

Announcements

- Guest Lecture on the following dates: Nov 09 (bodo.ai) Nov 11 (Jane Street) Nov 30 (Google: 04:30 pm). Attendance: **MANDATORY**.

X86_64 Operand Sizes

In previous assignments, we have only worked with 32-bit types. Now, you will need to modify your compiler to account for 64-bit types. Most of the x86-64 instructions specify the sizes of their operands, so you will want to use the `__Q` (rather than `__L`) variants.

The assembler will gladly let you use 32-bit instructions with your 64-bit data, but be careful: most 32-bit instructions will **zero the upper 32 bits of your registers**. Additionally, you need to be careful when reading or writing to memory. You can either allocate 8 bytes for every temp you spill, or worry about a non-linear mapping from temp number to memory offset.

The `==` and `!=` operations are particularly tricky in Lab 4 because they can be used with integers, Boolean values, and pointers. You'll notice that because pointers are 64 bits wide, you'll need to take the operand sizes into account when generating code for these operations.

Structs and Alignment: a 213 Review

Structs store a collection of fields. In Lab 4, our fields will either be primitive types, arrays, or other structs. Each primitive field of a struct must be aligned by the size of its type, and each struct/array field must be aligned to the size of its largest field. These rules can introduce padding if a larger field follows a smaller one. The size of the overall struct needs to be aligned with its largest field.

Most of the time, L4 programs will reference fields via the arrow notation (`->`), as we cannot place structs on the stack or manipulate a struct directly. That said, we do need the dot notation (`.`) for arrays of structs and expressions of the form `(*S).f`.

Checkpoint 0

Compute the byte offsets for each of the fields in the two struct definitions below.

```
1 struct A {
2     int* b;
3     int c;
4     struct A* d;
5 };
6
7 struct B {
8     int i;
9     struct A a;
10    int k;
11 }
```

Solution:

```
struct A (size: 24):
  b: 0    c: 8    d: 16
struct B (size: 40):
  i: 0    a: 8    k: 32
```

Checkpoint 1

Compute the byte offsets for each struct access in the following program. What information do we need about the structs for typechecking? For code generation?

```
12 int foo(struct B* x) {
13     x->a.d = alloc(struct A);
14     return x->a.d->c;
15 }
16
17 int main() {
18     struct B* b = alloc(struct B);
19     b->a.c += 15411;
20     return foo(b);
21 }
```

Solution:

```
[1.13 lhs] &(x->a.d)    x+8+16
[1.14]     x->a.d->c     (*(x+8+16)+8)
[1.19 rhs] b->a.c       *(b+8+8)
[1.19 lhs] &(b->a.c)    b+8+8
```

sizeof in 14

During code generation, you will need to know the sizes of the types your program uses. For pointers, this is pretty easy; for arrays, this is only slightly harder. Figuring out offsets for structs, however, can be non-trivial. You will want to write a recursive function for computing the sizes of types, since structs may be nested. With these two parts, you'll be able to compute struct offsets needed in code generation.

```
1 fun sizeof BOOL      => 4
2 | sizeof INT         => 4
3 | sizeof PTR t       => 8
4 | sizeof STRUCT id  => ???
```

Addressing Schemes

There are many ways to describe memory locations in x86 assembly:

Form	Address
(%rsp)	%rsp
a(%rsp)	%rsp + a
(%rsp, %rax)	%rsp + %rax
a(%rsp, %rax)	%rsp + %rax + a
a(%rsp,%rax,s)	%rsp + %rax*s + a

In the table above, %rsp and %rax represent registers, *a* is an arbitrary constant, and *s* can be 1, 2, 4, or 8. The last form is useful for array accesses, and the form a(%rsp) is useful for accessing fields of a struct. There are other forms, but they probably aren't necessary to know to implement your compiler.

Also, recall that when you use one of the above forms, the mov instruction dereferences the memory location while the lea instruction loads the address of the memory location. Both of these instructions will be useful in Lab 4.

Code Generation for Structs and Arrays

Once we have some sort of environment that maps struct fields to offsets, we can go about with actual code generation. From our example, alloc is straight forward, so let's focus on line 19.

We do not want to elaborate $b \rightarrow a.c += 15411$ to $b \rightarrow a.c = b \rightarrow a.c + 15411$, as in general, the left hand side may have side effects we cannot repeat. Instead, we must:

- (a) Compute the *address of* $b \rightarrow a.c$.
- (b) Compute the right side of the assignment (in this case, just 15411).
- (c) Check that the address is not NULL.
- (d) Dereference the address and add the right side to it.
- (e) Set the memory at the address to the result.

Checkpoint 2

Assuming b is in the register t_0 , write assembly that executes the statement $b \rightarrow a.c += 15411$.

Solution: (this is admittedly a weird amalgamation of abstract and concrete assembly, but I think it is useful to see it this way)

```
    lea 8(t0), t1 // &(b->a)
    lea 8(t1), t1 // &(b->a.c) [could also combine these two offsets]
    t2 <- 15411
    if t1 == 0 then goto L1 else L2
L1:
    raise(SIGUSR2)
L2:
    mov (t1), t3
    t3 <- t3 + t2
    mov t3, (t1)
```