

CS 245: Database System Principles

Notes 5: Hashing and More

Hector Garcia-Molina
(Some modifications by Chris Olston)


CS 245

Notes 5

1

Hashing

key \rightarrow $h(\text{key})$




CS 245

Notes 5

2

Two alternatives

(1) key \rightarrow $h(\text{key})$




CS 245

Notes 5

3

Two alternatives

(2) key \rightarrow $h(\text{key})$



- Alt (2) for "secondary" search key

CS 245

Notes 5

4

Example hash function

- Key = $\langle x_1 x_2 \dots x_n \rangle$ n byte character string
- Have b buckets
- h : add $x_1 + x_2 + \dots + x_n$
 - compute sum modulo b

CS 245

Notes 5

5

- This may not be best function ...
- Read Knuth Vol. 3 if you really need to select a good function.

Good hash function: \Rightarrow Expected number of keys/bucket is the same for all buckets

CS 245

Notes 5

6

Within a bucket:


- Do we keep keys sorted?
- Yes, if CPU time critical
& Inserts/Deletes not too frequent

CS 245

Notes 5

7

Next: example to illustrate inserts, overflows, deletes



CS 245


Notes 5

8

EXAMPLE 2 records/bucket

INSERT:

- h(a) = 1
- h(b) = 2
- h(c) = 1
- h(d) = 0
- h(e) = 1



CS 245


Notes 5

9

EXAMPLE: deletion

Delete:

- e
- f
- c



CS 245

Notes 5

10

Rule of thumb:

- Try to keep space utilization between 50% and 80%
Utilization = $\frac{\# \text{ keys used}}{\text{total } \# \text{ keys that fit}}$
- If < 50%, wasting space
- If > 80%, overflows significant
↳ depends on how good hash function is & on # keys/bucket

CS 245

Notes 5

11

How do we cope with growth?

- Overflows and reorganizations
- Dynamic hashing
 - Extensible
 - Linear


CS 245

Notes 5

12

Extensible hashing: two ideas

(a) Use i of b bits output by hash function




CS 245

Notes 5

13

(b) Use directory




CS 245

Notes 5

14

Example: $h(k)$ is 4 bits; 2 keys/bucket




CS 245

Notes 5

15

Example continued

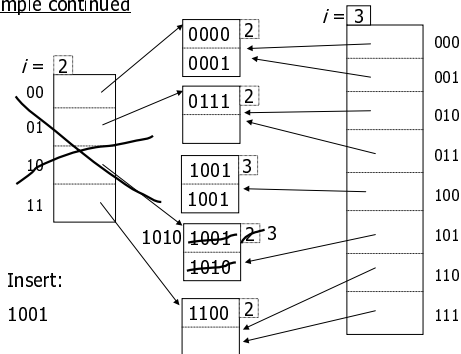


CS 245

Notes 5

16

Example continued



CS 245

Notes 5

17

Extensible hashing: deletion

- No merging of blocks
- Merge blocks and cut directory if possible (Reverse insert procedure)

CS 245

Notes 5

18

Deletion example:

- Run thru insert example in reverse!

Summary Extensible hashing


- ⊕ Can handle growing files
 - with less wasted space
 - with no full reorganizations
- ⊖ Indirection
(Not bad if directory in memory)
- ⊖ Directory doubles in size
(Now it fits, now it does not)

Linear hashing


- Another dynamic hashing scheme

Two ideas:


(a) Use i low order bits of hash



(b) File grows linearly




Example $b=4$ bits, $i=2$, 2 keys/bucket



- insert 0101
- can have overflow chains!


Rule If $h(k)[i] \leq m$, then look at bucket $h(k)[i]$ else, look at bucket $h(k)[i] - 2^{i-1}$

Example $b=4$ bits, $i=2$, 2 keys/bucket



- insert 0101

Example Continued: How to grow beyond this?



$i = 2 \rightarrow 3$

☛ When do we expand file?

- Keep track of: $\frac{\text{\# used slots}}{\text{total \# of slots}} = U$
- If $U > \text{threshold}$ then increase m
(and maybe i)

CS 245

Notes 5

25

Summary Linear Hashing


- ⊕ Can handle growing files
 - with less wasted space
 - with no full reorganizations
- ⊕ No indirection like extensible hashing
- ⊖ Can still have overflow chains

CS 245

Notes 5

26

Example: BAD CASE



CS 245

Notes 5

27

Summary

Hashing

- How it works
- Dynamic hashing
 - Extensible
 - Linear

CS 245

Notes 5

28

Next:

- Ordered Indexing vs Hash Indexing
- Index definition in SQL
- Multiple key access
- Index selection problem

CS 245

Notes 5

29

Ordered Indexing vs Hash Indexing

- Hashing good for probes given key
e.g.,
SELECT ...
FROM R
WHERE R.A = 5

CS 245

Notes 5

30

Ordered Indexing vs Hash Indexing

- ORDERED INDEXING (Including B-Trees) good for

Range Searches:

e.g.,
SELECT
FROM R
WHERE R.A > 5

CS 245

Notes 5

31

Index definition in SQL

- Create index_name on rel (attr)
- Create unique index name on rel (attr)
 - ↳ defines candidate key
- Drop INDEX name

CS 245

Notes 5

32

Note CANNOT SPECIFY TYPE OF INDEX
(e.g. B-tree, Hashing, ...)
OR PARAMETERS
(e.g. Load Factor, Size of Hash,...)

