

xv6 Overview

CS 202: Advanced Operating Systems

xv6

- xv6 is MIT's re-implementation of Unix v6

- Written in ANSI C
- Runs on RISC-V and x86
 - We will use the RISC-V version with the QEMU simulator
- Smaller than v6
- Preserve basic structure (processes, files, pipes. etc.)
- Runs on multicores
- Got paging support in 2011

*Ken Thompson &
Dennis Ritchie, 1975*

xv6

- To understand it, you'll need to read its source code
 - It's not that hard
 - Source code:
 - <https://github.com/rtenlab/xv6-riscv> (for our projects)
 - Forked from <https://github.com/mit-pdos/xv6-riscv>; to avoid unexpected updates during this course
 - Book/commentary
 - xv6: a simple, Unix-like teaching operating system
 - <https://pdos.csail.mit.edu/6.828/2022/xv6/book-riscv-rev3.pdf>

Why xv6?

- Why study an old OS instead of Linux, Solaris, or Windows?

1. Big enough
 - To illustrate basic OS design & implementation
2. Small enough
 - To be (relatively) easily understandable
3. Similar enough
 - To modern OSes
 - Once you've explored xv6, you will find your way inside kernels such as Linux

Why RISC-V?



RISC-V: The Free and Open RISC
Instruction Set Architecture

- RISC-V: open standard instruction set architecture (ISA) based on RISC principles
 - High quality, loyalty free, license free
 - Multiple proprietary and open-source core implementations
 - Supported by growing software ecosystem
 - Appropriate for all levels of computing system, from microcontrollers to supercomputers
- Fun to use toolchains for the new architecture

Apple shows interest in RISC-V chips, a competitor to iPhones' Arm tech

RISC-V chip technology could be used for tasks like AI and computer vision.

Intel Will Offer SiFive RISC-V CPUs on 7nm, Plans Own Dev Platform

By Joel Hruska on June 24, 2021 at 8:36 am | [Comments](#)

xv6 Structure

- Monolithic kernel
 - Provides services to running programs
- Processes uses system calls to access system services
- When a process call a system call
 - Execution will enter the kernel space
 - Perform the service
 - Return to the user space

xv6 System Calls

System call	Description
fork()	Create process
exit()	Terminate current process
wait()	Wait for a child process to exit
kill(pid)	Terminate process pid
getpid()	Return current process's id
sleep(n)	Sleep for n seconds
exec(filename, *argv)	Load a file and execute it
sbrk(n)	Grow process's memory by n bytes
open(filename, flags)	Open a file; flags indicate read-/write

xv6 System Calls (2)

System call	Description
read(fd, buf, n)	Read n bytes from an open file into buf
write(fd, buf, n)	Write n bytes to an open file
close(fd)	Release open file fd
dup(fd)	Duplicate fd
pipe(p)	Create a pipe and return fd's in p
chdir(dirname)	Change the current directory
mkdir(dirname)	Create a new directory
mknod(name, major, minor)	Create a device file
fstat(fd)	Return info about an open file
link(f1, f2)	Create another name (f2) for the file f1
unlink(filename)	Remove a file

xv6 kernel source files

- /kernel directory

bio.c	Disk block cache for the file system.
console.c	Connect to the user keyboard and screen.
entry.S	Very first boot instructions.
exec.c	exec() system call.
file.c	File descriptor support.
fs.c	File system.
kalloc.c	Physical page allocator.
kernelvec.S	Handle traps from kernel, and timer interrupts.
log.c	File system logging and crash recovery.
main.c	Control initialization of other modules during boot.
pipe.c	Pipes.
plic.c	RISC-V interrupt controller.
printf.c	Formatted output to the console.
proc.c	Processes and scheduling.
sleeplock.c	Locks that yield the CPU.
spinlock.c	Locks that don't yield the CPU.
start.c	Early machine-mode boot code.
string.c	C string and byte-array library.
swtch.S	Thread switching.
syscall.c	Dispatch system calls to handling function.
sysfile.c	File-related system calls.
sysproc.c	Process-related system calls.
trampoline.S	Assembly code to switch between user and kernel.
trap.c	C code to handle and return from traps and interrupts.
uart.c	Serial-port console device driver.
virtio_disk.c	Disk device driver.
vm.c	Manage page tables and address spaces.

Setup

- Toolchain
 - You need a RISC-V tool chain and QEMU for RISC-V
- Linux (Ubuntu 20.04)

