Department of Electrical Engineering and Computer Science

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

6.828 Operating System Engineering: Fall 2004

Quiz 2

All problems are open-ended questions. In order to receive credit you must answer the question as precisely as possible. You have 80 minutes to answer this quiz.

Write your name on this cover sheet AND at the bottom of each page of this booklet.

Some questions may be much harder than others. Read them all through first and attack them in the order that allows you to make the most progress. If you find a question ambiguous, be sure to write down any assumptions you make. Be neat. If we can't understand your answer, we can't give you credit!

THIS IS AN OPEN BOOK, OPEN NOTES QUIZ.

1-4 (xx/35)	5-6 (xx/15)	7-8 (xx/20)	9-11 (xx/15)	12 (xx/10)	13-14 (xx/5)	Total (xx/100)

Name:

I Labs

Ursula Unsafe is working on Lab 4, System Calls for Environment Creation. Here is her first version of the `sys_mem_map` system call:

```
//
// Map the page of memory at 'srcva' in srcenvid's address space
// at 'dstva' in dstenvid's address space with permission 'perm'.
// Perm has the same restrictions as in sys_mem_alloc, except
// that it also must not grant write access to a read-only
// page.
//
// Return 0 on success, < 0 on error.  Errors are:
// -E_BAD_ENV if srcenvid and/or dstenvid doesn't currently exist,
// or the caller doesn't have permission to change one of them.
// -E_INVAL if srcva >= UTOP or srcva is not page-aligned,
// or dstva >= UTOP or dstva is not page-aligned.
// -E_INVAL is srcva is not mapped in srcenvid's address space.
// -E_INVAL if perm is inappropriate (see sys_page_alloc).
// -E_INVAL if (perm & PTE_W), but srcva is read-only in srcenvid's
// address space.
// -E_NO_MEM if there's no memory to allocate the new page,
// or to allocate any necessary page tables.
static int
sys_mem_map(uint_t srcenvid, u_int srcva,
             uint_t dstenvid, u_int dstva,
             u_int perm)
{
        struct Env* srcenv, *dstenv;
        struct Page* pg;
        Pte* pte;
        int retval;

        if ((retval = envid2env(srcenvid, &srcenv, 1)) < 0
            || (retval = envid2env(dstenvid, &dstenv, 1)) < 0)
                return retval;
        if ((srcva & (BY2PG - 1)) || (dstva & (BY2PG - 1)))
                return -E_INVAL;
        if (!(pg = page_lookup(srcenv->env_pgdir, srcva, &pte)))
                return -E_INVAL;
        return page_insert(dstenv->env_pgdir, pg, dstva, perm);
}
```

1. [10 points]: Describe two different ways that Ursula's `sys_mem_map` would let a user environment inject arbitrary code into the kernel. In particular, give two code fragments, including calls to `sys_mem_map`, after which the statement "`strcpy((char*) UTEMP, "Ha ha!");`" would write the string "`Ha ha!`" at kernel virtual address 0xF0003000.

2. [10 points]: Describe one way that Ursula's `sys_mem_map` would let a user environment change arbitrary pages in *other* user environments. In particular, sketch out code steps (possibly using comments instead of actual C), including a call to `sys_mem_map`, after which the statement "`strcpy((char*) UTEMP, "Ha ha!");`" would write the string "`Ha ha!`" at environment `e`'s virtual address 0x80000. (Hint: Consider using `envs[e].env_cr3`.)

