

# CS107: Optimizations

---

**Guannan Wei**

guannan.wei@tufts.edu

February 26, 2026

Spring 2026

Tufts University

- Project 4 dues on **Sunday, March 1st at 11:59pm.**
- In-class midterm: next Thursday March 5
- Review session: next Tuesday
- Practice problems set has been released on Canvas

The goal of optimization is to rewrite a program being compiled to a new program that is simultaneously:

- behaviorally **equivalent** to the original one, and
- **better** in some respect - e.g., faster, smaller, more energy-efficient, etc.

The goal of optimization is to rewrite a program being compiled to a new program that is simultaneously:

- behaviorally **equivalent** to the original one, and
- **better** in some respect - e.g., faster, smaller, more energy-efficient, etc.

Optimizations can be broadly split into two classes:

- **machine-independent optimizations**, e.g. constant folding or dead code elimination, which are high-level and do not depend on the target architecture,
- **machine-dependent optimizations**, e.g. register allocation or instruction scheduling, which are low-level and depend on the target architecture.

In this lecture, we will examine a simple set of machine independent rewriting optimizations.

Most of them are relatively simple rewrite rules that transform the source program to a shorter, equivalent one.

- Constant propagation
- Constant folding
- Dead-code elimination
- Common subexpression elimination
- Eta reduction
- Algebraic laws
- Function inlining
- Contification

# The Importance Of IRs

The intermediate representation (IR) on which rewriting optimizations are performed can have a dramatic impact on their ease of implementation.

A rewriting optimization generally works in two steps:

- 1) the program is analyzed to find rewrite opportunities,
- 2) the program is rewritten based on the opportunities identified in the first step.

A good IR should make both steps as easy as possible. The following examples illustrate the importance of using the right IR.

## Case 1: Constant Propagation

To illustrate the impact of IR on the analysis step, consider the following program fragment in some imaginary IR:

```
x ← 7
```

Is it legal to perform constant propagation and blindly replace all later occurrences of  $x$  by 7?

## Case 1: Constant Propagation

To illustrate the impact of IR on the analysis step, consider the following program fragment in some imaginary IR:

```
x ← 7
```

Is it legal to perform constant propagation and blindly replace all later occurrences of  $x$  by 7?

It depends on the IR:

- If multiple assignments to the same variable are allowed, then additional (data-flow) analyses are required to answer the question, as  $x$  might be re-assigned later.
- However, if multiple assignments to the same variable are prohibited, then yes, all occurrences of  $x$  can be unconditionally replaced by 7!

## Other Simple Optimizations

Apart from constant propagation, many simple optimizations are made hard by the presence of multiple assignments to a single variable:

- common-subexpression elimination, which consists in avoiding the repeated evaluation of expressions,
- (simple) dead code elimination, which consists in removing assignments to variables whose value is not used later,
- etc.

In all cases, analyses are required to distinguish the various “versions” of a variable that appear in the program.

**Conclusion:** a good IR should not allow multiple assignments to a variable!

## Case 2: Inlining

Inlining (or in-line expansion) consists in replacing a call to a function by a copy of the function body, with parameters replaced by the actual arguments. It is a very important compiler optimization, as it often opens the door to further optimizations.

Some aspects of the IR can have an important impact on the implementation of inlining. To illustrate this, let us examine some problems that can occur when performing inlining directly on the AST - a choice that might seem reasonable at first sight.

## Naive Inlining: Problem #1

```
def printReturn(x: Int) = { printInt(x); x };  
def twice(y: Int) = y + y;  
def f(z: Int) = twice(printReturn(z));
```

Incorrect inlining of twice in f:

```
def f(z: Int) = printReturn(z) + printReturn(z)
```

## Naive Inlining: Problem #1

```
def printReturn(x: Int) = { printInt(x); x };  
def twice(y: Int) = y + y;  
def f(z: Int) = twice(printReturn(z));
```

Incorrect inlining of `twice` in `f`:

```
def f(z: Int) = printReturn(z) + printReturn(z)
```

`z` is printed twice!

Possible solution: bind actual parameters to variables (using a `Let`) to ensure that they are evaluated at *most* once.

