

Operating Systems

Grado en Informática. Course 2019-2020

Lab assignment 2: Memory

CONTINUE to code the shell started in previous lab assignments. The goal of this assignment is to understand how the memory of a process is organized. At this stage of development the shell will be able to map files into memory and to allocate memory to itself (either shared or private). In next assignments we will deal with creating processes and executing programs.

The shell will keep track (using a list) of all the memory blocks it allocates using the commands *asignar* and it deallocates with *desasignar*. For each block of memory allocated with those commands the shell must store its memory address (the address of the block), its size (the size of the block), the time it was allocated, the type of allocation (malloc, shared memory, mapped file), and other pieces of information depending on the type of allocation (name of file for mapped files, key for shared memory ...)

See the notes on the implementation of the list.

Values in that list must be coherent with what pmap shows for the shell process

In addition to the commands done in previous lab assignments, the shell has to implement the following commands

asignar allocates memory in the shell, and keeps the address of the block allocated in a list. If no arguments are given, it prints a list of the allocated memory blocks (that is, prints the list)

asignar -malloc [tam] The shell allocates *tam* bytes using *malloc* and shows the memory address returned by *malloc*. This address, together with *tam* and the time of the allocation, must be kept in the aforementioned list. If *tam* is not specified the command will show the list of addresses allocated with the **asignar -malloc** command.**malloc() requires an argument of size_t** Example:

```
-> asignar -malloc 100000000
allocated 100000000 at 0x7f6649f24010
-> asignar -malloc
0x7f6649f24010: size:100000000. malloc Mon Oct 26 20:07:05 2018
-> asignar -malloc 500000000
allocated 500000000 at 0x7f662c24d010
-> asignar -malloc
0x7f6649f24010: size:100000000. malloc Mon Oct 26 20:07:05 2018
```

```
0x7f662c24d010: size:500000000. malloc Mon Oct 26 20:08:17 2018
->
```

asignar -mmap fich [perm] Maps in memory the file *fich* (all of its length starting at offset 0) and shows the memory address where the file has been mapped. *perm* represents the mapping permissions (*rw*x format, without spaces). The address of the mapping, together with the size, the name of the file, the file descriptor, and the time of the mapping will be stored in the aforementioned list. If *fich* is not specified, the command will show the list of addresses allocated with the **asignar -mmap** command. Example

```
->asignar -mmap Shell.c rw
file Shell.c mapped at 0x7f6650436000
-> asignar -mmap Shell.c rwx
file Shell.c mapped at 0x7f6650424000
-> asignar -mmap
0x7f6650436000: size:71806. mmap Shell.c (fd:3) Mon Oct 26 20:10:07 2018
0x7f6650424000: size:71806. mmap Shell.c (fd:5) Mon Oct 26 20:10:17 2018
```

asignar -crearshared [cl] [tam] Gets shared memory of key *cl*, maps it in the process address space and shows the memory address where the shared memory has been mapped. That address, together with the key, the size of the shared memory block and the time of the mapping, will be stored in the aforementioned list. *cl* **IS THE KEY**, (*ftok* should not be used). It is assumed that key *cl* is not in use in the system so a new block of shared memory SHOULD BE CREATED, and an error must be reported if a block of key *cl* already exists. If either *cl* or *tam* are not specified, the command will show the list of addresses (and size, and time ...) allocated with the **asignar -createshared** and **asignar -shared** commands. **shared memory blocks can be of size_t size)**

