

Data on the Heap

num
bool
x char
x double

Next, lets add support for

- Data Structures

In the process of doing so, we will learn about

- Heap Allocation
- Run-time Tags
- High-order Func (Closures)

$\langle \text{env}, \text{code} \rangle$

Creating Heap Data Structures

We have already support for two primitive data types

data Ty

```
= TNumber -- e.g. 0,1,2,3,...  
| TBoolean -- e.g. true, false
```

we could add several more of course, e.g.

- Char
- Double or Float

etc. (you should do it!)

However, for all of those, the same principle applies, more or less

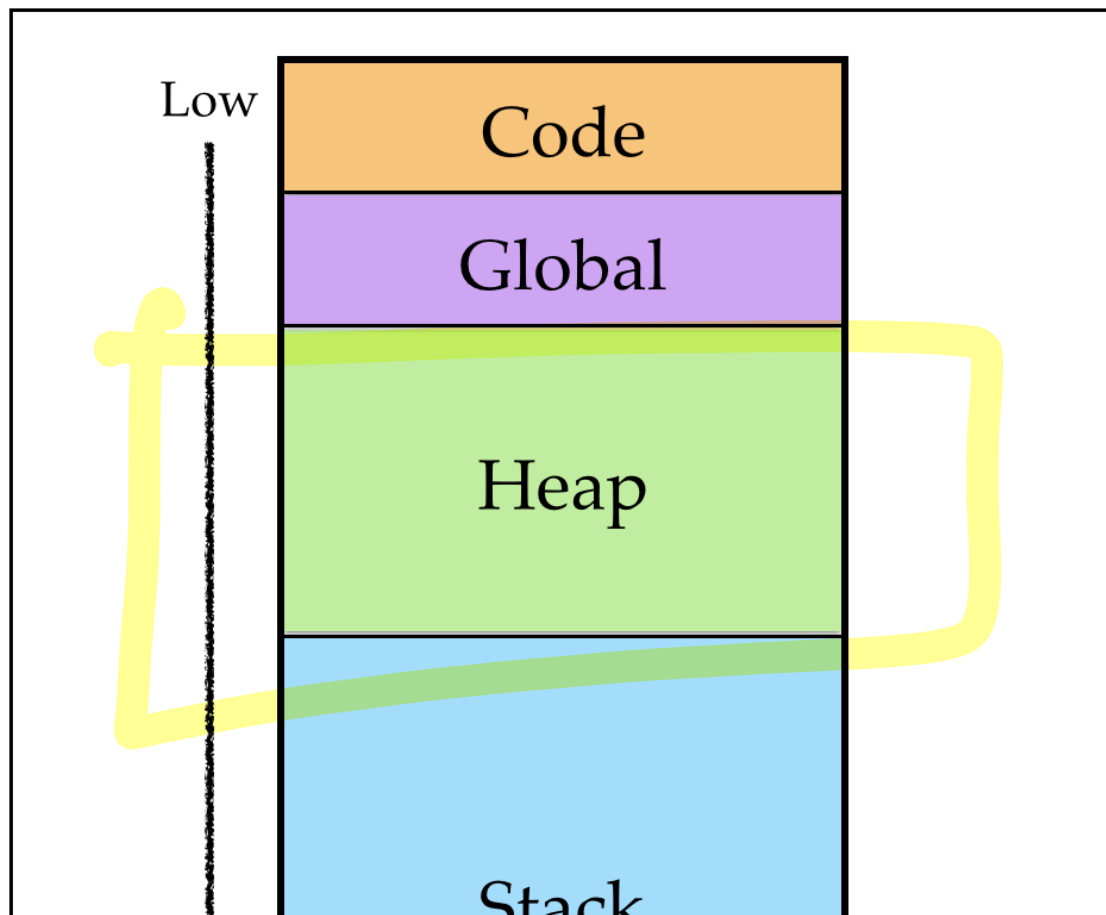
- As long as the data fits into a single word (8-bytes)

Instead, lets learn how to make **unbounded data structures**

- Lists
- Trees
- ...

which require us to put data on the **heap**

not just the stack that we've used so far.





