

10

The Essence of ML Type Inference

François Pottier and Didier Rémy

10.1 What Is ML?

The name ML appeared during the late seventies. It then referred to a general-purpose programming language that was used as a meta-language (whence its name) within the theorem prover LCF (Gordon, Milner, and Wadsworth, 1979). Since then, several new programming languages, each of which offers several different implementations, have drawn inspiration from it. So, what does ML stand for today?

For a semanticist, ML might stand for a programming language featuring first-class functions, data structures built out of products and sums, mutable memory cells called *references*, exception handling, automatic memory management, and a call-by-value semantics. This view encompasses the Standard ML (Milner, Tofte, and Harper, 1990) and Caml (Leroy, 2000) families of programming languages. We refer to it as *ML-the-programming-language*.

For a type theorist, ML might stand for a particular breed of type systems, based on the simply-typed λ -calculus, but extended with a simple form of polymorphism introduced by `let` declarations. These type systems have decidable type inference; their type inference algorithms strongly rely on first-order unification and can be made efficient in practice. Besides Standard ML and Caml, this view encompasses programming languages such as Haskell (Peyton Jones, 2003) and Clean (Brus, van Eekelen, van Leer, and Plasmeijer, 1987), whose semantics is rather different—indeed, it is nonstrict and pure (Sabry, 1998)—but whose type system fits this description. We refer to it as *ML-the-type-system*. It is also referred to as the *Hindley-Milner type discipline* in the literature.

Code for this chapter may be found on the book's web site.

For us, ML might also stand for the particular programming language whose formal definition is given and studied in this chapter. It is a core calculus featuring first-class functions, local definitions, and constants. It is equipped with a call-by-value semantics. By customizing constants and their semantics, one may recover data structures, references, and more. We refer to this particular calculus as *ML-the-calculus*.

Why study ML-the-type-system today, such a long time after its initial discovery? One may think of at least two reasons.

First, its treatment in the literature is often cursory, because it is considered either as a simple extension of the simply-typed λ -calculus (*TAPL*, Chapter 9) or as a subset of Girard and Reynolds' System F (*TAPL*, Chapter 23). The former view is supported by the claim that local (`let`) definitions, which distinguish ML-the-type-system from the simply-typed λ -calculus, may be understood as a simple textual expansion facility. However, this view tells only part of the story, because it fails to give an account of the *principal types* property enjoyed by ML-the-type-system, leads to a naive type inference algorithm whose time complexity is exponential not only in the worst case but in the common case, and breaks down when the language is extended with side effects, such as state or exceptions. The latter view is supported by the fact that every type derivation within ML-the-type-system is also a valid type derivation within an implicitly-typed variant of System F. Such a view is correct but again fails to give an account of type inference for ML-the-type-system, since type inference for System F is undecidable (Wells, 1999).

Second, existing accounts of type inference for ML-the-type-system (Milner, 1978; Damas and Milner, 1982; Tofte, 1988; Leroy, 1992; Lee and Yi, 1998; Jones, 1999) often involve heavy manipulations of type substitutions. Such a ubiquitous use of type substitutions is often quite obscure. Furthermore, actual implementations of the type inference algorithm do *not* explicitly manipulate substitutions; instead, they extend a standard first-order unification algorithm, where terms are updated in place as new equations are discovered (Huet, 1976; Martelli and Montanari, 1982). Thus, it is hard to tell, from these accounts, how to write an efficient type inference algorithm for ML-the-type-system. Yet, in spite of the increasing speed of computers, efficiency remains crucial when ML-the-type-system is extended with expensive features, such as Objective Caml's object types (Rémy and Vouillon, 1998), variant types (Garrigue, 1998), or polymorphic methods (Garrigue and Rémy, 1999).

Our emphasis on efficiency might come as a surprise, since type inference for ML-the-type-system is known to be DEXPTIME-complete (Kfoury, Tiuryn, and Urzyczyn, 1990; Mairson, Kanellakis, and Mitchell, 1991). In practice, however, most implementations of it behave well. This apparent contradiction may be explained by observing that types usually remain small and

that `let` constructs are never deeply nested towards the left. Indeed, under the assumption that types have bounded size and that programs have bounded “scheme depth,” type inference may be performed in quasi-linear time (McAllester, 2003). In ML-the-programming-language, algebraic data type definitions allow complex data structures to be described by concise expressions, such as “list X,” which helps achieve the bounded-type-size property.

In fact, in such favorable circumstances, even an inefficient algorithm may behave well. For instance, some deployed implementations of type inference for ML-the-type-system contain sources of inefficiency (see remark 10.1.21 on page 404) and do not operate in quasi-linear time under the bounded-type-size assumption. However, such implementations are put under greater stress when types become larger, a situation that occurs in some programs (Saha, Heintze, and Oliva, 1998) and also arises when large, transparent type expressions are used instead of algebraic data types, as in Objective Caml’s object-oriented fragment (Rémy and Vouillon, 1998).

For these reasons, we believe it is worth giving an account of ML-the-type-system that focuses on *type inference* and strives to be at once elegant and faithful to an efficient implementation, such as Rémy’s (1992a). In this presentation, we forego type substitutions and instead put emphasis on constraints, which offer a number of advantages.

First, constraints allow a modular presentation of type inference as the combination of a constraint generator and a constraint solver, allowing separate reasoning about *when* a program is correct and *how* to check whether it is correct. This perspective has long been standard in the setting of the simply-typed λ -calculus: see, for example, Wand (1987b) and *TAPL*, Chapter 22. In the setting of ML-the-type-system, such a decomposition is provided by the reduction of typability problems to acyclic semi-unification problems (Henglein, 1993; Kfoury, Tiuryn, and Urzyczyn, 1994). This approach, however, was apparently never used in production implementations of ML-the-programming-language. An experimental extension of SML/NJ with polymorphic recursion (Emms and LeiSS, 1996) did reduce type inference to a semi-unification problem. Semi-unification found applications in the closely related area of program analysis; see, for example, Fähndrich, Rehof, and Das (2000) and Birkedal and Tofte (2001). In this chapter, we give a constraint-based description of a “classic” implementation of ML-the-type-system, which is based on first-order unification and a mechanism for creating and instantiating principal *type schemes*.

Second, it is often natural to define and implement the solver as a constraint rewriting system. The constraint language allows reasoning not only about correctness—is every rewriting step meaning-preserving?—but also about low-level implementation details, since constraints are the data struc-

$x, y ::=$	<i>Identifiers:</i>	m	<i>Memory location</i>
z	<i>Variable</i>	$\lambda z.t$	<i>Function</i>
m	<i>Memory location</i>	$c v_1 \dots v_k$	<i>Data</i>
c	<i>Constant</i>	$c \in \mathcal{Q}^+ \wedge k \leq a(c)$	
$t ::=$	<i>Expressions:</i>	$c v_1 \dots v_k$	<i>Partial application</i>
x	<i>Identifier</i>	$c \in \mathcal{Q}^- \wedge k < a(c)$	
$\lambda z.t$	<i>Function</i>	$\mathcal{E} ::=$	<i>Evaluation Contexts:</i>
$t t$	<i>Application</i>	$[]$	<i>Empty context</i>
$\text{let } z = t \text{ in } t$	<i>Local definition</i>	$\mathcal{E} t$	<i>Left side of an application</i>
$v, w ::=$	<i>Values:</i>	$v \mathcal{E}$	<i>Right side of an application</i>
z	<i>Variable</i>	$\text{let } z = \mathcal{E} \text{ in } t$	<i>Local definition</i>

Figure 10-1: Syntax of ML-the-calculus

tures manipulated throughout the type inference process. For instance, describing unification in terms of *multi-equations* allows reasoning about the sharing of nodes in memory, which a substitution-based approach cannot account for. Last, constraints are more general than type substitutions, and allow smooth extensions of ML-the-type-system with recursive types, rows, subtyping, and more. These arguments are developed, for example, in Jouanaud and Kirchner (1991).

Before delving into the details of this new presentation of ML-the-type-system, it is worth recalling its standard definition. Thus, in what follows, we first define the syntax and operational semantics of ML-the-calculus, and equip it with a type system, known as *Damas and Milner's type system*.

ML-the-Calculus

The syntax of ML-the-calculus is defined in Figure 10-1. It is made up of several syntactic categories.

Identifiers group several kinds of names that may be referenced in a program: variables, memory locations, and constants. We let x and y range over identifiers. *Variables*—also called *program variables*, to avoid ambiguity—are names that may be bound to values using λ or let binding forms; in other words, they are names for function parameters or local definitions. We let z and f range over program variables. We sometimes write $_$ for a program variable that does not occur free within its scope: for instance, $\lambda_ .t$ stands for $\lambda z.t$, provided z is fresh for t . (We say that z is *fresh for* t when z does not oc-

cur free in τ .) *Memory locations* are names that represent memory addresses. They are used to model references (see Example 10.1.9 below). Memory locations never appear in *source programs*, that is, programs that are submitted to a compiler. They only appear during execution, when new memory blocks are allocated. *Constants* are fixed names for primitive values and operations, such as integer literals and integer arithmetic operations. Constants are elements of a finite or infinite set \mathcal{Q} . They are never subject to α -conversion, in contrast to variables and memory locations. Program variables, memory locations, and constants belong to distinct syntactic classes and may never be confused.

The set of constants \mathcal{Q} is kept abstract, so most of our development is independent of its concrete definition. We assume that every constant c has a nonnegative integer *arity* $a(c)$. We further assume that \mathcal{Q} is partitioned into subsets of *constructors* \mathcal{Q}^+ and *destructors* \mathcal{Q}^- . Constructors and destructors differ in that the former are used to *form* values, while the latter are used to *operate* on values.

- 10.1.1 EXAMPLE [INTEGERS]: For every integer n , one may introduce a nullary constructor \hat{n} . In addition, one may introduce a binary destructor $\hat{+}$, whose applications are written infix, so $\tau_1 \hat{+} \tau_2$ stands for the double application $\hat{+} \tau_1 \tau_2$ of the destructor $\hat{+}$ to the expressions τ_1 and τ_2 . \square

Expressions—also known as *terms* or *programs*—are the main syntactic category. Indeed, unlike procedural languages such as C and Java, functional languages, including ML-the-programming-language, suppress the distinction between expressions and statements. Expressions consist of identifiers, λ -*abstractions*, *applications*, and local definitions. The λ -abstraction $\lambda z. \tau$ represents the function of one parameter named z whose result is the expression τ , or, in other words, the function that maps z to τ . Note that the variable z is bound within the term τ , so (for instance) the notations $\lambda z_1. z_1$ and $\lambda z_2. z_2$ denote the same entity. The application $\tau_1 \tau_2$ represents the result of calling the function τ_1 with actual parameter τ_2 , or, in other words, the result of applying τ_1 to τ_2 . Application is left-associative, that is, $\tau_1 \tau_2 \tau_3$ stands for $(\tau_1 \tau_2) \tau_3$. The construct $\text{let } z = \tau_1 \text{ in } \tau_2$ represents the result of evaluating τ_2 after binding the variable z to τ_1 . Note that the variable z is bound within τ_2 , but not within τ_1 , so for instance $\text{let } z_1 = z_1 \text{ in } z_1$ and $\text{let } z_2 = z_1 \text{ in } z_2$ are the same object. The construct $\text{let } z = \tau_1 \text{ in } \tau_2$ has the same meaning as $(\lambda z. \tau_2) \tau_1$, but is dealt with in a more flexible way by ML-the-type-system. To sum up, the syntax of ML-the-calculus is that of the pure λ -calculus, extended with memory locations, constants, and the let construct.

Values form a subset of expressions. They are expressions whose evaluation is completed. Values include identifiers, λ -abstractions, and applications

of constants, of the form $c v_1 \dots v_k$, where k does not exceed c 's arity if c is a constructor, and k is smaller than c 's arity if c is a destructor. In what follows, we are often interested in closed values—ones that do not contain any free program variables. We use the meta-variables v and w for values.

- 10.1.2 EXAMPLE: The integer literals $\dots, \widehat{-1}, \widehat{0}, \widehat{1}, \dots$ are nullary constructors, so they are values. Integer addition $\widehat{+}$ is a binary destructor, so it is a value, and so is every partial application $\widehat{+} v$. Thus, both $\widehat{+} \widehat{1}$ and $\widehat{+} \widehat{+}$ are values. An application of $\widehat{+}$ to two values, such as $\widehat{2}\widehat{+}\widehat{2}$, is not a value. \square
- 10.1.3 EXAMPLE [PAIRS]: Let (\cdot, \cdot) be a binary constructor. If t_1 and t_2 are expressions, then the double application $(\cdot, \cdot) t_1 t_2$ may be called the *pair* of t_1 and t_2 , and may be written (t_1, t_2) . By the definition above, (t_1, t_2) is a value if and only if t_1 and t_2 are both values. \square

Stores are finite mappings from memory locations to closed values. A store μ represents what is usually called a heap, that is, a collection of values, each of which is allocated at a particular address in memory and may contain pointers to other elements of the heap. ML-the-programming-language allows overwriting the contents of an existing memory block—an operation sometimes referred to as a *side effect*. In the operational semantics, this effect is achieved by mapping an existing memory location to a new value. We write \emptyset for the empty store. We write $\mu[m \mapsto v]$ for the store that maps m to v and otherwise coincides with μ . When μ and μ' have disjoint domains, we write $\mu\mu'$ for their union. We write $dom(\mu)$ for the domain of μ and $range(\mu)$ for the set of memory locations that appear in its codomain.

The operational semantics of a pure language like the λ -calculus may be defined as a rewriting system on expressions. Because ML-the-calculus has side effects, however, we define its operational semantics as a rewriting system on *configurations*. A configuration t/μ is a pair of an expression t and a store μ . The memory locations in the domain of μ are *not* considered bound within t/μ , so, for instance, $m_1/(m_1 \mapsto \widehat{0})$ and $m_2/(m_2 \mapsto \widehat{0})$ denote distinct entities. (In the literature, memory locations are often considered bound inside configurations. This offers the advantage of making memory allocation a deterministic operation. However, there is still a need for non- α -convertible configurations: rules R-EXTEND and R-CONTEXT in Figure 10-2 cannot otherwise be correctly stated! Quite a few papers run into this pitfall.)

A configuration t/μ is *closed* if and only if t has no free program variables and every memory location that appears within t or within the range of μ is in the domain of μ . If t is a closed source program, its evaluation begins within an empty store—that is, with the configuration t/\emptyset . Because source programs do not contain memory locations, this configuration is closed. Furthermore, we shall see that closed configurations are preserved by reduction.

$(\lambda z. t) v \rightarrow [z \mapsto v]t \quad \text{(R-BETA)}$	$\frac{\tau/\mu \rightarrow \tau'/\mu' \quad \text{dom}(\mu'') \# \text{dom}(\mu') \quad \text{range}(\mu'') \# \text{dom}(\mu' \setminus \mu)}{\tau/\mu\mu'' \rightarrow \tau'/\mu'\mu''} \quad \text{(R-EXTEND)}$
$\text{let } z = v \text{ in } t \rightarrow [z \mapsto v]t \quad \text{(R-LET)}$	$\frac{\tau/\mu \rightarrow \tau'/\mu'}{\mathcal{E}[t]/\mu \rightarrow \mathcal{E}[t']/\mu'} \quad \text{(R-CONTEXT)}$
$\frac{\tau/\mu \xrightarrow{\delta} \tau'/\mu'}{\tau/\mu \rightarrow \tau'/\mu'} \quad \text{(R-DELTA)}$	

Figure 10-2: Semantics of ML-the-calculus

Note that, instead of separating expressions and stores, it is possible to make store fragments part of the syntax of expressions; this idea, proposed in Crank and Felleisen (1991), has also been used for the encoding of reference cells in process calculi.

A *context* is an expression where a single subexpression has been replaced with a *hole*, written $[\]$. *Evaluation contexts* form a strict subset of contexts. In an evaluation context, the hole is meant to highlight a point in the program where it is valid to apply a reduction rule. Thus, the definition of evaluation contexts determines a reduction strategy: it tells where and in what order reduction steps may occur. For instance, the fact that $\lambda z. [\]$ is not an evaluation context means that the body of a function is never evaluated—that is, not until the function is applied. The fact that $\tau \mathcal{E}$ is an evaluation context only if τ is a value means that, to evaluate an application $\tau_1 \tau_2$, one should fully evaluate τ_1 before attempting to evaluate τ_2 . More generally, in the case of a multiple application, it means that arguments should be evaluated from left to right. Of course, other choices could be made: for instance, defining $\mathcal{E} ::= \dots | \tau \mathcal{E} | \mathcal{E} v | \dots$ would enforce a right-to-left evaluation order, while defining $\mathcal{E} ::= \dots | \tau \mathcal{E} | \mathcal{E} \tau | \dots$ would leave the evaluation order unspecified, effectively allowing reduction to alternate between both subexpressions, and making evaluation nondeterministic (because side effects could occur in different order). The fact that $\text{let } z = v \text{ in } \mathcal{E}$ is not an evaluation context means that the body of a local definition is never evaluated—that is, not until the definition itself is reduced. We write $\mathcal{E}[t]$ for the expression obtained by replacing the hole in \mathcal{E} with the expression t .

Figure 10-2 defines first a relation \rightarrow between arbitrary configurations, then a relation \rightarrow between closed configurations. If $\tau/\mu \rightarrow \tau'/\mu$ holds for every store μ , then we write $\tau \rightarrow \tau'$ and say that the reduction is *pure*.

The semantics need not be deterministic. That is, a configuration may reduce to two different configurations. In fact, our semantics is deterministic

only if the relation $\xrightarrow{\delta}$, which is a parameter to our semantics, is itself deterministic. In practice, $\xrightarrow{\delta}$ is usually deterministic, up to α -conversion of memory locations. As explained above, the semantics could also be made nondeterministic by a different choice in the definition of evaluation contexts.

The key reduction rule is R-BETA, which states that a function application $(\lambda z. \tau) v$ reduces to the function body, namely τ , where every occurrence of the formal argument z has been replaced with the actual argument v . The λ construct, which prevented the function body τ from being evaluated, disappears, so the new term may (in general) be further reduced. Because ML-the-calculus adopts a *call-by-value* strategy, rule R-BETA is applicable only if the actual argument is a value v . In other words, a function cannot be invoked until its actual argument has been fully evaluated. Rule R-LET is very similar to R-BETA. Indeed, it specifies that $\text{let } z = v \text{ in } \tau$ has the same behavior, with respect to reduction, as $(\lambda z. \tau) v$. Substitution of a value for a program variable throughout a term is expensive, so R-BETA and R-LET are never implemented literally: they are only a simple specification. Actual implementations usually employ *runtime environments*, which may be understood as a form of *explicit substitutions* (Abadi, Cardelli, Curien, and Lévy, 1991; Hardin, Maranget, and Pagano, 1998). Note that our choice of a call-by-value reduction strategy has essentially no impact on the type system; the programming language Haskell, whose reduction strategy is known as *lazy* or *call-by-need*, also relies on the Hindley-Milner type discipline.

Rule R-DELTA describes the semantics of constants. It states that a certain relation $\xrightarrow{\delta}$ is a subset of \rightarrow . Of course, since the set of constants is unspecified, the relation $\xrightarrow{\delta}$ must be kept abstract as well. We require that, if $\tau/\mu \xrightarrow{\delta} \tau'/\mu'$ holds, then

- (i) τ is of the form $c v_1 \dots v_n$, where c is a destructor of arity n ; and
- (ii) $\text{dom}(\mu)$ is a subset of $\text{dom}(\mu')$.

Condition (i) ensures that δ -reduction concerns full applications of destructors only, and that these are evaluated in accordance with the call-by-value strategy. Condition (ii) ensures that δ -reduction may allocate new memory locations, but not deallocate existing locations. In particular, a “garbage collection” operator, which destroys unreachable memory cells, cannot be made available as a constant. Doing so would not make much sense anyway in the presence of R-EXTEND. Condition (ii) allows proving that, if τ/μ reduces (by \rightarrow) to τ'/μ' , then $\text{dom}(\mu)$ is also a subset of $\text{dom}(\mu')$; checking this is left as an exercise to the reader.

Rule R-EXTEND states that any valid reduction is also valid in a larger store. The initial and final stores μ and μ' in the original reduction are both ex-

tended with a new store fragment μ'' . The rule's second premise requires that the domain of μ'' be disjoint with that of μ' (and consequently, also with that of μ), so that the new memory locations are indeed undefined in the original reduction. (They may, however, appear in the image of μ .) The last premise ensures that the new memory locations in μ'' do not accidentally carry the same names as the locations allocated during the original reduction step, that is, the locations in $\text{dom}(\mu' \setminus \mu)$. The notation $A \# B$ stands for $A \cap B = \emptyset$.

Rule R-CONTEXT completes the definition of the operational semantics by defining \longrightarrow , a relation between closed configurations, in terms of \rightarrow . The rule states that reduction may take place not only at the term's root, but also deep inside it, provided the path from the root to the point where reduction occurs forms an evaluation context. This is how evaluation contexts determine an evaluation strategy. As a purely technical point, because \longrightarrow relates closed configurations only, we do not need to require that the memory locations in $\text{dom}(\mu' \setminus \mu)$ be fresh for \mathcal{E} ; indeed, every memory location that appears within \mathcal{E} must be a member of $\text{dom}(\mu)$.

- 10.1.4 EXAMPLE [INTEGERS, CONTINUED]: The operational semantics of integer addition may be defined as follows:

$$\hat{k}_1 \hat{+} \hat{k}_2 \xrightarrow{\delta} \widehat{k_1 + k_2} \quad (\text{R-ADD})$$

The left-hand term is the double application $\hat{+} \hat{k}_1 \hat{k}_2$, while the right-hand term is the integer literal \hat{k} , where k is the sum of k_1 and k_2 . The distinction between object level and meta level (that is, between \hat{k} and k) is needed here to avoid ambiguity. \square

- 10.1.5 EXAMPLE [PAIRS, CONTINUED]: In addition to the pair constructor defined in Example 10.1.3, we may introduce two destructors π_1 and π_2 of arity 1. We may define their operational semantics as follows, for $i \in \{1, 2\}$:

$$\pi_i (v_1, v_2) \xrightarrow{\delta} v_i \quad (\text{R-PROJ})$$

Thus, our treatment of constants is general enough to account for pair construction and destruction; we need not build these features explicitly into the language. \square

- 10.1.6 EXERCISE [BOOLEANS, RECOMMENDED, $\star\star$, \rightarrow]: Let `true` and `false` be nullary constructors. Let `if` be a ternary destructor. Extend the semantics with

$$\text{if true } v_1 v_2 \xrightarrow{\delta} v_1 \quad (\text{R-TRUE})$$

$$\text{if false } v_1 v_2 \xrightarrow{\delta} v_2 \quad (\text{R-FALSE})$$

Let us use the syntactic sugar `if t0 then t1 else t2` for the triple application of `if t0 t1 t2`. Explain why these definitions do not quite provide the expected behavior. Without modifying the semantics of `if`, suggest a new

definition of the syntactic sugar `if t0 then t1 else t2` that corrects the problem. \square

- 10.1.7 EXAMPLE [SUMS]: Booleans may in fact be viewed as a special case of the more general concept of *sum*. Let `inj1` and `inj2` be unary constructors, called respectively *left* and *right injections*. Let `case` be a ternary destructor, whose semantics is defined as follows, for $i \in \{1, 2\}$:

$$\text{case } (\text{inj}_i v) v_1 v_2 \xrightarrow{\delta} v_i v \quad (\text{R-CASE})$$

Here, the value `inji v` is being scrutinized, while the values v_1 and v_2 , which are typically functions, represent the two arms of a standard case construct. The rule selects an appropriate arm (here, v_i) based on whether the value under scrutiny was formed using a left or right injection. The arm v_i is executed and given access to the data carried by the injection (here, v). \square

- 10.1.8 EXERCISE [\star , \rightarrow]: Explain how to encode `true`, `false`, and the `if` construct in terms of sums. Check that the behavior of R-TRUE and R-FALSE is properly emulated. \square

- 10.1.9 EXAMPLE [REFERENCES]: Let `ref` and `!` be unary destructors. Let `:=` be a binary destructor. We write `t1 := t2` for the double application `:= t1 t2`. Define the operational semantics of these three destructors as follows:

$$\text{ref } v/\emptyset \xrightarrow{\delta} m/(m \mapsto v) \quad \text{if } m \text{ is fresh for } v \quad (\text{R-REF})$$

$$!m/(m \mapsto v) \xrightarrow{\delta} v/(m \mapsto v) \quad (\text{R-DEREF})$$

$$m := v/(m \mapsto v_0) \xrightarrow{\delta} v/(m \mapsto v) \quad (\text{R-ASSIGN})$$

According to R-REF, evaluating `ref v` allocates a fresh memory location m and binds v to it. The name m must be chosen fresh for v to prevent inadvertent capture of the memory locations that appear free within v . By R-DEREF, evaluating `!m` returns the value bound to the memory location m within the current store. By R-ASSIGN, evaluating `m := v` discards the value v_0 currently bound to m and produces a new store where m is bound to v . Here, the value returned by the assignment `m := v` is v itself; in ML-the-programming-language, it is usually a nullary constructor `()`, pronounced *unit*. \square

- 10.1.10 EXAMPLE [RECURSION]: Let `fix` be a binary destructor, whose operational semantics is:

$$\text{fix } v_1 v_2 \xrightarrow{\delta} v_1 (\text{fix } v_1) v_2 \quad (\text{R-FIX})$$

`fix` is a fixpoint combinator: it effectively allows recursive definitions of functions. Indeed, the construct `letrec f = $\lambda z.t_1$ in t_2` provided by ML-the-programming-language may be viewed as syntactic sugar for `let f = $\text{fix } (\lambda f.\lambda z.t_1)$ in t_2` . \square

- 10.1.11 EXERCISE [RECOMMENDED, ★★, →]: Assuming the availability of Booleans and conditionals, integer literals, subtraction, multiplication, integer comparison, and a fixpoint combinator, most of which were defined in previous examples, define a function that computes the factorial of its integer argument, and apply it to $\hat{3}$. Determine, step by step, how this expression reduces to a value. \square

It is straightforward to check that, if τ/μ reduces to τ'/μ' , then τ is not a value. In other words, values are irreducible: they represent completed computations. The proof is left as an exercise to the reader. The converse, however, does not hold: if the closed configuration τ/μ is irreducible with respect to \rightarrow , then τ is not necessarily a value. In that case, the configuration τ/μ is said to be *stuck*. It represents a *runtime error*, that is, a situation that does not allow computation to proceed, yet is not considered a valid outcome. A closed source program τ is said to *go wrong* if and only if the initial configuration τ/\emptyset reduces to a stuck configuration.

- 10.1.12 EXAMPLE: Runtime errors typically arise when destructors are applied to arguments of an unexpected nature. For instance, the expressions $\hat{\pi}_1 \hat{1} m$ and $\hat{\pi}_1 \hat{2}$ and $!\hat{3}$ are stuck, regardless of the current store. The program $\text{let } z = \hat{\pi}_1 \hat{\pi}_1 \text{ in } z + 1$ is not stuck, because $\hat{\pi}_1 \hat{\pi}_1$ is a value. However, its reduct through R-LET is $\hat{\pi}_1 \hat{\pi}_1 + 1$, which is stuck, so this program goes wrong. The primary purpose of type systems is to prevent such situations from arising. \square
- 10.1.13 REMARK: The configuration $!m/\mu$ is stuck if m is not in the domain of μ . In that case, however, $!m/\mu$ is not closed. Because we consider \rightarrow as a relation between closed configurations only, this situation cannot arise. In other words, the semantics of ML-the-calculus never allows the creation of *dangling pointers*. As a result, no particular precautions need be taken to guard against them. Several strongly typed programming languages do nevertheless allow dangling pointers in a controlled fashion (Tofte and Talpin, 1997; Walker, Crary, and Morrisett, 2000; DeLine and Fähndrich, 2001; Grossman, Morrisett, Jim, Hicks, Wang, and Cheney, 2002). \square

Damas and Milner's Type System

ML-the-type-system was originally defined by Milner (1978). Here, we reproduce the definition given a few years later by Damas and Milner (1982), which is written in a more standard style: typing judgments are defined inductively by a collection of typing rules. We refer to this type system as DM.

We must first define *types*. In DM, types are terms built out of *type constructors* and *type variables*. Furthermore, they are *first-order* terms: that is,

in the grammar of types, none of the productions *binds* a type variable. This situation is identical to that of the simply-typed λ -calculus.

We begin with several considerations concerning the specification of type constructors.

First, we do not wish to fix the set of type constructors. Certainly, since ML-the-calculus has functions, we need to be able to form an arrow type $T \rightarrow T'$ out of arbitrary types T and T' ; that is, we need a binary type constructor \rightarrow . However, because ML-the-calculus includes an unspecified set of constants, we cannot say much else in general. If constants include integer literals and integer operations, as in Example 10.1.1, then a nullary type constructor `int` is needed; if they include pair construction and destruction, as in Examples 10.1.3 and 10.1.5, then a binary type constructor \times is needed; etc.

Second, it is common to refer to the parameters of a type constructor by position, that is, by numeric index. For instance, when one writes $T \rightarrow T'$, it is understood that the type constructor \rightarrow has arity 2, that T is its first parameter, known as its *domain*, and that T' is its second parameter, known as its *codomain*. Here, however, we refer to parameters by names, known as *directions*. For instance, we define two directions *domain* and *codomain* and let the type constructor \rightarrow have arity $\{\text{domain}, \text{codomain}\}$. The extra generality afforded by directions is exploited in the definition of nonstructural subtyping (Example 10.2.9) and in the definition of rows (§10.8).

Last, we allow types to be classified using *kinds*. As a result, every type constructor must come not only with an arity, but with a richer *signature*, which describes the kinds of the types to which it is applicable and the kind of the type that it produces. A distinguished kind \star is associated with “normal” types, that is, types that are directly ascribed to expressions and values. For instance, the signature of the type constructor \rightarrow is $\{\text{domain} \mapsto \star, \text{codomain} \mapsto \star\} \Rightarrow \star$, because it is applicable to two normal types and produces a normal type. Introducing kinds other than \star allows viewing some types as ill-formed: this is illustrated, for instance, in §10.8. In the simplest case, however, \star is really the only kind, so the signature of a type constructor is nothing but its arity (a set of directions), and every term is a well-formed type, provided every application of a type constructor respects its arity.

- 10.1.14 DEFINITION: Let d range over a finite or denumerable set of *directions* and κ over a finite or denumerable set of *kinds*. Let \star be a distinguished kind. Let K range over partial mappings from directions to kinds. Let F range over a finite or denumerable set of *type constructors*, each of which has a *signature* of the form $K \Rightarrow \kappa$. The domain of K is called the *arity* of F , while κ is referred to as its *image kind*. We write κ instead of $K \Rightarrow \kappa$ when K is empty. Let \rightarrow be a type constructor of signature $\{\text{domain} \mapsto \star, \text{codomain} \mapsto \star\} \Rightarrow \star$. \square

The type constructors and their signatures collectively form a *signature* S . In the following, we assume that a fixed signature S is given and that every type constructor in it has *finite* arity, so as to ensure that types are machine representable. However, in §10.8, we shall explicitly work with several distinct signatures, one of which involves type constructors of denumerable arity.

A *type variable* is a name that is used to stand for a type. For simplicity, we assume that every type variable is branded with a kind, or in other words, that type variables of distinct kinds are drawn from disjoint sets. Each of these sets of type variables is individually subject to α -conversion: that is, renamings must preserve kinds. Attaching kinds to type variables is only a technical convenience; in practice, every operation performed during type inference preserves the property that every type is well-kinded, so it is not necessary to keep track of the kind of every type variable. It is only necessary to check that all types supplied by the programmer, within type declarations, type annotations, or module interfaces, are well-kinded.

10.1.15 DEFINITION: For every kind κ , let \mathcal{V}_κ be a disjoint, denumerable set of *type variables*. Let X, Y , and Z range over the set \mathcal{V} of all type variables. Let \bar{X} and \bar{Y} range over finite sets of type variables. We write $\bar{X}\bar{Y}$ for the set $\bar{X} \cup \bar{Y}$ and often write X for the singleton set $\{X\}$. We write $ftv(o)$ for the set of *free type variables* of an object o . \square

The set of types, ranged over by T , is the free many-kinded term algebra that arises out of the type constructors and type variables. Types are given by the following inductive definition:

10.1.16 DEFINITION: A *type* of kind κ is either a member of \mathcal{V}_κ , or a term of the form $F \{d_1 \mapsto T_1, \dots, d_n \mapsto T_n\}$, where F has signature $\{d_1 \mapsto \kappa_1, \dots, d_n \mapsto \kappa_n\} \Rightarrow \kappa$ and T_1, \dots, T_n are types of kind $\kappa_1, \dots, \kappa_n$, respectively. \square

As a notational convention, we assume that, for every type constructor F , the directions that form the arity of F are implicitly ordered, so that we may say that F has signature $\kappa_1 \otimes \dots \otimes \kappa_n \Rightarrow \kappa$ and employ the syntax $F T_1 \dots T_n$ for applications of F . Applications of the type constructor \rightarrow are written infix and associate to the right, so $T \rightarrow T' \rightarrow T''$ stands for $T \rightarrow (T' \rightarrow T'')$.

In order to give meaning to the free type variables of a type, or more generally, of a typing judgment, traditional presentations of ML-the-type-system, including Damas and Milner's, employ *type substitutions*. Most of our presentation avoids substitutions and uses constraints instead. However, we do need substitutions on a few occasions, especially when relating our presentation to Damas and Milner's.

10.1.17 DEFINITION: A *type substitution* θ is a total, kind-preserving mapping of type variables to types that is the identity everywhere but on a finite subset of \mathcal{V} ,

which we call the *domain* of θ and write $dom(\theta)$. The *range* of θ , which we write $range(\theta)$, is the set $ftv(\theta(dom(\theta)))$. A type substitution may naturally be viewed as a total, kind-preserving mapping of types to types. \square

If \vec{X} and \vec{T} are respectively a vector of distinct type variables and a vector of types of the same (finite) length such that, for every index i , X_i and T_i have the same kind, then $[\vec{X} \mapsto \vec{T}]$ denotes the substitution that maps X_i to T_i for every index i and is the identity elsewhere. The domain of $[\vec{X} \mapsto \vec{T}]$ is a subset of \vec{X} , the set underlying the vector \vec{X} . Its range is a subset of $ftv(\vec{T})$, where \vec{T} is the set underlying the vector \vec{T} . (These may be strict subsets; for instance, the domain of $[X \mapsto X]$ is the empty set, since this substitution is the identity.) Every substitution θ may be written under the form $[\vec{X} \mapsto \vec{T}]$, where $\vec{X} = dom(\theta)$. Then, θ is *idempotent* if and only if $\vec{X} \# ftv(\vec{T})$ holds.

As pointed out earlier, types are first-order terms. As a result, every type variable that appears within a type T appears *free* within T . Things become more interesting when we introduce *type schemes*. As its name implies, a type scheme may describe an entire family of types; this effect is achieved via *universal quantification* over a set of type variables.

10.1.18 DEFINITION: A type scheme S is an object of the form $\forall \vec{X}.T$, where T is a type of kind \star and the type variables \vec{X} are considered bound within T . Any type of the form $[\vec{X} \mapsto \vec{T}]T$ is called an *instance* of the type scheme $\forall \vec{X}.T$. \square

One may view the type T as the trivial type scheme $\forall \emptyset.T$, where no type variables are universally quantified, so types of kind \star may be viewed as a subset of type schemes. The type scheme $\forall \vec{X}.T$ may be viewed as a finite way of describing the possibly infinite family of its instances. Note that, throughout most of this chapter, we work with *constrained type schemes*, a generalization of DM type schemes (Definition 10.2.2).

Typing environments, or *environments* for short, are used to collect assumptions about an expression's free identifiers.

10.1.19 DEFINITION: An *environment* Γ is a finite ordered sequence of pairs of a program identifier and a type scheme. We write \emptyset for the empty environment and “;” for the concatenation of environments. An environment may be viewed as a finite mapping from program identifiers to type schemes by letting $\Gamma(x) = S$ if and only if Γ is of the form $\Gamma_1; x : S; \Gamma_2$, where Γ_2 contains no assumption about x . The set of *defined program identifiers* of an environment Γ , written $dpi(\Gamma)$, is defined by $dpi(\emptyset) = \emptyset$ and $dpi(\Gamma; x : S) = dpi(\Gamma) \cup \{x\}$. \square

To complete the definition of Damas and Milner's type system, there remains to define *typing judgments*. A typing judgment takes the form $\Gamma \vdash t : S$, where t is an expression of interest, Γ is an environment, which typically contains assumptions about t 's free program identifiers, and S is a type scheme.

$\frac{\Gamma(x) = S}{\Gamma \vdash x : S}$	(DM-VAR)	$\frac{\Gamma \vdash t_1 : S \quad \Gamma; z : S \vdash t_2 : T}{\Gamma \vdash \text{let } z = t_1 \text{ in } t_2 : T}$	(DM-LET)
$\frac{\Gamma; z : T \vdash t : T'}{\Gamma \vdash \lambda z. t : T \rightarrow T'}$	(DM-ABS)	$\frac{\Gamma \vdash t : T \quad \bar{X} \# \text{fv}(\Gamma)}{\Gamma \vdash t : \forall \bar{X}. T}$	(DM-GEN)
$\frac{\Gamma \vdash t_1 : T \rightarrow T' \quad \Gamma \vdash t_2 : T}{\Gamma \vdash t_1 t_2 : T'}$	(DM-APP)	$\frac{\Gamma \vdash t : \forall \bar{X}. T}{\Gamma \vdash t : [\bar{X} \mapsto \bar{T}]T}$	(DM-INST)

Figure 10-3: Typing rules for DM

Such a judgment may be read: *under assumptions* Γ , *the expression* t *has the type scheme* S . By abuse of language, it is sometimes said that t *has type* S . A typing judgment is *valid* (or *holds*) if and only if it may be derived using the rules that appear in Figure 10-3. An expression t is *well-typed* within the environment Γ if and only if there exists some type scheme S such that the judgment $\Gamma \vdash t : S$ holds; it is *ill-typed* within Γ otherwise.

Rule DM-VAR allows fetching a type scheme for an identifier x from the environment. It is equally applicable to program variables, memory locations, and constants. If no assumption concerning x appears in the environment Γ , then the rule isn't applicable. In that case, the expression x is ill-typed within Γ . Assumptions about constants are usually collected in a so-called *initial environment* Γ_0 . It is the environment under which closed programs are typechecked, so every subexpression is typechecked under some extension Γ of Γ_0 . Of course, the type schemes assigned by Γ_0 to constants must be consistent with their operational semantics; we say more about this later (§10.5). Rule DM-ABS specifies how to typecheck a λ -abstraction $\lambda z. t$. Its premise requires the body of the function, t , to be well-typed under an extra assumption that causes all free occurrences of z within t to receive a common type T . Its conclusion forms the arrow type $T \rightarrow T'$ out of the types of the function's formal parameter, T , and result, T' . It is worth noting that this rule always augments the environment with a type T —recall that, by convention, types form a subset of type schemes—but never with a nontrivial type scheme. Rule DM-APP states that the type of a function application is the codomain of the function's type, provided that the domain of the function's type is a valid type for the actual argument. Rule DM-LET closely mirrors the operational semantics: whereas the semantics of the local definition $\text{let } z = t_1 \text{ in } t_2$ is to augment the runtime environment by binding z to the value of t_1 prior to evaluating t_2 , the effect of DM-LET is to augment the typing envi-

ronment by binding z to a type scheme for τ_1 prior to typechecking τ_2 . Rule DM-GEN turns a type into a type scheme by universally quantifying over a set of type variables that do not appear free in the environment; this restriction is discussed in Example 10.1.20 below. Rule DM-INST, on the contrary, turns a type scheme into one of its instances, which may be chosen arbitrarily. These two operations are referred to as *generalization* and *instantiation*. The notion of type scheme and the rules DM-GEN and DM-INST are characteristic of ML-the-type-system: they distinguish it from the simply-typed λ -calculus.

10.1.20 EXAMPLE: It is unsound to allow generalizing type variables that appear free in the environment. For instance, consider the typing judgment $z : X \vdash z : X$ (1), which, according to DM-VAR, is valid. Applying an unrestricted version of DM-GEN to it, we obtain $z : X \vdash z : \forall X.X$ (2), whence, by DM-INST, $z : X \vdash z : Y$ (3). By DM-ABS and DM-GEN, we then have $\emptyset \vdash \lambda z.z : \forall XY.X \rightarrow Y$. In other words, the identity function has unrelated argument and result types! Then, the expression $(\lambda z.z) \hat{0} \hat{0}$, which reduces to the stuck expression $\hat{0} \hat{0}$, has type scheme $\forall Z.Z$. So, well-typed programs may cause runtime errors: the type system is unsound.

What happened? It is clear that the judgment (1) is correct only because the type assigned to z is the *same* in its assumption and in its right-hand side. For the same reason, the judgments (2) and (3)—the former of which may be written $z : X \vdash z : \forall Y.Y$ —are incorrect. Indeed, such judgments defeat the very purpose of environments, since they disregard their assumption. By universally quantifying over X in the right-hand side only, we break the connection between occurrences of X in the assumption, which remain free, and occurrences in the right-hand side, which become bound. This is correct only if there are in fact no free occurrences of X in the assumption. \square

10.1.21 REMARK: A naive implementation of DM-GEN would traverse the environment Γ in order to compute the set of its free type variables. However, the number of entries in Γ may be linear in the program size, so, even if types have bounded size, the time required by this computation may be linear in the program size. Since it is performed at every `let` node, this naive approach gives type inference *quadratic* time complexity. To avoid this pitfall, our constraint solver annotates every type variable with an integer rank, which allows telling, in constant time, whether it appears free in Γ (page 444). \square

It is a key feature of ML-the-type-system that DM-ABS may only introduce a type T , rather than a type scheme, into the environment. Indeed, this allows the rule's conclusion to form the arrow type $T \rightarrow T'$. If instead the rule were to introduce the assumption $z : S$ into the environment, then its conclusion would have to form $S \rightarrow T'$, which is not a well-formed type. In other words,

this restriction is necessary to preserve the stratification between types and type schemes. If we were to remove this stratification, thus allowing universal quantifiers to appear deep inside types, we would obtain an implicitly-typed version of System F (*TAPL*, Chapter 23). Type inference for System F is undecidable (Wells, 1999), while type inference for ML-the-type-system is decidable, as we show later, so this design choice has a rather drastic impact.

10.1.22 EXERCISE [RECOMMENDED, ★]: Build a type derivation for the expression $\lambda z_1. \text{let } z_2 = z_1 \text{ in } z_2$. \square

10.1.23 EXERCISE [RECOMMENDED, ★]: Let int be a nullary type constructor of signature \star . Let Γ_0 consist of the bindings $\hat{f} : \text{int} \rightarrow \text{int} \rightarrow \text{int}$ and $\hat{k} : \text{int}$, for every integer k . Can you find derivations of the following valid typing judgments? Which of these judgments are valid in the simply-typed λ -calculus, where $\text{let } z = t_1 \text{ in } t_2$ is syntactic sugar for $(\lambda z. t_2) t_1$?

$$\begin{aligned} \Gamma_0 &\vdash \lambda z. z : \text{int} \rightarrow \text{int} \\ \Gamma_0 &\vdash \lambda z. z : \forall X. X \rightarrow X \\ \Gamma_0 &\vdash \text{let } f = \lambda z. z \hat{1} \text{ in } f \hat{2} : \text{int} \\ \Gamma_0 &\vdash \text{let } f = \lambda z. z \text{ in } f f \hat{2} : \text{int} \end{aligned}$$

Show that the expressions $\hat{1} \hat{2}$ and $\lambda f. (f f)$ are ill-typed within Γ_0 . Could these expressions be well-typed in a more powerful type system? \square

DM enjoys a number of nice theoretical properties, which have practical implications.

First, it is sound: that is, well-typed programs do not go wrong. This essential property ensures that programs that are accepted by the typechecker may be compiled without runtime checks. Establishing this property requires (i) suitable hypotheses about the semantics of constants and the type schemes assigned to constants in the initial environment, and (ii) in the presence of side effects, a slight restriction of the syntax of `let` constructs, known as the *value restriction*.

Furthermore, there exists an algorithm that, given a (closed) environment Γ and a program t , tells whether t is well-typed with respect to Γ , and if so, produces a *principal* type scheme S . A principal type scheme is such that (i) it is valid, that is, $\Gamma \vdash t : S$ holds, and (ii) it is most general, that is, every judgment of the form $\Gamma \vdash t : S'$ follows from $\Gamma \vdash t : S$ by DM-INST and DM-GEN. (For the sake of simplicity, we have stated the properties of the type inference algorithm only in the case of a closed environment Γ ; the specification is slightly heavier in the general case.) This implies that type inference is decidable: the compiler need not require expressions to be annotated with types. The fact that, under a fixed environment Γ , all of the type information associated with

an expression t may be summarized in the form of a single, principal type scheme is also key to modular programming. Indeed, exporting a value out of a module requires explicitly assigning a type scheme to it as part of the module's signature. If the chosen type scheme is not principal, then part of the value's (hence, of the module's) potential for reuse is lost.

Road Map

Before proving the above claims, we first generalize our presentation by moving to a *constraint-based* setting. The necessary tools—the constraint language, its interpretation, and a number of constraint equivalence laws—are introduced in §10.2. In §10.3, we describe the standard constraint-based type system $HM(X)$ (Odersky, Sulzmann, and Wehr, 1999). We prove that, when constraints are made up of equations between free, finite terms, $HM(X)$ is a reformulation of DM. In the presence of a more powerful constraint language, $HM(X)$ is an extension of DM. In §10.4, we show that type inference may be viewed as a combination of constraint generation and constraint solving, as promised earlier. Then, in §10.5, we give a type soundness theorem. It is stated purely in terms of constraints, but—thanks to the results developed in the previous sections—applies equally to $HM(X)$ and DM.

Throughout this core material, the syntax and interpretation of constraints are left partly unspecified. Thus, the development is *parameterized* with respect to them—hence the unknown X in the name $HM(X)$. We really describe a family of constraint-based type systems, all of which share a common constraint generator and a common type soundness proof. Constraint solving, however, cannot be independent of X : on the contrary, the design of an efficient solver is heavily dependent on the syntax and interpretation of constraints. In §10.6, we consider constraint solving in the particular case where constraints are made up of equations interpreted in a free tree model, and define a constraint solver on top of a standard first-order unification algorithm.

The remainder of this chapter deals with extensions of the framework. In §10.7, we explain how to extend ML-the-calculus with a number of features, including products, sums, references, recursion, algebraic data types, and recursive types. Last, in §10.8, we extend the type language with *rows* and use them to assign polymorphic type schemes to operations on records and variants.

$\sigma ::=$	<i>type scheme:</i>	$C, D ::=$	<i>Syntactic sugar for constraints:</i>
$\forall \tilde{X}[C].T$		\dots	<i>As before</i>
$C, D ::=$	<i>constraint:</i>	$\sigma \leq T$	<i>Definition 10.2.3</i>
true	<i>truth</i>	let $x : \sigma$ in C	<i>Definition 10.2.3</i>
false	<i>falsity</i>	$\exists \sigma$	<i>Definition 10.2.3</i>
$P T_1 \dots T_n$	<i>predicate application</i>	def Γ in C	<i>Definition 10.2.4</i>
$C \wedge C$	<i>conjunction</i>	let Γ in C	<i>Definition 10.2.4</i>
$\exists \tilde{X}.C$	<i>existential quantification</i>	$\exists \Gamma$	<i>Definition 10.2.4</i>
def $x : \sigma$ in C	<i>type scheme introduction</i>		
$x \leq T$	<i>type scheme instantiation</i>		

Figure 10-4: Syntax of type schemes and constraints

10.2 Constraints

In this section, we define the syntax and logical meaning of constraints. Both are partly unspecified. Indeed, the set of *type constructors* (Definition 10.1.14) must contain at least the binary type constructor \rightarrow , but might contain more. Similarly, the syntax of constraints involves a set of so-called *predicates* on types, which we require to contain at least a binary *subtyping* predicate \leq , but might contain more. (The introduction of subtyping, which is absent in DM, has little impact on the complexity of our proofs, yet increases the framework's expressive power. When subtyping is not desired, we interpret the predicate \leq as equality.) The logical interpretation of type constructors and of predicates is left almost entirely unspecified. This freedom allows reasoning not only about Damas and Milner's type system, but also about a family of constraint-based extensions of it.

Syntax of Constraints

We now define the syntax of constrained type schemes and of constraints and introduce some extra constraint forms as syntactic sugar.

- 10.2.1 **DEFINITION:** Let P range over a finite or denumerable set of *predicates*, each of which has a *signature* of the form $\kappa_1 \otimes \dots \otimes \kappa_n \Rightarrow \cdot$, where $n \geq 0$. For every kind κ , let $=_\kappa$ and \leq_κ be distinguished predicates of signature $\kappa \otimes \kappa \Rightarrow \cdot$. \square
- 10.2.2 **DEFINITION:** The syntax of *type schemes* and *constraints* is given in Figure 10-4. It is further restricted by the following requirements. In the type scheme $\forall \tilde{X}[C].T$ and in the constraint $x \leq T$, the type T must have kind \star . In the con-

straint $P T_1 \dots T_n$, the types T_1, \dots, T_n must have kind $\kappa_1, \dots, \kappa_n$, respectively, if P has signature $\kappa_1 \otimes \dots \otimes \kappa_n \Rightarrow \cdot$. We write $\forall \bar{X}. T$ for $\forall \bar{X}[\text{true}]. T$, which allows viewing DM type schemes as a subset of constrained type schemes. \square

We write $T_1 =_{\kappa} T_2$ and $T_1 \leq_{\kappa} T_2$ for the binary predicate applications $=_{\kappa} T_1 T_2$ and $\leq_{\kappa} T_1 T_2$, and refer to them as equality and subtyping constraints, respectively. We often omit the subscript κ , so $T_1 = T_2$ and $T_1 \leq T_2$ are well-formed constraints whenever T_1 and T_2 have the same kind. By convention, \exists and def bind tighter than \wedge ; that is, $\exists \bar{X}. C \wedge D$ is $(\exists \bar{X}. C) \wedge D$ and $\text{def } x : \sigma \text{ in } C \wedge D$ is $(\text{def } x : \sigma \text{ in } C) \wedge D$. In $\forall \bar{X}[C]. T$, the type variables \bar{X} are bound within C and T . In $\exists \bar{X}. C$, the type variables \bar{X} are bound within C . The sets of free type variables of a type scheme σ and of a constraint C , written $ftv(\sigma)$ and $ftv(C)$, respectively, are defined accordingly. In $\text{def } x : \sigma \text{ in } C$, the identifier x is bound within C . The sets of free program identifiers of a type scheme σ and of a constraint C , written $fpi(\sigma)$ and $fpi(C)$, respectively, are defined accordingly. Note that x occurs free in the constraint $x \leq T$.

