind $\text{Rcd}(L)$, which is the kind of all record types whose labels are included in the label set L .

$K ::= \text{Type}$	kind of all types
$\text{Rcd}(L)$	kind of record types with (at most) labels L

The kinding relation is defined as follows:

$B \in \text{Type}$	(K-BASE)
$[] \in \text{Rcd}(\{\})$	(K-EMPTY)
$\frac{T \in \text{Type}}{[l : T] \in \text{Rcd}(\{l\})}$	(K-FIELD)
$\frac{S \in \text{Rcd}(L_1) \quad T \in \text{Rcd}(L_2) \quad L_1 \cap L_2 = \emptyset}{S \times T \in \text{Rcd}(L_1 \cup L_2)}$	(K-RCD)
$\frac{S \in K \quad T \in K}{S \vee T \in K}$	(K-UNION)
$\frac{T \in \text{Type}}{\text{set}(T) \in \text{Type}}$	(K-SET)
$\frac{T \in \text{Rcd}(L_1)}{T \in \text{Rcd}(L_1 \cup L_2)}$	(K-SUBSUMPTION-1)
$\frac{T \in \text{Rcd}(L)}{T \in \text{Type}}$	(K-SUBSUMPTION-2)

There are two important consequences of these rules. First, record kinds extend to the union type. For example, $([A : t] \times [B : t]) \times ([C : t] \vee [D : t])$ has kind $\text{Rcd}(\{A, B, C, D\})$. Second, the kinding rules require the labels in a concatenation of two record types to be disjoint. (However the union type constructor is not limited in the same way; $\text{Int} \vee \text{Str}$ and $\text{Int} \vee [a : \text{Str}]$ are well-kinded types.)

3.3 Subtyping

As usual, the subtype relation written $S <: T$ captures a principle of “safe substitutibility”: any element of S may safely be used in a context expecting an element of T .

For sets and records, the subtyping rules are the standard ones: $\text{set}(S) <: \text{set}(T)$ if $S <: T$ (e.g., a set of employees can be used as a set of people), and a record type S is a subtype of a record type T if S has more fields than T and the types of the common fields in S are subtypes of the corresponding fields in T . This effect is actually achieved by the combination of several rules below. This “exploded presentation” of record subtyping corresponds to our presentation of record types in terms of separate empty set, singleton, and concatenation constructors.

For union types, the subtyping rules are a little more interesting. First, we axiomatize the fact that $S \vee T$ is the least upper bound of S and T – that is, $S \vee T$ is above both S and T , and everything that is above both S and T is also above their union (rules S-UNION-UB and S-UNION-L below). We then have two rules (S-DIST-RCD and S-DIST-FIELD) showing how union distributes over records.

Formally, the subtype relation is the least relation on well-kinded types closed under the following rules.

$$\begin{array}{ll}
T <: T & (\text{S-REFL}) \\
\\
\frac{R <: S \quad S <: T}{R <: T} & (\text{S-TRANS}) \\
\\
[l : T] <: [] & (\text{S-RCD-FE}) \\
\\
S \times T <: S & (\text{S-RCD-RE}) \\
\\
S \times T <: T \times S & (\text{S-RCD-COMM}) \\
\\
S \times (T \times U) <: (S \times T) \times U & (\text{S-RCD-ASSOC}) \\
\\
S <: S \times [] & (\text{S-RCD-IDENT}) \\
\\
\frac{S <: T}{[l : S] <: [l : T]} & (\text{S-RCD-DF}) \\
\\
\frac{S_1 <: T_1 \quad S_2 <: T_2}{S_1 \times S_2 <: T_1 \times T_2} & (\text{S-RCD-DR}) \\
\\
\frac{S <: T}{\text{set}(S) <: \text{set}(T)} & (\text{S-SET}) \\
\\
\frac{R <: T \quad S <: T}{R \vee S <: T} & (\text{S-UNION-L}) \\
\\
S <: S \vee T & (\text{S-UNION-UB})
\end{array}$$

$$R \times (S \vee T) <: (R \times S) \vee (R \times T) \quad (\text{S-DIST-RCD})$$

$$[l : S \vee T] <: [l : S] \vee [l : T] \quad (\text{S-DIST-FIELD})$$

Note that we restrict the subtype relation to *well-kinded* types: S is never a subtype of T if either S or T is ill-kinded. (The typing rules will be careful only to “call” the subtype relation on types that are already known to be well kinded.)

If both $S <: T$ and $T <: S$, we say that S and T are *equivalent* and write $S \sim T$. Note, for example, that the distributive laws S-DIST-RCD and S-DIST-FIELD are actually equivalences: the other directions follow from the laws for union (plus transitivity). Also, note the absence of the “other” distributivity law for unions and records: $P \vee (Q \times R) \sim (P \vee Q) \times (P \vee R)$. This law doesn’t make sense here, because it violates the kinding constraint that products of record types can only be formed if the two types have disjoint label sets.