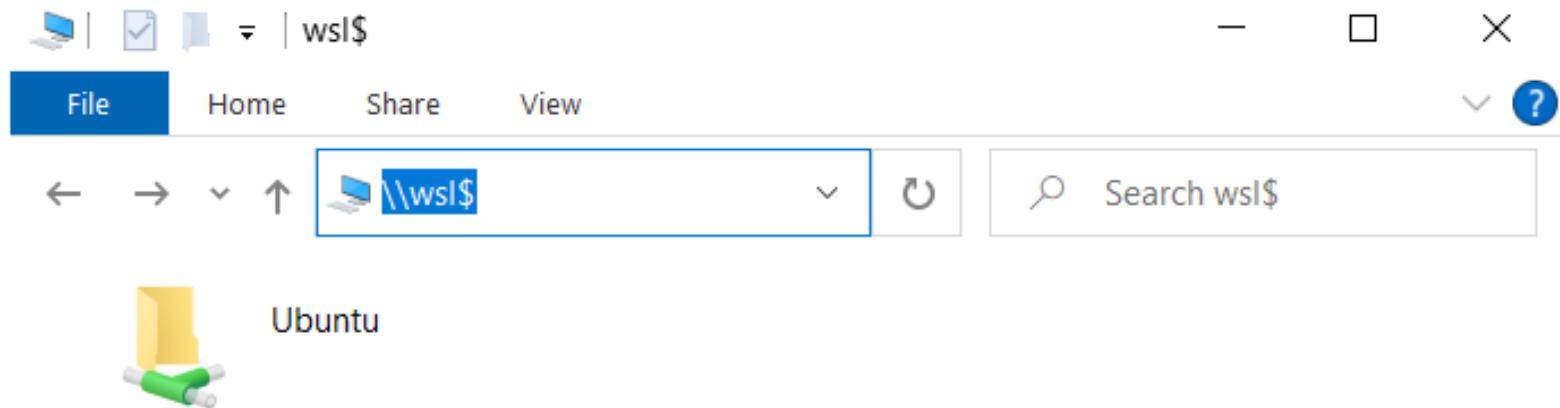
```
$ sudo apt update
$ sudo apt install git build-essential gdb-multiarch qemu-system-misc gcc-riscv64-linux-gnu binutils-riscv64-linux-gnu
```

- Windows
 - You can use Windows Subsystem for Linux (WSL) with Ubuntu 20.04
 - Unsure what version of Ubuntu you have? Open WSL terminal and type “lsb_release -a”
 - Follow the above Linux instruction for package installation

```
~$ lsb_release -a
No LSB modules are available.
Distributor ID: Ubuntu
Description:    Ubuntu 20.04 LTS
Release:        20.04
Codename:       focal
```

Setup

- Windows (cont')
 - All your WSL Linux files are accessible as \\wsl\$ in File Explorer
 - Exposed as network shared files
 - Your home directory is \\wsl\$\home<username>



Setup

- macOS

- Install developer tools:

```
$ xcode-select --install
```

- Install Homebrew (package manager)

```
$ /usr/bin/ruby -e "$(curl -fsSL https://raw.githubusercontent.com/Homebrew/install/master/install)"
```

- Install the RISC-V compiler toolchain:

```
$ brew tap riscv/riscv
$ brew install riscv-tools
```

- Update path; open `~/.bashrc` and add the following line

```
PATH=$PATH:/usr/local/opt/riscv-gnu-toolchain/bin
```

- Install QEMU

```
$ brew install qemu
```

Setup

- **Download xv6:**

```
$ git clone https://github.com/rtenlab/xv6-riscv
```

```
$ cd xv6-riscv
```

- **Compile and run xv6:**

```
$ make qemu
```

(“make” to compile only)

- **Exit from QEMU**

- Press **Ctrl+a** and then press **c** to get the QEMU console
- Then type “**quit**” to exit

```
xv6 kernel is booting
hart 2 starting
hart 1 starting
init: starting sh
$
```

(xv6 shell)

What does “hart” mean?

- In RISC-V, hart refers to a *hardware thread*
- xv6 boots on hart 0 and turns on other harts

Create a System Call

- Goal: create a system call “`sys_hello`” that call a kernel function that displays: “Hello from the kernel space!”
- To do that, open the following files and add the lines with “`// hello`” comment:

Create a System Call (2)

- kernel/syscall.h: define new syscall number

```
16  #define SYS_open    15
17  #define SYS_write   16
18  #define SYS_mknod   17
19  #define SYS_unlink  18
20  #define SYS_link    19
21  #define SYS_mkdir   20
22  #define SYS_close   21
23  #define SYS_hello   22 // hello
```

Create a System Call (3)

- kernel/syscall.c: update system call table

```
105  extern uint64 sys_write(void);
106  extern uint64 sys_uptime(void);
107  extern uint64 sys_hello(void); // hello: declaration
108
109 static uint64 (*syscalls[]) (void) = {
110     [SYS_fork]      sys_fork,
111     [SYS_exit]      sys_exit,
112     [SYS_wait]      sys_wait,
113     [SYS_pipe]      sys_pipe,
114     [SYS_read]      sys_read,
115     [SYS_kill]      sys_kill,
116     [SYS_exec]      sys_exec,
117     [SYS_fstat]     sys_fstat,
118     [SYS_chdir]     sys_chdir,
119     [SYS_dup]       sys_dup,
120     [SYS_getpid]    sys_getpid,
121     [SYS_sbrk]       sys_sbrk,
122     [SYS_sleep]     sys_sleep,
123     [SYS_uptime]    sys_uptime,
124     [SYS_open]       sys_open,
125     [SYS_write]     sys_write,
126     [SYS_mknod]     sys_mknod,
127     [SYS_unlink]    sys_unlink,
128     [SYS_link]      sys_link,
129     [SYS_mkdir]     sys_mkdir,
130     [SYS_close]     sys_close,
131     [SYS_hello]     sys_hello, // hello: syscall entry
132 };
133
```

