

# **MMAP and Page Fault Handling**

SWE3015

Sung-hun Kim



성균관대학교  
SUNGKYUNKWAN UNIVERSITY

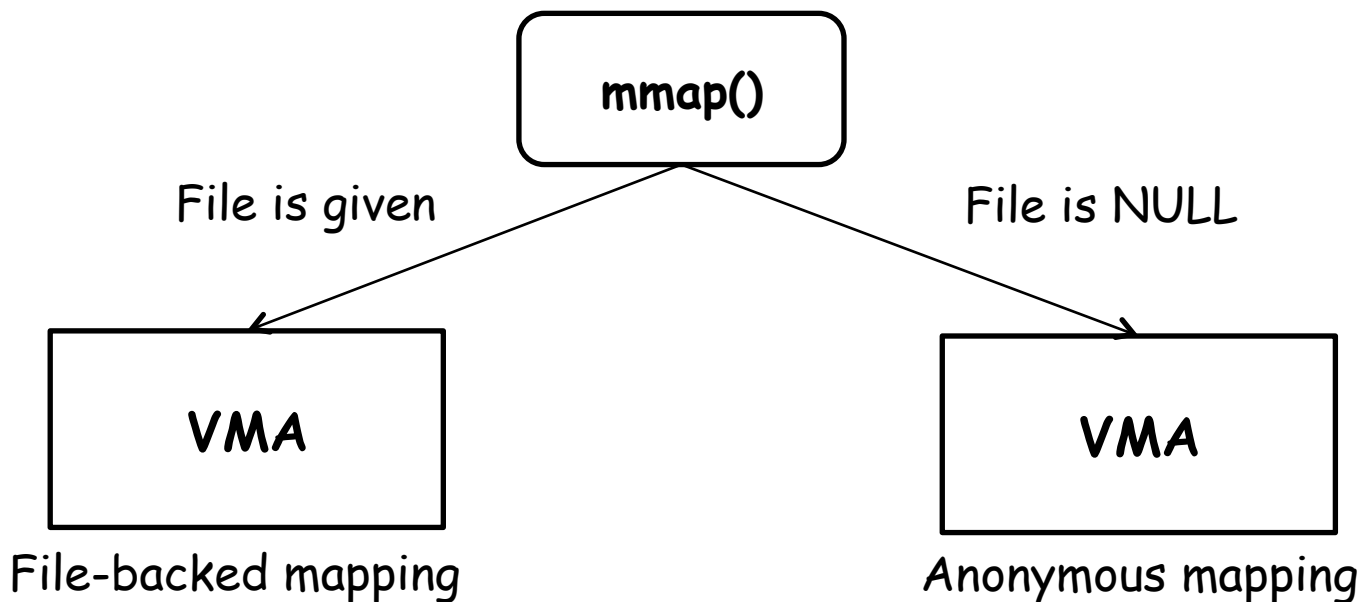
SWE3015: Operating System Project

# MMAP



- POSIX compliant function
  - Create a new mapping in the virtual address space of calling process

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





# MMAP Calling Path

mmap\_pgoff()

vm\_mmap\_pgoff()

down\_write(&mm->mmap\_sem);

do\_mmap\_pgoff()

up\_write(&mm->mmap\_sem);

mmap\_region()



# MMAP\_REGION(1/2)

- Main function for mmap call
  - Check *mmap* is calling for old mapping

```
vma = vma_merge(args);  
if (vma)  
    goto out;
```

- Make a new mapping (vma)

```
vma = kmem_cache_zalloc(vm_area_cachep, GFP_KERNEL);
```

```
vma->vm_mm = mm;  
vma->vm_start = addr;  
vma->vm_end = addr + len;  
vma->vm_flags = vm_flags;  
vma->vm_page_prot = vm_get_page_prot(vm_flags);  
vma->vm_pgoff = pgoff;  
INIT_LIST_HEAD(&vma->anon_vma_chain);
```



# MMAP\_REGION(2/2)

- Main function for mmap call (continued)
  - Call file's mmap function

```
if (file) {  
    vma->vm_file = get_file(file);  
    error = file->f_op->mmap(file, vma);  
    //...  
}
```

- You should implement mmap of file\_operations
  - And configure for newly allocated vma

```
static const struct file_operations fake_fops = {  
    .owner = THIS_MODULE,  
    .mmap = mmap_fake,  
};
```



