

CS107: Intermediate Representations and CPS

Guannan Wei

guannan.wei@tufts.edu

October 1, 2026

Fall 2026

Tufts University

A Taste Of MiniScala

Our implementation of MiniScala is already quite powerful.

- We can do essentially everything we can do in C
- Missing features (structs, strings): can be implemented as arrays

```
def gcd(x: Int, y: Int): Unit =  
  if (y == 0)  
    printInt(x)  
  else  
    gcd(y, y % x);  
gcd(2016, 714)
```

What are still missing?

- Higher-level features: nested first-class functions, objects, etc.
- Code quality: optimizations

What are still missing?

- Higher-level features: nested first-class functions, objects, etc.
- Code quality: optimizations

ASTs although directly correspond to the source code, they do not provide a good basis for compiling higher-level features and optimizations.

We need to introduce **intermediate representations** that are better suited for these tasks.

Introducing Intermediate Representations

- Like ASTs, intermediate representations (IR) are the data structures used by the compiler to represent the program being compiled.
- An intermediate representation is also a proper language, with its own syntax and semantics.

Introducing Intermediate Representations

- Like ASTs, intermediate representations (IR) are the data structures used by the compiler to represent the program being compiled.
- An intermediate representation is also a proper language, with its own syntax and semantics.
- Many analyses and transformations (e.g. optimizations) are substantially easier to perform on some IRs than on others, so choosing a good IR is crucial.
- Some non-trivial compilers use several IRs, and they tend to become more low-level as the code approaches its final form.

IR #1: MiniScala/CPS - A Functional IR

A functional IR is an intermediate representation that is close to a (very) simple functional programming language. Typical functional IRs have the following characteristics:

- all primitive operations (e.g. arithmetic operations) are performed on atomic values (variables or constants),
- the result of these operations is always named,
- variables cannot be re-assigned.

IR #1: MiniScala/CPS - A Functional IR


A functional IR is an intermediate representation that is close to a (very) simple functional programming language. Typical functional IRs have the following characteristics:

- all primitive operations (e.g. arithmetic operations) are performed on atomic values (variables or constants),
- the result of these operations is always named,
- variables cannot be re-assigned.

One popular functional IR is Continuation-Passing Style (CPS), and we will be using a CPS-based functional IR called MiniScala/CPS.


Some of these characteristics are shared with more mainstream IRs, like SSA used in GCC and LLVM.

Continuation-Passing Style (CPS)



¹Image credit: <https://matt.might.net/articles/cps-conversion/>

New Compiler Architecture



Intermediate Representations

The revised MiniScala compiler manipulates a total of **four** languages:

- 1) MiniScala the source language – the surface language in which programs are written,
- 2) MiniScala/AST – the parsed and desugared ASTs of MiniScala (Project 1-3),
- 3) **MiniScala/CPS** – the main intermediate language, on which optimizations are performed (Project 4-6),
- 4) ASM is the assembly language of the target (virtual) machine.

Intermediate Representations

The revised MiniScala compiler manipulates a total of **four** languages:

- 1) MiniScala the source language – the surface language in which programs are written,
- 2) MiniScala/AST – the parsed and desugared ASTs of MiniScala (Project 1-3),
- 3) **MiniScala/CPS** – the main intermediate language, on which optimizations are performed (Project 4-6),
- 4) ASM is the assembly language of the target (virtual) machine.

There are interpreters for the last three languages, serving as the “semantics” of their respective languages. They are useful to check that a program behaves in the same way as it undergoes transformation.

RABBIT:
A Compiler for SCHEME
(A Dialect of LISP)

A Study in
Compiler Optimization

Based on Viewing
LAMBDA as RENAME
and
PROCEDURE CALL as GOTO

using the techniques of
Macro Definition of Control and Environment Structures
Source-to-Source Transformation
Procedure Integration
and
Tail-Recursion

Guy Lewis Steele Jr.
Massachusetts Institute of Technology

May 1978

A Bit of History

A subset of SCHEME (rather than triples, for example) serves as the representation intermediate between the optimized SCHEME code and the final output code; code is expressed in this subset in the so-called continuation-passing style. As a subset of SCHEME, it enjoys the same theoretical properties; one could even apply the same optimizer used on the input code to the intermediate code. However, the subset is so chosen that all temporary quantities are made manifest as variables, and no control stack is needed to evaluate it. As a result, this apparently applicative representation admits an imperative interpretation which permits easy transcription to final imperative machine code. These qualities suggest that an applicative language like SCHEME is a better candidate for an UNCOL than the more imperative candidates proposed to date.

