


System Call operation (syscall)

Describe the steps of a system operation.



1. syscall invoked via a special CPU instruction that triggers a software trap (software interrupt)

ex. `int $0x80`: (used to invoke syscalls in intel-compatible processors)

- `lcall 7`: "execution domain specific" ← iBCS (Application binary interface)
- `lcall 27`: ← Solaris/x86 binaries
- `SYSENTER/SYSEXIT`: level 0

Alternative syscall software interrupt method for system call entry mechanism implemented available on Pentium II+ and compatible processors

2. Process making the syscall is interrupted.

3. Information needed to continue its execution later is saved (Process state saved)

4. Processor switches to higher privileged level

5. Processor determine the service being requested by the user - made by examining the process state and/or its stack

6. Execute the requested system call operation

7. Process making the syscall may be "put to sleep" if the syscall involves blocking I/O (Any time you have access to I/O)

(What happens to the process that makes a system call?)

8. When syscall completes process is "woken up" (if needed) (Moving back from block state to the ready state. One in ready state the CPU process will pick the process)

9. Original process state is restored. (Process saved will be taken back)

10. Processor switches back to lower (user) privilege level

11. Process return from syscall and continue execution

• When you make a normal function call, you save address on the stack, then you jump

What is the difference between a system call and a normal function call?

• System calls are similar but they are additional steps, invocation of special instructions that vary between OS

Process switching the service being provided for the user - where the program is in state and how it is executed

████████ (Self-Study)

Intermediate Library to Invoke syscalls

• To make it easier to invoke system calls, OS writers normally provide a library that sits between programs and system calls interface.

ex: libc: Overview of standard C libraries

glibc: GNU C Libraries

• This library provide wrapper routines

ex: truncate(): invokes the underlying

• Wrapper function is thin, doing little more than copying arguments to the right registers before invoking the system call, and then setting errno appropriately after the system call returns. (sometimes they do extra work before invoking system call)

• Wrappers hide the low-level detail of:

• Preparing arguments

• Passing arguments to kernel

• Switching to supervisor mode

(executing the `int0x` instruction, done by wrapper level routine)

• Fetching and returning results to application (Return back to user program, transition between kernel space to user space)

• Helps to reduce OS dependency and increase portability of programs.

Where does the library sit?

What is a wrapper function?

What does a wrapper

do?

Write the wrapper

function

(2) Write the wrapper

function

function

function

(3) Compiler use

the wrapper

function

function

function

function

function

function

function

function

function

function

function

function

function

function

function

function

Implementing System Calls

Steps in writing a system call

① Create an entry for the system call in the kernel's syscall-table.

- o User processor trapping to the kernel (through SYS-ENTER or int 0x80)
- find the syscall function by indexing into this table

② Write the system call code as a kernel function

- o Note: be careful when reading/writing to user space
- o Use copy-to-user or copy-from-user routines

③ Generate/Use a user-level system call stub

- o Hides the complexity of many system calls from user applications

sys ENTER / SYS EXIT Method

- o Newest and faster method to name system calls
- o Faster technique to enter and exit kernel mode that "int 0x80" software interrupt

• Plibuy: A name for some commands to be executed when you make an explicit request

- o used to avoid conflict with the file of the same name
- o used to improve performance.

- o uname -r
- o lproc cpuinfo

arch/x86/ia32/ia32_01.c

- o quad syscall (*32)
- o arch/x86/include/asm/unistd.h
- o #define __NR_00
- o arch/x86/include/asm/unistd.h
- o #define __NR_01_255
- o __syscall (NR_00, sys_00)

- o system call w/ no arg and return value (void)
- o syscall with one primitive argument

Definition of stub:

- o article containing only few sentence of text with its short and quick coverage of a subject

Ask what stub means

Which are the steps to implement a syscall

What is sysenter/sysexit methods?

What is Plibuy and what is it used for?

KERNEL Modules

- Allow code to be added to the kernel, dynamically
- Only those modules that are needed are loaded
- Enables independent development of drivers for different devices.
- Can insert modules while the system is working. This makes kernel faster since it is smaller, and only load modules when needed

Module Control

- insmod: command does insertion of a module into the linux kernel
- rmmod: command does removal of a module from the linux kernel
- lsmod: command shows the status of modules in the linux kernel

• Ext3 is a module (for example)

- You can compile something as module, some are going to be used in boot time to access to hardware, some will be used in the ram disk
- ramdisk will hold all the modules required in run-time ex. the disk driver or network driver in boot time.

