Advanced topics in programming languages

Michaelmas 2025

Partial evaluation

Jeremy Yallop jeremy.yallop@cl.cam.ac.uk

Partial evaluation basics

Partial evaluation

Partial evaluation

•000

dynamic inputs d

For program q, with

$$PE(q,s)=q_s$$
 such that $q_s(d)\equiv q(s,d)$

define partial evaluation:

static inputs s

Example: consider a parser with inputs g (grammar) and \bar{c} (string). We want

$$PE(\mathsf{parser}, g) = \mathsf{parser}_g \; \mathsf{such} \; \mathsf{that} \; \mathsf{parser}_g(\overline{c}) \equiv \mathsf{parser}(g, \overline{c})$$

Binding-time analysis (BTA)

Partial evaluation

Key idea: start with inputs s and d; assign each expression a *binding time*.



Reading

Two key analogies: BTA as type inference; BTA as abstract interpretation.

Binding-time analysis (BTA)

Partial evaluation

Key idea: start with inputs s and d; assign each expression a binding time.

```
let rec pow x n =
  if n = 0 then 1
  else x * (pow x (n - 1))
```

Reading

Two key **analogies**: BTA as type inference; BTA as abstract interpretation.

Partial evaluation

Partial evaluation

Reading

Take a static environment $\{ n \mapsto 3 \}$ and annotated program P:

let rec pow x n =

if n = 0 then 1 else x * (pow x (n - 1))

and β -reduce P to produce a (*more efficient*?) **specialized** program:

 $pow_3 x = x * (x * (x * 1))$

Check: $pow_3 x \equiv pow x 3$.

Binding-time improvements

Partial evaluation

Naive specialization can produce poor results. For example, in this program

$$s + (d + 1)$$

the sub-expression d+1 is dynamic because d is dynamic.

Binding-time improvements can bring static expressions together:

$$d + (s + 1)$$

Background reading

Partial evaluation

Neil D. Jones Carsten K. Gomard Peter Sestoft Partial Evaluation and Automatic Program Generation Partial Evaluation and Automatic Program Generation

N.D. Jones, C.K. Gomard, and P. Sestoft, With chapters by L.O. Andersen and T. Mogensen.

Prentice Hall International June 1993 ISBN 0-13-020249-5.

Online: https://www.itu.dk/people/sestoft/pebook/

Reading

HALL TONAL RIES IN PUTER RIENCE

C.A.R. HOARE SERIES EDITOR

Paper 1: Continuation-based partial evaluation

Partial

Reading

Continuation-Based Partial Evaluation

Julia L. Lawall
Computer Science Department
Brandeis University *
(Willes brandeis edu)

Olivier Danvy
Computer Science Department
Aarhus University **
(danvy@daimi.aan.dk)

Abstract

Binding-time improvements aim at making partial evaluations (a.k.a. program specialization) yield a better result. They have been achieved on the mostly by hand-transforming the source program. We observe that as they are better understood, these hand-transformations are progressively integrated into partial evaluation, thereby abeviating the need

Control-based binding-time improvements, for example, follow this nattern: they have evolved from ad-hor sourcelevel rewrites to a systematic source-level transformation into continuation-passing style (CPS). Recently, Bondool has explicitly integrated the CPS transformation into the enertalizer, thus narrly affertating the word for accepts basel CBS transformation. This CBS interestion is assessfully effective but were complex and ones beyond a simple CPS transformation. We show that it can be achieved directly by using the control operators shift and reset, which provide accost to the current continuation as a composable procedure. We automate, rependuce, and extend Rondorf's results. and describe how this approach scales on to hand-soiting partial-evaluation compilers. The first author has used this method to bootstrop the new release of Conselly nastial and nator Schiem. The control operators not only allow the nactial evalence to remain in direct style, but also can speed un martial evaluation significantly

1 Introduction

Partile evaluation is a program-transformation scheinpr for opcididing pergune [1, 23], it was developed in the sixties and sevendies [1, 25], drantically simplified in the eightine for purposes of eld-application [16], and it are notticed quantitatively and qualitatively. Quantitatively, pattle evaluation handle more and more programminglanguage features — types, higher-order procedure, data "webbree, Memochem 2024, (Arx. This mort was included as

"Withhere, Mannechmetts 02254, USA. This work was initiated as the Gregor Greshwise Instructor of Science & Technology in automate 1990; continued at Indiana University in full 1990, and was completed at Resolved University. It was partially supported by NSF under grant. CCR-9224433 and by ONN under grant N0044-03-1-015. "My Munkingshie, 8804 Auchias C, Demanie.

