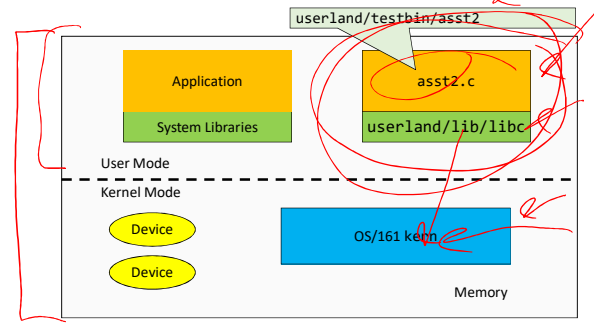


Assignment 2 tips

1

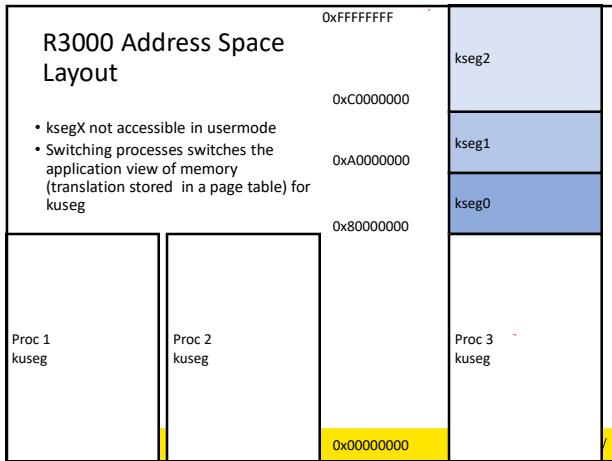
Structure of a Computer System



2

R3000 Address Space Layout

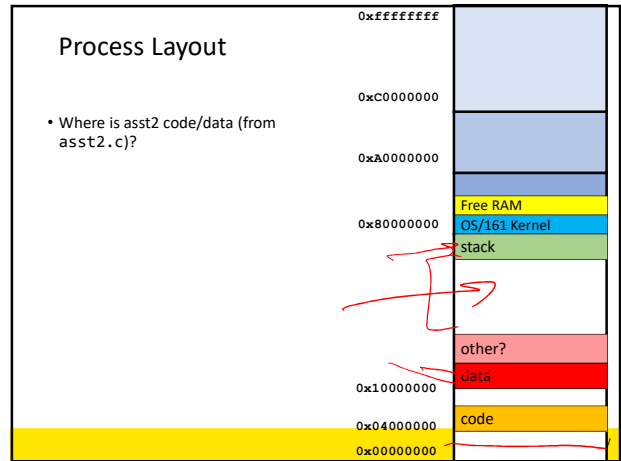
- ksegX not accessible in usermode
- Switching processes switches the application view of memory (translation stored in a page table) for kuseg



3

Process Layout

- Where is asst2 code/data (from asst2.c)?

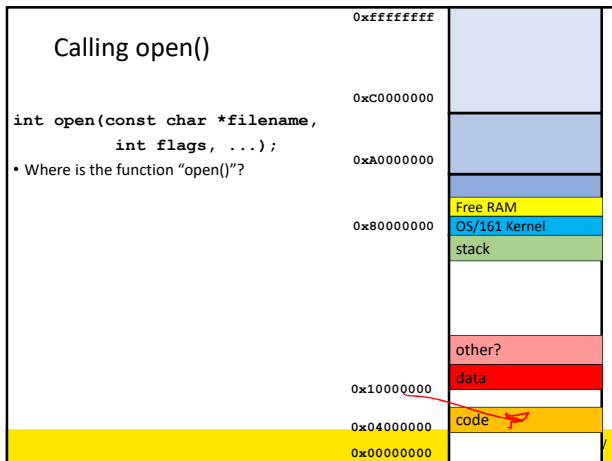


4

Calling open()

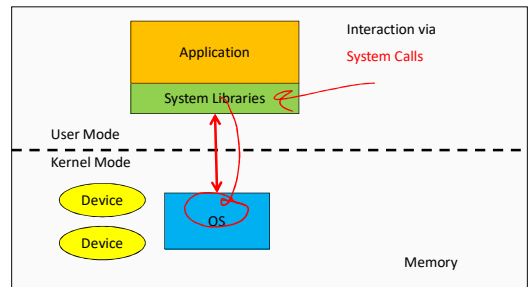
```
int open(const char *filename,
        int flags, ...);
```

- Where is the function "open()"?