The subtype relation includes explicit rules for associativity and commutativity of the operator $\times$. Also, it is easy to check that the associativity, commutativity and idempotence of $\vee$ follow directly from the rules given. We shall take advantage of this fluidity in the following by writing both records and unions in a compound, n -ary form:

$$\begin{aligned} [l_1 : T_1, \dots, l_n : T_n] &\stackrel{\text{def}}{=} [l_1 : T_1] \times \dots \times [l_n : T_n] \\ \vee(T_1, \dots, T_n) &\stackrel{\text{def}}{=} T_1 \vee \dots \vee T_n \end{aligned}$$

We often write compound unions using a simple comprehension notation. For example,

$$\bigvee (A \times B \mid A \in A_1 \vee A_2 \vee \dots \vee A_m \text{ and } B \in B_1 \vee B_2 \vee \dots \vee B_n)$$

denotes

$$A_1 \times B_1 \vee A_1 \times B_2 \vee \dots \vee A_1 \times B_n \vee A_2 \times B_1 \vee \dots \vee A_m \times B_n.$$

3.4 Properties of Subtyping

For proving properties of the subtype relation, it is convenient to work with types in a more constrained syntactic form:

3.4.1 Definition: The sets of *normal* (N) and *simple* (A) types are defined as follows:

$$\begin{aligned} N &::= \bigvee (A_1, \dots, A_n) \\ A &::= B \\ &\quad [l_1 : A_1, \dots, l_n : A_n] \\ &\quad \text{set}(N) \end{aligned}$$

Intuitively, a simple type is one in which unions only appear (immediately) inside of the `set` constructor; a normal type is a union of simple types. Note that every simple type is also normal. $\square$

The restricted form of normal and simple types can be exploited to give a much simpler subtyping relation, written $S \leq T$, in terms of the following “macro rules”:

$$B \leq B \quad (\text{SA-BASE})$$

$$\frac{N \leq M}{\text{set}(N) \leq \text{set}(M)} \quad (\text{SA-SET})$$

$$\frac{\{k_1, \dots, k_m\} \subseteq \{l_1, \dots, l_n\} \quad \text{for all } k_i \in \{k_1, \dots, k_m\}, A_{k_i} \leq B_{k_i}}{[l_1 : A_{l_1}, \dots, l_m : A_{l_m}] \leq [k_1 : B_{k_1}, \dots, k_m : B_{k_m}]} \quad (\text{SA-RCD})$$

$$\frac{\forall i \leq m. \exists j \leq n. A_i \leq B_j}{\bigvee(A_1, \dots, A_m) \leq \bigvee(B_1, \dots, B_n)} \quad (\text{SN-UNION})$$

3.4.2 Fact: $N \leq M$ is decidable. □

Proof: The macro rules can be read as a pair of algorithms, one for subtyping between simple types and one for subtyping between normal types. Both of these algorithms are syntax directed and obviously terminate on all inputs (all recursive calls reduce the size of the inputs). □

3.4.3 Lemma: $N \leq N$, for all N . □

3.4.4 Lemma: If $N \leq M$ and $M \leq L$ then $N \leq L$. □

Proof: By induction on the total size of L, M, N . First suppose that all of L, M, N are simple. The induction hypothesis is immediately satisfied for SA-BASE and SA-SET. For SA-RCD use the transitivity of set inclusion and induction on the appropriate subterms.

If at least one of L, M, N is (non-trivially) normal, use the transitivity of the functional relationship expressed by the SN-Union rule and induction on the appropriate subterms. □

To transfer the property of decidability from $\leq$ to $<$, we first show how any type may be converted to an equivalent type in *disjunctive normal form*.

3.4.5 Definition: The disjunctive normal form (dnf) of a type T is defined as follows:

$$\begin{aligned} \text{dnf}(\mathbf{B}) &= \mathbf{B} \\ \text{dnf}([\]) &= [\] \\ \text{dnf}(P \times Q) &= \bigvee(A_i \times B_j \mid A_i \in \text{dnf}(P), B_j \in \text{dnf}(Q)) & (a) \\ \text{dnf}([l : P]) &= \bigvee([l : A_i] \mid A_i \in \text{dnf}(P)) & (b) \\ \text{dnf}(P \vee Q) &= \text{dnf}(P) \vee \text{dnf}(Q) & (c) \\ \text{dnf}(\text{set}(P)) &= \text{set}(\text{dnf}(P)) & (d) \end{aligned}$$

3.4.6 Fact: $\text{dnf}(P) \sim P$. □

3.4.7 Fact: $\text{dnf}(P)$ is a normal type, for every type P . □

3.4.8 Fact: $N \leq M$ implies $N < M$ □

3.4.9 Lemma: $S < T$ iff $\text{dnf}(S) \leq \text{dnf}(T)$ □

Proof: ($\Leftarrow$) By 3.4.6 we have derivations of $S < \text{dnf}(S)$ and $\text{dnf}(T) < T$, and by 3.4.8 we have a derivation of $\text{dnf}(S) < \text{dnf}(T)$. Use transitivity to build a derivation of $S < T$.