## Naive Inlining: Problem #2

```
def first[T, U](t: T, u: U) = t;  
def printReturn(z: Int) = first[Int, Unit](z, printInt(z));
```

Incorrect inlining of first in printReturn :

```
def printReturn(z: Int) = z
```

## Naive Inlining: Problem #2

```
def first[T, U](t: T, u: U) = t;  
def printReturn(z: Int) = first[Int, Unit](z, printInt(z));
```

Incorrect inlining of first in printReturn :

```
def printReturn(z: Int) = z
```

z is not printed at all!!!

Possible solution: bind actual parameters to variables (using a Let) to ensure that they are evaluated at *least* once.

The two pitfalls presented earlier can be avoided by bindings actual arguments to variables (using a `Let`) before using them in the body of the inlined function.

However, a properly-designed IR can also avoid the problems altogether by ensuring that actual parameters are always atoms, i.e. variables or constants.

Conclusion: a good IR should only allow atomic arguments to functions.

Using the two very basic criteria we identified, we can evaluate the various classes of IRs we have seen:

- standard RTL/CFG is bad in that its variables are mutable; however, it allows only atoms as function arguments, which is good,
- RTL/CFG in SSA form, MiniScala/CPS and similar functional IRs are good in that their variables are immutable, and they only allow atoms as function arguments.

The rewriting optimizations for MiniScala/CPS are specified as a set of rewriting rules of the form  $T \mapsto_{opt} U$ .

These rules rewrite a MiniScala/CPS term  $T$  to an equivalent - but hopefully more efficient - term  $U$ .

## Optimization Contexts

Optimization rewriting rules cannot be applied anywhere, but only in specific locations. For example, it would be incorrect to try to rewrite the parameter list of a function. This constraint can be captured by specifying all the *contexts* in which it is valid to perform a rewrite, where a context is a term with a single hole denoted by  $\square$ .

The hole of a context  $C$  can be plugged with a term  $T$ , an operation written as  $C[T]$ .

### Example

If  $C$  is **if**  $\square$  **ct** **else**  $cf$ , then  $C[x < y]$  is **if**  $(x < y)$  **ct** **else**  $cf$ .

## Optimization Contexts

The optimization contexts for MiniScala/CPS are generated by the following *grammar*:

$$\begin{aligned} C_{opt} ::= & \square \\ & | \text{val}_l \ n = \mathbb{1}; C_{opt} \\ & | \text{val}_p \ n = p(n_1, \dots); C_{opt} \\ & | \text{def}_c \ c_1 \dots ; \text{def}_c \ c_i(n_{i,1}, \dots) = \{ C_{opt} \}; \dots; \text{def}_c \ c_k \dots; e \\ & | \text{def}_c \ c_1 \dots; C_{opt} \\ & | \text{def}_f \ f_1 \dots ; \text{def}_f \ f_i(c_i, n_{i,1}, \dots) = \{ C_{opt} \}; \dots; \text{def}_c \ f_k \dots; e \\ & | \text{def}_f \ f_1 \dots; C_{opt} \end{aligned}$$

A  $C_{opt}$  describes the surrounding of a term in which it is valid to perform an optimization rewrite.

Using the optimization rewriting rules (presented later) and the optimization contexts, it is possible to specify the optimization relation  $opt$  that rewrites a term to an optimized version:

$$C_{opt}[T] \Rightarrow_{opt} C_{opt}[U] \text{ where } T \mapsto_{opt} U$$

We distinguish two classes of rewriting rules:

- 1) **shrinking rules** rewrite a term to an equivalent but **smaller** one,
- 2) **non-shrinking rules** rewrite a term to an equivalent but **potentially larger** one.

## (Non-)shrinking rules

We distinguish two classes of rewriting rules:

- 1) **shrinking rules** rewrite a term to an equivalent but **smaller** one,
- 2) **non-shrinking rules** rewrite a term to an equivalent but **potentially larger** one.

Shrinking rules can be applied at will, possibly until the term is fully reduced, while non-shrinking rules cannot, as they could result in infinite expansion. Heuristics must be used to decide when to apply non-shrinking rules.

Except for (non-linear) inlining, all optimizations we will see are shrinking.

## Dead Code Elimination

Dead code elimination removes all code that neither performs side effects nor produces a value used later.

```
vall n1 = 42; // dead-code  
vall n2 = 10;  
n2
```

## Dead Code Elimination

Dead code elimination removes all code that neither performs side effects nor produces a value used later.

