Adding System Calls to OS/161

CSE 4001 Operating Systems Concepts

E. Ribeiro

January 26, 2022

Outline

1. Review: traps and system calls
2. Overview of steps to add system calls to OS/161
   - Kernel-level steps
   - User-level steps
   - Testing the system call
3. Kernel-level steps in detail
4. User-level steps in detail
5. Testing steps in detail

Review: System-call trapping mechanism

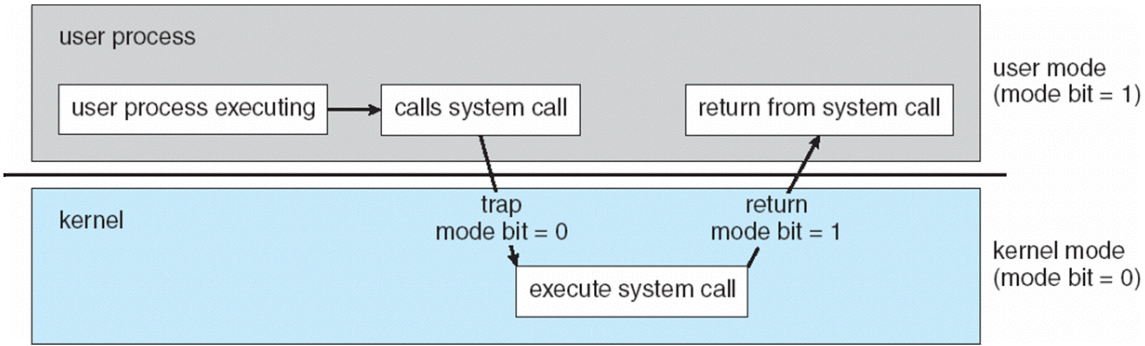

Figure adapted from Silberschatz, Galvin, and Gagne, 2009.

Review: System-call trapping mechanism in OS/161

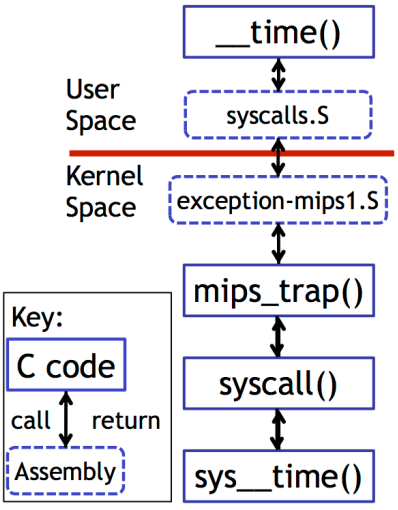

Kernel-level steps

1. Add the prototype of the system-call function to the header file: `kern/include/syscall.h`
2. The kernel-level implementation (e.g., `newsyscall.c`) goes into `kern/syscall/`
3. Add a new ID number for the system call. The new entry goes in the file `kern/include/kern/syscall.h`
4. Add a new branch in the switch-case statement in: `kern/arch/mips/syscall/syscall.c`
5. Add file entry definition for `syscall/newsyscall.c` in `kern/conf/conf.kern`

User-level steps

1. Add the user-level prototype of the system call to: `user/include/unistd.h`
2. Add the user-level test function. For this, create a new subdirectory directory `user/testbin/testnewsyscall/` and inside it add the test function (e.g., `testnewsyscall.c`).
3. Create a Makefile inside this subdirectory for building the test function. You can use one of the subdirectories as a template.
4. Add an entry to the new function to the top-level Makefile in `user/testbin`

Testing the new system call

1. Re-build the kernel
2. Start the new kernel (i.e., run sys161 kernel in the root directory)
3. At the OS161 prompt, use the p option (from OS161 menu) to run the test program, i.e., `p testbin/testnewsyscall`

Kernel-level steps

1 Prototype of the system call

1. Add the prototype of the system call to the header file: `kern/include/syscall.h`
2. At the end of the file, you will find prototypes for `sys_reboot()` and `sys__time()`.