Continuation-Passing Style (CPS)

What is continuation-passing style (CPS)?

A programming style or program representation where control is made explicit by passing “**the rest of the computation**” as an extra argument to functions.

Continuation-Passing Style (CPS)

What is continuation-passing style (CPS)?

A programming style or program representation where control is made explicit by passing “**the rest of the computation**” as an extra argument to functions.

Why compiling to CPS?

- CPS makes control flow explicit, which simplifies many analyses and transformations.
- Uniform representation of control flow: function calls, returns, jumps, exceptions, etc. are all represented uniformly as continuation invocations.

Continuations are used for several purposes in our compiler for MiniScala/CPS:

- 1) They represent **code blocks** which can be “jumped to” from several locations.
- 2) They represent the code to execute **after** a function call.
- 3) Control transfer (jump, return, break, etc.) is represented as invoking continuations.

Every function gets a continuation as argument, which it must invoke with its return value.

Code Example

To illustrate the differences between the various IRs, we will use a program fragment to compute and print the greatest common divisor (GCD) of 2016 and 714.

The MiniScala version of that fragment could be:

```
def gcd(x: Int, y: Int): Unit =  
  if (y == 0)  
    printInt(x)  
  else  
    gcd(y, y % x);  
  
gcd(2016, 714)
```

GCD in MiniScala/CPS

The MiniScala/CPS version of the GCD program fragment looks as follows:

// source code

```
def gcd(x: Int, y: Int): Unit =  
  if (y == 0)  
    printInt(x)  
  else  
    gcd(y, x % y);  
  
gcd(2016, 714)
```

// CPS IR

```
deff gcd(c, x, y) = {  
  defc ct() = { printInt(c, x) };  
  defc cf() = {  
    valp t = x % y;  
    gcd(c, y, t);  
  }  
  vall z = 0;  
  if (y == z) ct() else cf()  
};  
defc ce(x) = { vall z = 0; halt(z) };  
vall x = 2016; vall y = 714;  
gcd(ce, x, y)
```

- The grammar of MiniScala/CPS is defined as follows:

```
c, f, n ∈ <N> ::= <ident> // names
l ∈ <L> ::= <integer> | <char> | <boolean> | <unit> // literals
p ∈ <P> ::= '+' | '-' | '*' | '/' | '%' | ... // arith op
pb ∈ <C> ::= '<' | '<=' | '==' | '!=' | '>=' | '>' // boolean op

e ∈ <T> ::= vall <N> = <L>; <T> // terms
          | valp <N> = <P>(<N>, ...); <T>
          | defc <N>(<N>, ...) = { <T> }; <T>
          | deff <N>(<N>, ...) = { <T> }; <T>
          | <N>(<N>, ...)
          | if (<N> <C> <N>) <N>() else <N>()
          | halt(<N>)
```

- Bind the name n to the literal value $\mathfrak{l}$ in expression e :

`val $\mathfrak{l}$  n =  $\mathfrak{l}$ ; e`

The literal value can be an integer, a character, a boolean or the unit value.

- Bind the name n to the literal value $\mathfrak{l}$ in expression e :

$\text{val}_{\mathfrak{l}} n = \mathfrak{l}; e$

The literal value can be an integer, a character, a boolean or the unit value.

- Bind the name n to the result of the application of primitive p to the value of n_1 , ... in expression e :

$\text{val}_p n = p(n_1, \dots); e$

The primitive p cannot be a logical (i.e. boolean) primitive, as such primitives are only meant to be used in conditional expressions - see later.

- Define functions $f_1, \dots$ with arguments $n_{1,1}, \dots$ and **return continuation** $c_1, \dots$ in expression e .

`deff f1(c1, n1,1, ...) = { e1 }; deff f2 = ...; ...; e`

The functions can be mutually recursive. The return continuation c_1 takes a single argument: the return value. Applying it is interpreted as returning from the function.

- Define functions $f_1, \dots$ with arguments $n_{1,1}, \dots$ and **return continuation** $c_1, \dots$ in expression e .

`deff f1(c1, n1,1, ...) = { e1 }; deff f2 = ...; ...; e`

The functions can be mutually recursive. The return continuation c_1 takes a single argument: the return value. Applying it is interpreted as returning from the function.

- Apply function f to return continuation c and arguments $n_1, \dots$
`f(c, n1, ...)`

The name c must either be bound in an enclosing `defc` or be the name of the return continuation of the current function.