Here's a version of the user-level page fault handler you wrote in Lab 4.

```
// Page fault upcall entrypoint.
// This is where we ask the kernel
// to redirect us to whenever we cause a page fault in user space
// (see the call to sys_set_pgfault_handler in pgfault.c).
//
// When a page fault actually occurs,
// the kernel switches our ESP to point to the user exception stack
// if we're not already on the user exception stack,
// and then it pushes the following minimal trap frame
// onto our user exception stack:
//
// [ 5 spare words ]
// trap-time eip
// trap-time eflags
// trap-time esp
// tf_err (error code)
// fault_va <-- %esp
//
// We then have to save additional caller-saved registers
// and call up to the appropriate page fault handler in C code,
// pointed to by the global variable '_pgfault_handler' declared above.

.text
.globl _pgfault_upcall
_pgfault_upcall:
        // Save the caller-saved registers.
        movl %eax, 28(%esp)
        movl %ecx, 24(%esp)
        movl %edx, 20(%esp)

        // Call the C page fault handler.
        movl _pgfault_handler, %eax
        call *%eax

        // Push trap-time %eip and %eflags onto the trap-time stack.
        //
        // Explanation:
        //   We must prepare the trap-time stack for our eventual return to
        //   re-execute the instruction that faulted.
        //   Unfortunately, we can't return directly from this stack
        //   (the exception stack).  Why not?
        //   We can't call 'jmp', since that requires that we load the address
        //   into a register, and all registers must have their trap-time
        //   values after the return.
        //   We can't call 'ret' from the exception stack either, since if we
        //   did, %esp would have the wrong value.
        //   So instead, we push the trap-time %eip onto the *trap-time* stack!
        //   Below we'll switch to that stack and call 'ret', which will
        //   restore %esp to its pre-fault value.
        //   We'll also restore %eflags from the trap-time stack, in case an
        //   intervening instruction changes the flags.  ('ret' does not.)
        //
        //   In the case of a recursive fault on the exception stack,
        //   note that the two words we're pushing now will overlap with
        //   the current exception frame!
        //
        movl 8(%esp), %eax      // trap-time esp in eax
        subl $8, %eax           // add space for eip and eflags
        movl %eax, 8(%esp)

        movl 16(%esp), %ecx     // eip
        movl %ecx, 4(%eax)
        movl 12(%esp), %ecx     // eflags
        movl %ecx, 0(%eax)
```

Name:

```
// Restore the caller-saved registers.
movl 20(%esp), %edx
movl 24(%esp), %ecx
movl 28(%esp), %eax

// Switch back to the adjusted trap-time stack.
movl 8(%esp), %esp

// Restore eflags from the stack.
popf

// Return to re-execute the instruction that faulted.
ret
```

3. [5 points]: Fill in the blanks in the following exception stack layout, referring to the code above. The top two words are used only in case of a recursive fault (i.e., a fault in the page fault handler itself); additionally circle the two instructions above that will set these two words, and draw a box around the two instructions that will use them.

```
+------------------------+
|                        |
|     trap-time _____    |
+------------------------+   <-- 36(%esp)
|                        |
|     trap-time _____    |
+------------------------+   <-- 32(%esp)
|                        |
|     trap-time _____    |
+------------------------+   <-- 28(%esp)
|                        |
|     trap-time _____    |
+------------------------+   <-- 24(%esp)
|                        |
|     trap-time _____    |
+------------------------+   <-- 20(%esp)
|     trap-time eip      |
+------------------------+   <-- 16(%esp)
|     trap-time eflags   |
+------------------------+   <-- 12(%esp)
|     trap-time esp      |
+------------------------+   <-- 8(%esp)
|  tf_err (error code)   |
+------------------------+   <-- 4(%esp)
|       fault_va         |
+------------------------+   <-- %esp when handler is run
```

Name:

4. [10 points]: Fill in the following function, which implements simple shared memory.

```
// Allocate a page of memory at 'va' that is shared with environment 'e'.
// Do not return until environment 'e' has also called shmget() with
// similar arguments.
// Don't worry about race conditions.
// Assume that 'e' calls sys_ipc_recv only from shmget.
void shmget(void* va, envid_t e, int perm)
{
        assert(!((u_int) va & (BY2PG - 1)));

        // Is 'e' waiting in ipc_recv?
        if (envs[e].env_ipc_recving) {









        } else {
                // Receive a page from 'e' mapped at 'va'.
                u_int va_from_e = ipc_recv(0, va, 0);
                assert(va_from_e == (u_int) va);
        }
}
```