Stack vs. Heap

Pairs

While our *goal* is to get to lists and trees, the journey of a thousand miles begins with

a single step...

So! we will *begin* with the humble pair.

Construct

(2, 3)

length=2

(1, (2, 3))

$e[0]$

$e[1]$

Assignment

$e[0]$

(1, 2, 3, 7, (9, 12))

"length" = ?
Storage

$e[i]$

Pairs: Semantics (Behavior)

First, let's ponder what exactly we're trying to achieve.

We want to enrich our language with *two* new constructs:

- **Constructing** pairs, with a new expression of the form (e_0, e_1) where e_0 and e_1 are expressions.
- **Accessing** pairs, with new expressions of the form $e[0]$ and $e[1]$ which

$e[0]$ $e[1]$

evaluate to the first and second element of the tuple e respectively.

For example,

```
let t = (2, 3) in  
  t[0] + t[1]
```

should evaluate to 5.

Strategy

Next, let's informally develop a strategy for extending our language with pairs, implementing the above semantics. We need to work out strategies for:

1. **Representing** pairs in the machine's memory,

$(e_0, e_1) \longrightarrow \langle \text{asm} \rangle$

$e[0]$ \longrightarrow $\langle \text{asm} \rangle$

2. **Constructing** pairs (i.e. implementing (e_0, e_1) in assembly),
3. **Accessing** pairs (i.e. implementing $e[0]$ and $e[1]$ in assembly).

1. Representation

Recall that we represent all values: ([05-cobra.md/#option-2-use-a-tag-bit](https://ucsd-cse131.github.io/sp21/lectures/05-cobra.md/#option-2-use-a-tag-bit))

- Number like 0, 1, 2 ...
- Boolean like true, false

} 64

as a ~~single~~ **word** either

- 8 bytes on the stack, or
- a single register rax, rbx etc.

EXERCISE


What kinds of problems do you think might arise if we represent a pair $(2, 3)$ on the *stack* as:

```
|   |
| 3 |
| 2 |
| ... |
```

Let $t = (\underline{2}, \underline{3}), 4$ in

⋮

$t = (\underline{2}, 3), 4$
 $t = 2, (3, 4)$



$t[0][0]$

1, 2, 3, 4, 5

$\text{cons}(1, \text{cons}(2, \text{cons}(3, \text{cons}(4, \text{nil}))))$

(1, 2) 3

↓ ↓
(e_0, e_1)

$e[0], e[1]$

(1, (2, (3, (4, false))))

QUIZ

How many words would we need to store the tuple

(3, (4, 5))


↑ ↑
1. 1 word →

2. 2 words

3. 3 words

4. 4 words

5. 5 words



tup

def tail(l):
l[1]

def isNil():
l == False

def nil():
false

{ def cons(h,t):
(h,t)

def range(lo, hi):
if lo < hi:

cons(lo, range(lo+1, hi))
else:
nil()

build list

def length(l):

if isNil(l):

0

else:

1 + length(tail(l))


$l = \text{range}(0, 100)$
length(l)




Pointers

Every problem in computing can be solved by adding a level of indirection.

We will **represent a pair by a pointer** to a block of **two adjacent words** of memory.





Pairs on the heap

The above shows how the pair $(2, (3, (4, 5)))$ and its sub-pairs can be stored in the **heap** using pointers.

(4, 5) is stored by adjacent words storing

- 4 and
- 5

(3, (4, 5)) is stored by adjacent words storing

- 3 and
- a **pointer** to a heap location storing (4, 5)

(2, (3, (4, 5))) is stored by adjacent words storing

- 2 and
- a **pointer** to a heap location storing (3, (4, 5)).

A Problem: Numbers vs. Pointers?

How will we tell the difference between *numbers* and *pointers*?

That is, how can we tell the difference between

1. the *number* 5 and
2. a *pointer* to a block of memory (with address 5)?

