

Advanced Topics in Types and Programming Languages

Advanced Topics in
Types and Programming Lan-
guages

Benjamin C. Pierce, editor

The MIT Press
Cambridge, Massachusetts
London, England

Contents

Preface **1**

1 ML **2**

By François Pottier and Didier Rémy

| | | |
|------|---|-----|
| 1.1 | Preliminaries | 2 |
| 1.2 | What is ML? | 3 |
| 1.3 | Constraints | 20 |
| 1.4 | HM(X) | 40 |
| 1.5 | A purely constraint-based type system: PCB(X) | 52 |
| 1.6 | Constraint generation | 58 |
| 1.7 | Type soundness | 63 |
| 1.8 | Constraint solving | 72 |
| 1.9 | From ML-the-calculus to ML-the-programming-language | 86 |
| 1.10 | Universal quantification in constraints | 107 |
| 1.11 | Rows | 119 |

2 Logical Relations and a Case Study in Equivalence Checking
139

By Karl Cray

| | | |
|-----|-------------------------------|-----|
| 2.1 | The Equivalence Problem | 140 |
| 2.2 | Untyped Equivalence Checking | 141 |
| 2.3 | Type-Driven Equivalence | 143 |
| 2.4 | An Equivalence Algorithm | 144 |
| 2.5 | Completeness: A First Attempt | 147 |
| 2.6 | Logical Relations | 149 |
| 2.7 | A Monotone Logical Relation | 152 |
| 2.8 | The Main Lemma | 153 |

| | | |
|----------|---|------------|
| 2.9 | The Fundamental Theorem | 155 |
| 2.10 | Notes | 160 |
| 3 | <i>Typed Operational Reasoning</i> | 161 |
| | <i>By Andrew Pitts</i> | |
| 3.1 | Introduction | 161 |
| 3.2 | Motivating Examples | 163 |
| 3.3 | The Language | 169 |
| 3.4 | Contextual Equivalence | 175 |
| 3.5 | An Operationally-Based Logical Relation | 180 |
| 3.6 | Operational Extensionality | 187 |
| 3.7 | Notes | 193 |
| 4 | <i>Dependent Types</i> | 196 |
| | <i>By David Aspinall and Martin Hofmann</i> | |
| 4.1 | Pure first-order dependent types | 201 |
| 4.2 | Properties | 205 |
| 4.3 | Algorithmic typing and equality | 207 |
| 4.4 | Dependent sum types | 212 |
| 4.5 | The Calculus of Constructions | 214 |
| 4.6 | Relating abstractions: Pure Type Systems | 222 |
| 4.7 | Implementation | 224 |
| 4.8 | Further reading | 227 |
| 5 | <i>Effect Types and Region-based Memory Management</i> | 230 |
| | <i>By Fritz Henglein, Henning Makholm, and Henning Niss</i> | |
| 5.1 | Type-based program analysis | 230 |
| 5.2 | Value flow analysis | 231 |
| 5.3 | Effects | 241 |
| 5.4 | Region-based memory management | 245 |
| 5.5 | The Tofte-Talpin type system | 254 |
| 5.6 | Region inference | 263 |
| 5.7 | More powerful models for region-based memory management | 266 |
| 5.8 | Practical region-based memory management systems | 271 |
| 6 | <i>Substructural Type Systems</i> | 274 |
| | <i>By David Walker</i> | |
| 6.1 | Structural Properties | 275 |
| 6.2 | A Linear Type System | 277 |

| | | |
|----------|--|------------|
| 6.3 | Extensions and Variations | 287 |
| 6.4 | An Ordered Type System | 300 |
| 6.5 | Further Applications | 305 |
| 6.6 | Notes | 310 |
| 7 | <i>Proof-Carrying Code</i> | 313 |
| | <i>By George Necula</i> | |
| 7.1 | Overview of Proof Carrying Code | 314 |
| 7.2 | Formalizing the Safety Policy | 319 |
| 7.3 | Verification-Condition Generation | 323 |
| 7.4 | Soundness Proof | 336 |
| 7.5 | The Representation and Checking of Proofs | 340 |
| 7.6 | Proof Generation | 351 |
| 7.7 | PCC Beyond Types | 352 |
| 7.8 | Conclusion | 355 |
| 8 | <i>Typed Assembly Language</i> | 358 |
| | <i>By Greg Morrisett</i> | |
| 8.1 | TAL-0: Control-Flow-Safety | 359 |
| 8.2 | The TAL-0 Type System | 363 |
| 8.3 | TAL-1: Simple Memory-Safety | 372 |
| 8.4 | TAL-1 Changes to the Type System | 378 |
| 8.5 | Compiling to TAL-1 | 381 |
| 8.6 | Some Real World Issues | 384 |
| 8.7 | Scaling to Other Language Features | 386 |
| 8.8 | Conclusions | 392 |
| 9 | <i>Design Issues in Advanced Module Systems</i> | 393 |
| | <i>By Robert Harper and Benjamin C. Pierce</i> | |
| 9.1 | Basic Modularity | 394 |
| 9.2 | Phase Distinction and Phase Separation | 405 |
| 9.3 | Data Abstraction | 407 |
| 9.4 | Hierarchical Modularity | 416 |
| 9.5 | Families of Interfaces | 419 |
| 9.6 | Families of Modules | 423 |
| 9.7 | Advanced Topics | 433 |
| 9.8 | Relation to Existing Languages | 435 |
| 9.9 | History and Further Reading | 437 |
| 9.10 | Other stuff | 439 |

| | | |
|-----------|---|------------|
| 10 | <i>Type Definitions</i> | 440 |
| | <i>By Christopher A. Stone</i> | |
| 10.1 | Definitions in the Typing Context | 442 |
| 10.2 | Definitions in Modules | 457 |
| 10.3 | Singleton Kinds | 466 |
| 10.4 | Notes | 481 |
| A | <i>Notational Conventions</i> | 483 |
| A.1 | Metavariable Names | 483 |
| A.2 | Rule Naming Conventions | 483 |
| B | <i>Solutions to Selected Exercises</i> | 487 |
| | <i>References</i> | 536 |
| | <i>Index</i> | 565 |

1 *ML*

By François Pottier and Didier Rémy

1.1 Preliminaries

Names and renaming

Mathematicians and computer scientists use *names* to refer to arbitrary or unknown objects in the statement of a theorem, to refer to the parameters of a function, *etc.* Names are convenient because they are understandable by humans; nevertheless, they can be tricky. An in-depth treatment of the difficulties associated with names and renaming is beyond the scope of the present chapter: we encourage the reader to study Gabbay and Pitts’ excellent series of papers (Gabbay and Pitts, 2002; Pitts, 2002b). Here, we merely recall a few notions that are used throughout this chapter. Consider, for instance, an inductive definition of the abstract syntax of a simple programming language, the pure λ -calculus:

$$\mathfrak{t} ::= \mathfrak{z} \mid \lambda \mathfrak{z} . \mathfrak{t} \mid \mathfrak{t} \mathfrak{t}$$

Here, the *meta-variable* \mathfrak{z} ranges over an infinite set of *variables*—that is, names—while the meta-variable \mathfrak{t} ranges over *terms*. As usual in mathematics, we write “the variable \mathfrak{z} ” and “the term \mathfrak{t} ” instead of “the variable denoted by \mathfrak{z} ” and “the term denoted by \mathfrak{t} ”. The above definition states that a term may be a variable \mathfrak{z} , a pair of a variable and a term, written $\lambda \mathfrak{z} . \mathfrak{t}$, or a pair of terms, written $\mathfrak{t}_1 \mathfrak{t}_2$. However, this is not quite what we need. Indeed, according to

The (currently unfinished) code that accompanies this chapter may be found at <http://pauillac.inria.fr/~remy/mlrow/>. For space reasons, some material, including proofs, exercises, and more, has been left out of this version. In the future, a full version of this chapter that includes the missing material will be made available at the same address. In spite of these omissions, this chapter is still oversize with respect to Benjamin’s 100 page barrier: we currently have roughly 135 pages of text and 15 pages of solutions to exercises. We would appreciate comments and suggestions from the proofreaders as to how this chapter could be made shorter, without severely compromising its interest or readability.

this definition, the terms $\lambda z_1.z_1$ and $\lambda z_2.z_2$ are distinct, while we would like them to be a single mathematical object, because we intend $\lambda z.z$ to mean “the function that maps z to z ”—a meaning that is independent of the name z . To achieve this effect, we complete the above definition by stating that the construction $\lambda z.t$ *binds* z within t . One may also say that λz is a *binder* whose *scope* is t . Then, $\lambda z.t$ is no longer a pair: rather, it is an *abstraction* of the variable z within the term t . Abstractions have the property that the identity of the bound variable does not matter; that is, $\lambda z_1.z_1$ and $\lambda z_2.z_2$ are the same term. Informally, we say that terms are considered equal modulo α -*conversion*. Once the position and scope of binders are known, several standard notions follow, such as the set of *free variables* of a term t , written $fv(t)$, and the *capture-avoiding substitution* of a term t_1 for a variable z within a term t_2 , written $[z \mapsto t_1]t_2$. For conciseness, we write $fv(t_1, t_2)$ for $fv(t_1) \cup fv(t_2)$. A term is said to be *closed* when it has no free variables.

A *renaming* is a total bijective mapping from variables to variables that affects only a finite number of variables. The sole property of a variable is its identity, that is, the fact that it is distinct from other variables. As a result, at a global level, all variables are interchangeable: if a theorem holds in the absence of hypotheses about any particular variable, then any renaming of it holds as well. We often make use of this fact. When proving a theorem T , we say that a hypothesis H may be assumed *without loss of generality (w.l.o.g.)* if the theorem T follows from the theorem $H \Rightarrow T$ via a renaming argument, which is usually left implicit.

If \bar{z}_1 and \bar{z}_2 are sets of variables, we write $\bar{z}_1 \# \bar{z}_2$ as a shorthand for $\bar{z}_1 \cap \bar{z}_2 = \emptyset$, and say that \bar{z}_1 is *fresh* for \bar{z}_2 (or vice-versa). We say that \bar{z} is fresh for t if and only if $\bar{z} \# fv(t)$ holds.

In this chapter, we work with several distinct varieties of names: program variables, memory locations, and type variables, the latter of which may be further divided into *kinds*. We draw names of different varieties from disjoint sets, each of which is infinite.

1.2 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, 1979b). 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 crucially rely on first-order unification and can be made efficient in practice. In addition to Standard ML and Caml, this view encompasses programming languages such as Haskell (Hudak, Peyton Jones, Wadler, Boutel, Fairbairn, Fasel, Guzman, Hammond, Hughes, Johnsson, Kieburtz, Nikhil, Partain, and Peterson, 1992) and Clean (Brus, van Eekelen, van Leer, and Plasmeijer, 1987), whose semantics is rather different—indeed, it is pure and lazy—but whose type system fits this description. We refer to it as *ML-the-type-system*. It is also referred to as *Hindley and Milner’s type discipline* in the literature.

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, `let` declarations, 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 the `let` construct, which distinguishes ML-the-type-system from the simply-typed λ -calculus, may be understood as a simple textual expansion facility. However, this view only tells 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 naïve type inference algorithm whose time complexity is exponential, 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 implicit-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) usually involve heavy manipulations of type substitutions. Such an 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). 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) or polymorphic methods (Garrigue and Rémy, 1999).

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. To achieve these goals, 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. Such a decomposition allows reasoning separately about *when* a program is correct, on the one hand, and *how* to check whether it is correct, on the other hand. It has long been standard in the setting of the simply-typed λ -calculus (TAPL Chapter 22), but, to the best of our knowledge, has never been proposed for ML-the-type-system. Second, it is often natural to define and implement the solver as a constraint rewriting system. Then, 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 structures manipulated throughout the type inference process. For instance, describing unification in terms of *multi-equations* (Jouannaud and Kirchner, 1991) 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 describing many extensions of ML-the-type-system, among which extensions with recursive types, rows, subtyping, first-order unification under a mixed prefix, and more.

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

| | | | |
|-----------------------------------|-------------------------|---|-------------------------------------|
| $x, y ::=$ | <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> | | $c \in Q^+ \wedge k \leq a(c)$ |
| $t ::=$ | <i>Expressions:</i> | $c \ v_1 \ \dots \ v_k$ | <i>Partial application</i> |
| x | <i>Identifier</i> | | $c \in Q^- \wedge k < a(c)$ |
| $\lambda z. t$ | <i>Function</i> | $\mathcal{E} ::=$ | <i>Evaluation Contexts:</i> |
| $t \ t$ | <i>Application</i> | \square | <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 1-1: Syntax of ML-the-calculus

ML-the-calculus

The syntax of ML-the-calculus is defined in Figure 1-1. 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*—sometimes 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, λ_t stands for $\lambda z. t$, provided z is fresh for t . *Memory locations* are names that represent memory addresses. By convention, 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 Q . They are never subject to α -conversion. Program variables, memory locations, and constants belong to distinct syntactic classes and may never be confused.

The set of constants 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 Q is partitioned into subsets of *constructors* Q^+ and *destructors* Q^- . Constructors and destructors differ in that the former are used to *form* values, while the latter are used to

operate on values.

- 1.2.1 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 $\mathbf{t}_1 \hat{+} \mathbf{t}_2$ stands for the double application $\hat{+} \mathbf{t}_1 \mathbf{t}_2$ of the destructor $\hat{+}$ to the expressions \mathbf{t}_1 and \mathbf{t}_2 . \square

Expressions—also known as *program 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 include identifiers, λ -abstractions, applications, and local definitions. The λ -abstraction $\lambda \mathbf{z}. \mathbf{t}$ represents the function of one parameter named \mathbf{z} whose result is the expression \mathbf{t} , or, in other words, the function that maps \mathbf{z} to \mathbf{t} . Note that the variable \mathbf{z} is bound within the term \mathbf{t} , so (for instance) $\lambda \mathbf{z}_1. \mathbf{z}_1$ and $\lambda \mathbf{z}_2. \mathbf{z}_2$ are the same object. The *application* $\mathbf{t}_1 \mathbf{t}_2$ represents the result of calling the function \mathbf{t}_1 with actual parameter \mathbf{t}_2 , or, in other words, the result of applying \mathbf{t}_1 to \mathbf{t}_2 . Application is left-associative, that is, $\mathbf{t}_1 \mathbf{t}_2 \mathbf{t}_3$ stands for $(\mathbf{t}_1 \mathbf{t}_2) \mathbf{t}_3$. The construct $\text{let } \mathbf{z} = \mathbf{t}_1 \text{ in } \mathbf{t}_2$ represents the result of evaluating \mathbf{t}_2 after binding the variable \mathbf{z} to \mathbf{t}_1 . Note that the variable \mathbf{z} is bound within \mathbf{t}_2 , but not within \mathbf{t}_1 , so for instance $\text{let } \mathbf{z}_1 = \mathbf{z}_1 \text{ in } \mathbf{z}_1$ and $\text{let } \mathbf{z}_2 = \mathbf{z}_1 \text{ in } \mathbf{z}_2$ are the same object. The construct $\text{let } \mathbf{z} = \mathbf{t}_1 \text{ in } \mathbf{t}_2$ has the same meaning as $(\lambda \mathbf{z}. \mathbf{t}_2) \mathbf{t}_1$, 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 $\mathbf{c} \mathbf{v}_1 \dots \mathbf{v}_k$, where k does not exceed \mathbf{c} 's arity if \mathbf{c} is a constructor, and k is smaller than \mathbf{c} 's arity if \mathbf{c} is a destructor. In what follows, we are often interested in closed values, that is, values that do not contain any free program variables. We use the meta-variables \mathbf{v} and \mathbf{w} for values.

- 1.2.2 EXAMPLE: The integer literals $\dots, \widehat{-1}, \widehat{0}, \widehat{1}, \dots$ are nullary constructors, so they are values. Integer addition $\hat{+}$ is a binary destructor, so it is a value, and so is every partial application $\hat{+} \mathbf{v}$. Thus, both $\hat{+} \widehat{1}$ and $\hat{+} \hat{+}$ are values. An application of $\hat{+}$ to two values, such as $\widehat{2} \hat{+} \widehat{2}$, is not a value. \square
- 1.2.3 EXAMPLE [PAIRS]: Let (\cdot, \cdot) be a binary constructor. If \mathbf{t}_1 are \mathbf{t}_2 are expressions, then the double application $(\cdot, \cdot) \mathbf{t}_1 \mathbf{t}_2$ may be called the *pair* of \mathbf{t}_1 and \mathbf{t}_2 , and may be written $(\mathbf{t}_1, \mathbf{t}_2)$. By the definition above, $(\mathbf{t}_1, \mathbf{t}_2)$ is a value if and only if \mathbf{t}_1 and \mathbf{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 data structures,

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 $\text{dom}(\mu)$ for the domain of μ and $\text{range}(\mu)$ for the set of memory locations that appear in its codomain.

The operational semantics of a purely functional language, such as the pure 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 \mathfrak{t}/μ is a pair of an expression \mathfrak{t} and a store μ . The memory locations in the domain of μ are considered bound within \mathfrak{t}/μ , so (for instance) $m_1/(m_1 \mapsto \hat{0})$ and $m_2/(m_2 \mapsto \hat{0})$ are the same object. In what follows, we are often interested in *closed* configurations, that is, configurations \mathfrak{t}/μ such that \mathfrak{t} has no free program variables and every memory location that appears within \mathfrak{t} or within the range of μ is in the domain of μ . If \mathfrak{t} is a source program, its evaluation begins within an empty store—that is, with the configuration \mathfrak{t}/\emptyset . Because, by convention, source programs do not contain memory locations, this is a closed configuration. Furthermore, we shall see that all reducts of a closed configuration are closed as well. Please 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), is reminiscent of the encoding of reference cells in process calculi (Turner, 1995; Fournet and Gonthier, 1996).

A *context* is an expression where a single subexpression has been replaced with a *hole*, written \square . *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 $\lambda z.\square$ is not an evaluation context means that the body of a function is never evaluated—that is, not until the function is applied, see R-BETA below. The fact that $\mathfrak{t} \mathcal{E}$ is an evaluation context only if \mathfrak{t} is a value means that, to evaluate an application $\mathfrak{t}_1 \mathfrak{t}_2$, one should fully evaluate \mathfrak{t}_1 before attempting to evaluate \mathfrak{t}_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 $\mathcal{E} ::= \dots | \mathfrak{t} \mathcal{E} | \mathcal{E} v | \dots$ would enforce a right-to-left evaluation order, while defining $\mathcal{E} ::= \dots | \mathfrak{t} \mathcal{E} | \mathcal{E} \mathfrak{t} | \dots$ would leave the evaluation order unspecified, effectively allowing reduction to alternate between

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

Figure 1-2: Semantics of ML-the-calculus

both subexpressions, and making evaluation nondeterministic. The fact that $\text{let } z = \mathbf{v} \text{ in } \mathcal{E}$ 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, see R-LET below. We write $\mathcal{E}[\mathbf{t}]$ for the expression obtained by replacing the hole in \mathcal{E} with the expression \mathbf{t} .

Figure 1-2 defines first a relation \longrightarrow between configurations, then a relation \longrightarrow between *closed* configurations. If $\mathbf{t}/\mu \longrightarrow \mathbf{t}'/\mu'$ or $\mathbf{t}/\mu \xrightarrow{\delta} \mathbf{t}'/\mu'$ holds, then we say that the configuration \mathbf{t}/μ *reduces* to the configuration \mathbf{t}'/μ' ; the ambiguity involved in this definition is benign. If $\mathbf{t}/\mu \longrightarrow \mathbf{t}'/\mu$ holds for every store μ , then we write $\mathbf{t} \longrightarrow \mathbf{t}'$ and say that the reduction is *pure*.

The key reduction rule is R-BETA, which states that a function application $(\lambda z. \mathbf{t}) \mathbf{v}$ reduces to the function body, namely \mathbf{t} , where every occurrence of the formal argument z has been replaced with the actual argument \mathbf{v} . The λ construct, which prevented the function body \mathbf{t} 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 \mathbf{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 $\text{let } z = \mathbf{v} \text{ in } \mathbf{t}$ has the same behavior, with respect to reduction, as $(\lambda z. \mathbf{t}) \mathbf{v}$. We remark that 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). Please note that our choice of a call-by-value reduction strategy is fairly arbitrary, and has essentially no impact on the type system; the programming language Haskell (Hudak, Peyton Jones, Wadler, Boutel, Fairbairn, Fasel, Guzman, Hammond, Hughes, Johnsson, Kiebertz,

Nikhil, Partain, and Peterson, 1992), whose reduction strategy is known as *lazy* or *call-by-need*, also relies on Hindley and Milner’s type discipline.

Rule R-DELTA describes the semantics of constants. It merely states that a certain relation $\xrightarrow{\delta}$ is a subset of \longrightarrow . 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, which states that any valid reduction is also valid in a larger store. Condition (ii) allows proving that, if τ/μ reduces to τ'/μ' , then $\text{dom}(\mu)$ is a subset of $\text{dom}(\mu')$; this is left as an exercise to the reader.

- 1.2.4 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

- 1.2.5 EXAMPLE [PAIRS, CONTINUED]: In addition to the pair constructor defined in Example 1.2.3, 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

- 1.2.6 EXERCISE [BOOLEANS, RECOMMENDED, ★★]: Let `true` and `false` be nullary constructors. Let `if` be a ternary destructor. Extend the operational 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

- 1.2.7 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 $i \in \{1, 2\}$:

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

Here, the value `inji v` is being scrutinized, while the values `v1` and `v2`, which are typically functions, represent the two arms of a standard `case` construct. The rule selects an appropriate arm (here, `vi`) based on whether the value under scrutiny was formed using a left or right injection. The arm `vi` is executed and given access to the data carried by the injection (here, `v`). \square

- 1.2.8 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

- 1.2.9 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. Because configurations are considered equal up to α -conversion of memory locations, the choice of the name `m` is irrelevant, provided it is chosen fresh for `v`, so as 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 `v0` 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

- 1.2.10 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 = λz.t1 in t2` provided by ML-the-programming-language may be viewed as syntactic sugar for `let f = fix (λf.λz.t1) in t2`. \square

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 $\text{dom}(\mu)$.

- 1.2.11 EXERCISE [★★, RECOMMENDED, \rightarrow]: 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 \mathfrak{t}/μ reduces to \mathfrak{t}'/μ' , then \mathfrak{t} is not a value. In other words, values are irreducible: they represent a completed computation. The proof is left as an exercise to the reader. The converse, however, does not hold: if \mathfrak{t}/μ is irreducible with respect to \longrightarrow , then \mathfrak{t} is not necessarily a value. In that case, the configuration \mathfrak{t}/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 \mathfrak{t} is said to *go wrong* if and only if the configuration \mathfrak{t}/\emptyset reduces to a stuck configuration.

- 1.2.12 EXAMPLE: Runtime errors typically arise when destructors are applied to arguments of an unexpected nature. For instance, the expressions `+ 1 m` and `π1 2` and `!3` are stuck, regardless of the current store. The program `let z = $\hat{+} \hat{+}$ in z 1` is not stuck, because $\hat{+} \hat{+}$ is a value. However, its reduct through R-LET is $\hat{+} \hat{+} 1$, which is stuck, so this program goes wrong. The primary purpose of type systems is to prevent such situations from arising. \square
- 1.2.13 EXAMPLE: The configuration `!m/μ` is stuck if m is not in the domain of μ . In that case, however, `!m/μ` is not closed. Because we consider \longrightarrow 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; Crary, Walker, and Morrisett, 1999b; DeLine and Fähndrich, 2001; Grossman, Morrisett, Jim, Hicks, Wang, and Cheney, 2002a). \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 judgements are defined inductively by a collection of typing rules. We refer to this type system as DM.

To begin, we must define *types*. In DM, like in the simply-typed λ -calculus, types are first-order terms built out of *type constructors* and *type variables*. 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 1.2.1, then a nullary type constructor `int` is needed; if they include pair construction and destruction, as in Examples 1.2.3 and 1.2.5, then a binary type constructor \times is needed; and so on.

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 1.3.9) and in the definition of rows (Section 1.11).

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 terms as ill-formed types; this is illustrated, for instance, in Section 1.11. 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.

- 1.2.14 DEFINITION: Let d range over a finite or denumerable set of *directions*. Let κ range 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 referred to as 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* \mathcal{S} . In the following, we assume that a fixed signature \mathcal{S} is given and that every type constructor in it has *finite* arity, so as to ensure that types are machine representable. However, in Section 1.11, we shall explicitly work with several distinct signatures, some of which involve 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 user, within type declarations, type annotations, or module interfaces, are well-kinded.

- 1.2.15 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.

- 1.2.16 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 judgement, 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.

1.2.17 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. In the following, we write $\bar{X} \# \theta$ for $\bar{X} \# (dom(\theta) \cup range(\theta))$. We write $\theta \setminus \bar{X}$ for the restriction of θ outside \bar{X} , that is, the restriction of θ to $\mathcal{V} \setminus \bar{X}$. We sometimes let φ denote a type substitution. \square

If \bar{X} and \bar{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 $[\bar{X} \mapsto \bar{T}]$ denotes the substitution that maps X_i to T_i for every index i . The domain of $[\bar{X} \mapsto \bar{T}]$ is a subset of \bar{X} , the set underlying the vector \bar{X} . Its range is a subset of $ftv(\bar{T})$, where \bar{T} is the set underlying the vector \bar{T} . Every substitution θ may be written under the form $[\bar{X} \mapsto \bar{T}]$, where $\bar{X} = dom(\theta)$. Then, θ is *idempotent* if and only if $\bar{X} \# ftv(\bar{T})$ holds.

As pointed out earlier, types are first-order terms; that is, in the grammar of types, none of the productions *binds* a type variable. As a result, every type variable that appears within a type T appears *free* within T . This situation is identical to that of the simply-typed λ -calculus. 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.

1.2.18 DEFINITION: A type scheme S is an object of the form $\forall \bar{X}. T$, where T is a type of kind \star and the type variables \bar{X} are considered bound within 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 may be viewed as a subset of type schemes. The type scheme $\forall \bar{X}. T$ may be viewed as a finite way of describing the possibly infinite family of types of the form $[\bar{X} \mapsto \bar{T}]T$, where \bar{T} is arbitrary.

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

Figure 1-3: Typing rules for DM

Such types are called *instances* of the type scheme $\forall \bar{x}. T$. Note that, throughout most of this chapter, we work with *constrained type schemes*, a generalization of DM type schemes (Definition 1.3.2).

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

1.2.19 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 $\text{dpi}(\Gamma)$, is defined by $\text{dpi}(\emptyset) = \emptyset$ and $\text{dpi}(\Gamma; x : S) = \text{dpi}(\Gamma) \cup \{x\}$. \square

To complete the definition of Damas and Milner's type system, there remains to define *typing judgements*. A typing judgement 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. Such a judgement 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 judgement is *valid* (or *holds*) if and only if it may be derived using the rules that appear in Figure 1-3. An expression t is *well-typed* within the environment Γ if and only if some judgement of the form $\Gamma \vdash t : S$ 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 —can you prove it? 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 (Section 1.7). Rule DM-ABS specifies how to typecheck a λ -abstraction $\lambda z.t$. Its premise requires the body of the function, namely t , to be well-typed under an extra assumption, which 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, namely T , and result, namely 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. 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. 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* environment by binding z to a *type scheme* for t_1 prior to typechecking t_2 . 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 1.2.20 below. 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.

1.2.20 EXAMPLE: It is unsound to allow generalizing type variables that appear free in the environment. For instance, consider the typing judgement $z : X \vdash z : X$ **(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 judgement (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 judgements (2) and (3)—the former of which may be written $z : X \vdash z : \forall Y.Y$ —are incorrect. Indeed, such judgements 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

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.

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

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

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

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

1.2.23 EXERCISE [$\star\star$]: In fact, the rules shown in Figure 1-3 are not exactly Damas and Milner's original rules. In (Damas and Milner, 1982), the generalization and instantiation rules are:

$$\frac{\Gamma \vdash t : S \quad X \notin ftv(\Gamma)}{\Gamma \vdash t : \forall X. S} \quad (\text{DM-GEN}') \quad \frac{\Gamma \vdash t : \forall \bar{X}. T \quad \bar{Y} \# ftv(\forall \bar{X}. T)}{\Gamma \vdash t : \forall \bar{Y}. [\bar{X} \mapsto \bar{Y}] T} \quad (\text{DM-INST}')$$

where $\forall\bar{x}.S$ stands for $\forall\bar{x}.\bar{x}.T$ if S stands for $\forall\bar{x}.T$. Show that the combination of DM-GEN' and DM-INST' is equivalent to the combination of DM-GEN and DM-INST. \square

DM enjoys a number of nice theoretical properties, which have practical implications. First, under suitable hypotheses about the semantics of constants, about the type schemes that they receive in the initial environment, and—in the presence of side effects—under a slight restriction of the syntax of `let` constructs, it is possible to show that the type system is sound: that is, *well-typed (closed) programs do not go wrong*. This essential property ensures that programs that are accepted by the typechecker may be compiled without runtime checks. Furthermore, it is possible to show that 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 judgement 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 does not require expressions to be annotated with types. It also implies 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, which is very convenient.

Road map

Before proving the above claims, we first generalize our presentation by moving to a *constraint-based* setting. The necessary tools, namely the constraint language, its interpretation, and a number of constraint equivalence laws, are introduced in Section 1.3. In Section 1.4, we describe the standard constraint-based type system $HM(X)$ (Odersky, Sulzmann, and Wehr, 1999a; Sulzmann, Müller, and Zenger, 1999; Sulzmann, 2000). 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 Section 1.5, we propose an original reformulation of $HM(X)$, dubbed $PCB(X)$, whose distinctive feature is to exploit *type scheme introduction* and *instantiation* constraints. In Section 1.6, we show that, thanks to the extra expressive power afforded by these constraint forms, type inference may be viewed as a combination of constraint generation and constraint solving, as promised earlier. Indeed, we define a constraint generator and relate it with $PCB(X)$. Then, in Section 1.7, 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 $\text{PCB}(X)$, $\text{HM}(X)$, 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 names $\text{HM}(X)$ and $\text{PCB}(X)$. 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 Section 1.8, 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 Section 1.9, we explain how to extend ML-the-calculus with a number of features, including data structures, pattern matching, and type annotations. In Section 1.10, we extend the constraint language with universal quantification and describe a number of extra features that require this extension, including a different flavor of type annotations, polymorphic recursion, and first-class universal and existential types. Last, in Section 1.11, we extend the constraint language with *rows* and describe their applications, which include extensible variants and records.

1.3 Constraints

In this section, we define the syntax and logical meaning of constraints. Both are partly unspecified. Indeed, the set of *type constructors* (Definition 1.2.14) 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. Furthermore, 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.

Type constructors other than \rightarrow and predicates other than \leq will never explicitly appear in the definition of our constraint-based type systems, precisely because the definition is parametric with respect to them. They can (and usually do) appear in the type schemes assigned to constructors and destructors by the initial environment Γ_0 .

The introduction of subtyping has little impact on the complexity of our

| | | |
|--|-----------------------------------|--|
| $\sigma ::=$ | <i>type scheme:</i> | \emptyset |
| $\forall \bar{x}[C].T$ | | $x : \sigma$ |
| $C, D ::=$ | <i>constraint:</i> | $\Gamma; \Gamma$ |
| true | <i>truth</i> | $C, D ::=$ <i>Syntactic sugar for constraints:</i> |
| false | <i>falsity</i> | \dots <i>As before</i> |
| $P T_1 \dots T_n$ | <i>predicate application</i> | $\sigma \preceq T$ <i>Definition 1.3.3</i> |
| $C \wedge C$ | <i>conjunction</i> | $\text{let } x : \sigma \text{ in } C$ <i>Definition 1.3.3</i> |
| $\exists \bar{x}.C$ | <i>existential quantification</i> | $\exists \sigma$ <i>Definition 1.3.3</i> |
| $\text{def } x : \sigma \text{ in } C$ | <i>type scheme introduction</i> | $\text{def } \Gamma \text{ in } C$ <i>Definition 1.3.4</i> |
| $x \preceq T$ | <i>type scheme instantiation</i> | $\text{let } \Gamma \text{ in } C$ <i>Definition 1.3.4</i> |
| $\Gamma ::=$ | <i>Typing environments:</i> | $\exists \Gamma$ <i>Definition 1.3.4</i> |

Figure 1-4: Syntax of type schemes and constraints

proofs, yet increases the framework's expressive power. When subtyping is not desired, we interpret the predicate \leq as equality.

Syntax

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

1.3.1 **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 $n \geq 0$. Let \leq be a distinguished predicate of signature $\star \otimes \star \Rightarrow \cdot$. \square

1.3.2 **DEFINITION:** The syntax of *type schemes* and *constraints* is given in Figure 1-4. It is further restricted by the following requirements. In the type scheme $\forall \bar{x}[C].T$ and in the constraint $x \preceq T$, the type T must have kind \star . In the constraint $P T_1 \dots T_n$, 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 $T_1 \leq T_2$ for the binary predicate application $\leq T_1 T_2$, and call it a subtyping constraint. 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. Please note that x occurs free in the constraint $x \preceq T$.

We immediately introduce a number of derived constraint forms:

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

Constrained type schemes generalize Damas and Milner's type schemes, while our definition of instantiation constraints generalizes Damas and Milner's instance relation (Definition 1.2.18). Let us draw a comparison. First, Damas and Milner's instance relation yields a "yes/no" answer, 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, $\forall X.X \rightarrow X \preceq Y \rightarrow Z$ 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 $Y \leq Z$, or, if subtyping is interpreted as equality, to $Y = Z$. That is, $\sigma \preceq T'$ represents a condition on (the 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, so as to ensure that any supertype of an instance of σ is again an instance of σ (see rule C-EXTRANS in Figure 1-6 and Lemma 1.3.17). This is consistent with the purpose of subtyping, which is to allow supplying a subtype where a supertype is expected (TAPL Chapter 15). Third and last, every type scheme now carries a constraint. The constraint C , whose free type variables may or may not be members of \bar{X} , restricts the instances of the type scheme $\forall\bar{X}[C].T$. This is expressed in the instantiation constraint $\exists\bar{X}.(C \wedge T \preceq T')$, where the values that \bar{X} may assume are restricted by the requirement 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: they are exactly those of HM(X) (Odersky, Sulzmann, and Wehr, 1999a).

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 (2001b). In short, the construct $\text{def } x : \sigma$ in C binds the name x to the type scheme σ within the constraint C . If C contains a subconstraint of

the form $\mathbf{x} \preceq T$, where this occurrence of \mathbf{x} is free in C , then this subconstraint acquires the meaning $\sigma \preceq T$. Thus, the constraint $\mathbf{x} \preceq T$ is indeed an instantiation constraint, where the type scheme that is being instantiated is referred to by name. The constraint $\text{def } \mathbf{x} : \sigma \text{ in } C$ may be viewed as an *explicit substitution* of the type scheme σ for the name \mathbf{x} within C . Later (Section 1.5), 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 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 \mathbf{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, 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, 2001b). Fährdrich, Rehof, and Das’s instantiation constraints (2000) are also related, but may be recursive and are meant to be solved using a semi-unification procedure, as opposed to a unification algorithm extended with facilities for creating and instantiating type schemes, as in our case.

One consequence of introducing constraints inside type schemes is that some type schemes have no instances at all, or have instances only if a certain constraint holds. For instance, the type scheme $\sigma = \forall X[\text{bool} = \text{int}].X$, where the nullary type constructors *int* and *bool* have distinct interpretations, has no instances; that is, no constraint of the form $\sigma \preceq T'$ has a solution. The type scheme $\sigma = \forall Z[X = Y \rightarrow Z].Z$ has an instance only if $X = Y \rightarrow Z$ holds for some Z ; in other words, for every T' , $\sigma \preceq T'$ entails $\exists Z.(X = Y \rightarrow Z)$. (We define *entailment* on page 29.) We later prove that the constraint $\exists \sigma$ is equivalent to $\exists Z.\sigma \preceq Z$, where $Z \notin \text{ftv}(\sigma)$ (Exercise 1.3.23). That is, $\exists \sigma$ expresses the requirement that σ have an instance. Type schemes that do not have an instance indicate a type error, so in many situations, one wishes to avoid them; for this reason, we often use the constraint form *let* $\mathbf{x} : \sigma$ in C , which requires σ to have an instance and at the same time associates it with the name \mathbf{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 Section 1.3; we always use *let* instead.

- 1.3.4 **DEFINITION:** Environments Γ remain as in Definition 1.2.19, except DM type schemes \mathbf{S} are replaced with constrained type schemes σ . We write $dfpi(\Gamma)$ for $dpi(\Gamma) \cup fpi(\Gamma)$. We define $\text{def } \emptyset \text{ in } C = C$ and $\text{def } \Gamma; \mathbf{x} : \sigma \text{ in } C = \text{def } \Gamma \text{ in } \text{def } \mathbf{x} : \sigma \text{ in } C$. Similarly, we define $\text{let } \emptyset \text{ in } C = C$ and $\text{let } \Gamma; \mathbf{x} : \sigma \text{ in } C = \text{let } \Gamma \text{ in } \text{let } \mathbf{x} : \sigma \text{ in } C$. We define $\exists \emptyset = \text{true}$ and $\exists(\Gamma; \mathbf{x} : \sigma) = \exists \Gamma \wedge \text{def } \Gamma \text{ in } \exists \sigma$. \square

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

$$C ::= \square \mid C \mid C \wedge C \mid \exists \bar{x}. C \mid \text{def } \mathbf{x} : \sigma \text{ in } C \mid \text{def } \mathbf{x} : \forall \bar{x}. C. T \text{ in } C$$

The application of a constraint context \mathcal{C} to a constraint C , written $\mathcal{C}[C]$, is defined in the usual way. Because a 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 $\mathcal{C}[C]$. We write $dtv(\mathcal{C})$ and $dpi(\mathcal{C})$ for the sets of type variables and program identifiers, respectively, that \mathcal{C} may thus capture. We write $\text{let } \mathbf{x} : \forall \bar{x}. \mathcal{C}. T \text{ in } C$ for $\exists \bar{x}. \mathcal{C} \wedge \text{def } \mathbf{x} : \forall \bar{x}. \mathcal{C}. T \text{ in } C$. Being able to state such a definition is why we require multi-hole contexts. We let range over *existential constraint contexts*, defined by $\mathcal{X} ::= \square \mid \exists \bar{x}. \mathcal{X}$.

Meaning

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*:

- 1.3.5 **DEFINITION:** For every kind κ , let \mathcal{M}_κ be a nonempty set, whose elements are 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 of domain $dom(K)$ that map every $d \in dom(K)$ to an element of $\mathcal{M}_{K(d)}$. For every predicate P of signature $\kappa_1 \otimes \dots \otimes \kappa_n \Rightarrow \cdot$, let P denote a predicate on $\mathcal{M}_{\kappa_1} \times \dots \times \mathcal{M}_{\kappa_n}$. We require the predicate \leq on $\mathcal{M}_\star \times \mathcal{M}_\star$ to be a partial order. \square

For the sake of convenience, we abuse notation and write F for both the type constructor and its interpretation; similarly for predicates. We freely assume that a binary equality predicate, whose interpretation is equality on \mathcal{M}_κ , is available at every kind κ , so $T_1 = T_2$, where T_1 and T_2 have kind κ , is a well-formed constraint.

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 informally refer to the collection of these choices as X . Thus, the type systems $\text{HM}(X)$ and $\text{PCB}(X)$, described in Sections 1.4 and 1.5, are *parameterized* by X .

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

1.3.6 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

1.3.7 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* (resp. *regular*) trees t such that $t(\epsilon)$ has image kind κ : then, we have a *finite* (resp. *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. Please note that free finite tree models coincide with syntactic models, as defined in the previous example. \square

Rows (Section 1.11) are interpreted in a tree model, albeit not a free one. The following examples suggest different ways of interpreting the subtyping predicate.

1.3.8 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

1.3.9 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 and such that $-- = +$ and $\pm- = \pm\pm = \pm$. 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 (Kozen, Palsberg, and Schwartzbach., 1995) 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, and introducing, in addition to the type constructor \rightarrow , two type constructors \perp and \top of signature \star . 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.

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 \diamond , a type constructor *abs* of signature \diamond , a type constructor *pre* of signature $\{\text{content} \mapsto \star\} \Rightarrow \diamond$, and let $\text{pre} \leq \text{abs}$. This gives rise to a model where $\text{pre } t \leq \text{abs}$ holds for every $t \in \mathcal{M}_\star$. This form of subtyping is called *nonstructural* because comparable ground types may have different shapes, such as \perp and $\perp \rightarrow \top$, or $\text{pre } \top$ and abs . Nonstructural subtyping has been studied, for example, in (Kozen, Palsberg, and Schwartzbach., 1995; Palsberg, Wand, and O’Keefe, 1997; Pottier, 2001b; Niehren and Priesnitz, 2003). Section 1.11 says more about typechecking operations on records.

An important particular case arises when *any two type constructors related by \leq have the same arity*. 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 $\text{dom}(t_2)$ 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; 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

Our last example suggests a predicate other than equality and subtyping.

1.3.10 EXAMPLE [CONDITIONAL CONSTRAINTS]: Consider a nonstructural subtyping model. For every type constructor F of image kind κ and for every kind κ' , let $(F \leq \cdot \Rightarrow \cdot \leq \cdot)$ be a predicate of signature $\kappa \otimes \kappa' \otimes \kappa' \Rightarrow \cdot$. Thus, if T_0 has kind κ and T_1, T_2 have the same kind, then $F \leq T_0 \Rightarrow T_1 \leq T_2$ is a well-formed constraint, called a *conditional subtyping constraint*. Its interpretation is defined as follows: if $t_0 \in \mathcal{M}_\kappa$ and $t_1, t_2 \in \mathcal{M}_{\kappa'}$, then $F \leq t_0 \Rightarrow t_1 \leq t_2$ holds if and only if $F \leq t_0(\epsilon)$ implies $t_1 \leq t_2$. In other words, if t_0 's head symbol exceeds F according to the ordering on type constructors, then the subtyping constraint $t_1 \leq t_2$ must hold; otherwise, the conditional constraint holds vacuously. Conditional constraints have been studied *e.g.* in (Reynolds, 1969a; Heintze, 1993; Aiken, Wimmers, and Lakshman, 1994; Pottier, 2000; Su and Aiken, 2001). \square

Many other kinds of constraints exist; see *e.g.* (Comon, 1993).

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

As usual, the meaning of a constraint is a function of the meaning of its free type variables, which is given by a *ground assignment*. The meaning of free program identifiers may be defined as part of the constraint, if desired, using a *def* prefix, so it need not be given by a separate assignment.

1.3.11 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 . Whether a constraint C holds under a ground assignment ϕ , written $\phi \vdash C$ (read: ϕ *satisfies* C), is defined by the rules in Figure 1-5. A constraint C is *satisfiable* if and only if $\phi \vdash C$ holds for some ϕ . It is *false* if and only if $\phi \vdash \text{def } \Gamma \text{ in } C$ holds for *no* ground assignment ϕ and environment Γ . \square

Let us now explain the rules that define constraint satisfaction (Figure 1-5). They are syntax-directed: that is, to a given constraint, at most one rule applies. It is determined by the nature of the first construct that appears under a maximal *def* prefix. CM-TRUE states that a constraint of the form *def* Γ in *true* is a tautology, that is, holds under every ground assignment. No rule matches constraints of the form *def* Γ in *false*, which means that such constraints do not have a solution. CM-PREDICATE states that the meaning

| | |
|---|--|
| $\frac{\phi \vdash \text{def } \Gamma \text{ in true}}{\phi \vdash \text{def } \Gamma \text{ in } P \text{ T}_1 \dots \text{ T}_n} \text{ (CM-TRUE)}$ $\frac{P(\phi(\text{T}_1), \dots, \phi(\text{T}_n))}{\phi \vdash \text{def } \Gamma \text{ in } P \text{ T}_1 \dots \text{ T}_n} \text{ (CM-PREDICATE)}$ $\frac{\phi \vdash \text{def } \Gamma \text{ in } C_1 \quad \phi \vdash \text{def } \Gamma \text{ in } C_2}{\phi \vdash \text{def } \Gamma \text{ in } (C_1 \wedge C_2)} \text{ (CM-AND)}$ | $\frac{\phi[\vec{x} \mapsto \vec{t}] \vdash \text{def } \Gamma \text{ in } C \quad \vec{x} \# \text{ftv}(\Gamma)}{\phi \vdash \text{def } \Gamma \text{ in } \exists \vec{x}. C} \text{ (CM-EXISTS)}$ $\frac{\phi \vdash \text{def } \Gamma_1 \text{ in } \sigma \preceq T' \quad x \notin \text{dpi}(\Gamma_2)}{\phi \vdash \text{def } \Gamma_1; x : \sigma; \Gamma_2 \text{ in } x \preceq T'} \text{ (CM-INSTANCE)}$ |
|---|--|

Figure 1-5: Meaning of constraints

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 $\phi(\text{T}_i)$ is a ground type in \mathcal{M}_{κ_i} . By Definition 1.3.5, 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 Γ , which is natural, since a predicate application has no free program identifiers. CM-AND requires each of the conjuncts to be valid in isolation. The information in Γ is made available to each branch. CM-EXISTS allows the type variables \vec{x} to denote arbitrary ground types \vec{t} within C , independently of their image through ϕ . We implicitly require \vec{x} and \vec{t} to have matching kinds, so that $\phi[\vec{x} \mapsto \vec{t}]$ remains a kind-preserving ground assignment. The side condition $\vec{x} \# \text{ftv}(\Gamma)$ —which may always be satisfied by suitable α -conversion of the constraint $\exists \vec{x}. C$ —prevents free occurrences of the type variables \vec{x} within Γ from being unduly affected. CM-INSTANCE concerns constraints of the form $\text{def } \Gamma \text{ in } x \preceq T'$. The constraint $x \preceq T'$ is turned into $\sigma \preceq T'$, where, according to the second premise, σ is $\Gamma(x)$. Please recall that constraints of such a form were introduced in Definition 1.3.3. The environment Γ is replaced with a suitable prefix of itself, namely Γ_1 , so that the free program identifiers of σ retain their meaning.

It is intuitively clear that the constraints $\text{def } x : \sigma \text{ in } C$ 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 (informal) claim that the def form is an explicit substitution form.

It is possible for a constraint to be neither satisfiable nor false. Consider, for instance, the constraint $\exists Z. x \preceq Z$. Because the identifier x is free, CM-INSTANCE is not applicable, so the constraint is not satisfiable. Furthermore,

placing it within the context $\text{let } x : \forall X.X \text{ in } []$ makes it satisfied by every ground assignment, so it is not false. Here, the assertions “ C is satisfiable” and “ C is false” are opposite when $\text{fpi}(C) = \emptyset$ holds, whereas in a standard first-order logic, they always are.

In a judgement of the form $\phi \vdash C$, the ground assignment ϕ applies to the free type variables of C . This is made precise by the following statements. In the second one, \circ is composition and $\theta(C)$ is the capture-avoiding application of the type substitution θ to C .

- 1.3.12 LEMMA: If $\bar{x} \# \text{fv}(C)$ holds, then $\phi \vdash C$ and $\phi[\bar{x} \mapsto \bar{t}] \vdash C$ are equivalent. \square
- 1.3.13 LEMMA: $\phi \circ \theta \vdash C$ and $\phi \vdash \theta(C)$ are equivalent. \square

Reasoning with constraints

Because constraints lie at the heart of our treatment of ML-the-type-system, most of our proofs involve establishing logical properties of constraints, that is, *entailment* or *equivalence* assertions. Let us first define these notions.

- 1.3.14 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 environment Γ , $\phi \vdash \text{def } \Gamma$ in C_1 implies $\phi \vdash \text{def } \Gamma$ in 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

This definition measures the strength of a constraint by the set of pairs (ϕ, Γ) that satisfy it, and considers a constraint stronger if fewer such pairs satisfy it. 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. We remark that C is false 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* (resp. *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 $T_i \leq T'_i$ (resp. $T'_i \leq T_i$, $T_i = T'_i$). We let the reader check the following facts: (i) in an equality model, these three notions 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 (resp. contravariant, invariant), then every type constructor whose arity includes d is covariant (resp. contravariant,

invariant) with respect to d . 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). These properties are summed up by the following equivalence law:

$$T_1 \rightarrow T_2 \leq T'_1 \rightarrow T'_2 \equiv T'_1 \leq T_1 \wedge T_2 \leq T'_2 \quad (\text{C-ARROW})$$

Please note that this is a high-level requirement about the interpretation of types and of the subtyping predicate. In an equality free tree model, for instance, 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 $F_1 \vec{T}_1 \leq F_2 \vec{T}_2$ and $F_2 \vec{T}_2 \leq F_1 \vec{T}_1$ are false; then, we write $F_1 \bowtie F_2$. Please 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 each of F_1 and F_2 is 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 consider the type constructor \rightarrow isolated.

Entailment is preserved by arbitrary constraint contexts, as stated by the following theorem. As a result, constraint equivalence is a congruence. Throughout this chapter, these facts are often used implicitly.

1.3.15 THEOREM [CONGRUENCE]: $C_1 \Vdash C_2$ implies $\mathcal{C}[C_1] \Vdash \mathcal{C}[C_2]$. □

We now give a series of lemmas that provide useful entailment laws.

The following is a standard property of existential quantification.

1.3.16 LEMMA: $C \Vdash \exists \vec{X}. C$. □

The following lemma states that any supertype of an instance of σ is also an instance of σ .

1.3.17 LEMMA: $\sigma \preceq T \wedge T \leq T' \Vdash \sigma \preceq T'$. □

The next lemma gives another interesting simplification law.

1.3.18 LEMMA: $X \notin \text{ftv}(T)$ implies $\exists X.(X = T) \equiv \text{true}$. □

The following lemma states that, provided D is satisfied, the type T is an instance of the constrained type scheme $\forall \vec{X}[D].T$.

1.3.19 LEMMA: $D \Vdash \forall \bar{x}[D].T \preceq T$. □

This technical lemma helps justify Definition 1.3.21 below.

1.3.20 LEMMA: Let $Z \notin \text{ftv}(C, \sigma, T)$. Then, $C \Vdash \sigma \preceq T$ holds if and only if $C \wedge T \preceq Z \Vdash \sigma \preceq Z$ holds. □

It is useful to define what it means for a type scheme σ_1 to be *more general* than a type scheme σ_2 . Our informal intent is for $\sigma_1 \preceq \sigma_2$ to mean: *every instance of σ_2 is an instance of σ_1* . In Definition 1.3.3, we have introduced the constraint form $\sigma \preceq T$ as syntactic sugar. Similarly, one might wish to make $\sigma_1 \preceq \sigma_2$ a derived constraint form; however, this is impossible, because neither universal quantification nor implication are available in the constraint language. We can, however, exploit the fact that these logical connectives are implicit in entailment assertions by defining a judgement of the form $C \Vdash \sigma_1 \preceq \sigma_2$, whose meaning is: *under the constraint C , σ_1 is more general than σ_2* .

1.3.21 DEFINITION: We write $C \Vdash \sigma_1 \preceq \sigma_2$ if and only if $Z \notin \text{ftv}(C, \sigma_1, \sigma_2)$ implies $C \wedge \sigma_2 \preceq Z \Vdash \sigma_1 \preceq Z$. We write $C \Vdash \sigma_1 \equiv \sigma_2$ when both $C \Vdash \sigma_1 \preceq \sigma_2$ and $C \Vdash \sigma_2 \preceq \sigma_1$ hold. □

This notation is not ambiguous because the assertion $C \Vdash \sigma \preceq T$, whose meaning was initially given by Definitions 1.3.3 and 1.3.14, retains the same meaning under the new definition—this is shown by Lemma 1.3.20 above.

The next lemma provides a way of exploiting the ordering between type schemes introduced by Definition 1.3.21. It states that a type scheme occurs in *contravariant* position when it is within a `def` prefix. In other words, the more general the type scheme, the weaker the entire constraint.

1.3.22 LEMMA: $C \Vdash \sigma_1 \preceq \sigma_2$ implies $C \wedge \text{def } x : \sigma_2 \text{ in } D \Vdash \text{def } x : \sigma_1 \text{ in } D$. □

The following exercise generalizes this result to `let` forms.

1.3.23 EXERCISE [★★, \rightarrow]: Prove that $Z \notin \text{ftv}(\sigma)$ implies $\exists \sigma \equiv \exists Z.\sigma \preceq Z$. Explain why, as a result, $C \Vdash \sigma_1 \preceq \sigma_2$ implies $C \wedge \exists \sigma_2 \Vdash \exists \sigma_1$. Use this fact to prove that $C \Vdash \sigma_1 \preceq \sigma_2$ implies $C \wedge \text{let } x : \sigma_2 \text{ in } D \Vdash \text{let } x : \sigma_1 \text{ in } D$. □

The next lemma states that, modulo equivalence, the only constraint that constrains x without explicitly referring to it is false.

1.3.24 LEMMA: $C \Vdash x \preceq T$ and $x \notin \text{fpi}(C)$ imply $C \equiv \text{false}$. □

The following lemma states that the more universal quantifiers are present, the more general the type scheme.

1.3.25 LEMMA: $\text{let } x : \forall \bar{x}[C_1].T \text{ in } C_2 \Vdash \text{let } x : \forall \bar{x}\bar{y}[C_1].T \text{ in } C_2.$ \square

Conversely, and perhaps surprisingly, it is sometimes possible to *remove* some type variables from the universal quantifier prefix of a type scheme without compromising its generality. This is the case when the value of these type variables is *determined* in a unique way. In short, C determines \bar{Y} if and only if, given the values of $\text{ftv}(C) \setminus \bar{Y}$ and given that C holds, it is possible to reconstruct, in a unique way, the values of \bar{Y} .

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

Two concrete instances of determinacy, one of which is valid only in free tree models, are given by Lemma 1.8.7 on page 82. Determinacy is exploited by the equivalence law C-LETALL in Figure 1-6.

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.

1.3.27 THEOREM: All equivalence laws in Figure 1-6 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 it does 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 remark that the instances of a monotype are its supertypes, that is, by Definition 1.3.3, $T \preceq T'$ 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

| | | | |
|--|----------|--|--|
| $C_1 \wedge C_2$ | \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 $\bar{x} \# \text{ftw}(C)$ (C-EX*) |
| $(\exists \bar{X}. C_1) \wedge C_2$ | \equiv | $\exists \bar{X}. (C_1 \wedge C_2)$ | if $\bar{x} \# \text{ftw}(C_2)$ (C-EXAND) |
| $\exists Z. (\sigma \preceq Z \wedge Z \leq T')$ | \equiv | $\sigma \preceq T'$ | if $Z \notin \text{ftw}(\sigma, T')$ (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) |
| $\text{let } \Gamma \text{ in } C$ | \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 $\bar{x} \# \text{ftw}(\Gamma)$ (C-INEX) |
| $\text{let } \Gamma_1; \Gamma_2 \text{ in } C$ | \equiv | $\text{let } \Gamma_2; \Gamma_1 \text{ in } C$ | (C-LETLET) |
| $\text{let } x : \forall \bar{X}[C_1 \wedge C_2]. T \text{ in } C_3$ | \equiv | $C_1 \wedge \text{let } x : \forall \bar{X}[C_2]. T \text{ in } C_3$ | if $\bar{x} \# \text{ftw}(C_1)$ (C-LETAND) |
| $\text{let } \Gamma; x : \forall \bar{X}[C_1]. T \text{ in } C_2$ | \equiv | $\text{let } \Gamma; x : \forall \bar{X}[\text{let } \Gamma \text{ in } C_1]. T \text{ in } C_2$ | if $\bar{x} \# \text{ftw}(\Gamma)$ and $\text{dpi}(\Gamma) \# \text{fpi}(\Gamma)$ (C-LETDUP) |
| $\text{let } x : \forall \bar{X}[\exists \bar{Y}. C_1]. T \text{ in } C_2$ | \equiv | $\text{let } x : \forall \bar{X} \bar{Y}[C_1]. T \text{ in } C_2$ | if $\bar{Y} \# \text{ftw}(T)$ (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) |
| $\exists X. (T \leq X \wedge \text{let } x : X \text{ in } C)$ | \equiv | $\text{let } x : T \text{ in } C$ | if $x \notin \text{ftw}(T, C)$ (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 $\bar{x} \# \text{ftw}(\vec{T})$ (C-NAME) |
| $[\vec{X} \mapsto \vec{T}]C$ | \equiv | $\exists \bar{X}. (\vec{X} = \vec{T} \wedge C)$ | if $\bar{x} \# \text{ftw}(\vec{T})$ (C-NAMEEQ) |

Figure 1-6: Constraint equivalence laws

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 remark 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 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-LETEX 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 provides a restricted converse of Lemma 1.3.25. Together, C-LETEX and C-LETALL allow—in some situations only—to hoist existential quantifiers out of the *left*-hand side of a let form.

- 1.3.28 EXAMPLE: C-LETALL would be invalid without the condition that $\exists \bar{x}.C_1$ determines \bar{y} . Consider, for instance, the constraint $\text{let } x : \forall Y.Y \rightarrow Y \text{ in } (x \preceq \text{int} \rightarrow \text{int} \wedge x \preceq \text{bool} \rightarrow \text{bool})$ (1), where int and bool are incompatible nullary type constructors. By C-INID and C-IN*, it is equivalent to $\exists Y.(Y \rightarrow Y \leq \text{int} \rightarrow \text{int}) \wedge \exists Y.(Y \rightarrow Y \leq \text{bool} \rightarrow \text{bool})$, that is, true. Now, if C-LETALL was valid without its side-condition, then (1) would also be equivalent to $\exists Y.\text{let } x : Y \rightarrow Y \text{ in } (x \preceq \text{int} \rightarrow \text{int} \wedge x \preceq \text{bool} \rightarrow \text{bool})$, 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; when Y is not universally quantified, however, 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 . Lemma 1.3.25 states that, when deprived of its side-condition, C-LETALL is only an entailment law, as opposed to an equivalence law. 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 $\text{let } x : \forall X[\exists Y.X = Y \rightarrow Y].X \text{ in } (x \preceq \text{int} \rightarrow \text{int} \wedge x \preceq \text{bool} \rightarrow \text{bool})$.

Naturally, in the above examples, the side-condition “true determines Y ” does *not* hold: by Definition 1.3.26, it is equivalent to “two ground assignments that coincide outside Y must coincide on Y as well”, which is false as soon as \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 1.8.7. As a result, C-LETALL states that $\text{let } x : \forall XY[Y = \text{int}].Y \rightarrow X \text{ in } C$ **(1)** is equivalent to $\exists Y.\text{let } x : \forall X[Y = \text{int}].Y \rightarrow X \text{ in } C$ **(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. We remark that, by C-LETAND, **(2)** is equivalent to $\exists Y.(Y = \text{int} \wedge \text{let } x : \forall X.Y \rightarrow X \text{ in } C)$ **(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 analogue 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. 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. This immediately yields a proof of the following fact:

1.3.29 LEMMA: $C \Vdash D$ implies $\theta(C) \Vdash \theta(D)$. \square

It is important to stress that, because the effect of a type substitution may be emulated using equations, conjunction, and existential quantification, there is no need ever to employ type substitutions in the definition of a constraint-based type system—it is possible, instead, to express every concept in terms

of constraints. In this chapter, we follow this route, and use type substitutions only when dealing with the type system DM, whose historical formulation is substitution-based.

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 proof of several constraint equivalence laws. However, in the remainder of this chapter, we work with `let` forms exclusively and never employ the `def` construct. As a result, *it is possible, from here on, to discard `def` and pretend that `let` is primitive*. This change in perspective offers us a few extra properties, stated in the next two lemmas. First, every constraint that contains a false subconstraint must be false. Second, no satisfiable constraint has a free program identifier.

