

Operating Systems And Applications For Embedded Systems

Processes and Threads

Plan

Processes

- Process definition
- Creating a new process
- Output
- Running a different program


Threads

- Thread definition
- Creating a new thread
- Terminating a thread
- Compiling a program with threads
- Partitioning the problem
- Scheduling
- Further reading

Process definition

A process is a memory address space and a thread of execution, as shown in the following diagram. The address space is private to the process and so threads running in different processes. cannot access it. This memory separation is created by the memory management subsystem in the kernel, which keeps a memory page mapping for each process and re-programs the memory management unit on each context switch.

Creating a new process I



Creating a new process II

Here is a simple example, showing process creation and termination:

Listing 1: Listing

```
1 #include <stdio.h>
2 #include <stdlib.h>
3 #include <unistd.h>
4 #include <sys/types.h>
5 #include <sys/wait.h>
6 int main(void)
7 {
8     int pid;
9     int status;
10    pid = fork();
11    if (pid == 0) {
12        printf("I am the child, PID %d\n", getpid());
```

Creating a new process III

```
13     sleep(10);
14     exit(42);
15 } else if (pid > 0) {
16     printf("I am the parent, PID %d\n", getpid());
17     wait(&status);
18     printf("Child terminated, status %d\n",
19           WEXITSTATUS(status));
20 } else
21     perror("fork:");
22 return 0;
23 }
```

Output

I am the parent, PID 13851
I am the child, PID 13852
Child terminated with status 42

Running a different program I

1. `int execl(const char *path, const char *arg, ...);`
2. `int execlp(const char *file, const char *arg, ...);`
3. `int execlx(const char *path, const char *arg, ..., char * const envp[]);`
4. `int execv(const char *path, char *const argv[]);`
5. `int execvp(const char *file, char *const argv[]);`
6. `int execvpe(const char *file, char *const argv[], char *const envp[]);`

Running a different program II

Listing 2: Listing

```
1  #include <stdio.h>
2  #include <stdlib.h>
3  #include <string.h>
4  #include <unistd.h>
5  #include <sys/types.h>
6  #include <sys/wait.h>
7  int main(int argc, char *argv[])
8  {
9  char command_str[128];
10 int pid;
11 int child_status;
12 int wait_for = 1;
13 while (1) {
14     printf("sh> ");
```

Running a different program III

```
15 scanf("%s", command_str);
16 pid = fork();
17 if (pid == 0) {
18     /* child */
19     printf("cmd '%s'\n", command_str);
20     execl(command_str, command_str, (char *)NULL);
21     /* We should not return from execl, so only get
22     to this line if it failed */
23     perror("exec");
24     exit(1);
25 }
26 if (wait_for) {
27     waitpid(pid, &child_status, 0);
28     printf("Done, status %d\n", child_status);
29 }
```

Running a different program IV

```
30     }  
31     return 0;  
32 }
```

Thread definition I

A thread is a thread of execution within a process. All processes begin with one thread that runs the `main()` function and is called the main thread. You can create additional threads using the POSIX threads function `pthread_create(3)`, causing additional threads to execute in the same address space, as shown in the following diagram. Being in the same process, they share resources with each other. They can read and write the same memory and use the same file descriptors, and so communication between threads is easy, so long as you take care of the synchronization and locking issues.

Creating a new thread I

```
int pthread_create(pthread_t *thread, const pthread_attr_t *attr, void *(*start_routine)  
(void *), void *arg);
```

Creating a new thread II

Listing 3: Listing

```
1  #include <stdio.h>
2  #include <unistd.h>
3  #include <pthread.h>
4  #include <sys/syscall.h>
5  static void *thread_fn(void *arg)
6  {
7      printf("New thread started , PID %d TID %d\n",
8      getpid(), (pid_t)syscall(SYS_gettid));
9      sleep(10);
10     printf("New thread terminating\n");
11     return NULL;
12 }
13 int main(int argc, char *argv[])
14 {
```

Creating a new thread III

```
15 pthread_t t;  
16 printf("Main thread , PID %d TID %d\n",  
17 getpid(), (pid_t)syscall(SYS_gettid));  
18 pthread_create(&t, NULL, thread_fn, NULL);  
19 pthread_join(t, NULL);  
20 return 0;  
21 }
```

Terminating a thread

1. It reaches the end of its start_routine
2. It calls pthread_exit(3)
3. It is canceled by another thread calling pthread_cancel(3)
4. The process which contains the thread terminates, for example, because of a thread calling exit(3), or the process receiving a signal that is not handled, masked or ignored

Compiling a program with threads

The support for POSIX threads is part of the C library, in the library `libpthread.so`. When building a threaded program, you must add the switch `-pthread` at the compile and link stages.

Partitioning the problem

1. Keep tasks that have a lot of interaction.
Minimize overheads by keeping closely inter-operating threads together in one process.
2. Don't put all your threads in one basket.
On the other hand, try and keep components with limited interaction in separate processes, in the interests of resilience and modularity.
3. Don't mix critical and non-critical threads in the same process.
This is an amplification of Rule 2: the critical part of the system, which might be the machine control program, should be kept as simple as possible and written in a more rigorous way than other parts. It must be able to continue even if other processes fail. If you have real-time threads, they, by definition, must be critical and should go into a process by themselves.
4. Threads shouldn't get too intimate.
One of the temptations when writing a multi-threaded program is to intermingle the code and variables between threads because it is all in one program and easy to do. Don't keep threads modular with well-defined interactions.
5. Don't think that threads are for free.
It is very easy to create additional threads but there is a cost, not least in the additional synchronization necessary to coordinate their activities.

6. Threads can work in parallel.

Threads can run simultaneously on a multi-core processor, giving higher throughput. If you have a large computing job, you can create one thread per core and make maximum use of hardware. There are limits, of course. You can't create more threads than you have hardware.

Scheduling

1. A thread blocks by calling `sleep()` or in a blocking I/O call
2. A timeshare thread exhausts its time slice
3. An interrupt causes a thread to be unblocked, for example, because of I/O completing

Further reading

1. The Art of Unix Programming, by Eric Steven Raymond, Addison Wesley; (23 Sept. 2003) ISBN 978-0131429017
2. Linux System Programming, 2nd edition, by Robert Love, O'Reilly Media; (8 Jun. 2013) ISBN-10: 1449339530
3. Linux Kernel Development, 3rd edition by Robert Love, Addison-Wesley Professional; (July 2, 2010) ISBN-10: 0672329468
4. The Linux Programming Interface, by Michael Kerrisk, No Starch Press; (October 2010) ISBN 978-1-59327-220-3
5. UNIX Network Programming: v. 2: Interprocess Communications, 2nd Edition, by W. Richard Stevens, Prentice Hall; (25 Aug. 1998) ISBN-10: 0132974290
6. Programming with POSIX Threads, by Butenhof, David R, Addison-Wesley, Professional
7. Scheduling Algorithm for multiprogramming in a Hard-Real-Time Environment, by C. L. Liu and James W. Layland, Journal of ACM, 1973, vol 20, no 1, pp. 46-61

References



C. Simmonds.

Mastering Embedded Linux Programming.

Packt Publishing, 2015.

The End