Advanced topics in programming languages

Michaelmas 2025

Dependent types

Jeremy Yallop

jeremy.yallop@cl.cam.ac.uk

Dependent types: basics

Basics



What can depend on what? (e.g. what can appear as an argument in an application?)

Pattern matching

No dependencies involving types

Simple types

(all types are global)

Recursion

 $\mathbb{N} o (\mathbb{N} o \mathbb{N}) o \mathbb{N}$ $\lambda x \colon \mathbb{N}.\lambda f \colon (\mathbb{N} o \mathbb{N}).f x$ Terms depend on terms

Polymorphism

es

 $\forall \beta. (\forall \alpha. \alpha) \rightarrow \beta$ $\Lambda \beta. \lambda x: (\forall \alpha. \alpha). x[\beta]$ $\vdash \vdash \vdash$ Terms depend on types

Dependent types

Types depend on terms

 $\Pi(n:\mathbb{N}).\mathsf{Vec}\; n \to \mathsf{Vec}\; n$

 $\lambda n: \mathbb{N}.\lambda v: \text{Vec n.} v$

The Curry-Howard correspondence

Basics

Correspondence between simply-typed language and propositional logics:

Pattern

 $A
ightarrow B \simeq A
ightharpoonup B$ (functions and implications) $A
ightharpoonup B \simeq A
ightharpoonup B$ (products and conjunctions) $A + B \simeq A \lor B$ (sums and disjunctions)

matching

Correspondence between dependently-typed languages and predicate logics:

Recursion

 $(x:A) \to B \simeq \forall (x:A).B$ (functions and universal quantification) $\Sigma(x:A).B \simeq \exists (x:A).B$ (dependent pairs and existential quantification)

How should we **start** to design a dependently-typed language?
Foundation for constructive mathematics (Martin-Löf Type Theory)
Lambda calculus with fancy types (Calculus of Constructions)

Equalities



With dependent types we can form **types from terms**. Parameterise B by a term of type A:

Pattern matching $\Pi(x:A).B(x)$

Key Q: when are two types equal? (essential for type checking!)

Recursion

Is B(2+2) equal to B(4)?

Determining equality typically requires **normalization** (i.e. computation).

Readir

(Separate question: what equalities can we prove?)

Pattern matching

Pattern matching with simple types

else cons (head x) (append (tail xs) ys)

Simple branching reveals nothing to the type checker:

append xs ys = if empty xs then ys

append Nil ys = ys

append (Cons x xs) ys = Cons x (append xs ys)

Pattern matching \bullet \circ \circ

Either branch can access the head and tail.

Pattern matching exposes the value structure to the type checker:

Only the cons branch can access the head and tail.

Inductive families support indexing data types by terms:

Pattern matching



Recursion

the full type of vappend: {a : Type} \rightarrow {m : \mathbb{N} } \rightarrow {n : \mathbb{N} } \rightarrow Vect m a \rightarrow Vect n a \rightarrow Vect (m + n) a

Inductive families and pattern matching

Dependent matching may reveal something about another value:

Pattern matching



```
Matching the first vector with Nil tells us that m \equiv Z in the first branch
```

vappend : Vect m a \rightarrow Vect n a \rightarrow Vect (m + n) a

vappend (Cons x xs) ys = Cons x (vappend xs ys)

- so the return type in the first branch is Vect (Z + n) a \rightsquigarrow Vect n a
- 3. so ys has the appropriate type in the first branch

vappend Nil vs = vs

Inductive families and pattern matching

Rasics

Pattern matching

Dependent matching may reveal something about *another value*:

- Recursion
- 1. Matching the first vector with Ni1 tells us that n \equiv Z
- 2. so the type of ys is Vect Z b
- 3. and so Nil is the only possible constructor for ys

Recursion

Dependent types and termination

Basics

Ideally: all functions terminate.

Non-terminating functions can introduce logical inconsistency, e.g.:

Pattern matching

 $\begin{array}{lll} \mbox{circular} & : & \forall & (\mbox{A} & : \mbox{Type}) & \rightarrow & \mbox{A} \\ \mbox{circular} & \mbox{a} & = & \mbox{circular} & \mbox{a} \end{array}$

Recursion



or:

```
data Empty : Type where
   -- (no constructors)

loopy : Empty
loopy = loopy
```

Approximating termination

Pattern

Recursion

Problem: termination is undecidable, so we must approximate syntactically

Question: what to do with functions that are not structurally decreasing?

```
length : List a \rightarrow Int
structurally decreasing: length [] = 0
                     length (x:xs) = 1 + length xs
```

not (obviously)

```
quicksort :: List N \rightarrow List N
                    quicksort [] = []
structurally decreasing: quicksort (x:xs) = quicksort (filter (< x) xs) ++
                                     x : quicksort (filter (>= x) xs)
```