Each of the above corresponds to a *different* tuple

1. (4, 5) or
2. (4, (...)).

so its pretty crucial that we have a way of knowing *which* value it is.

$$t = (1, (2, 3))$$


$$t = 3$$

$$t = \text{false}$$

Tagging Pointers

As you might have guessed, we can extend our tagging mechanism (05-cobra.md/#option-2-use-a-tag-bit) to account for *pointers*.

Type	LSB
number	xx0
boolean	111
pointer	001



$$t = (1, (2, 3))$$

Pointers are
8-byte aligned

That is, for

- number the **last bit** will be 0 (as before),
- boolean the **last 3 bits** will be 111 (as before), and
- pointer the **last 3 bits** will be 001.

(We have 3-bits worth for tags, so have wiggle room for other primitive types.)

Address Alignment

As we have a **3 bit tag**

- leaving **$64 - 3 = 61$ bits** for the actual address

So actual addresses, written in binary, omitting trailing zeros, are of the form

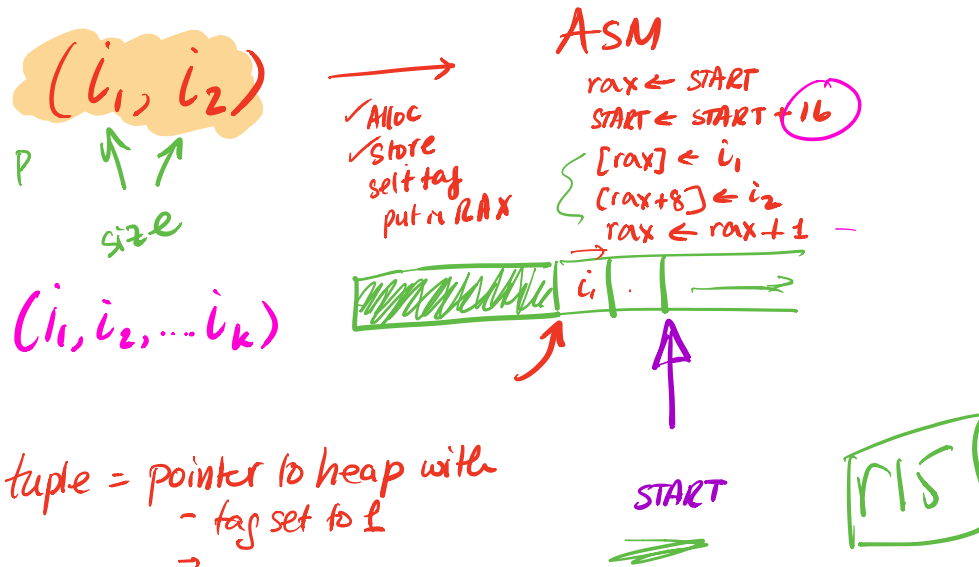
Binary	Decimal
0b00000000	0
0b00001000	8
0b00010000	16
0b00011000	24
0b00100000	32
...	

That is, the addresses are **8-byte aligned**.

Which is great because at each address, we have a pair, i.e. a **2-word = 16-byte block**, so the **next allocated address will also fall on an 8-byte boundary**.



- But ... what if we had 3-tuples? or 5-tuples? ...



tuple = pointer to heap with
 - tag set to i
 →

2. Construction

Next, let's look at how to implement pair **construction** that is, generate the assembly for expressions like:

(e_1, e_2)

To construct a pair (e_1, e_2) we

1. **Allocate** a new 2-word block, and getting the starting address at rax ,
2. **Copy** the value of e_1 (resp. e_2) into $[rax]$ (resp. $[rax + 8]$).

3. **Tag** the last bit of `rax` with `1`.

The resulting `eax` is the **value of the pair**

- The *last step* ensures that the value carries the proper tag.

ANF will ensure that `e1` and `e2` are immediate expressions (04-boa.md/#idea-immediate-expressions)

- will make the second step above straightforward.

EXERCISE How will we do ANF conversion for $(e1, e2)$?

Allocating Addresses

Lets use a **global** register `r15` to maintain the address of the **next free block** on the heap.

