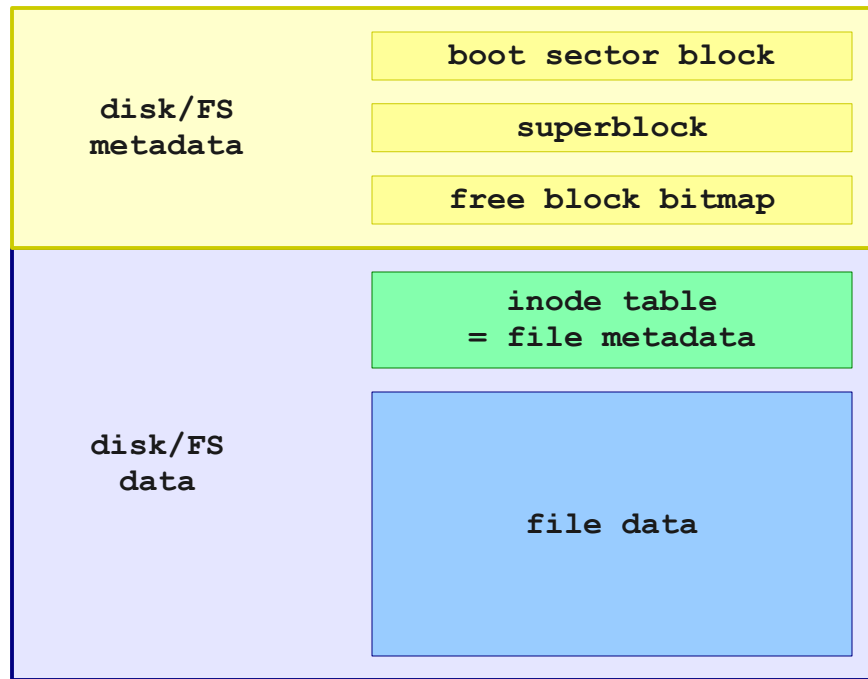

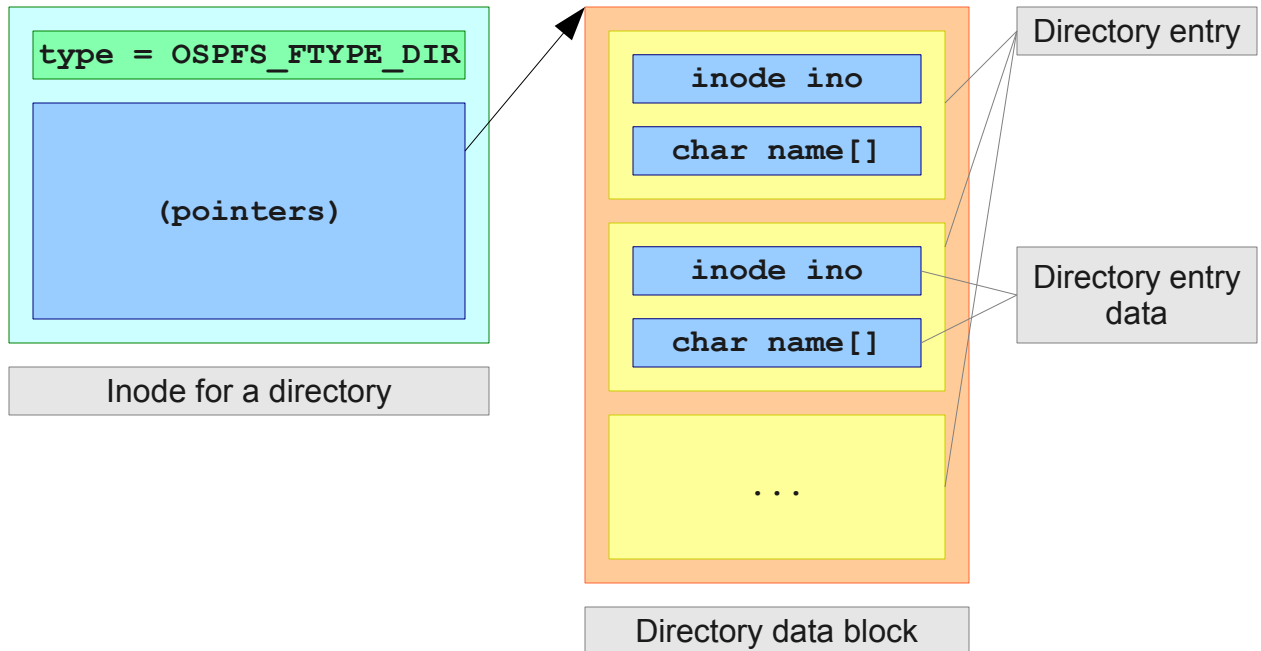


Lab 3	File System Overview	1/16
<div><div>What do file systems do?<ul style="list-style-type: none">– Organize storage (and devices)– Manages storage to implement the file interface– Gives names to files to be easy to access</div><div>Why do we need them?<ul style="list-style-type: none">– Difficult to deal with sectors directly or manually– Difficult to grow and shrink space directly– Difficult to remember sector names</div><div>⇒ Abstracts storage to be usable by humans</div></div>		





Lab 3	Allocation Problem	5/16
<ul style="list-style-type: none">– Disk has 3920140 sectors \approx 2 GB– Each block is 8 sectors– Average file uses 3 data blocks and one inode; no ind, ind² blocks– Each inode is 256 B– One boot sector block– One superblock <p>Q: How big should our inode table be, in blocks, to maximize the number of files the file system can hold?</p>		



Lab 3	Metadata for lab 3	7/16
<div><div><div>File system level: superblock</div><div><div>– os_magic</div><div># marker</div></div><div><div>– os_nblocks</div><div># blocks</div></div><div><div>– os_ninodes</div><div># inodes</div></div><div><div>– os_firstinob</div><div># offset</div><div>first inode block</div></div></div><div><div>File level: inode (regular files, dirs)</div><div><div>– oi_size</div><div>file size</div></div><div><div>– oi_ftype</div><div>file, dir, symlink</div></div><div><div>– oi_nlink</div><div># of hardlinks</div></div><div><div>– oi_mode</div><div>file permissions</div></div><div>(dir, ind, ind² block pointers)</div></div></div>		

Lab 3	Links I	8/16