• the boot time speed is almost the same every time.

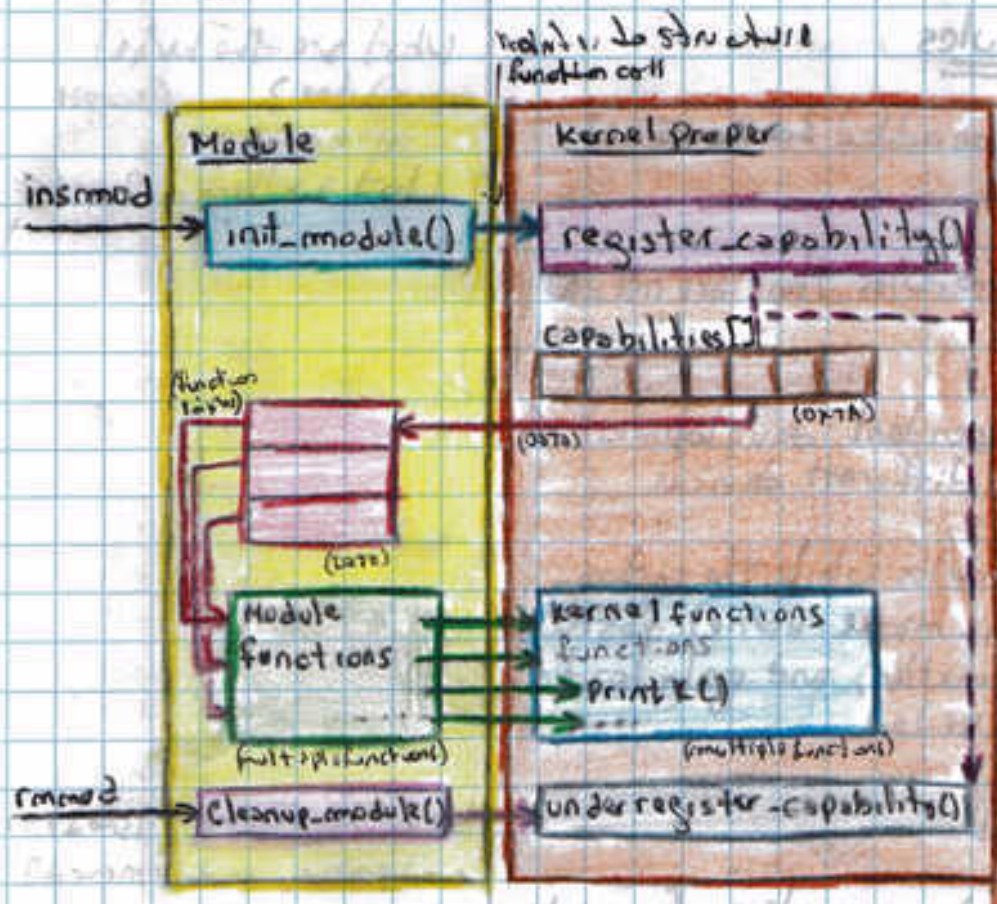
What are the kernel modules?

What do they do?

What is lsmod, lsmod, modprobe, and rmmod?

Where the modules are loaded to?

44 Workings of a Generic Module



- `init_module()`: introduce a new module to the system on which `init_module` will register some capabilities ("Here I am, here's what I can do")
- `register_capability()`: puts passing arguments ^(required) `init_module` into an internal data structure. Then the system defines a protocol for how applications get access to the information in the capabilities data structure through system calls.
- some functions are exported by the kernel's scheduling functions. Exporting means, that if some 2nd or 3rd party modules, they can see these modules and call them and forward.

How does a generic module work?
 capabilities[]
 (0370) (037A)
 (multiple functions)
 What is `init_module()`?
 What is `register_capability`?
 What does it mean
 exporting?
 What is `underregister_capability`?

