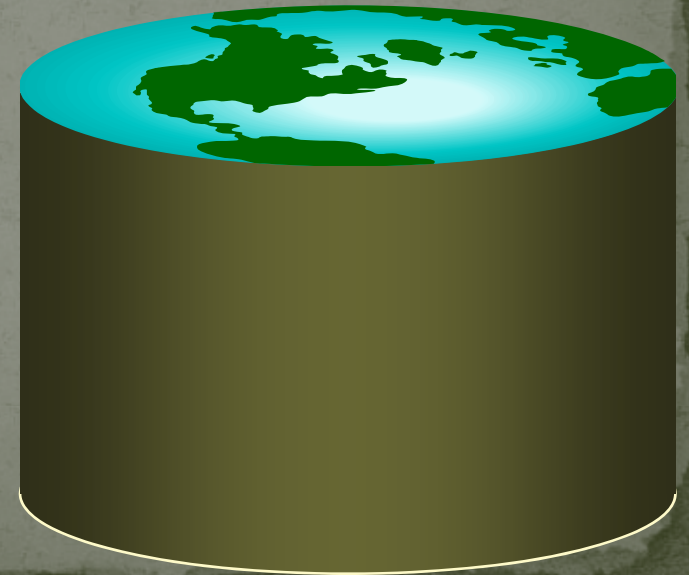


CS 505: Intermediate Topics to Database Systems

Instructor: Jinze Liu

Fall 2008



Hashing

2

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



Buckets
(typically 1
disk block)



Two Alternatives

3

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

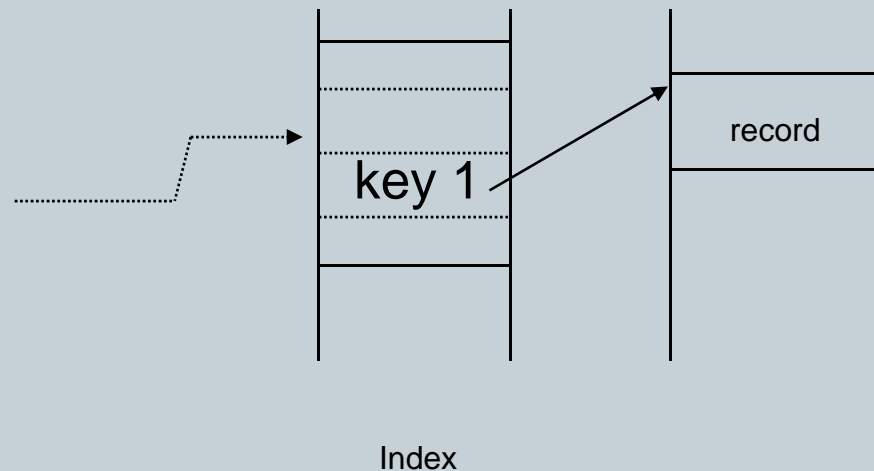
records

The diagram illustrates a process flow. On the left, the text '(1) key → h(key)' is followed by a dotted arrow that points to a vertical container. This container is divided into three sections: the top and bottom sections each contain three vertical dots, and the middle section is labeled 'records'. The entire diagram is set against a light blue background with a darker blue footer bar.

Two Alternatives

4

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



- Alt (2) for “secondary” search key

Example hash function

5

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

- ☒ 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

Within a bucket:

7

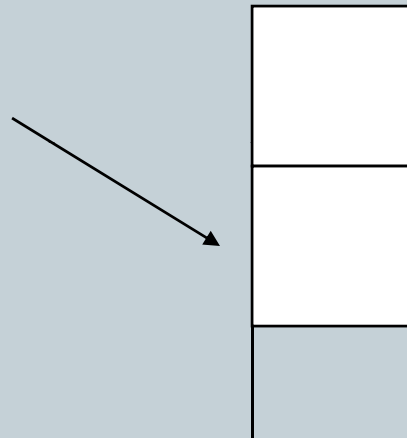
- Do we keep keys sorted?