Permission to copy without fee all or part of this materials granted provided that the ecclear are not made or distributed for direct commends advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copyring is by permission of the Association of Computing Machiney. To copy otherwise, or to republish, registres a fee analist specific permission.

structures, and so on. Qualitatively, these features need to be handled effectively. This is where binding-time improvements intervene [23, Chap. 12].

1.1 Binding-time improvements

The notion of hinding time arises naturally in partial evalnative since source programs are evaluated in two stages: at partial-evaluation time (statically) and at run time (dynumically). The parts of the source program that can be evaluated statically are referred to as "static" and the others as a formation.

ces as "dynamic".

Obviously, the more static parts there are in a source program, the better it gets specialized. A hinding-time improvement is a source-level transformation that enables more parts to be classified as static. Say that we want to partially evaluate the expression

x + (y - 1)

where we know that x is bound statically and y is bound dynamically. A naïve kinding-time analysis would classify both the subtraction and the addition to be dynamic, those in each case one of the operands is dynamic. Using the associativity and commutativity have of arithmetic, we can rewrite the expression on Solizons.

w 4 (e = 1)

The same naive hinding time analysis would now classify the substantion to be static (cine or will be known as an partialevaluation time and 1 is an immediate constant) and the difficion to be dynamic (cine or will not be known until run time). By rewriting the expression, we have addissed a binding-time improvement: the earth sideling-time analysis classifies more expressions as static, thus easiling the same related risids a better result.

1.2 Evaluation and partial evaluation Partial evaluation mimics evaluation — computing values

of static expressions, but residualizing (i.e., reconstructing) dynamic expressions, to peedes the specialized program. When an expression is residualized, the continuation of the partial evaluation of its components may differ from the continuation of their evaluation. This can cause a bose of static information. Consider the Scheme expression² "Control-based binding-time improvements [...] have evolved from ad-hoc source-level rewrites to a systematic source-level transformation into continuation-passing style (CPS).

"Recently, Bondorf has explicitly integrated the CPS transformation into the specializer, thus partly alleviating the need for source-level CPS transformation. This CPS integration is remarkably effective [...] We show that it can be achieved directly by using the control operators shift and reset [...]

"The control operators not only allow the partial evaluator to remain in direct style, but also can speed up partial evaluation significantly."

Partial evaluation

Eta-Expansion Does The Trick *

Olivier Danvy Karoline Malmkjær Jens Palsberg $\begin{array}{ccc} \mathbf{BRICS}^{\dagger} & \mathbf{MIT}^{\S} \\ \mathbf{Aarhus} \; \mathbf{University}^{\ddagger} & \end{array}$

 ${\rm May}\ 1996$

Abstract

Partial-evaluation follore has it that massaging one's source programs can make them specialize better. In Jones, Gomard, and Sestoft's recent textbook, a whole chapter is dedicated to listing such "binding-time improvements": nonstandard use of continuationpassing style, et-expansion, and a popular transformation called "The Trick". We provide a unified view of these binding-time improvements, from a typing perspective.

Just as a proper treatment of product values in partial evaluation requires partially static values, a proper treatment of digioint sums requires moving static contexts across dynamic case expressions. This requirement proceedly accounts for the nonstandard use of continuationpassing style encountered in partial evaluation. Eta-expansion thus seats as a uniform binding-time correction between values and contexts, be they of function type, product type, or disjoint-sum type. For the latter case, it enables STIP-Dricks

In this article, we extend Gomard and Jones's partial evaluator for the \(\text{-calcuts}, \) \text{-Mix, with products and disjoint sums we point out how eta-expansion for (finite) disjoint sums enables The Trici; we generalize our earlier work by identifying that eta-expansion can be obtained in the binding-time analysis simply by adding two coercion rules; and we specify and prove the correctness of our extension to

Keywords: Partial evaluation, binding-time analysis, program specialization, binding-time improvement, eta-expansion, static reduction. "Just as a proper treatment of product values in partial evaluation requires partially static values, a proper treatment of disjoint sums requires moving static contexts across dynamic case expressions. This requirement precisely accounts for the nonstandard use of continuation-passing style encountered in partial evaluation. Eta-expansion thus acts as a uniform binding-time coercion between values and contexts, be they of function type, product type, or disjoint-sum type. For the latter case, it enables "The Trick"."

