

CS 245: Database System Principles

Notes 5: Hashing and More

Hector Garcia-Molina


CS 245

Notes 5

1

Hashing

key \rightarrow h(key)




CS 245

Notes 5

2

Two alternatives

(1) key \rightarrow h(key)




CS 245

Notes 5

3

Two alternatives

(2) key \rightarrow h(key)



- Alt (2) for "secondary" search key

CS 245

Notes 5

4

Example hash function

- Key = 'x₁ x₂ ... x_n' 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: ☞ 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

Next: example to illustrate inserts, overflows, deletes



EXAMPLE 2 records/bucket

INSERT:


$h(a) = 1$

$h(b) = 2$

$h(c) = 1$

$h(d) = 0$

$h(e) = 1$




EXAMPLE: deletion

Delete:

e

f

c



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

How do we cope with growth?


- Overflows and reorganizations
- Dynamic hashing
 - Extensible
 - Linear

Extensible hashing: two ideas


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




(b) Use directory



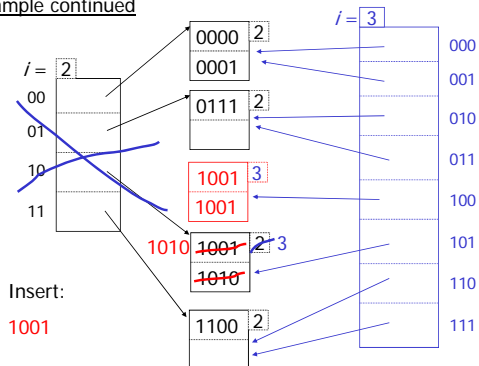
Example: $h(k)$ is 4 bits; 2 keys/bucket



Example continued



Example continued



Extensible hashing: deletion

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

Deletion example:

- Run thru insert example in reverse!


CS 245

Notes 5

19

Note: Still need overflow chains

- Example: many records with duplicate keys




CS 245

Notes 5

20

Solution: overflow chains



CS 245

Notes 5

21

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)

CS 245

Notes 5


22

Linear hashing

- Another dynamic hashing scheme

Two ideas:

(a) Use i low order bits of hash



(b) File grows linearly




CS 245

Notes 5

23

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



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

CS 245

Notes 5

24

Note

- In textbook, n is used instead of m
- $n = m + 1$

$n = 10$

$m = 01$ (max used block)

Future growth buckets

CS 245 Notes 5 25

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

- insert 0101

$n = 10$

$m = 01$ (max used block)

Future growth buckets

CS 245 Notes 5 26

Example Continued: How to grow beyond this?

$i = 2^3$

$m = 11$ (max used block)

Future growth buckets

CS 245 Notes 5 27

☒ 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 28

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 29

Example: BAD CASE

Very empty Very full

Need to move m here...
Would waste space...

CS 245 Notes 5 30

Summary

Hashing

- How it works
- Dynamic hashing
 - Extensible
 - Linear

CS 245

Notes 5

31

Next:

- Indexing vs Hashing
- Index definition in SQL
- Multiple key access

CS 245

Notes 5

32

Indexing vs Hashing

- Hashing good for probes given key

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

CS 245

Notes 5

33

Indexing vs Hashing

- INDEXING (Including B Trees) good for Range Searches:

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

CS 245

Notes 5

34

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

35

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

36

Note ATTRIBUTE LIST \Rightarrow MULTIKEY INDEX
(next)
e.g., `CREATE INDEX foo ON R(A,B,C)`

CS 245

Notes 5

37

Multi-key Index

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


CS 245

Notes 5

38

Strategy I:

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




CS 245

Notes 5

39

Strategy II:

- Use 2 Indexes; Manipulate Pointers



CS 245


Notes 5

40

Strategy III:

- Multiple Key Index

One idea:




CS 245

Notes 5

41

Example



CS 245

Notes 5

42

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


CS 245

Notes 5

43

Interesting application:

- Geographic Data



DATA:

- <X1,Y1, Attributes>
- <X2,Y2, Attributes>
- ⋮

CS 245

Notes 5

44

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

45

Example



- Search points near f
- Search points near b

CS 245

Notes 5

46

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

47

- Many types of geographic index structures have been suggested

- kd-Trees (very similar to what we described here)
- Quad Trees
- R Trees
- ...

CS 245

Notes 5

48

Two more types of multi key indexes


- Grid
- Partitioned hash

CS 245

Notes 5

49

Grid Index



CS 245

Notes 5

50

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 $< X_j$


CS 245

Notes 5

51

✉ 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

CS 245

Notes 5

52

Solution: Use Indirection



CS 245

Notes 5

53

With indirection:


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

CS 245

Notes 5

54


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

Partitioned hash function



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>

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

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

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 61

Summary

Post hashing discussion:

- Indexing vs. Hashing
- SQL Index Definition
- Multiple Key Access
 - Multi Key Index
 - Variations: Grid, Geo Data
 - Partitioned Hash

CS 245 Notes 5 62

Reading Chapter 5

- Skim the following sections:
 - Sections 14.3.6, 14.3.7, 14.3.8
[Second Ed: 14.6.6, 14.6.7, 14.6.8]
 - Sections 14.4.2, 14.4.3, 14.4.4
[Second Ed: 14.7.2, 14.7.3, 14.7.4]
- Read the rest

CS 245 Notes 5 63

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 ➔

CS 245 Notes 5 64