

CS 3204 Operating Systems

Lecture 31
Godmar Back



Announcements

- Project 3 due April 13
- Additional Office Hours today 3-4pm



CS 3204 Spring 2006

4/7/2006

2

Disks & Filesystems

Buffer Cache



Disk Caching – Buffer Cache

- How much memory should be dedicated for it?
- How should eviction be handled?
- How should prefetching be done?
- How should concurrent access be mediated?
 - How is consistency guaranteed? (All must go through buffer cache!)
- What write-back strategy should be used?



CS 3204 Spring 2006

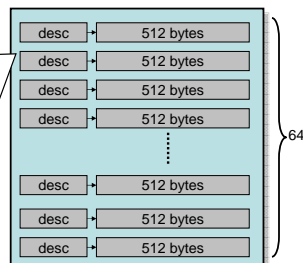
4/7/2006

4

Buffer Cache in Pintos

Cache Block Descriptor

- disk_sector_id, if in use
- dirty bit
- valid bit
- # of readers
- # of writers
- # of pending read/write requests
- lock to protect above variables
- signaling variables to signal availability changes
- usage information for eviction policy
- data (pointer or embedded)



CS 3204 Spring 2006

4/7/2006

5

A Buffer Cache Interface

```
// cache.h
struct cache_block; // opaque type
// reserve a block in buffer cache dedicated to hold this sector
// possibly evicting some other unused buffer
// either grant exclusive or shared access
struct cache_block * cache_get_block(disk_sector_t sector, bool exclusive);
// release access to cache block
void cache_put_block(struct cache_block *b);
// read cache block from disk, returns pointer to data
void *cache_read_block(struct cache_block *b);
// fill cache block with zeros, returns pointer to data
void *cache_zero_block(struct cache_block *b);
// mark cache block dirty (must be written back)
void cache_mark_block_dirty(struct cache_block *b);
// not shown: initialization, readahead, shutdown
```



CS 3204 Spring 2006

4/7/2006

6

Buffer Cache Rationale

Compare to buffer pool assignment in CS2604

Differences:

- Do not combine allocating a buffer (a resource management decision) with loading the data into the buffer from file (which is not always necessary)
- Provide a way for buffer user to say they're done with the buffer
- Provide a way to share buffer
- More efficient interface (opaque type instead of block idx saves lookup, constant size buffers)

```
class BufferPool { // (2) Buffer Passing
public:
    virtual void* getblock(int block) = 0;
    virtual void dirtyblock(int block) = 0;
    virtual int blocksize() = 0;
};
```



CS 3204 Spring 2006

4/7/2006

7

Buffer Cache Sizing



- Simple approach
 - Set aside part of physical memory for buffer cache/use rest for virtual memory pages as page cache – evict buffer/page from same pool
- Disadvantage: can't use idle memory of other pool - usually use unified cache subject to shared eviction policy
- Windows allows user to limit buffer cache size
- Problem:
 - Bad prediction of buffer caches accesses can result in poor VM performance (and vice versa)



CS 3204 Spring 2006

4/7/2006

8

Buffer Cache Replacement

- Similar to VM Page Replacement, differences:
 - Can do exact LRU (because user must call `cache_get_block(!)`)
 - But LRU hurts when long sequential accesses – should use MRU (most recently used) instead.
- Example reference string: ABCDABCDABCD, can cache 3:
 - LRU causes 12 misses, 0 hits, 9 evictions
 - How many misses/hits/evictions with MRU?
- Also: not all blocks are equally important, benefit from some hits more than from others



CS 3204 Spring 2006

4/7/2006

9

Buffer Cache Writeback Strategies

- Write-Through:
 - Good for floppy drive, USB stick
 - Poor performance – every write causes disk access
- (Delayed) Write-Back:
 - Makes individual writes faster – just copy & set bit
 - Absorbs multiple writes
 - Allows write-back in batches
- Problem: what if system crashes before you've written data back?
 - Trade-off: performance in no-fault case vs. damage control in fault case
 - If crash occurs, order of write-back can matter



CS 3204 Spring 2006

4/7/2006

10

Writeback Strategies (2)

- Must write-back on eviction (naturally)
- Periodically (every 30 seconds or so)
- When user demands:
 - `fsync(2)` writes back all modified data belonging to one file – database implementations use this
 - `sync(1)` writes back entire cache
- Some systems guarantee write-back on file close



CS 3204 Spring 2006

4/7/2006

11