Advanced topics in programming languages

Michaelmas 2025

Module systems

Jeremy Yallop

jeremy.yallop@cl.cam.ac.uk

Module systems basics

"A module is a function which produces environments of a particular signature when applied to argument instances of specified signatures."

Modules for Standard ML (1984)
David MacQueen

Structures and signatures

Basics



History

Reading

```
module IntSet =
struct
type elem = int
type t = elem list
let empty = []
let rec mem x = function
| [] \rightarrow (* \ldots \ldots)
end
```

```
module type SET = sig type elem type t val empty : t val mem : elem \rightarrow t \rightarrow bool (* \dots *) end
```

Ascribing signatures to structures (IntSet :SET) involves subtyping, including abstraction (turning concrete types into abstract types) instantiation (turning polymorphic types into concrete types) as well as width and depth subtyping (dropping and subtyping entries).

Structures and signatures

Basics

• 0 0 0 0

History

Reading

```
module IntSet = struct

type elem = int

type t = elem list

let empty = []

let rec mem x = function

| [] \rightarrow (* \ldots \ldots)

end
```

```
module type SET =

sig

type elem
type t

val empty : t

val mem : elem → t → bool

(* . . . *)
end
```

Ascribing signatures to structures (IntSet:SET) involves subtyping, including abstraction (turning concrete types into abstract types) instantiation (turning polymorphic types into concrete types) as well as width and depth subtyping (dropping and subtyping entries).

Structures and signatures

Basics

History

Reading

```
module IntSet =
struct
type elem = int
type t = elem list

[let empty = []
let rec mem x = function
| [] → (* ... *)
end
```

```
module type SET =

sig
   type elem
   type t

val empty: t
   val mem: elem → t → bool
   (* ... *)
end
```

Ascribing signatures to structures (IntSet:SET) involves subtyping, including abstraction (turning concrete types into abstract types)

instantiation (turning polymorphic types into concrete types)
as well as width and depth subtyping (dropping and subtyping entries).

Functors

Basics



History

Reading

```
a functor
module type ORDERED =
sig
  type t
 val compare : t \rightarrow t \rightarrow int
end
module MakeSet (Elem: ORDERED) =
struct
  type elem = Elem.t
  type t = elem list
  let mem = List.mem
  . . .
end
```

Functors: functions from modules to modules.

Abstract (and less abstract) types

Basics



History

Reading

```
a type for MakeSet
  module MakeSet (Elem: ORDERED) :
     SET with type elem = Elem.t
                       expanded signature
SET with type elem = Elem.t
type elem = Elem.t
     type t
     val mem : Elem.t \rightarrow t \rightarrow bool
     . . .
   end
```

In the type of mem: t is **abstract**, Elem.t is **shared**, bool is **concrete**.

Sharing as dependency

Basics



History

Reading

Module types involve various forms of **dependency**:

Dependency between types and values:

Dependency between arguments and results:

Sharing as dependency

Basics



History

Reading

Module types involve various forms of **dependency**:

Dependency between types and values:

```
module type ORDERED = sig type the val compare : t \rightarrow t \rightarrow int \ (* depends on t *) end
```

Dependency between arguments and results:

Sharing as dependency

Basics



History

Module types involve various forms of dependency:

Dependency between types and values:

```
\begin{array}{c} \text{module type ORDERED =} \\ \text{sig} \\ \text{type } \\ \text{val compare : } \\ \text{t} \rightarrow \text{ t} \rightarrow \text{ int (* depends on t *)} \\ \text{end} \end{array}
```

Dependency between arguments and results:

```
module MakeSet :

(Elem) ← ORDERED) →

SET with type elem = Elem.t (* depends on Elem.t *)
```

Reading

Higher-order modules

Basics



History

```
Using higher-order modules can lead to loss of type equalities:
```

```
module Apply (MakeSet : (Elem:ORDERED) → SET)
(Elem : ORDERED) = MakeSet(Elem)

module IS1 = Apply(MakeSet)(Int) (* IS1.t /= Int.t *)
module IS2 = MakeSet(Int) (* IS2.t == Int.t *)
```

Leroy's solution: extend the **path notation** to include applications

```
type t = MakeSet(Int).t
```

Reading

Module systems **history**

"In the case of constructions, we obtain the notion of a very high-level functional programming language, with complex polymorphism well-suited for module specification."

