

# THEORETICAL PEARL

## *CPS in little pieces: composing partial continuations*

DANIEL P. FRIEDMAN\* and AMR SABRY†

*Computer Science Department, Indiana University, Bloomington, IN 47405, USA*

---

### Abstract

This paper presents a new two-stage CPS algorithm. The first stage plants trivial partial continuations via a recursive-descent traversal and the second stage is a rewrite system that transforms all nontail calls into tail calls. The algorithm combines the metaphors of the Plotkin-style CPS transformation along with reduction in the  $\lambda$ -calculus.

---

### 1 Introduction

The CPS transformation is usually presented as a recursive-descent algorithm that constructs continuations on the fly (Plotkin, 1976; Fischer, 1993). Instead, we tease the transformation into two parts: a simple recursive-descent traversal that plants trivial partial continuations, and a rewriting system for composing and simplifying these partial continuations. Our development focuses on the call-by-value CPS transformation but should also hold for the call-by-name case. Partial continuations (also known as functional continuations or subcontinuations) were precisely introduced as a more general notion of continuations that can be composed and simplified (Felleisen *et al.*, 1988; Johnson & Duggan, 1988; Danvy & Filinski, 1990; Danvy & Filinski, 1992).

The remainder of this section illustrates these ideas with two examples written in the untyped call-by-value  $\lambda$ -calculus. The examples characterize the operational aspects of the two parts of the transformation, but the untyped framework fails to capture some essential invariants. Hence sections 2 and 3 formalize our algorithms in the context of a typed CPS language with partial continuations. Section 4 is some discussion of the two algorithms and the conclusion, section 5, puts our intuition into context.

If we work in the untyped call-by-value  $\lambda$ -calculus, the recursive-descent part of

\* Supported by the National Science Foundation under Grant No. CCR-9987458.

† Supported by the National Science Foundation under Grant No. CCR-0196063.

the transformation translates a term  $e$  to  $\lambda k.ke^*$  where:

$$\begin{aligned} x^* &= x \\ (\lambda x.e)^* &= \lambda k.\lambda x.ke^* \\ (e_1 e_2)^* &= e_1^* I e_2^* \end{aligned}$$

Every application takes  $I$  as a trivial partial continuation. Every  $\lambda$  takes an additional continuation argument that is immediately applied to the body. For example, the term  $x(\lambda y.y)z$  translates to  $\lambda k.k((x I_1 (\lambda k.\lambda y.ky)) I_2 z)$ , where we have subscripted the two occurrences of  $I$  for clarity.

The second part of the transformation processes each nontail call by composing its surrounding context with its partial continuation. For example, the call  $(x I_1 (\lambda k.\lambda y.ky))$  has  $k$  as a context and  $I_2$  as a partial continuation. Merging these produces  $\lambda k.(x I_1 (\lambda k.\lambda y.ky)) k z$ .

The remaining nontail call is in an application position expecting  $k$  and  $z$  as arguments. This context is merged with the partial continuation  $I_1$  to produce  $\lambda k.x(\lambda v.vkz)(\lambda k.\lambda y.ky)$ , which is the result of the traditional CPS transformation.

In general, several possible reductions on partial continuations may apply at a given point. The reductions will be restricted to ensure a left-to-right evaluation order of applications but can otherwise be applied in any order.

As another example, consider the term  $x(y(zw))$ . Applying the first algorithm produces  $\lambda k.k(xI_1(yI_2(zI_3w)))$ . One reduction sequence for the second algorithm is:

$$\begin{aligned} &\lambda k.k(xI_1(yI_2(zI_3w))) \\ \rightarrow &\lambda k.k(xI_1(z(\lambda v_2.yI_2v_2)w)) \\ \rightarrow &\lambda k.xk(z(\lambda v_2.yI_2v_2)w) \\ \rightarrow &\lambda k.z(\lambda v_1.xk(yI_2v_1))w \\ \rightarrow &\lambda k.z(\lambda v_1.y(\lambda v_2.xkv_2)v_1)w \end{aligned}$$

In the first line, the return value of  $z$  is passed to  $y$ , whose return value is then passed to  $x$ . Thus for example, it would be possible to evaluate the term in a context that binds  $z$  to a function that ignores its continuation and returns an `int`, binds  $y$  to a function taking an `int` as its second argument and returning a `bool`, and binds  $x$  to a function taking a `bool` as its second argument. After two reduction steps, in the third line, the return value of  $z$  is passed directly to  $x$  since the call to  $y$  has been absorbed as part of the continuation of  $z$ . Since  $z$  ignores its continuation and returns an `int`, the term is ill-typed. Given sensible restrictions on the free variables, the term becomes well-typed and reduction can be shown to preserve typing. To make this explicit, we move to a typed calculus.

## 2 Source and target typed languages

The types and terms of the source calculus are:

$$\begin{aligned} t &::= b \mid t \rightarrow t \\ e &::= x \mid \lambda x.e \mid ee \end{aligned}$$

where  $b$  ranges over an unspecified collection of base types.

