

Function Inlining

15-411/15-611 Compiler Design

Seth Copen Goldstein

April 16, 2026

Today

- Why code layout matters
- Basic-block ordering and fall-through
- Static heuristics
- Hot/cold splitting
- Function placement and alignment
- Using Profile Data
- Example

What is Function Inlining

Before

```
int add1(int x) {  
    return x + 1;  
}
```

```
foo() {  
    ...  
    y = add1(a);  
}
```

After

```
foo() {  
    ...  
    t = a + 1;  
    y = t;  
}
```

Replace a call with specialized copy of callee body

- Map actual arguments to formal parameters
- Splice the callee CFG into the caller CFG
- Replace returns with jumps to caller continuation
- Merge returned values

Why Function Inlining

- Eliminate call and return overhead
- Exposes constants, copies, and control-flow context across the call boundary
- Enables CCP, copy propagation, dead code elimination, and CFG simplification
- Can remove dead parameters and dead branches
- Allows later optimizations to “see through” abstraction boundaries

What to Inline

- Not everything!
- Code size and I-cache pressure
- Compile times
- Possible spills in hot code
- Recursive or mutually recursive functions
- Opportunities for optimization

Things to Keep in Mind

- Recursion?
- Callee size
- Hot or not? (needs profiling data)
- Constant arguments?
- varargs? Exceptions?
- directives
- Debug info

When to Inline

- AST?
- Tree-IR?
- Abstract Assembly?
- Pre- or Post-optimization?

When to Inline

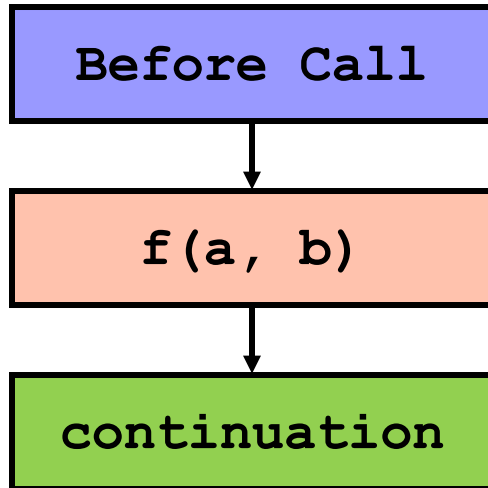
- After initial cleanup of AA
- Before any machine specific passes
- Inline
- Apply more optimizations
- Maybe inline again?

Basic Approach

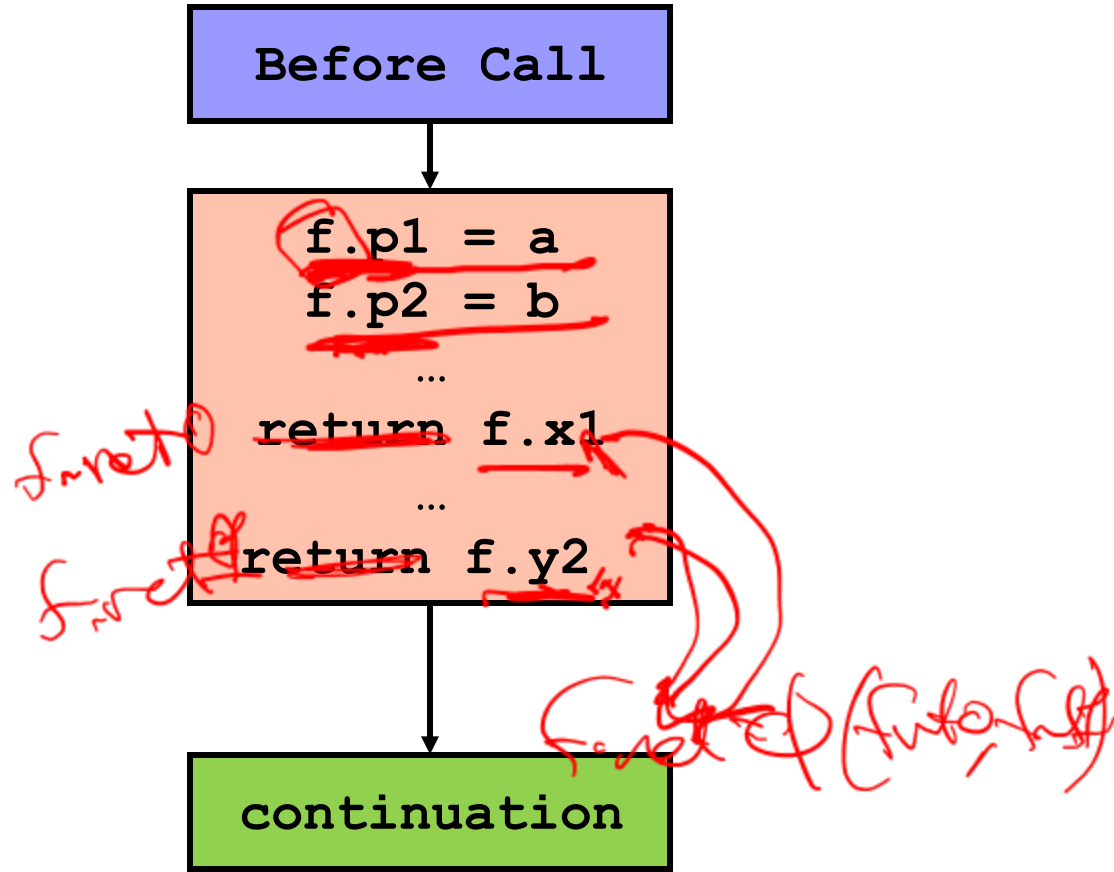
- Choose call site and callee
- Clone callee CFG into fresh caller blocks
- Map formals to actuals
- Create fresh SSA-named for callee clone
- Redirect all returns to caller continuation
- Merge return values (if needed)
- (if only 1 caller) delete callee
- Delete original call