Here are some header comments for reference.

```
// Allocate a page of memory and map it at 'va' with permission
// 'perm' in the address space of 'envid'.
// The page's contents are set to 0.
// If a page is already mapped at 'va', that page is unmapped as a
// side effect.
//
// perm -- PTE_U | PTE_P must be set, PTE_AVAIL | PTE_W may or may not be set,
//         but no other bits may be set.
//
// Return 0 on success, < 0 on error.  ...
int sys_page_alloc(u_int envid, u_int va, int perm);

// Try to send 'value' to the target env 'envid'.
// If va != 0, then also send page currently mapped at va,
// so that receiver gets a duplicate mapping of the same page.
//
// The send fails with a return value of -E_IPC_NOT_RECV if the
```

```
// target has not requested IPC with sys_ipc_recv.
//
// Otherwise, the send succeeds, and the target's ipc fields are
// updated as follows:
//    env_ipc_recving is set to 0 to block future sends
//    env_ipc_from is set to the sending envid
//    env_ipc_value is set to the 'value' parameter
// The target environment is marked runnable again.
//
// Return 0 on success, < 0 on error.
//
// If the sender sends a page but the receiver isn't asking for one,
// then no page mapping is transferred but no error occurs.
//
// srcva and perm should have the same restrictions as they had
// in sys_mem_map.
//
// Hint: you will find envid2env() useful.
int sys_ipc_can_send(u_int envid, u_int value,
                     u_int srcva, unsigned perm);

// Block until a value is ready.  Record that you want to receive,
// mark yourself not runnable, and then give up the CPU.
//
// Again, dstva should have the same restrictions as it had in
// sys_mem_map.  If it violates these restrictions, assume that it is
// zero.
static int sys_ipc_recv(u_int dstva)
```

II Brief paper questions

5. [10 points]: shell Answer the following question with respect to "*Es*: A shell with higher-order functions," by Haahr and Rakitzis. What will the following code print in the *es* shell?

```
fn xxx first rest {
    if {~ $#rest 0} {
        return $first
    }
    return <>{ xxx $rest } $first
}

echo <>{ xxx a b c }
```

6. [5 points]: Synchronization primitives Professor Dumbo wants to demonstrate a variant of the "ticket lock" to his class. Here's what he comes up with:

```
int nwaiting;

void acquire_lock() {
        atomic_inc(&nwaiting);
        while (nwaiting != 1)
                // wait for the other (nwaiting - 1) processes to finish
                pause(nwaiting - 1);
}

void release_lock() {
        atomic_dec(&nwaiting);
}
```

Give a specific scenario demonstrating that this implementaiton of the ticket lock is wrong.

III Microkernels

The paper "Improving IPC by Kernel Design," by Jochen Liedtke, describes a number of optimizations the author made to achieve good IPC performance in the L3 kernel. One such optimization involved the addition of new system calls. L3 contained two traditional calls for implementing IPC:

- **send** – send a message (asynchronously) to another process
- **receive** – wait for and return a message sent to a process

In addition, Liedtke added two more calls each of which combines the functionality of **send** and **receive**:

- **call**
- **reply & receive next**

7. [10 points]: Which of the following statements is true for the new **call/reply & receive next** interface, compared to equivalent code that issues a **receive** immediately following a **send** system call?

(Circle all that apply; there may be more than one answer.)

A. The new interface reduces the number of user-kernel crossings required for an IPC.

B. The new interface reduces the number of TLB flushes required during an IPC.

C. The new interface lets L3 reduce the number of scheduler queue manipulations required during an IPC.

D. When no processes are swapped out to disk, the new interface reduces the number of times the kernel must copy IPC arguments and results that don't fit in registers.

E. The new interface makes the scheduler less fair, and could even lead to starvation were it not for the "long time wakeup list."

To optimize data transfers during IPC, the kernel copies straight from the sender's address space into the receiver's, as follows. The kernel reserves a portion of virtual memory known as the communication window. The kernel also restricts each IPC argument to a maximum of 4 Megabytes. To copy into the recipient's address space, the kernel takes two page directory entries from the recipient's page directory and places them into the sender's at the virtual address of the communication window. (The kernel also clears the `PTE_U` bit, so that user code in the sender cannot access the memory mapped in the communications window.)

One concern is that, on the processor used by L3, it is expensive to flush TLB entries for a large virtual address range. (You can only flush one page translation at a time, or else the whole TLB.) Thus, Liedtke introduces the term *window clean*, to mean the TLB is guaranteed not to be caching any mappings for virtual addresses in the communication window.

8. [10 points]: Describe a scenario in which, when the kernel must copy data between address spaces, the TLB is *not* window clean. (Hint: feel free to use the old or the new system calls in your scenario.)

