

CMU SCS

Carnegie Mellon Univ.
Dept. of Computer Science
15-415 - Database Applications

Lecture #8 (R&G ch9)
Storing Data: Disks and Files

Faloutsos CMU SCS 15-415 #1

CMU SCS

Overview

- Memory hierarchy
- RAID (briefly)
- Disk space management
- Buffer management
- Files of records
- Page Formats
- Record Formats

Faloutsos CMU SCS 15-415 #2

CMU SCS

DBMS Layers:

Queries

Query Optimization and Execution

Relational Operators

Files and Access Methods

Buffer Management

Disk Space Management

DB

TODAY →

Faloutsos CMU SCS 15-415 #3

 CMU SCS

Leverage OS for disk/file management?

- Layers of abstraction are good ... but:

Faloutsos CMU SCS 15-415 #4

 CMU SCS

Leverage OS for disk/file management?

- Layers of abstraction are good ... but:
 - Unfortunately, OS often **gets in the way** of DBMS

Faloutsos CMU SCS 15-415 #5

 CMU SCS

Leverage OS for disk/file management?

- DBMS wants/needs to do things “its own way”
 - **Specialized prefetching**
 - **Control over buffer replacement policy**
 - LRU not always best (sometimes worst!!)
 - **Control over thread/process scheduling**
 - “Convoy problem”
 - Arises when OS scheduling conflicts with DBMS locking
 - **Control over flushing data to disk**
 - WAL protocol requires flushing log entries to disk

Faloutsos CMU SCS 15-415 #6

CMU SCS

Disks and Files

- DBMS stores information on disks.
 - but: disks are (relatively) VERY slow!
- Major implications for DBMS design!



Faloutsos CMU SCS 15-415 #7

CMU SCS

Disks and Files

- Major implications for DBMS design:
 - **READ**: disk -> main memory (RAM).
 - **WRITE**: reverse
 - Both are high-cost operations, relative to in-memory operations, so must be planned carefully!

Faloutsos CMU SCS 15-415 #8

CMU SCS

Why Not Store It All in Main Memory?

Faloutsos CMU SCS 15-415 #9

CMU SCS

Why Not Store It All in Main Memory?

- *Costs too much.*
 - disk: ~\$1/Gb; memory: ~\$100/Gb
 - High-end Databases today in the 10-100 TB range.
 - Approx 60% of the cost of a production system is in the disks.
- *Main memory is volatile.*
- *Note:* some specialized systems do store entire database in main memory.

Faloutsos CMU SCS 15-415 #10

CMU SCS

The Storage Hierarchy

Smaller, Faster

Bigger, Slower

Faloutsos CMU SCS 15-415 #11

CMU SCS

The Storage Hierarchy

Smaller, Faster

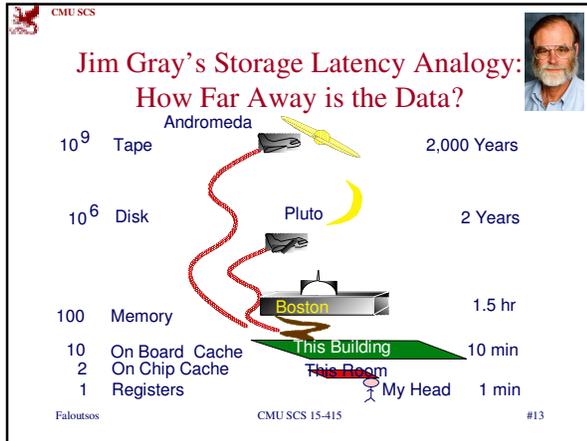
–Main memory (RAM) for currently used data.

–Disk for the main database (secondary storage).

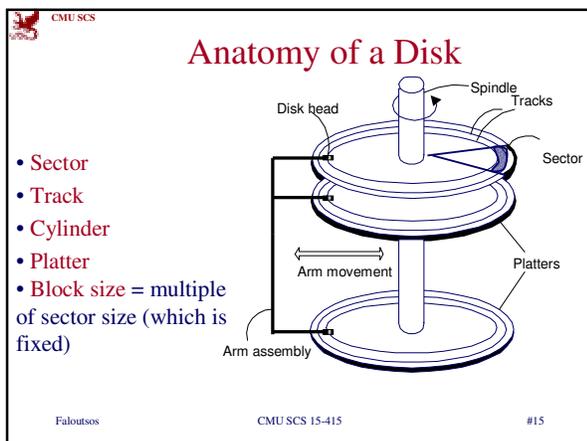
–Tapes for archiving older versions of the data (tertiary storage).