Rules:

$$\text{val}_l \ n = \perp; e \mapsto_{opt} e \quad [ \text{if } n \text{ is not free in } e ]$$

## Dead Code Elimination

Dead code elimination removes all code that neither performs side effects nor produces a value used later.

Rules:

```
vall n = l; e  
   $\mapsto_{opt}$  e [ if n is not free in e ]  
valp n = p(n1, ...); e  
   $\mapsto_{opt}$  e [ if n is not free in e and  
              p  $\notin$  { byte-read, byte-write, block-set } ]
```

## Dead Code Elimination

Dead code elimination removes all code that neither performs side effects nor produces a value used later.

Rules:

$\text{val}_l \ n = l; e$   
 $\mapsto_{opt} e$  [ if  $n$  is not free in  $e$  ]

$\text{val}_p \ n = p(n_1, \dots); e$   
 $\mapsto_{opt} e$  [ if  $n$  is not free in  $e$  and  
 $p \notin \{ \text{byte-read}, \text{byte-write}, \text{block-set} \}$  ]

$\text{def}_f \ f_1 \dots; \text{def}_f \ f_i \dots; \text{def}_f \ f_k \dots; e$   
 $\mapsto_{opt} \text{def}_f \ f_1; \dots; \text{def}_f \ f_{i-1} \dots; \text{def}_f \ f_{i+1}; \text{def}_f \ f_k \dots; e$   
[ if  $f_i$  is not free in  $\{ f_1, \dots, f_{i-1}, f_{i+1}, \dots, f_k, e \}$  ]

The rule for continuations is similar to the one for functions

## Dead Code Elimination

The rewriting rules to eliminate dead functions does not eliminate dead but mutually-recursive functions.

```
deff even(x: Int) = x == 0 || odd(x - 1);  
deff odd(x: Int) = x == 1 || even(x - 1);  
0
```

## Dead Code Elimination

The rewriting rules to eliminate dead functions does not eliminate dead but mutually-recursive functions.

```
deff even(x: Int) = x == 0 || odd(x - 1);  
deff odd(x: Int) = x == 1 || even(x - 1);  
0
```

To handle them, one must start from the main expression of the program, and identify all functions transitively reachable from it. All unreachable functions are dead.

The rule for continuations has the same problem, which admits a similar solution. The only difference is that, continuations being local to functions, the reachability analysis can start in function bodies.

Common subexpression elimination (CSE) avoids the repeated evaluation of a single expression.

Certainly

```
vall n1 = 42 + 10;  
vall n2 = 42 + 10;  
f(n1, n2)
```

can be optimized to

```
vall n1 = 42 + 10;  
f(n1, n1)
```

Common subexpression elimination (CSE) avoids the repeated evaluation of a single expression.

But also

```
vall n1 = 42 + 10;  
deff g() = {  
  vall n2 = 42 + 10;  
  f(n1, n2)  
}
```

should be optimized too.

How should we express the optimization rule for this case?

Common subexpression elimination (CSE) avoids the repeated evaluation of a single expression.

$$\text{val}_l \ n_1 = \perp; \ C_{opt} [ \text{val}_l \ n_2 = \perp; \ e \ ] \\ \mapsto_{opt} \ \text{val}_l \ n_1 = \perp; \ C_{opt} [ \ e[n_2 \mapsto n_1] \ ]$$

$$\text{val}_p \ n_1 = a_1 + a_2; \ C_{opt} [ \text{val}_p \ n_2 = a_1 + a_2; \ e \ ] \\ \mapsto_{opt} \ \text{val}_p \ n_1 = a_1 + a_2; \ C_{opt} [ \ e[n_2 \mapsto n_1] \ ]$$

etc.

Recall that

$$C_{opt} [T] \Rightarrow_{opt} C_{opt} [U] \ \text{where} \ T \mapsto_{opt} U$$

so we allow optimizing cases such as  $C_{opt}^1 [ \dots \ C_{opt}^2 [ \dots ] ]$ .

These simple rules for common subexpression elimination are also a bit too naive and miss some opportunities because of scoping.

For example, the common subexpression  $y + z$  below is not optimized by them:

```
defc c1() = {  
  valp n1 = y + z; ...  
};  
defc c2() = {  
  valp n2 = y + z; ...  
};  
...
```

Variants of the standard  $\eta$ -reduction can be performed to remove redundant definitions.

```
def g(x: Int) = ...  
def f(x: Int) = g(x) // f is an  $\eta$ -expansion of g  
f(42)
```

can be optimized to

```
def g(x: Int) = ...  
g(42)
```

## Eta Reduction

Variants of the standard  $\eta$ -reduction can be performed to remove redundant definitions.

