

CS 245: Database System Principles

Notes 03: Disk Organization

Hector Garcia-Molina

CS 245

Notes 3

1

Topics for today

- How to lay out data on disk
- How to move it to memory

CS 245

Notes 3

2

What are the data items we want to store?

- a salary
- a name
- a date
- a picture

CS 245

Notes 3

3

What are the data items we want to store?

- a salary
- a name
- a date
- a picture

⇒ What we have available: Bytes



CS 245

Notes 3

4

To represent:

- Integer (short): 2 bytes
e.g., 35 is

00000000 00100011

- Real, floating point
 n bits for mantissa, m for exponent....

CS 245

Notes 3

5

To represent:

- Characters
→ various coding schemes suggested,
most popular is ascii

Example:

A: 1000001
a: 1100001
5: 0110101
LF: 0001010

CS 245

Notes 3

6

To represent:

- Boolean
e.g., TRUE

1111	1111
------	------

FALSE

0000	0000
------	------
- Application specific
e.g., RED → 1 GREEN → 3
BLUE → 2 YELLOW → 4 ...

CS 245

Notes 3

7

To represent:

- Boolean
e.g., TRUE

1111	1111
------	------

FALSE

0000	0000
------	------
- Application specific
e.g., RED → 1 GREEN → 3
BLUE → 2 YELLOW → 4 ...

⇒ Can we use less than 1 byte/code?

Yes, but only if desperate...

CS 245

Notes 3

8

To represent:

- Dates
e.g.: - Integer, # days since Jan 1, 1900
- 8 characters, YYYYMMDD
- 7 characters, YYYYDDD
(not YYMMDD! Why?)
- Time
e.g. - Integer, seconds since midnight
- characters, HHMMSSFF

CS 245

Notes 3

9

To represent:

- String of characters
- Null terminated
e.g.,

c	a	t	⊗		
---	---	---	---	--	--
- Length given
e.g.,

3	c	a	t	⊗		
---	---	---	---	---	--	--
- Fixed length

CS 245

Notes 3

10

To represent:

- Bag of bits



CS 245

Notes 3

11

Key Point

- Fixed length items
- Variable length items
- usually length given at beginning

CS 245

Notes 3

12

Also

- Type of an item: Tells us how to interpret (plus size if fixed)


CS 245

Notes 3

13

Overview

Data Items



CS 245

Notes 3

14

Record - Collection of related data items (called FIELDS)

E.g.: Employee record:
name field,
salary field,
date-of-hire field, ...

CS 245

Notes 3

15

Types of records:

- Main choices:
 - FIXED vs VARIABLE FORMAT
 - FIXED vs VARIABLE LENGTH

CS 245

Notes 3

16

Fixed format

A SCHEMA (not record) contains following information

- # fields
- type of each field
- order in record
- meaning of each field

CS 245

Notes 3

17

Example: fixed format and length

Employee record

(1) E#, 2 byte integer

(2) E.name, 10 char.

(3) Dept, 2 byte code

} Schema

55 s m i t h 02

83 j o n e s 01

} Records

CS 245

Notes 3

18

Variable format


- Record itself contains format
"Self Describing"

CS 245

Notes 3

19

Example: variable format and length



Field name codes could also be strings, i.e. TAGS

CS 245

Notes 3

20

Variable format useful for:

- "sparse" records
- repeating fields
- evolving formats

.....→ But may waste space...

CS 245

Notes 3

21

- EXAMPLE:** var format record with repeating fields
Employee → one or more → children

3	E_name: Fred	Child: Sally	Child: Tom
---	--------------	--------------	------------

CS 245

Notes 3

22

Note: Repeating fields does not imply

- variable format, nor
- variable size

John	Sailing	Chess	--
------	---------	-------	----

CS 245

Notes 3

23

Note: Repeating fields does not imply

- variable format, nor
- variable size

John	Sailing	Chess	--
------	---------	-------	----

- Key is to allocate maximum number of repeating fields (if not used → null)


CS 245

Notes 3

24

☆ Many variants between fixed - variable format:

Example: Include record type in record



Record header - data at beginning that describes record


May contain:

- record type
- record length
- time stamp
- other stuff ...


Other interesting issues:

- Compression
 - within record - e.g. code selection
 - collection of records - e.g. find common patterns
- Encryption


Encrypting Records



Encrypting Records




Search Key in the Clear



- each record is [k,b]
- store [k, E(b)]
- can search for records with k=x

Encrypt Key



[E(1), E(b₁)]
 [E(2), E(b₂)]
 [E(3), E(b₃)]
 [E(4), E(b₄)]
 ...

- each record is [k,b]
- store [E(k), E(b)]
- can search for records with k=E(x)


CS 245

Notes 3

31

Issues

- Hard to do range queries
- Encryption not good
- Better to use encryption that does not always generate same cyphertext




simplification

CS 245

Notes 3

32

How Do We Search Now?



