

# Operating Systems ID1206

(ID2200/06 6hp)

## Exam

2019-01-11 14:00-18:00

### Instruction

- You are, besides writing material, only allowed to bring one self hand written A4 of notes. The notes are handed in and can not be reused.
- All answers should be written in these pages, use the space allocated after each question to write down your answer.
- Answers should be written in swedish or english.
- You should hand in the whole exam and the hand written page of notes. No additional pages should be handed in.

### Grades

The exam is divided into a number of questions where some are a bit harder than others. The harder questions are marked with a star *points\**, and will give you points for the higher grades. The exam is thus divided into basic points and points for higher grades. First of all make sure that you pass the basic points before engaging with the higher points.

Questions with multiple sub-questions are normally awarded 2p for all correct and 1p for one wrong answer.

Note that, of the 24 basic points only at most 22 are counted, the points for higher grades will not make up for lack of basic points. The limits for the grades are as follows:

- Fx: 12 basic points
- E: 13 basic points
- D: 16 basic points
- C: 20 basic points
- B: 22 basic points and 6 higher points
- A: 22 basic points and 10 higher points

The limits could be adjusted to lower values but not raised.

Name: \_\_\_\_\_ Persnr: \_\_\_\_\_

## 1 Processer

### 1.1 stack or heap [2 points]

What is done in the procedure below and where should **gurka** be allocated? Why? Complete the code so that **gurka** is allocated space.

```
void tomat(int *a, int *b) {
    // allocate room for gurka

    gurka = *a;
    *a = *b;
    *b = gurka;
}
```

### 1.2 fork() [2 points]

What is printed when we run the program below, what alternatives exist and why do we get this result?

```
#include <unistd.h>
#include <stdio.h>
#include <sys/wait.h>

void foo(int *x) {
    (*x)++;
}

int main() {
    int global = 17;
    int pid = fork();
    if(pid == 0) {
        foo(&global);
    } else {
        foo(&global);
        wait(NULL);
        printf("global = %d \n", global);
    }
    return 0;
}
```

Name: \_\_\_\_\_ Persnr: \_\_\_\_\_

### 1.3 a stack, a bottle and.. [2 points]

You have written the program below to examine what is on the stack.

```
void zot(unsigned long *stop ) {
    unsigned long r = 0x3;
    unsigned long *i;
    for(i = &r; i <= stop; i++){ printf("%p          0x%lx\n", i, *i); }
}

void foo(unsigned long *stop ) {
    unsigned long q = 0x2;
    zot(stop);
}

int main() {
    unsigned long p = 0x1;
    foo(&p);
back:
    printf(" p: %p \n", &p);
    printf(" back: %p \n", &&back);
    return 0;
}
```

This is the print out. Explain what is found at the locations indicated by arrows.

```
0x7ffca03d1748      0x3
0x7ffca03d1750      0x7ffca03d1750
0x7ffca03d1758      0xb93d7906926a7d00
0x7ffca03d1760      0x7ffca03d1790      <-----
0x7ffca03d1768      0x55cdac31d78c      <-----
0x7ffca03d1770      0x7ffca03d17d8
0x7ffca03d1778      0x7ffca03d17b0
0x7ffca03d1780      0x1
0x7ffca03d1788      0x2
0x7ffca03d1790      0x7ffca03d17c0
0x7ffca03d1798      0x55cdac31d7c2
0x7ffca03d17a0      0x55cdac31d810
0x7ffca03d17a8      0x12acac31d5f0
0x7ffca03d17b0      0x1
p: 0x7ffca03d17b0
back: 0x55cdac31d7c2
```


Name: \_\_\_\_\_ Persnr: \_\_\_\_\_

#### 1.4 the size of the block? [2 points]

When implementing *free()* we need to know the size of the block that should be freed. How do we know the size? Draw and explain in the figure below what an implementation could look like.

```
int *new_table(int elements) {
    return (int*)malloc(sizeof(int)*elements);
}

int main() {
    int *table = new_table(24);
    :
    table[0] = 29;
    table[1] = 13;
    table[2] = 87;
    :
    free(table);
    return 0;
}
```



Name: \_\_\_\_\_ Persnr: \_\_\_\_\_

### **1.5 IDTR [2 points\*]**

When executing the INT instruction, there is a jump to given position in the IDT. At the same time there is a transition from user mode to kernel mode. This is done to allow the operating system to access its own data structures. What prevents a user from changing IDTR so that it is referring to a table that it controls?

### **1.6 library call vs system call [2 points\*]**

An operating system that implements POSIX should provide specified functionality to a user process. Is this provided by system calls, library procedures or a combination of both? Explain the difference between system calls and procedure calls and which parts belong to the operating system.

Name: \_\_\_\_\_ Persnr: \_\_\_\_\_

## 2 Communication

### 2.1 a buffer [2 points]

We implement a buffer with one element as shown below (get() defined in similar way). We will have several threads that produce and consume values from the buffer. The buffer is protected by a lock and the threads are synchronized using a conditional variable. The program below does not work (the calls to the pthread procedures are not even legal), why does it not work (even if they were legal)?

```
#define TRUE 1
#define FALSE 0

volatile int buffer = 0;
volatile int empty = TRUE;

pthread_mutex_t lock;
pthread_cond_t signal;

void put(int value) {
    pthread_mutex_lock(&lock);
    while (TRUE) {
        if (empty) {
            buffer = value;
            empty = FALSE;
            pthread_cond_signal(&signal);
            pthread_mutex_unlock(&lock);
            break;
        } else {
            pthread_mutex_unlock(&lock);
            pthread_cond_wait(&signal);
        }
    }
}
```

