

CS 170, 1.24.24

Ryan Wenger



PART 1 – KOS

Origins of KOS

- KOS is the name of the kernel you will be building over the course of this quarter
- It's based on ULTRIX, a flavor of UNIX that was built and shipped in the 1990's by Digital Equipment Corporation (DEC) for their line of MIPS workstations, called the DECstation (though it was later ported to DEC's in-house VAX architecture)
- Crucially, ULTRIX and early Linux were around at the same time—resulting in the existence of a 32-bit Linux-hosted cross compiler (decstation-ultrix-gcc/g++/ld/etc.) that can build C programs to run on ULTRIX
 - This is GCC version 2.95, which supports nothing newer than the ANSI C89 standard. If your programs are behaving weirdly or you're getting nonsensical compiler errors, make sure you aren't doing anything too fancy in your test code.
 - GNU Binutils v2.9.1 are also present in the same directory as GCC (objdump, gas, ar, etc.)
- By implementing a subset of the ULTRIX system calls, you will be able to run such programs on our simulator to test the functionality of your own KOS kernel
 - See the provided Makefile in the lab handout



PART 2 - THE SIMULATOR

The simulator and your KOS kernel are *the same program.*

- When you build and run KOS on CSIL, **you are executing a regular x86_32 Linux program.** Linux sees *no difference* between your kernel and python, or firefox, or any C++ program you run yourself (other than the fact it's 32-bit).
- The simulator that we provide you is nothing more than a collection of .o files (some of which are packaged into .a archives). Inside of one of these is a main() function that eventually calls KOS()—the entry point of the kernel, which you should write.
- In a way, you can think of your KOS code as an external library (*dependency*, if you will) that is needed to run the simulator.
- Understanding this is crucial to understanding the most fundamental difference between KOS and any real-world kernel.

A Closer Look

Simulator

- Kernel does not run on hardware—in fact, your kernel is compiled using normal Linux gcc (in 32-bit mode)
- This means you can `#include <stdlib.h>` and call, for example, `malloc()`
- Helpfully, this means you can use `printf()` or GDB for debugging your kernel
 - Do not use `printf/scanf` to bypass the simulator's `ConsoleRead/Write` in your submission –you will not get credit and your kernel will behave incorrectly

Real Machine

- Kernel is responsible for managing hardware, and interfacing directly with firmware
- A real kernel must implement its own memory allocator—no code is accessible other than what is contained in the program

Key Things to Recognize in your Mental Model

- Abstractly, the goal of any kernel is to take the mishmash collection of asynchronous hardware devices that make up a machine and unite them into some synchronous-looking interface that can be used by a user's programs.
 - This interface is the kernel's System Call API
 - Luckily, you don't have to design this API, you just have to implement the one used by Linux/UNIX/ULTRIX/POSIX/etc.
- The simulator calls your code in 3 ways:
 - At startup
 - When an exception occurs (system call, segfault, etc)
 - On a device interrupt
- Crucially, on real hardware, the latter 2 of these can occur at any moment in the kernel's execution. **That is not the case with KOS.**
 - Because the simulator simply makes "library" calls out to your kernel, it can only call your exception and interrupt handlers when the kernel returns control to it (via `run_user_code()`)
 - Thus, it is impossible for your code to jump from the middle of a system call routine into an interrupt handler (unless you explicitly do so).



PART 2.5 - KT THREADS

Quick Summary of Using KT Threads

- If you're familiar with C++ coroutines or GO's "goroutines," then you understand kt threads
 - Like pthreads, they can be created, joined, suspended, and destroyed
 - Unlike pthreads, the OS on which they are running—Linux, *not* KOS— has no knowledge of their existence
 - Thus, no more than one piece of code is executing at any given time—no parallelism occurs
 - Likewise, no preemption can occur—your code will not halt in middle of one thread and transition to another unless you explicitly invoke `kt_join/yield/exit/etc.`



PART 3 - TESTING

TEST TEST TEST TEST TEST TEST TEST TEST

- As with lab0, testing KOS is absolutely imperative
- There are a handful of programs provided in Rich's cs170 directory on CSIL (~rich/cs170/test_execs), but you will *need* to write your own if you want to ensure a good score on the labs
- Even though KOS runs as an x86 Linux process, the simulator simulates a DECstation machine with a 32-bit MIPS R3000 CPU. You should use and/or modify the provided Makefile.xcomp in the handout to build your test programs
 - This will ensure all of the linker options and transformations necessary to run code on the simulator are applied to the output of the decstation-ultrix-gcc compiler
 - Note: the <unistd.h> header file, which typically contains UNIX system call declarations, is missing from the toolchain; however, the cross compiler "knows" about these functions, so you can call them without including any header files
 - Highly recommend compiling your test programs with decstation-ultrix-gcc and with the regular '/usr/bin/gcc -m32 -std=gnu89' so you can compare behavior between your kernel and Linux.



PART 4 - ARCHITECTING YOUR KERNEL

The Scheduler

- Your kernel's scheduler is the heart of the entire system—it decides what MIPS code gets run, when it gets run, and (later on) how long to run it for.
- The key thing to recognize is that, after KOS initially starts up, it doesn't actually *do* anything until a user space program invokes it (whether by system call or segfault/etc)
 - Thus, once it's done handling whatever actions have been generated by the user's program, it just waits for something else to happen.
- In future labs, the scheduler will have to juggle multiple processes (programs)—some of which may be waiting for I/O interrupts and thus cannot yet continue, some of which can continue and are waiting for their turn to execute, and (at most) exactly one who is currently executing MIPS code in the simulator.
 - Using a "run queue" to manage these processes can be very helpful

The Console

- The simulator's only means of sending information to/from your terminal ("Console") is the `console_read()` and `console_write()` functions. Each operates on 1 character at a time.
- Each of these is associated with an interrupt: a `ConsoleReadInt` occurs when there is data ready to be read by `console_read()`; `ConsoleWriteInt` occurs when the previous call to `console_write()` has completed and the console is ready to put out another character.
- See the cookbook/handout for more details

System Calls

- Because KOS must be able to handle multiple in-progress system calls (i.e. one process makes a syscall while another is still waiting for its own to complete), every system call handler should execute in its own kt thread.
 - This means arguments must be marshalled to and from each handler invocation. Consider writing a macro to streamline this, as you'll eventually end up with 19 different calls, and copy/pasting the same marshal/unmarshal chunk can get annoying
- Here's part of a set of macros my partner and I used when we took 170. If you're curious about them, come see me in office hours, or feel free to try and fill in the missing ones yourself.

```
/* Use this instead of an actual prototype in system call definitions
 * It matches with the auto declarations in the header so syscalls
 * are compatible with `kt_fork(void * (*) (void *))` and auto unmarshalls
 * the arguments into local variables. It also creates a local variable
 * called `caller` of type `PCB *`.
 *
 * ex: `kssize_t read(int fd, void *buf, ksize_t count)`
 *      becomes
 *      `SYSCALL_DEFINE(read, kssize_t, int fd, void *buf, ksize_t count)`
 */

#define SYSCALL_DECL(name)      void *sys_##name(void * __caller)

#define SYSCALL_DEFINE(name, ret_t, ...) \
    SYSCALL_DECL(name) \
    { \
        PCB *caller = __caller; \
        __VA_OPT__(__SYSCALL_ARG_UNMARSHALL(5, __VA_ARGS__))

#define SYSCALL_RETURN(caller, retval) \
    do { \
        return __syscall_return(caller, retval, 0); \
    } while(0)
```

Use the Cookbook

- There is a “cookbook” on the course page that provides nearly everything you need to design and implement the labs. Please check it out.