Every time we need a *new* block, we will:

1. **Copy** the current `r15` into `rax`

- Set the last bit to `1` to ensure proper tagging.
- `rax` will be used to fill in the values

2. **Increment** the value of `r15` by `16`

- Thus *allocating* `8` bytes (= `2` words) at the address in `rax`

Note that addresses stay 8-byte aligned (last 3 bits = 0) if we

- *Start* our blocks at an 8-byte boundary, and
- *Allocate* 16 bytes at a time,

NOTE: Your assignment will have *blocks of varying sizes*

- You will have to *maintain* the 8-byte alignment by *padding*

Example: Allocation

In the figure below, we have

- a source program on the left,
- the ANF equivalent next to it.

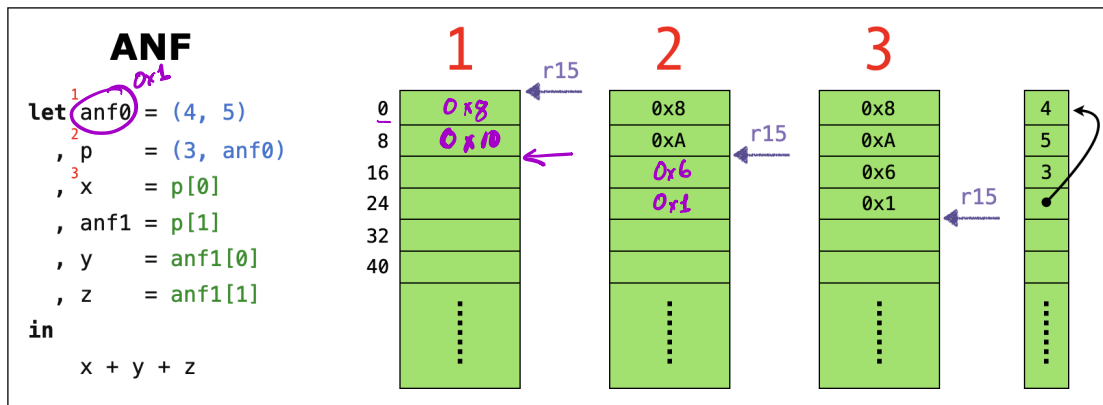


Source	ANF
<pre> let p = (3, (4, 5)) , x = p[0] , y = p[1][0] , z = p[1][1] in </pre>	<pre> 0 let anf0 = (4, 5) 1 , p = (3, anf0) 2 , x = p[0] , anf1 = p[1] y = anf1[0] </pre>



Example of Pairs

The figure below shows the how the heap and `r15` evolve at points 1, 2 and 3:




Allocating Pairs on the Heap

QUIZ

In the ANF version, `p` is the *second (local) variable* stored in the stack frame. What *value* gets moved into the *second stack slot* when evaluating the above program?

1. 0x3 \rightarrow 3
2. (3, (4, 5)) \rightarrow
3. 0x11 \rightarrow 17
4. 0x9 \rightarrow 9
5. 0x10 \rightarrow 16



3. Accessing

Finally, to **access** the elements of a pair

Lets compile $e[0]$ to get the first or $e[1]$ to get the second element

1. **Check** that immediate value e is a pointer
2. **Load** e into rbx
3. **Remove** the tag bit from rbx

