The University of Texas at Austin
CS 372H Introduction to Operating Systems: Honors: Spring 2010
Midterm

- This exam is **75 minutes**. Stop writing when "time" is announced at the end. *You must turn in your exam; we will not collect them.* Do not get up or pack up between 70 and 75 minutes. The instructor will leave the room 78 minutes after the exam begins and will not accept exams outside the room.

- There are **16 questions** in this booklet. Many can be answered very quickly. However, some questions may be harder than others, and some questions earn more points than others. You may want to skim all questions before starting.

- **This midterm is closed book and notes. You may not use electronics: phones, PDAs, calculators, etc.** You may refer to ONE two-sided 8.5x11" sheet with 10 point or larger Times New Roman font, 1 inch or larger margins, and a maximum of 60 lines per side. **Please leave your UT IDs out.**

- If you find a question unclear or ambiguous, be sure to write any assumptions you make.

- Follow the instructions: if they ask you to justify something, explain your reasoning and any important assumptions. **Write brief and precise answers. Rambling brain dumps will not work and will only waste time.** *Think before you start writing* so that you can describe an answer crisply. *Be neat. If we can't understand your answer, we can't give you credit!*

- To discourage guessing and brain dumps, we will, except where noted, give 25%-33% of the credit for any problem left *completely blank* (for example, 1 point for a 3 point question). If you attempt a problem, you start at zero points for the problem. Note that by *problem* we mean numbered questions for which a point total is listed. *Sub-problems* with no points listed are not eligible for this treatment. Thus, if you attempt any sub-problem, you may as well attempt the other sub-problems in the problem. The exception is the True/False problems, where wrong answers on individual items will cost more than leaving an item blank, regardless of the other items within the larger problem.

- Don't overthink the above. What you should do is: if you *think* you *might* know the answer, then answer the problem. If you *know* you *don't* know the answer, then leave it blank.

- Don't linger. If you know the answer, give it, and move on.

- **Write your name and UT EID on this cover sheet and on the bottom of every page of the exam.**

Do not write in the boxes below.

I (xx/28)	II (xx/14)	III (xx/17)	IV (xx/11)	Total (xx/70)

Name: **UT EID:**

I Short answer (28 points total)

1. [6 points]   Consider a user-level threading package running on a machine whose operating system does not offer kernel threads. Pat Q. Hacker intends to build a single application, using a single process, and is trying to decide whether to use a user-level threading package or to program in the event-driven style. Assume that the machine has two CPUs.

Circle True or False for each item below:

True / False If the user-level threading package is *cooperatively* (that is, non-preemptively) scheduled, then the user-level threading package can arrange to simultaneously run two threads, one on each processor.

True / False If the user-level threading package is *preemptively* scheduled, then the user-level threading package can arrange to simultaneously run two threads, one on each processor.

True / False Coding in the event-driven style will allow the process to simultaneously handle two events, one on each processor.

2. [4 points]   Now Pat Q. Hacker is using a system that has *kernel threads but only one processor*. Pat is writing an application and decides to structure that application as several kernel threads that share data. Pat reasons that, since the CPU can do only one thing at a time, there is no need to protect the shared data in the application code with synchronization objects like mutexes.

State whether Pat's reasoning is correct, and justify your answer *briefly* below:

3. [3 points]   We heard a fair bit in class about coarse-grained vs fine-grained locking. Tinus Lorvalds has written an operating system called Tinux, and Tinus is also taking CS 372H. Inspired by MCS locks, Tinus decides to replace the implementation of all of Tinux's spinlocks with MCS locks, leaving all code otherwise unmodified. (That is, Tinus replaces the implementation of `acquire()` and `release()` but does not otherwise change any of his code.) By doing so, Tinus does what to the granularity of locking in the kernel?

Select one:

  A With this change, Tinus makes Tinux's locking more coarse-grained.
  B With this change, Tinus has no effect on the granularity of locking in Tinux.
  C With this change, Tinus makes Tinux's locking more fine-grained.

Name: **UT EID:**

4. **[6 points]**   Consider a machine with 32 MB of RAM running an operating system with virtual memory and swapping. The OS's page replacement policy is: if, on a page fault, a process needs a new physical page in RAM, evict the page that has been in RAM in the *longest*, and write it to the disk if it is dirty. The machine owner notices that for some workloads, the operating system does a lot of disk writes, and the owner is unhappy about that. In response, the owner installs an extra 8 MB of RAM, and re-runs the workload.

Circle True or False for each item below:

True / False There are workloads for which the extra RAM will *decrease* the number of page faults.

