

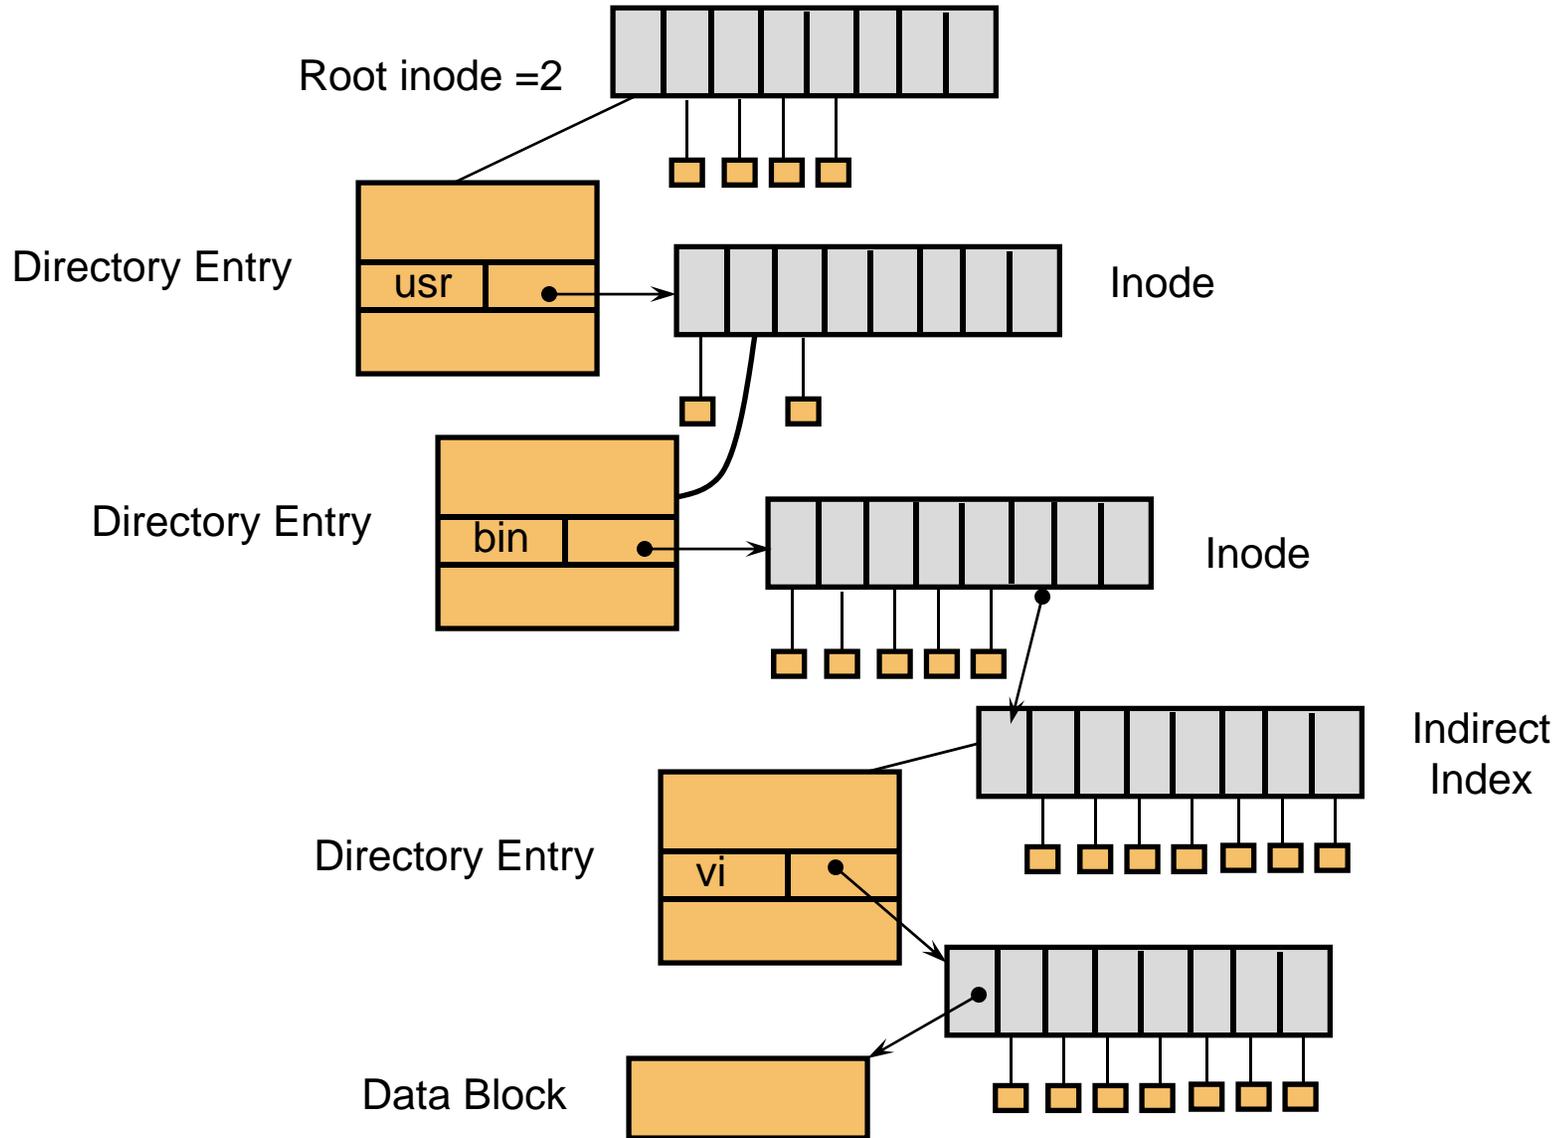
Announcements

- Reading Chapter 12
- Project #4 is Due Friday at 5:00 PM

Unix directories - links

- Each file has unique inode but it may have multiple directory entries in the same filesystem to reference inode
- Each directory entry creates a hard link of a filename to the file's inode
 - Number of links to file are kept in reference count variable in inode
 - If links are removed, file is deleted when number of links becomes zero
- **Symbolic or soft link**
 - Implemented as a file that contains a pathname
 - Symbolic links do not have an effect on inode reference count

File Lookup (/usr/bin/vi)



Using UNIX filesystem data structures

- Example: `find /usr/bin/vi`
 - from Leffler, McKusick, Karels and Quarterman
 - Search root directory of filesystem to find `/usr`
 - root directory inode is, by convention, stored in inode #2
 - inode shows *where data blocks are* for root directory - *these blocks* (not the inode itself) *must* be retrieved and searched for entry `user`
 - we discover that the directory `user's` inode is inode #4
 - Search `user` for `bin`
 - access blocks pointed to by inode #4 and search contents of blocks for entry that gives us `bin's` inode
 - we discover that `bin's` inode is inode #7
 - Search `bin` for `vi`
 - access blocks pointed to by inode #7 and search contents of block for an entry that gives us `vi's` inode
 - we discover that `vi's` inode is inode #7
 - Access inode #7 - this is `vi's` inode

How to Improve Speed?

- Use A Cache
- Name-to-Inode lookup
 - Hash on full path name
 - Find inode without and disk accesses on a hit

Mount System Call