성균관대학교  
SUNGKYUNKWAN UNIVERSITY

SWE3015: Operating System Project

# PAGE FAULT HANDLING



# Page Fault

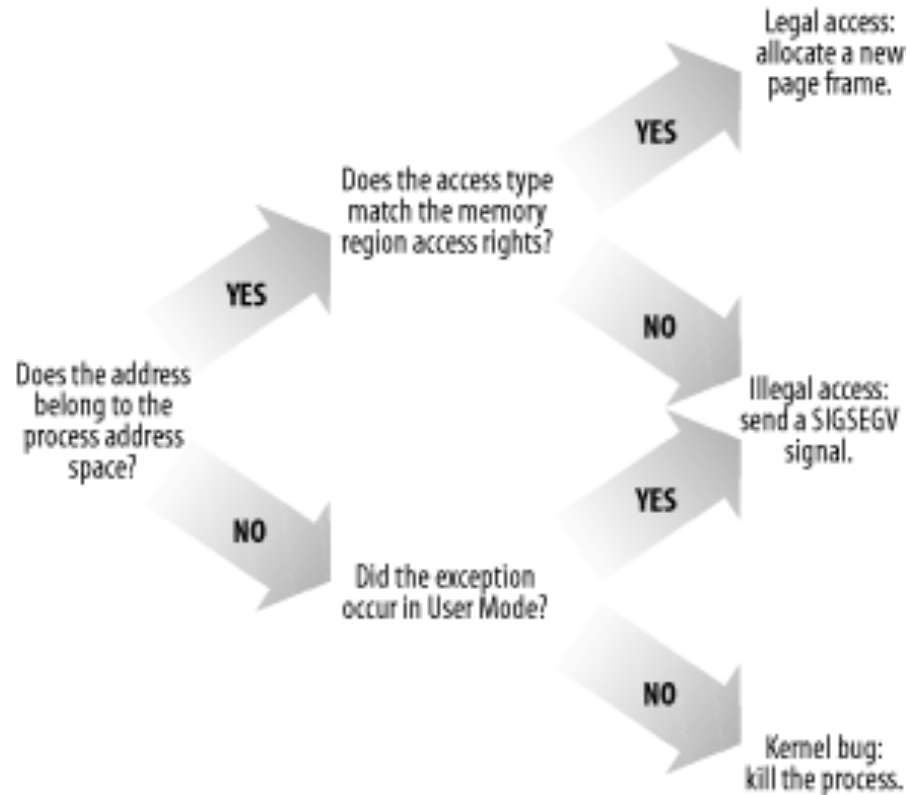
- A trap to the software raised by the hardware
  - when a program accesses a page that is mapped in the virtual address space,
  - but not loaded in physical memory
- Typical reasons are...
  - Demand paging
  - Copy on Write
  - User software bug (SIGSEGV)
  - Or even kernel bug (kernel “Oops”)