1.3.30 LEMMA: $\mathcal{C}[\text{false}] \equiv \text{false}$. □

1.3.31 LEMMA: If C is satisfiable, then $\text{fpi}(C) = \emptyset$. □

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 important 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 short, we aim to prove that every satisfiable constraint admits a *canonical solved form*, to show that this notion corresponds to the standard concept of a *most general unifier*, and to establish a few technical properties of most general unifiers.

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 Lemmas 1.3.32 and 1.3.39.

Under these new assumptions, the interpretation of equality coincides with its syntax: every equation that holds in the model is in fact a syntactic truism. The converse, of course, holds in every model.

1.3.32 LEMMA: If $\text{true} \Vdash T = T'$ holds, then T and T' coincide. □

In a syntactic model, ground types are finite trees. As a result, cyclic equations, such as $X = \text{int} \rightarrow X$, are false.

1.3.33 LEMMA: $X \in \text{ftv}(T)$ and $T \notin \mathcal{V}$ imply $(X = T) \equiv \text{false}$. □

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 Section 1.8 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.

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

Solved forms offer a convenient way of reasoning about constraints because every satisfiable constraint is equivalent to one. In other words, every constraint is equivalent to either a solved form or false. This property is established by the following lemma, whose proof provides a simple but effective procedure to rewrite a constraint to either a solved form or false.

1.3.35 LEMMA: Let $fpi(C) = \emptyset$. Then, C is equivalent to either a solved form or false. \square

Proof: We first establish that every conjunction of equations is equivalent to either a solved form or false. To do so, we present Robinson's unification algorithm (1971) as a rewriting system. The system's invariant is to operate on constraints of the form either $\bar{X} = \bar{T}; C$, where $\bar{X} \# ftv(\bar{T}, C)$ and the semicolon is interpreted as a distinguished conjunction, or false. We identify equations in C up to commutativity. The system is defined as follows:

$$\begin{array}{lll}
 \bar{X} = \bar{T}; & X = X \wedge C & \rightarrow \bar{X} = \bar{T}; C \\
 \bar{X} = \bar{T}; & F \bar{T}_1 = F \bar{T}_2 \wedge C & \rightarrow \bar{X} = \bar{T}; \bar{T}_1 = \bar{T}_2 \wedge C \\
 \bar{X} = \bar{T}; & F_1 \bar{T}_1 = F_2 \bar{T}_2 \wedge C & \rightarrow \text{false} \\
 & & \text{if } F_1 \neq F_2 \\
 \bar{X} = \bar{T}; & X = T \wedge C & \rightarrow \bar{X} = [X \mapsto T] \bar{T} \wedge X = T; [X \mapsto T] C \\
 & & \text{if } X \notin ftv(T) \\
 \bar{X} = \bar{T}; & X = T \wedge C & \rightarrow \text{false} \\
 & & \text{if } X \in ftv(T) \text{ and } T \notin \mathcal{V}
 \end{array}$$

It is straightforward to check that the above invariant is indeed preserved by the rewriting system. Let us check that constraint equivalence is also preserved. For the first rule, this is immediate. For the second and third rules, it

follows from the fact that we have assumed a free tree model; for the fourth rule, a consequence of C-EQ; for the last rule, a consequence of Lemma 1.3.33. Furthermore, the system is terminating; this is witnessed by an ordering where false is the least element and where constraints of the form $\vec{X} = \vec{T}; C$ are ordered lexicographically, first by the number of type variables that appear free within C , second by the size of C . Last, a normal form for this rewriting system must be of the form either $\vec{X} = \vec{T}; \text{true}$, where (by the invariant) $\vec{X} \# \text{ftv}(\vec{T})$ —that is, a solved form, or false.

Next, we show that the present lemma holds when C is built out of equations, conjunction, and existential quantification. Orienting C-EXAND from left to right yields a terminating rewriting system that preserves constraint equivalence. The normal form of C must be $\exists \bar{Y}.C'$, where C' is a conjunction of equations. By the previous result, C' is equivalent to either a solved form or false. Because solved forms are preserved by existential quantification and because $\exists \bar{Y}.\text{false}$ is false, the same holds of C .

Last, we establish the result in the general case. We assume $\text{fpi}(C) = \emptyset$ (1). Orienting C-INID and C-IN* from left to right yields a terminating rewriting system that preserves constraint equivalence. The normal form C' of C cannot contain any type scheme introduction forms; given (1), it cannot contain any instantiation forms either. Thus, C' is built out of equations, conjunction, and existential quantification only. By the previous result, it is equivalent to either a solved form or false, which implies that the same holds of C . \square

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

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

1.3.37 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 $[\vec{X} \mapsto \vec{T}]$.

1.3.38 LEMMA: A ground assignment ϕ satisfies a canonical solved form $\exists \bar{Y}.\vec{X} = \vec{T}$ if and only if there exists a ground assignment ϕ' such that $\phi(\vec{X}) = \phi'(\vec{T})$. As a result, every canonical solved form is satisfiable. \square

Proof: Let $\exists \bar{Y}.\vec{X} = \vec{T}$ be a canonical solved form. By CM-EXISTS and CM-PREDICATE, ϕ satisfies $\exists \bar{Y}.\vec{X} = \vec{T}$ if and only if there exists \vec{t} such that $\phi[\vec{Y} \mapsto \vec{t}](\vec{X}) = \phi[\vec{Y} \mapsto \vec{t}](\vec{T})$. Thanks to the hypotheses $\bar{X} \# \bar{Y}$ and $\text{ftv}(\vec{T}) \subseteq \bar{Y}$, this is equivalent to the existence of a ground assignment ϕ' such that $\phi(\vec{X}) = \phi'(\vec{T})$.

Thus, for every ground assignment ϕ' , $\phi'[\vec{X} \mapsto \phi'(\vec{T})]$ satisfies $\exists \bar{Y}.(\vec{X} = \vec{T})$, which proves that this constraint is satisfiable. \square

Together, Lemmas 1.3.37 and 1.3.38 imply that every solved form is satisfiable. Our interest in canonical solved forms stems from the following fundamental property, which provides a *syntactic* characterization of entailment between canonical solved forms: if $\exists \bar{Y}_1.(\vec{X} = \vec{T}_1)$ is more specific than $\exists \bar{Y}_2.(\vec{X} = \vec{T}_2)$, in a logical sense, then \vec{T}_1 refines \vec{T}_2 , in a syntactic sense. The converse also holds (can you prove it?), but is not needed here.

1.3.39 LEMMA: If $\exists \bar{Y}_1.(\vec{X} = \vec{T}_1) \Vdash \exists \bar{Y}_2.(\vec{X} = \vec{T}_2)$, where both sides are canonical solved forms, then there exists a type substitution φ such that $\vec{T}_1 = \varphi(\vec{T}_2)$. \square

As a corollary, we find that canonical solved forms are *unique* up to α -conversion and up to C-EX*, *provided* the set \bar{X} of their free type variables is fixed.

1.3.40 LEMMA: If the canonical solved forms $\exists \bar{Y}_1.(\vec{X} = \vec{T}_1)$ and $\exists \bar{Y}_2.(\vec{X} = \vec{T}_2)$ are equivalent, then there exists a renaming ρ such that $\vec{T}_1 = \rho(\vec{T}_2)$. \square

Please note that the fact that the canonical solved forms $\exists \bar{Y}_1.(\vec{X}_1 = \vec{T}_1)$ and $\exists \bar{Y}_2.(\vec{X}_2 = \vec{T}_2)$ are equivalent does *not* imply that \bar{X}_1 and \bar{X}_2 coincide. Consider, for example, the canonical solved forms `true` and $\exists Y.(X = Y)$, which by C-NAMEEQ are equivalent. One might wish to further restrict canonical solved forms by requiring \bar{X} to be the set of *essential* type variables of the constraint $\exists \bar{Y}.(\vec{X} = \vec{T})$, that is, the set of the type variables that appear free in *all* equivalent constraints. However, as far our technical development is concerned, it seems more convenient not to do so. Instead, we show that it is possible to explicitly restrict or extend \bar{X} when needed (Lemma 1.3.43).

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.

1.3.41 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 $\exists \bar{Y}.(\vec{X} = \vec{T})$, where $\bar{Y} = \text{ftv}(\vec{T})$. 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 1.3.31, 1.3.35, and 1.3.37, every satisfiable constraint admits a most general unifier. By Lemma 1.3.40, if $[\vec{X} \mapsto \vec{T}_1]$ and

$[\vec{x} \mapsto \vec{T}_2]$ are most general unifiers of C , then \vec{T}_1 and \vec{T}_2 coincide up to a renaming. Conversely, if $[\vec{x} \mapsto \vec{T}]$ is a most general unifier of C and if $\vec{x} \# \rho$ holds, then $[\vec{x} \mapsto \rho\vec{T}]$ is also a most general unifier of C ; indeed, these two substitutions correspond to α -equivalent canonical solved forms.

The following result relates the substitution θ to the canonical solved form $\exists\theta$, stating that every ground refinement of the former satisfies the latter.

1.3.42 LEMMA: $\theta(\exists\theta) \equiv \text{true}$. □

The following lemma offers two technical results: the domain of a most general unifier of C may be restricted so as to become a subset of $ftv(C)$; it may also be extended to include arbitrary fresh variables. The next lemma is a simple corollary.

1.3.43 LEMMA: Let θ be a most general unifier of C . If $\bar{z} \# ftv(C)$, then $\theta \setminus \bar{z}$ is also a most general unifier of C . If $\bar{z} \# \theta$, then there exists a most general unifier of C that extends θ and whose domain is $dom(\theta) \cup \bar{z}$. □

1.3.44 LEMMA: Let θ_1 and θ_2 be most general unifiers of C . Let $\bar{x} = dom(\theta_1) \cap dom(\theta_2)$. Then, $\theta_1(\bar{x})$ and $\theta_2(\bar{x})$ coincide up to a renaming. □

Our last technical result relates the most general unifiers of C with the most general unifiers of $\exists X.C$. It states that the former are extensions of the latter. Furthermore, under a few freshness conditions, *every* most general unifier of $\exists X.C$ may be extended to yield a most general unifier of C .

1.3.45 LEMMA: If θ is a most general unifier of C , then $\theta \setminus X$ is a most general unifier of $\exists X.C$. Conversely, if θ is a most general unifier of $\exists X.C$ and $X \# \theta$ and $ftv(\exists X.C) \subseteq dom(\theta)$, then there exists a type substitution θ' such that θ' extends θ , θ' is a most general unifier of C , and $dom(\theta') = dom(\theta) \cup X$. □

1.4 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, 1994a; 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 24. Its original description is due to Odersky, Sulzmann, and Wehr (1999a). Since then, it has been completed in a number of works (Sulzmann, Müller, and Zenger, 1999; Sulzmann, 2000; Pottier, 2001a;

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 section 1.3. Technically, our approach of constraints is more direct than that of (Odersky, Sulzmann, and Wehr, 1999a). We interpret constraints within a model, give conjunction and existential quantification their standard meaning, and derive a number of equivalence laws (Section 1.3). 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, for the sake of definiteness; 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 judgements to *have no free program identifiers*. Please note that this does not prevent them from containing let forms; we shall in fact exploit this feature when establishing an equivalence between $HM(X)$ and the type system presented in section 1.5, where the new constraint forms are effectively used.

The type system $HM(X)$ consists of a four-place *judgement* whose parameters are a constraint C , an environment Γ , an expression τ , and a type scheme σ . A judgement is written $C, \Gamma \vdash \tau : \sigma$ and is read: *under the assumptions C and Γ , the expression τ has type σ* . One may view C as an assumption about the judgement’s free type variables and Γ as an assumption about τ ’s free program identifiers. Please recall that Γ now contains *constrained* type schemes, and that σ is a *constrained* type scheme.

We would like the validity of a typing judgement to depend not on the

| | |
|--|---|
| $\frac{\Gamma(\mathbf{x}) = \sigma \quad C \Vdash \exists \sigma}{C, \Gamma \vdash \mathbf{x} : \sigma} \quad (\text{HMX-VAR})$ | $\frac{C \wedge D, \Gamma \vdash \mathbf{t} : \mathbb{T} \quad \bar{\mathbf{x}} \# \text{ftv}(C, \Gamma)}{C \wedge \exists \bar{\mathbf{x}}. D, \Gamma \vdash \mathbf{t} : \forall \bar{\mathbf{x}}[D]. \mathbb{T}} \quad (\text{HMX-GEN})$ |
| $\frac{C, (\Gamma; \mathbf{z} : \mathbb{T}) \vdash \mathbf{t} : \mathbb{T}'}{C, \Gamma \vdash \lambda \mathbf{z}. \mathbf{t} : \mathbb{T} \rightarrow \mathbb{T}'} \quad (\text{HMX-ABS})$ | $\frac{C, \Gamma \vdash \mathbf{t} : \forall \bar{\mathbf{x}}[D]. \mathbb{T}}{C \wedge D, \Gamma \vdash \mathbf{t} : \mathbb{T}} \quad (\text{HMX-INST})$ |
| $\frac{C, \Gamma \vdash \mathbf{t}_1 : \mathbb{T} \rightarrow \mathbb{T}' \quad C, \Gamma \vdash \mathbf{t}_2 : \mathbb{T}}{C, \Gamma \vdash \mathbf{t}_1 \mathbf{t}_2 : \mathbb{T}'} \quad (\text{HMX-APP})$ | $\frac{C, \Gamma \vdash \mathbf{t} : \mathbb{T} \quad C \Vdash \mathbb{T} \leq \mathbb{T}'}{C, \Gamma \vdash \mathbf{t} : \mathbb{T}'} \quad (\text{HMX-SUB})$ |
| $\frac{C, \Gamma \vdash \mathbf{t}_1 : \sigma \quad C, (\Gamma; \mathbf{z} : \sigma) \vdash \mathbf{t}_2 : \mathbb{T}}{C, \Gamma \vdash \text{let } \mathbf{z} = \mathbf{t}_1 \text{ in } \mathbf{t}_2 : \mathbb{T}} \quad (\text{HMX-LET})$ | $\frac{C, \Gamma \vdash \mathbf{t} : \sigma \quad \bar{\mathbf{x}} \# \text{ftv}(\Gamma, \sigma)}{\exists \bar{\mathbf{x}}. C, \Gamma \vdash \mathbf{t} : \sigma} \quad (\text{HMX-EXISTS})$ |

Figure 1-7: Typing rules for HM(X)

syntax, but only on the *meaning* of its constraint assumption. We enforce this point of view by considering judgements equal modulo equivalence of their constraint assumptions. In other words, the typing judgements $C, \Gamma \vdash \mathbf{t} : \sigma$ and $D, \Gamma \vdash \mathbf{t} : \sigma$ are considered identical when $C \equiv D$ holds. As a result, it does not make sense to analyze the syntax of a judgement's constraint assumption. A judgement is *valid*, or *holds*, if and only if it is derivable via the rules given in Figure 1-7. Please note that a valid judgement may involve an unsatisfiable constraint. A program \mathbf{t} is *well-typed* within the environment Γ if and only if a judgement of the form $C, \Gamma \vdash \mathbf{t} : \sigma$ holds for some *satisfiable* constraint C .

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 \mathbf{x} . The constraint C that appears in the conclusion must be strong enough to guarantee that σ has an instance; this is expressed by the second premise. This technical requirement is used in the proof of Lemma 1.4.1. 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 sub-derivation. We recall that the type \mathbb{T} may be viewed as the type scheme $\forall \emptyset[\text{true}]. \mathbb{T}$ (Definitions 1.2.18 and 1.3.2). As a result, types form a subset of type schemes, which implies that $\Gamma; \mathbf{z} : \mathbb{T}$ is a well-formed environment and $C, \Gamma \vdash \mathbf{t} : \mathbb{T}$ a well-formed typing judgement. To understand HMX-GEN, it is best to first consider the particular case where C is true. This yields the following, simpler rule:

$$\frac{D, \Gamma \vdash \mathbf{t} : \mathbb{T} \quad \bar{\mathbf{x}} \# \text{ftv}(\Gamma)}{\exists \bar{\mathbf{x}}. D, \Gamma \vdash \mathbf{t} : \forall \bar{\mathbf{x}}[D]. \mathbb{T}} \quad (\text{HMX-GEN}')$$

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 . Please 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 . The rule's conclusion carries the constraint $\exists \bar{X}.D$, thus recording the requirement that the newly formed type scheme should have an instance; again, this is used in the proof of Lemma 1.4.1. 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 constraint 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. 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 $\vec{X} \mapsto \vec{X}$. 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 the proof of Theorem 1.4.7 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. Please 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 judgements; it is a pleasant technical convenience. Indeed, because judgements are considered equal modulo constraint equivalence, constraints may be transparently *simplified* at any time. (By *simplifying* a constraint, we mean replacing it with an equiva-

lent 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 judgement $\mathbf{x} \leq \mathbf{y} \leq \mathbf{z}, \Gamma \vdash \mathbf{t} : \sigma$, but if it is known that \mathbf{y} does not appear free in Γ or σ , then HMX-EXISTS allows deriving $\exists \mathbf{y}.(\mathbf{x} \leq \mathbf{y} \leq \mathbf{z}), \Gamma \vdash \mathbf{t} : \sigma$, which is the same judgement as $\mathbf{x} \leq \mathbf{z}, \Gamma \vdash \mathbf{t} : \sigma$. Thus, an interesting simplification has been enabled. Please note that $\mathbf{x} \leq \mathbf{y} \leq \mathbf{z} \equiv \mathbf{x} \leq \mathbf{z}$ does *not* hold, while, according to C-EXTRANS, $\exists \mathbf{y}.(\mathbf{x} \leq \mathbf{y} \leq \mathbf{z}) \equiv \mathbf{x} \leq \mathbf{z}$ does.

We now establish a few simple properties of the type system $\text{HM}(X)$. Our first lemma is a minor technical property.

1.4.1 LEMMA: $C, \Gamma \vdash \mathbf{t} : \sigma$ implies $C \Vdash \exists \sigma$. □

The next lemma states that *strengthening* a judgement's constraint assumption preserves its validity. In other words, *weakening* a judgement preserves its validity. It is worth noting that in traditional presentations, which rely more heavily on type substitutions, the analogue of this result is a *type substitution* lemma; see for instance (Tofte, 1988, Lemma 2.7), (Leroy, 1992, Proposition 1.2), (Skalka and Pottier, 2002, Lemma 3.4). Here, the lemma further states that weakening a judgement does not alter the shape of its derivation, a useful property when reasoning by induction on type derivations.

1.4.2 LEMMA [WEAKENING]: If $C' \Vdash C$, then every derivation of $C, \Gamma \vdash \mathbf{t} : \sigma$ may be turned into a derivation of $C', \Gamma \vdash \mathbf{t} : \sigma$ with the same shape. □

Proof: The proof is by structural induction on a derivation of $C, \Gamma \vdash \mathbf{t} : \sigma$. In each proof case, we adopt the notations of Figure 1-7.

◦ *Case* HMX-VAR. The rule's conclusion is $C, \Gamma \vdash \mathbf{x} : \sigma$. Its premises are $\Gamma(\mathbf{x}) = \sigma$ **(1)** and $C \Vdash \exists \sigma$ **(2)**. By hypothesis, we have $C' \Vdash C$ **(3)**. By transitivity of entailment, **(3)** and **(2)** imply $C' \Vdash \exists \sigma$ **(4)**. By HMX-VAR, **(1)** and **(4)** yield $C', \Gamma \vdash \mathbf{x} : \sigma$.

◦ *Cases* HMX-ABS, HMX-APP, HMX-LET. By the induction hypothesis and by HMX-ABS, HMX-APP, or HMX-LET, respectively.

◦ *Case* HMX-GEN. The rule's conclusion is $C \wedge \exists \bar{x}.D, \Gamma \vdash \mathbf{t} : \forall \bar{x}[D].T$. Its premises are $C \wedge D, \Gamma \vdash \mathbf{t} : T$ **(1)** and $\bar{x} \# \text{ftv}(C, \Gamma)$ **(2)**. By hypothesis, we have $C' \Vdash C \wedge \exists \bar{x}.D$ **(3)**. We may assume, *w.l.o.g.*, $\bar{x} \# \text{ftv}(C')$ **(4)**. Applying the induction hypothesis to **(1)** and to the entailment assertion $C' \wedge C \wedge D \Vdash C \wedge D$, we obtain $C' \wedge C \wedge D, \Gamma \vdash \mathbf{t} : T$ **(5)**. By HMX-GEN, applied to **(5)**, **(2)** and **(4)**, we get $C' \wedge C \wedge \exists \bar{x}.D, \Gamma \vdash \mathbf{t} : \forall \bar{x}[D].T$ **(6)**. By **(3)** and C-DUP, the constraints $C' \wedge C \wedge \exists \bar{x}.D$ and C' are equivalent, so **(6)** is the goal $C', \Gamma \vdash \mathbf{t} : \forall \bar{x}[D].T$.

◦ *Case HMX-INST.* The rule's conclusion is $C \wedge D, \Gamma \vdash \tau : T$. Its premise is $C, \Gamma \vdash \tau : \forall \bar{x}[D].T$ **(1)**. By hypothesis, C' entails $C \wedge D$ **(2)**. Because **(2)** implies $C' \Vdash C$, the induction hypothesis may be applied to **(1)**, yielding $C', \Gamma \vdash \tau : \forall \bar{x}[D].T$ **(3)**. By HMX-INST, we obtain $C' \wedge D, \Gamma \vdash \tau : T$ **(4)**. Because **(2)** implies $C' \equiv C' \wedge D$, **(4)** is the goal $C', \Gamma \vdash \tau : T$.

◦ *Case HMX-SUB.* The rule's conclusion is $C, \Gamma \vdash \tau : T'$. Its premises are $C, \Gamma \vdash \tau : T$ **(1)** and $C \Vdash T \leq T'$ **(2)**. By hypothesis, we have $C' \Vdash C$ **(3)**. Applying the induction hypothesis to **(1)** and **(3)** yields $C', \Gamma \vdash \tau : T$ **(4)**. By transitivity of entailment, **(3)** and **(2)** imply $C' \Vdash T \leq T'$ **(5)**. By HMX-SUB, **(4)** and **(5)** yield $C', \Gamma \vdash \tau : T'$.

◦ *Case HMX-EXISTS.* The rule's conclusion is $\exists \bar{x}.C, \Gamma \vdash \tau : \sigma$. Its premises are $C, \Gamma \vdash \tau : \sigma$ **(1)** and $\bar{x} \# \text{ftv}(\Gamma, \sigma)$ **(2)**. By hypothesis, we have $C' \Vdash \exists \bar{x}.C$ **(3)**. We may assume, *w.l.o.g.*, $\bar{x} \# \text{ftv}(C')$ **(4)**. Applying the induction hypothesis to **(1)** and to the entailment assertion $C' \wedge C \Vdash C$, we obtain $C' \wedge C, \Gamma \vdash \tau : \sigma$ **(5)**. By HMX-EXISTS, **(5)** and **(2)** yield $\exists \bar{x}.(C' \wedge C), \Gamma \vdash \tau : \sigma$ **(6)**. By **(4)** and C-EXAND, the constraint $\exists \bar{x}.(C' \wedge C)$ is equivalent to $C' \wedge \exists \bar{x}.C$, which, by **(3)** and C-DUP, is equivalent to C' . Thus, **(6)** is the goal $C', \Gamma \vdash \tau : \sigma$. \square

We do not give a direct type soundness proof for $HM(X)$. Instead, in section 1.5, we prove that it is equivalent to another type system, which later is itself proven sound. A direct type soundness result, based on a denotational semantics, may be found in (Odersky, Sulzmann, and Wehr, 1999a). Another type soundness proof, which follows Wright and Felleisen's syntactic approach (1994b), 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 $HM(X)$

The presentation of $HM(X)$ given in Figure 1-7 has only four syntax-directed rules out of eight. It is a good specification of the type system, but it is far 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 1-8). It has the property that all judgements are of the form $C, \Gamma \vdash \tau : T$, rather than $C, \Gamma \vdash \tau : \sigma$. The following theorem states that the two presentations are indeed equivalent.

1.4.3 THEOREM: $C, \Gamma \vdash \tau : T$ is derivable via the rules of Figure 1-8 if and only if it is a valid $HM(X)$ judgement. \square

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

Figure 1-8: An alternate presentation of $\text{HM}(X)$

This theorem shows that the rule sets of Figures 1-7 and 1-8 derive the same monomorphic judgements, that is, the same judgements of the form $C, \Gamma \vdash t : T$. The fact that judgements 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; the first two exercises below shed some light on this issue.

- 1.4.4 EXERCISE [★★]: Show that both rule sets lead to the same set of *well-typed* programs. □
- 1.4.5 EXERCISE [★★]: Show that, if HMX-GEN is added to the rule set of Figure 1-8, then both rule sets derive exactly the same judgements. □
- 1.4.6 EXERCISE [★★★, \rightarrow]: 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 judgements 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 in (Clément, Despeyroux, Despeyroux, and Kahn, 1986). □

Relating $\text{HM}(X)$ with Damas and Milner's type system

In order to explain our interest in $\text{HM}(X)$, we wish to show that it is more general than Damas and Milner's type system. Since $\text{HM}(X)$ really is a *family* of type systems, we must make this statement more precise. First, every member of the $\text{HM}(X)$ family contains DM. Conversely, DM contains $\text{HM}(=)$, 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 judgements to $HM(X)$ judgements is essentially the identity: every valid DM judgement may be viewed as a valid $HM(X)$ judgement 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 $\forall \bar{X}[\text{true}].T$, so DM type schemes (resp. environments) form a subset of $HM(X)$ type schemes (resp. environments). Its proof is routine, except perhaps in the case of DM-INST, where it is shown how the effect of applying a substitution in DM is emulated by strengthening the current constraint in $HM(X)$.

1.4.7 THEOREM: If $\Gamma \vdash t : S$ holds in DM, then $\text{true}, \Gamma \vdash t : S$ holds in $HM(X)$. \square

Proof: The proof is by structural induction on a derivation of $\Gamma \vdash t : S$. In each proof case, we adopt the notations of Figure 1-3.

◦ *Case DM-VAR.* The rule's conclusion is $\Gamma \vdash x : S$. Its premise is $\Gamma(x) = S$ **(1)**. By definition and by C-EX*, the constraint $\exists S$ is equivalent to true . By applying HMX-VAR to (1) and to the assertion $\text{true} \Vdash \text{true}$, we obtain $\text{true}, \Gamma \vdash x : S$.

◦ *Cases DM-ABS, DM-APP, DM-LET.* By the induction hypothesis and by HMX-ABS, HMX-APP or HMX-LET, respectively.

◦ *Case DM-GEN.* The rule's conclusion is $\Gamma \vdash t : \forall \bar{X}.T$. Its premises are $\Gamma \vdash t : T$ **(1)** and $\bar{X} \# \text{ftv}(\Gamma)$ **(2)**. Applying the induction hypothesis to (1) yields $\text{true}, \Gamma \vdash t : T$ **(3)**. Furthermore, (2) implies $\bar{X} \# \text{ftv}(\text{true}, \Gamma)$ **(4)**. By HMX-GEN, (3) and (4) yield $\text{true}, \Gamma \vdash t : \forall \bar{X}[\text{true}].T$.

◦ *Case DM-INST.* The rule's conclusion is $\Gamma \vdash t : [\bar{X} \mapsto \vec{T}]T$. Its premise is $\Gamma \vdash t : \forall \bar{X}.T$ **(1)**. We may assume, *w.l.o.g.*, $\bar{X} \# \text{ftv}(\Gamma, \vec{T})$ **(2)**. Applying the induction hypothesis to (1) yields $\text{true}, \Gamma \vdash t : \forall \bar{X}[\text{true}].T$ **(3)**. By HMX-INST, (3) implies $\text{true}, \Gamma \vdash t : T$ **(4)**. By Lemma 1.4.2, we may weaken this judgement so as to obtain $\vec{X} = \vec{T}, \Gamma \vdash t : T$ **(5)**. Using C-EQ, C-EXTRANS, and C-EXAND, it is possible to establish $\vec{X} = \vec{T} \Vdash T = [\vec{X} \mapsto \vec{T}]T$ **(6)**. Applying HMX-SUB to (5) and (6), we find $\vec{X} = \vec{T}, \Gamma \vdash t : [\vec{X} \mapsto \vec{T}]T$ **(7)**. Last, (2) implies $\bar{X} \# \text{ftv}(\Gamma, [\vec{X} \mapsto \vec{T}]T)$ **(8)**. Applying HMX-EXISTS to (7) and (8), we obtain $\exists \bar{X}.(\vec{X} = \vec{T}), \Gamma \vdash t : [\vec{X} \mapsto \vec{T}]T$ **(9)**. By (2) and C-NAME, the constraint $\exists \bar{X}.(\vec{X} = \vec{T})$ is equivalent to true , so (9) is the goal. \square

We are now interested in proving that $HM(=)$, as defined above, is contained within DM. To this end, we must translate every $HM(=)$ judgement to a DM judgement. It quickly turns out that this is possible if the original judgement's constraint assumption is *satisfiable*.

We begin by explaining how an HM(=) is translated into a DM type scheme. Such a translation is made possible by 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 1.3.41), whose properties we make essential use of.

In fact, we must not only translate a type scheme, but also apply a type substitution to it. Instead of separating these steps, we perform both at once, and parameterize the translation by a type substitution θ . (It does not appear that separating them would help.) The definition of $\llbracket \sigma \rrbracket_\theta$ is somewhat involved: it is given in the statement of the following lemma, whose proof establishes that the definition is indeed well-formed.

1.4.8 LEMMA: Consider a type scheme σ and an idempotent type substitution θ such that $ftv(\sigma) \subseteq dom(\theta)$ **(1)** and $\exists \theta \Vdash \exists \sigma$ **(2)**. Write $\sigma = \forall \bar{X}[D].T$, where $\bar{X} \# \theta$ **(3)**. Then, there exists a type substitution θ' such that θ' extends θ , $dom(\theta')$ is $dom(\theta) \cup \bar{X}$, and θ' is a most general unifier of $\exists \theta \wedge D$. Let $\bar{Y} = ftv(\theta'(\bar{X})) \setminus range(\theta)$. Then, the *translation* of σ under θ , written $\llbracket \sigma \rrbracket_\theta$, is the DM type scheme $\forall \bar{Y}.\theta'(T)$. This is a well-formed definition. Furthermore, $ftv(\llbracket \sigma \rrbracket_\theta) \subseteq range(\theta)$ holds. \square

Proof: By (2), $\exists \theta$ is equivalent to $\exists \theta \wedge \exists \sigma$, which may be written $\exists \theta \wedge \exists \bar{X}.D$. By (3) and C-EXAND, this is $\exists \bar{X}.\exists \theta \wedge D$. Thus, because θ is a most general unifier of $\exists \theta$, θ is also a most general unifier of $\exists \bar{X}.\exists \theta \wedge D$ **(4)**. Furthermore, $ftv(\exists \bar{X}.\exists \theta \wedge D)$ is $ftv(\exists \theta \wedge \exists \sigma)$, which by definition of $\exists \theta$ and by (1) is a subset of $dom(\theta)$ **(5)**. By (4), (3), (5), and Lemma 1.3.45, there exists a type substitution θ' such that θ' extends θ **(6)** and θ' is a most general unifier of $\exists \theta \wedge D$ **(7)** and $dom(\theta') = dom(\theta) \cup \bar{X}$ **(8)**.

Let us now define $\bar{Y} = ftv(\theta'(\bar{X})) \setminus range(\theta)$ and $\llbracket \sigma \rrbracket_\theta = \forall \bar{Y}.\theta'(T)$. By (1), we have $ftv(T) \subseteq \bar{X} \cup dom(\theta)$. Applying θ' and exploiting (6), we find $ftv(\theta'(T)) \subseteq ftv(\theta'(\bar{X})) \cup range(\theta)$, which by definition of \bar{Y} may be written $ftv(\theta'(T)) \subseteq \bar{Y} \cup range(\theta)$. Subtracting \bar{Y} on each side, we find $ftv(\llbracket \sigma \rrbracket_\theta) \subseteq range(\theta)$ **(9)**.

To show that the definition of $\llbracket \sigma \rrbracket_\theta$ is valid, there remains to show that it does not depend on the choice of \bar{X} or θ' . To prove the former, it suffices to establish $\bar{X} \# ftv(\llbracket \sigma \rrbracket_\theta)$, which indeed follows from (3) and (9). As for the latter, because of the constraints imposed by (6), (7), and (8), and by Lemma 1.3.44, distinct choices of θ' may differ only by a renaming of $ftv(\theta'(\bar{X})) \setminus range(\theta)$, that is, \bar{Y} . So, we must check $\bar{Y} \# ftv(\llbracket \sigma \rrbracket_\theta)$, which holds by definition. \square

Please note that if σ is in fact a type T , where $ftv(T) \subseteq dom(\theta)$, then \bar{X} is empty, so θ' is θ , \bar{Y} is empty, and $\llbracket T \rrbracket_\theta = \theta(T)$. In other words, the translation of a type under θ is its image through θ . More generally, the translation of an unconstrained type scheme (that is, a type scheme whose constraint is **true**) is its image through θ , as stated by the following exercise.

1.4.9 EXERCISE [★★, \rightarrow]: Prove that $\llbracket \forall \bar{x}. T \rrbracket_\theta$, when defined, is $\theta(\forall \bar{x}. T)$. \square

The translation becomes more than a mere type substitution when applied to a nontrivial constrained type scheme. Some examples of this situation are given below.

1.4.10 EXAMPLE: Let $\sigma = \forall XY[X = Y \rightarrow Y].X$. Let θ be the identity substitution. The type scheme σ is closed and the constraint $\exists \sigma$ is equivalent to **true**, so $\llbracket \sigma \rrbracket_\theta$ is defined. We must find a type substitution θ' whose domain is XY and that is a most general unifier of $X = Y \rightarrow Y$. All such substitutions are of the form $[X \mapsto (Z \rightarrow Z), Y \mapsto Z]$, where Z is fresh. We have $ftv(\theta'(XY)) = Z$, whence $\llbracket \sigma \rrbracket_\theta = \forall Z.Z \rightarrow Z$. Note that the choice of Z does not matter, since it is bound in $\llbracket \sigma \rrbracket_\theta$. Roughly speaking, the effect of the translation was to replace the body X of the constrained type scheme with its most general solution under the constraint $X = Y \rightarrow Y$.

Let $\sigma = \forall XY_1[X = Y_1 \rightarrow Y_2].X$. Let $\theta = [Y_2 \mapsto Z_2]$. We have $ftv(\sigma) = Y_2 \subseteq dom(\theta)$. The constraint $\exists \sigma$ is equivalent to **true**, so $\llbracket \sigma \rrbracket_\theta$ is defined. We must find a type substitution θ' whose domain is XY_1Y_2 that extends θ and that is a most general unifier of $X = Y_1 \rightarrow Y_2$. All such substitutions are of the form $[X \mapsto (Z_1 \rightarrow Z_2), Y_1 \mapsto Z_1, Y_2 \mapsto Z_2]$, where Z_1 is fresh. We have $ftv(\theta'(XY_1)) \setminus range(\theta) = Z_1Z_2 \setminus Z_2 = Z_1$, whence $\llbracket \sigma \rrbracket_\theta = \forall Z_1.Z_1 \rightarrow Z_2$. The type variable Z_2 is *not* universally quantified—even though it appears in the image of X , which *was* universally quantified in σ —because Z_2 is the image of Y_2 , which was free in σ . \square

Before attacking the main theorem, let us establish a couple of technical properties of the translation. First, $\llbracket \sigma \rrbracket_\theta$ is insensitive to the behavior of θ outside $ftv(\sigma)$, a natural property, since our informal intent is for θ to be applied to σ .

1.4.11 LEMMA: If θ_1 and θ_2 coincide on $ftv(\sigma)$, then $\llbracket \sigma \rrbracket_{\theta_1}$ and $\llbracket \sigma \rrbracket_{\theta_2}$ are either both undefined, or both defined and identical. \square

Second, if $C \vdash \sigma \preceq T'$ holds, then the translations of σ and T' under a most general unifier of C are in Damas and Milner's instance relation. One might say, roughly speaking, that the instance relation is preserved by the translation.

1.4.12 LEMMA: Let $ftv(\sigma, T') \subseteq dom(\theta)$ **(1)** and $\exists \theta \vdash \exists \sigma$ **(2)**. Let $\exists \theta \vdash \sigma \preceq T'$ **(3)**. Then, $\theta(T')$ is an instance of the DM type scheme $\llbracket \sigma \rrbracket_\theta$. \square

Proof: Write $\sigma = \forall \bar{x}[D].T$, where $\bar{x} \# \theta$ **(4)** and $\bar{x} \# ftv(T')$ **(5)**. By (1), (2), and (4), one may define θ' , \bar{y} , and $\llbracket \sigma \rrbracket_\theta$ exactly as in the statement of Lemma 1.4.8. By (5) and Definition 1.3.3, (3) is synonymous with $\exists \theta \vdash \exists \bar{x}.(D \wedge$

$T = T'$). Reasoning in the same manner as in the first paragraph of the proof of Lemma 1.4.8, we find that there exists a type substitution θ'' such that θ'' extends θ , $\text{dom}(\theta'')$ is $\text{dom}(\theta) \cup \bar{X}$, and θ'' is a most general unifier of $\exists\theta \wedge D \wedge T = T'$.

We have $\text{dom}(\theta') = \text{dom}(\theta'')$ **(6)**. Furthermore, θ' is a most general unifier of $\exists\theta \wedge D$, while θ'' is a most general unifier of $\exists\theta \wedge D \wedge T = T'$, which implies $\exists\theta'' \Vdash \exists\theta'$ **(7)**. By Lemma 1.3.39, θ'' refines θ' . That is, there exists a type substitution φ such that θ'' is the restriction of $\varphi \circ \theta'$ to $\text{dom}(\theta) \cup \bar{X}$ **(8)**. We may require $\text{dom}(\varphi) \subseteq \text{range}(\theta) \cup \text{ftv}(\theta'(\bar{X}))$ **(9)** without compromising (8).

Consider $X \in \text{dom}(\theta)$. Because θ'' extends θ , we have $\theta''(X) = \theta(X)$ **(10)**. Furthermore, by (8), we have $\theta''(X) = (\varphi \circ \theta')(X) = (\varphi \circ \theta)(X)$ **(11)**. Using (10) and (11), we find $\theta(X) = \varphi(\theta(X))$. Because this holds for every $X \in \text{dom}(\theta)$, φ must be the identity over $\text{range}(\theta)$; that is, $\text{dom}(\varphi) \# \text{range}(\theta)$ **(12)** holds. Combining (9) and (12), we find $\text{dom}(\varphi) \subseteq \text{ftv}(\theta'(\bar{X})) \setminus \text{range}(\theta)$, that is, $\text{dom}(\varphi) \subseteq \bar{Y}$ **(13)**.

By construction of θ'' , we have $\exists\theta'' \Vdash T = T'$. By Lemma 1.3.29, this implies $\theta''(\exists\theta'') \Vdash \theta''(T) = \theta''(T')$, which by Lemma 1.3.42 may be read $\text{true} \Vdash \theta''(T) = \theta''(T')$. By Lemma 1.3.32, $\theta''(T)$ and $\theta''(T')$ coincide. Because by (1) $\text{ftv}(T)$ is a subset of $\text{dom}(\theta) \cup \bar{X}$ and by (8), the former may be written $\varphi(\theta'(T))$. By (1) and because θ'' extends θ , the latter is $\theta(T')$. Thus, we have $\varphi(\theta'(T)) = \theta(T')$. Together with (13), this establishes that $\theta(T')$ is an instance of $\forall\bar{Y}.\theta'(T)$, that is, $\llbracket\sigma\rrbracket_\theta$. \square

We extend the translation to environments as follows. $\llbracket\emptyset\rrbracket_\theta$ is \emptyset . If $\exists\theta \Vdash \exists\sigma$ holds, then $\llbracket\Gamma; x : \sigma\rrbracket_\theta$ is $\llbracket\Gamma\rrbracket_\theta; x : \llbracket\sigma\rrbracket_\theta$, otherwise it is $\llbracket\Gamma\rrbracket_\theta$. Notice that $\llbracket\Gamma\rrbracket_\theta$ contains fewer bindings than Γ , which ensures that bindings $x : \sigma$ for which $\exists\theta \Vdash \exists\sigma$ does not hold will not be used in the translation. Please note that $\llbracket\Gamma\rrbracket_\theta$ is defined when $\text{ftv}(\Gamma) \subseteq \text{dom}(\theta)$ holds.

We are now ready to prove the main theorem. Please note that, by requiring θ to be a most general unifier of C , we also require C to be satisfiable. Judgements that carry an unsatisfiable constraint cannot be translated.

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

Proof: Let us first remark that, by Lemma 1.4.1, we have $C \Vdash \exists\sigma$. This may be written $\exists\theta \Vdash \exists\sigma$, which guarantees that $\llbracket\sigma\rrbracket_\theta$ is defined. The proof is by structural induction on an $\text{HM}(=)$ typing derivation. We assume that the derivation is expressed in terms of the rules of Figure 1-8, but split HMD-LETGEN into HMX-LET and HMX-GEN for the sake of readability.

◦ *Case HMD-VARINST.* The rule's conclusion is $C \wedge D, \Gamma \vdash x : T$. By hypothesis, θ is a most general unifier of $C \wedge D$ **(1)**, and $\text{ftv}(T) \subseteq \text{dom}(\theta)$ **(2)**

holds. The rule's premise is $\Gamma(x) = \sigma$ **(3)**, where σ stands for $\forall \bar{x}[D].T$. By (1), we have $\exists \theta \equiv C \wedge D \Vdash D \Vdash \exists \bar{x}.D \equiv \exists \sigma$ **(4)**. Furthermore, we have $ftv(\sigma) \subseteq ftv(\Gamma) \subseteq dom(\theta)$ **(5)**. These facts show that $\llbracket \sigma \rrbracket_\theta$ is defined. Together with (3), this implies $\llbracket \Gamma \rrbracket_\theta(x) = \llbracket \sigma \rrbracket_\theta$. By DM-VAR, $\llbracket \Gamma \rrbracket_\theta \vdash x : \llbracket \sigma \rrbracket_\theta$ **(6)** follows. Now, by Lemma 1.3.19, we have $D \Vdash \sigma \preceq T$, which, combined with $\exists \theta \Vdash D$, yields $\exists \theta \Vdash \sigma \preceq T$ **(7)**. By (7), (4), (5), (2), and Lemma 1.4.12, we find that $\theta(T)$ is an instance of $\llbracket \sigma \rrbracket_\theta$. Thus, applying DM-INST to (6) yields $\llbracket \Gamma \rrbracket_\theta \vdash \mathbf{t} : \theta(T)$.

◦ *Case HMD-ABS.* The rule's conclusion is $C, \Gamma \vdash \lambda z. \mathbf{t} : T \rightarrow T'$. Its premise is $C, (\Gamma; z : T) \vdash \mathbf{t} : T'$. Applying the induction hypothesis to it yields $\llbracket \Gamma \rrbracket_\theta; z : \theta(T) \vdash \mathbf{t} : \theta(T')$. By DM-ABS, this implies $\llbracket \Gamma \rrbracket_\theta \vdash \lambda z. \mathbf{t} : \theta(T) \rightarrow \theta(T')$, that is, $\llbracket \Gamma \rrbracket_\theta \vdash \lambda z. \mathbf{t} : \theta(T \rightarrow T')$.

◦ *Case HMD-APP.* By an extension of $dom(\theta)$ to include $ftv(T)$, by the induction hypothesis, and by DM-APP.

◦ *Case HMX-LET.* By an extension of $dom(\theta)$ to include $ftv(\sigma)$, by the induction hypothesis, and by DM-LET.

◦ *Case HMX-GEN.* The rule's conclusion is $C \wedge \exists \sigma, \Gamma \vdash \mathbf{t} : \sigma$, where σ stands for $\forall \bar{x}[D].T$. By hypothesis, θ is a most general unifier of $C \wedge \exists \sigma$ **(1)**, and $ftv(\Gamma, \sigma) \subseteq dom(\theta)$ **(2)** holds. The rule's premises are $C \wedge D, \Gamma \vdash \mathbf{t} : T$ **(3)** and $\bar{x} \# ftv(C, \Gamma)$ **(4)**. We may further assume, *w.l.o.g.*, $\bar{x} \# \theta$ **(5)**. Given (1), (2), and (5), we may define θ' and \bar{y} exactly as in Lemma 1.4.8. Then, θ' is a most general unifier of $\exists \theta \wedge D$, that is, $C \wedge D$. Furthermore, $dom(\theta')$ is $dom(\theta) \cup \bar{x}$, which by (2) is a superset of $ftv(\Gamma, T)$. Thus, the induction hypothesis applies to θ' and to (3), yielding $\llbracket \Gamma \rrbracket_{\theta'} \vdash \mathbf{t} : \theta'(T)$. Because θ' extends θ , by (2) and by Lemma 1.4.11, this may be read $\llbracket \Gamma \rrbracket_\theta \vdash \mathbf{t} : \theta'(T)$ **(6)**. According to Lemma 1.4.8, we have $ftv(\llbracket \Gamma \rrbracket_\theta) \subseteq range(\theta)$, which by construction of \bar{y} implies $\bar{y} \# ftv(\llbracket \Gamma \rrbracket_\theta)$ **(7)**. By DM-GEN, (6) and (7) yield $\llbracket \Gamma \rrbracket_\theta \vdash \mathbf{t} : \forall \bar{y}. \theta'(T)$, that is, $\llbracket \Gamma \rrbracket_\theta \vdash \mathbf{t} : \llbracket \sigma \rrbracket_\theta$.

◦ *Case HMD-SUB.* The rule's conclusion is $C, \Gamma \vdash \mathbf{t} : T'$. By hypothesis, θ is a most general unifier of C **(1)**, and $ftv(\Gamma, T') \subseteq dom(\theta)$ **(2)** holds. The goal is $\llbracket \Gamma \rrbracket_\theta \vdash \mathbf{t} : \theta(T')$ **(3)**. The rule's premises are $C, \Gamma \vdash \mathbf{t} : T$ **(4)** and $C \Vdash T = T'$ **(5)**. We may assume, *w.l.o.g.*, $ftv(T) \# range(\theta)$ **(6)**. Then, by (6) and Lemma 1.3.43, we may extend the domain of θ , so as to achieve $ftv(T) \subseteq dom(\theta)$ **(7)**, without compromising (1) or (2) or affecting the goal (3). By (1), (2), and (7), the induction hypothesis applies to (4), yielding $\llbracket \Gamma \rrbracket_\theta \vdash \mathbf{t} : \theta(T)$ **(8)**. Now, thanks to (1), (5) may be read $\exists \theta \Vdash T = T'$, which by Lemmas 1.3.29 and 1.3.42 implies $true \Vdash \theta(T) = \theta(T')$. Then, Lemma 1.3.32 shows that $\theta(T)$ and $\theta(T')$ coincide. As a result, (8) is the goal (3).

◦ *Case HMD-EXISTS.* The rule's conclusion is $\exists \bar{x}. C, \Gamma \vdash \mathbf{t} : T$. By hypothesis, θ is a most general unifier of $\exists \bar{x}. C$ **(1)**, and $ftv(\Gamma, T) \subseteq dom(\theta)$ **(2)** holds. The

rule's premises are $C, \Gamma \vdash \tau : T$ **(3)** and $\bar{x} \# ftv(\Gamma, T)$. We may assume, *w.l.o.g.*, $\bar{x} \# \theta$ **(4)**. As in the previous case, we may extend the domain of θ to guarantee $ftv(\exists \bar{x}. C) \subseteq dom(\theta)$ **(5)**. By (1), (4), (5), and Lemma 1.3.45, there exists a type substitution θ' such that θ' extends θ **(6)** and θ' is a most general unifier of C . Applying the induction hypothesis to θ' and to (3) yields $[\Gamma]_{\theta'} \vdash \tau : \theta'(T)$. By (2), (6), and Lemma 1.4.11, this may be read $[\Gamma]_{\theta} \vdash \tau : \theta(T)$. \square

Together, Theorems 1.4.7 and 1.4.13 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 HM(=) is a constraint-based formulation of DM. Furthermore, Theorem 1.4.7 states that every member of the HM(X) family is an extension of DM. This explains our double interest in HM(X), as an alternate formulation of DM, which we believe is more pleasant, for reasons already discussed, and as a more expressive framework.

1.5 A purely constraint-based type system: PCB(X)

In the previous section, we have presented HM(X), an elegant constraint-based extension of Damas and Milner's type system. However, HM(X), as described there, suffers from a drawback. A typing judgement involves both a constraint, which represents an assumption about its free type variables, and an environment, which represents an assumption about its free program identifiers. At a let node, HMD-LETGEN turns a part of the current constraint, namely D , into a type scheme, namely $\forall \bar{x}[D].T$, and stores it into the environment. Then, at every occurrence of the let-bound variable, HMD-VARINST retrieves this type scheme from the environment and adds a copy of D back to the current constraint. In practice, 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, it appears that, in order to preserve efficiency, constraint generation and constraint simplification cannot be separated.

Of course, in practice, it is not difficult to intermix these phases, so the problem is not technical, but pedagogical. Indeed, we argued earlier that it is natural and desirable to separate them. *Type scheme introduction and elimination constraints*, which we introduced in Section 1.3 but did not use in the specification of HM(X), are intended as a means of solving this problem. In the present section, we exploit them to give a novel formulation of HM(X), which no longer requires copying constraints back and forth between the environment and the constraint assumption. In fact, the environment is

| | |
|--|--|
| $\frac{C \Vdash x \preceq T}{C \vdash x : T} \quad (\text{VAR})$ | $\frac{C_1 \vdash t_1 : T_1 \quad C_2 \vdash t_2 : T_2}{\text{let } z : \forall V[C_1].T_1 \text{ in } C_2 \vdash \text{let } z = t_1 \text{ in } t_2 : T_2} \quad (\text{LET})$ |
| $\frac{C \vdash t : T'}{\text{let } z : T \text{ in } C \vdash \lambda z.t : T \rightarrow T'} \quad (\text{ABS})$ | $\frac{C \vdash t : T}{C \wedge T \preceq T' \vdash t : T'} \quad (\text{SUB})$ |
| $\frac{C_1 \vdash t_1 : T \rightarrow T' \quad C_2 \vdash t_2 : T}{C_1 \wedge C_2 \vdash t_1 t_2 : T'} \quad (\text{APP})$ | $\frac{C \vdash t : T \quad \bar{x} \# \text{fv}(T)}{\exists \bar{x}. C \vdash t : T} \quad (\text{EXISTS})$ |

Figure 1-9: Typing rules for PCB(X)

suppressed altogether: taking advantage of the new constraint forms, we encode information about program identifiers within the constraint assumption.

Presentation

We now employ the full constraint language (Section 1.3). Typing judgements take the form $C \vdash t : T$, where C may have free type variables *and* free program identifiers. The rules that allow deriving such judgements appear in Figure 1-9. As before, we identify judgements up to constraint equivalence.

Let us review the rules. VAR states that x has type T under any constraint that entails $x \preceq T$. Note that we no longer consult the type scheme associated with x in the environment—indeed, there is no environment. Instead, we let the constraint assumption record the fact that the type scheme should admit T as one of its instances. Thus, in a judgement $C \vdash t : T$, any program identifier that occurs free within t typically also occurs free within C . ABS requires the body t of a λ -abstraction to have type T' under assumption C . Although no explicit assumption about z appears in the premise, C typically contains a number of instantiation constraints bearing on z , of the form $z \preceq T_i$. In the rule's conclusion, C is wrapped within the prefix $\text{let } z : T \text{ in } []$, where T is the type assigned to z . This effectively requires every T_i to denote a supertype of T , as desired. Please note that z does not occur free in the constraint $\text{let } z : T \text{ in } C$, which is natural, since it does not occur free in $\lambda z.t$. APP exhibits a minor stylistic difference with respect to HMX-APP: its constraint assumption is split between its premises. It is not difficult to prove that, when weakening holds (see Lemma 1.5.2 below), this choice does not affect the set of valid judgements. This new presentation encourages reading the rules in Figure 1-9 as the specification of an algorithm, which, given t and T , pro-

duces C such that $C \vdash \tau : T$ holds. In the case of APP, the algorithm invokes itself recursively for each of the two subexpressions, yielding the constraints C_1 and C_2 , then constructs their conjunction. LET is analogous to ABS: by wrapping C_2 within a let prefix, it gives meaning to the instantiation constraints bearing on z within C_2 . The difference is that z may now be assigned a type scheme, as opposed to a monotype. An appropriate type scheme is built in the most straightforward manner from the constraint C_1 and the type T_1 that describe τ_1 . All of the type variables that appear free in the left-hand premise are generalized, hence the notation $\forall\mathcal{V}[C_1].T_1$, which is a convenient shorthand for $\forall ftv(C_1, T_1)[C_1].T_1$. The side-condition that “type variables that occur free in the environment must not be generalized”, which was present in DM and HM(X), naturally disappears, since judgements no longer involve an environment. SUB again exhibits a minor stylistic difference with respect to HMX-SUB: the comments made about APP above apply here as well. EXISTS is essentially identical to HMX-EXISTS.

In the standard specification of HM(X), HMD-ABS and HMD-LETGEN accumulate information in the environment. Through the environment, this information is made available to HMD-VARINST, which retrieves and copies it. Here, instead, no information is explicitly transmitted. Where a program identifier is bound, a type scheme introduction constraint is built; where a program identifier is used, a type scheme instantiation constraint is produced. The two are related only by our definition of the meaning of constraints.

The reader may be puzzled by the fact that LET allows *all* type variables that occur free in its left-hand premise to be generalized. The following exercise sheds some light on this issue.

- 1.5.1 EXERCISE [\star , RECOMMENDED]: Build a type derivation for the expression $\lambda z_1. \text{let } z_2 = z_1 \text{ in } z_2$ within PCB(X). Draw a comparison with the solution of Exercise 1.2.21. \square

The following lemma is an analogue of Lemma 1.4.2.

- 1.5.2 LEMMA [WEAKENING]: If $C' \Vdash C$, then every derivation of $C \vdash \tau : T$ may be turned into a derivation of $C' \vdash \tau : T$ with the same shape. \square

Proof: The proof is by structural induction on a derivation of $C \vdash \tau : T$. In each proof case, we adopt the notations of Figure 1-9.

◦ *Case VAR.* By transitivity of entailment.

◦ *Case ABS.* The rule’s conclusion is $\text{let } z : T \text{ in } C \vdash \lambda z. \tau : T \rightarrow T'$ **(1)**. By hypothesis, we have $C' \Vdash \text{let } z : T \text{ in } C$ **(2)**. We may assume, *w.l.o.g.*, $z \notin fpi(C')$ **(3)**. The rule’s premise is $C \vdash \tau : T'$ **(4)**. Applying the induction hypothesis to **(4)** yields $C \wedge C' \vdash \tau : T'$, which by ABS implies $\text{let } z : T \text{ in } (C \wedge$

$C' \vdash \lambda z.t : T \rightarrow T'$ (5). By (3) and C-INAND*, let $z : T$ in $(C \wedge C')$ is equivalent to $(\text{let } z : T \text{ in } C) \wedge C'$, which by (2) and C-DUP is equivalent to C' . Thus, (5) is the goal $C' \vdash \lambda z.t : T \rightarrow T'$.

◦ *Case APP.* By applying the induction hypothesis to each premise, using the fact that $C' \Vdash C_1 \wedge C_2$ implies $C' \Vdash C_1$ and $C' \Vdash C_2$.

◦ *Case LET.* Analogous to the case of ABS. The induction hypothesis is applied to the second premise only.

◦ *Case SUB.* Analogous to the case of APP.

◦ *Case EXISTS.* See the corresponding case in the proof of Lemma 1.4.2. \square

Relating PCB(X) with HM(X)

Let us now provide evidence for our claim that PCB(X) is an alternate presentation of HM(X). The next two theorems define an effective translation from HM(X) to PCB(X) and back.

The first theorem states that if, within HM(X), t has type T under assumptions C and Γ , then, within PCB(X), t also has type T , under some assumption C' . The relationship $C \Vdash \text{let } \Gamma \text{ in } C'$ states that C entails the residual constraint obtained by confronting Γ , which provides information about the free program identifiers in t , with C' , which contains instantiation constraints bearing on these program identifiers. The statement requires C and Γ to have no free program identifiers, which is natural, since they are part of an HM(X) judgement. 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 1.7.3.

1.5.3 THEOREM: Let $C \Vdash \exists \Gamma$. Assume $fpi(C, \Gamma) = \emptyset$. If $C, \Gamma \vdash t : T$ holds in HM(X), then there exists a constraint C' such that $C' \vdash t : T$ holds in PCB(X) and C entails $\text{let } \Gamma \text{ in } C'$. \square

Proof: The proof is by structural induction on a derivation of $C, \Gamma \vdash t : T$. In each proof case, we adopt the notations of Figure 1-8.

◦ *Case HMD-VARINST.* The rule's conclusion is $C \wedge D, \Gamma \vdash x : T$. By hypothesis, we have $C \wedge D \Vdash \exists \Gamma$ (1) and $fpi(C, D, \Gamma) = \emptyset$ (2). The rule's premise is $\Gamma(x) = \forall \bar{X}[D].T$ (3). By VAR, we have $x \preceq T \vdash x : T$, so there remains to establish $C \wedge D \Vdash \text{let } \Gamma \text{ in } x \preceq T$ (4). By (3), (2), and C-INID, the constraint $\text{let } \Gamma \text{ in } x \preceq T$ is equivalent to $\text{let } \Gamma \text{ in } \forall \bar{X}[D].T \preceq T$, which, by (2) and C-IN*, is itself equivalent to $\exists \Gamma \wedge \forall \bar{X}[D].T \preceq T$ (5). By (1) and Lemma 1.3.19, $C \wedge D$ entails (5). We have established (4).

◦ *Case HMD-ABS.* The rule's conclusion is $C, \Gamma \vdash \lambda z.t : T \rightarrow T'$. Its premise is $C, (\Gamma; z : T) \vdash t : T'$ (1). The constraints $\exists \Gamma$ and $\exists(\Gamma; z : T)$ are equivalent,

so the induction hypothesis applies to (1) and yields a constraint C' such that $C' \vdash \mathfrak{t} : T'$ (2) and $C \Vdash \text{let } \Gamma; z : T \text{ in } C'$ (3). Applying ABS to (2) yields $\text{let } z : T \text{ in } C' \vdash \lambda z. \mathfrak{t} : T \rightarrow T'$. There remains to check that C entails $\text{let } \Gamma \text{ in } \text{let } z : T \text{ in } C'$ —but that is precisely (3).

◦ *Case HMD-APP.* The rule's conclusion is $C, \Gamma \vdash \mathfrak{t}_1 \mathfrak{t}_2 : T'$. Its premises are $C, \Gamma \vdash \mathfrak{t}_1 : T \rightarrow T'$ (1) and $C, \Gamma \vdash \mathfrak{t}_2 : T$ (2). Applying the induction hypothesis to (1) and (2), we obtain constraints C'_1 and C'_2 such that $C'_1 \vdash \mathfrak{t}_1 : T \rightarrow T'$ (3) and $C'_2 \vdash \mathfrak{t}_2 : T$ (4) and $C \Vdash \text{let } \Gamma \text{ in } C'_1$ (5) and $C \Vdash \text{let } \Gamma \text{ in } C'_2$ (6). By APP, (3) and (4) imply $C'_1 \wedge C'_2 \vdash \mathfrak{t}_1 \mathfrak{t}_2 : T'$. Furthermore, by C-INAND, (5) and (6) yield $C \Vdash \text{let } \Gamma \text{ in } C'_1 \wedge C'_2$.