MiniScala/CPS Local Continuations

- Define **local continuations** c_1 with arguments $n_1, \dots$ and body e_1 in e :

`defc  $c_1(n_1, \dots) = \{ e_1 \}; e$`

Interpretation: like a local function that never returns. e could contain other continuation definitions too.

MiniScala/CPS Local Continuations

- Define **local continuations** c_1 with arguments $n_1, \dots$ and body e_1 in e :

`defc c1( $n_1, \dots$ ) = {  $e_1$  }; e`

Interpretation: like a local function that never returns. e could contain other continuation definitions too.

- Apply continuation c to the value of $n_1, \dots$
 $c(n_1, \dots)$

The name c must either be bound by an enclosing `defc` or be the name of the return continuation of the current function.

- If c is a local continuation, $c(\dots)$ can be seen as a jump with arguments.
- If c refers to the current return continuation, $c(\dots)$ can be seen as a return from the current function, with the given return value.

MiniScala/CPS Control Constructs

- Test whether the condition p is true for the value of n_1 and n_2 , then apply continuation c_t if it is, or c_f if it isn't.

if (n_1 p_b n_2) $c_t()$ **else** $c_f()$

Both continuation c_t and c_f must have no parameter. The primitive p_b must be a logical primitive.

Note: **if** is a branching form of continuation invocation for parameterless continuations. It is therefore a conditional version of $c()$.

MiniScala/CPS Control Constructs

- Test whether the condition p is true for the value of n_1 and n_2 , then apply continuation c_t if it is, or c_f if it isn't.

if (n_1 p_b n_2) $c_t()$ **else** $c_f()$

Both continuation c_t and c_f must have no parameter. The primitive p_b must be a logical primitive.

Note: **if** is a branching form of continuation invocation for parameterless continuations. It is therefore a conditional version of $c()$.

- Halts program execution:

halt(n)

Exit with the value bound to n (which must be an integer).

To make MiniScala/CPS programs easier to read and write, we allow to use the prefix and infix notations for the primitive operations.

```
vall t = 1;  
valp x = -t + t
```

is equivalent to the actual representation:

```
vall t = 1;  
valp mt = -t;  
valp x = +(mt, t)
```

A crucial notion in MiniScala/CPS is that of **local continuation**.

A local continuation is similar to a (local) function but with the following restrictions:

- continuations are *not* “first class citizens”: they cannot be stored in variables or passed as arbitrary arguments. The only exception being the return continuation,
- continuations never return, and must therefore be invoked in tail position only.

These restrictions enable continuations to be compiled much more efficiently than normal functions. This is the only reason why continuations exist as a separate construct.

Scopes of Continuation

The scoping rules of MiniScala/CPS are mostly the “obvious ones”. The only exception is the rule for continuation variables, which are not visible in nested functions!

For example, in the following code:

```
defc c0(r) = { printInt(ce, r); }  
deff f(c1, x) = {  
  valp t = x + x;  
  c1(t)  
}
```

c_0 is not visible in the body of f !

This guarantees that continuations are truly local to the function that defines them, and can therefore be compiled efficiently.

GCD in MiniScala/CPS

The MiniScala/CPS version of the GCD program fragment looks as follows:

// source code

```
def gcd(x: Int, y: Int): Unit =  
  if (y == 0)  
    printInt(x)  
  else  
    gcd(y, x % y);  
  
gcd(2016, 714)
```

// CPS IR

```
deff gcd(c, x, y) = {  
  defc ct() = { printInt(c, x) };  
  defc cf() = {  
    valp t = x % y;  
    gcd(c, y, t);  
  }  
  vall z = 0;  
  if (y == z) ct() else cf()  
};  
defc ce(x) = { vall z = 0; halt(z) };  
vall x = 2016; vall y = 714;  
gcd(ce, x, y)
```

Translating MiniScala to MiniScala/CPS

The translation from MiniScala/AST to MiniScala/CPS is specified as a function denoted by $\llbracket \cdot \rrbracket$ taking two arguments:

- 1) T , the MiniScala/AST term to be translated,
- 2) C , the context, a MiniScala/CPS term containing a hole into which a name bound to the value of the translated term has to be plugged.

Translating MiniScala to MiniScala/CPS

The translation from MiniScala/AST to MiniScala/CPS is specified as a function denoted by $\llbracket \cdot \rrbracket$ taking two arguments:

- 1) T , the MiniScala/AST term to be translated,
- 2) C , the context, a MiniScala/CPS term containing a hole into which a name bound to the value of the translated term has to be plugged.