Yes, if CPU time critical
& Inserts/Deletes not too frequent

Inserts, overflows, deletes

8

$h(K)$



EXAMPLE 2 records/bucket

9

INSERT:

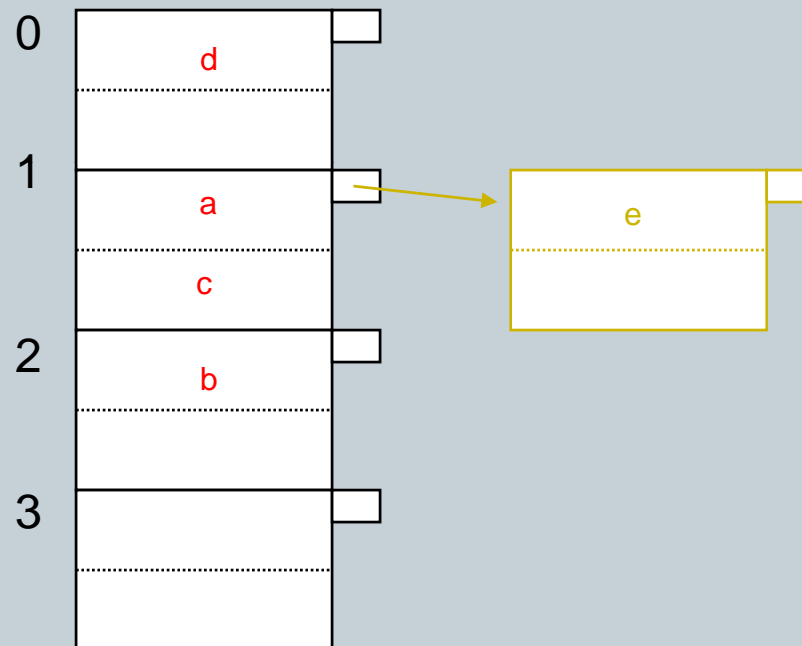
$h(a) = 1$

$h(b) = 2$

$h(c) = 1$

$h(d) = 0$

$h(e) = 1$



Example: Deletion

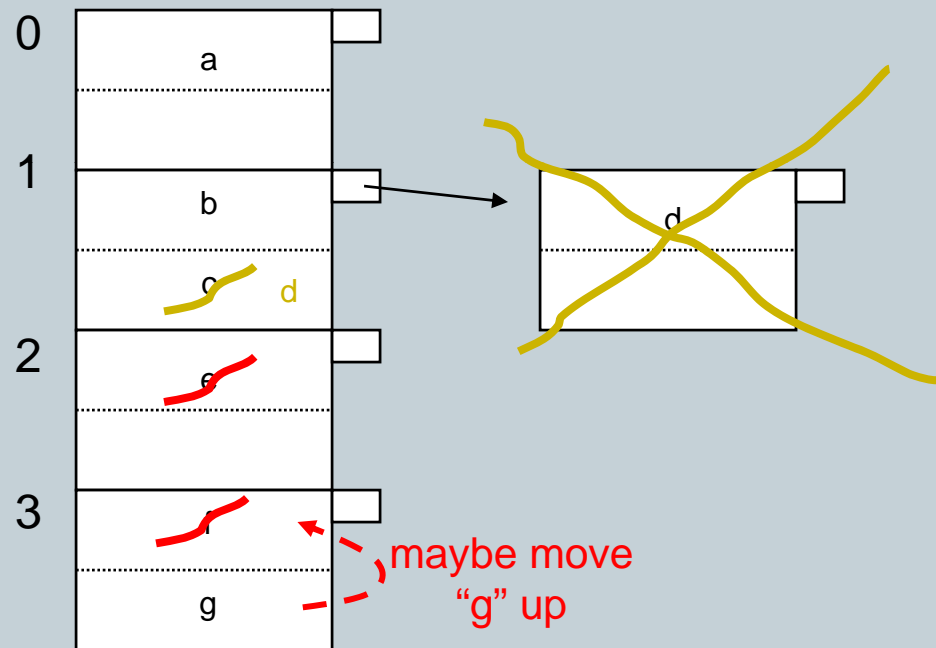
10

Delete:

e

f

c



Rule of thumb:

11

- Try to keep space utilization between 50% and 80%

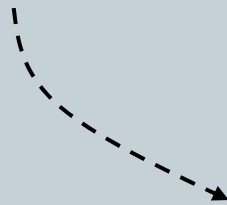
$$\text{Utilization} = \frac{\text{\# keys used}}{\text{total \# keys that fit}}$$

- If $< 50\%$, wasting space
- If $> 80\%$, overflows significant
 ↖ depends on how good hash function is & on # keys/bucket

How do we cope with growth?

12

- Overflows and reorganizations
- Dynamic hashing

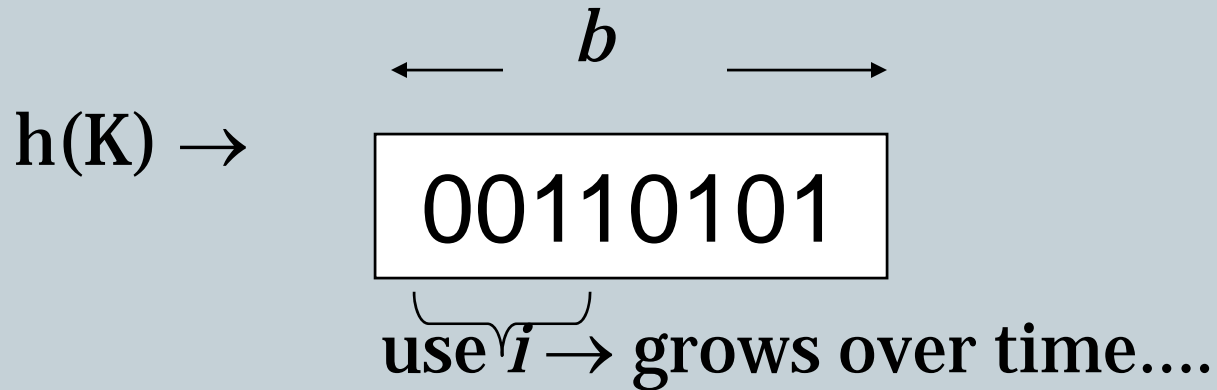


- Extensible
- Linear

Extensible hashing: two ideas

13

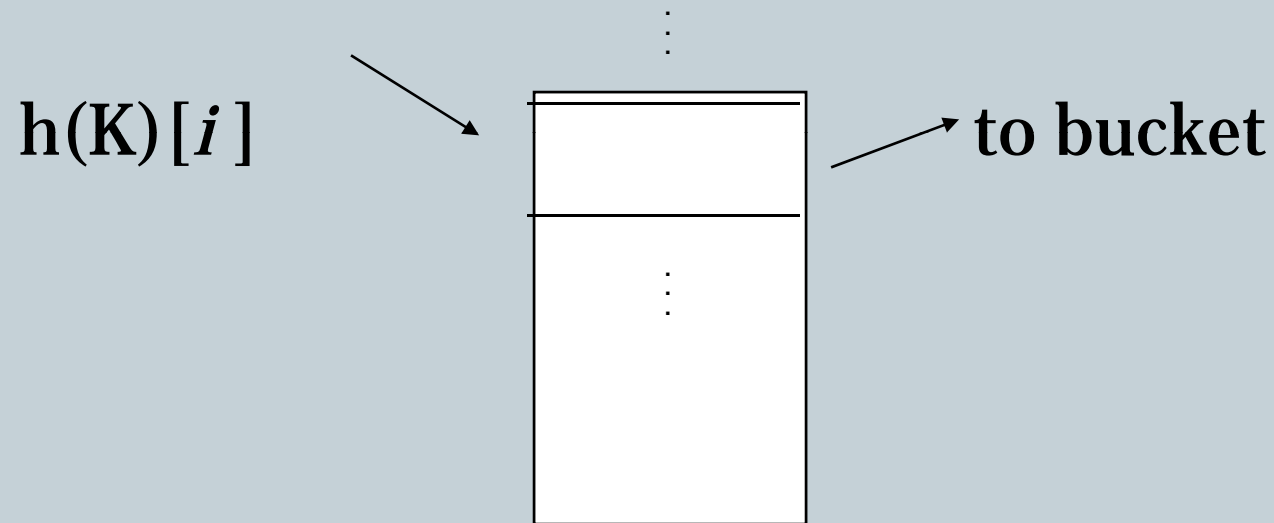
(a) Use i of b bits output by hash function