True / False There are workloads for which the extra RAM will *have no effect on* the number of page faults.

True / False There are workloads for which the extra RAM will *increase* the number of page faults.

5. **[2 points]**   Which of the following scheduling disciplines can lead to starving of processes? Assume a system in which processes do not synchronize with each other or otherwise coordinate.

Circle all that apply:

  A Strict priority

  B Round-robin

  C Lottery scheduling (assume each process has at least one ticket)

6. **[4 points]**   State one actual coding practice that developers follow to avoid deadlocks in their code. State your answer briefly (there is far more white space below than you need). If you cannot think of one, write, "hope and pray" to get 1 point, and do not write anything else below. A blank answer will receive 0 points.

Name: **UT EID:**

7. [3 points] Consider a JOS environment: `struct Env e;` which of the following installs e's address space as the current address space on the x86 CPU that runs JOS?

Select one, and note that the definitions below the choices may be helpful:

A `lcr3(e.env_cr3);`

B `lcr3(e.env_pgdir);`

C `env_pop_tf(&e.env_tf);`

D None of the above

Relevant definitions:

`lcr3(arg):` loads register `%cr3` with `arg`.

`env_pop_tf(tf):` restores saved register values and then calls `iret`.

`struct Env` is defined as follows:

```
struct Env {
        struct Trapframe env_tf;        // Saved registers
        LIST_ENTRY(Env) env_link;       // Free list link pointers
        envid_t env_id;                 // Unique environment identifier
        envid_t env_parent_id;          // env_id of this env's parent
        unsigned env_status;            // Status of the environment
        uint32_t env_runs;              // Number of times environment has run

        // Address space
        pde_t *env_pgdir;               // Kernel virtual address of page dir
        physaddr_t env_cr3;             // Physical address of page dir

};
```

Name: **UT EID:**

II   The readings (14 points total)

8. [4 points]   Answer ONE of the following two questions, **A** or **B**, neither of which we discussed in class (these are testing the textbook reading). The answers ought to be brief. If you don't know the answer, just leave it blank and keep going.

- **A** What is the classic *dining philosophers problem* described in the book? How many philosophers? What is the constraint? (You may answer this question by drawing a picture and then just stating a few words.)
- **B** In the context of deadlock prevention, what is the concept of a *safe state*, as explicated in the textbook?

Your answer here:

9. [4 points]   Liedtke writes, "IPC performance is the Master". He follows this creed seriously, going to great lengths to serve the goal of IPC performance in L3. Which of the following techniques does he use?

Circle all that apply:

- **A** Stack doubling
- **B** Changing the system call interface
- **C** Reducing the number of cycles needed to handle IPC below the theoretical minimum of 172
- **D** Using low-level bit and byte tricks, such as strategic placing of needed data within 32-bit words and packing along cache lines.
- **E** When a thread P sends an IPC to a thread Q, mapping the MSRs (model-specific registers) of P read-only in Q's address space.

Name: **UT EID:**

10. [4 points] There were two software errors described in the Therac-25 article. We discussed the first one in class (three threads, one of them ignored signals for eight seconds, etc.). This question is about the second software error. Recall from the paper that this error only happened when the operator positioned the turntable for field light operation, then attended to the patient, then pressed "set" (which was supposed to save the parameters and move the turntable out of field light position), and then pressed "beam on". Further recall that these circumstances did not always cause the error. The pseudocode below simplifies the Therac's control flow (as reported in the paper), but it leaves intact the identical bug reported in the paper:

```
uint8_t class3 = 0; /* 8-bit quantity */

while (1) {

        if (in field light position) {
              increment class3;
        }

        check whether operator pressed "set";

        if (operator pressed set) {
            if (class3 != 0) {
                move turntable out of field light position;
            }
            break;
        }
}
```

Explain the problem in the above pseudocode *briefly*. You do not have to explain the effect on the Therac-25's mechanics or on patients, or explain the error in detail, and you should ignore larger software engineering issues like that the code was poorly structured to begin with. Just explain the error briefly in the space below:

11. [2 points] Who wrote, "That's one hell of a good excuse for some of the brain-damages of minix. I can only hope (and assume) that Amoeba doesn't suck like minix does."?