The constraint true , which is always satisfied, mainly serves to indicate the absence of a nontrivial constraint, while false , which has no solution, may be understood as the indication of a type error. Composite constraints include conjunction and existential quantification, which have their standard meaning, as well as *type scheme introduction* and *type scheme instantiation* constraints, which are similar to Gustavsson and Svenningsson's constraint abstractions (2001). In order to be able to explain these last two forms, we must first introduce a number of derived constraint forms:

10.2.3 DEFINITION: Let σ be $\forall \bar{X}[D]. T$. If $\bar{X} \# ftv(T')$ holds, then $\sigma \leq T'$ (read: T' is an instance of σ) stands for the constraint $\exists \bar{X}. (D \wedge T \leq T')$. We write $\exists \sigma$ (read: σ has an instance) for $\exists \bar{X}. D$ and let $x : \sigma \text{ in } C$ for $\exists \sigma \wedge \text{def } x : \sigma \text{ in } C$. \square

Constrained type schemes generalize Damas and Milner's type schemes, while this definition of instantiation constraints generalizes Damas and Milner's notion of instance (Definition 10.1.18). Let us draw a comparison. First, Damas and Milner's instance relation is binary (given a type scheme S and a type T , either T is an instance of S , or it isn't), and is purely syntactic. For instance, the type $Y \rightarrow Z$ is *not* an instance of $\forall X. X \rightarrow X$ in Damas and Milner's sense, because Y and Z are distinct type variables. In our presentation, on the other hand, $\forall X. X \rightarrow X \leq Y \rightarrow Z$ is not an assertion; rather, it is a constraint, which by definition is $\exists X. (\text{true} \wedge X \rightarrow X \leq Y \rightarrow Z)$. We later prove that it is equivalent to $\exists X. (Y \leq X \wedge X \leq Z)$ and to $Y \leq Z$, and, if subtyping is interpreted as equality, to $Y = Z$. That is, $\sigma \leq T'$ represents a condition on (the ground types denoted by) the type variables in $ftv(\sigma, T')$ for T' to be an instance of σ , in a logical, rather than purely syntactic, sense. Second, the definition of instantiation

constraints involves subtyping, to ensure that any supertype of an instance of σ is again an instance of σ (see rule C-EXTRANS on page 418). This is consistent with the purpose of subtyping: to allow a subtype where a supertype is expected (*TAPL*, Chapter 15). Third and last, every type scheme σ is now of the form $\forall \tilde{X}[C].T$. The constraint C , whose free type variables may or may not be members of \tilde{X} , is meant to restrict the set of instances of the type scheme $\forall \tilde{X}[C].T$. This is evident in the instantiation constraint $\forall \tilde{X}[C].T \preceq T'$, which by Definition 10.2.3 stands for $\exists \tilde{X}.(C \wedge T \preceq T')$: the values that \tilde{X} may assume are restricted by the demand that C be satisfied. This requirement vanishes in the case of DM type schemes, where C is true. Our notions of constrained type scheme and of instantiation constraint are standard, coinciding with those of $HM(X)$ (Odersky, Sulzmann, and Wehr, 1999).

Let us now come back to an explanation of type scheme introduction and instantiation constraints. In brief, the construct $\text{def } x : \sigma \text{ in } C$ binds the name x to the type scheme σ within the constraint C . If C contains a subconstraint of the form $x \preceq T$, where this occurrence of x is free in C , then this subconstraint acquires the meaning $\sigma \preceq T$. Thus, the constraint $x \preceq T$ is indeed an instantiation constraint, where the type scheme that is being instantiated is referred to by name. The constraint $\text{def } x : \sigma \text{ in } C$ may be viewed as an *explicit substitution* of the type scheme σ for the name x within C . Later (§10.4), we use such explicit substitutions to supplant typing environments. That is, where Damas and Milner's type system augments the current typing environment (DM-ABS, DM-LET), we introduce a new *def* binding in the current constraint; where it looks up the current typing environment (DM-VAR), we employ an instantiation constraint. (The reader may wish to look ahead at Figure 10-9 on page 431.) The point is that it is then up to a constraint solver to choose a strategy for reducing explicit substitutions—for instance, one might wish to simplify σ before substituting it for x within C —whereas the use of environments in standard type systems such as DM and $HM(X)$ imposes an eager substitution strategy, which is inefficient and thus never literally implemented. The use of type scheme introduction and instantiation constraints allows separating constraint generation and constraint solving without compromising efficiency, or, in other words, without introducing a gap between the description of the type inference algorithm and its actual implementation. Although the algorithm that we plan to describe is not new (Rémy, 1992a), its description in terms of constraints is: to the best of our knowledge, the only close relative of our *def* constraints is to be found in Gustavsson and Svenningsson (2001). An earlier work that contains similar ideas is Müller (1994). Approaches based on semi-unification (Henglein, 1989, 1993) achieve a similar separation between constraint generation and constraint solving, but are based on a rather different constraint language.

In the type system of Damas and Milner, every type scheme S has a fixed, nonempty set of instances. In a constraint-based setting, things are more complex: given a type scheme σ and a type T , whether T is an instance of σ (that is, whether the constraint $\sigma \leq T$ is satisfied) depends on the meaning assigned to the type variables in $ftv(\sigma, T)$. Similarly, given a type scheme, whether *some* type is an instance of σ (that is, whether the constraint $\exists Z. \sigma \leq Z$, where Z is fresh for σ , is satisfied) depends on the meaning assigned to the type variables in $ftv(\sigma)$. Because we do not wish to allow forming type schemes that have no instances, we often use the constraint $\exists Z. \sigma \leq Z$. In fact, we later prove that it is equivalent to $\exists \sigma$, as defined above. We also use the constraint form $\text{let } x : \sigma \text{ in } C$, which requires σ to have an instance and at the same time associates it with the name x . Because the def form is more primitive, it is easier to work with at a low level, but it is no longer explicitly used after §10.2; we always use let instead.

10.2.4 DEFINITION: Environments Γ remain as in Definition 10.1.19, except DM type schemes S are replaced with constrained type schemes σ . The set of *free program identifiers* of an environment Γ , written $fpi(\Gamma)$, is defined by $fpi(\emptyset) = \emptyset$ and $fpi(\Gamma; x : \sigma) = fpi(\Gamma) \cup fpi(\sigma)$. We write $dfpi(\Gamma)$ for $dpi(\Gamma) \cup fpi(\Gamma)$. We define $\text{def } \emptyset \text{ in } C$ as C and $\text{def } \Gamma; x : \sigma \text{ in } C$ as $\text{def } \Gamma \text{ in } \text{def } x : \sigma \text{ in } C$. Similarly, we define $\text{let } \emptyset \text{ in } C$ as C and $\text{let } \Gamma; x : \sigma \text{ in } C$ as $\text{let } \Gamma \text{ in } \text{let } x : \sigma \text{ in } C$. We define $\exists \emptyset$ as true and $\exists(\Gamma; x : \sigma)$ as $\exists \Gamma \wedge \text{def } \Gamma \text{ in } \exists \sigma$. \square

In order to establish or express certain laws of equivalence between constraints, we need *constraint contexts*. A constraint context is a constraint with zero, one, or several *holes*, written $[\]$. The syntax of contexts is as follows:

$$\mathbb{C} ::= [\] \mid C \mid C \wedge C \mid \exists \bar{X}. C \mid \text{def } x : \sigma \text{ in } \mathbb{C} \mid \text{def } x : \forall \bar{X}[C]. T \text{ in } C$$

The application of a constraint context \mathbb{C} to a constraint C , written $\mathbb{C}[C]$, is defined in the usual way. Because a constraint context may have any number of holes, C may disappear or be duplicated in the process. Because a hole may appear in the scope of a binder, some of C 's free type variables and free program identifiers may become bound in $\mathbb{C}[C]$. We write $dTV(\mathbb{C})$ and $dPI(\mathbb{C})$ for the sets of type variables and program identifiers, respectively, that \mathbb{C} may thus capture. We write $\text{let } x : \forall \bar{X}[C]. T \text{ in } C$ for $\exists \bar{X}. C \wedge \text{def } x : \forall \bar{X}[C]. T \text{ in } C$. (Being able to state such a definition is why we require multi-hole contexts.) We let \bar{X} range over *existential constraint contexts*, defined by $\bar{X} ::= \exists \bar{X}. [\]$.

Meaning of Constraints

We have defined the syntax of constraints and given an informal description of their meaning. We now give a formal definition of the interpretation of constraints. We begin with the definition of a *model*:

- 10.2.5 DEFINITION: For every kind κ , let \mathcal{M}_κ be a nonempty set, whose elements are called the *ground types* of kind κ . In the following, t ranges over \mathcal{M}_κ , for some κ that may be determined from the context. For every type constructor F of signature $K \Rightarrow \kappa$, let F denote a total function from \mathcal{M}_K into \mathcal{M}_κ , where the indexed product \mathcal{M}_K is the set of all mappings T of domain $\text{dom}(K)$ that map every $d \in \text{dom}(K)$ to an element of $\mathcal{M}_{K(d)}$. For every predicate symbol P of signature $\kappa_1 \otimes \dots \otimes \kappa_n \Rightarrow \cdot$, let P denote a predicate on $\mathcal{M}_{\kappa_1} \times \dots \times \mathcal{M}_{\kappa_n}$. For every kind κ , we require the predicate $=_\kappa$ to be equality on \mathcal{M}_κ and the predicate \leq_κ to be a partial order on \mathcal{M}_κ . \square

For the sake of convenience, we abuse notation and write F for both the type constructor and its interpretation, and similarly for predicates.

By varying the set of type constructors, the set of predicates, the set of ground types, and the interpretation of type constructors and predicates, one may define an entire family of related type systems. We refer to the collection of these choices as X . Thus, the type system $\text{HM}(X)$, described in §10.3, is *parameterized by X* .

The following examples give standard ways of defining the set of ground types and the interpretation of type constructors.

- 10.2.6 EXAMPLE [SYNTACTIC MODELS]: For every kind κ , let \mathcal{M}_κ consist of the *closed* types of kind κ . Then, ground types are types that do not have any free type variables, and form the so-called *Herbrand universe*. Let every type constructor F be interpreted as itself. Models that define ground types and interpret type constructors in this manner are referred to as *syntactic*. \square
- 10.2.7 EXAMPLE [TREE MODELS]: Let a *path* π be a finite sequence of directions. The empty path is written ϵ and the concatenation of the paths π and π' is written $\pi \cdot \pi'$. Let a *tree* be a partial function t from paths to type constructors whose domain is nonempty and prefix-closed and such that, for every path π in the domain of t , if the type constructor $t(\pi)$ has signature $K \Rightarrow \kappa$, then $\pi \cdot d \in \text{dom}(t)$ is equivalent to $d \in \text{dom}(K)$ and, furthermore, for every $d \in \text{dom}(K)$, the type constructor $t(\pi \cdot d)$ has image kind $K(d)$. If π is in the domain of t , then the *subtree* of t rooted at π , written t/π , is the partial function $\pi' \mapsto t(\pi \cdot \pi')$. A tree is *finite* if and only if it has finite domain. A tree is *regular* if and only if it has a finite number of distinct subtrees. Every finite tree is thus regular. Let \mathcal{M}_κ consist of the *finite* (respectively *regular*) trees t such that $t(\epsilon)$ has image kind κ : then, we have a *finite* (respectively *regular*) *tree model*.

If F has signature $K \Rightarrow \kappa$, one may interpret F as the function that maps $T \in \mathcal{M}_K$ to the ground type $t \in \mathcal{M}_\kappa$ defined by $t(\epsilon) = F$ and $t/d = T(d)$ for $d \in \text{dom}(T)$, that is, the unique ground type whose head symbol is F and

whose subtree rooted at d is $T(d)$. Then, we have a *free* tree model. Note that free finite tree models coincide with syntactic models, as defined in the previous example. \square

Rows (§10.8) are interpreted in a tree model, albeit not a free one. The following examples suggest different ways of interpreting the subtyping predicate.

10.2.8 EXAMPLE [EQUALITY MODELS]: The simplest way of interpreting the subtyping predicate is to let \leq denote equality on every \mathcal{M}_κ . Models that do so are referred to as *equality models*. When no predicate other than equality is available, we say that the model is *equality-only*. \square

10.2.9 EXAMPLE [STRUCTURAL, NONSTRUCTURAL SUBTYPING]: Let a *variance* ν be a nonempty subset of $\{-, +\}$, written $-$ (*contravariant*), $+$ (*covariant*), or \pm (*invariant*) for short. Define the *composition* of two variances as an associative, commutative operation with $+$ as neutral element, \pm as absorbing element (that is, $\pm- = \pm+ = \pm\pm = \pm$), and such that $-- = +$. Now, consider a free (finite or regular) tree model, where every direction d comes with a fixed variance $\nu(d)$. Define the variance $\nu(\pi)$ of a path π as the composition of the variances of its elements. Let \leq be a partial order on type constructors such that (i) if $F_1 \leq F_2$ holds and F_1 and F_2 have signature $K_1 \Rightarrow \kappa_1$ and $K_2 \Rightarrow \kappa_2$, respectively, then K_1 and K_2 agree on the intersection of their domains and κ_1 and κ_2 coincide; and (ii) $F_0 \leq F_1 \leq F_2$ implies $\text{dom}(F_0) \cap \text{dom}(F_2) \subseteq \text{dom}(F_1)$. Let \leq^+ , \leq^- , and \leq^\pm stand for \leq , \geq , and $=$, respectively. Then, define the interpretation of subtyping as follows: if $t_1, t_2 \in \mathcal{M}_\kappa$, let $t_1 \leq t_2$ hold if and only if, for every path $\pi \in \text{dom}(t_1) \cap \text{dom}(t_2)$, $t_1(\pi) \leq^{\nu(\pi)} t_2(\pi)$ holds. It is not difficult to check that \leq is a partial order on every \mathcal{M}_κ . The reader is referred to Amadio and Cardelli (1993), Kozen, Palsberg, and Schwartzbach (1995), and Brandt and Henglein (1997) for more details about this construction. Models that define subtyping in this manner are referred to as *nonstructural subtyping models*.

A simple nonstructural subtyping model is obtained by: letting the directions *domain* and *codomain* be contra- and covariant, respectively; introducing, in addition to the type constructor \rightarrow , two type constructors \perp and \top of signature \star ; and letting $\perp \leq \rightarrow \leq \top$. This gives rise to a model where \perp is the least ground type, \top is the greatest ground type, and the arrow type constructor is, as usual, contravariant in its domain and covariant in its codomain. This form of subtyping is called *nonstructural* because comparable ground types may have different shapes: consider, for instance, \perp and $\perp \rightarrow \top$.

A typical use of nonstructural subtyping is in type systems for records. One may, for instance, introduce a covariant direction *content* of kind \star , a kind

\circ , a type constructor abs of signature \circ , a type constructor pre of signature $\{\text{content} \mapsto \star\} \Rightarrow \circ$, and let $\text{pre} \leq \text{abs}$. This gives rise to a model where $\text{pre } t \leq \text{abs}$ holds for every $t \in \mathcal{M}_*$. Again, comparable ground types may have different shapes: consider, for instance, $\text{pre } \top$ and abs . §10.8 says more about typechecking operations on records.

Nonstructural subtyping has been studied, for example, in Kozen, Palsberg, and Schwartzbach (1995), Palsberg, Wand, and O’Keefe (1997), Jim and Palsberg (1999), Pottier (2001b), Su et al. (2002), and Niehren and Priesnitz (2003).

An important particular case arises when any two type constructors related by \leq have the same arity (and thus also the same signatures). In that case, it is not difficult to show that any two ground types related by subtyping must have the same shape, that is, if $t_1 \leq t_2$ holds, then $\text{dom}(t_1)$ and $\text{dom}(t_2)$ must coincide. For this reason, such an interpretation of subtyping is usually referred to as *atomic* or *structural* subtyping. It has been studied in the finite (Mitchell, 1984, 1991b; Tiuryn, 1992; Pratt and Tiuryn, 1996; Frey, 1997; Rehof, 1997; Kuncak and Rinard, 2003; Simonet, 2003) and regular (Tiuryn and Wand, 1993) cases. Structural subtyping is often used in automated program analyses that enrich standard types with atomic annotations without altering their shape. \square

Many other kinds of constraints exist, which we lack space to list; see Comon (1994) for a short survey.

Throughout this chapter, we assume (unless otherwise stated) that the set of type constructors, the set of predicates, and the model—which, together, form the parameter X —are arbitrary, but fixed.

As usual, the meaning of a constraint is a function of the meaning of its free type variables and of its free program identifiers, which are respectively given by a *ground assignment* and a *ground environment*.

10.2.10 **DEFINITION:** A *ground assignment* ϕ is a total, kind-preserving mapping from \mathcal{V} into \mathcal{M} . Ground assignments are extended to types by $\phi(F T_1 \dots T_n) = F(\phi(T_1), \dots, \phi(T_n))$. Then, for every type T of kind κ , $\phi(T)$ is a ground type of kind κ .

A *ground type scheme* s is a set of ground types, which we require to be *upward-closed* with respect to subtyping: that is, $t \in s$ and $t \leq t'$ must imply $t' \in s$. A *ground environment* ψ is a partial mapping from identifiers to ground type schemes.

Because the syntax of type schemes and constraints is mutually recursive, so is their interpretation. The interpretation of a type scheme σ under a ground assignment ϕ and a ground environment ψ is a ground type scheme, written $(\phi, \psi)\sigma$. It is defined in Figure 10-5. The \uparrow is the upward closure

<p>Interpretation of type schemes:</p> $(\phi, \psi)(\forall \bar{X}[C].T) = \uparrow \{ \phi[\bar{X} \mapsto \bar{t}](T) ; \phi[\bar{X} \mapsto \bar{t}], \psi \models C \}$ <p>Interpretation of constraints:</p> $\frac{}{\phi, \psi \models \text{true}} \quad (\text{CM-TRUE})$ $\frac{P(\phi(T_1), \dots, \phi(T_n))}{\phi, \psi \models P T_1 \dots T_n} \quad (\text{CM-PREDICATE})$	$\frac{\phi, \psi \models C_1 \quad \phi, \psi \models C_2}{\phi, \psi \models C_1 \wedge C_2} \quad (\text{CM-AND})$ $\frac{\phi[\bar{X} \mapsto \bar{t}], \psi \models C}{\phi, \psi \models \exists \bar{X}. C} \quad (\text{CM-EXISTS})$ $\frac{\phi, \psi[x \mapsto (\phi, \psi)\sigma] \models C}{\phi, \psi \models \text{def } x : \sigma \text{ in } C} \quad (\text{CM-DEF})$ $\frac{\phi(T) \in \psi(x)}{\phi, \psi \models x \leq T} \quad (\text{CM-INSTANCE})$
--	--

Figure 10-5: Meaning of constraints

operator and \models is the constraint satisfaction predicate, defined next. The interpretation of a constraint C under a ground assignment ϕ and a ground environment ψ is a truth value, written $\phi, \psi \models C$ (read: ϕ and ψ *satisfy* C). The three-place predicate \models is defined by the rules in Figure 10-5. A constraint C is *satisfiable* if and only if $\phi, \psi \models C$ holds for some ϕ and ψ . It is *false* (or *unsatisfiable*) otherwise. \square

Let us now explain these definitions. The interpretation of the type scheme $\forall \bar{X}[C].T$ is a set of ground types, which we may refer to as the type scheme's *ground instances*. It contains the images of T under extensions of ϕ with new values for the universally quantified variables \bar{X} ; these values may be arbitrary, but must be such that the constraint C is satisfied. We implicitly require \bar{X} and \bar{t} to have matching kinds, so that $\phi[\bar{X} \mapsto \bar{t}]$ remains a kind-preserving ground assignment. This set is upward closed, so any ground type that lies above a ground instance of σ is also a ground instance of σ . This interpretation is standard; see, for example, Pottier (2001a).

The rules that define \models (Figure 10-5) are syntax-directed. CM-TRUE states that the constraint `true` is a tautology, that is, holds in every context. No rule matches the constraint `false`, which means that it holds in no context. CM-PREDICATE states that the meaning of a predicate application is given by the predicate's interpretation within the model. More specifically, if P 's signature is $\kappa_1 \otimes \dots \otimes \kappa_n \Rightarrow \cdot$, then, by well-formedness of the constraint, every T_i is of kind κ_i , so $\phi(T_i)$ is a ground type in \mathcal{M}_{κ_i} . By Definition 10.2.5, P denotes a predicate on $\mathcal{M}_{\kappa_1} \times \dots \times \mathcal{M}_{\kappa_n}$, so the rule's premise is mathematically well-formed. It is independent of ψ , which is natural, since a predicate application has no free program identifiers. CM-AND requires each of the conjuncts to be

valid in isolation. CM-EXISTS allows the type variables \vec{X} to denote arbitrary ground types \vec{t} within C , independently of their image through ϕ . CM-DEF deals with type scheme introduction constraints, of the form $\text{def } x : \sigma$ in C . It binds x , within C , to the ground type scheme currently denoted by σ . Last, CM-INSTANCE concerns type scheme instantiation constraints of the form $x \leq T$. Such a constraint is valid if and only if the ground type denoted by T is a member of the ground type scheme denoted by x .

It is possible to prove that the constraints $\text{def } x : \sigma$ in C and $[x \mapsto \sigma]C$ have the same meaning, where the latter denotes the capture-avoiding substitution of σ for x throughout C . As a matter of fact, it would have been possible to use this equivalence as a definition of the meaning of def constraints, but the present style is pleasant as well. This confirms our claim that the def form is an explicit substitution form.

Because constraints lie at the heart of our treatment of ML-the-type-system, most of our proofs involve establishing logical properties of constraints. These properties are usually not stated in terms of the satisfaction predicate \models , which is too low-level. Instead, we reason in terms of *entailment* or *equivalence* assertions. Let us first define these notions.

- 10.2.11 DEFINITION: We write $C_1 \Vdash C_2$, and say that C_1 *entails* C_2 , if and only if, for every ground assignment ϕ and for every ground environment ψ , the assertion $\phi, \psi \models C_1$ implies $\phi, \psi \models C_2$. We write $C_1 \equiv C_2$, and say that C_1 and C_2 are *equivalent*, if and only if $C_1 \Vdash C_2$ and $C_2 \Vdash C_1$ hold. \square

In other words, C_1 entails C_2 when C_1 imposes stricter requirements on its free type variables and program identifiers than C_2 does. Note that C is unsatisfiable if and only if $C \equiv \text{false}$ holds. It is straightforward to check that entailment is reflexive and transitive and that \equiv is indeed an equivalence relation.

We immediately exploit the notion of constraint equivalence to define what it means for a type constructor to be *covariant*, *contravariant*, or *invariant* with respect to one of its parameters. Let F be a type constructor of signature $\kappa_1 \otimes \dots \otimes \kappa_n \Rightarrow \kappa$. Let $i \in \{1, \dots, n\}$. F is covariant (respectively contravariant, invariant) with respect to its i^{th} parameter if and only if, for all types T_1, \dots, T_n and T'_i of appropriate kinds, the constraint $F T_1 \dots T_i \dots T_n \leq F T_1 \dots T'_i \dots T_n$ is equivalent to $T_i \leq T'_i$ (respectively $T'_i \leq T_i$, $T_i = T'_i$).

- 10.2.12 EXERCISE [\star , \rightsquigarrow]: Check the following facts: (i) in an equality model, covariance, contravariance, and invariance coincide; (ii) in an equality free tree model, every type constructor is invariant with respect to each of its parameters; and (iii) in a nonstructural subtyping model, if the direction d has been declared covariant (respectively contravariant, invariant), then every type con-

structor whose arity includes d is covariant (respectively contravariant, invariant) with respect to d . \square

In the following, we require the type constructor \rightarrow to be contravariant with respect to its domain and covariant with respect to its codomain—a standard requirement in type systems with subtyping (TAPL, Chapter 15). This requirement is summed up by the following equivalence law:

$$\tau_1 \rightarrow \tau_2 \leq \tau'_1 \rightarrow \tau'_2 \equiv \tau'_1 \leq \tau_1 \wedge \tau_2 \leq \tau'_2 \quad (\text{C-ARROW})$$

Note that this requirement bears on the interpretation of types and of the subtyping predicate. In an equality free tree model, by (i) and (ii) in the exercise above, it is always satisfied. In a nonstructural subtyping model, it boils down to requiring that the directions *domain* and *codomain* be declared contravariant and covariant, respectively. In the general case, we do not have any knowledge of the model and cannot formulate a more precise requirement. Thus, it is up to the designer of the model to ensure that C-ARROW holds.

We also exploit the notion of constraint equivalence to define what it means for two type constructors to be *incompatible*. Two type constructors F_1 and F_2 with the same image kind are incompatible if and only if all constraints of the form $F_1 \vec{\tau}_1 \leq F_2 \vec{\tau}_2$ and $F_2 \vec{\tau}_2 \leq F_1 \vec{\tau}_1$ are false. Note that in an equality free tree model, any two distinct type constructors are incompatible. In the following, we often indicate that a newly introduced type constructor must be *isolated*. We implicitly require that, whenever both F_1 and F_2 are isolated, F_1 and F_2 be incompatible. Thus, the notion of isolation provides a concise and modular way of stating a collection of incompatibility requirements. We require the type constructor \rightarrow to be isolated.

Reasoning with Constraints

In this section, we give a number of equivalence laws that are often useful and help understand the meaning of constraints. To begin, we note that entailment is preserved by arbitrary constraint contexts, as stated by the theorem below. As a result, constraint equivalence is a congruence. Throughout this chapter, these facts are often used implicitly.

10.2.13 THEOREM [CONGRUENCE]: $C_1 \Vdash C_2$ implies $\mathbb{C}[C_1] \Vdash \mathbb{C}[C_2]$. \square

Next, we define what it means for a constraint to determine a set of type variables. In brief, C determines \bar{Y} if and only if, given a ground assignment for $ftv(C) \setminus \bar{Y}$ and given that C holds, it is possible to reconstruct, in a unique way, a ground assignment for \bar{Y} . Determinacy appears in the equivalence law C-LETTALL on page 418 and is exploited by the constraint solver in §10.6.

- 10.2.14 DEFINITION: C determines \bar{Y} if and only if, for every environment Γ , two ground assignments that satisfy $\text{def } \Gamma$ in C and that coincide outside \bar{Y} must coincide on \bar{Y} as well. \square

We now give a toolbox of constraint equivalence laws. It is worth noting that they do not form a complete axiomatization of constraint equivalence; in fact, they cannot, since the syntax and meaning of constraints is partly unspecified.

- 10.2.15 THEOREM: All equivalence laws in Figure 10-6 hold. \square

Let us explain. C-AND and C-ANDAND state that conjunction is commutative and associative. C-DUP states that redundant conjuncts may be freely added or removed, where a conjunct is redundant if and only if it is entailed by another conjunct. Throughout this chapter, these three laws are often used implicitly. C-EXEX and C-EX* allow grouping consecutive existential quantifiers and suppressing redundant ones, where a quantifier is redundant if and only if the variables bound by it do not occur free within its scope. C-EXAND allows conjunction and existential quantification to commute, provided no capture occurs; it is known as a *scope extrusion* law. When the rule is oriented from left to right, its side-condition may always be satisfied by suitable α -conversion. C-EXTRANS states that it is equivalent for a type T to be an instance of σ or to be a supertype of some instance of σ . We note that the instances of a monotype are its supertypes, that is, by Definition 10.2.3, $T' \leq T$ and $T' \leq T$ are equivalent. As a result, specializing C-EXTRANS to the case where σ is a monotype, we find that $T' \leq T$ is equivalent to $\exists Z.(T' \leq Z \wedge Z \leq T)$, for fresh Z , a standard equivalence law. When oriented from left to right, it becomes an interesting simplification law: in a chain of subtyping constraints, an intermediate variable such as Z may be suppressed, provided it is local, as witnessed by the existential quantifier $\exists Z$. C-INID states that, within the scope of the binding $x : \sigma$, every free occurrence of x may be safely replaced with σ . The restriction to free occurrences stems from the side-condition $x \notin \text{dpi}(C)$. When the rule is oriented from left to right, its other side-conditions, which require the context $\text{let } x : \sigma \text{ in } C$ not to capture σ 's free type variables or free program identifiers, may always be satisfied by suitable α -conversion. C-IN* complements the previous rule by allowing redundant let bindings to be simplified. We note that C-INID and C-IN* provide a simple procedure for eliminating let forms. C-INAND states that the let form commutes with conjunction; C-INAND* spells out a common particular case. C-INEX states that it commutes with existential quantification. When the rule is oriented from left to right, its side-condition may always be satisfied by suitable α -conversion. C-LETLET states that let forms may commute, provided they bind distinct

$C_1 \wedge C_2$	\equiv	$C_2 \wedge C_1$		(C-AND)
$(C_1 \wedge C_2) \wedge C_3$	\equiv	$C_1 \wedge (C_2 \wedge C_3)$		(C-ANDAND)
$C_1 \wedge C_2$	\equiv	C_1	if $C_1 \Vdash C_2$	(C-DUP)
$\exists \bar{X}. \exists \bar{Y}. C$	\equiv	$\exists \bar{X} \bar{Y}. C$		(C-EXEX)
$\exists \bar{X}. C$	\equiv	C	if $\bar{X} \# \text{ftv}(C)$	(C-EX*)
$(\exists \bar{X}. C_1) \wedge C_2$	\equiv	$\exists \bar{X}. (C_1 \wedge C_2)$	if $\bar{X} \# \text{ftv}(C_2)$	(C-EXAND)
$\exists Z. (\sigma \preceq Z \wedge Z \preceq T)$	\equiv	$\sigma \preceq T$	if $Z \notin \text{ftv}(\sigma, T)$	(C-EXTRANS)
$\text{let } x : \sigma \text{ in } C[x \preceq T]$	\equiv	$\text{let } x : \sigma \text{ in } C[\sigma \preceq T]$		(C-INID)
			if $x \notin \text{dpi}(C)$ and $\text{dtv}(C) \# \text{ftv}(\sigma)$ and $\{x\} \cup \text{dpi}(C) \# \text{fpi}(\sigma)$	
$\text{let } \Gamma \text{ in } C$	\equiv	$\exists \Gamma \wedge C$	if $\text{dpi}(\Gamma) \# \text{fpi}(C)$	(C-IN*)
$\text{let } \Gamma \text{ in } (C_1 \wedge C_2)$	\equiv	$(\text{let } \Gamma \text{ in } C_1) \wedge (\text{let } \Gamma \text{ in } C_2)$		(C-INAND)
$\text{let } \Gamma \text{ in } (C_1 \wedge C_2)$	\equiv	$(\text{let } \Gamma \text{ in } C_1) \wedge C_2$	if $\text{dpi}(\Gamma) \# \text{fpi}(C_2)$	(C-INAND*)
$\text{let } \Gamma \text{ in } \exists \bar{X}. C$	\equiv	$\exists \bar{X}. \text{let } \Gamma \text{ in } C$	if $\bar{X} \# \text{ftv}(\Gamma)$	(C-INEX)
$\text{let } \Gamma_1; \Gamma_2 \text{ in } C$	\equiv	$\text{let } \Gamma_2; \Gamma_1 \text{ in } C$		(C-LETLET)
			if $\text{dpi}(\Gamma_1) \# \text{dpi}(\Gamma_2)$ and $\text{dpi}(\Gamma_2) \# \text{fpi}(\Gamma_1)$ and $\text{dpi}(\Gamma_1) \# \text{fpi}(\Gamma_2)$	
$\text{let } x : \forall \bar{X}[C_1 \wedge C_2]. T \text{ in } C_3$	\equiv	$C_1 \wedge \text{let } x : \forall \bar{X}[C_2]. T \text{ in } C_3$	if $\bar{X} \# \text{ftv}(C_1)$	(C-LETAND)
$\text{let } \Gamma; x : \forall \bar{X}[C_1]. T \text{ in } C_2$	\equiv	$\text{let } \Gamma; x : \forall \bar{X}[\text{let } \Gamma \text{ in } C_1]. T \text{ in } C_2$		(C-LETDUP)
			if $\bar{X} \# \text{ftv}(\Gamma)$ and $\text{dpi}(\Gamma) \# \text{fpi}(\Gamma)$	
$\text{let } x : \forall \bar{X}[\exists \bar{Y}. C_1]. T \text{ in } C_2$	\equiv	$\text{let } x : \forall \bar{X} \bar{Y}[C_1]. T \text{ in } C_2$	if $\bar{Y} \# \text{ftv}(T)$	(C-LETEX)
$\text{let } x : \forall \bar{X} \bar{Y}[C_1]. T \text{ in } C_2$	\equiv	$\exists \bar{Y}. \text{let } x : \forall \bar{X}[C_1]. T \text{ in } C_2$		(C-LETALL)
			if $\bar{Y} \# \text{ftv}(C_2)$ and $\exists \bar{X}. C_1$ determines \bar{Y}	
$\exists X. (T \preceq X \wedge \text{let } x : X \text{ in } C)$	\equiv	$\text{let } x : T \text{ in } C$	if $X \notin \text{ftv}(T, C)$	(C-LETSUB)
$\vec{X} = \vec{T} \wedge [\vec{X} \mapsto \vec{T}]C$	\equiv	$\vec{X} = \vec{T} \wedge C$		(C-EQ)
true	\equiv	$\exists \bar{X}. (\vec{X} = \vec{T})$	if $\bar{X} \# \text{ftv}(\vec{T})$	(C-NAME)
$[\vec{X} \mapsto \vec{T}]C$	\equiv	$\exists \bar{X}. (\vec{X} = \vec{T} \wedge C)$	if $\bar{X} \# \text{ftv}(\vec{T})$	(C-NAMEEQ)

Figure 10-6: Constraint equivalence laws

program identifiers and provided no free program identifiers are captured in the process. C-LETAND allows the conjunct C_1 to be moved outside of the constrained type scheme $\forall\bar{X}[C_1 \wedge C_2].T$, provided it does not involve any of the universally quantified type variables \bar{X} . When oriented from left to right, the rule yields an important simplification law: indeed, taking an instance of $\forall\bar{X}[C_2].T$ is less expensive than taking an instance of $\forall\bar{X}[C_1 \wedge C_2].T$, since the latter involves creating a copy of C_1 , while the former does not. C-LETDUP allows pushing a series of let bindings into a constrained type scheme, provided no capture occurs in the process. It is not used as a simplification law but as a tool in some proofs. C-LETTEX states that it does not make any difference for a set of type variables \bar{Y} to be existentially quantified inside a constrained type scheme or part of the type scheme's universal quantifiers. Indeed, in either case, taking an instance of the type scheme means producing a constraint where \bar{Y} is existentially quantified. C-LETALL states that it is equivalent for a set of type variables \bar{Y} to be part of a type scheme's universal quantifiers or existentially bound outside the let form, provided these type variables are determined. In other words, when a type variable is sufficiently constrained, it does not matter whether it is polymorphic or monomorphic. Together, C-LETTEX and C-LETALL allow, in some situations, hoisting existential quantifiers out of the left-hand side of a let form.

10.2.16 EXAMPLE: C-LETALL would be invalid without the condition that $\exists\bar{X}.C_1$ determines \bar{Y} . Consider, for instance, the constraint $\text{let } x : \forall Y.Y \rightarrow Y \text{ in } (x \leq \text{int} \rightarrow \text{int} \wedge x \leq \text{bool} \rightarrow \text{bool})$ (1), where int and bool are incompatible nullary type constructors. By C-INID and C-IN*, it is equivalent to $\forall Y.Y \rightarrow Y \leq \text{int} \rightarrow \text{int} \wedge \forall Y.Y \rightarrow Y \leq \text{bool} \rightarrow \text{bool}$ which, by Definition 10.2.3, means $\exists Y.(Y \rightarrow Y \leq \text{int} \rightarrow \text{int}) \wedge \exists Y.(Y \rightarrow Y \leq \text{bool} \rightarrow \text{bool})$, that is, true. Now, if C-LETALL was valid without its side-condition, then (1) would also be equivalent to $\exists Y.\text{let } x : Y \rightarrow Y \text{ in } (x \leq \text{int} \rightarrow \text{int} \wedge x \leq \text{bool} \rightarrow \text{bool})$, which by C-INID and C-IN* is $\exists Y.(Y \rightarrow Y \leq \text{int} \rightarrow \text{int} \wedge Y \rightarrow Y \leq \text{bool} \rightarrow \text{bool})$. By C-ARROW and C-EXTRANS, this is $\text{int} = \text{bool}$, that is, false. Thus, the law is invalid in this case. It is easy to see why: when the type scheme σ contains a $\forall Y$ quantifier, every instance of σ receives its own $\exists Y$ quantifier, making Y a distinct (local) type variable; but when Y is not universally quantified, all instances of σ share references to a single (global) type variable Y . This corresponds to the intuition that, in the former case, σ is polymorphic in Y , while in the latter case, it is monomorphic in Y . It is possible to prove that, when deprived of its side-condition, C-LETALL is only an entailment law, that is, its right-hand side entails its left-hand side. Similarly, it is in general invalid to hoist an existential quantifier out of the left-hand side of a let form. To see this, one may study the (equivalent) constraint $\text{let } x : \forall X[\exists Y.X = Y \rightarrow Y].X \text{ in } (x \leq \text{int} \rightarrow \text{int} \wedge x \leq \text{bool} \rightarrow \text{bool})$. Naturally, in the above examples, the side-condition “true determines Y ” does

not hold: by Definition 10.2.14, it is equivalent to “two ground assignments that coincide outside Y must coincide on Y as well,” which is false when \mathcal{M}_* contains two distinct elements, such as `int` and `bool` here.

There are cases, however, where the side-condition does hold. For instance, we later prove that $\exists X.Y = \text{int}$ determines Y ; see Lemma 10.6.7. As a result, C-LETALL states that $\text{let } x : \forall XY[Y = \text{int}].Y \rightarrow X$ in C **(1)** is equivalent to $\exists Y.\text{let } x : \forall X[Y = \text{int}].Y \rightarrow X$ in C **(2)**, provided $Y \notin \text{ftv}(C)$. The intuition is simple: because Y is forced to assume the value `int` by the equation $Y = \text{int}$, it makes no difference whether Y is or isn’t universally quantified. By C-LETAND, (2) is equivalent to $\exists Y.(Y = \text{int} \wedge \text{let } x : \forall X.Y \rightarrow X$ in $C)$ **(3)**. In an efficient constraint solver, simplifying (1) into (3) before using C-INID to eliminate the `let` form is worthwhile, since doing so obviates the need for copying the type variable Y and the equation $Y = \text{int}$ at every free occurrence of x inside C . \square

C-LETSUB is the analog of an environment strengthening lemma: roughly speaking, it states that, if a constraint holds under the assumption that x has type X , where X is some supertype of T , then it also holds under the assumption that x has type T . The last three rules deal with the equality predicate. C-EQ states that it is valid to replace equals with equals; note the absence of a side-condition. When oriented from left to right, C-NAME allows introducing fresh names \vec{X} for the types \vec{T} . As always, \vec{X} stands for a vector of distinct type variables; \vec{T} stands for a vector of the same length of types of appropriate kind. Of course, this makes sense only if the definition is not circular, that is, if the type variables \vec{X} do not occur free within the terms \vec{T} . When oriented from right to left, C-NAME may be viewed as a simplification law: it allows eliminating type variables whose value has been determined. C-NAMEEQ is a combination of C-EQ and C-NAME. It shows that applying an idempotent substitution to a constraint C amounts to placing C within a certain context.

So far, we have considered `def` a primitive constraint form and defined the `let` form in terms of `def`, conjunction, and existential quantification. The motivation for this approach was to simplify the (omitted) proofs of several constraint equivalence laws. However, in the remainder of this chapter, we work with `let` forms exclusively and never employ the `def` construct. This offers us an extra property: every constraint that contains a false subconstraint must be false.

10.2.17 LEMMA: $C[\text{false}] \equiv \text{false}$. \square

Reasoning with Constraints in an Equality-Only Syntactic Model

We have given a number of equivalence laws that are valid with respect to any interpretation of constraints, that is, within any model. However, an im-

portant special case is that of *equality-only syntactic models*. Indeed, in that specific setting, our constraint-based type systems are in close correspondence with DM. In brief, we aim to prove that every satisfiable constraint C such that $fpi(C) = \emptyset$ admits a *canonical solved form* and to show that this notion corresponds to the standard concept of a *most general unifier*. These results are exploited when we relate $HM(X)$ with Damas and Milner's system (p. 428).

Thus, let us now assume that constraints are interpreted in an equality-only syntactic model. Let us further assume that, for every kind κ , (i) there are at least *two* type constructors of image kind κ and (ii) for every type constructor F of image kind κ , there exists $t \in \mathcal{M}_\kappa$ such that $t(\epsilon) = F$. We refer to models that violate (i) or (ii) as *degenerate*; one may argue that such models are of little interest. The assumption that the model is nondegenerate is used in the proof of Theorem 10.3.7. Last, throughout the present subsection we manipulate only constraints that have *no free program identifiers*.

A *solved form* is a conjunction of equations, where the left-hand sides are distinct type variables that do not appear in the right-hand sides, possibly surrounded by a number of existential quantifiers. Our definition is identical to Lassez, Maher, and Marriott's solved forms (1988) and to Jouannaud and Kirchner's tree solved forms (1991), except we allow for prenex existential quantifiers, which are made necessary by our richer constraint language. Jouannaud and Kirchner also define *dag* solved forms, which may be exponentially smaller. Because we define solved forms only for proof purposes, we need not take performance into account at this point. The efficient constraint solver presented in §10.6 does manipulate graphs, rather than trees. Type scheme introduction and instantiation constructs cannot appear within solved forms; indeed, provided the constraint at hand has no free program identifiers, they can be expanded away. For this reason, their presence in the constraint language has no impact on the results contained in this section.

10.2.18 DEFINITION: A *solved form* is of the form $\exists \vec{Y}.(\vec{X} = \vec{T})$, where $\vec{X} \# ftv(\vec{T})$. \square

Solved forms offer a convenient way of reasoning about constraints because every satisfiable constraint is equivalent to one. This property is established by the following lemma.

10.2.19 LEMMA: Every constraint is equivalent to either a solved form or *false*. \square

It is possible to impose further restrictions on solved forms. A solved form $\exists \vec{Y}.(\vec{X} = \vec{T})$ is *canonical* if and only if its free type variables are exactly \vec{X} . This is stated, in an equivalent way, by the following definition.

10.2.20 DEFINITION: A *canonical solved form* is a constraint of the form $\exists \vec{Y}.(\vec{X} = \vec{T})$, where $ftv(\vec{T}) \subseteq \vec{Y}$ and $\vec{X} \# \vec{Y}$. \square

- 10.2.21 LEMMA: Every solved form is equivalent to a canonical solved form. \square

It is easy to describe the solutions of a canonical solved form: they are the ground refinements of the substitution $[\vec{X} \mapsto \vec{T}]$. Hence, every canonical solved form is satisfiable.

The following definition allows entertaining a dual view of canonical solved forms, either as constraints or as idempotent type substitutions. The latter view is commonly found in standard treatments of unification (Lassez, Maher, and Marriott, 1988; Jouannaud and Kirchner, 1991) and in classic presentations of ML-the-type-system.

- 10.2.22 DEFINITION: If $[\vec{X} \mapsto \vec{T}]$ is an idempotent substitution of domain \bar{X} , let $\exists[\vec{X} \mapsto \vec{T}]$ denote the canonical solved form $\exists \vec{Y}.(\vec{X} = \vec{T})$, where $\vec{Y} = \text{ftv}(\vec{T})$. An idempotent substitution θ is a *most general unifier* of the constraint C if and only if $\exists \theta$ and C are equivalent. \square

By definition, equivalent constraints admit the same most general unifiers. Many properties of canonical solved forms may be reformulated in terms of most general unifiers. By Lemmas 10.2.19 and 10.2.21, every satisfiable constraint admits a most general unifier.

10.3 HM(X)

Constraint-based type systems appeared during the 1980s (Mitchell, 1984; Fuh and Mishra, 1988) and were widely studied during the following decade (Curtis, 1990; Aiken and Wimmers, 1993; Jones, 1994; Smith, 1994; Palsberg, 1995; Trifonov and Smith, 1996; Fähndrich, 1999; Pottier, 2001b). We now present one such system, baptized HM(X) because it is a *parameterized* extension of Hindley and Milner's type discipline; the meaning of the parameter X was explained on page 411. Its original description is due to Odersky, Sulzmann, and Wehr (1999). Since then, it has been completed in a number of works including Müller (1998), Sulzmann, Müller, and Zenger (1999), Sulzmann (2000), Pottier (2001a), and Skalka and Pottier (2002). Each of these presentations introduces minor variations. Here, we follow Pottier (2001a), which is itself inspired by Sulzmann, Müller, and Zenger (1999).

Definition

Our presentation of HM(X) relies on the constraint language introduced in §10.2. Technically, our approach to constraints is less abstract than that of Odersky, Sulzmann, and Wehr (1999). We interpret constraints within a model, give conjunction and existential quantification their standard mean-

$\frac{\Gamma(x) = \sigma \quad C \Vdash \exists \sigma}{C, \Gamma \vdash x : \sigma} \quad (\text{HMX-VAR})$	$\frac{C \wedge D, \Gamma \vdash t : T \quad \tilde{X} \# \text{ftv}(C, \Gamma)}{C \wedge \exists \tilde{X}. D, \Gamma \vdash t : \forall \tilde{X}[D]. T} \quad (\text{HMX-GEN})$
$\frac{C, (\Gamma; z : T) \vdash t : T'}{C, \Gamma \vdash \lambda z. t : T \rightarrow T'} \quad (\text{HMX-ABS})$	$\frac{C, \Gamma \vdash t : \forall \tilde{X}[D]. T}{C \wedge D, \Gamma \vdash t : T} \quad (\text{HMX-INST})$
$\frac{C, \Gamma \vdash t_1 : T \rightarrow T' \quad C, \Gamma \vdash t_2 : T}{C, \Gamma \vdash t_1 t_2 : T'} \quad (\text{HMX-APP})$	$\frac{C, \Gamma \vdash t : T \quad C \Vdash T \leq T'}{C, \Gamma \vdash t : T'} \quad (\text{HMX-SUB})$
$\frac{C, \Gamma \vdash t_1 : \sigma \quad C, (\Gamma; z : \sigma) \vdash t_2 : T}{C, \Gamma \vdash \text{let } z = t_1 \text{ in } t_2 : T} \quad (\text{HMX-LET})$	$\frac{C, \Gamma \vdash t : \sigma \quad \tilde{X} \# \text{ftv}(\Gamma, \sigma)}{\exists \tilde{X}. C, \Gamma \vdash t : \sigma} \quad (\text{HMX-EXISTS})$

Figure 10-7: Typing rules for HM(X)

ing, and derive a number of equivalence laws (§10.2). Odersky et al., on the other hand, do not explicitly rely on a logical interpretation; instead, they axiomatize constraint equivalence, that is, they consider a number of equivalence laws as axioms. Thus, they ensure that their high-level proofs, such as type soundness and correctness and completeness of type inference, are independent of the low-level details of the logical interpretation of constraints. Their approach is also more general, since it allows dealing with other logical interpretations, such as “open-world” interpretations, where constraints are interpreted not within a fixed model, but within a family of extensions of a “current” model. In this chapter, we have avoided this extra layer of abstraction and given fixed meaning to constraints, making things somewhat simpler. However, the changes required to adopt Odersky et al.’s approach would not be extensive, since the forthcoming proofs do indeed rely mostly on constraint equivalence laws, rather than on low-level details of the logical interpretation of constraints.

Another slight departure from Odersky et al.’s work lies in the fact that we have enriched the constraint language with type scheme introduction and instantiation forms, which were absent in the original presentation of HM(X). To prevent this addition from affecting HM(X), we require the constraints that appear in HM(X) typing judgments to have no free program identifiers. Note that this does not prevent them from containing let forms.

The type system HM(X) consists of a four-place judgment whose parameters are a constraint C , an environment Γ , an expression t , and a type scheme σ . A judgment is written $C, \Gamma \vdash t : \sigma$ and is read: *under the assumptions C and Γ , the expression t has type σ* . One may view C as an assumption about

the judgment's free type variables and Γ as an assumption about τ 's free program identifiers. Recall that Γ now contains constrained type schemes, and that σ is a constrained type scheme.

We would like the validity of a typing judgment to depend not on the syntax, but only on the meaning of its constraint assumption. We enforce this point of view by considering judgments equal modulo equivalence of their constraint assumptions. In other words, the typing judgments $C, \Gamma \vdash \tau : \sigma$ and $D, \Gamma \vdash \tau : \sigma$ are considered identical when $C \equiv D$ holds. A judgment is *valid*, or *holds*, if and only if it is derivable via the rules given in Figure 10-7. Note that a valid judgment may involve an arbitrary constraint. A (closed) program τ is *well-typed* within the (closed) environment Γ if and only if a judgment of the form $C, \Gamma \vdash \tau : \sigma$ holds for some satisfiable constraint C . One might wonder why we do not make the apparently stronger requirement that $C \wedge \exists \sigma$ be satisfiable; however, by inspection of the typing rules, the reader may check that, if the above judgment is derivable, then $C \Vdash \exists \sigma$ holds, hence the two requirements are equivalent.

Let us now explain the rules. Like DM-VAR, HMX-VAR looks up the environment to determine the type scheme associated with the program identifier x . Its second premise plays a minor technical role: as noted in the previous paragraph, its presence helps simplify the definition of well-typedness. HMX-ABS, HMX-APP, and HMX-LET are identical to DM-ABS, DM-APP, and DM-LET, respectively, except that the assumption C is made available to every subderivation. We recall that the type T may be viewed as the type scheme $\forall \emptyset[\text{true}].T$ (Definitions 10.1.18 and 10.2.2). As a result, types form a subset of type schemes, which implies that $\Gamma; z : T$ is a well-formed environment and $C, \Gamma \vdash \tau : T$ a well-formed typing judgment. To understand HMX-GEN, it is best to first consider the particular case where C is true. This yields the following, simpler rule:

$$\frac{D, \Gamma \vdash \tau : T \quad \bar{X} \# \text{fv}(\Gamma)}{\exists \bar{X}. D, \Gamma \vdash \tau : \forall \bar{X}[D]. T} \quad (\text{HMX-GEN}')$$

The second premise is identical to that of DM-GEN: the type variables that are generalized must not occur free within the environment. The conclusion forms the type scheme $\forall \bar{X}[D].T$, where the type variables \bar{X} have become universally quantified, but are still subject to the constraint D . Note that the type variables that occur free in D may include not only \bar{X} , but also other type variables, typically free in Γ . HMX-GEN may be viewed as a more liberal version of HMX-GEN', whereby part of the current constraint, namely C , need not be copied if it does not concern the type variables that are being generalized, namely \bar{X} . This optimization is important in practice, because C may be very large. An intuitive explanation for its correctness is given by the con-

straint equivalence law C-LETAND, which expresses the same optimization in terms of let constraints. Because $HM(X)$ does not use let constraints, the optimization is hard-wired into the typing rule. As a last technical remark, let us point out that replacing $C \wedge \exists \bar{X}.D$ with $C \wedge D$ in HMX-GEN's conclusion would not affect the set of derivable judgments; this fact may be established using HMX-EXISTS and Lemma 10.3.1. HMX-INST allows taking an instance of a type scheme. The reader may be surprised to find that, contrary to DM-INST, it does not involve a type substitution. Instead, the rule merely drops the universal quantifier, which amounts to applying the identity substitution $\bar{X} \mapsto \bar{X}$. One should recall, however, that type schemes are considered equal modulo α -conversion, so it is possible to rename the type scheme's universal quantifiers prior to using HMX-INST. The reason why this provides sufficient expressive power appears in Exercise 10.3.2 below. The constraint D carried by the type scheme is recorded as part of the current constraint in HMX-INST's conclusion. The *subsumption* rule HMX-SUB allows a type T to be replaced at any time with an arbitrary supertype T' . Because both T and T' may have free type variables, whether $T \leq T'$ holds depends on the current assumption C , which is why the rule's second premise is an entailment assertion. An operational explanation of HMX-SUB is that it requires all uses of subsumption to be explicitly recorded in the current constraint. Note that HMX-SUB remains a useful and necessary rule even when subtyping is interpreted as equality: then, it allows exploiting the type equations found in C . Last, HMX-EXISTS allows the type variables that occur only within the current constraint to become existentially quantified. As a result, these type variables no longer occur free in the rule's conclusion; in other words, they have become local to the subderivation rooted at the premise. One may prove that the presence of HMX-EXISTS in the type system does not augment the set of well-typed programs, but does augment the set of valid typing judgments; it is a pleasant technical convenience. Indeed, because judgments are considered equal modulo constraint equivalence, constraints may be transparently simplified at any time. (By simplifying a constraint, we mean replacing it with an equivalent constraint whose syntactic representation is considered simpler.) Bearing this fact in mind, one finds that an effect of rule HMX-EXISTS is to enable more simplifications: because constraint equivalence is a congruence, $C \equiv D$ implies $\exists \bar{X}.C \equiv \exists \bar{X}.D$, but the converse does not hold in general. For instance, there is in general no way of simplifying the judgment $X \leq Y \leq Z, \Gamma \vdash t : \sigma$, but if it is known that Y does not appear free in Γ or σ , then HMX-EXISTS allows deriving $\exists Y.(X \leq Y \leq Z), \Gamma \vdash t : \sigma$, which is the same judgment as $X \leq Z, \Gamma \vdash t : \sigma$. Thus, an interesting simplification has been enabled. Note that $X \leq Y \leq Z \equiv X \leq Z$ does not hold, while, according to C-EXTRANS, $\exists Y.(X \leq Y \leq Z) \equiv X \leq Z$ does.