◦ *Case HMD-LETGEN.* The rule's conclusion is $C \wedge \exists \bar{x}. D, \Gamma \vdash \text{let } z = \mathfrak{t}_1 \text{ in } \mathfrak{t}_2 : T_2$. By hypothesis, we have $C \wedge \exists \bar{x}. D \Vdash \exists \Gamma$ (1) and $fpi(C, D, \Gamma) = \emptyset$ (2). The rule's premises are $C \wedge D, \Gamma \vdash \mathfrak{t}_1 : T_1$ (3) and $\bar{x} \# ftv(C, \Gamma)$ (4) and $C \wedge \exists \bar{x}. D, \Gamma' \vdash \mathfrak{t}_2 : T_2$ (5), where Γ' is $\Gamma; z : \forall \bar{x}[D]. T_1$. Applying the induction hypothesis to (3) yields a constraint C'_1 such that $C'_1 \vdash \mathfrak{t}_1 : T_1$ (6) and $C \wedge D \Vdash \text{let } \Gamma \text{ in } C'_1$ (7). By (1), (2), and C-IN*, we have $C \wedge \exists \bar{x}. D \Vdash \exists \Gamma'$. Thus, the induction hypothesis applies to (5) and yields a constraint C'_2 such that $C'_2 \vdash \mathfrak{t}_2 : T_2$ (8) and $C \wedge \exists \bar{x}. D \Vdash \text{let } \Gamma' \text{ in } C'_2$ (9). By LET, (6) and (8) imply $\text{let } z : \forall \mathcal{V}[C'_1]. T_1 \text{ in } C'_2 \vdash \text{let } z = \mathfrak{t}_1 \text{ in } \mathfrak{t}_2 : T_2$ (10). By Lemmas 1.3.25 and 1.5.2, (10) yields $\text{let } z : \forall \bar{x}[C'_1]. T_1 \text{ in } C'_2 \vdash \text{let } z = \mathfrak{t}_1 \text{ in } \mathfrak{t}_2 : T_2$ (11), where the universal quantification is over \bar{x} only. There remains to establish that $C \wedge \exists \bar{x}. D$ entails $\text{let } \Gamma; z : \forall \bar{x}[C'_1]. T_1 \text{ in } C'_2$ (12). By (4), (2), and C-LETDUP, the constraint (12) is equivalent to $\text{let } \Gamma; z : \forall \bar{x}[\text{let } \Gamma \text{ in } C'_1]. T_1 \text{ in } C'_2$. By (7), this constraint is entailed by $\text{let } \Gamma; z : \forall \bar{x}[C \wedge D]. T_1 \text{ in } C'_2$, which by (4) and C-LETAND, is equivalent to $C \wedge \text{let } \Gamma; z : \forall \bar{x}[D]. T_1 \text{ in } C'_2$, that is, $C \wedge \text{let } \Gamma' \text{ in } C'_2$. By (9), this constraint is entailed by $C \wedge \exists \bar{x}. D$.

◦ *Case HMD-SUB.* The rule's conclusion is $C, \Gamma \vdash \mathfrak{t} : T'$. Its premises are $C, \Gamma \vdash \mathfrak{t} : T$ (1) and $C \Vdash T \leq T'$ (2). Applying the induction hypothesis to (1) yields a constraint C' such that $C' \vdash \mathfrak{t} : T$ (3) and $C \Vdash \text{let } \Gamma \text{ in } C'$ (4). By SUB, (3) implies $C' \wedge T \leq T' \vdash \mathfrak{t} : T'$. There remains to establish $C \Vdash \text{let } \Gamma \text{ in } (C' \wedge T \leq T')$, which follows from (4) and (2) by C-INAND*.

◦ *Case HMD-EXISTS.* The rule's conclusion is $\exists \bar{x}. C, \Gamma \vdash \mathfrak{t} : T$. Its premises are $C, \Gamma \vdash \mathfrak{t} : T$ (1) and $\bar{x} \# ftv(\Gamma, T)$ (2). By hypothesis, we have $\exists \bar{x}. C \Vdash \exists \Gamma$, which by Lemma 1.3.16 implies $C \Vdash \exists \Gamma$. Thus, the induction hypothesis applies to (1) and yields a constraint C' such that $C' \vdash \mathfrak{t} : T$ (3) and $C \Vdash \text{let } \Gamma \text{ in } C'$ (4). By EXISTS, (3) and (2) imply $\exists \bar{x}. C' \vdash \mathfrak{t} : T$. There remains to establish $\exists \bar{x}. C \Vdash \text{let } \Gamma \text{ in } \exists \bar{x}. C'$. By congruence of entailment, (4) implies $\exists \bar{x}. C \Vdash \exists \bar{x}. \text{let } \Gamma \text{ in } C'$. The result follows by (2) and C-INEX. \square

The second theorem states that if, within $PCB(X)$, \mathfrak{t} has type T under assumption C , then, within $HM(X)$, \mathfrak{t} also has type T , under assumptions

let Γ in C and Γ . The idea is simple: the constraint C represents a combined assumption about the initial judgement's free type variables and free program identifiers. In $\text{HM}(X)$, these two kinds of assumptions must be maintained separately. So, we split them into a pair of an environment Γ , which may be chosen arbitrarily, provided it satisfies $\text{fpi}(C) \subseteq \text{dpi}(\Gamma)$ —that is, provided it defines all program variables of interest, and the residual constraint $\text{let } \Gamma \text{ in } C$, which has no free program identifiers, thus represents an assumption about the new judgement's type variables only. Distinct choices of Γ give rise to distinct $\text{HM}(X)$ judgements, which may be incomparable; this is related to the fact that ML-the-type-system does not have principal typings (Jim, 1995). Again, the hypothesis $\text{fpi}(\Gamma) = \text{fpi}(\text{let } \Gamma \text{ in } C) = \emptyset$ is natural, since we wish Γ and $\text{let } \Gamma \text{ in } C$ to appear in an $\text{HM}(X)$ judgement.

1.5.4 THEOREM: Assume $\text{fpi}(\Gamma) = \text{fpi}(\text{let } \Gamma \text{ in } C) = \emptyset$ and $C \not\equiv \text{false}$. If $C \vdash \tau : T$ holds in $\text{PCB}(X)$, then $\text{let } \Gamma \text{ in } C, \Gamma \vdash \tau : T$ holds in $\text{HM}(X)$. \square

Proof: The proof is by structural induction on a derivation of $C \vdash \tau : T$. In each proof case, we adopt the notations of Figure 1-9.

By Lemma 1.3.30, the hypothesis $C \not\equiv \text{false}$ is preserved whenever the induction hypothesis is invoked. It is explicitly used only in case VAR, where it guarantees that the identifier at hand is bound in Γ .

◦ *Case VAR.* The rule's conclusion is $C \vdash x : T$. Its premise is $C \Vdash x \preceq T$ **(1)**. By Lemma 1.3.24, (1) and the hypothesis $C \not\equiv \text{false}$ imply $x \in \text{fpi}(C)$. Because $\text{let } \Gamma \text{ in } C$ has no free program identifiers, this implies $x \in \text{dpi}(\Gamma)$, that is, the environment Γ must define x . Let $\Gamma(x) = \forall \bar{X}[D].T'$ **(2)**, where $\bar{X} \# \text{ftv}(\Gamma, T)$ **(3)**. By (2), HMD-VARINST , and HMD-SUB , we have $D \wedge T' \leq T, \Gamma \vdash x : T$. By (3) and HMD-EXISTS , this implies $\exists \bar{X}.(D \wedge T' \leq T), \Gamma \vdash x : T$ **(4)**. Now, by (3), the constraint $\exists \bar{X}.(D \wedge T' \leq T)$ may be written $\forall \bar{X}[D].T' \preceq T$ **(5)**. The hypothesis $\text{fpi}(\Gamma) = \emptyset$ implies $\text{fpi}(D) = \emptyset$ **(6)**. By (6), C-INID and C-IN^* , (5) is equivalent to $\text{let } \Gamma \text{ in } x \preceq T$. Thus, (4) may be written $\text{let } \Gamma \text{ in } x \preceq T, \Gamma \vdash x : T$. By (1), by congruence of entailment, and by Lemma 1.4.2, this implies $\text{let } \Gamma \text{ in } C, \Gamma \vdash x : T$.

◦ *Case ABS.* The rule's conclusion is $\text{let } z : T \text{ in } C \vdash \lambda z. \tau : T \rightarrow T'$. Its premise is $C \vdash \tau : T'$ **(1)**. Let Γ' stand for $\Gamma; z : T$. Applying the induction hypothesis to (1) yields $\text{let } \Gamma' \text{ in } C, \Gamma' \vdash \tau : T'$. By HMD-ABS , this implies $\text{let } \Gamma' \text{ in } C, \Gamma \vdash \lambda z. \tau : T \rightarrow T'$.

◦ *Case APP.* The rule's conclusion is $C_1 \wedge C_2 \vdash \tau_1 \tau_2 : T'$. Its premises are $C_1 \vdash \tau_1 : T \rightarrow T'$ and $C_2 \vdash \tau_2 : T$. Applying the induction hypothesis yields respectively $\text{let } \Gamma \text{ in } C_1, \Gamma \vdash \tau_1 : T \rightarrow T'$ and $\text{let } \Gamma \text{ in } C_2, \Gamma \vdash \tau_2 : T$, which by Lemma 1.4.2 and HMD-APP imply $\text{let } \Gamma \text{ in } (C_1 \wedge C_2), \Gamma \vdash \tau_1 \tau_2 : T'$.

◦ *Case LET.* The rule's conclusion is $\text{let } z : \forall V[C_1].T_1 \text{ in } C_2 \vdash \text{let } z =$

\mathfrak{t}_1 in $\mathfrak{t}_2 : T_2$. Its premises are $C_1 \vdash \mathfrak{t}_1 : T_1$ **(1)** and $C_2 \vdash \mathfrak{t}_2 : T_2$ **(2)**. Let $\bar{\mathfrak{x}}$ stand for $ftv(C_1, T_1)$. We may require, *w.l.o.g.*, $\bar{\mathfrak{x}} \# ftv(\Gamma, C_2)$ **(3)**. By hypothesis, we have $fpi(\Gamma) = \emptyset$ **(4)**. We also have $fpi(\text{let } \Gamma; z : \forall \mathcal{V}[C_1]. T_1 \text{ in } C_2) = \emptyset$, which implies $fpi(\text{let } \Gamma \text{ in } C_1) = \emptyset$. Thus, the induction hypothesis applies to (1) and yields $\text{let } \Gamma \text{ in } C_1, \Gamma \vdash \mathfrak{t}_1 : T_1$ **(5)**. Now, let σ stand for $\forall \bar{\mathfrak{x}}[\text{let } \Gamma \text{ in } C_1]. T_1$ and Γ' stand for $\Gamma; z : \sigma$. We have $fpi(\Gamma') = fpi(\text{let } \Gamma' \text{ in } C_2) = \emptyset$. Thus, the induction hypothesis applies to (2) and yields $\text{let } \Gamma' \text{ in } C_2, \Gamma' \vdash \mathfrak{t}_2 : T_2$ **(6)**. Let us now weaken (5) and (6) so as to make them suitable premises for HMD-LETGEN. Applying Lemma 1.4.2 to (5) yields $(\text{let } \Gamma' \text{ in } C_2) \wedge (\text{let } \Gamma \text{ in } C_1), \Gamma \vdash \mathfrak{t}_1 : T_1$ **(7)**. Applying Lemma 1.4.2 to (6) yields $(\text{let } \Gamma' \text{ in } C_2) \wedge \exists \bar{\mathfrak{x}}. (\text{let } \Gamma \text{ in } C_1), \Gamma' \vdash \mathfrak{t}_2 : T_2$ **(8)**. Last, (3) implies $\bar{\mathfrak{x}} \# ftv(\Gamma, \text{let } \Gamma' \text{ in } C_2)$ **(9)**. Applying HMD-LETGEN to (7), (9) and (8), we obtain $(\text{let } \Gamma' \text{ in } C_2) \wedge \exists \bar{\mathfrak{x}}. (\text{let } \Gamma \text{ in } C_1), \Gamma \vdash \text{let } z = \mathfrak{t}_1 \text{ in } \mathfrak{t}_2 : T_2$ **(10)**. Now, by (4), (3), and C-LETDUP, $\text{let } \Gamma' \text{ in } C_2$ is equivalent to $\text{let } \Gamma; z : \forall \bar{\mathfrak{x}}[C_1]. T_1 \text{ in } C_2$. Using this fact, as well as (3), C-INEX, and C-INAND, we find that the constraint $(\text{let } \Gamma' \text{ in } C_2) \wedge \exists \bar{\mathfrak{x}}. (\text{let } \Gamma \text{ in } C_1)$ is equivalent to $\text{let } \Gamma \text{ in } (\text{let } z : \forall \bar{\mathfrak{x}}[C_1]. T_1 \text{ in } C_2 \wedge \exists \bar{\mathfrak{x}}. C_1)$, which by definition of the let form, is itself equivalent to $\text{let } \Gamma; z : \forall \bar{\mathfrak{x}}[C_1]. T_1 \text{ in } C_2$. Last, by definition of $\bar{\mathfrak{x}}$, this constraint is $\text{let } \Gamma; z : \forall \mathcal{V}[C_1]. T_1 \text{ in } C_2$. Thus, (10) is the goal.

◦ *Case SUB*. The rule's conclusion is $C \wedge T \leq T' \vdash \mathfrak{t} : T'$. Its premise is $C \vdash \mathfrak{t} : T$ **(1)**. Applying the induction hypothesis to (1) yields $\text{let } \Gamma \text{ in } C, \Gamma \vdash \mathfrak{t} : T$ **(2)**. By Lemma 1.4.2 and HMD-SUB, (2) implies $(\text{let } \Gamma \text{ in } C) \wedge T \leq T', \Gamma \vdash \mathfrak{t} : T'$, which by C-INAND* may be written $\text{let } \Gamma \text{ in } (C \wedge T \leq T'), \Gamma \vdash \mathfrak{t} : T'$.

◦ *Case EXISTS*. The rule's conclusion is $\exists \bar{\mathfrak{x}}. C \vdash \mathfrak{t} : T$. Its premises are $C \vdash \mathfrak{t} : T$ **(1)** and $\bar{\mathfrak{x}} \# ftv(T)$ **(2)**. We may further require, *w.l.o.g.*, $\bar{\mathfrak{x}} \# ftv(\Gamma)$ **(3)**. Applying the induction hypothesis to (1) yields $\text{let } \Gamma \text{ in } C, \Gamma \vdash \mathfrak{t} : T$ **(4)**. Applying HMD-EXISTS to (2), (3), and (4), we find $\exists \bar{\mathfrak{x}}. \text{let } \Gamma \text{ in } C, \Gamma \vdash \mathfrak{t} : T$, which, by (3) and C-INEX, may be written $\text{let } \Gamma \text{ in } \exists \bar{\mathfrak{x}}. C, \Gamma \vdash \mathfrak{t} : T$. \square

As a corollary, we find that, for closed programs, the type systems $\text{HM}(X)$ and $\text{PCB}(X)$ coincide. In particular, a program is well-typed with respect to one if and only if it is well-typed with respect to the other. This supports the view that $\text{PCB}(X)$ is an alternate formulation of $\text{HM}(X)$.

1.5.5 **THEOREM:** Assume $fpi(C) = \emptyset$ and $C \not\equiv \text{false}$. Then, $C, \emptyset \vdash \mathfrak{t} : T$ holds in $\text{HM}(X)$ if and only if $C \vdash \mathfrak{t} : T$ holds in $\text{PCB}(X)$. \square

1.6 Constraint generation

We now explain how to reduce type inference problems for $\text{PCB}(X)$ to constraint solving problems. A type inference problem consists of an expression

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

Figure 1-10: Constraint generation

t and a type T of kind \star . The problem is to determine whether t is well-typed with type T , that is, whether there exists a satisfiable constraint C such that $C \vdash t : T$ holds. This formulation of the problem may seem to require an appropriate type T to be known in advance; this is not really the case, since T may be a type variable. A constraint solving problem consists of a constraint C . The problem is to determine whether C is satisfiable. To reduce a type inference problem (t, T) to a constraint solving problem, we must produce a constraint C that is both *sufficient* and *necessary* for $C \vdash t : T$ to hold. Below, we explain how to compute such a constraint, which we write $\llbracket t : T \rrbracket$. We check that it is indeed *sufficient* by proving $\llbracket t : T \rrbracket \vdash t : T$. That is, the constraint $\llbracket t : T \rrbracket$ is specific enough to guarantee that t has type T . We say that constraint generation is *sound*. We check that it is indeed *necessary* by proving that, for every constraint C , $C \vdash t : T$ implies $C \Vdash \llbracket t : T \rrbracket$. That is, every constraint that guarantees that t has type T is at least as specific as $\llbracket t : T \rrbracket$. We say that constraint generation is *complete*. Together, these properties mean that $\llbracket t : T \rrbracket$ is the *least specific* constraint that guarantees that t has type T .

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 $C \vdash t : T$ holds, then, by the completeness property, $C \Vdash \llbracket t : T \rrbracket$ holds, so $\llbracket t : T \rrbracket$ is satisfiable. Conversely, the soundness property states that $\llbracket t : T \rrbracket \vdash t : T$ holds, so, if $\llbracket t : T \rrbracket$ is satisfiable, then there exists a satisfiable constraint C such that $C \vdash t : T$ holds. In other words, t is well-typed with type T if and only if $\llbracket t : T \rrbracket$ is satisfiable.

The existence of such a constraint is the analogue 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 $\llbracket t : T \rrbracket$ is least specific in the sense that all valid constraints entail it. Earlier, we have established a connection between constraint entailment and refinement of type substitutions, in the specific case of equality constraints interpreted over a free algebra of finite types; see Lemma 1.3.39.

The constraint $\llbracket \mathbf{t} : \mathbf{T} \rrbracket$ is defined in Figure 1-10 by induction on the structure of the expression \mathbf{t} . We refer to these defining equations as the *constraint generation rules*. The definition is quite terse. It is perhaps even simpler than the declarative specification of $\text{PCB}(X)$ given in Figure 1-9; 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* in the following sense: *type variables that appear bound in an equation's right-hand side must not appear free 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. We remark that, 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 \mathbf{T} . However, this simplification is rendered invalid by the introduction of type annotations within expressions (page 102). Please note that we are able to state a *formal* freshness requirement. This is made possible by the fact that $\llbracket \mathbf{t} : \mathbf{T} \rrbracket$ has no free type variables other than those of \mathbf{T} , which in turn depends on our explicit use of existential quantification.

Let us now review the four equations. The first one simply mirrors VAR. The second one requires \mathbf{t} to have type X_2 under the hypothesis that \mathbf{z} has type X_1 , and forms the arrow type $X_1 \rightarrow X_2$; this corresponds to ABS. Here, X_1 and X_2 must be fresh type variables, because we cannot in general guess the expected types of \mathbf{z} and \mathbf{t} . The expected type \mathbf{T} is required to be a supertype of $X_1 \rightarrow X_2$; this corresponds to SUB. We must bind the fresh type variables X_1 and X_2 , so as to guarantee that the generated constraint is unique up to α -conversion. Furthermore, we must bind them *existentially*, because we intend the constraint solver to choose some appropriate value for them. This is justified by EXISTS. The third equation uses the fresh type variable X_2 to stand for the unknown type of \mathbf{t}_2 . The subexpression \mathbf{t}_1 is expected to have type $X_2 \rightarrow \mathbf{T}$. This corresponds to APP. The fourth equation, which corresponds to LET, is most interesting. It summons a fresh type variable X and produces $\llbracket \mathbf{t}_1 : X \rrbracket$. This constraint, whose sole free type variable is X , is the *least specific* constraint that must be imposed on X so as to make it a valid type for \mathbf{t}_1 . As a result, the type scheme $\forall X[\llbracket \mathbf{t}_1 : X \rrbracket].X$, abbreviated σ in the following, is a *principal* type scheme for \mathbf{t}_1 . There remains to place $\llbracket \mathbf{t}_2 : \mathbf{T} \rrbracket$ inside the context $\text{let } \mathbf{z} : \sigma \text{ in } \square$. Indeed, when placed inside this context, an instantiation constraint of the form $\mathbf{z} \preceq \mathbf{T}'$ acquires the meaning

$\sigma \preceq T'$, which by definition of σ and by Lemma 1.6.4 (see below) is equivalent to $\llbracket \mathfrak{t}_1 : T' \rrbracket$. Thus, the constraint produced by the fourth equation simulates a textual expansion of the `let` construct, whereby every occurrence of z would be replaced with \mathfrak{t}_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; duplication may take place during constraint solving only.

- 1.6.1 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 $\llbracket \mathfrak{t} : T \rrbracket$ is linear in the sum of the sizes of \mathfrak{t} and T . \square

We now establish several properties of constraint generation. We begin with soundness, whose proof is straightforward.

- 1.6.2 THEOREM [SOUNDNESS]: $\llbracket \mathfrak{t} : T \rrbracket \vdash \mathfrak{t} : T$. \square

Proof: By induction on the structure of \mathfrak{t} .

- *Case* x . The goal $x \preceq T \vdash x : T$ follows from VAR.
- *Case* $\lambda z.t$. By the induction hypothesis, we have $\llbracket \mathfrak{t} : X_2 \rrbracket \vdash \mathfrak{t} : X_2$. By ABS, this implies $\text{let } z : X_1 \text{ in } \llbracket \mathfrak{t} : X_2 \rrbracket \vdash \lambda z.t : X_1 \rightarrow X_2$. By SUB, this implies $\text{let } z : X_1 \text{ in } \llbracket \mathfrak{t} : X_2 \rrbracket \wedge X_1 \rightarrow X_2 \leq T \vdash \lambda z.t : T$. Lastly, because $X_1 X_2 \# \text{ftv}(T)$ holds, EXISTS applies and yields $\llbracket \lambda z.t : T \rrbracket \vdash \lambda z.t : T$.
- *Case* $\mathfrak{t}_1 \mathfrak{t}_2$. By the induction hypothesis, we have $\llbracket \mathfrak{t}_1 : X_2 \rightarrow T \rrbracket \vdash \mathfrak{t}_1 : X_2 \rightarrow T$ and $\llbracket \mathfrak{t}_2 : X_2 \rrbracket \vdash \mathfrak{t}_2 : X_2$. By APP, this implies $\llbracket \mathfrak{t}_1 : X_2 \rightarrow T \rrbracket \wedge \llbracket \mathfrak{t}_2 : X_2 \rrbracket \vdash \mathfrak{t}_1 \mathfrak{t}_2 : T$. Because $X_2 \notin \text{ftv}(T)$ holds, EXISTS applies and yields $\llbracket \mathfrak{t}_1 \mathfrak{t}_2 : T \rrbracket \vdash \mathfrak{t}_1 \mathfrak{t}_2 : T$.
- *Case* $\text{let } z = \mathfrak{t}_1 \text{ in } \mathfrak{t}_2$. By the induction hypothesis, we have $\llbracket \mathfrak{t}_1 : X \rrbracket \vdash \mathfrak{t}_1 : X$ and $\llbracket \mathfrak{t}_2 : T \rrbracket \vdash \mathfrak{t}_2 : T$. By LET, these imply $\text{let } z : \forall \mathcal{V} \llbracket \mathfrak{t}_1 : X \rrbracket.X \text{ in } \llbracket \mathfrak{t}_2 : T \rrbracket \vdash \text{let } z = \mathfrak{t}_1 \text{ in } \mathfrak{t}_2 : T$. Because $\text{ftv}(\llbracket \mathfrak{t}_1 : X \rrbracket)$ is X , the universal quantification on \mathcal{V} really bears on X alone. We have proved $\llbracket \text{let } z = \mathfrak{t}_1 \text{ in } \mathfrak{t}_2 : T \rrbracket \vdash \text{let } z = \mathfrak{t}_1 \text{ in } \mathfrak{t}_2 : T$. \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 $\llbracket \mathfrak{t} : T \rrbracket$ is *covariant* with respect to T . Roughly speaking, this means that enough subtyping constraints are generated to achieve completeness with respect to SUB.

- 1.6.3 LEMMA: $\llbracket \mathfrak{t} : T \rrbracket \wedge T \leq T'$ entails $\llbracket \mathfrak{t} : T' \rrbracket$. \square
- 1.6.4 LEMMA: $X \notin \text{ftv}(T)$ implies $\exists X. (\llbracket \mathfrak{t} : X \rrbracket \wedge X \leq T) \equiv \llbracket \mathfrak{t} : T \rrbracket$. \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. Then, fresh type variables need not be generated; one may directly use the arrow's domain and codomain instead.

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

We conclude with the completeness property.

1.6.6 THEOREM [COMPLETENESS]: if $C \vdash \mathbf{t} : T$, then $C \Vdash \llbracket \mathbf{t} : T \rrbracket$. \square

Proof: By induction on the derivation of $C \vdash \mathbf{t} : T$.

◦ *Case VAR.* The rule's conclusion is $C \vdash x : T$. Its premise is $C \Vdash x \preceq T$, which is also the goal.

◦ *Case ABS.* The rule's conclusion is $\text{let } z : T \text{ in } C \vdash \lambda z. \mathbf{t} : T \rightarrow T'$. Its premise is $C \vdash \mathbf{t} : T'$. By the induction hypothesis, we have $C \Vdash \llbracket \mathbf{t} : T' \rrbracket$. By congruence of entailment, this implies $\text{let } z : T \text{ in } C \Vdash \text{let } z : T \text{ in } \llbracket \mathbf{t} : T' \rrbracket$, which, by Lemma 1.6.5, may be written $\text{let } z : T \text{ in } C \Vdash \llbracket \lambda z. \mathbf{t} : T \rightarrow T' \rrbracket$.

◦ *Case APP.* The rule's conclusion is $C_1 \wedge C_2 \vdash \mathbf{t}_1 \mathbf{t}_2 : T'$. Its premises are $C_1 \vdash \mathbf{t}_1 : T \rightarrow T'$ and $C_2 \vdash \mathbf{t}_2 : T$. By the induction hypothesis, we have $C_1 \Vdash \llbracket \mathbf{t}_1 : T \rightarrow T' \rrbracket$ and $C_2 \Vdash \llbracket \mathbf{t}_2 : T \rrbracket$. Thus, $C_1 \wedge C_2$ entails $\llbracket \mathbf{t}_1 : T \rightarrow T' \rrbracket \wedge \llbracket \mathbf{t}_2 : T \rrbracket$, which, by C-NAMEEQ, may be written $\exists X_2. (X_2 = T \wedge \llbracket \mathbf{t}_1 : X_2 \rightarrow T' \rrbracket \wedge \llbracket \mathbf{t}_2 : X_2 \rrbracket)$, where $X_2 \notin \text{ftv}(T, T')$. Forgetting about the equation $X_2 = T$, we find that $C_1 \wedge C_2$ entails $\exists X_2. (\llbracket \mathbf{t}_1 : X_2 \rightarrow T' \rrbracket \wedge \llbracket \mathbf{t}_2 : X_2 \rrbracket)$, which is precisely $\llbracket \mathbf{t}_1 \mathbf{t}_2 : T' \rrbracket$.

◦ *Case LET.* The rule's conclusion is $\text{let } z : \forall \mathcal{V}[C_1]. T_1 \text{ in } C_2 \vdash \text{let } z = \mathbf{t}_1 \text{ in } \mathbf{t}_2 : T_2$. Its premises are $C_1 \vdash \mathbf{t}_1 : T_1$ and $C_2 \vdash \mathbf{t}_2 : T_2$. By the induction hypothesis, we have $C_1 \Vdash \llbracket \mathbf{t}_1 : T_1 \rrbracket$ and $C_2 \Vdash \llbracket \mathbf{t}_2 : T_2 \rrbracket$, which implies $\text{let } z : \forall \mathcal{V}[C_1]. T_1 \text{ in } C_2 \Vdash \text{let } z : \forall \mathcal{V}[\llbracket \mathbf{t}_1 : T_1 \rrbracket]. T_1 \text{ in } \llbracket \mathbf{t}_2 : T_2 \rrbracket$ (1).

Now, let us establish $\text{true} \Vdash \forall X[\llbracket \mathbf{t}_1 : X \rrbracket]. X \preceq \forall \mathcal{V}[\llbracket \mathbf{t}_1 : T_1 \rrbracket]. T_1$ (2). By definition, this requires proving $\exists \bar{X}_1. (\llbracket \mathbf{t}_1 : T_1 \rrbracket \wedge T_1 \leq Z) \Vdash \exists X. (\llbracket \mathbf{t}_1 : X \rrbracket \wedge X \leq Z)$ (3), where $\bar{X}_1 = \text{ftv}(T_1)$ and $Z \notin \bar{X}_1$ (4). By Lemma 1.6.3, (4), and C-EX*, the left-hand side of (3) entails $\llbracket \mathbf{t}_1 : Z \rrbracket$. By (4) and Lemma 1.6.4, the right-hand side of (3) is $\llbracket \mathbf{t}_1 : Z \rrbracket$. Thus, (3) holds, and so does (2).

By (2) and Lemma 1.3.22, we have $\text{let } z : \forall \mathcal{V}[\llbracket \mathbf{t}_1 : T_1 \rrbracket]. T_1 \text{ in } \llbracket \mathbf{t}_2 : T_2 \rrbracket \Vdash \text{let } z : \forall X[\llbracket \mathbf{t}_1 : X \rrbracket]. X \text{ in } \llbracket \mathbf{t}_2 : T_2 \rrbracket$ (5). By transitivity of entailment, (1) and (5) yield $\text{let } z : \forall \mathcal{V}[C_1]. T_1 \text{ in } C_2 \Vdash \llbracket \text{let } z = \mathbf{t}_1 \text{ in } \mathbf{t}_2 : T_2 \rrbracket$.

◦ *Case SUB.* The rule's conclusion is $C \wedge T \leq T' \vdash \mathbf{t} : T'$. Its premise is $C \vdash \mathbf{t} : T$. By the induction hypothesis, we have $C \Vdash \llbracket \mathbf{t} : T \rrbracket$, which implies $C \wedge T \leq T' \Vdash \llbracket \mathbf{t} : T \rrbracket \wedge T \leq T'$. By lemma 1.6.3 and by transitivity of entailment, we obtain $C \wedge T \leq T' \Vdash \llbracket \mathbf{t} : T' \rrbracket$.

◦ *Case EXISTS.* The rule's conclusion is $\exists \bar{X}. C \vdash \mathbf{t} : T$. Its premises are $C \vdash \mathbf{t} : T$ and $\bar{X} \# \text{ftv}(T)$ (1). By the induction hypothesis, we have $C \Vdash \llbracket \mathbf{t} : T \rrbracket$.

By congruence of entailment, this implies $\exists \bar{x}. C \Vdash \exists \bar{x}. \llbracket \mathfrak{t} : T \rrbracket$ **(2)**. Furthermore, **(1)** implies $\bar{x} \# \text{ftv}(\llbracket \mathfrak{t} : T \rrbracket)$ **(3)**. By **(3)** and C-EX*, **(2)** may be written $\exists \bar{x}. C \Vdash \llbracket \mathfrak{t} : T \rrbracket$. \square

1.7 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 1.4.7, 1.5.4, and 1.6.6 imply that type soundness also holds when well-typedness is defined with respect to the typing judgements of DM, HM(X), or PCB(X). We establish type soundness by following Wright and Felleisen's so-called *syntactic approach* (1994b). The approach consists in 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*. We immediately give two versions of this lemma: the former is stated in terms of PCB(X) judgements, while the latter is stated in terms of the constraint generation rules.

1.7.1 LEMMA [SUBSTITUTION]: $C \vdash \mathfrak{t} : T$ and $C_0 \vdash \mathfrak{t}_0 : T_0$ imply $\text{let } z_0 : \forall \bar{x}_0[C_0]. T_0 \text{ in } C \vdash [z_0 \mapsto \mathfrak{t}_0]\mathfrak{t} : T$. \square

Proof: The proof is by structural induction on the derivation of $C \vdash \mathfrak{t} : T$. In each proof case, we adopt the notations of Figure 1-9. We write σ_0 for $\forall \bar{x}_0[C_0]. T_0$. We refer to the hypothesis $C_0 \vdash \mathfrak{t}_0 : T_0$ as **(1)**. We assume, *w.l.o.g.*, $\bar{x}_0 \# \text{ftv}(C, T)$ **(2)** and $z_0 \notin \text{fpi}(\sigma_0)$ **(3)**.

◦ *Case VAR.* The rule's conclusion is $C \vdash x : T$ **(4)**. Its premise is $C \Vdash x \preceq T$ **(5)**. Two subcases arise.

Subcase x is z_0 . Applying SUB to **(1)** yields $C_0 \wedge T_0 \leq T \vdash \mathfrak{t}_0 : T$. By **(2)** and EXISTS, this implies $\exists \bar{x}_0. (C_0 \wedge T_0 \leq T) \vdash \mathfrak{t}_0 : T$ **(6)**. Furthermore, by **(2)** again, the constraint $\exists \bar{x}_0. (C_0 \wedge T_0 \leq T)$ is $\sigma_0 \preceq T$, which is equivalent to $\text{let } z_0 : \sigma_0 \text{ in } z_0 \preceq T$. As a result, **(6)** may be written $\text{let } z_0 : \sigma_0 \text{ in } x \preceq T \vdash [z_0 \mapsto \mathfrak{t}_0]x : T$ **(7)**.

Subcase x isn't z_0 . Then, $[z_0 \mapsto \mathfrak{t}_0]x$ is x . Thus, VAR yields $\exists \sigma_0 \wedge x \preceq T \vdash [z_0 \mapsto \mathfrak{t}_0]x : T$. By C-IN*, this may be read $\text{let } z_0 : \sigma_0 \text{ in } x \preceq T \vdash [z_0 \mapsto \mathfrak{t}_0]x : T$, that is, again **(7)**.

In either subcase, by (5), by congruence of entailment, and by Lemma 1.5.2, (7) implies $\text{let } z_0 : \sigma_0 \text{ in } C \vdash [z_0 \mapsto t_0]t : T$.

◦ *Case ABS*. The rule's conclusion is $\text{let } z : T \text{ in } C \vdash \lambda z. t : T \rightarrow T'$. Its premise is $C \vdash t : T'$ (8). We may assume, *w.l.o.g.*, that z is distinct from z_0 and does not occur free within t_0 or σ_0 (9). Applying the induction hypothesis to (8) yields $\text{let } z_0 : \sigma_0 \text{ in } C \vdash [z_0 \mapsto t_0]t : T'$, which, by ABS, implies $\text{let } z : T \text{ in } (\text{let } z_0 : \sigma_0 \text{ in } C) \vdash \lambda z. [z_0 \mapsto t_0]t : T \rightarrow T'$. By (9) and C-LETLET, this may be written $\text{let } z_0 : \sigma_0 \text{ in } (\text{let } z : T \text{ in } C) \vdash [z_0 \mapsto t_0](\lambda z. t) : T \rightarrow T'$.

◦ *Case APP*. By the induction hypothesis, by APP, and by C-INAND.

◦ *Case LET*. The rule's conclusion is $\text{let } z : \forall \bar{X}_1[C_1]. T_1 \text{ in } C_2 \vdash \text{let } z = t_1 \text{ in } t_2 : T_2$, where $\bar{X}_1 = \text{ftv}(C_1, T_1)$. Its premises are $C_1 \vdash t_1 : T_1$ (10) and $C_2 \vdash t_2 : T_2$ (11). We may assume, *w.l.o.g.*, that z is distinct from z_0 and does not occur free within t_0 or σ_0 (12). We may also assume, *w.l.o.g.*, $\bar{X}_1 \# \text{ftv}(\sigma_0)$ (13). Applying the induction hypothesis to (10) and (11) respectively yields $\text{let } z_0 : \sigma_0 \text{ in } C_1 \vdash [z_0 \mapsto t_0]t_1 : T_1$ (14) and $\text{let } z_0 : \sigma_0 \text{ in } C_2 \vdash [z_0 \mapsto t_0]t_2 : T_2$ (15). Applying LET to (14) and (15) produces $\text{let } z : \forall \mathcal{V}[\text{let } z_0 : \sigma_0 \text{ in } C_1]. T_1 \text{ in } \text{let } z_0 : \sigma_0 \text{ in } C_2 \vdash [z_0 \mapsto t_0](\text{let } z = t_1 \text{ in } t_2) : T_2$ (16). Now, we have

$$\begin{aligned} & \text{let } z_0 : \sigma_0; z : \forall \bar{X}_1[C_1]. T_1 \text{ in } C_2 \\ \equiv & \text{let } z_0 : \sigma_0; z : \forall \bar{X}_1[\text{let } z_0 : \sigma_0 \text{ in } C_1]. T_1 \text{ in } C_2 & (17) \\ \equiv & \text{let } z : \forall \bar{X}_1[\text{let } z_0 : \sigma_0 \text{ in } C_1]. T_1; z_0 : \sigma_0 \text{ in } C_2 & (18) \\ \Vdash & \text{let } z : \forall \mathcal{V}[\text{let } z_0 : \sigma_0 \text{ in } C_1]. T_1; z_0 : \sigma_0 \text{ in } C_2 & (19) \end{aligned}$$

where (17) follows from (13), (3), and C-LETDUP; (18) follows from (12) and C-LETLET; and (19) is by Lemma 1.3.25. Thus, applying Lemma 1.5.2 to (16) yields $\text{let } z_0 : \sigma_0; z : \forall \bar{X}_1[C_1]. T_1 \text{ in } C_2 \vdash [z_0 \mapsto t_0](\text{let } z = t_1 \text{ in } t_2) : T_2$.

◦ *Case SUB*. By the induction hypothesis, by SUB, and by C-INAND*.

◦ *Case EXISTS*. The rule's conclusion is $\exists \bar{X}. C \vdash t : T$. Its premises are $C \vdash t : T$ (20) and $\bar{X} \# \text{ftv}(T)$ (21). We may assume, *w.l.o.g.*, $\bar{X} \# \text{ftv}(\sigma_0)$ (22). Applying the induction hypothesis to (20) yields $\text{let } z_0 : \sigma_0 \text{ in } C \vdash [z_0 \mapsto t_0]t : T$, which, by (21) and EXISTS, implies $\exists \bar{X}. \text{let } z_0 : \sigma_0 \text{ in } C \vdash [z_0 \mapsto t_0]t : T$ (23). By (22) and C-INEX, (23) is $\text{let } z_0 : \sigma_0 \text{ in } \exists \bar{X}. C \vdash [z_0 \mapsto t_0]t : T$. \square

1.7.2 LEMMA: $\text{let } z : \forall \bar{X}[[t_2 : T_2]]. T_2 \text{ in } [[t_1 : T_1]]$ entails $[[z \mapsto t_2]t_1 : T_1]$. \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 made 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.

- 1.7.3 DEFINITION: Let Γ_0 be an environment whose domain is the set of constants \mathcal{Q} . We require $ftv(\Gamma_0) = \emptyset$, $fpi(\Gamma_0) = \emptyset$, and $\exists\Gamma_0 \equiv \text{true}$. We refer to Γ_0 as the *initial* typing environment. \square
- 1.7.4 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 `ref` M for the environment that maps m to `ref` $M(m)$ when m is in the domain of M . Assuming $dom(\mu)$ and $dom(M)$ 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 $dom(\mu)$. Under the same assumption, the constraint $\llbracket \mathfrak{t}/\mu : T/M \rrbracket$ is defined as $\llbracket \mathfrak{t} : T \rrbracket \wedge \llbracket \mu : M \rrbracket$. A configuration \mathfrak{t}/μ is *well-typed* if and only if there exist a type T and a store type M such that $dom(\mu) = dom(M)$ and the constraint `let` Γ_0 ; `ref` M in $\llbracket \mathfrak{t}/\mu : T/M \rrbracket$ is satisfiable. \square

The type `ref` T is the type of references (that is, memory locations) that store data of type T . It must be *invariant* in its parameter, reflecting the fact that references may be *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 $\llbracket \mathfrak{t}/\mu : T/M \rrbracket$ within the context `let` `ref` M in $\llbracket \cdot \rrbracket$, we give meaning to free occurrences of memory locations within $\llbracket \mathfrak{t}/\mu : T/M \rrbracket$, 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. Since no satisfiable constraint may have a free program identifier (Lemma 1.3.31), every well-typed configuration must be closed. The context `let` Γ_0 in $\llbracket \cdot \rrbracket$ gives meaning to occurrences of constants within $\llbracket \mathfrak{t}/\mu : T/M \rrbracket$.

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 \mathfrak{t} has type T under a certain constraint, then its reduct \mathfrak{t}' has type T under the same constraint. In terms of generated constraints, this statement becomes: `let` Γ_0 in $\llbracket \mathfrak{t} : T \rrbracket$ entails `let` Γ_0 in $\llbracket \mathfrak{t}' : T \rrbracket$.

Let us now return to the general case, where a store is present. Then, the statement of well-typedness for a configuration \mathfrak{t}/μ involves a store type M whose domain is that of μ . So, the statement of well-typedness for its reduct \mathfrak{t}'/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 $\text{dom}(M') \setminus \text{dom}(M)$ 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:

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

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

The relation \sqsubseteq is intended to express a connection between a configuration and its reduct. Thus, subject reduction may be stated as: $(\longrightarrow) \subseteq (\sqsubseteq)$, 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.

- 1.7.6 DEFINITION: We require (i) $(\xrightarrow{\delta}) \subseteq (\sqsubseteq)$; and (ii) 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. \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:

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

This expression reduces to `true $\hat{+}$ $\hat{1}$` , so it must not be well-typed. Yet, if natural type schemes are assigned to `ref`, `!`, and `:=` (see Example 1.9.5), 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).

1.7.7 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, as stated above, 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 (Section 1.9). 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 other improvements to the value restriction exist; see *e.g.* Exercise (Garrigue, 2002). Another frequent limitation of the value restriction are constructor functions, that is, functions that only build values, which are treated as ordinary functions and not as constructors, and their applications are not considered to be

values. 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 $c \ v$ bound to f is a constructor function (of arity 1), but $f \ w$ is treated as a regular application and cannot be generalized. Technically, the effect of the (strict) value restriction is summarized by the following result.

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

We are done with definitions and requirements. We now come to the bulk of the type soundness proof.

1.7.9 THEOREM [SUBJECT REDUCTION]: $(\longrightarrow) \subseteq (\sqsubseteq)$. \square

Proof: Because \longrightarrow and \longrightarrow are the smallest relations that satisfy the rules of Figure 1-2, it suffices to prove that \sqsubseteq satisfies these rules as well. We remark that if, for every type T , $\llbracket t : T \rrbracket \Vdash \llbracket t' : T \rrbracket$ holds, then $t/\mu \sqsubseteq t'/\mu$ holds. (Take $\bar{Y} = \emptyset$ and $M' = M$ and use the fact that entailment is a congruence to check that the conditions of Definition 1.7.5 are met.) We make use of this fact in cases R-BETA and R-LET below.

◦ *Case R-BETA.* We have

$$\begin{aligned} & \llbracket (\lambda z. t) \ v : T \rrbracket \\ \equiv & \exists X. (\llbracket \lambda z. t : X \rightarrow T \rrbracket \wedge \llbracket v : X \rrbracket) & \text{(1)} \\ \equiv & \exists X. (\text{let } z : X \ \text{in } \llbracket t : T \rrbracket \wedge \llbracket v : X \rrbracket) & \text{(2)} \\ \equiv & \exists X. \text{let } z : \forall \emptyset [\llbracket v : X \rrbracket]. X \ \text{in } \llbracket t : T \rrbracket & \text{(3)} \\ \Vdash & \llbracket [z \mapsto v] t : T \rrbracket & \text{(4)} \end{aligned}$$

where (1) is by definition of constraint generation; (2) is by Lemma 1.6.5; (3) is by C-LETAND; (4) is by Lemma 1.7.2 and C-EX*.

◦ *Case R-LET.* We have

$$\begin{aligned} & \llbracket \text{let } z = v \ \text{in } t : T \rrbracket \\ = & \text{let } z : \forall X [\llbracket v : X \rrbracket]. X \ \text{in } \llbracket t : T \rrbracket & \text{(1)} \\ \Vdash & \llbracket [z \mapsto v] t : T \rrbracket & \text{(2)} \end{aligned}$$

where (1) is by definition of constraint generation and (2) is by Lemma 1.7.2.

◦ *Case R-DELTA.* This case is exactly requirement (i) in Definition 1.7.6.

◦ *Case R-EXTEND.* Our hypotheses are $t/\mu \sqsubseteq t'/\mu'$ **(1)** and $\text{dom}(\mu'') \# \text{dom}(\mu')$ **(2)** and $\text{range}(\mu'') \# \text{dom}(\mu' \setminus \mu)$ **(3)**. Because $\text{dom}(\mu)$ must be a subset of $\text{dom}(\mu')$, it is also disjoint with $\text{dom}(\mu'')$. Our goal is $t/\mu\mu'' \sqsubseteq t'/\mu'\mu''$ **(4)**. Thus, let us introduce a type T and a store type of domain

$dom(\mu\mu'')$, or (equivalently) two store types M and M'' whose domains are respectively $dom(\mu)$ and $dom(\mu'')$. By (1), there exist type variables \bar{Y} and a store type M' such that $\bar{Y} \# ftv(\mathbb{T}, M)$ (5) and $ftv(M') \subseteq \bar{Y} \cup ftv(M)$ and $dom(M') = dom(\mu')$ and M' extends M (6) and let $\Gamma_0; \text{ref } M$ in $\llbracket \mathfrak{t}/\mu : \mathbb{T}/M \rrbracket \Vdash \exists \bar{Y}. \text{let } \Gamma_0; \text{ref } M' \text{ in } \llbracket \mathfrak{t}'/\mu' : \mathbb{T}/M' \rrbracket$. We may further require, *w.l.o.g.*, $\bar{Y} \# ftv(M'')$ (7). Let us now add the conjunct $\text{let } \Gamma_0; \text{ref } M \text{ in } \llbracket \mu'' : M'' \rrbracket$ to each side of this entailment assertion. On the left-hand side, by C-INAND and by Definition 1.7.4, we obtain $\text{let } \Gamma_0; \text{ref } M \text{ in } \llbracket \mathfrak{t}/\mu\mu'' : \mathbb{T}/MM'' \rrbracket$ (8). On the right-hand side, by (5), (7), C-EXAND, and C-INAND, we obtain $\exists \bar{Y}. \text{let } \Gamma_0 \text{ in } (\text{let } \text{ref } M' \text{ in } \llbracket \mathfrak{t}'/\mu' : \mathbb{T}/M' \rrbracket \wedge \text{let } \text{ref } M \text{ in } \llbracket \mu'' : M'' \rrbracket)$ (9). Now, recall that M' extends M (6) and, furthermore, (3) implies $fpi(\llbracket \mu'' : M'' \rrbracket) \# dpi(M' \setminus M)$ (10). By (10), C-INAND*, and C-INAND, (9) is equivalent to $\exists \bar{Y}. \text{let } \Gamma_0; \text{ref } M' \text{ in } (\llbracket \mathfrak{t}'/\mu' : \mathbb{T}/M' \rrbracket \wedge \llbracket \mu'' : M'' \rrbracket)$, that is, $\exists \bar{Y}. \text{let } \Gamma_0; \text{ref } M' \text{ in } \llbracket \mathfrak{t}'/\mu'\mu'' : \mathbb{T}/M'M'' \rrbracket$ (11). Thus, we have established that (8) entails (11). Let us now place this entailment assertion within the constraint context $\text{let } \text{ref } M'' \text{ in } \square$. On the left-hand side, because $fpi(\Gamma_0, M, M'') = \emptyset$ and $dpi(M'') \cap dpi(\Gamma_0, M) \subseteq dom(\mu'') \cap (Q \cup dom(\mu)) = \emptyset$, C-LETLET applies, yielding $\text{let } \Gamma_0; \text{ref } MM'' \text{ in } \llbracket \mathfrak{t}/\mu\mu'' : \mathbb{T}/MM'' \rrbracket$ (12). On the right-hand side, by (7), C-INEX, and by analogous reasoning, we obtain $\exists \bar{Y}. \text{let } \Gamma_0; \text{ref } M'M'' \text{ in } \llbracket \mathfrak{t}'/\mu'\mu'' : \mathbb{T}/M'M'' \rrbracket$ (13). Thus, (12) entails (13). Given (5), (7), given $ftv(M'M'') \subseteq \bar{Y} \cup ftv(MM'')$, and given that $M'M''$ extends MM'' , this establishes the goal (4).

◦ *Case R-CONTEXT.* The hypothesis is $\mathfrak{t}/\mu \sqsubseteq \mathfrak{t}'/\mu'$. The goal is $\mathcal{E}[\mathfrak{t}]/\mu \sqsubseteq \mathcal{E}[\mathfrak{t}']/\mu'$. Because \longrightarrow relates closed configurations only, we may assume that the configuration $\mathcal{E}[\mathfrak{t}]/\mu$ is closed, so the memory locations that appear free within \mathcal{E} are members of $dom(\mu)$. Let us now reason by induction on the structure of \mathcal{E} .

Subcase $\mathcal{E} = \square$. The hypothesis and the goal coincide.

Subcase $\mathcal{E} = \mathcal{E}_1 \mathfrak{t}_1$. The induction hypothesis is $\mathcal{E}_1[\mathfrak{t}]/\mu \sqsubseteq \mathcal{E}_1[\mathfrak{t}']/\mu'$ (1). Let us introduce a type \mathbb{T} and a store type M such that $dom(M) = dom(\mu)$. Consider the constraint $\text{let } \Gamma_0; \text{ref } M \text{ in } \llbracket \mathcal{E}[\mathfrak{t}]/\mu : \mathbb{T}/M \rrbracket$ (2). By definition of constraint generation, C-EXAND, C-INEX, and C-INAND, it is equivalent to

$$\exists \mathbb{X}. (\text{let } \Gamma_0; \text{ref } M \text{ in } \llbracket \mathcal{E}_1[\mathfrak{t}]/\mu : \mathbb{X} \rightarrow \mathbb{T}/M \rrbracket \wedge \text{let } \Gamma_0; \text{ref } M \text{ in } \llbracket \mathfrak{t}_1 : \mathbb{X} \rrbracket) \quad (3)$$

where $\mathbb{X} \not\subseteq ftv(\mathbb{T}, M)$ (4). By (1), there exist type variables \bar{Y} and a store type M' such that $\bar{Y} \# ftv(\mathbb{X}, \mathbb{T}, M)$ (5) and $ftv(M') \subseteq \bar{Y} \cup ftv(M)$ (6) and $dom(M') = dom(\mu')$ and M' extends M and (3) entails

$$\exists \mathbb{X}. (\exists \bar{Y}. \text{let } \Gamma_0; \text{ref } M' \text{ in } \llbracket \mathcal{E}_1[\mathfrak{t}']/\mu' : \mathbb{X} \rightarrow \mathbb{T}/M' \rrbracket \wedge \text{let } \Gamma_0; \text{ref } M \text{ in } \llbracket \mathfrak{t}_1 : \mathbb{X} \rrbracket) \quad (7).$$

We pointed out earlier that the memory locations that appear free in \mathfrak{t}_1 are members of $dom(M)$, which implies $\text{let } \text{ref } M \text{ in } \llbracket \mathfrak{t}_1 : \mathbb{X} \rrbracket \equiv \text{let } \text{ref } M' \text{ in } \llbracket \mathfrak{t}_1 :$

$\mathbf{x}]$ (8). By (5), C-EXAND, (8), C-INAND, and by definition of constraint generation, we find that (7) is equivalent to

$$\exists \bar{\mathbf{x}} \bar{\mathbf{y}}. \text{let } \Gamma_0; \text{ref } M' \text{ in } (\llbracket \mathcal{E}_1[\mathbf{t}'] : \mathbf{X} \rightarrow \mathbf{T} \rrbracket \wedge \llbracket \mathbf{t}_1 : \mathbf{X} \rrbracket \wedge \llbracket \mu' : M' \rrbracket) \quad (9).$$

(4), (5) and (6) imply $\mathbf{x} \notin \text{ftv}(M')$. Thus, by C-INEX and C-EXAND, (9) may be written

$$\exists \bar{\mathbf{y}}. \text{let } \Gamma_0; \text{ref } M' \text{ in } (\exists \mathbf{x}. (\llbracket \mathcal{E}_1[\mathbf{t}'] : \mathbf{X} \rightarrow \mathbf{T} \rrbracket \wedge \llbracket \mathbf{t}_1 : \mathbf{X} \rrbracket) \wedge \llbracket \mu' : M' \rrbracket),$$

which, by definition of constraint generation, is

$$\exists \bar{\mathbf{y}}. \text{let } \Gamma_0; \text{ref } M' \text{ in } \llbracket \mathcal{E}[\mathbf{t}']/\mu' : \mathbf{T}/M' \rrbracket \quad (10).$$

Thus, we have proved that (2) entails (10). By Definition 1.7.5, this establishes $\mathcal{E}[\mathbf{t}]/\mu \sqsubseteq \mathcal{E}[\mathbf{t}']/\mu'$.

Subcase $\mathcal{E} = \mathbf{v} \mathcal{E}_1$. Analogous to the previous subcase.

Subcase $\mathcal{E} = \text{let } \mathbf{z} = \mathcal{E}_1 \text{ in } \mathbf{t}_1$. The induction hypothesis is $\mathcal{E}_1[\mathbf{t}]/\mu \sqsubseteq \mathcal{E}_1[\mathbf{t}']/\mu'$ (1). This subcase is particularly interesting, because it is where **let**-polymorphism and side effects interact. In the previous two subcases, we relied on the fact that the $\exists \bar{\mathbf{y}}$ quantifier, which hides the types of the memory cells created by the reduction step, *commutes* with the connectives \exists and \wedge introduced by application contexts. However, it does not in general (left-)commute with the **let** connective (Example 1.3.28). Fortunately, under the value restriction, this subcase *never arises* (Lemma 1.7.8). By Definition 1.7.7, this subcase may arise only if all constants have pure semantics, which implies $\mu = \mu' = \emptyset$. Then, we have

$$\begin{aligned} & \text{let } \Gamma_0 \text{ in } \llbracket \mathcal{E}[\mathbf{t}] : \mathbf{T} \rrbracket \\ \equiv & \text{let } \Gamma_0; \mathbf{z} : \forall \mathbf{X} [\llbracket \mathcal{E}_1[\mathbf{t}] : \mathbf{X} \rrbracket]. \mathbf{X} \text{ in } \llbracket \mathbf{t}_1 : \mathbf{T} \rrbracket & (2) \\ \equiv & \text{let } \Gamma_0; \mathbf{z} : \forall \mathbf{X} [\text{let } \Gamma_0 \text{ in } \llbracket \mathcal{E}_1[\mathbf{t}] : \mathbf{X} \rrbracket]. \mathbf{X} \text{ in } \llbracket \mathbf{t}_1 : \mathbf{T} \rrbracket & (3) \\ \Vdash & \text{let } \Gamma_0; \mathbf{z} : \forall \mathbf{X} [\text{let } \Gamma_0 \text{ in } \llbracket \mathcal{E}_1[\mathbf{t}'] : \mathbf{X} \rrbracket]. \mathbf{X} \text{ in } \llbracket \mathbf{t}_1 : \mathbf{T} \rrbracket & (4) \\ \equiv & \text{let } \Gamma_0 \text{ in } \llbracket \mathcal{E}[\mathbf{t}'] : \mathbf{T} \rrbracket & (5) \end{aligned}$$

where (2) is by definition of constraint generation; (3) follows from $\text{ftv}(\Gamma_0) = \text{fpi}(\Gamma_0) = \emptyset$ and C-LETDUP; (4) follows from (1), specialized to the case of a pure semantics; and (5) is obtained by performing these steps in reverse. \square

1.7.10 EXERCISE [RECOMMENDED, ★★]: Try to carry out the last subcase of the above proof in the case of an impure semantics and in the absence of the value restriction. Find out why it fails. Show that it succeeds if $\bar{\mathbf{y}}$ is assumed to be empty. Use this fact to prove that generalization is still safe when restricted to *nonexpansive* expressions, provided (i) evaluating a nonexpansive expression

cannot cause new memory cells to be allocated, (ii) nonexpansive expressions are stable by substitution of values for variables, and (iii) nonexpansive expressions are preserved by reduction. \square

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

1.7.11 LEMMA: Let $\mathfrak{t}/\mu \longrightarrow \mathfrak{t}'/\mu'$. If \mathfrak{t}/μ is well-typed, then so is \mathfrak{t}'/μ' . \square

Proof: Assume $\mathfrak{t}/\mu \longrightarrow \mathfrak{t}'/\mu'$ (1) and \mathfrak{t}/μ is well-typed (2). By (2) and Definition 1.7.4, there exist a type \mathbb{T} and a store type M such that $\text{dom}(\mu) = \text{dom}(M)$ and the constraint $\text{let } \Gamma_0; \text{ref } M \text{ in } \llbracket \mathfrak{t}/\mu : \mathbb{T}/M \rrbracket$ (3) is satisfiable. By Theorem 1.7.9 and Definition 1.7.5, (1) implies that there exist a set of type variables \bar{Y} and a store type M' such that $\text{dom}(M') = \text{dom}(\mu')$ (4) and the constraint (3) entails $\exists \bar{Y}. \text{let } \Gamma_0; \text{ref } M' \text{ in } \llbracket \mathfrak{t}'/\mu' : \mathbb{T}/M' \rrbracket$ (5). Because (3) is satisfiable, so is (5), which implies that $\text{let } \Gamma_0; \text{ref } M' \text{ in } \llbracket \mathfrak{t}'/\mu' : \mathbb{T}/M' \rrbracket$ is satisfiable (6). By (4) and (6) and Definition 1.7.4, \mathfrak{t}'/μ' is well-typed. \square

Let us now establish the progress property.

1.7.12 LEMMA: If $\mathfrak{t}_1 \ \mathfrak{t}_2$ is well-typed, then \mathfrak{t}_1/μ and \mathfrak{t}_2/μ are well-typed. If $\text{let } \mathfrak{z} = \mathfrak{t}_1 \text{ in } \mathfrak{t}_2/\mu$ is well-typed, then \mathfrak{t}_1/μ is well-typed. \square

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

Proof: The proof is by induction on the structure of \mathfrak{t} .

- *Case* $\mathfrak{t} = \mathfrak{z}$. Well-typed configurations are closed: this case cannot occur.
- *Case* $\mathfrak{t} = m$. \mathfrak{t} is a value.
- *Case* $\mathfrak{t} = c$. By requirement (ii) of Definition 1.7.6.
- *Case* $\mathfrak{t} = \lambda \mathfrak{z}. \mathfrak{t}_1$. \mathfrak{t} is a value.

- *Case* $\mathfrak{t} = \mathfrak{t}_1 \ \mathfrak{t}_2$. By Lemma 1.7.12, \mathfrak{t}_1/μ is well-typed. By the induction hypothesis, either it is reducible, or \mathfrak{t}_1 is a value. If the former, by R-CONTEXT and because every context of the form $\mathcal{E} \ \mathfrak{t}_2$ is an evaluation context, the configuration \mathfrak{t}/μ is reducible as well. Thus, let us assume \mathfrak{t}_1 is a value. By Lemma 1.7.12, \mathfrak{t}_2/μ is well-typed. By the induction hypothesis, either it is reducible, or \mathfrak{t}_2 is a value. If the former, by R-CONTEXT and because every context of the form $\mathfrak{t}_1 \ \mathcal{E}$ —where \mathfrak{t}_1 is a value—is an evaluation context, the configuration \mathfrak{t}/μ is reducible as well. Thus, let us assume \mathfrak{t}_2 is a value. Let us now reason by cases on the structure of \mathfrak{t}_1 .

Subcase $\mathfrak{t}_1 = \mathfrak{z}$. Again, this subcase cannot occur.

