

NOTICE WARNING CONCERNING COPYRIGHT RESTRICTIONS:
The copyright law of the United States (title 17, U.S. Code) governs the making of photocopies or other reproductions of copyrighted material. Any copying of this document without permission of its author may be prohibited by law.

A Record Calculus Based on Symmetric Concatenation

Robert Harper Benjamin Pierce

August 21, 1990

CMU-CS-90-157₂

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Abstract

A number of different formulations of record calculi have been proposed. These may be broadly classified according to whether or not they are based on subsumption and whether they admit the specification of positive or negative information about record type variables. The systems considered by Wand (and their derivatives) may be classified as subsumption-free, negative information calculi. Early systems of bounded quantification are subsumption-based, and admit specification only of positive information. More recently, Cardelli and Mitchell have proposed a subsumption-based system featuring both positive and negative constraints, and, in earlier work, the authors have considered a subsumption-free variant. None of these systems copes well with the “symmetric merge” operation in which two records are considered mergeable only if they have no overlapping fields. To address this question, Cardelli and Mitchell suggest an extension to their subsumption-based calculus to admit mergeability constraints. In this paper we show that a subsumption-free calculus based only on mergeability constraints suffices not only for symmetric merge, but also for a wide class of other record operations. Moreover, the proposed calculus avoids some known difficulties with recursive types associated with positive-information calculi; in particular, the motivating examples for F-bounded quantification may be readily expressed without extending the language. A number of object-oriented features, such as those considered by Wand and Buneman and Ogori, may also be represented naturally in this calculus.

This research was supported in part by the Office of Naval Research and in part by the Defense Advanced Research Projects Agency (DOD), monitored by the Office of Naval Research under Contract N00014-84-K-0415, ARPA Order No. 5404.

The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of ONR, DARPA or the U.S. government.

510,7808
C28r
95-157
0.2

Keywords: Lambda calculus and related systems, Language theory, Programming, Type structure, Data types and structures, Records.

1 Introduction

Cardelli [2, 3] observed that certain aspects of *inheritance* in object-oriented languages can be understood in terms of inclusion relations among record types in a typed λ -calculus. These inclusions are defined formally as a *subtype* relation: a type t is a subtype of t' , written $t \leq t'$, if any member of t may safely be used in a context where a member of t' is expected. The fact that the type of an expression may always be promoted to a supertype is captured by the rule of *subsumption*:

$$\frac{\begin{array}{l} G \vdash e \in t \\ G \vdash t \leq t' \end{array}}{G \vdash e \in t'} \quad (\text{SUBSUMPTION})$$

Cardelli and Wegner [6] extended this idea to a powerful second-order type system combining Cardelli's ordering on types with type quantification [9, 22], using techniques developed by Mitchell [16]. Wand [24, 25] analyzed the concept of record inclusions in the context of ML type inference and introduced the notion of "row variables," which allow types to be given to terms involving a natural *record extension* operator. This work was refined by Jategaonkar and Mitchell [12, 13] and Stansifer [23].

Rémy [19] introduced the notion of *positive and negative information* about record fields and the intuition that increasing either positive or negative information — specifying that fields are either definitely present or definitely absent — gives more refined types. This intuition, formalized as an appropriate extension to the kind system, plus the restriction that the set of field labels is finite, enabled him to use ordinary unification as in ML [8] to do type inference for programs involving extensible records. Both Wand [27] and Rémy [20, 21] later extended this system to infinite label sets.

More recently, Cardelli and Mitchell [4, 5] discovered an elegant calculus of primitive record operations combining bounded quantification with positive and negative information about fields and generalizing Cardelli's original subtype ordering on fixed-length records. In this system, the preorder on types is used to encode both positive and negative information. For example, record extensions like $e|l=e'$ (" e extended with value e' at label l ") are only well formed when the field being added is not already present; to prevent run time type errors, the typing rule for $|$ must ensure that this is the case. Cardelli and Mitchell express this constraint in terms of the preorder by requiring that " e has some type τ that is a subtype of a type lacking l ." The *restriction* operator \backslash is used to increase negative information; for example, if τ is the record type $\{l_1:t_1\}$, then $\tau \backslash l_2$ is a subtype of τ .

In an earlier report [10], we set out to represent positive and negative information about record types as directly as possible, using explicit constraints rather than encoding them with a preorder structure on types. For example, this system expresses the well-typedness constraint on $e|l=e'$ as " e has type τ and τ lacks l ," where the judgement form " τ lacks l " is axiomatized explicitly. The result is a somewhat simpler but more verbose system with no preorder on types, where genericity over record types arises solely from *constrained quantification* of type variables: a quantified type " $\forall a \text{ lacks } L^- \text{ has } L^+:T^+. \tau$ " can be instantiated with any record type τ , so long as τ 's set of fields is disjoint from L^- and includes each $l_i:t_i$ in $L^+:T^+$. This line of development essentially amounts to reverse-engineering Cardelli and Mitchell's system back in the direction of Rémy's. In fact, during the early stages of their work Cardelli and Mitchell independently developed a similar system by extending the kind system of the polymorphic λ -calculus along the lines suggested by Rémy.

Other useful ways of manipulating record types are provided by the *merge* or *concatenate* operator $||$, which combines the fields of two existing records. There are several different forms of the merge operator, distinguished by what happens when the records e_1 and e_2 in a merge expression $e_1 || e_2$ have one or more fields in common. The *asymmetric merge* operator gives preference to the values from e_2 . The *symmetric merge* operator disallows merge expressions where the fields of e_1 and e_2 are not guaranteed to be disjoint. The *recursive merge* computes the values of common fields by recursively merging their values in e_1 and e_2 . Of these, the symmetric variant seems to us to be the most useful, since it makes the finest distinctions. This is the version we assume in the present paper, except where explicitly indicated.

A restricted form of merging can be defined directly. For example, $\{l_1:t_1, l_2:t_2\} || \{l_3:t_3, l_4:t_4\}$ can be rewritten simply as $\{l_1:t_1, l_2:t_2, l_3:t_3, l_4:t_4\}$. But in general — when the types of e_1 and e_2 may be variables, for example — typechecking with $||$ requires a substantial increase in the complexity of the type

system. To ensure that $\lambda x:a. \text{lambda } y : b.(x||y)$ is well typed, we need to guarantee that the type variables a and b are never instantiated to types with overlapping fields. This condition cannot be stated in terms of the previous forms of positive and negative information (“has” and “lacks” constraints), but must be provided explicitly by annotating type variables and quantifiers with *compatibility constraints*. Once this information is provided for type variables, we can define what it means for two arbitrary types τ_1 and τ_2 to be compatible, written $\tau_1 \# \tau_2$. The merge expression $e_1 || e_2$ is then well typed when $e_1 \in \tau_1$, $e_2 \in \tau_2$, and $\tau_1 \# \tau_2$. The type of this expression is the *merge type* $\tau_1 || \tau_2$.

Wand [27] studied type inference for an extension of ML with an asymmetric merge operator and showed how to encode class definitions similar to those found in object-oriented programming languages in this calculus. The asymmetric merge operation has the advantage that compatibility constraints are not necessary: any two records may be merged, with the rightmost one overriding. On the other hand, the use of symmetric merge allows the type checker to detect inadvertent clashes of labels, which can be useful in practice. Both systems have the property that the type checker must keep track of which component of a merge has a given field, which leads to problems in type reconstruction.

Ohuri and Buneman studied type inference in for database programming languages [18], and for object-oriented programming [17]. In their work on database programming languages they consider a type system with records and sets, with operations (such as the relational join) chosen to model typical database applications. To support ML-like type inference for this language, they introduced the notion of a *conditional type scheme*, in which a type variable may be constrained to range over records possessing certain components. These constraints are similar to the “has” constraints discussed above. (They also used another form of constraint relevant only to the relational join and projection operations.) In their work on object-oriented programming, they consider a form of record update operation in which a value of a given field may be overridden. Once again, their notion of conditional type scheme plays a central role. They also considered extensions to support object-oriented programming constructs similar to those suggested by Wand. In order to support inheritance, they introduced an *ad hoc* form of subsumption in connection with “self.” This introduced an additional form of condition on conditional type schemes; they exhibited a sound type inference algorithm for this extension.

Cardelli and Mitchell [5] sketched several increasingly ambitious formulations of recursive record concatenation as extensions to their calculus of operations on records. The most powerful of these requires the notion of “constrained, multiply bounded quantification,” where a finite set of type variables is bound simultaneously, with each constrained to be a subtype of some given type and certain subsets constrained to be compatible. For example, a function that takes two compatible records x_1 and x_2 , where x_1 has at least a field l_1 of type Int and x_2 has at least a field l_2 of type Int , and returns the result of merging x_1 and x_2 , can be written (altering their concrete syntax slightly)

$$\forall (a_1 \leq \{l_1 : \text{Int}\}, a_2 \leq \{l_2 : \text{Int}\}, a_1 \# a_2). \lambda x_1:a_1. \lambda x_2:a_2. x_1 || x_2.$$

In the present paper, we study a record calculus $\lambda^||$ based on a much more straightforward formulation of constrained quantification. In $\lambda^||$, each record type variable in a context G has a list of *compatibility assumptions* R , where the elements of R are record types and $a \# R$ asserts that $a \# r_i$ for each $r_i \in R$. The constrained type abstraction operator $\Lambda a \# R. e$ adds the assumption $a \# R$ to the context used to typecheck e . A type application $e[r]$ must check that r satisfies all of the constraints on the quantifier. For example, the function mentioned in the previous paragraph can be defined in $\lambda^||$ as follows:

$$\begin{aligned} &\forall a_1 \# l_1 : \text{Int}. \forall a_2 \# a_1, l_2 : \text{Int}. \\ &\lambda x_1 : a_1 || l_1 : \text{Int}. \lambda x_2 : a_2 || l_2 : \text{Int}. \\ &x_1 || x_2. \end{aligned}$$

This example underscores an important point: the form of constraint used in $\lambda^||$ can only be used to express *negative* information about record type variables. The function above takes two type variables, each of which *lacks* the appropriate field. To form the types expected for the records x_1 and x_2 on the two λ -abstractions, the missing fields are merged back into a_1 and a_2 . This kind of transformation from mixed positive and negative constraints on quantifiers to pure negative constraints can be carried out mechanically.

This has a very similar flavor to Wand’s treatment of merging with row variables [27]. In fact, we adopt the same point of view as Wand with regard to subsumption: rather than introduce a preorder on types that includes record extension as a special case, we prefer to use a form of quantification to capture the

possible extensions of a record type and use type application to choose the appropriate extension for a given context. However, in contrast to Wand, we are dealing with an explicitly-typed, second-order calculus with a symmetric, rather than asymmetric, merge operator. This leads to a somewhat different overall flavor, as we shall illustrate below.

Our central claim in this paper is that the straightforward formulation of constrained quantification embodied in λ^{\parallel} may be viewed as *primary*, in the sense that most of the examples motivating row variables, bounded quantification, and Cardelli and Mitchell's bounded quantification with positive and negative information can be expressed in a calculus based on this form of constrained quantifier, with no need for additional mechanisms like subsumption. Similarly, in Section 3 we sketch an extension $\lambda^{\parallel\mu}$ of λ^{\parallel} with recursive types and show that the examples motivating the extension from systems with bounded quantification and recursive types to systems with F-bounded quantification [1] can be expressed directly in $\lambda^{\parallel\mu}$.

In Section 2 we define the syntax and typing rules of λ^{\parallel} and sketch a proof of the decidability of typechecking. Section 3 illustrates the expressiveness of the system through a number of examples. Section 4 describes some promising avenues for further research. Section 5 offers concluding remarks. A complete listing of the typing rules for λ^{\parallel} appears in Appendix A; complete rules for the algorithmic formulation used to prove decidability are given in Appendix B.

2 Definition and Properties of λ^{\parallel}

2.1 Syntax

This section introduces the notational conventions used in the rest of the paper and defines the concrete syntax of λ^{\parallel} .

The metavariables q , x , s and t range over types (q , x , and s are used when the type in question is expected to be a record type); p ranges over primitive types; a and b range over record type variables; R and S range over constraint lists; e and f range over terms; x ranges over variables; l ranges over field labels.

Types:

t	::=	p	primitive
		a	record type variable
		Empty	empty record
		$l:t$	single-field record
		$x \setminus l$	restriction
		$r_1 \parallel r_2$	merge
		$t_1 \rightarrow t_2$	function space
		$\forall a \# R. t$	constrained type quantification

The record types are those built up from record type variables, **Empty**, and single-field records by applications of merge and restriction. Ordinary type quantification is omitted from this presentation, but could be added by considering general type variables and an associated quantifier.

Constraint lists:

R	::=	\diamond	empty constraint list
		x, R	nonempty constraint list

Note that we make no provision here for including “bare labels” in constraint lists. Formally, expressions like $\forall a \# l. t$ are taken as abbreviations for $\forall a \# (l:s). t$, where s is some dummy type. (See Section 3 for examples of this idiom.)

It is sometimes convenient to consider a list as a finite set and write, for example, $r \in R$.

Terms:		
e	$::=$	x
		$\lambda x:t. e$
		$e_1 e_2$
		empty
		$l=e$
		$e \setminus l$
		$e_1 \parallel e_2$
		$e.l$
		$\Lambda a\#R. e$
		$e[t]$
		variable
		abstraction
		application
		empty record
		single-field record
		restriction
		merge
		selection
		constrained type abstraction
		constrained type application

Free and bound variables are defined in the usual way, with the proviso that in the case of type quantification and abstraction, the variable a is *not* considered bound in the constraint list R . Terms and types are identified up to renaming of bound variables. The notation $[t/a]t'$ denotes capture-avoiding substitution of t for free occurrences of a in t' ; similarly, $[e/x]e'$ denotes capture-avoiding substitution of e for free occurrences of x in e' .