[E(1, abc), E(b₁)]
 [E(2, dhe), E(b₂)]
 [E(3, nft), E(b₃)]
 [E(2, lkz), E(b₄)]
 ...

- each record is [k,b]
- store [E(k, rand), E(b)]
- can search for records with k=E(x,???)

CS 245

Notes 3

33

Solution?

- Develop new decryption function:
 $D(f(k_1), E(k_2, \text{rand}))$ is true if $k_1=k_2$


CS 245

Notes 3

34

Solution?

- Develop new decryption function:
 $D(f(k_1), E(k_2, \text{rand}))$ is true if $k_1=k_2$



[E(1, abc), E(b₁)]
 [E(2, dhe), E(b₂)]
 [E(3, nft), E(b₃)]
 [E(2, lkz), E(b₄)]
 ...

CS 245

Notes 3

35

Issues?


- Cannot do non-equality predicates
- Hard to build indexes

CS 245

Notes 3


36

Next: placing records into blocks



CS 245 Notes 3 37

Next: placing records into blocks



CS 245 Notes 3 38

Options for storing records in blocks:

- (1) separating records
- (2) spanned vs. unspanned
- (3) sequencing
- (4) indirection

CS 245 Notes 3 39

(1) Separating records



- (a) no need to separate - fixed size recs.
- (b) special marker
- (c) give record lengths (or offsets)
 - within each record
 - in block header


CS 245 Notes 3 40

(2) Spanned vs. Unspanned

- Unspanned: records must be within one block




- Spanned



CS 245 Notes 3 41

With spanned records:



need indication of partial record "pointer" to rest

need indication of continuation (+ from where?)

CS 245 Notes 3 42

Spanned vs. unspanned:

- Unspanned is much simpler, but may waste space...
- Spanned essential if record size > block size

CS 245

Notes 3

43

(3) Sequencing

- Ordering records in file (and block) by some key value

Sequential file (\Rightarrow sequenced)

CS 245

Notes 3

44

Why sequencing?

Typically to make it possible to efficiently read records in order
(e.g., to do a merge-join — discussed later)


CS 245

Notes 3


45

Sequencing Options

(a) Next record physically contiguous



(b) Linked




CS 245

Notes 3

46

Sequencing Options

(c) Overflow area




CS 245

Notes 3

47

Sequencing Options

(c) Overflow area



CS 245

Notes 3

48

(4) Indirection

- How does one refer to records?




CS 245

Notes 3

49

(4) Indirection

- How does one refer to records?



Many options:

Physical \longleftrightarrow Indirect

CS 245

Notes 3

50

☆ Purely Physical

E.g., Record Address or ID = $\left\{ \begin{array}{l} \text{Device ID} \\ \text{Cylinder \#} \\ \text{Track \#} \\ \text{Block \#} \\ \text{Offset in block} \end{array} \right\}$ Block ID


CS 245

Notes 3

51

☆ Fully Indirect

E.g., Record ID is arbitrary bit string



CS 245

Notes 3

52

Tradeoff

Flexibility \longleftrightarrow Cost
to move records of indirection
(for deletions, insertions)

CS 245

Notes 3

53

Physical \longleftrightarrow Indirect


↑
Many options
in between ...

CS 245

Notes 3

54

Example: Indirection in block



CS 245

Notes 3

55

Block header - data at beginning that describes block

May contain:

- File ID (or RELATION or DB ID)
- This block ID
- Record directory
- Pointer to free space
- Type of block (e.g. contains recs type 4; is overflow, ...)
- Pointer to other blocks "like it"
- Timestamp ...

CS 245

Notes 3

56

Options for storing records in blocks:

- (1) separating records
- (2) spanned vs. unspanned
- (3) sequencing
- (4) indirection


CS 245

Notes 3

57

Case Study: salesforce.com

- salesforce.com provides CRM services
- salesforce customers are *tenants*
- Tenants run apps and DBMS as service



CS 245

Notes 3

58

Options for Hosting

- Separate DBMS per tenant
- One DBMS, separate tables per tenant
- One DBMS, shared tables

CS 245

Notes 3

59

Tenants have similar data

tenant 1:

customer	A	B	C	D	E	F
a1	b1	c1	d1	e1	-	
a2	b2	c2	-	e2	f2	

tenant 2:

customer	A	B	C	D	G
a3	b3	c2	-	-	
a1	b1	c1	-	g1	
a4	-	-	d1		

CS 245

Notes 3

60

salesforce.com solution

customer	tenant	A	B	C
1	1	a1	b1	c1
1	1	a2	b2	c2
2	2	a3	b3	c2
2	2	a1	b1	c1

← fixed schema for all tenants

cust-other	tenant	A	f1	v1	f2	v2 ...
1	1	a1	D	d1	E	e1
1	1	a2	E	e2	F	f2
2	2	a1	G	g1		
3	3	a4	D	d1		

← var schema for all tenants

CS 245 Notes 3 61

Other Topics

- (1) Insertion/Deletion
- (2) Buffer Management
- (3) Comparison of Schemes

CS 245 Notes 3 62

Deletion

Block

The diagram shows a rectangular box labeled 'Block'. Inside the box, on the left side, is a smaller rectangle labeled 'Rx'. A horizontal arrow points from the right side of the 'Rx' box towards the right edge of the 'Block' box, indicating a cursor or pointer.