The Calculus of Constructions (1988) Thierry Coquand and Gérard Huet

Modules and dependent types

Basics

History

•

Reading

1974	Towards a theory of type structure (Reynolds)
------	---

1985 Abstract types have existential type (Mitchell & Plotkin)

1986 Using dependent types to express modular structure (MacQueen)

1988 The Calculus of Constructions (Coquand & Huet)

1990 (Higher-order modules and the phase distinction (Harper, Mitchell & Moggi)

A type-theoretic approach to higher-order modules with sharing (Harper & Lillibridge)

2010 F-ing modules

1994

(Rossberg, Russo & Dreyer)

dependent types

Reading

Background reading (optional)

Basic:

§11 (*Related work and discussion*) of *F-ing modules*, extended version (Rossberg, Russo, Drever, 2015)

11 Related work and discussion

The literature on ML module semantics is voluminous and varied. We will therefore focus on the most closely related work. A more detailed history of various accounts of ML-style modules can be found in Chapter 2 of Russo's thesis (1998; 2003).

Existential types for ADTs. Mitchell & Plotkin (1988) were the first to connect the informal notion of "abstract type" to the existential types of System F. In F, values

Chapter 1 (The Design Space of ML Modules) of *Understanding and Evolving the ML Module System* (Dreyer, 2005)

Chapter 1

The Design Space of ML Modules

What is the ML module system? It is difficult to say. There are several dialects of the ML language, and while the module systems of the elidacts are cretainly for more alike than not, there are important and rather subtle differences among them, particularly with regard to the semantics of data sabstraction. The goal of Part I of this thesis is to offer a new vay of understanding these differences, and to derive from that understanding a unifying module system that harmonizes and improves on the existing designs.

In this chapter, I will give an overview of the existing ML module system design space. I begin in Section 1.1 by developing a simple example—a module implementing sets—that establishes some basic terminology and illustrates some of the key features shared by all the modern variants of the ML module system. Then, in Section 1.2, I describe several dialects that represent key points in the design space, and discuss the major axes along which they differ.

Reading

History



Paper 1: Translucent sums

History

Reading

A Type-Theoretic Approach to Higher-Order Modules with Sharing*

Robert Harper† Mark Lillibridge[‡] School of Computer Science Carnegie Mellon University Pittsburgh, PA 15213-3891

Abstract

The design of a module system for constructing and main taining large programs is a difficult task that raises a number of theoretical and practical issues. A fundamental issue is the management of the flow of information between program units at compile time via the notion of an interface. Experience has shown that felly openes interfaces are switness! to use in practice since too much information is hidden, and that fully transparent interfaces lead to excessive intendependencies, creating problems for maintenance and separate compilation. The "sharing" enerifications of Standard MI. address this issue by allowing the programmer to specify equational relationships between types in separate modules. but are not expressive enough to allow the programmer complets control over the propagation of type information between modules.

These problems are addressed from a type-theoretic viewpoint by considering a calculus based on Girard's system F... The calculus differs from those considered in receions studies by relying exclusively on a new form of weak sumtype to propagate information at compile-time, in contrast to approaches based on strong sums which roly on substitotion. The new force of sum tune allows for the specification of constional, so well as type and kind, information in interfaces. This provides complete control over the proparation of compile-time information between program units and is sufficient to encode in a straightforward way most uses of type sharing specifications in Standard ML. Medules are treated as "first-class" citizens, and therefore the exetem comparts bishes only modules and some chiest estant-d

This work was sponsored by the Advanced Research Projects Agency CSTO under the title "The Ear Desiret, Admined Develorement of Systems Software ARPA Coder No. 9313, issued by ESD/AVS under Contract No. F19038-91-C-0168 | Electronic mail address: rebtos.cau.eds. Electronic mail address: #410cs.csu.edu

Parmission to conv without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial adventage, the ACM copyright notice and the programming idioms: the language may be easily restricted to "second-class" modules found in MI-like languages.

1 Introduction

Modularity is an essential technique for developing and maintaining large software systems [46, 24, 36]. Most modern programming languages provide some form of module system that supports the construction of large systems from a collection of separately-defined program units [7, 8, 26, 32]. A fundamental problem is the management of the tension between the need to treat the components of a large system in relative isolation (for both conceptual and pragmatic reasons) and the need to combine these components into a coherent whole. In turnical cases this problem is addressed by equipping each module with a well-defined interface that mediates all across to the module and requiring that interfaces he enforced at system link time

The Standard ML (SML) module system [17, 32] is a particularly interesting design that has proved to be useful in the development of large software systerms [2, 1, 3, 11, 13]. The main constituents of the SML module system are signatures, structures, and functors, with the latter two sometimes called modules. A structure is a program unit defining a collection of types, exceptions, values, and structures (known as sufstructures of the structure). A functor may be thought of as a "parameterized structure", a first-order function mapping structures to structures. A signature is an interface describing the constituents of a structure - the types values exceptions and structures that it defines along with their kinds, types, and interfaces. See Fig. ure 1 for an illustrative example of the use of the SML module system: a number of sources are smileble for further examples and information [15, 39]

A crucial feature of the SML module system is the notion of ton at mind which allows for the modification

"The calculus differs from those considered in previous studies by relying exclusively on a new form of weak sum type to propagate information at compile-time, in contrast to approaches based on strong sums which rely on substitution [...]