A pleasant property of $\text{HM}(X)$ is that strengthening a judgment's constraint assumption (that is, weakening the judgment itself) preserves its validity. It is worth noting that in traditional presentations, which rely more heavily on type substitutions, the analog of this result is a *type substitution* lemma; see for instance Tofte (1988), Lemma 2.7; Rémy (1992a), Lemma 1; Leroy (1992), Proposition 1.2; and Skalka and Pottier (2002), Lemma 3.4. Here, the lemma further states that weakening a judgment does not alter the shape of its derivation, a useful property when reasoning by induction on type derivations.

10.3.1 LEMMA [WEAKENING]: If $C' \Vdash C$, then every derivation of $C, \Gamma \vdash \tau : \sigma$ may be turned into a derivation of $C', \Gamma \vdash \tau : \sigma$ with the same shape. \square

10.3.2 EXERCISE [RECOMMENDED, ★★]: In some presentations of $\text{HM}(X)$, HMX-INST is replaced with the following variant:

$$\frac{C, \Gamma \vdash \tau : \forall \vec{X}[D].T \quad C \Vdash [\vec{X} \mapsto \vec{T}]D}{C, \Gamma \vdash \tau : [\vec{X} \mapsto \vec{T}]T} \quad (\text{HMX-INST}')$$

Show that $\text{HMX-INST}'$ is admissible in our presentation of $\text{HM}(X)$ —that is, if its premise is derivable according to the rules of Figure 10-7, then so is its conclusion. Thus, the choice between HMX-INST and $\text{HMX-INST}'$ is only stylistic: it makes no difference in the system's expressive power. Because HMX-INST is more elementary, choosing it simplifies some proofs. \square

10.3.3 EXERCISE [★]: Give a derivation of $\text{true}, \emptyset \vdash \lambda z.z : \text{int} \rightarrow \text{int}$. Give a derivation of $\text{true}, \emptyset \vdash \lambda z.z : \forall X.X \rightarrow X$. Check that the former judgment also follows from the latter via $\text{HMX-INST}'$ (Exercise 10.3.2), and determine which derivation of $\text{true}, \emptyset \vdash \lambda z.z : \text{int} \rightarrow \text{int}$ this path gives rise to. \square

We do not give a direct type soundness proof for $\text{HM}(X)$. Instead, in the forthcoming sections, we prove that well-typedness in $\text{HM}(X)$ is equivalent to the satisfiability of a certain constraint and use that characterization as a basis for our type soundness proof. A direct type soundness result, based on a denotational semantics, may be found in Odersky, Sulzmann, and Wehr (1999). Another type soundness proof, which follows Wright and Felleisen's syntactic approach (1994), appears in Skalka and Pottier (2002). Last, a hybrid approach, which combines some of the advantages of the previous two, is given in Pottier (2001a).

An Alternate Presentation of $\text{HM}(X)$

The presentation of $\text{HM}(X)$ given in Figure 10-7 has only four syntax-directed rules out of eight. It is a good specification of the type system, but it is far

$\frac{\Gamma(x) = \forall \tilde{X}[D].T}{C \wedge D, \Gamma \vdash x : T} \quad (\text{HMD-VARINST})$	$\frac{C \wedge D, \Gamma \vdash t_1 : T_1 \quad \tilde{X} \# \text{ftv}(C, \Gamma)}{C \wedge \exists \tilde{X}.D, (\Gamma; z : \forall \tilde{X}[D].T_1) \vdash t_2 : T_2} \quad (\text{HMD-LETGEN})$
$\frac{C, (\Gamma; z : T) \vdash t : T'}{C, \Gamma \vdash \lambda z. t : T \rightarrow T'} \quad (\text{HMD-ABS})$	$\frac{C, \Gamma \vdash t : T \quad C \Vdash T \leq T'}{C, \Gamma \vdash t : T'} \quad (\text{HMD-SUB})$
$\frac{C, \Gamma \vdash t_1 : T \rightarrow T' \quad C, \Gamma \vdash t_2 : T}{C, \Gamma \vdash t_1 t_2 : T'} \quad (\text{HMD-APP})$	$\frac{C, \Gamma \vdash t : T \quad \tilde{X} \# \text{ftv}(\Gamma, T)}{\exists \tilde{X}. C, \Gamma \vdash t : T} \quad (\text{HMD-EXISTS})$

Figure 10-8: An alternate presentation of $HM(X)$

from an algorithmic description. As a first step towards such a description, we provide an alternate presentation of $HM(X)$, where generalization is performed only at `let` expressions and instantiation takes place only at references to program identifiers (Figure 10-8). This presentation only has two non-syntax-directed rules, making it sometimes easier to reason about. It has the property that all judgments are of the form $C, \Gamma \vdash t : T$, rather than $C, \Gamma \vdash t : \sigma$. The following theorem states that the two presentations are indeed equivalent.

10.3.4 **THEOREM:** $C, \Gamma \vdash t : T$ is derivable via the rules of Figure 10-8 if and only if it is a valid $HM(X)$ judgment. \square

This theorem shows that the rule sets of Figures 10-7 and 10-8 derive the same monomorphic judgments, that is, the same judgments of the form $C, \Gamma \vdash t : T$. The fact that judgments of the form $C, \Gamma \vdash t : \sigma$, where σ is not a monotype, cannot be derived using the new rule set is a technical simplification, without deep significance.

10.3.5 **EXERCISE [***, +]:** Show that it is possible to simplify the presentation of Damas and Milner's type system in an analogous manner. That is, define an alternate set of typing rules for DM, which allows deriving judgments of the form $\Gamma \vdash t : T$; then, show that this new rule set is equivalent to the previous one, in the same sense as above. Which auxiliary properties of DM does your proof require? A solution is given by Clement, Despeyroux, Despeyroux, and Kahn (1986). \square

Relating $HM(X)$ with Damas and Milner's Type System

In order to explain our interest in $HM(X)$, we wish to show that it is more general than Damas and Milner's type system. Since $HM(X)$ really is a family of type systems, we must make this statement more precise. First, every member of the $HM(X)$ family contains DM. Conversely, DM contains $HM(=)$, the constraint-based type system obtained by specializing $HM(X)$ to the setting of an equality-only syntactic model.

The first of these assertions is easy to prove because the mapping from DM judgments to $HM(X)$ judgments is essentially the identity: every valid DM judgment may be viewed as a valid $HM(X)$ judgment under the trivial assumption $true$. This statement relies on the fact that the DM type scheme $\forall \bar{X}.T$ is identified with the constrained type scheme $\forall \bar{X}[true].T$, so DM type schemes (respectively environments) form a subset of $HM(X)$ type schemes (respectively environments). Its proof is easy and relies on Exercise 10.3.2.

10.3.6 THEOREM: If $\Gamma \vdash \tau : S$ holds in DM, then $true, \Gamma \vdash \tau : S$ holds in $HM(X)$. \square

We are now interested in proving that $HM(=)$, as defined above, is contained within DM. To this end, we must translate every $HM(=)$ judgment to a DM judgment. It turns out that this is possible if the original judgment's constraint assumption is satisfiable. The translation relies on the fact that the definition of $HM(=)$ assumes an equality-only syntactic model. Indeed, in that setting, every satisfiable constraint admits a most general unifier (Definition 10.2.22), whose properties we make essential use of.

Unfortunately, due to lack of space, we cannot give the details of this translation, which are fairly involved. Let us merely say that, given a type scheme σ and an idempotent type substitution θ such that $ftv(\sigma) \subseteq dom(\theta)$ and $\exists \theta \Vdash \exists \sigma$ hold, the translation of σ under θ is a DM type scheme, written $\llbracket \sigma \rrbracket_\theta$. Its meaning is intended to be the same as that of the $HM(X)$ type scheme $\theta(\sigma)$. For instance, under the identity substitution, the translation of the $HM(X)$ type scheme $\forall XY[X = Y \rightarrow Y].X$ is the DM type scheme $\forall Z.Z \rightarrow Z$. The translation is extended to environments in such a way that $\llbracket \Gamma \rrbracket_\theta$ is defined when $ftv(\Gamma) \subseteq dom(\theta)$ holds. We are now ready to state the main theorem.

10.3.7 THEOREM: Let $C, \Gamma \vdash \tau : \sigma$ hold in $HM(=)$. Let θ be a most general unifier of C such that $ftv(\Gamma, \sigma) \subseteq dom(\theta)$. Then, $\llbracket \Gamma \rrbracket_\theta \vdash \tau : \llbracket \sigma \rrbracket_\theta$ holds in DM. \square

Note that, by requiring θ to be a most general unifier of C , we also require C to be satisfiable. Judgments that carry an unsatisfiable constraint cannot be translated.

Together, Theorems 10.3.6 and 10.3.7 yield a precise correspondence between DM and $HM(=)$: there exists a compositional translation from each to

the other. In other words, they may be viewed as two equivalent formulations of a single type system. One might also say that $\text{HM}(=)$ is a constraint-based formulation of DM. Furthermore, Theorem 10.3.6 states that every member of the $\text{HM}(X)$ family is an extension of DM. This explains our double interest in $\text{HM}(X)$, as an alternate formulation of DM, which we believe is more pleasant for reasons already discussed, and as a more expressive framework.

10.4 Constraint Generation

We now explain how to reduce type inference problems for $\text{HM}(X)$ to constraint solving problems. A type inference problem consists of a type environment Γ , an expression t , and a type T of kind \star . The problem is to determine whether there exists a satisfiable constraint C such that $C, \Gamma \vdash t : T$ holds. A constraint solving problem consists of a constraint C . The problem is to determine whether C is satisfiable. To reduce a type inference problem (Γ, t, T) to a constraint solving problem, we must produce a constraint C that is both *sufficient* and *necessary* for $C, \Gamma \vdash t : T$ to hold. Below, we explain how to compute such a constraint, which we write $\llbracket \Gamma \vdash t : T \rrbracket$. We check that it is indeed *sufficient* by proving $\llbracket \Gamma \vdash t : T \rrbracket, \Gamma \vdash t : T$. That is, the constraint $\llbracket \Gamma \vdash t : T \rrbracket$ is specific enough to guarantee that t has type T under environment Γ . We say that constraint generation is *sound*. We check that it is indeed *necessary* by proving that, for every constraint C , the validity of $C, \Gamma \vdash t : T$ implies $C \Vdash \llbracket \Gamma \vdash t : T \rrbracket$. That is, every constraint that guarantees that t has type T under environment Γ is at least as specific as $\llbracket \Gamma \vdash t : T \rrbracket$. We say that constraint generation is *complete*. Together, these properties mean that $\llbracket \Gamma \vdash t : T \rrbracket$ is the *least specific* constraint that guarantees that t has type T under environment Γ .

We now see how to reduce a type inference problem to a constraint solving problem. Indeed, if there exists a satisfiable constraint C such that $C, \Gamma \vdash t : T$ holds, then, by the completeness property, $C \Vdash \llbracket \Gamma \vdash t : T \rrbracket$ holds, so $\llbracket \Gamma \vdash t : T \rrbracket$ is satisfiable. Conversely, by the soundness property, if $\llbracket \Gamma \vdash t : T \rrbracket$ is satisfiable, then we have a satisfiable constraint C such that $C, \Gamma \vdash t : T$ holds. In other words, t is well-typed with type T under environment Γ if and only if $\llbracket \Gamma \vdash t : T \rrbracket$ is satisfiable.

The reader may be somewhat puzzled by the fact that our formulation of the type inference problem requires an appropriate type T to be known in advance, whereas the very purpose of type inference seems to consist in *discovering* the type of t ! In other words, we have made T an *input* of the constraint generation algorithm, instead of an *output*. Fortunately, this causes no loss of generality, because it is possible to let T be a type variable X , cho-

sen fresh for Γ . Then, the constraint produced by the algorithm will contain information about X . This is the point of the following exercise.

- 10.4.1 EXERCISE [RECOMMENDED, ★]: Let $X \notin \text{ftv}(\Gamma)$. Show that, if there exist a satisfiable constraint C and a type T such that $C, \Gamma \vdash t : T$ holds, then there exists a satisfiable constraint C' such that $C', \Gamma \vdash t : X$ holds. Conclude that, given a closed environment Γ and an arbitrary type variable X , the term t is well-typed within Γ if and only if $\llbracket \Gamma \vdash t : X \rrbracket$ is satisfiable. \square

This shows that providing T as an input to the constraint generation procedure is not essential. We adopt this style because it is convenient. A somewhat naive alternative would be to provide Γ and t only, and to have the procedure return both a constraint C and a type T (Sulzmann, Müller, and Zenger, 1999). It turns out that this does not quite work, because C and T may mention “fresh” variables, which we must be able to quantify over, if we are to avoid an informal treatment of “freshness.” Thus, the true alternative is to provide Γ and t only and to have the procedure return a *type scheme* σ (Bourdoncle and Merz, 1997; Bonniot, 2002).

The existence of a sound and complete constraint generation procedure is the analog of the existence of *principal type schemes* in classic presentations of ML-the-type-system (Damas and Milner, 1982). Indeed, a principal type scheme is least specific in the sense that all valid types are substitution instances of it. Here, the constraint $\llbracket \Gamma \vdash t : T \rrbracket$ is least specific in the sense that all valid constraints entail it. More about principal types and principal typings may be found in Jim (1996) and Wells (2002).

How do we perform constraint generation? A standard approach (Sulzmann, Müller, and Zenger, 1999; Bonniot, 2002) is to define $\llbracket \Gamma \vdash t : T \rrbracket$ by induction on the structure of t . At every `let` node, following HMD-LETGEN, part of the current constraint, namely D , is turned into a type scheme, namely $\forall \bar{X}[D].T$, which is used to extend the environment. Then, at every occurrence of the program variable that was bound at this `let` node, following HMD-VARINST, this type scheme is retrieved from the environment, and a copy of D is added back to the current constraint. If such an approach is adopted, it is important to simplify the type scheme $\forall \bar{X}[D].T$ before it is stored in the environment, because it would be inefficient to copy an unsimplified constraint. In other words, in an efficient implementation of this standard approach, constraint generation and constraint simplification cannot be separated.

Type scheme introduction and elimination constraints, which we introduced in §10.2 but did not use in the specification of $\text{HM}(X)$, are intended as a means of solving this problem. By extending our vocabulary, we are able to achieve the desired separation between constraint generation, on the one hand, and constraint solving and simplification, on the other hand, without

$$\begin{aligned}
\llbracket x : T \rrbracket &= x \leq T \\
\llbracket \lambda z. t : T \rrbracket &= \exists X_1 X_2. (\text{let } z : X_1 \text{ in } \llbracket t : X_2 \rrbracket \wedge X_1 \rightarrow X_2 \leq T) \\
\llbracket t_1 \ t_2 : T \rrbracket &= \exists X_2. (\llbracket t_1 : X_2 \rightarrow T \rrbracket \wedge \llbracket t_2 : X_2 \rrbracket) \\
\llbracket \text{let } z = t_1 \text{ in } t_2 : T \rrbracket &= \text{let } z : \forall X [\llbracket t_1 : X \rrbracket]. X \text{ in } \llbracket t_2 : T \rrbracket
\end{aligned}$$

Figure 10-9: Constraint generation

compromising efficiency. Indeed, by exploiting these new constraint forms, we may define a constraint generation procedure whose time and space complexity is linear, because it no longer involves copying subconstraints back and forth between the environment and the constraint that is being generated. (It is then up to the constraint solver to perform simplification and copying, if and when necessary.) In fact, the environment is suppressed altogether: we define $\llbracket t : T \rrbracket$ by induction on the structure of t —notice the absence of the parameter Γ . Then, the constraint $\llbracket \Gamma \vdash t : T \rrbracket$ discussed above becomes syntactic sugar for $\text{let } \Gamma \text{ in } \llbracket t : T \rrbracket$. We now employ the full constraint language: the program identifiers that appear free in t may also appear free in $\llbracket t : T \rrbracket$, as part of instantiation constraints. They become bound when $\llbracket t : T \rrbracket$ is placed within the context $\text{let } \Gamma \text{ in } []$. A similar approach to constraint generation appears in Müller (1994).

The defining equations for $\llbracket t : T \rrbracket$ appear in Figure 10-9. We refer to them as the *constraint generation rules*. The definition is quite terse and certainly simpler than the declarative specification of $\text{HM}(X)$ given in Figure 10-7; yet, we prove below that the two are equivalent.

Before explaining the definition, we state the requirements that bear on the type variables X_1 , X_2 , and X , which appear bound in the right-hand sides of the second, third, and fourth equations. These type variables must have kind \star . They must be chosen distinct (that is, $X_1 \neq X_2$ in the second equation) and fresh for the objects that appear on the left-hand side—that is, the type variables that appear bound in an equation’s right-hand side must not occur free in the term and type that appear in the equation’s left-hand side. Provided this restriction is obeyed, different choices of X_1 , X_2 , and X lead to α -equivalent constraints—that is, to the same constraint, since we identify objects up to α -conversion—which guarantees that the above equations make sense. Since expressions do not have free type variables, the freshness requirement may be simplified to: type variables that appear bound in an equation’s right-hand side must not appear free in T . However, this simplification would be rendered invalid by the introduction of open type annotations

within expressions. Note that we are able to state a precise (as opposed to informal) freshness requirement. This is made possible by the fact that $\llbracket \tau : T \rrbracket$ has no free type variables other than those of T , which in turn depends on our explicit use of existential quantification to limit the scope of auxiliary variables.

Let us now review the four equations. The first equation may be read: x has type T if and only if T is an instance of the type scheme associated with x . Note that we no longer consult the type scheme associated with x in the environment—indeed, there is no environment. Instead, we merely generate an instantiation constraint, where x appears free. (For this reason, every program identifier that occurs free within τ typically also occurs free within $\llbracket \tau : T \rrbracket$.) This constraint acquires its full meaning when it is later placed within a context of the form $\text{let } x : \sigma \text{ in } []$. This equation roughly corresponds to HMD-VARINST. The second equation may be read: $\lambda z. \tau$ has type T if and only if, for some X_1 and X_2 , (i) under the assumption that z has type X_1 , τ has type X_2 , and (ii) T is a supertype of $X_1 \rightarrow X_2$. Here, the types associated with z and τ must be fresh type variables, namely X_1 and X_2 , because we cannot in general guess them. These type variables are *bound* so as to guarantee that the generated constraint is unique up to α -conversion. They are *existentially bound* because we intend the constraint solver to discover their value. Condition (i) is expressed by the subconstraint $\text{let } z : X_1 \text{ in } \llbracket \tau : X_2 \rrbracket$. This makes sense as follows. The constraint $\llbracket \tau : X_2 \rrbracket$ typically contains a number of instantiation constraints bearing on z , of the form $z \leq T_i$. By wrapping it within the context $\text{let } z : X_1 \text{ in } []$, we effectively require every T_i to be a supertype of X_1 . Note that z does not occur free in the constraint $\text{let } z : X_1 \text{ in } \llbracket \tau : X_2 \rrbracket$, which is necessary for well-formedness of the definition, since it does not occur free in $\lambda z. \tau$. This equation roughly corresponds to HMD-EXISTS, HMD-ABS, and HMD-SUB. The third equation may be read: $\tau_1 \tau_2$ has type T if and only if, for some X_2 , τ_1 has type $X_2 \rightarrow T$ and τ_2 has type X_2 . Here, the fresh type variable X_2 stands for the unknown type of τ_2 . This equation roughly corresponds to HMD-APP. The last equation, which roughly corresponds to HMD-LETGEN, may be read: $\text{let } z = \tau_1 \text{ in } \tau_2$ has type T if and only if, under the assumption that z has every type X such that $\llbracket \tau_1 : X \rrbracket$ holds, τ_2 has type T . As in the case of λ -abstractions, the instantiation constraints bearing on z that appear within $\llbracket \tau_2 : T \rrbracket$ are given a meaning via a *let* prefix. The difference is that z may now be assigned a type scheme, as opposed to a monotype. An appropriate type scheme is built as follows. The constraint $\llbracket \tau_1 : X \rrbracket$ is the least specific constraint that must be imposed on the fresh type variable X so as to make it a valid type for τ_1 . In other words, τ_1 has every type X such that $\llbracket \tau_1 : X \rrbracket$ holds, and none other. That is, the type scheme $\forall X[\llbracket \tau_1 : X \rrbracket].X$, abbreviated σ in the following, is a *principal* type scheme for τ_1 . It is inter-

esting to note that there is no question of which type variables to generalize. Indeed, by construction, no type variables other than X may appear free in $\llbracket \tau_1 : X \rrbracket$, so we cannot generalize more variables. On the other hand, it is valid to generalize X , since it does not appear free anywhere else. This interesting simplification is inspired by Sulzmann, Müller, and Zenger (1999), where a similar technique is used. Now, what happens when $\llbracket \tau_2 : T \rrbracket$ is placed inside the context $\text{let } z : \sigma \text{ in } []$? When placed inside this context, an instantiation constraint of the form $z \leq T'$ acquires the meaning $\sigma \leq T'$, which by definition of σ and by Lemma 10.4.6 (see below) is equivalent to $\llbracket \tau_1 : T' \rrbracket$. Thus, the constraint produced by the fourth equation simulates a textual expansion of the `let` construct, where every occurrence of z would be replaced with τ_1 . Thanks to type scheme introduction and instantiation constraints, however, this effect is achieved without duplication of source code or constraints. In other words, constraint generation has linear time and space complexity.

10.4.2 EXERCISE [\star, \rightarrow]: Define the *size* of an expression, of a type, and of a constraint, viewed as abstract syntax trees. Check that the size of $\llbracket \tau : T \rrbracket$ is linear in the sum of the sizes of τ and T . \square

10.4.3 EXERCISE [RECOMMENDED, \star, \rightarrow]: Compute and simplify, as best as you can, the constraint $\llbracket \text{let } f = \lambda z. z \text{ in } f f : T \rrbracket$. \square

We now state several properties of constraint generation. We begin with soundness, whose statement was explained above.

10.4.4 THEOREM [SOUNDNESS]: $\text{let } \Gamma \text{ in } \llbracket \tau : T \rrbracket, \Gamma \vdash \tau : T$. \square

The following lemmas are used in the proof of the completeness property and in a number of other occasions. The first two state that $\llbracket \tau : T \rrbracket$ is *covariant* with respect to T . Roughly speaking, this means that enough subtyping constraints are generated to achieve completeness with respect to HMD-SUB.

10.4.5 LEMMA: $\llbracket \tau : T \rrbracket \wedge T \leq T' \Vdash \llbracket \tau : T' \rrbracket$. \square

10.4.6 LEMMA: $X \notin \text{ftv}(T)$ implies $\exists X. (\llbracket \tau : X \rrbracket \wedge X \leq T) \equiv \llbracket \tau : T \rrbracket$. \square

The next lemma gives a simplified version of the second constraint generation rule, in the specific case where the expected type is an arrow type. Thus, fresh type variables need not be generated; one may directly use the arrow's domain and codomain instead.

10.4.7 LEMMA: $\llbracket \lambda z. \tau : T_1 \rightarrow T_2 \rrbracket$ is equivalent to $\text{let } z : T_1 \text{ in } \llbracket \tau : T_2 \rrbracket$. \square

We conclude with the completeness property. The theorem states that if, within $\text{HM}(X)$, τ has type T under assumptions C and Γ , then C must be at

least as specific as $\text{let } \Gamma \text{ in } \llbracket \mathbf{t} : \mathbb{T} \rrbracket$. The statement requires C and Γ to have no free program identifiers, which is natural, since they are part of an $\text{HM}(X)$ judgment. The hypothesis $C \Vdash \exists \Gamma$ excludes the somewhat pathological situation where Γ contains constraints not apparent in C . This hypothesis vanishes when Γ is the initial environment; see Definition 10.5.2.

- 10.4.8 THEOREM [COMPLETENESS]: Let $C \Vdash \exists \Gamma$. Assume $\text{fpi}(C, \Gamma) = \emptyset$. If $C, \Gamma \vdash \mathbf{t} : \mathbb{T}$ holds in $\text{HM}(X)$, then C entails $\text{let } \Gamma \text{ in } \llbracket \mathbf{t} : \mathbb{T} \rrbracket$. \square

10.5 Type Soundness

We are now ready to establish type soundness for our type system. The statement that we wish to prove is sometimes known as *Milner's slogan*: “Well-typed programs do not go wrong” (Milner, 1978). Below, we define well-typedness in terms of our constraint generation rules, for the sake of convenience, and establish type soundness with respect to that particular definition. Theorems 10.3.6 and 10.4.8 imply that type soundness also holds when well-typedness is defined with respect to the typing judgments of DM or $\text{HM}(X)$. We establish type soundness by following Wright and Felleisen's so-called *syntactic approach* (1994). The approach consists of isolating two independent properties. *Subject reduction*, whose exact statement will be given below, implies that well-typedness is preserved by reduction. *Progress* states that no stuck configuration is well-typed. It is immediate to check that, if both properties hold, then no well-typed program can reduce to a stuck configuration. Subject reduction itself depends on a key lemma, usually known as a (term) *substitution lemma*. Here is a version of this lemma, stated in terms of the constraint generation rules.

- 10.5.1 LEMMA: $\text{let } z : \forall \bar{X}[\llbracket \mathbf{t}_2 : \mathbb{T}_2 \rrbracket]. \mathbb{T}_2 \text{ in } \llbracket \mathbf{t}_1 : \mathbb{T}_1 \rrbracket$ entails $\llbracket [z \mapsto \mathbf{t}_2] \mathbf{t}_1 : \mathbb{T}_1 \rrbracket$. \square

Before going on, let us give a few definitions and formulate several requirements. First, we must define an initial environment Γ_0 , which assigns a type scheme to every constant. A couple of requirements must be established to ensure that Γ_0 is consistent with the semantics of constants, as specified by $\xrightarrow{\delta}$. Second, we must extend constraint generation and well-typedness to configurations, as opposed to programs, since reduction operates on configurations. Last, we must formulate a restriction to tame the interaction between side effects and let -polymorphism, which is unsound if unrestricted.

- 10.5.2 DEFINITION: Let Γ_0 be an environment whose domain is the set of constants \mathcal{Q} . We require $\text{ftv}(\Gamma_0) = \emptyset$, $\text{fpi}(\Gamma_0) = \emptyset$, and $\exists \Gamma_0 \equiv \text{true}$. We refer to Γ_0 as the *initial typing environment*. \square

10.5.3 **DEFINITION:** Let ref be an isolated, invariant type constructor of signature $\star \Rightarrow \star$. A *store type* M is a finite mapping from memory locations to types. We write $\text{ref } M$ for the environment that maps every $m \in \text{dom}(M)$ to $\text{ref } M(m)$. Assuming $\text{dom}(\mu)$ and $\text{dom}(M)$ coincide, the constraint $\llbracket \mu : M \rrbracket$ is defined as the conjunction of the constraints $\llbracket \mu(m) : M(m) \rrbracket$, where m ranges over $\text{dom}(\mu)$. Under the same assumption, the constraint $\llbracket \tau/\mu : T/M \rrbracket$ is defined as $\llbracket \tau : T \rrbracket \wedge \llbracket \mu : M \rrbracket$. A closed configuration τ/μ is *well-typed* if and only if there exist a type T and a store type M such that $\text{dom}(\mu) = \text{dom}(M)$ and the constraint $\text{let } \Gamma_0; \text{ref } M \text{ in } \llbracket \tau/\mu : T/M \rrbracket$ is satisfiable. \square

The type $\text{ref } T$ is the type of references (that is, memory locations) that store data of type T (*TAPL*, Chapter 13). It must be *invariant* in its parameter, reflecting the fact that references may be both read and written.

A store is a complex object: it may contain values that indirectly refer to each other via memory locations. In fact, it is a representation of the graph formed by objects and pointers in memory, which may contain cycles. We rely on store types to deal with such cycles. In the definition of well-typedness, the store type M imposes a constraint on the contents of the store—the value $\mu(m)$ must have type $M(m)$ —but also plays the role of a hypothesis: by placing the constraint $\llbracket \tau/\mu : T/M \rrbracket$ within the context $\text{let } \text{ref } M \text{ in } []$, we give meaning to free occurrences of memory locations within $\llbracket \tau/\mu : T/M \rrbracket$, and stipulate that it is valid to assume that m has type $M(m)$. In other words, we essentially view the store as a large, mutually recursive binding of locations to values. The context $\text{let } \Gamma_0 \text{ in } []$ gives meaning to occurrences of constants within $\llbracket \tau/\mu : T/M \rrbracket$.

We now define a relation between configurations that plays a key role in the statement of the subject reduction property. The point of subject reduction is to guarantee that well-typedness is preserved by reduction. However, such a simple statement is too weak to be amenable to inductive proof. Thus, for the purposes of the proof, we must be more specific. To begin, let us consider the simpler case of a pure semantics, that is, a semantics without stores. Then, we must state that if an expression τ has type T under a certain constraint, then its reduct τ' has type T under the same constraint. In terms of generated constraints, this statement becomes: $\text{let } \Gamma_0 \text{ in } \llbracket \tau : T \rrbracket$ entails $\text{let } \Gamma_0 \text{ in } \llbracket \tau' : T \rrbracket$. Let us now return to the general case, where a store is present. The statement of well-typedness for a configuration τ/μ now involves a store type M whose domain is that of μ . So, the statement of well-typedness for its reduct τ'/μ' must involve a store type M' whose domain is that of μ' , which is larger if allocation occurred. The types of existing memory locations must not change: we must request that M and M' agree on $\text{dom}(M)$, that is, M' must extend M . Furthermore, the types assigned to new memory locations in

$dom(M') \setminus dom(M)$ might involve new type variables, that is, variables that do not appear free in M or T . We must allow these variables to be hidden—that is, existentially quantified—otherwise the entailment assertion cannot hold. These considerations lead us to the following definition:

- 10.5.4 DEFINITION: $\tau/\mu \sqsubseteq \tau'/\mu'$ holds if and only if, for every type T and for every store type M such that $dom(\mu) = dom(M)$, there exist a set of type variables \bar{Y} and a store type M' such that $\bar{Y} \# ftv(T, M)$ and $ftv(M') \subseteq \bar{Y} \cup ftv(M)$ and $dom(M') = dom(\mu')$ and M' extends M and

$$\begin{aligned} & \text{let } \Gamma_0; \text{ref } M \text{ in } \llbracket \tau / \mu : T / M \rrbracket \\ \Vdash \exists \bar{Y}. \text{let } \Gamma_0; \text{ref } M' \text{ in } \llbracket \tau' / \mu' : T / M' \rrbracket. \end{aligned}$$

The relation \sqsubseteq is intended to express a connection between a configuration and its reduct. Thus, subject reduction may be stated as: $(\rightarrow) \subseteq (\sqsubseteq)$, that is, \sqsubseteq is indeed a conservative description of reduction. \square

We have introduced an initial environment Γ_0 and used it in the definition of well-typedness, but we haven't yet ensured that the type schemes assigned to constants are an adequate description of their semantics. We now formulate two requirements that relate Γ_0 with $\xrightarrow{\delta}$. They are specializations of the subject reduction and progress properties to configurations that involve an application of a constant. They represent proof obligations that must be discharged when concrete definitions of \mathcal{Q} , $\xrightarrow{\delta}$, and Γ_0 are given.

- 10.5.5 DEFINITION: We require (i) $(\xrightarrow{\delta}) \subseteq (\sqsubseteq)$; and (ii) if the configuration $c \ v_1 \dots v_k / \mu$ (where $k \geq 0$) is well-typed, then either it is reducible, or $c \ v_1 \dots v_k$ is a value. \square

The last point that remains to be settled before proving type soundness is the interaction between side effects and let -polymorphism. The following example illustrates the problem:

$$\text{let } r = \text{ref } \lambda z.z \text{ in } \text{let } _ = (r := \lambda z.(z \hat{+} \hat{1})) \text{ in } !r \text{ true}$$

This expression reduces to $\text{true} \hat{+} \hat{1}$, so it must not be well-typed. Yet, if natural type schemes are assigned to ref , $!$, and $:=$ (see Example 10.7.5), then it is well-typed with respect to the rules given so far, because r receives the polymorphic type scheme $\forall X.\text{ref } (X \rightarrow X)$, which allows writing a function of type $\text{int} \rightarrow \text{int}$ into r and reading it back with type $\text{bool} \rightarrow \text{bool}$. The problem is that let -polymorphism simulates a textual duplication of the let -bound expression $\text{ref } \lambda z.z$, while the semantics first reduces it to a value m , causing a new binding $m \mapsto \lambda z.z$ to appear in the store, then duplicates the address m .

The new store binding is not duplicated: both copies of m refer to the same memory cell. For this reason, generalization is unsound in this case, and must be restricted. Many authors have attempted to come up with a sound type system that accepts all pure programs and remains flexible enough in the presence of side effects (Tofte, 1988; Leroy, 1992). These proposals are often complex, which is why they have been abandoned in favor of an extremely simple syntactic restriction, known as the *value restriction* (Wright, 1995).

- 10.5.6 DEFINITION: A program satisfies the *value restriction* if and only if all subexpressions of the form $\text{let } z = t_1 \text{ in } t_2$ are in fact of the form $\text{let } z = v_1 \text{ in } t_2$. In the following, we assume that either all constants have pure semantics, or all programs satisfy the value restriction. \square

Put slightly differently, the value restriction states that only values may be generalized. This eliminates the problem altogether, since duplicating values does not affect a program's semantics. Note that any program that does not satisfy the value restriction can be turned into one that does and has the same semantics: it suffices to change $\text{let } z = t_1 \text{ in } t_2$ into $(\lambda z.t_2) t_1$ when t_1 is not a value. Of course, such a transformation may cause the program to become ill-typed. In other words, the value restriction causes some perfectly safe programs to be rejected. In particular, in its above form, it prevents generalizing applications of the form $c v_1 \dots v_k$, where c is a destructor of arity k . This is excessive, because many destructors have pure semantics; only a few, such as `ref`, allocate new mutable storage. Furthermore, we use pure destructors to encode numerous language features (§10.7). Fortunately, it is easy to relax the restriction to allow generalizing not only values, but also a more general class of *nonexpansive* expressions, whose syntax guarantees that such expressions cannot allocate new mutable storage (that is, expand the domain of the store). The term *nonexpansive* was coined by Tofte (1988). Nonexpansive expressions may include applications of the form $c t_1 \dots t_k$, where c is a pure destructor of arity k and t_1, \dots, t_k are nonexpansive. Experience shows that this slightly relaxed restriction is acceptable in practice. Some limitations remain: for instance, *constructor functions* (that is, functions that do not allocate mutable storage and build a value) are regarded as ordinary functions, so their applications are considered potentially expansive, even though a naked constructor application would be a value and thus considered nonexpansive. For instance, in the expression $\text{let } f = c v \text{ in } \text{let } z = f w \text{ in } t$, where c is a constructor of arity 2, the partial application $c v$, to which the name f is bound, is a constructor function (of arity 1). The program variable z cannot receive a polymorphic type scheme, because $f w$ is not a value, even though it has the same semantic meaning as $c v w$, which is a value. A recent improvement to the value restriction (Garrigue, 2004)

provides a partial remedy. Technically, the effect of the value restriction (as stated in Definition 10.5.6) is summarized by the following result.

- 10.5.7 LEMMA: Under the value restriction, the production $\mathcal{E} ::= \text{let } z = \mathcal{E} \text{ in } \mathfrak{t}$ may be suppressed from the grammar of evaluation contexts (Figure 10-1) without altering the operational semantics. \square

We are finished with definitions and requirements. Let us now turn to the type soundness results.

- 10.5.8 THEOREM [SUBJECT REDUCTION]: $(\rightarrow) \subseteq (\sqsubseteq)$. \square

Subject reduction ensures that well-typedness is preserved by reduction.

- 10.5.9 COROLLARY: Let $\mathfrak{t}/\mu \rightarrow \mathfrak{t}'/\mu'$. If \mathfrak{t}/μ is well-typed, then so is \mathfrak{t}'/μ' . \square

Let us now state the progress property.

- 10.5.10 THEOREM [PROGRESS]: If \mathfrak{t}/μ is well-typed, then either it is reducible, or \mathfrak{t} is a value. \square

We may now conclude:

- 10.5.11 THEOREM [TYPE SOUNDNESS]: Well-typed source programs do not go wrong. \square

Recall that this result holds only if the requirements of Definition 10.5.5 are met. In other words, some proof obligations remain to be discharged when concrete definitions of \mathcal{Q} , δ , and Γ_0 are given. This is illustrated by several examples in §10.7 and §10.8.

10.6 Constraint Solving

We have introduced a parameterized constraint language, given equivalence laws describing the interaction between its logical connectives, and exploited them to prove theorems about type inference and type soundness, which are valid independently of the nature of primitive constraints—the so-called predicate applications. However, there would be little point in proposing a parameterized constraint solver, because much of the difficulty of designing an efficient constraint solver lies precisely in the treatment of primitive constraints and in its interaction with `let`-polymorphism. In this section, we focus on constraint solving in the setting of an *equality-only free tree model*. Thus, the constraint solver developed here allows performing type inference for $\text{HM}(=)$ (that is, for Damas and Milner's type system) and for its extension with recursive types. Of course, some of its mechanisms may be useful in other settings. The program analysis and type inference literature abounds

with constraint-based systems of all kinds; a short list of papers that put particular emphasis on constraint solving is Aiken and Wimmers (1992), Henglein (1993), Niehren, Müller, and Podelski (1997), Fähndrich (1999), Melski and Reps (2000), Müller, Niehren, and Treinen (2001), Pottier (2001b), Nielson, Nielson, and Seidl (2002), McAllester (2002; 2003), and Simonet (2003).

We begin with a rule-based presentation of a standard, efficient first-order unification algorithm. This yields a constraint solver for a subset of the constraint language, except for the type scheme introduction and instantiation forms. On top of it, we build a full constraint solver, which corresponds to the code that accompanies this chapter.

Unification

Unification is the process of solving equations between terms. It was first introduced by Robinson (1971), but his original algorithm could be very inefficient. Efficient algorithms, which perform unification in quasi-linear time, were independently proposed by Martelli and Montanari (1976; 1982) and by Huet (1976, Chapter 5). Both algorithms rely on a data structure that efficiently solves the *union-find* problem (Tarjan, 1975). Martelli and Montanari's algorithm performs unification in topological (top-down) order, and is thus restricted to the acyclic case, that is, to the case where equations are interpreted in a syntactic model. In this specific case, unification may actually be performed in truly linear time (Paterson and Wegman, 1978). On the other hand, Huet's algorithm is able to deal with cyclic structures. The acyclicity check is postponed until the very end of the solving process if equations are interpreted within a syntactic model, or omitted altogether if working within a regular tree model. Except for the final acyclicity check, Huet's algorithm is incremental. Furthermore, it is simple; we present a version of it here. Knight (1989) and Baader and Siekmann (1994) also describe Huet's algorithm, and provide further historical background and references.

Following Jouannaud and Kirchner (1991), we specify the algorithm as a (nondeterministic) system of constraint rewriting rules. As suggested above, it is almost the same for finite and regular tree models; only one rule, which implements the *occurs check*, must be removed in the latter case. In other words, the algorithm works with possibly cyclic data structures and does not rely in an essential way on the occurs check. In order to more closely reflect the behavior of the actual algorithm, and in particular the union-find data structure, we modify the syntax of constraints by replacing equations with multi-equations—equations involving an arbitrary number of types, as opposed to exactly two.

- 10.6.1 DEFINITION: Let there be, for every kind κ and for every $n \geq 1$, a predicate $=_{\kappa}^n$, of signature $\kappa^n \Rightarrow \cdot$, whose interpretation is (n -ary) equality. The predicate constraint $=_{\kappa}^n T_1 \dots T_n$ is written $T_1 = \dots = T_n$, and called a *multi-equation*. We consider the constraint true as a multi-equation of length 0 and let ϵ range over all multi-equations. In the following, we identify multi-equations up to permutations of their members, so a multi-equation ϵ of kind κ may be viewed as a finite *multiset* of types of kind κ . We write $\epsilon = \epsilon'$ for the multi-equation obtained by concatenating ϵ and ϵ' . \square

Thus, we are interested in the following subset of the constraint language:

$$U ::= \text{true} \mid \text{false} \mid \epsilon \mid U \wedge U \mid \exists \bar{X}.U$$

Equations are replaced with multi-equations; no other predicates are available. Type scheme introduction and instantiation forms are absent.

- 10.6.2 DEFINITION: A multi-equation is *standard* if and only if its variable members are distinct and it has at most one nonvariable member. A constraint U is *standard* if and only if every multi-equation inside U is standard and every variable that occurs (free or bound) in U is a member of at most one multi-equation inside U . (Note that to be a member of ϵ implies, but is not equivalent to, to occur free in ϵ .) \square

A union-find algorithm maintains equivalence classes (that is, disjoint sets) of variables, and associates with each class a *descriptor*, which in our case is either absent or a nonvariable term. Thus, a *standard* constraint represents a state of the union-find algorithm. A constraint that is not standard may be viewed as a superposition of a state of the union-find algorithm, on the one hand, and of control information, on the other hand. For instance, a multi-equation of the form $\epsilon = T_1 = T_2$, where ϵ is made up of distinct variables and T_1 and T_2 are nonvariable terms, may be viewed, roughly speaking, as the equivalence class $\epsilon = T_1$, together with a pending request to solve $T_1 = T_2$ and to update the class's descriptor accordingly. Because multi-equations encode both state and control, our specification of the unification algorithm remains rather abstract. It would be possible to give a lower-level description, where state (standard conjunctions of multi-equations) and control (pending binary equations) are distinguished.

- 10.6.3 DEFINITION: Let U be a conjunction of multi-equations. Y is *dominated* by X with respect to U (written: $Y <_U X$) if and only if U contains a conjunct of the form $X = F \bar{T} = \epsilon$, where $Y \in \text{fv}(\bar{T})$. U is *cyclic* if and only if the graph of $<_U$ exhibits a cycle. \square

$(\exists \vec{X}. U_1) \wedge U_2 \rightarrow \exists \vec{X}. (U_1 \wedge U_2)$	(S-EXAND)
$\text{if } \vec{X} \# \text{ftv}(U_2)$	
$X = \epsilon \wedge X = \epsilon' \rightarrow X = \epsilon = \epsilon'$	(S-FUSE)
$X = X = \epsilon \rightarrow X = \epsilon$	(S-STUTTER)
$F \vec{X} = F \vec{T} = \epsilon \rightarrow \vec{X} = \vec{T} \wedge F \vec{X} = \epsilon$	(S-DECOMPOSE)
$F T_1 \dots T_i \dots T_n = \epsilon \rightarrow \exists X. (X = T_i \wedge F T_1 \dots X \dots T_n = \epsilon)$	(S-NAME-1)
$\text{if } T_i \notin \mathcal{V} \wedge X \notin \text{ftv}(T_1, \dots, T_n, \epsilon)$	
$F \vec{T} = F' \vec{T}' = \epsilon \rightarrow \text{false}$	(S-CLASH)
$\text{if } F \neq F'$	
$T \rightarrow \text{true}$	(S-SINGLE)
$U \wedge \text{true} \rightarrow U$	(S-TRUE)
$U \rightarrow \text{false}$	(S-CYCLE)
$\text{if the model is syntactic and } U \text{ is cyclic}$	
$\mathcal{U}[\text{false}] \rightarrow \text{false}$	(S-FAIL)
$\text{if } \mathcal{U} \neq []$	

Figure 10-10: Unification

The specification of the unification algorithm consists of a set of constraint rewriting rules, given in Figure 10-10. Rewriting is performed modulo α -conversion, modulo permutations of the members of a multi-equation, modulo commutativity and associativity of conjunction, and under an arbitrary context. The specification is nondeterministic: several rule instances may be simultaneously applicable.

S-EXAND is a directed version of C-EXAND, whose effect is to float up all existential quantifiers. In the process, all multi-equations become part of a single conjunction, possibly causing rules whose left-hand side is a conjunction of multi-equations, namely S-FUSE and S-CYCLE, to become applicable. S-FUSE identifies two multi-equations that share a common variable X , and fuses them. The new multi-equation is not necessarily standard, even if the two original multi-equations were. Indeed, it may have repeated variables or contain two nonvariable terms. The purpose of the next few rules, whose left-hand side consists of a single multi-equation, is to deal with these situations. S-STUTTER eliminates redundant variables. It only deals with vari-

ables, as opposed to terms of arbitrary size, so as to have constant time cost. The comparison of nonvariable terms is implemented by S-DECOMPOSE and S-CLASH. S-DECOMPOSE decomposes an equation between two terms whose head symbols match. It produces a conjunction of equations between their subterms, namely $\vec{X} = \vec{T}$. Only one of the two terms remains in the original multi-equation, which may thus become standard. The terms \vec{X} are copied: there are two occurrences of \vec{X} on the right-hand side. For this reason, we require them to be type variables, as opposed to terms of arbitrary size. (We slightly abuse notation by using \vec{X} to denote a vector of type variables whose elements are not necessarily distinct.) By doing so, we allow explicit reasoning about sharing: since a variable represents a pointer to an equivalence class, we explicitly specify that only pointers, not whole terms, are copied. As a result of this decision, S-DECOMPOSE is not applicable when both terms at hand have a nonvariable subterm. S-NAME-1 remedies this problem by introducing a fresh variable that stands for one such subterm. When repeatedly applied, S-NAME-1 yields a unification problem composed of so-called *small terms* only—that is, where sharing has been made fully explicit. S-CLASH complements S-DECOMPOSE by dealing with the case where two terms with different head symbols are equated; in a free tree model, such an equation is false, so failure is signaled. S-SINGLE and S-TRUE suppress multi-equations of size 1 and 0, respectively, which are tautologies. S-CYCLE is the occurs check: it signals failure if the constraint is cyclic. It is applicable only in the case of syntactic unification, that is, when ground types are finite trees. It is a global check: its left-hand side is an entire conjunction of multi-equations. S-FAIL propagates failure; \mathcal{U} ranges over unification constraint contexts.

The constraint rewriting system in Figure 10-10 enjoys the following properties. First, rewriting is strongly normalizing, so the rules define a (non-deterministic) algorithm. Second, rewriting is meaning-preserving. Third, every normal form is either false or of the form $\exists \vec{X}.U$, where U is satisfiable. The latter two properties indicate that the algorithm is indeed a constraint solver.

- 10.6.4 LEMMA: The rewriting system \rightarrow is strongly normalizing. □
- 10.6.5 LEMMA: $U_1 \rightarrow U_2$ implies $U_1 \equiv U_2$. □
- 10.6.6 LEMMA: Every normal form is either false or of the form $\mathcal{X}[U]$, where \mathcal{X} is an existential constraint context, U is a standard conjunction of multi-equations and, if the model is syntactic, U is acyclic. These conditions imply that U is satisfiable. □

A Constraint Solver

On top of the unification algorithm, we now define a constraint solver. Its specification is independent of the rules and strategy employed by the unification algorithm. However, the structure of the unification algorithm's normal forms as well as the logical properties of multi-equations are exploited when performing generalization, that is, when creating and simplifying type schemes. Like the unification algorithm, the constraint solver is specified in terms of a reduction system. However, the objects that are subject to rewriting are not just constraints: they have more complex structure. Working with such richer states allows distinguishing the solver's external language—namely, the full constraint language, which is used to express the problem that one wishes to solve—and an internal language, introduced below, which is used to describe the solver's private data structures. In the following, C and D range over external constraints, that is, constraints that were part of the solver's input. External constraints are to be viewed as abstract syntax trees, subject to no implicit laws other than α -conversion. As a simplifying assumption, we require external constraints not to contain any occurrence of false—otherwise the problem at hand is clearly false. *Internal* data structures include unification constraints U , as previously studied, and *stacks*, whose syntax is as follows:

$$S ::= [] \mid S[[] \wedge C] \mid S[\exists \bar{X}. []] \mid S[\text{let } x : \forall \bar{X}[[]]. T \text{ in } C] \mid S[\text{let } x : \sigma \text{ in } []]$$

In the second and fourth productions, C is an external constraint. In the last production, we require σ to be of the form $\forall \bar{X}[U]. X$, and we demand $\exists \sigma \equiv \text{true}$. Every stack may be viewed as a one-hole constraint context (page 410); indeed, one may interpret $[]$ as the empty context and $\cdot[.]$ as context composition, which replaces the hole of its first context argument with its second context argument. A stack may also be viewed, literally, as a list of *frames*. Frames may be added and deleted at the inner end of a stack, that is, near the hole of the constraint context that it represents. We refer to the four kinds of frames as *conjunction*, *existential*, *let*, and *environment* frames, respectively. A *state* of the constraint solver is a triple $S; U; C$ where S is a stack, U is a unification constraint, and C is an external constraint. The state $S; U; C$ is to be understood as a representation of the constraint $S[U \wedge C]$, that is, the constraint obtained by placing both U and C within the hole of the constraint context S . The notion of α -equivalence between states is defined accordingly. In particular, one may rename type variables in $dtv(S)$, provided U and C are renamed as well. In brief, the three components of a state play the following roles. C is an external constraint that the solver intends to examine next. U

is the internal state of the underlying unification algorithm; one might think of it as the knowledge that has been obtained so far. S tells where the type variables that occur free in U and C are bound, associates type schemes with the program variables that occur free in C , and records what should be done after C is solved. The solver's initial state is usually of the form $[\]; \text{true}; C$, where C is the external constraint that one wishes to solve, that is, whose satisfiability one wishes to determine. If the constraint to be solved is of the form $\text{let } \Gamma_0 \text{ in } C$, and if the type schemes that appear within Γ_0 meet the requirements that bear on environment frames, as defined above, then it is possible to pick $\text{let } \Gamma_0 \text{ in } [\]; \text{true}; C$ as an initial state. For simplicity, we make the (unessential) assumption that states have no free type variables.