Words of a generic module (continue)

- Modules are a bunch of code that doesn't need to be running
- If you decide the modules are not required by the system then remove them the kernel will clean if not done it may call function that doesn't exist in the kernel and crash the kernel

Device Classification

Character (Char) devices: the way in which an application program gains access to the services of a character device driver is through the open() system call. a serial port you can type on is a character

Block devices

- Difference with character is that you can move back and forth within the stream of characters
- Data of block devices is usually transferred in one or more blocks at a time and pre-fetched and cached
- Block devices generally have a file system associated with them whereas character devices don't

Network devices

- They differ in that they handle "packets" of data for multiple protocol clients rather than a "stream" of data for a single client
- it requires a different interface between kernel and the device driver

Character Devices

Do modules need to be running?

Character devices are those devices which are the ones that are used to describe each of those devices.

Block Devices

Block devices are those devices which are used to transfer data in one or more blocks at a time and pre-fetched and cached.

Network Devices

Network devices are those devices which are used to transfer data in packets for multiple protocol clients rather than a stream of data for a single client.

Character Devices

- Abstraction: stream of bytes
- Support: `open()`, `close()`, `read()`, and `write()`
- Typical ex: `/dev/console`, `/dev/ttyS0`
- Oddball ex: video frame grabbers: devices that captures individual, digital still frames from an analog video signal as a digital video stream
- A difference to a block device, we cannot go back^{forth} within a stream of characters. For example: in the serial device we can't go back to reread a character

Block Devices

- Abstraction: array of storage blocks
- Support: `open()`, `close()`, `read()`, `write()`, `seek()`, `mount`
- Typical Examples: `/dev/hda`, `/dev/sda`, `/dev/cram`
- Block devices can go back and forth in a stream of characters

Network Devices

- Abstraction: data packets
- Support: API for sending and receiving packets
- Support: multiple protocols and streams on one device
- Typical example: `eth0`, `eth1`,
- Network devices usually have no corresponding file system representation in Linux
 - ex: no `/dev/eth0` in file system
 - Instead use `shin/ifconfig` to check wireless network devices

Device classification

• Devices that don't fit in among char, block, or network devices

o USB

- Arguably a category by itself
- other devices operation can be emulated on top of USB such as keyboard, mouse, printer, scanner, hard disks, camera, flash drives, etc

o SCSI - Like USB, SCSI can be used in large range of devices

- most commonly used for hard disks and tape storage
- Also for scanner, CD-ROM, printers, etc

o Firewire

- Most recent serial bus interface standard
- replacing SCSI in many applications due to lower cost
- ex: Camcorders, iPods

o I20 - typically for low-end embedded system

- very simple bus specification

Filesystem modules

- Are software drivers
- Are not device drivers
- serve as a layer between user API and block devices
- Are intended to be device independent
- generally block devices have a filesystem associated with them whereas character devices don't
- proc filesystem serves as a window into the kernel and its device drivers

40 Implementing Kernel Modules

Kernel Modules

- Allow code to be added to the kernel, dynamically
- Only those modules that are needed are loaded
- Reduce kernel size
- Enables independent development of drivers for different devices

Module Control

• insmod `_____ko`

- Inserts a module (dynamic linker, call system)
- Internally, makes a call to `sys_init_module`
- Call `vmalloc()` to allocate kernel memory (Allocate enough space for memory for module)
- Copies module binary to memory
- Resolves any kernel preferences (by priority)
- via kernel symbol table (there is a table called symbol table with all the exported functions and kernel variables, it will check any call called by the linker provided by the kernel)

• modprobe `_____ko`

- Same as `insmod`, except it also loads any other modules that `_____ko` references
- Since you may have a perfect match with the table for the module, the module may fail. `modprobe` take care of intermodules dependencies.
 - If you want to look at intermodules dependencies, use `lsmod`. The modules name on the left are the one used all the modules on the right side.

self study

Module Control (continue)

- rmmod
 - remove a module
 - fails if module is still in use
 - need to remove all the dependencies before being able to remove the module
- lsmod
 - tell what modules are currently load
 - internally read /proc/modules.

What does rmmod do?

What does lsmod do?

Thing to remember

- Modules can call exported kernel functions such as `printk`, `kmalloc`, etc.
- There can't call functions that are not exported by the kernel
- Modules should not include standard libraries such as `stdio.h`, `stdlib.h`, etc.
- NOTE: Very careful to not having bad pointers, if kernel access to invalid memory then kernel crashes
- Segmentation fault may be harmless in user space but a kernel fault can crash the entire system
- Version Dependency:
 - Module should be recompiled for each version of kernel that is linked to.
 - If you compile in one version of the kernel, it may not work with another version of the kernel.

