CS61 Section Notes
Week 4 (Fall 2011)

Outline for this week:

Processor Architecture
- Logic gates
- Instruction encoding
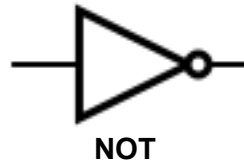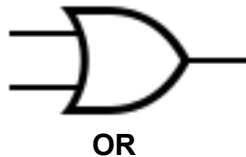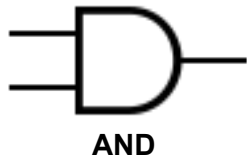- Sequential processing with stages
- Pipelining

Program Optimization
- Code motion
- Memory accesses
- Loop unrolling
- Pipelining
- Demo of efficiency gains

All page numbers cited below refer to the course textbook, *Computer Systems: A Programmer's Perspective, Second Edition* by Randal E. Bryant and David R. O'Hallaron.

Processor Architecture

1. Logic gates

| AND | OR | NOT |

Warm up: Fill in the truth table below for NAND.

To the right, design a circuit for bit-level NAND using OR and NOT gates.

a	b	a NAND b
0	0	
0	1	
1	0	
1	1	

In class, we introduced a circuit for XOR. Suppose you want to implement a word-level equality circuit using XOR rather than from bit-level equality circuits. Design such a circuit for a 32-bit word consisting of 32 bit-level XOR circuits and two additional logic gates.[1] (Assume that you can have up to 32 inputs to AND and OR gates.)

What determines the length of time to evaluate a circuit?

2. Instruction encoding

Consider the instruction set, register identifiers, and function codes on the next page. These are from the Y86 instruction set architecture introduced in the book. It is NOT important to memorize the details of Y86, which is inspired by IA32 but simpler and reduced in design. The point of the following exercises is simply to get a taste of instruction encoding and decoding.

Determine the byte encoding of the following instruction sequence. Using the byte encodings that you compute, also determine the addresses of the instructions below, as well as a hypothetical next instruction. The first instruction (`irmovl $15, %ebx`) starts at address `0x100`.[2]

```
0x100  :_____| irmovl $15, %ebx  # Load 15 into %ebx



0x_____:_____| rrmovl %ebx, %ecx # Copy 15 into %ecx



0x_____: next_encoding          | next_instruction
```

_____

[1]Adapted from Practice Problem 4.9, p. 356.
[2]Adapted from Practice Problem 4.1, p. 341.

Instruction set.

Byte	0	1	2	3	4	5

halt   `0 0`

nop   `1 0`

rrmovl rA, rB   `2 0 rA rB`

irmovl V, rB   `3 0 F rB    V`

rmmovl rA, D(rB)   `4 0 rA rB    D`

mrmovl D(rB), rA   `5 0 rA rB    D`

OPl rA, rB   `6 fn rA rB`

jXX Dest   `7 fn    Dest`

cmovXX rA, rB   `2 fn rA rB`

call Dest   `8 0    Dest`

ret   `9 0`

pushl rA   `A 0 rA F`

popl rA   `B 0 rA F`

Register identifiers.

Number	Register
0	%eax
1	%ecx
2	%edx
3	%ebx
4	%esp
5	%ebp
6	%esi
7	%edi
F	No register

Function codes.

Operation	Code
addl	60
subl	61
andl	62
xorl	63

Branches	Code
jmp	70
jle	71
jl	72
je	73
jne	74
jge	75
jg	76

Moves	Code
rrmovl	20
cmovle	21
cmovl	22
cmove	23
cmovne	24
cmovge	25
cmovg	26

For the byte sequence below, determine the instruction sequence it encodes. We show the starting address, then a colon, then the byte sequence.[3]

```
0x400: 6113 7300 0400 0000
```

3. Sequential processing with stages[4]

Consider a sequential processor that executes in stages. Some of the hardware units for the stages have state that may be updated during the execution of an instruction. Indicate which of the following hardware units hold state that might be updated:

Program counter (PC) _____   Arithmetic logic unit (ALU) _____

Condition code register (CC) _____   Data memory _____

Register file _____   Instruction memory _____

The sequential processor introduced in class and the textbook breaks the processing of a single instruction into six stages: Fetch, Decode, Execute, Memory, Write back, and PC update. Consider the processing of a single instruction. Do the stages update state in a serial fashion, i.e. do the stages update state in some order? When do these updates occur?

Principle: The processor never needs to read back the state updated by an instruction in order to complete the processing of this instruction. For example, some instructions (integer operations) set condition codes, and some instructions (jumps) read condition codes, but no instruction both sets and reads condition codes. This way, the condition codes are always up-to-date before another instruction needs to read them. (p. 380-1)

_____

[3]Adapted from Practice Problem 4.2, p. 341.
[4]Material and language from Section 4.3.

4. Pipelining

The latency of a stage is the time required to execute that stage. Throughput is the rate at which stages are executed, and so latency is equal to the inverse of throughput. E.g., If a car wash has a throughput of 2 cars per minute, its latency is ½ minutes = 30 seconds.

Suppose a pipelined processor executes N stages in parallel. Label the stages $s_1, s_2, \ldots, s_N$ such that $t_1 < t_2 < \ldots < t_N$, where t_i is the latency of stage i. What is the throughput of the processor when executing all N stages in parallel?

Pipelining divides a computation into separate stages separated by pipeline registers. What happens as we divide the computation into increasingly shorter stages?