The solver consists of a (nondeterministic) state rewriting system, given in Figure 10-11. Rewriting is performed modulo α -conversion. S-UNIFY makes the unification algorithm a component of the constraint solver, and allows the current unification problem U to be solved at any time. Rules S-Ex-1 to S-Ex-4 float existential quantifiers out of the unification problem into the stack and through the stack up to the nearest enclosing let frame, if there is any, or to the outermost level, otherwise. Their side-conditions prevent capture of type variables, and can always be satisfied by suitable α -conversion of the left-hand state. If $S; U; C$ is a normal form with respect to these five rules, then U must be either false or a conjunction of standard multi-equations, and every type variable in $\text{dte}(S)$ must be either universally quantified at a let frame or existentially bound at the outermost level. (Recall that, by assumption, states have no free type variables.) In other words, provided these rules are applied in an eager fashion, there is no need for existential frames to appear in the machine representation of stacks. Instead, it suffices to maintain, at every let frame and at the outermost level, a list of the type variables that are bound at this point and, conversely, to annotate every type variable in $\text{dte}(S)$ with an integer *rank*, which allows telling, in constant time, where the variable is bound: type variables of rank 0 are bound at the outermost level, and type variables of rank $k \geq 1$ are bound at the k^{th} let frame down in the stack S . The code that accompanies this chapter adopts this convention. Ranks were initially described in Rémy (1992a) and have also been studied by McAllester (2003).

Rules S-SOLVE-EQ to S-SOLVE-LET encode an analysis of the structure of the third component of the current state. There is one rule for each possible case, except false, which by assumption cannot arise, and true, which is dealt with further on. S-SOLVE-EQ discovers an equation and makes it available to the unification algorithm. S-SOLVE-ID discovers an instantiation constraint $x \preceq T$ and replaces it with $\sigma \preceq T$, where the type scheme $\sigma = S(x)$ is the type

$S; U; C$	\rightarrow	$S; U'; C$ if $U \rightarrow U'$	(S-UNIFY)
$S; \exists \bar{X}. U; C$	\rightarrow	$S[\exists \bar{X}. []]; U; C$ if $\bar{X} \# \text{ftv}(C)$	(S-EX-1)
$S[(\exists \bar{X}. S') \wedge D]; U; C$	\rightarrow	$S[\exists \bar{X}. (S' \wedge D)]; U; C$ if $\bar{X} \# \text{ftv}(D)$	(S-EX-2)
$S[\text{let } x : \forall \bar{X}[\exists \bar{Y}. S']. T \text{ in } D]; U; C$	\rightarrow	$S[\text{let } x : \forall \bar{X}\bar{Y}[S]. 'T \text{ in } D]; U; C$ if $\bar{Y} \# \text{ftv}(T)$	(S-EX-3)
$S[\text{let } x : \sigma \text{ in } \exists \bar{X}. S']; U; C$	\rightarrow	$S[\exists \bar{X}. \text{let } x : \sigma \text{ in } S']; U; C$ if $\bar{X} \# \text{ftv}(\sigma)$	(S-EX-4)
$S; U; T_1 = T_2$	\rightarrow	$S; U \wedge T_1 = T_2; \text{true}$	(S-SOLVE-EQ)
$S; U; x \leq T$	\rightarrow	$S; U; S(x) \leq T$	(S-SOLVE-ID)
$S; U; C_1 \wedge C_2$	\rightarrow	$S[[] \wedge C_2]; U; C_1$	(S-SOLVE-AND)
$S; U; \exists \bar{X}. C$	\rightarrow	$S[\exists \bar{X}. []]; U; C$ if $\bar{X} \# \text{ftv}(U)$	(S-SOLVE-EX)
$S; U; \text{let } x : \forall \bar{X}[D]. T \text{ in } C$	\rightarrow	$S[\text{let } x : \forall \bar{X}[[]]. T \text{ in } C]; U; D$ if $\bar{X} \# \text{ftv}(U)$	(S-SOLVE-LET)
$S[[] \wedge C]; U; \text{true}$	\rightarrow	$S; U; C$	(S-POP-AND)
$S[\text{let } x : \forall \bar{X}[[]]. T \text{ in } C]; U; \text{true}$	\rightarrow	$S[\text{let } x : \forall \bar{X}X[[]]. X \text{ in } C];$ $U \wedge X = T; \text{true}$ if $X \notin \text{ftv}(U, T) \wedge T \notin \mathcal{V}$	(S-NAME-2)
$S[\text{let } x : \forall \bar{X}Y[[]]. X \text{ in } C]; Y = Z = \epsilon \wedge U; \text{true}$	\rightarrow	$S[\text{let } x : \forall \bar{X}Y[[]]. \theta(X) \text{ in } C];$ $Y \wedge Z = \theta(\epsilon) \wedge \theta(U); \text{true}$ if $Y \neq Z \wedge \theta = [Y \mapsto Z]$	(S-COMPRESS)
$S[\text{let } x : \forall \bar{X}Y[[]]. X \text{ in } C]; Y = \epsilon \wedge U; \text{true}$	\rightarrow	$S[\text{let } x : \forall \bar{X}[[]]. X \text{ in } C]; \epsilon \wedge U; \text{true}$ if $Y \notin X \cup \text{ftv}(\epsilon, U)$	(S-UNNAME)
$S[\text{let } x : \forall \bar{X}\bar{Y}[[]]. X \text{ in } C]; U; \text{true}$	\rightarrow	$S[\exists \bar{Y}. \text{let } x : \forall \bar{X}[[]]. X \text{ in } C]; U; \text{true}$ if $\bar{Y} \# \text{ftv}(C) \wedge \exists \bar{X}. U$ determines \bar{Y}	(S-LETALL)
$S[\text{let } x : \forall \bar{X}[[]]. X \text{ in } C]; U_1 \wedge U_2; \text{true}$	\rightarrow	$S[\text{let } x : \forall \bar{X}[U_2]. X \text{ in } []]; U_1; C$ if $\bar{X} \# \text{ftv}(U_1) \wedge \exists \bar{X}. U_2 \equiv \text{true}$	(S-POP-LET)
$S[\text{let } x : \sigma \text{ in } []]; U; \text{true}$	\rightarrow	$S; U; \text{true}$	(S-POP-ENV)

Figure 10-11: A constraint solver

scheme carried by the nearest environment frame that defines x in the stack S . It is defined as follows:

$$\begin{aligned}
S[[] \wedge C](x) &= S(x) \\
S[\exists \bar{X}.[]](x) &= S(x) \quad \text{if } \bar{X} \# \text{ftv}(S(x)) \\
S[\text{let } y : \forall \bar{X}[[]].T \text{ in } C](x) &= S(x) \quad \text{if } \bar{X} \# \text{ftv}(S(x)) \\
S[\text{let } y : \sigma \text{ in } []](x) &= S(x) \quad \text{if } x \neq y \\
S[\text{let } x : \sigma \text{ in } []](x) &= \sigma
\end{aligned}$$

If $x \in \text{dpi}(S)$ does not hold, then $S(x)$ is undefined and the rule is not applicable. If it does hold, then the rule may always be made applicable by suitable α -conversion of the left-hand state. Recall that, if σ is of the form $\forall \bar{X}[U].X$, where $\bar{X} \# \text{ftv}(T)$, then $\sigma \leq T$ stands for $\exists \bar{X}.(U \wedge X = T)$. The process of constructing this constraint is informally referred to as “taking an instance of σ .” In the worst case, it is just as inefficient as textually expanding the corresponding `let` construct in the program’s source code, and leads to exponential time complexity. In practice, however, the unification constraint U is often compact because it was simplified before the environment frame `let $x : \sigma$ in $[]$` was created, which explains why the solver usually performs well. (The creation of environment frames, performed by `S-POP-LET`, is discussed below.) `S-SOLVE-AND` discovers a conjunction. It arbitrarily chooses to explore the left branch first, and pushes a conjunction frame onto the stack, so as to record that the right branch should be explored afterwards. `S-SOLVE-EX` discovers an existential quantifier and enters it, creating a new existential frame to record its existence. Similarly, `S-SOLVE-LET` discovers a `let` form and enters its left-hand side, creating a new `let` frame to record its existence. The choice of examining the left-hand side first is not arbitrary. Indeed, examining the right-hand side first would require creating an environment frame—but environment frames must contain simplified type schemes of the form $\forall \bar{X}[U].X$, whereas the type scheme $\forall \bar{X}[D].T$ is arbitrary. In other words, our strategy is to simplify type schemes prior to allowing them to be copied by `S-SOLVE-ID`, so as to avoid any duplication of effort. The side-conditions of `S-SOLVE-EX` and `S-SOLVE-LET` may always be satisfied by suitable α -conversion of the left-hand state.

Rules `S-SOLVE-EQ` to `S-SOLVE-LET` may be referred to as *forward* rules, because they “move down into” the external constraint, causing the stack to grow. This process stops when the external constraint at hand becomes true. Then part of the work has been finished, and the solver must examine the stack in order to determine what to do next. This task is performed by the last series of rules, which may be referred to as *backward* rules, because they “move back out,” causing the stack to shrink and possibly scheduling new external constraints for examination. These rules encode an analysis of the

structure of the innermost stack frame. There are three cases, corresponding to conjunction, let, and environment frames. The case of existential stack frames need not be considered, because rules S-Ex-2 to S-Ex-4 allow either fusing them with let frames or floating them up to the outermost level, where they shall remain inert. S-POP-AND deals with conjunction frames. The frame is popped, and the external constraint that it carries is scheduled for examination. S-POP-ENV deals with environment frames. Because the right-hand side of the let construct at hand has been solved—that is, turned into a unification constraint U —it cannot contain an occurrence of x . Furthermore, by assumption, $\exists\sigma$ is true. Thus, this environment frame is no longer useful: it is destroyed. The remaining rules deal with let frames. Roughly speaking, their purpose is to change the state $S[\text{let } x : \forall\tilde{X}[\]].T \text{ in } C; U; \text{true}$ into $S[\text{let } x : \forall\tilde{X}[U].T \text{ in } [\]]; \text{true}; C$, that is, to turn the current unification constraint U into a type scheme, turn the let frame into an environment frame, and schedule the right-hand side of the let construct (that is, the external constraint C) for examination. In fact, the process is more complex, because the type scheme $\forall\tilde{X}[U].T$ must be *simplified* before becoming part of an environment frame. The simplification process is described by rules S-NAME-2 to S-POP-LET. In the following, we refer to type variables in \tilde{X} as *young* and to type variables in $\text{dtr}(S) \setminus \tilde{X}$ as *old*. The former are the universal quantifiers of the type scheme that is being created; the latter contain its free type variables.

S-NAME-2 ensures that the body T of the type scheme that is being created is a type variable, as opposed to an arbitrary term. If it isn't, then it is replaced with a fresh variable X , and the equation $X = T$ is added so as to recall that X stands for T . Thus, the rule moves the term T into the current unification problem, where it potentially becomes subject to S-NAME-1. This ensures that sharing is made explicit everywhere. S-COMPRESS determines that the (young) type variable Y is an alias for the type variable Z . Then, every free occurrence of Y other than its defining occurrence is replaced with Z . In an actual implementation, this occurs transparently when the union-find algorithm performs *path compression* (Tarjan, 1975, 1979). We note that the rule does not allow substituting a younger type variable for an older one; indeed, that would make no sense, since the younger variable could then possibly escape its scope. In other words, in implementation terms, the union-find algorithm must be slightly modified so that, in each equivalence class, the representative element is always a type variable with minimum rank. S-UNNAME determines that the (young) type variable Y has no occurrences other than its defining occurrence in the current type scheme. (This occurs, in particular, when S-COMPRESS has just been applied.) Then, Y is suppressed altogether. In the particular case where the remaining multi-equation ϵ has cardinal 1,

it may then be suppressed by S-SINGLE. In other words, the combination of S-UNNAME and S-SINGLE is able to suppress young unused type variables as well as the term that they stand for. This may, in turn, cause new type variables to become eligible for elimination by S-UNNAME. In fact, assuming the current unification constraint is acyclic, an inductive argument shows that every young type variable may be suppressed unless it is dominated either by X or by an old type variable. (In the setting of a regular tree model, it is possible to extend the rule so that young cycles that are not dominated either by X or by an old type variable are suppressed as well.) S-LETALL is a directed version of C-LETALL. It turns the young type variables \bar{Y} into old variables. How to tell whether $\exists \bar{X}.U$ determines \bar{Y} is discussed later (see Lemma 10.6.7). Why S-LETALL is an interesting and important rule will be explained shortly. S-POP-LET is meant to be applied when the current state has become a normal form with respect to S-UNIFY, S-NAME-2, S-COMPRESS, S-UNNAME, and S-LETALL, that is, when the type scheme that is about to be created is fully simplified. It splits the current unification constraint into two components U_1 and U_2 , where U_1 is made up entirely of *old* variables, as expressed by the side-condition $\bar{X} \# \text{ftv}(U_1)$, and U_2 constrains *young* variables only, as expressed by the side-condition $\exists \bar{X}.U_2 \equiv \text{true}$. Note that U_2 may still contain free occurrences of old type variables, so the type scheme $\forall \bar{X}[U_2].X$ that appears on the right-hand side is not necessarily closed. It is not obvious why such a decomposition must exist; Lemma 10.6.10 proves that it does. Let us say for now that S-LETALL plays a role in guaranteeing its existence, whence comes part of its importance. Once the decomposition $U_1 \wedge U_2$ is obtained, the behavior of S-POP-LET is simple. The unification constraint U_1 concerns old variables only, that is, variables that are not quantified in the current let frame; thus, it need not become part of the new type scheme and may instead remain part of the current unification constraint. This is justified by C-LETAND and C-INAND* and corresponds to the difference between HMX-GEN' and HMX-GEN discussed in §10.3. The unification constraint U_2 , on the other hand, becomes part of the newly built type scheme $\forall \bar{X}[U_2].X$. The property $\exists \bar{X}.U_2 \equiv \text{true}$ guarantees that the newly created environment frame meets the requirements imposed on such frames. Note that the more type variables are considered old, the larger U_1 may become, and the smaller U_2 . This is another reason why S-LETALL is interesting: by allowing more variables to be considered old, it decreases the size of the type scheme $\forall \bar{X}[U_2].X$, making it cheaper to instantiate.

To complete our description of the constraint solver, there remains to explain how to decide when $\exists \bar{X}.U$ determines \bar{Y} , since this predicate occurs in the side-condition of S-LETALL. The following lemma describes two important situations where, by examining the structure of an equation, it is possible to

discover that a constraint C determines some of its free type variables \bar{Y} (Definition 10.2.14). In the first situation, the type variables \bar{Y} are equated with or dominated by a distinct type variable X that occurs free in C . In that case, because the model is a free tree model, the values of the type variables \bar{Y} are determined by the value of X : they are subtrees of it at specific positions. For instance, $X = Y_1 \rightarrow Y_2$ determines Y_1Y_2 , while $\exists Y_1.(X = Y_1 \rightarrow Y_2)$ determines Y_2 . In the second situation, the type variables \bar{Y} are equated with a term T , all of whose type variables are free in C . Again, the value of the type variables \bar{Y} is then determined by the values of the type variables $ftv(T)$. For instance, $X = Y_1 \rightarrow Y_2$ determines X , while $\exists Y_1.(X = Y_1 \rightarrow Y_2)$ does not. In the second situation, no assumption is in fact made about the model. (Note that $X = Y_1 \rightarrow Y_2$ determines Y_1Y_2 and determines X , but does not simultaneously determine XY_1Y_2 .)

- 10.6.7 LEMMA: Let $\bar{X} \# \bar{Y}$. Assume either ϵ is $X = \epsilon'$, where $X \notin \bar{X}\bar{Y}$ and $\bar{Y} \subseteq ftv(\epsilon')$, or ϵ is $\bar{Y} = T = \epsilon'$, where $ftv(T) \# \bar{X}\bar{Y}$. Then, $\exists \bar{X}.(C \wedge \epsilon)$ determines \bar{Y} . \square

Thanks to Lemma 10.6.7, an efficient implementation of S-LETALL comes to mind. The problem is, given a constraint $\exists \bar{X}.U$, where U is a standard conjunction of multi-equations, to determine the greatest subset \bar{Y} of \bar{X} such that $\exists(\bar{X} \setminus \bar{Y}).U$ determines \bar{Y} . By the first part of the lemma, it is safe for \bar{Y} to include all members of \bar{X} that are directly or indirectly dominated (with respect to U) by some free variable of $\exists \bar{X}.U$. Those can be found, in time linear in the size of U , by a top-down traversal of the graph of $<_U$. By the second part of the lemma, it is safe to close \bar{Y} under the closure law $X \in \bar{X} \wedge (\forall Y \ Y <_U \ X \Rightarrow Y \in \bar{Y}) \Rightarrow X \in \bar{Y}$. That is, it is safe to also include all members of \bar{X} whose descendants (with respect to U) have already been found to be members of \bar{Y} . This closure computation may be performed, again in linear time, by a bottom-up traversal of the graph of $<_U$. When U is acyclic, it is possible to show that this procedure is complete, that is, does compute the greatest subset \bar{Y} that meets our requirement.

The above discussion has shown that when Y and Z are equated, if Y is young and Z is old, then S-LETALL allows making Y old as well. If binding information is encoded in terms of integer ranks, as suggested earlier, then this remark may be formulated as follows: when Y and Z are equated, if the rank of Y exceeds that of Z , then it may be decreased so that both ranks match. As a result, it is possible to attach ranks with multi-equations, rather than with variables. When two multi-equations are fused, the smaller rank is kept. This treatment of ranks is inspired by Rémy (1992a); see the resolution rule FUSE, as well as the simplification rules PROPAGATE and REALIZE, in that paper.

Let us now state the properties of the constraint solver. First, the reduction system is terminating, so it defines an algorithm.

10.6.8 LEMMA: The reduction system \rightarrow is strongly normalizing. \square

Second, every rewriting step preserves the meaning of the constraint that the current state represents. We recall that the state $S;U;C$ is meant to represent the constraint $S[U \wedge C]$.

10.6.9 LEMMA: $S;U;C \rightarrow S';U';C'$ implies $S[U \wedge C] \equiv S'[U' \wedge C']$. \square

Last, we classify the normal forms of the reduction system:

10.6.10 LEMMA: A normal form for the reduction system \rightarrow is one of (i) $S;U;x \leq T$, where $x \notin \text{dpi}(S)$; (ii) $S;\text{false};\text{true}$; or (iii) $X;U;\text{true}$, where X is an existential constraint context and U a satisfiable conjunction of multi-equations. \square

In case (i), the constraint $S[U \wedge C]$ has a free program identifier x . In other words, the source program contains an unbound program identifier. Such an error could of course be detected prior to constraint solving, if desired. In case (ii), the unification algorithm failed. By Lemma 10.2.17, the constraint $S[U \wedge C]$ is then false. In case (iii), the constraint $S[U \wedge C]$ is equivalent to $X[U]$, where U is satisfiable, so it is satisfiable as well. If the initial constraint is closed, case (i) cannot arise, while cases (ii) and (iii) respectively denote failure and success. Thus, Lemmas 10.6.9 and 10.6.10 indeed prove that the algorithm is a constraint solver.

10.6.11 REMARK: Type inference for ML-the-calculus is DEXPTIME-complete (Kfoury, Tiuryn, and Urzyczyn, 1990; Mairson, Kanellakis, and Mitchell, 1991). Thus, our constraint solver cannot run any faster, asymptotically. This cost is essentially due to `let`-polymorphism, which requires a constraint to be duplicated at every occurrence of a `let`-bound variable (S-SOLVE-ID). In order to limit the amount of duplication to a bare minimum, it is important that rule S-LETALL be applied before S-POP-LET, allowing variables and constraints that need not be duplicated to be shared. We have observed that algorithms based on this strategy behave remarkably well in practice (Rémy, 1992a). In fact, McAllester (2003) has proved that they have linear time complexity, provided the size of type schemes and the (left-) nesting depth of `let` constructs are bounded. Unfortunately, many implementations of type inference for ML-the-programming-language do not behave as efficiently as the algorithm presented here. Some spend an excessive amount of time in computing the set of nongeneralizable type variables; some do not treat types as dags, thus losing precious sharing information; others perform the expensive occurs check after every unification step, rather than only once at every `let` construct, as suggested here (S-POP-LET). \square

10.7 From ML-the-Calculus to ML-the-Language

In this section, we explain how to extend the framework developed so far to accommodate operations on values of base type (such as integers), pairs, sums, references, and recursive function definitions. Then, we describe algebraic data type definitions. Last, the issues associated with recursive types are briefly discussed. For space reasons, exceptions are not discussed; the reader is referred to (*TAPL*, Chapter 14).

Simple Extensions

Introducing new constants and extending $\xrightarrow{\delta}$ and Γ_0 appropriately allows adding many features of ML-the-programming-language to ML-the-calculus. In each case, it is necessary to check that the requirements of Definition 10.5.5 are met, that is, to ensure that the new initial environment faithfully reflects the nature of the new constants as well as the behavior of the new reduction rules. Below, we describe several such extensions in isolation. The first exercise establishes a technical result that is useful in the next exercises.

- 10.7.1 EXERCISE [RECOMMENDED, ★]: Let Γ_0 contain the binding $c : \forall \bar{X}. T_1 \rightarrow \dots \rightarrow T_n \rightarrow T$. Prove $\text{let } \Gamma_0 \text{ in } \llbracket c \ t_1 \ \dots \ t_n : T' \rrbracket$ equivalent to $\text{let } \Gamma_0 \text{ in } \exists \bar{X}. (\bigwedge_{i=1}^n \llbracket t_i : T_i \rrbracket \wedge T \leq T')$. \square
- 10.7.2 EXERCISE [INTEGERS, RECOMMENDED, ★★]: Integer literals and integer addition have been introduced and given an operational semantics in Examples 10.1.1, 10.1.2, and 10.1.4. Let us now introduce an isolated type constructor int of signature \star and extend the initial environment Γ_0 with the bindings $\hat{n} : \text{int}$, for every integer n , and $\hat{+} : \text{int} \rightarrow \text{int} \rightarrow \text{int}$. Check that these definitions meet the requirements of Definition 10.5.5. \square
- 10.7.3 EXERCISE [PAIRS, ★★, ⇨]: Pairs and pair projections have been introduced and given an operational semantics in Examples 10.1.3 and 10.1.5. Let us now introduce an isolated type constructor \times of signature $\star \otimes \star \Rightarrow \star$, covariant in both of its parameters, and extend the initial environment Γ_0 with the following bindings:

$$\begin{aligned} (\cdot, \cdot) &: \forall XY. X \rightarrow Y \rightarrow X \times Y \\ \pi_1 &: \forall XY. X \times Y \rightarrow X \\ \pi_2 &: \forall XY. X \times Y \rightarrow Y \end{aligned}$$

Check that these definitions meet the requirements of Definition 10.5.5. \square

- 10.7.4 EXERCISE [SUMS, ★★, ⇨]: Sums have been introduced and given an operational semantics in Example 10.1.7. Let us now introduce an isolated type constructor $+$ of signature $\star \otimes \star \Rightarrow \star$, covariant in both of its parameters, and extend

the initial environment Γ_0 with the following bindings:

$$\begin{aligned} \text{inj}_1 &: \forall XY. X \rightarrow X + Y \\ \text{inj}_2 &: \forall XY. Y \rightarrow X + Y \\ \text{case} &: \forall XYZ. (X + Y) \rightarrow (X \rightarrow Z) \rightarrow (Y \rightarrow Z) \rightarrow Z \end{aligned}$$

Check that these definitions meet the requirements of Definition 10.5.5. \square

- 10.7.5 EXERCISE [REFERENCES, ★★★]: References have been introduced and given an operational semantics in Example 10.1.9. The type constructor `ref` has been introduced in Definition 10.5.3. Let us now extend the initial environment Γ_0 with the following bindings:

$$\begin{aligned} \text{ref} &: \forall X. X \rightarrow \text{ref } X \\ ! &: \forall X. \text{ref } X \rightarrow X \\ := &: \forall X. \text{ref } X \rightarrow X \rightarrow X \end{aligned}$$

Check that these definitions meet the requirements of Definition 10.5.5. \square

- 10.7.6 EXERCISE [RECURSION, RECOMMENDED, ★★★, →]: The fixpoint combinator `fix` has been introduced and given an operational semantics in Example 10.1.10. Let us now extend the initial environment Γ_0 with the following binding:

$$\text{fix}: \forall XY. ((X \rightarrow Y) \rightarrow (X \rightarrow Y)) \rightarrow X \rightarrow Y$$

Check that these definitions meet the requirements of Definition 10.5.5. Recall how the `letrec` syntactic sugar was defined in Example 10.1.10, and check that this gives rise to the following constraint generation rule:

$$\begin{aligned} &\text{let } \Gamma_0 \text{ in } \llbracket \text{letrec } f = \lambda z. t_1 \text{ in } t_2 : T \rrbracket \\ \equiv &\text{ let } \Gamma_0 \text{ in let } f : \forall XY. [\text{let } f : X \rightarrow Y; z : X \text{ in } \llbracket t_1 : Y \rrbracket]. X \rightarrow Y \text{ in } \llbracket t_2 : T \rrbracket \end{aligned}$$

Note the somewhat peculiar structure of this constraint: the program variable `f` is bound twice in it, with different type schemes. The constraint requires all occurrences of `f` within t_1 to be assigned the monomorphic type $X \rightarrow Y$. This type is generalized and turned into a type scheme before inspecting t_2 , however, so every occurrence of `f` within t_2 may receive a different type, as usual with `let`-polymorphism. A more powerful way of typechecking recursive function definitions, proposed by (Mycroft, 1984) and known as *polymorphic recursion*, allows the types of occurrences of `f` within t_1 to be possibly distinct instances of a single type scheme. However, type inference for this extension is equivalent to semi-unification (Henglein, 1993), which has been proved undecidable (Kfoury, Tiuryn, and Urzyczyn, 1993). Hence, type inference must either require type annotations or rely on a semi-algorithm. \square

In the exercises above, we have considered a number of extensions (integers, booleans, pairs, etc.) in isolation. We have checked that each of them preserves type soundness. Unfortunately, this does not in general imply that their combination preserves type soundness. In fact, it is possible to prove that these extensions are independent in a suitable sense and that independent extensions may be safely combined. Unfortunately, we lack space to further explain these notions.

Algebraic Data Types

Exercises 10.7.3 and 10.7.4 have shown how to extend the language with binary, anonymous products and sums. These constructs are quite general but still have several shortcomings. First, they are only binary, while we would like to have k -ary products and sums, for arbitrary $k \geq 0$. Such a generalization is of course straightforward. Second, more interestingly, their components must be referred to by numeric index (as in “extract the second component of the pair”), rather than by name (“extract the component named y ”). In practice, it is crucial to use names, because they make programs more readable and more robust in the face of changes. One could introduce a mechanism that allows defining names as syntactic sugar for numeric indices. That would help a little, but not much, because these names would not appear in types, which would still be made of anonymous products and sums. Third, in the absence of recursive types, products and sums do not have sufficient expressiveness to allow defining unbounded data structures, such as lists. Indeed, it is easy to see that every value whose type T is composed of base types (`int`, `bool`, etc.), products, and sums must have bounded size, where the bound $|T|$ is a function of T . More precisely, up to a constant factor, we have $|\text{int}| = |\text{bool}| = 1$, $|T_1 \times T_2| = 1 + |T_1| + |T_2|$, and $|T_1 + T_2| = 1 + \max(|T_1|, |T_2|)$. The following example describes another facet of the same problem.

- 10.7.7 **EXAMPLE:** A list is either empty, or a pair of an element and another list. So, it seems natural to try and encode the type of lists as a sum of some arbitrary type (say, `unit`) on the one hand, and of a product of some element type and of the type of lists itself on the other hand. With this encoding in mind, we can go ahead and write code—for instance, a function that computes the length of a list:

$$\text{letrec length} = \lambda l. \text{case } l \ (\lambda_ \hat{0}) \ (\lambda z. \hat{1} \hat{+} \text{length} \ (\pi_2 \ z))$$

We have used integers, pairs, sums, and the `letrec` construct introduced in the previous section. The code analyzes the list `l` using a `case` construct.

If the left branch is taken, the list is empty, so 0 is returned. If the right branch is taken, then z becomes bound to a pair of some element and the tail of the list. The latter is obtained using the projection operator π_2 . Its length is computed using a recursive call to `length` and incremented by 1. This code makes perfect sense. However, applying the constraint generation and constraint solving algorithms eventually leads to an equation of the form $X = Y + (Z \times X)$, where X stands for the type of `!`. This equation accurately reflects our encoding of the type of lists. However, in a syntactic model, it has no solution, so our definition of `length` is ill-typed. It is possible to adopt a free regular tree model, thus introducing *equirecursive* types into the system (*TAPL*, Chapter 20); however, there are good reasons not to do so (see the section on Recursive Types on p. 459). \square

To work around this problem, ML-the-programming-language offers *algebraic data type* definitions, whose elegance lies in the fact that, while representing only a modest theoretical extension, they do solve the three problems mentioned above. An algebraic data type may be viewed as an *abstract type* that is declared to be *isomorphic* to a (k -ary) product or sum type with named components. The type of each component is declared, as well, and may refer to the algebraic data type that is being defined: thus, algebraic data types are *isorecursive* (*TAPL*, Chapter 20). In order to allow sufficient flexibility when declaring the type of each component, algebraic data type definitions may be parameterized by a number of type variables. Last, in order to allow the description of complex data structures, it is necessary to allow several algebraic data types to be defined at once; the definitions may then be mutually recursive. In fact, in order to simplify this formal presentation, we assume that all algebraic data types are defined at once at the beginning of the program. This decision is, of course, at odds with modular programming but will not otherwise be a problem.

In the following, D ranges over a set of *data types*. We assume that data types form a subset of type constructors. We require each of them to be isolated and to have image kind \star . Furthermore, ℓ ranges over a set \mathcal{L} of *labels*, which we use both as data constructors and as record labels. An *algebraic data type definition* is either a *variant type* definition or a *record type* definition, whose respective forms are

$$D\vec{X} \approx \sum_{i=1}^k \ell_i : T_i \quad \text{and} \quad D\vec{X} \approx \prod_{i=1}^k \ell_i : T_i.$$

In either case, k must be nonnegative. If D has signature $\vec{k} \Rightarrow \star$, then the type variables \vec{X} must have kind \vec{k} . Every T_i must have kind \star . We refer to \vec{X} as the *parameters* and to \vec{T} (the vector formed by T_1, \dots, T_k) as the *components*

of the definition. The parameters are bound within the components, and the definition must be closed, that is, $ftv(\vec{T}) \subseteq \vec{X}$ must hold. Last, for an algebraic data type definition to be valid, the behavior of the type constructor D with respect to subtyping must match its definition. This requirement is clarified below.

- 10.7.8 **DEFINITION:** Consider an algebraic data type definition whose parameters and components are respectively \vec{X} and \vec{T} . Let \vec{X}' and \vec{T}' be their images under an arbitrary renaming. Then, $D\vec{X} \leq D\vec{X}' \Vdash \vec{T} \leq \vec{T}'$ must hold. \square

Because it is stated in terms of an entailment assertion, the above requirement bears on the interpretation of subtyping. The idea is, since $D\vec{X}$ is declared to be isomorphic to (a sum or a product of) \vec{T} , whenever two types built with D are comparable, their unfoldings should be comparable as well. The reverse entailment assertion is not required for type soundness, and it is sometimes useful to declare algebraic data types that do not validate it—so-called *phantom types* (Fluet and Pucella, 2002). Note that the requirement may always be satisfied by making the type constructor D invariant in all of its parameters. Indeed, in that case, $D\vec{X} \leq D\vec{X}'$ entails $\vec{X} = \vec{X}'$, which must entail $\vec{T} = \vec{T}'$ since \vec{T}' is precisely $[\vec{X} \mapsto \vec{X}']\vec{T}$. In an equality free tree model, every type constructor is naturally invariant, so the requirement is trivially satisfied. In other settings, however, it is often possible to satisfy the requirement of Definition 10.7.8 while assigning D a less restrictive variance. The following example illustrates such a case.

- 10.7.9 **EXAMPLE:** Let `list` be a data type of signature $\star \Rightarrow \star$. Let `Ni1` and `Cons` be data constructors. Then, the following is a definition of `list` as a variant type:

$$\text{list } X \approx \Sigma (\text{Ni1} : \text{unit}; \text{Cons} : X \times \text{list } X)$$

Because data types form a subset of type constructors, it is valid to form the type `list` X in the right-hand side of the definition, even though we are still in the process of defining the meaning of `list`. In other words, data type definitions may be recursive. However, because \approx is not interpreted as equality, the type `list` X is not a recursive type: it is nothing but an application of the unary type constructor `list` to the type variable X . To check that the definition of `list` satisfies the requirement of Definition 10.7.8, we must ensure that

$$\text{list } X \leq \text{list } X' \Vdash \text{unit} \leq \text{unit} \wedge X \times \text{list } X \leq X' \times \text{list } X'$$

holds. This assertion is equivalent to $\text{list } X \leq \text{list } X' \Vdash X \leq X'$. To satisfy the requirement, it is sufficient to make `list` a covariant type constructor, that is, to define subtyping in the model so that $\text{list } X \leq \text{list } X' \equiv X \leq X'$ holds.

Let *tree* be a data type of signature $\star \Rightarrow \star$. Let *root* and *sons* be record labels. Then, the following is a definition of *tree* as a record type:

$$\text{tree } X \approx \Pi (\text{root} : X; \text{sons} : \text{list} (\text{tree } X))$$

This definition is again recursive, and relies on the previous definition. Because *list* is covariant, it is straightforward to check that the definition of *tree* is valid if *tree* is made a covariant type constructor as well. \square

A *prologue* is a set of algebraic data type definitions, where each data type is defined at most once and where each data constructor or record label appears at most once. A *program* is a pair of a prologue and an expression. The effect of a prologue is to enrich the programming language with new constants. That is, a variant type definition extends the operational semantics with several injections and a case construct, as in Example 10.1.7. A record type definition extends it with a record formation construct and several projections, as in Examples 10.1.3 and 10.1.5. In either case, the initial typing environment Γ_0 is extended with information about these new constants. Thus, algebraic data type definitions might be viewed as a simple configuration language that allows specifying in which instance of ML-the-calculus the expression that follows the prologue should be typechecked and interpreted. Let us now give a precise account of this phenomenon.

To begin, suppose the prologue contains the definition $D\vec{X} \approx \sum_{i=1}^k \ell_i : T_i$. Then, for each $i \in \{1, \dots, k\}$, a constructor of arity 1, named ℓ_i , is introduced. Furthermore, a destructor of arity $k + 1$, named case_D , is introduced. When $k > 0$, it is common to write $\text{case } \mathfrak{t} [\ell_i : \mathfrak{t}_i]_{i=1}^k$ for the application $\text{case}_D \mathfrak{t} \mathfrak{t}_1 \dots \mathfrak{t}_n$. The operational semantics is extended with the following reduction rules, for $i \in \{1, \dots, k\}$:

$$\text{case } (\ell_i \mathfrak{v}) [\ell_j : \mathfrak{v}_j]_{j=1}^k \xrightarrow{\delta} \mathfrak{v}_i \mathfrak{v} \quad (\text{R-ALG-CASE})$$

For each $i \in \{1, \dots, k\}$, the initial environment is extended with the binding $\ell_i : \forall \vec{X}. T_i \rightarrow D\vec{X}$. It is further extended with the binding $\text{case}_D : \forall \vec{X}. D\vec{X} \rightarrow (T_1 \rightarrow Z) \rightarrow \dots \rightarrow (T_k \rightarrow Z) \rightarrow Z$.

Now, suppose the prologue contains the definition $D\vec{X} \approx \prod_{i=1}^k \ell_i : T_i$. Then, for each $i \in \{1, \dots, k\}$, a destructor of arity 1, named ℓ_i , is introduced. Furthermore, a constructor of arity k , named make_D , is introduced. It is common to write $\mathfrak{t}.\ell$ for the application $\ell \mathfrak{t}$ and, when $k > 0$, to write $\{\ell_i = \mathfrak{t}_i\}_{i=1}^k$ for the application $\text{make}_D \mathfrak{t}_1 \dots \mathfrak{t}_k$. The operational semantics is extended with the following reduction rules, for $i \in \{1, \dots, k\}$:

$$(\{\ell_j = \mathfrak{v}_j\}_{j=1}^k).\ell_i \xrightarrow{\delta} \mathfrak{v}_i \quad (\text{R-ALG-PROJ})$$

For each $i \in \{1, \dots, k\}$, the initial environment is extended with the binding $\ell_i : \forall \vec{X}. D\vec{X} \rightarrow T_i$. It is further extended with the binding $\text{make}_D : \forall \vec{X}. T_1 \rightarrow \dots \rightarrow T_k \rightarrow D\vec{X}$.

10.7.10 **EXAMPLE:** The effect of defining list (Example 10.7.9) is to make Nil and Cons data constructors of arity 1 and to introduce a binary destructor $\text{case}_{\text{list}}$. The definition also extends the initial environment as follows:

$$\begin{aligned} \text{Nil} &: \forall X. \text{unit} \rightarrow \text{list } X \\ \text{Cons} &: \forall X. X \times \text{list } X \rightarrow \text{list } X \\ \text{case}_{\text{list}} &: \forall XZ. \text{list } X \rightarrow (\text{unit} \rightarrow Z) \rightarrow (X \times \text{list } X \rightarrow Z) \rightarrow Z \end{aligned}$$

Thus, the value $\text{Cons}(\hat{0}, \text{Nil}())$, an integer list of length 1, has type list int . A function that computes the length of a list may now be written as follows:

$$\text{letrec length} = \lambda l. \text{case } l \text{ [Nil : } \lambda _ . \hat{0} \mid \text{Cons : } \lambda z. \hat{1} \hat{+} \text{length } (\pi_2 z)]$$

Recall that this notation is syntactic sugar for

$$\text{letrec length} = \lambda l. \text{case}_{\text{list}} l (\lambda _ . \hat{0}) (\lambda z. \hat{1} \hat{+} \text{length } (\pi_2 z))$$

The difference with the code in Example 10.7.7 appears minimal: the case construct is now annotated with the data type list . As a result, the type inference algorithm employs the type scheme assigned to $\text{case}_{\text{list}}$, which is derived from the definition of list , instead of the type scheme assigned to the anonymous case construct, given in Exercise 10.7.4. This is good for a couple of reasons. First, the former is more informative than the latter, because it contains the type T_i associated with the data constructor ℓ_i . Here, for instance, the generated constraint requires the type of z to be $X \times \text{list } X$ for some X , so a good error message would be given if a mistake was made in the second branch, such as omitting the use of π_2 . Second, and more fundamentally, the code is now well-typed, even in the absence of recursive types. In Example 10.7.7, a cyclic equation was produced because case required the type of l to be a sum type and because a sum type carries the types of its left and right branches as subterms. Here, $\text{case}_{\text{list}}$ requires l to have type $\text{list } X$ for some X . This is an abstract type: it does not explicitly contain the types of the branches. As a result, the generated constraint no longer involves a cyclic equation. It is, in fact, satisfiable; the reader may check that length has type $\forall X. \text{list } X \rightarrow \text{int}$, as expected. \square

Example 10.7.10 stresses the importance of using declared, abstract types, as opposed to anonymous, concrete sum or product types, in order to obviate the need for recursive types. The essence of the trick lies in the fact that the type schemes associated with operations on algebraic data types implicitly fold and unfold the data type's definition. More precisely, let us recall the type scheme assigned to the i^{th} injection in the setting of (k -ary) anonymous sums: it is $\forall X_1 \dots X_k. X_i \rightarrow X_1 + \dots + X_k$, or, more concisely, $\forall X_1 \dots X_k. X_i \rightarrow \sum_{i=1}^k X_i$.

By instantiating each X_i with T_i and generalizing again, we find that a more specific type scheme is $\forall \vec{X}. T_i \rightarrow \sum_{i=1}^k T_i$. Perhaps this could have been the type scheme assigned to ℓ_i ? Instead, however, it is $\forall \vec{X}. T_i \rightarrow D\vec{X}$. We now realize that the latter type scheme not only reflects the operational behavior of the i^{th} injection but also folds the definition of the algebraic data type D by turning the anonymous sum $\sum_{i=1}^k T_i$ —which forms the definition’s right-hand side—into the parameterized abstract type $D\vec{X}$ —which is the definition’s left-hand side. Conversely, the type scheme assigned to case_D unfolds the definition. The situation is identical in the case of record types: in either case, constructors fold, destructors unfold. In other words, occurrences of data constructors and record labels in the code may be viewed as explicit instructions for the typechecker to fold or unfold an algebraic data type definition. This mechanism is characteristic of isorecursive types.

- 10.7.11 EXERCISE [\star , \rightarrow]: For a fixed k , check that all of the machinery associated with k -ary anonymous products—that is, constructors, destructors, reduction rules, and extensions to the initial typing environment—may be viewed as the result of a single algebraic data type definition. Conduct a similar check in the case of k -ary anonymous sums. \square
- 10.7.12 EXERCISE [$\star\star\star$, \rightarrow]: Check that the above definitions meet the requirements of Definition 10.5.5. \square
- 10.7.13 EXERCISE [$\star\star\star$, \rightarrow]: For the sake of simplicity, we have assumed that all data constructors have arity one. If desired, it is possible to accept variant data type definitions of the form $D\vec{X} \approx \sum_{i=1}^k \ell_i : \vec{T}_i$, where the arity of the data constructor ℓ_i is the length of the vector \vec{T}_i , and may be an arbitrary nonnegative integer. This allows, for instance, altering the definition of `list` so that the data constructors `Nil` and `Cons` are respectively nullary and binary. Make the necessary changes in the above definitions and check that the requirements of Definition 10.5.5 are still met. \square

One significant drawback of algebraic data type definitions resides in the fact that a label ℓ cannot be *shared* by two distinct variant or record type definitions. Indeed, every algebraic data type definition extends the calculus with new constants. Strictly speaking, our presentation does not allow a single constant c to be associated with two distinct definitions. Even if we did allow such a collision, the initial environment would contain two bindings for c , one of which would then hide the other. This phenomenon arises in actual implementations of ML—the-programming-language, where a new algebraic data type definition may hide some of the data constructors or record labels introduced by a previous definition. An elegant solution to this lack of expressiveness is discussed in §10.8.

Recursive Types

We have shown that specializing $HM(X)$ with an equality-only syntactic model yields $HM(=)$, a constraint-based formulation of Damas and Milner's type system. Similarly, it is possible to specialize $HM(X)$ with an equality-only free regular tree model, yielding a constraint-based type system that may be viewed as an extension of Damas and Milner's type discipline with recursive types. This flavor of recursive types is sometimes known as *equirecursive*, since cyclic *equations*, such as $X = X \rightarrow X$, are then satisfiable. Our theorems about type inference and type soundness, which are independent of the model, remain valid. The constraint solver described in §10.6 may be used in the setting of an equality-only free regular tree model; the only difference with the syntactic case is that the occurs check is no longer performed.

Note that, although ground types are regular, types remain finite objects: their syntax is unchanged. The μ notation commonly employed to describe recursive types may be emulated using type equations: for instance, the notation $\mu X.X \rightarrow X$ corresponds, in our constraint-based approach, to the type scheme $\forall X[X = X \rightarrow X].X$.

Although recursive types come for free, as explained above, they have not been adopted in mainstream programming languages based on ML-the-type-system. The reason is pragmatic: experience shows that many nonsensical expressions are well-typed in the presence of recursive types, whereas they are not in their absence. Thus, the gain in expressiveness is offset by the fact that many programming mistakes are detected later than otherwise possible. Consider, for instance, the following OCaml session:

```
ocaml -rectypes
# let rec map f = function
  | [] -> []
  | x :: l -> (map f x) :: (map f l);;
val map : 'a -> ('b list as 'b) -> ('c list as 'c) = <fun>
```

This nonsensical version of `map` is essentially useless, yet well-typed. Its principal type scheme, in our notation, is $\forall XYZ[Y = \text{list } Y \wedge Z = \text{list } Z].X \rightarrow Y \rightarrow Z$. In the absence of recursive types, it is ill-typed, since the constraint $Y = \text{list } Y \wedge Z = \text{list } Z$ is then false.

The need for equirecursive types is usually suppressed by the presence of algebraic data types, which offer isorecursive types, in the language. Yet, they are still necessary in some situations, such as in Objective Caml's extensions with objects (Rémy and Vouillon, 1998) or polymorphic variants (Garrigue, 1998, 2000, 2002), where recursive object or variant types are commonly inferred. In order to allow recursive object or variant types while still rejecting the above version of `map`, Objective Caml's constraint solver implements a

selective occurs check, which forbids cycles unless they involve the type constructors $\langle \cdot \rangle$ or $[\cdot]$ respectively associated with objects and variants. The corresponding model is a tree model where every infinite path down a tree must encounter the type constructor $\langle \cdot \rangle$ or $[\cdot]$ infinitely often.

10.8 Rows

In §10.7, we have shown how to extend ML-the-programming-language with algebraic data types, that is, variant and record type definitions, which we now refer to as *simple*. This mechanism has a severe limitation: two distinct definitions must define incompatible types. As a result, one cannot hope to write code that uniformly operates over variants or records of different shapes, because the type of such code is not even expressible.

For instance, it is impossible to express the type of the *polymorphic record access* operation, which retrieves the value stored at a particular field ℓ inside a record, regardless of which other fields are present. Indeed, if the label ℓ appears with type T in the definition of the simple record type $D\vec{X}$, then the associated record access operation has type $\forall \vec{X}. D\vec{X} \rightarrow T$. If ℓ appears with type T' in the definition of another simple record type, say $D'\vec{X}'$, then the associated record access operation has type $\forall \vec{X}'. D'\vec{X}' \rightarrow T'$; and so on. The most precise type scheme that subsumes all of these incomparable type schemes is $\forall XY. X \rightarrow Y$. It is, however, not a sound type scheme for the record access operation. Another powerful operation whose type is currently not expressible is *polymorphic record extension*, which copies a record and stores a value at field ℓ in the copy, possibly creating the field if it did not previously exist, again regardless of which other fields are present. (If ℓ was known to previously exist, the operation is known as *polymorphic record update*.)

In order to assign types to polymorphic record operations, we must do away with record type definitions: we must replace named record types, such as $D\vec{X}$, with structural record types that provide a direct description of the record's domain and contents. (Following the analogy between a record and a partial function from labels to values, we use the word *domain* to refer to the set of fields that are defined in a record.) For instance, a product type is structural: the type $T_1 \times T_2$ is the (undeclared) type of pairs whose first component has type T_1 and whose second component has type T_2 . Thus, we wish to design record types that behave very much like product types. In doing so, we face two orthogonal difficulties. First, as opposed to pairs, records may have different domains. Because the type system must statically ensure that no undefined field is accessed, information about a record's domain must be made part of its type. Second, because we suppress record type definitions,

labels must now be predefined. However, for efficiency and modularity reasons, it is impossible to explicitly list every label in existence in every record type.

In what follows, we explain how to address the first difficulty in the simple setting of a finite set of labels. Then we introduce *rows*, which allow dealing with an infinite set of labels, and address the second difficulty. We define the syntax and logical interpretation of rows, study the new constraint equivalence laws that arise in their presence, and extend the first-order unification algorithm with support for rows. Then we review several applications of rows, including polymorphic operations on records, variants, and objects, and discuss alternatives to rows.

Because our interest is in typechecking and type inference issues, we do not address the compilation issue: how does one efficiently compile polymorphic records or polymorphic variants? A few relevant papers are Pugh and Weddell (1990), Ohori (1995), and Garrigue (1998). The problem of optimizing message dispatch in object-oriented languages, which has received considerable attention in the literature, is related.

Records with Finite Carrier

Let us temporarily assume that \mathcal{L} is finite. In fact, for the sake of definiteness, let us assume that \mathcal{L} is the three-element set $\{\ell_a, \ell_b, \ell_c\}$.

To begin, let us consider only *full* records, whose domain is exactly \mathcal{L} —in other words, tuples indexed by \mathcal{L} . To describe them, it is natural to introduce a type constructor Π of signature $\star \otimes \star \otimes \star \Rightarrow \star$. The type $\Pi \tau_a \tau_b \tau_c$ represents all records where the field ℓ_a (respectively ℓ_b, ℓ_c) contains a value of type τ_a (respectively τ_b, τ_c). Note that Π is nothing but a product type constructor of arity 3. The basic operations on records, namely *creation* of a record out of a default value, which is stored into every field, *update* of a particular field (say, ℓ_b), and *access* to a particular field (say, ℓ_b), may be assigned the following type schemes:

$$\begin{aligned} \{\cdot\} &: \forall X.X \rightarrow \Pi X X X \\ \{\cdot \text{ with } \ell_b = \cdot\} &: \forall X_a X_b X'_b X_c. \Pi X_a X_b X_c \rightarrow X'_b \rightarrow \Pi X_a X'_b X_c \\ \cdot.\{\ell_b\} &: \forall X_a X_b X_c. \Pi X_a X_b X_c \rightarrow X_b \end{aligned}$$

Here, polymorphism allows updating or accessing a field without knowledge of the types of the other fields. This flexibility stems from the key property that all record types are formed using a single Π type constructor.

This is fine, but in general, the domain of a record is not necessarily \mathcal{L} : it may be a subset of \mathcal{L} . How may we deal with this fact while maintaining the above key property? A naive approach consists of encoding arbitrary records

in terms of full records, using the standard algebraic data type option, whose definition is $\text{option } X \approx \text{pre } X + \text{abs}$. We use *pre* for present and *abs* for absent: indeed, a field that is defined with value v is encoded as a field with value *pre* v , while an undefined field is encoded as a field with value *abs*. Thus, an arbitrary record whose fields, if present, have types T_a , T_b , and T_c , respectively, may be encoded as a full record of type Π (*option* T_a) (*option* T_b) (*option* T_c). This naive approach suffers from a serious drawback: record types still contain no domain information. As a result, field access must involve a dynamic check, so as to determine whether the desired field is present; in our encoding, this corresponds to the use of `caseoption`.

To avoid this overhead and increase programming safety, we must move this check from runtime to compile time. In other words, we must make the type system aware of the difference between *pre* and *abs*. To do so, we replace the definition of *option* by two separate algebraic data type definitions, namely $\text{pre } X \approx \text{pre } X$ and $\text{abs} \approx \text{abs}$. In other words, we introduce a unary type constructor *pre*, whose only associated data constructor is *pre*, and a nullary type constructor *abs*, whose only associated data constructor is *abs*. Record types now contain domain information; for instance, a record of type Π *abs* (*pre* T_b) (*pre* T_c) must have domain $\{\ell_b, \ell_c\}$. Thus, the type of a field tells whether it is defined. Since the type *pre* has no data constructors other than *pre*, the accessor pre^{-1} , whose type is $\forall X. \text{pre } X \rightarrow X$, and which allows retrieving the value stored in a field, cannot fail. Thus, the dynamic check has been eliminated.