"Modules are treated as "first-class" citizens. and therefore the system supports higher-order modules and some object-oriented programming idioms"

Paper 2: Applicative functors

History

Reading

Permission to come without fee all or part of this material is remeasion to copy without tee all or part of this material is granted provided that the copies are not made or distributed for

Applicative functors and fully transparent higher-order modules

Xavier Lerov INRIA

B.P. 105 Roccusmourt 28153 La Chenny France Zavier Lerov@inria.fr

ML module system: the "dot notation" as elimination con-

struct for abstract types [3, 4] and the notion of type sharing

frameworks remain unaccounted for in the abstract

approach, such as structure sharing and the "fully

transparent" behavior of higher-order functors predicted

by the operational approach [13]. Also, even though the

abstract anneouch is syntactic in nature and therefore

highly compatible with separate compilation [10], code

fragments with free functor identifiers could be supported hetter (see section 2.4 for an example). MacOneen [13, 1]

claims that the problem with higher-order functors is

serious enough to invalidate the abstract approach

and justify the recourse to complicated stamp-based descriptions of higher-order functors and constant

two of these problems (fully transparent higher-order func-

toes and supposet for non-closed rode (exements) in a syntac-

tic framework derived from [10]. It relies on a modification of the bahavior of functors (parameterized modules). In

Standard ML and other models based on type generativity.

a functor defining an abstract type returns a different type each time it is applied. We say that functors are assers-

tive. In this work, we consider functors as ambigation if the

functor is applied twice to provably arnal arguments, the

two abstract types returned remain compatible. Functors therefore map equals to equals, which enables equational

reasoning on functor applications during type-checking. In

turn. This allows more precise signatures for higher-order

ple of a module system that ensures type abstraction (the

representation independence properties still hold) without respecting strict type generativity (some applications of a

given functor may return new types while others return com-

patible types). In this approach, type abstraction mecha-

nisms are considered from a semantic point of view (how

to make programs robust with respect to changes of im-

plementations?) rather than from an operational point of

view (when are two structurally identical types compati-

ble?). This work illustrates the additional everywhere

Applicative functors are also interesting as an evam-

functors, thereby solving the full transparency problem.

The work presented in this paper is an attempt to solve

correlation mechanisms

Unfortunately, some features described by operational

A

We present a variant of the Standard ML module system where norameterized abstract types (i.e. functors returning penerative types) map provably equal arguments to compatible abstract types, instead of generating distinct types at each application as in Standard ML. This extension solves the full transparency problem (how to give syntactic signatures for higher order functors that courses exactly their propagation of type equations), and also provides better support for non-closed code fragments.

1 Introduction

Most modern programming languages provide support for type abstraction: the important programming technique where a named type t is covinged with operations f.r.... then the concrete implementation of t is hidden, leaving an abstract type t that can only be accessed through the operations f, g, . . . Type abstraction provides fundamental typing support for modularity, since it enables a type-checker to catch violations of the modular structure of

Type abstraction is usually implemented through generative data type declarations: to make a type t abstract, the tron-shorter consenter a new tree t incompatible with any other type, including types with the same structure. From this, it is tempting to explain type abstraction in terms of concrativity of type declarations and say for instance that 'a type is abstract because it is created each time its definition is evaluated". The Definition of Standard ML D4. 81 formalizes this approach as a calculus over type stamps that defines when "need" trace are expensed and when "old" types are propagated. This approach is adequate for specifying a type-checker, but too low-level and operational in nature to help understanding type abstraction and report

about programs using it Independently, Mitchell and Plotkin [16] have proposed a more abstract. has operational account of two abstraction hand on a parallel with evictorial quantification in logic. Instead of operational intuitions about two concrativity, this approach uses a precise semantic characterization: representation independence [17, 15], to show that type abstruction is enforced. This abstract approach has since been extended to account for the main features of the Standard

The remainder of this paper is organized as follows. Section 2 introduces informally the application compantics of functors and the main technical devices that implement it. Section 3 formalizes a calculus with applicative functors granted provised that the copies are not make to the direct commercial advantage, the ACM copyright notice and the title of the number of one and the title of the number of one and the Section 4 shows that the representation independence propesty still holds, and section 5 that higher-order functors are "We present a variant of the Standard ML module system where parameterized abstract types [...] map provably equal arguments to compatible abstract types, instead of generating distinct types at each application as in Standard ML.