Bigger, Slower

Faloutsos CMU SCS 15-415 #12



- ### Disks
- Secondary storage device of choice.
 - Main advantage over tapes: *random access* vs. *sequential*.
 - Data is stored and retrieved in units called *disk blocks* or *pages*.
 - Unlike RAM, time to retrieve a disk page varies depending upon location on disk.
 - relative placement of pages on disk is important!
- Faloutsos CMU SCS 15-415 #14



CMU SCS

Accessing a Disk Page

- Time to access (read/write) a disk block:
 - .
 - .
 - .

Faloutsos CMU SCS 15-415 #16

CMU SCS

Accessing a Disk Page

- Time to access (read/write) a disk block:
 - *seek time*: moving arms to position disk head on track
 - *rotational delay*: waiting for block to rotate under head
 - *transfer time*: actually moving data to/from disk surface

Faloutsos CMU SCS 15-415 #17

CMU SCS

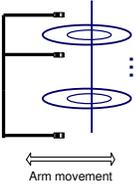
Seek Time

The diagram shows a disk head moving across tracks. The graph plots Time on the y-axis (with markers x, 3x to 20x) against Cylinders Traveled on the x-axis (with markers 1 and N). Three curves are shown: A? (steepest), B? (middle), and C? (shallowest). A double-headed arrow labeled 'Arm movement' indicates the distance between tracks.

Faloutsos CMU SCS 15-415 #18

CMU SCS

Seek Time



Time

3x to 20x

x

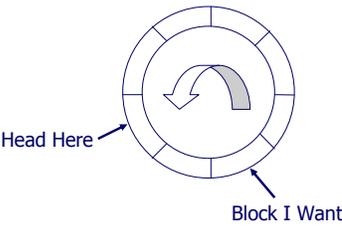
1 N

Cylinders Traveled

Faloutsos CMU SCS 15-415 #19

CMU SCS

Rotational Delay



Head Here

Block I Want

Faloutsos CMU SCS 15-415 #20

CMU SCS

Accessing a Disk Page

- Relative times?
 - seek time:
 - rotational delay:
 - transfer time:

Faloutsos CMU SCS 15-415 #21

CMU SCS

Accessing a Disk Page

- Relative times?
 - *seek time*: about 1 to 20msec
 - *rotational delay*: 0 to 10msec
 - *transfer time*: < 1msec per 4KB page

Seek

Rotate

transfer

Faloutsos CMU SCS 15-415 #22

CMU SCS

Seek time & rotational delay dominate

- Key to lower I/O cost: **reduce seek/rotation delays!**
- Also note: For shared disks, much time spent waiting in queue for access to arm/controller

Seek

Rotate

transfer

Faloutsos CMU SCS 15-415 #23

CMU SCS

Arranging Pages on Disk

- “*Next*” block concept:
 - blocks on same track, followed by
 - blocks on same cylinder, followed by
 - blocks on adjacent cylinder
- Accessing ‘next’ block is cheap
- An important optimization: pre-fetching
 - See R&G page 323

Faloutsos CMU SCS 15-415 #24

CMU SCS

Rules of thumb...

1. Memory access much faster than disk I/O (~ 1000x)
- “Sequential” I/O faster than “random” I/O (~ 10x)

write on blackboard

Faloutsos CMU SCS 15-415 #25

CMU SCS

Overview

- Memory hierarchy
- RAID (briefly)
- Disk space management
- Buffer management
- Files of records
- Page Formats
- Record Formats

Faloutsos CMU SCS 15-415 #26

CMU SCS

Disk Arrays: RAID

Logical Physical

- Benefits:
 - Higher throughput (via data “striping”)
 - Longer MTTF (via redundancy)

Faloutsos CMU SCS 15-415 #27

CMU SCS

Overview

- Memory hierarchy
- RAID (briefly)
- Disk space management
- Buffer management
- Files of records
- Page Formats
- Record Formats

Faloutsos CMU SCS 15-415 #28

CMU SCS

Disk Space Management

- Lowest layer of DBMS software manages space on disk
- Higher levels call upon this layer to:
 - allocate/de-allocate a page
 - read/write a page
- Best if requested pages are stored **sequentially** on disk! Higher levels don't need to know if/how this is done, nor how free space is managed.