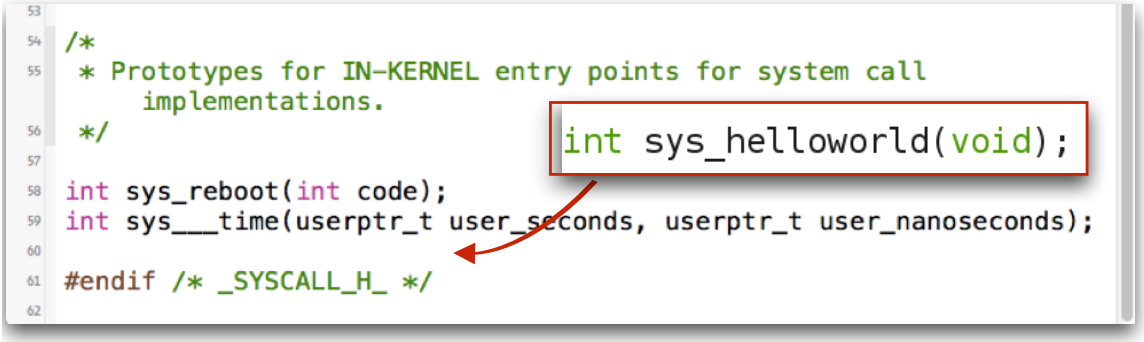

2 Kernel-level implementation

1. The kernel-level implementation goes into `kern/syscall`. This directory contains an example of a system call, i.e., `time_syscalls.c`.
2. Here, create a program called `simple_syscall.c`, and implement your system call in it.

```
int sys_helloworld(void){
        return kprintf("Hello World!\n");
}
```

3 Create the ID number for the new system call

1. The OS needs to know the ID number of the system call
2. Add a new entry to the file `kern/include/kern/syscall.h`

```
98  //#define SYS_setpgid     41
99  //#define SYS_getsid      42
100 //#define SYS_setsid      43
101 //                              (userlevel debugging)
102 //#define SYS_ptrace      44
103
104 //                              -- File-handle-related --
105
106
107 #define SYS_open            45
108 #define SYS_pipe            46
109 #define SYS_dup             47
110 #define SYS_dup2            48
```

4 Add a new branch in the switch-case statement in: kern/arch/mips/syscall/syscall.c

```
case SYS_helloworld:
    err = sys_helloworld();
    break;
```

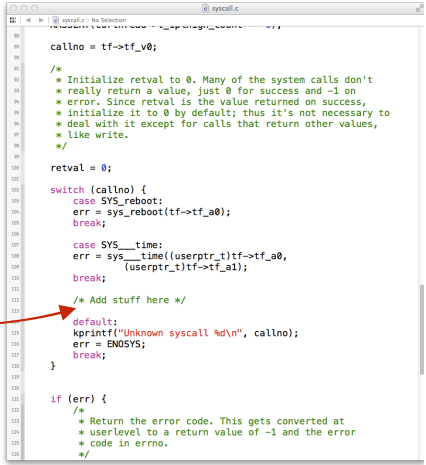

Note how user-level input parameters are passed to kernel-level functions via the trapframe.

5 Add file-entry definition to `config.kern`

```
358
359 file      vfs/devnull.c
360
361 #
362 # System call layer
363 # (You will probably want to add stuff here while doing the basic system
364 # calls assignment.)
365 #
366
367 file      syscall/loadelf.c
368 file      syscall/runprogram.c
369 file      syscall/time_syscalls.c
370
371 #
372 # Startup and initialization
373 #
374
375 file      startup/main.c
376 file      startup/menu.c
377
378 ########################################
379 #                                      #
380 #             Filesystems              #
381 #                                      #
382 ########################################
```

User-level steps

1. Add the user-level prototype of the system call to: `userland/include/unistd.h`

2. Add the user-level test function.

For this, create a new subdirectory directory `user/testbin/testnewsyscall/` and inside it add the test function (e.g., `testnewsyscall.c`).

helloworldtest.c

```c
#include <unistd.h>

int
main()
{
    helloworld();
    return 0;
}
```

```
# Makefile for helloworldtest

TOP=../../..
.include "$(TOP)/mk/os161.config.mk"

PROG=helloworldtest
SRCS=helloworldtest.c
BINDIR=/testbin

.include "$(TOP)/mk/os161.prog.mk"
```