The metavariable T ranges over type contexts — finite sequences of declarations of the form $a\#R$ with no type variable declared twice. The metavariable G ranges over term contexts — finite sequences of declarations of the form $x:t$ with no variable mentioned twice.

If l is a label and r is a well-formed record type, then $r.l$ is defined to be the type associated with label l in r , if any. More precisely, the partial function $r.l$ is defined by induction on the structure of r as follows. (We write $r.l \uparrow$ for “ $r.l$ is undefined” and $r.l \downarrow$ for “ $r.l$ is defined.”)

$$\begin{aligned}
(l:t).l &= t \\
(r \setminus l').l &= r.l \quad \text{when } l \neq l' \\
(r_1 \parallel r_2).l &= r_1.l \quad \text{if } r_1.l \downarrow \text{ and } r_2.l \uparrow \\
(r_1 \parallel r_2).l &= r_2.l \quad \text{if } r_2.l \downarrow \text{ and } r_1.l \uparrow \\
r.l \uparrow &\quad \text{otherwise.}
\end{aligned}$$

2.2 Typing Rules

λ^{\parallel} is a formal system for deriving judgements of the following forms:

Well-formed type context:	$T \text{ ok}$
Well-formed term context:	$T \vdash G \text{ ok}$
Well-formed type:	$T \vdash t \text{ type}$
Well-formed record type:	$T \vdash r \text{ record}$
Well-formed constraint list:	$T \vdash R \text{ ok}$
Constraint list satisfaction:	$T \vdash r \# R$
Compatible types:	$T \vdash r_1 \# r_2$
Equivalent types:	$t_1 \sim t_2$
Equivalent constraint sets:	$R_1 \sim R_2$
Well-formed term:	$T; G \vdash e \in t$

The complete set of rules for λ^{\parallel} appears in Appendix A. A representative selection of the rules appears in Figures 2 to 5. The rules are summarized in the remainder of this section.

The rules for context formation are the expected ones (see Appendix A). A type context $T, a\#R$ is well-formed if a is not already declared in T and R is well-formed in T . A term context $G, x:t$ is well-formed in T if x is not already declared in G and if t is a well-formed type in T .

A constraint list R is well formed in T iff each component type expression r in R is a well-formed record type in T .

The rules for formation of primitive types and function types are as usual. A quantified type $\forall a\#R. t$ is well-formed in T if R is a well-formed constraint list in T , and t is a well-formed type in $T, r\#R$ (rule WFT-ALL). Every record type is a typerule WFT-REC).

The formation rules for record types are summarized in Figure 1. A record type variable is a record type in T provided that it is declared in T (rule WFR-VAR). The empty record is always well-formed (rule

$$\begin{array}{c}
\frac{T_1, a \# R, T_2 \text{ ok}}{T_1, a \# R, T_2 \vdash a \text{ record}} \quad (\text{WFR-VAR}) \\
\\
\frac{T \vdash t \text{ type}}{T \vdash l:t \text{ record}} \quad (\text{WFR-BASE}) \\
\\
\frac{T \vdash r \text{ record} \quad r.l \downarrow}{T \vdash r \setminus l \text{ record}} \quad (\text{WFR-RESTR}) \\
\\
\frac{T \text{ ok}}{T \vdash \text{Empty record}} \quad (\text{WFR-EMPTY}) \\
\\
\frac{T \vdash r_1 \# r_2}{T \vdash r_1 \parallel r_2 \text{ record}} \quad (\text{WFR-MERGE})
\end{array}$$

Figure 1: Selected well-formedness rules for types

$$\begin{array}{c}
\frac{T; G \vdash e \in t \quad T \vdash t' \text{ type} \quad t \sim t'}{T; G \vdash e \in t'} \quad (\text{WTT-EQ}) \\
\\
\frac{T, a \# R; \Gamma \vdash e \in t}{T; G \vdash \Lambda a \# R. e \in \forall a \# R. t} \quad (\text{WTT-ALL-I}) \\
\\
\frac{T; G \vdash e \in \forall a \# R. t \quad T \vdash r \# R}{T; G \vdash e[x] \in [x/a]t} \quad (\text{WTT-ALL-E})
\end{array}$$

Figure 2: Selected typing rules

WFR-EMPTY). The single-field record $l:t$ is well-formed in T if t is a well-formed type in T (rule WFR-BASE). The restriction $r \setminus l$ is well formed in T if r is well formed in T and $r.l$ is defined (rule WFR-RESTR). This means, in particular, that no expression of the form $a \setminus l$ where a is a variable is well formed, since $a.l$ is never defined. This reflects the fact that compatibility constraints on a variable a are negative in character and cannot be used to postulate that all instances of a *have* any particular fields, only that they *lack* certain fields. This also implies that if $r \setminus l$ is a well-formed record expression, then the restriction may be eliminated (see below for a discussion of type equivalence). Finally, a merge $r_1 \parallel r_2$ is well-formed in T if r_1 and r_2 are well-formed in T , and, moreover, r_1 and r_2 are compatible in T . Informally, this means that r_1 lacks whatever fields r_2 has and vice versa, so that the merge is well-defined. It is important to realize that the compatibility relation must be axiomatized relative to a context since, in the presence of constrained quantification, it is not possible to determine from the syntactic form of record expressions alone whether they are compatible. (See below for further discussion of the compatibility relation.)

The typing rules for λ^{\parallel} terms are organized along standard lines. (See Figure 2.) Typing is invariant under type equality (rule WTT-EQ). The rules for variables, lambda abstractions, and application are standard. The empty record is always well-formed. A single-field record $l=e$ has type $l:t$ in T if e has type t in T . The restriction $e \setminus l$ has type $r \setminus l$ in T provided that e has type r in T and $r.l \downarrow$: we may restrict only on a field that e actually possesses. The merge $e_1 \parallel e_2$ has type $r_1 \parallel r_2$ in T provided that e_1 has type r_1 in T and e_2 has type r_2 in T and r_1 and r_2 are compatible in T . In other words, we may not merge two records unless they are non-overlapping. To achieve the effect of overriding, one must restrict the fields to be overridden before forming the merge. The selection $e.l$ has type $r.l$ in T if e has type r in T and $r.l \downarrow$. By the definition of $r.l$, the type of $e.l$ is unique if it well-formed.

The abstraction $\Lambda a \# R. e$ has type $\forall a \# R. t$ provided that e has type t in $T, a \# R$, which also entails that R is a well-formed constraint set in T (rule WTT-ALL-I). It is important to realize that the constraint list

$$\begin{array}{c}
\frac{T \vdash r \# s \quad r \sim r' \quad s \sim s'}{T \vdash r' \# s'} \quad (\text{CMP-EQ}) \\
\\
\frac{T \vdash r \# s}{T \vdash s \# r} \quad (\text{CMP-SYMM}) \\
\\
\frac{T \vdash r \# l:t \quad T \vdash t' \text{ type}}{T \vdash r \# l:t'} \quad (\text{CMP-BASE}) \\
\\
\frac{T_1, a \# R, T_2 \text{ ok} \quad r_i \in R}{T_1, a \# R, T_2 \vdash a \# r_i} \quad (\text{CMP-TVAR}) \\
\\
\frac{T \vdash r \# (s1||s2)}{T \vdash r \# s_i} \quad (\text{CMP-MERGE-E}) \\
\\
\frac{T \vdash s1 \# s2 \quad T \vdash r \# s1 \quad T \vdash r \# s2}{T \vdash r \# (s1||s2)} \quad (\text{CMP-MERGE-I}) \\
\\
\frac{l \neq l' \quad T \vdash l:t \text{ record} \quad T \vdash l':t' \text{ record}}{T \vdash l:t \# l':t'} \quad (\text{CMP-BASE/BASE}) \\
\\
\frac{T \vdash r \text{ record}}{T \vdash r \# \text{Empty}} \quad (\text{CMP-EMPTY})
\end{array}$$

Figure 3: Selected compatibility rules

R cannot be replaced by a single record type r , for two related reasons. First, it is necessary to postulate that a variable be compatible with a number of different record types. For example, if we are to merge a variable a with the base records $l:t$ and $l':t'$, then a must be constrained to be compatible with both of these records, which is to say that all instances of a must not contain l or l' fields. Second, the constraint list R cannot be collapsed into a single record type consisting of the merge of the component record types r of R because the records in R need not be mutually compatible. In particular, we may wish to postulate that a variable a is compatible with two other variables b and c so that the merges $a||b$ and $a||c$ are both well-formed. This in no way entails that b and c must be compatible, and it would be a serious restriction to require that they be so.

The type application $e[r]$ is well formed in T if e has type $\forall a \# R. t$ in T , r is a record in T , and r is compatible with each record type in R (relative to T). In other words, r must satisfy the constraints associated with the quantifier in order for the type application to be sensible. When this is the case, the type of $e[r]$ is as expected, namely $[r/a]t$ (rule WTT-ALL-E). It will turn out that the formulation of the system ensures that if r satisfies the constraints in R relative to T , then r is a record type in T . To support general parametric polymorphism, we would have to extend the system with a separate form of quantifier that quantifies over all types. We omit this extension for the sake of simplicity.

We now turn to the definition of compatibility. Informally, $T \vdash r \# s$ holds iff every instance r' and s' of r and s obtained by substituting record types for type variables consistently with the constraints in T is such that whatever fields r' has, s' lacks, and conversely, whatever fields s' has, r' lacks. (The notion of “consistently with the constraints in T ” itself makes use of the notion of compatibility, but for terms with fewer type variables.) This description raises several important points. First, the definition of compatibility is relative to a type context T . Second, what matters about r' and s' is the singleton records that occur within them, and, moreover, it is only the *labels*, not the types of these singletons that matter for compatibility. In other words $l:t$ is compatible with $l':t'$ provided only that l is different from l' , and conversely $l:t$ is incompatible with $l:t'$ even if t and t' are identical. Third, the definition of compatibility should respect type equivalence: equal types (as we will see below) have the same fields.

$r \parallel \text{Empty} \sim r$	(EQ-MERGE-UNIT)
$r_1 \parallel (r_2 \parallel r_3) \sim (r_1 \parallel r_2) \parallel r_3$	(EQ-MERGE-ASSOC)
$r_1 \parallel r_2 \sim r_2 \parallel r_1$	(EQ-MERGE-COMM)
$r \setminus l \setminus l' \sim r \setminus l' \setminus l$	(EQ-RESTR/RESTR')
$(l:t) \setminus l \sim \text{Empty}$	(EQ-BASE/RESTR)
$\frac{r_1 \setminus l \downarrow}{(r_1 \parallel r_2) \setminus l \sim (r_1 \setminus l \parallel r_2)}$	(EQ-MERGE/RESTR)
$\frac{R \sim R' \quad t \sim t'}{\forall a \# R. t \sim \forall a \# R'. t'}$	(EQ-CONG-ALL)

Figure 4: Selected type equivalence rules

$r, (r', R) \sim r', (r, R)$	(CEQ-SWAP)
$\text{Empty}, R \sim R$	(CEQ-EMPTY)
$(r_1 \parallel r_2), R \sim r_1, (r_2, R)$	(CEQ-MERGE)
$r, (r, R) \sim r, R$	(CEQ-DUPL)
$l:t, R \sim l:t', R$	(CEQ-BASE)

Figure 5: Selected constraint list equivalence rules

These observations lead to the following formalization of the compatibility relation (see Figure 3). Compatibility is symmetric (rule **CMP-SYMM**) and respects type equivalence (rule **CMP-EQ**). If a record is compatible with a single-field record $l:t$, then it is compatible with every single-field record $l:t'$ (rule **CMP-BASE**). A type variable is compatible with the records in its constraint list (rule **CMP-TVAR**). A record is compatible with a merge iff it is compatible with each of the components of the merge (rules **CMP-MERGE-I** and **CMP-MERGE-E**). Two single-field records are compatible if they have different labels (rule **CMP-BASE/BASE**). Every record is compatible with the empty record (rule **CMP-EMPTY**). Finally, a record type satisfies a constraint list iff it is compatible with every element of the list.

It is worth noting that there are no rules for compatibility of restrictions. It will be possible to prove that every well-formed record is equivalent to a unique restriction-free record; hence compatibility of restrictions is covered by the general rule of invariance of compatibility under equivalence (rule **CMP-EQ**).

Finally, we consider equivalence of types and constraint sets. (A selection of the equivalence rules appears in Figures 4 and 5.) We begin with type equivalence. The relation $t \sim t'$ is an equivalence relation and is a congruence with respect to the type-forming operations. The axioms for record type equivalence are as follows. Restrictions commute with one another (rule **EQ-RESTR/RESTR'**) (in a well-formed expression l will be distinct from l' , but there is no reason to insist on that in the rule.) Restriction of a singleton $l:t$ on

\perp results in the empty record type (rule EQ-BASE-RESTR). Restriction distributes to the appropriate (given EQ-SYMM and EQ-CONG-RESTR) component of a merge (rule EQ-MERGE-RESTR). Merge is associative (rule EQ-MERGE-ASSOC), **Empty** is a unit of merge (rule EQ-MERGE-UNIT), and merge is commutative (rule EQ-MERGE-COMM). This last rule is a reflection of the fact that in a well-formed merge there can be no overlapping fields.