asignar -shared] [cl] Gets shared memory of key *cl*, maps it in the process address space and shows the memory address where the shared memory has been mapped. That address, together with the key, the size of the shared memory block and the time of the mapping, will be stored in the aforementioned list. *cl* **IS THE KEY**, (*ftok* should not be used). It is assumed that key *cl* is already in use in the system so a new block of shared memory MUST NOT BE CREATED, and an error must be reported if a block of key *cl* does not exist. If *cl* is not specified, the command will show the list of addresses (and size, and time ...) allocated with the **asignar -createshared** and **asignar -shared** commands.

```

-> asignar -shared 15
Cannot allocate: No such file or directory
-> asignar -createshared 15 300000000
Allocated shared memory (key 15) at 0x7f661a432000
-> asignar -createshared 15 200000000
Cannot allocate: File exists
-> asignar -shared 15
Allocated shared memory (key 15) at 0x7f6608617000
-> asignar -shared
0x7f661a432000: size:300000000. shared memory (key 15) Mon Oct 26 20:12:4
0x7f6608617000: size:300000000. shared memory (key 15) Mon Oct 26 20:12:5
-> asignar
0x7f6649f24010: size:100000000. malloc Mon Oct 26 20:07:05 2018
0x7f662c24d010: size:500000000. malloc Mon Oct 26 20:08:17 2018
0x7f6650436000: size:71806. mmap Shell.c (fd:3) Mon Oct 26 20:10:07 2018
0x7f6650424000: size:71806. mmap Shell.c (fd:5) Mon Oct 26 20:10:17 2018
0x7f661a432000: size:300000000. shared memory (key 15) Mon Oct 26 20:12:4
0x7f6608617000: size:300000000. shared memory (key 15) Mon Oct 26 20:12:5

```

desasignar deallocates one of the memory blocks allocated with the command *allocate* and removes it from the list. If no arguments are given, it prints a list of the allocated memory blocks (that is, prints the list)

desasignar -malloc [tam] The shell deallocates one of the blocks of size *tam* that has been allocated with the command *asignar -malloc*. If no such block exists or if *tam* is not specified, the command will show the list of addresses allocated with the **asignar -malloc** command. Should there be more than one block of size *tam* it deallocates ONLY one of them (any). Example:

```

-> desasignar -malloc 100000000
deallocated 100000000 at 0x7f6649f24010
- desasignar -malloc
0x7f662c24d010: size:500000000. malloc Mon Oct 26 20:08:17 2015
->

```

desasignar -mmap fich Unmaps and closes the file *fich* and removes the address where it was mapped from the list. If *fich* has been mapped several times, only one of the mappings will be undone. If the file *fich* is not mapped by the process or if *fich* is not specified, the command will show the list of addresses (and size, and time ...) allocated with the **asignar -mmap** command.

desasignar [-shared] [cl] Detaches the shared memory block with key *cl* from the process' address space and eliminates its address from the list. If shared memory block with key *cl* has been attached several

times, ONLY one of them is detached. *cl* **IS THE KEY**, *ftok* should not be used. If *cl* is not specified, the command will show the list of addresses (and size, and time ...) allocated with the **asignar -createshared** and **asignar -shared** commands.

desasignar addr Deallocates *addr* (it searches in the list how it was allocated, and proceeds accordingly) and removes it from the list. If *addr* is not in the list or if *addr* is not supplied the command will show all the addresses (and size, and time ...) allocated with the **asignar -malloc**, **asignar -mmap**, **asignar -sharednew** and **asignar shared** commands. This is equivalent to *desasignar -malloc*, *desasignar -shared* or *desasignar -mmap* depending on *addr*

```
-> desasignar
0x7f6649f24010: size:1000000000. malloc   Mon Oct 26 20:07:05 2018
0x7f662c24d010: size:5000000000. malloc   Mon Oct 26 20:08:17 2018
0x7f6650436000: size:71806. mmap Shell.c (fd:3) Mon Oct 26 20:10:07 2018
0x7f6650424000: size:71806. mmap Shell.c (fd:5) Mon Oct 26 20:10:17 2018
0x7f661a432000: size:3000000000. shared memory (key 15) Mon Oct 26 20:12:4
0x7f6608617000: size:3000000000. shared memory (key 15) Mon Oct 26 20:12:5
-> desasignar -malloc 1000000000
block at address 0x7f6649f24010 deallocated (malloc)
-> desasignar -shared 15
block at address 0x7f661a432000 deallocated (shared)
-> desasignar -mmap Shell.c
block at address 0x7f6650436000 deallocated (mmap)
-> desasignar 0x7f662c24d010
block at address 0x7f662c24d010 deallocated (malloc)
-> desasignar 0x7f662c24d010
0x7f6650424000: size:71806. mmap Shell.c (fd:5) Mon Oct 26 20:10:17 2018
0x7f6608617000: size:3000000000. shared memory (key 15) Mon Oct 26 20:12:5
```

borrarkey cl Removes the shared memory region of key *cl*. NOTHING GETS UNMAPPED: this is just a call to *shmctl(id, IPC_RMID...)*

mem [-malloc] [-shared] [-mmap] [-all] Shows addresses inside the process memory space. If no arguments are given, it prints the memory addresses of three program functions, three extern (global) variables and three automatic (local) variables.