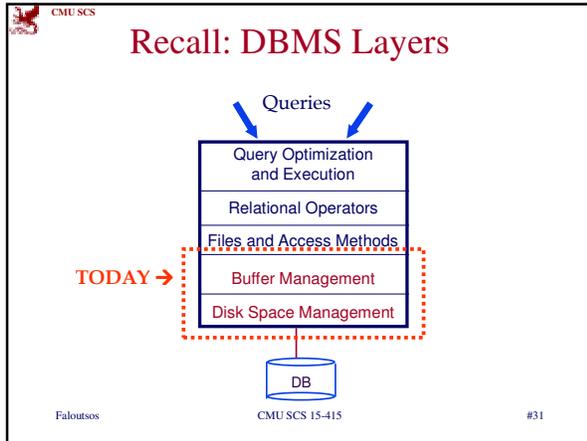
Faloutsos CMU SCS 15-415 #29

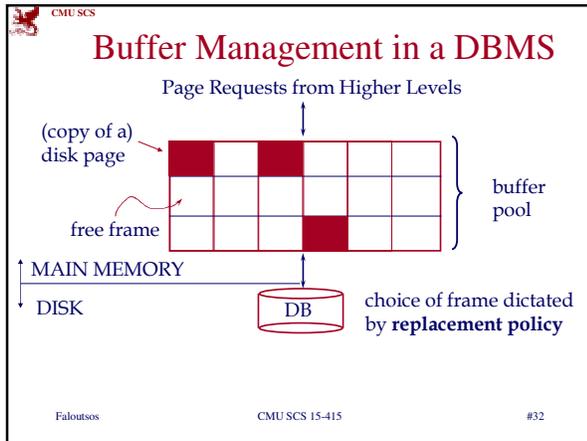
CMU SCS

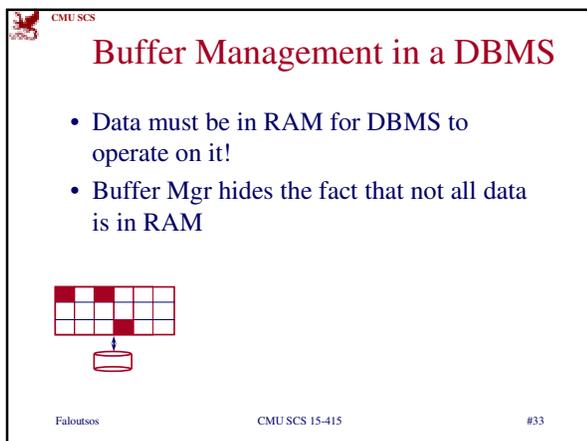
Overview

- Memory hierarchy
- RAID (briefly)
- Disk space management
- Buffer management
- Files of records
- Page Formats
- Record Formats

Faloutsos CMU SCS 15-415 #30







CMU SCS

When a Page is Requested ...

Buffer pool information table contains:
 <frame#, pageid, pin_count, dirty-bit>

- If requested page is not in pool:
 - Choose an (un-pinned) frame for *replacement*
 - If frame is “dirty”, write it to disk
 - Read requested page into chosen frame
- *Pin* the page and return its address

Faloutsos CMU SCS 15-415 #34

CMU SCS

When a Page is Requested ...

- If requests can be predicted (e.g., sequential scans)
- then pages can be pre-fetched several pages at a time!

Faloutsos CMU SCS 15-415 #35

CMU SCS

More on Buffer Management

- When done, requestor of page must
 - unpin it, and
 - indicate whether page has been modified: *dirty* bit
- Page in pool may be requested many times:
 - *pin count*
- if *pin count* = 0 (“unpinned”), page is candidate for replacement

Faloutsos CMU SCS 15-415 #36

CMU SCS

More on Buffer Management

- CC & recovery may entail additional I/O when a frame is chosen for replacement. (*Write-Ahead Log* protocol; more later.)

Faloutsos CMU SCS 15-415 #37

CMU SCS

Buffer Replacement Policy

- Frame is chosen for replacement by a *replacement policy*:
 - Least-recently-used (LRU), MRU, Clock, etc.
- Policy can have big impact on # of I/O's; depends on the *access pattern*.

Faloutsos CMU SCS 15-415 #38

CMU SCS

LRU Replacement Policy

- *Least Recently Used (LRU)*
 - for each page in buffer pool, keep track of time last *unpinned*
 - replace the frame which has the oldest (earliest) time
 - very common policy: intuitive and simple
- Problems?