The presence of partial continuations in the CPS language complicates the types of the language. As Filinski argues (1999; 1996), the main challenge in typing

partial continuations is in defining a rich enough answer type in which all types can be embedded and projected. Our typed framework adapts the development of Sabry (1996) for a language with control primitives that can express partial continuations, and has similarities to Filinski's presentation but is considerably simpler since we are not concerned with the operational interpretation of the embeddings and projections. The types of the CPS language include a type  $o$  of answers, as well as polymorphic types:

$$u ::= b \mid o \mid u \rightarrow u \mid \alpha \mid \forall \alpha.u$$

Each partial continuation will return a value of type  $o$ . To compose partial continuations with regular functions, two polymorphic constants are included to mediate between the type  $o$  and the types expected by the context of the partial continuation. The set of CPS terms is:

$$\begin{array}{ll} \text{(Values)} & W ::= x \mid \lambda k.K \mid \#^u A \\ \text{(Partial Continuations)} & K ::= k \mid \lambda x.A \mid WK \mid @^u \\ \text{(Partial Answers)} & A ::= KW \\ \text{(Programs)} & P ::= \lambda k.A \end{array}$$

Ignoring the polymorphic constants prompt ( $\#$ ) and abort ( $@$ ) for the moment, the CPS language is traditional (Sabry & Felleisen, 1993). Values are either variables or continuation transformers. Continuations are either variables, functions mapping values to answers, or the result of transforming another continuation. All continuations must have type  $u \rightarrow o$  for some  $u$ . Answers are produced by applying continuations and must be of type  $o$ . Entire programs abstract over the single free continuation variable  $k$  and return an answer.

The constant  $@$  has type  $\forall \alpha.o \rightarrow o$ , and must be applied to a type  $u$  before it is used. It embeds a value into the answer type that is common to all continuations. The constant  $\#$  has type  $\forall \alpha.o \rightarrow \alpha$ , and must also be first applied to a type  $u$ . It projects values that were embedded into the answer type back to their original type. Both constants mediate between answers and values as will be apparent in the next section when we present the modified recursive-descent algorithm. If these constants are never used, then all calls are of the form  $WKW$ , i.e. all calls are tail calls.

As usual, the semantics of the CPS language is given by  $\beta$  and  $\eta$  reductions. In addition, the prompt and abort at identical types are stipulated to be inverses:

$$\begin{array}{ll} \#^u(@^u W) & \rightarrow W \\ @^u(\#^u A) & \rightarrow A \\ (\lambda k.K_1)K_2 & \rightarrow K_1[k := K_2] \\ (\lambda x.A)W & \rightarrow A[x := W] \\ \lambda k.Wk & \rightarrow W \quad \text{where } k \notin FV(W) \\ \lambda x.Kx & \rightarrow K \quad \text{where } x \notin FV(K) \end{array}$$

All but the  $\beta$  and  $\eta$  reductions on source terms are considered to be administrative. As expected, the rules are type-preserving.

### 3 Composing partial continuations

In this section, we modify the recursive-descent algorithm to use the type judgments of the source language. By introducing prompts and aborts at the appropriate types in the output of the translation, it becomes evident that both phases of the new CPS algorithm are type-preserving.

As usual (Meyer & Wand, 1985), the CPS translation on types is the following:

$$\begin{aligned} b^* &= b \\ (t_1 \rightarrow t_2)^* &= (t_2^* \rightarrow o) \rightarrow t_1^* \rightarrow o \end{aligned}$$

The revised recursive-descent algorithm treats variables and procedures as in the untyped case. For applications, the trivial partial continuation  $I$  is replaced with an  $@$  that embeds the return type of the application into the answer type for continuations. The entire application is surrounded with a  $\#$  that extracts the value from the answer type:

$$\begin{aligned} x^* &= x \\ (\lambda x.e)^* &= \lambda k.\lambda x.ke^* \\ (e_1^{t_2 \rightarrow t} e_2)^* &= \#^{t^*}(e_1^* @^t e_2^*) \end{aligned}$$

The type  $t_2$  mentioned in the left-hand side of the last equation does not appear in the right-hand side because we are only interested in the return type.

It is interesting that even the trivial planting of partial continuations produces terms that appear to have been transformed to CPS (as far as the types are concerned).

#### Proposition 3.1

If  $\Gamma \vdash e : t$  then  $\Gamma^* \vdash e^* : t^*$ .

Then the second pass can be described as “Do any of the administrative reductions on CPS terms and any of the following two reductions for composing partial continuations”:

$$\begin{aligned} K(\#^u(WK_1W_1)) &\rightarrow W(\lambda v.K(\#^u(K_1v)))W_1 && \text{(lift arg)} \\ &\text{where } K \text{ is not of the form } (\#^uA')K' \\ (\#^u(WK_1W_1))K_2W_2 &\rightarrow W(\lambda v.(\#^u(K_1v))K_2W_2)W_1 && \text{(lift fun)} \\ &\text{where } u = (u_1 \rightarrow o) \rightarrow u_2 \rightarrow o \end{aligned}$$

The first rule (*lift arg*) has  $\#^u(WK_1W_1)$  as the *argument* to  $K$  and the second rule (*lift fun*) has  $\#^u(WK_1W_1)$  as the *function* applied to  $K_2$  and  $W_2$ . To enforce a left-to-right evaluation order of applications, the first rule is restricted to the case where the continuation has already been composed with the caller’s continuation. In the first rule, the invocation of  $W$  is a nontail call, but in the right-hand side the invocation of  $W$  is a tail call. Thus, we reduce the left-hand expression by one nontail call. Similarly, for the second rule, since the invocation of  $W$  is a nontail call, which becomes a tail call.

