

# Memory mapping

# [ DEMO Time ]

---


- Hacking the virtual memory system of Linux
  - A journey through the page directories and tables of Linux kernel 32bits

By Renato Mancuso



# [ Two level page table hierarchy ]

virtual address (32 bits) → addresses a byte in VM

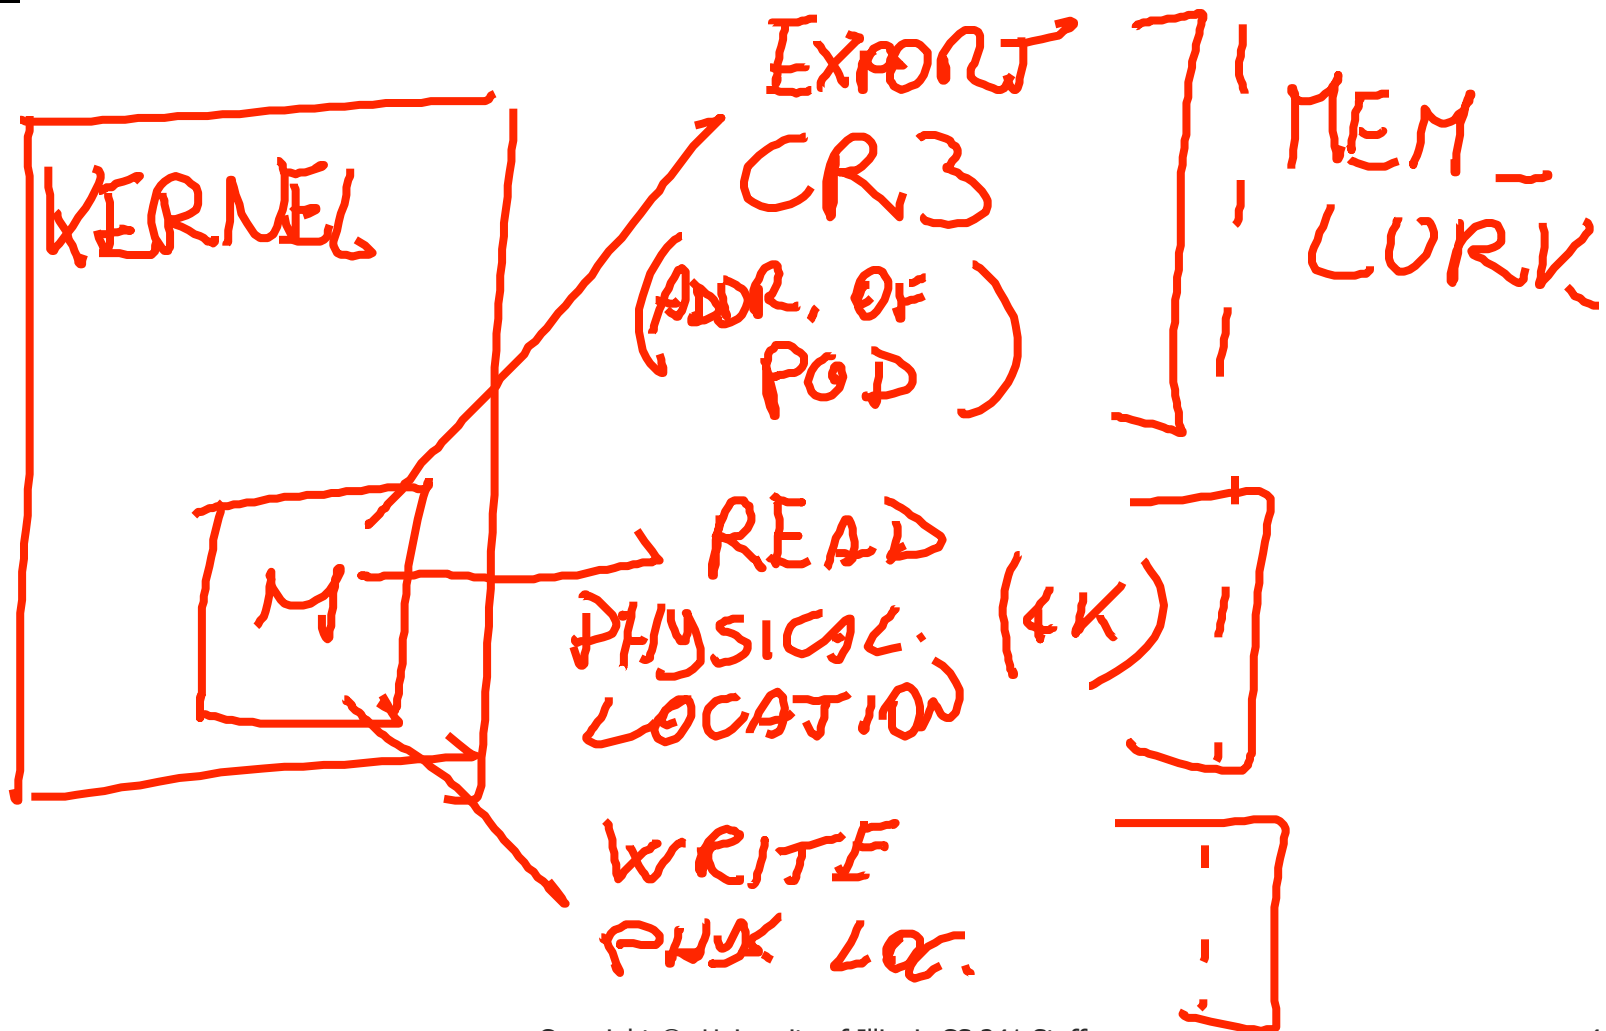


addresses a byte in Physical Mem.

\*32 bits aligned onto a 4-KByte boundary.