Extensible hashing: two ideas

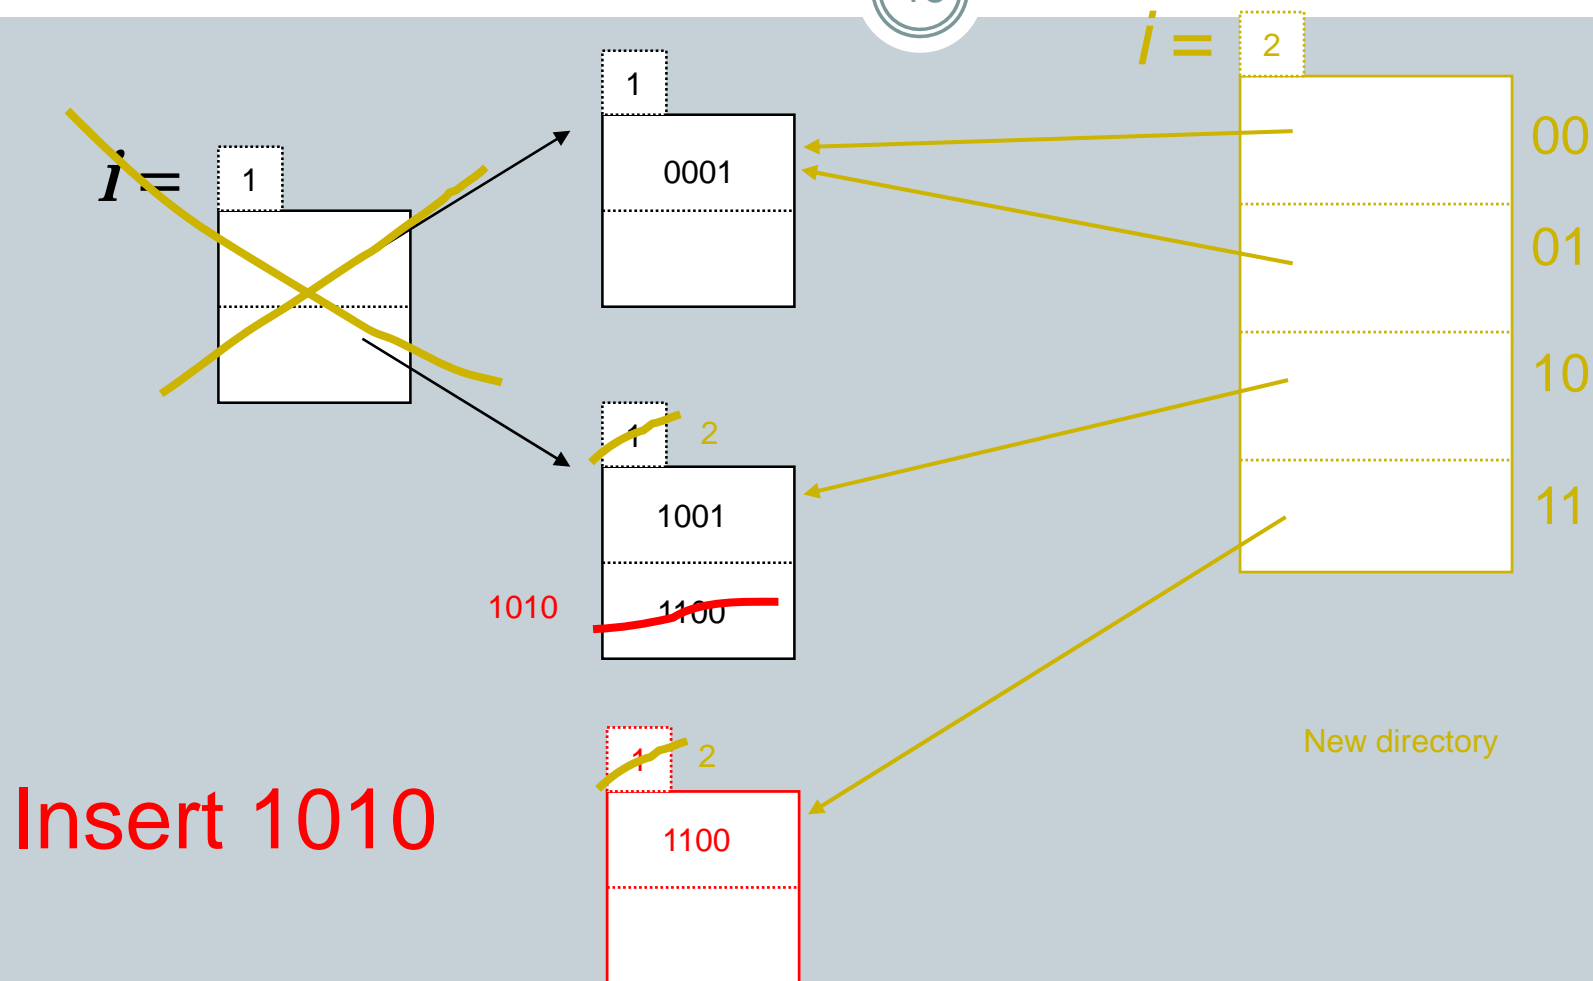
14

(b) Use directory