The presence of the constrained quantifier requires that we say something about equivalence of constraint lists (rule EQ-CONG-ALL). Here the guiding principle is that equivalent constraint sets should be satisfied by the same record types. In addition to being an equivalence relation and compatible with the “cons” operation on constraint lists, we require that order and multiplicity be irrelevant (rules CEQ-SWAP and CEQ-DUPL) so that constraint lists are essentially finite sets. We may also “flatten” merges (rule CEQ-MERGE), eliminate **Empty** (rule CEQ-EMPTY), and ignore the types of single-field records (rule CEQ-BASE). As we will see below, these equivalences are sufficient to ensure that constraint sets may be normalized into the form $l_1, \dots, l_n, a_1, \dots, a_k$ where the l_i 's are labels and the a_i 's are record type variables. This is consistent with the intuition that the relevant properties of a record, for the purposes of consistency checking, are its atomic components. Note that it will not do to simply *define* constraint sets this way: substitutions of record types for record type variables can disturb this form by instantiating variables to complex record types. The constraint set equivalence rules are needed to restore this simple form.

It should be stressed that there is considerable leeway in the choice of type and constraint list equivalence. The choices made above are guided by the intended interpretation of record types, and by algorithmic considerations (we impose enough equivalences to achieve well-behaved normal forms.) Since we postulate the invariance of typing under type equivalence, the class of well-typed terms increases as more types are identified. It is therefore desirable to impose whatever equivalences may be consistently added to the system, constrained only by informal (at this stage) considerations about models and implementations of the calculus.

This completes our summary of λ^{\parallel} .

2.3 Properties

In this section we establish the decidability of type checking for λ^{\parallel} .

First, we prove that equivalence of well-formed types and constraint lists is decidable. The main complication is due to the associative and commutative axioms for record type equality and the permutation and duplicate-elimination axioms for constraint lists. However, these may be handled by applying Huet's method of “confluence modulo an equivalence relation” [11]. The idea is to segregate the “proper” reduction axioms from the “pure” equational axioms and to show that the full equivalence relation may be characterized as “pure” equivalence of normal forms. A byproduct of this development is a useful characterization of normal forms.

Second, we consider the decidability of the compatibility relation for well-formed, restriction-free record types. The main complication here arises from the rules of symmetry, invariance under change of single-field record type, and respect for equivalence. The algorithm is presented as an “almost syntax-directed” set of inference rules that are readily shown to be decidable, and that may be proved equivalent (in a suitable sense) to the full compatibility relation. The algorithm relies on the characterization of normal forms mentioned above, and is therefore sensitive to the definition of equivalence of record types and constraint sets.

Finally, we give a type synthesis algorithm that computes, given a context and a term, a representative of the equivalence class of possible types for that term. This algorithm proceeds along standard lines, making use of the algorithms for checking equivalence and compatibility.

2.3.1 Technical Lemmas

Lemma 2.3.1.1: A derivation of the well-formedness of a type, record type, or constraint list includes subderivations of the well-formedness of all of its subphrases.

Lemma 2.3.1.2:

1. $a \setminus 1$ is never well formed for any variable a , regardless of context.

2. $\text{Empty}\backslash 1$ is never well-formed.
3. $r\backslash 1\backslash 1'$ is well-formed only if $1 \neq 1'$ and r is not a single-field record.
4. $1:t\backslash 1'$ is well-formed only if $1 = 1'$.

Lemma 2.3.1.3: If r_1 is defined and $r \sim s$, then s_1 is defined and $r_1 \sim s_1$.

Lemma 2.3.1.4: Assume that $T \vdash r \# s$. Then it is not the case that both r_1 and s_1 are defined.

Corollary 2.3.1.5: If $r||s$ is well formed, then r_1 defined implies s_1 undefined, and vice versa.

Lemma 2.3.1.6: $\vdash T_1, a\#R, T_2 \text{ ok}$ implies $T_1 \vdash R \text{ ok}$ as a subderivation.

Lemma 2.3.1.7: $T \vdash r \# s$ implies $T \vdash r \text{ record}$ and $T \vdash s \text{ record}$.

Lemma 2.3.1.8: If $T \vdash r \# R$ and $r \sim r'$ and $R \sim R'$, then $T \vdash r' \# R'$.

Lemma 2.3.1.9: Let J stand for an arbitrary judgement. If $T_1, a\#R, T_2 \vdash J$, and $T_1 \vdash r\#R$, then $T_1, [r/a]T_2 \vdash [r/a]J$.

2.3.2 Confluence of reduction modulo an equivalence relation

Our proof of the confluence of type and constraint list normalization uses a technique due to Huet [11]. In this section we briefly recapitulate Huet's main theorem.

Definition 2.3.2.1: The metavariables M, N, P, Q , and U here range over an unspecified set of phrases (in the next section they will be either types or constraint lists).

Definition 2.3.2.2: $\mathcal{V}(M)$ is the set of variables appearing free in a phrase M .

Definition 2.3.2.3: A phrase M is linear iff all of its variable occurrences are distinct.

Definition 2.3.2.4: A rewrite rule is a pair of terms.

Definition 2.3.2.5: A critical pair $\langle P, Q \rangle$ for two rewriting rules $\langle M_1, N_1 \rangle$ and $\langle M_2, N_2 \rangle$ is a most general common instance of the two rules, found by superposing M_2 on a non-variable subphrase U of M_1 . P is the result of rewriting the subphrase U of the common instance according to the rule $\langle M_2, N_2 \rangle$. Q is the result of rewriting the whole common instance according to the rule $\langle M_1, N_1 \rangle$.

Definition 2.3.2.6: An equational theory $\langle \mathcal{R}, \mathcal{E} \rangle$ consists of two sets of rewriting rules: a "reduction part" \mathcal{R} and an "equivalence part" \mathcal{E} .

Definition 2.3.2.7: Let $\langle \mathcal{R}, \mathcal{E} \rangle$ be an equational theory. The one-step reduction relation induced by \mathcal{R} is denoted by $\rightarrow_{\mathcal{R}}^1$, and the one-step reduction relation induced by \mathcal{E} is denoted by $\approx_{\mathcal{E}}^1$. The reflexive, transitive closure of $\rightarrow_{\mathcal{R}}^1$ is denoted by $\rightarrow_{\mathcal{R}}$, and the reflexive, symmetric, and transitive closure of $\approx_{\mathcal{E}}^1$ is denoted by $\approx_{\mathcal{E}}$.

The subscripts are omitted when they are clear from context.

Definition 2.3.2.8: The set of critical pairs of a reduction relation \mathcal{R} is

$$\{(P, Q) \mid (P, Q) \text{ is a critical pair for } \langle M_1, N_1 \rangle, \langle M_2, N_2 \rangle \in \mathcal{R}\}.$$

Definition 2.3.2.9: Let $\langle \mathcal{R}, \mathcal{E} \rangle$ be an equational theory. The set of critical pairs of \mathcal{E}/\mathcal{R} is

$$\begin{aligned} & \{(P, Q) \mid (P, Q) \text{ is a critical pair for } \langle \mathbb{M}_1, \mathbb{M}_2 \rangle \in \mathcal{R} \text{ and } \langle \mathbb{N}_1, \mathbb{N}_2 \rangle \in (\mathcal{E} \cup \mathcal{E}^{-1})\} \\ \cup & \{(P, Q) \mid (P, Q) \text{ is a critical pair for } \langle \mathbb{M}_1, \mathbb{M}_2 \rangle \in (\mathcal{E} \cup \mathcal{E}^{-1}) \text{ and } \langle \mathbb{N}_1, \mathbb{N}_2 \rangle \in \mathcal{R}\}, \end{aligned}$$

where

$$\mathcal{E}^{-1} = \{ \langle \mathbb{N}, \mathbb{M} \rangle \mid \langle \mathbb{M}, \mathbb{N} \rangle \in \mathcal{E} \}.$$

Definition 2.3.2.10: The set of critical pairs of a theory $\langle \mathcal{R}, \mathcal{E} \rangle$ is the union of the set of critical pairs of \mathcal{R} and the set of critical pairs of \mathcal{E}/\mathcal{R} .

Definition 2.3.2.11: \mathbb{M}^* denotes an arbitrary \mathcal{R} -normal form of \mathbb{M} .

Proposition 2.3.2.12: (Huet's Theorem 3.3) Let $\langle \mathcal{R}, \mathcal{E} \rangle$ be an equational theory such that

1. for all $\langle \mathbb{M}, \mathbb{N} \rangle \in \mathcal{R}$, $\mathcal{V}(\mathbb{N}) \subseteq \mathcal{V}(\mathbb{M})$ and \mathbb{M} is linear;
2. for all $\langle \mathbb{M}, \mathbb{N} \rangle \in \mathcal{E}$, $\mathcal{V}(\mathbb{M}) = \mathcal{V}(\mathbb{N})$;
3. $\rightarrow^1 \circ \approx$ is noetherian, where \approx is the reflexive, symmetric, transitive closure of \approx^1 .

The theory $\langle \mathcal{R}, \mathcal{E} \rangle$ is confluent iff for all its critical pairs $\langle P, Q \rangle$ we have $P^* \approx Q^*$. When $\langle \mathcal{R}, \mathcal{E} \rangle$ is confluent, we say that “ \mathcal{R} is confluent modulo \mathcal{E} .” Then $\mathbb{M} =_{\mathcal{R} \cup \mathcal{E}} \mathbb{N}$ iff $\mathbb{M}^* \approx \mathbb{N}^*$, where $=_{\mathcal{R} \cup \mathcal{E}}$ is the equivalence relation induced by \mathcal{R} and \mathcal{E} .

2.3.3 Equivalence of Types and Constraint Lists

We now apply Huet's theory to the confluence of a type normalization procedure for λ^{ll} .

Definition 2.3.3.1: For the algorithmic version of the system, constraint lists are extended with “bare labels”:

$$\begin{array}{ll} \mathbb{R} ::= & \diamond \quad \text{empty constraint list} \\ & | \tau, \mathbb{R} \quad \text{constraint list beginning with a type} \\ & | \mathbb{1}, \mathbb{R} \quad \text{constraint list beginning with a bare label} \end{array}$$

We extend the relation $T \vdash \tau \# \mathbb{R}$ to constraint lists by defining $T \vdash \tau \# \mathbb{1}, \mathbb{R}$ to mean $T \vdash \tau \# \mathbb{R}$ and $T \vdash \tau \# \mathbb{1} : \tau$ for any type τ .

Definition 2.3.3.2: In this section, it is convenient to be able to refer to phrases from certain related syntactic categories uniformly:

1. The metavariable α ranges over atoms: variables, base records, and bare labels.
2. The metavariables ρ , σ , and τ range over types and bare labels.
3. The metavariable ϕ ranges over type variables and bare labels.
4. The metavariable \mathbb{C} ranges over types and constraint lists.

Definition 2.3.3.3: Let \rightarrow^1 denote the following reduction relation on record types and constraint lists, compatibly extended to all type expressions:

$$\text{Empty} \parallel r \rightarrow^1 r \quad (\text{A-RED-EMPTY-1})$$

$$r \parallel \text{Empty} \rightarrow^1 r \quad (\text{A-RED-EMPTY-2})$$

$$l:t \setminus l \rightarrow^1 \text{Empty} \quad (\text{A-RED-RESTR/BASE})$$

$$(r \parallel s) \setminus l \rightarrow^1 (r \setminus l) \parallel s \quad \text{if } r _l \downarrow \quad (\text{A-RED-RESTR/MERGE-1})$$

$$(r \parallel s) \setminus l \rightarrow^1 r \parallel (s \setminus l) \quad \text{if } s _l \downarrow \quad (\text{A-RED-RESTR/MERGE-2})$$

$$\rho \parallel \sigma, R \rightarrow^1 \rho, \sigma, R \quad (\text{A-CRED-MERGE})$$

$$\text{Empty}, R \rightarrow^1 R \quad (\text{A-CRED-EMPTY})$$

$$l:t, R \rightarrow^1 l, R \quad (\text{A-CRED-BASE})$$

Definition 2.3.3.4: Let \approx^1 denote the following reduction relation on record types and constraint lists, compatibly extended to all type expressions:

$$r \setminus l \setminus l' \approx^1 r \setminus l' \setminus l \quad (\text{A-EQ-RESTR/RESTR'})$$

$$r \parallel s \approx^1 s \parallel r \quad (\text{A-EQ-COMM})$$

$$r \parallel (s \parallel t) \approx^1 (r \parallel s) \parallel t \quad (\text{A-EQ-ASSOC})$$

$$\phi, \phi, R \approx^1 \phi, R \quad (\text{A-CEQ-DUPL})$$

$$\rho, \sigma, R \approx^1 \sigma, \rho, R \quad (\text{A-CEQ-COMM})$$

Let \approx denote the reflexive, symmetric, transitive closure of \approx^1 , that is, the equivalence relation generated by \approx^1 .

Let \equiv denote the reflexive, symmetric, transitive closure of $\rightarrow^1 \cup \approx^1$, that is, the equivalence relation generated by \rightarrow^1 and \approx^1 .

Theorem 2.3.3.5:

1. If $T \vdash r$ record and $r \rightarrow^1 r'$, then $T \vdash r'$ record.
2. If $T \vdash t$ type and $t \rightarrow^1 t'$, then $T \vdash t'$ type.
3. If $T \vdash R$ ok, and $R \rightarrow^1 R'$, then $T \vdash R'$ ok.

Proof: By induction on derivations, making use of the technical lemmas in Section 2.3.1. □

Definition 2.3.3.6: The “outer size” of a type r is defined as follows:

$$\begin{aligned}
\text{outer-size}(a) &= 1 \\
\text{outer-size}(\text{Empty}) &= 1 \\
\text{outer-size}(l:t) &= 1 \\
\text{outer-size}(r \setminus l) &= \text{outer-size}(r) + 1 \\
\text{outer-size}(r_1 || r_2) &= \text{outer-size}(r_1) + \text{outer-size}(r_2) + 1 \\
\text{outer-size}(p) &= 0 \\
\text{outer-size}(t_1 \rightarrow t_2) &= 0 \\
\text{outer-size}(\forall a \# R. t) &= 0.
\end{aligned}$$

Definition 2.3.3.7: The rank of a constraint list is defined as follows:

$$\begin{aligned}
\text{rank}(\diamond) &= 0 \\
\text{rank}(l, R) &= \text{rank}(R) \\
\text{rank}(r, R) &= \text{rank}(r) + \text{rank}(R).
\end{aligned}$$

The rank of a type is the sum

$$\text{rank}(t) = \text{restr-count}(t) + \text{base-count}(t) + \text{empty-count}(t) + \text{merge-count}(t),$$

where

- $\text{restr-count}(t)$ is the sum, over all occurrences of the form $s \setminus l$ in t , of $\text{outer-size}(s)$,
- $\text{base-count}(t)$ is the number of base record phrases $l:t'$ in t ,
- $\text{empty-count}(t)$ is the number of occurrences of **Empty** in t , and
- $\text{merge-count}(t)$ is the number of occurrences of $||$ in t .

