

**Program Execution: Direct Execution Model**

Operating System		Process + Thread
1. Create entry for process and thread 2. Allocate memory for process and thread 3. Load program into memory 4. Set up stack (argv/argc) 5. Clear registers 6. Call main()	⇒	(OS has CPU control.)
(OS does not have CPU control.)		7. Run main() 8. return from main()
9. Free memory for process 10. Remove process from process list	⇐	(OS has CPU control.)

What is the problem with this model?

**Protection Levels:**

We can develop multiple levels of protection so that the process must request access to certain resources from the Operating System. This way, there is a “handoff” between the OS and the application:

Operating System		Process + Thread
1. Process Init 2. return-from-trap	⇒	(OS has CPU control.)
(OS does not have CPU control.)		3. Run main() 4. Makes a system call ...calls <b>trap-to-OS</b>
5. Trap Handler ...do syscall work... 6. return-from-trap	⇐	(OS has CPU control.)
(OS does not have CPU control.)	⇒	...execution continues...

**Execution: Kernel and User Modes**

There are many different “protection levels” in modern systems:

- Ring 0 (“kernel mode”):
- Ring 1 and 2:
- Ring 3 (“user mode”):
- Ring -1 (VT-x):

**Trapping to the OS: More than Just System Calls**

There are several mechanisms to regain CPU control from an application back to the OS:

1. **System Calls:**
- 2.
- 3.

**Additional Reading:** “Operating Systems: Three Easy Pieces”  
Ch. 6: 6 Direct Execution (<https://pages.cs.wisc.edu/~remzi/OSTEP/>)

**Five-State Thread Model**

When the operating system has control over the CPU and needs to decide what program to run, it must maintain a model of all threads within the CPU.

We commonly refer to the “state” of a thread as part of the five-state model:

08/thread-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", argv[0]);
22         return 1;
23     }
24
25     const int countTo = atoi(argv[1]);
26     /* [...error checking...] */
27
28     const int thread_ct = atoi(argv[2]);
29     /* [...error checking...] */
30
31     // Create threads:
32     int i;
33     pthread_t tid[thread_ct];
34     for (i = 0; i < thread_ct; i++) {
35         pthread_create(&tid[i], NULL,
36                       thread_start, (void *)&countTo);
37     }
38
39     // Join threads:
40     for (i = 0; i < thread_ct; i++) {
41         pthread_join(tid[i], NULL);
42     }
43
44     // Display result:
45     printf("Final Result: %d\n", ct);
46     return 0;
47 }
```

## Multiple Threads and Synchronization

In the program to the left, we launch a number of different threads that will count up together in parallel!

**Q1:** What do we expect when we run this program?

**Q2:** What is the output of this program when it's running as:

`./count 100 2`

**Q3:** What is the output of this program when it's running as:

`./count 100 16`

**Q4:** What is the output of this program when it's running as:

`./count 10000000 2`

**Q5:** What is the output of this program when it's running as:

`./count 10000000 16`

**Q6:** What is going on???