Program Optimization

Ideally, a compiler should be able to take any code and generate the most efficient machine-level program with the correct behavior. In reality, compilers can only perform limited transformations like code motion, strength reduction, loop unrolling, etc. Even these transformations can be thwarted by *optimization blockers* – aspects of of the program's behavior that depend strongly on the execution environment.

Consider the following C code...

```
struct vector {
    int length;
    int *data;
}
typedef struct vector *vec_ptr;
int get_vec_length(vec_ptr v);
int *get_vec_data(vec_ptr v);

/* Multiply the elements of a vector. */
void multiply1(vec_ptr v, int *dest) {
    int i;
    int *data = get_vec_data(v);

    *dest = 1;
    for (i = 0; i < get_vec_length(v); i++) {
        *dest = *dest * data[i];
    }
}
```

Below, we will consider what optimizations are possible here.

1. Code Motion - Move code around to reduce the number of times it executes.
The call `get_vec_length(v)` can be moved outside the loop, because we know the length of the vector won't change from one iteration to the next.

```
void multiply2(vec_ptr v, int *dest) {
      int i;
      int *data = get_vec_data(v);
      int length = get_vec_length(v);

      *dest = 1;
      for (i = 0; i < length; i++) {
            *dest = *dest * data[i];
      }
}
```

Could `gcc -01` make this optimization? If not, why not, and could we change anything to help?

2. Reduce Memory Access
Introduce a local variable, `x`, to store intermediate computations, and copy the final result at the end to `*dest`.

```
void multiply3(vec_ptr v, int *dest) {
      int i;
      int *data = get_vec_start(v);
      int length = vec_length(v);
      int x = 1;

      for (i = 0; i < length; i++) {
            x = x * data[i];
      }
      *dest = x;
}
```

Why is this an optimization?

Could `gcc` make this optimization? If not, why not?

3. Loop Unrolling - Decrease the number of iterations by doing more work per iteration.

Consider the following implementation of `multiply` (assume `length % 2 == 0`).

```
void multiply4(vec_ptr v, int *dest) {
    int i;
    int *data = get_vec_start(v);
    int length = vec_length(v);
    int x = 1;
    int y = 1;

    for (i = 0; i < length; i += 2) {
        x = x * data[i];
        y = y * data[i + 1];
    }
    *dest = x * y;
}
```

This implementation is even more efficient! Why? (There are two principal reasons.)

Wouldn't it be even better to unroll the loop more, say to `i += 20`?

A Practical Note - Loop unrolling and taking advantage of pipelining are examples of optimizations that can be very hard for human programmers to implement effectively in practice. Additionally, we usually can rely on our compiler to do these types of optimizations for us!

Interactive Optimization Demo

At this point, please open up your laptops and follow the instructions below, substituting your username for user where appropriate. First, ssh to your CS 61 VM. Copy the file from Prof Chong's home directory to your own:

```
$ cp ~stephenchong/optimization.tar ~
```

You should now have a file called optimization.tar in your home directory. Let's extract the files in the tar.

```
$ tar xvf optimization.tar
$ cd optimization
```

This code is a function that implements a simple substitution cipher. It's pretty silly code, and there are dummy functions intended to let you practice various optimizations in the file. Get in small groups and discuss the types of optimizations you might implement!

Take a look at the file `benchmark.c`, in particular the function `run_benchmarks()` at the bottom of the file. As you can see, this function benchmarks a series of implementations of the substitution cipher: `cipher_orig`, `cipher_better`, `cipher_faster`, and `cipher_fast`. Let's run the benchmarks.

To run the benchmarks without compiler optimization (i.e., with `-O0`), type:
```
$ make run
```

To run the benchmarks with compiler optimization (i.e., with `-O3`), type:
```
$ make run-opt
```

Initially, all four of the cipher functions do the same thing, so running it should give very similar results for all four trials.

Take a look at the code for `cipher_orig`. Do not edit this function: it is used to provide a baseline for the improvements you'll be making. You'll be editing the functions: `cipher_better`, `encode_char_better`, `cipher_faster`, and `cipher_fast` (in that order). Follow the instructions in the comments in `benchmark.c` and from your TF. Feel free to look in `support.c` to inspect the main driver routine, including timing, etc.

When you're done implementing your optimizations, come back together and discuss what you tried, what worked, what didn't, and what things the compiler optimizations took care of for you.

A couple of final take-away points:
• Readability/maintainability is often more important than minor optimizations.
• The compiler is often smarter than you are.
• Low-level things like loop unrolling are often architecture-dependent, and should be done only with extreme caution! Nevertheless, it's good to understand how the CPU works.
• Be sure you know what you're optimizing and why before you start. It's easy to waste days only to realize that all your hard work was unimportant in the grander scheme of things.
• Testing is even more key. It doesn't matter how fast your code runs if it doesn't work!

EXTRA BONUS SECTION

Program Optimization, continued

4. Pipelining - Take advantage of how the machine pipelines instructions.

Pipelining, the process of starting the next instruction before the previous one has finished, can be extremely useful, especially when you're looking at expensive operations like multiplication, but it won't work if the operations need to be executed in sequence.

Consider the following C code...

```
void copy_array(int *src, int*dest, int n) {
     int i;
     for (i = 0; i < n; i++)
          dest[i] = src[i];
}
int a[1000] = {0, 1, 2, ..., 999};
```

It turns out `copy_array(a+1, a, 999);` executes about twice as fast as `copy_array(a, a+1, 999);`. Why is this?