4. Copy the value in `[rbx]` (resp. `[rbx + 8]`) into `rbx`.

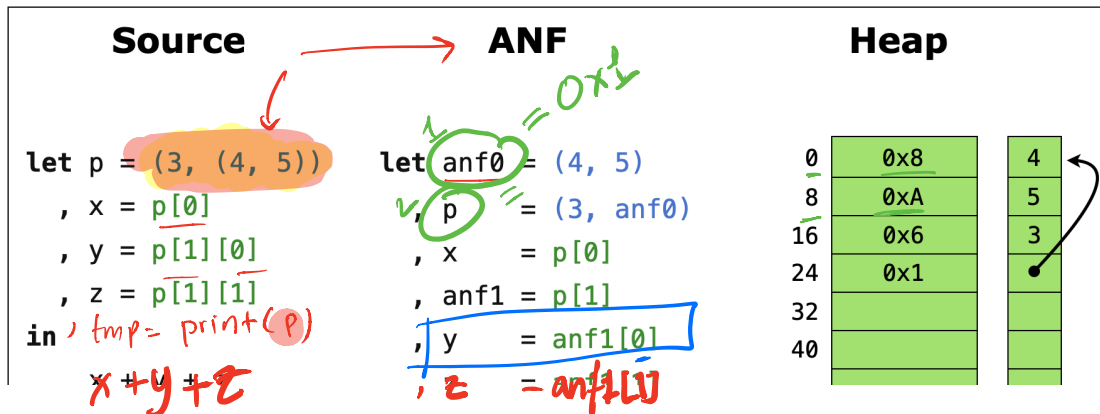
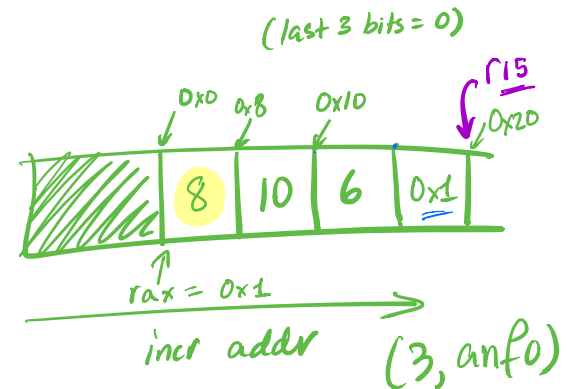
```

mov rax, (anf1)
-sub rax, 1
mov rax, [rax + 8.1]

```

^a
Example: Access


Here is a snapshot of the heap after the pair(s) are allocated.



Heap

0	0x8	4
8	0xA	5
16	0x6	3
24	0x1	
32		
40		

y	0x8
anf1	0x1
x	0x6
p	0x11
anf0	0x1



Allocating Pairs on the Heap

Lets work out how the values corresponding to x , y and z in the example above get stored on the stack frame in the course of evaluation.


Variable	Hex Value	Value
anf0	0x001	ptr 0
p	0x011	ptr 16
x	0x006	num 3
anf1	0x001	ptr 0
y	0x008	num 4
z	0x00A	num 5
anf2	0x00E	num 7
result	0x018	num 12

Plan

(1, (2, (3, 4)))

Pretty pictures are well and good, time to build stuff!

As usual, lets continue with our recipe:

1. Run-time 
2. Types
3. Transforms

We've already built up intuition of the *strategy* for implementing tuples. Next, lets look at how to implement each of the above.

Run-Time

We need to extend the run-time (`c-bits/main.c`) in two ways.

1. **Allocate** a chunk of space on the heap and pass in start address to `our_code` .
2. **Print** pairs properly.

Allocation

The first step is quite easy we can use `calloc` as follows:

```
int main(int argc, char** argv) {
    int* HEAP = calloc(HEAP_SIZE, sizeof (int));
    long result = our_code_starts_here(HEAP);
    print(result);
    return 0;
}
```

*Where does 'HEAP' live?
in our code....*

The above code, (A) `r15` (B) `rdi`

1. **Allocates** a big block of contiguous memory (starting at `HEAP`), and
2. **Passes** this address in to `our_code`.

Now, `our_code` needs to, at the beginning start with instructions that

- copy the parameter (in `rdi`) into global pointer (`r15`) *→ 8-byte?*
- and then bump it up at each allocation.

alloc
print ✓



Printing

To print pairs, we must **recursively traverse** pointers

tuple = 001
bool = 111

- until we hit number or boolean.