mem -malloc shows the list of addresses (and size, and time ...) allocated with the **asignar -malloc**

mem -shared shows the list of addresses (and size, and time ...) allocated with the **asignar -createshared** and **asignar -shared** commands.

mem -mmap shows the list of addresses (and size, and time ...) allocated with the **asignar -mmap**

mem -all shows the list of addresses (and size, and time ...) allocated with the **asignar -malloc**, **asignar -mmap**, **asignar -createshared** and **asignar -shared** commands.

volcar *addr* [*cont*] Shows the contents of *cont* bytes starting at memory address *addr*. If *cont* is not specified, it shows 25 bytes. For each byte it prints (at different lines) its hex value and its associate char (a blank if it is a non-printable character). It prints 25 bytes per line. **addr SHOULD NOT BE CHECKED FOR VALIDITY**, so, this command could produce segmentation fault should *addr* were not valid.

```
->volcar 0xb8019000 300
```

```
# i n c l u d e < u n i s t d . h > # i n c l
23 69 6E 63 6C 75 64 65 20 3C 75 6E 69 73 74 64 2E 68 3E 0A 23 69 6E 63 6C
u d e < s t d i o . h > # i n c l u d e < s
75 64 65 20 3C 73 74 64 69 6F 2E 68 3E 0A 23 69 6E 63 6C 75 64 65 20 3C 73
t r i n g . h > # i n c l u d e < s t d l i b
74 72 69 6E 67 2E 68 3E 0A 23 69 6E 63 6C 75 64 65 20 3C 73 74 64 6C 69 62
. h > # i n c l u d e < s y s / t y p e s . h
2E 68 3E 0A 23 69 6E 63 6C 75 64 65 20 3C 73 79 73 2F 74 79 70 65 73 2E 68
> # i n c l u d e < s y s / s t a t . h > #
3E 0A 23 69 6E 63 6C 75 64 65 20 3C 73 79 73 2F 73 74 61 74 2E 68 3E 0A 23
```

llenar *addr* [*cont*] [*byte*] Fills *cont* bytes of memory starting at address *addr* with the value '*byte*'. If '*byte*' is not specified, the value of 65 (0x42 or capital A) is assumed, and if *cont* is not specified, we'll use a default value of 128. **addr SHOULD NOT BE CHECKED FOR VALIDITY**, so, this command could produce segmentation fault should *addr* were not valid.

Example:

```
-> llenar 0xb8019000 1500 0x42
```

recursiva n. Calls a recursive function passing the integer *n* as its parameter. This recursive function receives the number of times it has to call itself. This function has two variables: an automatic array of 2048 bytes and a static array of the same size. It does the following

- prints the value of the received parameter as well as its memory address.
- prints the address of the static array.
- prints the address of the automatic array.
- decrements *n* (its parameter) and if *n* is greater than 0 it calls itself.

A possible coding for this function:

```
void doRecursiva (int n)
{
    char automatico[TAMANO];
    static char estatico[TAMANO];

    printf ("parametro n:%d en %p\n",n,&n);
    printf ("array estatico en:%p \n",estatico);
    printf ("array automatico en %p\n",automatico);

    n--;
    if (n>0)
        recursiva(n);
}
```

rfich fich addr cont Reads (using ONE *read* system call) *cont* bytes from file *fich* into memory address *addr*. If *cont* is not specified ALL of *fich* is read onto memory address *addr*. Depending on the value of *addr* a segmentation fault could be produced.

wfich [-o] fich addr cont Writes (using ONE *write* system call) *cont* bytes from memory address *addr* into file *fich*. If file *fich* does not exist it gets created; if it already exists it is not overwritten unless “-o” (overwrite) is specified.