Lemma 2.3.3.8: The relation $\rightarrow^1 \circ \approx$ is Noetherian when restricted to well-formed expressions. That is, there are no infinite reduction sequences of the form

$$\dots \rightarrow^1 \approx \dots \rightarrow^1 \approx \dots \rightarrow^1 \dots$$

where each of the terms in the sequence is well-formed.

Proof: Applying any of the \rightarrow^1 rules to a type or constraint list strictly reduces its rank, while applying any of the \approx^1 rules in either direction leaves the rank unchanged. \square

Lemma 2.3.3.9:

1. If $T \vdash (r \setminus l || s) \setminus l'$ record and $T \vdash (r || s) \setminus l \setminus l'$ record, then $((r \setminus l || s) \setminus l')^* \approx ((r || s) \setminus l \setminus l')^*$.
2. If $T \vdash a, b, (a || b), R$ ok and $T \vdash (a || b), R$ ok, then $(a, b, (a || b), R)^* \approx ((a || b), R)^*$.
3. If $T \vdash a, b, R$ ok and $T \vdash (a || b), (a || b), R$ ok, then $(a, b, R)^* \approx ((a || b), (a || b), R)^*$.
4. If $T \vdash \tau, \rho, \sigma, R$ ok and $T \vdash (\rho || \sigma), \tau, R$ ok, then $(\tau, \rho, \sigma, R)^* \approx ((\rho || \sigma), \tau, R)^*$.

Lemma 2.3.3.10: For every critical pair $\langle C_1, C_2 \rangle$ of $\langle \rightarrow^1, \approx^1 \rangle$, we have $C_1^* \approx C_2^*$.

Proof: By a tedious verification, using Lemma 2.3.3.9 for the interesting cases. \square

Theorem 2.3.3.11: The restriction of \rightarrow^1 to well-typed terms is confluent modulo \approx .

Corollary 2.3.3.12:

1. If \mathbf{r} and \mathbf{s} are well-formed records, then $\mathbf{r} \equiv \mathbf{s}$ iff there exist \mathbf{r}' and \mathbf{s}' in normal form such that $\mathbf{r} \rightarrow \mathbf{r}'$ and $\mathbf{s} \rightarrow \mathbf{s}'$ and $\mathbf{r}' \approx \mathbf{s}'$.
2. If \mathbf{t} and \mathbf{u} are well-formed types, then $\mathbf{t} \equiv \mathbf{u}$ iff there exist \mathbf{t}' and \mathbf{u}' such that $\mathbf{t} \rightarrow \mathbf{t}'$ and $\mathbf{u} \rightarrow \mathbf{u}'$ and $\mathbf{t}' \approx \mathbf{u}'$.
3. If \mathbf{R} and \mathbf{S} are well-formed constraint lists, then $\mathbf{R} \approx \mathbf{S}$ iff there exist \mathbf{R}' and \mathbf{S}' such that $\mathbf{R} \rightarrow \mathbf{R}'$ and $\mathbf{S} \rightarrow \mathbf{S}'$ and $\mathbf{R}' \approx \mathbf{S}'$.

Corollary 2.3.3.13: Normal forms are unique up to \approx .

If \mathbf{r} is a well-formed record, then \mathbf{r}^* denotes one of its normal forms computed by applying the \rightarrow rules in some canonical order; similarly for types and constraint lists.

Theorem 2.3.3.14:

1. Let \mathbf{r} be a well-formed record type in normal form. Then \mathbf{r} has the form

$$\mathbf{a}_1 \parallel \mathbf{a}_2 \parallel \dots \parallel \mathbf{a}_n \parallel \mathbf{l}_1:\mathbf{t}_1 \parallel \dots \parallel \mathbf{l}_k:\mathbf{t}_k,$$

where the \mathbf{a}_i 's are distinct variables, the \mathbf{l}_j 's are distinct labels, the \mathbf{t}_j 's are in normal form, and n and k are greater than or equal to 0 (when both are 0, the normal form is **Empty**). By "has the form ..." we mean that \mathbf{r} has the above form up to associativity and commutativity of \parallel .

2. Let \mathbf{R} be a well-formed constraint list. Then \mathbf{R} has the form

$$\mathbf{a}_1, \dots, \mathbf{a}_n, \mathbf{l}_1, \dots, \mathbf{l}_k$$

where the \mathbf{a}_i 's are all distinct, the \mathbf{l}_j 's are all distinct, and n and k are greater than or equal to 0. By "has the form ..." we mean that \mathbf{R} is a list of variables and labels in some order.

Note that as a consequence, a well-formed normal form is restriction-free. This implies that every well-formed record expression reduces to a well-formed restriction-free record expression.

Lemma 2.3.3.15: The rule

$$\mathbf{r}, \mathbf{r}, \mathbf{R} \equiv \mathbf{r}, \mathbf{R} \quad (\text{A-CEQ-FULL-DUPL})$$

is admissible when $\mathbf{r}, \mathbf{r}, \mathbf{R}$ and \mathbf{r}, \mathbf{R} are well formed.

Proof sketch: By construction. Consider a reduction of \mathbf{r}, \mathbf{R} to normal form. By Theorem 2.3.3.14, $(\mathbf{r}, \mathbf{R})^*$ consists only of type variables and labels. We may therefore use rules A-CEQ-DUPL and A-CEQ-MERGE to duplicate the sublist of type variables and labels that come from \mathbf{r} . By replaying the reduction from \mathbf{r}, \mathbf{R} to $(\mathbf{r}, \mathbf{R})^*$ twice in reverse, we obtain $\mathbf{r}, \mathbf{r}, \mathbf{R} \equiv (\mathbf{r}, \mathbf{R})^*$. The result follows by transitivity of \equiv . \square

Theorem 2.3.3.16: The relations \sim and \equiv coincide on well-typed terms.

Proof: It is easy to see that every \rightarrow^1 and \approx^1 axiom is a valid equivalence, and conversely (using Lemma 2.3.3.15 in the case of EQ-DUPL), that every equivalence axiom holds in \equiv . \square

Corollary 2.3.3.17: The relation \sim is decidable for well-formed expressions.

Proof: To decide whether or not $\mathbf{r} \sim \mathbf{s}$, given that \mathbf{r} and \mathbf{s} are well-formed, reduce each to their normal forms and test whether or not $\mathbf{r}^* \sim \mathbf{s}^*$. Given the characterization of normal forms given above, this reduces to re-grouping the atomic record types in some standard order. Similar methods apply to constraint sets and types. \square

$$\begin{array}{c}
T \vdash r \# \Rightarrow \text{Empty} \quad (\text{A-CMP-EMPTY-1}) \\
T \vdash \text{Empty} \# \Rightarrow r \quad (\text{A-CMP-EMPTY-2}) \\
\frac{T \vdash s_1 \# \Rightarrow r \quad T \vdash s_2 \# \Rightarrow r}{T \vdash (s_1 \parallel s_2) \# \Rightarrow r} \quad (\text{A-CMP-MERGE-I-1}) \\
\frac{T \vdash r \# \Rightarrow s_1 \quad T \vdash r \# \Rightarrow s_2}{T \vdash r \# \Rightarrow (s_1 \parallel s_2)} \quad (\text{A-CMP-MERGE-I-2}) \\
\frac{l \neq l'}{T \vdash l:t \# \Rightarrow l':t'} \quad (\text{A-CMP-BASE/BASE}) \\
\frac{a \in S^*}{T_1, a\#R, T_2, b\#S, T_3 \vdash a \# \Rightarrow b} \quad (\text{A-CMP-TVAR/TVAR-1}) \\
\frac{a \in S^*}{T_1, a\#R, T_2, b\#S, T_3 \vdash b \# \Rightarrow a} \quad (\text{A-CMP-TVAR/TVAR-2}) \\
\frac{l \in R^*}{T_1, a\#R, T_2 \vdash a \# \Rightarrow l:t} \quad (\text{A-CMP-BASE/TVAR-1}) \\
\frac{l \in R^*}{T_1, a\#R, T_2 \vdash l:t \# \Rightarrow a} \quad (\text{A-CMP-BASE/TVAR-2})
\end{array}$$

Figure 6: Algorithmic compatibility rules

2.3.4 Compatibility Checking

The compatibility checking algorithm is presented as a collection of inference rules for deriving judgements of the form $T \vdash r \# \Rightarrow s$ where r and s are restriction-free, well-formed record types. The rules appear in Figure 6. To prove that these rules define an algorithm for compatibility checking, we show that they are sound and complete with respect to the declarative formulation, and that we may effectively decide whether or not a derivation exists in accordance with these rules.

The soundness of the algorithm is stated by the following theorem:

Theorem 2.3.4.1:

1. If $T \vdash r$ record and $T \vdash r'$ record and $T \vdash r \# \Rightarrow r'$, then $T \vdash r \# r'$.
2. If $T \vdash r$ record and $T \vdash R$ ok and $T \vdash r \# \Rightarrow R$, then $T \vdash r \# R$.

The proof relies on the following facts:

Lemma 2.3.4.2: If $T \vdash R$ ok and $R \rightarrow^1 R'$ and $T \vdash r \# R'$, then $T \vdash r \# R$.

Lemma 2.3.4.3: If $\vdash T$ ok and $T = T_1, a\#R, T_2, b\#S, T_3$ and $a \in S^*$, then $T \vdash a \# b$.

Lemma 2.3.4.4: If $\vdash T$ ok and $T = T_1, a\#R, T_2$ and $l \in R^*$ and $T \vdash l:t$ record, then $T \vdash a \# l:t$.

The completeness of the algorithm is stated by the following theorem:

Theorem 2.3.4.5:

1. If $T \vdash r \# s$, then $T \vdash r^* \# \Rightarrow s^*$.
2. If $T \vdash r \# R$, then $T \vdash r^* \# \Rightarrow R^*$.

The proof relies on a number of preliminary results.

Lemma 2.3.4.6: The relation computed by the algorithm is symmetric on normal forms:

1. Suppose that r and s are well-formed record types in normal form and that $r \approx r'$ and $s \approx s'$. Then $T \vdash r \# \Rightarrow s$ iff $T \vdash r' \# \Rightarrow s'$.
2. Suppose that r and s are well-formed record types in normal form. If $T \vdash r \# \Rightarrow s$, then $T \vdash s \# \Rightarrow r$.

Proof:

1. Since r and s are in normal form, r' and s' must also be, and r' and s' must be a commutative, associative re-grouping of r and s .
2. A simple induction on derivations. Box

Lemma 2.3.4.7: Suppose that r and t are well-formed normal forms. Then $T \vdash r \# \Rightarrow 1:t$ implies $T \vdash r \# 1:t'$ for any well-formed type t' .

Proof: Consider the form of r , and note that the property holds for atomic record types. □

Lemma 2.3.4.8: If $T \vdash s_1 || s_2$ record and $T \vdash r^* \# \Rightarrow s_1^*$ and $T \vdash r^* \# \Rightarrow s_2^*$, then $T \vdash r^* \# \Rightarrow (s_1 || s_2)^*$.

Proof: By a straightforward induction on a reduction sequence from $s_1 || s_2$ to $(s_1 || s_2)^*$. □

Definition 2.3.4.9: Let $\text{atoms}(r^*)$ be

1. the empty set if r^* is `Empty`,
2. otherwise, the result of replacing each atom of the form $1:t$ by the bare label 1 in the set of atoms of the normal form of r .

Lemma 2.3.4.10: If $T \vdash R$ ok and $r \in R$, then $\text{atoms}(r^*) \subseteq R^*$.

Proof: If r^* is `Empty`, the result is immediate. So assume r^* is not `Empty`. Since the relation \rightarrow is confluent modulo \approx , we may focus on any reduction path from R to a \rightarrow -normal form R'' such that $R'' \approx R^*$. Begin by reducing r to r^* in R , giving R' . Then flatten r^* in R' to its constituent atoms, giving R'' , and perform all the base reductions $1:t \rightarrow 1$ in the residuals of r in R'' , giving R''' . Observe that $\text{atoms}(r^*) \subseteq R'''$. Finally, rewrite R''' to R^* using the axioms for \approx , noting that no bare label or type variable is dropped from R''' (considered as a set) during any of these reductions. □

Lemma 2.3.4.11: Let $T = T_1, a \# R, T_2$, and suppose that $\vdash T$ ok and $r \in R$ (so that $T \vdash a \# r$.) Then $T \vdash a \# \Rightarrow r^*$.