We need the restriction imposed by (*lift arg*). Consider the term  $xy(zw)$ . After the first algorithm, we have  $\lambda k.k(\#(x@y))@(\#(z@w))$ , where we ignore the types of  $\#$  and  $@$ . The first reduction yields  $\lambda k.(\#(x@y))(\lambda v.k(\#(@v)))(\#(z@w))$ . Now

consider (*lift arg*) without the condition. To determine the next redex, we would have instances of both lifting rules on the same subterm. This would violate our promise of left-to-right evaluation.

*Proposition 3.2*

For each rule if the left-hand side has type  $t$ , then so does the right-hand side.

#### 4 Discussion

Our goal is to make all calls be tail calls. In other words, we want all calls to be of the form  $(WKW_1)$ . We can see that the rules (*lift arg*) and (*lift fun*) do just that. The only difference in these two rules in the right-hand side is what the newly created  $K$  looks like. In both cases, however, the new  $K$  is formed by composing the context of  $(WK_1W_1)$  from the left-hand side and  $K_1$ . Clearly, each lift rule removes one nontail call and the CPS target language guarantees that no additional nontail calls are introduced. Therefore, our goal is met.

When we first formalized the algorithm, it contained three rules, but by re-arranging the argument order, so that the continuation was the first instead of the second argument, one of the three rules became a special case of one of the two remaining rules. In addition, we got the *let*-optimization as a byproduct. (By *let*-optimization, we mean an additional rule for translating source judgments that special-cases the application rule when  $e_1$  is a  $\lambda$ -expression. The inclusion of the administrative reductions obviates this special-casing.)

#### 5 Conclusion

From an intuitive perspective, we can see that the abort partial continuation absorbs the surrounding context up to the prompt, a bit at a time until the entire surrounding context has been consumed. Unlike conventional (Plotkin-style) CPS algorithms, which do not introduce partial continuations, ours has this bit of overhead. But these partial continuations disappear from the target-language output if they are not present in the source language.

Because each transformation not only preserves types but also preserves correctness, we know that once the first algorithm completes, we can apply the CPS semantics to reduce to the same final value after each reduction step of the second algorithm. As a result, we know that once no more lift rules apply, not only are we free of nontail calls, but we know that we have correctly preserved the semantics of the original program.

#### Acknowledgements

We are grateful for the comments of an anonymous referee.

#### References

- Danvy, O. & Filinski, A. (1990) Abstracting control. *Proceedings of the 1990 ACM Conference on LISP and Functional Programming*, pp. 151–160.

- Danvy, O. & Filinski, A. (1992) Representing control: A study of the CPS transformation. *Mathematical Structures in Computer Science*, **2**(4), 361–391.
- Felleisen, M., Friedman, D. P., Kohlbecker, E. & Duba, B. (1987) A syntactic theory of sequential control. *Theor. Comput. Sci.*, **52**(3), 205–237. (Preliminary version: Reasoning with continuations. *Proceedings 1st IEEE Symposium on Logic in Computer Science*, 1986.)
- Felleisen, M., Wand, M., Friedman, D. P. & Duba, B. F. (1988) Abstract continuations: A mathematical semantics for handling full functional jumps. *Proceedings 1988 ACM Conference on Lisp and Functional Programming*, pp. 52–62.
- Filinski, A. (1996) *Controlling effects*. PhD thesis, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA.
- Filinski, A. (1999) Representing layered monads. *Conference Record of POPL 99: 26th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pp. 175–188. ACM.
- Fischer, M. J. (1993) Lambda-calculus schemata. *Lisp & Symbolic Computation*, **6**(3/4), 259–288.
- Johnson, G. F. & Duggan, D. (1988) Stores and partial continuations as first-class objects in a language and its environment. *Conference Record of the 15th Annual ACM Symposium on Principles of Programming Languages*, pp. 158–168.
- Meyer, A. R. & Wand, M. (1985) Continuation semantics in typed lambda-calculi. *Proceedings Workshop Logics of Programs: Lecture Notes in Computer Science 193*, pp. 219–224.
- Plotkin, G. (1976) Call-by-name, call-by-value and the  $\lambda$ -calculus. *Theor. Comput. Sci.*, **1**(1), 125–159.
- Sabry, A. (1996) *Note on axiomatizing the semantics of control operators*. Technical report CIS-TR-96-03, Department of Computer and Information Science, University of Oregon.
- Sabry, A. & Felleisen, M. (1993) Reasoning about programs in continuation-passing style. *Lisp & Symbolic Computation*, **6**(3–4), 289–360.