Paper 1: termination

Basics

Pattern matching

Recursion

Reading



The Size-Change Principle for Program Termination

Chin Soon Lee'
Department of Computer
Science and Software
Engineering
The University of Western
Australia
Neclands 6907
Western Australia
Ieecs @cs.uwa.edu.au

Neil D. Jones Datalogisk Institut University of Copenhagen Universitetsparken 1 DK-2100 Copenhagen Denmark neil@diku dk Amir M. Ben-Amram Academic College of Tel-Aviv-Yaffo 4 Antokolsky Street

Tel-Aviv 64044 Israel amirben@mta.ac.il

ABSTRACT

The 'directions bermination' principle for a finite oter functional log may with well-fromded data his a prepriam terminative and all highest is easy infinite call augment. John ing program control flew, would note an admitte decreat in

permuter the changes derivable from program symbol. The est of limiter cell reprinces that follow program flow under the tree pushed to remote indicate descends in a security rest, representable by a 10 debt is meaning to Alpinthum for an assumiter can be used to decide streeting as permutation. We also give a direct significant operating on "directioning graphs" political the primary to instantial.

Compared to other member in the literature, termine that maybe haved in the stretchess principle to supervisingly among an other partial policy of the stretchest control, indirect method in control, the control of the stretchest control, the stretchest control of the stretchest

We establish the problems intrinue or explains. This term is as to sumprissiply high, complete for exacts, in special the simplicity of the principle, exacts bardons is proved by a reduction for millionism program termination. An interesting commence the same hardons result applies to many other unityees found in the termination and quasitermination in directors.

This research was done white visiting DIKU.

Categories and Subject Descriptors

D.3.4 [Nottween Engineering]: 5: fiven/P m gram Verfiniting; D.3.4 [Programming Languages]: Processor, P.3.2 [Engine and Meminings of Programs]: 5 profifying and Verifying and Reasoning about Programs [-2.2] [Engine and Meminings of Programs]: 5 cm active of Programming.

Keywords

Termination, program analysis, omega automaton, PSPACEcompleteness, portial evaluation.

1. INTRODUCTION 1.1 Motivation

There may remove to study outs matic methods to

- Program windration; typically deductive methods are used to slow put it is correctness, the inquise apput specification is as solided provided the program terminates followed by a separate proof a farmination [11].
- grams, or one imported form a possibly untraktweethy contest.

 Offeral interest: termination has been studied in fields including functional programming. No. 8 gir ope year no-
- including functional programming [3], he gic program ming [7, 18, 18, 18, 14], term rewriting systems [3 24] and partial evaluation. Discussion of related werk appears at the end of this paper.
- the control of program values, but in ther more subtle.

 Use in partial enclusion this is a step towards a
- Use in partial cachasism this is a step towards a binding-time ambyts that will guarantee termination of program specialization [2, 2, 3, 3] and still also an accepts bly high degree of specialization in an affine guitted evaluator and as finally [3].

We emphasize here a careful and precise formulation of a simple but powerful principle to decide termination. It is away to this clear statement of the termination criterion "[A] program terminates on all inputs if every infinite call sequence (following program control flow) would cause an infinite descent in some data values."

"The set of infinite all sequences that follow program flow and can be recognized as causing infinite descent is an ω -regular set, representable by a Büchi automaton"

"There are many reasons to study automatic methods to prove program termination, including: Program verification [...] Interesting analysis: termination is not just an "abstract interpretation" [...] Use in partial evaluation"

Pattern matching

Recursion

Reading



IDRIS — Systems Programming Meets Full Dependent Types

Edwin C. Brady

School of Computer Science, University of St Andrews, St Andrews, Scotland.

Ernail: ab@cs.at-andrews.ac.uk

Abstract

Security of the security of th

Cetegories and Subject Descriptors D.3.2 [Programming Languages]: Language Classifications—Applicative (functional) Languages; C.2.2 [Computer-Communication Networks]: Network

General Terms Languages, Verification

Kerwoods Dependent Types, Data Description

1. Introduction

Systems opforuser, such as an operating system or a nativest-stack, misculative enveryable we do on a computer, whether that comparing it is a dictable matchine; a service, a mobile phone, or any embedded strice. It is a flowest well that made subserve opiosition convolved as a promising approach to ensuring the corrections of sufficiency with high bred variations tooks who for coll [31 and Apal. 157] with high bred variations tooks who for coll [32 and Apal. 157] of the convenience of the convenience

Permittion to make digital or head region of all or part of this work for permittid characters are it pramed without first permitted that expens on an email or distributed for profit or communical advantage and that copies hear this notice and the full clustion on the for page. It copy or otherwise, to republish, to post evenes or to redistribute to both, requires prior specific permitted for fac. EAFVII. James 95, 2011. Audit Trans. USA. formats precisely, but it does not attempt to store concrete data

compactly or disturble.