Proof: By induction on the form of r^* , using Lemma 2.3.4.10. □

It remains to show that the relation axiomatized by the rules of Figure 6 is decidable. This follows from two observations. First, the rules are “almost syntax-directed” in the sense that there is a pair of rules for each form of restriction-free record type, one the symmetric version of the other. Second, in the cases where two rules apply to a given pair of types, it is immaterial which rule is applied. To see this we have to consider four combinations. If both r and s are `Empty`, either rule yields the same result (success). If one of r and s is `Empty` and the other is a merge, then no matter how we proceed the result is the same (success). In the case that both r and s are merges, we have no choice but to apply `A-CMP-MERGE-I-1` and `A-CMP-MERGE-I-2` to break down both merges into their components; the order in which we do this has no effect on the result. We may therefore check compatibility by decomposing merges in an arbitrary order until we reach atoms, then apply the appropriate rules. If any derivation exists, this procedure will construct it, and we can effectively decide whether any rule applies.

$$\begin{array}{c}
T_1, a \# R, T_2 \vdash a \Rightarrow \text{record} \quad (\text{A-WFR-VAR}) \\
\\
\frac{T \vdash r_1 \Rightarrow \text{record} \quad T \vdash r_2 \Rightarrow \text{record} \quad T \vdash r_1 \#^{\Rightarrow} r_2}{T \vdash r_1 \parallel r_2 \Rightarrow \text{record}} \quad (\text{A-WFR-MERGE}) \\
\\
\frac{T; G \vdash e_1 \Rightarrow r_1 \quad T; G \vdash e_2 \Rightarrow r_2 \quad T \vdash r_1^* \#^{\Rightarrow} r_2^*}{T; G \vdash e_1 \parallel e_2 \Rightarrow r_1 \parallel r_2} \quad (\text{A-WTT-MERGE}) \\
\\
\frac{T; G \vdash e \Rightarrow \forall a \# R. t \quad T \vdash r \Rightarrow \text{record} \quad T \vdash r^* \#^{\Rightarrow} R^*}{T; G \vdash e[r] \Rightarrow \{r/a\}t} \quad (\text{A-WTT-ALL-E})
\end{array}$$

Figure 7: Selected type synthesis rules

2.3.5 Type Synthesis

The type checking algorithm for λ^{\parallel} is given in terms of a type synthesis procedure that constructs a “canonical” type for a given expression in a given context. This procedure is described by a formal system for deriving judgements of the form $T; G \vdash e \Rightarrow t$, together with a number of auxiliary judgements of a similar form. As with the compatibility checker, we show that this formal system defines a type checking algorithm by proving that it is sound and complete with respect to the definition of λ^{\parallel} , and that we may effectively decide whether or not a derivation exists. A representative set of rules from the definition of the type synthesis algorithm is given in Figure 7. The complete set of rules appears in Appendix B.

The judgements derived by the algorithm are:

$$T \vdash R \Rightarrow \text{ok} \quad T \vdash t \Rightarrow \text{type} \quad T \vdash r \Rightarrow \text{record} \quad T; G \vdash e \Rightarrow t.$$

Note that in rules A-WTT-MERGE and A-WTT-ALL-E, the record types and constraint lists are normalized before checking compatibility.

The soundness and completeness of the algorithm is stated by the following theorems:

Theorem 2.3.5.1: (Soundness)

1. If $T \vdash T \text{ ok}$ and $T \vdash r \Rightarrow \text{record}$, then $T \vdash r \text{ record}$.
2. If $T \vdash T \text{ ok}$ and $T \vdash t \Rightarrow \text{type}$, then $T \vdash t \text{ type}$.
3. If $T \vdash T \text{ ok}$ and $T \vdash R \Rightarrow \text{ok}$, then $T \vdash R \text{ ok}$.
4. If $T \vdash G \text{ ok}$ and $T; G \vdash e \Rightarrow t$, then $T; G \vdash e \in t$.

Theorem 2.3.5.2: (Completeness)

1. If $T \vdash r \text{ record}$, then $T \vdash r \Rightarrow \text{record}$.
2. If $T \vdash t \text{ type}$, then $T \vdash t \Rightarrow \text{type}$.
3. If $T \vdash R \text{ ok}$, then $T \vdash R \Rightarrow \text{ok}$.
4. If $T \vdash e \in t$, then $T \vdash e \Rightarrow t'$ for some t' such that $t' \sim t$.

Both may be proved by induction on derivations.

For decidability, we have only to note that the relevant instances of compatibility and conversion checking are decidable, and that the rules are syntax-directed.

3 Examples

In this section we develop some examples illustrating the expressiveness of λ^{\parallel} . Section 3.1 presents some short examples and compares them to their formulations in related record calculi. Section 3.2 shows how Wand’s model of a simple object-oriented programming language with method inheritance and a notion of *self* may be emulated in λ^{\parallel} . Section 3.3 shows that in the presence of recursive types, quantifiers carrying negative constraints are sufficient to satisfactorily express the examples motivating the extension of bounded quantification to F-bounded quantification. Section 3.4 shows how some other primitive record operations may be added to the language.

3.1 Simple Record Manipulation

A function that accepts any record not already possessing an *l* field and adds the field *l*=5 can be written in λ^{\parallel} as:

$$\begin{aligned} &\Lambda a \# l. \\ &\quad \lambda x : a. \\ &\quad \quad x \parallel l = 5 \\ &\in \forall a \# l. a \rightarrow (a \parallel l : \text{Int}) \end{aligned}$$

A similar function that accepts any record *with* an *l* field and overrides it with the field *l*=5 can also be expressed:

$$\begin{aligned} &\Lambda b. \Lambda a \# l. \\ &\quad \lambda x : (a \parallel l : b). \\ &\quad \quad (x \setminus l) \parallel l = 5 \\ &\in \forall b. \forall a \# l. (a \parallel l : b) \rightarrow (a \setminus l) \parallel l : \text{Int} \end{aligned}$$

(Strictly speaking, this function can only accept a record whose *l* field has some *record* type, because for simplicity we have omitted ordinary type quantification. The empty constraint list on the first quantifier is dropped, by convention, to reduce clutter, but it should not be confused with the usual unconstrained quantifier.)

Unlike the calculi of Rémy, Wand, and Cardelli and Mitchell, λ^{\parallel} is unable to express the function that takes any record whatsoever and gives it an *l* field with value 5.

Functions that perform “deep updates” of fields within fields of records, preserving all the type information about the unmodified fields, can be expressed in λ^{\parallel} . For example:

$$\begin{aligned} &\forall a \# l_a. \\ &\quad \forall b \# l_b. \\ &\quad \quad \lambda x : (b \parallel l_b : (a \parallel l_a : \text{Bool})). \\ &\quad \quad \quad (x \setminus l_b) \parallel l_b = ((x.l_b) \setminus l_a \parallel l_a = (\text{not } x.l_b.l_a)) \end{aligned}$$

Cardelli and Mitchell’s formulation of this function, which relies on *extraction types* like *a.l*, is somewhat shorter.

3.2 Object-Oriented Programming

Wand [24, 25, 27] has shown how a simple form of object-oriented programming may be modeled in an extension of ML [14, 15] with record concatenation and recursive types. An analogous construction can be made in λ^{\parallel} .

We assume the existence of a number of extensions of the basic type system, including polymorphism (quantification over arbitrary types), fixed point combinators, recursive types, and base types such as *int* and *bool* (and associated operations). Of these, only recursive types present any serious problems. Recursive

types are denoted $\mu a. t(a)$, where $t(a)$ is an arbitrary type expression, possibly involving the variable a . Let $\lambda^{||\mu}$ stand for the extension of $\lambda^{||}$ with recursive types. Type equivalence for $\lambda^{||\mu}$ is defined to be the greatest fixed point of a monotone operator Φ on binary relations on types determined by the inference rules for type equivalence given in Appendix A, augmented by the axiom $\mu a. t(a) \sim t(\mu a. t(a))$. This characterization of type equality captures the informal notion that recursively-defined types stand for their “infinite unrollings” as regular trees. In the examples we make use of the following useful characterization of type equivalence: to show that $s \sim t$, it suffices to assume that $s \sim t$ and prove that $s \Phi(\sim) t$. Although this method is useful in the examples, we have not studied the decidability of this notion of type equivalence, and hence the decidability of type checking for $\lambda^{||\mu}$ remains open.

Following Wand, we consider first a very simple notion of class that does not include a notion of **self** or provide for inheritance. In this setting classes are functions from a parameter type, representing instance-specific values, to some record type, representing the methods. Specifically,

```
class x:t methods l1=e1, ..., ln=en
```

stands for the expression

```
 $\lambda x:t. (l_1=e_1 \parallel \dots \parallel l_n = e_n)$ 
```

which has type

```
 $t \rightarrow (l_1:t_1 \parallel \dots \parallel l_n:t_n)$ 
```

provided that t is a type, $x:t \vdash e_i \in t_i$, and all the l_i 's are distinct. Let C be such a class. If e is a term of type t , then $C(e)$ is a record of type $l_1:t_1 \parallel \dots \parallel l_n:t_n$, so that $C(e).l_i$ selects the method associated with the instance of C determined by e . Following Wand we write

```
new C e
```

for the application $C(e)$.

Wand's list-processing example may be readily translated into this setting as follows. First, we introduce some abbreviations:

```
cellPlus[r,s,t] = r || null:Bool || car:s || cdr:t
cell[s,t]       = cellPlus[Empty,s,t]
listPlus[r,s]   =  $\mu a. cellPlus[r,s,a]$ 
list[s]         = listPlus[Empty,s].
```

Note that

```
listPlus[r,s] ~ cellPlus[r,s,listPlus[s]]
              ~ r || null:Bool || car:s || cdr:listPlus[r,s]
```

and that

```
list[s] ~ cell[s,list[s]]
        ~ null:Bool || car:s || cdr:list[s].
```

These equations will be important for type checking the examples below.

Define

```
nil =  $\Lambda s,t. \lambda x:l. (null=true \parallel car=error[s] \parallel cdr =error[t])$ 
      $\in \forall s,t. l \rightarrow cell[s,t]$ .
```

Note that **nil** is parameterized by the types of the cells. We assume a special constant **error** $\in \forall t. t$ assigning an “error element” to each type. (There are other possible solutions. In a richer calculus, evaluating **nil.car** could raise an exception, for example.)

Define

```
cons =  $\Lambda s,t. \lambda(h,t):s \times t. (null=false \parallel car=h \parallel cdr=t)$ 
       $\in \forall s,t. (s \times t) \rightarrow cell[s,t]$ .
```

Now consider

```
map =  $\Lambda s,t. \Lambda r \# null, car, cdr. \lambda f:s \rightarrow t.$ 
      fixu  $\lambda map'u. \lambda l:listPlus[r,s].$ 
        if l.null
          then new nil[t,list[t]]()
          else new cons[t][list[t]](f l.car)(map' l.cdr)
```

where $u = \text{listPlus}[r, s] \rightarrow \text{list}[t]$.

It is easy to see that map has type

$$\forall s, t. \forall r \neq \text{null}, \text{car}, \text{cdr}. (s \rightarrow t) \rightarrow \text{listPlus}[r, s] \rightarrow \text{list}[s]$$

under some obvious assumptions about typing conditionals and taking account of the equations between recursive types noted above.

Extending the example to a simple, self-less form of multiple inheritance presents no particular difficulties:

$$\text{class } x:t \text{ inherits } f_1, \dots, f_k \text{ methods } l_1=e_1, \dots, l_n=e_n \text{ end}$$

stands for

$$\lambda x:t. (f_1 \parallel \dots \parallel f_k \parallel l_1=e_1 \parallel \dots \parallel l_n=e_n)$$

which has type

$$t \rightarrow (F_1 \parallel \dots \parallel F_k \parallel l_1:t_1 \parallel \dots \parallel l_n:t_n)$$

provided that t is well formed, $x:t \vdash f_j \in F_j$, $x:t \vdash e_i \in t_i$, and the following compatibility conditions are satisfied:

1. all l_i 's are distinct,
2. $F_i \neq F_j$ whenever $i \neq j$,
3. $F_i \neq l_j$ for each F_i and each label l_j .

These conditions are not enforced in Wand's system since he prefers the asymmetric form of merge in which the rightmost occurrence of a field in a record type determines the value assigned to that field. We prefer instead that the type checker ensure that no clashes are possible, and insist that the programmer explicitly eliminate the field that he or she wishes to override using the restriction operation.

The extension to model "self" presents some interesting problems for typing. We begin by considering a form of **self** that does not behave correctly in the presence of inheritance and then consider how to combine **self** and inheritance. The simple-minded approach to **self** proceeds by defining

$$\text{class } x:s \text{ methods } l_1=e_1, \dots, l_n=e_n \text{ end}$$

to be

$$\lambda x:s. \lambda \text{self}:t. (l_1=e_1 \parallel \dots \parallel l_n=e_n)$$

which has type $s \rightarrow t \rightarrow t$, where $t = (l_1:t_1 \parallel \dots \parallel l_n:t_n)$, provided that t is a type and $x:t, \text{self}:t \vdash e_i : t_i$.

Class instantiation is performed by taking a fixed point, as described in Wand's paper. If C is a class of the form just described, then

$$\text{new } C \ e$$

is defined to be

$$\text{fix}_t (C(e)),$$

where C has type $s \rightarrow t \rightarrow t$ and e has type s .

This simple-minded approach interacts badly with inheritance in two ways. First, the type assigned to **self** is required to coincide with the type of the body of the class definition. This forces us to define at class definition time all methods that are used by any method of the class. This precludes the possibility of "mixing in" the missing methods at a later stage, a significant feature of object-oriented programming. Second, the type assigned to **self** does not take account of the fact that the class may occur as the superclass of some subclass which may define not only the "missing" methods, but also some additional methods not relevant to the definition of the superclass. In short, **self** should refer to the entire subclass, and not just the superclass or even just the subclass of the superclass required to achieve "closure."

This means that a class definition must have a "floating" notion of **self** that can be instantiated at any subclass consistent with the requirements imposed on **self** and with the methods that are given for **self** as part of the class definition. Following Wand, we take

$$\text{class } x:t \text{ inherits } f_1, \dots, f_k \text{ methods } l_1=e_1, \dots, l_n=e_n \text{ end}$$

to stand for (roughly)

$$\lambda x:t. \text{fix}_v (\lambda s:v. (f_1(s) \parallel \dots \parallel f_2(s) \parallel l_1=e_1 \parallel \dots \parallel l_n=e_n))$$

where v is a suitable type expression. This ensures that the f_i 's and the new class share a common sense of **self**.