Name: **UT EID:**

III Monitors and condition variables (17 points total)

12. [5 points]   This question concerns the bounded buffer producer/consumer code, using mutexes and condition variables, that we saw in lecture. This code is reproduced on the next page, but it's not important for you to read it carefully; the question references the important line numbers. Recall that `cond_wait()` (see line 12) takes two arguments: a mutex and a condition variable. Further recall that this function releases the mutex and sleeps *atomically*. Then, when the thread is woken by `cond_signal()` (line 32), it acquires the mutex. Your task here is to show why the two steps that are atomic must indeed be atomic. You will answer this specific question: **if lines 11–12 were instead the following, then what could go wrong?**

```
while (count == BUFFER_SIZE) {
    release(&mutex); // no guarantee of atomicity between this and next line
    wait_for_signal(&nonfull);
    acquire(&mutex);
}
```

Answer this question by (a) giving an interleaving in which something undesirable happens and (b) stating what the undesirable thing is. Include line numbers where possible. We have started the answer for you:

```
producer executes "release(&mutex)"
....
```

```
1      Mutex mutex;
2      Cond nonempty;
3      Cond nonfull;
4
5      void producer (void *ignored) {
6          for (;;) {
7              /* next line produces an item and puts it in nextProduced */
8              nextProduced = means_of_production();
9
10             acquire(&mutex);
11             while (count == BUFFER_SIZE)
12                 cond_wait(&nonfull, &mutex);
13
14             buffer [in] = nextProduced;
15             in = (in + 1) % BUFFER_SIZE;
16             count++;
17             cond_signal(&nonempty, &mutex);
18             release(&mutex);
19         }
20     }
21
22     void consumer (void *ignored) {
23         for (;;) {
24
25             acquire(&mutex);
26             while (count == 0)
27                 cond_wait(&nonempty, &mutex);
28
29             nextConsumed = buffer[out];
30             out = (out + 1) % BUFFER_SIZE;
31             count--;
32             cond_signal(&nonfull, &mutex);
33             release(&mutex);
34
35             /* next line abstractly consumes the item */
36             consume_item(nextConsumed);
37         }
38     }
39
```

13. [6 points] Say we have a monitor M that houses a mutex, called `mutex` (naturally), and three condition variables, `cv1`, `cv2`, and `cv3`. That is, the monitor's definition starts like this:

```
class M {
    Mutex mutex;
    Cond  cv1;
    Cond  cv2;
    Cond  cv3;
    ....
};
```

Assume that the monitor is implemented correctly ("correct" here means that it is built according to our coding standards, and it is both safe and live) and that it never calls `cond_broadcast()`. Pat Q. Hacker hates wasting memory, so suggests making the following replacements inside the monitor code:

- Replace all three condition variables with a single condition variable, `single_cv`.
- Replace any call to `cond_wait(&cv1, &mutex)`, `cond_wait(&cv2, &mutex)`, or `cond_wait(&cv3, &mutex)` with a call to `cond_wait(&single_cv, &mutex)`.
- Replace any call to `cond_signal(&cv1, &mutex)`, `cond_signal(&cv2, &mutex)`, or `cond_signal(&cv3, &mutex)` with a call to `cond_signal(&single_cv, &mutex)`.

Can the three bulleted changes, taken together, detract from the monitor's *safety*? Explain your answer *briefly*. (Recall that a safe program is, for our purposes, one that will never do anything bad; it has no race conditions, for example.)

Can the three bulleted changes, taken together, detract from the monitor's *liveness*? Explain your answer *briefly*. (Recall that a live program is one that is guaranteed to make progress in all cases.)

Name: **UT EID:**

14. [6 points] Now assume that the third bullet is instead the following:

- Replace any call to `cond_signal(&cv1, &mutex)`, `cond_signal(&cv2, &mutex)`, or `cond_signal(&cv3, &mutex)` with a call to `cond_broadcast(&single_cv, &mutex)`. *Read that carefully.*

Can the three bulleted changes, relative to the original monitor, detract from the monitor's *safety*? Explain your answer *briefly*. (Recall that a safe program is, for our purposes, one that will never do anything bad; it has no race conditions, for example.)

Can the three bulleted changes, relative to the original monitor, detract from the monitor's *liveness*? Explain your answer *briefly*. (Recall that a live program is one that is guaranteed to make progress in all cases.)

IV Spinlocks, plus one last question (11 points total)

15. **[10 points]** In this problem, you will implement a multiple-reader, single-writer lock as a *spinlock*. Here is the description:

```
struct sharedlock {
    int value; // when the lock is created, value is initialized to 0
};
```

- It allows multiple readers OR one single writer, and there are four functions:
  reader_acquire(struct sharedlock*),
  reader_release(struct sharedlock*),
  writer_acquire(struct sharedlock*),
  writer_release(struct sharedlock*).
  We have given you the first of these, and your task is to write the last three of these. ***Each of these three functions only needs to be a single line of code.***

- When the lock is unlocked (no readers or writers holding the lock), its value is 0.