($\Rightarrow$) By induction on the height of the derivation of $S < T$. We consider the final rule in the derivation. By induction we assume we can build a derivation of the normal forms for the antecedents, and now we consider all possible final rules.

We start with the axioms.

(S-REFL) By reflexivity of $\leq$: (3.4.3).

(S-RCD-FE) $dnf([l : T]) = [l : dnf(T)]$, and $[l : dnf(T)] \leq []$ by SA-RCD.

(S-RCD-RE) $dnf(S \times T) = \bigvee (S_i \times T_j \mid S_i \in dnf(S), T_j \in dnf(T))$. Now $dnf(S_i) \times dnf(T_j) \leq dnf(S_i)$ by SA-RCD, and the result follows from SN-UNION.

(S-RCD-COMM) If $dnf(S)$ and $dnf(T)$ are simple then $dnf(S) \times dnf(T) \leq dnf(T) \times dnf(S)$ by SA-RCD. If not, use SN-UNION first.

(S-RCD-ASSOC) As for S-RCD-COMM.

(S-RCD-IDENT) As for S-RCD-COMM.

(S-UNION-UB) By SN-UNION.

$$\begin{aligned}
 \text{(S-DIST-RCD)} \quad dnf(R \times (S \vee T)) &= \bigvee (R_i \times U_j \mid R_i \in dnf(R), U_j \in dnf(S \vee T)) \\
 &= \bigvee (R_i \times S_j \mid R_i \in dnf(R), U_j \in dnf(S)) \vee \\
 &\quad \bigvee (R_i \times T_k \mid R_i \in dnf(R), T_k \in dnf(T)) \\
 &= dnf((R \times S) \vee (R \times T))
 \end{aligned}$$

$$\begin{aligned}
 \text{(S-DIST-FIELD)} \quad dnf([l : S \vee T]) &= \bigvee ([l : U_i] \mid U_i \in dnf(S \vee T)) \\
 &= \bigvee ([l : S_i] \mid S_i \in dnf(S)) \vee \bigvee ([l : T_i] \mid T_i \in dnf(T)) \\
 &= dnf([l : S]) \vee dnf([l : T]) \\
 &= dnf([l : S] \vee [l : T])
 \end{aligned}$$

Now for the inference rules. The premises for all the rules are of the form $S <: T$ and our inductive hypothesis is that for the premises of the final rule we have obtained a derivation using SA-* and SN-UNION rules of the corresponding $dnf(S) \leq dnf(T)$. Without loss of generality we may assume that the final rule in the derivation of each such premise is SN-UNION. We examine the remaining inference rules.

(S-TRANS) By Lemma 3.4.4.

(S-RCD-DF) Since $dnf(S) \leq dnf(T)$ was derived by SN-UNION we know that for each $A_i \in dnf(S)$ there is a $B_j \in dnf(T)$ such that $A_i \leq B_j$. Therefore, for each such A_i , we may use SA-RCD to derive $[l : A_i] \leq [l : B_j]$. These derivations may be combined using SN-UNION to obtain a derivation of $dnf([l : S]) \leq dnf([l : T])$.

(S-RCD-DR) For each $A_{i_1}^1 \in dnf(S_1)$ and each $A_{i_2}^2 \in dnf(S_2)$ there exist $B_{j_1}^1 \in dnf(T_1)$ and $B_{j_2}^2 \in dnf(T_2)$ such that we have a derivations of $A_{i_1}^1 \leq B_{j_1}^1$ and $A_{i_2}^2 \leq B_{j_2}^2$. For each such pair we can therefore use SA-RCD to derive $A_{i_1}^1 \times A_{i_2}^2 \leq B_{j_1}^1 \times B_{j_2}^2$ and then use SN-UNION to derive $dnf(S_1 \times S_2) \leq dnf(T_1 \times T_2)$.

(S-SET) Immediate, by SA-SET.

(S-UNION-L) For each $A_i \in dnf(R)$ there is a $C_j \in dnf(T)$ such that $A_i \leq C_j$ and for each $B_k \in dnf(S)$ there is a $C_l \in dnf(T)$ such that $B_k \leq C_l$. From these $dnf(R \vee S) \leq dnf(T)$ can be derived directly using SN-UNION. $\square$

3.4.10 Theorem: The subtype relation is decidable. $\square$

Proof: Immediate from Lemmas 3.4.9 and 3.4.2. $\square$

We do not yet have any results on the *complexity* of checking subtyping (or equivalence). (The proof strategy we have adopted here leads to an algorithm with running time exponential in the size of its inputs.)