Formally:

```
defc c1(...) = { e1 }; ...;  
defc ci(n1, ...) = { d(n1, ...) }; // d is a continuation  
defc ck(...) = { ek }; e  
↳opt defc c1(...) = { e1[ci ↦ d] }; ...; defc ck(...) = { ek[ci ↦ d] };  
      e[ci ↦ d]
```

```
deff f1(...) = { e1 }; ...;  
deff fi(c, n1, ...) = { g(c, n1, ...) }; // g is a function  
deff fk(...) = { ek }; e  
↳opt deff f1(...) = { e1[fi ↦ g] }; ...; deff fk(...) = { ek[fi ↦ g] };  
      e[fi ↦ g]
```

## Constant folding (1)

Constant folding replaces a constant expression by its value.

For example,

```
vall n = 42;  
vall m = 10;  
valp k = n + m;  
...
```

can be optimized to

```
vall n = 42;  
vall m = 10;  
vall k = 52;  
...
```

## Constant folding (1)

Constant folding replaces a constant expression by its value.

Example for addition:

$$\begin{aligned} & \text{val}_l n_1 = \lrcorner 1; \\ & \mathcal{C}_{opt}^1 [ \text{val}_l n_2 = \lrcorner 2; \\ & \quad \mathcal{C}_{opt}^2 [ \text{val}_p n_3 = n_1 + n_2; e ] ] \\ \mapsto_{opt} & \\ & \text{val}_l n_1 = \lrcorner 1; \\ & \mathcal{C}_{opt}^1 [ \text{val}_l n_2 = \lrcorner 2; \\ & \quad \mathcal{C}_{opt}^2 [ \text{val}_l n_3 = (\lrcorner 1 + \lrcorner 2); e ] ] \end{aligned}$$

Similar rules exist for other arithmetic primitives.

## Constant folding (2)

Primitives appearing in conditional expressions are also amenable to constant folding, for example:

```
if (n == n) ct else cf  $\mapsto_{opt}$  ct()
```

```
if (n != n) ct else cf  $\mapsto_{opt}$  cf()
```

```
vall n1 = l1;  
  Copt1 [ vall n2 = l2;  
    Copt2 [ if (n1 == n2) ct else cf ] ]
```

```
 $\mapsto_{opt}$   
vall n1 = l1;  
  Copt1 [ vall n2 = l2;  
    Copt2 [ ct() ] ] [ if l1 = l2 ]
```

## Neutral/Absorbing Elements

Uses of neutral and absorbing elements of arithmetic primitives can be simplified. For multiplication, this is expressed by the following rules:

$$\begin{array}{l} \mathbf{val}_l \ n_1 = 0; C_{opt} [ \mathbf{val}_p \ n = n_1 * n_2; e ] \\ \mapsto_{opt} \ \mathbf{val}_l \ n_1 = 0; C_{opt} [ e [n \mapsto n_1] ] \end{array}$$

$$\begin{array}{l} \mathbf{val}_l \ n_2 = 0; C_{opt} [ \mathbf{val}_p \ n = n_1 * n_2; e ] \\ \mapsto_{opt} \ \mathbf{val}_l \ n_2 = 0; C_{opt} [ e [n \mapsto n_2] ] \end{array}$$

$$\begin{array}{l} \mathbf{val}_l \ n_1 = 1; C_{opt} [ \mathbf{val}_p \ n = n_1 * n_2; e ] \\ \mapsto_{opt} \ \mathbf{val}_l \ n_1 = 1; C_{opt} [ e [n \mapsto n_2] ] \end{array}$$

$$\begin{array}{l} \mathbf{val}_l \ n_2 = 1; C_{opt} [ \mathbf{val}_p \ n = n_1 * n_2; e ] \\ \mapsto_{opt} \ \mathbf{val}_l \ n_2 = 1; C_{opt} [ e [n \mapsto n_1] ] \end{array}$$

Similar rules exist for other arithmetic operators.

# Block Primitives

Block primitives are harder to optimize, because block elements can be modified.