To complete the definition of our encoding, we now define operations on arbitrary records in terms of operations on full records. To distinguish between the two, we write the former with angle braces, instead of curly braces. The *empty record* $\langle \rangle$, where all fields are undefined, may be defined as $\{\text{abs}\}$. *Extension* at a particular field (say, ℓ_b) $\langle \cdot \text{ with } \ell_b = \cdot \rangle$ is defined as $\lambda r. \lambda z. \{r \text{ with } \ell_b = \text{pre } z\}$. *Access* at a particular field (say, ℓ_b) $\cdot.\langle \ell_b \rangle$ is defined as $\lambda z. \text{pre}^{-1} z. \{\ell_b\}$. It is straightforward to check that these operations have the following principal type schemes:

$$\begin{aligned} \langle \rangle &: \Pi \text{abs abs abs} \\ \langle \cdot \text{ with } \ell_b = \cdot \rangle &: \forall X_a X_b X'_b X_c. \Pi X_a X_b X_c \rightarrow X'_b \rightarrow \Pi X_a (\text{pre } X'_b) X_c \\ \cdot.\langle \ell_b \rangle &: \forall X_a X_b X_c. \Pi X_a (\text{pre } X_b) X_c \rightarrow X_b \end{aligned}$$

It is important to notice that the type schemes associated with extension and access at ℓ_b are polymorphic in X_a and X_c , which now means that these operations are insensitive, not only to the type, but also to the presence or absence of the fields ℓ_a and ℓ_c . Furthermore, extension is polymorphic in X_b , which means that it is insensitive to the presence or absence of the field ℓ_b in its argument. The subterm $\text{pre } X'_b$ in its result type reflects the fact that

ℓ_b is defined in the extended record. Conversely, the subterm $\text{pre } X_b$ in the type of the access operation reflects the requirement that ℓ_b be defined in its argument.

Our encoding of arbitrary records in terms of full records was carried out for pedagogical purposes. In practice, no such encoding is necessary: the *data* constructors pre and abs have no machine representation, and the compiler is free to lay out records in memory in an efficient manner. The encoding is interesting, however, because it provides a natural way of introducing the *type* constructors pre and abs , which play an important role in our treatment of polymorphic record operations.

Once we forget about the encoding, the arguments of the type constructor Π are expected to be either type variables or formed with pre or abs , while, conversely, the type constructors pre and abs are not intended to appear anywhere else. It is possible to enforce this invariant using kinds. In addition to \star , let us introduce the kind \circ of *field types*. Then, let us adopt the following signatures: $\text{pre} : \star \Rightarrow \circ$, $\text{abs} : \circ$, and $\Pi : \circ \otimes \circ \otimes \circ \Rightarrow \star$.

- 10.8.1 EXERCISE [RECOMMENDED, \star , \rightarrow]: Check that the three type schemes given above are well-kinded. What is the kind of each type variable? \square
- 10.8.2 EXERCISE [RECOMMENDED, $\star\star$]: Our Π types contain information about every field, regardless of whether it is defined: we encode definedness information within the type of each field, using the type constructors pre and abs . A perhaps more natural approach would be to introduce a family of record type constructors, indexed by the subsets of \mathcal{L} , so that the types of records with different domains are formed with different constructors. For instance, the empty record would have type $\{\}$; a record that defines the field ℓ_a only would have a type of the form $\{\ell_a : \tau_a\}$; a record that defines the fields ℓ_b and ℓ_c only would have a type of the form $\{\ell_b : \tau_b; \ell_c : \tau_c\}$; and so on. Assuming that the type discipline is Damas and Milner's (that is, assuming an equality-only syntactic model), would it be possible to assign satisfactory type schemes to polymorphic record access and extension? Would it help to equip record types with a nontrivial subtyping relation? \square

Records with Infinite Carrier

The treatment of records described above is not quite satisfactory, from practical and theoretical points of view. First, in practice, the set \mathcal{L} of all record labels that appear within a program could be very large. Because every record type is just as large as \mathcal{L} itself, even if the record that it describes only has a few fields, this is unpleasant. Furthermore, in a modular setting, the set of all record labels that appear within a program cannot be determined until link

time, so it is still unknown at compile time, when each compilation unit is separately typechecked. As a result, it may only be assumed to be a subset of the infinite set of all syntactically valid record labels. Resolving these issues requires coming up with a treatment of records that does not become more costly as \mathcal{L} grows and that, in fact, allows \mathcal{L} to be infinite. Thus, from here on, let us assume that \mathcal{L} is infinite.

As in the previous section, we first concentrate on full records, whose domain is exactly \mathcal{L} . The case of arbitrary records, whose domain is a subset of \mathcal{L} , will then follow in the same manner, by using the type constructors `pre` and `abs` to encode domain information.

Of course, even though we have assumed that \mathcal{L} is infinite, we must ensure that every record has a finite representation. We choose to restrict our attention to records that are *almost constant*, that is, records where all fields but a finite number contain the same value. Every such record may be defined in terms of two primitive operations, namely (i) *creation* of a constant record out of a value; for instance, `{false}` is the record where every field contains the value `false`; and (ii) *update* of a record at a particular field; for instance, `{{false} with ℓ = 1}` carries the value 1 at field ℓ and the value `false` at every other field. As usual, an *access* operation allows retrieving the contents of a field. Thus, the three primitive operations are the same as in the previous subsection, only in the setting of an infinite number of fields.

If we were to continue as before, we would now introduce a type constructor Π , equipped with an infinite family of type parameters. Because types must remain finite objects, we cannot do so. Instead, we must find a finite (and economical) representation of such an infinite family of types. This is precisely the role played by *rows*.

A row is a type that denotes a function from labels to types or, equivalently, a family of types indexed by labels. Its domain is \mathcal{L} —the row is then *complete*—or a cofinite subset of \mathcal{L} —the row is then *incomplete*. (A subset of \mathcal{L} is cofinite if and only if its complement is finite. Incomplete rows are used only as building blocks for complete rows.) Because rows must admit a finite representation, we build them out of two syntactic constructions, namely (i) construction of a *constant row* out of a type; for instance, the notation `∂bool` denotes a row that maps every label in its domain to `bool`; and (ii) strict *extension* of an incomplete row; for instance, `(ℓ : int ; ∂bool)` denotes a row that maps ℓ to `int` and every other field in its domain to `bool`. Formally, `∂` is a unary type constructor, while, for every label ℓ , `(ℓ : · ; ·)` is a binary type constructor. These two constructions are reminiscent of the two operations used above to build records. There are, however, a couple of subtle but important differences. First, `∂T` may be a complete or incomplete row. Second, `(ℓ : T ; T')` is defined only if ℓ is not in the domain of the row T' , so this

construction is strict extension, not update. These aspects are made clear by a *kinding discipline*, to be introduced later on.

It is possible for two syntactically distinct rows to denote the same function from labels to types. For instance, according to the intuitive interpretation of rows given above, the three complete rows $(\ell : \text{int} ; \partial\text{bool})$, $(\ell : \text{int} ; (\ell' : \text{bool} ; \partial\text{bool}))$, and $(\ell' : \text{bool} ; (\ell : \text{int} ; \partial\text{bool}))$ denote the same total function from labels to types. In the following, we define the logical interpretation of types in such a way that the interpretations of these three rows in the model are indeed equal.

We may now make the record type constructor Π a unary type constructor, whose parameter is a row. Then, (say) $\Pi (\ell : \text{int} ; \partial\text{bool})$ is a record type, and we intend it to be a valid type for the record $\{\{\text{false}\} \text{ with } \ell = 1\}$. The basic operations on records may be assigned the following type schemes:

$$\begin{aligned} \{\cdot\} &: \forall X. X \rightarrow \Pi (\partial X) \\ \{\cdot \text{ with } \ell = \cdot\} &: \forall XX'Y. \Pi (\ell : X ; Y) \rightarrow X' \rightarrow \Pi (\ell : X' ; Y) \\ \cdot.\{\ell\} &: \forall XY. \Pi (\ell : X ; Y) \rightarrow X \end{aligned}$$

These type schemes are reminiscent of those given above. However, in the previous section, the size of the type schemes was linear in the cardinal of \mathcal{L} , whereas here it is constant, even though \mathcal{L} is infinite. This is made possible by the fact that record types no longer list all labels in existence; instead, they use rows. In the type scheme assigned to record creation, the constant row ∂X is used to indicate that all fields have the same type in the newly created record. In the next two type schemes, the row $(\ell : X_\ell ; X)$ is used to separate the type X_ℓ , which describes the contents of the field ℓ , and the row X , which collectively describes the contents of all other fields. Here, the type variable X stands for an arbitrary row; it is often referred to as a *row variable*. The ability of quantifying over row and type variables alike confers great expressiveness to the type system.

We have explained, in an informal manner, how rows allow typechecking operations on full records, in the setting of an infinite set of labels. We return to this issue in Example 10.8.25. To deal with the case of arbitrary records, whose domain is finite, we rely on the field type constructors `pre` and `abs`, as explained previously. We return to this point in Example 10.8.30. In the following, we give a formal exposition of rows. We begin with their syntax and logical interpretation. Then we give some new constraint equivalence laws, which characterize rows, and allow extending our first-order unification algorithm with support for rows. We conclude with several illustrations of the use of rows and some pointers to related work.

Syntax of Rows

In the following, the set of labels \mathcal{L} is considered denumerable. We let L range over finite subsets of \mathcal{L} . When $\ell \notin L$ holds, we write $\ell.L$ for $\{\ell\} \uplus L$. Before explaining how the syntax of types is enriched with rows, we introduce *row kinds*, whose grammar is as follows:

$$s ::= \text{Type} \mid \text{Row}(L)$$

Row kinds help distinguish between three kinds of types, namely ordinary types, complete rows, and incomplete rows. While ordinary types are used to describe expressions, complete or incomplete rows are used only as building blocks for ordinary types. For instance, the record type $\Pi (\ell : \text{int} ; \partial\text{bool})$, which was informally introduced above, is intended to be an ordinary type, that is, a type of row kind *Type*. Its subterm $(\ell : \text{int} ; \partial\text{bool})$ is a complete row, that is, a type of row kind $\text{Row}(\emptyset)$. Its subterm ∂bool is an incomplete row, whose row kind is $\text{Row}(\{\ell\})$. Intuitively, a row of kind $\text{Row}(L)$ denotes a family of types whose domain is $\mathcal{L} \setminus L$. In other words, L is the set of labels that the row does not define. The purpose of row kinds is to outlaw meaningless types, such as $\Pi (\text{int})$, which makes no sense because the argument to the record type constructor Π should be a (complete) row, or $(\ell : \tau_1 ; \ell : \tau_2 ; \partial\text{bool})$, which makes no sense because no label may occur twice within a row.

Let us now define the syntax of types in the presence of rows. As usual, it is given by a signature S (Definition 10.1.14), which lists all type constructors together with their signatures. Here, for the sake of generality, we do not wish to give a fixed signature S . Instead, we give a procedure that builds S out of two simpler signatures, referred to as S_0 and S_1 . The input signature S_0 lists the type constructors that have nothing to do with rows, such as \rightarrow , \times , int , etc. The input signature S_1 lists the type constructors that allow a row to be a subterm of an ordinary type, such as the record type constructor Π . In a type system equipped with extensible variant types or with object types, there might be several such type constructors; see the sections on Polymorphic Variants (p. 483) and Other Applications of Rows (p. 486). Without loss of generality, we assume that all type constructors in S_1 are unary. The point of parameterizing the definition of S over S_0 and S_1 is to make the construction more general: instead of defining a fixed type grammar featuring rows, we wish to explain how to enrich an arbitrary type grammar with rows.

In the following, we let G (respectively H) range over the type constructors in S_0 (respectively S_1). We let κ range over the kinds involved in the definition of S_0 and S_1 , and refer to them as *basic kinds*. We let F range over the type constructors in S . The kinds involved in the definition of S are *composite kinds*, that is, pairs of a basic kind κ and a row kind s , written $\kappa.s$.

This allows the kind discipline enforced by S to reflect that enforced by S_0 and S_1 and also to impose restrictions on the structure and use of rows, as suggested above. For the sake of conciseness, we write $K.s$ for the mapping $(d \mapsto K(d).s)^{d \in \text{dom}(K)}$ and $(K \Rightarrow \kappa).s$ for the (composite) kind signature $K.s \Rightarrow \kappa.s$. (In other words, we let $.s$ distribute over basic signatures.) We use symmetric notations to build a composite kind signature out of a basic kind and a row kind signature.

10.8.3 DEFINITION: The signature S is defined as follows:

$F \in \text{dom}(S)$	Signature	Conditions
G^s	$(K \Rightarrow \kappa).s$	$(G : K \Rightarrow \kappa) \in S_0$
H	$K.\text{Row}(\emptyset) \Rightarrow \kappa.\text{Type}$	$(H : K \Rightarrow \kappa) \in S_1$
$\partial^{\kappa,L}$	$\kappa.(\text{Type} \Rightarrow \text{Row}(L))$	
$\ell^{\kappa,L}$	$\kappa.(\text{Type} \otimes \text{Row}(\ell.L) \Rightarrow \text{Row}(L))$	$\ell \notin L$

We sometimes refer to S as the *row extension of S_0 with S_1* . □

Examples 10.8.7 and 10.8.8 suggest common choices of S_0 and S_1 and give a perhaps more concrete-looking definition of the grammar of types that they determine. First, however, let us explain the definition. The type constructors that populate S come in four varieties: they may be (i) taken from S_0 , (ii) taken from S_1 , (iii) a unary row constructor ∂ , or (iv) a binary row constructor $(\ell : \cdot ; \cdot)$. Let us review and explain each case.

Let us first consider case (i) and assume, for the time being, that s is *Type*. Then, for every type constructor G in S_0 , there is a corresponding type constructor G^{Type} in S . For instance, S_0 must contain an arrow type constructor \rightarrow , whose signature is $\{\text{domain} \mapsto \star, \text{codomain} \mapsto \star\} \Rightarrow \star$. Then, S contains a type constructor $\rightarrow^{\text{Type}}$, whose signature is $\{\text{domain} \mapsto \star.\text{Type}, \text{codomain} \mapsto \star.\text{Type}\} \Rightarrow \star.\text{Type}$. Thus, $\rightarrow^{\text{Type}}$ is a binary type constructor whose parameters and result must have basic kind \star and must have row kind *Type*; in other words, they must be ordinary types, as opposed to complete or incomplete rows. The family of all type constructors of the form G^{Type} , where G ranges over S_0 , forms a copy of S_0 at row kind *Type*: one might say, roughly speaking, that S contains S_0 . This is not surprising, since our purpose is to enrich the existing signature S_0 with syntax for rows.

Perhaps more surprising is the existence of the type constructor G^s , for every G in S_0 , and for every row kind s . For instance, for every L , S contains a type constructor $\rightarrow^{\text{Row}(L)}$, whose signature is $\{\text{domain} \mapsto \star.\text{Row}(L), \text{codomain} \mapsto \star.\text{Row}(L)\} \Rightarrow \star.\text{Row}(L)$. Thus, $\rightarrow^{\text{Row}(L)}$ is a binary type constructor whose parameters and result must have basic kind \star and must have row kind $\text{Row}(L)$. In other words, this type constructor maps a pair of rows that have a common domain to a row with the same domain. Recall that a row is to be interpreted

as a family of types. Our intention is that $\rightarrow^{Row(L)}$ maps two families of types to a family of arrow types. This is made precise in the next subsection. One should point out that the type constructors G^s , with $s \neq Type$, are required only in some advanced applications of rows; Examples 10.8.28 and 10.8.39 provide illustrations. They are not used when assigning types to the usual primitive operations on records, namely creation, update, and access (Examples 10.8.25 and 10.8.30).

Case (ii) is simple: it simply means that S contains S_1 . It is only worth noting that every type constructor H maps a parameter of row kind $Row(\emptyset)$ to a result of row kind $Type$, that is, a complete row to an ordinary type. Thanks to this design choice, the type $\Pi (int^{Type})$ is invalid: indeed, int^{Type} has row kind $Type$, while Π expects a parameter of row kind $Row(\emptyset)$.

Cases (iii) and (iv) introduce new type constructors that were not present in S_0 or S_1 and allow forming rows. They were informally described in the previous subsection. First, for every κ and L , there is a *constant row constructor* $\partial^{\kappa,L}$. Its parameter must have row kind $Type$, while its result has row kind $Row(L)$; in other words, this type constructor maps an ordinary type to a row. It is worth noting that the row thus built may be complete or incomplete; for instance, $\partial^{*,\emptyset} bool$ is a complete row, and may be used, for example, to build the type $\Pi (\partial^{*,\emptyset} bool)$, while $\partial^{*,\{\ell\}} bool$ is an incomplete row, and may be used, for example, to build the type $\Pi (\ell : int ; \partial^{*,\{\ell\}} bool)$. Second, for every κ , L , and $\ell \notin L$, there is a *row extension constructor* $\ell^{\kappa,L}$. We usually write $\ell^{\kappa,L} : T_1 ; T_2$ for $\ell^{\kappa,L} T_1 T_2$ and let this symbol be right associative so as to recover the familiar list notation for rows. According to the definition of S , if T_2 has row kind $Row(\ell.L)$, then $\ell^{\kappa,L} : T_1 ; T_2$ has row kind $Row(L)$. Thanks to this design choice, the type $(\ell^{*,L} : T_1 ; \ell^{*,L} : T_2 ; \partial^{*,\ell.L} bool)$ is invalid; indeed, the outer ℓ expects a parameter of row kind $Row(\ell.L)$, while the inner ℓ produces a type of row kind $Row(L)$.

The superscripts carried by the type constructors G , ℓ , and ∂ in the signature S make all kind information explicit, obviating the need for assigning several kinds to a single type constructor. In practice, however, we often drop the superscripts and use *unannotated* types. No ambiguity arises because, given a type expression T of known kind, it is possible to reconstruct all superscripts in a unique manner. This is the topic of the next example and exercises.

10.8.4 EXAMPLE [ILL-KINDED TYPES]: Assume that S_0 contains type constructors int and \rightarrow , whose signatures are respectively \star and $\star \otimes \star \Rightarrow \star$, and that S_1 contains a type constructor Π , whose signature is $\star \Rightarrow \star$.

The unannotated type $X \rightarrow \Pi(X)$ is invalid. Indeed, because Π 's image row kind is $Type$, the arrow must be \rightarrow^{Type} . Thus, the leftmost occurrence of X

must have row kind *Type*. On the other hand, because Π expects a parameter of row kind $Row(\emptyset)$, its rightmost occurrence must have row kind $Row(\emptyset)$ —a contradiction. The unannotated type $X \rightarrow \Pi(\partial X)$ is, however, valid, provided X has kind $\star.Type$. In fact, it is the type of the primitive record creation operation.

The unannotated type $(\ell : T ; \ell : T ; T')$ is also invalid: there is no way of reconstructing the missing superscripts so as to make it valid. Indeed, the row $(\ell : T' ; T')$ must have row kind $Row(L)$ for some L that does not contain ℓ . However, the context where it occurs requires it to also have row kind $Row(L)$ for some L that does contain ℓ . This makes it impossible to reconstruct consistent superscripts.

Any type of the form $\Pi(\Pi(T))$ is invalid, because the outer Π expects a parameter of row kind $Row(\emptyset)$, while the inner Π constructs a type of row kind *Type*. This is an intentional limitation: unlike those of S_0 , the type constructors of S_1 are not lifted to every row kind s . (If they were, we would be led to work not only with rows of ordinary types, but also with rows of rows, rows of rows of rows, and so on. Rémy (1990) explores this avenue.) \square

10.8.5 EXERCISE [RECOMMENDED, \star]: Consider the unannotated type

$$X \rightarrow \Pi(\ell : \text{int} ; (Y \rightarrow \partial X)).$$

Can you guess the kind of the type variables X and Y , as well as the missing superscripts, so as to ensure that this type has kind $\star.Type$? \square

10.8.6 EXERCISE [$\star\star\star$, \rightarrow]: Propose a kind checking algorithm that, given an unannotated type T , given the kind of T , and given the kind of all type variables that appear within T , ensures that T is well-kinded, and reconstructs the missing superscripts within T . Next, propose a kind inference algorithm that, given an unannotated type T , discovers the kind of T and the kind of all type variables that appear within T so as to ensure that T is well-kinded. \square

We have given a very general definition of the syntax of types. In this view, types, ranged over by the meta-variable T , encompass both “ordinary” types and rows: the distinction between the two is established only via the kind system. In the literature, however, it is common to establish this distinction by letting distinct meta-variables, say T and R , range over ordinary types and rows, respectively, so as to give the syntax a more concrete aspect. The next two examples illustrate this style and suggest common choices for S_0 and S_1 .

10.8.7 EXAMPLE: Assume that there is a single basic kind \star , that S_0 consists of the arrow type constructor \rightarrow , whose signature is $\star \otimes \star \Rightarrow \star$, and that S_1 consists

of the record type constructor Π , whose signature is $\star \Rightarrow \star$. Then, the composite kinds are $\star.Type$ and $\star.Row(L)$, where L ranges over the finite subsets of \mathcal{L} . Let us employ T (respectively R) to range over types of the former (respectively latter) kind, and refer to them as ordinary types (respectively rows). Then, the syntax of types, as defined by the signature S , may be presented under the following form:

$$\begin{aligned} T & ::= X \mid T \rightarrow T \mid \Pi R \\ R & ::= X \mid R \rightarrow R \mid (\ell : T ; R) \mid \partial T \end{aligned}$$

Ordinary types T include ordinary type variables (that is, type variables of kind $\star.Type$), arrow types (where the type constructor \rightarrow is really \rightarrow^{Type}), and record types, which are formed by applying the record type constructor Π to a row. Rows R include row variables (that is, type variables of kind $\star.Row(L)$ for some L), arrow rows (where the row constructor \rightarrow is really $\rightarrow^{Row(L)}$ for some L), row extension (whereby a row R is extended with an ordinary type T at a certain label ℓ), and constant rows (formed out of an ordinary type T). It would be possible to also introduce a syntactic distinction between ordinary type variables and row variables, if desired.

Such a presentation is rather pleasant, because the syntactic segregation between ordinary types and rows makes the syntax less ambiguous. It does not allow getting rid of the kind system, however: (row) kinds are still necessary to keep track of the domain of every row. \square

10.8.8 EXAMPLE: Assume that there are two basic kinds \star and \circ , that S_0 consists of the type constructors \rightarrow , abs , and pre , whose respective signatures are $\star \otimes \star \Rightarrow \star$, \circ , and $\star \Rightarrow \circ$, and that S_1 consists of the record type constructor Π , whose signature is $\circ \Rightarrow \star$. Then, the composite kinds are $\star.Type$, $\star.Row(L)$, $\circ.Type$, and $\circ.Row(L)$, where L ranges over the finite subsets of \mathcal{L} . Let us employ T^\star , R^\star , T° , and R° , respectively, to range over types of these four kinds. Then, the syntax of types, as defined by the signature S , may be presented under the following form:

$$\begin{aligned} T^\star & ::= X \mid T^\star \rightarrow T^\star \mid \Pi R^\circ \\ R^\star & ::= X \mid R^\star \rightarrow R^\star \mid (\ell : T^\star ; R^\star) \mid \partial T^\star \\ T^\circ & ::= X \mid \text{abs} \mid \text{pre } T^\star \\ R^\circ & ::= X \mid \text{abs} \mid \text{pre } R^\star \mid (\ell : T^\circ ; R^\circ) \mid \partial T^\circ \end{aligned}$$

Ordinary types T^\star are as in the previous example, except the record type constructor Π must now be applied to a row of field types R° . Rows R^\star are unchanged. Field types T° include field type variables (that is, type variables of kind $\circ.Type$) and applications of the type constructors abs and pre (which are really abs^{Type} and pre^{Type}). Field rows R° include field row variables (that

is, type variables of kind $\circ.\text{Row}(L)$ for some L), applications of the row constructors abs and pre (which are really $\text{abs}^{\text{Row}(L)}$ and $\text{pre}^{\text{Row}(L)}$ for some L), row extension, and constant rows, where row components are field types T° .

In many basic applications of rows, $\text{abs}^{\text{Row}(L)}$ and $\text{pre}^{\text{Row}(L)}$ are never required: that is, they do not appear in the type schemes that populate the initial environment. (Applications where they are required appear in Pottier [2000].) In that case, they may be removed from the syntax. Then, the nonterminal R^* becomes unreachable from the nonterminal T^* , which is the grammar's natural entry point, so it may be removed as well. In that simplified setting, the syntax of types and rows becomes:

$$\begin{aligned} T^* &::= X \mid T^* \rightarrow T^* \mid \Pi R^\circ \\ T^\circ &::= X \mid \text{abs} \mid \text{pre } T^* \\ R^\circ &::= X \mid (\ell : T^\circ ; R^\circ) \mid \partial T^\circ \end{aligned}$$

This is the syntax found in some introductory accounts of rows (Rémy, 1989; Pottier, 2000). \square

Meaning of Rows

We now give meaning to the type grammar defined in the previous section by interpreting it within a model. We choose to define a regular tree model, but alternatives exist; see Remark 10.8.12 below. In this model, every type constructor whose image row kind is *Type* (that is, every type constructor of the form G^{Type} or H) is interpreted as itself, as in a free tree model. However, every application of a type constructor whose image row kind is $\text{Row}(L)$ for some L receives special treatment: it is interpreted as a family of types indexed by $\mathcal{L} \setminus L$, which we encode as an infinitely branching tree. To serve as the root label of this tree, we introduce, for every κ and for every L , a symbol L^κ , whose arity is $\mathcal{L} \setminus L$. More precisely,

10.8.9 DEFINITION: The model, which consists of a set $\mathcal{M}_{\kappa,s}$ for every κ and s , is the regular tree algebra that arises out the following signature:

Symbol	Signature	Conditions
G	$(K \Rightarrow \kappa).\text{Type}$	$(G : K \Rightarrow \kappa) \in S_0$
H	$K.\text{Row}(\emptyset) \Rightarrow \kappa.\text{Type}$	$(H : K \Rightarrow \kappa) \in S_1$
L^κ	$\kappa.(\text{Type}^{\mathcal{L} \setminus L} \Rightarrow \text{Row}(L))$	

The first two lines in this signature coincide with the definitions of G^{Type} and H in the signature S . Indeed, as stated above, we intend to interpret

these type constructors in a syntactic manner, so each of them must have a counterpart in the model. The third line introduces the symbols L^κ hinted at above.

According to this signature, if t is a ground type of kind $\kappa.Type$ (that is, an element of $\mathcal{M}_{\kappa.Type}$), then its head symbol $t(\epsilon)$ must be of the form G or H . If t is a ground type of kind $\kappa.Row(L)$, then its head symbol must be L^κ , and its immediate subtrees, which are indexed by $\mathcal{L} \setminus L$, are ground types of kind $\kappa.Type$; in other words, the ground row t is effectively a family of ordinary ground types indexed by $\mathcal{L} \setminus L$. Thus, our intuition that rows denote infinite families of types is made literally true.

We have defined the model; there remains to explain how types are mapped to elements of the model.

10.8.10 DEFINITION: The interpretation of the type constructors that populate S is defined as follows.

1. Let $(G : K \Rightarrow \kappa) \in S_0$. Then, G^{Type} is interpreted as the function that maps $T \in \mathcal{M}_{K.Type}$ to the ground type $t \in \mathcal{M}_{\kappa.Type}$ defined by $t(\epsilon) = G$ and $t/d = T(d)$ for every $d \in dom(K)$. This is a syntactic interpretation.
2. Let $(H : K \Rightarrow \kappa) \in S_1$. Then, H is interpreted as the function that maps $T \in \mathcal{M}_{K.Row(\emptyset)}$ to the ground type $t \in \mathcal{M}_{\kappa.Type}$ defined by $t(\epsilon) = H$ and $t/d = T(d)$ for every $d \in dom(K)$. (Because H is unary, there is exactly one such d .) This is also a syntactic interpretation.
3. Let $(G : K \Rightarrow \kappa) \in S_0$. Then, $G^{Row(L)}$ is interpreted as the function that maps $T \in \mathcal{M}_{K.Row(L)}$ to the ground type $t \in \mathcal{M}_{\kappa.Row(L)}$ defined by $t(\epsilon) = L^\kappa$ and $t(\ell) = G$ and $t/(\ell \cdot d) = T(d)/\ell$ for every $\ell \in \mathcal{L} \setminus L$ and $d \in dom(K)$. Thus, when applied to a family of rows, the type constructor $G^{Row(L)}$ produces a row where every component has head symbol G . This definition may sound quite technical; its effect is summed up in a simpler fashion by the equations C-Row-GD and C-Row-GL in the next section.
4. Interpret $\partial^{\kappa,L}$ as the function that maps $t_1 \in \mathcal{M}_{\kappa.Type}$ to the ground type $t \in \mathcal{M}_{\kappa.Row(L)}$ defined by $t(\epsilon) = L^\kappa$ and $t/\ell = t_1$ for every $\ell \in \mathcal{L} \setminus L$. Note that t/ℓ does not depend on ℓ : t is a constant ground row.
5. Let $\ell \notin L$. Then, $\ell^{\kappa,L}$ is interpreted as the function that maps $(t_1, t_2) \in \mathcal{M}_{\kappa.Type} \times \mathcal{M}_{\kappa.Row(\ell.L)}$ to the ground type $t \in \mathcal{M}_{\kappa.Row(L)}$ defined by $t(\epsilon) = L^\kappa$ and $t/\ell = t_1$ and $t/\ell' = t_2(\ell')$ for every $\ell' \in \mathcal{L} \setminus \ell.L$. This definition is precisely row extension; indeed, the ground row t maps ℓ to t_1 and coincides with the ground row t_2 at every other label ℓ' . \square

Defining a model and an interpretation allows our presentation of rows to fit within the formalism proposed earlier in this chapter (§10.2). It also provides a basis for the intuition that rows denote infinite families of types. From a formal point of view, the model and its interpretation allow proving several constraint equivalence laws concerning rows, which are given and discussed in the next subsection. Of course, it is also possible to accept these equivalence laws as axioms and give a purely syntactic account of rows without relying on a model; this is how rows were historically dealt with (Rémy, 1993).

- 10.8.11 **REMARK:** We have not defined the interpretation of the subtyping predicate, because much of the material that follows is independent of it. One common approach is to adopt a nonstructural definition of subtyping (Example 10.2.9), where every L^k is considered covariant in every direction, and where the variances and relative ordering of all other symbols (G and H) are chosen at will, subject to the restrictions associated with nonstructural subtyping and to the conditions necessary to ensure type soundness.

Recall that the arrow type constructor \rightarrow is contravariant in its domain and covariant in its codomain. The record type constructor Π is usually covariant. These properties are exploited in proofs of the subject reduction theorem. The type constructors \rightarrow and Π are usually incompatible. This property is exploited in proofs of the progress theorem. In the case of Example 10.8.7, because no type constructors other than \rightarrow and Π are present, these conditions imply that there is no sensible way of interpreting subtyping other than equality. In the case of Example 10.8.8, two sensible interpretations of subtyping exist: one is equality, while the other is the nonstructural subtyping order obtained by letting $\text{pre} \leq \text{abs}$. In the former interpretation, abs means “definitely absent,” while in the latter, it means “possibly absent.” \square

- 10.8.12 **REMARK:** The model proposed above is a regular tree model. Of course, it is possible to adopt a finite tree model instead. Furthermore, other interpretations of rows are possible: for instance, Fähndrich (1999) extends the set constraints formalism with rows. In his model, an ordinary type is interpreted as a set of values, while a row is interpreted as a set of functions from labels to values. While the definition of the model may vary, the key point is that the characteristic laws of rows, which we discuss next, hold in the model. \square

Reasoning with Rows

The interpretation presented in the previous section was designed to support the intuition that a row denotes an infinite family of types, indexed by labels, that the row constructor $\ell : \cdot ; \cdot$ denotes row extension, and that the row constructor ∂ denotes the creation of a constant row. From a formal point of

$(\ell_1 : \mathbb{T}_1 ; \ell_2 : \mathbb{T}_2 ; \mathbb{T}_3) = (\ell_2 : \mathbb{T}_2 ; \ell_1 : \mathbb{T}_1 ; \mathbb{T}_3)$	(C-Row-LL)
$\partial \mathbb{T} = (\ell : \mathbb{T} ; \partial \mathbb{T})$	(C-Row-DL)
$G \partial \mathbb{T}_1 \dots \partial \mathbb{T}_n = \partial (G \mathbb{T}_1 \dots \mathbb{T}_n)$	(C-Row-GD)
$G (\ell : \mathbb{T}_1 ; \mathbb{T}'_1) \dots (\ell : \mathbb{T}_n ; \mathbb{T}'_n) = (\ell : G \mathbb{T}_1 \dots \mathbb{T}_n ; G \mathbb{T}'_1 \dots \mathbb{T}'_n)$	(C-Row-GL)

Figure 10-12: Equational reasoning with rows

view, the definition of the model and interpretation may be exploited to establish some reasoning principles concerning rows. These principles take the form of equations between types (Figure 10-12) and constraint equivalence laws (Figure 10-13), which we now explain and prove.

- 10.8.13 **REMARK:** As stated earlier, we omit the superscripts of row constructors. We also omit the side conditions that concern the kind of the type variables (X) and type meta-variables (\mathbb{T}) involved. Thus, each equation in Figure 10-12 really stands for the (infinite) family of equations obtained by reconstructing the missing kind information in a consistent way. For instance, the second equation may be read $\partial^{\ell.L} \mathbb{T} = (\ell^{\kappa.L} : \mathbb{T} ; \partial^L \mathbb{T})$, where $\ell \notin L$ and \mathbb{T} has kind $\kappa.Type$. □
- 10.8.14 **EXERCISE [RECOMMENDED, ★, →]:** Reconstruct all of the missing kind information in the equations of Figure 10-12. □
- 10.8.15 **REMARK:** There is a slight catch with the unannotated version of the second equation in Figure 10-12: its left-hand side admits strictly more kinds than its right-hand side, because the former has row kind $Row(L)$ for every L , while the latter has row kind $Row(L)$ for every L such that $\ell \notin L$ holds. As a result, while replacing the unannotated term $(\ell : \mathbb{T} ; \partial \mathbb{T})$ with $\partial \mathbb{T}$ is always valid, the converse is not: replacing the unannotated term $\partial \mathbb{T}$ with $(\ell : \mathbb{T} ; \partial \mathbb{T})$ is valid only if it does not result in an ill-kinded term. □

The first equation in Figure 10-12 states that rows are equal up to commutation of labels. For the equation to be well-kinded, the labels ℓ_1 and ℓ_2 must be distinct. The equation holds under our interpretation because extension of a ground row at ℓ_1 and extension of a ground row at ℓ_2 commute. The second equation states that $\partial \mathbb{T}$ maps every label within its domain to \mathbb{T} , that is, $\partial^L \mathbb{T}$ maps every label $\ell \notin L$ to \mathbb{T} . This equation holds because $\partial \mathbb{T}$ is interpreted as a constant row. The last two equations deal with the relationship between the row constructors G and the ordinary type constructor G . Indeed, notice that their left-hand sides involve $G^{Row(L)}$ for some L , while their right-hand sides involve G^{Type} . Both equations state that it is equivalent to apply

$(\ell_1 : T_1 ; T'_1) = (\ell_2 : T_2 ; T'_2)$	$\equiv \exists X. (T'_1 = (\ell_2 : T_2 ; X) \wedge T'_2 = (\ell_1 : T_1 ; X))$ if $X \# \text{ftv}(T_1, T'_1, T_2, T'_2) \wedge \ell_1 \neq \ell_2$	(C-MUTATE-LL)
$\partial T = (\ell : T' ; T'')$	$\equiv T = T' \wedge \partial T = T''$	(C-MUTATE-DL)
$G T_1 \dots T_n = \partial T$	$\equiv \exists X_1 \dots X_n. (G X_1 \dots X_n = T \wedge \bigwedge_{i=1}^n (T_i = \partial X_i))$ if $X_1 \dots X_n \# \text{ftv}(T_1, \dots, T_n, T)$	(C-MUTATE-GD)
$G T_1 \dots T_n = (\ell : T ; T')$	$\equiv \exists X_1 \dots X_n, X'_1 \dots X'_n. (G X_1 \dots X_n = T \wedge$ $G X'_1 \dots X'_n = T' \wedge$ $\bigwedge_{i=1}^n (T_i = (\ell : X_i ; X'_i)))$ if $X_1 \dots X_n, X'_1 \dots X'_n \# \text{ftv}(T_1, \dots, T_n, T, T')$	(C-MUTATE-GL)

Figure 10-13: Constraint equivalence laws involving rows

$G^{Row(L)}$ at the level of rows or to apply G^{Type} at the level of types. Our interpretation of $G^{Row(L)}$ was designed to give rise to these equations; indeed, the application of $G^{Row(L)}$ to n ground rows (where n is the arity of G) is interpreted as a pointwise application of G^{Type} to the rows' components (item 3 of Definition 10.8.10). Their use is illustrated in Examples 10.8.28 and 10.8.39.

10.8.16 LEMMA: Each of the equations in Figure 10-12 is equivalent to true. \square

The four equations in Figure 10-12 show that two types with distinct head symbols may denote the same element of the model. In other words, in the presence of rows, the interpretation of types is no longer free: an equation of the form $T_1 = T_2$, where T_1 and T_2 have distinct head symbols, is not necessarily equivalent to false. In Figure 10-13, we give several constraint equivalence laws, known as *mutation* laws, that concern such “heterogeneous” equations, and, when viewed as rewriting rules, allow solving them. To each equation in Figure 10-12 corresponds a mutation law. The soundness of the mutation law, that is, the fact that its right-hand side entails its left-hand side, follows from the corresponding equation. The completeness of the mutation law, that is, the fact that its left-hand side entails its right-hand side, holds by design of the model.

10.8.17 EXERCISE [RECOMMENDED, \star , \rightarrow]: Reconstruct all of the missing kind information in the laws of Figure 10-13. \square

Let us now review the four mutation laws. For the sake of brevity, in the following informal explanation, we assume that a ground assignment ϕ that

satisfies the left-hand equation is fixed, and write “the ground type T ” for “the ground type $\phi(T)$.” C-MUTATE-LL concerns an equation between two rows, which are both given by extension but exhibit distinct head labels ℓ_1 and ℓ_2 . When this equation is satisfied, both of its members must denote the same ground row. Thus, the ground row T'_1 must map ℓ_2 to the ground type T_2 , while, symmetrically, the ground row T'_2 must map ℓ_1 to the ground type T_1 . This may be expressed by two equations of the form $T'_1 = (\ell_2 : T_2 ; \dots)$ and $T'_2 = (\ell_1 : T_1 ; \dots)$. Furthermore, because the ground rows T'_1 and T'_2 must agree on their common labels, the ellipses in these two equations must denote the same ground row. This is expressed by letting the two equations share a fresh, existentially quantified row variable X . C-MUTATE-DL concerns an equation between two rows, one of which is given as a constant row, the other of which is given by extension. Then, because the ground row ∂T maps every label to the ground type T , the ground type T' must coincide with the ground type T , while the ground row T'' must map every label in its domain to the ground type T . This is expressed by the equations $T = T'$ and $\partial T = T''$. C-MUTATE-GD and C-MUTATE-GL concern an equation between two rows, one of which is given as an application of a row constructor G , the other of which is given either as a constant row or by extension. Again, the laws exploit the fact that the ground row $G T_1 \dots T_n$ is obtained by applying the type constructor G , pointwise, to the ground rows T_1, \dots, T_n . If, as in C-MUTATE-GD, it coincides with the constant ground row ∂T , then every T_i must itself be a constant ground row, of the form ∂X_i , and T must coincide with $G X_1 \dots X_n$. C-MUTATE-GL is obtained in a similar manner.

10.8.18 LEMMA: Each of the equivalence laws in Figure 10-13 holds. □

Solving Equality Constraints in the Presence of Rows

We now extend the unification algorithm given in §10.6 with support for rows. The extended algorithm is intended to solve unification problems where the syntax and interpretation of types are as defined in the discussions above of the syntax (p. 466) and meaning (p. 471) of rows. Its specification consists of the original rewriting rules of Figure 10-10, minus S-CLASH, which is removed and replaced with the rules given in Figure 10-14. Indeed, S-CLASH is no longer valid in the presence of rows: not all distinct type constructors are incompatible.

The extended algorithm features four mutation rules, which are in direct correspondence with the mutation laws of Figure 10-13, as well as a weakened version of S-CLASH, dubbed S-CLASH', which applies when neither S-DECOMPOSE nor the mutation rules are applicable. (Let us point out that, in

$(\ell_1 : X_1 ; X'_1) = (\ell_2 : T_2 ; T'_2) = \epsilon$	$\rightarrow \exists X. (X'_1 = (\ell_2 : T_2 ; X) \wedge T'_2 = (\ell_1 : X_1 ; X))$	
	$\wedge (\ell_1 : X_1 ; X'_1) = \epsilon$	(S-MUTATE-LL)
	if $\ell_1 \neq \ell_2$	
$\partial X = (\ell : T ; T') = \epsilon$	$\rightarrow X = T \wedge \partial X = T' \wedge \partial X = \epsilon$	(S-MUTATE-DL)
$G T_1 \dots T_n = \partial X = \epsilon$	$\rightarrow \exists X_1 \dots X_n. (G X_1 \dots X_n = X \wedge \bigwedge_{i=1}^n (T_i = \partial X_i))$	
	$\wedge \partial X = \epsilon$	(S-MUTATE-GD)
$G T_1 \dots T_n = (\ell : X ; X') = \epsilon$	$\rightarrow \exists X_1 \dots X_n, X'_1 \dots X'_n. (G X_1 \dots X_n = X \wedge$	
	$G X'_1 \dots X'_n = X' \wedge$	
	$\bigwedge_{i=1}^n (T_i = (\ell : X_i ; X'_i)))$	
	$\wedge (\ell : X ; X') = \epsilon$	(S-MUTATE-GL)
$F \bar{T} = F' \bar{T}' = \epsilon$	$\rightarrow \text{false}$	(S-CLASH')
	if $F \neq F'$ and none of the four rules above applies	

Figure 10-14: Row unification (changes to Figure 10-10)

S-DECOMPOSE, the meta-variable F ranges over all type constructors in the signature S , so that S-DECOMPOSE is applicable to multi-equations of the form $\partial X = \partial T = \epsilon$ or $(\ell : X ; X') = (\ell : T ; T') = \epsilon$.) Three of the mutation rules may allocate fresh type variables, which must be chosen fresh for the rule's left-hand side. The four mutation rules paraphrase the four mutation laws very closely. Two minor differences are (i) the mutation rules deal with multi-equations, as opposed to equations; and (ii) any subterm that appears more than once on the right-hand side of a rule is required to be a type variable, as opposed to an arbitrary type. Neither of these features is specific to rows: both may be found in the definition of the standard unification algorithm (Figure 10-10), where they help reason about sharing.

- 10.8.19 EXERCISE [\star , \rightsquigarrow]: Check that the rewriting rules in Figure 10-14 preserve well-kindedness. Conclude that, provided its input constraint is well-kinded, the unification algorithm needs not keep track of kinds. \square

The properties of the unification algorithm are preserved by this extension, as witnessed by the next three lemmas. Note that the termination of reduction is ensured only when the initial unification problem is well-kinded. The ill-kinded unification problem $X = (\ell_1 : T ; Y) \wedge X = (\ell_2 : T ; Y)$, where ℓ_1 and ℓ_2 are distinct, illustrates this point.

- 10.8.20 LEMMA: The rewriting system \rightarrow is strongly normalizing. \square

- 10.8.21 LEMMA: $U_1 \rightarrow U_2$ implies $U_1 \equiv U_2$. □
- 10.8.22 LEMMA: Every normal form is either false or of the form $X[U]$, where X is an existential constraint context, U is a standard conjunction of multi-equations and, if the model is syntactic, U is acyclic. These conditions imply that U is satisfiable. □

The time complexity of standard first-order unification is quasi-linear. What is, then, the time complexity of row unification? Only a partial answer is known. In practice, the algorithm given in this chapter is extremely efficient and appears to behave just as well as standard unification. In theory, the complexity of row unification remains unexplored and forms an interesting open issue.

- 10.8.23 EXERCISE [***, →]: The unification algorithm presented above, although very efficient in practice, does *not* have linear or quasi-linear time complexity. Find a family of unification problems U_n such that the size of U_n is linear with respect to n and the number of steps required to reach its normal form is quadratic with respect to n . □
- 10.8.24 REMARK: Mutation is a common technique for solving equations in a large class of non-free algebras that are described by *syntactic theories* (Kirchner and Klay, 1990). The equations of Figure 10-12 happen to form a syntactic presentation of an equational theory. Thus, it is possible to derive a unification algorithm out of these equations in a systematic way (Rémy, 1993). Here, we have presented the same algorithm in a direct manner, without relying on the apparatus of syntactic theories. □

Operations on Records

We now illustrate the use of rows for typechecking operations on records. We begin with full records; our treatment follows Rémy (1992b).

- 10.8.25 EXAMPLE [FULL RECORDS]: As before, let us begin with full records, whose domain is exactly \mathcal{L} . The primitive operations are record creation $\{\cdot\}$, update $\{\cdot$ with $\ell = \cdot\}$, and access $\cdot.\{\ell\}$.

Let $<$ denote a fixed strict total order on row labels. For every set of labels L of cardinal n , let us introduce a $(n + 1)$ -ary constructor $\{\}_L$. We use the following syntactic sugar: we write $\{\ell_1 = \mathbf{t}_1; \dots; \ell_n = \mathbf{t}_n; \mathbf{t}\}$ for the application $\{\}_L \mathbf{t}_{i_1} \dots \mathbf{t}_{i_n} \mathbf{t}$, where $L = \{\ell_1, \dots, \ell_n\} = \{\ell_{i_1}, \dots, \ell_{i_n}\}$ and $\ell_{i_1} < \dots < \ell_{i_n}$ holds. The use of the total order $<$ makes the meaning of record expressions independent of the order in which fields are defined; in particular, it allows fixing the order in which $\mathbf{t}_1, \dots, \mathbf{t}_n$ are evaluated. We abbreviate the record

value $\{\ell_1 = v_1; \dots; \ell_n = v_n; v\}$ as $\{V; v\}$, where V is the finite function that maps ℓ_i to v_i for every $i \in \{1, \dots, n\}$.

The operational semantics of the above three operations may now be defined in the following straightforward manner. First, record creation $\{\cdot\}$ is precisely the unary constructor $\{\}_\emptyset$. Second, for every $\ell \in \mathcal{L}$, let update $\{\cdot \text{ with } \ell = \cdot\}$ and access $\cdot.\{\ell\}$ be destructors of arity 1 and 2, respectively, equipped with the following reduction rules:

$$\begin{array}{lll} \{\{V; v\} \text{ with } \ell = v'\} & \xrightarrow{\delta} & \{V[\ell \mapsto v']; v\} & \text{(R-UPDATE)} \\ \{V; v\}.\{\ell\} & \xrightarrow{\delta} & V(\ell) & (\ell \in \text{dom}(V)) \quad \text{(R-ACCESS-1)} \\ \{V; v\}.\{\ell\} & \xrightarrow{\delta} & v & (\ell \notin \text{dom}(V)) \quad \text{(R-ACCESS-2)} \end{array}$$

In these rules, $V[\ell \mapsto v]$ stands for the function that maps ℓ to v and coincides with V at every other label, while $V(\ell)$ stands for the image of ℓ through V . Because these rules make use of the syntactic sugar defined above, they are, strictly speaking, rule schemes: each of them really stands for the infinite family of rules that would be obtained if the syntactic sugar was eliminated.

Let us now define the syntax of types as in Example 10.8.7. Let the initial environment Γ_0 contain the following bindings:

$$\begin{array}{l} \{\}_{\{\ell_1, \dots, \ell_n\}} : \forall X_1 \dots X_n X. X_1 \rightarrow \dots \rightarrow X_n \rightarrow X \rightarrow \Pi (\ell_1 : X_1; \dots; \ell_n : X_n; \partial X) \\ \quad \text{where } \ell_1 < \dots < \ell_n \\ \{\cdot \text{ with } \ell = \cdot\} : \forall X X' Y. \Pi (\ell : X; Y) \rightarrow X' \rightarrow \Pi (\ell : X'; Y) \\ \cdot.\{\ell\} : \forall X Y. \Pi (\ell : X; Y) \rightarrow X \end{array}$$

Note that, in particular, the type scheme assigned to record creation $\{\cdot\}$ is $\forall X. X \rightarrow \Pi (\partial X)$. As a result, these bindings are exactly as stated in the discussion of records with infinite carrier (p. 463).

To illustrate how these definitions work together, let us consider the program $\{\{0\} \text{ with } \ell_1 = \text{true}\}.\{\ell_2\}$, which builds a record, extends it at ℓ_1 , then accesses it at ℓ_2 . Can we build an $\text{HM}(X)$ type derivation for it, under the constraint true and the initial environment Γ_0 ? To begin, by looking up Γ_0 and using HMx-INST , we find that $\{\cdot\}$ has type $\text{int} \rightarrow \Pi (\partial \text{int})$. Thus, assuming that 0 has type int , the expression $\{0\}$ has type $\Pi (\partial \text{int})$. Indeed, this expression denotes a record all of whose fields hold an integer value. Then, by looking up Γ_0 and using HMx-INST , we find that $\{\cdot \text{ with } \ell_1 = \cdot\}$ has type $\Pi (\ell_1 : \text{int}; \partial \text{int}) \rightarrow \text{bool} \rightarrow \Pi (\ell_1 : \text{bool}; \partial \text{int})$. May we immediately use HMx-APP to typecheck the application of $\{\cdot \text{ with } \ell_1 = \cdot\}$ to $\{0\}$? Unfortunately, no, because there is an apparent mismatch between the expected type $\Pi (\ell_1 : \text{int}; \partial \text{int})$ and the effective type $\Pi (\partial \text{int})$. To work around this problem, let us recall that, by C-ROW-DL , the equation $\Pi (\partial \text{int}) = \Pi (\ell_1 : \text{int}; \partial \text{int})$ is equivalent to true . Thus, HMx-SUB allows proving that $\{0\}$ has type $\Pi (\ell_1 : \text{int}; \partial \text{int})$.

Assuming that `true` has type `bool`, we may now apply HMX-APP and deduce

$$\text{true}, \Gamma_0 \vdash \{\{0\} \text{ with } \ell_1 = \text{true}\} : \Pi (\ell_1 : \text{bool} ; \partial \text{int}).$$

We let the reader check that, in a similar manner involving C-ROW-DL, C-ROW-LL, and HMX-SUB, one may prove that $\{\{0\} \text{ with } \ell_1 = \text{true}\}.\{\ell_2\}$ has type `int`, provided ℓ_1 and ℓ_2 are distinct. \square

- 10.8.26 EXERCISE [★★, →]: Unfold the definition of the constraint $\text{let } \Gamma_0 \text{ in } \llbracket \{\{0\} \text{ with } \ell_1 = \text{true}\}.\{\ell_2\} : X \rrbracket$, which states that X is a valid type for the above program. Assuming that subtyping is interpreted as equality, simulate a run of the constraint solver (§10.6), extended with support for rows, so as to solve this constraint. Check that the solved form is equivalent to $X = \text{int}$. \square
- 10.8.27 EXERCISE [★★★]: Check that the definitions of Example 10.8.25 meet the requirements of Definition 10.5.5. \square
- 10.8.28 EXAMPLE [RECORD APPLICATION]: Let us now introduce a more unusual primitive operation on full records. This operation accepts two records, the first of which is expected to hold a function in every field and produces a new record, whose contents are obtained by applying, pointwise, the functions in the first record to the values in the second record. In other words, this new primitive operation lifts the standard application combinator (which may be defined as $\lambda f.\lambda z.f z$), pointwise, to the level of records. For this reason, we refer to it as `rapply`. Its operational semantics is defined by making it a binary destructor and equipping it with the following reduction rules:

$$\begin{aligned} \text{rapply } \{V; v\} \{V'; v'\} &\xrightarrow{\delta} \{V V'; v v'\} && \text{(R-APPLY-1)} \\ \text{rapply } \{V; v\} \{V'; v'\} &\xrightarrow{\delta} \text{rapply } \{V; v\} \{V'[\ell \mapsto v']; v'\} && \text{(R-APPLY-2)} \\ &&& \text{if } \ell \in \text{dom}(V) \setminus \text{dom}(V') \\ \text{rapply } \{V; v\} \{V'; v'\} &\xrightarrow{\delta} \text{rapply } \{V[\ell' \mapsto v]; v\} \{V'; v'\} && \text{(R-APPLY-3)} \\ &&& \text{if } \ell' \in \text{dom}(V') \setminus \text{dom}(V) \end{aligned}$$

In the first rule, $V V'$ is defined only if V and V' have a common domain; it is then defined as the function that maps ℓ to the expression $V(\ell) V'(\ell)$. The second and third rules, which are symmetric, deal with the case where some field is explicitly defined in one input record but not in the other; in that case, the field is made explicit by creating a copy of the record's default value.