Faloutsos CMU SCS 15-415 #39

CMU SCS

LRU Replacement Policy

- Problem: *Sequential flooding*
 - LRU + repeated sequential scans.
 - # *buffer frames* < # *pages in file* means each page request causes an I/O. MRU much better in this situation (but not in all situations, of course).

Faloutsos CMU SCS 15-415 #40

CMU SCS

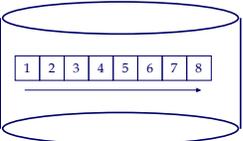
Sequential Flooding – Illustration

LRU:

BUFFER POOL			
102	116	242	105

MRU:

BUFFER POOL			
102	116	242	105



Repeated scan of file ...

Faloutsos CMU SCS 15-415 #41

CMU SCS

Sequential Flooding – Illustration

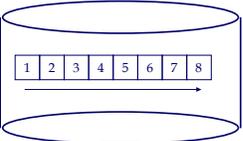
LRU:

BUFFER POOL			
1	2	3	4

will not re-use these pages;

MRU:

BUFFER POOL			
4	116	242	105



Repeated scan of file ...

Faloutsos CMU SCS 15-415 #42

CMU SCS

Other policies?

- LRU is often good - but needs timestamps and sorting on them
- something easier to maintain?

Faloutsos CMU SCS 15-415 #43

CMU SCS

“Clock” Replacement Policy

Main ideas:

- Approximation of LRU.
- Instead of maintaining & sorting time-stamps, find a ‘reasonably old’ frame to evict.
- How? by round-robin, and marking each frame - frames are evicted the second time they are visited.
- Specifically:

Faloutsos CMU SCS 15-415 #44

CMU SCS

“Clock” Replacement Policy

- Arrange frames into a cycle, store one “reference bit” per frame
- When pin count goes to 0, reference bit set on (= ‘one life left’ - not ready for eviction yet)
- When replacement necessary, get the next frame that has reference-bit = 0

The diagram shows a circle representing a cycle of frames. Three points on the circle are labeled A(1), B(1), and C(1), with an arrow pointing from A(1) to B(1) to C(1) and back to A(1). A fourth point, D(0), is also on the circle. The numbers in parentheses represent the reference bit for each frame.

Faloutsos CMU SCS 15-415 #45

CMU SCS

“Clock” Replacement Policy

```

do {
  if (pincount == 0 && ref bit is off)
    choose current page for replacement;
  else if (pincount == 0 && ref bit is on)
    turn off ref bit;
    advance current frame;
} until a page is chosen for replacement;

```

Faloutsos CMU SCS 15-415 #46

CMU SCS

Summary

- Buffer manager brings pages into RAM.
- Very important for performance
 - Page stays in RAM until released by requestor.
 - Written to disk when frame chosen for replacement (which is sometime after requestor releases the page).
 - Choice of frame to replace based on *replacement policy*.
 - Good to *pre-fetch* several pages at a time.

Faloutsos CMU SCS 15-415 #47

CMU SCS

Overview

- Memory hierarchy
- RAID (briefly)
- Disk space management
- Buffer management
- Files of records
- Page Formats
- Record Formats

Faloutsos CMU SCS 15-415 #48

CMU SCS

Files

- FILE: A collection of pages, each containing a collection of records.
- Must support:
 - insert/delete/modify record
 - read a particular record (specified using *record id*)
 - scan all records (possibly with some conditions on the records to be retrieved)

Faloutsos CMU SCS 15-415 #49

CMU SCS

Alternative File Organizations

Several alternatives (w/ trade-offs):

- Heap files: Suitable when typical access is a file scan retrieving all records.
- Sorted Files:
- Index File Organizations: } later

Faloutsos CMU SCS 15-415 #50

CMU SCS

Files of records

- Heap of pages
 - as linked list or
 - directory of pages

Faloutsos CMU SCS 15-415 #51

CMU SCS

Heap File Using Lists

- The header page id and Heap file name must be stored someplace.
- Each page contains 2 `pointers` plus data.

Faloutsos CMU SCS 15-415 #52

CMU SCS

Heap File Using Lists

- Any problems?

Faloutsos CMU SCS 15-415 #53

CMU SCS

Heap File Using a Page Directory

Faloutsos CMU SCS 15-415 #54

CMU SCS

Heap File Using a Page Directory

- The entry for a page can include the number of free bytes on the page.
- The directory is a collection of pages; linked list implementation is just one alternative.
– Much smaller than linked list of all HF pages!