Since our system is explicitly-typed, we must take explicit account of the variability of the type of **self** by introducing suitable constrained type abstractions. The main idea is to assign to abstract with respect to a record type variable a representing an arbitrary "extension" to the type of record comprising the object, and to assign to **self** a type of the form $a \parallel l_1:t_1 \parallel \dots \parallel l_n:t_n$, where the $l_i:t_i$'s include both the types of the methods actually defined by the class and the types of the methods assumed by the new methods. By choosing the type a appropriately we may instantiate **self** to be any superclass of the class, as required for inheritance.

A class definition of the form

$$\text{class } x:t \text{ methods } l_1=e_1, \dots, l_n=e_n \text{ end}$$

is represented by a λ -term of the form

$$\begin{aligned} \lambda x:t. \\ \Lambda a \# l_1, \dots, l_n. \\ \lambda s:(a \parallel l_1:t_1 \parallel \dots \parallel l_n:t_n). \\ (l_{i_1}=e_{i_1} \parallel \dots \parallel l_{i_r}=e_{i_r}), \end{aligned}$$

where $\{i_1, \dots, i_r\} \subseteq 1 \dots n$. This expression has type

$$\begin{aligned} t \rightarrow \forall a \# l_1, \dots, l_n. \\ (a \parallel l_1:t_1 \parallel \dots \parallel l_n:t_n) \\ \rightarrow (l_{i_1}:t_{i_1} \parallel \dots \parallel l_{i_r}:t_{i_r}) \end{aligned}$$

provided that t type and $x:t$, $a \# l_1, \dots, l_n$ and $s:a \parallel l_1:t_1 \parallel \dots \parallel l_n:t_n \vdash e_j \in t_j$.

To build a subclass of such a class C , we proceed as follows. Let us suppose that we intend to build a class D , and instantiate it at some value v . Therefore we must at least define the methods assumed in the definition of C , and perhaps some additional methods particular to D . Thus we will want to use a definition of the form

$$\text{class } x:t \text{ inherits new } C \text{ e methods } l_{j_1}=e_{j_1} \parallel \dots \parallel l_{j_q}=e_{j_q} \parallel k_1=d_1 \parallel \dots \parallel k_p=d_p \text{ end}$$

where $l_{i_1}, \dots, l_{i_r}, l_{j_1}, \dots, l_{j_q} = l_1 \dots l_n$ (so that the l_j 's supply the missing methods and the k_i 's are new methods associated with D). This corresponds to the following lambda term:

$$\begin{aligned} \lambda x:t. \Lambda a \# l_1, \dots, l_n, k_1, \dots, k_p. \\ \lambda \text{self}:(a \parallel l_1:t_1 \parallel \dots \parallel l_n:t_n \parallel k_1:s_1 \parallel \dots \parallel k_p:s_p). \\ C(e)[a \parallel k_1:s_1 \parallel \dots \parallel k_p:s_p](\text{self}) \\ \parallel l_{j_1}=e_{j_1} \parallel \dots \parallel l_{j_q}=e_{j_q} \parallel k_1=d_1 \parallel \dots \parallel k_p=d_p \end{aligned}$$

which has type

$$\begin{aligned} t \rightarrow \forall a \# l_1, \dots, l_n, k_1, \dots, k_p. \\ (a \parallel l_1:t_1 \parallel \dots \parallel l_n:t_n \parallel k_1:s_1 \parallel \dots \parallel k_p:s_p) \\ \rightarrow (l_1:t_1 \parallel \dots \parallel l_n:t_n \parallel k_1:s_1 \parallel \dots \parallel k_p:s_p) \end{aligned}$$

so that $\lambda x:t. \text{fix}(D(x)[\text{Empty}]) \in t \rightarrow l_1:t_1 \parallel \dots \parallel k_p:s_p$ (assuming that the l_j 's supplied the missing methods).

Here is a concrete example, taken from Wand. Define the class A to be

$$\begin{aligned} \lambda x:\text{int}. \\ \Lambda a \# \text{sum}, n. \\ \lambda \text{self}:(a \parallel \text{sum}:\text{int} \parallel n:\text{int}). \\ \text{sum}=(x + \text{self}.n). \end{aligned}$$

with type

$$\text{int} \rightarrow \forall a \# \text{sum}, n. (a \parallel \text{sum}:\text{int} \parallel n:\text{int}) \rightarrow (\text{sum}:\text{int}).$$

This class cannot be instantiated since it lacks a method for n . Define the class B to be a subclass of A that provides a method for n :

```

λy:int.
  Λa#sum,n.
    λself:(a|sum:int|n:int).
      A(5)[a](self) || n=y

```

with type

$$\text{int} \rightarrow \forall a\#sum,n. (a|\text{sum:int}|n:\text{int}) \rightarrow (\text{sum:int}|n:\text{int}).$$

Then $B(3)[\text{Empty}] \in \text{sum:int}|n:\text{int} \rightarrow \text{sum:int}|n:\text{int}$, so that $\text{fix}(B(3)[\text{Empty}]).\text{sum} = 8$.

We may also define a class C that extends A with a method for n and a method for m , as follows:

```

λy:int.
  Λa#sum,n,m.
    λself:(a|sum:int|n:int|m:int).
      A(5)[a|m:int](self) || n=y || m=10

```

so that $\text{fix}(C(3)[\text{Empty}])$ makes sense, and is an object whose sum and n fields agree with B , and which also has an m field with value 10.

It is worth noting that the translations of the object-oriented programming idioms into our calculus are to be understood only informally since the target expression involves information not apparent in the source expression. Moreover, it seems reasonable to impose additional typing constraints on classes beyond those that would arise out of their representation in $\lambda^{||\mu}$. For example, it seems sensible to insist that the type ascribed to self be consistent with the type of the body of the class. This suggests that in a more complete development of these ideas, special syntax for class operations would be needed.

One major difference between this model of object-oriented programming and Wand's is that Wand's record concatenation operator is asymmetric while that of $\lambda^{||\mu}$ is symmetric. This difference represents an important methodological commitment in the models of object oriented programming that arise from the two type systems.

Any language with multiple inheritance must address the issue of what happens when a class definition inherits the same method from two or more superclasses. One solution, adopted in many existing languages, is to choose one of the superclasses — say, the one mentioned furthest to the right in the subclass' definition — as the “principal superclass,” resolving all conflicts in favor of this one. Wand's model implements this kind of conflict resolution scheme.

Another possibility, preferred by some designers of strongly typed object-oriented languages, is simply to disallow class definitions where the same method is inherited ambiguously from more than one superclass. When describing a class whose superclasses have overlapping sets of method names, the programmer must explicitly specify which superclass each ambiguous method comes from. Our model can be used to implement this kind of conflict resolution strategy: the typing rules for the merge operator ensure that each method comes from exactly one superclass; the restriction operator is used to hide all but one of the inherited instances of a method.

3.3 F-Bounded Quantification

The ABEL group at HP-Labs has argued convincingly that “bounded quantification does not provide the same degree of flexibility in the presence of recursive types as it does for non-recursive types.” [1, 7]. They propose an extended notion, called *F-bounded quantification*, where, rather than just pure bounded quantifiers of the form $\forall a \leq r. t$, they allow quantifiers of the form $\forall a \leq F(a). t$, where F is a function from types to types.

In this section, we show that an analog of F -bounded quantification is *already* available in $\lambda^{||\mu}$.

In one class of situations, involving recursive record types where the recursion variable appears in negative positions, Canning et al. show that the pure type system of Cardelli and Wegner [6] does not allow functions to be applied to a variety of values for which they make semantic sense. For example, define the type

$$\text{PartialOrder} = \mu \text{po}. \{\text{leq} : \text{po} \rightarrow \text{Bool}\}$$

and assume we are given a function for computing the minimum of two values of any “subclass” of PartialOrder :

$\text{min} \in \forall a \leq \text{PartialOrder}. a \rightarrow a \rightarrow a.$

One of the types that we would like to be able to pass to min is

$\text{Number} = \mu \text{num}. \{\text{leq} : \text{num} \rightarrow \text{Bool}, \text{otherstuff} : \text{t}\}.$

But by the usual rule for subtyping on recursive types,

$$\frac{G, a \leq b \vdash s \leq t \quad \begin{array}{l} a \text{ free only in } s \\ b \text{ free only in } t \end{array}}{G \vdash \mu a. s \leq \mu b. t} \quad (\text{LEQ-REC})$$

it is not the case that $\text{Number} \leq \text{PartialOrder}$.

F-bounded quantification can be used to redefine min so that it can be applied to elements of Number as well as PartialOrder . Define a type function

$\text{FPartialOrder}(t) = \{\text{leq} : t \rightarrow \text{Bool}\}$

and check that $\text{Number} \leq \text{FPartialOrder}(\text{Number})$. Now write

$\text{min} = \Lambda a \leq \text{FPartialOrder}(a).$
 $\quad \lambda x : a. \lambda y : a.$
 $\quad \quad \text{if } x.\text{leq}(y) \text{ then } x \text{ else } y$
 $\in \forall a \leq \text{FPartialOrder}(a). a \rightarrow a \rightarrow a.$

[What if the instance of a is supposed to involve b ?

To express min in λ^{II} so that it can be applied to elements of Number , we write:

$\text{PartialOrder} = \mu \text{po}. \text{leq} : \text{po} \rightarrow \text{Bool}$
 $\text{min} \in \forall a \# \text{leq}. \mu b. (a \parallel \text{leq} : b \rightarrow \text{Bool})$
 $\quad \rightarrow \mu b. (a \parallel \text{leq} : b \rightarrow \text{Bool})$
 $\quad \rightarrow \mu b. (a \parallel \text{leq} : b \rightarrow \text{Bool})$
 $\text{Number} = \mu \text{num}. (\text{leq} : \text{num} \rightarrow \text{Bool} \parallel \text{otherstuff} : \text{t})$
 $\text{five} \in \text{Number}.$

To type the application

min five five

we need to restrict away the leq field from Number :

$\text{min } [\text{Number} \backslash \text{leq}] \text{ five five}.$

The type application is well formed because $\text{Number} \backslash \text{leq}$ certainly lacks leq . To apply this term to five , we need to know that

$\text{Number} \sim \mu b. (\text{Number} \backslash \text{leq}) \parallel \text{leq} : b \rightarrow \text{Bool},$

that is

$\text{Number} \sim \mu b. (\text{otherstuff} : [\text{Number}/\text{num}] \text{t}) \parallel \text{leq} : b \rightarrow \text{Bool},$

which may be readily verified using the proof technique suggested above. Specifically, let Number' be the right-hand side of the above equation, and assume that $\text{Number} \sim \text{Number}'$. To show that this is consistent, it suffices to unroll Number and Number' , and show that the corresponding fields of the resulting record types are equivalent. But this follows immediately from the assumption and the reflexivity of equivalence. The reasoning for the second application is identical.

The other class of situations where bounded quantification appears to be inadequate are those involving positive occurrences of the recursion variable in a recursively defined record type. Here again, pure bounded quantification prevents functions from being applied to a variety of values for which they make sense.

For example, define a recursive record type

$\text{Movable} = \mu \text{mv}. \{\text{move} : \text{Real} \rightarrow \text{Real} \rightarrow \text{mv}\}$

and a function

$\text{translate} = \lambda x : \text{Movable}. x.\text{move } 1.0 \ 1.0.$

Since `translate` is intended to operate on values possessing any “subclass” of “class” `Movable`, we want to show that it has type $\forall a \leq \text{Movable}. a \rightarrow a$. But it does not; the best that can be said for it is that it has type $\forall a \leq \text{Movable}. a \rightarrow \text{Movable}$, which is no better than $\text{Movable} \rightarrow \text{Movable}$. So if

`Point = μ pnt. {x:void \rightarrow Real, y:void \rightarrow Real, move:Real \rightarrow Real \rightarrow pnt}`

and

`mypoint \in Point,`

then the application

`translate mypoint`

will have type `Movable`, not `Point`.

Using F-bounded quantification, a `translate` function that will map elements of `Point` back into `Point` can be written as follows. Define the type function

`FMovable(t) = {move : Real \rightarrow Real \rightarrow t}.`

and write

`translate = Λ a \leq FMovable(a).
 λ x:a.
 x.move 1.0 1.0.`

Then since `Point \leq FMovable(Point)`, the application `translate [Point] mypoint` is well typed and has result type `Point`.

To express `translate` in λ^{\parallel} so that it can be applied to `mypoint`, we write:

`translate = Λ a#move.
 λ x:(μ b. a||(move : Real \rightarrow Real \rightarrow b)).
 x.move 1.0 1.0
 \in \forall a#move. (μ b. a||(move : Real \rightarrow Real \rightarrow b)) \rightarrow (μ b. a||(move : Real \rightarrow Real \rightarrow b)).`

Again, we apply `translate` to `mypoint` by first restricting away the `move` field of its type:

`translate [Point\move] mypoint.`

3.4 Language Extensions

This section shows how some other primitive record operations may be added to λ^{\parallel} .

