

From Hardware to Software: Moving to the OS!

Up until now, everything we have discussed has been focused on the direct interaction with hardware. Today, we escape hardware to explore **controlling** and **using** the hardware!

Operating System (OS):

[1]:

[2]:

[3]:

Apps
OS
Hardware

Running a Program

A _____ is an instance of a running program.
- **NOT** the binary source code; it's an instance of it running.

A _____ provides two key abstractions:

1. [Memory]:

2. [Execution]:

Pressing the Power Button

When you press the power button on **ANY** computing device, not much is available:

- Your _____ after being off (ex: does not contain your Operating System, that's on your disk!)
- Goals of a computer at boot:
 - 1.
 - 2.

To do this, almost all modern systems perform three tasks:

1. [POST]:

2. [Bootloader]:

3. [Process #1]:

Once one process exists, it can spawn new processes through the **fork()** command. You can explore every process on a system:

- Linux command: **ps**
- ...options for **all** processes with details: **ps -aef**

Multiprogramming

On a modern computer, there are dozens of different processes running simultaneously -- but only a few CPUs.

- In the period of microseconds, the OS rapidly switches between all processes to **allow each process to run on one or more of the CPUs**.
- When the OS swaps out one process from one CPU and allows a new process to run, this is called a _____.

Context Switching

What is required during a context switch?

[CPU]:

[Caches]:

[Page Table]:

...overall cost?

A Process' Execution Unit: Threads

Each process contains one or more _____ that can run concurrently with other threads:

[1]: Main Thread:

[2]: Uses for Additional Threads:

Threads vs. Processes

	Threads with a Process	Processes
Overhead		
Context Switching		
Virtual Memory		
CPUs		

Creating Threads in C

The pthread library is the POSIX thread library allowing you to create additional threads beyond the main thread.

Creating a new thread is a complex call with four arguments:

```
int pthread_create(  
    pthread_t *thread,           /* thread struct */  
    const pthread_attr_t *attr, /* usually NULL */  
    void *(*start_routine) (void *), /* start func */  
    void *arg                    /* thread start arg */  
);
```

The start_routine variable is a function pointer and requires the argument to be a function with the prototype:

```
void *_____ (void *ptr);
```

...you can use any name for the function name.

Example 1: Launching Fifteen Threads

```
threads/fifteen.c
1 #include <stdio.h>
2 #include <pthread.h>
3 #include <stdlib.h>
4
5 const int num_threads = 15;
6
7 void *thread_start(void *ptr) {
8     int id = *((int *)ptr);
9     printf("Thread %d running...\n", id);
10    return NULL;
11 }
12
13 int main(int argc, char *argv[]) {
14     // Create threads:
15     int i;
16     pthread_t tid[num_threads];
17     for (i = 0; i < num_threads; i++) {
18         pthread_create(&tid[i], NULL,
19                       thread_start, (void *)&i);
19     }
20
21     printf("Done!\n");
22     return 0;
23 }
```

Q1: What is the expected output of this program?

Q2: What actually happens?

Q3: What do we know about threads in C?

Example 2: Joining Threads

```
threads/fifteen-join.c
13 int main(int argc, char *argv[]) {
14     // Create threads:
15     int i;
16     pthread_t tid[num_threads];
17     for (i = 0; i < num_threads; i++) {
18         int *val = malloc(sizeof(int));
19         *val = i;
20         pthread_create(&tid[i], NULL,
21                       thread_start, (void *)val);
21     }
22
23     // Joining Threads
24     for (i = 0; i < num_threads; i++) {
25         pthread_join(tid[i], NULL);
26     }
27
28     printf("Done!\n");
29     return 0;
30 }
```

In the above program, we use `pthread_join`. This call will block the CPU from running the program further until the specified thread has **finished and returned**.

Q4: What happens in this program?

Q5: Does the order vary each time we run it? What is happening?

Q6: What can we say about the relationship between “Done” and “Thread %d running...” lines?

threads/count.c

```
5 int ct = 0;
6
7 void *thread_start(void *ptr) {
8     int countTo = *((int *)ptr);
9
10    int i;
11    for (i = 0; i < countTo; i++) {
12        ct = ct + 1;
13    }
14
15    return NULL;
16 }
17
18 int main(int argc, char *argv[]) {
19     // Parse Command Line:
20     if (argc != 3) {
21         printf("Usage: %s <countTo> <thread count>\n",
22             argv[0]);
23         return 1;
24     }
25
26     const int countTo = atoi(argv[1]);
27     if (countTo == 0) { printf("Valid `countTo` is
28         required.\n"); return 1; }
29
30     const int thread_ct = atoi(argv[2]);
31     if (thread_ct == 0) { printf("Valid thread count is
32         required.\n"); return 1; }
33
34     // Create threads:
35     int i;
36     pthread_t tid[thread_ct];
37     for (i = 0; i < thread_ct; i++) {
38         pthread_create(&tid[i], NULL,
39             thread_start, (void *)&countTo);
40     }
41
42     // Join threads:
43     for (i = 0; i < thread_ct; i++) {
44         pthread_join(tid[i], NULL);
45     }
46
47     // Display result:
48     printf("Final Result: %d\n", ct);
49     return 0;
50 }
```

Q7: What do we expect when we run this program?

Q8: What is the output of this program when it's running as:
./count 100 2

Q9: What is the output of this program when it's running as:
./count 100 16

Q10: What is the output of this program when it's running as:
./count 10000000 2

Q11: What is the output of this program when it's running as:
./count 10000000 16

Q12: What is going on???