This function is written in a “mixfix” notation, as follows:

$\llbracket T \rrbracket C$

The translation function must return a MiniScala/CPS term.

The **translation context** C is a MiniScala/CPS term representing the partial translation of the MiniScala expression surrounding the one being translated.

- This term contains a single hole, written $\square$, representing the currently unknown *name* that will be bound to the value of the expression being translated.
- The hole of a context C must eventually be filled with some name n , written as $C[n]$.

For example, $f(\square)[m]$ denotes plugging name m into the context $f(\square)$, resulting in the term $f(m)$.

Context Representation

Translation rules often build contexts that include other contexts. One (fictional) example to translate application $e_1(e_2)$ could be:

$$\llbracket e_1 \rrbracket (\llbracket e_2 \rrbracket \square(\square))$$

Context Representation

Translation rules often build contexts that include other contexts. One (fictional) example to translate application $e_1 (e_2)$ could be:

$$\llbracket e_1 \rrbracket (\llbracket e_2 \rrbracket \square(\square))$$

In such situations, using the anonymous hole ($\square$) is ambiguous. To resolve the ambiguity, we **represent contexts as meta-functions** taking a single (named) argument. The above is therefore written as:

$$\llbracket e_1 \rrbracket \lambda v_1 (\llbracket e_2 \rrbracket \lambda v_2 (v_1(v_2)))$$

where v_1/v_2 are the meta-variables representing the translation result of e_1/e_2 .

Context Representation

Translation rules often build contexts that include other contexts. One (fictional) example to translate application $e_1 (e_2)$ could be:

$$\llbracket e_1 \rrbracket (\llbracket e_2 \rrbracket \square(\square))$$

In such situations, using the anonymous hole ($\square$) is ambiguous. To resolve the ambiguity, we **represent contexts as meta-functions** taking a single (named) argument. The above is therefore written as:

$$\llbracket e_1 \rrbracket \lambda v_1 (\llbracket e_2 \rrbracket \lambda v_2 (v_1 (v_2)))$$

where v_1/v_2 are the meta-variables representing the translation result of e_1/e_2 .

With such a representation of contexts, filling the hole of context C with name n is performed by **meta-function application**, written $C[n]$.

The Translation In Scala

In Scala (our meta-language of MiniScala), the translation function $\llbracket \cdot \rrbracket$ is defined as a function with the following signature:

```
def MiniScalaToCPS(t: MSTree, C: Symbol => CPSTree): CPSTree
```

In the body of that function, plugging the context C with a name (i.e. a `Symbol`) bound to a Scala value named n is performed by Scala function application:

$C(n)$

To clarify the presentation, MiniScala terms appear in orange, MiniScala/CPS terms in purple, and meta-terms in black.

MiniScala/AST to MiniScala/CPS Translation (1)

Note: in the following expressions, all underlined names are fresh.

$\llbracket n \rrbracket \ C = C[n]$
where n is an identifier for immutable variable

$\llbracket l \rrbracket \ C = \text{val } \underline{n} = l; C[n]$
where l is a literal

MiniScala/AST to MiniScala/CPS Translation (1)

Note: in the following expressions, all underlined names are fresh.

$\llbracket n \rrbracket \ C = C[n]$
where n is an identifier for immutable variable

$\llbracket l \rrbracket \ C = \text{val}_{\underline{l}} \ \underline{n} = l; C[n]$
where l is a literal

$\llbracket \text{val } n_1 = e_1; e \rrbracket \ C =$
 $\llbracket e_1 \rrbracket (\lambda v \ (\text{val}_p \ n_1 = \text{id}(v); \llbracket e \rrbracket \ C))$

MiniScala/AST to MiniScala/CPS Translation (1)

Note: in the following expressions, all underlined names are fresh.