The “consistent update” operator `e upd l=e'` takes a record `e` with an `l` and replaces the value of this field with `e'`, which must have exactly the same type as the present contents of the field. This operation cannot be expressed directly in λ^{\parallel} , but it can be provided easily by adding a new typing rule:

$$\frac{T; G \vdash e \in r \quad r.l \downarrow \quad T; G \vdash e' \in t \quad t \sim r.l}{T; G \vdash e \text{ upd } l=e' \in r} \quad (\text{WTT-UPD})$$

Similarly, field renaming can be provided by adding two new rules:

$$\frac{T \vdash r \text{ record} \quad r.l_1 \downarrow \quad r.l_2 \uparrow}{T; G \vdash r[l_1 \rightarrow l_2] \text{ record}} \quad (\text{WFT-RENAME})$$

$$\frac{r.l_1 \downarrow \quad r.l_2 \uparrow}{T; G \vdash r[l_1 \rightarrow l_2] \sim (r \setminus l_1) \parallel l_2:r.l_1} \quad (\text{EQ-RENAME})$$

At the level of values, syntactic sugar suffices for the field renaming operation:

$$e[l_1 \rightarrow l_2] = e \setminus l_1 \parallel l_1=(e.l_1).$$

The identity function on all records with *only* an `l` field requires ordinary type quantification:

$$\Lambda a. \Lambda b \#(l:a). \lambda x:b. x.l.$$

4 Future Work

The research described here suggests a number of intriguing paths for future work. Among those that we would like to pursue are:

Record kinds: At one stage in the development of λ^{\parallel} , we attempted to classify record types into *kinds* according to the record types with which they can safely be merged. Rather than “ $r \# R$,” we wrote “ $r : \text{rec}\#R$,” meaning “ r has the kind describing records compatible with the types in R .”

This reorganization has the advantage of allowing record type quantification to be unified cleanly with ordinary type quantification. Also, a form of subsumption — already implicit in λ^{\parallel} in the fact that a record type abstraction may be applied to a record type that happens to satisfy *more* constraints than those required by the quantifier — can be made explicit in this reformulation by introducing a “weaker than” ordering on record kinds.

In the end, the formulation used in this paper appeared to have a more tractable proof theory, so we chose to stick with it for this stage of our investigation. But the elegance of the formulation in terms of kinds remains attractive.

Alternative primitives: The definition of λ^{\parallel} is strongly committed to a particular choice of primitives: symmetric merge and “checked” restriction, where a record can only be restricted at a field that it definitely possesses. Substituting a different formulation of either of these operators would require making some changes to the definition of the system, and probably major changes to the algorithms for typechecking. Nevertheless, we believe that the methodology used to develop λ^{\parallel} could be effectively applied to similar systems with different — perhaps more useful — sets of primitives.

Type reconstruction: The prenex fragment of λ^{\parallel} , in which quantifiers are restricted to appear only outermost in type phrases, can be viewed as an explicitly typed version of Wand’s calculus with row variables, with the difference that the merge operator is symmetric and the restriction operator is checked. From this analogy, we believe that the problem of finding a well-typed term whose erasure is the same as a given untyped term in this fragment (type reconstruction) may be decidable, though a type reconstruction algorithm would almost certainly have exponential complexity because, like Wand’s, it would be inferring not principal types but finite sets of principal types.

Recursive types: In order to fully justify the examples inspired by F-bounded quantification in Section 3.3, the decidability of typechecking for $\lambda^{\parallel\mu}$ needs to be established. Typechecking for polymorphic λ -calculi in the presence of μ -types is beginning to be understood, but the complex equivalence relation already defined on types appears to complicate matters substantially.

Notation for object-oriented programming: From the examples inspired by Wand in Section 3.2, it appears that a type system like this one could be used as the basis of a fairly expressive object-oriented programming language. It would be interesting to try to formalize the informal “class” notation used in these examples and describe a rigorous translation into $\lambda^{\parallel\mu}$.

5 Conclusions

The decision to annotate quantifiers with purely positive information, with purely negative information, or with a mixture of positive and negative information is an important point of variation among calculi of record operations. Ordinary bounded quantification [6], F-bounded quantification [1], and the systems of Ohori and Buneman [17, 18] are positive-information systems. Cardelli and Mitchell’s calculus [4, 5] and our earlier “symmetric system” [10] are mixed positive and negative. Wand’s system of row variables [27] and λ^{\parallel} are pure negative-information systems.

The differences among these classes of systems are particularly clear in the presence of recursive types. In positive-information systems, one seems to need something like F-bounded quantification: in the case of our “symmetric system” this would be a something like a recursive “has” constraint. In the negative setting one does not need an analogue of F-bounded quantification, since one can explicitly quantify over

the “rest” of the fields in a record type. This leads to the observation that, in mixed systems like Cardelli and Mitchell’s, the negative-information fragment can be used to directly express the examples motivating F-bounded quantification. Indeed, the construction given in Section 3.3 can be carried out almost verbatim in Cardelli and Mitchell’s calculus.

6 Acknowledgements

We are grateful for productive discussions with Luca Cardelli and Didier Rémy.

A Complete Typing Rules

A.1 Judgement Forms

Well-formed type context:	$T \text{ ok}$
Well-formed term context:	$T \vdash G \text{ ok}$
Well-formed type:	$T \vdash t \text{ type}$
Well-formed record type:	$T \vdash r \text{ record}$
Well-formed constraint list:	$T \vdash R \text{ ok}$
Constraint list satisfaction:	$T \vdash r \# R$
Compatible types:	$T \vdash r_1 \# r_2$
Equivalent types:	$t_1 \sim t_2$
Equivalent constraint sets:	$R_1 \sim R_2$
Well-formed term:	$T; G \vdash e \in t$

A.2 Well-formed type contexts

$$\diamond \text{ ok} \quad (\text{WFCT-EMPTY})$$

$$\frac{T \vdash R \text{ ok}}{T, a \# R \text{ ok}} \quad (\text{WFCT-TVAR})$$

A.3 Well-formed constraint lists

$$\frac{T \text{ ok}}{T \vdash \diamond \text{ ok}} \quad (\text{WFCL-NIL})$$

$$\frac{T \vdash r \text{ record} \quad T \vdash R \text{ ok}}{T \vdash r, R \text{ ok}} \quad (\text{WFCL-CONS})$$

A.4 Well-formed term contexts

$$\frac{T \text{ ok}}{T \vdash \diamond \text{ ok}} \quad (\text{WFC-EMPTY})$$

$$\frac{T \vdash G \text{ ok} \quad T \vdash t \text{ type}}{T \vdash G, x:t \text{ ok}} \quad (\text{WFC-VAR})$$

A.5 Well-formed types:

$\frac{T \text{ ok}}{T \vdash p \text{ type}}$	(WFT-PRIM)
$\frac{T \vdash t_1 \text{ type} \quad T \vdash t_2 \text{ type}}{T \vdash t_1 \rightarrow t_2 \text{ type}}$	(WFT-ARROW)
$\frac{T, a \# R \vdash t \text{ type}}{T \vdash \forall a \# R. t \text{ type}}$	(WFT-ALL)
$\frac{T \vdash r \text{ record}}{T \vdash r \text{ type}}$	(WFT-REC)

A.6 Well-formed record types

$\frac{T_1, a \# R, T_2 \text{ ok}}{T_1, a \# R, T_2 \vdash a \text{ record}}$	(WFR-VAR)
$\frac{T \vdash t \text{ type}}{T \vdash l:t \text{ record}}$	(WFR-BASE)
$\frac{T \vdash r \text{ record} \quad r.l \downarrow}{T \vdash r.l \text{ record}}$	(WFR-RESTR)
$\frac{T \text{ ok}}{T \vdash \text{Empty record}}$	(WFR-EMPTY)
$\frac{T \vdash r_1 \# r_2}{T \vdash r_1 r_2 \text{ record}}$	(WFR-MERGE)

A.7 Constraint list satisfaction

$\frac{T \vdash r \text{ record}}{T \vdash r \# \circ}$	(CMP-LIST-NIL)
$\frac{T \vdash r \# r_i \quad T \vdash r \# R}{T \vdash r \# r_i, R}$	(CMP-LIST-CONS)

A.8 Compatibility

$\frac{T \vdash r \# s \quad r \sim r' \quad s \sim s'}{T \vdash r' \# s'}$	(CMP-EQ)
$\frac{T \vdash r \# s}{T \vdash s \# r}$	(CMP-SYMM)
$\frac{T \vdash r \# l:t \quad T \vdash t' \text{ type}}{T \vdash r \# l:t'}$	(CMP-BASE)
$\frac{T_1, a \# R, T_2 \text{ ok} \quad r_i \in R}{T_1, a \# R, T_2 \vdash a \# r_i}$	(CMP-TVAR)
$\frac{T \vdash r \# (s_1 s_2)}{T \vdash r \# s_i}$	(CMP-MERGE-E)

$$\frac{T \vdash s_1 \# s_2 \quad T \vdash r \# s_1 \quad T \vdash r \# s_2}{T \vdash r \# (s_1 || s_2)} \quad (\text{CMP-MERGE-I})$$

$$\frac{l \neq l' \quad T \vdash l:t \text{ record} \quad T \vdash l':t' \text{ record}}{T \vdash l:t \# l':t'} \quad (\text{CMP-BASE/BASE})$$

$$\frac{T \vdash r \text{ record}}{T \vdash r \# \text{Empty}} \quad (\text{CMP-EMPTY})$$

A.9 Constraint list equivalence

$$R \sim R \quad (\text{CEQ-REFL})$$

$$\frac{R \sim R'}{R' \sim R} \quad (\text{CEQ-SYMM})$$

$$\frac{R \sim R' \quad R' \sim R''}{R \sim R''} \quad (\text{CEQ-TRANS})$$

$$\frac{R \sim R' \quad r \sim r'}{r, R \sim r', R'} \quad (\text{CEQ-INNER})$$

$$r, (r', R) \sim r', (r, R) \quad (\text{CEQ-SWAP})$$

$$\text{Empty}, R \sim R \quad (\text{CEQ-EMPTY})$$

$$(r_1 || r_2), R \sim r_1, (r_2, R) \quad (\text{CEQ-MERGE})$$

$$r, (r, R) \sim r, R \quad (\text{CEQ-DUPL})$$

$$l:t, R \sim l:t', R \quad (\text{CEQ-BASE})$$

A.10 Type equivalence

$$r || \text{Empty} \sim r \quad (\text{EQ-MERGE-UNIT})$$

$$r_1 || (r_2 || r_3) \sim (r_1 || r_2) || r_3 \quad (\text{EQ-MERGE-ASSOC})$$

$$r_1 || r_2 \sim r_2 || r_1 \quad (\text{EQ-MERGE-COMM})$$

$$r \setminus l \setminus l' \sim r \setminus l' \setminus l \quad (\text{EQ-RESTR/RESTR'})$$

$$(l:t) \setminus l \sim \text{Empty} \quad (\text{EQ-BASE/RESTR})$$

$$\frac{r_1 \setminus l \setminus l'}{(r_1 || r_2) \setminus l \sim (r_1 \setminus l \setminus l' || r_2)} \quad (\text{EQ-MERGE/RESTR})$$

A.11 Type equivalence – congruence rules:

$t \sim t$	(Eq-REFL)
$\frac{t' \sim t}{t \sim t'}$	(Eq-SYMM)
$\frac{t \sim t' \quad t' \sim t''}{t \sim t''}$	(Eq-TRANS)
$\frac{t_1 \sim t'_1 \quad t_2 \sim t'_2}{t_1 \rightarrow t_2 \sim t'_1 \rightarrow t'_2}$	(Eq-CONG-ARROW)
$\frac{R \sim R' \quad t \sim t'}{\forall a \# R. t \sim \forall a \# R'. t'}$	(Eq-CONG-ALL)
$\frac{t \sim t'}{l:t \sim l:t'}$	(Eq-CONG-BASE)
$\frac{r \sim r'}{r \setminus l \sim r' \setminus l}$	(Eq-CONG-RESTR)
$\frac{r_1 \sim r'_1 \quad r_2 \sim r'_2}{r_1 \parallel r_2 \sim r'_1 \parallel r'_2}$	(Eq-CONG-MERGE)

A.12 Well-typed terms

$\frac{T; G \vdash e \in t \quad T \vdash t' \text{ type} \quad t \sim t'}{T; G \vdash e \in t'}$	(WTT-EQ)
$\frac{T \vdash G_1, x:t, G_2 \text{ ok}}{T; G_1, x:t, G_2 \vdash x \in t}$	(WTT-VAR)
$\frac{T; G, x:t \vdash e \in t'}{T; G \vdash \lambda x:t. e \in t \rightarrow t'}$	(WTT-ARROW-I)
$\frac{T; G \vdash e_1 \in t \rightarrow t' \quad T; G \vdash e_2 \in t}{T; G \vdash e_1 e_2 \in t'}$	(WTT-ARROW-E)
$\frac{T; G \vdash e \in t}{T; G \vdash l=e \in l:t}$	(WTT-BASE)
$\frac{T; G \vdash e \in r \quad r \cdot l \downarrow}{T; G \vdash e \setminus l \in r \setminus l}$	(WTT-RESTR)
$\frac{T \vdash G \text{ ok}}{T; G \vdash \text{empty} \in \text{Empty}}$	(WTT-EMPTY)
$\frac{T; G \vdash e_1 \in r_1 \quad T; G \vdash e_2 \in r_2 \quad T \vdash r_1 \# r_2}{T; G \vdash e_1 \parallel e_2 \in r_1 \parallel r_2}$	(WTT-MERGE)
$\frac{T; G \vdash e \in r \quad r \cdot l \downarrow}{T; G \vdash e.l \in r.l}$	(WTT-SELECT)
$\frac{T, a \# R; G \vdash e \in t}{T; G \vdash \Lambda a \# R. e \in \forall a \# R. t}$	(WTT-ALL-I)
$\frac{T; G \vdash e \in \forall a \# R. t \quad T \vdash r \# R}{T; G \vdash e[r] \in [r/a]t}$	(WTT-ALL-E)

B Complete Typechecking Rules (Algorithmic Version)