Subcase $\mathfrak{t}_1 = m$. Because \mathfrak{t}/μ is well-typed, a constraint of the form $\text{let } \Gamma_0; \text{ref } M \text{ in } (\exists \mathfrak{X}. (m \preceq \mathfrak{X} \rightarrow \mathbb{T} \wedge \llbracket \mathfrak{t}_2 : \mathfrak{X} \rrbracket) \wedge \llbracket \mu : M \rrbracket)$ must be satisfiable. This implies that m is a member of $\text{dom}(M)$ and that the constraint

$\text{ref } M(m) \leq \mathbf{x} \rightarrow \mathbf{T}$ is satisfiable. Because the type constructors ref and \rightarrow are incompatible, this is a contradiction. So, this subcase cannot occur.

Subcase $\mathfrak{t}_1 = \lambda z. \mathfrak{t}'_1$. By R-BETA, \mathfrak{t}/μ is reducible.

Subcase $\mathfrak{t}_1 = c \, v_1 \dots v_k$. Then, \mathfrak{t} is of the form $c \, v_1 \dots v_{k+1}$. The result follows by requirement (ii) of Definition 1.7.6.

◦ *Case* $\mathfrak{t} = \text{let } z = \mathfrak{t}_1 \text{ in } \mathfrak{t}_2$. By Lemma 1.7.12, \mathfrak{t}_1/μ is well-typed. By the induction hypothesis, either \mathfrak{t}_1/μ is reducible, or \mathfrak{t}_1 is a value. If the former, by R-CONTEXT and because every context of the form $\text{let } z = \mathcal{E} \text{ in } \mathfrak{t}_2$ is an evaluation context, the configuration \mathfrak{t}/μ is reducible as well. If the latter, then \mathfrak{t}/μ is reducible by R-LET. \square

We may now conclude:

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

Proof: We say that a source program \mathfrak{t} is well-typed if and only if the configuration \mathfrak{t}/\emptyset is well-typed, that is, if and only if $\exists \mathbf{x}. \text{let } \Gamma_0 \text{ in } [\mathfrak{t} : \mathbf{x}] \equiv \text{true}$ holds. By Lemma 1.7.11, all reducts of \mathfrak{t}/\emptyset are well-typed. By Theorem 1.7.13, none is stuck. \square

Let us recall that this result holds only if the requirements of Definition 1.7.6 are met. In other words, some proof obligations remain to be discharged when concrete definitions of \mathcal{Q} , $\xrightarrow{\delta}$, and Γ_0 are given. This is illustrated by several examples in the next section.

1.8 Constraint solving

We have introduced a parameterized constraint language, given equivalence laws that describe 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 precisely lies in the treatment of primitive constraints and in its interaction with **let**-polymorphism. For this reason, 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 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. Other constraint solvers used in program analysis or type inference are described *e.g.* in (Aiken and Wimmers, 1992; 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).

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, deprived of 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. We now present a unification algorithm due to Huet (1976) as a (nondeterministic) system of constraint rewriting rules. The specification is almost the same in the case of *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* terms, and does not rely in an essential way on the occurs check. In order to accurately reflect the behavior of the actual algorithm, which relies on a *union-find* data structure (Tarjan, 1975), we modify the syntax of constraints by replacing equations with *multi-equations*. A multi-equation is an equation that involves an arbitrary number of types, as opposed to exactly two.

- 1.8.1 **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. 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 *multi-set* 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.

- 1.8.2 **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 . \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 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 unification is rather high-level. It would be possible to give a lower-level description, where state (standard conjunctions of multi-equations) and control (pending binary equations) are distinguished.

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

The specification of the unification algorithm consists of a set of constraint rewriting rules, given in Figure 1-11. 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 variables, 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

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

Figure 1-11: Unification

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 explicitly 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-SINGLE is restricted to nonvariable terms so as not to break the property that every variable is a member

of exactly one multi-equation (Definition 1.8.2). **S-CYCLE** is the occurs check: that is, 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 1-11 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 $\exists \bar{x}.U$, where U is satisfiable. The latter two properties indicate that the algorithm is indeed a constraint solver.

- 1.8.4 LEMMA: The rewriting system \rightarrow is strongly normalizing. □
- 1.8.5 LEMMA: $U_1 \rightarrow U_2$ implies $U_1 \equiv U_2$. □
- 1.8.6 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*. Stacks form a subset of constraint contexts, defined on page 24. Their 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 $\exists \sigma \equiv \text{true}$. A stack may be viewed 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 $S[U \wedge C]$. 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 short, 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. 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 1-12. 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 may always be satisfied by suitable α -conversion of the left-hand state. If $S; U; C$ is a normal form with respect to the above five rules, then every type variable in $dtv(S)$ is 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 $dtv(S)$ 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

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

Figure 1-12: A constraint solver

initially described in (Rémy, 1992a), and also appear in (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 scheme carried by the nearest environment frame that defines x in the stack S . It is defined as follows:

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

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. Please recall that, if σ is of the form $\forall \bar{x}[U].X$, where $\bar{x} \# \text{ftv}(T)$, then $\sigma \preceq T$ 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 ”. It involves taking fresh copies of the type variables \bar{x} , of the unification constraint U , and of the body X . In the worst case, this process is just as inefficient as textually expanding the corresponding `let` construct in the program’s source code, and leads to exponential time complexity (Mairson, Kanellakis, and Mitchell, 1991). In practice, however, the unification constraint U is often compact, because it was simplified before the environment frame `let x : σ in \Box` was created. which is 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 $S[\text{let } x : \forall\bar{x}[\square].T \text{ in } C]; U; \text{true}$ into $S[\text{let } x : \forall\bar{x}[U].T \text{ in } \square]; \text{true}; C$, 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 $\forall\bar{x}[U].T$ 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 \bar{x} as *young* and to type variables in $\text{dte}(S) \setminus \bar{x}$ as *old*. The former are the universal quantifiers of the type scheme that is being created; the latter are 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), provided we are careful never to create a link from a variable to a variable of higher rank. This requires making the unification algorithm aware of ranks, but is otherwise

easily achieved. 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 1.8.7). 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 $\bar{X} \# ftv(U_1)$ —and U_2 constrains *young* variables only—as expressed by the side-condition $\exists \bar{X}.U_2 \equiv \text{true}$. Please 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; the proof of Lemma 1.8.11 sheds more light on this issue. Let us say, for now, that S-LETALL plays a role in guaranteeing its existence, whence 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* (see the proof of Lemma 1.8.10) and corresponds to the difference between HMX-GEN' and HMX-GEN discussed in Section 1.4. 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. Please 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 take instances of.

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 1.3.26). 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 $Y_1 Y_2$, 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 free 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 $ftv(T)$ —indeed, the term T itself defines a function that maps the latter to the former. 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. Please note that $X = Y_1 \rightarrow Y_2$ determines $Y_1 Y_2$ and determines X , but does *not* simultaneously determine $XY_1 Y_2$.

1.8.7 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

Proof: Let $\bar{X} \# \bar{Y}$ **(1)**. Let $\phi \vdash \text{def } \Gamma$ in $\exists \bar{X}.(C \wedge \epsilon)$ **(2)** and $\phi' \vdash \text{def } \Gamma$ in $\exists \bar{X}.(C \wedge \epsilon)$ **(3)**, where ϕ and ϕ' coincide outside of \bar{Y} . We may assume, *w.l.o.g.*, $\bar{X} \# ftv(\Gamma)$ **(4)**. By (2), (4), CM-EXISTS, and CM-AND, we obtain $\phi_1 \vdash \text{def } \Gamma$ in ϵ **(5)**, where ϕ and ϕ_1 coincide outside \bar{X} . By CM-PREDICATE, (5) implies that all members of ϵ have the same image through ϕ_1 . Similarly, exploiting (3) and (4), we find that all members of ϵ have the same image through ϕ'_1 , where ϕ' and ϕ'_1 coincide outside \bar{X} . Now, we claim that ϕ_1 and ϕ'_1 coincide on \bar{Y} . Once the claim is established, by (1), there follows that ϕ and ϕ' must coincide on \bar{Y} as well, which is the goal. So, there only remains to establish the claim; we distinguish two subcases.

Subcase ϵ is $X = \epsilon'$ and $X \notin \bar{X}\bar{Y}$ **(6)** and $\bar{Y} \subseteq ftv(\epsilon')$ **(7)**. Because ϕ_1 and ϕ'_1 coincide outside $\bar{X}\bar{Y}$ and by (6), we have $\phi_1(X) = \phi'_1(X)$. As a result, all members of ϵ' have the same image through ϕ_1 and ϕ'_1 . In a free tree model, where decomposition is valid, a simple inductive argument shows that ϕ_1 and ϕ'_1 must coincide on $ftv(\epsilon')$, hence—by (7)—also on \bar{Y} .

Subcase ϵ is $\bar{Y} = T = \epsilon'$ and $ftv(T) \# \bar{X}\bar{Y}$ **(8)**. Because ϕ_1 and ϕ'_1 coincide outside $\bar{X}\bar{Y}$ and by (8), we have $\phi_1(T) = \phi'_1(T)$. Thus, for every $Y \in \bar{Y}$, we have $\phi_1(Y) = \phi_1(T) = \phi'_1(T) = \phi'_1(Y)$. That is, ϕ_1 and ϕ'_1 coincide on \bar{Y} . \square

Thanks to Lemma 1.8.7, a straightforward 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 \prec_U . By the second part of the lemma, it is safe to close \bar{Y} under the closure law $\mathbf{X} \in \bar{X} \wedge (\forall \mathbf{Y} \ \mathbf{Y} \prec_U \mathbf{X} \Rightarrow \mathbf{Y} \in \bar{Y}) \Rightarrow \mathbf{X} \in \bar{Y}$. 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 \prec_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. This is the topic of the following exercise.

- 1.8.8 EXERCISE [★★★, \rightarrow]: Assuming U is acyclic, prove that the above procedure computes the greatest subset \bar{Y} of \bar{X} such that $\exists(\bar{X} \setminus \bar{Y}).U$ determines \bar{Y} . In the setting of a regular tree model, exhibit a satisfiable constraint U such that the above procedure is incomplete. Can you define a complete procedure in that setting? \square

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.

S-SOLVE-LET and S-NAME-2 to S-POP-LET are unnecessarily complex when x is assigned a *monotype* T , rather than an arbitrary type scheme $\forall \bar{X}[D].T$. In that case, the combined effect of these rules may be obtained directly via the following two new rules, which may be implemented in a more efficient way:

$$\begin{aligned}
 S; U; \text{let } x : T \text{ in } C &\rightarrow S[\exists x.[]]; U \wedge x = T; \text{let } x : X \text{ in } C && \text{(S-NAME-2-MONO)} \\
 & \text{if } x \notin \text{fv}(U, T, C) \wedge T \notin \mathcal{V} \\
 S; U; \text{let } x : X \text{ in } C &\rightarrow S[\text{let } x : X \text{ in } []]; U; C && \text{(S-SOLVE-LET-MONO)}
 \end{aligned}$$

If T isn't a variable, it is replaced with a fresh variable X , together with the equation $x = T$. This corresponds to the effect of S-NAME-2. Then, we directly

create an environment frame for \mathbf{x} , without bothering to create and discard a `let` frame, since there is no way the type scheme \mathbf{X} may be further simplified.

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

1.8.9 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 $S[U \wedge C]$.

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

Proof: By examination of every rule.

◦ *Case S-UNIFY.* By Lemma 1.8.5.

◦ *Case S-EX-1, S-EX-2, S-SOLVE-EX.* By C-EXAND.

◦ *Case S-EX-3.* By C-LET $\bar{\text{E}}\text{X}$.

◦ *Case S-EX-4.* By C-INEX.

◦ *Case S-SOLVE-EQ, S-POP-AND.* By C-DUP.

◦ *Case S-SOLVE-ID.* Because σ is of the form $\forall \bar{\mathbf{X}}[U].\mathbf{X}$, we have $fpi(\sigma) = \emptyset$. The result follows by C-INID.

◦ *Case S-SOLVE-AND.* By C-ANDAND.

◦ *Case S-SOLVE-LET.* By C-LETAND.

◦ *Case S-NAME-2.* By Definition 1.3.21 and C-NAMEEQ, $\mathbf{x} \notin ftv(U, \mathbf{T})$ implies $\text{true} \Vdash \forall \bar{\mathbf{X}}[U].\mathbf{T} \equiv \forall \bar{\mathbf{X}}\mathbf{X}[U \wedge \mathbf{x} = \mathbf{T}].\mathbf{x}$. The result follows by Lemma 1.3.22.

◦ *Case S-COMPRESS.* Let $\theta = [Y \mapsto Z]$. By Definition 1.3.21 and C-NAMEEQ, $Y \neq Z$ implies $\text{true} \Vdash \forall \bar{\mathbf{X}}\mathbf{Y}[Y = Z = \epsilon \wedge U].\mathbf{x} \equiv \forall \bar{\mathbf{X}}\mathbf{Y}[Y \wedge Z = \theta(\epsilon) \wedge \theta(U)].\theta(\mathbf{x})$. The result follows by Lemma 1.3.22.

◦ *Case S-UNNAME.* Using Lemma 1.3.18, it is straightforward to check that $Y \notin ftv(\epsilon)$ implies $\exists \mathbf{Y}.\mathbf{Y} = \epsilon \equiv \epsilon$. The result follows by C-EXAND and C-LET $\bar{\text{E}}\text{X}$.

◦ *Case S-LETALL.* By C-LETALL.

◦ *Case S-POP-LET.* By C-LETAND and C-INAND*.

◦ *Case S-POP-ENV.* By C-IN*, recalling that $\exists \sigma$ must be true. \square

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

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

Proof: Because, by definition, $S; U; \text{false}$ is not a valid state, a normal form for S-SOLVE-EQ, S-SOLVE-ID, S-SOLVE-AND, S-SOLVE-EX, and S-SOLVE-LET must be either an instance of the left-hand side of S-SOLVE-ID, with $\mathbf{x} \notin \text{dpi}(S)$, which corresponds to case (i), or of the form $S; U; \text{true}$. Let us consider the latter case. Because $S; U; \text{true}$ is a normal form with respect to S-UNIFY, by Lemma 1.8.6, U must be either false or of the form $\mathcal{X}[U']$, where U' is a standard conjunction of multi-equations and, if the model is syntactic, U' is acyclic. The former case corresponds to (ii); thus, let us consider the latter case. Because $S; \mathcal{X}[U']; \text{true}$ is a normal form with respect to S-EX-1, the context \mathcal{X} must in fact be empty, and U' is U . If S is an existential constraint context, then we are in situation (iii). Otherwise, because $S; U; \text{true}$ is a normal form with respect to S-EX-2, S-EX-3, and S-EX-4, the stack S does not end with an existential frame. Because $S; U; \text{true}$ is a normal form with respect to S-POP-AND and S-POP-ENV, S must then be of the form $S'[\text{let } \mathbf{x} : \forall \bar{\mathbf{x}}[\square]. \text{T in } C]$. Because $S; U; \text{true}$ is a normal form with respect to S-NAME-2, T must be a type variable \mathbf{x} . Let us write U as $U_1 \wedge U_2$, where $\bar{\mathbf{x}} \# \text{ftv}(U_1)$, and where U_1 is maximal for this criterion. Then, consider a multi-equation $\epsilon \in U$. By the first part of Lemma 1.8.7, if one variable member of ϵ is free (that is, outside $\bar{\mathbf{x}}$), then $\exists \bar{\mathbf{x}}. U$ determines all other variables in $\text{ftv}(\epsilon)$. Because $S; U; \text{true}$ is a normal form with respect to S-LETALL, all variables in $\text{ftv}(\epsilon)$ must then be free (that is, outside $\bar{\mathbf{x}}$). By definition of U_1 , this implies $\epsilon \in U_1$. By contraposition, for every multi-equation $\epsilon \in U_2$, all variable members of ϵ are in $\bar{\mathbf{x}}$. Furthermore, let us recall that U_2 is a standard conjunction of multi-equations and, if the model is syntactic, U_2 is acyclic. We let the reader check that this implies $\exists \bar{\mathbf{x}}. U_2 \equiv \text{true}$; the proof is a slight generalization of the last part of that of Lemma 1.8.6. Then, $S; U; \text{true}$ is reducible via S-POP-LET. This is a contradiction, so this last case cannot arise. \square

In case (i), the constraint $S[U \wedge C]$ has a free program identifier \mathbf{x} , so it is not satisfiable. 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 1.3.30, the constraint $S[U \wedge C]$ is then false. In case (iii), the constraint $S[U \wedge C]$ is equivalent to $\mathcal{X}[U]$, where U is satisfiable, so it is satisfiable as well. Thus, each of the three classes of normal forms may be immediately identified as denoting success or failure. Thus, Lemmas 1.8.10 and 1.8.11 indeed prove that the algorithm is a constraint solver.

1.9 From ML-the-calculus to ML-the-programming-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 more complex extensions, namely algebraic data type definitions, pattern matching, and type annotations. Last, the issues associated with recursive types are briefly discussed. Exceptions are not discussed; the reader is referred to (TAPL Chapter 14).

Simple extensions

Many features of ML-the-programming-language may be introduced into ML-the-calculus by introducing new constants and extending $\xrightarrow{\delta}$ and Γ_0 appropriately. In each case, it is necessary to check that the requirements of Definition 1.7.6 are met, that is, 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.

- 1.9.1 EXERCISE [INTEGERS, RECOMMENDED, ★★]: Integer literals and integer addition have been introduced and given an operational semantics in Examples 1.2.1, 1.2.2 and 1.2.4. 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 1.7.6. \square
- 1.9.2 EXERCISE [BOOLEANS, RECOMMENDED, ★★, \rightarrow]: Booleans and conditionals have been introduced and given an operational semantics in Exercise 1.2.6. Introduce an isolated type constructor `bool` to represent Boolean values and explain how to extend the initial environment. Check that your definitions meet the requirements of Definition 1.7.6. What is the constraint generation rule for the syntactic sugar `if t_0 then t_1 else t_2` ? \square
- 1.9.3 EXERCISE [PAIRS, ★★, \rightarrow]: Pairs and pair projections have been introduced and given an operational semantics in Examples 1.2.3 and 1.2.5. 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 1.7.6. \square

- 1.9.4 EXERCISE [SUMS, ★★, \rightarrow]: Sums have been introduced and given an operational semantics in Example 1.2.7. 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 1.7.6. \square

- 1.9.5 EXERCISE [REFERENCES, ★★★]: References have been introduced and given an operational semantics in Example 1.2.9. The type constructor `ref` has been introduced in Definition 1.7.4. 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 1.7.6. \square

- 1.9.6 EXERCISE [RECURSION, RECOMMENDED, ★★★]: The fixpoint combinator `fix` has been introduced and given an operational semantics in Example 1.2.10. 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 1.7.6. Recall how the `letrec` syntactic sugar was defined in Example 1.2.10, 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 is discussed in Section 1.10 (page 113). \square

Algebraic data types

Exercises 1.9.3 and 1.9.4 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 “please 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.

- 1.9.7 **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:

```
letrec length = λl.case 1 (λ_0) (λz.1 + length (π2 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 `l`. 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 (page 106). \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 a signature of the form $\vec{\kappa} \Rightarrow \star$. Furthermore, ℓ ranges over a set \mathcal{L} of *labels*, which we use indifferently 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 $\vec{\kappa} \Rightarrow \star$, then the type variables \vec{X} must have kind $\vec{\kappa}$. 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, $fv(\vec{T}) \subseteq \vec{X}$ 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.

1.9.8 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

The above requirement bears on the definition of subtyping in the model. 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 1.9.8 while assigning D a less restrictive variance. The following example illustrates such a case.

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

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

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 1.9.8, 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

1.9.10 EXERCISE [$\star\star$, \rightarrow]: Consider a nonrecursive algebraic data type definition, where the variance of every type constructor that appears on the right-hand side is known. Can you systematically determine, for each of the parameters, the least restrictive variance that makes the definition valid? Generalize this procedure to the case of recursive and mutually recursive algebraic data type definitions. \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 1.2.7. A record type definition extends it with a record formation construct and several projections, as in Examples 1.2.3 and 1.2.5. 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 $\text{case}_{\mathbb{D}}$, 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 $\text{case}_{\mathbb{D}} \mathfrak{t} \mathfrak{t}_1 \dots \mathfrak{t}_n$. 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 $\text{case}_{\mathbb{D}} : \forall \vec{X}. D\vec{X} \rightarrow (T_1 \rightarrow Z) \rightarrow \dots (T_k \rightarrow Z) \rightarrow Z$.

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 $\text{make}_{\mathbb{D}}$, 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 $\text{make}_{\mathbb{D}} \mathfrak{t}_1 \dots \mathfrak{t}_k$. 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 $\text{make}_{\mathbb{D}} : \forall \vec{X}. T_1 \rightarrow \dots \rightarrow T_k \rightarrow D\vec{X}$.

1.9.11 EXAMPLE: The effect of defining list (Example 1.9.9) 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:

```
letrec length = λl.case 1 [ Nil : λ_· $\hat{0}$  | Cons : λz.  $\hat{1}$  + length ( $\pi_2$  z) ]
```

Recall that this notation is syntactic sugar for

```
letrec length = λl.caselist 1 (λ_· $\hat{0}$ ) (λz.  $\hat{1}$  + length ( $\pi_2$  z))
```

The difference with the code in Example 1.9.7 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 `caselist`, which is derived from the definition of `list`, instead of the type scheme assigned to the anonymous `case` construct, given in Exercise 1.9.4. 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 1.9.7, 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, instead, `caselist` requires l to have type `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 1.9.11 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 this 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 `caseD` *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 type-checker to fold or unfold an algebraic data type definition. This mechanism is characteristic of *isorecursive* types.

- 1.9.12 EXERCISE [★, ⇨]: 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
- 1.9.13 EXERCISE [★★★, ⇨]: Check that the above definitions meet the requirements of Definition 1.7.6. \square
- 1.9.14 EXERCISE [★★★, ⇨]: For sake of simplicity, we have assumed that data constructors are always of arity one. It is indeed possible to allow data constructors of any arity and define variants as $D\vec{X} \approx \sum_{i=1}^k \ell_i : \vec{T}_i$. For instance, the definition of list could then be $\text{list } X \approx \Sigma(\text{Nil}; \text{Cons} : X \times \text{list } X)$ and for instance $\text{Cons}(\hat{0}, \text{Nil})$ would be a list value. Make the necessary changes in the definitions above and check that they still meet the requirements of Definition 1.7.6. \square

In this formal presentation of algebraic data types, we have assumed that all algebraic data type definitions are known before the program is typechecked. This simplifying assumption is forced on us by the fact that we interpret constraints in a fixed model, that is, we assume a fixed universe of types. In practice, programming languages have *module systems*, which allow distinct modules to have distinct, partial views of the universe of types. Then, it becomes possible for each module to come with its own data type definitions. Interestingly, it is even possible, in principle, to split the definition of a single data type over several modules, yielding *extensible* algebraic data types. For instance, module A might declare the existence of a parameterized variant type $D\vec{X}$, without giving its components. Later on, module B might define a component $\ell : T$, where $\text{ftv}(T) \subseteq \vec{X}$. Such a definition makes ℓ a unary constructor with type scheme $\forall \vec{X}. T \rightarrow D\vec{X}$, as before. It becomes impossible, however, to introduce a destructor case_D , because the definition of an extensible variant type can never be assumed to be complete—other, unknown modules might extend it further. To compensate for its absence, one may supplement every constructor ℓ with a destructor ℓ^{-1} , whose semantics is given by $\ell^{-1}(\ell \ v) \ v_1 \ v_2 \xrightarrow{\delta} v_1 \ v$ and $\ell^{-1}(\ell' \ v) \ v_1 \ v_2 \xrightarrow{\delta} v_2 \ (\ell' \ v)$ when $\ell \neq \ell'$, and whose type scheme is $\forall \vec{X}. D\vec{X} \rightarrow (T \rightarrow Z) \rightarrow (D\vec{X} \rightarrow Z) \rightarrow Z$. When

pattern matching is available, ℓ^{-1} may in fact be defined in the language. ML-the-programming-language does not offer extensible algebraic data types as a language feature, but does have one built-in extensible variant type, namely the type `exn` of exceptions. Thus, it is possible to define new constructors for the type `exn` within any module. The price of this extra flexibility is that no exhaustive case analysis on values of type `exn` is possible.

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 become inaccessible. 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 Section 1.11.

Pattern matching

Our presentation of products, sums and algebraic data types has remained within the setting of ML-the-calculus: that is, data structures have been built out of constructors, while the case analysis and record access operations have been viewed as destructors. Some syntactic sugar has been used to recover standard notations. The language is now expressive enough to allow defining and manipulating complex data structures, such as lists and trees. Yet, experience shows that programming in such a language is still somewhat cumbersome. Indeed, case analysis and record access are low-level operations: the former allows inspecting a tag and branching, while the latter allows dereferencing a pointer. In practice, one often needs to carry out more complex tasks, such as determining whether a data structure has a certain shape or whether two data structures have comparable shapes. Currently, the only way to carry out these tasks is to program an explicit sequence of low-level operations. It would be much preferable to extend the language so that it becomes directly possible to describe shapes, called *patterns*, and so that checking whether a patterns *matches* a value becomes an elementary operation. ML-the-programming-language offers this feature, called *pattern matching*. Although pattern matching may be added to ML-the-calculus by introducing a family of destructors, we rather choose to extend the calculus with a new `match` construct, which subsumes the existing `let` construct. This approach appears somewhat simpler and more powerful. We now carry out this

| $p ::=$ | Patterns: | Pattern matching |
|-------------------|-------------------------------------|--|
| $_$ | Wildcard | $[_ \mapsto v] = \emptyset$ |
| z | Variable | |
| $c p_1 \dots p_k$ | Data $c \in Q^+ \wedge k = a(c)$ | $[c p_1 \dots p_k \mapsto c v_1 \dots v_k]$ $= [p_1 \mapsto v_1] \otimes \dots \otimes [p_k \mapsto v_k]$ |
| $p \wedge p$ | Conjunction | $[p_1 \wedge p_2 \mapsto v] = [p_1 \mapsto v] \otimes [p_2 \mapsto v]$ |
| $p \vee p$ | Disjunction | $[p_1 \vee p_2 \mapsto v] = [p_1 \mapsto v] \oplus [p_2 \mapsto v]$ |

Figure 1-13: Patterns and pattern matching

extension.

Let us first define the syntax of patterns (Figure 1-13) and describe (informally, for now) which values they match. To a pattern p , we associate a set of *defined program variables* $dpi(p)$, whose definition appears in the text that follows. The pattern p is well-formed if and only if $dpi(p)$ is defined. To begin, the wildcard $_$ is a pattern, which matches every value and binds no variables. We let $dpi(_) = \emptyset$. Although the wildcard may be viewed as an anonymous variable, and we have done so thus far, it is now simpler to view it as a distinct pattern. A program variable z is also a pattern, which matches every value and binds z to the matched value. We let $dpi(z) = \{z\}$. Next, if c is a constructor of arity k , then $c p_1 \dots p_k$ is a pattern, which matches $c v_1 \dots v_k$ when p_i matches v_i for every $i \in \{1, \dots, k\}$. We let $dpi(c p_1 \dots p_k) = dpi(p_1) \uplus \dots \uplus dpi(p_k)$. That is, the pattern $c p_1 \dots p_k$ is well-formed when p_1, \dots, p_k define *disjoint* sets of variables. This condition rules out *nonlinear* patterns such as (z, z) . Defining the semantics of such a pattern would require a notion of equality at every type, which introduces various complications, so it is commonly considered ill-formed. The pattern $p_1 \wedge p_2$ matches all values that both p_1 and p_2 match. It is commonly used with p_2 a program variable: then, it allows examining the shape of a value and binding a name to it at the same time. Again, we define $dpi(p_1 \wedge p_2) = dpi(p_1) \uplus dpi(p_2)$. The pattern $p_1 \vee p_2$ matches all values that either p_1 or p_2 matches. We define $dpi(p_1 \vee p_2) = dpi(p_1) \cup dpi(p_2)$. That is, the pattern $p_1 \vee p_2$ is well-formed when p_1 and p_2 define the *same* variables. Thus, $(inj_1 z) \vee (inj_2 z)$ is a well-formed pattern, which binds z to the component of a binary sum, without regard for its tag. However, $(inj_1 z_1) \vee (inj_2 z_2)$ is ill-formed, because one cannot statically predict whether it defines z_1 or z_2 .

Let us now formally define whether a pattern p matches a value v and how the variables in $dpi(p)$ become bound to values in the process. This is done by introducing a *generalized substitution*, written $[p \mapsto v]$, which is either

| | | |
|--|--|--|
| $t ::= \dots$ $\text{match } t \text{ with } (p_i \cdot t_i)_{i=1}^k$ | <p style="text-align: center;"><i>Expressions:</i></p> | <p style="text-align: center;"><i>Reduction rules</i></p> $\text{match } v \text{ with } (p_i \cdot t_i)_{i=1}^k \longrightarrow \bigoplus_{i=1}^k [p_i \mapsto v]t_i$ <p style="text-align: right;">(R-MATCH)</p> |
| $\mathcal{E} ::= \dots$ $\text{match } \mathcal{E} \text{ with } (p_i \cdot t_i)_{i=1}^k$ | <p style="text-align: center;"><i>Evaluation Contexts:</i></p> | |

Figure 1-14: Extended syntax and semantics of ML-the-calculus

undefined or a substitution of values for the program variables in $dpi(\mathbf{p})$. If the former, then \mathbf{p} does not match \mathbf{v} . If the latter, then \mathbf{p} matches \mathbf{v} and, for every $z \in dpi(\mathbf{p})$, the variable z becomes bound to the value $[\mathbf{p} \mapsto \mathbf{v}]z$. Of course, when \mathbf{p} is a variable z , the generalized substitution $[z \mapsto \mathbf{v}]$ is defined and coincides with the substitution $[z \mapsto v]$, which justifies our abuse of notation. To construct generalized substitutions, we use two simple combinators. First, when $dpi(\mathbf{p}_1)$ and $dpi(\mathbf{p}_2)$ are disjoint, $[\mathbf{p}_1 \mapsto \mathbf{v}_1] \otimes [\mathbf{p}_2 \mapsto \mathbf{v}_2]$ stands for the set-theoretic union of $[\mathbf{p}_1 \mapsto \mathbf{v}_1]$ and $[\mathbf{p}_2 \mapsto \mathbf{v}_2]$, if both are defined, and is undefined otherwise. We use this combinator to ensure that \mathbf{p}_1 matches \mathbf{v}_1 and \mathbf{p}_2 matches \mathbf{v}_2 and to combine the two corresponding sets of bindings. Second, when o_1 and o_2 are two possibly undefined mathematical objects that belong to the same space when defined, $o_1 \oplus o_2$ stands for o_1 , if it is defined, and for o_2 otherwise—that is, \oplus is an angelic choice operator with a left bias. In particular, when $dpi(\mathbf{p}_1)$ and $dpi(\mathbf{p}_2)$ coincide, $[\mathbf{p}_1 \mapsto \mathbf{v}_1] \oplus [\mathbf{p}_2 \mapsto \mathbf{v}_2]$ stands for $[\mathbf{p}_1 \mapsto \mathbf{v}_1]$, if it is defined, and for $[\mathbf{p}_2 \mapsto \mathbf{v}_2]$ otherwise. We use this combinator to ensure that \mathbf{p}_1 matches \mathbf{v}_1 or \mathbf{p}_2 matches \mathbf{v}_2 and to retain the corresponding set of bindings. The full definition of generalized substitutions, which relies on these combinators, appears in Figure 1-13. It reflects the informal presentation of the previous paragraph.

Once patterns and pattern matching are defined, it is straightforward to extend the syntax and operational semantics of ML-the-calculus. We enrich the syntax of expressions with a new construct, $\text{match } t \text{ with } (p_i \cdot t_i)_{i=1}^k$, where $k \geq 1$. It consists of a term t and a nonempty, ordered list of clauses, each of which is composed of a pattern p_i and a term t_i . The syntax of evaluation contexts is extended as well, so that the term t that is being examined is first reduced to a value v . The operational semantics is extended with a new rule, R-MATCH, which states that $\text{match } v \text{ with } (p_i \cdot t_i)_{i=1}^k$ reduces to $[p_i \mapsto v]t_i$, where i is the least element of $\{1, \dots, k\}$ such that p_i matches v_i . Technically, $\bigoplus_{i=1}^k [p_i \mapsto v]t_i$ stands for $[p_1 \mapsto v]t_1 \oplus \dots \oplus [p_k \mapsto v]t_k$, so that the reduct is the first term that is defined in this sequence.

As far as semantics is concerned, the `match` construct may be viewed as a

generalization of the `let` construct. Indeed, `let z = t1 in t2` may now be viewed as syntactic sugar for `match t1 with z . t2`, that is, a `match` construct with a single clause and a variable pattern. Then, R-LET becomes a special case of R-MATCH.

It is pleasant to introduce some more syntactic sugar. We write $\lambda(\mathbf{p}_i . \mathbf{t}_i)_{i=1}^k$ for $\lambda \mathbf{z} . \text{match } \mathbf{z} \text{ with } (\mathbf{p}_i . \mathbf{t}_i)_{i=1}^k$, where \mathbf{z} is fresh for $(\mathbf{p}_i . \mathbf{t}_i)_{i=1}^k$. Thus, it becomes possible to define functions by cases—a common idiom in ML-the-programming-language.

- 1.9.15 EXAMPLE: Using pattern matching, a function that computes the length of a list (Example 1.9.11) may now be written as follows:

$$\text{letrec length} = \lambda(\text{Nil } _ . \hat{0} \mid \text{Cons } (_, \mathbf{z}) . \hat{1} \hat{+} \text{length } \mathbf{z})$$

The second pattern matches a nonempty list and binds \mathbf{z} to its tail at the same time, obviating the need for an explicit application of π_2 . \square

- 1.9.16 EXERCISE [★★, RECOMMENDED, \rightarrow]: Under the above definition of `length`, consider an application of `length` to the list `Cons($\hat{0}$, Nil())`. After eliminating the syntactic sugar, determine by which reduction sequence this expression reduces to a value. \square

Before we can proceed and extend the type system to deal with the new `match` construct, we must make two mild extensions to the syntax and meaning of constraints. First, if σ is $\forall \bar{\mathbf{x}}[C].\mathbf{T}$, where $\bar{\mathbf{x}} \# \text{ftv}(\mathbf{T}')$, then $\mathbf{T}' \preceq \sigma$ stands for the constraint $\exists \bar{\mathbf{x}}.(C \wedge \mathbf{T}' \leq \mathbf{T})$. This relation is identical to the instance relation (Definition 1.3.3), except the direction of subtyping is reversed. We extend the syntax of constraints with instantiation constraints of the form $\mathbf{T} \preceq \mathbf{x}$ and define their meaning by adding a symmetric counterpart of CM-INSTANCE. We remark that, when subtyping is interpreted as equality, the relations $\sigma \preceq \mathbf{T}$ and $\mathbf{T} \preceq \sigma$ coincide, so this extension is unnecessary in that particular case. Second, we extend the syntax of environments so that several successive bindings may *share* a set of quantifiers and a constraint. That is, we allow writing $\forall \bar{\mathbf{x}}[C].(\mathbf{x}_1 : \mathbf{T}_1; \dots; \mathbf{x}_k : \mathbf{T}_k)$ for $\mathbf{x}_1 : \forall \bar{\mathbf{x}}[C].\mathbf{T}_1; \dots; \mathbf{x}_k : \forall \bar{\mathbf{x}}[C].\mathbf{T}_k$. From a theoretical standpoint, this is little more than syntactic sugar; however, in practice, it is useful to implement this new idiom literally, since it avoids unnecessary copying of the constraint C .

Let us now extend the type system. For the sake of brevity, we extend the constraint generation rules only. Of course, it would also be possible to define corresponding extensions of the rule-based type systems shown earlier, namely DM, HM(X), and PCB(X). We begin by defining a constraint $\llbracket \mathbf{T} : \mathbf{p} \rrbracket$ that represents a necessary and sufficient condition for values of type \mathbf{T} to be acceptable inputs for the pattern \mathbf{p} . Its free type variables are a subset of

$$\begin{aligned}
\llbracket \mathbf{T} : _ \rrbracket &= \text{true} \\
\llbracket \mathbf{T} : \mathbf{z} \rrbracket &= \mathbf{T} \preceq \mathbf{z} \\
\llbracket \mathbf{T} : \mathbf{c} \mathbf{p}_1 \dots \mathbf{p}_k \rrbracket &= \exists \vec{\mathbf{X}}. (\vec{\mathbf{X}} \rightarrow \mathbf{T} \preceq \mathbf{c} \wedge \bigwedge_{i=1}^k \llbracket \mathbf{X}_i : \mathbf{p}_i \rrbracket) \\
\llbracket \mathbf{T} : \mathbf{p}_1 \wedge \mathbf{p}_2 \rrbracket &= \llbracket \mathbf{T} : \mathbf{p}_1 \rrbracket \wedge \llbracket \mathbf{T} : \mathbf{p}_2 \rrbracket \\
\llbracket \mathbf{T} : \mathbf{p}_1 \vee \mathbf{p}_2 \rrbracket &= \llbracket \mathbf{T} : \mathbf{p}_1 \rrbracket \wedge \llbracket \mathbf{T} : \mathbf{p}_2 \rrbracket \\
\llbracket \text{match } \mathbf{t} \text{ with } (\mathbf{p}_i . \mathbf{t}_i)_{i=1}^k : \mathbf{T} \rrbracket &= \bigwedge_{i=1}^k \text{let } \forall \mathbf{X} \vec{\mathbf{X}}_i. [\llbracket \mathbf{t} : \mathbf{X} \rrbracket \wedge \text{let } \vec{\mathbf{z}}_i : \vec{\mathbf{X}}_i \text{ in } \llbracket \mathbf{X} : \mathbf{p}_i \rrbracket]. (\vec{\mathbf{z}}_i : \vec{\mathbf{X}}_i) \text{ in } \llbracket \mathbf{t}_i : \mathbf{T} \rrbracket \\
&\quad \text{where } \vec{\mathbf{z}}_i = \text{dpi}(\mathbf{p}_i)
\end{aligned}$$

Figure 1-15: Constraint generation for patterns and pattern matching

$ftv(\mathbf{T})$, while its free program identifiers are either constructors or program variables bound by \mathbf{p} . It is defined in the upper part of Figure 1-15. The first rule states that a wildcard matches values of arbitrary type. The second and third rules govern program variables and constructor applications in *patterns*. They are identical to the rules that govern these constructs in *expressions* (page 59), except that the direction of subtyping is reversed. In the absence of subtyping, they would be entirely identical. We write $\vec{\mathbf{X}}$ for $\mathbf{X}_1 \dots \mathbf{X}_k$ and $\vec{\mathbf{X}} \rightarrow \mathbf{T}$ for $\mathbf{X}_1 \rightarrow \dots \rightarrow \mathbf{X}_k \rightarrow \mathbf{T}$. As usual, the type variables $\mathbf{X}_1, \dots, \mathbf{X}_k$ must have kind \star and must be distinct and fresh for the equation's left-hand side. The last two rules simply distribute the type \mathbf{T} to both subpatterns. It is easy to check that $\llbracket \mathbf{T} : \mathbf{p} \rrbracket$ is *contravariant* in \mathbf{T} :

1.9.17 LEMMA: $\mathbf{T}' \leq \mathbf{T} \wedge \llbracket \mathbf{T} : \mathbf{p} \rrbracket$ entails $\llbracket \mathbf{T}' : \mathbf{p} \rrbracket$. □

This property reflects the fact that \mathbf{T} represents the type of an *input* for the pattern \mathbf{p} . Compare it with Lemma 1.6.3.

1.9.18 EXAMPLE: Consider the pattern $\text{Cons}(_, \mathbf{z})$, which appears in Example 1.9.15. We have

$$\begin{aligned}
&\llbracket \mathbf{T} : \text{Cons}(_, \mathbf{z}) \rrbracket \\
&\equiv \exists \mathbf{Z}_1. (\llbracket \mathbf{Z}_1 \rightarrow \mathbf{T} : \text{Cons} \rrbracket \wedge \llbracket \mathbf{Z}_1 : (_, \mathbf{z}) \rrbracket) \\
&\equiv \exists \mathbf{Z}_1. (\mathbf{Z}_1 \rightarrow \mathbf{T} \preceq \text{Cons} \wedge \exists \mathbf{Z}_2 \mathbf{Z}_3. (\llbracket \mathbf{Z}_2 \rightarrow \mathbf{Z}_3 \rightarrow \mathbf{Z}_1 : (\cdot, \cdot) \rrbracket \wedge \llbracket \mathbf{Z}_2 : _ \rrbracket \wedge \llbracket \mathbf{Z}_3 : \mathbf{z} \rrbracket)) \\
&\equiv \exists \mathbf{Z}_1 \mathbf{Z}_2 \mathbf{Z}_3. (\mathbf{Z}_1 \rightarrow \mathbf{T} \preceq \text{Cons} \wedge \mathbf{Z}_2 \rightarrow \mathbf{Z}_3 \rightarrow \mathbf{Z}_1 \preceq (\cdot, \cdot) \wedge \mathbf{Z}_3 \preceq \mathbf{z})
\end{aligned}$$

where $\mathbf{Z}_1, \mathbf{Z}_2, \mathbf{Z}_3$ are fresh for \mathbf{T} . Let us now place this constraint within the scope of the initial environment, which assigns type schemes to the constructors Cons and (\cdot, \cdot) , and within the scope of a binding of \mathbf{z} to some type \mathbf{T}' .

We find

$$\begin{aligned}
& \text{let } \Gamma_0 \text{ in let } z : T' \text{ in } \llbracket T : \text{Cons } (_, z) \rrbracket \\
\equiv & \exists Z_1 Z_2 Z_3. (\exists X. (Z_1 \rightarrow T \leq X \times \text{list } X \rightarrow \text{list } X) \wedge \\
& \quad \exists Y_1 Y_2. (Z_2 \rightarrow Z_3 \rightarrow Z_1 \leq Y_1 \rightarrow Y_2 \rightarrow Y_1 \times Y_2) \wedge Z_3 \leq T') \\
\equiv & \exists X. (T \leq \text{list } X \wedge \text{list } X \leq T')
\end{aligned}$$

where the final simplification relies mainly on C-ARROW, on the corresponding rule for products, and on C-EXTRANS, and is left as an exercise to the reader. Thus, the constraint states that the pattern matches values that have type $\text{list } X$ (equivalently, values whose type T is a subtype of $\text{list } X$), for some undetermined element type X , and binds z to values of type $\text{list } X$ (equivalently, values whose type T' is a supertype of $\text{list } X$). \square

The above example seems to indicate that the constraint generation rules for patterns make some sense. Still, the careful reader may be somewhat puzzled by the third rule, which, compared to its analogue for expressions, reverses the direction of subtyping, but does not reverse the direction of instantiation. Indeed, in order for this rule to make sense, and to be sound, we must formulate a requirement concerning the type schemes assigned to constructors.

1.9.19 DEFINITION: A constructor c is *invertible* if and only if, when \vec{X} and \vec{X}' have length $a(c)$, the constraint $\text{let } \Gamma_0 \text{ in } (\vec{X}' \rightarrow T \leq c \wedge c \leq \vec{X} \rightarrow T)$ entails $\vec{X} \leq \vec{X}'$. In the following, we assume patterns contain invertible constructors only. \square

Intuitively, when c is invertible, it is possible to recover the type of every v_i from the type of $c v_1 \dots v_k$, a crucial property for pattern matching to be possible. Please note that, if $\Gamma_0(c)$ is monomorphic, then c is invertible. The following lemma identifies another important class of invertible constructors.

1.9.20 LEMMA: The constructors of algebraic data types are invertible. \square

Proof: Let c be a constructor introduced by the definition of an algebraic data type D . Let $k = a(c)$. Then, the type scheme $\Gamma_0(c)$ is of the form $\forall \vec{Y}. \vec{T} \rightarrow D \vec{Y}$, where \vec{Y} are the parameters of the definition and \vec{T} , a vector of length k , consists of *some of* the definition's components. (More precisely, \vec{T} contains just one component in the case of variant types and contains all components in the case of record types.) Let \vec{X} and \vec{X}' have length k . Let $\forall \vec{Y}_1. \vec{T}_1 \rightarrow D \vec{Y}_1$ and $\forall \vec{Y}_2. \vec{T}_2 \rightarrow D \vec{Y}_2$ be two α -equivalent forms of the type scheme $\Gamma_0(c)$, with $\vec{Y}_1 \# \vec{Y}_2$ and $\vec{Y}_1 \vec{Y}_2 \# \text{ftv}(\vec{X}, \vec{X}', T)$. The constraint $\text{let } \Gamma_0 \text{ in } (\vec{X}' \rightarrow T \leq c \wedge c \leq \vec{X} \rightarrow T)$ is, by definition, equivalent to $\vec{X}' \rightarrow T \leq \Gamma_0(c) \wedge \Gamma_0(c) \leq \vec{X} \rightarrow T$, that is, $\exists \vec{Y}_1. (\vec{X}' \rightarrow T \leq \vec{T}_1 \rightarrow D \vec{Y}_1) \wedge \exists \vec{Y}_2. (\vec{T}_2 \rightarrow D \vec{Y}_2 \leq \vec{X} \rightarrow T)$. By C-EXAND and C-ARROW, this may be written $\exists \vec{Y}_1 \vec{Y}_2. (D \vec{Y}_2 \leq T \leq D \vec{Y}_1 \wedge \vec{X} \leq \vec{T}_2 \wedge \vec{T}_1 \leq \vec{X}')$. Now,

by Definition 1.9.8, $D\bar{V}_2 \leq D\bar{V}_1$ entails $\bar{T}_2 \leq \bar{T}_1$, so the previous constraint entails $\exists\bar{Y}_1\bar{Y}_2.(\bar{X} \leq \bar{X}')$, that is, $\bar{X} \leq \bar{X}'$. \square

An important class of *noninvertible* constructors are those associated with existential type definitions (page 118), where not all quantifiers of the type scheme $\Gamma_0(c)$ are parameters of the type constructor D . For instance, under the definition $D \approx \ell : \exists X.X$, the type scheme associated with ℓ is $\forall X.X \rightarrow D$. Then, it is easy to check that ℓ is not invertible. This reflects the fact that it is not possible to recover the type of v from the type of ℓv —which must be D in any case—and explains why existential types require special treatment.

We are now ready to associate a constraint generation rule with the `match` construct. It is given in the lower part of Figure 1-15. In the rule's right-hand side, we write \bar{z}_i for the program variables bound by the pattern p_i , and we write \bar{X}_i for a vector of type variables of the same length. The type variables $X\bar{X}_i$ must have kind \star , must be pairwise distinct and must not appear free in the rule's left-hand side. Let us now explain the rule. Its right-hand side is a conjunction, where each conjunct deals with one clause of the `match` construct, requiring t_i to have type T under certain assumptions about the program variables \bar{z}_i bound by the pattern p_i . There remains to explain how these assumptions are built. First, as in the case of a `let` construct, we summon a fresh type variable X and produce $\llbracket t : X \rrbracket$, the least specific constraint that guarantees t has type X . Then, reflecting the operational semantics, which feeds (the value produced by) t into the pattern p_i , we feed the type X into p_i and produce $\text{let } \bar{z}_i : \bar{X}_i \text{ in } \llbracket X : p_i \rrbracket$, a constraint that guarantees that \bar{X}_i is a correct vector of type assumptions for the program variables \bar{z}_i (see Example 1.9.18). This explains why we may place $\llbracket T : t_i \rrbracket$ within the scope of $(\bar{z}_i : \bar{X}_i)$. There remains to point out that, as in the case of the `let` construct, *every* assignment of ground types to $X\bar{X}_i$ that satisfies the constraint $\llbracket t : X \rrbracket \wedge \text{let } \bar{z}_i : \bar{X}_i \text{ in } \llbracket X : p_i \rrbracket$ is acceptable, so it is valid to universally quantify these type variables. This allows the program variables \bar{z}_i to receive polymorphic type schemes when t itself has polymorphic type.

1.9.21 EXERCISE [\star , RECOMMENDED]: We have previously suggested viewing `let z = t1 in t2` as syntactic sugar for `match t1 with z. t2`, and shown that the operational semantics validates this view. Check that it is also valid from a typing perspective. \square

The `match` constraint generation rule, if implemented literally, takes k copies of the constraint $\llbracket t : X \rrbracket$. When k is greater than 1, this compromises the linear time and space complexity of constraint generation. To remedy this problem, one may modify the rule as follows: replace every copy of $\llbracket t : X \rrbracket$ with $z \preceq X$ and place the constraint within the context `let z : $\forall X[\llbracket t : X \rrbracket].X$ in []`, where z is

a fresh program variable. It is not difficult to check that the logical meaning of the constraint is not affected and that a linear behavior is recovered. In practice, solving the new constraint requires taking instances of the type scheme $\forall X[[\mathbf{t} : X]].X$, which essentially requires copying $[[\mathbf{t} : X]]$ again—however, an efficient solver may now *simplify* this subconstraint before duplicating it.

The following lemma is a key to establishing subject reduction for R-MATCH. It relies on the requirement that constructors be invertible.

- 1.9.22 LEMMA: Assume $[p \mapsto v]$ is defined and maps \vec{z} to \vec{w} , where $\bar{z} = dpi(p)$. Let $\vec{z} : \vec{T}$ be an arbitrary monomorphic environment of domain \bar{z} . Then, let Γ_0 in $([[v : T]] \wedge \text{let } \vec{z} : \vec{T} \text{ in } [[T : p]])$ entails let Γ_0 in $[[\vec{w} : \vec{T}]]$. \square

We now prove that our extension of ML-the-calculus with pattern matching enjoys subject reduction. We only state that R-MATCH preserves types, and leave the new subcase of R-CONTEXT, where the evaluation context involves a `match` construct, to the reader. For this subcase to succeed, the value restriction (Definition 1.7.7) must be extended to require that either all constants have pure semantics or all `match` constructs are in fact of the form `match v with (pi . ti)i=1k`.

- 1.9.23 THEOREM [SUBJECT REDUCTION]: (R-MATCH) \subseteq (\square). \square

- 1.9.24 EXERCISE [★★★, \rightarrow]: For the sake of simplicity, we have omitted the production `ref p` from the syntax of patterns. The pattern `ref p` matches every memory location whose content (with respect to the current store) is matched by `p`. Determine how the previous definitions and proofs must be extended in order to accommodate this new production. \square

The progress property does *not* hold in general: for instance, `match Nil with (Cons z . z)` is well-typed (with type $\forall X.X$) but is stuck. In actual implementations of ML-the-programming-language, such errors are dynamically detected. This may be considered a weakness of ML-the-type-system. Fortunately, however, it is often possible to statically prove that a particular `match` construct is *exhaustive* and cannot go wrong. Indeed, if `match v with (pi . ti)i=1k` is well-typed, then for every $i \in \{1, \dots, k\}$, the constraint `let Γ_0 in $([[v : X]] \wedge \exists \vec{x}.\text{let } \vec{z}_i : \vec{X} \text{ in } [[X : p_i]])$` , where \vec{z}_i are the program variables bound by `pi`, must be satisfiable; that is, `v` must have some type that is an acceptable input for `pi`. This fact yields information about `v`, from which it may be possible to derive that `v` must match one of the patterns `pi`.

- 1.9.25 EXAMPLE: Let $k = 2$, $p_1 = \text{Nil } _$, and $p_2 = \text{Cons } (z_1, z_2)$. Then, the constraints `let Γ_0 in $\exists \vec{x}.\text{let } \vec{z}_i : \vec{X} \text{ in } [[X : p_i]]$` , for $i \in \{1, 2\}$, are both equivalent (after simplification, when $i = 2$) to $\exists Z.X \leq \text{list } Z$. Because the type constructor `list` is isolated, every closed value `v` whose type `X` satisfies this constraint

must be an application of `Nil` or `Cons`. If the latter, because `Cons` has type $\forall X. X \times \text{list } X \rightarrow \text{list } X$, and because the type constructor \times is isolated, the argument to `Cons` must be a pair. We conclude that `v` must match either `p1` or `p2`, which guarantees that this `match` construct is exhaustive and its evaluation cannot go wrong. \square

It is beyond the scope of this chapter to give more details about the check for exhaustiveness. The reader is referred to (Sekar, Ramesh, and Ramakrishnan, 1995; Le Fessant and Maranget, 2001).

Type annotations

So far, we have been interested in a very pure, and extreme, form of type inference. Indeed, in ML-the-calculus, expressions contain no explicit type information whatsoever: it is entirely inferred. In practice, however, it is often useful to insert *type annotations* within expressions, because they provide a form of machine-checked documentation. Type annotations are also helpful when attempting to trace the cause of a type error: by supplying the type-checker with (supposedly) correct type information, one runs a better chance of finding a type inconsistency near an actual programming mistake.

When type annotations are allowed to contain type variables, one must be quite careful about *where* (at which program point) and *how* (existentially or universally) these variables are bound. Indeed, the meaning of type annotations cannot be made precise without settling these issues. In what follows, we first explain how to introduce type annotations whose type variables are bound *locally* and existentially. We show that extending ML-the-calculus with such limited type annotations is again a simple matter of introducing new constants. Then, we turn to a more general case, where type variables may be explicitly existentially introduced *at any program point*. We defer the discussion of *universally* bound type variables to Section 1.10.

Let a *local existential type annotation* $\exists \bar{X}. T$ be a pair of a set of type variables \bar{X} and a type T , where T has kind \star , \bar{X} is considered bound within T , and \bar{X} contains $fv(T)$. For every such annotation, we introduce a new unary destructor $(\cdot : \exists \bar{X}. T)$. Such a definition is valid only because a type annotation must be *closed*, that is, does not have any free type variables. We write $(\tau : \exists \bar{X}. T)$ for the application $((\cdot : \exists \bar{X}. T)) \tau$. Since a type annotation does not affect the meaning of a program, the new destructor has identity semantics:

$$(\mathbf{v} : \exists \bar{X}. T) \xrightarrow{\delta} \mathbf{v} \quad (\text{R-ANNOTATION})$$

Its type scheme, however, is not that of the identity, namely $\forall X. X \rightarrow X$: instead, it is less general, so that annotating an expression *restricts* its type. Indeed,

we extend the initial environment Γ_0 with the binding

$$(\cdot : \exists \bar{X}. T) : \forall \bar{X}. T \rightarrow T$$

- 1.9.26 EXERCISE [★]: Check that $\forall \bar{X}. T \rightarrow T$ is an *instance* of $\forall X. X \rightarrow X$ in Damas and Milner’s sense, that is, the former is obtained from the latter via the rule DM-INST’ given in Exercise 1.2.23. Does this allow arguing that the type scheme assigned to $(\cdot : \exists \bar{X}. T)$ is sound? Check that the above definitions meet the requirements of Definition 1.7.6. \square

Although inserting a type annotation does not change the semantics of the program, it does affect constraint generation, hence type inference. We let the reader check that, assuming $\bar{X} \# ftv(t, T')$, the following derived constraint generation rule holds:

$$\text{let } \Gamma_0 \text{ in } \llbracket (t : \exists \bar{X}. T) : T' \rrbracket \equiv \text{let } \Gamma_0 \text{ in } \exists \bar{X}. (\llbracket t : T \rrbracket \wedge T \leq T')$$

So far, expressions cannot have free type variables, so the hypothesis $\bar{X} \# ftv(t)$ may seem superfluous. However, we shall soon allow expressions to contain type annotations with free type variables, so we prefer to make this condition explicit now. According to this rule, the effect of the type annotation is to force the expression t to have type T , for *some* choice of the type variables \bar{X} . As usual in type systems with subtyping, the expression’s final type T' may then be an arbitrary supertype of this particular instance of T . When subtyping is interpreted as equality, T' and T are equated by the constraint, so this constraint generation rule may be read: *a valid type for $(t : \exists \bar{X}. T)$ must be of the form T , for some choice of the type variables \bar{X} .*

- 1.9.27 EXAMPLE: In DM extended with integers, the expression $(\lambda z. z : \text{int} \rightarrow \text{int})$ has most general type $\text{int} \rightarrow \text{int}$, even though the underlying identity function has most general type $\forall X. X \rightarrow X$, so the annotation restricts its type. The expression $(\lambda z. z \hat{=} \hat{=} : \exists X. X \rightarrow X)$ has type $\text{int} \rightarrow \text{int}$, which is also the most general type of the underlying function, so the annotation acts merely as documentation in this case. Note that the type variable X is instantiated to int by the constraint solver. The expression $(\lambda z. (z \hat{=} \hat{=})) : \exists X. X \rightarrow \text{int}$ has type $(\text{int} \rightarrow \text{int}) \rightarrow \text{int}$ because the underlying function has type $(\text{int} \rightarrow Y) \rightarrow Y$, which successfully unifies with $X \rightarrow \text{int}$ by instantiating X to $\text{int} \rightarrow \text{int}$ and Y to int . Last, the expression $(\lambda z. (z \hat{=} \hat{=})) : \exists X. \text{int} \rightarrow X$ is ill-typed—even though the underlying expression is well-typed—because the equation $(\text{int} \rightarrow Y) \rightarrow Y = \text{int} \rightarrow X$ is unsatisfiable. \square

1.9.28 EXAMPLE: In DM extended with pairs, the expression $\lambda z_1.\lambda z_2.((z_1 : \exists X.X), (z_2 : \exists X.X))$ has most general type $\forall XY.X \rightarrow Y \rightarrow X \times Y$. In other words, the two occurrences of X do not represent the same type. Indeed, one could just as well have written $\lambda z_1.\lambda z_2.((z_1 : \exists X.X), (z_2 : \exists Y.Y))$. If one wishes z_1 and z_2 to receive the same type, one must lift the type annotations and merge them above the pair constructor, as follows: $\lambda z_1.\lambda z_2.((z_1, z_2) : \exists X.X \times X)$. In the process, the type constructor \times has appeared in the annotation, causing its size to increase. \square

The above example reveals a limitation of this style of type annotations: by requiring every type annotation to be closed, we lose the ability for two separate annotations to *share* a type variable. Yet, such a feature is sometimes desirable. If the two annotations where sharing is desired are distant in the code, it may be awkward to lift and merge them into a single annotation; so, more expressive power is sometimes truly needed.

Thus, we are lead to consider more general type annotations, of the form $(\tau : T)$, where T has kind \star , and where the type variables that appear within T are considered *free*, so that distinct type annotations may refer to shared type variables. For this idea to make sense, however, it is still necessary to specify where these type variables are bound. We do so using expressions of the form $\exists \bar{X}.\tau$. Such an expression binds the type variables \bar{X} within the expression τ , so that all free occurrences of X (where $X \in \bar{X}$) in type annotations inside τ stand for the same type. Thus, we break the simple type annotation construct $(\cdot : \exists \bar{X}.T)$ into two more elementary constituents, namely *existential type variable introduction* $\exists \bar{X}.\cdot$ and *type constraint* $(\cdot : T)$. Note that both are new forms of expressions; neither can be encoded by adding new constants to the calculus, because it is not possible to assign *closed* type schemes to them.

Technically, allowing expressions to contain type variables requires some care. Several constraint generation rules employ auxiliary type variables, which become bound in the generated constraint. These type variables may be chosen in an arbitrary way, provided they do not appear free in the rule's left-hand side—a side-condition intended to avoid inadvertent capture. So far, this side-condition could be read: the auxiliary type variables used to form the constraint $\llbracket \tau : T \rrbracket$ must not appear free within T . Now, since type annotations may contain free type variables, the side-condition becomes: *the auxiliary type variables used to form $\llbracket \tau : T \rrbracket$ must not appear free within τ or T .*

With this extended side-condition in mind, our original constraint generation rules remain unchanged. We add two new rules to describe how the new expression forms affect constraint generation:

$$\begin{aligned} \llbracket \exists \bar{X}.\tau : T \rrbracket &= \exists \bar{X}.\llbracket \tau : T \rrbracket && \text{provided } \bar{X} \# ftv(T) \\ \llbracket (\tau : T) : T' \rrbracket &= \llbracket \tau : T \rrbracket \wedge T \leq T' \end{aligned}$$

The effect of these rules is simple. The construct $\exists \bar{x}. \mathfrak{t}$ is an indication to the constraint generator that the type variables \bar{x} , which may occur free within type annotations inside \mathfrak{t} , should be existentially bound at this point. The side-condition $\bar{x} \# \text{ftv}(\mathsf{T})$ ensures that quantifying over \bar{x} in the generated constraint does not capture type variables in the expected type T . It can always be satisfied by α -conversion of the expression $\exists \bar{x}. \mathfrak{t}$. The construct $(\mathfrak{t} : \mathsf{T})$ is an indication to the constraint generator that the expression \mathfrak{t} should have type T , and it is treated as such by generating the subconstraint $\llbracket \mathfrak{t} : \mathsf{T} \rrbracket$. The expression's type may be an arbitrary supertype of T , hence the auxiliary constraint $\mathsf{T} \leq \mathsf{T}'$.

- 1.9.29 **EXAMPLE:** In DM extended with pairs, the expression $\lambda z_1. \lambda z_2. \exists \bar{x}. ((z_1 : \mathsf{X}), (z_2 : \mathsf{X}))$ has most general type $\forall \mathsf{X}. \mathsf{X} \rightarrow \mathsf{X} \rightarrow \mathsf{X} \times \mathsf{X}$. Indeed, the constraint generated for this expression contains the pattern $\exists \bar{x}. (\llbracket z_1 : \mathsf{X} \rrbracket \wedge \llbracket z_2 : \mathsf{X} \rrbracket \wedge \dots)$, which causes z_1 and z_2 to receive the same type. Note that this style is more flexible than that employed in Example 1.9.28, where we were forced to use a single, monolithic type annotation to express this sharing constraint. \square
- 1.9.30 **REMARK:** In practice, a type variable is usually represented as a memory cell in the typechecker's heap. So, one cannot say that the source code contains type variables; rather, it contains *names* that are meant to stand for type variables. Let us write X for such a name, and T for a type made of type constructors and names, rather than of type constructors and type variables. Then, our new expression forms are really $\exists \bar{X}. \mathfrak{t}$ and $(\mathfrak{t} : T)$. When the constraint generator enters the scope of an introduction form $\exists \bar{X}. \mathfrak{t}$, it allocates a vector of fresh type variables \bar{x} , and augments an internal environment with the bindings $\bar{X} \mapsto \bar{x}$. Because the type variables are fresh, the side-condition of the first constraint generation rule above is automatically satisfied. When the constraint generator finds a type annotation $(\mathfrak{t} : T)$, it looks up the internal environment to translate the type annotation T into an internal type T —which fails if T contains a name that is not in scope—and applies the second constraint generation rule above. \square
- 1.9.31 **EXERCISE [★★, \rightarrow]:** Let $\bar{x} \supseteq \text{ftv}(\mathsf{T})$ and $\bar{x} \# \text{ftv}(\mathfrak{t})$. Check that the constraints $\llbracket (\mathfrak{t} : \exists \bar{x}. \mathsf{T}) : \mathsf{T}' \rrbracket$ and $\llbracket \exists \bar{x}. (\mathfrak{t} : \mathsf{T}) : \mathsf{T}' \rrbracket$ are equivalent. In other words, the *local* type annotations introduced earlier may be expressed in terms of the more complex constructs described above. \square
- 1.9.32 **EXERCISE [★★, \rightarrow]:** One way of giving identity semantics to our new type annotation constructs is to *erase* them altogether prior to execution. Give an inductive definition of $\llbracket \mathfrak{t} \rrbracket$, the expression obtained by removing all type annotation constructs from the expression \mathfrak{t} . Check that $\llbracket \mathfrak{t} : \mathsf{T} \rrbracket$ entails $\llbracket \llbracket \mathfrak{t} \rrbracket : \mathsf{T} \rrbracket$ and explain why this is sufficient to ensure type soundness. \square

It is interesting to study how explicit introduction of existentially quantified type variables interacts with `let`-polymorphism. The source of their interaction lies in the difference between the constraints `let z : $\forall \bar{X}[\exists X.C_1].T$` in C_2 and `$\exists X.$ let z : $\forall \bar{X}[C_1].T$` in C_2 , which was explained in Example 1.3.28. In the former constraint, every free occurrence of `z` inside C_2 causes a copy of `$\exists X.C_1$` to be taken, thus creating its own fresh copy of `X`. In the latter constraint, on the other hand, every free occurrence of `z` inside C_2 produces a copy of C_1 . All such copies *share* references to `X`, because its quantifier was not duplicated. In the former case, one may say that the type scheme assigned to `z` is *polymorphic* with respect to `X`, while in the latter case it is *monomorphic*. As a result, the placement of type variable introduction expressions with respect to `let` bindings in the source code is meaningful: introducing a type variable *outside* of a `let` construct *prevents* it from being generalized.

- 1.9.33 **EXAMPLE:** In DM extended with integers and Booleans, the program `let f = $\exists X.$ $\lambda z.(z : X)$ in (f 0, f true)` is well-typed. Indeed, the type scheme assigned to `f` is `$\forall X.X \rightarrow X$` . However, the program `$\exists X.$ let f = $\lambda z.(z : X)$ in (f 0, f true)` is ill-typed. Indeed, the type scheme assigned to `f` is `$X \rightarrow X$` ; then, no value of `X` satisfies the constraints associated with the applications `f 0` and `f true`. The latter behavior is observed in Objective Caml, where type variables are *implicitly* introduced at the outermost level of expressions:

```
# let f z = (z:'a) in (f 0, f true);;
This expression has type bool but is here used with type int
```

More details about the treatment of type annotations in Standard ML, Objective Caml, and Haskell are given on page 113. \square

- 1.9.34 **EXERCISE [\star , \rightarrow]:** Determine which constraints are generated for the two programs in Example 1.9.33. Check that the former is indeed well-typed, while the latter is ill-typed. \square

Recursive types

We have shown that specializing $\text{HM}(X)$ with an equality-only syntactic model yields $\text{HM}(=)$, a constraint-based formulation of Damas and Milner's type system. Similarly, it is possible to specialize $\text{HM}(X)$ 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 Section 1.8 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.

Please 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 variant 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 object-oriented extension (Rémy and Vouillon, 1998), where recursive object types are commonly inferred. In order to allow recursive object types while still rejecting the above variant of `map`, Objective Caml's constraint solver implements a selective occurs check, which forbids cycles unless they involve the type constructor $\langle \cdot \rangle$ associated with objects. The corresponding model is a tree model where every infinite path down a tree must encounter the type constructor $\langle \cdot \rangle$ infinitely often.

1.10 Universal quantification in constraints

The constraint logic studied so far allows a set of variables \bar{X} to be existentially quantified within a formula C . The resulting formula $\exists \bar{X}.C$ receives its standard meaning: it requires C to hold for *some* \bar{X} . However, we currently have no way of requiring a formula C to hold for *all* \bar{X} . Is it possible to extend our

logic with universal quantification? If so, what are the new possibilities offered by this extension, in terms of type inference? The present section proposes some answers to these questions.

It is worth noting that, although the standard notation for type schemes involves the symbol \forall , type scheme introduction and instantiation constraints do not allow an encoding of universal quantification. Indeed, a universal quantifier in a type scheme is very much like an existential quantifier in a constraint: this is suggested, for instance, by Definition 1.3.3 and by C-LET \bar{E} X.

Constraints

We extend the syntax of constraints as follows:

$$C ::= \dots \mid \forall \bar{x}. C$$

Universally quantified variables are often referred to as *rigid*, while existentially quantified variables are known as *flexible*. The logical interpretation of constraints (Figure 1-5) is extended as follows:

$$\frac{\forall \bar{t} \quad \phi[\bar{x} \mapsto \bar{t}] \vdash \text{def } \Gamma \text{ in } C \quad \bar{x} \# \text{ftv}(\Gamma)}{\phi \vdash \text{def } \Gamma \text{ in } \forall \bar{x}. C} \quad (\text{CM-FORALL})$$

We let the reader check that none of the results established in Section 1.3 are affected by this addition. Furthermore, the extended constraint language enjoys the following properties.

- 1.10.1 LEMMA: $\forall \bar{x}. C \Vdash C$. Conversely, $\bar{x} \# \text{ftv}(C)$ implies $C \Vdash \forall \bar{x}. C$. □
- 1.10.2 LEMMA: $\bar{x} \# \text{ftv}(C_2)$ implies $\forall \bar{x}. (C_1 \wedge C_2) \equiv (\forall \bar{x}. C_1) \wedge C_2$. □
- 1.10.3 LEMMA: $\forall \bar{x}. \forall \bar{y}. C \equiv \forall \bar{x} \bar{y}. C$. □
- 1.10.4 LEMMA: Let $\bar{x} \# \bar{y}$. Then, $\exists \bar{x}. \forall \bar{y}. C$ entails $\forall \bar{y}. \exists \bar{x}. C$. Conversely, if $\exists \bar{y}. C$ determines \bar{x} , then $\forall \bar{y}. \exists \bar{x}. C$ entails $\exists \bar{x}. \forall \bar{y}. C$. □

Constraint solving

We briefly explain how to extend the constraint solver described in Section 1.8 with support for universal quantification. (Thus, we again assume an equality-only free tree model.) Constraint solving in the presence of equations and of existential and universal quantifiers is known as *unification under a mixed prefix*. It is a particular case of the decision problem for the first-order theory of equality on trees; see *e.g.* (Comon and Lescanne, 1989). Extending our solver is straightforward: in fact, the treatment of universal quantifiers turns

$$\begin{array}{ll}
S; U; \forall \bar{x}. C & \rightarrow S[\forall \bar{x}. \square]; U; C & \text{(S-SOLVE-ALL)} \\
& \text{if } \bar{x} \# \text{ftw}(U) \\
S[\forall \bar{x}. \exists \bar{y} \bar{z}. \square]; U; \text{true} & \rightarrow S[\exists \bar{y}. \forall \bar{x}. \exists \bar{z}. \square]; U; \text{true} & \text{(S-ALLEX)} \\
& \text{if } \bar{x} \# \bar{y} \wedge \exists \bar{x} \bar{z}. U \text{ determines } \bar{y} \\
S[\forall \bar{x} \bar{x}. \exists \bar{y}. \square]; U; \text{true} & \rightarrow \text{false} & \text{(S-ALL-FAIL-1)} \\
& \text{if } \bar{x} \notin \bar{y} \wedge \bar{x} \prec_U^* \bar{z} \wedge \bar{z} \notin \bar{x} \bar{y} \\
S[\forall \bar{x} \bar{x}. \exists \bar{y}. \square]; \bar{x} = \tau = \epsilon \wedge U; \text{true} & \rightarrow \text{false} & \text{(S-ALL-FAIL-2)} \\
& \text{if } \bar{x} \notin \bar{y} \wedge \tau \notin \mathcal{V} \\
S[\forall \bar{x}. \exists \bar{y}. \square]; U_1 \wedge U_2; \text{true} & \rightarrow S; U_1; \text{true} & \text{(S-POP-ALL)} \\
& \text{if } \bar{x} \bar{y} \# \text{ftw}(U_1) \wedge \exists \bar{y}. U_2 \equiv \text{true}
\end{array}$$

Figure 1-16: Solving universal constraints

out to be surprisingly analogous to that of `let` constraints. To begin, we extend the syntax of stacks with so-called *universal frames*:

$$S ::= \dots \mid S[\forall \bar{x}. \square]$$

Because existential quantifiers cannot, in general, be hoisted out of universal quantifiers, rules S-EX-1 to S-EX-4 now allow floating them up to the nearest enclosing `let` or *universal* frame, if any, or to the outermost level, otherwise. Thus, in our machine representation of stacks, where rules S-EX-1 to S-EX-4 are applied in an eager fashion, every universal frame carries a list of the type variables that are existentially bound immediately after it, and integer ranks count not only `let` frames, but also universal frames.

The solver's specification is extended with the rules in Figure 1-16. S-SOLVE-ALL, a forward rule, discovers a universal constraint and enters it, creating a new universal frame to record its existence. S-ALLEX exploits Lemma 1.10.4 to hoist existential quantifiers out of the universal frame. It is analogous to S-LETALL, and its implementation may rely upon the same procedure (Exercise 1.8.8). The next two rules detect failure conditions. S-ALL-FAIL-1 states that the constraint $\forall \bar{x}. \exists \bar{y}. U$ is false if the rigid variable \bar{x} is directly or indirectly *dominated* by a *free* variable \bar{z} . Indeed, the value of \bar{x} is then determined by that of \bar{z} —but a universally quantified variable ranges over *all* values, so this is a contradiction. In such a case, \bar{x} is commonly said to *escape its scope*. S-ALL-FAIL-2 states that the same constraint is false if \bar{x} is equated with a *nonvariable* term. Indeed, the value of \bar{x} is then

partially determined, since its head constructor is known, which again contradicts its universal status. Last, S-POP-ALL splits the current unification constraint into two components U_1 and U_2 , where U_1 is made up entirely of *old* variables and U_2 constrains *young* variables only. This decomposition is analogous to that performed by S-POP-LET. Then, it is not difficult to check that $\forall \bar{X}. \exists \bar{Y}. (U_1 \wedge U_2)$ is equivalent to U_1 . So, the universal frame, as well as U_2 , are discarded, and the solver proceeds by examining whatever remains on top of the stack S .

It is possible to further extend the treatment of universal frames with two rules analogous to S-COMPRESS and S-UNNAME. In practice, this improves the solver's efficiency, and makes it easier to share code between the treatment of *let* frames and that of universal frames.

It is interesting to remark that, as far as the underlying unification algorithm is concerned, there is no difference between existentially and universally quantified type variables. The algorithm solves whatever equations are presented to it, without inquiring about the status of their variables. Equations that lead to failure, because a rigid variable escapes its scope or is equated with a nonvariable term, are detected only when the universal frame is exited. A perhaps more common approach is to *mark* rigid variables as such, allowing the unification algorithm to signal failure as soon as one of the two error conditions is encountered. In this approach, a rigid variable may successfully unify only with itself or with flexible variables fresher than itself. It is often called a *Skolem constructor* in the literature (Läufer and Odersky, 1994; Shields and Peyton Jones, 2002). An interesting variant of this approach appears in Dowek, Hardin, Kirchner and Pfenning's treatment of (higher-order) unification (1995; 1998), where flexible variables are represented as ordinary variables, while rigid variables are encoded using De Bruijn indices.

The properties of our constraint solver are preserved by this extension: it is possible to prove that Lemmas 1.8.9, 1.8.10, and 1.8.11 remain valid.

Type annotations, continued

In Section 1.9, we introduced the expression form $(\mathfrak{t} : \exists \bar{X}. T)$, allowing an expression \mathfrak{t} to be annotated with a type T whose free variables \bar{X} are *locally* and *existentially* bound. It is now natural to introduce the symmetric expression form $(\mathfrak{t} : \forall \bar{X}. T)$, where T has kind \star , \bar{X} is bound within T , and \bar{X} contains *ftv*(T), as before. Its constraint generation rule is as follows:

$$\llbracket (\mathfrak{t} : \forall \bar{X}. T) : T' \rrbracket = \forall \bar{X}. \llbracket \mathfrak{t} : T \rrbracket \wedge \exists \bar{X}. (T \leq T') \quad \text{provided } \bar{X} \# \text{ftv}(\mathfrak{t}, T')$$

The first conjunct requires \mathfrak{t} to have type T for *all* values of \bar{X} . Here, the type variables \bar{X} are *universally* bound, as expected. The second conjunct requires

T' to be *some* instance of the universal annotation $\forall \bar{X}.T$. Since T' is only a monotype, it seems difficult to think of another sensible way of constraining T' . For this reason, the type variables \bar{X} are still *existentially* bound in the second conjunct. This makes the interpretation of the universal quantifier in type annotations a bit more complex than that of the existential quantifier. For instance, when subtyping is interpreted as equality, the constraint generation rule may be read: *a valid type for $(\tau : \forall \bar{X}.T)$ is of the form T , for some choice of the type variables \bar{X} , provided τ has type T for all choices of \bar{X} .*

We remark that $(\tau : \forall \bar{X}.T)$ must be a new expression form: it cannot be encoded by adding new constants to the calculus—whereas $(\tau : \exists \bar{X}.T)$ could—because none of the existing constraint generation rules produce universally quantified constraints. Like all type annotations, it has identity semantics.

What is the use of universal type annotations, compared with existential type annotations? When a type variable is existentially bound, the typechecker is free to assign it whatever value makes the program well-typed. As a result, the expressions $(\lambda z.z \hat{+} \hat{1} : \exists X.X \rightarrow X)$ and $(\lambda z.z : \exists X.X \rightarrow X)$ are both well-typed: X is assigned `int` in the former case, and remains undetermined in the latter. However, it is sometimes useful to be able to insist that an expression should be polymorphic. This effect is naturally achieved by using a universally bound type variable. Indeed, $(\lambda z.z \hat{+} \hat{1} : \forall X.X \rightarrow X)$ is ill-typed, because $\forall X.(X = \text{int})$ is false, while $(\lambda z.z : \forall X.X \rightarrow X)$ is well-typed.

1.10.5 EXERCISE [★]: Write down the constraints $\exists Z. \llbracket (\lambda z.z \hat{+} \hat{1} : \forall X.X \rightarrow X) : Z \rrbracket$ and $\exists Z. \llbracket (\lambda z.z : \forall X.X \rightarrow X) : Z \rrbracket$, which tell whether these expressions are well-typed. Check that the former is false, while the latter is satisfiable. \square

A universal type annotation, as defined above, is nothing but a (closed) Damas-Milner type scheme. Thus, the new construct $(\tau : \forall \bar{X}.T)$ gives us the ability to ensure that the expression τ admits the type scheme $\forall \bar{X}.T$. This feature is exploited at the module level in ML-the-programming-language, where it is necessary to check that the inferred type for a module component τ is more general than the type scheme S that appears in the module's signature. In our view, this process simply consists in ensuring that $(\tau : S)$ is well-typed.

In Section 1.9, we have pointed out that local (that is, closed) type annotations offer limited expressiveness, because they cannot share type variables. To lift this limitation, we have introduced the expression forms $\exists \bar{X}.\tau$ and $(\tau : T)$. The former binds the type variables \bar{X} within τ , making them available for use in type annotations, and instructs the constraint generator to existentially quantify them at this point. The latter requires τ to have T . It is natural to proceed in the same manner in the case of universal type annotations. We now introduce the expression form $\forall \bar{X}.\tau$, which also binds \bar{X} within τ , but comes

with a different constraint generation rule:

$$\llbracket \forall \bar{X}. \mathfrak{t} : \mathsf{T} \rrbracket = \forall \bar{X}. \exists Z. \llbracket \mathfrak{t} : Z \rrbracket \wedge \exists \bar{X}. \llbracket \mathfrak{t} : \mathsf{T} \rrbracket \quad \text{provided } \bar{X} \# \text{ftv}(\mathsf{T}) \wedge Z \notin \text{ftv}(\mathfrak{t})$$

This rule is a bit more complex than that associated with the expression form $\exists \bar{X}. \mathfrak{t}$. Again, this is due to the fact that we do not wish to overconstrain T . The first exercise below shows that a more naïve version of the rule does not yield the desired behavior. The second exercise shows that this version does. The third exercise clarifies an efficiency concern.

- 1.10.6 EXERCISE [★]: Assume that $\llbracket \forall \bar{X}. \mathfrak{t} : \mathsf{T} \rrbracket$ is defined as $\forall \bar{X}. \llbracket \mathfrak{t} : \mathsf{T} \rrbracket$, provided $\bar{X} \# \text{ftv}(\mathsf{T})$. Write down the constraint $\llbracket \forall \bar{X}. (\lambda z. z : \bar{X} \rightarrow \bar{X}) : Z \rrbracket$. Can you describe its solutions? Does it have the intended meaning? \square
- 1.10.7 EXERCISE [★★]: Let $\bar{X} \supseteq \text{ftv}(\mathsf{T})$ and $\bar{X} \# \text{ftv}(\mathfrak{t})$. Check that the constraints $\llbracket (\mathfrak{t} : \forall \bar{X}. \mathsf{T}) : \mathsf{T}' \rrbracket$ and $\llbracket \forall \bar{X}. (\mathfrak{t} : \mathsf{T}) : \mathsf{T}' \rrbracket$ are equivalent. In other words, *local* universal type annotations may also be expressed in terms of the more complex constructs described above. \square
- 1.10.8 EXERCISE [★★★★, \rightarrow]: The constraint generation rule that appears above compromises the linear time and space complexity of constraint generation, because it duplicates the term \mathfrak{t} . It is possible to avoid this problem, but this requires a slight generalization of the constraint language. Let us write let $\mathfrak{x} : \forall \underline{\bar{X}} \bar{Y}[C_1]. \mathsf{T}$ in C_2 for $\forall \bar{X}. \exists \bar{Y}. C_1 \wedge \text{def } \mathfrak{x} : \forall \underline{\bar{X}} \bar{Y}[C_1]. \mathsf{T}$ in C_2 . In this extended let form, the underlined variables \bar{X} are interpreted as *rigid*, instead of *flexible*, while checking that C_1 is satisfiable. However, the type scheme associated with \mathfrak{x} is not affected. *Check that the above constraint generation rule may now be written as follows:*

$$\llbracket \forall \bar{X}. \mathfrak{t} : \mathsf{T} \rrbracket = \text{let } \mathfrak{x} : \forall \underline{\bar{X}} Z[\llbracket \mathfrak{t} : Z \rrbracket]. Z \text{ in } \mathfrak{x} \preceq \mathsf{T} \quad \text{provided } Z \notin \text{ftv}(\mathfrak{t})$$

Roughly speaking, the new rule forms a most general type scheme for \mathfrak{t} , ensures that the type variables \bar{X} are unconstrained in it, and checks that T is an instance of it. Furthermore, it does not duplicate \mathfrak{t} . To complete the exercise, *extend the specification of the constraint solver (Figures 1-12 and 1-16), as well as its implementation, to deal with this extension of the constraint language.* \square

To conclude, let us once again stress that, if T has free type variables, the effect of the type annotation $(\mathfrak{t} : \mathsf{T})$ depends on *how* and *where* they are bound. The effect of *how* stems from the fact that binding a type variable universally, rather than existentially, leads to a stricter constraint. Indeed, we let the reader check that $\llbracket \forall \bar{X}. \mathfrak{t} : \mathsf{T} \rrbracket$ entails $\llbracket \exists \bar{X}. \mathfrak{t} : \mathsf{T} \rrbracket$, while the converse

does not hold in general. The effect of *where* has been illustrated, in the case of existentially bound type variables, in Section 1.9. It is due, in that case, to the fact that `let` and \exists do not commute. In the case of universally bound type variables, it may be imputed to the fact that \forall and \exists do not commute. For instance, $\lambda z.\forall X.(z : X)$ is ill-typed, because *inside the λ -abstraction*, the program variable `z` cannot be said to have every type. However, $\forall X.\lambda z.(z : X)$ is well-typed, because the identity function does have type $X \rightarrow X$ for every X .

- 1.10.9 EXERCISE [★]: Write down the constraints $\exists Z. \llbracket \lambda z.\forall X.(z : X) : Z \rrbracket$ and $\exists Z. \llbracket \forall X.\lambda z.(z : X) : Z \rrbracket$, which tell whether these expressions are well-typed. Is the former satisfiable? Is the latter? \square

In Standard ML and Objective Caml, the type variables that appear in type annotations are *implicitly* bound. That is, there is no syntax in the language for the constructs $\exists \bar{x}.\tau$ and $\forall \bar{x}.\tau$. When a type annotation $(\tau : T)$ contains a free type variable X , a fixed convention tells how and where X is bound. In Standard ML, X is *universally* bound at the nearest `val` binding that encloses all related occurrences of X (Milner, Tofte, and Harper, 1990). The 1997 revision of Standard ML (Milner, Tofte, Harper, and MacQueen, 1997b) slightly improves on this situation by allowing type variables to be *explicitly* introduced at `val` bindings. However, they still must be universally bound. In Objective Caml, X is *existentially* bound at the nearest enclosing toplevel `let` binding; this behavior seems to be presently undocumented. We argue that (i) allowing type variables to be implicitly introduced is confusing; and (ii) for expressiveness, both universal and existential quantifiers should be made available to programmers. Surprisingly, these language design and type inference issues seem to have received little attention in the literature, although they have most likely been “folklore” for a long time. Peyton Jones and Shields (2003) study these issues in the context of Haskell, and concur with (i). Concerning (ii), they seem to think that the language designer must choose between existential and universal type variable introduction forms—which they refer to as “type-sharing” and “type-lambda”—whereas we point out that they may and should coexist.

Polymorphic recursion

Example 1.2.10 explains how the `letrec` construct found in ML-the-programming-language may be viewed as an application of the constant `fix`, wrapped inside a normal `let` construct. Exercise 1.9.6 shows that this gives rise to a somewhat restrictive constraint generation rule: generalization occurs only *after* the application of `fix` is typechecked. In other words, in `letrec f = $\lambda z.\tau_1$ in τ_2` , all occurrences of `f` within τ_1 must have the same

(monomorphic) *type*. This restriction is sometimes a nuisance, and seems unwarranted: if the function that is being defined is polymorphic, it should be possible to use it at different types even inside its own definition. Indeed, Mycroft (1984) extended Damas and Milner’s type system with a more liberal treatment of recursion, commonly known as *polymorphic recursion*. The idea is to only request occurrences of f within τ_1 to have the same *type scheme*. Hence, they may have different *types*, all of which are instances of a common type scheme. It was later shown that well-typedness in Mycroft’s extended type system is undecidable (Henglein, 1993; Kfoury, Tiuryn, and Urzyczyn, 1993). To work around this stumbling block, one solution is to use a semi-algorithm, falling back to monomorphic recursion if it does not succeed or fail in reasonable time. Although such a solution might be appealing in the setting of an automated program analysis, it is less so in the setting of a programmer-visible type system, because it may become difficult to understand why a program is ill-typed. Thus, we describe a simpler solution, which consists in requiring the programmer to explicitly supply a type scheme for f . This is an instance of a *mandatory* type annotation.

To begin, we must change the status of `fix`, because if `fix` remains a constant, then f must remain λ -bound and cannot receive a polymorphic type scheme. We turn `fix` into a language construct, which binds a program variable f , and annotates it with a DM type scheme. The syntax of values and expressions is thus extended as follows:

$$v ::= \dots \mid \text{fix } f : S.\lambda z.t \quad t ::= \dots \mid \text{fix } f : S.\lambda z.t$$

Please note that f is bound within $\lambda z.t$. The operational semantics is extended as follows.

$$(\text{fix } f : S.\lambda z.t) v \longrightarrow (\text{let } f = \text{fix } f : S.\lambda z.t \text{ in } \lambda z.t) v \quad (\text{R-FIX}')$$

The type annotation S plays no essential role in the reduction; it is merely preserved. It is now possible to define `letrec` $f : S = \lambda z.t_1$ `in` t_2 as syntactic sugar for `let` $f = \text{fix } f : S.\lambda z.t_1$ `in` t_2 .

We now give a constraint generation rule for `fix`:

$$\llbracket \text{fix } f : S.\lambda z.t : T \rrbracket = \text{let } f : S \text{ in } \llbracket \lambda z.t : S \rrbracket \wedge S \preceq T$$

The left-hand conjunct requires the function $\lambda z.t$ to have type scheme S , under the assumption that f has type S . Thus, it is now possible for different occurrences of f within t to receive different types. If S is $\forall \bar{X}.T$, where $\bar{X} \# ftv(t)$, then we write $\llbracket t : S \rrbracket$ for $\forall \bar{X}.\llbracket t : T \rrbracket$. Indeed, checking the validity of a polymorphic type annotation—be it mandatory, as is the case here, or optional, as was previously the case—requires a universally quantified constraint. The right-hand conjunct merely constrains T to be an instance of S .

Given the definition of `letrec f : S = λz.t1 in t2` as syntactic sugar, the above rule leads to the following derived constraint generation rule for `letrec`:

$$\llbracket \text{letrec } f : S = \lambda z.t_1 \text{ in } t_2 : T \rrbracket = \text{let } f : S \text{ in } (\llbracket \lambda z.t_1 : S \rrbracket \wedge \llbracket t_2 : T \rrbracket)$$

This rule is arguably quite natural. The program variable `f` is assigned the type scheme `S` throughout its scope, that is, both inside and outside of the function's definition. The function `λz.t1` must itself have type scheme `S`. Last, `t2` must have type `T`, as in every `let` construct.

- 1.10.10 EXERCISE [★★]: Prove that the derived constraint generation rule above is indeed valid. □

It is straightforward to prove that the extended language still enjoys subject reduction. The proof relies on the following lemma: if `t` has type scheme `S`, then every instance of `S` is also a valid type for `t`.

- 1.10.11 LEMMA: $\llbracket t : S \rrbracket \wedge S \preceq T \Vdash \llbracket t : T \rrbracket$. □

- 1.10.12 THEOREM [SUBJECT REDUCTION]: (R-FIX') $\subseteq (\sqsubseteq)$. □

The programming language Haskell (Hudak, Peyton Jones, Wadler, Boutel, Fairbairn, Fasel, Guzman, Hammond, Hughes, Johnsson, Kieburtz, Nikhil, Partain, and Peterson, 1992) offers polymorphic recursion. Interesting details about its typing rules may be found in (Jones, 1999).

It is worth pointing out that some restricted instances of type inference in the presence of polymorphic recursion are decidable. This is typically the case in certain program analyses, where a type derivation for the program is already available, and the goal is only to infer extra atomic annotations, such as binding time or strictness properties. Several papers that exploit this idea are (Dussart, Henglein, and Mossin, 1995a; Jensen, 1998; Rehof and Fähndrich, 2001).

Universal types

ML-the-type-system enforces a strict stratification between types and type schemes, or, in other words, allows only prenex universal quantifiers inside types. We have pointed out earlier that there is good reason to do so: type inference for ML-the-type-system is decidable, while type inference for System F, which has no such restriction, is undecidable. Yet, this restriction comes at a cost in expressiveness: it prevents higher-order functions from accepting polymorphic function arguments, and forbids storing polymorphic functions inside data structures. Fortunately, it is in fact possible to circumvent the problem by requiring the programmer to supply additional type information.

The approach that we are about to describe is reminiscent of the way algebraic data type definitions allow circumventing the problems associated with equirecursive types (Section 1.9). Because we do not wish to extend the syntax of types with universal types of the form $\forall \bar{Y}.T$, we instead allow *universal type definitions*, of the form

$$D \vec{X} \approx \forall \bar{Y}.T$$

where D still ranges over data types. If D has signature $\vec{\kappa} \Rightarrow \star$, then the type variables \vec{X} must have kind $\vec{\kappa}$. The type T must have kind \star . The type variables \bar{X} and \bar{Y} are considered bound within T , and the definition must be closed, that is, $ftv(T) \subseteq \bar{X}\bar{Y}$ must hold. Last, the variance of the type constructor D must match its definition—a requirement stated as follows:

1.10.13 DEFINITION: Let $D \vec{X} \approx \forall \bar{Y}.T$ and $D \vec{X}' \approx \forall \bar{Y}'.T'$ be two α -equivalent instances of a single universal type definition, such that $\bar{Y} \# ftv(T')$ and $\bar{Y}' \# ftv(T)$. Then, $D \vec{X} \leq D \vec{X}' \Vdash \forall \bar{Y}'.\exists \bar{Y}.T \leq T'$ must hold. \square

This requirement is analogous to that found in Definition 1.9.8. The idea is, if $D \vec{X}$ and $D \vec{X}'$ are comparable, then their unfoldings $\forall \bar{Y}.T$ and $\forall \bar{Y}'.T'$ should be comparable as well. The comparison between them is expressed by the constraint $\forall \bar{Y}'.\exists \bar{Y}.T \leq T'$, which may be read: *every instance of $\forall \bar{Y}'.T'$ is (a supertype of) an instance of $\forall \bar{Y}.T$* . Again, when subtyping is interpreted as equality, the requirement of Definition 1.10.13 is always satisfied; it becomes nontrivial only in the presence of true subtyping.

The effect of the universal type definition $D \vec{X} \approx \forall \bar{Y}.T$ is to enrich the programming language with a new construct:

$$v ::= \dots \mid \text{pack}_{\mathfrak{D}} v \quad t ::= \dots \mid \text{pack}_{\mathfrak{D}} t \quad \mathcal{E} ::= \dots \mid \text{pack}_{\mathfrak{D}} \mathcal{E}$$

and with a new unary destructor $\text{open}_{\mathfrak{D}}$. Their operational semantics is as follows:

$$\text{open}_{\mathfrak{D}} (\text{pack}_{\mathfrak{D}} v) \xrightarrow{\delta} v \quad (\text{R-OPEN-ALL})$$

Intuitively, $\text{pack}_{\mathfrak{D}}$ and $\text{open}_{\mathfrak{D}}$ are the two coercions that witness the isomorphism between $D \vec{X}$ and $\forall \bar{Y}.T$. The value $\text{pack}_{\mathfrak{D}} v$ behaves exactly like v , except it is marked, as a hint to the typechecker. As a result, the mark must be removed using $\text{open}_{\mathfrak{D}}$ before the value can be used.

What are the typing rules for $\text{pack}_{\mathfrak{D}}$ and $\text{open}_{\mathfrak{D}}$? In System F, they would receive types $\forall \bar{X}.(\forall \bar{Y}.T) \rightarrow D \vec{X}$ and $\forall \bar{X}.D \vec{X} \rightarrow \forall \bar{Y}.T$, respectively. However, neither of these is a valid type scheme: both exhibit a universal quantifier under an arrow.

In the case of $\text{pack}_{\mathfrak{D}}$, which has been made a language construct rather than a constant, we work around the problem by embedding this universal

quantifier in the constraint generation rule:

$$\llbracket \text{pack}_D \mathfrak{t} : T' \rrbracket = \exists \bar{X}. (\llbracket \mathfrak{t} : \forall \bar{Y}. T \rrbracket \wedge D \bar{X} \leq T')$$

The rule implicitly requires that \bar{X} be fresh for the left-hand side and that $D \bar{X} \approx \forall \bar{Y}. T$ be (an α -variant of) the definition of D . The left-hand conjunct requires \mathfrak{t} to have type scheme $\forall \bar{Y}. T$. The notation $\llbracket \mathfrak{t} : S \rrbracket$ was defined on page 114. The right-hand conjunct states that a valid type for $\text{pack}_D \mathfrak{t}$ is (a supertype of) $D \bar{X}$.

We deal with open_D as follows. Provided $\bar{X} \# \bar{Y}$, we extend the initial environment Γ_0 with the binding $\text{open}_D : \forall \bar{X} \bar{Y}. D \bar{X} \rightarrow T$. We have simply hoisted the universal quantifier outside of the arrow—a valid isomorphism in System F.

The proof of the subject reduction theorem must be extended with the following new case:

1.10.14 THEOREM [SUBJECT REDUCTION]: (R-OPEN-ALL) \subseteq (\sqsubseteq). □

Proof: We have

$$\begin{aligned} & \text{let } \Gamma_0 \text{ in } \llbracket \text{open}_D (\text{pack}_D v) : T_0 \rrbracket \\ \equiv & \text{let } \Gamma_0 \text{ in } \exists Z. (\text{open}_D \leq Z \rightarrow T_0 \wedge \llbracket \text{pack}_D v : Z \rrbracket) & (1) \\ \equiv & \text{let } \Gamma_0 \text{ in } \exists Z. (\exists \bar{X}' \bar{Y}'. (D \bar{X}' \rightarrow T' \leq Z \rightarrow T_0) \wedge \exists \bar{X}. (\llbracket v : \forall \bar{Y}. T \rrbracket \wedge D \bar{X} \leq Z)) & (2) \\ \equiv & \text{let } \Gamma_0 \text{ in } \exists \bar{X} \bar{X}' \bar{Y}'. (\llbracket v : \forall \bar{Y}. T \rrbracket \wedge D \bar{X} \leq D \bar{X}' \wedge T' \leq T_0) & (3) \\ \Vdash & \text{let } \Gamma_0 \text{ in } \exists \bar{X} \bar{Y} \bar{X}' \bar{Y}'. (\llbracket v : \forall \bar{Y}. T \rrbracket \wedge T \leq T' \wedge T' \leq T_0) & (4) \\ \Vdash & \text{let } \Gamma_0 \text{ in } \exists \bar{X} \bar{Y} \bar{X}' \bar{Y}'. \llbracket v : T_0 \rrbracket & (5) \\ \equiv & \text{let } \Gamma_0 \text{ in } \llbracket v : T_0 \rrbracket & (6) \end{aligned}$$

where (1) is by definition of constraint generation for applications and for constants; Z is fresh; (2) is by definition of constraint generation for pack_D and open_D , where $D \bar{X} \approx \forall \bar{Y}. T$ and $D \bar{X}' \approx \forall \bar{Y}'. T'$ are two α -equivalent instances of the definition of D ; \bar{X} , \bar{Y} , \bar{X}' , and \bar{Y}' are fresh and satisfy $\bar{Y} \# \text{ftv}(T')$ and $\bar{Y}' \# \text{ftv}(T)$; (3) is by C-EXAND, C-ARROW, and C-EXTRANS, which allows eliminating Z ; (4) is by Definition 1.10.13, Lemma 1.10.1, and C-EXAND; (5) is by Lemmas 1.10.11 and 1.6.3; (6) is by C-EX*. □

The proof of (R-CONTEXT) \subseteq (\sqsubseteq) must also be extended with a new subcase, corresponding the new production $\mathcal{E} ::= \dots \mid \text{pack}_D \mathcal{E}$. If the language is pure, this is straightforward. In the presence of side effects, however, this subcase fails, because universal and existential quantifiers in constraints do not commute. The problem is then avoided by restricting pack_D to values, as in Definition 1.7.7.

This approach to extending ML-the-type-system with universal (or existential—see below) types has been studied in (Läufer and Odersky, 1994;

Rémy, 1994; Odersky and Läufer, 1996; Shields and Peyton Jones, 2002). Läufer and Odersky have suggested combining universal or existential type declarations with algebraic data type definitions. This allows suppressing the cumbersome pack_D and open_D constructs; instead, one simply uses the standard syntax for constructing and deconstructing variants and records.

Existential types

Existential types (TAPL Chapter 24) are close cousins of universal types, and may be introduced into ML-the-type-system in the same manner. Actually, existential types have been introduced in ML-the-type-system before universal types. We give a brief description of this extension, insisting mainly on the differences with the case of universal types.

We now allow *existential type definitions*, of the form $D\vec{X} \approx \exists \bar{Y}.T$. The conditions required of a well-formed definition are unchanged, except the variance requirement, which is dual:

1.10.15 DEFINITION: Let $D\vec{X} \approx \exists \bar{Y}.T$ and $D\vec{X}' \approx \exists \bar{Y}'.T'$ be two α -equivalent instances of a single existential type definition, such that $\bar{Y} \# \text{ftv}(T')$ and $\bar{Y}' \# \text{ftv}(T)$. Then, $D\vec{X} \leq D\vec{X}' \Vdash \forall \bar{Y}. \exists \bar{Y}'. T \leq T'$ must hold. \square

The effect of this existential type definition is to enrich the programming language with a new unary constructor pack_D and with a new construct: $t ::= \dots \mid \text{open}_D t$ and $\mathcal{E} ::= \dots \mid \text{open}_D \mathcal{E} t \mid \text{open}_D v \mathcal{E}$. Their operational semantics is as follows:

$$\text{open}_D (\text{pack}_D v_1) v_2 \longrightarrow v_2 v_1 \quad (\text{R-OPEN-EX})$$

In the literature, the second argument of open_D is often required to be a λ -abstraction $\lambda z.t$, so the construct becomes $\text{open}_D t (\lambda z.t)$, often written $\text{open}_D t$ as z in t .

Provided $\bar{X} \# \bar{Y}$, we extend the initial environment Γ_0 with the binding $\text{pack}_D : \forall \bar{X}\bar{Y}.T \rightarrow D\vec{X}$. The constraint generation rule for open_D is as follows:

$$\llbracket \text{open}_D t_1 t_2 : T' \rrbracket = \exists \bar{X}. (\llbracket t_1 : D\vec{X} \rrbracket \wedge \llbracket t_2 : \forall \bar{Y}.T \rightarrow T' \rrbracket)$$

The rule implicitly requires that \bar{X} be fresh for the left-hand side, that \bar{Y} be fresh for T' , and that $D\vec{X} \approx \forall \bar{Y}.T$ be (an α -variant of) the definition of D . The left-hand conjunct simply requires t_1 to have type $D\vec{X}$. The right-hand conjunct states that the function t_2 must be prepared to accept an argument of type T , for *any* \bar{Y} , and produce a result of the expected type T' . In other words, t_2 must be a polymorphic function.

The type scheme of existential pack_D resembles that of universal open_D , while the constraint generation rule for existential open_D is a close cousin

of that for universal pack_D . Thus, the duality between universal and existential types is rather strong. The main difference lies in the fact that the existential open_D construct is *binary*, rather than unary, so as to limit the scope of the newly introduced type variables \bar{Y} . The duality may be better understood by studying the encoding of existential types in terms of universal types (Reynolds, 1983b).

As expected, R-OPEN-EX preserves types.

1.10.16 THEOREM [SUBJECT REDUCTION]: (R-OPEN-EX) \subseteq (\sqsubseteq). □

1.10.17 EXERCISE [★★, \rightarrow]: Prove Theorem 1.10.16. The proof is analogous, although not identical, to that of Theorem 1.10.14. □

In the presence of side effects, the new production $\mathcal{E} ::= \dots \mid \text{open}_D v \mathcal{E}$ is problematic. The standard workaround is to restrict the second argument to open_D to be a value.

1.11 Rows

In Section 1.9, 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 $D\bar{X}$, then the associated record access operation has type $\forall\bar{X}.D\bar{X} \rightarrow T$. If ℓ appears with type T' in the definition of another simple record type, say $D'\bar{X}'$, then the associated record access operation has type $\forall\bar{X}'.D'\bar{X}' \rightarrow T'$; 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 $D\bar{X}$, 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.

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 $\{\ell_a, \ell_b, \ell_c\}$.

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 *record* of signature $\star \otimes \star \otimes \star \Rightarrow \star$. The type record $T_a T_b T_c$ represents all records where the field ℓ_a (resp. ℓ_b , ℓ_c) contains a value of type T_a (resp. T_b , T_c). Please note that *record* 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:

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

Here, polymorphism allows updating or accessing a field without knowledge of the types of the other fields. This flexibility is made possible by the property that all record types are formed using a single record 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 naïve approach consists in 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 $\text{pre } v$, 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 record (option T_a) (option T_b) (option T_c). This naïve 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 `case_option`.

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 record 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 $\lambda x. \lambda z. \{\mathbf{r} \text{ with } \ell_b = \text{pre } z\}$. *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:

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

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 \mathbf{X}_b , which means that it is insensitive to the presence or absence of the field ℓ_b in its argument. The subterm $\text{pre } \mathbf{X}'_b$ in its result type reflects the fact that ℓ_b is defined in the extended record. Conversely, the subterm $\text{pre } \mathbf{X}_b$ 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.

We remark that, in our encoding, the arguments of the type constructor `record` 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 \diamond of *field types*. Then, let us adopt the following signatures: `pre`: $\star \Rightarrow \diamond$, `abs`: \diamond , and `record`: $\diamond \otimes \diamond \otimes \diamond \Rightarrow \star$.

- 1.11.1 EXERCISE [\star , RECOMMENDED, \Rightarrow]: Check that the three type schemes given above are well-kinded. What is the kind of each type variable? \square
- 1.11.2 EXERCISE [$\star\star$, RECOMMENDED, \Rightarrow]: Our record 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 $\{\ell_a : T_a\}$; a record that defines the fields ℓ_b and ℓ_c only would have a type of the form $\{\ell_b : T_b; \ell_c : T_c\}$; 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

Finite records are insufficient both from practical and theoretical points of view. In practice, the set of labels could become very large, making the type of every record as large as the set of labels itself, even if only a few labels are

actually defined. In principle, the set of labels could even be infinite. Actually, in modular programs the whole set of labels may not be known in advance, which amounts in some way to working with an infinite set of labels. Thus, records must be drawn from an infinite set of labels—whether their domains are finite or infinite. Still, we can restrict our attention to records that are almost constant, that is, records where only a finite number of fields differ. With this restriction, full records (defined everywhere) can always be built by giving explicit definitions for a finite number of fields and a default value for all other fields, as in the finite case. For instance, the record $\{\{\{\text{false}\}$ with $\ell = 1\}$ with $\ell' = \text{true}\}$ is the record equal to `true` on field ℓ' , to 1 on field ℓ , and to `false` on any other field.

Types of records are functions from labels to types, called *rows*. However, for sake of generality, we use a unary type constructor, say Π , as an indirection between rows and record types. Moreover, we further restrict our attention to the case where rows are also almost constant. (The fact that the property holds for record values does not imply that it also holds for record types, for the default value of some record could have a polymorphic type, and one could wish to see each field with a different instance of this polymorphic type. So this is a true restriction, but a reasonable one.) Thus, rows can also be represented by giving explicit types for a finite number of fields and a default type for all other fields. We write ∂T the row whose type is T on every field, and $(\ell : T ; T')$ the row whose type is T on field ℓ and T' on other fields. Formally, ∂ is a unary type constructor and ℓ is a family of binary type constructors, written with syntactic sugar $(\ell : \cdot ; \cdot)$. For example, $\Pi(\ell : \text{bool} ; (\ell' : \text{int} ; \partial \text{bool}))$ is a record type that describes records whose field ℓ carries a value of type `bool`, field ℓ' carries a value of type `int`, and all other fields carry values of type `bool`. In fact, this is a sound type for the record defined above. In fact, the type $\Pi(\ell' : \text{int} ; (\ell : \text{bool} ; \partial \text{bool}))$ should also be a sound type for this record, since the order in which fields are specified should not matter. We actually treat both types as equivalent. Furthermore, the row $(\ell : \text{bool} ; \partial \text{bool})$, which stands for `bool` on field ℓ and ∂bool everywhere else, must also be equivalent to ∂bool , which stands for `bool` everywhere.

A record type may also contain type variables. For instance, the record $\lambda z.\{z\}$ that maps any value v to a record with the default value v has type $X \rightarrow \Pi(\partial X)$. Projections of this record on any field will return a value of the same type X . By comparison, the function that reads some field ℓ of its (record) argument has type $\Pi(\ell : X ; Y) \rightarrow X$: this says that the argument must be a record where field ℓ has type X and other fields may have *any* type. Variable Y is called a *row variable*, since it can be instantiated to any row. For instance, Y can be instantiated to $(\ell' : \text{int} ; Y')$ and as a result this function can be applied to the record above. Conversely, the row ∂X , which is equal to $(\ell' : X ; X)$, can

only be instantiated to rows of the form ∂T , which are equal to $(\ell' : T ; T)$, that is, to constant rows.

Syntax of row types

Let \mathcal{L} be a denumerable collection of labels. We write $\ell.L$ for $\{\ell\} \uplus L$, which implies $\ell \notin L$. We first introduce kinds, so as to distinguish rows such as $(\ell : \text{int} ; \partial \text{bool})$ from *basic types*, such as int or $\text{int} \rightarrow \text{int}$.

- 1.11.3 DEFINITION [ROW KINDS]: Let *row kinds* be composed of a particular kind *Type* and the collection of kinds $\text{Row}(L)$ where L ranges over finite subsets of \mathcal{L} . We use letter s to range over row kinds. \square

Intuitively, a row of kind $\text{Row}(L)$ is a function of domain $\mathcal{L} \setminus L$ to types. That is, L specifies the set of labels that the row *must not* define. For instance, the (basic) type $\Pi(\ell : \text{int} ; \mathbf{X})$ has kind *Type*, the row $(\ell : \text{int} ; \mathbf{X})$ has kind $\text{Row}(\emptyset)$ provided \mathbf{X} has kind $\text{Row}(\{\ell\})$.

To remain abstract the definition of rows is parameterized by a signature \mathcal{S}_0 for building basic types and a signature \mathcal{S}_1 for building rows. From those two signatures, we then define a new signature \mathcal{S} that completely specifies the set of types. However, the signature \mathcal{S} must superimposed row kinds on top of the (basic) kinds of the two input signatures \mathcal{S}_0 and \mathcal{S}_1 . We use product signatures to enlighten this process. More precisely, we build a product signature from two signatures $K \Rightarrow \kappa$ and $K' \Rightarrow \kappa'$ with the following notations: we write $\kappa.\kappa'$ for the pair (κ, κ') ; $K.\kappa$ for the mapping $(d \mapsto K(d).\kappa)^{d \in \text{dom}(K)}$; $(K \Rightarrow \kappa).\kappa'$ for the kind signature $K.\kappa \Rightarrow \kappa.\kappa'$; and symmetrically, we write $\kappa.K'$ and $\kappa.(K' \Rightarrow \kappa')$. The signature \mathcal{S} reuses the same input type constructors as \mathcal{S}_0 and \mathcal{S}_1 , but at different row kinds. We use superscripts to provide copies of type constructors at different kinds, and thus avoid overloading of kinds.

- 1.11.4 DEFINITION [ROW EXTENSION OF A SIGNATURE]: Let \mathcal{S}_0 and \mathcal{S}_1 be signatures where all symbols of \mathcal{S}_1 are unary. The row extension of \mathcal{S}_0 with \mathcal{S}_1 is the signature \mathcal{S} defined as follows where κ ranges over basic kinds (those used in \mathcal{S}_0 and \mathcal{S}_1) and s ranges over row kinds:

| $F \in \text{dom}(\mathcal{S})$ | Signature | Conditions |
|---------------------------------|---|--|
| G^s | $(K \Rightarrow \kappa).s$ | $(G : K \Rightarrow \kappa) \in \mathcal{S}_0$ |
| H | $K.\text{Row}(\emptyset) \Rightarrow \kappa.\text{Type}$ | $(H : K \Rightarrow \kappa) \in \mathcal{S}_1$ |
| $\partial^{\kappa,L}$ | $\kappa.(\text{Type} \Rightarrow \text{Row}(L))$ | |
| $\ell^{\kappa,L}$ | $\kappa.(\text{Type} \otimes \text{Row}(\ell.L) \Rightarrow \text{Row}(L))$ | $\ell \notin L$ |

\square

We usually write $\ell^{\kappa,L} : \tau_1 ; \tau_2$ instead of $\ell^{\kappa,L} \tau_1 \tau_2$ and let this symbol be right-associative. We often drop the superscripts of type constructors since,

for any type expression T , superscripts can be unambiguously recovered from the kind of T .

- 1.11.5 EXAMPLE: Let us assume there is a single basic kind \star and that \mathcal{S}_1 contain a unique type constructor Π (hence of kind $\star \Rightarrow \star$). An example of row type is $X_0 \rightarrow \Pi(\ell_1 : G ; (Y \rightarrow \partial X_0))$. With all superscripts annotations, this type is

$$X_0 \rightarrow^{\star, Type} \Pi(\ell_1^{\star, Row(\emptyset)} : G^{Type} ; (Y \rightarrow^{\star, Row(\{\ell_1\})} \partial^{\star, Row(\{\ell_1\})} X_0)).$$

Intuitively, this is the type of a function that takes a value of type X_0 and returns a record where field ℓ_1 has type G and all other fields are functions that given a value of an arbitrary type would returns a value of (the same) type X_0 . An instance of this type is $X_0 \rightarrow \Pi(\ell_1 : G ; ((\ell_2 : Y_2 ; Y') \rightarrow (\ell_2 : X_0 ; \partial X_0)))$, obtained by instantiating row variable Y and by expanding the constant row ∂X_0 . As shown below, this type is actually equivalent to $X_0 \rightarrow \Pi(\ell_1 : G ; \ell_2 : Y_2 \rightarrow X_0 ; (Y' \rightarrow \partial X_0))$, by distributivity of type constructor \rightarrow other type constructor ℓ_2 . Please, note again the difference between Y , which is a row variable that can expand to different type variables on different labels, and ∂X , which is a constant row that expands to the same type variable X on all labels. \square

- 1.11.6 EXAMPLE [ILL-KINDED EXPRESSION]: Under the assumptions of the previous example, the expression $X \rightarrow \Pi(X)$ is not a row type, since variable X cannot simultaneously be of row kind $Type$ and $Row(\emptyset)$ as required by its two occurrences, from left to right respectively. The expression $(\ell : X ; \ell : X' ; X'')$ is also ill-kinded, for the inner expression $(\ell : X' ; X'')$ of row kind $Row(L)$ with $\ell \notin L$ cannot also have row kind $Row(\{\ell\})$, as required by its occurrence in the whole expression. Indeed, row kinds prohibit multiple definitions of the same label, as well as using rows in place of basic types and conversely. Notice that $\Pi(\Pi(X))$ is also ill-formed, since type constructors of \mathcal{S}_1 are not lifted to row kinds and thus cannot appear in rows, except under the type constructor ∂ , hence as basic types. \square
- 1.11.7 EXERCISE [★★★, \rightarrow]: Design an algorithm that infers superscripts of type constructors of a type expression from its kind. Can the kind of the expression be inferred as well? Can you give an algorithm to check that type expressions are well-kinded when both the superscripts of type constructors and the kinds of the whole type expressions are omitted? \square

Meaning of rows

As mentioned above, a row of kind $Row(L)$ stands for a function from $\mathcal{L} \setminus L$ to types. Actually, it is simpler to represent this function explicitly as an

infinitely branching tree in the model. For this purpose, we use a collection of constructors L of (infinite but denumerable) arity $\mathcal{L} \setminus L$.

- 1.11.8 DEFINITION [ROW MODEL]: Let \mathcal{S} be the row extension of a signature \mathcal{S}_0 with a signature \mathcal{S}_1 . Let $\mathcal{S}_{\mathcal{M}}$ be the following signature, where κ ranges over basic kinds and L ranges over finite subsets of \mathcal{L} :

| $F \in \text{dom}(\mathcal{S}_{\mathcal{M}})$ | Signature | Conditions |
|---|---|--|
| G | $(K \Rightarrow \kappa). \text{Type}$ | $(G : K \Rightarrow \kappa) \in \mathcal{S}_0$ |
| H | $K. \text{Row}(\emptyset) \Rightarrow \kappa. \text{Type}$ | $(H : K \Rightarrow \kappa) \in \mathcal{S}_1$ |
| L^κ | $\kappa. (\text{Type}^{\mathcal{L} \setminus L} \Rightarrow \text{Row}(L))$ | |

Let \mathcal{M}_κ consist of the regular trees t built over the signature $\mathcal{S}_{\mathcal{M}}$ such that $t(\epsilon)$ has image kind κ . We interpret a type constructor F of signature $K \Rightarrow \kappa.s$ as a function that maps $T \in \mathcal{M}_K$ to $t \in \mathcal{M}_{\kappa.s}$ defined by cases on F and on the basic kind κ .

| $F \in \mathcal{S}$ | $t(\epsilon)$ | For $d \in \text{dom}(K)$ and $\ell \in \mathcal{L} \setminus L, \ell \neq \ell_0$. |
|-----------------------|---------------|--|
| G^{Type} | G | $t/d = T(d)$ |
| H | H | $t/1 = T(1)$ |
| $G^{\text{Row}(L)}$ | L^κ | $t(\ell) = G \wedge t/(\ell \cdot d) = T(d)/\ell$ |
| $\partial^{\kappa,L}$ | L^κ | $t/\ell = T(1)$ |
| $\ell_0^{\kappa,L}$ | L^κ | $t/\ell_0 = T(1) \wedge t/\ell = T(2)/\ell$ |

In the presence of subtyping, we let type constructors G and H behave in $\mathcal{S}_{\mathcal{M}}$ as in \mathcal{S}_0 and \mathcal{S}_1 and type constructors L^κ be isolated and covariant in every position. Models that define ground types and interpret type constructors in this manner are referred to as *row* models. \square

Reasoning with row types

In this section, we assume a subtyping model. All reasoning principles also apply to the equality-only model, which is a subcase of the subtyping model.

The meaning of rows has been carefully defined so as to be independent of some syntactical choices. In particular, the order in which the types of significant fields have been declared leaves the meaning of rows unchanged. This is formally stated by the following Lemma.

- 1.11.9 LEMMA: The equations of Figure 1-17 are equivalent to true. \square

Proof: Each equation can be considered independently. It suffices to see that any ground assignment ϕ sends both sides of the equation to the same element

$$\begin{aligned}
(\ell_1 : \mathbf{T}_1 ; \ell_2 : \mathbf{T}_2 ; \mathbf{T}_3) &= (\ell_2 : \mathbf{T}_2 ; \ell_1 : \mathbf{T}_1 ; \mathbf{T}_3) && \text{(C-Row-LL)} \\
\partial \mathbf{T} &= (\ell : \mathbf{T} ; \partial \mathbf{T}) && \text{(C-Row-DL)} \\
\partial(G \mathbf{T}_1 \dots \mathbf{T}_n) &= G \partial \mathbf{T}_1 \dots \partial \mathbf{T}_n && \text{(C-Row-DG)} \\
G(\ell : \mathbf{T}_1 ; \mathbf{T}'_1) \dots (\ell : \mathbf{T}_n ; \mathbf{T}'_n) &= (\ell : G \mathbf{T}_1 \dots \mathbf{T}_n ; G \mathbf{T}'_1 \dots \mathbf{T}'_n) && \text{(C-Row-GL)}
\end{aligned}$$

Figure 1-17: Equational reasoning for row types.

$$\begin{aligned}