The syntax of types remains as in Example 10.8.25. We extend the initial environment Γ_0 with the following binding:

$$\text{rapply} : \forall XY. \Pi (X \rightarrow Y) \rightarrow \Pi X \rightarrow \Pi Y$$

To understand this type scheme, recall that the principal type scheme of the standard application combinator (which may be defined as $\lambda f.\lambda z.f z$) is $\forall XY.(X \rightarrow Y) \rightarrow X \rightarrow Y$. The type scheme assigned to `rapply` is very similar; the most visible difference is that both arguments, as well as the result, are now wrapped within the record type constructor Π . A more subtle, yet essential change is that X and Y are now row variables: their kind is $\star.\text{Row}(\emptyset)$. As a result, the leftmost occurrence of the arrow constructor is really $\rightarrow^{\text{Row}(\emptyset)}$. Thus, we are exploiting the presence of type constructors of the form G^s , with $s \neq \text{Type}$, in the signature S .

To illustrate how these definitions work together, let us consider the program `rapply {ℓ = not; succ} {ℓ = true; 0}`, where the terms `not` and `succ` are assumed to have types $\text{bool} \rightarrow \text{bool}$ and $\text{int} \rightarrow \text{int}$, respectively. Can we build an $\text{HM}(X)$ type derivation for it, under the constraint `true` and the initial environment Γ_0 ? To begin, it is straightforward to derive that the record `{ℓ = not; succ}` has type $\Pi (\ell : \text{bool} \rightarrow \text{bool} ; \partial(\text{int} \rightarrow \text{int}))$ (1). In order to use `rapply`, however, we must prove that this record has a type of the form $\Pi (R_1 \rightarrow R_2)$, where R_1 and R_2 are rows. This is where C-ROW-GD and C-ROW-GL (Figure 10-12) come into play. Indeed, by C-ROW-GD, the type $\partial(\text{int} \rightarrow \text{int})$ may be written $\partial\text{int} \rightarrow \partial\text{int}$. So, (1) may be written $\Pi (\ell : \text{bool} \rightarrow \text{bool} ; \partial\text{int} \rightarrow \partial\text{int})$ (2), which by C-ROW-GL may be written $\Pi ((\ell : \text{bool} ; \partial\text{int}) \rightarrow (\ell : \text{bool} ; \partial\text{int}))$ (3). Thus, HMx-SUB allows deriving that the record `{ℓ = not; succ}` has type (3). We let the reader continue and conclude that the program has type $\Pi (\ell : \text{bool} ; \partial\text{int})$ under the constraint `true` and the initial environment Γ_0 .

This example illustrates a very important use of rows, namely to lift an operation on ordinary values so as to turn it into a pointwise operation on records. Here, we have chosen to lift the standard application combinator, giving rise to `rapply` on records. The point is that, thanks to the expressive power of rows, we were also able to lift the standard combinator's type scheme in the most straightforward manner, giving rise to a suitable type scheme for `rapply`. \square

10.8.29 EXERCISE [★★★, →]: Check that the definitions of Example 10.8.28 meet the requirements of Definition 10.5.5. \square

The previous examples have illustrated the use of rows to typecheck operations on full records. Let us now move to records with finite domain. As explained in the discussion above of records with finite carrier (p. 461), they may be either encoded in terms of full records, or given a direct definition. The latter approach is illustrated below.

10.8.30 EXAMPLE [FINITE RECORDS]: For every set of labels L of cardinal n , let us introduce a n -ary constructor $\langle \rangle_L$. We define the notations $\langle \ell_1 = \tau_1; \dots; \ell_n = \tau_n \rangle$

and $\langle V \rangle$, where V is a finite mapping of labels to values, in a manner similar to that of Example 10.8.25.

The three primitive operations on finite records, namely the empty record $\langle \rangle$, extension $\langle \cdot \text{ with } \ell = \cdot \rangle$, and access $\cdot.\langle \ell \rangle$, may be defined as follows. First, the empty record $\langle \rangle$ is precisely the nullary constructor $\langle \rangle_\emptyset$. Second, for every $\ell \in \mathcal{L}$, let extension $\langle \cdot \text{ with } \ell = \cdot \rangle$ and access $\cdot.\langle \ell \rangle$ be destructors of arity 1 and 2, respectively, equipped with the following reduction rules:

$$\begin{aligned} \langle \langle V \rangle \text{ with } \ell = v \rangle &\xrightarrow{\delta} \langle V[\ell \mapsto v] \rangle && \text{(R-EXTEND)} \\ \langle V \rangle.\langle \ell \rangle &\xrightarrow{\delta} V(\ell) && (\ell \in \text{dom}(V)) \quad \text{(R-ACCESS)} \end{aligned}$$

Let us now define the syntax of types as in Example 10.8.8. Let the initial environment Γ_0 contain the following bindings:

$$\begin{aligned} \langle \rangle_{\{\ell_1, \dots, \ell_n\}} &: \forall X_1 \dots X_n. X_1 \rightarrow \dots \rightarrow X_n \rightarrow \Pi (\ell_1 : \text{pre } X_1; \dots; \ell_n : \text{pre } X_n; \partial \text{abs}) \\ &\text{where } \ell_1 < \dots < \ell_n \\ \langle \cdot \text{ with } \ell = \cdot \rangle &: \forall X X' Y. \Pi (\ell : X; Y) \rightarrow X' \rightarrow \Pi (\ell : \text{pre } X'; Y) \\ \cdot.\langle \ell \rangle &: \forall X Y. \Pi (\ell : \text{pre } X; Y) \rightarrow X \end{aligned}$$

Note that, in particular, the type scheme assigned to the empty record $\langle \rangle$ is $\Pi (\partial \text{abs})$. \square

- 10.8.31 EXERCISE [RECOMMENDED, \star , \rightarrow]: Reconstruct all of the missing kind information in the type schemes given in Example 10.8.30. \square
- 10.8.32 EXERCISE [RECOMMENDED, $\star\star$, \rightarrow]: Give an encoding of finite records in terms of full records, along the lines of the discussion of records with finite carrier (p. 461). Check that the principal type schemes associated, via the encoding, with the three operations on finite records are precisely those given in Example 10.8.30. \square
- 10.8.33 EXERCISE [RECOMMENDED, \star]: The extension operation, as defined above, may either change the value of an existing field or create a new field, depending on whether the field ℓ is or isn't present in the input record. This flavor is known as free extension. Can you define a strict flavor of extension that is not applicable when the field ℓ already exists? Can you define (free and strict flavors of) a restriction operation that removes a field from a record? \square
- 10.8.34 EXERCISE [RECOMMENDED, \star]: Explain why, when $\text{pre} \leq \text{abs}$ holds, subsumption allows a record with more fields to be supplied in a context where a record with fewer fields is expected. This phenomenon is often known as *width subtyping*. Explain why such is not the case when subtyping is interpreted as equality. \square
- 10.8.35 EXERCISE [$\star\star\star$, \rightarrow]: Check that the definitions of Example 10.8.30 meet the requirements of Definition 10.5.5. \square

Polymorphic Variants

So far, we have emphasized the use of rows for flexible typechecking of operations on records. The record type constructor Π expects one parameter, which is a row; informally speaking, one might say that it is a product constructor of infinite arity. It appears natural to also define sums of infinite arity. This may be done by introducing a new unary type constructor Σ , whose parameter is a row.

As in the case of records, we use a nullary type constructor abs and a unary type constructor pre in order to associate information with every row label. Thus, for instance, the type $\Sigma (\ell_1 : \text{pre } T_1 ; \ell_2 : \text{pre } T_2 ; \partial \text{abs})$ is intended to contain values of the form $\ell_1 v_1$, where v_1 has type T_1 , or of the form $\ell_2 v_2$, where v_2 has type T_2 . The type constructors abs and pre are not the same type constructors as in the case of records. In particular, their subtyping relationship, if there is one, is reversed. Indeed, the type $\Sigma (\ell_1 : \text{pre } T_1 ; \ell_2 : \text{abs} ; \partial \text{abs})$ is intended to contain only values of the form $\ell_1 v_1$, where v_1 has type T_1 , so it is safe to make it a subtype of the above type; in other words, it is safe to allow $\text{abs} \leq \text{pre } T_2$. In spite of this, we keep the names abs and pre by tradition.

The advantages of this approach over algebraic data types are the same as in the case of records. The namespace of data constructors becomes global, so it becomes possible for two distinct sum types to share data constructors. Also, the expressiveness afforded by rows allows assigning types to new operations, such as *filtering* (see below), which allows functions that perform case analysis to be incrementally extended with new cases. One disadvantage is that it becomes more difficult to understand what it means for a function defined by pattern matching to be exhaustive; this issue is, however, out of the scope of this chapter.

- 10.8.36 **EXAMPLE [POLYMORPHIC VARIANTS]:** For every label $\ell \in \mathcal{L}$, let us introduce a unary constructor ℓ and a ternary destructor $[\ell : \cdot | \cdot] \cdot$. We refer to the former as a *data constructor*, and to the latter as a *filter*. Let us also introduce a unary destructor $[\]$. We equip these destructors with the following reduction rules:

$$\begin{aligned} [\ell : v | v'] (\ell w) &\xrightarrow{\delta} v w && \text{(R-FILTER-1)} \\ [\ell : v | v'] (\ell' w) &\xrightarrow{\delta} v' (\ell' w) && \text{if } \ell \neq \ell' \quad \text{(R-FILTER-2)} \end{aligned}$$

Let us define the syntax of types as follows. Let there be two basic kinds \star and \bullet . Let S_0 consist of the type constructors \rightarrow , abs , and pre , whose respective signatures are $\star \otimes \star \Rightarrow \star$, \bullet , and $\star \Rightarrow \bullet$. Let S_1 consist of the record type constructor Σ , whose signature is $\bullet \Rightarrow \star$. Note the similarity with the case of records (Example 10.8.8).

Subtyping is typically interpreted in one of two ways. One is equality. The other is the nonstructural subtyping order obtained by letting \rightarrow be contravariant in its domain and covariant in its codomain, Σ be covariant, \rightarrow and Σ be incompatible, and letting $\text{abs} \leq \text{pre}$. Compare this definition with the case of records (Remark 10.8.11).

To complete the setup, let the initial environment Γ_0 contain the following bindings:

$$\begin{aligned} \ell \cdot & : \quad \forall XY. X \rightarrow \Sigma (\ell : \text{pre } X ; Y) \\ [\ell : \cdot \mid \cdot] \cdot & : \quad \forall XX'YY'. (X \rightarrow Y) \rightarrow (\Sigma (\ell : X' ; Y') \rightarrow Y) \rightarrow \Sigma (\ell : \text{pre } X ; Y') \rightarrow Y \\ [] & : \quad \forall X. \Sigma (\partial \text{abs}) \rightarrow X \end{aligned}$$

The first binding means, in particular, that if v has type T , then a value of the form ℓv has type $\Sigma (\ell : \text{pre } T ; \partial \text{abs})$. This is a sum type with only one branch labeled ℓ , hence a very precise type for this value. However, it is possible to instantiate the row variable Y with rows other than ∂abs . For instance, the value ℓv also has type $\Sigma (\ell : \text{pre } T ; \ell' : \text{pre } T' ; \partial \text{abs})$. This is a sum type with two branches, hence a somewhat less precise type, but still a valid one for this value. It is clear that, through this mechanism, the value ℓv admits an infinite number of types. The point is that, if v has type T and v' has type T' , then both ℓv and $\ell' v'$ have type $\Sigma (\ell : \text{pre } T ; \ell' : \text{pre } T' ; \partial \text{abs})$, so they may be stored together in a homogeneous data structure, such as a list.

Filters are used to perform case analysis on variants, that is, on values of a sum type. According to R-FILTER-1 and R-FILTER-2, a filter $[\ell : v \mid v']$ is a function that expects an argument of the form $\ell' w$ and reduces to $v w$ if ℓ' is ℓ and to $v' (\ell' w)$ otherwise. Thus, a filter defines a two-way branch, where the label of the data constructor at hand determines which branch is taken. The expressive power of filters stems from the fact that they may be organized in a sequence, so as to define a multi-way branch. The inert filter $[\]$, which does not have a reduction rule, serves as a terminator for such sequences. For instance, the composite filter $[\ell : v \mid [\ell' : v' \mid [\]]]$, which may be abbreviated as $[\ell : v \mid \ell' : v']$, may be applied either to a value of the form ℓw , yielding $v w$, or to a value of the form $\ell' w'$, yielding $v' w'$. Applying it to a value w whose head symbol is not ℓ or ℓ' would lead to the term $[\] w$, which is stuck, since $[\]$ does not have a reduction rule.

For the type system to be sound, we must ensure that every application of the form $[\] w$ is ill-typed. This is achieved by the third binding above: the domain type of $[\]$ is $\Sigma (\partial \text{abs})$, a sum type with zero branches, which contains no values. The return type of $[\]$ may be chosen at will, which is fine; since it can never be invoked, it can never return. The second binding above means that, if v accepts values of type T and v' accepts values of type $\Sigma (\ell : T'' ; T')$, then the filter $[\ell : v \mid v']$ accepts values of type $\Sigma (\ell : \text{pre } T ; T')$. Note that

any choice of T'' will do, including, in particular, abs . In other words, it is okay if v' does not accept values of the form ℓw . Indeed, by definition of the semantics of filters, it will never be passed such a value. \square

10.8.37 EXERCISE [***, +]: Check that the definitions of Example 10.8.36 meet the requirements of Definition 10.5.5. \square

10.8.38 REMARK: It is interesting to study the similarity between the type schemes assigned to the primitive operations on polymorphic variants and those assigned to the primitive operations on records (Example 10.8.30). The type of $[]$ involves the complete row ∂abs , just like the empty record $\langle \rangle$. The type of $[\ell : \cdot \mid \cdot] \cdot$ is pretty much identical to the type of record extension $\langle \cdot \text{ with } \ell = \cdot \rangle$, provided the three continuation arrows $\rightarrow Y$ are dropped. Last, the type of the data constructor ℓ is strongly reminiscent of the type of record access $\cdot.\langle \ell \rangle$. With some thought, this is hardly a surprise. Indeed, records and variants are *dual*: it is possible to encode the latter in terms of the former and vice-versa. For instance, in the encoding of variants in terms of records, a function defined by cases is encoded as a record of ordinary functions, in continuation-passing style. Thus, the encoding of $[]$ is $\lambda f.f \langle \rangle$, the encoding of $[\ell : v \mid v']$ is $\lambda f.f \langle v' \text{ with } \ell = v \rangle$, and the encoding of ℓv is $\lambda r.r.\langle \ell \rangle v$. The reader is encouraged to study the type schemes that arise out of this encoding and how they relate to the type schemes given in Example 10.8.36. \square

10.8.39 EXAMPLE [FIRST-CLASS MESSAGES]: In a programming language equipped with both records and variants, it is possible to make the duality between these two forms of data explicit by extending the language with a primitive operation $\#$ that turns a record of ordinary functions into a single function, defined by cases. More precisely, $\#$ may be introduced as a binary destructor, whose reduction rule is

$$\# v (\ell w) \xrightarrow{\delta} v.\langle \ell \rangle w \quad (\text{R-SEND})$$

What type may we assign to such an operation? In order to simplify the answer, let us assume that we are dealing with full records (Example 10.8.25) and full variants; that is, we have a single basic kind \star , and do not employ abs and pre . Then, a suitable type scheme would be

$$\forall XY.\Pi (X \rightarrow \partial Y) \rightarrow \Sigma X \rightarrow Y$$

In other words, this operation accepts a record of functions, all of which have the same return type Y , but may have arbitrary domain types, which are given by the row X . It produces a function that accepts a parameter of sum type ΣX

and returns a result of type Y . The fact that the row X appears both in the Σ type and in the Π type reflects the operational semantics. Indeed, according to R-SEND, the label ℓ carried by the value ℓw is used to extract, out of the record v , a function, which is then applied to w . Thus, the domain type of the function stored at ℓ within the record v should match the type of w . In other words, at every label, the domain of the contents of the record and the contents of the sum should be type compatible. This is encoded by letting a single row variable X stand for both of these rows. Note that the arrow in $X \rightarrow \partial Y$ is really $\rightarrow^{Row(\emptyset)}$; once again, we are exploiting the presence of type constructors of the form G^s , with $s \neq Type$, in the signature S .

If the record of functions v is viewed as an *object*, and if the variant ℓw is viewed as a *message* ℓ carrying a parameter w , then R-SEND may be understood as (*first-class message dispatch*), a common feature of object-oriented languages. (The *first-class* qualifier refers to the fact that the message name ℓ is not statically fixed, but is discovered at runtime.) The issue of type inference in the presence of such a feature has been studied by Nishimura (1998), Müller and Nishimura (1998), and Pottier (2000). These papers address two issues that are not dealt with in the above example, namely (i) accommodating finite (as opposed to full) record and variants and (ii) allowing distinct methods to have distinct result types. This is achieved via the use of subtyping and of some form of conditional constraints. \square

10.8.40 EXERCISE [***, +]: Check that the definitions of Example 10.8.39 meet the requirements of Definition 10.5.5. \square

The name *polymorphic variants* stems from the highly polymorphic type schemes assigned to the operations on variants (Example 10.8.36). A row-based type system for polymorphic variants was first proposed by Rémy (1989). A somewhat similar, constraint-based type system for polymorphic variants was then studied by Garrigue (1998; 2000; 2002) and implemented by him as part of the programming language Objective Caml.

Other Applications of Rows

Typechecking records and variants is the best-known application of rows. Many variations of it are conceivable, some of which we have illustrated, such as the choice between full and finite records and variants. However, rows may also be put to other uses, of which we now list a few.

First, since objects may be viewed as records of functions, at least from a typechecking point of view, rows may be used to typecheck object-oriented languages in a structural style (Wand, 1994; Rémy, 1994). This is, in particular, the route followed in Objective Caml (Rémy and Vouillon, 1998). There,

an object type consists of a row of method types, and gives the object's interface. Such a style is considered structural, as opposed to the style adopted by many popular object-oriented languages, such as C++, Java, and C#, where an object type consists of the name of its class. Thanks to rows, method invocation may be assigned a polymorphic type scheme, similar to that of record access (Example 10.8.30), making it possible to invoke a specific method (say, ℓ) without knowing which class the receiver object belongs to.

Rows may also be used to encode sets of properties within types or to encode type refinements, with applications in type-based program analysis. Some instances worth mentioning are soft typing (Cartwright and Fagan, 1991; Wright and Cartwright, 1994), exception analysis (Leroy and Pessaux, 2000; Pottier and Simonet, 2003), and static enforcement of an access control policy (Pottier, Skalka, and Smith, 2001). BANE (Fähndrich, 1999), a versatile program analysis toolkit, also implements a form of rows.

Variations on Rows

A type system may be said to have rows, in a broad sense, if mappings from labels to types may be (i) defined incrementally, via some syntax for extending an existing mapping with information about a new label and (ii) abstracted by a type variable. In this chapter, which follows Rémy's ideas (1993; 1992a; 1992b), the former feature is provided by the row constructors $(\ell : \cdot ; \cdot)$, while the latter is provided by the existence of row variables, that is, type variables of row kind $Row(L)$ for some L . There are, however, type systems that provide (i) and (ii) while departing significantly from the one presented here. These systems differ mainly in how they settle some important design choices:

1. Does a row denote a finite or an infinite mapping from labels to types?
2. Is a row with duplicate labels considered well-formed? If not, by which mechanism is it ruled out?

In Rémy's approach, every row denotes an infinite (in fact, cofinite) mapping from labels to types. The type constructors `abs` and `pre` are used to encode domain information within field types. A row with duplicate labels, such as $(\ell : T_1 ; \ell : T_2 ; T_3)$, is ruled out by the kind system. Below, we mention a number of type systems that make different design choices.

The first use of rows for typechecking operations on records, including record extension, is due to Wand (1987a; 1988). In Wand's approach, rows denote finite mappings. Furthermore, rows with duplicate labels are considered legal; row extension is interpreted as function extension, so that, if a label occurs twice, the later occurrence takes precedence. This leads to a difficulty in

the constraint solving process: the constraint $(\ell : T_1 ; R_1) = (\ell : T_2 ; R_2)$ entails $T_1 = T_2$, but does not entail $R_1 = R_2$, because R_1 and R_2 may have different domains—indeed, their domains may differ at ℓ . Wand’s proposed solution (1988) introduces a four-way disjunction, because each of R_1 and R_2 may or may not define ℓ . This gives type inference exponential time complexity.

Later work (Berthomieu, 1993; Berthomieu and le Moniès de Sagazan, 1995) interprets rows as infinite mappings but sticks with Wand’s interpretation of row extension as function extension, so that duplicate labels are allowed. The constraint solving algorithm rewrites the problematic constraint $(\ell : T_1 ; R_1) = (\ell : T_2 ; R_2)$ to $(T_1 = T_2) \wedge (R_1 =_{\{\ell\}} R_2)$, where the new predicate $=_L$ is interpreted as row equality *outside* L . Of course, the entire constraint solver must then be extended to deal with constraints of the form $T_1 =_L T_2$. The advantage of this approach over Wand’s lies in the fact that no disjunctions are ever introduced, so that the time complexity of constraint solving apparently remains polynomial.

Several other works make opposite choices, sticking with Wand’s interpretation of rows as finite mappings but forbidding duplicate labels. No kind discipline is imposed: some other mechanism is used to ensure that duplicate labels do not arise. In Jategaonkar and Mitchell (1988) and Jategaonkar (1989), somewhat ad hoc steps are taken to ensure that, if the row $(\ell : T ; X)$ appears anywhere within a type derivation, then X is never instantiated with a row that defines ℓ . In Gaster and Jones (1996), Gaster (1998), and Jones and Peyton Jones (1999), explicit constraints prevent duplicate labels from arising. This line of work uses *qualified types* (Jones, 1994), a constraint-based type system that bears strong similarity with $HM(X)$. For every label ℓ , a unary predicate $\cdot \text{ lacks } \ell$ is introduced; roughly speaking, the constraint $R \text{ lacks } \ell$ is considered to hold if the (finite) row R does not define the label ℓ . The constrained type scheme assigned to record access is

$$\cdot.\langle \ell \rangle : \forall XY[Y \text{ lacks } \ell].\Pi (\ell : X ; Y) \rightarrow X.$$

The constraint $Y \text{ lacks } \ell$ ensures that the row $(\ell : X ; Y)$ is well-formed. Although interesting, this approach is not as expressive as that described in this chapter. For instance, although it accommodates record update (where the field being modified is known to exist in the initial record) and strict record extension (where the field is known not to initially exist), it cannot express a suitable type scheme for free record extension, where it is not known whether the field initially exists. This approach has been implemented as the “Trex” extension to Hugs (Jones and Peterson, 1999).

It is worth mentioning a line of type systems (Ohori and Buneman, 1988, 1989; Ohori, 1995) that do not have rows, because they lack feature (i) above, but are still able to assign a polymorphic type scheme to record access. One

might explain their approach as follows. First, these systems are equipped with ordinary, structural record types, of the form $\{\ell_1 : T_1; \dots; \ell_n : T_n\}$. Second, for every label ℓ , a binary predicate $\cdot \text{ has } \ell : \cdot$ is available. The idea is that the constraint $T \text{ has } \ell : T'$ holds if and only if T is a record type that contains the field $\ell : T'$. Then, record access may be assigned the constrained type scheme

$$\cdot.\langle\ell\rangle : \forall XY[X \text{ has } \ell : Y].X \rightarrow Y.$$

This technique also accommodates a restricted form of record update, where the field being written must initially exist and must keep its initial type; it does not, however, accommodate any form of record extension, because of the absence of row extension in the syntax of types. Although the papers cited above employ different terminology, we believe it is fair to view them as constraint-based type systems. In fact, Odgersky, Sulzmann, and Wehr (1999) prove that Ohori's system (1995) may be viewed as an instance of $\text{HM}(X)$. Sulzmann (2000) proposes several extensions of it, also presented as instances of $\text{HM}(X)$, which accommodate record extension and concatenation using new, ad hoc constraint forms in addition to $\cdot \text{ has } \ell$.

In the *label-selective λ -calculus* (Garrigue and Ait-Kaci, 1994; Furuse and Garrigue, 1995), the arrow type constructor carries a label, and arrows that carry distinct labels may commute, so as to allow labeled function arguments to be supplied in any order. Some of the ideas that underlie this type system are closely related to rows.

Pottier (2003) describes an instance of $\text{HM}(X)$ where rows are not part of the syntax of types: equivalent expressive power is obtained via an extension of the constraint language. The idea is to work with constraints of the form $R_1 \leq_L R_2$, where L may be finite or cofinite, and to interpret such a constraint as row subtyping *inside* L . In this approach, no new type variables need be allocated during constraint solving; contrast this with S-MUTATE-LL , S-MUTATE-GD , and S-MUTATE-GL in Figure 10-14. One benefit is to simplify the complexity analysis; another is to yield insights that lead to generalizations of rows.

Even though rows were originally invented with type inference in mind, they are useful in explicitly typed languages as well; indeed, other approaches to typechecking operations on records appear quite complex (Cardelli and Mitchell, 1991).

2. Adding a conditional module expression destroys the phase distinction, because the types in a conditional module, e.g.

`if ... then (Nat::*) else (Unit::*),`

depends on the run-time value of the test.

- 10.1.22 SOLUTION: Within Damas and Milner's type system, we have:

$$\frac{\frac{\text{DM-VAR } \frac{}{z_1 : X \vdash z_1 : X} \quad \frac{}{z_1 : X; z_2 : X \vdash z_2 : X} \text{DM-VAR}}{\text{DM-LET } \frac{}{z_1 : X \vdash \text{let } z_2 = z_1 \text{ in } z_2 : X}}}{\text{DM-ABS } \frac{}{\emptyset \vdash \lambda z_1. \text{let } z_2 = z_1 \text{ in } z_2 : X \rightarrow X}}$$

Note that, because X occurs free within the environment $z_1 : X$, it is impossible to apply DM-GEN to the judgment $z_1 : X \vdash z_1 : X$ in a nontrivial way. For this reason, z_2 cannot receive the type scheme $\forall X.X$, and the whole expression cannot receive type $X \rightarrow Y$, where X and Y are distinct.

- 10.1.23 SOLUTION: It is straightforward to prove that the identity function has type $\text{int} \rightarrow \text{int}$:

$$\frac{\frac{}{\Gamma_0; z : \text{int} \vdash z : \text{int}} \text{DM-VAR}}{\Gamma_0 \vdash \lambda z. z : \text{int} \rightarrow \text{int}} \text{DM-ABS}$$

In fact, nothing in this type derivation depends on the choice of int as the type of z . Thus, we may just as well use a type variable X instead. Furthermore, after forming the arrow type $X \rightarrow X$, we may employ DM-GEN to quantify universally over X , since X no longer appears in the environment.

$$\frac{\frac{\text{DM-VAR } \frac{}{\Gamma_0; z : X \vdash z : X}}{\text{DM-ABS } \frac{}{\Gamma_0 \vdash \lambda z. z : X \rightarrow X} \quad X \notin \text{ftv}(\Gamma_0)}}{\text{DM-GEN } \frac{}{\Gamma_0 \vdash \lambda z. z : \forall X. X \rightarrow X}}$$

It is worth noting that, although the type derivation employs an arbitrary type variable X , the final typing judgment has no free type variables. It is thus independent of the choice of X . In the following, we refer to the above type derivation as Δ_0 .

Next, we prove that the successor function has type $\text{int} \rightarrow \text{int}$ under the initial environment Γ_0 . We write Γ_1 for $\Gamma_0; z : \text{int}$, and make use of DM-VAR implicit.

$$\frac{\frac{\text{DM-APP } \frac{\Gamma_1 \vdash \hat{z} : \text{int} \rightarrow \text{int} \rightarrow \text{int} \quad \Gamma_1 \vdash z : \text{int}}{\Gamma_1 \vdash \hat{z} z : \text{int} \rightarrow \text{int}} \quad \Gamma_1 \vdash \hat{1} : \text{int}}{\text{DM-APP } \frac{}{\Gamma_1 \vdash z \hat{z} \hat{1} : \text{int}}}}{\text{DM-ABS } \frac{}{\Gamma_0 \vdash \lambda z. z \hat{z} \hat{1} : \text{int} \rightarrow \text{int}}}$$

In the following, we refer to the above type derivation as Δ_1 . We may now build a derivation for the third typing judgment. We write Γ_2 for $\Gamma_0; f : \text{int} \rightarrow \text{int}$.

$$\frac{\Delta_1 \quad \frac{\Gamma_2 \vdash f : \text{int} \rightarrow \text{int} \quad \Gamma_2 \vdash \hat{2} : \text{int}}{\Gamma_2 \vdash f \hat{2} : \text{int}} \text{DM-APP}}{\Gamma_0 \vdash \text{let } f = \lambda z. z \hat{1} \text{ in } f \hat{2} : \text{int}} \text{DM-LET}$$

To derive the fourth typing judgment, we re-use Δ_0 , which proves that the identity function has polymorphic type $\forall X. X \rightarrow X$. We write Γ_3 for $\Gamma_0; f : \forall X. X \rightarrow X$. By DM-VAR and DM-INST, we have both $\Gamma_3 \vdash f : (\text{int} \rightarrow \text{int}) \rightarrow (\text{int} \rightarrow \text{int})$ and $\Gamma_3 \vdash f : \text{int} \rightarrow \text{int}$. Thus, we may build the following derivation:

$$\frac{\Delta_0 \quad \frac{\Gamma_3 \vdash f : (\text{int} \rightarrow \text{int}) \rightarrow (\text{int} \rightarrow \text{int}) \quad \Gamma_3 \vdash f : \text{int} \rightarrow \text{int}}{\Gamma_3 \vdash f f : \text{int} \rightarrow \text{int}} \text{DM-APP} \quad \Gamma_3 \vdash \hat{2} : \text{int}}{\Gamma_0 \vdash \text{let } f = \lambda z. z \text{ in } f f \hat{2} : \text{int}} \text{DM-LET}$$

The first and third judgments are valid in the simply-typed λ -calculus, because they use neither DM-GEN nor DM-INST, and use DM-LET only to introduce the *monomorphic* binding $f : \text{int} \rightarrow \text{int}$ into the environment. The second judgment, of course, is not: because it involves a nontrivial type scheme, it is not even a well-formed judgment in the simply-typed λ -calculus. The fourth judgment is well-formed, but *not* derivable, in the simply-typed λ -calculus. This is because f is used at two incompatible types, namely $(\text{int} \rightarrow \text{int}) \rightarrow (\text{int} \rightarrow \text{int})$ and $\text{int} \rightarrow \text{int}$, inside the expression $f f \hat{2}$. Both of these types are *instances* of $\forall X. X \rightarrow X$, the type scheme assigned to f in the environment Γ_3 .

By inspection of the rules, a derivation of $\Gamma_0 \vdash \hat{1} : T$ must begin with an instance of DM-VAR, of the form $\Gamma_0 \vdash \hat{1} : \text{int}$. It may be followed by an arbitrary number of instances of the sequence (DM-GEN; DM-INST), turning int into a type scheme of the form $\forall \bar{X}. \text{int}$, then back to int . Thus, T must be int . Because int is not an arrow type, there follows that the application $\hat{1} \hat{2}$ cannot be well-typed under Γ_0 . In fact, because this expression is stuck, it cannot be well-typed in a sound type system.

The expression $\lambda f. (f f)$ is ill-typed in the simply-typed λ -calculus, because no type T may coincide with a type of the form $T \rightarrow T'$: indeed, T would be a subterm of itself. In DM, this expression is ill-typed as well, but the proof of this fact is slightly more complex. One must point out that, because f is λ -bound, it must be assigned a type T (as opposed to a type scheme) in the environment. Furthermore, one must note that DM-GEN is not applicable (except in a trivial way) to the judgment $\Gamma_0; f : T \vdash f : T$, because all of the

type variables in the type T appear free in the environment $\Gamma_0; f : T$. Once these points are made, the proof is the same as in the simply-typed λ -calculus.

It is important to note that the above argument crucially relies on the fact that f is λ -bound and must be assigned a *type*, as opposed to a type scheme. Indeed, we have proved earlier in this exercise that the self-application $f f$ is well-typed when f is λ -bound and is assigned the type scheme $\forall X.X \rightarrow X$. For the same reason, $\lambda f.(f f)$ is well-typed in an implicitly-typed variant of System F. It also relies on the fact that types are *finite*: indeed, $\lambda f.(f f)$ is well-typed in an extension of the simply-typed λ -calculus with recursive types, where the equation $T = T \rightarrow T'$ has a solution.

Later, we will develop a type inference algorithm for ML-the-type-system and prove that it is correct and complete. Then, to prove that a term is ill-typed, it will be sufficient to simulate a run of the algorithm and to check that it reports a failure.

- 10.3.2 SOLUTION: Our hypotheses are $C, \Gamma \vdash t : \forall \vec{X}[D].T$ (1) and $C \Vdash [\vec{X} \mapsto \vec{T}]D$ (2). We may also assume, *w.l.o.g.*, $\vec{X} \# \text{ftv}(C, \Gamma, \vec{T})$ (3). By HMX-INST and (1), we have $C \wedge D, \Gamma \vdash t : T$, which by Lemma 10.3.1 yields $C \wedge D \wedge \vec{X} = \vec{T}, \Gamma \vdash t : T$ (4). Now, we claim that $\vec{X} = \vec{T} \Vdash T \leq [\vec{X} \mapsto \vec{T}]T$ (5) holds; the proof appears in the next paragraph. Applying HMX-SUB to (4) and to (5), we obtain $C \wedge D \wedge \vec{X} = \vec{T}, \Gamma \vdash t : [\vec{X} \mapsto \vec{T}]T$ (6). By C-EQ and by (2), we have $C \wedge \vec{X} = \vec{T} \Vdash D$, so (6) may be written $C \wedge \vec{X} = \vec{T}, \Gamma \vdash t : [\vec{X} \mapsto \vec{T}]T$ (7). Last, (3) implies $\vec{X} \# \text{ftv}(\Gamma, [\vec{X} \mapsto \vec{T}]T)$ (8). Applying rule HMX-EXISTS to (7) and (8), we get $\exists \vec{X}.(C \wedge \vec{X} = \vec{T}), \Gamma \vdash t : [\vec{X} \mapsto \vec{T}]T$ (9). By C-NAMEEQ and by (3), $\exists \vec{X}.(C \wedge \vec{X} = \vec{T})$ is equivalent to C , hence (9) is the goal $C, \Gamma \vdash t : [\vec{X} \mapsto \vec{T}]T$.

There now remains to establish (5). One possible proof method is to unfold the definition of \Vdash and reason by structural induction on T . Here is another, axiomatic approach. Let Z be fresh for T, \vec{X} , and \vec{T} . By reflexivity of subtyping and by C-EXTRANS, we have $\text{true} \equiv T \leq T \equiv \exists Z.(T \leq Z \wedge Z \leq T)$, which by congruence of \equiv and by C-EXAND implies $\vec{X} = \vec{T} \equiv \exists Z.(T \leq Z \wedge \vec{X} = \vec{T} \wedge Z \leq T)$ (10). Furthermore, by C-EQ, we have $(\vec{X} = \vec{T} \wedge Z \leq T) \equiv (\vec{X} = \vec{T} \wedge Z \leq [\vec{X} \mapsto \vec{T}]T) \Vdash (Z \leq [\vec{X} \mapsto \vec{T}]T)$ (11). Combining (10) and (11) yields $\vec{X} = \vec{T} \Vdash \exists Z.(T \leq Z \wedge Z \leq [\vec{X} \mapsto \vec{T}]T)$, which by C-EXTRANS may be read $\vec{X} = \vec{T} \Vdash T \leq [\vec{X} \mapsto \vec{T}]T$.

- 10.3.3 SOLUTION: The simplest possible derivation of $\text{true}, \emptyset \vdash \lambda z.z : \text{int} \rightarrow \text{int}$ is syntax-directed. It closely resembles the Damas-Milner derivation given in Exercise 10.1.23.

$$\frac{\frac{}{\text{true}, z : \text{int} \vdash z : \text{int}} \text{HMX-VAR}}{\text{true}, \emptyset \vdash \lambda z.z : \text{int} \rightarrow \text{int}} \text{HMX-ABS}$$

As in Exercise 10.1.23, we may use a type variable X instead of the type int ,

then employ HMX-GEN to quantify universally over X .

$$\frac{\frac{\frac{}{\text{true}, z : X \vdash z : X} \text{HMX-VAR}}{\text{true}, \emptyset \vdash \lambda z.z : X \rightarrow X} \text{HMX-ABS} \quad X \# \text{ftv}(\text{true}, \emptyset)}{\text{true}, \emptyset \vdash \lambda z.z : \forall X[\text{true}].X \rightarrow X} \text{HMX-GEN}}$$

The validity of this instance of HMX-GEN relies on the equivalence $\text{true} \wedge \text{true} \equiv \text{true}$ and on the fact that judgments are identified up to equivalence of their constraint assumptions.

If we now wish to instantiate X with int , we may use HMX-INST' as follows:

$$\frac{\text{true}, \emptyset \vdash \lambda z.z : \forall X[\text{true}].X \rightarrow X \quad \text{true} \Vdash [X \mapsto \text{int}]\text{true}}{\text{true}, \emptyset \vdash \lambda z.z : \text{int} \rightarrow \text{int}} \text{HMX-INST'}$$

This is not, strictly speaking, an $\text{HM}(X)$ derivation, since HMX-INST' is not part of the rules of Figure 10-7. However, since the proof of Lemma 10.3.1 and the solution of Exercise 10.3.2 are constructive, it is possible to exhibit the $\text{HM}(X)$ derivation that underlies it. We find:

$$\frac{\frac{\frac{\frac{\frac{}{Y = \text{int}, z : X \vdash z : X} \text{HMX-VAR}}{Y = \text{int}, \emptyset \vdash \lambda z.z : X \rightarrow X} \text{HMX-ABS}}{Y = \text{int}, \emptyset \vdash \lambda z.z : \forall X.X \rightarrow X} \text{HMX-GEN}}{Y = \text{int}, \emptyset \vdash \lambda z.z : Y \rightarrow Y} \text{HMX-INST}}{Y = \text{int} \Vdash Y \rightarrow Y \leq \text{int} \rightarrow \text{int}} \text{HMX-SUB}}{\exists Y.(Y = \text{int}), \emptyset \vdash \lambda z.z : \text{int} \rightarrow \text{int}} \text{HMX-EXISTS}}$$

Since $\exists Y.(Y = \text{int})$ is equivalent to true , the conclusion is indeed the desired judgment.

- 10.4.1 SOLUTION: Let $X \notin \text{ftv}(\Gamma)$ **(1)**. Assume that there exist a satisfiable constraint C and a type T such that $C, \Gamma \vdash \tau : T$ **(2)** holds. Thanks to (1), we find that, up to a renaming of C and T , we may further assume $X \notin \text{ftv}(C, T)$ **(3)**. Then, applying Lemma 10.3.1 to (2), we obtain $C \wedge T = X, \Gamma \vdash \tau : T$, which by HMX-SUB yields $C \wedge T = X, \Gamma \vdash \tau : X$ **(4)**. Furthermore, by (3) and C-NAMEEQ, we have $\exists X.(C \wedge T = X) \equiv C$. Because C is satisfiable, this implies that $C \wedge T = X$ is satisfiable as well. As a result, we have found a satisfiable constraint C' such that $C', \Gamma \vdash \tau : X$ holds.

Now, assume Γ is closed and X is arbitrary. Then, (1) holds, so the previous paragraph proves that, if τ is well-typed within Γ , then there exists a satisfiable constraint C' such that $C', \Gamma \vdash \tau : X$ holds. By the completeness property,

we must then have $C' \Vdash \llbracket \Gamma \vdash \mathbf{t} : X \rrbracket$. Since C' is satisfiable, this implies that $\llbracket \Gamma \vdash \mathbf{t} : X \rrbracket$ is satisfiable as well. Conversely, if $\llbracket \Gamma \vdash \mathbf{t} : X \rrbracket$ is satisfiable, then, by the soundness property, \mathbf{t} is well-typed within Γ .

10.7.1 SOLUTION: We have

$$\begin{aligned} & \text{let } \Gamma_0 \text{ in } \llbracket c \ \mathbf{t}_1 \ \dots \ \mathbf{t}_n : T' \rrbracket \\ \equiv & \text{let } \Gamma_0 \text{ in } \exists Z_1 \dots Z_n. (\bigwedge_{i=1}^n \llbracket \mathbf{t}_i : Z_i \rrbracket \wedge c \leq Z_1 \rightarrow \dots \rightarrow Z_n \rightarrow T') \end{aligned} \quad (1)$$

$$\equiv \text{let } \Gamma_0 \text{ in } \exists Z_1 \dots Z_n \bar{X}. (\bigwedge_{i=1}^n \llbracket \mathbf{t}_i : Z_i \rrbracket \wedge T_1 \rightarrow \dots \rightarrow T_n \rightarrow T \leq Z_1 \rightarrow \dots \rightarrow Z_n \rightarrow T') \quad (2)$$

$$\equiv \text{let } \Gamma_0 \text{ in } \exists \bar{X}. (\bigwedge_{i=1}^n \llbracket \mathbf{t}_i : T_i \rrbracket \wedge T \leq T') \quad (3)$$

where (1) is by definition of constraint generation; (2) is by C-INID; (3) is by C-ARROW, C-EXAND, and by Lemma 10.4.6.

10.7.2 SOLUTION: We must first ensure that R-ADD respects \sqsubseteq (Definition 10.5.4). Since the rule is pure, it is sufficient to establish that let Γ_0 in $\llbracket \hat{n}_1 \hat{+} \hat{n}_2 : T \rrbracket$ entails let Γ_0 in $\llbracket \widehat{n_1 + n_2} : T \rrbracket$. In fact, we have

$$\begin{aligned} & \text{let } \Gamma_0 \text{ in } \llbracket \hat{n}_1 \hat{+} \hat{n}_2 : T \rrbracket \\ \equiv & \text{let } \Gamma_0 \text{ in } (\llbracket \hat{n}_1 : \text{int} \rrbracket \wedge \llbracket \hat{n}_2 : \text{int} \rrbracket \wedge \text{int} \leq T) \end{aligned} \quad (1)$$

$$\equiv \text{let } \Gamma_0 \text{ in } (\text{int} \leq \text{int} \wedge \text{int} \leq \text{int} \wedge \text{int} \leq T) \quad (2)$$

$$\equiv \text{int} \leq T \quad (3)$$

$$\equiv \text{let } \Gamma_0 \text{ in } \llbracket \widehat{n_1 + n_2} : T \rrbracket \quad (4)$$

where (1) and (2) are by Exercise 10.7.1; (3) is by C-IN* and by reflexivity of subtyping; (4) is by Exercise 10.7.1 again.

Second, we must check that if the configuration $c \ \mathbf{v}_1 \ \dots \ \mathbf{v}_k / \mu$ (where $k \geq 0$) is well-typed, then either it is reducible, or $c \ \mathbf{v}_1 \ \dots \ \mathbf{v}_k$ is a value.

We begin by checking that every value that is well-typed with type int is of the form \hat{n} . Indeed, suppose that let $\Gamma_0; \text{ref } M$ in $\llbracket \mathbf{v} : \text{int} \rrbracket$ is satisfiable. Then, \mathbf{v} cannot be a program variable, for a well-typed value must be closed. \mathbf{v} cannot be a memory location m , for otherwise $\text{ref } M(m) \leq \text{int}$ would be satisfiable—but the type constructors ref and int are incompatible. \mathbf{v} cannot be $\hat{+}$ or $\hat{+} \ \mathbf{v}'$, for otherwise $\text{int} \rightarrow \text{int} \rightarrow \text{int} \leq \text{int}$ or $\text{int} \rightarrow \text{int} \leq \text{int}$ would be satisfiable—but the type constructors \rightarrow and int are incompatible. Similarly, \mathbf{v} cannot be a λ -abstraction. Thus, \mathbf{v} must be of the form \hat{n} , for it is the only case left.

Next, we note that, according to the constraint generation rules, if the configuration $c \ \mathbf{v}_1 \ \dots \ \mathbf{v}_k / \mu$ is well-typed, then a constraint of the form let $\Gamma_0; \text{ref } M$ in $(c \leq X_1 \rightarrow \dots \rightarrow X_k \rightarrow T \wedge \llbracket \mathbf{v}_1 : X_1 \rrbracket \wedge \dots \wedge \llbracket \mathbf{v}_k : X_k \rrbracket)$ is satisfiable. We now reason by cases on c .

◦ *Case c is \hat{n} .* Then, $\Gamma_0(c)$ is int . Because the type constructors int and \rightarrow are incompatible with each other, this implies $k = 0$. Since \hat{n} is a constructor, the expression is a value.

◦ *Case c is $\hat{+}$.* We may assume $k \geq 2$, because otherwise the expression is a value. Then, $\Gamma_0(c)$ is $\text{int} \rightarrow \text{int} \rightarrow \text{int}$, so, by C-ARROW, the above constraint entails $\text{let } \Gamma_0; \text{ref } M \text{ in } (X_1 \leq \text{int} \wedge X_2 \leq \text{int} \wedge \llbracket v_1 : X_1 \rrbracket \wedge \llbracket v_2 : X_2 \rrbracket)$, which, by Lemma 10.4.5, entails $\text{let } \Gamma_0; \text{ref } M \text{ in } (\llbracket v_1 : \text{int} \rrbracket \wedge \llbracket v_2 : \text{int} \rrbracket)$. Thus, v_1 and v_2 are well-typed with type int . By the remark above, they must be integer literals \hat{n}_1 and \hat{n}_2 . As a result, the configuration is reducible by R-ADD.

10.7.5 SOLUTION: We must first ensure that R-REF, R-DEREF and R-ASSIGN respect \sqsubseteq (Definition 10.5.4).

◦ *Case R-REF.* The reduction is $\text{ref } v / \emptyset \rightarrow m / (m \mapsto v)$, where $m \notin \text{fpi}(v)$ (1). Let T be an arbitrary type. According to Definition 10.5.4, the goal is to show that there exist a set of type variables \tilde{Y} and a store type M' such that $\tilde{Y} \# \text{ftv}(T)$ and $\text{ftv}(M') \subseteq \tilde{Y}$ and $\text{dom}(M') = \{m\}$ and let Γ_0 in $\llbracket \text{ref } v : T \rrbracket$ entails $\exists \tilde{Y}. \text{let } \Gamma_0; \text{ref } M' \text{ in } \llbracket m / (m \mapsto v) : T / M' \rrbracket$. Now, we have

$$\begin{aligned} & \text{let } \Gamma_0 \text{ in } \llbracket \text{ref } v : T \rrbracket \\ \equiv & \exists Y. \text{let } \Gamma_0 \text{ in } (\text{ref } Y \leq T \wedge \llbracket v : Y \rrbracket) & (2) \\ \equiv & \exists Y. \text{let } \Gamma_0; \text{ref } M' \text{ in } (m \leq T \wedge \llbracket v : M'(m) \rrbracket) & (3) \\ \equiv & \exists Y. \text{let } \Gamma_0; \text{ref } M' \text{ in } \llbracket m / (m \mapsto v) : T / M' \rrbracket & (4) \end{aligned}$$

where (2) is by Exercise 10.7.1 and by C-INEX; (3) assumes M' is defined as $m \mapsto Y$, and follows from (1), C-INID and C-IN*; and (4) is by definition of constraint generation.

◦ *Case R-DEREF.* The reduction is $!m / (m \mapsto v) \rightarrow v / (m \mapsto v)$. Let T be an arbitrary type and let M be a store type of domain $\{m\}$. We have

$$\begin{aligned} & \text{let } \Gamma_0; \text{ref } M \text{ in } \llbracket !m / (m \mapsto v) : T / M \rrbracket \\ \equiv & \text{let } \Gamma_0; \text{ref } M \text{ in } \exists Y. (\text{ref } M(m) \leq \text{ref } Y \wedge Y \leq T \wedge \llbracket v : M(m) \rrbracket) & (1) \\ \equiv & \text{let } \Gamma_0; \text{ref } M \text{ in } \exists Y. (M(m) = Y \wedge Y \leq T \wedge \llbracket v : M(m) \rrbracket) & (2) \\ \equiv & \text{let } \Gamma_0; \text{ref } M \text{ in } (M(m) \leq T \wedge \llbracket v : M(m) \rrbracket) & (3) \\ \Vdash & \text{let } \Gamma_0; \text{ref } M \text{ in } (\llbracket v : T \rrbracket \wedge \llbracket v : M(m) \rrbracket) & (4) \\ \equiv & \text{let } \Gamma_0; \text{ref } M \text{ in } \llbracket v / (m \mapsto v) : T / M \rrbracket & (5) \end{aligned}$$

where (1) is by Exercise 10.7.1 and by C-INID; (2) follows from C-EXTRANS and from the fact that ref is an invariant type constructor; (3) is by C-NAMEEQ; (4) is by Lemma 10.4.5 and C-DUP; and (5) is again by definition of constraint generation.

◦ *Case R-ASSIGN.* The reduction is $m := v / (m \mapsto v_0) \rightarrow v / (m \mapsto v)$. Let T

be an arbitrary type and let M be a store type of domain $\{m\}$. We have

$$\begin{aligned}
& \text{let } \Gamma_0; \text{ref } M \text{ in } \llbracket m := v / (m \mapsto v_0) : T/M \rrbracket \\
\Vdash & \text{let } \Gamma_0; \text{ref } M \text{ in } \llbracket m := v : T \rrbracket & (1) \\
\equiv & \text{let } \Gamma_0; \text{ref } M \text{ in } \exists Z. (\text{ref } M(m) \leq \text{ref } Z \wedge \llbracket v : Z \rrbracket \wedge Z \leq T) & (2) \\
\equiv & \text{let } \Gamma_0; \text{ref } M \text{ in } \exists Z. (M(m) = Z \wedge Z \leq T \wedge \llbracket v : Z \rrbracket) & (3) \\
\equiv & \text{let } \Gamma_0; \text{ref } M \text{ in } (M(m) \leq T \wedge \llbracket v : M(m) \rrbracket) & (4) \\
\Vdash & \text{let } \Gamma_0; \text{ref } M \text{ in } \llbracket v / (m \mapsto v) : T/M \rrbracket & (5)
\end{aligned}$$

where (1) is by definition of constraint generation; (2) is by Exercise 10.7.1 and C-INID; (3) follows from the fact that `ref` is an invariant type constructor; (4) is by C-NAMEEQ; and (5) is obtained as in the previous case.