5

Structure of a Computer System



6

open()? 0xffffffff

```
int open(const char *filename,
        int flags, ...);
```

- Where is "open()"s implementation?
- By convention, it's called sys_open() in the kernel.

This is what you are implementing in ASST2

0xc0000000	
0xa0000000	
0x80000000	Free RAM
0x40000000	OS/161 Kernel
0x00000000	stack
0x10000000	other?
0x04000000	data
0x00000000	code

7

Argument passing

```
#include <unistd.h>
int reboot(int code);
```

Description
reboot reboots or shuts down the system. The specific action depends on the code passed:

- RB_REBOOT** The system is rebooted.
- RB_HALT** The system is halted.
- RB_POWEROFF** The system is powered off.

Return Values
On success, reboot does not return. On error, -1 is returned, and errno is set according to the error encountered.

8

<p>Convention for kernel entry</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <tr><td>ra</td></tr><tr><td>fp</td></tr><tr><td>sp</td></tr><tr><td>gp</td></tr><tr><td>k1</td></tr><tr><td>k0</td></tr><tr><td>s7</td></tr><tr><td>...</td></tr><tr><td>s0</td></tr><tr><td>t9</td></tr><tr><td>t0</td></tr><tr><td>a3</td></tr><tr><td>a2</td></tr><tr><td>a1</td></tr><tr><td>a0</td></tr><tr><td>v1</td></tr><tr><td>v0</td></tr><tr><td>AT</td></tr><tr><td>zero</td></tr> </table>	ra	fp	sp	gp	k1	k0	s7	...	s0	t9	t0	a3	a2	a1	a0	v1	v0	AT	zero	<p>Preserved →</p> <p>Preserved for C calling convention →</p> <p>Preserved →</p> <p>Success? →</p> <p>Result →</p>	<p>Convention for kernel exit</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <tr><td>ra</td></tr><tr><td>fp</td></tr><tr><td>sp</td></tr><tr><td>gp</td></tr><tr><td>k1</td></tr><tr><td>k0</td></tr><tr><td>s7</td></tr><tr><td>...</td></tr><tr><td>s0</td></tr><tr><td>t9</td></tr><tr><td>t0</td></tr><tr><td>a3</td></tr><tr><td>a2</td></tr><tr><td>a1</td></tr><tr><td>a0</td></tr><tr><td>v1</td></tr><tr><td>v0</td></tr><tr><td>AT</td></tr><tr><td>zero</td></tr> </table>	ra	fp	sp	gp	k1	k0	s7	...	s0	t9	t0	a3	a2	a1	a0	v1	v0	AT	zero
ra																																								
fp																																								
sp																																								
gp																																								
k1																																								
k0																																								
s7																																								
...																																								
s0																																								
t9																																								
t0																																								
a3																																								
a2																																								
a1																																								
a0																																								
v1																																								
v0																																								
AT																																								
zero																																								
ra																																								
fp																																								
sp																																								
gp																																								
k1																																								
k0																																								
s7																																								
...																																								
s0																																								
t9																																								
t0																																								
a3																																								
a2																																								
a1																																								
a0																																								
v1																																								
v0																																								
AT																																								
zero																																								

SysCall No.

9

```
struct trapframe {
    u_int32_t tf_vaddr; /* vaddr register */
    u_int32_t tf_status; /* status register */
    u_int32_t tf_cause; /* cause register */
    u_int32_t tf_lo;
    u_int32_t tf_hi;
    u_int32_t tf_ra; /* Saved register 31 */
    u_int32_t tf_a0; /* Saved register 1 (a0) */
    u_int32_t tf_v0; /* Saved register 2 (v0) */
    /* etc. */
    u_int32_t tf_v1;
    u_int32_t tf_a0;
    u_int32_t tf_a1;
    u_int32_t tf_a2;
    u_int32_t tf_a3;
    u_int32_t tf_t0;
    u_int32_t tf_t7;
    u_int32_t tf_a0;
    u_int32_t tf_a7;
    u_int32_t tf_t8;
    u_int32_t tf_t9;
    u_int32_t tf_k0;
    u_int32_t tf_k1;
    u_int32_t tf_gp;
    u_int32_t tf_sp;
    u_int32_t tf_a8;
    u_int32_t tf_epc; /* coprocessor 0 epc register */
};
```