Name some kernel functions.

Do modules include standard libraries

What happens when there is a segmentation fault at kernel level?

What does it mean to update the kernel?

- Take care of concurrencies. When program regular program, you have to flow from the top to the bottom, however, in kernel program you may not have a sequential execution.
- After a module is called, it will put information into memory so it can be used by the kernel any time.
- Be very careful with data structures, if both functions access to a data struct at the same time without hierarchy access, you may encounter a problem, that is why we use locks.

What would happen if more than one function would try to access the same data structure

Concurrency Issues

- Many processes could try to access your module concurrently so different parts of your module may be active at the same time.
- The difference between simultaneous process and concurrently process:

What are the difference between simultaneous and concurrent

Ex: let say you have one CPU, but you can run 100 processes that looks like they run at the same time, this is concurrency.

b. But simultaneous, there are in a table of execution, two processes can be running at the same time in different CPUs therefore simultaneous means at the same time (this is related with threads)

What is the difference between simultaneous and concurrent

- Device interrupts can trigger Interrupt Service Routines (ISR)
- ISRs may access common data that your module uses as well (these can be tricky)

What is ISR?

Concurrency Issues (continue)

- kernel timers can concurrently execute with your module and access common data.
- kernel's timers are normally timer schedules to execute in certain time.
- these kernel's timers can be triggered at the same time, so it can create some bad conditions especially with multi-processor system.

What are the kernel timers?
How do the kernel timers work?

- You may have symmetric multi-processor (SMP) system, so multiple processors may be executing your module code simultaneously (not just concurrently) therefore

What do your module have to manage?

- Therefore, your module (and most kernel code in general) should be re-entrant - capable of correctly executing in more than one context simultaneously.
- The correctness of the function is not get affected because is called many times at the same time. Same piece of code can get called again from a different context.

Can you have concurrency with a single CPU?

- Even with a single CPU you can have concurrency because you can have a CPU interruption at the same time have a I/O interruption.

Can you have concurrency with a single CPU?

When can you have concurrency with a single CPU?

52 Version Numbering

- Kernel components have inter-dependencies
- Version numbers help to track these dependencies
- Well written modules have minimal dependencies
 - They should be backwards compatible with all versions.
- Linux kernel version numbers:
 - Major, minor, release (eg. 2.6.20)
 - Odd numbered minors are for developers only
 - Even numbers are supposed to be stable enough for general use

GNU General Public License (GPL)

- Basis for all GNU software development, includes linux
 - MIT licence, you can mod. Bin. code and not release
- GPL 2, if you take a little code and mod. it to call it your own, you must release.
- GPL Allow users to modify software as see the need
- GPL 3 is more restricted
- "Infects" transitively through modification and distribution by multiple parties
- Device drivers need not be licensed under GPL, but the mainstream ones are

Memory Management (CHP 4)


- Ideally programmers want memory that is
 - o large
 - o fast
 - o not volatile

What does ideally programmer wants?

- Memory hierarchy
 - o Small amount of fast, expensive memory (cache)
 - o Some medium-speed, medium price memory
 - o Gigabytes of slow, cheap disk storage
- Memory management handles the memory hierarchy

How is the memory hierarchy?

Typical Memory Hierarchy



Basic Memory Management

• Mono programming without swapping or paging
↳ No memory abstraction - simple physical memory

• Three ways of organizing memory:


- o OS at the bottom of memory ram
- o OS at the top of memory ram
- o OS device drivers at top of memory ram and rest of the OS in RAM below

Which are the three ways of organizing memory with basic memory management?

• Generally only one process at a time can be running

How many processes could be running generally?

54 a) OS at bottom of memory in RAM (Random Access Mem)



• Used on mainframes and minicomputers

you have OS at the bottom in RAM, what kind of computer would normally work with this? OS managed?

b) OS in ROM (Read-Only Memory) at top of memory



• Handheld computers and embedded systems

c) Device drivers at top of memory (ROM), and the rest of OS at RAM down below



• Early personal computers

• One way to get some parallelism in a system with no memory abstraction is to program with multiple threads.