... at least in SQL...

CS 245

Notes 5

33

Note ATTRIBUTE LIST ⇒ MULTIKEY INDEX
(next)

e.g., CREATE INDEX foo ON R(A,B,C)

CS 245

Notes 5

34

Multi-key Index

Motivation: Find records where
DEPT = "Toy" AND SAL > 50k


CS 245

Notes 5

35

Strategy I:

- Use one index, say Dept.
- Get all Dept = "Toy" records and check their salary




CS 245

Notes 5

36


Strategy II:

- Use 2 Indexes; Manipulate Pointers

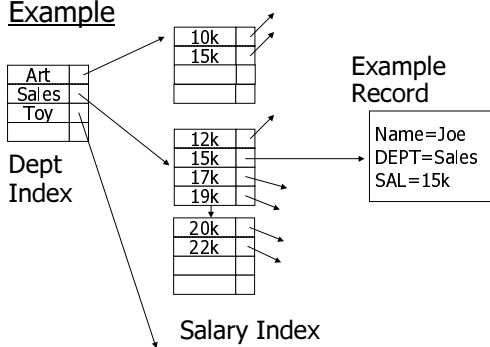


Strategy III:

- Multiple Key Index




Example



Efficient B-Tree Implementation

Index entries are pairs (dept, salary)
Relies on total ordering over pairs:




For which queries is this index good?

- Find RECs Dept = "Sales" \wedge SAL=20k
- Find RECs Dept = "Sales" \wedge SAL \geq 20k
- Find RECs Dept = "Sales"
- Find RECs SAL = 20k

Interesting application:

- Geographic ("spatial") Data



DATA:
<X1,Y1, Attributes>
<X2,Y2, Attributes>
⋮

Queries:

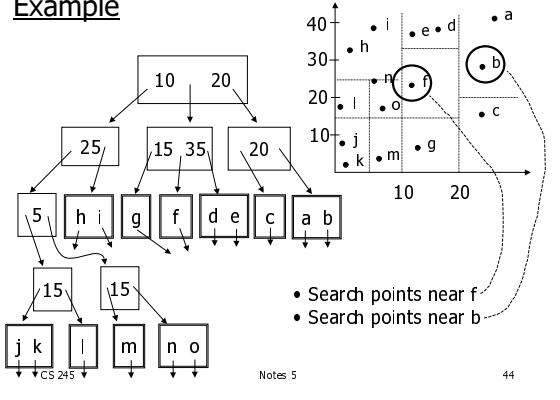
- What city is at $\langle X_i, Y_i \rangle$?
- What is within 5 miles from $\langle X_i, Y_i \rangle$?
- Which is closest point to $\langle X_i, Y_i \rangle$?

CS 245

Notes 5

43

Example



Notes 5

44

Queries

- Find points with $Y_i > 20$
- Find points with $X_i < 5$
- Find points "close" to $i = \langle 12, 38 \rangle$
- Find points "close" to $b = \langle 7, 24 \rangle$

CS 245

Notes 5

45

- Many types of spatial index structures have been suggested

- Quad Trees
- R Trees
- K-D Trees
- ...

CS 245

Notes 5

46

Three more types of multi key indexes


- Grid
- R-Tree
- Partitioned hash

CS 245

Notes 5

47

Grid Index



CS 245

Notes 5


48

CLAIM

- Can quickly find records with
 - key 1 = $V_i \wedge$ Key 2 = X_j
 - key 1 = V_i
 - key 2 = X_j
- And also ranges...
 - E.g., key 1 $\geq V_i \wedge$ key 2 $< X_j$

☞ But there is a catch with Grid Indexes!


- How is Grid Index stored on disk?