The structured form of the macro rules can be used to derive several *inversion properties*, which will be useful later in reasoning about the typing relation.

3.4.11 Corollary: If $S <: \text{set}(T_1)$, then $S = \text{set}(S_1)$, with $S_1 <: T_1$. □

3.4.12 Corollary: If $W <: U$, with

$$\begin{aligned} W &= [l_1 : W_1] \times \dots \times [l_m : W_m] \times \dots \times [l_n : W_n] \\ U &= [l_1 : U_1] \times \dots \times [l_m : U_m], \end{aligned}$$

then $W_k <: U_k$ for each $k \leq m$. □

Proof: From the definition of disjunctive normal forms, we know that $\text{dnf}(W) = \bigvee([l_1 : W_{i1}] \times \dots \times [l_m : W_{im}] \times \dots \times [l_n : W_{in}] \mid W_{i1} \dots W_{im} \dots W_{in} \in \text{dnf}(W_1) \dots \text{dnf}(W_m) \dots \text{dnf}(W_n))$ and $\text{dnf}(U) = \bigvee([l_1 : U_{j1}] \times \dots \times [l_m : U_{jm}] \mid U_{j1} \dots U_{jm} \in \text{dnf}(U_1) \dots \text{dnf}(U_m))$. By SN-UNION,

$$\begin{aligned} \text{for each } A_i &= [l_1 : W_{i1}] \times \dots \times [l_m : W_{im}] \times \dots \times [l_n : W_{in}] \in \text{dnf}(W) \\ \text{there is some } B_j &= [l_1 : U_{j1}] \times \dots \times [l_m : U_{jm}] \in \text{dnf}(U) \\ \text{with } A_i &\leq B_j. \end{aligned}$$

This derivation must be an instance of SM-RCD, with $W_{ik} \leq U_{jk}$. In other words, for each $W_{ik} \in \text{dnf}(W_k)$ there is some $U_{jk} \in \text{dnf}(U_k)$ with $W_{ik} \leq U_{jk}$. By SN-UNION, $\text{dnf}(W_k) \leq \text{dnf}(U_k)$. The desired result, $W_k <: U_k$ now follows by Lemma 3.4.9. □

3.4.13 Corollary: If S is a simple type and $S <: T_1 \vee T_2$, then either $S <: T_1$ or else $S <: T_2$. □

3.5 Terms

The sets of programs, functions, and patterns are described by the following grammar:

$e ::= b$	base value
x	variable
$[l_1 = e_1, \dots, l_n = e_n]$	record construction
$e_1 \# e_2$	record concatenation
$\text{case } e \text{ of } f$	pattern matching
$\{e_1, \dots, e_n\}$	set
$e_1 \text{ union } e_2$	union of sets
$\text{collect } e_1 \text{ where } p \leftarrow e_2$	set comprehension
$p ::= x : T$	variable pattern (typecase)
$[l_1 = p_1, \dots, l_n = p_n]$	record pattern
$x \text{ as } f$	function nested in pattern
$f ::= p \Rightarrow e$	base function
$f_1 \mid f_2$	compound function

3.6 Typing

The typing rules are quite standard.

Expressions ($\Gamma \vdash e \in T$)

$$\Gamma \vdash b \in B$$

(T-BASE)

$$\begin{array}{c}
\Gamma \vdash x \in \Gamma(x) \quad (\text{T-VAR}) \\
\\
\frac{\Gamma \vdash e_i \in T_i \quad \text{all the } l_i \text{ are distinct}}{\Gamma \vdash [l_1 = e_1, \dots, l_n = e_n] \in [l_1 : T_1] \times \dots \times [l_n : T_n]} \quad (\text{T-RCD}) \\
\\
\frac{\Gamma \vdash e_1 \in T_1 \quad \Gamma \vdash e_2 \in T_2 \quad T_1 \times T_2 \in K}{\Gamma \vdash e_1 \# e_2 \in T_1 \times T_2} \quad (\text{T-CONCAT}) \\
\\
\frac{\Gamma \vdash f \in S \rightarrow T \quad \Gamma \vdash e \in R \quad R <: S}{\Gamma \vdash \text{case } e \text{ of } f \in T} \quad (\text{T-CASE}) \\
\\
\frac{\Gamma \vdash e_i \in T_i \quad \text{for each } i}{\Gamma \vdash \{e_1, \dots, e_n\} \in \text{set}(T_1 \vee \dots \vee T_n)} \quad (\text{T-SET}) \\
\\
\frac{\Gamma \vdash e_1 \in \text{set}(T_1) \quad \Gamma \vdash e_2 \in \text{set}(T_2)}{\Gamma \vdash e_1 \text{ union } e_2 \in \text{set}(T_1 \vee T_2)} \quad (\text{T-UNION}) \\
\\
\frac{\Gamma \vdash e_2 \in \text{set}(S) \quad \Gamma \vdash p \in U \Rightarrow \Gamma' \quad S <: U \quad \Gamma, \Gamma' \vdash e_1 \in \text{set}(T)}{\Gamma \vdash \text{collect } e_1 \text{ where } p \leftarrow e_2 \in \text{set}(T)} \quad (\text{T-COLLECT})
\end{array}$$