CS 245 Notes 3 63

Options:

- (a) Immediately reclaim space
- (b) Mark deleted

CS 245 Notes 3 64

Options:

- (a) Immediately reclaim space
- (b) Mark deleted
 - May need chain of deleted records (for re-use)
 - Need a way to mark:
 - special characters
 - delete field
 - in map

CS 245 Notes 3 65


☆ As usual, many tradeoffs...

- How expensive is to move valid record to free space for immediate reclaim?
- How much space is wasted?
 - e.g., deleted records, delete fields, free space chains,...

CS 245 Notes 3 66

Concern with deletions

Dangling pointers



CS 245

Notes 3

67

Solution #1: Do not worry

CS 245

Notes 3

68

Solution #2: Tombstones

E.g., Leave "MARK" in map or old location

CS 245


Notes 3

69

Solution #2: Tombstones

E.g., Leave "MARK" in map or old location

• Physical IDs



CS 245

Notes 3

70

Solution #2: Tombstones

E.g., Leave "MARK" in map or old location

• Logical IDs

map

ID	LOC
7788	

Never reuse
ID 7788 nor
space in map...

CS 245

Notes 3

71

Insert

Easy case: records not in sequence

- Insert new record at end of file or in deleted slot
- If records are variable size, not as easy...

CS 245

Notes 3

72

Insert

Hard case: records in sequence

- If free space "close by", not too bad...
- Or use overflow idea...

CS 245

Notes 3

73


Interesting problems:

- How much free space to leave in each block, track, cylinder?
- How often do I reorganize file + overflow?

CS 245

Notes 3

74



CS 245

Notes 3

75

Buffer Management


- DB features needed
 - Why LRU may be bad
 - Pinned blocks
 - Forced output
 - Double buffering
 - Swizzling
- Read
Textbook!
- in Notes02

CS 245

Notes 3

76

Swizzling




CS 245

Notes 3

77

Swizzling



CS 245

Notes 3

78

Row vs Column Store

- So far we assumed that fields of a record are stored contiguously (row store)...
- Another option is to store like fields together (column store)

CS 245

Notes 3

79

Row Store

- Example: Order consists of
 - id, cust, prod, store, price, date, qty

id1	cust1	prod1	store1	price1	date1	qty1
-----	-------	-------	--------	--------	-------	------

id2	cust2	prod2	store2	price2	date2	qty2
-----	-------	-------	--------	--------	-------	------

id3	cust3	prod3	store3	price3	date3	qty3
-----	-------	-------	--------	--------	-------	------

CS 245

Notes 3

80

Column Store

- Example: Order consists of
 - id, cust, prod, store, price, date, qty

id1	cust1	id1	prod1	id1	price1	qty1
id2	cust2	id2	prod2	id2	price2	qty2
id3	cust3	id3	prod3	id3	price3	qty3
id4	cust4	id4	prod4	id4	price4	qty4
...

ids may or may not be stored explicitly

CS 245

Notes 3

81

Row vs Column Store

- Advantages of Column Store
 - more compact storage (fields need not start at byte boundaries)
 - efficient reads on data mining operations
- Advantages of Row Store
 - writes (multiple fields of one record) more efficient
 - efficient reads for record access (OLTP)

CS 245

Notes 3

82

Interesting paper to read:

- Mike Stonebreaker, Elizabeth (Betty) O'Neil, Pat O'Neil, Xuedong Chen, et al. "C-Store: A Column-oriented DBMS," Presented at the 31st VLDB Conference, September 2005.
- http://www.cs.umb.edu/%7Eeponeil/vldb05_cstore.pdf

CS 245

Notes 3

83

Comparison

- There are 10,000,000 ways to organize my data on disk...


Which is right for me?

CS 245

Notes 3

84

Issues:



☆ To evaluate a given strategy, compute following parameters:

- > space used for expected data
- > expected time to
 - fetch record given key
 - fetch record with next key
 - insert record
 - append record
 - delete record
 - update record
 - read all file
 - reorganize file


Example

How would you design Megatron 3000 storage system? (for a relational DB, low end)

- Variable length records?
- Spanned?
- What data types?
- Fixed format?
- Record IDs ?
- Sequencing?
- How to handle deletions?

Summary

- How to lay out data on disk



Next

How to find a record quickly,
given a key