3. Modify the top-level makefile.

Add an entry to the new function to the top-level Makefile in `user/testbin/`

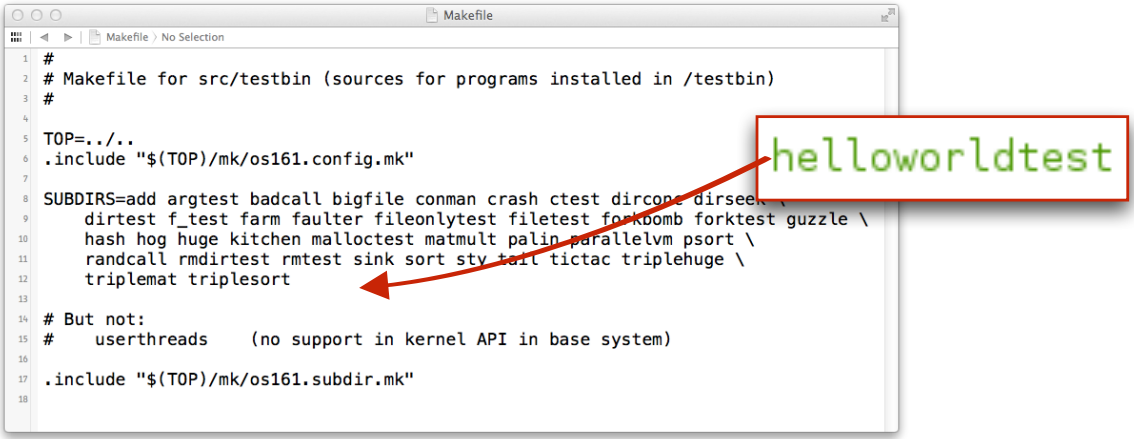

Directory tree showing main changes that need to be made

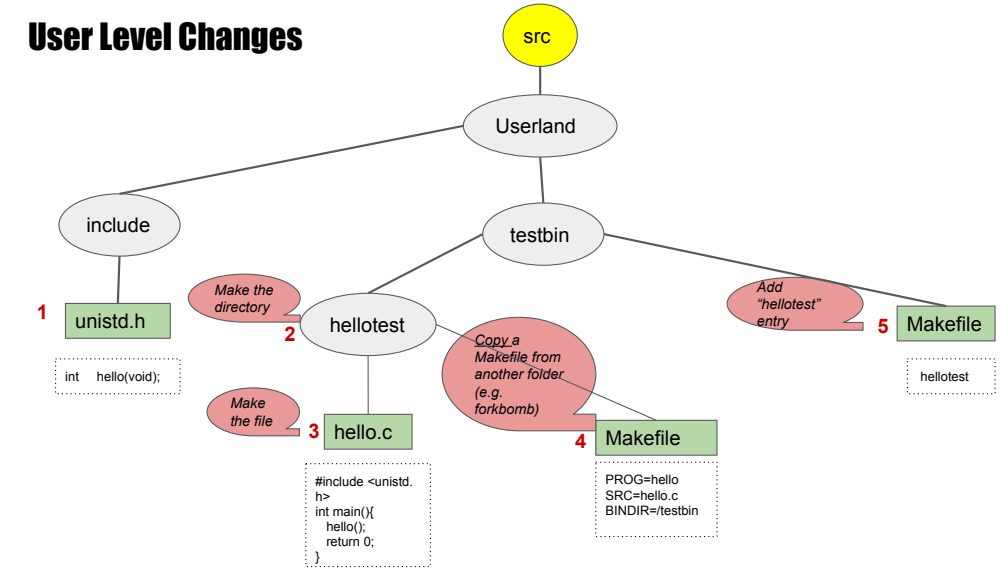

Directory tree showing main changes that need to be made

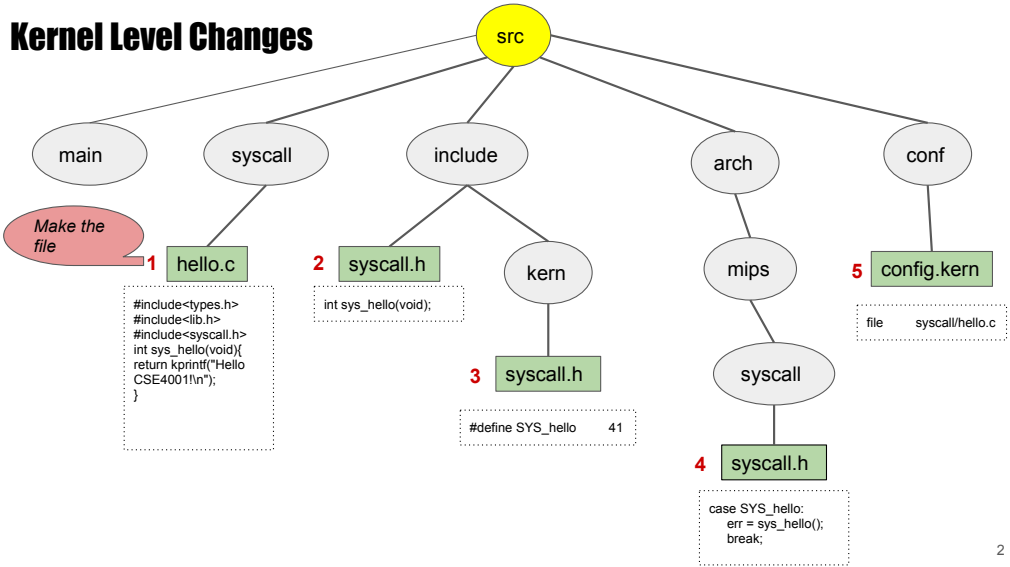

Testing the system call

Testing the system call

1. Inside the root folder, run the command `sys161 kernel`.
2. In the os161 terminal, run the command `p testbin/[name]` where your [name] is the name of your program.

Hellotest Program:

```
OS/161 kernel [? for menu]: p testbin/hellotest
Operation took 0.000145920 seconds
OS/161 kernel [? for menu]: syscall: #40, args 0 0 0 0
Hello World!
syscall: #3, args 0 0 0 0
Thread testbin/hellotest exiting due to 0 with value 0
```