We can check if a value is a pair by looking at its last 3 bits:

```
int isPair(int p) {
    return (p & 0x00000007) == 0x00000001;
}
```

We can use the above test to recursively print (word)-values:

```
void print(long val) {
    if(val & 0x1 == 0) { // val is a number
        printf("%ld", val >> 1);
    }
    else if(val == CONST_TRUE) {           // val is true
        printf("true");
    }
    else if(val == CONST_FALSE) {         // val is false
        printf("false");
    }
    else if(val & 7 == 1) {
        long* valp = (long *) (val - 1); // extract address
        printf("(");
        print(*valp);                       // print first element
        printf(", ");
        print(*(valp + 1));                 // print second element
        printf(")");
    }
    else {
        printf("Unknown value: %#010x", val);
    }
}
```

$$e_1 [e_2]$$

Types

Next, lets move into our compiler, and see how the **core types** need to be extended.

Source $(e_1, e_2) \longrightarrow \text{Pair } e_1 \ e_2$

We need to extend the source Expr with support for tuples

```

data Expr a  $\longrightarrow e[0]$   $\longrightarrow \text{GetItem } e \ \text{First}$ 
  = ...
x | Pair (Expr a) (Expr a) a -- ^ construct a pair
x | GetItem (Expr a) Field a -- ^ access a pair's element

```

In the above, Field is

$e[1] \longrightarrow \text{GetItem } e \ \text{Second}$

$\text{Tuple } [\text{Expr } a] \ a$

| GetItem (Expr a) Int → "static"

data Field

= First -- ^ access first element of pair

| Second -- ^ access second element of pair

(Expr a) → "dynamic"

NOTE: Your assignment will generalize pairs to **n-ary tuples** using

- Tuple [Expr a] representing (e_1, \dots, e_n)
- GetItem (Expr a) (Expr a) representing $e_1[e_2]$

Dynamic Types

Let us extend our **dynamic types** Ty see (05-cobra.md/#types) to include pairs:

```
data Ty = TNumber | TBoolean | TPair
```

↓ ↓ ↓
0x0 0x111 0x001

Assembly

The assembly Instruction are changed minimally; we just need access to r15 which will hold the value of the *next* available memory block:

```
data Register
```

```
= ...
```

```
| R15
```

Transforms

Our code must take care of three things:

1. **initialize** $r15$ to allow heap allocation,
2. **Construct** pairs, *compileEnv* *Pair/Tuple*
3. **Access** pairs. *compileEnv* *GetItem*

The latter two will be pointed out as cases in `anf` and `compileEnv`

- Tuple
- GetItem

Pair e_1 e_2

*ANF = like any
Prim2*

GetItem e₁ e₂

Initialize

We need to **initialize** `r15` with the **start position** of the heap

- passed in as `rdi` by the run-time.

How shall we get a hold of this position?

To do so, `our_code` starts off with a prelude

```
prelude :: [Instruction]
prelude =
  [ IMov (Reg R15) (Reg RDI)
  ]
```

-- copy param (HEAP) off rdi

Is that it?

mov r15, rdi


- ① Find the gap
add the gap as pad
- ② zero out last bits

QUIZ

Is `r15` 8-byte aligned?

A. Yes

B. No *x*



Ensuring alignment

64



```
prelude :: [Instruction]
```

```
prelude =
```

```
[ IMov (Reg RAX) (HexConst 0xFFFFFFFF)    -- setup regMask
  , IShl (Reg RAX) (Const 32)
  , IOr  (Reg RAX) (HexConst 0xFFFFFFFF8)
  , IMov (Reg R15) (Reg RDI)                -- copy param (HEAP) of
f rdi
  , IAdd (Reg R15) (Const 8)                -- add 8 and mask 3 bit
s to ensure
  , IAnd (Reg R15) (Reg RAX)                -- 8-byte aligned
]
```

← rax :=

1. **Copy** the value off the (parameter) stack, and
2. **Adjust** the value to ensure the value is 8-byte aligned.

```
mov r15, rdi
```

```
add r15, 8
```


```
and r15, 0x111...000
```

61-bits

QUIZ

Why add 8 to r15? What would happen if we removed that operation?

- A. r15 would not be 8-byte aligned?
- B. r15 would point into the stack?
- C. r15 would not point into the heap?
- D. r15 would not have enough space to write 2 bytes?