(\ell_1 : \mathbf{T}_1 ; \mathbf{T}'_1) = (\ell_2 : \mathbf{T}_2 ; \mathbf{T}'_2) &\equiv \exists \mathbf{X}. (\mathbf{T}'_1 = (\ell_2 : \mathbf{T}_2 ; \mathbf{X}) \wedge \mathbf{T}'_2 = (\ell_1 : \mathbf{T}_1 ; \mathbf{X})) && \text{(C-MUTE-LL)} \\
&\quad \text{if } \mathbf{X} \# \text{ftv}(\mathbf{T}_1, \mathbf{T}'_1, \mathbf{T}_2, \mathbf{T}'_2) \wedge \ell_1 \neq \ell_2 \\
(\ell : \mathbf{T} ; \mathbf{T}') = G \mathbf{T}'_i &\equiv \exists (\mathbf{X}_i, \mathbf{X}'_i)^I. (\mathbf{T} = G \mathbf{X}_i^I \wedge \mathbf{T}' = G \mathbf{X}'_i^I \wedge \mathbf{T}_i = (\ell : \mathbf{X}_i ; \mathbf{X}'_i)) && \text{(C-MUTE-LG)} \\
&\quad \text{if } (\mathbf{X}_i, \mathbf{X}'_i)^I \# \text{ftv}(\mathbf{T}, \mathbf{T}', \mathbf{T}'_i) \\
\partial \mathbf{T} = G \mathbf{T}'_i &\equiv \exists \mathbf{X}_i^I. (\mathbf{T} = G \mathbf{X}_i^I \wedge (\mathbf{T}_i = \partial \mathbf{X}_i)^I) && \text{(C-MUTE-DG)} \\
&\quad \text{if } \mathbf{X}_i^I \# \text{ftv}(\mathbf{T}, \mathbf{T}'_i) \\
\partial \mathbf{T} = (\ell : \mathbf{T}' ; \mathbf{T}'') &\equiv \mathbf{T} = \mathbf{T}' \wedge \partial \mathbf{T} = \mathbf{T}'' && \text{(C-MUTE-DL)}
\end{aligned}$$

Figure 1-18: Constraint equivalence laws for rows.

in the model, which follows directly from the meaning of row types. Notice that this fact only depends on the semantics of types and not on the semantics of the subtyping predicate. \square

It follows from those equations that type constructors ℓ , ∂ , and G are never isolated, each equation exhibiting a pair of compatible top symbols. Variances and incompatible pairs of type constructors depend on the signature $\mathcal{S}_0 \uplus \mathcal{S}_1$. However, it is not difficult to see that type constructors ∂ and ℓ are logically covariant in all directions and that the logical variance of types constructors G of $\text{dom}(\mathcal{S}_0 \uplus \mathcal{S}_1)$ correspond to their syntactic variance, which, in most cases, will allow the decomposition of equations with the same top symbols. Moreover, an equation between two terms whose top symbols form one of the four compatible pairs derived from the equations of Figure 1-17 holds only if immediate subexpressions can be “conciliated” in some way. There is a transformation quite similar to decomposition, called *mutation*, that mimics the equations for rows (Figure 1-17) and described by the rules of Figure 1-18. For sake of readability and conciseness, we write \mathbf{T}_i^I instead of $\mathbf{T}_i^{i \in I}$.

1.11.10 LEMMA [MUTATION]: All equivalence laws in Figure 1-18 hold. \square

Proof:

◦ *Case C-MUTE-LL*: Let $\mathbf{x} \# ftv(\mathbf{T}_1, \mathbf{T}'_1, \mathbf{T}_2, \mathbf{T}'_2)$ **(1)** and $\ell_1 \neq \ell_2$. Let $Row(L)$ be the row kind of this equation. Let ϕ be a ground assignment that validates the constraint $(\ell_1 : \mathbf{T}_1 ; \mathbf{T}'_1) = (\ell_2 : \mathbf{T}_2 ; \mathbf{T}'_2)$. That is, ϕ sends all terms of the multi-equation to the same ground type t of row kind $Row(L)$. Moreover, the row-term semantics implies that t satisfies $t(\epsilon) = L$, $t/\ell_1 = \phi(\mathbf{T}_1) = \phi(\mathbf{T}'_2)/\ell_1$, $t/\ell_2 = \phi(\mathbf{T}'_1)/\ell_2 = \phi(\mathbf{T}_2)$, and $t/\ell = \phi(\mathbf{T}'_2)/\ell = \phi(\mathbf{T}'_1)/\ell$ for all $\ell \in \mathcal{L} \setminus \ell_1, \ell_2, L$ **(2)**. Let t' be the tree defined by $t'(\epsilon) = \ell_1, \ell_2, L$ and $t'/\ell = t/\ell$ for all $\ell \in \mathcal{L} \setminus \ell_1, \ell_2, L$. By construction and **(2)**, $\phi[\mathbf{x} \mapsto t']$ satisfies both equations $\mathbf{T}'_1 = (\ell_2 : \mathbf{T}_2 ; \mathbf{x})$ and $\mathbf{T}'_2 = (\ell_1 : \mathbf{T}_1 ; \mathbf{x})$. Thus by CM-EXISTS and **(1)**, ϕ satisfies $\exists \mathbf{x}. \mathbf{T}'_1 = (\ell_2 : \mathbf{T}_2 ; \mathbf{x}) \wedge \mathbf{T}'_2 = (\ell_1 : \mathbf{T}_1 ; \mathbf{x})$. Conversely, we have the entailment:

$$\begin{aligned} & \exists \mathbf{x}. (\mathbf{T}'_1 = (\ell_2 : \mathbf{T}_2 ; \mathbf{x}) \wedge \mathbf{T}'_2 = (\ell_1 : \mathbf{T}_1 ; \mathbf{x})) \\ \equiv & \exists \mathbf{x}. ((\ell_1 : \mathbf{T}_1 ; \mathbf{T}'_1) = (\ell_1 : \mathbf{T}_1 ; \ell_2 : \mathbf{T}_2 ; \mathbf{x}) \wedge \\ & (\ell_2 : \mathbf{T}_2 ; \mathbf{T}'_2) = (\ell_2 : \mathbf{T}_2 ; \ell_1 : \mathbf{T}_1 ; \mathbf{x})) \end{aligned} \quad \mathbf{(3)}$$

$$\vdash \exists \mathbf{x}. (\ell_1 : \mathbf{T}_1 ; \mathbf{T}'_1) = (\ell_2 : \mathbf{T}_2 ; \mathbf{T}'_2) \quad \mathbf{(4)}$$

$$\equiv (\ell_1 : \mathbf{T}_1 ; \mathbf{T}'_1) = (\ell_2 : \mathbf{T}_2 ; \mathbf{T}'_2) \quad \mathbf{(5)}$$

(3) follows by covariance of ℓ_1 and ℓ_2 ; **(4)** by C-ROW-LL and transitivity of equivalence; **(5)** by C-EX* and **(1)**.

◦ *Cases C-MUTE-LG, C-MUTE-DG, and C-MUTE-DL*: The reasoning is similar to the case C-Mute-LL. \square

Solving row constraints in an equality model

In this section, we extend the constraint solver for the equality-only free tree model (Figure 1-11), so as to handle rows. We thus assume an equality-only model.

Mutation is a common technique to solve equations in a large class of non-free algebras that are described by *syntactic theories* (Kirchner and Klay, 1990). The equations of Figure 1-17 happen to be a syntactic presentation of an equational theory, from which a unification algorithm could be automatically derived (Rémy, 1993). We recover the same transformation rules directly, without using results on syntactic theories.

The following lemma shows that all pairs of distinct type constructors for which there is no mutation rule are in fact incompatible.

1.11.11 LEMMA: All symbols $H \in \mathcal{S}_1$ are isolated. Furthermore for every pair of distinct type constructors $G_1, G_2 \in dom(\mathcal{S}_0 \uplus \mathcal{S}_1)$, and every row kind s , we have $G_1^s \not\bowtie G_2^s$. \square

Proof: Terms of the form $H\vec{T}$ are interpreted by ground types with H at occurrence ϵ , and conversely the only interpretations of types with H at occurrence ϵ are terms of the form $H\vec{T}$. Hence, no ground assignment can ever

$$\begin{aligned}
(\ell_1 : \mathbf{X}_1 ; \mathbf{X}'_1) = (\ell_2 : \mathbf{T}_2 ; \mathbf{T}'_2) = \epsilon &\rightarrow \exists \mathbf{Y}. (\mathbf{X}'_1 = (\ell_2 : \mathbf{T}_2 ; \mathbf{Y}) \wedge \mathbf{T}'_2 = (\ell_1 : \mathbf{X}_1 ; \mathbf{Y})) \\
&\quad \wedge (\ell_1 : \mathbf{X}_1 ; \mathbf{X}'_1) = \epsilon && \text{(S-MUTE-LL)} \\
&\quad \text{if } \mathbf{Y} \# \text{ftv}(\mathbf{X}_1, \mathbf{X}'_1, \mathbf{T}_2, \mathbf{T}'_2) \wedge \ell_1 \neq \ell_2 \\
(\ell : \mathbf{X} ; \mathbf{X}') = G \mathbf{T}_i^{i \in I} = \epsilon &\rightarrow \exists (\mathbf{Y}_i, \mathbf{Y}'_i)^{i \in I}. (\mathbf{X} = G \mathbf{Y}_i^{i \in I} \wedge \mathbf{X}' = G \mathbf{Y}'_i^{i \in I} \wedge \mathbf{T}_i = (\ell : \mathbf{Y}_i ; \mathbf{Y}'_i)) \\
&\quad \wedge (\ell : \mathbf{X} ; \mathbf{X}') = \epsilon && \text{(S-MUTE-LG)} \\
&\quad \text{if } (\mathbf{Y}_i, \mathbf{Y}'_i)^{i \in I} \# \text{ftv}(\mathbf{X}, \mathbf{X}', \mathbf{T}_i^{i \in I}) \\
\partial \mathbf{X} = G \mathbf{T}_i^{i \in I} = \epsilon &\rightarrow \exists \mathbf{Y}_i^{i \in I}. (\mathbf{X} = G \mathbf{Y}_i^{i \in I} \wedge (\mathbf{T}_i = \partial \mathbf{Y}_i)^{i \in I}) \\
&\quad \wedge \partial \mathbf{X} = \epsilon && \text{(S-MUTE-DG)} \\
&\quad \text{if } \mathbf{Y}_i^{i \in I} \# \text{ftv}(\mathbf{X}, \mathbf{T}_i^{i \in I}) \\
\partial \mathbf{X} = (\ell : \mathbf{T} ; \mathbf{T}') = \epsilon &\rightarrow \mathbf{X} = \mathbf{T} \wedge \partial \mathbf{X} = \mathbf{T}' \wedge \partial \mathbf{X} = \epsilon && \text{(S-MUTE-DL)} \\
G \vec{\mathbf{T}} = G' \vec{\mathbf{T}}' = \epsilon &\rightarrow \text{false} && \text{(S-CLASH-I)} \\
&\quad \text{if } F \bowtie F'
\end{aligned}$$

Figure 1-19: Unification addendum for row types

send $H\vec{\mathbf{T}}$ and $F\vec{\mathbf{T}}'$ to the same ground term when $F \neq H$ and, as a result, H is isolated.

Let G_1 and G_2 be two type constructors of \mathcal{S}_0 . For $s = \text{Type}$, the interpretations of terms of the form $G_1^s \vec{\mathbf{T}}$ and $G_2^s \vec{\mathbf{T}}'$ are ground types with symbols G_1 and G_2 at occurrence ϵ , respectively. Hence they cannot be made equal under any ground assignment. For $s = \text{Row}(L)$, the interpretations of terms of the form $G_1^s \vec{\mathbf{T}}$ and $G_2^s \vec{\mathbf{T}}'$ are ground types with constructor L at occurrence ϵ and constructors G_1 and G_2 at occurrence 1, respectively. Hence they cannot be made equal under any ground assignment. \square

Any other combination of type constructors forms a compatible pair, as illustrated by equations of Figure 1-17 and can be transformed by mutation rules of Figure 1-18. The constraint solver for row-terms is the relation \rightarrow^\dagger defined by the rewriting rules of Figure 1-11, except Rule S-CLASH, which is replaced by the set of rules of Figure 1-19.

1.11.12 LEMMA: The rewriting system \rightarrow^\dagger is strongly normalizing. \square

Please, note that the termination of \rightarrow^\dagger relies on types being well-kinded. In particular, \rightarrow^\dagger would not terminate on the ill-kinded system of equations $\mathbf{X} = \ell : \mathbf{T} ; \mathbf{X}' \wedge \mathbf{X}' = \ell' : \mathbf{T} ; \mathbf{X}$.

1.11.13 LEMMA: If $U \rightarrow^\dagger U'$, then $U \equiv U'$. \square

Proof: It suffices to check the property independently for each rule defining \rightarrow^\dagger . The proof for rules of Figure 1-11 but S-CLASH remain valid for row terms. For S-DECOMPOSE, it follows by the invariance of all type constructors, which is preserved for row terms. For rule S-CLASS-I it follows by Lemma 1.11.11 and for mutation rules, it follows by Lemma 1.11.10. \square

Although reductions \rightarrow are not sound for row types, hence \rightarrow cannot be used for computation over row types, it has some interest. In particular, the following property shows that normal forms for row types are identical to normal forms for regular types.

1.11.14 LEMMA: A system U in normal form for \rightarrow^\dagger is also in normal form for \rightarrow . \square

Proof: The only rule of \rightarrow that is not in \rightarrow^\dagger is S-CLASH. Thus, it suffices to observe that if Rule S-CLASH would be applicable, then either Rule S-CLASS-I or a mutation rule would be applicable as well. \square

As a corollary, Lemma 1.8.6 extends to row types.

Operations on records

We now illustrate the use of rows for typechecking operations on records. The simplest application of rows are full records with infinite carrier. Records types are expressed with rows instead of a simple product type. The basic operations are the same as in the finite case, that is, *creation*, *polymorphic update*, and *polymorphic access*, but labels are now taken from an infinite set. However, creation and polymorphic update, which where destructors are now taken as constructors and used to represent records as association lists. The access of a record \mathbf{v} at a field ℓ is obtained by linearly searching \mathbf{v} for a definition of field ℓ and returning this definition, or returning the default value if no definition has been found for ℓ .

1.11.15 EXAMPLE [FULL RECORDS]: We assume a unique basic kind \star and a unique covariant isolated type constructor Π in \mathcal{S}_1 . Let $\{\cdot\}$ be a unary constructor, $(\{\cdot \text{ with } \cdot = \ell\})^{\ell \in \mathcal{L}}$ a collection of binary constructors, and $(\ell.\{\cdot\})^{\ell \in \mathcal{L}}$ a collection of unary destructors with the following reduction rules:

$$\begin{array}{lll} \{\mathbf{v}\}.\{\ell\} & \xrightarrow{\delta} & \mathbf{v} & \text{(RD-DEFAULT)} \\ \{\mathbf{w} \text{ with } \ell = \mathbf{v}\}.\{\ell\} & \xrightarrow{\delta} & \mathbf{v} & \text{(RD-FOUND)} \\ \{\mathbf{w} \text{ with } \ell' = \mathbf{v}\}.\{\ell\} & \xrightarrow{\delta} & \mathbf{w}.\{\ell\} & \text{if } \ell \neq \ell' \quad \text{(RD-FOLLOW)} \end{array}$$

Let the initial environment Γ_0 contain the following bindings

$$\begin{array}{l} \{\cdot\} : \quad \forall \mathbf{X}. \mathbf{X} \rightarrow \Pi(\partial \mathbf{X}) \\ \{\cdot \text{ with } \ell = \cdot\} : \quad \forall \mathbf{X} \mathbf{X}' \mathbf{Y}. \Pi(\ell : \mathbf{X} ; \mathbf{Y}) \rightarrow \mathbf{X}' \rightarrow \Pi(\ell : \mathbf{X}' ; \mathbf{Y}) \\ \ell.\{\ell\} : \quad \forall \mathbf{X} \mathbf{Y}. \Pi(\ell : \mathbf{X} ; \mathbf{Y}) \rightarrow \mathbf{X} \end{array}$$

-
- 1.11.16 EXERCISE [FULL RECORDS, ★★★, \rightarrow]: Check that these definitions meet the requirements of Definition 1.7.6. □
- 1.11.17 EXERCISE [FIELD EXCHANGE, ★★]: Add an operation to permute two fields of a record: give the reduction rules and the typing assumptions and check that the requirements of Definition 1.7.6 are preserved. □
- 1.11.18 EXERCISE [NORMAL FORMS FOR RECORDS, ★★★]: Record values may contain repetitions. For instance, $\{\{\mathbf{w} \text{ with } \ell = \mathbf{v}\} \text{ with } \ell = \mathbf{v}'\}$ is a value that is in fact observationally equivalent to $\{\mathbf{w} \text{ with } \ell = \mathbf{v}'\}$. So are the two record values $\{\{\mathbf{w} \text{ with } \ell = \mathbf{v}\} \text{ with } \ell' = \mathbf{v}'\}$ and $\{\{\mathbf{w} \text{ with } \ell' = \mathbf{v}'\} \text{ with } \ell = \mathbf{v}\}$ when $\ell' \neq \ell$. Modify the semantics, so that two record values of the same type have similar structure and records do not contain inaccessible values. □
- 1.11.19 EXERCISE [MAP APPLY, ★★]: Add a binary operation `map` such that the expressions $(\text{map } \mathbf{v} \ \mathbf{w}).\{\ell\}$ and $\mathbf{v}.\{\ell\} \ \mathbf{w}.\{\ell\}$ reduce to the same value for every label ℓ . □
- 1.11.20 EXERCISE [★, \rightarrow]: So far, full records are almost constants. This condition is not necessary for values, but only for types. As an example, introduce a primitive record, that is a nullary record constructor, that maps every label to a distinct integer. Give its typing assumption and review the semantics of records. □

As opposed to full records, standard records are partial and their domains are finite (but with infinite carrier) and statically determined from their types. Standard records can be built by extending the empty record on a finite number of fields. We refer to such records as records with *polymorphic extension*. Record with polymorphic extension can be obtained by means of encoding into full records, much as in the finite case.

- 1.11.21 EXAMPLE [ENCODING OF POLYMORPHIC EXTENSION]: Reusing the two type definitions `abs` and `pre` that have introduced for the finite case, we let `abs` encodes an undefined field and `pre` encode a field with value \mathbf{v} . We use the following syntactic sugar with their meaning and principal types:

$$\begin{aligned}
 \langle \rangle &\stackrel{\text{def}}{=} \{\text{abs}\} \\
 &: \Pi(\partial\text{abs}) \\
 \langle \cdot \text{ with } \ell = \cdot \rangle &\stackrel{\text{def}}{=} \lambda \mathbf{v} . \lambda \mathbf{w} . \{\mathbf{w} \text{ with } \ell = \text{pre } \mathbf{v}\} \\
 &: \forall \mathbf{X} \mathbf{X}' \mathbf{Y} . \Pi(\ell : \mathbf{X} ; \mathbf{Y}) \rightarrow \mathbf{X}' \rightarrow \Pi(\ell : \text{pre } \mathbf{X}' ; \mathbf{Y}) \\
 \cdot \langle \ell \rangle &\stackrel{\text{def}}{=} \lambda \mathbf{v} . \text{pre}^{-1}(\mathbf{v} . \{\ell\}) \\
 &: \forall \mathbf{X} \mathbf{Y} . \Pi(\ell : \text{pre } \mathbf{X} ; \mathbf{Y}) \rightarrow \mathbf{X}
 \end{aligned}$$

□

- 1.11.22 EXERCISE [RECOMMENDED, ★]: Extension may actually override an existing field. Can you define a version polymorphic extension that prevents this situation from happening? Add an operation that hide some particular field of a record. □

Extensible records can also be implemented directly, without encoding into full records. In fact, this requires only a tiny variation on full records.

- 1.11.23 EXAMPLE [RECORDS WITH POLYMORPHIC EXTENSION]: Let \star and \diamond be two basic kinds. Let the basic signature \mathcal{S}_0 contain (in addition to \rightarrow) the covariant isolated type constructors pre of kind $\star \Rightarrow \diamond$ and abs of kind \diamond . In the presence of subtyping, we may assume $\text{pre} \leq \text{abs}$. Let \mathcal{S}_1 contain the unique covariant isolated type constructor Π of kind $\diamond \Rightarrow \star$. Let $\langle \cdot \rangle$ be a unary constructor, $(\cdot \text{ with } \cdot = \ell)^{\ell \in \mathcal{L}}$ be a binary constructor, and $(\ell.\{\cdot\})^{\ell \in \mathcal{L}}$ be a unary destructor with the following reduction rules:

$$\begin{array}{ll} \langle \mathbf{w} \text{ with } \ell = \mathbf{v} \rangle . \langle \ell \rangle & \xrightarrow{\delta} \mathbf{v} & \text{(ER-FOUND)} \\ \langle \mathbf{w} \text{ with } \ell' = \mathbf{v} \rangle . \langle \ell \rangle & \xrightarrow{\delta} \mathbf{w} . \{ \ell \} & \text{if } \ell \neq \ell' \quad \text{(ER-FOLLOW)} \end{array}$$

Let Γ_0 contain the following typing assumptions:

$$\begin{array}{l} \langle \cdot \rangle : \Pi(\partial \text{abs}) \\ \langle \cdot \text{ with } \ell = \cdot \rangle : \forall \mathbf{X} \mathbf{X}' \mathbf{Y}. \Pi(\ell : \mathbf{X} ; \mathbf{Y}) \rightarrow \mathbf{X}' \rightarrow \Pi(\ell : \text{pre } \mathbf{X}' ; \mathbf{Y}) \\ \cdot . \langle \ell \rangle : \forall \mathbf{X} \mathbf{Y}. \Pi(\ell : \text{pre } \mathbf{X} ; \mathbf{Y}) \rightarrow \mathbf{X} \end{array}$$

□

Notice that the typing assumptions obtained from the direct approach are identical to those obtained via the encoding into full records in Example 1.11.21.

- 1.11.24 EXERCISE [★★★★★, \rightarrow]: Prove the equivalence between the direct semantics and the semantics via the encoding into records with a default. □
- 1.11.25 EXERCISE [RECOMMENDED, ★★, \rightarrow]: Prove that type soundness for extensible records hold in both the subtyping model and equality-only models. □
- 1.11.26 EXERCISE [RECOMMENDED, ★, \rightarrow]: Check that in the subtyping a record with more fields can be used in place of records with fewer fields. Check that this is not the case in the equality-only model. □

- 1.11.27 EXAMPLE [REFINEMENT OF RECORD TYPES]: In an equality-only model, records with more fields cannot be used in place of records with fewer fields. However, this may be partially recovered by a small refinement of the structure of types. The presence of fields can actually be split from their types, thus enabling some polymorphism over the presence of fields while type of fields themselves remains fixed. Let \circ be a new basic kind. Let type constructors abs and pre be both of kind \circ and let \cdot be a new type constructor of kind $\circ \otimes \star \Rightarrow \diamond$. Let Γ_0 contain the following typing assumptions:

$$\begin{aligned} \langle \rangle &: \forall X. \Pi(\partial(\text{abs} \cdot X)) \\ \langle \cdot \text{ with } \ell = \cdot \rangle &: \forall ZXX'Y. \Pi(\ell : X ; Y) \rightarrow X' \rightarrow \Pi(\ell : Z \cdot X' ; Y) \\ \cdot, \langle \ell \rangle &: \forall XY. \Pi(\ell : \text{pre} \cdot X ; Y) \rightarrow X \end{aligned}$$

The semantics of records remain unchanged. The new signature strictly generalizes the previous one (strictly more programs can be typed) while preserving type soundness. Here is a program that can now be typed and that could not be typed before:

(if a then $\langle \langle \rangle \text{ with } \ell' = \text{true} \rangle \text{ with } \ell = 1 \rangle$ else $\langle \langle \rangle \text{ with } \ell = 2 \rangle \rangle, \langle \ell \rangle$

Notice however, that when a present field is forgotten, the type of the field is not. Therefore two records defining the same field but with incompatible types can still not be mixed, which is possible in the subtyping model. \square

- 1.11.28 EXAMPLE [REFINED SUBTYPING]: The previous refinement for an equality-only model is not much interesting in the case of a subtyping model.

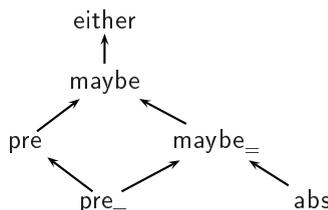
The subtyping assumption $\text{pre} \leq \text{abs}$ makes abs play the role of \top for fields. That is, abs encodes the absence of information and not the information of absence. In other words, a value whose field ℓ has type abs may either be undefined or defined on field ℓ ; in the latter case, the fact that field ℓ is actually defined has just been forgotten. Thus, types only provides a lower approximation of the actual domain of records. This is a lost of accuracy by comparison with the equality-only model, where a record domain is known from its type. As a result, some optimizations in the representation of records that are only possible when the exact domain of a record is statically known are lost.

Fortunately, there is a way to recover such accuracy. A conservative solution could of course to drop the inequality $\text{pre} \leq \text{abs}$. Notice that this would still be more expressive than using an equality model since, for instance $\Pi(\ell : \text{pre} (\top_1 \rightarrow \top_2) ; \top) \leq \Pi(\ell : \text{pre} \top ; \top)$ would still hold, as long as $\rightarrow \leq \top$ does hold. This solution is known as depth-only subtyping for records, while the previous one provided both depth and width record subtyping. Conversely, one could also keep width subtyping and disallow depth subtyping, by preserving

the relation $\text{pre} \leq \text{abs}$ while requiring pre to be invariant; in this case, presence of fields can be forgotten as a whole, but the types of fields cannot be weakened as long as fields remain visible.

Another more interesting solution consists in introducing another type constructor either of signature \diamond and assuming that $\text{pre} \leq \text{either}$ and $\text{abs} \leq \text{either}$ (but $\text{pre} \not\leq \text{abs}$). Here, either plays the role of \top for fields and means either *present* (and forgotten) or *absent*, while abs really means *absent*. The accuracy of typechecking can be formally stated as the fact that a record value of type $\Pi(\ell : \text{abs} ; \text{T})$ cannot define field ℓ . \square

1.11.29 EXAMPLE [MIXED SUBTYPING]: It is tempting to mix all variations of Example 1.11.28 together. As a first attempt, we may assume that the basic signature \mathcal{S}_0 contains covariant type constructors pre and maybe and invariant type constructors $\text{pre}_=$ and $\text{maybe}_=$, all of kind $\star \Rightarrow \diamond$ and two type constructors abs and either of kind \diamond , and that the subtype ordering \leq is defined by the following diagram:



Intuitively, we wish that $\text{pre}_=$ and $\text{maybe}_=$ be *logically* invariant, pre and maybe be *logically* covariant, and the equivalences $\text{pre}_= T \leq \text{maybe}_= T' \equiv T = T'$ and

$$\text{pre}_= T \leq \text{pre} T' \equiv \text{pre} T \leq \text{maybe} T' \equiv \text{maybe}_= T \leq \text{maybe} T' \equiv T \leq T' \quad (1)$$

simultaneously hold. However, (1) requires, for instance, type constructors $\text{pre}_=$ and pre to have the same direction, which is not currently possible since they do not have the same variance. Interestingly, this restriction may be relaxed by assigning variances of directions on a per type constructor basis and define structural subtyping accordingly (See Exercise 1.11.30). Then, replacing all occurrences of pre by $\text{pre}_=$ in Γ_0 preserves type soundness and allows for both accurate record types and flexible subtyping in the same setting. \square

1.11.30 EXERCISE [RELAXED VARIANCES, ★★★, \rightarrow]: Let \emptyset be allowed as a new variance, let extend the composition of variances defined in Example 1.3.9 with $\nu\emptyset = \emptyset$, and let \leq^\emptyset stands for the full relation on type constructors. Let each type constructor F of signature $K \Rightarrow \kappa$ now come with a mapping $\vartheta(F)$ from $\text{dom}(K)$ to variances. Let $\vartheta(t, t', \pi)$ be the variance of two

ground types t and t' at a path π recursively defined by $\vartheta(t, t', d \cdot \pi) = (\vartheta(t(\epsilon))(d) \cap \vartheta(t'(\epsilon))(d)) \vartheta(t/d, t'/d, \pi)$ and $\vartheta(t, t', \epsilon) = +$. Then define the interpretation of subtyping as follows: if $t, t' \in \mathcal{M}_\kappa$, let $t \leq t'$ hold if and only if for all path $\pi \in \text{dom}(t) \cap \text{dom}(t')$, $t(\pi) \leq^{\vartheta(t, t', \pi)} t'(\pi)$ holds.

Check that the relation \leq remains a partial ordering. Check that a type constructor whose direction d has been syntactically declared covariant (respectively contravariant, invariant) is still logically covariant (respectively contravariant, invariant) in d . \square

Record concatenation

Record concatenation takes two records and combines them into a new record whose fields are taken from whatever argument defines them. Of course, there is an ambiguity when the two records do not have disjoint domains and a choice should be made to disambiguate such cases. *Symmetric* concatenation let concatenation be undefined in this case (Harper and Pierce, 1991), while *asymmetric* concatenation let one-side (usually the right side) always take priority. Despite a rather simple semantics, record concatenation remains hard to type (with either a strict or a priority semantics). Solutions to type inference for record concatenation may be found, for instance, in (Wand, 1989; Rémy, 1992; Pottier, 2000).

Polymorphic variants

Variants can be defined via algebraic data-type definitions. However, as fields for records, variant tags are taken from a relatively small, finite collection of labels and two variant definitions will have incompatible types. Thus, to remain compatible, two variants must chose their tag among a larger collection that is a superset of all the possible tags of either variant. In general, this reduces the accuracy of types and forces useless dynamic checks for tags that could otherwise be known not to occur. Extensible variants (page 93) allow to work with an arbitrary large collection of tags, but do not improve accuracy. Polymorphic variants refers to a more precise typechecking mechanism for variants, where types more accurately describes the tags that may actually occur. They allow to build values of sum types out of a large, potentially infinite predefined set of tags and call polymorphic functions to explore them. As for record, this problem could be tackled by first considering polymorphic operations over variants built from a finite set of tags and total variants with an infinite set of tag independently and then by combining both approaches together. We propose a direct solution by a simple analogy with records.

Indeed, type constructor `pre` can be used to distinguish a (finite) set of tags

that the variant may actually carry, from other tags that are certain not to occur and typed with `abs`. For example, a variant $\ell.v$, built from a value v with a constructor tag ℓ of arity one, may be assigned the principal type scheme $\forall X. \Sigma(\ell : \text{pre } T; X)$ where T is the type of v . The unary type constructor Σ is used to coerce rows to variant types—thus, variant types and record types may share the same inner row structure and be simply distinguished by their top symbol. An instance of this polymorphic type is $\forall X. \Sigma(\ell : \text{pre } T; \text{abs})$, which tells that the variant must have been built with tag ℓ and no other tag, thus retaining exact information about the shape of the value. Another instance of the variant polymorphic type is $\Sigma(\ell : \text{pre } T; \ell' : \text{pre } T'; \text{abs})$. Indeed, it is sound to assume that the value might also have been built with some other tag ℓ' , even if we know that this is not actually the case. Interestingly, both values $\ell.v$ and $\ell'.v$ have this type and can be mixed at this type.

We use filters to explore variants. A filter $[\ell : v \mid v']$ is a function that expects a variant argument, thus of the form $\ell'.w$. It then proceeds with either $v \ w$, if $\ell' = \ell$, or $v' \ w$ otherwise. The type of this filter is $\Sigma(\ell : \text{pre } T; T') \rightarrow T''$ where T is the type of values accepted by v , $\Sigma(\ell : T''; T')$ is the type of values accepted by v' , and T'' is the type of values returned by both v and v' . Any type T'' would do, including, in particular, `abs`. Indeed, when w is passed to v' , it is known not to have tag ℓ , so the behavior of v' on ℓ does not matter. The null filter $[]$ can be used for v' . This filter should actually never be applied, which we ensure by assigning $[]$ the type $\forall X. \Sigma(\partial \text{abs}) \rightarrow X$, for no variant value has type $\Sigma(\partial \text{abs})$. For instance, the filter $[\ell : v_\ell \mid [\ell' : v_{\ell'} \mid []]]$, which may be abbreviated as $[\ell : v_\ell \mid \ell' : v_{\ell'}]$ can be applied to either $\ell.v$ or $\ell'.v'$. The following example formalizes polymorphic variants.

1.11.31 **EXAMPLE [POLYMORPHIC VARIANTS]:** Let \star and \diamond be two basic kinds. Let \mathcal{S} contain in addition to the arrow type constructor the two type constructors `pre` of kind $\star \Rightarrow \diamond$ and `abs` of kind \diamond . In the presence of subtyping we may assume $\text{abs} \leq \text{pre}$. Let \mathcal{S}_1 contain the unique covariant isolated type constructor Σ of kind $\diamond \Rightarrow \star$. Let Γ_0 be composed of unary constructors $(\ell. \cdot)^{\ell \in \mathcal{L}}$ and primitives $[]$ of arity 0 and $([\ell : \cdot \mid \cdot])^{\ell \in \mathcal{L}}$ of arity 3, given with the following reduction rules:

$$\begin{array}{ll} [\ell : v \mid v'] \ell.w & \xrightarrow{\delta} v \ w & (\text{EV-FOUND}) \\ [\ell : v \mid v'] \ell'.w & \xrightarrow{\delta} v' \ w & \text{if } \ell \neq \ell' \quad (\text{EV-FOLLOW}) \end{array}$$

and contain the following typing assumptions:

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

□

- 1.11.32 EXERCISE [SOUNDNESS FOR EXTENSIBLE VARIANTS, ★★★, ⇔]: Prove type soundness for extensible variants in both equality-only and subtyping models. □

Other applications of rows

Polymorphic records and variants are the most well-known applications of rows. Besides the many variations on their presentations—we have only illustrated some of them—there are several other interesting applications of rows.

Since objects can be viewed as record-of-functions, at least from a typing point of view, rows can also be used to type structural objects (Wand, 1994; Rémy, 1994; Rémy and Vouillon, 1998) and provide, in particular, polymorphic method invocation. This is the key to typechecking objects in Objective Caml (Rémy and Vouillon, 1998). First-class messages (Nishimura, 1998; Müller and Nishimura, 1998; Pottier, 2000) combine records and variants in an interesting way: while filters over variant types enforce all branches to have the same return type, first-class messages treat filters as records of functions (also called objects) rather than functions from a variant type to a shared return type. A message is an element of a variant type. The application of an object to a message, that is of a record of functions to a variant type, selects from the record the branch labeled with the same tag as the message and applies it to the content of the message, much as pattern matching. However, these applications are typechecked more accurately by first restricting the domain of the record to the set of tags that the message may possibly carry, and thus other branches and in particular their return type are left unconstrained.

Row types may also represent set of properties within types or type refinements and be used in type systems for program analysis. Two examples worth mentioning are their application to soft-typing (Cartwright and Fagan, 1991; Wright and Cartwright, 1994) and typechecking of uncaught exceptions (Leroy and Pessaux, 2000).

The key to rows is to decompose the set of row labels into a class of finite partitions that is closed by some operations. Here, those partitions are composed of singleton labels and co-finite sets of labels; the operations are merging (or conversely splitting) a singleton label and a co-finite set of labels. Other decompositions are possible, for instance, one could imagine to consider labels in a two-dimensional space. More generally, labels might also be given internal structure, for instance, one might consider automata as labels. Notice also that record types are stratified, since rows, that is, expressions of kind $Row(L)$, may not themselves contain records—constructors of S_1 are only given the image row kind *Type*. This restriction can be partially

relaxed leading to rows of increasing degrees (Rémy, 1992b) . . . and complexity! Yet more intriguing are typed-indexed rows where labels are themselves types (Shields and Meijer, 2001).

Alternatives to rows

The original idea of using rows to describe types of extensible records is due to Wand (Wand, 1987, 1988). A key simplification to row types is to make them total functions from labels to types and encode definiteness explicitly in the structure of fields, for instance with `pre` and `abs` type constructors, as presented here. This decomposition reduces the resolution of unification constraints to a simple equational reasoning (Rémy, 1993, 1992a). Other approaches that do not treat rows as total functions seem more *ad hoc* and have often hard-wired restrictions (Jategaonkar and Mitchell, 1988; Ogori and Buneman, 1989; Berthomieu, 1993; Ogori, 1999). Among these partial solutions, (Ogori, 1999) is quite interesting for its overall simplicity in the case where polymorphic access alone is required. Rows and fields may also be represented within ad-hoc type constraints rather than terms and equality (or subtyping) constraints. For example, qualified types use the predicates $(T \text{ has } \ell : T')$ and $(T \text{ lacks } \ell)$ to mean that field ℓ of row T is defined with type T' or undefined, respectively (Jones, 1994b; Odersky, Sulzmann, and Wehr, 1999b). These constraints are in fact equivalent in our equality-model to $\exists X.T = (\ell:\text{pre } T' ; X)$ and $\exists X.T = (\ell:\text{abs } ; X)$, respectively. Record typechecking has also been widely studied in the presence of subtyping. Usually, record subtyping is given meaning directly and not via rows. While these solutions are quite expressive, thanks to subtyping, they still suffer from their nonstructural treatment of record types and cannot type row extension. Thus, even in subtyping models the use of rows increases expressiveness, and is usually a simplification as well. The subtyping model can then also take advantage of the possibility of enriching type constructors `pre` and `abs` with more structure and relate them via subtyping (Pottier, 2000). Notice, that even though rows have been introduced for type inference, they seem to be beneficial to explicitly typed languages as well since even other advanced solutions (Cardelli and Mitchell, 1991; Cardelli, 1992) are limited.

Rules of Figure 1-19 are *one way* of solving row type constraints. In a model with subtyping constraints, a more direct closure-based resolution may be more appropriate (Pottier, 2003).

B *Solutions to Selected Exercises*

1.2.6 SOLUTION: The definition does not behave as expected, because `if` is a destructor, whose arguments—according to the call-by-value semantics of ML-the-calculus—are evaluated *before* `R-TRUE` or `R-FALSE` is allowed to fire. As a result, the semantics of the expression `if t0 then t1 else t2` is to evaluate *both* `t1` and `t2` before choosing one of them. Since these expressions may have side effects (for instance, they may fail to terminate, or update a reference), this semantics is undesirable. The desired evaluation order can be obtained by placing `t1` and `t2` within closures, which delays their evaluation, then invoking the closure returned by the conditional, forcing its body to be evaluated. In other words, the expression `if t0 then t1 else t2` should now be viewed as syntactic sugar for `if t0 (λz.t1) (λz.t2) \hat{O}` . The choice of the constant \hat{O} is arbitrary, since it is discarded; any value would do.

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

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

Please note that, because `X` occurs free within the environment `z1 : X`, it is impossible to apply `DM-GEN` to the judgement `z1 : X ⊢ z1 : X` in a nontrivial way. For this reason, `z2` cannot receive the type scheme $\forall X.X$, and the whole expression cannot receive type $X \rightarrow Y$, where `X` and `Y` are distinct.

1.2.22 SOLUTION: It is straightforward to prove that the identity function has type `int → 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 it no longer appears in the environment.

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

It is worth noting that, although the type derivation employs an arbitrary type variable `X`, the final typing judgement has no free 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 uses of DM-VAR implicit.

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

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

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

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

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

The first and third judgements 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 judgement, of course, is not: because it involves a nontrivial type scheme, it is not even a well-formed judgement in the simply-typed λ -calculus. The fourth judgement 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 DM-VAR, DM-GEN, and DM-INST, it is straightforward to see that, if $\Gamma_0 \vdash \hat{1} : T$ is derivable, then T must be int . Since int is not an arrow type, 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 then be a subterm of itself. For the same reason, this expression is ill-typed in DM as well. Indeed, it is not difficult to check that the presence of DM-GEN and DM-INST makes no difference: DM-GEN cannot generalize T as long as the binding $f : T$ appears in the environment, and DM-INST can only instantiate T to T itself. Thus, the self-application $f f$ is well-typed in DM *only* if f is let -bound, as opposed to λ -bound. The argument crucially relies on the fact that f must be assigned a *monotype*. Indeed, the expression $\lambda f. (f f)$ is well-typed in an implicitly-typed variant of System F: one of its types is $(\forall X. X \rightarrow X) \rightarrow (\forall X. X \rightarrow X)$. It also relies on the fact that types are *finite*: indeed, this expression 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.

- 1.2.23 SOLUTION: It is clear that the effect of DM-GEN may be obtained by a series of successive applications of DM-GEN'. Conversely, consider an instance of DM-GEN', whose premises are $\Gamma \vdash t : S$ (1) and $X \notin \text{ftv}(\Gamma)$ (2). Let us write $S = \forall \bar{X}. T$, where $\bar{X} \# \text{ftv}(\Gamma)$ (3). Applying DM-INST to (1) and to the identity substitution yields $\Gamma \vdash t : T$ (4). Applying DM-GEN to (4), (2) and (3) yields $\Gamma \vdash t : \forall X \bar{X}. T$, that is, $\Gamma \vdash t : \forall X. S$. Thus, the effect of DM-GEN' may be obtained by DM-INST and DM-GEN.

It is clear that DM-INST is a particular case of DM-INST' where \bar{Y} is empty. Conversely, consider an instance of DM-INST', whose premises are $\Gamma \vdash t : \forall \bar{X}. T$ (1) and $\bar{Y} \# \text{ftv}(\forall \bar{X}. T)$ (2). Let ρ be a renaming that exchanges \bar{Y} with \bar{Z} , where $\bar{Z} \# \text{ftv}(\forall \bar{Y}. [\bar{X} \mapsto \bar{T}']T)$ (3) and $\bar{Z} \# \text{ftv}(\Gamma)$ (4). Applying DM-INST to (1) yields $\Gamma \vdash t : [\bar{X} \mapsto \rho \bar{T}']T$ (5). Applying DM-GEN to (5) and (4) yields $\Gamma \vdash t : \forall \bar{Z}. [\bar{X} \mapsto \rho \bar{T}']T$, that is, $\Gamma \vdash t : \forall \rho \bar{Y}. [\bar{X} \mapsto \rho \bar{T}']T$ (6). Now, by (2) and (3), we have $[\bar{X} \mapsto \rho \bar{T}']T = \rho([\bar{X} \mapsto \bar{T}']T)$, so (6) may be written $\Gamma \vdash t : \forall \rho \bar{Y}. \rho([\bar{X} \mapsto \bar{T}']T)$, that is, $\Gamma \vdash t : \rho(\forall \bar{Y}. [\bar{X} \mapsto \bar{T}']T)$ (7). By (3), this is exactly $\Gamma \vdash t : \forall \bar{Y}. [\bar{X} \mapsto \bar{T}']T$. Thus, the effect of DM-INST' may be obtained by DM-INST and DM-GEN.

- 1.4.4 SOLUTION: Let us recall that a program t is well-typed if and only if a judgement of the form $C, \Gamma \vdash t : \sigma$, where C is satisfiable, holds. Let us show that it is in fact possible, without loss of generality, to require σ to be a monotype.

Assume $C, \Gamma \vdash t : \sigma$ (1) is derivable within HM(X). Let us write $\sigma = \forall \bar{X}[D]. T$, where $\bar{X} \# \text{ftv}(C)$ (2). Applying Lemma 1.4.1 to (1) yields $C \Vdash \exists \bar{X}. D$ (3). By HM-INST, (1) implies $C \wedge D, \Gamma \vdash t : T$ (4). By (3), we have $C \equiv C \wedge \exists \bar{X}. D \equiv \exists \bar{X}. (C \wedge D)$. Because C is satisfiable, this implies that $C \wedge D$

is satisfiable as well. Thus, the judgement (4), which involves the monotype \mathbb{T} , witnesses that \mathfrak{t} is well-typed.

We have shown that a program \mathfrak{t} is well-typed if and only if a judgement of the form $C, \Gamma \vdash \mathfrak{t} : \mathbb{T}$, where C is satisfiable, holds. Thus, by Theorems ?? and ??, well-typedness is the same for both rule sets.

1.4.5 SOLUTION: By Theorem ??, every rule in Figure 1-8 is admissible in $\text{HM}(X)$. Of course, so is HM-GEN . So, every judgement that is derivable via the rules of Figure 1-8 and HM-GEN is a valid $\text{HM}(X)$ judgement.

Conversely, assume $C, \Gamma \vdash \mathfrak{t} : \sigma$ (1) holds in $\text{HM}(X)$. We must show that it is derivable via the rules of Figure 1-8 and HM-GEN . Let us write $\sigma = \forall \bar{X}[D].\mathbb{T}$, where $\bar{X} \# \text{ftv}(C, \Gamma)$ (2). By HM-INST and (1), the judgement $C \wedge D, \Gamma \vdash \mathfrak{t} : \mathbb{T}$ (3) holds in $\text{HM}(X)$. This judgement involves a monotype, so, by Theorem ??, it is derivable via the rules of Figure 1-8. Furthermore, from (3) and (2), HM-GEN allows deriving $C \wedge \exists \sigma, \Gamma \vdash \mathfrak{t} : \sigma$ (4). Applying Lemma 1.4.1 to (1) yields $C \Vdash \exists \sigma$, so the judgement (4) may be written $C, \Gamma \vdash \mathfrak{t} : \sigma$. We have shown that (1) is derivable via the rules of Figure 1-8 and HM-GEN . In fact, it is possible to apply HM-GEN only once, at the end of the derivation.

1.5.1 SOLUTION: Within the type system $\text{PCB}(X)$, we have

$$\frac{\frac{\frac{}{z_1 \preceq Z \vdash z_1 : Z} \text{VAR}}{\text{let } z_2 : \forall Z[z_1 \preceq Z].Z \text{ in } z_2 \preceq Y \vdash \text{let } z_2 = z_1 \text{ in } z_2 : Y} \text{LET}}{\text{let } z_1 : X; z_2 : \forall Z[z_1 \preceq Z].Z \text{ in } z_2 \preceq Y \vdash \lambda z_1. \text{let } z_2 = z_1 \text{ in } z_2 : X \rightarrow Y} \text{ABS}}$$

The type variable Z , which occurs free in the left-hand instance of VAR , is generalized. However, z_2 does *not* receive the type scheme $\forall Z.Z$, which, as suggested earlier, is unsound; instead, it receives the *constrained* type scheme $\forall Z[z_1 \preceq Z].Z$. The latter is more restrictive than the former: indeed, the former claims that z_2 has every type, while the latter only claims that every valid type for z_1 is also a valid type for z_2 . Let us now examine the constraint $\text{let } z_1 : X; z_2 : \forall Z[z_1 \preceq Z].Z \text{ in } z_2 \preceq Y$, which appears at the root of the derivation. By C-INID and C-IN^* , it is equivalent to $\text{let } z_1 : X \text{ in } \exists Z.(z_1 \preceq Z \wedge Z \leq Y)$ and to $\exists Z.(X \leq Z \wedge Z \leq Y)$, which by C-EXTRANS is equivalent to $X \leq Y$. Thus, the judgement at the root of the above derivation may be written $X \leq Y \vdash \lambda z_1. \text{let } z_2 = z_1 \text{ in } z_2 : X \rightarrow Y$. In other words, the expression $\text{let } z_2 = z_1 \text{ in } z_2$ has type $X \rightarrow Y$ only under the assumption that X is a subtype of Y , which is sound. Even though LET allows unrestricted generalization of type variables, it remains sound, because the type scheme that it produces typically has *free program identifiers*, such as $\forall Z[z_1 \preceq Z].Z$ above.

1.7.10 SOLUTION: Let $\mathcal{E} = \text{let } z = \mathcal{E}_1 \text{ in } \mathfrak{t}_1$ and $\mathcal{E}_1[\mathfrak{t}]/\mu \sqsubseteq \mathcal{E}_1[\mathfrak{t}']/\mu'$ (1). Then,

$$\begin{aligned} & \text{let } \Gamma_0; \text{ref } M \text{ in } \llbracket \mathcal{E}[\mathfrak{t}]/\mu : T/M \rrbracket \\ & \equiv \text{let } \Gamma_0; \text{ref } M \text{ in } ((\text{let } z : \forall X[\llbracket \mathcal{E}_1[\mathfrak{t}] : X \rrbracket].X \text{ in } \llbracket \mathfrak{t}_1 : T \rrbracket) \wedge \llbracket \mu : M \rrbracket) \end{aligned} \quad (2)$$

$$\equiv \text{let } \Gamma_0; \text{ref } M; z : \forall X[\llbracket \mathcal{E}_1[\mathfrak{t}]/\mu : X/M \rrbracket].X \text{ in } \llbracket \mathfrak{t}_1 : T \rrbracket \quad (3)$$

$$\equiv \text{let } \Gamma_0; \text{ref } M; z : \forall X[\text{let } \Gamma_0; \text{ref } M \text{ in } \llbracket \mathcal{E}_1[\mathfrak{t}]/\mu : X/M \rrbracket].X \text{ in } \llbracket \mathfrak{t}_1 : T \rrbracket \quad (4)$$

$$\Vdash \text{let } \Gamma_0; \text{ref } M; z : \forall \bar{X}\bar{Y}[\text{let } \Gamma_0; \text{ref } M' \text{ in } \llbracket \mathcal{E}_1[\mathfrak{t}']/\mu' : X/M' \rrbracket].X \text{ in } \llbracket \mathfrak{t}_1 : T \rrbracket \quad (5)$$

where (2) is by definition of constraint generation, where $X \notin \text{ftv}(T, M)$ (6); (3) is by (6), C-LETAND, and by definition of constraint generation; (4) is by (6) and C-LETDUP; (5) follows from (1) and C-LETEX, for some \bar{Y} and M' such that $\bar{Y} \# \text{ftv}(X, M)$ (7) and $\text{ftv}(M') \subseteq \bar{Y} \cup \text{ftv}(M)$ (8) and $\text{dom}(M') = \text{dom}(\mu')$ and M' extends M . Note that (6), (7) and (8) imply $X \notin \text{ftv}(M')$ (9).

At this point, the type variables \bar{Y} , which determine the types of the newly allocated store cells, are universally quantified in the type scheme assigned to z , which is undesirable. We are stuck, because we cannot in general apply C-LETALL to hoist $\exists \bar{Y}$ out of the let constraint. Let us now assume that, by some external means, we are guaranteed $\bar{Y} = \emptyset$ (10). Then, we may proceed as follows:

$$\equiv \text{let } \Gamma_0; \text{ref } M'; z : \forall X[\text{let } \Gamma_0; \text{ref } M' \text{ in } \llbracket \mathcal{E}_1[\mathfrak{t}']/\mu' : X/M' \rrbracket].X \text{ in } \llbracket \mathfrak{t}_1 : T \rrbracket \quad (11)$$

$$\equiv \text{let } \Gamma_0; \text{ref } M' \text{ in } \llbracket \mathcal{E}[\mathfrak{t}']/\mu' : T/M' \rrbracket \quad (12)$$

where (11) follows from the fact the the memory locations that appear free in $\llbracket \mathfrak{t}_1 : T \rrbracket$ are members of $\text{dom}(\mu)$, thus are not members of $\text{dom}(M') \setminus \text{dom}(M)$; (12) is obtained by performing the steps that lead to (4) in reverse.

The requirement that \bar{Y} be empty, that is, $\text{ftv}(M) = \text{ftv}(M')$, is classic (Tofte, 1988). How is it enforced? Assume that the left-hand side of every let construct is required to be a non-expansive expression. By assumptions (ii) and (iii), this invariant is preserved by reduction. So, $\mathcal{E}_1[\mathfrak{t}]$ must be non-expansive, which, by assumption (i), guarantees that the reduction step does not allocate new memory cells. Then, μ' is μ , so M' is M .

1.9.1 SOLUTION: We must first ensure that R-ADD respects \sqsubseteq (Definition 1.7.5). Since the rule is pure, it is sufficient to establish that let Γ_0 in $\llbracket \hat{k}_1 \hat{+} \hat{k}_2 : T \rrbracket$ entails let Γ_0 in $\llbracket \widehat{k_1 + k_2} : T \rrbracket$. In fact, we have

$$\text{let } \Gamma_0 \text{ in } \llbracket \hat{k}_1 \hat{+} \hat{k}_2 : T \rrbracket$$

$$\equiv \text{let } \Gamma_0 \text{ in } \exists XY.(\hat{+} \preceq X \rightarrow Y \rightarrow T \wedge \hat{k}_1 \preceq X \wedge \hat{k}_2 \preceq Y) \quad (1)$$

$$\equiv \exists XY.(\text{int} \rightarrow \text{int} \rightarrow \text{int} \leq X \rightarrow Y \rightarrow T \wedge \text{int} \leq X \wedge \text{int} \leq Y) \quad (2)$$

$$\equiv \exists XY.(X = \text{int} \wedge Y = \text{int} \wedge \text{int} \leq T) \quad (3)$$

$$\equiv \text{int} \leq T \quad (4)$$

$$\equiv \text{let } \Gamma_0 \text{ in } \llbracket \widehat{k_1 + k_2} : T \rrbracket \quad (5)$$

where (1) is by definition of constraint generation; (2) is by definition of Γ_0 , by C-INID and C-IN*; (3) is by C-ARROW and by antisymmetry of subtyping; (4) is by C-EXAND and C-NAME; (5) is again by definition of Γ_0 , by C-INID and C-IN*, and by definition of constraint generation.

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 begin by checking that every value that is well-typed with type `int` is of the form \hat{k} . Indeed, suppose that $\text{let } \Gamma_0; \text{ref } M \text{ in } \llbracket v : \text{int} \rrbracket$ is satisfiable. Then, v cannot be a program variable, for a well-typed value must be closed. 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. v cannot be $\hat{\dagger}$ or $\hat{\dagger} \ v'$, 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, v cannot be a λ -abstraction. Thus, v must be of the form \hat{k} , for it is the only case left.

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 $\text{let } \Gamma_0; \text{ref } M \text{ in } (c \preceq X_1 \rightarrow \dots \rightarrow X_k \rightarrow T \wedge \llbracket v_1 : X_1 \rrbracket \wedge \dots \wedge \llbracket v_k : X_k \rrbracket)$ is satisfiable. We now reason by cases on c .

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

◦ *Case* c is $\hat{\dagger}$. 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 1.6.3, 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{k}_1 and \hat{k}_2 . As a result, the configuration is reducible by R-ADD.

1.9.5 SOLUTION: We must first ensure that R-REF, R-DEREF and R-ASSIGN respect \sqsubseteq (Definition 1.7.5).

◦ *Case* R-REF. The reduction is $\text{ref } v / \emptyset \longrightarrow m / (m \mapsto v)$, where $m \notin \text{fpi}(v)$ (1). Let T be an arbitrary type. According to Definition 1.7.5, the goal is to show that there exist a set of type variables \bar{Y} and a store type M' such that $\bar{Y} \# \text{ftv}(T)$ and $\text{ftv}(M') \subseteq \bar{Y}$ and $\text{dom}(M') = \{m\}$ and let Γ_0 in $\llbracket \text{ref } v : T \rrbracket$

entails $\exists \bar{Y}. \text{let } \Gamma_0; \text{ref } M' \text{ in } \llbracket m/(m \mapsto \bar{v}) : \mathbb{T}/M' \rrbracket$. Now, we have

$$\begin{aligned}
& \text{let } \Gamma_0 \text{ in } \llbracket \text{ref } \bar{v} : \mathbb{T} \rrbracket \\
\equiv & \text{let } \Gamma_0 \text{ in } \exists XY. (Y \rightarrow \text{ref } Y \leq X \rightarrow \mathbb{T} \wedge \llbracket \bar{v} : X \rrbracket) \quad (2) \\
\equiv & \exists Y. \text{let } \Gamma_0 \text{ in } (\text{ref } Y \leq \mathbb{T} \wedge \llbracket \bar{v} : Y \rrbracket) \quad (3) \\
\equiv & \exists Y. \text{let } \Gamma_0; \text{ref } M' \text{ in } (m \preceq \mathbb{T} \wedge \llbracket \bar{v} : M'(m) \rrbracket) \quad (4) \\
\equiv & \exists Y. \text{let } \Gamma_0; \text{ref } M' \text{ in } \llbracket m/(m \mapsto \bar{v}) : \mathbb{T}/M' \rrbracket \quad (5)
\end{aligned}$$

where (2) is by definition of constraint generation and by definition of $\Gamma_0(\text{ref})$; (3) is by C-ARROW, Lemma 1.6.4, and C-INEX; (4) assumes M' is defined as $m \mapsto Y$, and follows from (1), C-INID and C-IN*; and (5) is by definition of constraint generation.

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

$$\begin{aligned}
& \text{let } \Gamma_0; \text{ref } M \text{ in } \llbracket !m/(m \mapsto \bar{v}) : \mathbb{T}/M \rrbracket \\
\equiv & \text{let } \Gamma_0; \text{ref } M \text{ in } \exists XY. (\text{ref } Y \rightarrow Y \leq X \rightarrow \mathbb{T} \wedge m \preceq X \wedge \llbracket \bar{v} : M(m) \rrbracket) \quad (1) \\
\equiv & \text{let } \Gamma_0; \text{ref } M \text{ in } \exists XY. (\text{ref } M(m) \leq X \leq \text{ref } Y \wedge Y \leq \mathbb{T} \wedge \llbracket \bar{v} : M(m) \rrbracket) \quad (2) \\
\equiv & \text{let } \Gamma_0; \text{ref } M \text{ in } \exists Y. (M(m) = Y \wedge Y \leq \mathbb{T} \wedge \llbracket \bar{v} : M(m) \rrbracket) \quad (3) \\
\equiv & \text{let } \Gamma_0; \text{ref } M \text{ in } (M(m) \leq \mathbb{T} \wedge \llbracket \bar{v} : M(m) \rrbracket) \quad (4) \\
\Vdash & \text{let } \Gamma_0; \text{ref } M \text{ in } (\llbracket \bar{v} : \mathbb{T} \rrbracket \wedge \llbracket \bar{v} : M(m) \rrbracket) \quad (5) \\