# Page Fault Handling

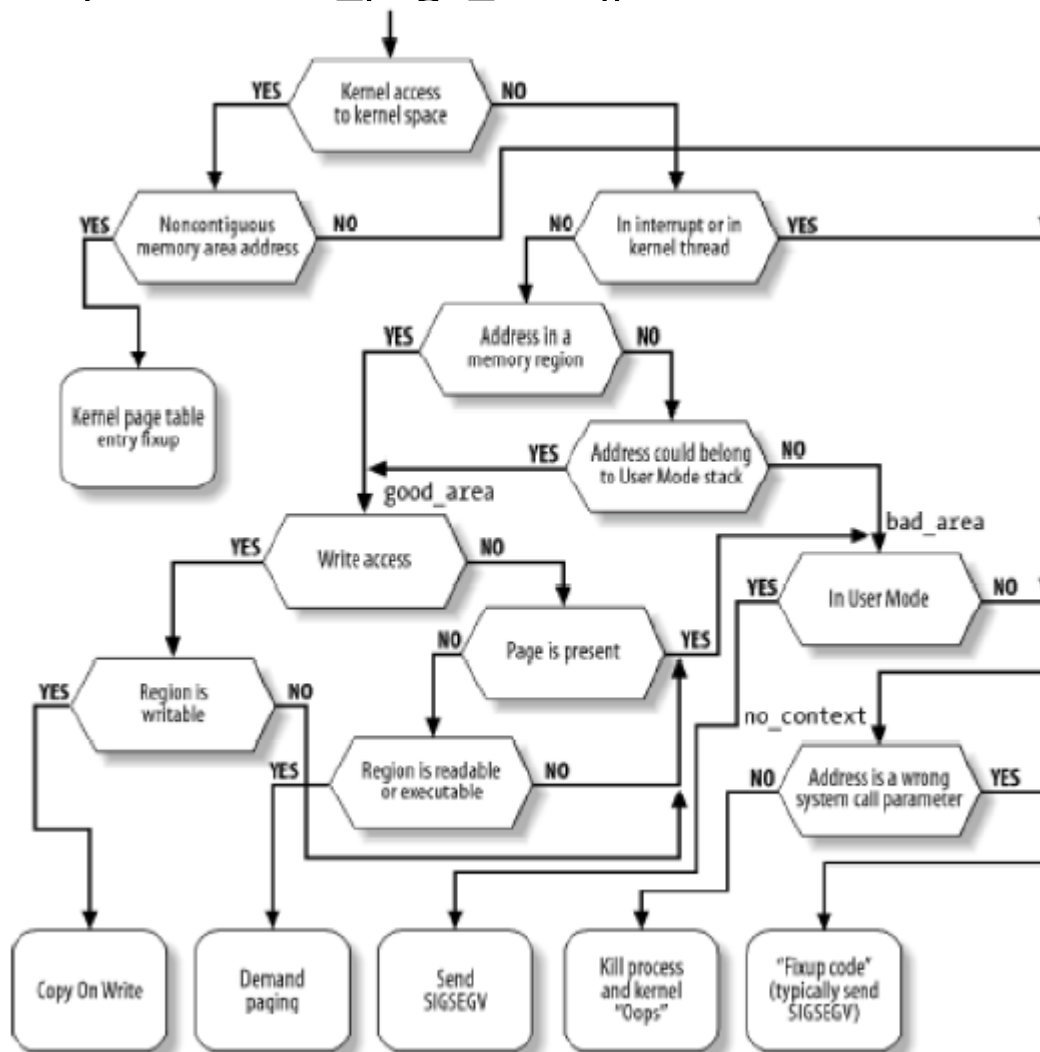
- `<arch/x86/mm/fault.c> do_page_fault()`





# Page Fault Handling

- <arch/x86/mm/fault.c> do\_page\_fault()





# Page Fault Handler Calling Path

- arch/x86/mm/fault.c

do\_page\_fault()

\_\_do\_page\_fault()

find\_vma()

handle\_mm\_fault()



# Page Fault Handler Calling Path

- mm/memory.c

handle\_mm\_fault()

\_\_handle\_mm\_fault()

pte\_offset\_map(), handle\_pte\_fault()

do\_linear\_fault()

do\_[read | cow | linear]\_fault()

\_\_do\_fault()

vma->vm\_ops->fault()



# \_\_handle\_mm\_fault

- Find pte
  - pgd\_offset(), pud\_alloc(), pmd\_alloc()
  - pte\_offset\_map() -> find location of pte in pmd
- Handle fault
  - handle\_pte\_fault(mm, vma, address, pte, pmd, flags)



# do\_[read|cow|shared]\_fault

- Call pre-define fault handler
  - `__do_fault(vma, address, pgoff, flags, &fault_page);`
- Make *pte* entry
  - Acquire lock for *pmd*
  - Make *pte*
    - `pte = pte_offset_map_lock(mm, pmd, address, &ptl);`
    - Get the location of faulted page table entry
  - Insert *pte*
    - `do_set_pte(vma, address, fault_page, pte, true, false);`



# \_\_do\_fault

- Key data structure: vm\_fault

```
struct vm_fault vmf;  
  
vmf.virtual_address = (void __user *)(address & PAGE_MASK);  
vmf.pgoff = pgoff;  
vmf.flags = flags;  
vmf.page = NULL;  
  
ret = vma->vm_ops->fault(vma, &vmf);  
  
//...  
  
*page = vmf.page;
```



성균관대학교  
SUNGKYUNKWAN UNIVERSITY

# Questions