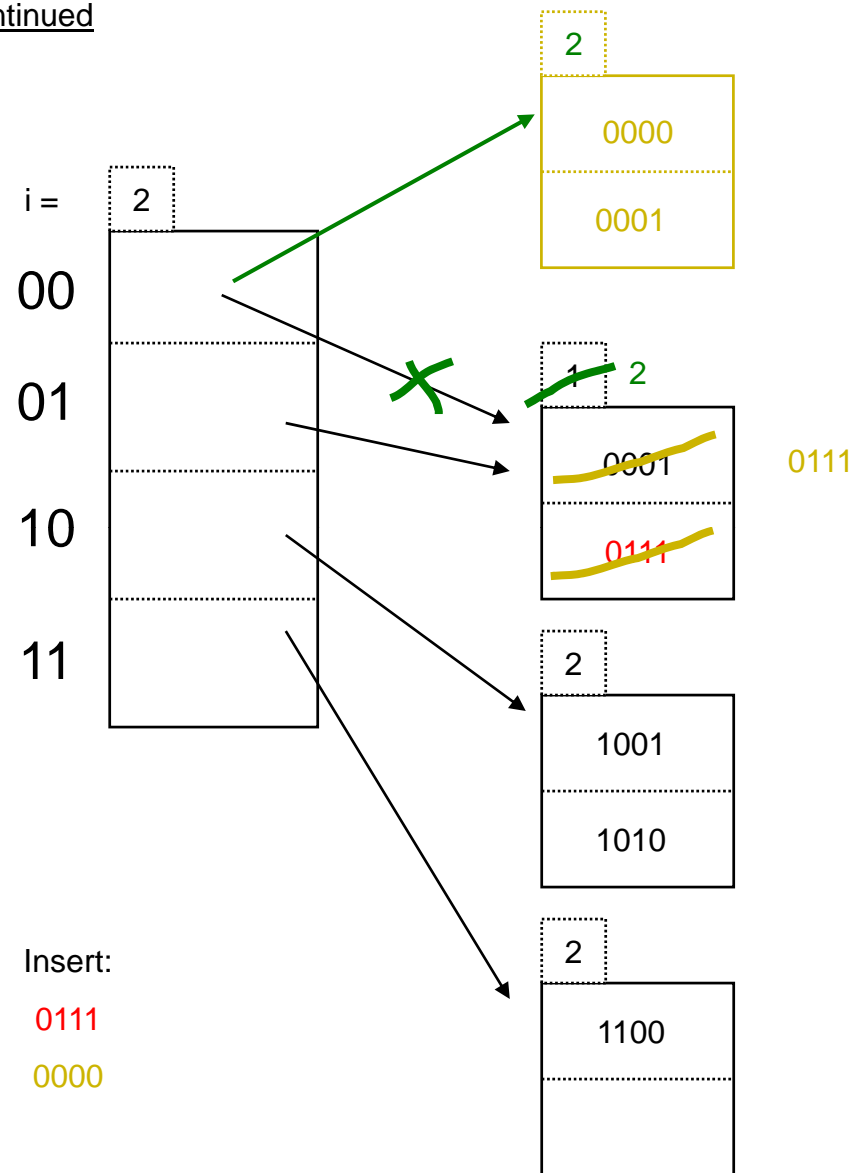


Example : $h(k)$ is 4 bits; 2 keys/bucket

