

Section 2: Processes

February 1, 2019

Contents

1	Vocabulary	2
2	Warmup	3
2.1	Hello World	3
3	Problems	3
3.1	Forks	3
3.2	Stack Allocation	4
3.3	Heap Allocation	4
3.4	Slightly More Complex Heap Allocation	5
3.5	Simple Wait	6
3.6	Fork and File Descriptors	7
3.7	Exec	7
3.8	Exec + Fork	8
3.9	Implementing fork() efficiently (Design)	8

1 Vocabulary

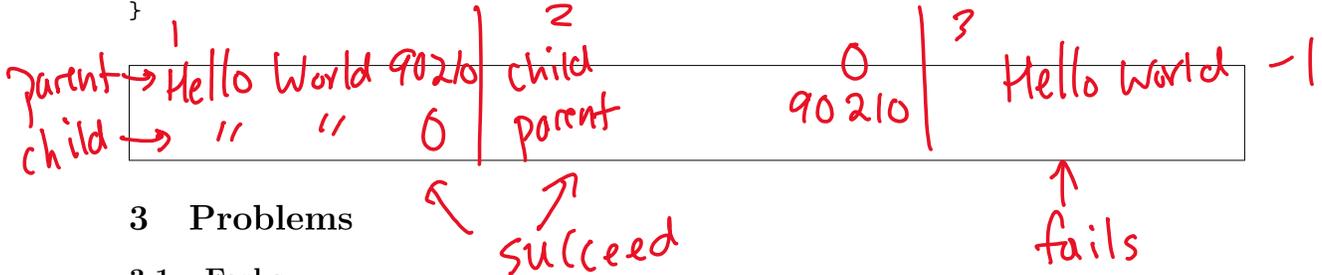
- **process** - a process is an instance of a computer program that is being executed. It consists of an address space and one or more threads of control.
- **address space** - The address space for a process is the set of memory addresses that it can use. The address space for each process is private and cannot be accessed by other processes unless it is shared.
- **stack** - The stack is the memory set aside as scratch space for a thread of execution. When a function is called, a block is reserved on the top of the stack for local variables and some book-keeping data. When that function returns, the block becomes unused and can be used the next time a function is called. The stack is always reserved in a LIFO (last in first out) order; the most recently reserved block is always the next block to be freed.
- **heap** - The heap is memory set aside for dynamic allocation. Unlike the stack, there's no enforced pattern to the allocation and deallocation of blocks from the heap; you can allocate a block at any time and free it at any time.
- **fork** - A C function that calls the fork syscall that creates a new process by duplicating the calling process. The new process, referred to as the child, is an exact duplicate of the calling process (except for a few details, read more in the man page). Both the newly created process and the parent process return from the call to fork. On success, the PID of the child process is returned in the parent, and 0 is returned in the child. On failure, -1 is returned in the parent, no child process is created.
- **wait** - A class of C functions that call syscalls that are used to wait for state changes in a child of the calling process, and obtain information about the child whose state has changed. A state change is considered to be: the child terminated; the child was stopped by a signal; or the child was resumed by a signal.
- **exit code** - The exit status or return code of a process is a 1 byte number passed from a child process (or callee) to a parent process (or caller) when it has finished executing a specific procedure or delegated task
- **exec** - The exec() family of functions replaces the current process image with a new process image. The initial argument for these functions is the name of a file that is to be executed.

2 Warmup

2.1 Hello World

What can C print in the below code? Assume the child's PID is 90210
 (Hint: There is more than one correct answer)

```
int main() {
    pid_t pid = fork();
    printf("Hello World: %d\n", pid);
}
```

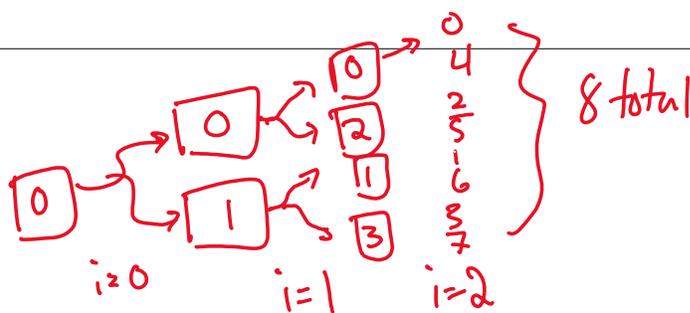


3 Problems

3.1 Forks

How many new processes are created in the below program assuming calls to fork succeeds?