\equiv & \text{let } \Gamma_0; \text{ref } M \text{ in } \llbracket \bar{v}/(m \mapsto \bar{v}) : \mathbb{T}/M \rrbracket \quad (6)
\end{aligned}$$

where (1) is by definition of constraint generation and by definition of $\Gamma_0(!)$; (2) is by C-ARROW and C-INID; (3) follows from C-EXTRANS and from the fact that ref is an invariant type constructor; (4) is by C-NAMEEQ; (5) is by Lemma 1.6.3 and C-DUP; and (6) is again by definition of constraint generation.

◦ *Case R-ASSIGN.* The reduction is $m := \bar{v}/(m \mapsto \bar{v}_0) \longrightarrow \bar{v}/(m \mapsto \bar{v})$. Let \mathbb{T} be an arbitrary type and let M be a store type of domain $\{m\}$. We have

$$\begin{aligned}
& \text{let } \Gamma_0; \text{ref } M \text{ in } \llbracket m := \bar{v}/(m \mapsto \bar{v}_0) : \mathbb{T}/M \rrbracket \\
\Vdash & \text{let } \Gamma_0; \text{ref } M \text{ in } \llbracket m := \bar{v} : \mathbb{T} \rrbracket \quad (1) \\
\equiv & \text{let } \Gamma_0; \text{ref } M \text{ in } \exists XYZ. (\text{ref } Z \rightarrow Z \rightarrow Z \leq X \rightarrow Y \rightarrow \mathbb{T} \wedge m \preceq X \wedge \llbracket \bar{v} : Y \rrbracket) \quad (2) \\
\equiv & \text{let } \Gamma_0; \text{ref } M \text{ in } \exists XYZ. (\text{ref } M(m) \leq X \leq \text{ref } Z \wedge Z \leq \mathbb{T} \wedge \llbracket \bar{v} : Y \rrbracket \wedge Y \leq Z) \quad (3) \\
\equiv & \text{let } \Gamma_0; \text{ref } M \text{ in } \exists Z. (M(m) = Z \wedge Z \leq \mathbb{T} \wedge \llbracket \bar{v} : Z \rrbracket) \quad (4) \\
\equiv & \text{let } \Gamma_0; \text{ref } M \text{ in } (M(m) \leq \mathbb{T} \wedge \llbracket \bar{v} : M(m) \rrbracket) \quad (5) \\
\Vdash & \text{let } \Gamma_0; \text{ref } M \text{ in } \llbracket \bar{v}/(m \mapsto \bar{v}) : \mathbb{T}/M \rrbracket \quad (6)
\end{aligned}$$

where (1) and (2) are by definition of constraint generation; (3) is by C-ARROW and C-INID; (4) is by C-EXTRANS, Lemma 1.6.4, and from the fact that ref is an invariant type constructor; (5) is by C-NAMEEQ; and (6) 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 1.9.1 for details of a similar proof.

We begin by checking that every value that is well-typed with a type of the form $\text{ref } T$ 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 $\text{let } \Gamma_0; \text{ref } M \text{ in } (c \preceq X_1 \rightarrow \dots \rightarrow X_k \rightarrow T \wedge \llbracket v_1 : X_1 \rrbracket \wedge \dots \wedge \llbracket v_k : X_k \rrbracket)$ 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 $k \geq 1$, because otherwise the expression is a value. Then, 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 1.6.3, 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 $\text{ref } Y$. 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 $\text{ref } v_1 \ \dots \ v_k / \mu$ 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.

1.9.6 SOLUTION: We must first ensure that R-FIX respects \sqsubseteq (Definition 1.7.5). Since the rule is pure, it is sufficient to establish that $\text{let } \Gamma_0 \text{ in } \llbracket \text{fix } v_1 \ v_2 : T \rrbracket$ entails $\text{let } \Gamma_0 \text{ in } \llbracket v_1 (\text{fix } v_1) \ v_2 : T \rrbracket$. Let C stand for the constraint $\text{fix } \preceq ((X \rightarrow Y) \rightarrow (X \rightarrow Y)) \rightarrow X \rightarrow Y \wedge Y \leq T \wedge \llbracket v_1 : (X \rightarrow Y) \rightarrow (X \rightarrow Y) \rrbracket \wedge \llbracket v_2 : X \rrbracket$. We have

$$\begin{aligned}
 & \text{let } \Gamma_0 \text{ in } \llbracket \text{fix } v_1 \ v_2 : T \rrbracket \\
 \equiv & \text{let } \Gamma_0 \text{ in } \exists X_1 X_2. (\text{fix } \preceq X_1 \rightarrow X_2 \rightarrow T \wedge \llbracket v_1 : X_1 \rrbracket \wedge \llbracket v_2 : X_2 \rrbracket) & (1) \\
 \equiv & \text{let } \Gamma_0 \text{ in } \exists X_1 X_2 XY. (((X \rightarrow Y) \rightarrow (X \rightarrow Y)) \rightarrow X \rightarrow Y \leq X_1 \rightarrow X_2 \rightarrow T \\
 & \wedge \llbracket v_1 : X_1 \rrbracket \wedge \llbracket v_2 : X_2 \rrbracket) & (2) \\
 \equiv & \text{let } \Gamma_0 \text{ in } \exists XY. (Y \leq T \wedge \llbracket v_1 : (X \rightarrow Y) \rightarrow (X \rightarrow Y) \rrbracket \wedge \llbracket v_2 : X \rrbracket) & (3) \\
 \equiv & \text{let } \Gamma_0 \text{ in } \exists XY. C & (4)
 \end{aligned}$$

where (1) is by definition of constraint generation; (2) is by definition of $\Gamma_0(\text{fix})$; (3) is by C-ARROW and Lemma 1.6.4; (4) is by definition of $\Gamma_0(\text{fix})$. By Theorem 1.6.2 and WEAKEN, the judgements $C \vdash v_1 : (X \rightarrow Y) \rightarrow$

$(X \rightarrow Y)$ and $C \vdash v_2 : X$ hold. By VAR, WEAKEN, APP, and SUB, it follows that $C \vdash v_1 (\text{fix } v_1) v_2 : T$ holds. By Theorem 1.6.6, this implies $C \Vdash \llbracket v_1 (\text{fix } v_1) v_2 : T \rrbracket$. By congruence of entailment and by C-EX*, (4) entails $\text{let } \Gamma_0 \text{ in } \llbracket v_1 (\text{fix } v_1) v_2 : T \rrbracket$.

Second, we must check that if the configuration $\text{fix } v_1 \dots v_k / \mu$ (where $k \geq 0$) is well-typed, then either it is reducible, or $\text{fix } v_1 \dots v_k$ is a value. This is immediate, for it is a value when $k < 2$, and it is reducible by R-FIX when $k \geq 2$.

We now recall that the construct $\text{letrec } f = \lambda z. t_1 \text{ in } t_2$ provided by ML-the-programming-language may be viewed as syntactic sugar for $\text{let } f = \text{fix } (\lambda f. \lambda z. t_1) \text{ in } t_2$, and set forth to discover the constraint generation rule that arises out of such a definition. We have

$$\begin{aligned} & \text{let } \Gamma_0 \text{ in } \llbracket \text{fix } (\lambda f. \lambda z. t_1) : T \rrbracket \\ \equiv & \text{let } \Gamma_0 \text{ in } \exists Z. (\text{fix } \preceq Z \rightarrow T \wedge \llbracket \lambda f. \lambda z. t_1 : Z \rrbracket) & (1) \\ \equiv & \text{let } \Gamma_0 \text{ in } \exists XY. (X \rightarrow Y \leq T \wedge \llbracket \lambda f. \lambda z. t_1 : (X \rightarrow Y) \rightarrow (X \rightarrow Y) \rrbracket) & (2) \\ \equiv & \text{let } \Gamma_0 \text{ in } \exists XY. (X \rightarrow Y \leq T \wedge \text{let } f : X \rightarrow Y; z : X \text{ in } \llbracket t_1 : Y \rrbracket) & (3) \end{aligned}$$

where (1) is by definition of constraint generation; (2) is by definition of $\Gamma_0(\text{fix})$, by C-ARROW, and by Lemma 1.6.4; and (3) follows from Lemma 1.6.5. This allows us to write

$$\begin{aligned} & \text{let } \Gamma_0 \text{ in } \llbracket \text{let } f = \text{fix } (\lambda f. \lambda z. t_1) \text{ in } t_2 : T \rrbracket \\ \equiv & \text{let } \Gamma_0; f : \forall Z [\llbracket \text{fix } (\lambda f. \lambda z. t_1) : Z \rrbracket]. Z \text{ in } \llbracket t_2 : T \rrbracket & (4) \\ \equiv & \text{let } \Gamma_0; f : \forall Z [\exists XY. (X \rightarrow Y \leq Z \wedge \text{let } f : X \rightarrow Y; z : X \text{ in } \llbracket t_1 : Y \rrbracket)]. Z \text{ in } \llbracket t_2 : T \rrbracket & (5) \\ \equiv & \text{let } \Gamma_0; f : \forall XY [\text{let } f : X \rightarrow Y; z : X \text{ in } \llbracket t_1 : Y \rrbracket]. X \rightarrow Y \text{ in } \llbracket t_2 : T \rrbracket & (6) \end{aligned}$$

where (4) is by definition of constraint generation; (5) follows from C-LETDUP and from the previous series of equivalences; (6) is by C-LETTEX, C-EXTRANS and Lemma 1.3.22.

1.9.21 SOLUTION: We have

$$\begin{aligned} & \llbracket \text{match } t_1 \text{ with } z. t_2 : T \rrbracket \\ \equiv & \text{let } \forall XX' [\llbracket t_1 : X \rrbracket \wedge \text{let } z : X' \text{ in } \llbracket X : z \rrbracket]. (z : X') \text{ in } \llbracket t_2 : T \rrbracket & (1) \\ \equiv & \text{let } z : \forall X' [\exists X. (\llbracket t_1 : X \rrbracket \wedge X \leq X')]. X' \text{ in } \llbracket t_2 : T \rrbracket & (2) \\ \equiv & \text{let } z : \forall X' [\llbracket t_1 : X' \rrbracket]. X' \text{ in } \llbracket t_2 : T \rrbracket & (3) \\ \equiv & \llbracket \text{let } z = t_1 \text{ in } t_2 : T \rrbracket & (4) \end{aligned}$$

where (1) is by definition of constraint generation for match; (2) is by definition of constraint generation for patterns, by C-INID, C-IN*, and C-LETTEX; (3) is by Lemma 1.6.4; (4) is by definition of constraint generation for let.

1.9.26 SOLUTION: The type scheme $\forall \bar{X}. T \rightarrow T$ may be written $\forall \bar{X}. [X \mapsto T](X \rightarrow X)$. Furthermore, $\bar{X} \# \forall X. X \rightarrow X$ holds. Thus, $\forall \bar{X}. T \rightarrow T$ is an instance of $\forall X. X \rightarrow X$

in the sense of DM-INST'. Since DM-INST' is an admissible rule for the type system DM, and since it is clear that the identity function $\lambda z.z$ has type $\forall X.X \rightarrow X$, it must also have type $\forall \bar{X}.T \rightarrow T$. (A more direct proof of this fact would not be difficult.) So, the destructor $(\cdot : \exists \bar{X}.T)$ has not only identity semantics, but also an identity type. This shows that our definitions are sound.

Let us now check requirement (i) of Definition 1.7.6. Since R-ANNOTATION is pure, it suffices to show that $\text{let } \Gamma_0 \text{ in } \llbracket (\nu : \exists \bar{X}.T) : T' \rrbracket$ entails $\text{let } \Gamma_0 \text{ in } \llbracket \nu : T' \rrbracket$. Now, we have

$$\begin{aligned} & \text{let } \Gamma_0 \text{ in } \llbracket (\nu : \exists \bar{X}.T) : T' \rrbracket \\ \equiv & \text{let } \Gamma_0 \text{ in } \exists \bar{X}\bar{X}. (T \rightarrow T \leq X \rightarrow T' \wedge \llbracket \nu : X \rrbracket) & (1) \\ \equiv & \text{let } \Gamma_0 \text{ in } \exists \bar{X}\bar{X}. (X \leq T \leq T' \wedge \llbracket \nu : X \rrbracket) & (2) \\ \Vdash & \text{let } \Gamma_0 \text{ in } \llbracket \nu : T' \rrbracket & (3) \end{aligned}$$

where (1) is by definition of constraint generation and by definition of $\Gamma_0((\cdot : \exists \bar{X}.T))$; (2) is by C-ARROW; and (3) follows from Lemma 1.6.3 and C-EX*.

1.10.5 SOLUTION: We have

$$\begin{aligned} & \text{let } \Gamma_0 \text{ in } \exists Z. \llbracket (\lambda z.z \hat{+} \hat{1} : \forall X.X \rightarrow X) : Z \rrbracket \\ \equiv & \text{let } \Gamma_0 \text{ in } \exists Z. (\forall X. \llbracket \lambda z.z \hat{+} \hat{1} : X \rightarrow X \rrbracket \wedge \exists X. (X \rightarrow X \leq Z)) & (1) \\ \equiv & \text{let } \Gamma_0 \text{ in } \forall X. \text{let } z : X \text{ in } \llbracket z \hat{+} \hat{1} : X \rrbracket & (2) \\ \equiv & \forall X. (\text{int} \rightarrow \text{int} \rightarrow \text{int} \leq X \rightarrow \text{int} \rightarrow X) & (3) \\ \equiv & \forall X. (X = \text{int}) & (4) \\ \equiv & \text{false} & (5) \end{aligned}$$

where (1) is by definition of constraint generation for universal type annotations; (2) is obtained by restricting the scope of $\exists Z$ to the second conjunct, then dropping the latter altogether, since it is equivalent to *true*, and by Lemma 1.6.5; (3) is obtained by definition of constraint generation, by definition of $\Gamma_0(\hat{+})$ and of $\Gamma_0(\hat{1})$, and by a few simple equivalence laws; (4) follows from C-ARROW and antisymmetry of subtyping; (5) follows from the fact that *int* and (say) $\text{int} \rightarrow \text{int}$ have distinct interpretations, since the type constructors *int* and \rightarrow are incompatible. On the other hand, we have

$$\begin{aligned} & \text{let } \Gamma_0 \text{ in } \exists Z. \llbracket (\lambda z.z : \forall X.X \rightarrow X) : Z \rrbracket \\ \equiv & \text{let } \Gamma_0 \text{ in } \forall X. \text{let } z : X \text{ in } \llbracket z : X \rrbracket & (1) \\ \equiv & \forall X. (X \leq X) & (2) \\ \equiv & \text{true} & (3) \end{aligned}$$

where (1) is obtained as above; (2) by definition of constraint generation, C-INID and C-IN*; (3) is by reflexivity of subtyping.

- 1.10.6 SOLUTION: Under the naïve constraint generation rule for universal type variable introduction, the constraint $\llbracket \forall X. (\lambda z. z : X \rightarrow X) : Z \rrbracket$ is equivalent to $\forall X. (\llbracket \lambda z. z : X \rightarrow X \rrbracket \wedge X \rightarrow X \leq Z)$. Since the first conjunct is a tautology, this is in turn equivalent to $\forall X. (X \rightarrow X \leq Z)$. In a nondegenerate free term model where subtyping is interpreted as equality, this constraint is unsatisfiable. In a non-structural subtyping model equipped with a least type \perp and a greatest type \top , it is equivalent to $\perp \rightarrow \top \leq Z$. This is a pretty restrictive constraint: since no value has type \perp , a function whose type is (a supertype of) $\perp \rightarrow \top$ cannot ever be invoked at runtime. This situation is clearly unsatisfactory. Checking that $\forall X. \llbracket \lambda z. z : X \rightarrow X \rrbracket$ holds was indeed part of our intent, but constraining Z to be a supertype of $X \rightarrow X$ for *every* X was not.
- 1.10.7 SOLUTION: Let $\bar{x} \supseteq ftv(\top)$ **(1)** and $\bar{x} \# ftv(\tau)$ **(2)**. We may assume, *w.l.o.g.*, $\bar{x} \# ftv(\top')$ **(3)**. By (1), (2), (3), and by definition of constraint generation for local universal type annotations, $\llbracket (\tau : \forall \bar{x}. \top) : \top' \rrbracket$ is well-defined and is $\forall \bar{x}. \llbracket \tau : \top \rrbracket \wedge \exists \bar{x}. (\top \leq \top')$ **(4)**. By (3) and by definition of constraint generation for introduction of universal type variables and for general type annotations, $\llbracket \forall \bar{x}. (\tau : \top) : \top' \rrbracket$ is $\forall \bar{x}. \exists Z. (\llbracket \tau : \top \rrbracket \wedge \top \leq Z) \wedge \exists \bar{x}. (\llbracket \tau : \top \rrbracket \wedge \top \leq \top')$, where Z is fresh, which we may immediately simplify to $\forall \bar{x}. \llbracket \tau : \top \rrbracket \wedge \exists \bar{x}. (\llbracket \tau : \top \rrbracket \wedge \top \leq \top')$ **(5)**. Using C-EXAND and Lemma 1.10.1, it is straightforward to check that (4) and (5) are equivalent.
- 1.10.9 SOLUTION: We have

$$\begin{aligned} & \exists Z. \llbracket \lambda z. \forall X. (z : X) : Z \rrbracket \\ \Vdash & \exists Z_1 Z_2. \text{let } z : Z_1 \text{ in } \llbracket \forall X. (z : X) : Z_2 \rrbracket & \text{(1)} \\ \Vdash & \exists Z_1. \forall X. (Z_1 \leq X) & \text{(2)} \end{aligned}$$

where (1) is by definition of constraint generation for λ -abstractions, dropping the constraint that relates Z , Z_1 , and Z_2 ; (2) is by definition of constraint generation for universal type variable introduction, this time dropping information about Z_2 . Now, in a nondegenerate equality model, the constraint (2) is equivalent to false. In fact, for (2) to be satisfiable, the interpretation of subtyping must admit a least element \perp . We now see that $\llbracket \lambda z. \forall X. (z : X) : Z \rrbracket$ is a very restrictive constraint. Indeed, it requires z to have every type at once. Because z is λ -bound—hence monomorphic—it must in fact have type \perp . On the other hand, we have

$$\begin{aligned} & \exists Z. \llbracket \forall X. \lambda z. (z : X) : Z \rrbracket \\ \equiv & \forall X. \exists Z. \llbracket \lambda z. (z : X) : Z \rrbracket & \text{(1)} \\ \equiv & \forall X. \exists Z Z_1 Z_2. (Z_1 \leq X \wedge X \leq Z_2 \wedge Z_1 \rightarrow Z_2 \leq Z) & \text{(2)} \\ \equiv & \text{true} & \text{(3)} \end{aligned}$$

where (1) is by definition of constraint generation for universal type variable introduction, dropping the second conjunct, which is entailed by the first; (2) is by Lemma 1.6.5, by definition of constraint generation for general type annotations, and by a few simple equivalence laws; (3) follows from C-NAMEEQ and the witness substitution $[Z_1 \mapsto X, Z_2 \mapsto X, Z \mapsto (X \rightarrow X)]$.

1.10.10 SOLUTION: We have

$$\begin{aligned}
& \llbracket \text{letrec } f : S = \lambda z. t_1 \text{ in } t_2 : T \rrbracket \\
\equiv & \text{let } f : \forall X [\llbracket \text{fix } f : S. \lambda z. t_1 : X \rrbracket]. X \text{ in } \llbracket t_2 : T \rrbracket & (1) \\
\equiv & \text{let } f : \forall X [\text{let } f : S \text{ in } \llbracket \lambda z. t_1 : S \rrbracket \wedge S \preceq X]. X \text{ in } \llbracket t_2 : T \rrbracket & (2) \\
\equiv & \text{let } f : S \text{ in } \llbracket \lambda z. t_1 : S \rrbracket \wedge \text{let } f : \forall X [S \preceq X]. X \text{ in } \llbracket t_2 : T \rrbracket & (3) \\
\equiv & \text{let } f : S \text{ in } (\llbracket \lambda z. t_1 : S \rrbracket \wedge \llbracket t_2 : T \rrbracket) & (4)
\end{aligned}$$

where (1) is by definition of the `letrec` syntactic sugar and by the definition of constraint generation for `let` constructs; we have $X \notin \text{ftv}(S, t_1)$; (2) is by definition of constraint generation for `fix`; (3) is by C-LETAND; (4) follows from the equivalence between the type schemes $\forall X [S \preceq X]. X$ and S —which itself is a direct consequence of C-EXTRANS—and from C-INAND.

1.11.16 SOLUTION: We reason simultaneously in both the subtyping model and the equal-only model, that is, we only rely on properties that are valid in both models.

We must first ensure that rules RD-DEFAULT, RD-FOUND, and RD-FOLLOW respect (Definition 1.7.5).

◦ *Case* RD-DEFAULT. The reduction is $\{v\}.\{\ell\} \xrightarrow{\delta} v$, which is pure. Therefore, it is sufficient to establish that $\text{let } \Gamma_0 \text{ in } \llbracket \{v\}.\{\ell\} : T \rrbracket$ entails $\text{let } \Gamma_0 \text{ in } \llbracket v : T \rrbracket$. In fact, we have:

$$\begin{aligned}
& \text{let } \Gamma_0 \text{ in } \llbracket \{v\}.\{\ell\} : T \rrbracket \\
\equiv & \text{let } \Gamma_0 \text{ in } \exists XY. (\cdot.\{\ell\} \preceq X \rightarrow T \wedge \{\cdot\} \preceq Y \rightarrow X \wedge \llbracket v : Y \rrbracket) & (1) \\
\equiv & \text{let } \Gamma_0 \text{ in } \exists XY. (\exists X_1 X_2. (\Pi(\ell : X_1 ; X_2) \rightarrow X_1 \leq X \rightarrow T) \\
& \quad \wedge \exists Y_1. (Y_1 \rightarrow \Pi(\partial Y_1) \leq Y \rightarrow X) \wedge \llbracket v : Y \rrbracket) \\
\vdash & \text{let } \Gamma_0 \text{ in } \exists X_2 Y. (\partial Y \leq (\ell : X_1 ; X_2) \wedge X_1 \leq X \wedge \llbracket v : Y \rrbracket) & (3) \\
\vdash & \text{let } \Gamma_0 \text{ in } \exists Y. (Y \leq X_1 \wedge X_1 \leq X \wedge \llbracket v : Y \rrbracket) & (4) \\
\vdash & \text{let } \Gamma_0 \text{ in } \llbracket v : T \rrbracket & (5)
\end{aligned}$$

where (1) is by definition of constraint generation; (2) is by definition of Γ_0 , C-INID; (3) by variances of Π , ℓ , and \rightarrow , C-AND, C-EX*, C-EXAND; (4) by C-ROW-DL and covariance of ℓ ; (5) by Lemma 1.6.3.

◦ *Case* RD-FOUND: The reduction is $\{w \text{ with } \ell = v\}.\{\ell\} \xrightarrow{\delta} v$. It suffices to establish $\text{let } \Gamma_0 \text{ in } \llbracket \{w \text{ with } \ell = v\}.\{\ell\} : T \rrbracket$ entails $\text{let } \Gamma_0 \text{ in } \llbracket v : T \rrbracket$. In fact, we

have:

$$\begin{aligned}
& \text{let } \Gamma_0 \text{ in } \llbracket \{\mathbf{w} \text{ with } \ell = \mathbf{v}\}.\{\ell\} : \mathbf{T} \rrbracket \\
\equiv & \text{ let } \Gamma_0 \text{ in } \exists \mathbf{X} \mathbf{Y} \mathbf{Y}'. (\cdot.\{\ell\} \preceq \mathbf{X} \rightarrow \mathbf{T} \wedge \{\cdot \text{ with } \ell = \cdot\} \preceq \mathbf{Y} \rightarrow \mathbf{Y}' \rightarrow \mathbf{X} \wedge \\
& \qquad \qquad \qquad \wedge \llbracket \mathbf{w} : \mathbf{Y} \rrbracket \wedge \llbracket \mathbf{v} : \mathbf{Y}' \rrbracket) \quad (1) \\
\equiv & \text{ let } \Gamma_0 \text{ in } \exists \mathbf{X} \mathbf{Y} \mathbf{Y}'. (\exists \mathbf{X}_1 \mathbf{X}_2. (\Pi(\ell : \mathbf{X}_1 ; \mathbf{X}_2) \rightarrow \mathbf{X}_1 \leq \mathbf{X} \rightarrow \mathbf{T}) \\
& \qquad \qquad \qquad \wedge \exists \mathbf{Y}_1 \mathbf{Y}_2 \mathbf{Y}_3. (\Pi(\ell : \mathbf{Y}_1 ; \mathbf{Y}_3) \rightarrow \mathbf{Y}_2 \rightarrow \Pi(\ell : \mathbf{Y}_2 ; \mathbf{Y}_3) \leq \mathbf{Y} \rightarrow \mathbf{Y}' \rightarrow \mathbf{X}) \\
& \qquad \qquad \qquad \wedge \llbracket \mathbf{w} : \mathbf{Y} \rrbracket \wedge \llbracket \mathbf{v} : \mathbf{Y}' \rrbracket) \quad (2) \\
\Vdash & \text{ let } \Gamma_0 \text{ in } \exists \mathbf{Y}' \mathbf{X}_1 \mathbf{Y}_2. (\mathbf{Y}' \leq \mathbf{Y}_2 \wedge \mathbf{Y}_2 \leq \mathbf{X}_1 \wedge \mathbf{X}_1 \leq \mathbf{T} \wedge \llbracket \mathbf{v} : \mathbf{Y}' \rrbracket) \quad (3) \\
\Vdash & \text{ let } \Gamma_0 \text{ in } \llbracket \mathbf{v} : \mathbf{T} \rrbracket \quad (4)
\end{aligned}$$

where (1) is by definition of constraint generation; (2) is by definition of Γ_0 , C-INID; (3) by variances of Π , ℓ , and \rightarrow , C-AND, C-EX*, C-EXAND; (4) by Lemma 1.6.3.

◦ *Case* RD-FOLLOW The proof is similar to the previous case.

We must now check that if the configuration $F \mathbf{v}_1 \dots \mathbf{v}_k / \mu$ is well-typed, then either it is reducible, or it is a value.

We begin by checking that every value that is well-typed with type $\Pi \mathbf{T}$ is a record value, that is, either of the form $\{\mathbf{v}'\}$ or $\{\mathbf{v}' \text{ with } \ell' = \mathbf{v}'\}$. Indeed, suppose that $\text{let } \Gamma_0 \text{ in } \llbracket \mathbf{v} : \Pi \mathbf{T} \rrbracket$ 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 $\text{ref } M(m) \leq \Pi \mathbf{T}$ would be satisfiable—but the top type constructors ref and Π are incompatible (since Π is isolated); \mathbf{v} cannot be a partial application of a constructor or a primitive, nor a λ -abstraction, since otherwise $\mathbf{T}' \rightarrow \mathbf{T}'' \leq \Pi \mathbf{T}$ would be satisfiable but the top type constructors \rightarrow and Π are incompatible (since they are both isolated); thus \mathbf{v} must either be of the form $\{\mathbf{v}'\}$ or $\{\mathbf{w} \text{ with } \ell = \mathbf{v}'\}$, for these are the only left cases.

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

◦ *Case* \mathbf{c} is $\{\cdot\}$. We may assume $k \geq 2$, since otherwise, the expression is a value. Then $\Gamma_0(\mathbf{c})$ is $\forall \mathbf{X} \mathbf{Y}. \mathbf{X} \rightarrow \Pi(\partial \mathbf{X})$, so by C-INID and C-ARROW the above constraint entails $\exists \mathbf{X}. (\Pi(\partial \mathbf{X}) \leq \mathbf{X}_2 \rightarrow \dots \rightarrow \mathbf{T})$, which by C-CLASS-I entails **false** since \rightarrow and Π are incompatible. Thus, this case cannot occur.

◦ *Case* \mathbf{c} is $\{\cdot \text{ with } \ell = \cdot\}$. Similar to the previous case.

◦ *Case* \mathbf{c} is $\cdot.\{\ell\}$. We may assume $k \geq 1$, since otherwise, the expression is a value. Then $\Gamma_0(\mathbf{c})$ is $\forall \mathbf{X} \mathbf{Y}. \Pi(\ell : \mathbf{X} ; \mathbf{Y}) \rightarrow \mathbf{X}$, so by C-INID and C-ARROW the above constraint entails $\text{let } \Gamma_0; \text{ref } M \text{ in } (\exists \mathbf{X} \mathbf{Y}. (\mathbf{X}_1 \leq \Pi(\ell : \mathbf{X} ; \mathbf{Y})) \wedge \llbracket \mathbf{v}_1 : \mathbf{X}_1 \rrbracket)$, which by lemma 1.6.3 entails $\text{let } \Gamma_0; \text{ref } M \text{ in } \exists \mathbf{X} \mathbf{Y}. \llbracket \mathbf{v}_1 : \Pi(\ell : \mathbf{X} ; \mathbf{Y}) \rrbracket$. Thus \mathbf{v}_1 is a record value, that is, either of the form $\{\mathbf{v}'\}$ and the configuration is reducible

to v' or of the form $\{v'' \text{ with } \ell' = v'\}$ and the configuration is reducible to either v' or $v''.\{\ell\}$.

- 1.11.17 SOLUTION: We add a collection of destructors $.[\ell_1 \leftrightarrow \ell_2]$ of arity 1 for all pairs of distinct labels, with the following semantics:

$$\begin{aligned} \{v\}[\ell_1 \leftrightarrow \ell_2] &\xrightarrow{\delta} v \\ \{w \text{ with } \ell = v\}[\ell_1 \leftrightarrow \ell_2] &\xrightarrow{\delta} \{w[\ell_1 \leftrightarrow \ell_2] \text{ with } \ell = v\} && \text{if } \ell \notin \{\ell_1, \ell_2\} \\ \{w \text{ with } \ell = v\}[\ell_1 \leftrightarrow \ell_2] &\xrightarrow{\delta} \{w[\ell_1 \leftrightarrow \ell_2] \text{ with } \bar{\ell} = v\} && \text{if } \{\ell, \bar{\ell}\} = \{\ell_1, \ell_2\} \end{aligned}$$

The initial environment Γ_0 must be extended with the following typing assumption:

$$.[\ell_1 \leftrightarrow \ell_2] : \forall X_1 X_2 Y. \Pi(\ell_1 : X_1 ; \ell_2 : X_2 ; Y) \rightarrow \Pi(\ell_1 : X_2 ; \ell_2 : X_1 ; Y)$$

We must then check subjection reduction for the new primitive. Since we only added a constructor, it suffices to check progress for the new primitive, that is, that well-typed expressions of the form $.[\ell_1 \leftrightarrow \ell_2]v_1 \dots v_n$ are either value or can be further reduced. Both parts are easy and similar to the corresponding parts in Exercise 1.11.16.

- 1.11.18 SOLUTION: There are several solutions. One of them is to assume a fixed total ordering on row-labels, and to retain as constructors only $\ell^{\kappa, L}$ such that $\ell < L$, that is $\ell < \ell'$ for all $\ell' \in L$; other constants $\ell^{\kappa, L}$ such that $\ell \not< L$ are moved from constructors to the status of destructors with the following collection of reduction rules:

$$\{\{w \text{ with } \ell' = v'\} \text{ with } \ell = v\} \xrightarrow{\delta} \{\{w \text{ with } \ell = v\} \text{ with } \ell' = v'\} \quad \text{(RD-TRANSPOSE)}$$

for all labels ℓ and ℓ' such that $\ell' < \ell$ and

$$\{\{w \text{ with } \ell = v'\} \text{ with } \ell = v\} \xrightarrow{\delta} \{w \text{ with } \ell = v\} \quad \text{(RD-DISCARD)}$$

for all labels ℓ . It is now obvious that values are in normal forms, in the sense that explicit fields are never repeated and are always listed in order. Typing rules need not be changed, so requirement (i) of Definition 1.7.6 still holds. Requirement (ii) need to be checked, in particular, for the new primitives ℓ^L , which we leave to the reader (the proof for $.\{\ell\}$ should hold unchanged).

- 1.11.19 SOLUTION: Let map have type $\Pi(X \rightarrow Y) \rightarrow \Pi(X) \rightarrow \Pi(Y)$, and the following reduction rules in the semantics with normal forms:

$$\begin{aligned} \text{map } \{v' \text{ with } \ell = v\} w &\xrightarrow{\delta} \{\text{map } v' w \text{ with } \ell = v (w.\{\ell\})\} \\ \text{map } v \{w' \text{ with } \ell = w\} &\xrightarrow{\delta} \{\text{map } v w' \text{ with } \ell = (v.\{\ell\}) w\} \\ \text{map } \{v\} \{w\} &\xrightarrow{\delta} \{v w\} \end{aligned}$$

- 1.11.22 SOLUTION: To ensure that the field is not present in the argument of extension, it suffices to restrict its the typing assumptions as follows:

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

To remove an existing field, we can use the following syntactic sugar:

$$\begin{aligned} \cdot \setminus \ell &\stackrel{\text{def}}{=} \lambda v. \{v \text{ with } \ell = \text{abs}\} \\ &: \forall XY. \Pi(\ell : X ; Y) \rightarrow \Pi(\ell : \text{abs} ; Y) \end{aligned}$$

The following weaker typing assumption could also be used to ensure that the field is always present before removal:

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

- 1.11.25 SOLUTION: The proof is similar to 1.11.16 but slightly more complex because we must also check that labels are defined when accessed, and with subtyping.

We reason simultaneously in both the subtyping model or the equal-only model, that is, we only rely on properties that are valid in both models.

We must first ensure that rules RE-FOUND and RE-FOLLOW respect (Definition 1.7.5).

- *Case* RE-FOUND: See Exercise ???. In line ???, field ℓ is pre X_1 instead of X_1 and pre Y_2 instead of Y_2 and step ??? also uses covariance of pre.

- *Case* RE-FOLLOW The proof is similar.

We must then check that if the configuration $F \nu_1 \dots \nu_k / \mu$ is well-typed, then either it is reducible, or it is a value.

We begin by checking that every value that is well-typed with type ΠT is a record value, that is, either of the form $\langle \rangle$ or $\langle v'' \text{ with } \ell' = v' \rangle$. See Exercise 1.11.16.

Next, we note that, according to the constraint generation rules, if the configuration $c \nu_1 \dots \nu_k / \mu$ is well-typed, then a constraint of the form $\text{let } \Gamma_0; \text{ref } M \text{ in } (c \preceq X_1 \rightarrow \dots \rightarrow X_k \rightarrow T \wedge \llbracket \nu_1 : X_1 \rrbracket \wedge \dots \wedge \llbracket \nu_k : X_k \rrbracket)$ is satisfiable. We now reason by cases on c .

- *Case* c is $\langle \rangle$ or $\langle \cdot \text{ with } \ell = \cdot \rangle$. See Exercise 1.11.16.

- *Case* c is $\cdot \langle \ell \rangle$. We may assume $k \geq 1$, since otherwise, the expression is a value. Then $\Gamma_0(c)$ is $\forall XY. \Pi(\ell : \text{pre } X ; Y) \rightarrow X$, so by C-INID and C-ARROW the above constraint entails $\text{let } \Gamma_0; \text{ref } M \text{ in } (\exists XY. (X_1 \leq \Pi(\ell : \text{pre } X ; Y)) \wedge \llbracket \nu_1 : X_1 \rrbracket)$, which by lemma 1.6.3 entails $\text{let } \Gamma_0; \text{ref } M \text{ in } \exists XY. \llbracket \nu_1 : \Pi(\ell : \text{pre } X ; Y) \rrbracket$. Thus ν_1 is a record value, that is, either of the form $\langle \rangle$ or $\langle v'' \text{ with } \ell = v' \rangle$. In fact, the former case cannot occur, since $\text{let } \Gamma_0; \text{ref } M \text{ in } \exists XY. \llbracket \langle \rangle : \Pi(\ell : \text{pre } X ; Y) \rrbracket$ entails $\exists XY \Pi(\partial \text{abs}) \leq \Pi(\ell : \text{pre } X ; Y)$ by C-INID and C-IN*, which in turns

entails $\exists \mathbf{x}. \text{abs} \leq \text{pre } \mathbf{x}$ by C-Row-DL and covariance of Π and ℓ . However, this constraint is equivalent to false, because $\phi(\text{abs}) \leq \phi(\text{pre } \mathbf{x})$ does not hold in any ground assignment ϕ . Thus \mathbf{v}_1 is $\langle \mathbf{v}'' \text{ with } \ell' = \mathbf{v}' \rangle$ and the configuration is reducible to \mathbf{v}' if ℓ' is ℓ or \mathbf{v}'' otherwise.

References

- Martín Abadi, 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.
- Amal Ahmed and David Walker. The logical approach to stack typing. In *ACM Workshop on Types in Compilation*, New Orleans, Louisiana, January 2003.
- Alfred V. Aho, Ravi Sethi, and Jeffrey D. Ullman. *Compilers: Principles, Techniques, and Tools*. Addison Wesley, 1986.
- Alexander Aiken, Manuel Fähndrich, and Raph Levien. Better static memory management: Improving region-based analysis of higher-order languages. In *Programming Language Design and Implementation*, volume 30(6), pages 174–185, 18–21 June 1995. URL <http://www.cs.berkeley.edu/~aiken/publications/papers/pldi95.ps>.
- Alexander Aiken and Edward L. Wimmers. Solving systems of set constraints (extended abstract). In *IEEE Symposium on Logic in Computer Science (LICS)*, pages 329–340, Santa Cruz, California, 22–25 June 1992. IEEE Computer Society Press.
- Alexander Aiken and Edward L. Wimmers. Type inclusion constraints and type inference. In *ACM Symposium on Functional Programming Languages and Computer Architecture (FPCA)*, pages 31–41, 1993.
- Alexander S. Aiken, Edward L. Wimmers, and T. K. Lakshman. Soft typing with conditional types. In *ACM Symposium on Principles of Programming Languages (POPL)*, Portland, Oregon, pages 163–173, January 1994.
- Thorsten Altenkirch. *Constructions, Inductive Types and Strong Normalization*. PhD thesis, LFCS, University of Edinburgh, 1993. URL <http://www.lfcs.informatics.ed.ac.uk/reports/93/ECS-LFCS-93-279/index.html>.
- Roberto M. Amadio and Luca Cardelli. Subtyping recursive types. In *POPL* POPL (a), pages 104–118.

- Torben Amtoft, Flemming Nielson, and Hanne Riis Nielson. *Type and Effect Systems: Behaviours for Concurrency*. IC Press, 1999.
- Andrew W. Appel and Amy P. Felty. A semantic model of types and machine instructions for proof-carrying code. In *ACM Symposium on Principles of Programming Languages (POPL)*, Boston, Massachusetts, pages 243–253, January 2000.
- David Aspinall. Subtyping with Singleton Types. In *Proc. Computer Science Logic (CSL '94)*, 1995. In LNCS 933.
- Lennart Augustsson. Cayenne — a language with dependent types. In *International Conference on Functional Programming (ICFP)*, Baltimore, Maryland, USA, pages 239–250, 1998.
- Henk P. Barendregt. *The Lambda Calculus*. North Holland, revised edition, 1984.
- Henk P. Barendregt. Introduction to generalized type systems. *Journal of Functional Programming*, 1(2):125–154, 1991.
- Henk P. Barendregt. Lambda calculi with types. In Abramsky, Gabbay, and Maibaum, editors, *Handbook of Logic in Computer Science*, volume II. Oxford University Press, 1992.
- Erik Barendsen and Sjaak Smetsers. Conventional and uniqueness typing in graph rewrite systems. In Shyamasundar, editor, *Foundations of Software Technology and Theoretical Computer Science*, number 761 in LNCS, pages 41–51, Bombay, India, 1993. Springer-Verlag.
- Bruno Barras, 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 (Institut National de Recherche en Informatique et en Automatique), France, 1997.
- Lujo Bauer, Andrew W. Appel, and Edward W. Felten. Mechanisms for secure modular programming in java. Technical Report TR-603-99, 1999. URL citeseer.nj.nec.com/bauer99mechanisms.html.
- Stephan Bellantoni and Stephan Cook. A new recursion-theoretic characterization of polytime functions. *Computational Complexity*, 2(2):97–110, 1992.
- Stephan Bellantoni, K.-H. Niggl, and H. Schwichtenberg. Higher type recursion, ramification and polynomial time. *Annals of Pure and Applied Logic*, 104:17–30, 2000.
- Stefano Berardi. 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.
- Bernard Berthomieu. 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.

- G. M. Bierman, A. M. Pitts, and C. V. Russo. Operational properties of Lily, a polymorphic linear lambda calculus with recursion. In *Fourth International Workshop on Higher Order Operational Techniques in Semantics, Montréal*, volume 41 of *Electronic Notes in Theoretical Computer Science*. Elsevier, September 2000. URL <http://www.elsevier.nl/locate/entcs/volume41.html>.
- L. Birkedal and R. W. Harper. Constructing interpretations of recursive types in an operational setting. *Information and Computation*, 155:3–63, 1999.
- Lars Birkedal and Mads Tofte. A constraint-based region inference algorithm. *Theoretical Computer Science*, 258:299–392, 2001. URL <http://www.it-c.dk/people/birkedal/papers/conria.ps.gz>.
- Lars Birkedal, Mads Tofte, and Magnus Vejlstrup. From region inference to von Neumann machines via region representation inference. In *Principles of Programming Languages*, pages 171–183, New York, NY, USA, 21–24 January 1996. ACM Press. ISBN 0-89791-769-3. URL <http://www.it-c.dk/people/birkedal/papers/reginm.ps.gz>.
- Matthias Blume. *The SML/NJ Compilation and Library Manager*, May 2002. URL <http://www.smlnj.org/doc/CM/index.html>.
- Matthias Blume and Andrew W. Appel. Hierarchical modularity. *ACM Transactions on Programming Languages and Systems*, 21(4):813–847, 1999. URL citeseer.nj.nec.com/blume98hierarchical.html.
- Kim B. Bruce. Typing in object-oriented languages: Achieving expressibility and safety, 1995. Available through <http://www.cs.williams.edu/simkim>.
- Kim B. Bruce, Luca Cardelli, and Benjamin C. Pierce. Comparing object encodings. In *Theoretical Aspects of Computer Software (TACS), Sendai, Japan*, 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.
- T. H. Brus, M. C. J. D. van Eekelen, M. O. van Leer, and M. J. Plasmeijer. Clean: A language for functional graph rewriting. In G. Kahn, editor, *Functional Programming Languages and Computer Architecture*, pages 364–384. Springer-Verlag, Berlin, DE, 1987. ISBN 3-540-18317-5. Lecture Notes in Computer Science 274; Proceedings of Conference held at Portland, OR.
- Rod Burstall and Butler Lampson. A kernel language for abstract data types and modules. In G. Kahn, D. MacQueen, and G. Plotkin, editors, *Semantics of Data Types*, volume 173 of *Lecture Notes in Computer Science*, pages 1–50. Springer-Verlag, 1984.
- Rod Burstall, David MacQueen, and Donald Sannella. HOPE: An experimental applicative language. In *Proceedings of the 1980 LISP Conference*, pages 136–143, Stanford, California, 1980. Stanford University.
- Cristiano Calcagno. Stratified operational semantics for safety and correctness of region calculus. In *POPL* POPL (b), pages 155–165. ISBN 1-58113-336-7. URL [ftp://ftp.disi.unige.it/person/CalcagnoC/regions.ps](http://ftp.disi.unige.it/person/CalcagnoC/regions.ps).

- Cristiano Calcagno, Simon Helsen, and Peter Thiemann. Syntactic type soundness results for the region calculus. *Information & Computation*, 173(2):199–221, 2002. URL <http://www.informatik.uni-freiburg.de/~helsen/calcagno-helsen-thiemann-iandc-2001.ps.gz>.
- Luca Cardelli. A polymorphic λ -calculus with Type:Type. Research report 10, DEC/Compaq Systems Research Center, May 1986.
- Luca Cardelli. Phase distinctions in type theory. unpublished manuscript, 1988a.
- Luca Cardelli. Typechecking dependent types and subtypes. In M. Boscarol, L. Carlucci Aiello, and G. Levi, editors, *Foundations of Logic and Functional Programming, Workshop Proceedings, Trento, Italy, (Dec. 1986)*, volume 306 of *Lecture Notes in Computer Science*, pages 45–57. Springer-Verlag, 1988b.
- Luca Cardelli. Extensible records in a pure calculus of subtyping. Research report 81, DEC/Compaq Systems Research Center, January 1992. Also in Carl A. Gunter and John C. Mitchell, editors, *Theoretical Aspects of Object-Oriented Programming: Types, Semantics, and Language Design*, MIT Press, 1994.
- Luca Cardelli. Program fragments, linking, and modularization. In *Conference Record of POPL'97: The 24th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 266–277, Paris, France, January 1997. ACM Press.
- Luca Cardelli, Jim Donahue, Mick Jordan, Bill Kalso, and Greg Nelson. The modular 3 type system. In *Sixteenth ACM Symposium on Principles of Programming Languages (POPL)*, pages 202–212, Austin, TX, January 1989.
- Luca Cardelli and Xavier Leroy. Abstract types and the dot notation. Technical Report 56, Digital Equipment Corporation Systems Research Center, Palo Alto, CA, March 1990.
- Luca Cardelli 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, June 1990. Also available as DEC/Compaq SRC Research Report 55, Feb. 1990.
- Luca Cardelli and John Mitchell. Operations on records. *Mathematical Structures in Computer Science*, 1:3–48, 1991. Also in Carl A. Gunter and John 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 MFPS '89, Springer LNCS volume 442.
- J. Cartmell. Generalised algebraic theories and contextual categories. *Annals of Pure and Applied Logic*, 32:209–243, 1986.
- R. Cartwright and M. Fagan. Soft typing. In *Proceedings of the SIGPLAN '91 Conference on Programming Language Design and Implementation*, pages 278–292, June 1991. Also available as SIGPLAN Notices 26(6) June 1991.

- Iliano Cervesato and Frank Pfenning. A linear logical framework. In *Information and Computation*, July 2000. To appear.
- Sagar Chaki, Sriram K. Rajamani, and Jakob Rehof. Types as models: Model checking message-passing programs. In *Principles of Programming Languages*, pages 45–57, New York, NY, USA, 16–18 January 2002. ACM Press. ISBN 1-58113-450-9.
- Jawahar Chirimar, 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.
- Alonzo Church. The calculi of lambda-conversion. *The Annals of Mathematical Studies*, 6, 1941.
- Alonzo Church. The weak theory of implication. *Kontrolliertes Denken: Untersuchungen zum Logikkalkül und zur Logik der Einzelwissenschaften*, pages 22–37, 1951.
- Dominique Clément, Joëlle Despeyroux, Thierry Despeyroux, and Gilles Kahn. A simple applicative language: Mini-ML. In *ACM Symposium on Lisp and Functional Programming (LFP)*, pages 13–27, August 1986.
- Christopher Colby, 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. ISSN 0362-1340.
- Hubert Comon. Constraints in term algebras (short survey). In *Conference on Algebraic Methodology and Software Technology (AMAST)*, Workshops in Computing. Springer-Verlag, 1993.
- Hubert Comon and Pierre Lescanne. Equational problems and disunification. *Journal of Symbolic Computation*, 7:371–425, 1989.
- R. L. Constable, S.F. Allen, H.M. Bromley, W.R. Cleaveland, J.F. Cremer, R.W. Harper, D.J. Howe, T.B. Knoblock, N.P. Mendler, P. Panangaden, J.T. Sasaki, and S.F. Smith. *Implementing Mathematics with the Nuprl Development System*. Prentice-Hall, NJ, 1986. URL <http://www.nuprl.org/book/doc.html>.
- Catarina Coquand. The AGDA proof system homepage, 1998. At <http://www.cs.chalmers.se/~catarina/agda/>.
- Thierry Coquand. An analysis of Girard’s paradox. In *Proceedings, Symposium on Logic in Computer Science IEE* (1986), pages 227–236.
- Thierry Coquand. An algorithm for testing conversion in type theory. In Gérard Huet and Gordon Plotkin, editors, *Logical Frameworks*, pages 255–279. Cambridge University Press, 1991a.
- Thierry Coquand. An Algorithm for Testing Conversion in Type Theory. In Gérard Huet and G. Plotkin, editors, *Logical frameworks*, pages 255–277. Cambridge University Press, 1991b.

- Thierry Coquand. Pattern matching with dependent types. In *Proceedings of the Workshop on Types for Proofs and Programs, Baastad*. Informal proceedings available by ftp from `ftp://ftp.cs.chalmers.se/pub/cs-reports/baastad.92/proc.ps.Z`, 1992.
- Thierry Coquand and Gérard Huet. The Calculus of Constructions. *Information and Computation*, 76(2/3):95–120, February/March 1988.
- Erik Crank and Matthias Felleisen. Parameter-passing and the lambda calculus. In *ACM Symposium on Principles of Programming Languages (POPL)*, pages 233–244, January 1991.
- Karl Crary. Toward a foundational typed assembly language. In *ACM Symposium on Principles of Programming Languages (POPL)*, New Orleans, Louisiana, pages 198–212, January 2003.
- Karl Crary, Robert Harper, and Sidd Puri. What is a recursive module? In *SIGPLAN '99 Conference on Programming Language Design and Implementation (PLDI)*, pages 50–63, Atlanta, GA, 1999a. ACM SIGPLAN.
- Karl Crary, David Walker, and Greg Morrisett. Typed memory management in a calculus of capabilities. In *ACM Symposium on Principles of Programming Languages (POPL)*, pages 262–275, January 1999b.
- Karl Crary, Stephanie Weirich, and J. Gregory Morrisett. Intensional polymorphism in type-erasure semantics. In *International Conference on Functional Programming (ICFP)*, Baltimore, Maryland, USA, pages 301–312, 1998.
- Pavel Curtis. *Constrained Quantification in Polymorphic Type Analysis*. PhD thesis, Cornell University, February 1990.
- Luis Damas 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.
- Nicolas G. de Bruijn. 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.
- Robert DeLine 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 56–69, June 2001.
- Robert DeLine and Manuel Fähndrich. Enforcing high-level protocols in low-level software. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, pages 59–69, June 2001.
- James Donahue and Alan Demers. Data types are values. *ACM Transactions on Programming Languages and Systems*, 7(3):426–445, July 1985.
- Kosta Došen and Peter Schroeder-Heister, editors. *Substructural Logics*, chapter A historical introduction to substructural logics, pages 1–30. Oxford University Press, 1993.

- Gilles Dowek, Thérèse Hardin, and Claude Kirchner. Higher order unification via explicit substitutions. Research Report 2709, INRIA, November 1995.
- Gilles Dowek, Thérèse Hardin, Claude Kirchner, and Frank Pfenning. Unification via explicit substitutions: the case of higher-order patterns. Research Report 3591, INRIA, December 1998.
- Derek Dreyer, Karl Cray, and Robert Harper. A type system for higher-order modules. In *POPL 2003: Proceedings of the 30th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 236–249, New Orleans, January 2003.
- Dirk Dussart, Fritz Henglein, and Christian Mossin. Polymorphic recursion and subtype qualifications: Polymorphic binding-time analysis in polynomial time. In Alan Mycroft, editor, *Static Analysis Symposium (SAS)*, volume 983 of *Lecture Notes in Computer Science*, pages 118–135. Springer-Verlag, September 1995a.
- Dirk Dussart, Fritz Henglein, and Christian Mossin. Polymorphic recursion and subtype qualifications: Polymorphic binding-time analysis in polynomial time. In Alan Mycroft, editor, *Static Analysis Symposium*, volume 983 of *Lecture Notes in Computer Science*, pages 118–135, Heidelberg, Germany, 25–27 September 1995b. Springer-Verlag. ISBN 3-540-60360-3.
- Thomas Erhard. A categorical semantics of constructions. In *Symposium on Logic in Computer Science*, pages 264–273, July 1988.
- Manuel Fähndrich, Jakob Rehof, and Manuvir Das. Scalable context-sensitive flow analysis using instantiation constraints. In *Programming Language Design and Implementation*, volume 35(5) of *SIGPLAN Notices*, pages 253–263, New York, NY, USA, 18–21 June 2000. ACM Press. ISBN 1-58113-199-2.
- M. Felleisen and R. Hieb. The revised report on the syntactic theories of sequential control and state. *Theoretical Computer Science*, 103:235–271, 1992.
- Manuel Fähndrich. BANE: *A Library for Scalable Constraint-Based Program Analysis*. PhD thesis, University of California at Berkeley, 1999.
- Manuel Fähndrich, 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*, June 2000.
- Matthew Flatt and Matthias Felleisen. Units: Cool modules for HOT languages. In *Proceedings of the ACM SIGPLAN '98 Conference on Programming Language Design and Implementation*, pages 236–248, 1998. URL citeseer.nj.nec.com/flatt98unit.html.
- Matthew Fluet and Riccardo Pucella. Phantom types and subtyping. pages 448–460, August 2002.
- Cédric Fournet and Georges Gonthier. The reflexive chemical abstract machine and the join-calculus. In *Principles of Programming Languages*, January 1996.

- Alexandre Frey. Satisfying subtype inequalities in polynomial space. In Pascal Van Hentenryck, editor, *International Symposium on Static Analysis (SAS)*, number 1302 in Lecture Notes in Computer Science, pages 265–277. Springer-Verlag, September 1997.
- You-Chin Fuh and Prateek Mishra. Type inference with subtypes. In H. Ganzinger, editor, *European Symp. on Programming (ESOP)*, volume 300 of *Lecture Notes in Computer Science*, pages 94–114. Springer-Verlag, 1988.
- Murdoch J. Gabbay and Andrew M. Pitts. A new approach to abstract syntax with variable binding. *Formal Aspects of Computing*, 13(3–5):341–363, July 2002.
- Jacques Garrigue. Relaxing the value restriction. Draft., November 2002.
- Jacques Garrigue and Didier Rémy. Extending ML with semi-explicit higher-order polymorphism. *Information and Computation*, 155(1):134–169, 1999.
- David Gay and Alexander Aiken. Language support for regions. In *Programming Language Design and Implementation*, volume 36(5) of *SIGPLAN Notices*, pages 70–80, New York, NY, USA, 20–22 June 2001. ACM Press. ISBN 1-58113-414-2. URL <http://www.cs.berkeley.edu/~dgay/papers/pldi01.ps>.
- Giorgio Ghelli and Benjamin Pierce. Bounded existentials and minimal typing. *Theoretical Computer Science*, 193:75–96, 1998. TR version originally announced on Types list in Summer 1992.
- David K. Gifford and John M. Lucassen. Integrating functional and imperative programming. In *LISP and Functional Programming*, pages 28–38, New York, NY, USA, 4–6 August 1986. ACM Press.
- J.-Y. Girard. *Interprétation fonctionnelle et élimination des coupures dans l'arithmétique d'ordre supérieur*. PhD thesis, Université Paris VII, 1972a. Thèse de doctorat d'état.
- Jean-Yves Girard. *Interprétation Fonctionnelle et Élimination des Coupures dans l'Arithmétique d'Ordre Supérieure*. PhD thesis, Université Paris VII, 1972b.
- Jean-Yves Girard. Linear logic. *Theoretical Computer Science*, 50:1–102, 1987.
- Jean-Yves Girard. Light linear logic. *Information and Computation*, 143, 1998.
- Neal Glew. Type dispatch for named hierarchical types. In *International Conference on Functional Programming (ICFP), Paris, France*, pages 172–182, 1999.
- GNU. GNU C library, version 2.2.5, 2001. URL http://www.gnu.org/manual/glibc-2.2.5/html_mono/libc.html.
- Healdene Goguen. *A Typed Operational Semantics for Type Theory*. PhD thesis, LFCS, University of Edinburgh, 1994. URL <http://www.lfcs.informatics.ed.ac.uk/reports/94/ECS-LFCS-94-304/index.html>. Report ESC-LFCS-94-304.
- A. D. Gordon. *Functional Programming and Input/Output*. Distinguished Dissertations in Computer Science. Cambridge University Press, 1994.

- A. D. Gordon. Bisimilarity as a theory of functional programming. In *Eleventh Conference on the Mathematical Foundations of Programming Semantics, New Orleans, 1995*, volume 1 of *Electronic Notes in Theoretical Computer Science*. Elsevier, 1995.
- A. D. Gordon. 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.
- Andrew D. Gordon and Alan Jeffrey. Authenticity by typing for security protocols. In *Proc. 14th IEEE Computer Security Foundations Workshop (CSFW 2001), Cape Breton*, pages 145–159, 2001a.
- Andrew D. Gordon and Alan Jeffrey. Typing correspondence assertions for communication protocols. *Electronic Notes in Theoretical Computer Science*, 45:22 pages, 2001b. <http://www.elsevier.nl/locate/entcs/volume45.html>.
- Andrew D. Gordon and Alan Jeffrey. Types and effects for asymmetric cryptographic protocols. In *Proc. 15th IEEE Computer Security Foundations Workshop (CSFW 2002), Cape Breton*, pages 77–91, 2002.
- Michael Gordon, Robin Milner, and Christopher Wadsworth. *Edinburgh LCF: A Mechanized Logic of Computation*, volume 78 of *Lecture Notes in Computer Science*. Springer-Verlag, 1979a.
- Michael J. Gordon, Robin Milner, and Christopher P. Wadsworth. *Edinburgh LCF*. Springer-Verlag LNCS 78, 1979b.
- Dan Grossman, 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)*, pages 282–293, June 2002a.
- Dan Grossman, Greg Morrisett, Trevor Jim, Michael Hicks, Yanling Wang, and James Cheney. Region-based memory management in Cyclone. In *Proceedings of the ACM SIGPLAN 2002 Conference on Programming language design and implementation (PLDI'02)*, pages 282–293. ACM Press, 2002b.
- Jörgen Gustavsson and Josef Svenningsson. Constraint abstractions. In Olivier Danvy and Andrzej Filinski, editors, *Programs as Data Objects*, volume 2053 of *LNCS*, pages 63–83, Heidelberg, Germany, 21–23 May 2001a. Springer-Verlag. ISBN 3-540-42068-1.
- Jörgen Gustavsson and Josef Svenningsson. Constraint abstractions. In *Symposium on Programs as Data Objects*, volume 2053 of *Lecture Notes in Computer Science*. Springer-Verlag, May 2001b.
- Jr. Guy L. Steele. *Common Lisp the Language*. Digital Press, 1990.
- Niels Hallenberg, Martin Elsman, and Mads Tofte. Combining region inference and garbage collection. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'02)*. ACM Press, June 2002. Berlin, Germany.

- Thomas Hallgren and Aarne Ranta. An extensible proof text editor (abstract). In *Logic for Programming and Automated Reasoning (LPAR'2000)*, pages 70–84. Springer-Verlag LNCS/LNAI. 1955, 2000.
- Nadeem Hamid, 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.
- David R. Hanson. Fast allocation and deallocation of memory based on object lifetimes. *Software—Practice and Experience*, 20(1):5–12, 1990.
- R. Harper, F. Honsell, and G. Plotkin. A framework for defining logics. *JACM*, 40(1):143–184, 1993a. URL <http://www.cs.cmu.edu/~fp/elf-papers/jacm93.dvi.gz>.
- Robert Harper. On equivalence and canonical forms in the LF type theory. (Submitted for publication.), August 2002.
- Robert Harper, Furio Honsell, and Gordon Plotkin. A framework for defining logics. *Journal of the ACM*, 40(1):143–184, 1993b. Summary in LICS'87.
- Robert Harper and Mark Lillibridge. A type-theoretic approach to higher-order modules with sharing. In *ACM Symposium on Principles of Programming Languages (POPL), Portland, Oregon*, pages 123–137, January 1994.
- Robert Harper and John C. Mitchell. On the type structure of Standard ML. *ACM Transactions on Programming Languages and Systems*, 15(2):211–252, April 1993.
- Robert Harper, John C. Mitchell, and Eugenio Moggi. Higher-order modules and the phase distinction. In *Seventeenth ACM Symposium on Principles of Programming Languages (POPL)*, San Francisco, CA, January 1990a.
- Robert Harper, 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 1990b.
- Robert Harper and Benjamin Pierce. A record calculus based on symmetric concatenation. In *ACM Symposium on Principles of Programming Languages (POPL), Orlando, Florida*, pages 131–142, January 1991. Extended version available as Carnegie Mellon Technical Report CMU-CS-90-157.
- Robert Harper and Robert Pollack. Type checking with universes. *Theoretical Computer Science*, 89:107–136, 1991.
- Robert Harper and Chris Stone. A type-theoretic interpretation of Standard ML. In Gordon Plotkin, Colin Stirling, and Mads Tofte, editors, *Proof, Language, and Interaction: Essays in Honor of Robin Milner*. MIT Press, 2000a.