Second, we must check that if the configuration $c \ v_1 \ \dots \ v_k / \mu$ (where $k \geq 0$) is well-typed, then either it is reducible, or $c \ v_1 \ \dots \ v_k$ is a value. We only give a sketch of this proof; see the solution to Exercise 10.7.2 for details of a similar proof.

We begin by checking that every value that is well-typed with a type of the form `ref T` is a memory location. This assertion relies on the fact that the type constructor `ref` is isolated.

Next, we note that, according to the constraint generation rules, if the configuration $c \ v_1 \ \dots \ v_k / \mu$ is well-typed, then a constraint of the form $\text{let } \Gamma_0; \text{ref } M \text{ in } (c \leq X_1 \rightarrow \dots \rightarrow X_k \rightarrow T \wedge \llbracket v_1 : X_1 \rrbracket \wedge \dots \wedge \llbracket v_k : X_k \rrbracket)$ is satisfiable. We now reason by cases on c .

◦ *Case c is `ref`.* If $k = 0$, then the expression is a value; otherwise, it is reducible by R-REF.

◦ *Case c is `!`.* We may assume $k \geq 1$; otherwise the expression is a value. By definition of $\Gamma_0(!)$, the above constraint entails $\text{let } \Gamma_0; \text{ref } M \text{ in } \exists Y. (\text{ref } Y \rightarrow Y \leq X_1 \rightarrow \dots \rightarrow X_k \rightarrow T \wedge \llbracket v_1 : X_1 \rrbracket)$, which, by C-ARROW, Lemma 10.4.5, and C-INEX, entails $\exists Y. \text{let } \Gamma_0; \text{ref } M \text{ in } \llbracket v_1 : \text{ref } Y \rrbracket$. Thus, v_1 is well-typed with a type of the form `ref Y`. By the remark above, v_1 must be a memory location m . Furthermore, because every well-typed configuration is closed, m must be a member of $\text{dom}(\mu)$. As a result, the configuration `ref $v_1 \ \dots \ v_k / \mu$` is reducible by R-DEREF.

◦ *Case c is `:=`.* We may assume $k \geq 2$, because otherwise the expression is a value. As above, we check that v_1 must be a memory location and a member of $\text{dom}(\mu)$. Thus, the configuration is reducible by R-ASSIGN.

10.8.2 SOLUTION: The record access operation $\cdot.\langle \ell_b \rangle$ may be given the type scheme $\forall X_b. \{\ell_b : X_b\} \rightarrow X_b$. However, this type scheme isn't satisfactory, because it allows accessing ℓ_b *only* in records where ℓ_a and ℓ_c are undefined. The type scheme $\forall X_a X_b. \{\ell_a : X_a; \ell_b : X_b\} \rightarrow X_b$ is also a valid type scheme for $\cdot.\langle \ell_b \rangle$,

but allows accessing ℓ_b only in records where ℓ_a is defined and ℓ_c is not. To sum up, a satisfactory description of $\cdot.\langle\ell_b\rangle$ requires a whole *family* of type schemes, none of which is principal (more general than the others). A similar problem arises with record extension $\langle\cdot$ with $\ell_b = \cdot\rangle$.

A potential solution is to equip record types with a subtyping relationship, so that (say) both $\{\ell_a : \tau_a; \ell_b : \tau_b\}$ and $\{\ell_a : \tau_a; \ell_b : \tau_b; \ell_c : \tau_c\}$ are subtypes of $\{\ell_b : \tau_b\}$. Then, $\forall X_b. \{\ell_b : X_b\} \rightarrow X_b$ becomes a satisfactory type scheme for the record access operation $\cdot.\langle\ell_b\rangle$. Indeed, the operation is now applicable to any record that admits a type of the form $\{\ell_b : \tau_b\}$, that is, thanks to subtyping, to any record where ℓ_b is defined, regardless of which other fields are defined.

However, this is only half a solution, because there still is a problem with record extension. The type scheme $\forall X_b. \{\ell_b : X_b\} \rightarrow \{\ell_b : X_b\}$ is valid, and makes record extension applicable to any record where ℓ_b is defined, which is good. The trouble is with its return type: it states that only ℓ_b may be safely assumed to be defined in the new record. In other words, it causes static information about *all fields other than ℓ_b* to be lost. Addressing this dramatic loss of precision is one of the key motivations for introducing rows.

- 10.8.5 SOLUTION: We let the reader check that X must have kind $\star.Type$ and Y must have kind $\star.Row(\{\ell\})$. The type with all superscripts made explicit is

$$X \rightarrow^{Type} \Pi (\ell^{\star, Row(\emptyset)} : \text{int}^{Type} ; (Y \rightarrow^{Row(\{\ell\})} \partial^{\star, Row(\{\ell\})} X)).$$

In this case, because the type constructor Π occurs on the right-hand side of the toplevel arrow, it is possible to guess that the type must have kind $\star.Type$. There are cases where it is not possible to guess the kind of a type, because it may have several kinds; consider, for instance, ∂int .

- 10.8.27 SOLUTION: For the sake of generality, we perform the proof in the presence of subtyping, that is, we do not assume that subtyping is interpreted as equality. We formulate some hypotheses about the interpretation of subtyping: the type constructors $(\ell : \cdot ; \cdot)$, ∂ , and Π must be covariant; the type constructors \rightarrow and Π must be isolated.

We begin with a preliminary fact: *if the domain of V is $\{\ell_1, \dots, \ell_n\}$, where $\ell_1 < \dots < \ell_n$, then the constraint let Γ_0 in $\llbracket \{V; v\} : T \rrbracket$ is equivalent to let Γ_0 in $\exists Z_1 \dots Z_n Z. (\bigwedge_{i=1}^n \llbracket V(\ell_i) : Z_i \rrbracket \wedge \llbracket v : Z \rrbracket \wedge \Pi (\ell_1 : Z_1; \dots; \ell_n : Z_n; \partial Z) \leq T$). We let the reader check this fact using the constraint generation rules, the definition of Γ_0 and rule C-INID, and the above covariance hypotheses. We note that, by C-ROW-LL, the above constraint is invariant under a permutation of the labels ℓ_1, \dots, ℓ_n , so the above fact still holds when the hypothesis $\ell_1 < \dots < \ell_n$ is removed.*

We now prove that rules R-UPDATE, R-ACCESS-1, and R-ACCESS-2 enjoy subject reduction (Definition 10.5.4). Because the store is not involved, the goal is to establish that $\text{let } \Gamma_0 \text{ in } \llbracket \mathfrak{t} : \mathbb{T} \rrbracket$ entails $\text{let } \Gamma_0 \text{ in } \llbracket \mathfrak{t}' : \mathbb{T} \rrbracket$, where \mathfrak{t} is the redex and \mathfrak{t}' is the reduct.

◦ *Case R-UPDATE.* We have:

$$\begin{aligned} & \text{let } \Gamma_0 \text{ in } \llbracket \{\mathbb{V}; \mathfrak{v}\} \text{ with } \ell = \mathfrak{v}' \rrbracket : \mathbb{T} \\ \equiv & \text{let } \Gamma_0 \text{ in } \exists X X' Y. (\llbracket \{\mathbb{V}; \mathfrak{v}\} : \Pi (\ell : X; Y) \rrbracket \wedge \llbracket \mathfrak{v}' : X' \rrbracket \wedge \Pi (\ell : X'; Y) \leq \mathbb{T}) \quad (1) \\ \equiv & \text{let } \Gamma_0 \text{ in } \exists X X' Y Z_1 \dots Z_n Z. (\bigwedge_{i=1}^n \llbracket \mathbb{V}(\ell_i) : Z_i \rrbracket \wedge \llbracket \mathfrak{v} : Z \rrbracket \quad (2) \\ & \wedge \Pi (\ell_1 : Z_1; \dots; \ell_n : Z_n; \partial Z) \leq \Pi (\ell : X; Y) \\ & \wedge \llbracket \mathfrak{v}' : X' \rrbracket \wedge \Pi (\ell : X'; Y) \leq \mathbb{T}) \end{aligned}$$

where (1) is by Exercise 10.7.1, and (2) follows from the preliminary fact and from C-EXAND, provided $\{\ell_1, \dots, \ell_n\}$ is the domain of \mathbb{V} . We now distinguish two subcases:

Subcase $\ell \in \text{dom}(\mathbb{V})$. We may assume, *w.l.o.g.*, that ℓ is ℓ_1 . Then, by our covariance hypotheses, the subconstraint in the second line of (2) entails $(\ell_2 : Z_2; \dots; \ell_n : Z_n; \partial Z) \leq Y$, which in turn entails $\Pi (\ell_1 : X'; \ell_2 : Z_2; \dots; \ell_n : Z_n; \partial Z) \leq \Pi (\ell : X'; Y)$. By transitivity of subtyping, the subconstraint in the second and third lines of (2) entails $\Pi (\ell_1 : X'; \ell_2 : Z_2; \dots; \ell_n : Z_n; \partial Z) \leq \mathbb{T}$. By this remark and by C-EX*, (2) entails

$$\text{let } \Gamma_0 \text{ in } \exists X' Z_2 \dots Z_n Z. (\llbracket \mathfrak{v}' : X' \rrbracket \wedge \bigwedge_{i=2}^n \llbracket \mathbb{V}(\ell_i) : Z_i \rrbracket \wedge \llbracket \mathfrak{v} : Z \rrbracket \wedge \Pi (\ell_1 : X'; \ell_2 : Z_2; \dots; \ell_n : Z_n; \partial Z) \leq \mathbb{T}) \quad (3)$$

which by our preliminary fact is precisely $\text{let } \Gamma_0 \text{ in } \llbracket \{\mathbb{V}[\ell \mapsto \mathfrak{v}']; \mathfrak{v}\} : \mathbb{T} \rrbracket$.

Subcase $\ell \notin \text{dom}(\mathbb{V})$. By C-ROW-DL and C-ROW-LL, the term $(\ell_1 : Z_1; \dots; \ell_n : Z_n; \partial Z)$ may be replaced with $(\ell : Z; \ell_1 : Z_1; \dots; \ell_n : Z_n; \partial Z)$. Thus, reasoning as in the previous subcase, we find that (2) entails

$$\text{let } \Gamma_0 \text{ in } \exists X' Z_1 \dots Z_n Z. (\llbracket \mathfrak{v}' : X' \rrbracket \wedge \bigwedge_{i=1}^n \llbracket \mathbb{V}(\ell_i) : Z_i \rrbracket \wedge \llbracket \mathfrak{v} : Z \rrbracket \wedge \Pi (\ell_1 : X'; \ell_1 : Z_1; \dots; \ell_n : Z_n; \partial Z) \leq \mathbb{T}) \quad (4)$$

which by our preliminary fact is precisely $\text{let } \Gamma_0 \text{ in } \llbracket \{\mathbb{V}[\ell \mapsto \mathfrak{v}']; \mathfrak{v}\} : \mathbb{T} \rrbracket$.

◦ *Cases R-ACCESS-1, R-ACCESS-2.* We have:

$$\begin{aligned} & \text{let } \Gamma_0 \text{ in } \llbracket \{\mathbb{V}; \mathfrak{v}\}. \{\ell\} : \mathbb{T} \rrbracket \\ \equiv & \text{let } \Gamma_0 \text{ in } \exists XY. (\llbracket \{\mathbb{V}; \mathfrak{v}\} : \Pi (\ell : X; Y) \rrbracket \wedge X \leq \mathbb{T}) \quad (1) \\ \equiv & \text{let } \Gamma_0 \text{ in } \exists XY Z_1 \dots Z_n Z. (\bigwedge_{i=1}^n \llbracket \mathbb{V}(\ell_i) : Z_i \rrbracket \wedge \llbracket \mathfrak{v} : Z \rrbracket \quad (2) \\ & \wedge \Pi (\ell_1 : Z_1; \dots; \ell_n : Z_n; \partial Z) \leq \Pi (\ell : X; Y) \\ & \wedge X \leq \mathbb{T}) \end{aligned}$$

where (1) is by Exercise 10.7.1, and (2) follows from the preliminary fact and from C-EXAND, provided $\{\ell_1, \dots, \ell_n\}$ is the domain of \mathbb{V} . We now distinguish two subcases:

Subcase $\ell \in \text{dom}(V)$, i.e., (R-ACCESS-1). We may assume, *w.l.o.g.*, that ℓ is ℓ_1 . Then, by our covariance hypotheses, the subconstraint in the second line of (2) entails $Z_1 \leq X$. By transitivity of subtyping, by Lemma 10.4.5, and by C-EX*, we find that (2) entails $\text{let } \Gamma_0 \text{ in } \llbracket V(\ell) : T \rrbracket$.

Subcase $\ell \notin \text{dom}(V)$, i.e., (R-ACCESS-2). By C-ROW-DL and C-ROW-LL, the term $(\ell_1 : Z_1; \dots; \ell_n : Z_n; \partial Z)$ may be replaced with $(\ell : Z; \ell_1 : Z_1; \dots; \ell_n : Z_n; \partial Z)$. Thus, reasoning as in the previous subcase, we find that (2) entails $\text{let } \Gamma_0 \text{ in } \llbracket v : T \rrbracket$.

Before attacking the proof of the progress property, let us briefly check that every value v that is well-typed with type ΠT must be a record value, that is, must be of the form $\{V; w\}$. Indeed, assume that $\text{let } \Gamma_0; \text{ref } M \text{ in } \llbracket v : \Pi T \rrbracket$ is satisfiable. Then, v cannot be a program variable, for a well-typed value must be closed. Furthermore, v cannot be a memory location m , because $\text{ref } M(m) \leq \Pi T$ is unsatisfiable: indeed, the type constructors ref and Π are incompatible (recall that Π is isolated). Similarly, v cannot be a partially applied constant or a λ -abstraction, because $T' \rightarrow T'' \leq \Pi T$ is unsatisfiable. Thus, v must be a fully applied constructor. Since the only constructors in the language are the record constructors $\{\}_L$, v must be a record value. (If there were other constructors in the language, they could be ruled out as well, provided their return types are incompatible with Π .)

We must now prove that if the configuration $c v_1 \dots v_k / \mu$ is well-typed, then either it is reducible, or $c v_1 \dots v_k$ is a value. By the well-typedness hypothesis, a constraint of the form $\text{let } \Gamma_0; \text{ref } M \text{ in } \llbracket c v_1 \dots v_k : T \rrbracket$ is satisfiable.

- *Case c is $\{\}_L$.* If k is less than or equal to $n + 1$, where n is the cardinal of L , then $c v_1 \dots v_k$ is a value. Otherwise, unfolding the above constraint, we find that it cannot be satisfiable, because Π and \rightarrow are incompatible; this yields a contradiction.

- *Case c is $\{\cdot \text{ with } \ell = \cdot\}$.* Analogous to the next case.

- *Case c is $\cdot.\{\ell\}$.* If $k = 0$, then $c v_1 \dots v_k$ is a value. Assume $k \geq 1$. Then, the constraint $\text{let } \Gamma_0; \text{ref } M \text{ in } \llbracket c v_1 : T \rrbracket$ is satisfiable. By Exercise 10.7.1, this implies that $\text{let } \Gamma_0; \text{ref } M \text{ in } \llbracket v_1 : \Pi(\ell : X; Y) \rrbracket$ is satisfiable. Thus, v_1 must be a record value, and the configuration is reducible by R-ACCESS-1 or R-ACCESS-2.

10.8.33 SOLUTION: To make extension strict, it suffices to restrict its binding in the initial environment Γ_0 , as follows:

$$\langle \cdot \text{ with } \ell = \cdot \rangle : \forall XY. \Pi(\ell : \text{abs}; Y) \rightarrow X \rightarrow \Pi(\ell : \text{pre } X; Y).$$

The new binding, which is less general than the former, requires the field ℓ to be absent in the input record. The operational semantics need not be modified, since strict extension coincides with free extension when it is defined.

Defining the operational semantics of (free) restriction is left to the reader. Its binding in the initial environment should be:

$$\cdot \setminus \langle \ell \rangle : \forall XY. \Pi (\ell : X ; Y) \rightarrow \Pi (\ell : \text{abs} ; Y)$$

In principle, there is no need to guess this binding: it may be discovered through the encoding of finite records in terms of full records (10.8.32). Strict restriction, which requires the field to be present in the input record, may be assigned the following type scheme:

$$\cdot \setminus \langle \ell \rangle : \forall XY. \Pi (\ell : \text{pre } X ; Y) \rightarrow \Pi (\ell : \text{abs} ; Y)$$

- 10.8.34 SOLUTION: The informal sentence “supplying a record with more fields in a context where a record with fewer fields is expected” may be understood as “providing an argument of type $\Pi (\ell : \text{pre } T ; T')$ to a function whose domain type is $\Pi (\ell : \text{abs} ; T')$,” or, more generally, as “writing a program whose well-typedness requires some constraint of the form $\Pi (\ell : \text{pre } T ; T') \leq \Pi (\ell : \text{abs} ; T')$ to be satisfiable.” Now, in a nonstructural subtyping order where $\text{pre} \leq \text{abs}$ holds, such a constraint is equivalent to true. On the opposite, if subtyping is interpreted as equality, then such a constraint is equivalent to false. In other words, it is the law $\text{pre } T \leq \text{abs} \equiv \text{true}$ that gives rise to width subtyping.

It is worth drawing a comparison with the way width subtyping is defined in type systems that do not have rows. In such type systems, a record type is of the form $\{\ell_1 : T_1 ; \dots ; \ell_n : T_n\}$. Let us forget about the types T_1, \dots, T_n , because they describe the contents of fields, not their presence, and are thus orthogonal to the issue at hand. Then, a record type is a set $\{\ell_1, \dots, \ell_n\}$, and width subtyping is obtained by letting subtyping coincide with (the reverse of) set containment. In a type system that exploits rows, on the other hand, a record type is a total mapping from row labels to either pre or abs . (Because we are ignoring T_1, \dots, T_n , let us temporarily imagine that pre is a nullary type constructor.) The above record type is then written $\{\ell_1 : \text{pre} ; \dots ; \ell_n : \text{pre} ; \partial \text{abs}\}$. In other words, a set is now encoded as its characteristic function. Width subtyping is obtained by letting $\text{pre} \leq \text{abs}$ and by lifting this ordering, pointwise, to rows (which corresponds to our convention that rows are covariant).

References

- Abadi, Martín, Luca Cardelli, Pierre-Louis Curien, and Jean-Jacques Lévy. Explicit substitutions. *Journal of Functional Programming*, 1(4):375–416, 1991. Summary in *ACM Symposium on Principles of Programming Languages (POPL)*, San Francisco, California, 1990.
- Adams, Rolf, Walter Tichy, and Annette Weinert. The cost of selective recompilation and environment processing. *ACM Transactions on Software Engineering and Methodology*, 3(1):3–28, January 1994.
- Ahmed, Amal, Limin Jia, and David Walker. Reasoning about hierarchical storage. In *IEEE Symposium on Logic in Computer Science (LICS)*, Ottawa, Canada, pages 33–44, June 2003.
- Ahmed, Amal and David Walker. The logical approach to stack typing. In *ACM SIGPLAN Workshop on Types in Language Design and Implementation (TLDI)*, New Orleans, Louisiana, pages 74–85, January 2003.
- Aho, Alfred V., Ravi Sethi, and Jeffrey D. Ullman. *Compilers: Principles, Techniques, and Tools*. Addison-Wesley, Reading, Massachusetts, 1986.
- Aiken, Alexander, Manuel Fähndrich, and Raph Levien. Better static memory management: Improving region-based analysis of higher-order languages. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, La Jolla, California, pages 174–185, June 1995.
- Aiken, Alexander, Jeffrey S. Foster, John Kodumal, and Tachio Terauchi. Checking and inferring local non-aliasing. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, San Diego, California, pages 129–140, June 2003.
- Aiken, Alexander and Edward L. Wimmers. Solving systems of set constraints. In *IEEE Symposium on Logic in Computer Science (LICS)*, Santa Cruz, California, pages 329–340, June 1992.

- Aiken, Alexander and Edward L. Wimmers. Type inclusion constraints and type inference. In *ACM Symposium on Functional Programming Languages and Computer Architecture (FPCA), Copenhagen, Denmark*, pages 31–41, June 1993.
- Altenkirch, Thorsten. *Constructions, Inductive Types and Strong Normalization*. PhD thesis, Laboratory for Foundations of Computer Science, University of Edinburgh, Edinburgh, Scotland, 1993.
- Amadio, Roberto M. and Luca Cardelli. Subtyping recursive types. *ACM Transactions on Programming Languages and Systems*, 15(4):575–631, 1993. Summary in *ACM Symposium on Principles of Programming Languages (POPL), Orlando, Florida*, pp. 104–118; also DEC/Compaq Systems Research Center Research Report number 62, August 1990.
- Amtoft, Torben, Flemming Nielson, and Hanne Riis Nielson. *Type and Effect Systems: Behaviours for Concurrency*. Imperial College Press, 1999.
- Ancona, Davide and Elena Zucca. A theory of mixin modules: Basic and derived operators. *Mathematical Structures in Computer Science*, 8(4):401–446, August 1998.
- Ancona, Davide and Elena Zucca. A calculus of module systems. *Journal of Functional Programming*, 12(2):91–132, March 2002.
- Appel, Andrew W. Foundational proof-carrying code. In *IEEE Symposium on Logic in Computer Science (LICS), Boston, Massachusetts*, pages 247–258, June 2001.
- Appel, Andrew W. and Amy P. Felty. A semantic model of types and machine instructions for proof-carrying code. In *ACM SIGPLAN–SIGACT Symposium on Principles of Programming Languages (POPL), Boston, Massachusetts*, pages 243–253, January 2000.
- Aspinall, David. Subtyping with singleton types. In *International Workshop on Computer Science Logic (CSL), Kazimierz, Poland*, volume 933 of *Lecture Notes in Computer Science*, pages 1–15. Springer-Verlag, September 1994.
- Aspinall, David and Martin Hofmann. Another type system for in-place update. In *European Symposium on Programming (ESOP), Grenoble, France*, volume 2305 of *Lecture Notes in Computer Science*, pages 36–52. Springer-Verlag, April 2002.
- Augustsson, Lennart. Cayenne—A language with dependent types. In *ACM SIGPLAN International Conference on Functional Programming (ICFP), Baltimore, Maryland*, pages 239–250, 1998.
- Baader, Franz and Jörg Siekmann. Unification theory. In D. M. Gabbay, C. J. Hogger, and J. A. Robinson, editors, *Handbook of Logic in Artificial Intelligence and Logic Programming*, volume 2, *Deduction Methodologies*, pages 41–125. Oxford University Press, 1994.
- Baker, Henry G. Lively linear Lisp—look ma, no garbage! *ACM SIGPLAN Notices*, 27(8):89–98, 1992.
- Barendregt, Henk P. *The Lambda Calculus*. North Holland, revised edition, 1984.
- Barendregt, Henk P. Introduction to generalized type systems. *Journal of Functional Programming*, 1(2):125–154, 1991.

- Barendregt, Henk P. Lambda calculi with types. In S. Abramsky, D. M. Gabbay, and T. Maibaum, editors, *Handbook of Logic in Computer Science*, volume 2, *Computational Structures*. Oxford University Press, 1992.
- Barendsen, Erik and Sjaak Smetsers. Conventional and uniqueness typing in graph rewrite systems. In *Foundations of Software Technology and Theoretical Computer Science (FSTTCS), Bombay, India*, volume 761 of *Lecture Notes in Computer Science*, pages 41–51. Springer-Verlag, December 1993.
- Barras, Bruno, Samuel Boutin, Cristina Cornes, Judicael Courant, Jean-Christophe Filliatre, Eduardo Gimenez, Hugo Herbelin, Gerard Huet, Cesar Munoz, Chetan Murthy, Catherine Parent, Christine Paulin-Mohring, Amokrane Saibi, and Benjamin Werner. The Coq proof assistant reference manual: Version 6.1. Technical Report RT-0203, INRIA, 1997.
- Bauer, Lujo, Andrew W. Appel, and Edward W. Felten. Mechanisms for secure modular programming in Java. Technical Report TR-603-99, Princeton University, 1999.
- Bellantoni, Stephan and Stephan Cook. A new recursion-theoretic characterization of polytime functions. *Computational Complexity*, 2(2):97–110, 1992.
- Bellantoni, Stephan, K.-H. Niggl, and H. Schwichtenberg. Higher type recursion, ramification and polynomial time. *Annals of Pure and Applied Logic*, 104:17–30, 2000.
- Berardi, Stefano. Towards a mathematical analysis of the Coquand-Huet calculus of constructions and the other systems in Barendregt's cube. Technical report, Department of Computer Science, CMU, and Dipartimento Matematica, Università di Torino, 1988.
- Berthomieu, Bernard. Tagged types: A theory of order sorted types for tagged expressions. Research Report 93083, LAAS, 7, avenue du Colonel Roche, 31077 Toulouse, France, March 1993.
- Berthomieu, Bernard and Camille le Monières de Sagazan. A calculus of tagged types, with applications to process languages. In *Workshop on Types for Program Analysis (TPA)*, informal proceedings, pages 1–15, May 1995.
- Biagioni, Edoardo, Nicholas Haines, Robert Harper, Peter Lee, Brian G. Milnes, and Eliot B. Moss. Signatures for a protocol stack: A systems application of Standard ML. In *ACM Symposium on Lisp and Functional Programming (LFP)*, Orlando, Florida, pages 55–64, June 1994.
- Bierman, G. M., A. M. Pitts, and C. V. Russo. Operational properties of Lily, a polymorphic linear lambda calculus with recursion. In *Workshop on Higher Order Operational Techniques in Semantics (HOOTS)*, Montréal, Québec, volume 41 of *Electronic Notes in Theoretical Computer Science*. Elsevier, September 2000.
- Birkedal, Lars and Robert W. Harper. Constructing interpretations of recursive types in an operational setting. *Information and Computation*, 155:3–63, 1999.
- Birkedal, Lars and Mads Tofte. A constraint-based region inference algorithm. *Theoretical Computer Science*, 258:299–392, 2001.

- Birkedal, Lars, Mads Tofte, and Magnus Vejlsturp. From region inference to von Neumann machines via region representation inference. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, St. Petersburg Beach, Florida, pages 171–183, 1996.
- Blume, Matthias. *The SML/NJ Compilation and Library Manager*, May 2002. Available from <http://www.smlnj.org/doc/CM/index.html>.
- Blume, Matthias and Andrew W. Appel. Hierarchical modularity. *ACM Transactions on Programming Languages and Systems*, 21(4):813–847, 1999.
- Bonniot, Daniel. Type-checking multi-methods in ML (a modular approach). In *International Workshop on Foundations of Object-Oriented Languages (FOOL)*, informal proceedings, January 2002.
- Bourdoncle, François and Stephan Merz. Type-checking higher-order polymorphic multi-methods. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, Paris, France, pages 302–315, January 1997.
- Bracha, Gilad and William R. Cook. Mixin-based inheritance. In *ACM SIGPLAN Conference on Object Oriented Programming: Systems, Languages, and Applications (OOPSLA)/European Conference on Object-Oriented Programming (ECOOP)*, Ottawa, Ontario, pages 303–311, October 1990.
- Brandt, Michael and Fritz Henglein. Coinductive axiomatization of recursive type equality and subtyping. In *International Conference on Typed Lambda Calculi and Applications (TLCA)*, Nancy, France, volume 1210 of *Lecture Notes in Computer Science*, pages 63–81. Springer-Verlag, April 1997. Full version in *Fundamenta Informaticae*, 33:309–338, 1998.
- Breazu-Tannen, Val, Thierry Coquand, Carl Gunter, and Andre Scedrov. Inheritance as implicit coercion. *Information and Computation*, 93(1):172–221, July 1991. Also in C. A. Gunter and J. C. Mitchell, editors, *Theoretical Aspects of Object-Oriented Programming: Types, Semantics, and Language Design*, MIT Press, 1994.
- Bruce, Kim B. Typing in object-oriented languages: Achieving expressibility and safety, 1995. Available through <http://www.cs.williams.edu/~kim>.
- Bruce, Kim B. *Foundations of Object-Oriented Languages: Types and Semantics*. MIT Press, 2002.
- Bruce, Kim B., Luca Cardelli, Giuseppe Castagna, the Hopkins Objects Group (Jonathan Eifrig, Scott Smith, Valery Trifonov), Gary T. Leavens, and Benjamin Pierce. On binary methods. *Theory and Practice of Object Systems*, 1(3):221–242, 1996.
- Bruce, Kim B., Luca Cardelli, and Benjamin C. Pierce. Comparing object encodings. In *International Symposium on Theoretical Aspects of Computer Software (TACS)*, September 1997. An earlier version was presented as an invited lecture at the Third International Workshop on Foundations of Object Oriented Languages (FOOL 3), July 1996; full version in *Information and Computation*, 155(1–2):108–133, 1999.

- de Bruijn, Nicolas G. A survey of the project AUTOMATH. In J. P. Seldin and J. R. Hindley, editors, *To H. B. Curry: Essays in Combinatory Logic, Lambda Calculus, and Formalism*, pages 589–606. Academic Press, 1980.
- Brus, Tom, Marko van Eekelen, Maarten van Leer, and Marinus Plasmeijer. Clean: A language for functional graph rewriting. In *ACM Symposium on Functional Programming Languages and Computer Architecture (FPCA), Portland, Oregon*, volume 274 of *Lecture Notes in Computer Science*, pages 364–384. Springer-Verlag, September 1987.
- Burstall, Rod and Butler Lampson. A kernel language for abstract data types and modules. In *International Symposium on Semantics of Data Types, Sophia-Antipolis, France*, volume 173 of *Lecture Notes in Computer Science*, pages 1–50. Springer-Verlag, June 1984.
- Burstall, Rod, David MacQueen, and Donald Sannella. HOPE: an experimental applicative language. In *ACM Symposium on Lisp and Functional Programming (LFP), Stanford, California*, pages 136–143, August 1980.
- Calcagno, Cristiano. Stratified operational semantics for safety and correctness of region calculus. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL), London, England*, pages 155–165, 2001.
- Calcagno, Cristiano, Simon Helsen, and Peter Thiemann. Syntactic type soundness results for the region calculus. *Information and Computation*, 173(2):199–221, 2002.
- Cardelli, Luca. A polymorphic λ -calculus with Type:Type. Research report 10, DEC/Compaq Systems Research Center, May 1986.
- Cardelli, Luca. Phase distinctions in type theory, 1988a. Manuscript, available from <http://www.luca.demon.co.uk>.
- Cardelli, Luca. Typechecking dependent types and subtypes. In *Foundations of Logic and Functional Programming, Trento, Italy*, (December, 1986), volume 306 of *Lecture Notes in Computer Science*, pages 45–57. Springer-Verlag, 1988b.
- Cardelli, Luca. Program fragments, linking, and modularization. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL), Paris, France*, pages 266–277, January 1997.
- Cardelli, Luca, James Donahue, Mick Jordan, Bill Kalsow, and Greg Nelson. The Modula-3 type system. In *Proceedings of the Sixteenth Annual ACM Symposium on Principles of Programming Languages*, pages 202–212, January 1989.
- Cardelli, Luca and Xavier Leroy. Abstract types and the dot notation. In *IFIP TC2 Working Conference on Programming Concepts and Methods*. North Holland, 1990. Also appeared as DEC/Compaq SRC technical report 56.
- Cardelli, Luca and Giuseppe Longo. A semantic basis for Quest. *Journal of Functional Programming*, 1(4):417–458, October 1991. Summary in *ACM Conference on Lisp and Functional Programming*, pp. 30–43, 1990. Also available as DEC/Compaq SRC Research Report 55, Feb. 1990.

- Cardelli, Luca and John Mitchell. Operations on records. *Mathematical Structures in Computer Science*, 1:3–48, 1991. Also in C. A. Gunter and J. C. Mitchell, editors, *Theoretical Aspects of Object-Oriented Programming: Types, Semantics, and Language Design*, MIT Press, 1994; available as DEC/Compaq Systems Research Center Research Report #48, August, 1989; and in the *Proceedings of Workshop on the Mathematical Foundations of Programming Semantics (MFPS)*, New Orleans, Louisiana, Springer LNCS, volume 442, pp. 22–52, 1989.
- Cartmell, John. Generalised algebraic theories and contextual categories. *Annals of Pure and Applied Logic*, 32:209–243, 1986.
- Cartwright, Robert and Mike Fagan. Soft typing. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, Toronto, Ontario, pages 278–292, June 1991.
- Cervesato, Iliano, Joshua S. Hodas, and Frank Pfenning. Efficient resource management for linear logic proof search. *Theoretical Computer Science*, 232(1–2):133–163, February 2000.
- Cervesato, Iliano and Frank Pfenning. A linear logical framework. *Information and Computation*, 179(1):19–75, November 2002.
- Chaki, Sagar, Sriram K. Rajamani, and Jakob Rehof. Types as models: Model checking message-passing programs. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, Portland, Oregon, pages 45–57, 2002.
- Chirimar, Jawahar, Carl A. Gunter, and Jon G. Riecke. Reference counting as a computational interpretation of linear logic. *Journal of Functional Programming*, 6(2): 195–244, March 1996.
- Christiansen, Morten Voetmann and Per Velschow. Region-based memory management in Java. Master's thesis, University of Copenhagen, Department of Computer Science, 1998.
- Church, Alonzo. *The Calculi of Lambda Conversion*. Princeton University Press, 1941.
- Church, Alonzo. The weak theory of implication. *Kontrolliertes Denken: Untersuchungen zum Logikkalkül und zur Logik der Einzelwissenschaften*, pages 22–37, 1951.
- Clement, Dominique, Joelle Despeyroux, Thierry Despeyroux, and Gilles Kahn. A simple applicative language: Mini-ML. In *ACM Symposium on Lisp and Functional Programming (LFP)*, Cambridge, Massachusetts, pages 13–27, August 1986.
- Colby, Christopher, Peter Lee, George C. Necula, Fred Blau, Mark Plesko, and Kenneth Cline. A certifying compiler for Java. *ACM SIGPLAN Notices*, 35(5):95–107, May 2000.
- Comon, Hubert. Constraints in term algebras (short survey). In *Conference on Algebraic Methodology and Software Technology (AMAST)*, June, 1993, Workshops in Computing, pages 97–108. Springer-Verlag, 1994.
- Constable, Robert L., Stuart F. Allen, Mark Bromley, Rance Cleaveland, James F. Cremer, Robert W. Harper, Douglas J. Howe, Todd B. Knoblock, Paul Mendler, Prakash Panangaden, James T. Sasaki, and Scott F. Smith. *Implementing Mathematics with the NuPRL Proof Development System*. Prentice-Hall, Englewood Cliffs, NJ, 1986.

- Coquand, Catarina. The AGDA proof system homepage, 1998. <http://www.cs.chalmers.se/~catarina/agda/>.
- Coquand, Thierry. An analysis of Girard's paradox. In *IEEE Symposium on Logic in Computer Science (LICS), Cambridge, Massachusetts*, pages 227-236, June 1986.
- Coquand, Thierry. An algorithm for testing conversion in type theory. In G. Huet and G. Plotkin, editors, *Logical Frameworks*, pages 255-279. Cambridge University Press, 1991.
- Coquand, Thierry. Pattern matching with dependent types. In *Workshop on Types for Proofs and Programs (TYPES), Båstad, Sweden*, informal proceedings. Available from <ftp://ftp.cs.chalmers.se/pub/cs-reports/baastad.92/proc.ps.Z>, June 1992.
- Coquand, Thierry and Gérard Huet. The calculus of constructions. *Information and Computation*, 76(2-3):95-120, February/March 1988.
- Coquand, Thierry, Randy Pollack, and Makoto Takeyama. A logical framework with dependently typed records. In *International Conference on Typed Lambda Calculi and Applications (TLCA), Valencia, Spain*, volume 2701 of *Lecture Notes in Computer Science*, pages 105-119. Springer-Verlag, June 2003.
- Courant, Judicaël. Strong normalization with singleton types. In *Workshop on Intersection Types and Related Systems (ITRS), Copenhagen, Denmark*, volume 70 of *Electronic Notes in Theoretical Computer Science*. Elsevier, July 2002.
- Crank, Erik and Matthias Felleisen. Parameter-passing and the lambda calculus. In *ACM Symposium on Principles of Programming Languages (POPL), Orlando, Florida*, pages 233-244, January 1991.
- Crary, Karl. A simple proof technique for certain parametricity results. In *ACM SIGPLAN International Conference on Functional Programming (ICFP), Paris, France*, pages 82-89, September 1999.
- Crary, Karl. Toward a foundational typed assembly language. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL), New Orleans, Louisiana*, pages 198-212, January 2003.
- Crary, Karl, Robert Harper, and Sidd Puri. What is a recursive module? In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, pages 50-63, May 1999.
- Crary, Karl, Stephanie Weirich, and Greg Morrisett. Intensional polymorphism in type-erasure semantics. In *ACM SIGPLAN International Conference on Functional Programming (ICFP), Baltimore, Maryland*, pages 301-312, 1998. Full version in *Journal of Functional Programming*, 12(6), Nov. 2002, pp. 567-600.
- Curtis, Pavel. *Constrained Quantification in Polymorphic Type Analysis*. PhD thesis, Cornell University, Ithaca, New York, February 1990.
- van Daalen, Diederik T. *The Language Theory of Automath*. PhD thesis, Technische Hogeschool Eindhoven, Eindhoven, The Netherlands, 1980.

- Damas, Luis and Robin Milner. Principal type schemes for functional programs. In *ACM Symposium on Principles of Programming Languages (POPL)*, Albuquerque, New Mexico, pages 207–212, 1982.
- Danvy, Olivier. Functional unparsing. *Journal of Functional Programming*, 8(6):621–625, 1998.
- DeLine, Rob and Manuel Fähndrich. Enforcing high-level protocols in low-level software. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, Snowbird, Utah, pages 59–69, June 2001.
- Donahue, James and Alan Demers. Data types are values. *ACM Transactions on Programming Languages and Systems*, 7(3):426–445, July 1985.
- Došen, Kosta. A historical introduction to substructural logics. In K. Došen and P. Schroeder-Heister, editors, *Substructural Logics*, pages 1–30. Oxford University Press, 1993.
- Dreyer, Derek, Karl Crary, and Robert Harper. A type system for higher-order modules. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, New Orleans, Louisiana, pages 236–249, New Orleans, January 2003.
- Dussart, Dirk, Fritz Henglein, and Christian Mossin. Polymorphic recursion and sub-type qualifications: Polymorphic binding-time analysis in polynomial time. In *International Symposium on Static Analysis (SAS)*, Paris, France, volume 983 of *Lecture Notes in Computer Science*, pages 118–135. Springer-Verlag, July 1995.
- Emms, Martin and Hans LeiSS. Extending the type checker for SML by polymorphic recursion—A correctness proof. Technical Report 96-101, Centrum für Informations- und Sprachverarbeitung, Universität München, 1996.
- Erhard, Thomas. A categorical semantics of constructions. In *IEEE Symposium on Logic in Computer Science (LICS)*, Edinburgh, Scotland, pages 264–273, July 1988.
- Fähndrich, Manuel. *BANE: A Library for Scalable Constraint-Based Program Analysis*. PhD thesis, University of California at Berkeley, Berkeley, California, 1999.
- Fähndrich, Manuel and Rob DeLine. Adoption and focus: Practical linear types for imperative programming. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, Berlin, Germany, pages 13–24, June 2002.
- Fähndrich, Manuel, Jakob Rehof, and Manuvir Das. Scalable context-sensitive flow analysis using instantiation constraints. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, Vancouver, British Columbia, Canada, pages 253–263, June 2000.
- Felleisen, Matthias and Robert Hieb. A revised report on the syntactic theories of sequential control and state. *Theoretical Computer Science*, 103(2):235–271, 1992.
- Fisher, Kathleen and John H. Reppy. The design of a class mechanism for Moby. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, Atlanta, Georgia, pages 37–49, May 1999.

- Flanagan, Cormac and Shaz Qadeer. A type and effect system for atomicity. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI), San Diego, California*, pages 338–349, June 2003.
- Flatt, Matthew and Matthias Felleisen. Units: Cool modules for HOT languages. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI), Montréal, Québec*, pages 236–248, 1998.
- Fluet, Matthew. Monadic regions. In *Workshop on Semantics, Program Analysis and Computing Environments for Memory Management (SPACE)*, informal proceedings, January 2004.
- Fluet, Matthew and Riccardo Pucella. Phantom types and subtyping. In *IFIP International Conference on Theoretical Computer Science (TCS)*, pages 448–460, August 2002.
- Foster, Jeffrey S., Tachio Terauchi, and Alex Aiken. Flow-sensitive type qualifiers. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI), Berlin, Germany*, pages 1–12, June 2002.
- Frey, Alexandre. Satisfying subtype inequalities in polynomial space. In *International Symposium on Static Analysis (SAS), Paris, France*, volume 1302 of *Lecture Notes in Computer Science*, pages 265–277. Springer-Verlag, September 1997.
- Fuh, You-Chin and Prateek Mishra. Type inference with subtypes. In *European Symposium on Programming (ESOP), Nancy, France*, volume 300 of *Lecture Notes in Computer Science*, pages 94–114. Springer-Verlag, March 1988.
- Furuse, Jun P. and Jacques Garrigue. A label-selective lambda-calculus with optional arguments and its compilation method. RIMS Preprint 1041, Kyoto University, October 1995.
- Garcia, Ronald, Jaakko Jarvi, Andrew Lumsdaine, Jeremy Siek, and Jeremiah Willcock. A comparative study of language support for generic programming. In *ACM SIGPLAN Conference on Object Oriented Programming: Systems, Languages, and Applications (OOPSLA), Anaheim, California*, pages 115–134, October 2003.
- Garrigue, Jacques. Programming with polymorphic variants. In *ACM SIGPLAN Workshop on ML*, informal proceedings, September 1998.
- Garrigue, Jacques. Code reuse through polymorphic variants. In *Workshop on Foundations of Software Engineering (FOSE)*, November 2000.
- Garrigue, Jacques. Simple type inference for structural polymorphism. In *International Workshop on Foundations of Object-Oriented Languages (FOOL)*, informal proceedings, January 2002.
- Garrigue, Jacques. Relaxing the value restriction. In *International Symposium on Functional and Logic Programming (FLOPS), Nara, Japan*, volume 2998 of *Lecture Notes in Computer Science*, pages 196–213. Springer-Verlag, April 2004.
- Garrigue, Jacques and Hassan Ait-Kaci. The typed polymorphic label-selective lambda-calculus. In *ACM SIGPLAN–SIGACT Symposium on Principles of Programming Languages (POPL), Portland, Oregon*, pages 35–47, 1994.

- Garrigue, Jacques and Didier Rémy. Extending ML with semi-explicit higher-order polymorphism. *Information and Computation*, 155(1):134–169, 1999.
- Gaster, Benedict R. *Records, variants and qualified types*. PhD thesis, University of Nottingham, Nottingham, England, July 1998.
- Gaster, Benedict R. and Mark P. Jones. A polymorphic type system for extensible records and variants. Technical Report NOTTCS-TR-96-3, Department of Computer Science, University of Nottingham, November 1996.
- Gay, David and Alexander Aiken. Language support for regions. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI), Snowbird, Utah*, pages 70–80, June 2001.
- Ghelli, Giorgio and Benjamin Pierce. Bounded existentials and minimal typing, 1992. Circulated in manuscript form. Full version in *Theoretical Computer Science*, 193(1–2):75–96, February 1998.
- Gifford, David K. and John M. Lucassen. Integrating functional and imperative programming. In *ACM Symposium on Lisp and Functional Programming (LFP), Cambridge, Massachusetts*, pages 28–38, August 1986.
- Girard, Jean-Yves. *Interprétation fonctionnelle et élimination des coupures de l'arithmétique d'ordre supérieur*. Thèse d'état, University of Paris VII, 1972. Summary in J. E. Fenstad, editor, *Scandinavian Logic Symposium*, pp. 63–92, North-Holland, 1971.
- Girard, Jean-Yves. Linear logic. *Theoretical Computer Science*, 50:1–102, 1987.
- Girard, Jean-Yves. Light linear logic. *Information and Computation*, 143:175–204, 1998.
- Girard, Jean-Yves, Yves Lafont, and Paul Taylor. *Proofs and Types*, volume 7 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 1989.
- Glew, Neal. Type dispatch for named hierarchical types. In *ACM SIGPLAN International Conference on Functional Programming (ICFP), Paris, France*, pages 172–182, 1999.
- GNU. GNU C library, version 2.2.5, 2001. Available from http://www.gnu.org/manual/glibc-2.2.5/html_mono/libc.html.
- Goguen, Healfdene. *A Typed Operational Semantics for Type Theory*. PhD thesis, LFCS, University of Edinburgh, Edinburgh, Scotland, 1994. Report ESC-LFCS-94-304.
- Gordon, Andrew D. Bisimilarity as a theory of functional programming. In *Workshop on the Mathematical Foundations of Programming Semantics (MFPS), New Orleans, Louisiana*, volume 1 of *Electronic Notes in Theoretical Computer Science*. Elsevier, April 1995.
- Gordon, Andrew D. Operational equivalences for untyped and polymorphic object calculi. In A. D. Gordon and A. M. Pitts, editors, *Higher-Order Operational Techniques in Semantics*, Publications of the Newton Institute, pages 9–54. Cambridge University Press, 1998.

- Gordon, Andrew D. and Alan Jeffrey. Authenticity by typing for security protocols. In *IEEE Computer Security Foundations Workshop (CSFW), Cape Breton, Nova Scotia*, pages 145–159, 2001a.
- Gordon, Andrew D. and Alan Jeffrey. Typing correspondence assertions for communication protocols. In *Workshop on the Mathematical Foundations of Programming Semantics (MFPS), Aarhus, Denmark*, volume 45 of *Electronic Notes in Theoretical Computer Science*, pages 379–409. Elsevier, May 2001b.
- Gordon, Andrew D. and Alan Jeffrey. Types and effects for asymmetric cryptographic protocols. In *IEE Computer Security Foundations Workshop (CSFW), Cape Breton, Nova Scotia*, pages 77–91, 2002.
- Gordon, Andrew D. and Don Syme. Typing a multi-language intermediate code. In *ACM SIGPLAN–SIGACT Symposium on Principles of Programming Languages (POPL), London, England*, pages 248–260, January 2001.
- Gordon, Michael J., Robin Milner, and Christopher P. Wadsworth. *Edinburgh LCF*, volume 78 of *Lecture Notes in Computer Science*. Springer-Verlag, 1979.
- Gough, John. *Compiling for the .NET Common Language Runtime*. .NET series. Prentice Hall, 2002.
- Grossman, Dan, Greg Morrisett, Trevor Jim, Michael Hicks, Yanling Wang, and James Cheney. Region-based memory management in Cyclone. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI), Berlin, Germany*, pages 282–293, 2002.
- Gustavsson, Jörgen and Josef Svenningsson. Constraint abstractions. In *Symposium on Programs as Data Objects (PADO), Aarhus, Denmark*, volume 2053 of *Lecture Notes in Computer Science*, pages 63–83. Springer-Verlag, May 2001.
- Hallenberg, Niels, Martin Elsman, and Mads Tofte. Combining region inference and garbage collection. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI), Berlin, Germany*, pages 141–152, June 2002.
- Hallgren, Thomas and Aarne Ranta. An extensible proof text editor (abstract). In *International Conference on Logic for Programming and Automated Reasoning (LPAR), Reunion Island*, volume 1955 of *Lecture Notes in Computer Science*, pages 70–84. Springer-Verlag, 2000.
- Hamid, Nadeem, Zhong Shao, Valery Trifonov, Stefan Monnier, and Zhaozhong Ni. A syntactic approach to foundational proof-carrying code. In *IEEE Symposium on Logic in Computer Science (LICS)*, pages 89–100, July 2002.
- Hanson, David R. Fast allocation and deallocation of memory based on object lifetimes. *Software—Practice and Experience*, 20(1):5–12, 1990.
- Hardin, Thérèse, Luc Maranget, and Bruno Pagano. Functional runtimes within the lambda-sigma calculus. *Journal of Functional Programming*, 8(2):131–172, March 1998.
- Harper, Robert, Furio Honsell, and Gordon Plotkin. A framework for defining logics. *Journal of the ACM*, 40(1):143–184, 1993. Summary in *IEEE Symposium on Logic in Computer Science (LICS), Ithaca, New York*, 1987.

- Harper, Robert and Mark Lillibridge. A type-theoretic approach to higher-order modules with sharing. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, Portland, Oregon, pages 123–137, January 1994.
- Harper, Robert and John C. Mitchell. On the type structure of Standard ML. *ACM Transactions on Programming Languages and Systems*, 15(2):211–252, April 1993. An earlier version appeared in *ACM Symposium on Principles of Programming Languages (POPL)*, San Diego, California, under the title “The Essence of ML” (Mitchell and Harper), 1988.
- Harper, Robert, John C. Mitchell, and Eugenio Moggi. Higher-order modules and the phase distinction. In *ACM Symposium on Principles of Programming Languages (POPL)*, San Francisco, California, pages 341–354, January 1990.
- Harper, Robert and Frank Pfenning. On equivalence and canonical forms in the LF type theory. *ACM Transactions on Computational Logic*, 2004. To appear. An earlier version is available as Technical Report CMU-CS-00-148, School of Computer Science, Carnegie Mellon University.
- Harper, Robert and Robert Pollack. Type checking with universes. *Theoretical Computer Science*, 89:107–136, 1991.
- Harper, Robert and Christopher Stone. A type-theoretic interpretation of Standard ML. In G. Plotkin, C. Stirling, and M. Tofte, editors, *Proof, Language and Interaction: Essays in Honour of Robin Milner*. MIT Press, 2000.
- Heintze, Nevin. Set based analysis of ML programs. In *ACM Symposium on Lisp and Functional Programming (LFP)*, Orlando, Florida, pages 306–317, June 1994.
- Heintze, Nevin. Control-flow analysis and type systems. In *International Symposium on Static Analysis (SAS)*, Glasgow, Scotland, volume 983 of *Lecture Notes in Computer Science*, pages 189–206. Springer-Verlag, 1995.
- Helsen, Simon and Peter Thiemann. Syntactic type soundness for the region calculus. In *Workshop on Higher Order Operational Techniques in Semantics (HOOTS)*, Montréal, Québec, volume 41(3) of *Electronic Notes in Theoretical Computer Science*, pages 1–20. Elsevier, September 2000.
- Helsen, Simon and Peter Thiemann. Polymorphic specialization for ML. *ACM Transactions on Programming Languages and Systems*, 26(4):652–701, July 2004.
- Henglein, Fritz. *Polymorphic Type Inference and Semi-Unification*. PhD thesis, Rutgers University, April 1989. Available as NYU Technical Report 443, May 1989, from New York University, Courant Institute of Mathematical Sciences, Department of Computer Science, 251 Mercer St., New York, NY 10012, USA.
- Henglein, Fritz. Type inference with polymorphic recursion. *ACM Transactions on Programming Languages and Systems*, 15(2):253–289, 1993.
- Henglein, Fritz, Henning Makhholm, and Henning Niss. A direct approach to control-flow sensitive region-based memory management. In *ACM SIGPLAN International Conference on Principles and Practice of Declarative Programming (PPDP)*, Firenze, Italy, pages 175–186, September 2001.