<ul style="list-style-type: none">– Often useful to make a file available in multiple directories<ul style="list-style-type: none">– For example, I have my .emacs in a dir that I back up– Also useful for sharing code between projects<ul style="list-style-type: none">– For example, project FOO has write access to com/example/FOO and BAR wants to use com/example/FOO/FooClass.java that project FOO maintains <p>Issues with duplication:</p> <ol style="list-style-type: none">1. Wasted space2. Updates of multiple copies		

Lab 3	Links II	9/16
<p data-bbox="256 344 477 378">Solution: links</p> <p data-bbox="256 415 943 449">Hard links: every inode has an nlinks counter</p> <ul data-bbox="350 487 1230 735" style="list-style-type: none"><li data-bbox="350 487 873 520">– When file is first created, set to 1<li data-bbox="350 558 1138 592">– Increment whenever we link from another location<li data-bbox="350 630 1230 663">– Decrement whenever we delete a filename for that inode<ul data-bbox="444 701 1029 735" style="list-style-type: none"><li data-bbox="444 701 1029 735">– When nlinks = 0, can free disk space <p data-bbox="256 842 1122 875">Q: can't (in ext4, anyway) hard link a directory—why not?</p>		

symlinks:

- Contains a pathname**
- When we open the symlink:**
 - The OS looks up the pathname**
 - If valid, it opens the file at that pathname**
 - If we move, delete, rename the referenced file, we must manually update the symlink**
- Q: It's okay to symlink directories—why?**
 - When we delete a symlink, we delete the link, but not the linked file!**

Loose C++ analogy: hard links = references, symlinks = pointers

Filesystems don't need to be just interfaces to disks:

- /dev
 - devices on the system
 - /dev/urandom: (pseudo) random byte pool
 - /dev/stdout: the standard output
 - symlinked to /proc/self/fd/1
- /proc
 - interface to kernel objects
 - /proc/cpuinfo
 - info about the CPUs, including [bogomips](#)

Filesystems can store data indirectly to disks:


- sshfs**
 - Mounts a remote directory over SSH (secure)**
 - A lot easier than using scp for moving a lot of files**
- GMailFS**
 - Stores data as email messages on a Gmail account**
 - Free space, but violates terms of service**

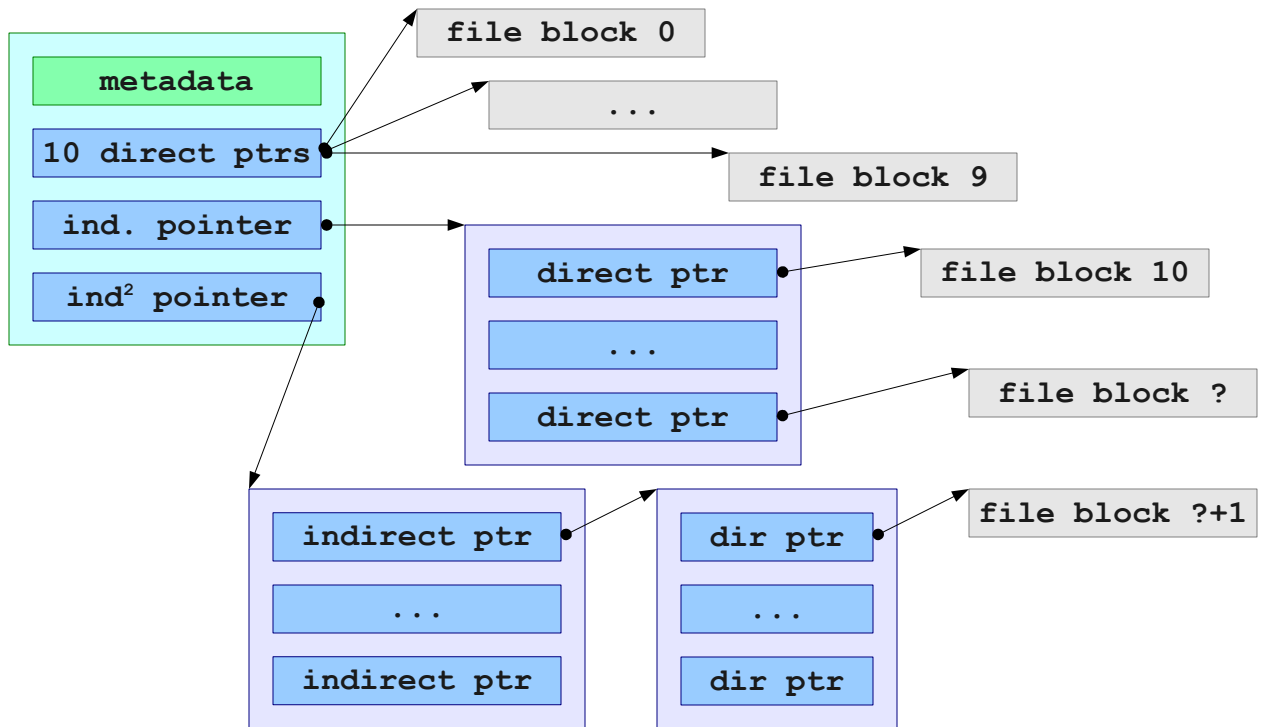
Filesystems can add functionality:

- Easy revision control
 - Files have versions
 - Current version
 - Can roll back to previous versions
 - Perhaps using `ioctl()`

Q: What extra information would we need to implement revision control in a file system?

Some examples of file systems that add functionality are [here](#)





d	i	i^2	Meaning
x	-1	-1	referenced by direct pointer x in inode
x	0	-1	ref'd by direct ptr x in inode's indirect block
x	y	0	ref'd by direct ptr x in the (indirect block ref'd by ptr y in inode's indirect ² block)