# OVERALL STRUCTURE



# [ TRANSITION ]

TARGET: 0x0804A01C

RD: 7840 OFFSET

CR3: 0x25E13000

0x0804A01C

↳ 0x20 OFFSET → 0x80



# [ PJ ADDRESS ]

PJ:  $0x B1843067$   
                                 
          ADDR.        FLAGS

PJ OFFSET:  $0x 04A \xrightarrow{\times 4} 0x 128$

PAGE ADDR.:  $0x 83B20067$   
    
                                  FLAGS

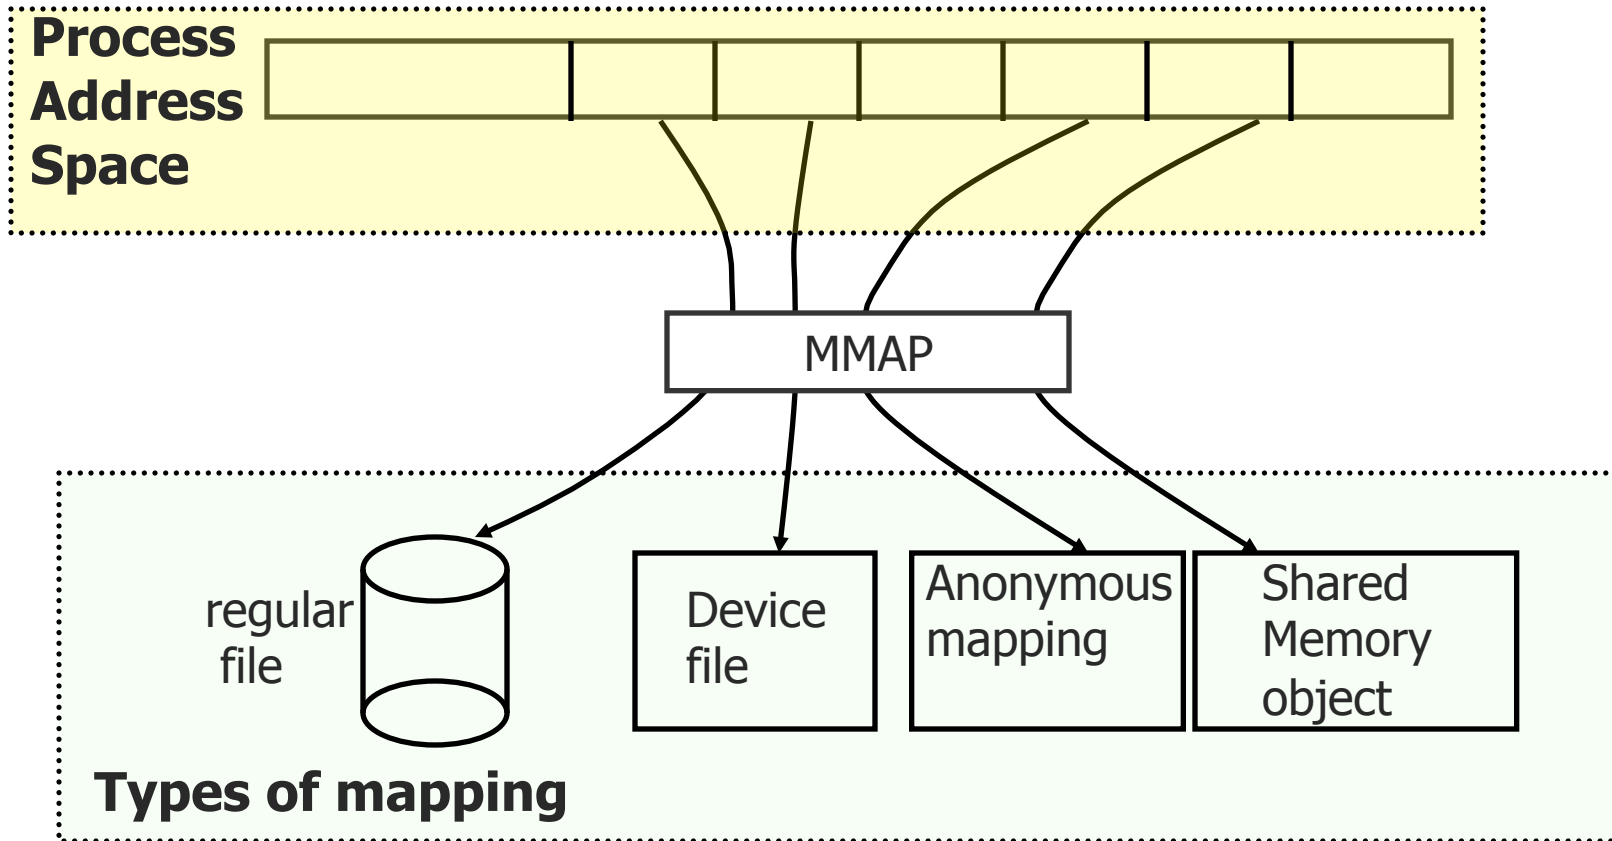


# [ Concept of memory mapping ]

- If the virtual memory sub-system is integrated with the file-system, it enables a simple and efficient mechanism to load programs and data into memory
- If disk I/O requires the transfer of **large amounts of data (one or more pages)**, mmap significantly speeds up I/O by mapping a disk file directly into user-space memory
  - It does not suffer the overhead of syscalls like read/write
  - User-process has direct access to kernel disk cache



# MMAP: a powerful syscall



- MMAP is used for mapping differing sorts of objects





# MMAP

memory address of the mapping

read/write/exec permissions for the mapping


MAP\_SHARED, MAP\_PRIVATE, MAP\_ANONYMOUS

```
#include <sys/mman.h>
```

```
void * mmap (void *addr, size_t len, int prot, int flags, int fd, off_t offset);
```

