

EECS 483 Lecture 3

Let-bindings and simple stack allocations

Eric Giovannini

September 14, 2022

Based on notes available at:

https://maxsnew.com/teaching/eecs-483-fa22/lec_let-and-stack_notes.html

Recap

So far, our language was pretty simple:

`<expr>`: ***NUMBER***

...with abstract syntax

```
type expr = int64
```

... and the compiler simply generated a mov instruction to place the integer into RAX

Refactoring the Compiler

When given a number, say 483, we generate the following assembly:

```
section .text
global start_here
start_here:
    mov RAX, 483
    ret
```



Only this line corresponds to our input program! The others are scaffolding.

Refactoring (continued)

Split the compiler into multiple parts, as follows:

```
fn instrs_to_string(is: &[Instr]) -> String {  
    /* do something to get a string of assembly */  
}
```

```
/* compile_to_instrs is responsible for compiling just a single  
expression, and does not care about the surrounding scaffolding */
```

```
fn compile_to_instrs(e: &Exp) -> Vec<Instr> {  
    vec![ Instr::Mov(Reg::Rax, *e)]  
}
```

Refactoring (continued)

```
/* compile_to_string surrounds a compiled program by whatever
scaffolding is needed */
fn compile_to_string(e: &Exp) -> String {
    Ok(format!( "\
        section .text
        global start_here
        start_here:
        {}" ,
        instrs_to_string(&compile_to_instrs(e))))
}
```

Refactoring (continued)

Now `compile_to_string` will remain the same, while `compile_to_instrs` will grow to accommodate more expression forms.

Growing the language

Things to consider when we add a new feature:

1. Its impact on the *concrete syntax* of the language
2. Examples using the new enhancements, so we build intuition of them
3. Its impact on the *abstract syntax* and *semantics* of the language
4. Any new or changed *transformations* needed to process the new forms
5. Executable *tests* to confirm the enhancement works as intended

New feature: Adding and Subtracting 1

Let's add support for instructions that add and subtract 1 from their argument.

We will consider each of the five items from the previous slide in the context of these new instructions.

Concrete Syntax

$\langle expr \rangle$:

| *NUMBER*

| **add1** ($\langle expr \rangle$)

| **sub1** ($\langle expr \rangle$)

Examples

Concrete Syntax	Answer
42	42
add1 (42)	43
sub1 (42)	41
sub1 (add1 (add1 (42)))	43

Abstract Syntax

```
pub enum Exp {  
    Num(i64),  
    Add1(Box<Exp>),  
    Sub1(Box<Exp>),  
}
```

Semantics: evaluate argument to a number, then add or subtract one from it

Transformations

New assembly instruction:

```
add <dest>, <val>
```

Increment the destination by the right-side value

Transformations (continued)

New definition of Instr:

```
enum Instr {  
    ...  
    Add(Reg, i32) /* Increment the left-hand reg by the  
value of the right-hand immediate */  
    // In x86 only 32-bit literals can be on the right side  
of an add instruction  
}
```

Example: compiling `add1(42)`

Two steps:

1. Load 42 into RAX
2. Add 1 to RAX

Resulting assembly:

```
mov RAX, 42
```

```
add RAX, 1
```

Another example

Compile `sub1(add1(add1(42)))`

How to handle subtraction? Just add -1.

```
mov RAX, 42
add RAX, 1
add RAX, 1
add RAX, -1
```

Notice that each piece of the program corresponds to a related piece of the assembly!

Important Observation: Compositionality

Our translations are **compositional**: a translation of a composite expression is just a function of the translations of its constituent parts!

This makes writing the compilation function easy; we can use recursion

Testing the Feature

After implementing the code for the feature, we should now test that it works as expected: we should get the handwritten assembly we expect.

Adding let

Growing the language

Things to consider when we add a new feature:

1. Its impact on the *concrete syntax* of the language
2. Examples using the new enhancements, so we build intuition of them
3. Its impact on the *abstract syntax* and *semantics* of the language
4. Any new or changed *transformations* needed to process the new forms
5. Executable *tests* to confirm the enhancement works as intended

Concrete Syntax for Let

$\langle \text{expr} \rangle$: . . .

| *IDENTIFIER*

| **let** *IDENTIFIER* **=** $\langle \text{expr} \rangle$ **in** $\langle \text{expr} \rangle$

Abstract Syntax for Let

```
enum Exp {  
    ...  
    Id(String),  
    Let(String, Box<Exp>, Box<Exp>)  
}
```

Examples

```
let x = 5 in add1(x)
```

=> 6

```
let x = 483 in (let y = add1(x) in add1(y))
```

=> 485

```
let x = (let y = add1(5) in add1(y)) in add1(x)
```

=> 8

Semantics: Writing an Interpreter for the New Language

Same as before:

- Numbers evaluate to themselves
- Adding or subtracting one should evaluate the expression and then add/subtract one from the result

But what about identifiers and let-bindings?