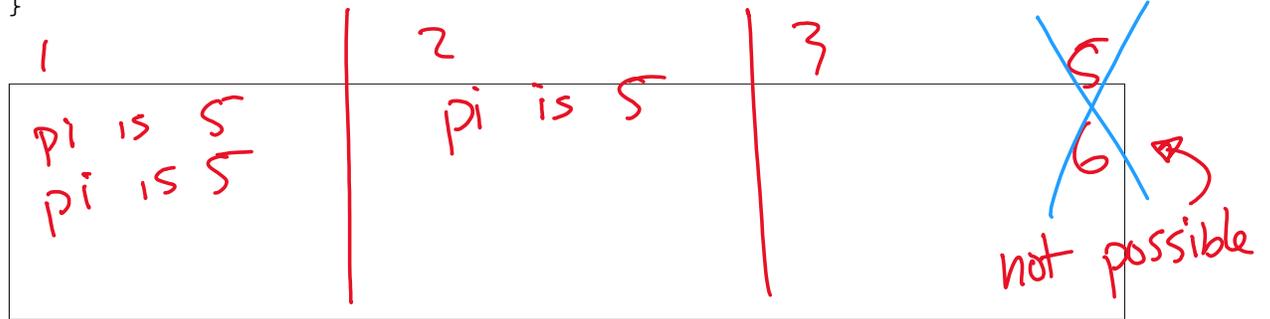
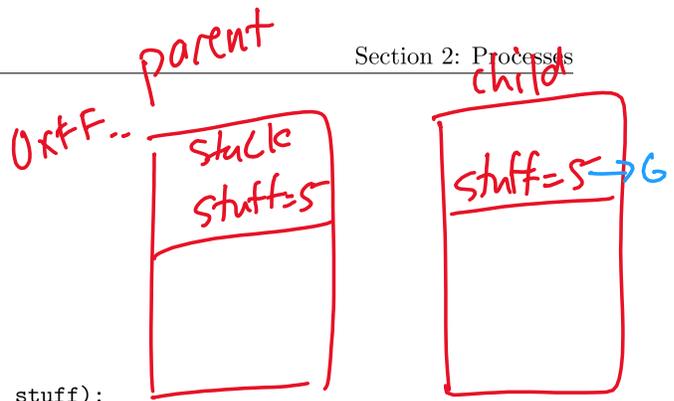
```
int main(void)
{
    for (int i = 0; i < 3; i++) {
        pid_t pid = fork();
    }
}
```



3.2 Stack Allocation

What can C print?

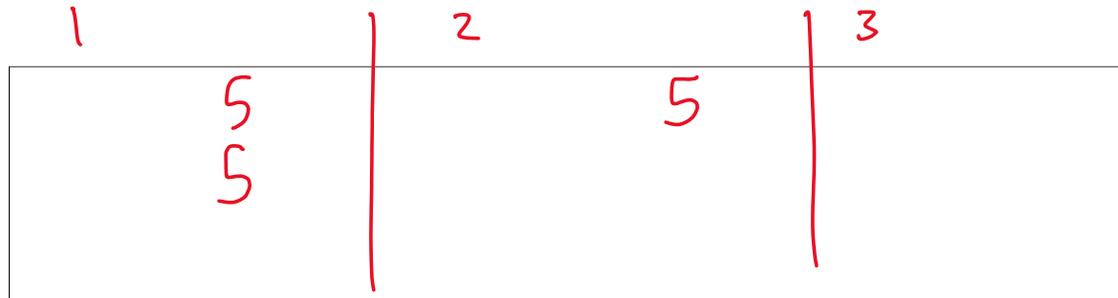
```
int main(void)
{
    int stuff = 5;
    pid_t pid = fork();
    printf("The last digit of pi is %d\n", stuff);
    if (pid == 0)
        stuff = 6;
}
```



3.3 Heap Allocation

What can C print?

```
int main(void)
{
    int* stuff = malloc(sizeof(int)*1);
    *stuff = 5;
    pid_t pid = fork();
    printf("The last digit of pi is %d\n", *stuff);
    if (pid == 0)
        *stuff = 6;
}
```