- Henglein, Fritz and Christian Mossin. Polymorphic binding-time analysis. In *European Symposium on Programming (ESOP), Edinburgh, Scotland*, volume 788 of *Lecture Notes in Computer Science*, pages 287–301. Springer-Verlag, April 1994.
- Hirschowitz, Tom and Xavier Leroy. Mixin modules in a call-by-value setting. In *European Symposium on Programming (ESOP), Grenoble, France*, pages 6–20, April 2002.
- Hoare, C. A. R. Proof of correctness of data representation. *Acta Informatica*, 1: 271–281, 1972.
- Hofmann, Martin. A mixed modal/linear lambda calculus with applications to bellantoni-cook safe recursion. In *International Workshop on Computer Science Logic (CSL), Aarhus, Denmark*, pages 275–294, August 1997a.
- Hofmann, Martin. Syntax and semantics of dependent types. In A. M. Pitts and P. Dybjer, editors, *Semantics and Logic of Computation*, pages 79–130. Cambridge University Press, 1997b.
- Hofmann, Martin. Linear types and non-size-increasing polynomial time computation. In *IEEE Symposium on Logic in Computer Science (LICS), Trento, Italy*, pages 464–473, June 1999.
- Hofmann, Martin. Safe recursion with higher types and BCK-algebra. *Annals of Pure and Applied Logic*, 104(1-3):113–166, 2000.
- Honsell, Furio, Ian A. Mason, Scott F. Smith, and Carolyn L. Talcott. A variable typed logic of effects. *Information and Computation*, 119(1):55–90, 1995.
- Howard, William A. Hereditarily majorizable functionals of finite type. In A. S. Troelstra, editor, *Metamathematical Investigation of Intuitionistic Arithmetic and Analysis*, volume 344 of *Lecture Notes in Mathematics*, pages 454–461. Springer-Verlag, Berlin, 1973.
- Howard, William A. The formulas-as-types notion of construction. In J. P. Seldin and J. R. Hindley, editors, *To H. B. Curry: Essays on Combinatory Logic, Lambda Calculus, and Formalism*, pages 479–490. Academic Press, 1980. Reprint of 1969 article.
- Howe, Douglas J. Proving congruence of bisimulation in functional programming languages. *Information and Computation*, 124(2):103–112, 1996.
- Huet, Gérard. *Résolution d'équations dans les langages d'ordre 1, 2, ..., ω* . Thèse de Doctorat d'Etat, Université de Paris 7, Paris, France, 1976.
- Igarashi, Atsushi and Naoki Kobayashi. A generic type system for the Pi-calculus. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL), London, England*, pages 128–141, January 2001.
- Igarashi, Atsushi and Naoki Kobayashi. Resource usage analysis. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL), Portland, Oregon*, pages 331–342, January 2002.

- Igarashi, Atsushi and Benjamin C. Pierce. Foundations for virtual types. In *European Conference on Object-Oriented Programming (ECOOP), Lisbon, Portugal*, June 1999. Also in informal proceedings of the *Workshop on Foundations of Object-Oriented Languages (FOOL)*, January 1999. Full version in *Information and Computation*, 175(1): 34–49, May 2002.
- Ishtiaq, Samin and Peter O’Hearn. BI as an assertion language for mutable data structures. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL), London, England*, pages 14–26, January 2001.
- Jacobs, Bart. *Categorical Logic and Type Theory*. Studies in Logic and the Foundations of Mathematics 141. Elsevier, 1999.
- Jategaonkar, Lalita A. ML with extended pattern matching and subtypes. Master’s thesis, Massachusetts Institute of Technology, August 1989.
- Jategaonkar, Lalita A. and John C. Mitchell. ML with extended pattern matching and subtypes (preliminary version). In *ACM Symposium on Lisp and Functional Programming (LFP), Snowbird, Utah*, pages 198–211, Snowbird, Utah, July 1988.
- Jensen, Kathleen and Niklaus Wirth. *Pascal User Manual and Report*. Springer-Verlag, second edition, 1975.
- Jim, Trevor. What are principal typings and what are they good for? In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL), St. Petersburg Beach, Florida*, pages 42–53, 1996.
- Jim, Trevor, J. Greg Morrisett, Dan Grossman, Michael W. Hicks, James Cheney, and Yanling Wang. Cyclone: A safe dialect of C. In *General Track: USENIX Annual Technical Conference*, pages 275–288, June 2002.
- Jim, Trevor and Jens Palsberg. Type inference in systems of recursive types with subtyping, 1999. Manuscript, available from <http://www.cs.purdue.edu/homes/palsberg/draft/jim-palsberg99.pdf>.
- Johann, Patricia. A generalization of short-cut fusion and its correctness proof. *Higher-Order and Symbolic Computation*, 15(4):273–300, 2002.
- Jones, Mark P. *Qualified Types: Theory and Practice*. Cambridge University Press, 1994.
- Jones, Mark P. Using parameterized signatures to express modular structure. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL), St. Petersburg Beach, Florida*, January 21–24, 1996.
- Jones, Mark P. Typing Haskell in Haskell. In *ACM Haskell Workshop*, informal proceedings, October 1999.
- Jones, Mark P. and John C. Peterson. The Hugs 98 user manual, 1999. Available from <http://www.haskell.org/hugs/>.
- Jones, Mark P. and Simon Peyton Jones. Lightweight extensible records for Haskell. In *ACM Haskell Workshop*, informal proceedings, October 1999.

- Jouannaud, Jean-Pierre and Claude Kirchner. Solving equations in abstract algebras: a rule-based survey of unification. In J.-L. Lassez and G. Plotkin, editors, *Computational Logic: Essays in honor of Alan Robinson*, pages 257–321. MIT Press, 1991.
- Jouvelot, Pierre and David Gifford. Algebraic reconstruction of types and effects. In *ACM Symposium on Principles of Programming Languages (POPL)*, Orlando, Florida, pages 303–310, January 1991.
- Jouvelot, Pierre and David K. Gifford. Reasoning about continuations with control effects. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, Portland, Oregon, pages 218–226, June 1989.
- Jung, Achim and Allen Stoughton. Studying the fully abstract model of PCF within its continuous function model. In *International Conference on Typed Lambda Calculi and Applications (TLCA)*, Utrecht, The Netherlands, volume 664 of *Lecture Notes in Computer Science*, pages 230–244. Springer-Verlag, March 1993.
- Jutting, L.S. van Benthem, James McKinna, and Robert Pollack. Checking algorithms for Pure Type Systems. In *International Workshop on Types for Proofs and Programs (TYPES)*, Nijmegen, The Netherlands, May 1993, volume 806 of *Lecture Notes in Computer Science*, pages 19–61. Springer-Verlag, 1994.
- Kfoury, Assaf J., Jerzy Tiuryn, and Pawel Urzyczyn. ML typability is DEXPTIME-complete. In *Colloquium on Trees in Algebra and Programming (CAAP)*, Copenhagen, Denmark, volume 431 of *Lecture Notes in Computer Science*, pages 206–220. Springer-Verlag, May 1990.
- Kfoury, Assaf J., Jerzy Tiuryn, and Pawel Urzyczyn. The undecidability of the semi-unification problem. *Information and Computation*, 102(1):83–101, January 1993.
- Kfoury, Assaf J., Jerzy Tiuryn, and Pawel Urzyczyn. An analysis of ML typability. *Journal of the ACM*, 41(2):368–398, March 1994.
- Kirchner, Claude and Francis Klay. Syntactic theories and unification. In *IEEE Symposium on Logic in Computer Science (LICS)*, Philadelphia, Pennsylvania, pages 270–277, June 1990.
- Knight, Kevin. Unification: a multidisciplinary survey. *ACM Computing Surveys*, 21(1):93–124, March 1989.
- Kobayashi, Naoki. Quasi-linear types. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, San Antonio, Texas, pages 29–42, January 1999.
- Kozen, Dexter, Jens Palsberg, and Michael I. Schwartzbach. Efficient recursive subtyping. *Mathematical Structures in Computer Science*, 5(1):113–125, 1995.
- Kuncak, Viktor and Martin Rinard. Structural subtyping of non-recursive types is decidable. In *IEEE Symposium on Logic in Computer Science (LICS)*, Ottawa, Canada, pages 96–107, June 2003.
- Lafont, Yves. The linear abstract machine. *Theoretical Computer Science*, 59:157–180, 1988.

- Lambek, Joachim. The mathematics of sentence structure. *American Mathematical Monthly*, 65:154–170, 1958.
- Lampson, Butler and Rod Burstall. Pebble, a kernel language for modules and abstract data types. *Information and Computation*, 76:278–346, February/March 1988.
- Lassen, Søren Bøgh. *Relational Reasoning about Functions and Nondeterminism*. PhD thesis, Department of Computer Science, University of Aarhus, Aarhus, Denmark, 1998.
- Lassez, Jean-Louis, Michael J. Maher, and Kim G. Marriott. Unification revisited. In J. Minker, editor, *Foundations of Deductive Databases and Logic Programming*, pages 587–625. Morgan Kaufmann, 1988.
- Lee, Oukseh and Kwangkeun Yi. Proofs about a folklore let-polymorphic type inference algorithm. *ACM Transactions on Programming Languages and Systems*, 20(4):707–723, July 1998.
- Lervang, Daniel. Stratified functional programs and computational complexity. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, Charleston, South Carolina, pages 325–333, January 1993.
- Leroy, Xavier. Polymorphic typing of an algorithmic language. Research Report 1778, INRIA, October 1992.
- Leroy, Xavier. Manifest types, modules and separate compilation. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, Portland, Oregon, pages 109–122, January 1994.
- Leroy, Xavier. Applicative functors and fully transparent higher-order modules. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, San Francisco, California, pages 142–153, January 1995.
- Leroy, Xavier. A syntactic theory of type generativity and sharing. *Journal of Functional Programming*, 6(5):667–698, September 1996.
- Leroy, Xavier. The Objective Caml system: Documentation and user’s manual, 2000. With Damien Doligez, Jacques Garrigue, Didier Rémy, and Jérôme Vouillon. Available from <http://caml.inria.fr>.
- Leroy, Xavier and François Pessaux. Type-based analysis of uncaught exceptions. *ACM Transactions on Programming Languages and Systems*, 22(2):340–377, March 2000. Summary in *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, San Antonio, Texas, 1999.
- Lillibridge, Mark. *Translucent Sums: A Foundation for Higher-Order Module Systems*. PhD thesis, School of Computer Science, Carnegie Mellon University, Pittsburgh, Pennsylvania, May 1997.
- Lindholm, Tim and Frank Yellin. *The Java Virtual Machine Specification*. The Java Series. Addison-Wesley, Reading, MA, January 1997.
- Liskov, Barbara. A history of CLU. *ACM SIGPLAN Notices*, 28(3):133–147, 1993.

- Loader, Ralph. Finitary PCF is not decidable. *Theoretical Computer Science*, 266(1–2): 341–364, September 2001.
- Lucassen, John M. *Types and Effects towards the Integration of Functional and Imperative Programming*. PhD thesis, Massachusetts Institute of Technology, Cambridge, Massachusetts, August 1987. Technical Report MIT-LCS-TR-408.
- Lucassen, John M. and David K. Gifford. Polymorphic effect systems. In *ACM Symposium on Principles of Programming Languages (POPL)*, San Diego, California, pages 47–57, 1988.
- Luo, Zhaohui. *Computation and Reasoning: A Type Theory for Computer Science*. Number 11 in International Series of Monographs on Computer Science. Oxford University Press, 1994.
- Luo, Zhaohui and Robert Pollack. The LEGO proof development system: A user's manual. Technical Report ECS-LFCS-92-211, University of Edinburgh, May 1992.
- MacQueen, David. Modules for Standard ML. In *ACM Symposium on Lisp and Functional Programming (LFP)*, Austin, Texas, pages 198–207, 1984.
- MacQueen, David. Using dependent types to express modular structure. In *ACM Symposium on Principles of Programming Languages (POPL)*, St. Petersburg Beach, Florida, pages 277–286, January 1986.
- MacQueen, David B. and Mads Tofte. A semantics for higher-order functors. In *European Symposium on Programming (ESOP)*, Edinburgh, Scotland, volume 788 of *Lecture Notes in Computer Science*, pages 409–423. Springer-Verlag, April 1994.
- Magnusson, Lena and Bengt Nordström. The ALF proof editor and its proof engine. In *International Workshop on Types for Proofs and Programs (TYPES)*, Nijmegen, The Netherlands, May, 1993, volume 806 of *Lecture Notes in Computer Science*, pages 213–237. Springer-Verlag, 1994.
- Mairson, Harry G., Paris C. Kanellakis, and John C. Mitchell. Unification and ML type reconstruction. In J.-L. Lassez and G. Plotkin, editors, *Computational Logic: Essays in Honor of Alan Robinson*, pages 444–478. MIT Press, 1991.
- Makholm, Henning. Region-based memory management in Prolog. Master's thesis, University of Copenhagen, Department of Computer Science, March 2000. Technical Report DIKU-TR-00/09.
- Makholm, Henning. *A Language-Independent Framework for Region Inference*. PhD thesis, University of Copenhagen, Department of Computer Science, Copenhagen, Denmark, 2003.
- Makholm, Henning and Kostis Sagonas. On enabling the WAM with region support. In *International Conference on Logic Programming (ICLP)*, volume 2401 of *Lecture Notes in Computer Science*, pages 163–178. Springer-Verlag, July 2002.
- Martelli, Alberto and Ugo Montanari. Unification in linear time and space: A structured presentation. Internal Report B76-16, Istituto di Elaborazione delle Informazioni, Consiglio Nazionale delle Ricerche, Pisa, July 1976.

- Martelli, Alberto and Ugo Montanari. An efficient unification algorithm. *ACM Transactions on Programming Languages and Systems*, 4(2):258–282, 1982.
- Martin-Löf, Per. *Intuitionistic Type Theory*. Bibliopolis, 1984.
- Mason, Ian A., Scott F. Smith, and Carolyn L. Talcott. From operational semantics to domain theory. *Information and Computation*, 128(1):26–47, 1996.
- Mason, Ian A. and Carolyn L. Talcott. Equivalence in functional languages with effects. *Journal of Functional Programming*, 1:287–327, 1991.
- McAllester, David. On the complexity analysis of static analyses. *Journal of the ACM*, 49(4):512–537, July 2002.
- McAllester, David. A logical algorithm for ML type inference. In *International Conference on Rewriting Techniques and Applications (RTA), Valencia, Spain*, volume 2706 of *Lecture Notes in Computer Science*, pages 436–451. Springer-Verlag, June 2003.
- McBride, Conor. *Dependently Typed Functional Programs and their Proofs*. PhD thesis, LFCS, University of Edinburgh, Edinburgh, Scotland, 2000.
- McBride, Conor and James McKinna. The view from the left. *Journal of Functional Programming*, 14(1):69–111, 2004.
- McKinna, James and Robert Pollack. Pure Type Systems formalized. In *International Conference on Typed Lambda Calculi and Applications (TLCA), Utrecht, The Netherlands*, volume 664 of *Lecture Notes in Computer Science*, pages 289–305. Springer-Verlag, March 1993.
- Melski, David and Thomas Reps. Interconvertibility of a class of set constraints and context-free language reachability. *Theoretical Computer Science*, 248(1–2), November 2000.
- Milner, Robin. A theory of type polymorphism in programming. *Journal of Computer and System Sciences*, 17:348–375, August 1978.
- Milner, Robin, Mads Tofte, and Robert Harper. *The Definition of Standard ML*. MIT Press, 1990.
- Milner, Robin, Mads Tofte, Robert Harper, and David MacQueen. *The Definition of Standard ML*, Revised edition. MIT Press, 1997.
- Minamide, Yasuhiko. A functional representation of data structures with a hole. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL), San Diego, California*, pages 75–84, January 1998.
- Minamide, Yasuhiko, Greg Morrisett, and Robert Harper. Typed closure conversion. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL), St. Petersburg Beach, Florida*, pages 271–283, January 1996.
- Miquel, Alexandre. *Le calcul des constructions implicite: syntaxe et sémantique*. PhD thesis, University Paris 7, Paris, France, 2001.
- Mitchell, John C. Coercion and type inference. In *ACM Symposium on Principles of Programming Languages (POPL), Salt Lake City, Utah*, pages 175–185, January 1984.

- Mitchell, John C. Representation independence and data abstraction. In *ACM Symposium on Principles of Programming Languages (POPL)*, St. Petersburg Beach, Florida, pages 263–276, January 1986.
- Mitchell, John C. On the equivalence of data representations. In V. Lifschitz, editor, *Artificial Intelligence and Mathematical Theory of Computation: Papers in Honor of John McCarthy*, pages 305–330. Academic Press, 1991a.
- Mitchell, John C. Type inference with simple subtypes. *Journal of Functional Programming*, 1(3):245–286, July 1991b.
- Mitchell, John C. *Foundations for Programming Languages*. MIT Press, 1996.
- Mitchell, John C. and Gordon D. Plotkin. Abstract types have existential types. *ACM Transactions on Programming Languages and Systems*, 10(3):470–502, 1988. Summary in *ACM Symposium on Principles of Programming Languages (POPL)*, New Orleans, Louisiana, 1985.
- Moggi, Eugenio. Computational lambda-calculus and monads. In *IEEE Symposium on Logic in Computer Science (LICS)*, Asilomar, California, pages 14–23, June 1989. Full version, titled *Notions of Computation and Monads*, in *Information and Computation*, 93(1), pp. 55–92, 1991.
- Moh, Shaw-Kwei. The deduction theorems and two new logical systems. *Methodos*, 2: 56–75, 1950.
- Mohring, Christine. Algorithm development in the calculus of constructions. In *IEEE Symposium on Logic in Computer Science (LICS)*, Cambridge, Massachusetts, pages 84–91, June 1986.
- Monnier, Stefan, Bratin Saha, and Zhong Shao. Principled scavenging. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, Snowbird, Utah, pages 81–91, June 2001.
- Morrisett, Greg, Karl Cray, Neal Glew, and David Walker. Stack-based typed assembly language. *Journal of Functional Programming*, 12(1):43–88, January 2002.
- Morrisett, Greg, David Walker, Karl Cray, and Neal Glew. From System-F to typed assembly language. *ACM Transactions on Programming Languages and Systems*, 21(3):527–568, May 1999.
- Mossin, Christian. *Flow Analysis of Typed Higher-Order Programs*. PhD thesis, University of Copenhagen, Department of Computer Science, Copenhagen, Denmark, 1997. Also available as Technical Report DIKU-TR-97/1.
- Müller, Martin. A constraint-based recast of ML-polymorphism. In *International Workshop on Unification*, June 1994. Also available as Technical Report 94-R-243, CRIN, Nancy, France.
- Müller, Martin. Notes on HM(X), August 1998. Available from <http://www.ps.uni-sb.de/~mmueller/papers/HMX.ps.gz>.
- Müller, Martin, Joachim Niehren, and Ralf Treinen. The first-order theory of ordering constraints over feature trees. *Discrete Mathematics and Theoretical Computer Science*, 4(2):193–234, 2001.

- Müller, Martin and Susumu Nishimura. Type inference for first-class messages with feature constraints. In *Asian Computer Science Conference (ASIAN), Manila, The Philippines*, volume 1538 of *Lecture Notes in Computer Science*, pages 169–187. Springer-Verlag, December 1998.
- Mycroft, Alan. Polymorphic type schemes and recursive definitions. In *International Symposium on Programming, Toulouse, France*, volume 167 of *Lecture Notes in Computer Science*, pages 217–228, Toulouse, France, April 1984. Springer-Verlag.
- Necula, George C. Proof-carrying code. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL), Paris, France*, pages 106–119, January 1997.
- Necula, George C. *Compiling with Proofs*. PhD thesis, Carnegie Mellon University, Pittsburgh, Pennsylvania, September 1998. Technical report CMU-CS-98-154.
- Necula, George C. Translation validation for an optimizing compiler. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI), Vancouver, British Columbia, Canada*, pages 83–94, June 2000.
- Necula, George C. and Peter Lee. Safe kernel extensions without run-time checking. In *USENIX Symposium on Operating Systems Design and Implementation (OSDI), Seattle, Washington*, pages 229–243, October 1996.
- Necula, George C. and Peter Lee. The design and implementation of a certifying compiler. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI), Montréal, Québec*, pages 333–344, June 1998a.
- Necula, George C. and Peter Lee. Efficient representation and validation of logical proofs. In *IEEE Symposium on Logic in Computer Science (LICS), Indianapolis, Indiana*, pages 93–104, June 1998b.
- Niehren, Joachim, Martin Müller, and Andreas Podelski. Inclusion constraints over non-empty sets of trees. In *Theory and Practice of Software Development (TAPSOFT), Lille, France*, volume 1214 of *Lecture Notes in Computer Science*, pages 217–231. Springer-Verlag, April 1997.
- Niehren, Joachim and Tim Priesnitz. Non-structural subtype entailment in automata theory. *Information and Computation*, 186(2):319–354, November 2003.
- Nielson, Flemming and Hanne Riis Nielson. From CML to its process algebra. *Theoretical Computer Science*, 155:179–219, 1996.
- Nielson, Flemming, Hanne Riis Nielson, and Christopher L. Hankin. *Principles of Program Analysis*. Springer-Verlag, 1999.
- Nielson, Flemming, Hanne Riis Nielson, and Helmut Seidl. A succinct solver for ALFP. *Nordic Journal of Computing*, 9(4):335–372, 2002.
- Nielson, Hanne Riis and Flemming Nielson. Higher-order concurrent programs with finite communication topology. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL), Portland, Oregon*, pages 84–97, January 1994.

- Nishimura, Susumu. Static typing for dynamic messages. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, San Diego, California, pages 266–278, 1998.
- Niss, Henning. *Regions are Imperative: Unscoped Regions and Control-Flow Sensitive Memory Management*. PhD thesis, University of Copenhagen, Department of Computer Science, Copenhagen, Denmark, 2002.
- Nöcker, Erick and Sjaak Smetsers. Partially strict non-recursive data types. *Journal of Functional Programming*, 3(2):191–215, 1993.
- Nöcker, Erick G. M. H., Sjaak E. W. Smetsers, Marko C. J. D. van Eekelen, and Marinus J. Plasmeijer. Concurrent clean. In *Symposium on Parallel Architectures and Languages Europe, Volume I: Parallel Architectures and Algorithms (PARLE)*, Eindhoven, The Netherlands, volume 505 of *Lecture Notes in Computer Science*, pages 202–219. Springer-Verlag, June 1991.
- Odersky, Martin. Observers for linear types. In *European Symposium on Programming (ESOP)*, Rennes, France, volume 582 of *Lecture Notes in Computer Science*, pages 390–407. Springer-Verlag, February 1992.
- Odersky, Martin, Vincent Cremet, Christine Rockl, and Matthias Zenger. A nominal theory of objects with dependent types. In *International Workshop on Foundations of Object-Oriented Languages (FOOL)*, informal proceedings, 2003.
- Odersky, Martin, Martin Sulzmann, and Martin Wehr. Type inference with constrained types. *Theory and Practice of Object Systems*, 5(1):35–55, 1999. Summary in *International Workshop on Foundations of Object-Oriented Languages (FOOL)*, informal proceedings, 1997.
- O’Hearn, Peter. On bunched typing. *Journal of Functional Programming*, 13(4):747–796, 2003.
- O’Hearn, Peter and David Pym. The logic of bunched implications. *Bulletin of Symbolic Logic*, 5(2):215–244, 1999.
- Ohori, Atsushi. A polymorphic record calculus and its compilation. *ACM Transactions on Programming Languages and Systems*, 17(6):844–895, November 1995.
- Ohori, Atsushi and Peter Buneman. Type inference in a database programming language. In *ACM Symposium on Lisp and Functional Programming (LFP)*, Snowbird, Utah, pages 174–183, July 1988.
- Ohori, Atsushi and Peter Buneman. Static type inference for parametric classes. In *Conference on Object Oriented Programming: Systems, Languages, and Applications (OOPSLA)*, New Orleans, Louisiana, pages 445–456, October 1989. Also in C. A. Gunter and J. C. Mitchell, editors, *Theoretical Aspects of Object-Oriented Programming: Types, Semantics, and Language Design*, MIT Press, 1994.
- Orlov, Ivan E. The calculus of compatibility of propositions (in Russian). *Matematicheskii Sbornik*, 35:263–286, 1928.

- Owre, Sam, Sreeranga Rajan, John M. Rushby, Natarajan Shankar, and Mandayam K. Srivas. PVS: Combining specification, proof checking, and model checking. In *International Conference on Computer Aided Verification (CAV), New Brunswick, New Jersey*, volume 1102 of *Lecture Notes in Computer Science*, pages 411–414. Springer-Verlag, July 1996.
- Palsberg, Jens. Efficient inference of object types. *Information and Computation*, 123 (2):198–209, 1995.
- Palsberg, Jens. Type-based analysis and applications. In *ACM SIGPLAN-SIGSOFT Workshop on Program Analysis for Software Tools and Engineering (PASTE), Snowbird, Utah*, pages 20–27, June 2001.
- Palsberg, Jens and Patrick O’Keefe. A type system equivalent to flow analysis. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL), San Francisco, California*, pages 367–378, 1995.
- Palsberg, Jens and Michael Schwartzbach. Type substitution for object-oriented programming. In *ACM SIGPLAN Conference on Object Oriented Programming: Systems, Languages, and Applications (OOPSLA)/European Conference on Object-Oriented Programming (ECOOP), Ottawa, Ontario*, volume 25(10) of *ACM SIGPLAN Notices*, pages 151–160, October 1990.
- Palsberg, Jens and Michael I. Schwartzbach. *Object-Oriented Type Systems*. Wiley, 1994.
- Palsberg, Jens, Mitchell Wand, and Patrick M. O’Keefe. Type inference with non-structural subtyping. *Formal Aspects of Computing*, 9:49–67, 1997.
- Parnas, David. The criteria to be used in decomposing systems into modules. *Communications of the ACM*, 14(1):221–227, 1972.
- Paterson, Michael S. and Mark N. Wegman. Linear unification. *Journal of Computer and System Sciences*, 16:158–167, 1978.
- Paulin-Mohring, Christine. Extracting F_ω ’s programs from proofs in the calculus of constructions. In *ACM Symposium on Principles of Programming Languages (POPL), Austin, Texas*, pages 89–104, January 1989.
- Petersen, Leaf, Perry Cheng, Robert Harper, and Chris Stone. Implementing the TILT internal language. Technical Report CMU-CS-00-180, Department of Computer Science, Carnegie Mellon University, 2000.
- Petersen, Leaf, Robert Harper, Karl Crary, and Frank Pfenning. A type theory for memory allocation and data layout. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL), New Orleans, Louisiana*, pages 172–184, January 2003.
- Peyton Jones, Simon. Special issue: Haskell 98 language and libraries. *Journal of Functional Programming*, 13, January 2003.
- Pfenning, Frank and Rowan Davies. A judgmental reconstruction of modal logic. *Mathematical Structures in Computer Science*, 11(4):511–540, 2001.

- Pfenning, Frank and Carsten Schürmann. Algorithms for equality and unification in the presence of notational definitions. In T. Altenkirch, W. Naraschewski, and B. Reus, editors, *International Workshop on Types for Proofs and Programs (TYPES), Kloster Irsee, Germany*, volume 1657 of *Lecture Notes in Computer Science*. Springer-Verlag, 1998.
- Pierce, Benjamin C. *Types and Programming Languages*. MIT Press, 2002.
- Pierce, Benjamin C. and David N. Turner. Object-oriented programming without recursive types. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL), Charleston, South Carolina*, pages 299–312, January 1993.
- Pitts, Andrew M. Relational properties of domains. *Information and Computation*, 127:66–90, 1996.
- Pitts, Andrew M. Existential types: Logical relations and operational equivalence. In *International Colloquium on Automata, Languages and Programming (ICALP), Aalborg, Denmark*, volume 1443 of *Lecture Notes in Computer Science*, pages 309–326. Springer-Verlag, 1998.
- Pitts, Andrew M. Parametric polymorphism and operational equivalence. *Mathematical Structures in Computer Science*, 10:321–359, 2000.
- Pitts, Andrew M. Operational semantics and program equivalence. In G. Barthe, P. Dybjer, and J. Saraiva, editors, *Applied Semantics, Advanced Lectures*, volume 2395 of *Lecture Notes in Computer Science, Tutorial*, pages 378–412. Springer-Verlag, 2002.
- Pitts, Andrew M. and Ian D. B. Stark. Observable properties of higher order functions that dynamically create local names, or: What’s new? In *International Symposium on Mathematical Foundations of Computer Science, Gdańsk, Poland*, volume 711 of *Lecture Notes in Computer Science*, pages 122–141. Springer-Verlag, 1993.
- Pitts, Andrew M. and Ian D. B. Stark. Operational reasoning for functions with local state. In A. D. Gordon and A. M. Pitts, editors, *Higher-Order Operational Techniques in Semantics*, Publications of the Newton Institute, pages 227–273. Cambridge University Press, 1998.
- Plotkin, Gordon D. Lambda-definability and logical relations. Memorandum SAI-RM-4, University of Edinburgh, Edinburgh, Scotland, October 1973.
- Plotkin, Gordon D. LCF considered as a programming language. *Theoretical Computer Science*, 5:223–255, 1977.
- Plotkin, Gordon D. Lambda-definability in the full type hierarchy. In J. P. Seldin and J. R. Hindley, editors, *To H. B. Curry: Essays on Combinatory Logic, Lambda Calculus and Formalism*, pages 363–373. Academic Press, 1980.
- Plotkin, Gordon D. and Martin Abadi. A logic for parametric polymorphism. In *International Conference on Typed Lambda Calculi and Applications (TLCA), Utrecht, The Netherlands*, volume 664 of *Lecture Notes in Computer Science*, pages 361–375. Springer-Verlag, March 1993.

- Polakow, Jeff and Frank Pfenning. Natural deduction for intuitionistic non-commutative linear logic. In *International Conference on Typed Lambda Calculi and Applications (TLCA), L'Aquila, Italy*, volume 1581 of *Lecture Notes in Computer Science*, pages 295–309. Springer-Verlag, April 1999.
- Poll, Erik. Expansion Postponement for Normalising Pure Type Systems. *Journal of Functional Programming*, 8(1):89–96, 1998.
- Pollack, Robert. *The Theory of LEGO: A Proof Checker for the Extended Calculus of Constructions*. PhD thesis, University of Edinburgh, Edinburgh, Scotland, 1994.
- Popkorn, Sally. *First Steps in Modal Logic*. Cambridge University Press, 1994.
- Pottier, François. A versatile constraint-based type inference system. *Nordic Journal of Computing*, 7(4):312–347, November 2000.
- Pottier, François. A semi-syntactic soundness proof for HM(X). Research Report 4150, INRIA, March 2001a.
- Pottier, François. Simplifying subtyping constraints: a theory. *Information and Computation*, 170(2):153–183, November 2001b.
- Pottier, François. A constraint-based presentation and generalization of rows. In *IEEE Symposium on Logic in Computer Science (LICS), Ottawa, Canada*, pages 331–340, June 2003.
- Pottier, François and Vincent Simonet. Information flow inference for ML. *ACM Transactions on Programming Languages and Systems*, 25(1):117–158, January 2003.
- Pottier, François, Christian Skalka, and Scott Smith. A systematic approach to static access control. In *European Symposium on Programming (ESOP), Genova, Italy*, volume 2028 of *Lecture Notes in Computer Science*, pages 30–45. Springer-Verlag, April 2001.
- Pratt, Vaughan and Jerzy Tiuryn. Satisfiability of inequalities in a poset. *Fundamenta Informaticae*, 28(1-2):165–182, 1996.
- Pugh, William and Grant Weddell. Two-directional record layout for multiple inheritance. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI), White Plains, New York*, pages 85–91, June 1990.
- Rajamani, Sriram K. and Jakob Rehof. A behavioral module system for the pi-calculus. In *International Symposium on Static Analysis (SAS), Paris, France*, volume 2126 of *Lecture Notes in Computer Science*, pages 375–394. Springer-Verlag, July 2001.
- Rajamani, Sriram K. and Jakob Rehof. Conformance checking for models of asynchronous message passing software. In *International Conference on Computer Aided Verification (CAV), Copenhagen, Denmark*, pages 166–179, July 2002.
- Rehof, Jakob. Minimal typings in atomic subtyping. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL), Paris, France*, pages 278–291, January 1997.
- Rehof, Jakob and Manuel Fähndrich. Type-based flow analysis: From polymorphic subtyping to CFL reachability. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL), London, England*, pages 54–66, 2001.

- Reid, Alastair, Matthew Flatt, Leigh Stoller, Jay Lepreau, and Eric Eide. Knit: Component composition for systems software. In *USENIX Symposium on Operating Systems Design and Implementation (OSDI), San Diego, California*, pages 347–360, October 2000.
- Rémy, Didier. Typechecking records and variants in a natural extension of ML. In *ACM Symposium on Principles of Programming Languages (POPL), Austin, Texas*, pages 242–249, January 1989. Long version in C. A. Gunter and J. C. Mitchell, editors, *Theoretical Aspects of Object-Oriented Programming: Types, Semantics, and Language Design*, MIT Press, 1994.
- Rémy, Didier. *Algèbres Touffues. Application au Typage Polymorphe des Objets Enregistrés dans les Langages Fonctionnels*. PhD thesis, Université Paris VII, 1990.
- Rémy, Didier. Extending ML type system with a sorted equational theory. Research Report 1766, Institut National de Recherche en Informatique et Automatique, Rocquencourt, BP 105, 78 153 Le Chesnay Cedex, France, 1992a.
- Rémy, Didier. Projective ML. In *ACM Symposium on Lisp and Functional Programming (LFP), San Francisco, California*, pages 66–75, June 1992b.
- Rémy, Didier. Syntactic theories and the algebra of record terms. Research Report 1869, Institut National de Recherche en Informatique et Automatique, Rocquencourt, BP 105, 78 153 Le Chesnay Cedex, France, 1993.
- Rémy, Didier. Programming objects with ML-ART: An extension to ML with abstract and record types. In *International Symposium on Theoretical Aspects of Computer Software (TACS), Sendai, Japan*, volume 789 of *Lecture Notes in Computer Science*, pages 321–346. Springer-Verlag, April 1994.
- Rémy, Didier and Jérôme Vouillon. Objective ML: An effective object-oriented extension to ML. *Theory And Practice of Object Systems*, 4(1):27–50, 1998. Summary in *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL), Paris, France*, 1997.
- van Renesse, Robbert, Kenneth P. Birman, Mark Hayden, Alexey Vaysburd, and David Karr. Building adaptive systems using Ensemble. *Software: Practice and Experience*, 28(9):963–979, August 1998.
- Restall, Greg. *An Introduction to Substructural Logics*. Routledge, February 2000.
- Restall, Greg. Relevant and substructural logics. In D. Gabbay and J. Woods, editors, *Handbook of the History and Philosophy of Logic*, volume 6, *Logic and the Modalities in the Twentieth Century*. Elsevier, 2005. To appear.
- Reynolds, John C. Automatic computation of data set definitions. In *Information Processing 68, Edinburgh, Scotland*, volume 1, pages 456–461. North Holland, 1969.
- Reynolds, John C. Towards a theory of type structure. In *Colloque sur la Programmation, Paris, France*, volume 19 of *Lecture Notes in Computer Science*, pages 408–425. Springer-Verlag, 1974.

- Reynolds, John C. Syntactic control of interference. In *ACM Symposium on Principles of Programming Languages (POPL)*, Tucson, Arizona, pages 39–46, January 1978. Reprinted in O’Hearn and Tennent, *ALGOL-like Languages*, vol. 1, pages 273–286, Birkhäuser, 1997.
- Reynolds, John C. Types, abstraction, and parametric polymorphism. In R. E. A. Mason, editor, *Information Processing 83, Paris, France*, pages 513–523. Elsevier, 1983.
- Reynolds, John C. Syntactic control of interference, part 2. Report CMU-CS-89-130, Carnegie Mellon University, April 1989.
- Reynolds, John C. Intuitionistic reasoning about shared mutable data structure. In J. Davies, A. W. Roscoe, and J. Woodcock, editors, *Millennial Perspectives in Computer Science: Proceedings of the 1999 Oxford-Microsoft Symposium in honour of Sir Tony Hoare*. Palgrave Macmillan, 2000.
- Robinson, J. Alan. Computational logic: The unification computation. *Machine Intelligence*, 6:63–72, 1971.
- Ross, Douglas T. The AED free storage package. *Communications of the ACM*, 10(8): 481–492, 1967.
- Russo, Claudio V. *Types for Modules*. PhD thesis, Edinburgh University, Edinburgh, Scotland, 1998. LFCS Thesis ECS-LFCS-98-389.
- Russo, Claudio V. Non-dependent types for standard ML modules. In *ACM SIGPLAN International Conference on Principles and Practice of Declarative Programming (PPDP)*, Paris France, pages 80–97, September 1999.
- Russo, Claudio V. Recursive structures for Standard ML. In *ACM SIGPLAN International Conference on Functional Programming (ICFP)*, Firenze, Italy, pages 50–61, September 2001.
- Sabry, Amr. What is a purely functional language? *Journal of Functional Programming*, 8(1):1–22, January 1998.
- Saha, Bratin, Nevin Heintze, and Dino Oliva. Subtransitive CFA using types. Technical Report YALEU/DCS/TR-1166, Yale University, Department of Computer Science, October 1998.
- Sangiorgi, Davide and David. *The π -Calculus: a Theory of Mobile Processes*. Cambridge University Press, 2001.
- Sannella, Donald, Stefan Sokolowski, and Andrzej Tarlecki. Toward formal development of programs from algebraic specifications: Parameterisation revisited. *Acta Informatica*, 29(8):689–736, 1992.
- Schneider, Fred B. Enforceable security policies. *ACM Transactions on Information and System Security*, 3(1):30–50, February 2000.
- Schwartz, Jacob T. Optimization of very high level languages (parts I and II). *Computer Languages*, 1(2–3):161–194, 197–218, 1975.

- Seldin, Jonathan. Curry's anticipation of the types used in programming languages. In *Proceedings of the Annual Meeting of the Canadian Society for History and Philosophy of Mathematics, Toronto, Ontario*, pages 143–163, May 2002.
- Semmelroth, Miley and Amr Sabry. Monadic encapsulation in ML. In *ACM SIGPLAN International Conference on Functional Programming (ICFP), Paris, France*, pages 8–17, September 1999.
- Sestoft, Peter. Replacing function parameters by global variables. In *ACM Symposium on Functional Programming Languages and Computer Architecture (FPCA), London, England*, pages 39–53, September 1989. Also available as University of Copenhagen, Department of Computer Science Technical Report 88-7-2.
- Sestoft, Peter. Moscow ML homepage, 2003. <http://www.dina.dk/~sestoft/mosml.html>.
- Severi, Paula and Erik Poll. Pure type systems with definitions. In *International Symposium on Logical Foundations of Computer Science (LFCS), St. Petersburg, Russia*, volume 813 of *Lecture Notes in Computer Science*, pages 316–328. Springer-Verlag, September 1994.
- Shao, Zhong. An overview of the FLINT/ML compiler. In *ACM SIGPLAN Workshop on Types in Compilation (TIC), Amsterdam, The Netherlands*, June 1997.
- Shao, Zhong. Typed cross-module compilation. In *ACM SIGPLAN International Conference on Functional Programming (ICFP), Baltimore, Maryland*, pages 141–152, September 1998.
- Shao, Zhong. Transparent modules with fully syntactic signatures. In *ACM SIGPLAN International Conference on Functional Programming (ICFP), Paris, France*, pages 220–232, September 1999.
- Shao, Zhong, Christopher League, and Stefan Monnier. Implementing typed intermediate languages. In *ACM SIGPLAN International Conference on Functional Programming (ICFP), Baltimore, Maryland*, pages 313–323, September 1998.
- Shivers, Olin. Control flow analysis in Scheme. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI), Atlanta, Georgia*, pages 164–174, June 1988.
- Shivers, Olin. *Control-Flow Analysis of Higher-Order Languages or Taming Lambda*. PhD thesis, Carnegie Mellon University, Pittsburgh, Pennsylvania, May 1991.
- Simonet, Vincent. Type inference with structural subtyping: a faithful formalization of an efficient constraint solver. In *Asian Symposium on Programming Languages and Systems (APLAS), Beijing, China*, pages 283–302, November 2003.
- Skalka, Christian and François Pottier. Syntactic type soundness for $HM(X)$. In *Workshop on Types in Programming (TIP), Dagstuhl, Germany*, volume 75 of *Electronic Notes in Theoretical Computer Science*. Elsevier, July 2002.
- Smith, Frederick, David Walker, and Greg Morrisett. Alias types. In *European Symposium on Programming (ESOP), Berlin, Germany*, volume 1782 of *Lecture Notes in Computer Science*, pages 366–381. Springer-Verlag, April 2000.

- Smith, Geoffrey S. Principal type schemes for functional programs with overloading and subtyping. *Science of Computer Programming*, 23(2-3):197-226, December 1994.
- Smith, Jan, Bengt Nordström, and Kent Petersson. *Programming in Martin-Löf's Type Theory: An Introduction*. Oxford University Press, 1990.
- Statman, Richard. Logical relations and the typed λ -calculus. *Information and Control*, 65(2-3):85-97, May-June 1985.
- Steele, Guy L., Jr. *Common Lisp: The Language*. Digital Press, 1990.
- Stone, Christopher A. *Singleton Kinds and Singleton Types*. PhD thesis, Carnegie Mellon University, Pittsburgh, Pennsylvania, August 2000.
- Stone, Christopher A. and Robert Harper. Deciding type equivalence in a language with singleton kinds. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, Boston, Massachusetts, pages 214-227, January 2000.
- Stone, Christopher A. and Robert Harper. Extensional equivalence and singleton types. 2005. To appear.
- Streicher, Thomas. *Semantics of Type Theory*. Springer-Verlag, 1991.
- Su, Zhendong, Alexander Aiken, Joachim Niehren, Tim Priesnitz, and Ralf Treinen. The first-order theory of subtyping constraints. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, Portland, Oregon, pages 203-216, January 2002.
- Sulzmann, Martin. *A General Framework for Hindley/Milner Type Systems with Constraints*. PhD thesis, Yale University, Department of Computer Science, New Haven, Connecticut, May 2000.
- Sulzmann, Martin, Martin Müller, and Christoph Zenger. Hindley/Milner style type systems in constraint form. Research Report ACRC-99-009, University of South Australia, School of Computer and Information Science, July 1999.
- Sumii, Eijiro and Benjamin C. Pierce. A bisimulation for type abstraction and recursion. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, Long Beach, California, 2005.
- Sun. *Java™ 2 Platform Micro Edition (J2ME™) Technology for Creating Mobile Devices—White Paper*. Sun Microsystems, May 2000. Available from <http://java.sun.com/products/kvm/wp/KVMwp.pdf>.
- Tait, William W. Intensional interpretations of functionals of finite type I. *Journal of Symbolic Logic*, 32(2):198-212, June 1967.
- Talcott, C. Reasoning about functions with effects. In A. D. Gordon and A. M. Pitts, editors, *Higher Order Operational Techniques in Semantics*, Publications of the Newton Institute, pages 347-390. Cambridge University Press, 1998.
- Talpin, Jean-Pierre and Pierre Jouvelot. Polymorphic type, region and effect inference. *Journal of Functional Programming*, 2(2):245-271, 1992.

- Talpin, Jean-Pierre and Pierre Jouvelot. The type and effect discipline. *Information and Computation*, 111:245–296, 1994.
- Tarditi, David, Greg Morrisett, Perry Cheng, Christopher Stone, Robert Harper, and Peter Lee. TIL: A type-directed optimizing compiler for ML. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI), Philadelphia, Pennsylvania*, pages 181–192, May 1996.
- Tarjan, Robert Endre. Efficiency of a good but not linear set union algorithm. *Journal of the ACM*, 22(2):215–225, April 1975.
- Tarjan, Robert Endre. Applications of path compression on balanced trees. *Journal of the ACM*, 26(4):690–715, October 1979.
- Terlouw, J. Een nadere bewijstheoretische analyse van GSTTs. Manuscript, University of Nijmegen, Netherlands, 1989.
- Thorup, Kresten Krab. Genericity in Java with virtual types. In *European Conference on Object-Oriented Programming (ECOOP), Jyväskylä, Finland*, volume 1241 of *Lecture Notes in Computer Science*, pages 444–471. Springer-Verlag, June 1997.
- Tiuryn, Jerzy. Subtype inequalities. In *IEEE Symposium on Logic in Computer Science (LICS), Santa Cruz, California*, pages 308–317, June 1992.
- Tiuryn, Jerzy and Mitchell Wand. Type reconstruction with recursive types and atomic subtyping. In *Theory and Practice of Software Development (TAPSOFT), Orsay, France*, volume 668 of *Lecture Notes in Computer Science*, pages 686–701. Springer-Verlag, April 1993.
- Tofte, Mads. *Operational Semantics and Polymorphic Type Inference*. PhD thesis, Computer Science Department, Edinburgh University, Edinburgh, Scotland, 1988.
- Tofte, Mads and Lars Birkedal. A region inference algorithm. *ACM Transactions on Programming Languages and Systems*, 20(4):724–767, 1998.
- Tofte, Mads, Lars Birkedal, Martin Elsman, and Niels Hallenberg. Region-based memory management in perspective. In *ACM SIGPLAN International Conference on Principles and Practice of Declarative Programming (PPDP), Firenze, Italy*, pages 175–186, September 2001a.
- Tofte, Mads, Lars Birkedal, Martin Elsman, Niels Hallenberg, Tommy Højfeldt Olesen, and Peter Sestoft. Programming with regions in the ML Kit (for version 4). Technical report, IT University of Copenhagen, October 2001b.
- Tofte, Mads and Jean-Pierre Talpin. Implementing the call-by-value lambda-calculus using a stack of regions. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL), Portland, Oregon*, January 1994.
- Tofte, Mads and Jean-Pierre Talpin. Region-based memory management. *Information and Computation*, 132(2):109–176, February 1997.
- Torgersen, Mads. Virtual types are statically safe. In *International Workshop on Foundations of Object-Oriented Languages (FOOL), informal proceedings*, January 1998.

- Trifonov, Valery and Scott Smith. Subtyping constrained types. In *International Symposium on Static Analysis (SAS)*, Aachen, Germany, volume 1145 of *Lecture Notes in Computer Science*, pages 349–365. Springer-Verlag, September 1996.
- Turner, David N. and Philip Wadler. Operational interpretations of linear logic. *Theoretical Computer Science*, 227:231–248, 1999. Special issue on linear logic.
- Turner, David N., Philip Wadler, and Christian Mossin. Once upon a type. In *ACM Symposium on Functional Programming Languages and Computer Architecture (FPCA) San Diego, California*, pages 1–11, June 1995.
- Vouillon, Jerome and Paul-André Melliès. Semantic types: A fresh look at the ideal model for types. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, Venice, Italy, pages 52–63, 2004.
- Wadler, Philip. Theorems for free! In *ACM Symposium on Functional Programming Languages and Computer Architecture (FPCA)*, London, England, pages 347–359, September 1989.
- Wadler, Philip. Linear types can change the world. In *IFIP TC 2 Working Conference on Programming Concepts and Methods*, Sea of Galilee, Israel, pages 546–566, April 1990.
- Wadler, Philip. The marriage of effects and monads. *ACM Transactions on Computational Logic*, 4(1):1–32, 2003.
- Wahbe, Robert, Steven Lucco, Thomas E. Anderson, and Susan L. Graham. Efficient software-based fault isolation. In *ACM Symposium on Operating Systems Principles (SOSP)*, Asheville, North Carolina, pages 203–216, December 1993.
- Walker, David, Karl Cray, and Greg Morrisett. Typed memory management via static capabilities. *ACM Transactions on Programming Languages and Systems*, 22(4): 701–771, July 2000.
- Walker, David and Greg Morrisett. Alias types for recursive data structures. In *ACM SIGPLAN Workshop on Types in Compilation (TIC)*, Montréal, Québec, September, 2000, volume 2071, pages 177–206. Springer-Verlag, 2001.
- Walker, David and Kevin Watkins. On regions and linear types. In *ACM SIGPLAN International Conference on Functional Programming (ICFP)*, Firenze, Italy, pages 181–192, September 2001.
- Wand, Mitchell. Complete type inference for simple objects. In *IEEE Symposium on Logic in Computer Science (LICS)*, Ithaca, New York, pages 37–44, June 1987a.
- Wand, Mitchell. A simple algorithm and proof for type inference. *Fundamenta Informaticae*, 10:115–122, 1987b.
- Wand, Mitchell. Corrigendum: Complete type inference for simple objects. In *IEEE Symposium on Logic in Computer Science (LICS)*, Edinburgh, Scotland, page 132, 1988.
- Wand, Mitchell. Type inference for objects with instance variables and inheritance. In C. A. Gunter and J. C. Mitchell, editors, *Theoretical Aspects of Object-Oriented*

- Programming: Types, Semantics, and Language Design*, pages 97–120. MIT Press, 1994.
- Wang, Daniel C. and Andrew W. Appel. Type-preserving garbage collectors. In *ACM SIGPLAN–SIGACT Symposium on Principles of Programming Languages (POPL)*, London, England, pages 166–178, January 2001.
- Wansbrough, Keith and Simon Peyton Jones. Once upon a polymorphic type. In *ACM SIGPLAN–SIGACT Symposium on Principles of Programming Languages (POPL)*, San Antonio, Texas, pages 15–28, January 1999.
- Wells, Joe B. Typability and type checking in system F are equivalent and undecidable. *Annals of Pure and Applied Logic*, 98(1–3):111–156, 1999.
- Wells, Joe B. The essence of principal typings. In *International Colloquium on Automata, Languages and Programming (ICALP)*, volume 2380 of *Lecture Notes in Computer Science*, pages 913–925. Springer-Verlag, 2002.
- Werner, Benjamin. *Une Théorie des Constructions Inductives*. PhD thesis, Université Paris 7, Paris, France, May 1994.
- Wirth, Niklaus. *Systematic Programming: An Introduction*. Prentice Hall, 1973.
- Wirth, Niklaus. *Programming in Modula-2*. Texts and Monographs in Computer Science. Springer-Verlag, 1983.
- Wright, Andrew K. Simple imperative polymorphism. *Lisp and Symbolic Computation*, 8(4):343–355, 1995.
- Wright, Andrew K. and Robert Cartwright. A practical soft type system for Scheme. In *ACM Symposium on Lisp and Functional Programming (LFP)*, Orlando, Florida, pages 250–262, June 1994. Full version available in *ACM Transactions on Programming Languages and Systems*, 19(1):87–52, January 1997.
- Wright, Andrew K. and Matthias Felleisen. A syntactic approach to type soundness. *Information and Computation*, 115(1):38–94, November 1994.
- Xi, Hongwei. *Dependent Types in Practical Programming*. PhD thesis, Carnegie Mellon University, Pittsburgh, Pennsylvania, 1998.
- Xi, Hongwei and Robert Harper. A dependently typed assembly language. In *ACM SIGPLAN International Conference on Functional Programming (ICFP)*, Firenze, Italy, pages 169–180, September 2001.
- Xi, Hongwei and Frank Pfenning. Dependent types in practical programming. In *ACM SIGPLAN–SIGACT Symposium on Principles of Programming Languages (POPL)*, San Antonio, Texas, pages 214–227, January 1999.
- Zenger, Christoph. Indexed types. *Theoretical Computer Science*, 187:147–165, 1997.
- Zwanenburg, Jan. Pure type systems with subtyping. In *International Conference on Typed Lambda Calculi and Applications (TLCA)*, L'Aquila, Italy, volume 1581 of *Lecture Notes in Computer Science*, pages 381–396. Springer-Verlag, April 1999.