Environments

We need to track the meaning of each identifier. We will do so using an **environment**.

Possible choices for the type of the environment:

- Match each identifier to the expression it was bound to
 - Environment type: $[(\&\text{str}, \text{Exp})]$
 - *Lazy* behavior
- Match each identifier to the result of evaluating that expression
 - Environment type: $[(\&\text{str}, \text{i64})]$
 - *Eager* behavior

Lazy vs Eager Evaluation

In **lazy** evaluation, an identifier is evaluated to a result on an as-needed basis.

In **eager** evaluation, an expression is fully evaluated before it is bounded to an identifier, and is subsequently never evaluated again.

In this language, the choice of eager versus lazy makes no difference in terms of the result or the performance.

But in more complicated languages, the choice can make a difference in terms of performance or even the result!

Scope

Scope tells us which names are available for use within a given expression.

Our convention for scope: the program `let x = e1 in e2` means that `x` can be used in `e2`, but not in `e1`.

Is this code valid?

```
let x = add1(x) in x
```

No, because `x` is not in scope in `add1(x)`!

(If the language supported recursion, this kind of definition could be sensible, but even if it did, in this particular example, there would be no solution, i.e., no `x` such that `x = x + 1`.)

Important Convention

What is the result of the following code:

```
let x = 1 in let x = 2 in x
```

Choices: 1, 2, or error

Our convention: answer = 2

“Inner bindings shadow outer ones.”

Interpreter Demo

(Look at example Rust code)

The Stack

Compiling the New Language

How can we compile programs in our updated language?

- No notion of identifier names or environments in the assembly language
- One register is not enough, since we may need to track multiple names at once. (In fact, no fixed number of registers would be enough.)

Solution

Insight #1: Broaden our notion of a name

In the interpreter, a name was used to map to a value or expression.

In reality, any unique identifier will suffice, and all values will need to exist in memory at runtime.

So now, instead of “a name is a string”, we should think “a name is a memory address”.

Insight #2

While compiling, we can maintain an environment of type `Vec<&str, Address>`.

- When we compile a **let-binding**, we can extend this environment with new addresses for new identifiers.
- When we compile an **identifier**, we look up the relevant address.

This environment is not needed at runtime!

New question: how do we assign addresses to identifiers?


Memory Layout

Conceptually, memory is an array of bytes, addressed from 0 to 2^{64} (assuming a 64-bit machine).

There are restrictions on what addresses can be used.

The typical memory layout for a program is shown on the next slide.

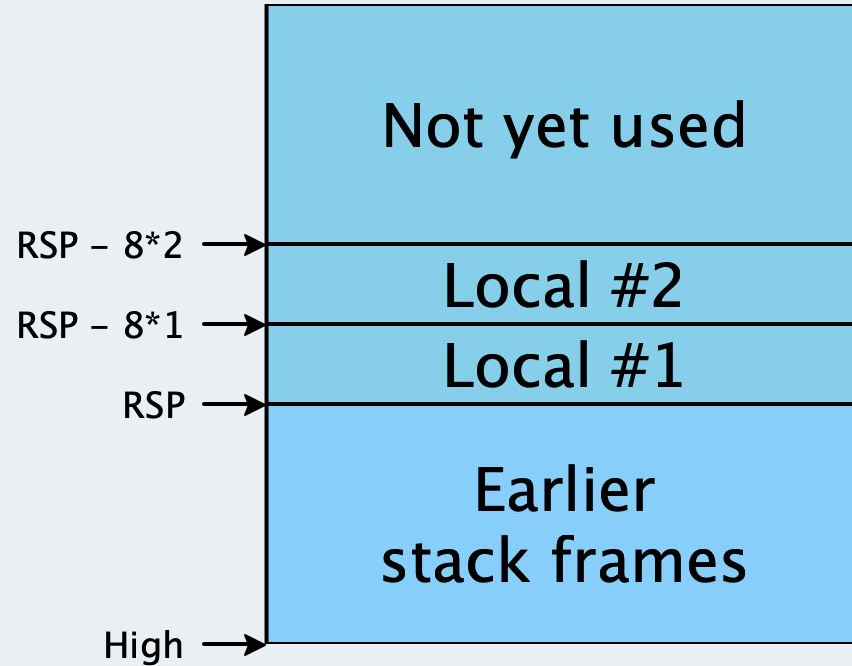
Memory Layout (continued)



Sections of Program Memory

- Code/text segment: includes the program machine code
- Global segment: global data available throughout the program's execution
- Heap: memory that is dynamically allocated as the program runs
- **Stack: used as the program calls and returns from functions**

Stack Layout



Stack Layout (continued)

- The stack is divided into stack frames, with each function in progress getting its own frame.
- Each stack frame can be used freely by its corresponding function.
- When the function returns, its stack frame is freed for use by future calls.
- The RSP register contains the address where the current stack frame begins.