NOTES ON LIST IMPLEMETATION

- the implementations of list should consist of the data types and the access funtions. A `list.h` and a `list.c` should be created. `list.h` should contain all the types definition neccessary for the list as well as the functions prototypes. `list.c` contains the implementation of the list functions. `list.c` must be compiled separately and a sentence to include `list.h` should be present in the main shell program.
- four list implementations are to be considered:
 - 0) **linked list**: The list is composed of dynamically allocated nodes. Each node has some item of information and a pointer to the following node. The list itself is a pointer to the first node, when the list is empty this pointer is NULL, so creating the list is asigning NULL to the list pointer, thus the functions **CreateList**, **InsertElement** and **RemoveElement** must receive the list by reference as they may have (case of inserting or removing the first element) to modify the list.
 - 1) **linked list with head node**: Similar to the linked list except

that the list itself is a pointer to a *empty* (with no information) first node. Creating the list is allocating this first element (head node). **CreateList** must receive the list by reference whereas **InsertElement** and **RemoveElement** can receive the list by value.

- 2) **array:** Elements in the list are stored in a statically allocated array of nodes, so the list type is a pointer to a structure containing the array of nodes and optionally one or more integers (depending on the implementation: `nextin` and `nextout` indexes, counter ...). For the purpose of this lab assignment we can assume the array dimension to be 4096 (which should be declared a named constant, and thus easily modifiable).
- 3) **array of pointers;** Elements in the list are allocated dynamically, and the list keeps track of all its elements using an array of pointers, (this can be either a NULL terminated array or we can use additional integers). This array is statically allocated and for the purpose of this lab assignment we can assume the array dimension to be 4096 which should be declared a named constant, and thus easily modifiable.

- **EACH WORKGROUP MUST DO THE LIST IMPLEMENTATION RESULTING FROM THE FOLLOWING PROCEDURE:** Add the two last digits of the D.N.I. of the workgroup components and calculate its module 4, that will yield the implementation to use. Should one (or more) component of the workgroup have no D.N.I., then the ascii code of its capitalized family name initial should be used instead.

- **example 1:** workgroup components D.N.I.s are 55555581 and 55555507, so the implementation to use will be given by $(81 + 07) \% 4$, so this group will have to use implementation 0, `linked list`
- **example 2:** workgroup componets are D.N.I. 55555581 and Donald Trump (who, as of now, does not have a valid D.N.I., to the best of our knowledge) so, as the ascii code for the **T** is 84, this group would have to use implementation $(81 + 84) \% 4 = 1$. `linked list with head node`

- This program should compile cleanly (produce no warnings even when compiling with `gcc -Wall`)
- **NO RUNTIME ERROR WILL BE ALLOWED** (**segmentation, bus error ...**), unless where explicitly spcified. Programs with runtime errors will yield no score.
- This program can have no memory leaks (memory blocks allocated with

the **asignar -malloc** command are not taken into account)

- When the program cannot perform its task (for whatever reason, for example, lack of privileges) it should inform the user
- All input and output is done through the standard input and output

Information on the system calls and library functions needed to code this program is available through `man`: (*open, read, write, close, shmget, shmat, shmdt, shmctl, mmap, munmap, malloc, free...*).

WORK SUBMISSION

- Work must be done in pairs.
- The source code will be submitted to the subversion repository under a directory named **P2**
- The name of the main program will be **shell.c**, the list implementation will use files *list.h* and *list.c*. A **Makefile** must be supplied so that the program (and all of its modules) can be compiled with just **make**
- Only one of the members of the workgroup will submit the source code. The names, logins and **D.N.I.** of all the members of the group should be in the source code of the main program (at the top of the file)

DEADLINE: NOVEMBER FRIDAY 22ND, 2019, 23:00H

ASSESSMENT: FOR EACH PAIR, IT WILL BE DONE IN ITS CORRESPONDING GROUP, DURING THE LAB CLASSES

CLUES

The following functions could guide the implementation of *asignar -shared*, *asignar -createshared*, *asignar -mmap*, *rfich* and *wfich*. It is assumed that