- How to attach a file system into a name space?
- Simple Idea:
 - use letters C, D, E, etc.
 - use volume names (VMS) – fixed length string
- Better Idea:
 - Allow attachment at arbitrary points in namespace
 - Designate one tree as the “root” file system
 - Others are attached to the root
- Mount used in:
 - UNIX
 - Windows (NTFS mount points)
 - GeekOS

Log Structured File Systems

- Key Idea
 - Use transactions like model for filesystem updates
- Write data to a log (also called a journal)
 - Records meta data changes
 - Records data blocks written
 - File operation is committed once it is to the log
 - Partial updates to log are lost on failure
- Next Step
 - Eliminate the filesystem and just keep the log
 - Requires a process called a cleaner
 - Copies old data from log to head of log to allow compaction

NTFS

- File system may
 - Be a partition (fraction) of a disk
 - May span multiple disks
- Clusters
 - Group sectors into a larger group (typically 4KB)
 - Logical cluster numbers (0...N) describe where a cluster is
- File consists of a set of attributes
 - Attributes
 - arbitrary sized
 - Linear ordering from 0...n
 - Examples
 - Filename
 - File data
 - Security
 - Mac Resource fork

NTFS Files

- Each file is stored in an entry in the Master File Table (MFT)
 - Each entry 1-4KB
 - Small attributes stored directly in MFT
 - Larger attributes are stored in one or more extents (contiguous clusters on the disk)
- Special Files
 - MFT – file 0
 - Copy of first 16 entries in MFT
 - Log file – log of changes to file system
 - Attribution definition table
 - Root directory
 - Bitmap free list
 - Boot file (must be at a standard disk address)
 - Bad cluster file