- Robert Harper and Christopher Stone. A type-theoretic interpretation of Standard ML. In Gordon Plotkin, Colin Stirling, and Mads Tofte, editors, *Proof, Language and Interaction: Essays in Honour of Robin Milner*. MIT Press, 2000b.
- Nevin Heintze. Set based analysis of ML programs. Technical Report CMU-CS-93-193, Carnegie Mellon University, School of Computer Science, July 1993.

- Simon Helsen and Peter Thiemann. Syntactic type soundness for the region calculus. In Alan Jeffrey, editor, *ACM Workshop on Higher Order Operational Techniques in Semantics*, volume 41(3) of *Electronic Notes in Theoretical Computer Science*, pages 1–20. Elsevier, September 2000. URL <http://www.elsevier.nl/locate/entcs/volume41.html>.
- Fritz Henglein. Type inference with polymorphic recursion. *ACM Transactions on Programming Languages and Systems*, 15(2):253–289, 1993.
- Fritz Henglein, Henning Makholm, and Henning Niss. A direct approach to control-flow sensitive region-based memory management. In *Proceedings of the 3rd International ACM SIGPLAN Conference on Principles and Practice of Declarative Programming (PPDP)*, pages 175–186, Firenze, Italy, September 2001. ACM Press. URL <http://www.diku.dk/~hniss/publications/ppdp2001-abstract.html>.
- Fritz Henglein and Christian Mossin. Polymorphic binding-time analysis. In Donald Sannella, editor, *5th European Symposium on Programming*, volume 788 of *Lecture Notes in Computer Science*, pages 287–301, Heidelberg, Germany, 11–13 April 1994. Springer-Verlag. ISBN 3-540-57880-3.
- Tom Hirschowitz and Xavier Leroy. Mixin modules in a call-by-value setting. In *European Symposium on Programming*, pages 6–20, 2002. URL citeseer.nj.nec.com/article/hirschowitz02mixin.html.
- C. A. R. Hoare. Proof of correctness of data representation. *Acta Informatica*, 1: 271–281, 1972.
- Martin Hofmann. Syntax and semantics of dependent types. In P. Dybjer and A. Pitts, editors, *Semantics of Logics of Computation*, chapter 3. Cambridge University Press, 1997. URL <http://www.mathematik.th-darmstadt.de/~mh/cupart.dvi.gz>.
- Martin Hofmann. Linear types and non-size-increasing polynomial time computation. In *Logic in computer science*, pages 464–473, Los Alamitos, CA, June 1999.
- Martin Hofmann. Safe recursion with higher types and BCK-algebra. *Annals of Pure and Applied Logic*, 2000.
- F. Honsell, I. A. Mason, S. F. Smith, and C. L. Talcott. A variable typed logic of effects. *Information and Computation*, 119(1):55–90, 1995.
- William A. Howard. Hereditarily majorizable functionals of finite type. In Anne Sjerp Troelstra, editor, *Metamathematical Investigation of Intuitionistic Arithmetic and Analysis*, volume 344 of *Lecture Notes in Mathematics*, pages 454–461. Springer-Verlag, Berlin, 1973. Appendix.
- William A. Howard. 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, New York, 1980. Reprint of 1969 article.

- D. J. Howe. Proving congruence of bisimulation in functional programming languages. *Information and Computation*, 124(2):103–112, 1996.
- Paul Hudak, S. Peyton Jones, P. Wadler, B. Boutel, J. Fairbairn, J. Fasel, M. M. Guzman, K. Hammond, J. Hughes, T. Johnsson, D. Kieburtz, R. Nikhil, W. Par-tain, and J. Peterson. Report on the programming language Haskell, version 1.2. *SIGPLAN Notices*, 27(5), May 1992.
- Gérard Huet. *Résolution d'équations dans les langages d'ordre 1,2, ..., ω* . Thèse de Doctorat d'Etat, Université de Paris 7 (France), 1976.
- Proceedings, Symposium on Logic in Computer Science*, Cambridge, Massachusetts, 16–18 June 1986. IEEE Computer Society.
- Atsushi Igarashi and Naoki Kobayashi. A generic type system for the pi-calculus. In POPL POPL (b), pages 128–141. ISBN 1-58113-336-7.
- Atsushi Igarashi and Benjamin C. Pierce. Foundations for virtual types. In *European Conference on Object-Oriented Programming (ECOOP)*, 1999. Also in informal proceedings of the Sixth International Workshop on Foundations of Object-Oriented Languages (FOOL). Full version to appear in *Information and Computation*.
- Bart Jacobs. *Categorical Logic and Type Theory*. Studies in Logic and the Foundations of Mathematics 141. North Holland, Elsevier, 1999.
- Lalita A. Jategaonkar and John C. Mitchell. ML with extended pattern matching and subtypes (preliminary version). In *Proceedings of the ACM Conference on Lisp and Functional Programming*, pages 198–211, Snowbird, Utah, July 1988.
- Kathleen Jensen and Niklaus Wirth. *Pascal User Manual and Report*. Springer-Verlag, second edition, 1975.
- Thomas Jensen. Inference of polymorphic and conditional strictness properties. In *ACM Symposium on Principles of Programming Languages (POPL)*, pages 209–221. ACM Press, January 1998. <http://www.irisa.fr/lande/jensen/papers/pop198.ps>.
- Trevor Jim. What are principal typings and what are they good for? Technical Report MIT/LCS TM-532, Massachusetts Institute of Technology, August 1995.
- Trevor Jim, J. Greg Morrisett, Dan Grossman, Michael W. Hicks, James Cheney, and Yanling Wang. Cyclone: A safe dialect of C. In *Proceedings of the General Track: 2002 USENIX Annual Technical Conference*, pages 275–288. USENIX Association, 2002.
- P. Johann. A generalization of short-cut fusion and its correctness proof. *Higher-Order and Symbolic Computation*, 15(4):273–300, 2002.
- Mark P. Jones. *Qualified Types: Theory and Practice*. Cambridge University Press, November 1994a.
- Mark P. Jones. *Qualified Types: Theory and Practice*. Cambridge University Press, Cambridge, England, 1994b.

- Mark P. Jones. Using parameterized signatures to express modular structure. In *Conference Record of the 23rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL'96)*, St. Petersburg, Florida, 21–24, 1996. ACM Press. URL citeseer.nj.nec.com/jones96using.html.
- Mark P. Jones. Typing haskell in haskell. In *Haskell Workshop*, October 1999.
- Jean-Pierre Jouannaud and Claude Kirchner. Solving equations in abstract algebras: a rule-based survey of unification. In Jean-Louis Lassez and Gordon Plotkin, editors, *Computational Logic. Essays in honor of Alan Robinson*, chapter 8, pages 257–321. MIT Press, 1991.
- Pierre Jouvelot and David K. Gifford. Reasoning about continuations with control effects. In *Programming Language Design and Implementation*, volume 24(7), pages 218–226, 21–23 June 1989.
- Pierre Jouvelot and David K. Gifford. Algebraic reconstruction of types and effects. In *POPL POPL (a)*, pages 303–310.
- A. Jung and A. Stoughton. Studying the fully abstract model of PCF within its continuous function model. In M. Bezem and J.M. Groote, editors, *Proc. Typed Lambda Calculi and Applications (TLCA)*, volume 664 of *Lecture Notes in Computer Science*, pages 230–244. Springer-Verlag, 1993.
- L.S. van Benthem Jutting, James McKinna, and Robert Pollack. Checking algorithms for Pure Type Systems. In Henk Barendregt and Tobias Nipkow, editors, *Proceedings of the International Workshop on Types for Proofs and Programs*, pages 19–61, Nijmegen, The Netherlands, May 1994. Springer-Verlag LNCS 806.
- Assaf J. Kfoury, Jerzy Tiuryn, and Pawel Urzyczyn. Type reconstruction in the presence of polymorphic recursion. *ACM Transactions on Programming Languages and Systems*, 15(2):290–311, April 1993.
- Claude Kirchner and F. Klay. Syntactic theories and unification. In *Proceedings 5th IEEE Symposium on Logic in Computer Science, Philadelphia (Pa., USA)*, pages 270–277, June 1990.
- Dexter Kozen, Jens Palsberg, and Michael I. Schwartzbach. Efficient recursive subtyping. *Mathematical Structures in Computer Science*, 5(1):113–125, 1995.
- Viktor Kuncak and Martin Rinard. Structural subtyping of non-recursive types is decidable. In *IEEE Symposium on Logic in Computer Science (LICS)*, June 2003.
- Yves Lafont. The linear abstract machine. *Theoretical Computer Science*, 59:157–180, 1988.
- Joachim Lambek. The mathematics of sentence structure. *American Mathematical Monthly*, 65:154–170, 1958.
- B. Lampson and R. Burstall. Pebble, a kernel language for modules and abstract data types. *Information and Computation*, 76:278–346, February/March 1988.
- S. B. Lassen. *Relational Reasoning about Functions and Nondeterminism*. PhD thesis, Department of Computer Science, University of Aarhus, 1998.

- Jean-Louis Lassez, Michael J. Maher, and Kim G. Marriott. Unification revisited. In Jack Minker, editor, *Foundations of Deductive Databases and Logic Programming*, chapter 15, pages 587–625. Morgan Kaufmann, 1988.
- Konstantin Läufer and Martin Odersky. Polymorphic type inference and abstract data types. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 16(5):1411–1430, September 1994. Summary in *Phoenix Seminar and Workshop on Declarative Programming*, Nov. 1991.
- Fabrice Le Fessant and Luc Maranget. Optimizing pattern matching. In *International Conference on Functional Programming (ICFP), Firenze, Italy*. ACM Press, 2001.
- Oukseh Lee and Kwangkeun Yi. Proofs about a folklore let-polymorphic type inference algorithm. *ACM Transactions on Programming Languages and Systems*, 20(4):707–723, 1998.
- Daniel Leivant. Stratified functional programs and computational complexity. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 325–333, jan 1993.
- Xavier Leroy. Polymorphic typing of an algorithmic language. Research Report 1778, INRIA, October 1992.
- Xavier Leroy. Manifest types, modules, and separate compilation. In *Proceedings of the Twenty-first Annual ACM Symposium on Principles of Programming Languages, Portland*. ACM, January 1994a.
- Xavier Leroy. Manifest types, modules and separate compilation. In *ACM Symposium on Principles of Programming Languages (POPL), Portland, Oregon*, pages 109–122, January 1994b.
- Xavier Leroy. Applicative functors and fully transparent higher-order modules. In *Conference Record of POPL '95: ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 142–153, San Francisco, CA, January 1995a.
- Xavier Leroy. Applicative functors and fully transparent higher-order modules. In *Proceedings of the Twenty-Second ACM Symposium on Principles of Programming Languages (POPL), Portland, Oregon*, pages 142–153, San Francisco, California, January 1995b.
- Xavier Leroy. The Objective Caml system: Documentation and user's guide. Available at <http://pauillac.inria.fr/ocaml/htmlman/>, 1996a.
- Xavier Leroy. A syntactic theory of type generativity and sharing. *Journal of Functional Programming*, 6(5):667–698, 1996b.
- Xavier Leroy. A syntactic theory of type generativity and sharing. *Journal of Functional Programming*, 6(5):667–698, September 1996c.
- Xavier Leroy. 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>.

- Xavier Leroy 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 Symposium on Principles of Programming Languages (POPL)*, San Antonio, Texas, 1999.
- Mark Lillibridge. *Translucent Sums: A Foundation for Higher-Order Module Systems*. PhD thesis, Carnegie Mellon University, Pittsburgh, PA, December 1996.
- Mark Lillibridge. *Translucent Sums: A Foundation for Higher-Order Module Systems*. PhD thesis, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, May 1997.
- Tim Lindholm and Frank Yellin. *The Java Virtual Machine Specification*. The Java Series. Addison-Wesley, Reading, MA, USA, January 1997. ISBN 0-201-63452-X. URL <http://www.aw.com/cp/javaseries.html>.
- Barbara Liskov. A history of CLU. *ACM SIGPLAN Notices*, 28(3):133–147, 1993.
- Ralph Loader. Finitary PCF is not decidable. *Theoretical Computer Science*, 266(1-2):341–364, September 2001.
- John M. Lucassen and David K. Gifford. Polymorphic effect systems. In *Principles of Programming Languages*, pages 47–57, New York, NY, USA, January 1988. ACM Press.
- Z. Luo. *Computation and Reasoning: A Type Theory for Computer Science*. Number 11 in International Series of Monographs on Computer Science. Oxford University Press, 1994.
- Zhaohui Luo and Robert Pollack. The LEGO proof development system: A user's manual. Technical Report ECS-LFCS-92-211, University of Edinburgh, May 1992.
- David MacQueen. Modules for Standard ML. In *1984 ACM Conference on LISP and Functional Programming*, pages 198–207, 1984.
- David MacQueen. Using dependent types to express modular structure. In *Thirteenth ACM Symposium on Principles of Programming Languages (POPL)*, 1986.
- David B. MacQueen. Using dependent types to express modular structure. In *ACM Symposium on Principles of Programming Languages (POPL)*, St. Petersburg Beach, Florida, 1986.
- Lena Magnusson and Bengt Nordström. The ALF proof editor and its proof engine. In *Types for Proofs and Programs*, volume 806, pages 213–237. Springer-Verlag LNCS 806, 1994.
- Harry G. Mairson, Paris C. Kanellakis, and John C. Mitchell. *Computational Logic: Essays in Honor of Alan Robinson*, chapter Unification and ML type reconstruction, pages 444–478. MIT Press, 1991.
- Henning Makhholm. Region-based memory management in Prolog. Master's thesis, Department of Computer Science, University of Copenhagen (DIKU), March 2000. URL <ftp://ftp.diku.dk/diku/semantics/papers/D-421.ps.gz>. DIKU Technical Report 00/09.

- Henning Makholm and Kostis Sagonas. On enabling the WAM with region support. In Peter J. Stuckey, editor, *Logic Programming*, volume 2401 of *Lecture Notes in Computer Science*, pages 163–178, Heidelberg, Germany, 29 July–1 August 2002. Springer-Verlag. ISBN 3-540-43930-7.
- Per Martin-Löf. *Intuitionistic Type Theory*. Bibliopolis, 1984.
- I. A. Mason, S. F. Smith, and C. L. Talcott. From operational semantics to domain theory. *Information and Computation*, 128(1):26–47, 1996.
- I. A. Mason and C. L. Talcott. Equivalence in functional languages with effects. *Journal of Functional Programming*, 1:287–327, 1991.
- David McAllester. On the complexity analysis of static analyses. *Journal of the ACM*, 49(4):512–537, July 2002.
- David McAllester. A logical algorithm for ML type inference. Manuscript, March 2003.
- Conor McBride. *Dependently Typed Functional Programs and their Proofs*. PhD thesis, LFCS, University of Edinburgh, 2000. URL <http://www.lfcs.informatics.ed.ac.uk/reports/00/ECS-LFCS-00-419/index.html>.
- Conor McBride and James McKinna. The view from the left. Submitted, 2002.
- James McKinna and Robert Pollack. Pure Type Systems formalized. In M. Bezem and J. F. Groote, editors, *Proceedings of the International Conference on Typed Lambda Calculi and Applications*, pages 289–305. Springer-Verlag LNCS 664, March 1993.
- David Melski and Thomas Reps. Interconvertibility of a class of set constraints and context-free language reachability. *Theoretical Computer Science*, 248(1–2), November 2000.
- R. Milner, M. Tofte, R. Harper, and D. MacQueen. *The Definition of Standard ML (Revised)*. MIT Press, 1997a.
- Robin Milner. A theory of type polymorphism in programming. *Journal of Computer and System Sciences*, 17(3):348–375, December 1978.
- Robin Milner, Mads Tofte, and Robert Harper. *The Definition of Standard ML*. MIT Press, 1990.
- Robin Milner, Mads Tofte, Robert Harper, and David MacQueen. *The Definition of Standard ML (Revised)*. MIT Press, 1997b.
- Robin Milner, Mads Tofte, Robert Harper, and David MacQueen. *The Definition of Standard ML (Revised)*. MIT Press, 1997c.
- Yasuhiko Minamide. A functional representation of data structures with a hole. In *ACM Symposium on Principles of Programming Languages (POPL)*, San Diego, California, pages 75–84, January 1998.
- Yasuhiko Minamide, Greg Morrisett, and Robert Harper. Typed closure conversion. In *ACM Symposium on Principles of Programming Languages (POPL)*, St. Petersburg Beach, Florida, pages 271–283, January 1996.

- J. C. Mitchell. 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.
- J. C. Mitchell and G. D. Plotkin. Abstract types have existential types. *ACM Transactions on Programming Languages and Systems*, 10:470–502, 1988a.
- John C. Mitchell. Coercion and type inference (summary). In *ACM Symposium on Principles of Programming Languages (POPL)*, Salt Lake City, Utah, pages 175–185, January 1984.
- John C. Mitchell. Type inference with simple subtypes. *Journal of Functional Programming*, 1(3):245–286, July 1991b.
- John C. Mitchell. *Foundations for Programming Languages*. MIT Press, 1996.
- John C. Mitchell and Gordon Plotkin. Abstract types have existential type. *ACM Transactions on Programming Languages and Systems*, 10(3):470–502, 1988b.
- Eugenio Moggi. Notions of computation and monads. *Information and Computation*, 93(1):55–92, July 1991. Presented at LICS '89.
- Shaw-Kwei Moh. The deduction theorems and two new logical systems. *Methodos*, 2:56–75, 1950.
- Christine Mohring. Algorithm development in the calculus of constructions. In *Proceedings, Symposium on Logic in Computer Science IEE* (1986), pages 84–91.
- Stefan Monnier, 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.
- Greg Morrisett, Karl Cray, Neal Glew, and David Walker. Stack-based typed assembly language. *Journal of Functional Programming*, 12(1):43–88, January 2002.
- Greg Morrisett, 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.
- Christian Mossin. *Flow Analysis of Typed Higher-Order Programs*. PhD thesis, DIKU, University of Copenhagen, Copenhagen, Denmark, 1997. URL <http://www.diku.dk/research/published/97-1.ps.gz>. Technical Report DIKU-TR-97/1.
- Martin Müller, 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.
- Martin Müller and Susumu Nishimura. Type inference for first-class messages with feature constraints. In Jieh Hsiang and Atsushi Ohori, editors, *Asian Computer Science Conference (ASIAN 98)*, volume 1538 of *LNCS*, pages 169–187, Manila, The Philippines, December 1998. Springer-Verlag.
- Alan Mycroft. Polymorphic type schemes and recursive definitions. In M. Paul and B. Robinet, editors, *Proceedings of the International Symposium on Programming*, volume 167 of *LNCS*, pages 217–228, Toulouse, France, April 1984. Springer.

- George C. Necula. Proof-carrying code. In *ACM Symposium on Principles of Programming Languages (POPL)*, Paris, France, pages 106–119, 15–17 January 1997.
- George C. Necula. *Compiling with Proofs*. PhD thesis, Carnegie Mellon University, September 1998. Also available as CMU-CS-98-154.
- George C. Necula. Translation validation for an optimizing compiler. In *ACM SIGPLAN '00 Conference on Programming Language Design and Implementation (PDLI)*, pages 83–94, Vancouver, BC, Canada, 18–21 June 2000. ACM SIGPLAN.
- George C. Necula and Peter Lee. Safe kernel extensions without run-time checking. In *2nd Symposium on Operating Systems Design and Implementation (OSDI '96)*, October 28–31, 1996, Seattle, WA, pages 229–243, Berkeley, CA, USA, October 1996. USENIX press.
- George C. Necula and Peter Lee. Efficient representation and validation of logical proofs. pages 93–104. IEEE Computer Society Press, 1998.
- Joachim Niehren, Martin Müller, and Andreas Podelski. Inclusion constraints over non-empty sets of trees. volume 1214 of *Lecture Notes in Computer Science*, pages 217–231. Springer-Verlag, April 1997.
- Joachim Niehren and Tim Priesnitz. Non-structural subtype entailment in automata theory. *Information and Computation*, 2003. To appear.
- Flemming Nielson and Hanne Riis Nielson. From CML to its process algebra. *Theoretical Computer Science*, 155:179–219, 1996.
- Flemming Nielson, Hanne Riis Nielson, and C. L. Hankin. *Principles of Program Analysis*. Springer, 1999.
- Flemming Nielson, Hanne Riis Nielson, and Helmut Seidl. A succinct solver for ALFP. 9(4):335–372, 2002. <http://www.informatik.uni-trier.de/~seidl/papers/succinct.pdf>.
- Hanne Riis Nielson and Flemming Nielson. Higher-order concurrent programs with finite communication topology. In *Principles of Programming Languages*, pages 84–97, New York, NY, USA, 1994. ACM Press. ISBN 0-89791-636-0.
- Susumu Nishimura. Static typing for dynamic messages. In *Proceedings of the 25th ACM Symposium on Principles of Programming Languages*, pages 266–278, New York, 1998. ACM Press.
- Henning Niss. *Regions are Imperative: Unscoped Regions and Control-Flow Sensitive Memory Management*. PhD thesis, Department of Computer Science, University of Copenhagen (DIKU), 2002.
- E.G.J.M.H. Nocker and J.E.W. Smetsers. Partially strict non-recursive data types. *Journal of Functional Programming*, 3(2):191–215, 1993.
- E.G.J.M.H. Nocker, J.E.W. Smetsers, M.C.J.D. van Eekelen, and M.J. Plasmeijer. Concurrent CLEAN. In Leeuwen and Rem, editors, *Parallel Architectures and Languages Europe*, number 505 in LNCS, pages 202–219. Springer-Verlag, 1991.

- Martin Odersky, Vincent Cremet, Christine Rockl, and Matthias Zenger. A nominal theory of objects with dependent types. In *Workshop on Foundations of Object-Oriented Languages (FOOL), informal proceedings*, 2003.
- Martin Odersky and Konstantin Läufer. Putting type annotations to work. In *ACM Symposium on Principles of Programming Languages (POPL), St. Petersburg Beach, Florida*, pages 54–67, St. Petersburg, Florida, January 21–24, 1996. ACM Press.
- Martin Odersky, Martin Sulzmann, and Martin Wehr. Type inference with constrained types. *Theory and Practice of Object Systems*, 5(1):35–55, 1999a.
- Martin Odersky, Martin Sulzmann, and Martin Wehr. Type inference with constrained types. *Theory and Practice of Object Systems*, 5(1):35–55, 1999b. Summary in *Workshop on Foundations of Object-Oriented Languages (FOOL), informal proceedings*, 1997.
- Peter O’Hearn and David Pym. The logic of bunched implications. *Bulletin of Symbolic Logic*, 5(2):215–244, 1999.
- Atsushi Ohori. A polymorphic record calculus and its compilation. *ACM Transactions on Programming Languages and Systems*, 17(6):844–895, 1999.
- Atsushi Ohori and Peter Buneman. Static type inference for parametric classes. In *ACM Symposium on Object Oriented Programming: Systems, Languages, and Applications (OOPSLA)*, pages 445–456, October 1989. Also in Carl A. Gunter and John C. Mitchell, editors, *Theoretical Aspects of Object-Oriented Programming: Types, Semantics, and Language Design*, MIT Press, 1994.
- I. E. Orlov. The calculus of compatibility of propositions (in russian). *Matematicheskii Sbornik*, 35:263–286, 1928.
- Jens Palsberg. Efficient inference of object types. *Information and Computation*, 123(2):198–209, 1995.
- Jens Palsberg and Patrick O’Keefe. A type system equivalent to flow analysis. In *Principles of Programming Languages*, pages 367–378, New York, NY, USA, 1995. ACM Press. ISBN 0-89791-692-1.
- Jens Palsberg and Michael Schwartzbach. Type substitution for object-oriented programming. In N. Meyrowitz, editor, *Proc. Conf. Object-Oriented Programming: Systems, Languages, and Applications and European Conf. on Object-Oriented Programming*, pages 151–160, Ottawa, Canada, October 1990. ACM Press.
- Jens Palsberg and Michael Schwartzbach. *Object-oriented Type Systems*. John Wiley & Sons, 1994.
- Jens Palsberg, Mitchell Wand, and Patrick M. O’Keefe. Type inference with non-structural subtyping. *Formal Aspects of Computing*, 9:49–67, 1997.
- Christine Paulin-Mohring. 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.

- Leaf Petersen, Robert Harper, Karl Crary, and Frank Pfenning. A type theory for memory allocation and data layout. In *ACM Symposium on Principles of Programming Languages (POPL)*, New Orleans, Louisiana, pages 172–184, January 2003.
- Simon Peyton Jones. Special issue: Haskell 98 language and libraries. *Journal of Functional Programming*, 13, January 2003.
- Simon Peyton Jones and Mark Shields. Lexically-scoped type variables. Submitted to ICFP'03., March 2003.
- Frank Pfenning and Carsten Schürmann. Algorithms for equality and unification in the presence of notational definitions. In T. Altenkirch, W. Naraschewski, and B. Reus, editors, *Types for Proofs and Programs*, number LNCS 1657. Springer-Verlag, 1998.
- Benjamin C. Pierce and David N. Turner. Object-oriented programming without recursive types. In *ACM Symposium on Principles of Programming Languages (POPL)*, Charleston, South Carolina, January 1993.
- A. M. Pitts. Relational properties of domains. *Information and Computation*, 127: 66–90, 1996.
- A. M. Pitts. Existential types: Logical relations and operational equivalence. In K. G. Larsen, S. Skyum, and G. Winskel, editors, *Automata, Languages and Programming, 25th International Colloquium, ICALP'98, Aalborg, Denmark, July 1998, Proceedings*, volume 1443 of *Lecture Notes in Computer Science*, pages 309–326. Springer-Verlag, Berlin, 1998.
- A. M. Pitts. Parametric polymorphism and operational equivalence. *Mathematical Structures in Computer Science*, 10:321–359, 2000.
- A. M. Pitts. 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, 2002a. International Summer School, APPSEM 2000, Caminha, Portugal, September 9–15, 2000.
- A. M. Pitts and I. 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*. 1996. To appear.
- Andrew M. Pitts. Nominal logic, a first order theory of names and binding. *Information and Computation*, 2002b. To appear.
- Andrew M. Pitts and Ian D. B. Stark. Observable properties of higher order functions that dynamically create local names, or: What's new? In *Mathematical Foundations of Computer Science, Proc. 18th Int. Symp., Gdańsk, 1993*, volume 711 of *Lecture Notes in Computer Science*, pages 122–141. Springer-Verlag, Berlin, 1993.
- G. D. Plotkin. Lambda-definability and logical relations. Memorandum SAI-RM-4, School of Artificial Intelligence, University of Edinburgh, 1973.

- G. D. Plotkin. LCF considered as a programming language. *Theoretical Computer Science*, 5:223–255, 1977.
- G. D. Plotkin and M. Abadi. A logic for parametric polymorphism. In M. Bezem and J. F. Groote, editors, *Typed Lambda Calculus and Applications*, volume 664 of *Lecture Notes in Computer Science*, pages 361–375. Springer-Verlag, Berlin, 1993.
- Gordon D. Plotkin. Lambda-definability in the full type hierarchy. In Jonathan P. Seldin and J. Roger Hindley, editors, *To H. B. Curry: Essays on Combinatory Logic, Lambda Calculus and Formalism*, pages 363–373. Academic Press, London, 1980.
- Jeff Polakow and Frank Pfenning. Natural deduction for intuitionistic non-commutative linear logic. In *International Conference on Typed Lambda Calculi*, number 1581 in LNCS, pages 295–309. Springer-Verlag, April 1999.
- Jeff Polakow and Frank Pfenning. Properties of terms in continuation-passing style in an ordered logical framework. In *Workshop on Logical Frameworks and Meta-Languages*, Santa Barbara, June 2000.
- Erik Poll. Expansion Postponement for Normalising Pure Type Systems. *Journal of Functional Programming*, 8(1):89–96, 1998.
- Robert Pollack. *The Theory of LEGO: A Proof Checker for the Extended Calculus of Constructions*. PhD thesis, University of Edinburgh, 1994.
- POPL. *Principles of Programming Languages*, New York, NY, USA, January 1991a. ACM Press.
- POPL. *Principles of Programming Languages*, New York, NY, USA, 17–19 January 2001b. ACM Press. ISBN 1-58113-336-7.
- François Pottier. A versatile constraint-based type inference system. *Nordic Journal of Computing*, 7(4):312–347, November 2000.
- François Pottier. A semi-syntactic soundness proof for $HM(X)$. Research Report 4150, INRIA, March 2001a.
- François Pottier. Simplifying subtyping constraints: a theory. *Information and Computation*, 170(2):153–183, November 2001b.
- François Pottier. A constraint-based presentation and generalization of rows. In *Eighteenth Annual IEEE Symposium on Logic In Computer Science (LICS'03)*, Ottawa, Canada, June 2003. URL <http://pauillac.inria.fr/~fpottier/publis/fpottier-lics03.ps.gz>.
- Sriram K. Rajamani and Jakob Rehof. A behavioral module system for the pi-calculus. In *SAS 01: Static Analysis*, LNCS 2126, pages 375–394. Springer-Verlag, 2001.
- Sriram K. Rajamani and Jakob Rehof. Conformance checking for models of asynchronous message passing software. In Ed Brinksma and Kim Guldstrand Larsen, editors, *Computer Aided Verification*, volume 2404 of *Lecture Notes in Computer*

- Science*, pages 166–179, Heidelberg, Germany, July 27–31 2002. Springer-Verlag. ISBN 3-540-43997-8.
- Jakob Rehof. Minimal typings in atomic subtyping. In *ACM Symposium on Principles of Programming Languages (POPL)*, Paris, France, pages 278–291, January 1997.
- Jakob Rehof and Manuel Fähndrich. Type-based flow analysis: From polymorphic subtyping to CFL reachability. In *POPL* POPL (b), pages 54–66. ISBN 1-58113-336-7.
- Jakob Rehof and Manuel Fähndrich. Type-based flow analysis: From polymorphic subtyping to CFL-reachability. In *ACM Symposium on Principles of Programming Languages (POPL)*, pages 54–66, January 2001.
- Alastair Reid, Matthew Flatt, Leigh Stoller, Jay Lepreau, and Eric Eide. Knit: Component composition for systems software. In *Proc. of the 4th Operating Systems Design and Implementation (OSDI)*, pages 347–360, 2000. URL citeseer.nj.nec.com/reid00knit.html.
- Didier Rémy. 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.
- Didier Rémy. Projective ML. In *ACM Symposium on Lisp and Functional Programming (LFP)*, pages 66–75, 1992b.
- Didier Rémy. Typing record concatenation for free. In *ACM Symposium on Principles of Programming Languages (POPL)*, Albuquerque, New Mexico, January 1992. Also in Carl A. Gunter and John C. Mitchell, editors, *Theoretical Aspects of Object-Oriented Programming: Types, Semantics, and Language Design*, MIT Press, 1994.
- Didier Rémy. 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.
- Didier Rémy. Programming objects with ML-ART: An extension to ML with abstract and record types. In Masami Hagiya and John C. Mitchell, editors, *International Symposium on Theoretical Aspects of Computer Software (TACS)*, pages 321–346, Sendai, Japan, April 1994. Springer-Verlag.
- Didier Rémy 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 Symposium on Principles of Programming Languages (POPL)*, Paris, France, 1997.
- Greg Restall. *An introduction to substructural logics*. Routledge, January 2000.
- Greg Restall. Handbook of the history and philosophy of logic. To appear., 2001.
- J. C. Reynolds. Towards a theory of type structure. In *Paris Colloquium on Programming*, volume 19 of *Lecture Notes in Computer Science*, pages 408–425. Springer-Verlag, Berlin, 1974a.

- J. C. Reynolds. Types, abstraction and parametric polymorphism. In R. E. A. Mason, editor, *Information Processing 83*, pages 513–523. North-Holland, Amsterdam, 1983a.
- John C. Reynolds. Automatic computation of data set definitions. In A. J. H. Morrell, editor, *Information Processing 68*, volume 1, pages 456–461. North Holland, 1969a.
- John C. Reynolds. Automatic computation of data set definitions. In A. J. H. Morrell, editor, *Information Processing 68*, volume 1, pages 456–461, Edinburgh, Scotland, 1969b. North Holland.
- John C. Reynolds. Towards a theory of type structure. In *Colloq. sur la Programmation*, volume 19 of *Lecture Notes in Computer Science*, pages 408–423. Springer-Verlag, 1974b.
- John C. Reynolds. Syntactic control of interference. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 39–46, Tucson, 1978.
- John C. Reynolds. Types, abstraction, and parametric polymorphism. In R. E. A. Mason, editor, *Information Processing 83*, pages 513–523, Amsterdam, 1983b. Elsevier Science Publishers B. V. (North-Holland).
- John C. Reynolds. Syntactic control of interference, part 2. In *International Colloquium on Automata, Languages and Programming*, July 1989.
- J. Alan Robinson. Computational logic: The unification computation. *Machine Intelligence*, 6:63–72, 1971.
- Douglas T. Ross. The AED free storage package. *Communications of the ACM*, 10(8):481–492, 1967.
- Claudio V. Russo. *Types for Modules*. PhD thesis, Edinburgh University, Edinburgh, Scotland, 1998. LFCS Thesis ECS–LFCS–98–389.
- Claudio V. Russo. Recursive structures for standard ml. In *Proc. Sixth ACM SIGPLAN International Conference on Functional Programming (ICFP '01)*, pages 50–61, Florence, Italy, September 2001.
- Fred B. Schneider. Enforceable security policies. *ACM Transactions on Information and System Security*, 3(1):30–50, February 2000.
- Jacob T. Schwartz. Optimization of very high level languages (parts I and II). *Computer Languages*, 1(2 & 3):161–194, 197–218, 1975.
- R. C. Sekar, R. Ramesh, and I. V. Ramakrishnan. Adaptive pattern matching. *SIAM Journal on Computing*, 24(6):1207–1234, December 1995.
- Miley Semmelroth and Amr Sabry. Monadic encapsulation in ML. In *International Conference on Functional Programming*, pages 8–17, 1999. URL citeseer.nj.nec.com/semmelroth99monadic.html.
- Peter Sestoft. Moscow ml. URL <http://www.dina.dk/~sestoft/mosml.html>.
- Peter Sestoft. Replacing function parameters by global variables. Technical Report 88-7-2, DIKU, University of Copenhagen, October 1988. SE88.

- Peter Sestoft. Replacing function parameters by global variables. In *Proc. Functional Programming Languages and Computer Architecture (FPCA), London, England*, pages 39–53. ACM Press, September 1989.
- Paula Severi and Erik Poll. Pure type systems with definitions. In *Proceedings of Logical Foundations of Computer Science (LFCS)*, pages 316–328. Springer-Verlag, 1994. LNCS volume 813.
- Zhong Shao. An overview of the FLINT/ML compiler. In *Proc. 1997 ACM SIGPLAN Workshop on Types in Compilation (TIC'97)*, Amsterdam, The Netherlands, June 1997.
- Zhong Shao, Christopher League, and Stefan Monnier. Implementing typed intermediate languages. In *Proceedings of the 1998 ACM SIGPLAN International Conference on Functional Programming*, pages 313–323, Baltimore, MD, September 1998. ACM SIGPLAN.
- Mark Shields and Erik Meijer. Type-indexed rows. In *ACM Symposium on Principles of Programming Languages (POPL), London, England*, pages 261–275, January 2001.
- Mark B. Shields and Simon Peyton Jones. First class modules for Haskell. In *Workshop on Foundations of Object-Oriented Languages (FOOL), informal proceedings*, pages 28–40, January 2002.
- Olin Shivers. Control flow analysis in Scheme. In *Programming Language Design and Implementation*, volume 23(7), pages 164–174, 22–24 June 1988.
- Olin Shivers. *Control-Flow Analysis of Higher-Order Languages or Taming Lambda*. PhD thesis, Carnegie Mellon University, May 1991.
- Vincent Simonet. Type inference with structural subtyping: A faithful formalization of an efficient constraint solver. Submitted for publication, March 2003.
- Christian Skalka and François Pottier. Syntactic type soundness for $HM(X)$. In *Workshop on Types in Programming (TIP'02)*, volume 75 of *Electronic Notes in Theoretical Computer Science*, July 2002.
- Frederick Smith, David Walker, and Greg Morrisett. Alias types. In Gert Smolka, editor, *Ninth European Symposium on Programming*, volume 1782 of *Lecture Notes in Computer Science*, pages 366–381. Springer-Verlag, April 2000.
- Geoffrey S. Smith. Principal type schemes for functional programs with overloading and subtyping. *Science of Computer Programming*, 23(2–3):197–226, December 1994.
- Jan Smith, Bengt Nordström, and Kent Petersson. *Programming in Martin-Löf's Type Theory. An Introduction*. Oxford University Press, 1990.
- R. Statman. Logical relations and the typed lambda calculus. *Information and Control*, 65:85–97, 1985a.
- Richard Statman. Logical relations and the typed λ -calculus. *Information and Control*, 65(2–3):85–97, May–June 1985b.

- Christopher A. Stone. *Singleton Kinds and Singleton Types*. PhD thesis, Carnegie Mellon University, Pittsburgh, PA, August 2000.
- Christopher A. Stone and Robert Harper. Deciding type equivalence in a language with singleton kinds. In *ACM Symposium on Principles of Programming Languages (POPL)*, Boston, Massachusetts, Boston, January 2000a.
- Christopher A. Stone and Robert Harper. Deciding type equivalence in a language with singleton kinds. In *Twenty Seventh ACM Symposium on Principles of Programming Languages (POPL)*, pages 214–227, Boston, January 2000b.
- Thomas Streicher. *Semantics of Type Theory*. Springer-Verlag, 1991.
- Zhendong Su and Alexander Aiken. Entailment with conditional equality constraints. In *European Symp. on Programming (ESOP)*, volume 2028 of *Lecture Notes in Computer Science*, pages 170–189, April 2001.
- Martin Sulzmann. *A general framework for Hindley/Milner type systems with constraints*. PhD thesis, Yale University, Department of Computer Science, May 2000.
- Martin Sulzmann, 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.
- William W. Tait. Intensional interpretations of functionals of finite type I. *Journal of Symbolic Logic*, 32(2):198–212, June 1967.
- C. Talcott. 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.
- Jean-Pierre Talpin and Pierre Jouvelot. Polymorphic type, region and effect inference. *Journal of Functional Programming (JFP)*, 2(2), 1992.
- Jean-Pierre Talpin and Pierre Jouvelot. The type and effect discipline. *Information and Computation*, 111:245–296, 1994.
- David Tarditi, 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 21–24 1996.
- Robert Endre Tarjan. Efficiency of a good but not linear set union algorithm. *Journal of the ACM*, 22(2):215–225, April 1975.
- Robert Endre Tarjan. Applications of path compression on balanced trees. *Journal of the ACM*, 26(4):690–715, October 1979.
- J. Terlouw. Een nadere bewijstheoretische analyse van GSTTs. Manuscript, University of Nijmegen, Netherlands, 1989.
- Kresten Krab Thorup. Genericity in Java with virtual types. In *European Conference on Object-Oriented Programming (ECOOP)*, volume 1241 of *Lecture Notes in Computer Science*, pages 444–471, Jyväskylä, Finland, June 1997. Springer-Verlag.

- Jerzy Tiurny and Mitchell Wand. Type reconstruction with recursive types and atomic subtyping. In *Proceedings of TAPSOFT '93*, volume 668 of *Lecture Notes in Computer Science*, pages 686–701. Springer-Verlag, April 1993.
- Mads Tofte. *Operational Semantics and Polymorphic Type Inference*. PhD thesis, University of Edinburgh, 1988.
- Mads Tofte and Lars Birkedal. A region inference algorithm. *ACM Transactions on Programming Languages and Systems*, 20(4):724–767, 1998. URL http://www.itu.dk/research/mlkit/kit_general/toplas98.ps.gz.
- Mads Tofte, Lars Birkedal, Martin Elsman, and Niels Hallenberg. Region-based memory management in perspective. To appear, 2003.
- Mads Tofte, Lars Birkedal, Martin Elsman, Niels Hallenberg, Tommy Højfeld Olesen, and Peter Sestoft. Programming with regions in the ML Kit (for version 4). Technical report, IT University of Copenhagen, October 2001.
- Mads Tofte, Lars Birkedal, Martin Elsman, Niels Hallenberg, Tommy Højfeld Olesen, Peter Sestoft, and Peter Bertelsen. Programming with regions in the ML Kit (for version 3). Technical Report DIKU-TR-98/25, Department of Computer Science, University of Copenhagen (DIKU), 1998. URL <http://www.it-c.dk/research/mlkit/kit3/manual.ps.gz>.
- Mads Tofte and Jean-Pierre Talpin. Implementing the call-by-value lambda-calculus using a stack of regions. In *ACM Symposium on Principles of Programming Languages (POPL)*, Portland, Oregon, January 1994.
- Mads Tofte and Jean-Pierre Talpin. Region-based memory management. *Information and Computation*, 132(2):109–176, 1 February 1997.
- Mads Torgersen. Virtual types are statically safe. In *Proceedings of the 5th Workshop on Foundations of Object-Oriented Languages (FOOL)*, San Diego, CA, January 1998.
- Valery Trifonov and Scott Smith. Subtyping constrained types. In *Static Analysis Symposium (SAS)*, volume 1145 of *Lecture Notes in Computer Science*, pages 349–365. Springer-Verlag, September 1996.
- David N. Turner. *The Polymorphic Pi-calculus: Theory and Implementation*. PhD thesis, University of Edinburgh, 1995.
- David N. Turner and Philip Wadler. Operational interpretations of linear logic. *Theoretical Computer Science*, 227:231–248, 1999. Special issue on linear logic.
- David N. Turner, Philip Wadler, and Christian Mossin. Once upon a type. In *ACM International Conference on Functional Programming and Computer Architecture*, San Diego, CA, June 1995.
- Diederik T. van Daalen. *The Language Theory of Automath*. PhD thesis, Technische Hogeschool Eindhoven, 1980.
- Per Velschow and Morten Voetmann Christensen. Region-based memory management in Java. Master's thesis, Department of Computer Science, University of

- Copenhagen (DIKU), 1998. URL <http://www.worldonline.dk/~voet/thesis.ps.gz>.
- Philip Wadler. Linear types can change the world! In M. Broy and C. Jones, editors, *Programming Concepts and Methods*, Sea of Galilee, Israel, April 1990. North Holland. IFIP TC 2 Working Conference.
- Philip Wadler. The marriage of effects and monads. In *Proceedings of the ACM SIGPLAN International Conference on Functional Programming (ICFP '98)*, volume 34(1), pages 63–74, 1999. URL citeseer.nj.nec.com/article/wadler98marriage.html. Journal version submitted to ACM Transactions on Computational Logic (2003).
- R. Wahbe, S. Lucco, T. E. Anderson, and S. L. Graham. Efficient software-based fault isolation. In *14th ACM Symposium on Operating Systems Principles*, pages 203–216. ACM, December 1993.
- David Walker, Karl Crary, and Greg Morrisett. Typed memory management in a calculus of capabilities. *ACM Transactions on Programming Languages and Systems*, 22(4):701–771, July 2000a.
- David Walker, Karl Crary, and Greg Morrisett. Typed memory management via static capabilities. *ACM Transactions on Programming Languages and Systems*, 22(4):701–771, July 2000b. URL <http://www.cs.princeton.edu/~dpw/capabilities-toplas.pdf>.
- David Walker and Greg Morrisett. Alias types for recursive data structures. *Lecture Notes in Computer Science*, 2071:177+, 2001.
- David Walker and Kevin Watkins. On linear types and regions. In *International Conference on Functional Programming*, Florence, September 2001a. ACM Press.
- David Walker and Kevin Watkins. On regions and linear types. In *6th International Conference on Functional Programming*, pages 181–192, New York, NY, USA, 3–5 September 2001b. ACM Press. ISBN 1-58113-415-0. URL <http://www.cs.princeton.edu/~dpw/papers/lr.pdf>.
- Mitchell Wand. Complete type inference for simple objects. In *Proceedings of the IEEE Symposium on Logic in Computer Science*, Ithaca, NY, June 1987.
- Mitchell Wand. Corrigendum: Complete type inference for simple objects. In *Proceedings of the IEEE Symposium on Logic in Computer Science*, 1988.
- Mitchell Wand. Type inference for record concatenation and multiple inheritance. In *Fourth Annual IEEE Symposium on Logic in Computer Science*, pages 92–97, Pacific Grove, CA, June 1989.
- Mitchell Wand. Type inference for objects with instance variables and inheritance. In Carl A. Gunter and John C. Mitchell, editors, *Theoretical Aspects of Object-Oriented Programming: Types, Semantics, and Language Design*, pages 97–120. MIT Press, 1994.
- Daniel C. Wang and Andrew W. Appel. Type-preserving garbage collectors. In *ACM Symposium on Principles of Programming Languages (POPL)*, London, England, pages 166–178, January 2001.

- Keith Wansbrough and Simon Peyton Jones. Once upon a polymorphic type. In *Twenty-sixth ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 15–28, San Antonio, January 1999.
- J. B. Wells. Typability and type checking in system F are equivalent and undecidable. *Annals of Pure and Applied Logic*, 98(1–3):111–156, 1999.
- Benjamin Werner. *Une Th'eorie des Constructions Inductives*. PhD thesis, L'Universite Paris, 1994.
- Niklaus Wirth. *Systematic Programming: An Introduction*. Prentice Hall, 1973.
- Niklaus Wirth. *Programming in Modula-2*. Texts and Monographs in Computer Science. Springer-Verlag, 1983.
- A. K. Wright and R. Cartwright. A practical soft type system for scheme. In *Proceedings of the 1994 ACM Conference on Lisp and Functional Programming*. ACM, June 1994. Also available as LISP Pointers VII(3) July-September 1994.
- A. K. Wright and M. Felleisen. A syntactic approach to type soundness. *Information and Computation*, 115:38–94, 1994a.
- Andrew K. Wright. Simple imperative polymorphism. *Lisp and Symbolic Computation*, 8(4):343–356, December 1995.
- Andrew K. Wright and Matthias Felleisen. A syntactic approach to type soundness. *Information and Computation*, 115(1):38–94, November 1994b.
- Hongwei Xi and Robert Harper. A dependently typed assembly language. In *International Conference on Functional Programming (ICFP), Firenze, Italy*, 2001.
- Jan Zwanenburg. Pure type systems with subtyping. In J.-Y. Girard, editor, *Typed Lambda Calculus and Applications (TLCA)*, pages 381–396. Springer-Verlag, 1999. Lecture Notes in Computer Science, volume 1581.

Index

- $\exists\sigma$ σ has an instance, 22
- $dfpi(\cdot)$ defined and free program variable identifiers, 24
- $dpi(\cdot)$ defined program identifiers, 16
- \emptyset empty environment, 16
- \Vdash entailment, 29
- Γ environment, 16
- $fpi(\cdot)$ free program identifiers, 22
- $ftv(\cdot)$ free type variables, 14, 21
- ϕ ground assignment, 27
- $\not\sim$ incompatible, 30
- $\cdot \preceq \cdot$ instance of, 21, 22, 31
- let \cdot in \cdot let constraints, 22
- \leq subtyping predicate, 20
- T type, 14
- Γ type environments, 24
- S type scheme, 15

- contravariant, 26, 29
- covariant, 26, 29

- entailment, 29
- environment, 16

- ground assignment, 27

- incompatible, 30
- instance of
 - type scheme, 16, 21, 22, 31
- invariant, 26, 29
- isolated, 30

- metavariables
 - naming conventions, 483

- naming conventions for metavariables and rules, 483

- rule
 - naming conventions, 483

- rules
 - A-ABS, 282
 - A-APP, 282
 - A-BOOL, 282
 - A-IF, 282
 - A-LVAR, 282
 - A-PAIR, 282
 - A-SPLIT, 282
 - A-UVAR, 282
 - ADD, 362
 - ABS, 53
 - APP, 53
 - BEQ-EQ, 362
 - BEQ-NEQ, 362
 - BETA-ABS, 206
 - BETA-ALL, 217
 - BETA-APP1, 206
 - BETA-APP2, 206
 - BETA-APPABS, 206
 - BETA-PAIR1, 213
 - BETA-PAIR2, 213
 - BETA-PROJ, 213
 - BETA-PROJPAIR, 213
 - C-ARROW, 30

C-DEF, 461
 CM-AND, 28
 CM-EXISTS, 28
 CM-FORALL, 108
 CM-INSTANCE, 28
 CM-PREDICATE, 28
 CM-TRUE, 28
 COMMIT, 376
 CTX-DEF, 444
 CTX-EMPTY, 443
 CTX-KIND, 443
 CTX-TYPE, 443
 E-APP', 301
 E-APP, 285
 E-APPRC, 301
 E-ARRAY, 297
 E-BOOL, 285
 E-CTXT, 285
 E-DEC1, 301
 E-DEC2, 301
 E-FREE, 297
 E-FUN', 301
 E-FUN, 285
 E-FUNRC, 301
 E-IF1, 285
 E-IF2, 285
 E-INC, 301
 E-LENGTH, 297
 E-MAPP1, 462
 E-MAPPABS, 462
 E-MLET, 462
 E-MLETV, 462
 E-MPAIR1, 462
 E-MPAIR2, 462
 E-MPAIRBETA1, 462
 E-MPAIRBETA2, 462
 E-MPROJ, 462
 E-MPROJV, 462
 E-MSEAL, 462
 E-MAPP2, 462
 E-PAPP, 295
 E-PFUN, 295
 E-PAIR', 301
 E-PAIR, 285
 E-PAIRRC, 301
 E-QAPP, 295
 E-QFUN, 295
 E-SPLIT', 301
 E-SPLIT, 285
 E-SPLITRC, 301
 E-SWAP, 297
 E-TLET, 444
 E-TERM, 462
 EXISTS, 53
 I-FUNCTOR, 460
 I-OPAQUE, 460
 I-PAIR, 460
 I-TERM, 460
 I-TRANSP, 460
 JMP, 362
 K-ABS, 443
 K-ALL, 443, 469
 K-APP, 202, 443, 469
 K-ARROW, 443
 K-CONV, 202
 K-FN-ETA, 469
 K-FN, 469
 K-FST, 469
 K-INT, 443, 469
 K-MPROJ, 461
 K-PAIR-ETA, 469
 K-PAIR, 469
 K-PI, 202
 K-SIGMA, 212
 K-SINTRO, 469
 K-SND, 469
 K-SUB, 469
 K-TVARDEF, 444
 K-VAR, 202, 443, 469
 K-PRF, 215
 K-PROP, 215
 KA-APP, 208
 KA-PI, 208
 KA-PRF, 217
 KA-PROP, 217
 KA-SIGMA, 214
 KA-VAR, 208
 LD-S, 376

- LD-U, 376
 LET, 53
 M-ABS1, 294
 M-ABS2, 294
 M-ABS, 461
 M-APPLY, 461
 M-APPLYV, 463
 M-EMPTY, 278, 302
 M-FST, 461
 M-LIN1, 278, 302
 M-LIN2, 278, 302
 M-LTOR, 302
 M-ORDL, 302
 M-ORDR, 302
 M-PAIR, 461
 M-SEAL, 461
 M-SELF-PAIR1, 461
 M-SELF-PAIR2, 461
 M-SELF, 461
 M-SND, 461
 M-SNDV, 463
 M-SUB, 461
 M-TERM, 461
 M-TOP, 302
 M-TYPE, 461
 M-UN, 278, 302
 M-VAR, 461
 MALLOC, 376
 MOV-1, 376
 MOV, 362
 Q-ABS, 140, 203, 443, 469
 Q-ALL, 469
 Q-APP, 140, 149, 203, 443, 469
 Q-APPABS, 472
 Q-ARROW, 443
 Q-BETA-FST, 472
 Q-BETA-PROD1, 159
 Q-BETA-PROD2, 159
 Q-BETA-SND, 472
 Q-BETA, 140, 155, 203, 443
 Q-DEF, 444
 Q-ELIM, 469
 Q-ETA-FN, 472
 Q-ETA-PAIR, 472
 Q-ETA, 203
 Q-EXT-PROD, 159
 Q-EXT, 140
 Q-FN-EXT, 469
 Q-FORALL, 443
 Q-FST, 469
 Q-MPROJ, 461
 Q-PAIR-EXT, 469
 Q-PAIR, 159, 469
 Q-PROJ1, 159, 212
 Q-PROJ2, 159, 212
 Q-REFL, 140, 203, 443, 469
 Q-SND, 469
 Q-SUB, 469
 Q-SURJPAIR, 212
 Q-SYM, 203, 443, 469
 Q-SYMM, 140
 Q-TRANS, 140, 203, 443, 469
 Q-UNIT-WEAK, 144
 Q-UNIT, 143
 QA-ABS, 209
 QA-ALL-E, 217
 QA-APP, 209
 QA-NE-PAIR, 214
 QA-NABS1, 209
 QA-NABS2, 209
 QA-PAIR-NE, 214
 QA-PAIR, 214
 QA-VAR, 209
 QA-WH, 209
 QAN-NORMAL, 147
 QAN-REDUCE, 147
 QAP-APP, 147
 QAP-CONST, 147
 QAP-PROJ1, 159
 QAP-PROJ2, 159
 QAP-VAR, 147
 QAR-APP, 147
 QAR-BETA-PROD1, 159
 QAR-BETA-PROD2, 159
 QAR-BETA, 147
 QAR-PROJ1, 159
 QAR-PROJ2, 159
 QAT-ARROW, 147

QAT-BASE, 147
 QAT-ONE, 147
 QAT-PROD, 159
 QK-*, 468
 QK-PI, 203, 468
 QK-REFL, 203
 QK-SIGMA, 468
 QK-SING, 468
 QK-SYM, 203
 QK-TRANS, 203
 QKA-PI-PRF, 217
 QKA-PI, 209
 QKA-PRF-PI, 217
 QKA-PRF, 217
 QKA-STAR, 209
 QR-ABS, 143, 447
 QR-ALL, 447
 QR-APP, 143, 447
 QR-ARROW, 447
 QR-BETA, 143, 447
 QR-DEF, 447
 QR-ETA, 143
 QR-REFL, 143, 447
 QT-ALL-E, 217
 QT-ALL, 215
 QT-APP, 203
 QT-PI, 203
 QT-REFL, 203
 QT-SYM, 203
 QT-TRANS, 203
 QTA-APP, 209
 QTA-PI, 209
 QTA-SIGMA, 214
 QTA-VAR, 209
 R-ADD, 10
 R-ALG-CASE, 91
 R-ALG-PROJ, 91
 R-ANNOTATION, 102
 R-APPABS, 172
 R-ASSIGN, 11
 R-BETA, 9
 R-CASE, 11
 R-CONDFALSE, 172
 R-CONDTRUE, 172
 R-CONTEXT, 9
 R-DELTA, 9
 R-DEREF, 11
 R-EXTEND, 9
 R-FALSE, 11
 R-FIX', 114
 R-FIX, 11
 R-LET, 9
 R-MATCH, 96
 R-OP, 172
 R-OPEN-ALL, 116
 R-OPEN-EX, 118
 R-PROJ, 10
 R-PROJRCd, 172
 R-REF, 11
 R-TAPPTABS, 172
 R-TRUE, 10
 R-UNPACKPACK, 172
 S-CONS, 172
 S-CONSVAl, 172
 S-NIL, 172
 S-NILVAL, 172
 S-RED, 172
 S-SEQ, 172
 SALLOC, 376
 SFREE, 376
 SI-FORGET, 460
 SI-OPAQUE, 460
 SI-PI, 460
 SI-SIGMA, 460
 SI-TERM, 460
 SI-TRANSP, 460
 SK-*, 468
 SK-FORGET, 468
 SK-PI, 468
 SK-SIGMA, 468
 SK-SING, 468
 ST-S, 376
 ST-U, 376
 SUB, 53
 T-ABS, 140, 202, 223, 276, 280, 294
 T-ALL, 215
 T-APP, 140, 170, 202, 223, 276,
 280, 294

- T-ARRAY, 297
- T-BOOL, 280
- T-BROKENVAR, 279
- T-CASE, 289
- T-CONST, 140, 170
- T-CONV, 202, 223
- T-DEC, 300
- T-EMPTYS, 286
- T-FREE, 297
- T-FUN, 170
- T-IF, 170, 280
- T-INC, 300
- T-INL, 289
- T-INR, 289
- T-LENGTH, 297
- T-MOD-LET, 461
- T-MOD-PROJ, 461
- T-NEXTRCS, 300
- T-NEXTLINS, 286
- T-NEXTUNS, 286
- T-OP, 170
- T-PABS, 294
- T-PAPP, 294
- T-PACK, 170
- T-PAIR, 159, 212, 280
- T-PI, 223
- T-PROG, 286
- T-PROJ1, 159, 212
- T-PROJ2, 159, 212
- T-PROJ, 170
- T-QABS, 294
- T-QAPP, 294
- T-RCD, 170
- T-ROLL, 289
- T-SEQ, 170
- T-SPLIT, 280
- T-STAR, 223
- T-SWAP, 297
- T-TFUN, 289
- T-TLET, 444
- T-TABS, 170
- T-TAPP, 170
- T-UNIT, 143
- T-UNPACK, 170
- T-UNROLL, 289
- T-VAR, 140, 170, 202, 223, 276, 280
- TA-ABS, 208
- TA-APP, 208
- TA-PAIR, 214
- TA-PROJ1, 214
- TA-PROJ2, 214
- TA-VAR, 208
- TERM, 172
- VAR, 53
- WFA-EMPTY, 209
- WFA-PI, 208
- WFA-STAR, 208
- WFA-TM, 209
- WFA-TY, 209
- WH-APP1, 208
- WH-APPABS, 208
- WH-PROJ, 213
- WH-PROJPAIR, 213
- WK-*, 468
- WK-PI, 468
- WK-SIGMA, 468
- WK-SING, 468
- WF-PI, 202
- WF-STAR, 202
- ANDEL, 322
- ANDER, 322
- ANDI, 322
- CONS, 323
- DM-ABS, 16
- DM-APP, 16
- DM-GEN', 18
- DM-GEN, 16
- DM-INST', 18
- DM-INST, 16
- DM-LET, 16
- DM-VAR, 16
- E-APP1, 232
- E-APP2, 232
- E-BETA, 232
- E-CASE, 262
- E-CASECONS, 262
- E-CASENIL, 262
- E-CONS1, 262

E-CONS2, 262
 E-CONALLOC, 262
 E-FIXBETA, 232
 E-IF, 232
 E-IFFALSE, 232
 E-IFTRUE, 232
 E-NEW, 239
 E-NEWBETA, 239
 E-TAG, 234
 E-TAGBETA, 234
 E-UNTAG, 234
 E-UNTAGBETA, 234
 HMD-ABS, 46
 HMD-APP, 46
 HMD-EXISTS, 46
 HMD-LETGEN, 46
 HMD-SUB, 46
 HMD-VARINST, 46
 HMX-ABS, 42
 HMX-APP, 42
 HMX-EXISTS, 42
 HMX-GEN', 42
 HMX-GEN, 42
 HMX-INST, 42
 HMX-LET, 42
 HMX-SUB, 42
 HMX-VAR, 42
 IMPE, 322
 IMPI, 322
 INIT, 354
 MEM0, 322
 MEM1, 322
 NEXT, 323
 NIL, 323
 PTRADDR, 323
 RD-DISCARD, 501
 RD-TRANSPose, 501
 RE-APP1, 246
 RE-APP2, 246
 RE-BETA, 246
 RE-CLOS, 246
 RE-DEALLOC, 246
 RE-FIX, 246
 RE-FIXBETA, 246
 RE-IF, 246
 RE-IFFALSE, 246
 RE-IFTRUE, 246
 RE-LETREG, 246
 RE-RAPP, 246
 RE-RBETA, 246
 RE-RCLOS, 246
 READ, 354
 S-ADD, 364
 S-BEQ, 364
 S-COMMIT, 380
 S-GEN, 364
 S-HEAP, 364
 S-INST, 364
 S-INT, 364
 S-JMP, 364
 S-LAB, 364
 S-LDS, 380
 S-LDU, 380
 S-MACH, 364
 S-MALLOC, 380
 S-MOV-1, 380
 S-MOV, 364
 S-PACK, 387
 S-REG, 364
 S-REGFILE, 364
 S-SALLOC, 380
 S-SEQ, 364
 S-SFREE, 380
 S-STC, 380
 S-STU, 380
 S-TUPLE, 379
 S-UNPACK, 387
 S-UPTR, 379
 S-VAL, 364
 SEL, 323
 SEND, 354
 SET, 323
 T-ABS, 232, 238
 T-APP, 232, 238
 T-BOOL, 232, 238
 T-BOOLSUB, 234, 238
 T-FIX, 232, 238
 T-FUNSUB, 234, 238

- T-IF, 232, 238
- T-LABELSUB, 234, 238
- T-NEWUN SOUND, 239
- T-SUBTYPE, 234
- T-TAG, 234, 238
- T-TAGVALSUB, 234, 238
- T-TAGVALUE, 234, 238
- T-UNTAG, 234, 238
- T-VAR, 232, 238
- TE-ABS, 243
- TE-APP, 243
- TE-AT, 243
- TE-BOOL, 243
- TE-BOOLSUB, 243
- TE-CELL, 243
- TE-EFFECTSUB, 243
- TE-FIX, 243
- TE-FROM, 243
- TE-FUNSUB, 243
- TE-IF, 243
- TE-LABELSUB, 243
- TE-NEW, 243
- TE-SUB, 243
- TE-TAGVALSUB, 243
- TE-VAR, 243
- THIS, 323
- TT-ABS, 255
- TT-APP, 255
- TT-ASSIGN, 262
- TT-BOOL, 255
- TT-CASE, 262
- TT-CLOS, 255
- TT-CONS, 262
- TT-CONSCELL, 262
- TT-DEREF, 261
- TT-EGEN, 255
- TT-EINST, 255
- TT-FIX, 255
- TT-IF, 255
- TT-LETREG, 255
- TT-NIL, 262
- TT-RABS, 255
- TT-RAPP, 255
- TT-RCLOS, 255
- TT-REF, 261
- TT-TGEN, 255
- TT-TINST, 255
- TT-VAR, 255
- UPD, 323
- satisfiable, 27
- standard multi-equation, 73
- substitution
 - type, 15
- type, 14
 - equirecursive, 106
 - isorecursive, 89
 - recursive, 106–107
- type scheme, 15