However, some blocks used by the compiler, e.g. to implement closures, are known to be constant once initialized. This makes rewritings like the following possible:

```
valp b = block-alloc-k(s);  
  Copt1 [ valp t = block-set(b, i, v);  
    Copt2 [ valp n = block-get(b, i); e ] ]
```

$\mapsto_{opt}$

```
valp b = block-alloc-k(s);  
  Copt1 [ valp t = block-set(b, i, v);  
    Copt2 [ e[n  $\mapsto$  v] ] ]
```

[ when tag k identifies a block that is not modified after initialization, e.g. a closure block ]

MiniScala/CPS contains the following block primitives:

- `block-alloc-n(size)`
- `block-tag(block)`
- `block-size(block)`
- `block-get(block, index)`
- `block-set(block, index, value)`

Informally describe some rewriting optimizations that could be performed on these primitives, and in which conditions they could be performed.

## Shrinking Inlining

A non-recursive continuation or function that is applied exactly once - i.e. used in a linear fashion - can always be inlined without making the code grow:

```
deff f1...; deff fi(ci, ni,1, ...) = { ei }; ...; deff fk...;
  Copt[ fi(c, m1, ...) ]
```

$\mapsto_{opt}$

```
deff f1...; deff fi-1...; deff fi+1...; ...; deff fk...;
  Copt[ ei[ci ↦ c, ni,1 ↦ m1] ]
[ when fi is not free in Copt, e1, ..., ek ]
```

The rule for continuations is similar.

Non-linear inlining can also be performed trivially in MiniScala/CPS, either for continuation or for functions (illustrated here):

```
...; deff fi(ci, ni,1, ...) = { ei }; ...;  
Copt[ fi(c, m1, ...) ]
```

$\mapsto_{opt}$

```
...; deff fi(ci, ni,1, ...) = { ei }; ...;  
Copt[ ei[ci  $\mapsto$  c, ni,1  $\mapsto$  m1] ]
```

To preserve the uniqueness of names, fresh versions of bound names should be created during general inlining.

## Inlining Heuristics (1)

The problem of these rules is that they are not shrinking and rewriting does not even terminate with recursive continuations or functions.

Since non-shrinking inlining cannot be performed indiscriminately, heuristics are used to decide whether a candidate function should be inlined at a given call site.

## Inlining Heuristics (1)

The problem of these rules is that they are not shrinking and rewriting does not even terminate with recursive continuations or functions.

Since non-shrinking inlining cannot be performed indiscriminately, heuristics are used to decide whether a candidate function should be inlined at a given call site.

These heuristics typically combine several factors, like:

- the **size** of the candidate function — smaller ones should be inlined more eagerly than bigger ones,
- the **number of times the candidate is called** in the whole program — a function called only a few times should be inlined,

to be continued...

## Inlining Heuristics (2)

- the nature of the candidate — not much gain can be expected from the inlining of a recursive function,
- the kind of arguments passed to the candidate, and/or the way these are used in the candidate — constant arguments could lead to further reductions in the inlined candidate, especially if it combines them with other constants,
- etc.

Imagine an imperative intermediate language equipped with a return statement to return from the current function to its caller.

1. What would happen when inlining a function containing such a return statement?

Imagine an imperative intermediate language equipped with a return statement to return from the current function to its caller.

1. What would happen when inlining a function containing such a return statement?
2. In our CPS IR, how would like return statement look like? Does it suffer from the same problem as in the previous question?

Contification is the name generally given to an optimization that transforms functions into (local) continuations.

When applicable, this transformation is interesting because it transforms expensive functions (that are compiled as closures) to inexpensive continuations (that are compiled as blocks).

## Contification Example

Contification can for example transform the loop function in the MiniScala example below to a local continuation, leading to efficient compiled code.

```
def fact(x: Int) = {  
  def loop(i: Int, res: Int): Int = {  
    if (i > x)  
      res  
    else  
      loop(i + 1, res * i)  
  };  
  loop(x, 1)  
}
```

A MiniScala/CPS function is contifiable if and only if it always *returns to the same location*, because then it does not need a return continuation.

A MiniScala/CPS function is contifiable if and only if it always *returns to the same location*, because then it does not need a return continuation.

- For a non-recursive function, this condition is satisfied if and only if that function is only used in AppF nodes, in function position, and always passed the same return continuation.
- For recursive functions, the condition is slightly more involved - see later.