"This extension solves the full transparency problem (how to give syntactic signatures for higher-order functors that express exactly their propagation of type equations)"

Paper 3: F-ing modules

Basics

History

Reading

F-ing Modules

Andreas Rossberg MPI-SWS rossberg@mpi-sws.org

Claudio V. Russo Microsoft Research crusso@microsoft.com Derek Dreyer MPI-SWS drever@mpi-sws.org

Abstract

Mit modules are a powerful language mechanism for decomposing programs into resounds composents. Undermady, they also have a reputation for being "complete," and requiring famy type theory that is mody quage to non-expert. While the reputation is care that have been developed in the process of studying modules, we aim here to dominanted that it is undescribe. To do so, we give a make the composition of the process of studying modules, we aim here to dominante that it is undescribe. To do so, we give a Mit. Mar module taperquise to a studying through the consisting time to the composition of the control of the control of the decided experiences to a studying time to the control of the Mit. Mar module taperquise to a studying time to the control time to the control of the control of the control of the module approximate to a studying time to the control of the study plant. It programments we therefore yielder that the control of the study plant. It programments we therefore yielder that the control of the control

Our module language supports the usual second-class modules with Standard ML-style generative function and local module with Standard ML-style generative function study of our approach, we further extend the language with the ability to package modules as first-class values—a very simple extension, as it turns out. Our approach also scales to handle OCam1-style applicative functor semantics, but the details are significatedly more subtle, so we have their pre-

sentation to a future, expanded version of this paper.

Lastly, we report on our experience using the "locally nameless" approach in order to mechanize the soundness of our elaboration sententials in Con-

Categories and Subject Descriptors D.3.1 [Programming Languages]: Formal Definitions and Theory; D.3.3 [Programming Languages]: Language Constructs and Festures—Modules, Abstract data types; F.3.2 [Logics and Menoings of Programs]: Semantics of Programming Languages—Operational semantics; F.3.3 [Logics and Menoings of Programs]: Studies of Program Constructs—Proc structure

General Terms Languages, Design, Theory

Keywords Type systems, ML modules, abstract data types, existential types, System F, elaboration, first-class modules

1. Introduction

Modularity is essential to the development and maintenance of large programs. Although most modern languages support modular programming and code reuse in one form or another, the languages

in the ML family employ a particularly expressive style of module system. The key features shared by all the dialects of the ML module system are their support for hierarchical namespace management (via structurer), a fine-grained variety of interfaces (via structurer) algebraturer). Given-side data shittarction (via functional)

and implementor-tide data abstraction twis notings; Unfortunately, which the utility of Mi. modules is not in disputs, they have neartheless acquired a reputation for being 'complet'. Simms Physion Jones, in an off-cold PVM'. 2003 Keynote edit. Simms Physion Jones, in the Cold PVM'. 2003 Keynote "high power, but poor power/cost exist". He contrast, he licend lastell—extended with urious "new"y per yearm extension to a Food Cortina with alloy wheeled. Although we disagree with Psyton Jones' amoning mankey; it seems, hand our conventions

noe complex for mere mortals to understand is sadly prodominant. Why; is this six Ane MI. modules routhy more difficult to program/implement/understand than other ambitious modularity mechanisms, such as GIIC's type clauses with per quality corrections [44] or Jura's clauses with generics and windowsh [45]? We think not callending this is deviously a findamentally abspected upsticities, Other can extrainly desire its colorant particular extra contraction of the con

"complexity" complaint.

Rather, we believe the problem is that the literature on the arRather, we believe the problem is that the literature on the arthan the problem is a most instance, and the problem is an occurrent, and the control of the control of the Arthur of the Control of the Contr

In response to this problem, Dreyer, Crary and Harper [9] developed a unifying type theory, in which previous systems can be understood as sublanguages that selectively include different features. Although formally and conceptually elegant, their unifying "Our elaboration defines the meaning of module expressions by a straightforward, compositional translation into vanilla System F_{ω} [...] We thereby show that ML modules are merely a particular mode of use of System F_{ω} . [...]

"[T]he previous [translations] all start from a pre-existing dependently-typed module language and show how to compile it down to F_{ω} [...] [O]ur approach is simpler and more accessible to someone who already understands F_{ω} and does not want to learn a new dependent type system just in order to understand the semantics of ML modules."

Writing suggestions

Basics

Abstract types

How do approaches to abstract types differ between designs?

Separate compilation

How do ML-style modules systems support separate compilation?

History

Higher-order functors

Are higher-order functors practically important?

Importance of sharing

What is the role and significance of sharing specifications?

Dependent types vs polymorphism

Are modules better approached via dependent types or polymorphism?