• limited use since people often want is unrelated

• programs to be running at the same time

• Threads abstraction does not provide this

to do parallelism. Describe one way to do parallelism in a system with no memory abstraction.

self study

(writing) and the other side of the page

Running Multiple Programs Without a Memory Abstraction

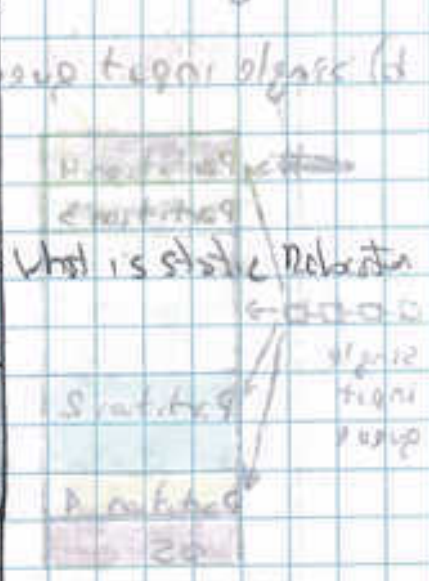
- OS has to save the entire content of memory to a disk file, then bring in and run the next program.
- As long as there is only one program at the time in memory, there are no conflicts (concept called swapping)

- using special hardware it is possible to run multiple programs concurrently using a technique known as static relocation.

↳ Static Relocation: When a program was loaded at address 16,384, the constant 16,384 was added to every program address during load process. While works fine, it is not a very general solution and slow down loading.

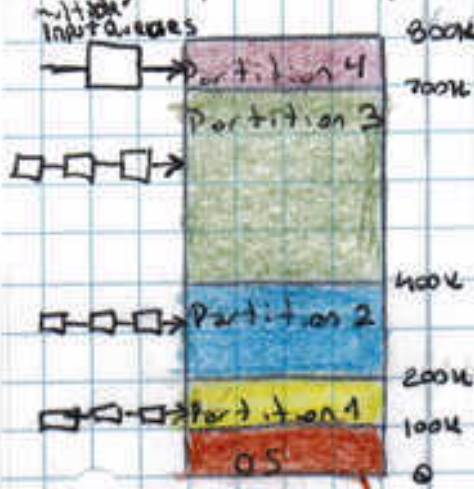
How can we run multiple programs without a memory abstraction?

What is static relocation?



Multiprogramming with Fixed Partitions

a) Separate input queues for each partition (not)



- Different size partitions
- Called continuous Multiprogramming with Fixed number of Tables or OS/MFT
- Once partition has taken certain size then it remains at that size
- Jobs arrive it is placed in the input queue in the smallest partition that will accommodate it.

How does separate input queues work with fixed partitions?

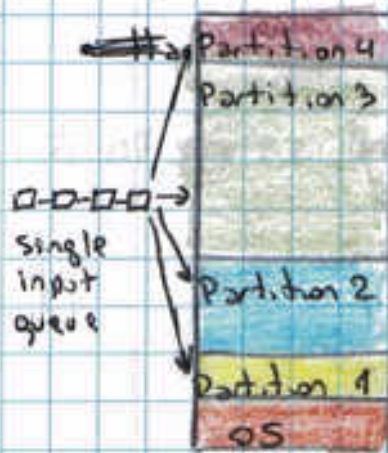
56 Multi-programming with fixed Partitions (continue)

Drawbacks: (disadvantage)

1. As the partition sizes are fixed, any space not used by a particular job is lost
2. It may not be easy to state how big a partition a particular job needs
3. If job is placed in (say) queue three it may be prevented from running by other jobs waiting (and using) that partition

What are the drawbacks of separate input queues?

b) Single input queue feature



- Single input queue where all jobs are held
- When a partition is free, we search the queue looking for the first job that fits into the partition
- Alternative is to look for the largest job that fits into the partition

What are the advantages of single input queue?

Advantage: Not wasted large partition on small job

Disadvantage: Smaller jobs are discriminated against



self study

Relocation and Protection


When introducing multiprogramming, two problems to address:

1. Relocation: When a program is run it does not know in advance what location it will be loaded at; therefore, the program cannot simply generate static addresses (i.e. jump instructions). Instead, they must be made relative to where the program has to be loaded.
 ↳ Address location of variables - code routines cannot be absolute.

2. Protection: Once you have two programs in memory at the same time there is a danger that one program can write to the address space of another program which is dangerous and should be avoided.
 ↳ Must keep a program out of other processes partition

What is Relocation?

What is protection?



• Relative address in original program binary