

Exercise session (Processes)

Operating Systems – EDA093/DIT401



UNIVERSITY OF
GOTHENBURG

Question 1

Describe the actions taken by a kernel to context-switch between processes

See slides for lecture 2...

Question 2

What is printed by this program?

CHILD 0
CHILD -1
CHILD -4
CHILD -9
CHILD -16
PARENT 0
PARENT 1
PARENT 2
PARENT 3
PARENT 4

```
1 #include <stdio.h>
2 #include <sys/types.h>
3 #include <unistd.h>
4
5 #define SIZE 5
6
7 int nums[SIZE] = {0,1,2,3,4};
8
9 int main()
10 {
11     int i;
12     pid_t pid;
13     pid = fork();
14     if (pid == 0) {
15         for (i = 0; i < SIZE; i++) {
16             nums[i] *= -i;
17             printf("CHILD %d\n", nums[i]);
18         }
19     }
20     else if (pid > 0) {
21         wait(NULL);
22         for (i = 0; i < SIZE; i++)
23             printf("PARENT: %d\n", nums[i]);
24     }
25     return 0;
26 }
```

Question 3

What is printed by this program?

child: pid = 0
child: pid1 = X
parent: pid = X
parent: pid1 = Y

Notice:

X is the same for child and parent

Y is different from X

X > 1

Y > 1

The parent is not waiting for the child to print. Messages can be printed in different order...

```
1 #include <stdio.h>
2 #include <sys/types.h>
3 #include <unistd.h>
4
5 int main()
6 {
7     pid_t pid, pid1;
8     pid = fork();
9     if (pid < 0) {
10         fprintf(stderr, "Fork Failed");
11     }
12     else if (pid == 0) {
13         pid1 = getpid();
14         printf("child: pid = %d", pid)
15         printf("child: pid1 = %d", pid1)
16     }
17     else if (pid > 0) {
18         pid1 = getpid();
19         printf("parent: pid = %d", pid)
20         printf("parent: pid1 = %d", pid1)
21         wait(NULL);
22     }
23     return 0;
24 }
```

Question 4

Consider a multiprogrammed system with degree of 5. If each process spends 40% of its time waiting for I/O, what will be the CPU utilization?

$$\text{CPU utilization} = 1 - p^n = 1 - 0.4^5 = 0.99$$

where p is the probability for a process to be waiting for I/O.

Question 5

To use cache memory, the main memory is divided into cache lines, typically 32 or 64 bytes long. An entire cache line is cached at once. What is the advantage of caching an entire line instead of a single byte or word at a time?

Question 5 - Solution

- Locality of reference: if one location is read then the probability of accessing nearby locations next is very high, particularly the following memory locations. So, by caching an entire cache line, the probability of a cache hit next is increased.
- Modern hardware can do a block transfer of 32 or 64 bytes into a cache line much faster than reading the same data as individual words.