$\llbracket n \rrbracket$ $C = C[n]$
where n is an identifier for immutable variable

$\llbracket l \rrbracket$ $C = \text{val}_{\underline{l}} \underline{n} = l; C[n]$
where l is a literal

$\llbracket \text{val } n_1 = e_1; e \rrbracket C =$
 $\llbracket e_1 \rrbracket (\lambda v (\text{val}_p n_1 = \text{id}(v); \llbracket e \rrbracket C))$

$\llbracket \text{var } n_1 = e_1; e \rrbracket C =$
 $\text{val}_{\underline{l}} \underline{s} = 1;$
 $\text{val}_p n_1 = \text{block-alloc-var}(s);$
 $\text{val}_{\underline{l}} \underline{z} = 0;$
 $\llbracket e_1 \rrbracket (\lambda v (\text{val}_p \underline{d} = \text{block-set}(n_1, z, v); \llbracket e \rrbracket C))$

MiniScala/AST to MiniScala/CPS Translation (2)

```
[[ n1 = e1 ]] C =  
  vall z = 0;  
  [[e1]](λv (valp d = block-set(n1, z, v); C[v] ))
```

```
[[n]] C =  
  vall z = 0;  
  valp v = block-get(n, z); C[v]  
where n is an identifier for mutable variable
```

MiniScala/AST to MiniScala/CPS Translation (2)

```
[[ n1 = e1 ]] C =  
  vall z = 0;  
  [[e1]](λv (valp d = block-set(n1, z, v); C[v] ))
```

```
[[n]] C =  
  vall z = 0;  
  valp v = block-get(n, z); C[v]  
  where n is an identifier for mutable variable
```

```
[[ def f1(n1,1: _, ...) = e1; def ... ; e ]] C =  
  deff f1(c, n1,1, ...) = {  
    [[e1]](λv(c(v)))  
  };  
  deff ... ;  
  [[e]] C
```

MiniScala/AST to MiniScala/CPS Translation (3)

```
[[ e(e1, e2, ...) ]] c =  
  [[e]](λv([e1](λv1([e2](λv2(...  
    defc c(r) = { c[r] };  
    v(c, v1, v2 ...))))))
```

MiniScala/AST to MiniScala/CPS Translation (3)

```
[[ e(e1, e2, ...) ] C =  
  [[e]](λv([e1]](λv1([e2]](λv2(...  
    defc c(r) = { C[r] };  
    v(c, v1, v2 ...))))))
```

```
[[ if (p(e1, ...)) e2 else e3 ] C =  
  defc c(r) = { C[r] };  
  defc ct() = { [[e2]](λv2(c(v2))) };  
  defc cf() = { [[e3]](λv3(c(v3))) };  
  [[e1]](λv1(... (if (p(v1 ...)) ct else cf)))  
  where p is a logical primitive
```

MiniScala/AST to MiniScala/CPS Translation (4)

```
[[ if (e1) e2 else e3 ]] c =  
  defc c(r) = { c[r] };  
  defc ct() = { [[e2]](λv2(c(v2))) };  
  defc cf() = { [[e3]](λv3(c(v3))) };  
  vall f = false;  
  [[e1]](λv1(if (v1 ≠ f) ct else cf))
```

MiniScala/AST to MiniScala/CPS Translation (4)

```
[[ if (e1) e2 else e3 ]] C =  
  defc c(r) = { C[r] };  
  defc ct() = { [[e2]](λv2(c(v2))) };  
  defc cf() = { [[e3]](λv3(c(v3))) };  
  vall f = false;  
  [[e1]](λv1(if (v1 ≠ f) ct else cf))  
  
[[ while (e1) e2; e3 ]] C =  
  defc loop() = {  
    defc c() = { [[e3]] C };  
    defc ct() = { [[e2]](λv(loop())) };  
    vall f = false;  
    [[e1]](λv(if (v ≠ f) ct else c))  
  };  
  loop()
```

MiniScala/AST to MiniScala/CPS Translation (5)

$\llbracket p(e_1, e_2, \dots) \rrbracket C =$
 $\llbracket \text{if } (p(e_1, e_2, \dots)) \text{ true else false} \rrbracket C$
 where p is a logical primitive

$\llbracket p(e_1, e_2, \dots) \rrbracket C =$
 left as an exercise
 where p is not a logical primitive

In which context should a complete program be translated?

The simplest answer is a context that halts execution with an exit code of 0 (no error), that is:

```
 $\lambda v(\text{val}_l\ z = 0; \text{halt}(z))$ 
```

An alternative would be to do something with the value v produced by the whole program, e.g. use it as the exit code instead of 0, print it, etc.

Example

Let k_0 be the initial context $\lambda v(\text{val}_l z = 0; \text{halt}(z))$.

```
[[ f(g(1, 2)) ]] k_0
{ f(g(1, 2)) is a function application }
= [[f]](λv([g(1, 2)](λv_1
  def_c c_1(r) = { k_0[r] };
  v(c_1, v_1))))
= ???
```