Partial

The Essence of LR Parsing

Michael Sperber Peter Thiemann Wilhelm-Schickard-Institut für Informatik

Universität Tühingen Sand 13. D-72076 Tübingen, Germany

Partial evaluation can turn a general paner into a parser eralization is necessary to prevent infinite specialization. emerator. The emerated parsers surpass those produced. Thus, the parsers modified for good specialization retain by traditional paper generators in speed and compactness. We use an inherently functional approach to implement general LR(k) parsers and specialize them using the partial evaluator Similia. The functional implementation troduces the basic concents of LR parsing along with a of LR parsing allows for concise implementation of the non-deterministic functional algorithm for it. Section 3 about the parameters and province only designation of an experimental agreement of the section o changes to achieve good specialization results. In contrast, a traditional, stack-based implementation of a general LR. binding-time improvements made to it. Section 4 deconservation and structural changes to make it senting one formulation and Scheme implement amenable to satisfactory specialization.

1 Introduction

We necessit two inherently functional implementations of general LR(k) parsers: a direct-style first-order textbook version and one using continuation-passing style (CPS) for state transitions. Neither requires the handling of an ex- 2 . LR Bassins plicit parsing stack

These parsers, when specialized with respect to a gram- 2.1 Notational Proliminaries may and lookahead k, yield efficient residual parages. To achieve need results with offline partial evaluation, only a small number of changes to the general parsers are neces-

- binding-time contexts
- secolling loops over lists to discard appeared computations.
- for the CPS-based pursus

We describe the most important applications of the above improvements. The generated parsers are comnext and are either about as feet or feeter then there are sented by Mossin [12]. His traditional stack-based name in the text are implicitly elements of P. requires substantial changes in the representation of the thermore since the name stack is a static data structure under dynamic control, specialization suffers from termina-and $\hat{\Rightarrow}$ denotes the reflexive and transitive closure of \Rightarrow

annmach as it does not deal with explicit stacks at all. For the CPS approach, it is immediately obvious where earthe structure of their anorstors and most of their simplic-

The paper is organized as follows: the first section inmentation of the algorithm in Scheme, and describes the tation of functional LR parsing using CPS, again with a description of the binding-time improvements made to it. Section 5 describes some additional features which can be added to the parsing algorithms. Section 6 gives the results of practical experiments. Sec. 7 discusses related work, and Sec 8 constates

We use mainly standard notation for context-free gram-

many blooming the definition of hundres which follows is expecific to the functional interpretation of LR pageing A context-free grammar is tuple G = (N.T.P.S). · some standard binding-time improvements, notably is the set of nonterminals. T the set of terminals. S. c. N. some applications of "The Trick" [9] as well as some the start symbol, V = T U N the set of grammar symdunlisation of procedures which occur in multiple bols, and P the set of productions of the form A -- a for a nonterminal A and a sequence a of grammar symbols. Additionally V* denotes the set of sequences of grammar eymbols—analogoush T' and N'

c is the empty sequence. | | | is the length of sequence | prevention of infinite specialization of LR transitions
 Purthermore, aⁿ denotes a sequence of k copies of a, and , is the secuence consisting of the first k terminals in A. Some letters are by default assumed to be elements of Some retters are by occurrent assume T certain sees: $A,B,C,E \in \mathbb{N}; \xi,\rho,\sigma,\tau \in \mathbb{T}^*; x,y,z \in \mathbb{T}$ $\alpha, \beta, \gamma, \delta, \nu, \mu \in V$, and $X, Y, Z \in V$. All grammar rules G includes the derives relation on on V* with

$\alpha \rightarrow \beta \implies \alpha = \delta A \times \Delta \beta = \delta u \times \Delta A \rightarrow u$

der dynamic control, specialization somes internation and ϕ carrows the remarked ϕ can be called ϕ and ϕ carrows the remarked ϕ can be called ϕ and ϕ can be called ϕ can be called where $\alpha_{i-1} \Rightarrow \alpha_i$ for $1 \le i \le n$. Lettmost-sumbol rearriling #4 is a relation defined as

 $B\alpha \Rightarrow \delta\beta : \Leftrightarrow B \rightarrow \delta \wedge \delta \neq c$

"Partial evaluation can turn a general parser into a parser generator. The generated parsers surpass those produced by traditional parser generators in speed and compactness. [...]

"The functional implementation of LR parsing allows for concise implementation of the algorithms themselves and requires only straightforward changes to achieve good specialization results. In contrast, a traditional, stack-based implementation of a general LR parser requires significant structural changes to make it amenable to satisfactory specialization."

Writing suggestions

Partial evaluation

Binding-time improvements

How useful are CPS conversion and eta expansion in practice? Are there any other generally-useful binding-time improvements

Applicability and limitations

How widely applicable is partial evaluation in practice? What kind of performance improvements might we expect?

Compilation

Might it be practical to use partial evaluation as a compilation technique?

Demise

What happened to partial evaluation as a research field?