Construct

To construct a pair (v1, v2) we directly implement the above strategy (07-egg-eater.md/#2-construction):

```

compileEnv env (Tuple v1 v2)
  = pairAlloc
  ++ pairCopy First
  ++ pairCopy Second
  ++ setTag RAX TPair
  
```

(Tuple v1 v2)
tupleAlloc (n+1) -- 1. allocate pair, resulting a ddr in `rax`
 -- 2. copy first value into slot
 -- 3. copy second value into slot
 -- 3. set the tag-bits of `rax`

Lets look at each step in turn.

add RAX, 1

pairAlloc =
mov rax, r15
add r15, 16

tupleAlloc k =
mov rax, r15
*add r15, 8*k*

pairCopy fld arg =
mov rbx, arg
mov [rax+off], rbx *[rax+off]*

where
off = field off fld *[rbp ± -*

Allocate

r15

To allocate, we just copy the current pointer `r15` and increment by 16 bytes,

- accounting for two 8-byte blocks for each element.

```
pairAlloc :: Asm
pairAlloc
  = [ IMov (Reg RAX) (Reg R15)    -- copy current "free address" `esi`
      `into` `eax`
      , IAdd (Reg RAX) (Const 16) -- increment `esi` by 8
      ]
```

Exercise How would you make this work for n -tuples?

Copy

We copy an Arg into a Field by

- saving the Arg into a helper register `rbx`,
- copying `rbx` into the field's slot on the heap.

```
pairCopy :: Field -> Arg -> Asm
pairCopy fld arg
  = [ IMov (Reg RBX) ← arg
    → IMov (pairAddr fld) (Reg RBX)
      ]
```

Recall, the field's slot is either `[rax]` or `[rax + 8]` depending on whether the field is First or Second.

QUIZ

What shall we fill in for `_1` and `_2`?

```
pairAddr :: Field -> Arg
pairAddr First = RegOffset 1 RAX
pairAddr Second = RegOffset -1 RAX
```

~~A. 0 and 1~~

✓ B. 0 and -1

$v_0 \quad v_1 \quad v_2 \quad v_3$
 \downarrow
 0 -1 -2 -3

RegOff 3 RBP \rightarrow $[RBP - 3*i]$

C. 1 and 2

D. -1 and -2

E. huh?

Tag

Finally, we set the tag bits of `rax` by using `typeTag TPair` which is defined

```
setTag :: Register -> Asm
setTag r = [ IAdd (Reg r) (HexConst 0x1) ]
```


$e[0]$ $e[1]$

Access

To access tuples, lets update `compileEnv` with the strategy above:

```

compileExpr env (GetItem e fld)
  = assertType env e TPair
  ++ [ IMov (Reg RAX) (immArg env e) ]
  ++ unsetTag RAX
  ++ [ IMov (Reg RAX) (pairAddr fld) ]
  lot to eax

```

-- 1. check that *e* is a (pair) pointer

-- 2. load pointer into *eax*

-- 3. remove tag bit to get a ddress

-- 4. copy value from resp. slot to *eax*

we remove the tag bits by doing the opposite of `setTag` namely:

```

unsetTag :: Register -> Asm
unsetTag r = ISub (Reg RAX) (HexConst 0x1)

```

N-ary Tuples

That's it! Let's take our compiler out for a spin, by using it to write some interesting programs!

First, let's see how to generalize pairs to allow for

- triples (e1, e2, e3)
- quadruples (e1, e2, e3, e4)
- pentuples (e1, e2, e3, e4, e5)

$(e_1, (e_2, e_3))$

$(e_1, (e_2, (e_3, (e_4, -))))$


and so on.

We just need a library of functions in our new egg language to


- **Construct** such tuples, and
- **Access** their fields.

Constructing Tuples


We can write a small set of functions to **construct** tuples (up to some given size):



```
def tup3(x1, x2, x3):  
    (x1, (x2, x3))
```



```
def tup4(x1, x2, x3, x4):  
    (x1, (x2, (x3, x4)))
```



```
def tup5(x1, x2, x3, x4, x5):  
    (x1, (x2, (x3, (x4, x5))))
```

Accessing Tuples

We can write a single function to access tuples of any size.