Name: \_\_\_\_\_ Persnr: \_\_\_\_\_

## 2.2 pipes [2 points]

The program below opens a pipe and iterates a number of times (ITERATIONS) where each iteration sends a number (BURST) of messages ("0123456789"). We need to handle the situation where the receiving process will not keep up with the sender; how do we implement flow-control to avoid buffer overflow?

```
int main() {
    int mode = S_IRUSR | S_IWUSR | S_IRGRP | S_IROTH;
    mkfifo("sesame", mode);
    // add flow control

    int flag = O_WRONLY;
    int pipe = open("sesame", flag);

    /* produce quickly */
    for(int i = 0; i < ITERATIONS; i++) {
        for(int j = 0; j < BURST; j++) {
            write(pipe, "0123456789", 10);
            // add flow control

        }
        printf("producer burst %d done\n", i);
    }
    printf("producer done\n");
}
```

## 2.3 SOCK\_WHAT [2 points\*]

When you create a socket you can choose to create a *SOCK\_STREAM* or *SOCK\_DGRAM*. Which properties differ and when is it an advantage to choose one over the other.

Name: \_\_\_\_\_ Persnr: \_\_\_\_\_

### 3 Scheduling

#### 3.1 Bonnie Tylor [2 points]

Assume that we have a scheduler that implements *shortest time to completion first* or as it is also called *preemptive shortest job first*. We have four jobs described below as  $\langle \text{arrive at, execution time} \rangle$  in ms. Draw a time diagram and specify the turnaround time for each of the jobs.


- J1 :  $\langle 0,40 \rangle$
- J2 :  $\langle 0,30 \rangle$
- J3 :  $\langle 10,10 \rangle$
- J4 :  $\langle 20,30 \rangle$

J1:

J2:

J3:

J4:






Name: \_\_\_\_\_ Persnr: \_\_\_\_\_

### 3.2 random not that wrong [2 points]

When benchmarking an implementation for swapping one found that a completely random algorithm worked better than expected. In the graph below you can see the ratio of page hits and page references when the size of the memory is changed. The process that runs uses 100 pages and the different runs shows how the hit ratio varies with the size of the memory.

One would expect 50% hit ratio when the memory is half of what the process needs, but the hit ratio is better than this. What could a plausible explanation to the observed behavior be?



Name: \_\_\_\_\_ Persnr: \_\_\_\_\_

### **3.3 priority inversion [2 points\*]**

When one uses a scheduler with strict priority and uses locks at the same time one can encounter some problems. Describe how a high priority process can wait forever on a lower priority process although the scheduler guarantees that higher priority processes always have precedence before lower priority processes.

Name: \_\_\_\_\_ Persnr: \_\_\_\_\_

## 4 Virtual memory

### 4.1 approximate what? [2 points]

The clock algorithm, that is used to swap pages from memory, is described as an approximation. What is it that it tries to approximate? Describe a scenario when it does not do the right choice because of the approximation.

### 4.2 implement the buddy algorithm [2 points]

Assume that you should implement the buddy-algorithm for memory management. To help you a function that locates the buddy of a block of size  $k$  is given. Assume that all blocks are tagged as either free or taken and have a field that gives the size of the block. Free blocks also have two pointers that link the block in a double linked list of free blocks of its size.

Assume that you should free a block and have found its buddy - what do you have to check before you can coalesce the block with its buddy? If the blocks can be coalesced, what are the operations that you need to perform.

Name: \_\_\_\_\_ Persnr: \_\_\_\_\_

### **4.3 x86\_32 addressing 512 byte pages [2 points\*]**

Assume that someone wants to create a x86-architecture for embedded systems where a large virtual address space is not essential. You're asked to propose a virtual memory architecture.

The word length is 32 bits and we should of course use a paged virtual memory. One requirement is that the page size is 512 bytes so this will set some constraints on your solution.

How do we divide a virtual address using a hierarchical page table based on a page size of 512 bytes? Give an example on how a virtual address should be decoded; the virtual address space need not be fully 32 bits.

Name: \_\_\_\_\_ Persnr: \_\_\_\_\_

## 5 File systems and storage

### 5.1 which goes where [2 points]

A file has many properties; where do we find the below listed properties? Connect the properties to the left with the correct locations to the right. Several properties might be found at the same location but each property is only found at one location (many to one).

<b>property</b>	<b>location</b>
current write position •	• open file table entry
text name of file •	• inode of the file
size of the file •	• data block of the file
mapping of <i>stdin</i> •	• directory that holds link to file
	• file descriptor table

### 5.2 different types [2 points]

A directory can hold links of different types, describe the type of the following five links:

```
drwxrwxr-x 2 johanmon johanmon 4096 dec 21 17:43 bar.doc
-rwxrwxr-x 1 johanmon johanmon 8464 dec 21 22:25 cave
lrwxrwxrwx 1 johanmon johanmon    7 dec 21 17:43 foo.pdf -> ./gurka
-rw-rw-r-- 1 johanmon johanmon    7 dec 21 17:42 gurka
prw-r--r-- 1 johanmon johanmon    0 dec 21 22:25 sesame
```

Name: \_\_\_\_\_ Persnr: \_\_\_\_\_

### 5.3 journal based fs [2 points\*]

Assume that we have a journal based file system. Which of the following statements below is/are correct and which ones is/are false - explain why.

- It is important the the transactions are checkpointed in the same order as they were created.
- A transaction is considered to be valid ones it is written to the journal.
- In a restart after a crash, we must be careful not to checkpoint a transaction twice.