file descriptor of object to map


**Mmap** (addr, len, prot, flags, fd, offset)



*Backing Store*

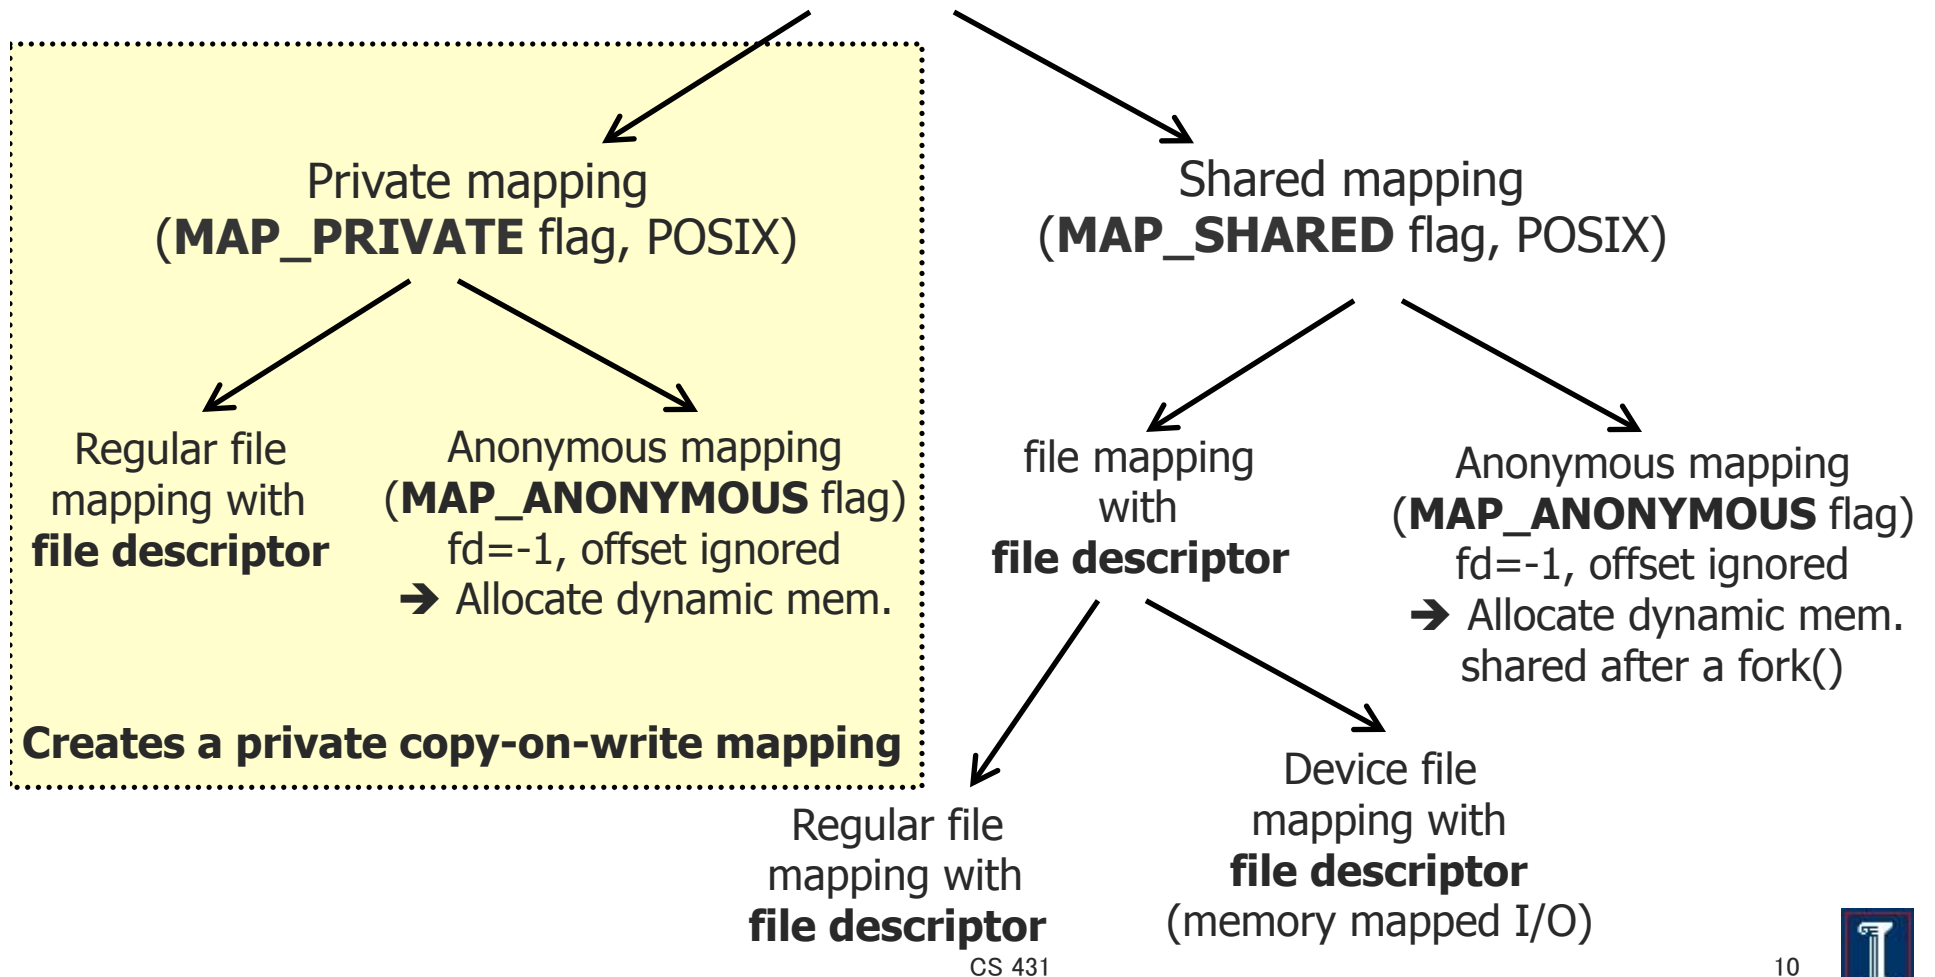
**Memory**

len



# [Types of mapping with MMAP]

## MMAP



# [ Private vs Shared mapping ]

## ■ MAP\_PRIVATE

- Updates to the mapping are not visible to other processes mapping the same file, and are not carried through to the underlying file.

## ■ MAP\_SHARED

- Updates to the mapping are visible to other processes that map a shared file, and are carried through to the underlying file. The file may not actually be updated until **msync** or **munmap()** is called
- Further discussion later in the semester when covering inter-process communication (IPC)



# [ Portable use of MMAP ]

- Set `addr = NULL`
  - the kernel chooses the address at which to create the mapping
  - MMAP always returns a page aligned address
- Virtual memory always allocates entire pages
  - `offset` must be a multiple of the page size
  - `len` must be a multiple of the page size



# [ MUNMAP ]

```
int munmap(void *addr, size_t length);
```

- **munmap()** system call
  - It deletes the mappings for the specified address range.
  - It can unmap a smaller number of pages among those allocated by mmap (partial unmapping)
  - **addr** argument must be page aligned
  - **len** must be a multiple of the page size



# Example: private mapping of regular file

```
#include <sys/mman.h>
#include <fcntl.h>
#include <stdio.h>
#include <unistd.h>
#define PAGESIZE 4096
int main(int argc, char **argv)
{
    int fd;
    char string[] = "CS241 takeaway: mmap can be used to map files in memory";
    char *ptr;

    //open a regular file, write a string, and map it into memory
    fd = open(argv[1], O_RDWR|O_CREAT, S_IRWXU);
    write(fd, string, sizeof(string) - 1);