## Non-Recursive Contification

The contification of the non-recursive function  $f$  is given by:

```
def $f$  f(c, n1, ...) = { e };  
Copt1 [ Copt2 [ f(c0, n1,1, ...), f(c0, n2,1, ...), ... ] ]
```

$\mapsto_{opt}$

```
Copt1 [ def $c$  m(n1, ...) = { e[c  $\mapsto$  c0] };  
Copt2 [ m(n1,1, ...), m(n2,1, ...), ... ] ]
```

where  $f$  does not appear free in either  $C_{opt}^1$  or  $C_{opt}^2$ , and  $C_{opt}^2$  is the smallest (multi-hole) context enclosing all applications of  $f$ .

It ensures that the scope of  $m$  is as small as possible, and therefore that  $m$  obeys the scoping rules for continuations.

In this rule,  $c_0$  is the (single) return continuation that is passed to function  $f$ .

A set of mutually-recursive functions  $F = \{ f_1, \dots, f_n \}$  is contifiable, which we write  $\text{Cnt}(F)$ , if and only if every function in  $F$  is always used in one of the following two ways:

1. applied to a common return continuation, or
2. called in tail position by a function in  $F$ .

Intuitively, this ensures that all functions in  $F$  eventually return through the same continuation.

## Example

As an example, functions `even` and `odd` in the MiniScala/CPS translation of the following MiniScala term are contifiable:

```
deff even(x: Int) = x == 0 || odd(x - 1);  
deff odd(x: Int)  = x == 1 || even(x - 1);  
even(12)
```

$\text{Cnt}(F = \{ \mathbf{even}, \mathbf{odd} \})$  is satisfied since:

- the single use of **odd** is a tail call from **even**  $\in F$ ,
- **even** is tail-called from **odd**  $\in F$  and called with the continuation of the `LetRec` statement - the common return continuation  $c_0$  for this example.

## Recursive Contification

Given a set of mutually-recursive functions

```
deff f1...; ...; deff fk...;  
e
```

the condition  $\text{Cnt}(F)$  for some  $F \subseteq \{ f_1, \dots, f_k \}$  can only be true if all the non-tail calls to functions in  $F$  appear either:

- in the term  $e$ , or
- in the body of exactly one function  $f_i \notin F$ .

Therefore, two separate rewriting rules must be defined, one per case.

## Recursive Contification #1

When all non tail calls to functions in  $F = \{ f_1, \dots, f_i \}$  appear in the body of the `LetRec`, and  $\text{Cnt}(F)$  holds, contification is performed by the following rewriting:

$$\mathbf{def}_f f_1(c_1, n_{1,1}, \dots) = \{ e_1 \}; \dots; \mathbf{def}_f f_k \dots;$$
$$C_{opt}[ e ]$$
$$\mapsto_{opt}$$
$$\mathbf{def}_f f_{i+1}(c_{i+1}, n_{i+1,1}, \dots) = \{ e_{i+1} \}; \dots; \mathbf{def}_f f_k \dots;$$
$$C_{opt}[ \mathbf{def}_c m_1(n_{1,1}, \dots) = \{ e_1^*[c_1 \mapsto c_0] \}; \mathbf{def}_c \dots; e^* ]$$

[ where  $f_1, \dots, f_i$  do not appear free in  $C_{opt}$  and  $e$  is minimal ]

Note: the term  $t^*$  is  $t$  with all applications of contified functions transformed to continuation applications.  $c_0$  is the continuation passed to the original call.

## Recursive Contification #2

When all non tail calls to functions in  $F = \{ f_1, \dots, f_i \}$  appear in the body of the function  $f_k$ , and  $\text{Cnt}(F)$  holds, contification is performed by the following rewriting:

```
deff f1(c1, n1,1, ...) = { e1 }; ...;  
deff fk(ck, nk,1, ...) = { Copt[ ek ] }; e
```

$\mapsto_{opt}$

```
deff fi+1(ci+1, ni+1,1, ...) = { ei+1 }; ...;  
deff fk(ck, nk,1, ...) = {  
  Copt[ defc m1(n1,1, ...) = { e1*[c1  $\mapsto$  c0] }; defc ...; ek* ]  
}; e
```

where  $f_1, \dots, f_i$  do not appear free in  $C_{opt}$  and  $e_k$  is minimal.

- A number of simple optimizations, and how the choice of IR can make them easier or harder to implement.
- How to specify optimizations as rewrite rules, and the notion of optimization contexts.

**Reminder: Project 4!**