So the below code

```
let yuple = (10, (20, (30, (40, (50, false)))))) in
```

```
get(yuple, 0) = 10
```

```
get(yuple, 1) = 20
```

```
get(yuple, 2) = 30
```

```
get(yuple, 3) = 40
```

```
get(yuple, 4) = 50
```

```
def tup3(x1, x2, x3):
```

```
  (x1, (x2, x3))
```

```
def tup5(x1, x2, x3, x4, x5):
```

```
  (x1, (x2, (x3, (x4, x5))))
```

```
let t = tup5(1, 2, 3, 4, 5) in
```

```
  , x0 = print(get(t, 0))
```

```
  , x1 = print(get(t, 1))
```

```
  , x2 = print(get(t, 2))
```

```
  , x3 = print(get(t, 3))
```

```
  , x4 = print(get(t, 4))
```

```
in
```

99

should print out:

0

1

2

3

4

99

How shall we write it?

```
def get(t, i):  
    TODO-IN-CLASS
```

QUIZ

Using the above “library” we can write code like:

```
let quad = tup4(1, 2, 3, 4) in  
  get(quad, 0) + get(quad, 1) + get(quad, 2) + get(quad, 3)
```

What will be the result of compiling the above?

1. Compile error
2. Segmentation fault
3. Other run-time error
4. 4
5. 10

QUIZ

Using the above “library” we can write code like:

```
def get(t, i):  
    if i == 0:  
        t[0]  
    else:  
        get(t[1],i-1)  
  
def tup3(x1, x2, x3):  
    (x1, (x2, (x3, false)))  
  
let quad = tup3(1, 2, 3) in  
    get(quad, 0) + get(quad, 1) + get(quad, 2) + get(quad, 3)
```

What will be the result of compiling the above?

1. Compile error
2. Segmentation fault
3. Other run-time error
4. 4
5. 10

Lists

Once we have pairs, we can start encoding **unbounded lists**.

To build a list, we need two constructor functions:

```
def empty():
    false

def cons(h, t):
    (h, t)
..
```

We can now encode lists as:

```
```python
cons(1, cons(2, cons(3, cons(4, empty()))))
```

## *Access*

To **access** a list, we need to know

1. Whether the list is `isEmpty`, and
2. A way to access the `head` and the `tail` of a non-empty list.

```
def isEmpty(l):
 l == empty()

def head(l):
 l[0]

def tail(l):
 l[1]
```

## *Examples*

We can now write various functions that build and operate on lists, for example, a function to generate the list of numbers between  $i$  and  $j$

```
def range(i, j):
 if (i < j):
 cons(i, range(i+1, j))
 else:
 empty()
```

range(1, 5)

which should produce the result

(1,(2,(3,(4,false))))

and a function to sum up the elements of a list:

```
def sum(xs):
 if (isEmpty(xs)):
 0
 else:
 head(xs) + sum(tail(xs))
```

sum(range(1, 5))

which should produce the result 10.

## *Recap*

We have a pretty serious language now, with:

- **Data Structures**

which are implemented using

- **Heap Allocation**
- **Run-time Tags**

which required a bunch of small but subtle changes in the

- runtime and compiler

In your assignment, you will add *native* support for n-ary tuples, letting the programmer write code like:


`(e1, e2, e3, ..., en)` # constructing tuples of arbitrary arity

`e1[e2]` # allowing expressions to be used as fields


Next, we'll see how to

- use the “tuple” mechanism to implement **higher-order functions** and
- reclaim unused memory via **garbage collection**.

---

 (<https://ucsd-cse131.github.io/sp21/feed.xml>)

 (<https://twitter.com/ranjitjhala>)

 (<https://plus.google.com/u/0/106612421534244742464>)

 (<https://github.com/ucsd-cse131/sp21>)

Copyright © Ranjit Jhala 2016–21. Generated by Hakyll (<http://jaspervdj.be/hakyll>),  
template by Armin Ronacher (<http://lucumr.pocoo.org>), Please suggest fixes here.  
(<http://github.com/ucsd-cse131/sp21>)