IV Disco

Answer questions in this section with respect to the paper "Disco: Running Commodity Operating Systems on Scalable Multiprocessors" by Bugnion, Devine, and Rosenblum.

9. [5 points]: On page 11, the authors write, "The execution of the operating system in general and critical section in particular is slower on disco, which increases the contention for semaphores and spinlocks." What is the intuition for why contention increases, and what experiment demonstrates it?

Name:

10. [5 points]: In Figure 5, why does "kernel time" decrease when running on Disco?

11. [5 points]: In Figure 6, why does Irix data grow? Why does buffer cache stay the same size for the M bar?

V File systems

12. [5 points]: After finishing a Window systems for JOS, Austin Ports is evaluating whether to use soft-updates to improve the performance of the JOS file system. He keeps the layout and the basic file structures the same, but he would like to arrange that the JOS file server writes as many disk blocks asynchronously to the disk as possible, while maintaining file system consistency. Is there a scenario in which the file server must undo file system modifications when writing blocks out to disk to ensure file system consistency? If so, draw a picture with a dependency. (The next page shows the JOS file system structures.)

13. [5 points]: What file system calls would add more dependencies? (Explain briefly why.)

```
// File nodes (both in-memory and on-disk)

// Bytes per file system block - same as page size
#define BY2BLK          BY2PG
#define BIT2BLK         (BY2BLK*8)


// Maximum size of a filename (a single path component), including null
#define MAXNAMELEN      128

// Maximum size of a complete pathname, including null
#define MAXPATHLEN      1024

// Number of (direct) block pointers in a File descriptor
#define NDIRECT         10
#define NINDIRECT       (BY2BLK/4)

#define MAXFILESIZE     (NINDIRECT*BY2BLK)

#define BY2FILE         256
struct File
  union
    struct
        uint8_t f_name[MAXNAMELEN];     // filename
        uint32_t f_size;                // file size in bytes
        uint32_t f_type;                // file type
        uint32_t f_direct[NDIRECT];
        uint32_t f_indirect;

        struct File *f_dir;             // valid only in memory
    ;
    uint8_t f_pad[BY2FILE];             // make sizeof(struct File) == BY2FILE
  ;
;

#define FILE2BLK        (BY2BLK/sizeof(struct File))

// File types
#define FTYPE_REG               0       // Regular file
#define FTYPE_DIR               1       // Directory


// File system super-block (both in-memory and on-disk)

#define FS_MAGIC        0x68286097      // Everyone's favorite OS class

struct Super
        uint32_t s_magic;       // Magic number: FS_MAGIC
        uint32_t s_nblocks;     // Total number of blocks on disk
        struct File s_root;     // Root directory node
;
```

```
/* from fs/fs.h */
int file_open(char *path, struct File **pfile);
int file_get_block(struct File *f, u_int blockno, void **pblk);
int file_set_size(struct File *f, u_int newsize);
void file_close(struct File *f);
int file_remove(char *path);
void fs_init(void);
int file_dirty(struct File *f, u_int offset);
void fs_sync(void);
void file_flush(struct File*);

/* from inc/fd.h */
struct Dev

        int dev_id;
        char *dev_name;
        int (*dev_read)(struct Fd*, void*, u_int, u_int);
        int (*dev_write)(struct Fd*, const void*, u_int, u_int);
        int (*dev_close)(struct Fd*);
        int (*dev_stat)(struct Fd*, struct Stat*);
        int (*dev_seek)(struct Fd*, u_int);
        int (*dev_trunc)(struct Fd*, u_int);
;

struct Fd

        u_int fd_dev_id;
        u_int fd_offset;
        u_int fd_omode;
;

struct Stat

        char st_name[MAXNAMELEN];
        u_int st_size;
        u_int st_isdir;
        struct Dev *st_dev;
;

struct Filefd

        struct Fd f_fd;
        u_int f_fileid;
        struct File f_file;
;
```

VI 6.828

We love to have your suggestions for improving 6.828. Please, answer the following question. (Any answer, except no answer, will receive full credit!)

14. [2 points]: If you could change one aspect of 6.828, what would it be?

15. [3 points]: Are there any topics you'd like to see added to the class, or any topics you'd like to see removed?

End of Quiz 2