Functions ($\Gamma \vdash f \in S \rightarrow T$)

$$\begin{array}{c}
\frac{\Gamma \vdash p \in S \Rightarrow \Gamma' \quad \Gamma, \Gamma' \vdash e \in T}{\Gamma \vdash p \Rightarrow e \in S \rightarrow T} \quad (\text{TF-PAT}) \\
\\
\frac{\Gamma \vdash f_1 \in S_1 \rightarrow T_1 \quad \Gamma \vdash f_2 \in S_2 \rightarrow T_2}{\Gamma \vdash f_1 \Rightarrow f_2 \in S_1 \vee S_2 \rightarrow T_1 \vee T_2} \quad (\text{TF-ALT})
\end{array}$$

Patterns ($\Gamma \vdash p \in T \Rightarrow \Gamma'$)

$$\begin{array}{c}
\frac{T \in K}{\Gamma \vdash x : T \in T \Rightarrow x : T} \quad (\text{TP-VAR}) \\
\\
\frac{\Gamma \vdash p_i \in T_i \Rightarrow \Gamma'_i \quad \text{the } \Gamma'_i \text{ all have disjoint domains}}{\Gamma \vdash [l_1 = p_1, \dots, l_n = p_n] \in [l_1 : T_1] \times \dots \times [l_n : T_n] \Rightarrow \Gamma'_1, \dots, \Gamma'_n} \quad (\text{TP-RCD}) \\
\\
\frac{\Gamma \vdash f \in S \rightarrow T}{\Gamma \vdash x \Rightarrow f \in S \Rightarrow x : T} \quad (\text{TP-AS})
\end{array}$$

3.7 Properties of Typing

3.7.1 Proposition: The typing relation is decidable. $\square$

Proof: Immediate from the decidability of subtyping and the syntax-directedness of the typing rules. $\square$

3.7.2 Definition: We say that a substitution σ *satisfies* a context Σ if they have the same domain and, for each x in their common domain, we have $\vdash \sigma(x) \in S_x$ for some S_x with $S_x <: \Sigma(x)$. $\square$

3.7.3 Definition: We say that a typing context Γ *refines* another context Γ' , written $\Gamma <: \Gamma'$, if their domains are the same and, for each $x \in \text{dom}(\Gamma)$, we have $\Gamma(x) <: \Gamma'(x)$. $\square$

3.7.4 Fact [Narrowing]: If $\Gamma \vdash e \in T$ and $\Gamma <: \Gamma'$, then $\Gamma' \vdash e \in T$. $\square$

3.7.5 Lemma [Substitution preserves typing]:

1. If $\Sigma \models \sigma$ and $\Sigma, \Delta \vdash e \in Q$ then $\Delta \vdash \sigma(e) \in P$, for some $P <: Q$.
2. If $\Sigma \models \sigma$ and $\Sigma, \Delta \vdash f \in S \rightarrow Q$ then $\Delta \vdash \sigma(f) \in S \rightarrow P$, for some $P <: Q$.
3. If $\Sigma \models \sigma$ and $\Sigma, \Delta \vdash p \in U \Rightarrow \Delta'$ then $\Delta \vdash \sigma(p) \in U \Rightarrow \Delta''$, for some $\Delta'' <: \Delta$. $\square$

Proof: By simultaneous induction on derivations. The arguments are all straightforward, using previously established facts. (For the second property, note that substitution into a pattern only affects functions that may be embedded in the pattern, since all other variables mentioned in the pattern are binding occurrences. Moreover, by our conventions about names of bound variables, we must assume that the variables bound in an expression, function, or pattern are distinct from those defined by σ .) $\square$

3.8 Evaluation

The operational semantics of our language is again quite standard: we define a relation $e \Downarrow v$, read “(closed) expression e evaluates to result v ,” by a collection of syntax-directed rules embodying a simple abstract machine.