3.4 Slightly More Complex Heap Allocation

What does C print in this case ?

```
void printTenNumbers(int *arr)
{
    int i;
    printf("\n");
    for(i=0; i<10; i++) {
        printf("%d",arr[i]);
    }
    exit(0);
}

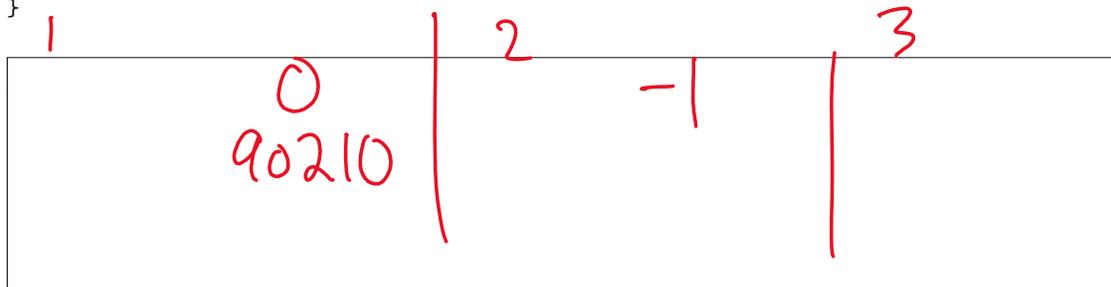
int main()
{
    int *arr, i;
    arr = (int *) malloc (sizeof(int));

    arr[0] = 0;
    for(i=1; i<10; i++) {
        arr = (int *) realloc( arr, (i+1) * sizeof(int));
        arr[i] = i;
        if (i == 7) {
            pid_t pid = fork();
            if (pid == 0) {
                printTenNumbers(arr);
            }
        }
    }
    printTenNumbers(arr);
}
```

3.5 Simple Wait

What can C print? Assume the child PID is 90210.

```
int main(void)
{
    pid_t pid = fork();
    int exit;
    if (pid != 0) {
        wait(&exit);
    }
    printf("Hello World\n: %d\n", pid);
}
```



What is the equivalent program using the `waitpid` function instead of `wait`?

$\text{wait}(\&\text{exit}) == \text{waitpid}(-1, \&\text{exit}, 0)$

↑
FLAG1 | FLAG2

3.6 Fork and File Descriptors

What will be stored in the file "output.txt" ?

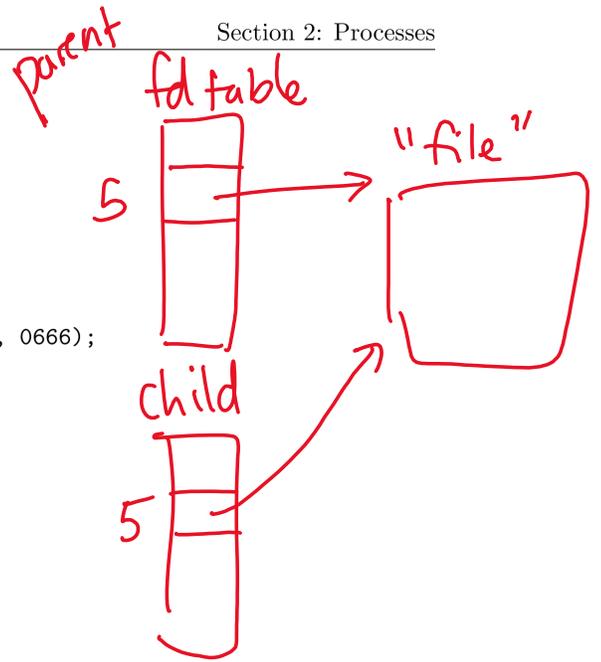
```
int main(void)
{
    int fd;

    fd = open("output.txt", O_CREAT|O_TRUNC|O_WRONLY, 0666);

    if(!fork()) {
        write(fd, "hello ", 6);
        exit(0);
    } else {
        int status;

        wait(&status);
        write(fd, "world\n", 6);
    }
}
```

fd=5



hello world → output.txt

world fail

3.7 Exec

What will C print?

```
int main(void)
{
    char** argv = (char**) malloc(3*sizeof(char*));
    argv[0] = "/bin/ls";
    argv[1] = ".";
    argv[2] = NULL;
    for (int i = 0; i < 10; i++) {
        printf("%d\n", i);
        if (i == 3)
            execv("/bin/ls", argv);
    }
}
```

*1
2
3
< output of ls >*

*} → if (i==3)
pid_t pid = fork ();
if (pid==0)
execv(...)*

*ls output {
1
2
3
4
5
6
7
8*

3.8 Exec + Fork

How would I modify the above program using fork so it both prints the output of `ls` and all the numbers from 0 to 9 (order does not matter)? You may not remove or reorder lines from the original program; only add statements (and use fork!).

3.9 Implementing fork() efficiently (Design)

Remember fork() makes the child process's address space exactly the same as its parent's. If you were designing an OS, list some steps you would take to make this address space copy more efficient ?