Allocating Identifiers on the Stack

With the above knowledge, the task of assigning addresses is more concrete.

We are given addresses on the stack at $RSP - 8 * 1$, $RSP - 8 * 2$, ... $RSP - 8 * i$

We need to allocate a number to each identifier so that identifiers needed simultaneously are mapped to different numbers.

Naive Allocation Algorithm

Give every unique binding its own unique integer, i.e., every binder gets its own stack slot.

Implementation: keep a global mutable counter of the number of variables we have seen, and a global table mapping names to counters.

Naive Allocation: Examples

```
let x = 10      /* [] */  
in add1(x)     /* [ x --> 1 ] */
```

```
let x = 10      /* [] */  
in let y = add1(x) /* [x --> 1] */  
in let z = add1(y) /* [y --> 2, x --> 1] */  
in add1(z)     /* [z --> 3, y --> 2, x --> 1] */
```


Naive Allocation: Examples

```
    let a = 10                                /* [] */
in let c =    let b = add1(a)                 /* [a --> 1] */
           in let d = add1(b)                 /* [b --> 2, a --> 1] */
           in add1(b)                         /* [d --> 3, b --> 2, a --> 1] */
in add1(c)                                    /* [c --> 4, d --> 3, b --> 2, a --> 1] */
```

Problems with this Approach

Wastes space (see last line of last example where neither b nor d are in scope, but their stack slots are still reserved).

Difficult to test because of the use of **mutable state!**

Another Attempt

Observation: as we enter the bodies of let-expressions, only the bindings of those particular let-expressions are in scope; everything else is unavailable.

We can trace a straight-line path from any given let-body out through its parents to the outermost expression of a given program.

So, we only need to maintain uniqueness among the variables on those paths!

Resulting Assembly Code

```
let a = 10
```

```
in let c = let b = add1(a)
```

```
    in let d = add1(b)
```

```
        in add1(b)
```

```
in add1(c)
```

```
mov RAX, 10
```

```
mov [RSP - 8*1], RAX
```

```
mov RAX, [RSP - 8*1]
```

```
add RAX, 1
```

```
mov [RSP - 8*2], RAX
```

```
mov RAX, [RSP - 8*2]
```

```
add RAX, 1
```

```
mov [RSP - 8*3], RAX
```

```
mov RAX, [RSP - 8*2]
```

```
add RAX, 1
```

```
mov [RSP - 8*2], RAX
```

```
mov RAX, [RSP - 8*2]
```

```
add RAX, 1
```

Allocating Identifiers on the Stack

Extending our Transformations

```
enum Reg {
    Rax,
    Rsp,
}

// Represents the address [ reg + 8 * offset]
struct MemRef {
    reg: Reg,
    offset: i32,
}

enum Arg64 {
    Reg(Reg),
    Imm(i64),
    Mem(MemRef)
}
```

Extending our Transformations

```
enum MovArgs {  
    ToReg(Reg, Arg64),  
    ToMem(MemRef, Reg32),  
}
```

```
enum Instr {  
    Mov(MovArgs),  
    Add(Reg, i32),  
}
```


Looking up an Identifier in an Environment

```
fn get(env: &[(String, i32)], x: &str) -> Option<i32> {
    for (y,n) in env.iter().rev() {
        if &x == y {
            return Some(*n);
        }
    }
    None
}
```

Notice that we search in reverse, because we push pairs on as we descend into the expression.

Compiling Let-bindings

```
fn compile(e: &Exp mut env: Vec<(String, i32)>) -> Result<Vec<Instr>, String>
{
    match e {
        Let(x, e1, e2) => {
            let mut is = compile(e1, env.clone())?;
            let offset = ... // Calculate the offset from env
            env.push((String::from(x), offset));
            is.push(Instr::Mov(MovArgs::ToMem(MemRef { reg: Reg::Rsp, offset:
offset }, Reg::Rax)));
            is.extend(compile(e2, env))?;
            Ok(is)
        }
    }
}
```

Thank you!