3.8.1 Definition: We will use the metavariables v and w to range over *values* – closed expressions not involving case, union, concatenation, or collect. $\square$

$$v ::= \mathbf{b} \quad \begin{array}{l} [l_1 = v_1, \dots, l_n = v_n] \\ \{v_1, \dots, v_n\} \end{array}$$

We write $\bar{v}$ as shorthand for a set of values $v_1, \dots, v_n$. $\square$

3.8.2 Definition: A *substitution* σ is a finite function from variables to values. When σ_1 and σ_2 have disjoint domains, we write $\sigma_1 + \sigma_2$ for their combination. $\square$

Reduction ($e \Downarrow v$, for closed terms e)

$$\mathbf{b} \Downarrow \mathbf{b} \quad (\text{E-BASE})$$

$$\frac{e_i \Downarrow v_i}{[l_1 = e_1, \dots, l_n = e_n] \Downarrow [l_1 = v_1, \dots, l_n = v_n]} \quad (\text{E-RCD})$$

$$\frac{e_1 \Downarrow [l_1 = v_1, \dots, l_m = v_m] \quad e_2 \Downarrow [j_1 = w_1, \dots, j_n = w_n] \quad \{l_1, \dots, l_m\} \cap \{j_1, \dots, j_n\} = \emptyset}{e_1 \# e_2 \Downarrow [l_1 = v_1, \dots, l_m = v_m, j_1 = w_1, \dots, j_n = w_n]} \quad (\text{E-CONCAT})$$

$$\frac{e \Downarrow v \quad \text{match}(v, f) \Rightarrow v'}{\text{case } f \text{ of } e \Downarrow v'} \quad (\text{E-CASE})$$

$$\frac{e_i \Downarrow v_i}{\{e_1, \dots, e_n\} \Downarrow \{v_1, \dots, v_n\}} \quad (\text{E-SET})$$

$$\frac{e_1 \Downarrow \{\bar{v}_1\} \quad e_2 \Downarrow \{\bar{v}_2\}}{e_1 \text{ union } e_2 \Downarrow \{\bar{v}_1 \cup \bar{v}_2\}} \quad (\text{E-UNION})$$

$$\frac{\begin{array}{c} e_2 \Downarrow \{v_1, \dots, v_n\} \\ \text{for each } i, \quad \text{match}(v_i, p) \Rightarrow \sigma_i \text{ and } \sigma_i(e_1) \Downarrow \{\bar{w}_i\} \end{array}}{\text{collect } e_1 \text{ where } p \leftarrow e_2 \Downarrow \{\bar{w}_1 \cup \dots \cup \bar{w}_n\}} \quad (\text{E-COLLECT})$$

Function matching ($\text{match}(v, f) \Rightarrow v'$)

$$\frac{\text{match}(v, p) \Rightarrow \sigma \quad \sigma(e) \Downarrow v'}{\text{match}(v, p \Rightarrow e) \Rightarrow v'} \quad (\text{EF-PAT})$$

$$\frac{\text{match}(v, f_1) \Rightarrow v'}{\text{match}(v, f_1 \mid f_2) \Rightarrow v'} \quad (\text{EF-ALT1})$$

$$\frac{\neg(\text{match}(v, f_1)) \quad \text{match}(v, f_2) \Rightarrow v'}{\text{match}(v, f_1 \mid f_2) \Rightarrow v'} \quad (\text{EF-ALT2})$$

Matching ($\text{match}(v, p) \Rightarrow \Gamma$)

$$\frac{\vdash v \in T}{\text{match}(v, x : T) \Rightarrow x = v} \quad (\text{EP-VAR})$$

$$\frac{\text{match}(v_i, p_i) \Rightarrow \sigma_i \quad \text{the } \sigma_i \text{ have disjoint domains}}{\text{match}([l_1 = v_1, \dots, l_m = v_m, \dots, l_n = v_n], [l_1 = p_1, \dots, l_m = p_m]) \Rightarrow \sigma_1 + \dots + \sigma_m} \quad (\text{EP-RCD})$$

$$\frac{\text{match}(v, f) \Rightarrow v'}{\text{match}(v, x \Rightarrow f) \Rightarrow x = v'} \quad (\text{EP-AS})$$

3.9 Properties of Evaluation

3.9.1 Fact: If v is a value and $\vdash v \in V$, then V is a simple type. □

3.9.2 Theorem [Subject reduction]:

1. If $e \Downarrow v$
 $\vdash e \in Q$,
then $\vdash v \in V$
 $V <: Q$.
2. If $\text{match}(v, f) \Rightarrow v'$
 $\vdash f \in U \rightarrow V$
 $\vdash v \in W$
 $W <: U$,
then $\vdash v' \in X$
 $X <: V$.

3. If $match(v, p) \Rightarrow \sigma$
 $\vdash v \in W$
 $\vdash p \in U \Rightarrow \Sigma$
 $W <: U$,
then $\Sigma \models \sigma$. □

Proof: By simultaneous induction on evaluation derivations.

1. Straightforward, using part (2) of the induction hypothesis for the interesting case (E-CASE).
2. Consider the final rule in the given derivation.