- When there are one or more readers holding the lock (that is, multiple threads have completed reader_acquire() but have not called reader_release()), the lock's value equals the number of readers

- When the lock is held by a writer (i.e., a thread has made it past writer_acquire() but has not called writer_release()), its value is -1.

- We are unconcerned here with fairness, efficiency, or starvation; just write something that is safe and that eventually allows a waiting thread, reader or writer, to make progress, even though a waiting writer may have to wait until there are no readers.

- Assume that the lock is never acquired by an interrupt handler, so you don't need to worry about enabling and disabling interrupts. You may also assume that the hardware provides sequential consistency.

You will likely need to call two atomic primitives. We describe them below. (We also include their pseudocode and inline assembly implementations in the appendix at the end of the exam. However, you do not need to check this appendix material to do the problem.)

- int cmpxchg_val(int* addr, int oldval, int newval), which we saw in lecture. It is an atomic operation. It compares oldval to *addr, and if the two are equal, it sets *addr = newval. It returns the old contents of *addr.

- void atomic_decrement(int* arg), which atomically performs *arg = *arg - 1.

We repeat, each of the three remaining functions requires only a line of code.

A final note: if you are stuck on this problem, recall that, in lecture, we saw a multiple-reader, single-writer lock implemented as a *monitor*. That code may be useful inspiration, and we include it in the appendix section at the end of the exam. However, you do not need to consult it to do the problem; it's there only because you may find it useful.

Name: **UT EID:**

```
// we are giving you the code for the first of the four functions:
void reader_acquire(struct sharedlock* lock) {
     int curr_val;
     while (1) {

          // spin while a writer owns the lock
          while ((curr_val = lock->value) == -1) {}

          assert(curr_val >= 0);

          // try to atomically increment the count, based on our best
          // guess of how many readers there had been. if we were
          // wrong, keep looping. if we got it right, then we
          // succeeded in incrementing the count atomically, and we
          // can proceed.
          if (cmpxchg_val(&lock->value, curr_val, curr_val + 1) == curr_val)
              break;
     }
     // lock->value now contains curr_val + 1
}

void reader_release(struct sharedlock* lock) {
     // your code here: only needs to be one line




}




void writer_acquire(struct sharedlock* lock) {
     // your code here: only needs to be one line




}


void writer_release(struct sharedlock* lock) {
     // your code here: only needs to be one line




}
```

There is one more page and one further question for you to complete.

Name: UT EID:

16. [1 points] Last question: What is the recommended completion date for lab 4A? For reference, today is Thursday, March 11, 2010.

End of Midterm

Enjoy Spring Break!!

A   Implementation of atomic primitives

cmpxchg_val()

```
/* pseudocode */
int cmpxchg_val(int* addr, int oldval, int newval) {
     LOCK: // remember, this is pseudocode
     int was = *addr;
     if (*addr == oldval)
          *addr = newval;
     return was;
}

/* inline assembly */
int cmpxchg_val(int* addr, int oldval, int newval) {
    int was;
    asm volatile("lock cmpxchg %3, %0"
         : "+m" (*addr), "=a" (was)
         : "a" (oldval), "r" (newval)
         : "cc");
    return was;
}
```

atomic_decrement()

```
/* pseudocode */
void atomic_decrement(int* arg) {
    LOCK: // remember, this is pseudocode
    *arg = *arg - 1;
}

/* inline assembly */
void atomic_decrement(int* arg) {
    asm volatile("lock decl %0" : "+m" (*arg) : "m" (arg));
}
```

Name: **UT EID:**

Name: **UT EID:**

B   Shared lock code from lecture

```
struct sharedlock {
  int i;
  Mutex mutex;
  Cond c;
};

void AcquireExclusive (sharedlock *sl) {
  acquire(&sl->mutex);
  while (sl->i) {
    cond_wait(&sl->c, &sl->mutex);
  }
  sl->i = -1;
  release(&sl->mutex);
}

void AcquireShared (sharedlock *sl) {
  acquire(&sl->mutex);
  while (sl->i < 0) {
    cond_wait(&sl->c, &sl->mutex);
  }
  sl->i++;
  release(&sl->mutex);
}

void ReleaseShared (sharedlock *sl) {
  acquire(&sl->mutex);
  if (!--sl->i)
    cond_signal(&sl->c, &sl->mutex);
  release(&sl->mutex);
}

void ReleaseExclusive (sharedlock *sl) {
  acquire(&sl->mutex);
  sl->i = 0;
  cond_broadcast(&sl->c, &sl->mutex);
  release(&sl->mutex);
}
```

Name: **UT EID:**