Problem:

- Need regularity so we can compute position of $\langle V_i, X_j \rangle$ entry


Solution: Use Indirection



With indirection:

- Grid can be regular without wasting space
- We do have price of indirection

Can also index grid on value ranges




Grid files

- ⊕ Good for multiple-key search
- ⊖ Space, management overhead (nothing is free)
- ⊖ Need partitioning ranges that evenly split keys


R-Trees

Extension of B-Tree idea

B-Tree: partition 1-D space into line segments



R-Tree: partition n-D space into regions (typically rectangles)




CS 245

Notes 5

55

Example R-Tree



CS 245

Notes 5

56

R-Tree vs. B-Tree

- Multi-key access
- Partitions not disjoint (may overlap)
 - May have to search multiple branches!
- Try to choose partitions to minimize amount of overlap
 - When initially build index
 - When inserting new keys

CS 245

Notes 5

57

R-Tree vs. Grid files

- ⊕ Dynamically selects/adjusts partitions
- ⊕ No storage for empty regions
- ⊖ Overlap: search multiple partitions


CS 245

Notes 5

58

Partitioned hash function

Idea:



CS 245

Notes 5

59

EX:

h1(toy)	= 0	000	
h1(sales)	= 1	001	<Fred>
h1(art)	= 1	010	
⋮		011	
h2(10k)	= 01	100	
h2(20k)	= 11	101	<Joe> <Sally>
h2(30k)	= 01	110	
h2(40k)	= 00	111	
⋮			

Insert → <Fred,toy,10k>, <Joe,sales,10k>
<Sally,art,30k>

CS 245

Notes 5

60

h1(toy)	=0	000	<Fred>
h1(sales)	=1	001	<Joe><Jan>
h1(art)	=1	010	<Mary>
.		011	
h2(10k)	=01	100	<Sally>
h2(20k)	=11	101	
h2(30k)	=01	110	<Tom><Bill>
h2(40k)	=00	111	<Andy>
.			
.			

- Find Emp. with Dept. = Sales \wedge Sal=40k

CS 245 Notes 5 61

h1(toy)	=0	000	<Fred>
h1(sales)	=1	001	<Joe><Jan>
h1(art)	=1	010	<Mary>
.		011	
h2(10k)	=01	100	<Sally>
h2(20k)	=11	101	
h2(30k)	=01	110	<Tom><Bill>
h2(40k)	=00	111	<Andy>
.			
.			

- Find Emp. with Sal=30k

look here

CS 245 Notes 5 62

h1(toy)	=0	000	<Fred>
h1(sales)	=1	001	<Joe><Jan>
h1(art)	=1	010	<Mary>
.		011	
h2(10k)	=01	100	<Sally>
h2(20k)	=11	101	
h2(30k)	=01	110	<Tom><Bill>
h2(40k)	=00	111	<Andy>
.			
.			

- Find Emp. with Dept. = Sales

look here

CS 245 Notes 5 63

Index Selection

- Problem: many indexing options, how to choose?
 - Primary index column (sequencing column)
 - Secondary index columns
 - Dense/sparse
 - Index types
 - B-trees, hashing, multi-key, and more...

CS 245 Notes 5 64

Index Selection in Practice

- Until recently, most DBMS's only supported B-trees (some multi-column)
 - Hash indexes rare
- New generation of DBMS's:
 - Focus on domain-specific indexes
 - E.g., images, text, spatial data, multimedia

CS 245 Notes 5 65

The best index selection depends on the query workload :

- Reads vs. writes
- Which columns accessed
- Selection predicates: equality, range, ...
- Multiple selection predicates common?
- Joins vs. single-table queries
- Etc...

CS 245 Notes 5 66

Automatic Index Selection

- Problem –
 - Given: database, workload
 - Find: optimal set of indexes to build
- Can't the computer do this for us?
- Current research topic:
 - E.g., [Chaudhuri and Narasayya, VLDB'97]

CS 245

Notes 5

67

Data Independence

- Nice feature of modern DBMS's: Physical Data Independence
- User doesn't need to be aware of physical design (i.e., don't have to rewrite queries)
 - (Also: Logical Data Independence
 - Logical view facilities shield users from changes to schema

CS 245

Notes 5

68

In the old days ...

- "Pre-relational" or "legacy" systems
 - No physical data independence
 - The application designer wrote the data access code
 - e.g., scan table, traverse index, ...
 - Add/remove an index → have to change the code ☹
 - Data distribution changes → access method in use may not be the most efficient ☹

CS 245

Notes 5

69

Relational Database "Revolution"



- Physical data independence achieved via declarative query language interface
 - Say what you want, not how to get it
 - Rely on query optimizer to select a fast access method
 - Query optimization covered later in this class...

CS 245

Notes 5

70

Summary

Post hashing discussion:

- Ordered Indexing vs. Hash Indexing
- SQL Index Definition
- Multiple Key Access
 - Multi Key Index
 - Variations: Grid, Geo Data
 - Partitioned Hash
- Index Selection Problem
- Physical Data Independence

CS 245

Notes 5

71

Reading Chapter 5

- Skim the following sections:
 - 5.3.3, 5.3.4, 5.3.5, 5.3.6
 - 5.4.2, 5.4.3, 5.4.4
- Read the rest

CS 245

Notes 5

72

The **BIG** picture....

- Chapters 2 & 3: Storage, records, blocks...
- Chapter 4 & 5: Access Mechanisms
 - Indexes
 - B trees
 - Hashing
 - Multi key
- Chapter 6 & 7: Query Processing