Case EF-PAT: $f = p \Rightarrow e$
 $match(v, p) \Rightarrow \sigma$
 $\sigma(e) \Downarrow v'$

From $\vdash f \in U \rightarrow V$, we know $\vdash p \in U \Rightarrow \Sigma$ and $\Sigma \vdash e \in V$. By part (3) of the induction hypothesis, $\Sigma \models \sigma$. By Lemma 3.7.5, $\vdash \sigma(e) \in V$. Now, by the induction hypothesis, $\vdash v' \in X$ and $X <: V$, as required.

Case EF-ALT1: $f = f_1 \mid f_2$
 $match(v, f_1) \Rightarrow v'$

From rule TF-ALT, we see that $\vdash f_1 \in U_1 \rightarrow V_1$ and $\vdash f_2 \in U_2 \rightarrow V_2$, with $U = U_1 \vee U_2$ and $V = V_1 \vee V_2$. The induction hypothesis yields $\vdash v' \in X$ with $X <: V_1$, from which the result follows immediately by S-UNION.

Case EF-ALT2: $f = f_1 \mid f_2$
 $\neg(match(v, f_1))$
 $match(v, f_2) \Rightarrow v'$

Similar.

3. Consider the final rule in the given derivation.

Case EP-VAR: $p = (x : T)$
 $\vdash v \in T$
 $\sigma = (x = v)$
 $\Sigma = (x : T)$

Immediate.

Case EP-RCD: $v = [l_1 = v_1, \dots, l_m = v_m, \dots, l_n = v_n]$
 $p = [l_1 = p_1, \dots, l_m = p_m]$
 $match(v_i, p_i) \Rightarrow \sigma_i$
 $\sigma = \sigma_1 + \dots + \sigma + m$

From T-RCD, we have $W = [l_1 : W_1] \times \dots \times [l_n : W_n]$ and $\vdash v_i \in W_i$ for each i . Similarly, by TP-RCD, we have $U = [l_1 : U_1] \times \dots \times [l_m : U_m]$ with $\vdash p_i \in U_i \Rightarrow \Sigma_i$. Finally, by Corollary 3.4.12, we see that $W_i <: U_i$. Now, by the induction hypothesis, $\Sigma_i \models \sigma_i$ for each i . But this means that $\Sigma \models \sigma$, as required.

Case EP-As: $p = x \Rightarrow f$
 $match(v, f) \Rightarrow v'$
 $\sigma = (x = v')$

By TP-As, $\vdash f \in U \rightarrow V$ and $\Sigma = (x : V)$. By part (2) of the induction hypothesis, $\vdash v' \in X$ for some $X <: V$. So $(x = v') \models (x : V)$ by the definition of satisfaction. □

3.9.3 Theorem [Safety]:

1. If $\vdash e \in T$, then $e \Downarrow v$ for some v . (That is, the evaluation of a closed, well-typed expression cannot lead to a match-failure or otherwise “get stuck.”)
2. If $\vdash f \in S \rightarrow T$ and $\vdash v \in R \prec S$, then $\text{match}(v, f) \Rightarrow v'$ with $\vdash v' \in T' \prec T$.
3. If $\vdash p \in U \Rightarrow \Gamma'$ and $\vdash v \in S \prec U$, then $\text{match}(v, p) \Rightarrow \sigma$ with $\Gamma' \models \sigma$. □

Proof: Straightforward induction on derivations. □

4 Conclusions

We have described a type system that may be of use in checking programs or queries that apply to semistructured data. Unlike other approaches to the problem, it is a “relaxed” version of a conventional system that can handle the kinds of irregular types that occur in semistructured data.

Although we have established the basic properties of the type system, a good deal of work remains to be done. First, there are some extensions that we do not see as problematic. These include:

- Both strict and relaxed unions. (In the former case the two types are constrained to be equivalent.) Similarly, one can imagine strict and relaxed case expressions.
- Equality. Both “absolute” equality and “equality at type T ” fit with this scheme.
- A “top” type. Such a type would be completely dynamic and would be analyzed by typecase expressions. One could also add type inspection primitives along the lines described for Amber [Carrk]
- An “otherwise” or “fall-through” branch in case expressions.

A number of more significant problems also remain to be addressed.