B.1 Judgement forms

Well-formed type:	$T \vdash t \Rightarrow \text{type}$
Well-formed record type:	$T \vdash r \Rightarrow \text{record}$
Well-formed constraint list:	$T \vdash R \Rightarrow \text{ok}$
Constraint list satisfaction:	$T \vdash r \# \Rightarrow R$
Compatible types:	$T \vdash r_1 \# \Rightarrow r_2$
Well-formed term:	$T; G \vdash e \Rightarrow t$

B.2 Type and constraint list normalization

$$\text{Empty} \parallel r \rightarrow^1 r \quad (\text{A-RED-EMPTY-1})$$

$$r \parallel \text{Empty} \rightarrow^1 r \quad (\text{A-RED-EMPTY-2})$$

$$l:t \setminus l \rightarrow^1 \text{Empty} \quad (\text{A-RED-RESTR/BASE})$$

$$(r \parallel s) \setminus l \rightarrow^1 (r \setminus l) \parallel s \quad \text{if } r _l \downarrow \quad (\text{A-RED-RESTR/MERGE-1})$$

$$(r \parallel s) \setminus l \rightarrow^1 r \parallel (s \setminus l) \quad \text{if } s _l \downarrow \quad (\text{A-RED-RESTR/MERGE-2})$$

$$\rho \parallel \sigma, R \rightarrow^1 \rho, \sigma, R \quad (\text{A-CRED-MERGE})$$

$$\text{Empty}, R \rightarrow^1 R \quad (\text{A-CRED-EMPTY})$$

$$l:t, R \rightarrow^1 l, R \quad (\text{A-CRED-BASE})$$

$$r \setminus l \setminus l' \approx^1 r \setminus l' \setminus l \quad (\text{A-EQ-RESTR/RESTR'})$$

$$r \parallel s \approx^1 s \parallel r \quad (\text{A-EQ-COMM})$$

$$r \parallel (s \parallel t) \approx^1 (r \parallel s) \parallel t \quad (\text{A-EQ-ASSOC})$$

$$\phi, \phi, R \approx^1 \phi, R \quad (\text{A-CEQ-DUPL})$$

$$\rho, \sigma, R \approx^1 \sigma, \rho, R \quad (\text{A-CEQ-COMM})$$

B.3 Well-formed constraint lists

$$\frac{T \Rightarrow \text{ok}}{T \vdash \circ \Rightarrow \text{ok}} \quad (\text{A-WFCL-NIL})$$

$$\frac{T \vdash r \Rightarrow \text{record} \quad T \vdash R \Rightarrow \text{ok}}{T \vdash r, R \Rightarrow \text{ok}} \quad (\text{A-WFCL-CONS})$$

B.4 Well-formed types

$$\begin{array}{l} T \vdash p \Rightarrow \text{type} \quad (\text{A-WFT-PRIM}) \\ \\ \frac{T \vdash t_1 \Rightarrow \text{type} \quad T \vdash t_2 \Rightarrow \text{type}}{T \vdash t_1 \rightarrow t_2 \Rightarrow \text{type}} \quad (\text{A-WFT-ARROW}) \\ \\ \frac{T, a \# R \vdash t \Rightarrow \text{type} \quad T \vdash R \Rightarrow \text{ok}}{T \vdash \forall a \# R. t \Rightarrow \text{type}} \quad (\text{A-WFT-ALL}) \\ \\ \frac{T \vdash r \Rightarrow \text{record}}{T \vdash r \Rightarrow \text{type}} \quad (\text{A-WFT-REC}) \end{array}$$

B.5 Well-formed record types

$$\begin{array}{l} T_1, a \# R, T_2 \vdash a \Rightarrow \text{record} \quad (\text{A-WFR-VAR}) \\ \\ \frac{T \vdash t \Rightarrow \text{type}}{T \vdash l:t \Rightarrow \text{record}} \quad (\text{A-WFR-BASE}) \\ \\ \frac{T \vdash r \Rightarrow \text{record} \quad r.l \downarrow}{T \vdash r \setminus l \Rightarrow \text{record}} \quad (\text{A-WFR-RESTR}) \\ \\ T \vdash \text{Empty} \Rightarrow \text{record} \quad (\text{A-WFR-EMPTY}) \\ \\ \frac{T \vdash r_1 \Rightarrow \text{record} \quad T \vdash r_2 \Rightarrow \text{record} \quad T \vdash r_1 \#^\Rightarrow r_2}{T \vdash r_1 \parallel r_2 \Rightarrow \text{record}} \quad (\text{A-WFR-MERGE}) \end{array}$$

B.6 Constraint list satisfaction

$$\begin{array}{l} \frac{T \vdash r \Rightarrow \text{record}}{T \vdash r \#^\Rightarrow \diamond} \quad (\text{A-CMP-LIST-NIL}) \\ \\ \frac{T \vdash r \#^\Rightarrow r_i \quad T \vdash r \#^\Rightarrow R}{T \vdash r \#^\Rightarrow r_i, R} \quad (\text{A-CMP-LIST-CONS}) \end{array}$$

B.7 Compatibility

$$\begin{array}{l} T \vdash r \#^\Rightarrow \text{Empty} \quad (\text{A-CMP-EMPTY-1}) \\ \\ T \vdash \text{Empty} \#^\Rightarrow r \quad (\text{A-CMP-EMPTY-2}) \\ \\ \frac{T \vdash s_1 \#^\Rightarrow r \quad T \vdash s_2 \#^\Rightarrow r}{T \vdash (s_1 \parallel s_2) \#^\Rightarrow r} \quad (\text{A-CMP-MERGE-I-1}) \\ \\ \frac{T \vdash r \#^\Rightarrow s_1 \quad T \vdash r \#^\Rightarrow s_2}{T \vdash r \#^\Rightarrow (s_1 \parallel s_2)} \quad (\text{A-CMP-MERGE-I-2}) \\ \\ \frac{l \neq l'}{T \vdash l:t \#^\Rightarrow l':t'} \quad (\text{A-CMP-BASE/BASE}) \\ \\ \frac{a \in S^*}{T_1, a \# R, T_2, b \# S, T_3 \vdash a \#^\Rightarrow b} \quad (\text{A-CMP-TVAR/TVAR-1}) \end{array}$$

$$\frac{a \in S^*}{T_1, a\#R, T_2, b\#S, T_3 \vdash b \# \Rightarrow a} \quad (\text{A-CMP-TVAR/TVAR-2})$$

$$\frac{l \in R^*}{T_1, a\#R, T_2 \vdash a \# \Rightarrow l:t} \quad (\text{A-CMP-BASE/TVAR-1})$$

$$\frac{l \in R^*}{T_1, a\#R, T_2 \vdash l:t \# \Rightarrow a} \quad (\text{A-CMP-BASE/TVAR-2})$$

B.8 Well-typed terms

$$T; G_1, x:t, G_2 \vdash x \Rightarrow t \quad (\text{A-WTT-VAR})$$

$$\frac{T \vdash t \Rightarrow \text{type} \quad T; G, x:t \vdash e \Rightarrow t'}{T; G \vdash \lambda x:t. e \Rightarrow t \rightarrow t'} \quad (\text{A-WTT-ARROW-I})$$

$$\frac{T; G \vdash e_1 \Rightarrow t \rightarrow t' \quad T; G \vdash e_2 \Rightarrow t'' \quad t'' \equiv t}{T; G \vdash e_1 e_2 \Rightarrow t'} \quad (\text{A-WTT-ARROW-E})$$

$$\frac{T; G \vdash e \Rightarrow t}{T; G \vdash l=e \Rightarrow l:t} \quad (\text{A-WTT-BASE})$$

$$\frac{T; G \vdash e \Rightarrow r \quad r.l \downarrow}{T; G \vdash e \setminus l \Rightarrow r \setminus l} \quad (\text{A-WTT-RESTR})$$

$$T; G \vdash \text{empty} \Rightarrow \text{Empty} \quad (\text{A-WTT-EMPTY})$$

$$\frac{T; G \vdash e_1 \Rightarrow r_1 \quad T; G \vdash e_2 \Rightarrow r_2 \quad T \vdash r_1^* \# \Rightarrow r_2^*}{T; G \vdash e_1 \parallel e_2 \Rightarrow r_1 \parallel r_2} \quad (\text{A-WTT-MERGE})$$

$$\frac{T; G \vdash e \Rightarrow r \quad r.l \downarrow}{T; G \vdash e.l \Rightarrow r.l} \quad (\text{A-WTT-SELECT})$$

$$\frac{T \vdash R \text{ ok} \quad T, a\#R; G \vdash e \Rightarrow t}{T; G \vdash \Lambda a\#R. e \Rightarrow \text{Va}\#R. t} \quad (\text{A-WTT-ALL-I})$$

$$\frac{T; G \vdash e \Rightarrow \text{Va}\#R. t \quad T \vdash r \Rightarrow \text{record} \quad T \vdash r^* \# \Rightarrow R^*}{T; G \vdash e[r] \Rightarrow [r/a]t} \quad (\text{A-WTT-ALL-E})$$

References

- [1] Peter Canning, William Cook, Walter Hill, Walter Olthoff, and John Mitchell. F-bounded quantification for object-oriented programming. In *Fourth International Conference on Functional Programming Languages and Computer Architecture*, pages 273–280, September 1989.
- [2] Luca Cardelli. A semantics of multiple inheritance. In G. Kahn, D. MacQueen, and G. Plotkin, editors, *Semantics of Data Types*, volume 173 of *Lecture Notes in Computer Science*, pages 51–67. Springer-Verlag, 1984.
- [3] Luca Cardelli. A semantics of multiple inheritance. *Information and Computation*, 76:138–164, 1988.
- [4] Luca Cardelli and John Mitchell. Operations on records. In *Proceedings of Fifth International Conference on Mathematical Foundations of Programming Language Semantics*, Tulane University, New Orleans, March 1989. To appear.
- [5] Luca Cardelli and John C. Mitchell. Operations on records. Research report 48, Digital Equipment Corporation, Systems Research Center, August 1989.
- [6] Luca Cardelli and Peter Wegner. On understanding types, data abstraction, and polymorphism. *Computing Surveys*, 17(4), December 1985.
- [7] William R. Cook, Walter L. Hill, and Peter S. Canning. Inheritance is not subtyping. In *Seventeenth Annual ACM Symposium on Principles of Programming Languages*, pages 125–135, San Francisco, CA, January 1990.
- [8] Luis Damas and Robin Milner. Principal type schemes for functional programs. In *Proceedings of the 9th ACM Symposium on Principles of Programming Languages*, pages 207–212. ACM SIGPLAN/SIGACT, 1982.
- [9] Jean-Yves Girard. *Interprétation fonctionnelle et élimination des coupures de l'arithmétique d'ordre supérieur*. PhD thesis, Université Paris VII, 1972.
- [10] Robert W. Harper and Benjamin C. Pierce. Extensible records without subsumption. Technical Report CMU-CS-90-102, School of Computer Science, Carnegie Mellon University, February 1990.
- [11] Gérard Huet. Confluent reductions: Abstract properties and applications to term rewriting systems. *Journal of the Association for Computing Machinery*, 27(4):797–821, October 1980.
- [12] Lalita A. Jategaonkar. ML with extended pattern matching and subtypes. Master's thesis, MIT, August 1989.
- [13] Lalita A. Jategaonkar and John C. Mitchell. ML with extended pattern matching and subtypes (preliminary version). In *Proceedings of the ACM Conference on Lisp and Functional Programming*, pages 198–211, Snowbird, Utah, July 1988.
- [14] Robin Milner. A theory of type polymorphism in programming. *Journal of Computer and System Sciences*, 17:348–375, August 1978.
- [15] Robin Milner, Mads Tofte, and Robert Harper. *The Definition of Standard ML*. MIT Press, 1990.
- [16] John C. Mitchell. Coercion and type inference (summary). In *Proc. 11th ACM Symp. on Principles of Programming Languages*, pages 175–185, January 1984.
- [17] Atsushi Ohori and Peter Buneman. Static type inference for parametric classes. In *OOPSLA '89: Object-Oriented Programming Systems, Languages, and Applications, Conference Proceedings*, pages 445–456, October 1989.

- [18] Atsushi Ohori and Peter Buneman. Type inference in a database programming language. In *1988 ACM Conference on Lisp and Functional Programming*, pages 174–183, Snowbird, Utah, July 1989. Revised version, September, 1988.
- [19] Didier Rémy. Typechecking records and variants in a natural extension of ML. In *Proceedings of the Sixteenth Annual ACM Symposium on Principles of Programming Languages*, Austin, pages 242–249. ACM, January 1989.
- [20] Didier Rémy. *Algèbres Touffues. Application au Typage Polymorphe des Objets Enregistrements dans les Langages Fonctionnels*. PhD thesis, Université Paris VII, 1990.
- [21] Didier Rémy. Typechecking records in a natural extension of ML. Submitted to TOPLAS, June 1990.
- [22] John Reynolds. Towards a theory of type structure. In *Proc. Colloque sur la Programmation*, pages 408–425, New York, 1974. Springer-Verlag LNCS 19.
- [23] Ryan Stansifer. Type inference with subtypes. In *Proceedings of the Fifteenth ACM Symposium on Principles of Programming Languages*, pages 88–97, San Diego, CA, January 1988.
- [24] Mitchell Wand. Complete type inference for simple objects. In *Proceedings of the IEEE Symposium on Logic in Computer Science*, Ithaca, NY, June 1987.
- [25] Mitchell Wand. Corrigendum: Complete type inference for simple objects. In *Proceedings of the IEEE Symposium on Logic in Computer Science*, 1988.
- [26] Mitchell Wand. Type inference for objects with instance variables and inheritance. Technical Report NU-CCS-89-2, College of Computer Science, Northeastern University, 1989.
- [27] Mitchell Wand. Type inference for record concatenation and multiple inheritance. In *Fourth Annual IEEE Symposium on Logic in Computer Science*, pages 92–97, Pacific Grove, CA, June 1989.