Create a System Call (4)

- kernel/sysproc.c: define syscall function

```
93  uint64 sys_hello(void) // hello syscall definition
94  {
95      int n;
96      argint(0, &n);
97      print_hello(n);
98      return 0;
99 }
```

- kernel/proc.c: new kernel function

```
685  // hello: printing hello msg
686  void print_hello(int n)
687  {
688      printf("Hello from the kernel space %d\n", n);
689 }
```

Create a System Call (5)

- kernel/defs.h

```
84 // proc.c
85 int cpuid(void);
86 void exit(int);
87 int fork(void);
88 int growproc(int);
89 void proc_mapstacks(pagetable_t);
90 pagetable_t proc_pagetable(struct proc *);
91 void proc_freepagetable(pagetable_t, uint64);
92 int kill(int);
93 int killed(struct proc* );
94 void setkilled(struct proc* );
95 struct cpu* mycpu(void);
96 struct cpu* getmycpu(void);
97 struct proc* myproc();
98 void procinit(void);
99 void scheduler(void) __attribute__((noreturn));
100 void sched(void);
101 void sleep(void*, struct spinlock* );
102 void userinit(void);
103 int wait(uint64);
104 void wakeup(void* );
105 void yield(void);
106 int either_copyout(int user_dst, uint64 dst, void *src, uint64 len);
107 int either_copyin(void *dst, int user_src, uint64 src, uint64 len);
108 void procdump(void);
109 void print_hello(int); // hello
```

Create a System Call (6)

- Update user-space syscall interface
- user/usys.pl

```
36 entry("sbrk");
37 entry("sleep");
38 entry("uptime");
39 # hello syscall for user
40 entry("hello");
41
```

- user/user.h

```
4 // system calls
5 int fork(void);
6 int exit(int) __attribute__((noreturn));
7 int wait(int*);
8 int pipe(int*);
9 int write(int, const void*, int);
10 int read(int, void*, int);
11 int close(int);
12 int kill(int);
13 int exec(char*, char**);
14 int open(const char*, int);
15 int mknod(const char*, short, short);
16 int unlink(const char*);
17 int fstat(int fd, struct stat*);
18 int link(const char*, const char*);
19 int mkdir(const char*);
20 int chdir(const char*);
21 int dup(int);
22 int getpid(void);
23 char* sbrk(int);
24 int sleep(int);
25 int uptime(void);
26 int hello(int); // hello
```

Test a System Call

1. Write a user program: Create “**test.c**” file in the user directory of “**xv6-riscv**” (user/test.c)

```
1  #include "kernel/types.h"
2  #include "kernel/stat.h"
3  #include "user/user.h"
4
5  int main(int argc, char *argv[])
6  {
7      int n = 0;
8      if (argc >= 2) n = atoi(argv[1]);
9
10     printf("Say hello to kernel %d\n", n);
11     hello(n);
12     exit(0);
13 }
```

Test a System Call (2)

2. Edit “Makefile” and append “\$U/_test\” to UPROGS

```
118 UPROGS=\  
119     $U/_cat\  
120     $U/_echo\  
121     $U/_forktest\  
122     $U/_grep\  
123     $U/_init\  
124     $U/_kill\  
125     $U/_ln\  
126     $U/_ls\  
127     $U/_mkdir\  
128     $U/_rm\  
129     $U/_sh\  
130     $U/_stressfs\  
131     $U/_usertests\  
132     $U/_grind\  
133     $U/_wc\  
134     $U/_zombie\  
135     $U/_test\
```

Test a System Call (3)

3. Type:

```
$ make qemu
```

4. After xv6 boots, type:

```
$ test
```

```
xv6 kernel is booting

hart 1 starting
hart 2 starting
init: starting sh
$ test 123
Say hello to kernel 123
Hello from the kernel space 123
$
```

How to use GDB

- To run Qemu with GDB, you need to open another terminal at the same xv6-riscv folder.
- In the first terminal, type:

```
$ make qemu-gdb
```

- In the second terminal, type:

```
$ gdb-multiarch -q -iex "set auto-load safe-path ."
```

```
~/xv6-riscv$ gdb-multiarch -q -iex "set auto-load safe-path ."
The target architecture is assumed to be riscv:rv64
warning: No executable has been specified and target does not support
determining executable automatically. Try using the "file" command.
0x0000000000001000 in ?? ()
(gdb) continue
```

Use “break <address>” to set a breakpoint
Type “continue” to run until breakpoint

MacOS: If gdb-multiarch doesn’t exist, try “riscv64-unknown-elf-gdb”

Change # of CPUs

- By default, xv6 is compiled for three CPUs. To change the number of CPUs, edit Makefile:

```
156 ifndef CPUS
157 # CPUS := 3
158 CPUS := 1
159 endif
```

- Unless otherwise mentioned, we will use a single-core system, so change to “ CPUS := 1 ”

```
xv6 kernel is booting
init: starting sh
$
```

No other harts (cores) after this change