The Layout Problem

Caller before



After Inlining



Maintaining SSA

- Formals to actuals
 - A map for each call site: callee names -> caller values
- Cloned variables/temps
 - Generate fresh names
 - Populate map of callee names -> in caller names
- Multiple return sites
 - Create phi at caller continuation to merge values

Two Approaches




- Clone then repair
- Maintain SSA as you clone
 - Rename as you clone
 - Explicitly insert phi to merge return values

Choosing

- Good Candidates:
 - Small callee
 - Hot call edge (in loop?)
 - Arguments are constants
 - One (or a few) call sites
 - Potentially dead-code/branch simplification
 - Callee is leaf function
- Bad Candidates:
 - Large callee
 - Cold call edge
 - Caller trace already hot and large
 - Potential spill increase

Example

- Call site: 1 constant arg
- Callee:
 - Called once
 - Potential dead code
 - Small
- Track args
- Track returns

```
int f(int p, int q) {  
    int t = p + q;  
    if (q < 0)  
        return t;   
    else  
        return p - q;   
}  
  
int g(int x) {  
    int y = f(x, 4);   
    return y * 2;  
}
```

SSA

```
int f(int p, int q) {
    int t = p + q;
    if (q < 0)
        return t;
    else
        return p - q;
}
```

```
int g(int x) {
    int y = f(x, 4);
    return y * 2;
}
```

```
int f(int p0, int q0) {
    t0 = p0 + q0
    if (q0 < 0) F1 else F2
F1:
    return t0;
F2:
    t1 = p0 - q0
    return t1
}

int g(int x0) {
    y0 = f(x0, 4)
    t0 = y0 * 2
    return t0
}
```

Build Parameter Map

Map:

p0 -> x0

q0 -> 4

```
int f(int p0, int q0) {  
    t0 = p0 + q0  
    if (q0 < 0) F1 else F2  
F1:  
    return t0;  
F2:  
    t1 = p0 - q0  
    return t1  
}  
  
int g(int x0) {  
    y0 = f(x0, 4)  
    t0 = y0 * 2  
    return t0  
}
```

Create Continuation

Map:

p0 -> x0

q0 -> 4

```
int f(int p0, int q0) {
    t0 = p0 + q0
    if (q0 < 0) F1 else F2
F1:
    return t0;
F2:
    t1 = p0 - q0
    return t1
}
```

```
int g(int x0) {
    clone(f(x0, 4))
f.cont:
    f.ret =  $\varphi?$ 
    y0 = f.ret
    t0 = y0 * 2
    return t0
}
```

Clone with fresh variables

Map:

p0 -> x0

q0 -> 4

t0 -> f.t0

t1 -> f.t1

```
int f(int p0, int q0) {  
    t0 = p0 + q0  
    if (q0 < 0) F1 else F2  
F1:  
    return t0;  
F2:  
    t1 = p0 - q0  
    return t1  
}
```

```
int g(int x0) {  
    clone(f(x0, 4))  
f.cont:  
    f.ret =  $\varphi$ ?  
    y0 = f.ret  
    t0 = y0 * 2  
    return t0  
}
```

```
f.t0 = x0 + 4  
if (4 < 0) f.F1 else f.F2
```

f.F1:

```
f.ret0 = t0  
goto f.cont
```

f.F2:

```
f.t1 = x0 - 4  
f.ret0 = f.t1  
goto f.cont
```

f.cont:

```
f.ret =  $\varphi$ (f.ret0, f.ret1)
```

Replace call

```
int g(int x0) {  
    f.t0 = x0 + 4  
    if (4 < 0) f.F1 else f.F2  
f.F1:  
    f.ret0 = t0  
    goto f.cont  
f.F2:  
    f.t1 = x0 - 4  
    f.ret0 = f.t1  
    goto f.cont  
  
f.cont:  
    f.ret =  $\phi$ (f.ret0, f.ret1)  
    y0 = f.ret  
    t0 = y0 * 2  
    return t0  
}
```

Post optimizations

```
int g(int x0) {  
  f.t0 = x0 + 4  
  if (4 < 0) f.F1 else f.F2  
f.F1:  
  f.ret0 = t0  
  goto f.cont  
f.F2:  
  f.t1 = x0 - 4  
  f.ret0 = f.t1  
  goto f.cont  
  
f.cont:  
  f.ret =  $\phi$ (f.ret0, f.ret1)  
  y0 = f.ret  
  t0 = y0 * 2  
  return t0  
}
```

```
int g(int x0) {  
  
f.F2:  
  f.t1 = x0 - 4  
  
  
  t0 = f.t1 * 2  
  return t0  
}
```

Summary

- At call site evaluate potential inlining
- Clone with fresh variables
- Then, perform optimizations