The compaction of the c

1.1 Contributions

The main contribution of this paper is to demonstrate that a high level dependently typed language is capable of implementing and ventifying code at a low level. We achieve this in the following specific ways:

- We describe the distinctive features of IDRIS which allow integration of low level systems programming constructs with higher level programs verified by type checking (Section 2).
 We done have no effective features features betterforce are less than the programming of the construction of the constru
- we salor now an enterior recognity transcoin inferrable can be embedded in a dependently typed language (Section 2.6).
 We introduce a serious systems application where a programing language meets program verification, and implement in fully: a benury data description language, which we use to describe ICMP and IP beaders receively excession the data layers.
- We show how to tackle some of the awkward problems which can arise in practice when implementing a dependently typed application. These problems include:
- Dealing with foreign functions which may have more specific inputs and outputs than their C types might suggest — e.g. we might know that an integer may be within a specific range.
 Sain's/stip proof obligations which arise due to piving that and functions precise types. As far as possible, we would like proof obligations to be solved sattermically, and proof expiriences should not inferior units. Terrature's condob/(vile)

"This paper describes the use of a dependently typed programming language, IDRIS, for specifying and verifying properties of low-level systems programs, taking network packet processing as an extended example."

"Our motivation is the need for systems software verification — programs such as operating systems, device drivers and network protocol implementations which are required for the correct operation of a computer system. Therefore it is important to consider not only how to verify software, but also how to do so without compromising on efficiency, and how to interoperate with concrete data as it is represented in the machine or on a network wire"

Why Dependent Types Matter

Thorsten Altenkirch Conor McBride

James McKinna

The University of Nottingham {txa.ctm}@cs.nott.ac.uk The University of St Andrews james.mckinna@st-andrews.ac.uk

Abstract

We exhibit the rationale behind the design of Epigram, a dependently typed programming language and interactive program development system, using refinements of a well known program—merge sort—as a running example. We discuss its relationship with other proposals to introduce aspects of dependent types into functional programming languages and sketch owns to take for further you it, this says.

1. Introduction

Types mater. That's what they're for—to classify data with respect to criteria which matter tow they should be stored in memory, whether they can be safely passed as injust to a given operation, even who is allowed to see them. Dependent types are types expressed in terms of data, explicitly relating their inhabitants to that data. Such charge they can be expressed on the trainer of what matters should talk. While conventional type systems allow us to validate our programs with respect to a fixed set of criteria, dependent types are much more flexible, they realize a continuum of precision from the basic assertions we are used to expect from types up to a complete specification of the program's behaviour. It is the programmer's choice to what degree he wants to exploit the expressiveness of such a powerful type disciplier. While the price for formally centraled software may be high, it is good to know that we can pay it in installments and that we are free to decide how far we want to go. Dependent types reduce certification to type checking, hence they provide a means to convince others that the assertions we make about our grouns are correct. Dependently type regions are correct. Dependently type deprograms are, by their antive, proof

Functional programmers have started to incorporate many supects of dependent types into novel type systems generalized algebraic data rypes and singleton rypes. Indeed, we share Sheard's vision [She04] of closing the semantic gap between programs and their properties. While Sheard's language Timega approaches this goal by an evolutionary step from current functional languages like Haskell, we are proposing a more radical departure with Egisjam, exploiting what we have learner from proof development tools like LEGO and COQ.

Epigram is a full dependently typed programming language defined by McBride and McKinna [MM04], drawing on experience with the LEGO system. McBride has implemented a prototype which is available together with basic documentation [McB04, McB05] from the Epigram homepage.¹ The prototype implements most of the features discussed in this article, and we are continuing to develop it to close the remaining "Dependent types [...] provide a means to convince others that the assertions we make about our programs are correct. Dependently typed programs are, by their nature, proof carrying code."

"Epigram can also typecheck and evaluate incomplete programs with unfinished sections sitting in *sheds*, [···], where the typechecker is forbidden to tread."

"Exploiting the expressivity of dependent types in a practicable way involves a wide range of challenges in the development of the theory, the design of language, the engineering of tools and the pragmatics of programming."

recuisioi



Writing suggestions

Basics

Termination

Is termination-checking practical for real-world programs?

Pattern matching

Efficiency

Are dependent types an impediment or an aid to efficiency?



How might dependent types change the way we think about programming?

Recursion

Radicalism

Do dependent types require radically new ways of programming?



Adoption

What might impede adoption of dependently-typed languages?