    ptr = (char*) mmap(NULL, PAGESIZE, PROT_READ|PROT_WRITE, MAP_PRIVATE, fd, 0);
    close(fd);

    printf("pointer to memory mapped file: %p \n", ptr);
    printf("%s \n", ptr);
    ptr[2]='4'; ptr[3]='3'; ptr[4]='1'; // it triggers a copy-on-write
    printf("%s \n", ptr);
}
```



# Example: shared mapping of device file

```
#include <sys/mman.h>
#include <fcntl.h>
#include <stdio.h>
#include <unistd.h>
#define PAGESIZE 4096
int main(int argc, char **argv)
{
    int fd;
    char string[]="CS241 takeaway:mmap can be used to map dev. files in memory";
    char *ptr;

    //open a raw UNUSED disk partition, write a string, and map it into memory
    fd = open("/dev/sda3", O_RDWR, S_IRWXU);
    write(fd, string, sizeof(string) - 1);
    ptr = (char*) mmap(NULL, PAGESIZE, PROT_READ|PROT_WRITE, MAP_SHARED, fd, 0);
    close(fd);

    printf("pointer to memory mapped file: %p \n", ptr);
    printf("%s \n", ptr);
    ptr[2]='4'; ptr[3]='3'; ptr[4]='1'; // it modifies the disk partition cache
    printf("%s \n", ptr);
}
```

# Example: shared mapping of device file

# copy the first 256bytes of disk partition /dev/sda3 to sda3.bin and check its content!

```
sudo dd bs=256 count=1 if=/dev/sda3 of=./sda3.bin
```

```
ghex sda3.bin
```

The screenshot shows the GHex application window titled "sda3.bin - GHex". The main area displays a hex dump of the file's contents. The right-hand pane shows a text view of the data, which appears to be a FAT32 boot sector message. The bottom of the window contains various conversion and display options.

Signed 8 bit:	0	Signed 32 bit:	0	Hexadecimal:	00
Unsigned 8 bit:	0	Unsigned 32 bit:	0	Octal:	000
Signed 16 bit:	0	Float 32 bit:	0.000000e+00	Binary:	00000000
Unsigned 16 bit:	0	Float 64 bit:	0.000000e+00	Stream Length:	8

Show little endian decoding  Show unsigned and float as hexadecimal

Offset: 0xF0