- The functions to perform the tasks '*Guardar En Direcciones de Memoria Shared*'. '*Listar Direcciones de Memoria Shared*'... must be implemented as part of the assignment
- *Cmd_AsignarCreateShared* and *Cmd_AsignarMmap* receive as parameter the array resulting from splitting the input string discarding the first element (they are called passing **arg+1** as parameter), and may require additional parameters.

```
/*  
/  
/  
void * ObtenerMemoriaShmget (key_t clave, size_t tam)  
{
```



```

void * p;
int aux,id,flags=0777;
struct shmid_ds s;

if (tam) /*si tam no es 0 la crea en modo exclusivo */
    flags=flags | IPC_CREAT | IPC_EXCL;
    /*si tam es 0 intenta acceder a una ya creada*/
if (clave==IPC_PRIVATE) /*no nos vale*/
    {errno=EINVAL; return NULL;}

if ((id=shmget(clave, tam, flags))== -1)
    return (NULL);

if ((p=shmat(id,NULL,0))==(void*) -1){
    aux=errno; /*si se ha creado y no se puede mapear*/
    if (tam) /*se borra */
        shmctl(id,IPC_RMID,NULL);
    errno=aux;
    return (NULL);
}
shmctl (id,IPC_STAT,&s);

/* Guardar En Direcciones de Memoria Shared (p, s.shm_segsz, clave.....);*/
return (p);
}

void Cmd_AsignarCreateShared (char *arg[])
{
    key_t k;
    size_t tam=0;
    void *p;

    if (arg[0]==NULL || arg[1]==NULL)
        {/*Listar Direcciones de Memoria Shared */;/* return;}

    k=(key_t) atoi(arg[0]);

    if (arg[1]!=NULL)
        tam=(size_t) atoll(arg[1]);

    if ((p=ObtenerMemoriaShmget(k,tam))==NULL)
        perror ("Imposible obtener memoria shmget");
    else

```

```

        printf ("Memoria de shmget de clave %d asignada en %p\n",k,p);
    }

```

```

/*****
*****/

```

```

void * MmapFichero (char * fichero, int protection)
{
    int df, map=MAP_PRIVATE,modo=O_RDONLY;
    struct stat s;
    void *p;

    if (protection&PROT_WRITE) modo=O_RDWR;

    if (stat(fichero,&s)==-1 || (df=open(fichero, modo))== -1)
        return NULL;

    if ((p=mmap (NULL,s.st_size, protection,map,df,0))==MAP_FAILED)
        return NULL;

    /*Guardar Direccion de Mmap (p, s.st_size,fichero,df.....);*/
    return p;
}

```

```

void Cmd_AsignarMmap (char *arg[])
{
    char *perm;
    void *p;
    int protection=0;

    if (arg[0]==NULL)
        {/*Listar Direcciones de Memoria mmap;*/ return;}

    if ((perm=arg[1])!=NULL && strlen(perm)<4) {
        if (strchr(perm,'r')!=NULL) protection|=PROT_READ;
        if (strchr(perm,'w')!=NULL) protection|=PROT_WRITE;
        if (strchr(perm,'x')!=NULL) protection|=PROT_EXEC;
    }
    if ((p=MmapFichero(arg[0],protection))==NULL)
        perror ("Imposible mapear fichero");
}

```

```

        else
            printf ("fichero %s mapeado en %p\n", arg[0], p);
    }

#define LEERCOMPLETO ((ssize_t)-1)
ssize_t LeerFichero (char *fich, void *p, ssize_t n) /*n=-1 indica que se lea todo*/
{
    ssize_t  nleidos,tam=n;
    int df, aux;
    struct stat s;

    if (stat (fich,&s)==-1 || (df=open(fich,O_RDONLY))==-1)
        return ((ssize_t)-1);

    if (n==LEERCOMPLETO)
        tam=(ssize_t) s.st_size;

    if ((nleidos=read(df,p, tam))==-1){
        aux=errno;
        close(df);
        errno=aux;
        return ((ssize_t)-1);
    }
    close (df);

    return (nleidos);
}

/*****/
/*****/
void Cmd_borrakey (char *args[])
{
    key_t clave;
    int id;
    char *key=args[0];

    if (key==NULL || (clave=(key_t) strtoul(key,NULL,10))==IPC_PRIVATE){
        printf ("rmkey clave_valida\n");
        return;
    }
    if ((id=shmget(clave,0,0666))==-1){
        perror ("shmget: imposible obtener memoria compartida");
    }
}

```

```
        return;
    }
    if (shmctl(id,IPC_RMID,NULL)==-1)
        perror ("shmctl: imposible eliminar memoria compartida\n");
}
```