- **Complexity.** The obvious method of checking whether two types are equivalent or whether one is a subtype of the other involves first reducing both to disjunctive normal form. As we have observed, this process may be exponential in the size of the two type expressions. We conjecture that equivalence (and subtyping) can be checked faster, but we have not been able to show this.
Even if these problems turn out to be intractable in general, it does not necessarily mean that this approach to typing semistructured data is pointless. Type inference in ML, for example, is known to be exponential [KTU94], yet the forms of ML programs that are the cause of this complexity never occur in practice. Here, it may be the case that types that only have “small” differences will not give rise to expensive transformations.
- **Recursive types.** The proof of the decidability of subtyping (3.4.10) works by induction on the derivation tree of a type, which is closely related to the structure of the type. We do not know whether the same result holds in the presence of recursive types.
- **Relationship with other typing schemes.** There may be some relationship between the typing scheme proposed here and those mentioned earlier [NAM97, Ali99] that work by inferring structure from semi-structured data. Simulation, for example, gives rise to something like a subtyping relationship [BDFS97]; but it is not clear what would give rise to union types.
- **Applications.** Finally, we would like to think that a system like this could be of practical benefit. We mentioned that there is a group of biological data formats that are all derived from a common basic format. We should also mention that the pattern matching constructs introduced in section 2.2, independently of any typing issues, might be used to augment other query languages such as XML-QL [DFF⁺] that exploit pattern matching.

5 Acknowledgements

The idea of using untagged unions to type semistructured data started when Peter Buneman was on a visiting research fellowship provided by the Japanese Society for the Promotion of Science. He is grateful to Atsushi Ohori for stimulating discussions. Benjamin Pierce is supported by the National Science Foundation under CAREER grant CCR-9701826.

References

- [AH87] Serge Abiteboul and Richard Hull. IFO: A formal semantic database model. *ACM Transactions on Database Systems*, 12(4):525–565, December 1987.
- [Ali99] Alin Deutsch and Mary Fernandez and Dan Suciu. Storing semistructured data with STORED. In *Proceedings of ACM SIGMOD International Conference on Management of Data*, June 1999.
- [AQM⁺96] S. Abiteboul, D. Quass, J. McHugh, J. Widom, and J. Wiener. The lorel query language for semistructured data. *Journal on Digital Libraries*, 1(1), 1996.
- [BDCd95] Franco Barbanera, Mariangiola Dezani-Ciancaglini, and Ugo de’Liguoro. Intersection and union types: Syntax and semantics. *Information and Computation*, 119(2):202–230, June 1995.
- [BDFS97] P. Buneman, S. Davidson, M. Fernandez, and D. Suciu. Adding structure to unstructured data. In *Proc. ICDT*, 1997.
- [BDHS96] P. Buneman, S. Davidson, G. Hillebrand, and D. Suciu. A query language and optimization techniques for unstructured data. In *ACM-SIGMOD*, pages 505–516, 1996.
- [BLS⁺94] P. Buneman, L. Libkin, D. Suciu, V. Tannen, and L. Wong. Comprehension syntax. *SIGMOD Record*, 23(1):87–96, March 1994.
- [BNTW95] Peter Buneman, Shamim Naqvi, Val Tannen, and Limsoon Wong. Principles of programming with complex objects and collection types. *Theoretical Computer Science*, 149(1):3–48, September 1995.
- [Carrk] L. Cardelli. Amber. In B. Robinet G. Cousineau, P.L. Curien, editor, *Combinators and Functional programming languages*, page 1986. Springer-Verlag, New-York.
- [CM94] M. Consens and T. Milo. Optimizing queries on files. In *Proc. ACM Sigmod, Minneapolis*, 1994.
- [Dam94] Flemming M. Damm. Subtyping with union types, intersection types and recursive types. In Masami Hagiya and John C. Mitchell, editors, *Theoretical Aspects of Computer Software*, volume 789 of *Lecture Notes in Computer Science*, pages 687–706. Springer-Verlag, April 1994.
- [DCdP96] M. Dezani-Ciancaglini, U. de’Liguoro, and A. Piperno. Filter models for conjunctive-disjunctive λ -calculi. *Theoretical Computer Science*, 170(1-2):83–128, December 1996.
- [DFF⁺] A. Deutsch, M. Fernandez, D. Florescu, A. Levy, and D. Suciu. Xml-ql: A query language for xml. <http://www.w3.org/TR/NOTE-xml-ql>.
- [Hay91] Susumu Hayashi. Singleton, union and intersection types for program extraction. In T. Ito and A. R. Meyer, editors, *Theoretical Aspects of Computer Software (Sendai, Japan)*, number 526 in *Lecture Notes in Computer Science*, pages 701–730. Springer-Verlag, September 1991. Full version in *Information and Computation*, 109(1/2):174-210, 1994.
- [KTU94] A. J. Kfoury, J. Tiuryn, and P. Urzyczyn. An analysis of ML typability. *Journal of the ACM*, 41(2):368–398, March 1994.

- [NAM97] S. Nestorov, S. Abiteboul, and R. Motwani. Inferring structure in semistructured data. In *Proceedings of the Workshop on Management of Semi-structured Data*, 1997. Available from <http://www.research.att.com/~suciu/workshop-papers.html>.
- [Pie91] Benjamin C. Pierce. Programming with intersection types, union types, and polymorphism. Technical Report CMU-CS-91-106, Carnegie Mellon University, February 1991.