Faloutsos CMU SCS 15-415 #55

CMU SCS

Overview

- Memory hierarchy
- RAID (briefly)
- Disk space management
- Buffer management
- Files of records
- Page Formats
- Record Formats

Faloutsos CMU SCS 15-415 #56

CMU SCS

Page Formats

- fixed length records
- variable length records

Faloutsos CMU SCS 15-415 #57

CMU SCS

Page Formats

Important concept: *rid* == record id

Q0: why do we need it?

Q1: How to mark the location of a record?

Q2: Why not its byte offset in the file?

Faloutsos CMU SCS 15-415 #58

CMU SCS

Page Formats

Important concept: *rid* == record id

Q0: why do we need it?

A0: eg., for indexing

Q1: How to mark the location of a record?

A1: rid = record id = page-id & slot-id

Q2: Why not its byte offset in the file?

A2: too much re-organization on ins/del.

Faloutsos CMU SCS 15-415 #59

CMU SCS

Fixed length records

- Q: How would you store them on a page/file?

Faloutsos CMU SCS 15-415 #60

CMU SCS

Fixed length records

- Q: How would you store them on a page/file?
- A1: How about:

'Packed'

slot #1
slot #2
slot #N
free space
N
number of full slots

Faloutsos CMU SCS 15-415 #61

CMU SCS

Fixed length records

- A1: How about: **BUT**: On insertion/deletion, we have too much to reorganize/update

'Packed'

slot #1
slot #2
slot #N
free space
N
number of full slots

Faloutsos CMU SCS 15-415 #62

CMU SCS

Fixed length records

- What would you do?

Faloutsos CMU SCS 15-415 #63

CMU SCS

Fixed length records

- Q: How would you store them on a page/file?
- A2: Bitmaps

free slots

slot #1
slot #2
...
slot #N

page header

1 0 M

Faloutsos CMU SCS 15-415 #64

CMU SCS

Variable length records

- Q: How would you store them on a page/file?

occupied records

page header

Faloutsos CMU SCS 15-415 #65

CMU SCS

Variable length records

- Q: How would you store them on a page/file?

occupied records

- pack them
- keep ptrs to them

page header

slot directory

other info (# slots etc)

Faloutsos CMU SCS 15-415 #66

CMU SCS

Variable length records

- Q: How would you store them on a page/file?

occupied records

- pack them
- keep ptrs to them
- mark start of free space

Faloutsos CMU SCS 15-415 #67

CMU SCS

Variable length records

- Q: How would you store them on a page/file?

occupied records

- how many disk accesses to insert a record?
- to delete one?

Faloutsos CMU SCS 15-415 #68

CMU SCS

Variable length records

- SLOTTED PAGE organization - popular.

occupied records

Faloutsos CMU SCS 15-415 #69

CMU SCS

Overview

- Memory hierarchy
- RAID (briefly)
- Disk space management
- Buffer management
- Files of records
- Page Formats
- Record Formats

Faloutsos CMU SCS 15-415 #70

CMU SCS

Formats of records

- Fixed length records
 - How would you store them?
- Variable length records

Faloutsos CMU SCS 15-415 #71

CMU SCS

Record Formats: Fixed Length

F1 F2 F3 F4

← L1 → L2 L3 L4

Base address (B) Address = B+L1+L2

- Information about field types same for all records in a file; stored in *system catalogs*.
- Finding *i*'th field done via arithmetic.

Faloutsos CMU SCS 15-415 #72

CMU SCS

Formats of records

- Fixed length records: straightforward - store info in catalog
- Variable length records: encode the length of each field
 - ?
 - ?

Faloutsos CMU SCS 15-415 #73

CMU SCS

Formats of records

- Fixed length records: straightforward - store info in catalog
- Variable length records: encode the length of each field
 - store its length or
 - use a field delimiter

Faloutsos CMU SCS 15-415 #74

CMU SCS

Variable Length records

- Two alternative formats (# fields is fixed):

Fields Delimited by Special Symbols

Array of Field Offsets

Pros and cons?

Faloutsos CMU SCS 15-415 #75

CMU SCS

Variable Length records

- Two alternative formats (# fields is fixed):

Fields Delimited by Special Symbols

Array of Field Offsets

Offset approach: usually superior (direct access to i-th field)

Faloutsos CMU SCS 15-415 #76

CMU SCS

Conclusions

- Memory hierarchy
- Disks: (>1000x slower) - thus
 - pack info in blocks
 - try to fetch nearby blocks (sequentially)
- Buffer management: very important
 - LRU, MRU, Clock, etc
- Record organization: Slotted page

Faloutsos CMU SCS 15-415 #77
