

Indexing: Part III

CPS 216
Advanced Database Systems

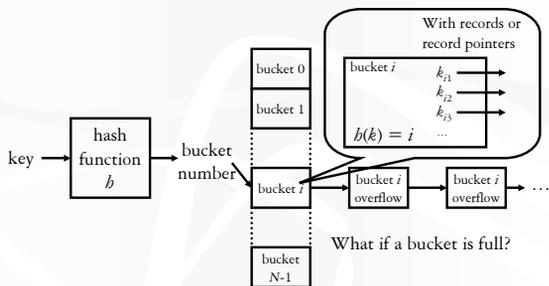
Announcements (February 10)

2

- ❖ Reading assignments
 - Query processing survey (due next Monday)
- ❖ Homework #2 will be assigned this Thursday
- ❖ Recitation session this Friday
- ❖ Midterm and course project proposal in 3½ weeks

Static hashing

3

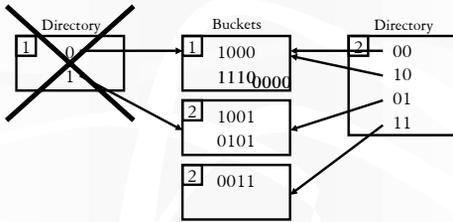


Does it make sense to use a hash-based index as a sparse index on a sorted table?

Extensible hashing example (slide 2)

7

❖ Insert 1110, 0000



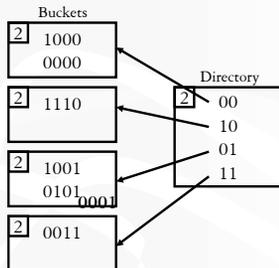
❖ Split again

- No directory doubling this time

Extensible hashing example (slide 3)

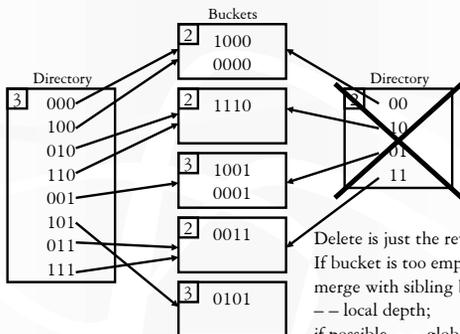
8

❖ Insert 0001



Extensible hashing example (slide 4)

9



Delete is just the reverse:
 If bucket is too empty,
 merge with sibling bucket,
 -- local depth;
 if possible, -- global depth
 and half the directory

Summary of extensible hashing 10

- ❖ Pros
 - Handles growing files
 - No full reorganization
- ❖ Cons

Linear hashing (VLDB 1980) 11

- ❖ Grow only when utilization exceeds a given threshold
- ❖ No extra indirection
 - Some extra math to figure out the right bucket

Insert 0101
Threshold exceeded; grow!

| | |
|--------------|--------------|
| 0 | 1 |
| 0000 1010 | 1111 0101 |

$i = 1$ Number of bits in use = $\lceil \log_2 n \rceil$
 $n = 2$ Number of primary buckets

Linear hashing example (slide 2) 12

- ❖ Grows linearly (hence the name)
- ❖ Always split the $(n - 2^{\lceil \log_2 n \rceil})$ -th bucket (0-based index)
 - Intuitively, the first bucket with the lowest depth
 - Not necessarily the bucket being inserted into!

Insert 0001 Insert 1100
Threshold exceeded; grow!

| | | |
|--------------|--------------|------|
| 00 | 1 | 10 |
| 0000 1100 | 1111 0101 | 1010 |

$i = 2$
 $n = 3$

0001

Linear hashing example (slide 3)

13

Insert 1110
Threshold exceeded; grow!

| 00 | 01 | 10 | 11 |
|--------------|--------------|--------------|------|
| 0000 1100 | 0001 0101 | 1010 1110 | 1111 |

$i = 2$
 $n = 4$

Linear hashing example (slide 4)

14

❖ Look up 1110

- Bucket 110 (6-th bucket) is not here
- Then look in the $(6 - 2^{\lfloor \log_2 n \rfloor})$ -th bucket (= 2nd)

| 000 | 01 | 10 | 11 | 100 |
|--------------|--------------|--------------|------|------|
| 0000 1100 | 0001 0101 | 1010 1110 | 1111 | 1100 |

$i = 3$
 $n = 5$

Summary of linear hashing

15

❖ Pros

- Handles growing files
- No full reorganization
- No extra level of indirection

❖ Cons

Hashing versus B-trees

16