By creating a pointer to here of type struct trapframe *, we can access the user's saved registers as normal variables within 'C'

epc
s8
sp
gp
k1
k0
t9
t8
...
at
ra
hi
lo
cause
status
vaddr

10

```
syscall(struct trapframe *tf)
{
    callno = tf->tf_v0;
    retval = 0;

    switch (callno) {
        case SYS_reboot:
            err = sys_reboot(tf->tf_a0);
            break;

        /* Add stuff here */

        default:
            kprintf("Unknown syscall %d\n", callno);
            err = ENOSYS;
            break;
    }
}
```

11

```
if (err) {
    tf->tf_v0 = err;
    tf->tf_a3 = 1; /* signal an error */
}
else {
    /* Success. */
    tf->tf_v0 = retval;
    tf->tf_a3 = 0; /* signal no error */
}

tf->tf_epc += 4;
```

12

```

int open(const char *filename, int flags);
int open(const char *filename, int flags, mode_t mode);
int close(int fd);
ssize_t read(int fd, void *buf, size_t buflen);
ssize_t write(int fd, const void *buf, size_t nbytes);
int dup2(int oldfd, int newfd);
off_t lseek(int fd, off_t pos, int whence);

```

Handwritten notes: *vop, 90, a2, a3, start*

13

lseek() Offset

```

uint64_t offset;
int whence;
off_t retval64;

join32to64(tf->tf_a2, tf->tf_a3, &offset);
copyin((userptr_t)tf->tf_sp + 16, &whence, sizeof(int));
split64to32(retval64, &tf->tf_v0, &tf->tf_v1);

```

14

Pointers

- What about the first argument to open()
 - It's a string?
- What are the problems with accessing a string (i.e. user-specified region of memory)?

15

Copy in/out(str)

```

int copyin(const userptr_t usersrc, void *dest, size_t len);
int copyout(const void *src, userptr_t userdest, size_t len);
int copyinstr(const userptr_t usersrc, char *dest, size_t len, size_t *got);
int copyoutstr(const char *src, userptr_t userdest, size_t len, size_t *got);

```

16

Buffers – e.g. read()

- Kernel framework for safely handling buffers
 - Does error/range/validity checking for you

```

ssize_t read(int fd, void *buf, size_t buflen);

struct iovec {
    union {
        userptr_t iov_base; /* user-supplied pointer */
        void *iov_kbase; /* kernel-supplied pointer */
    };
    size_t iov_len; /* length of data */
};

```

17

VFS READ

A macro with sanity checking

```
VOP_READ(vn, uio)
```

Invokes a function point of following prototype:

```
int (*vop_read)(struct vnode *file, struct uio *uio);
```

What are the arguments?

18

UIO

```
/* Source/destination. */
enum uio_seg {
    UIO_USERSPACE, /* User process code. */
    UIO_USERSPACE, /* User process data. */
    UIO_SYSSPACE, /* Kernel. */
};

struct uio {
    struct iovec *uio_iov; /* Data blocks */
    unsigned uio_iovcnt; /* Number of iovecs */
    off_t uio_offset; /* Desired offset into object */
    size_t uio_resid; /* Remaining amt of data to xfer */
    enum uio_seg uio_segflg; /* What kind of pointer we have */
    enum uio_rw uio_rw; /* Whether op is a read or write */
    struct addressspace *uio_space; /* Address space for user pointer */
};
```

19 UNSW

19

Sample Helper function

```
uio_init(struct iovec *iov, struct uio *u, userptr_t buf,
size_t len, off_t offset, enum uio_rw rw)
{
    iov->iov_ubase = buf;
    iov->iov_len = len;
    u->uio_iov = iov;
    u->uio_iovcnt = 1;
    u->uio_offset = offset;
    u->uio_resid = len;
    u->uio_segflg = UIO_USERSPACE;
    u->uio_rw = rw;
    u->uio_space = proc_getas();
}
```

20 UNSW

20

System call implementation

1. sys_open() → 1. vfs_open()
• copyinstr()
2. sys_close() → 2. vfs_close()
3. sys_read() → 3. VOP_READ()
4. sys_write() → 4. VOP_WRITE()
5. sys_lseek() → 5. VOP_ISSEEKABLE()
6. sys_dup2() → 6. VOP_STAT()

21 UNSW

21