

Linux

CS 256/456

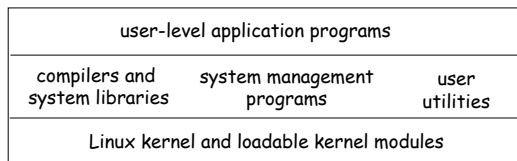
Dept. of Computer Science, University of Rochester

History

- Linux is a modern, open-source operating system that is mostly POSIX-compliant.
- First developed as a small but self-contained kernel in 1991 by Linus Torvalds, with the major design goal of UNIX compatibility.
- Collaboration by many users all around the world, corresponding almost exclusively over the Internet.

Linux Kernel/System/Distribution

- Kernel
 - the OS code that runs on privileged mode
- System
 - essential system components, but runs in user mode
 - compilers, system libraries
- Linux distribution
 - extra system-installation and management utilities
 - precompiled and ready-to-install tools & packages
 - popular distributions: Redhat/Fedora, Debian, SuSE, Caldera, ...



Processes and Threads

- Linux uses the same internal representation for processes and threads; a thread is simply a new process that happens to share the same address space as its parent.
 - getpid() semantics?
- A distinction is only made when a new thread is created by the **clone** system call.
 - a process is a task with its own entirely new context (including address space)
 - a thread is a task with its own identity, but not an dedicated address space
- The **clone** system call allows fine-grained control over exactly what is shared between two threads.
 - open files, memory space, page tables

Linux Task Scheduling

- Linux uses two task-scheduling classes:
 - time-sharing and real-time
- A prioritized, epoch-based algorithm for time-sharing
 - Each task has a static priority (default=20) and a dynamic quantum
 - Scheduling is prioritized based on quantum at the beginning of each epoch; quantum of the running task decrements by one at every clock tick
 - Recrediting when no runnable tasks have any quantum (end of an epoch):
$$\text{initial quantum in new epoch} = \frac{\text{remaining quantum}}{2} + \text{priority}$$
 - The initial process quantum at its first epoch is priority.
 - This quantum crediting system automatically prioritizes interactive or I/O-bound tasks.

4/24/2006

CSC 256/456 - Spring 2006

5

Linux Task Scheduling: O(1) Scheduler

- Linux O(1) scheduler
 - the scheduling overhead is constant, which is independent of the number of processes in the system
- Main operations in Linux scheduler
 - schedule(), quantum recalculation, wakeup()
- Using two priority arrays
 - one for active array, one for those whose quantum has been used up (called "expired")
 - array index indicates the priority
 - O(1) operation to find the highest priority in a task array: use array index bitmap and BSFL

4/24/2006

CSC 256/456 - Spring 2006

6

Interrupt Handling

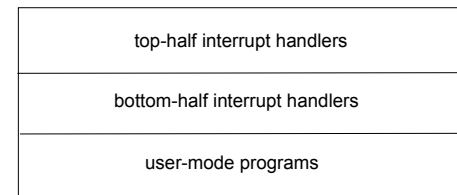
- Interrupt handling is usually non-reentrant
 - new interrupts are disabled during the handling of an old interrupt
- Linux's kernel allows long interrupt service routines to run without having interrupts disabled for too long
- Interrupt service routines are separated into a *top half* (urgent) and a *bottom half* (not so urgent).
 - The top half runs with interrupts disabled.
 - The bottom half is run later, with all interrupts enabled.
 - Bottom halves run one by one (they do not interrupt each other).

4/24/2006

CSC 256/456 - Spring 2006

7

Interrupt Protection Levels



- Each level may be interrupted by code running at a higher level, but will never be interrupted by code running at the same or a lower level.
- where is system call handling?

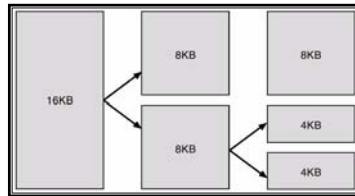
4/24/2006

CSC 256/456 - Spring 2006

8

Managing Physical Memory

- Keep track of free memory?
- Linux page allocator can allocate ranges of physically-contiguous pages on request.
- The allocator uses a *buddy-heap* algorithm to keep track of available physically-contiguous memory regions.
 - A free region list is maintained for each region size: 4KB, 8KB, 16KB, 32KB,
 - A large region can be split into multiple smaller regions if necessary



4/24/2006

CSC 256/456 - Spring 2006

9

Memory Page Replacement

- All memory pages are managed together
 - stack/heap/code, ...
 - file system buffer cache
- Memory pages are managed in two LRU lists: active list and inactive list
 - each LRU list is managed using a CLOCK (second-chance) LRU approximation
 - pages evicted from the active list go to the inactive list; pages evicted from the inactive list are out of the system
 - pages in the inactive list may be promoted to the active list under certain circumstance

4/24/2006

CSC 256/456 - Spring 2006

10

Ext2fs File System

- Disks are divided into contiguous block groups
 - the hope is that there is not much seeking within each block group (but no guarantee)
 - there is a section for inodes in each block group
 - the FS tries to keep inodes and corresponding file blocks in the same block group
- Ext2fs uses allocation policies designed to place logically adjacent blocks of a file into physically adjacent blocks on disk
 - with the help of the free block bitmap
- Ext3fs supporting file system journaling

4/24/2006

CSC 256/456 - Spring 2006

11

The Linux /proc File System

- The **proc** file system does not store data, rather, its contents are computed on demand according to user file I/O requests.
 - When data is read from one of these files, **proc** collects the appropriate information, formats it into text form and places it into the requesting process's read buffer.
- /proc is not a unique feature on Linux

4/24/2006

CSC 256/456 - Spring 2006

12

Prefetching and I/O Scheduling

- File prefetching/read-ahead
 - prefetching sequentially when the I/O access is considered as sequential
 - how to detect sequential pattern?
- Disk I/O scheduling
 - an elevator-style seek-reduction scheduling
 - non-work conserving scheduling: anticipatory scheduling
 - deadline to prevent starvation

Kernel Modules

- Sections of kernel code that can be compiled, loaded, and unloaded independent of the rest of the kernel.
- A kernel module may typically implement a device driver, a file system, or a networking protocol.
- Kernel modules allow a Linux system to be set up with a standard, minimal kernel, without any extra device drivers built in.
- Different from user-provided kernel extensions?

Disclaimer

- Parts of the lecture slides contain original work of Abraham Silberschatz, Peter B. Galvin, Greg Gagne, Andrew S. Tanenbaum, and Gary Nutt. The slides are intended for the sole purpose of instruction of operating systems at the University of Rochester. All copyrighted materials belong to their original owner(s).