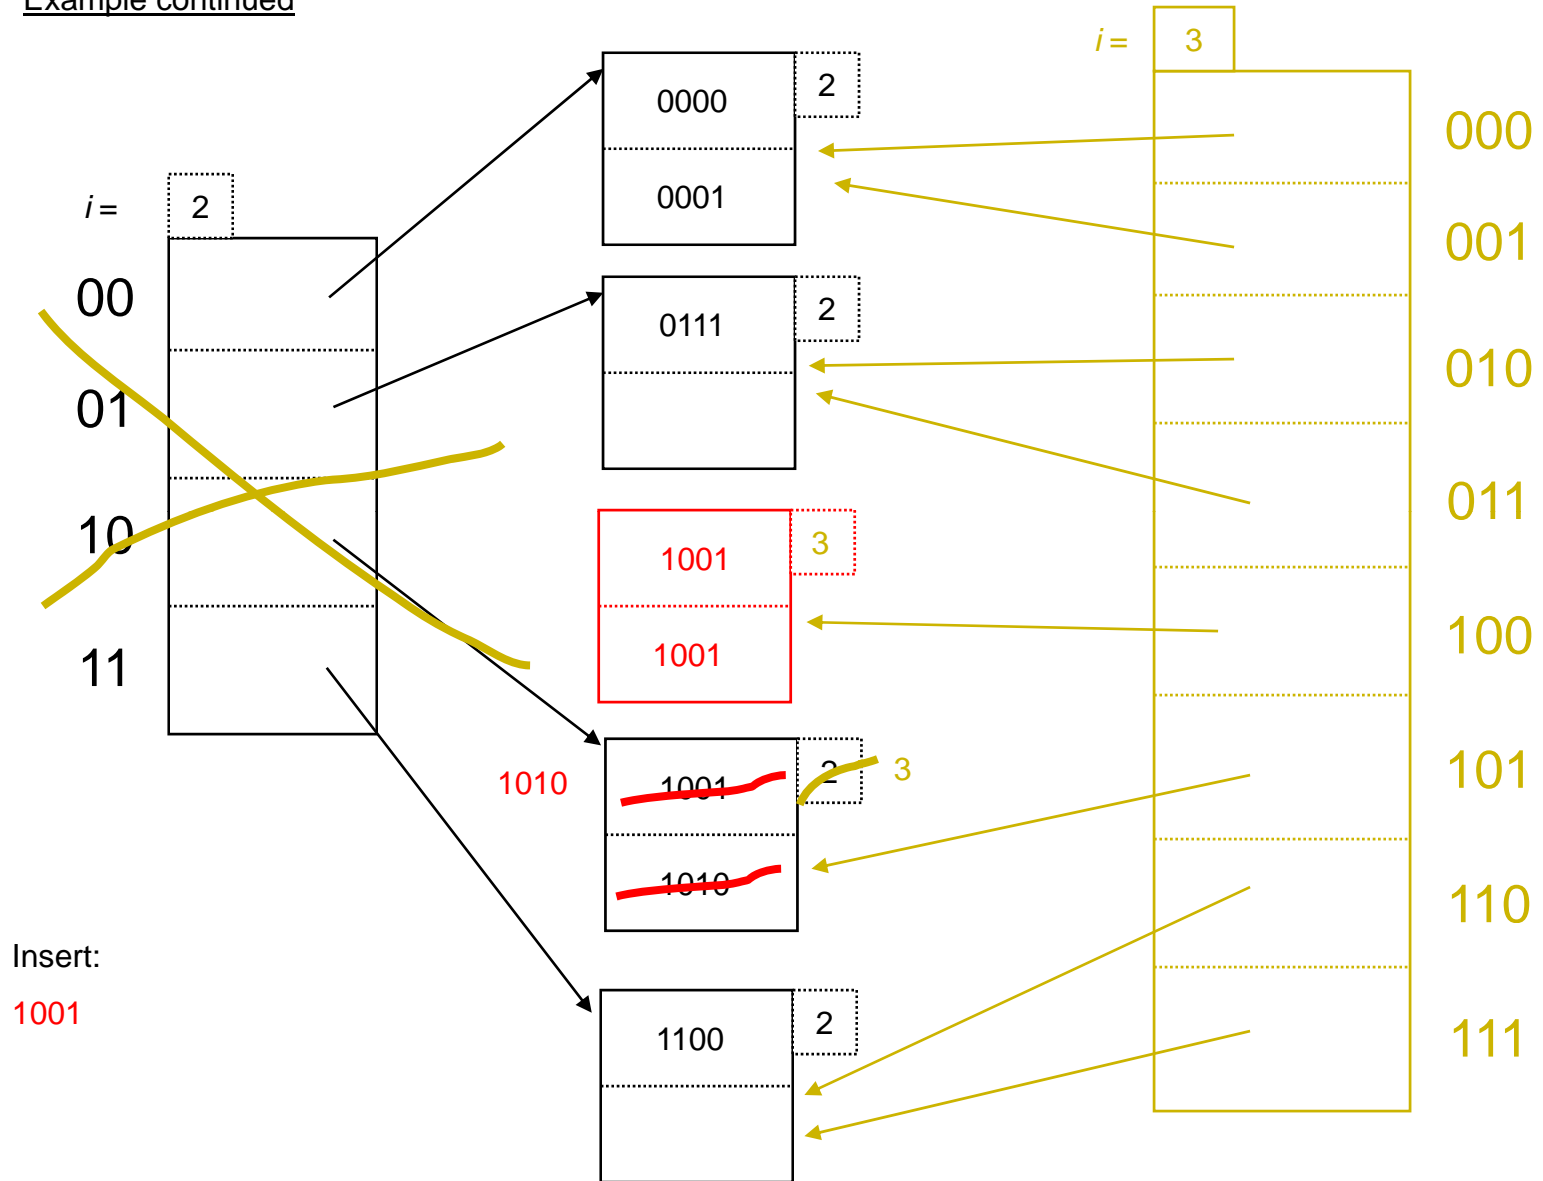
15



Example continued



Example continued



Extensible hashing: deletion

18

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

Deletion example:

19

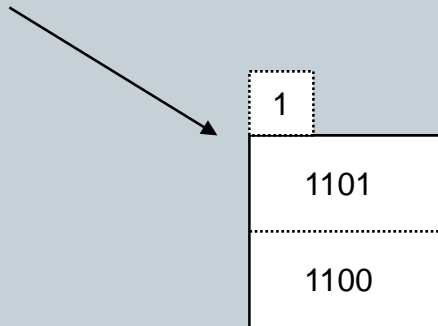
- **Run thru insert example in reverse!**

Note: Still need overflow chains

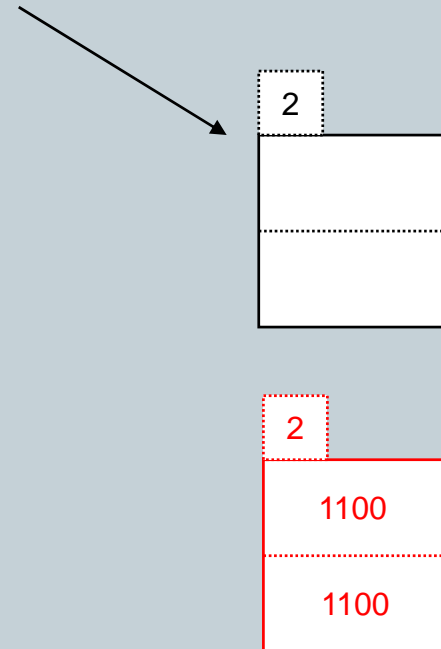
20

- Example: many records with duplicate keys

insert 1100



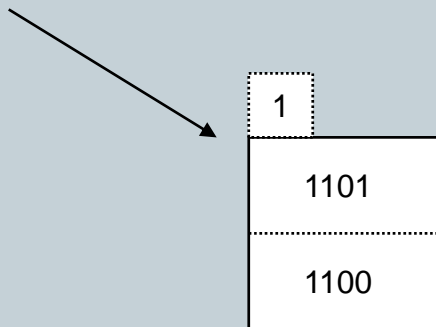
if we split:



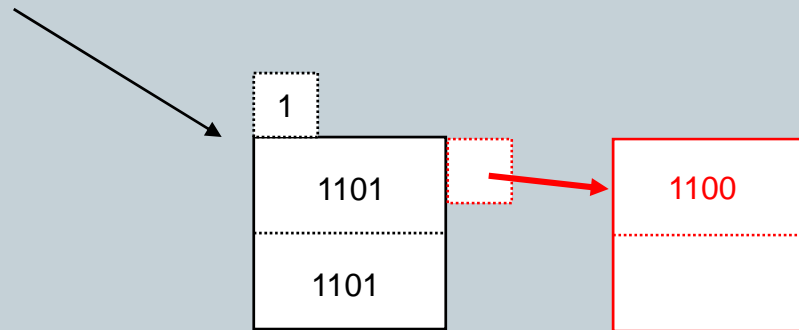
Solution: overflow chains

21

insert 1100



add overflow block:



Extensible hashing

22

- + 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)

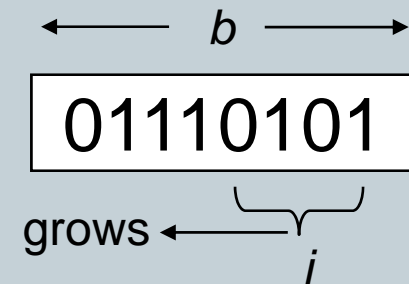
Linear hashing

23

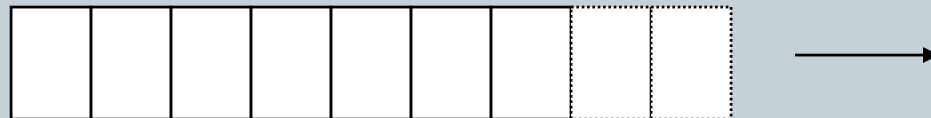
- Another dynamic hashing scheme

Two ideas:

(a) Use i low order bits of hash

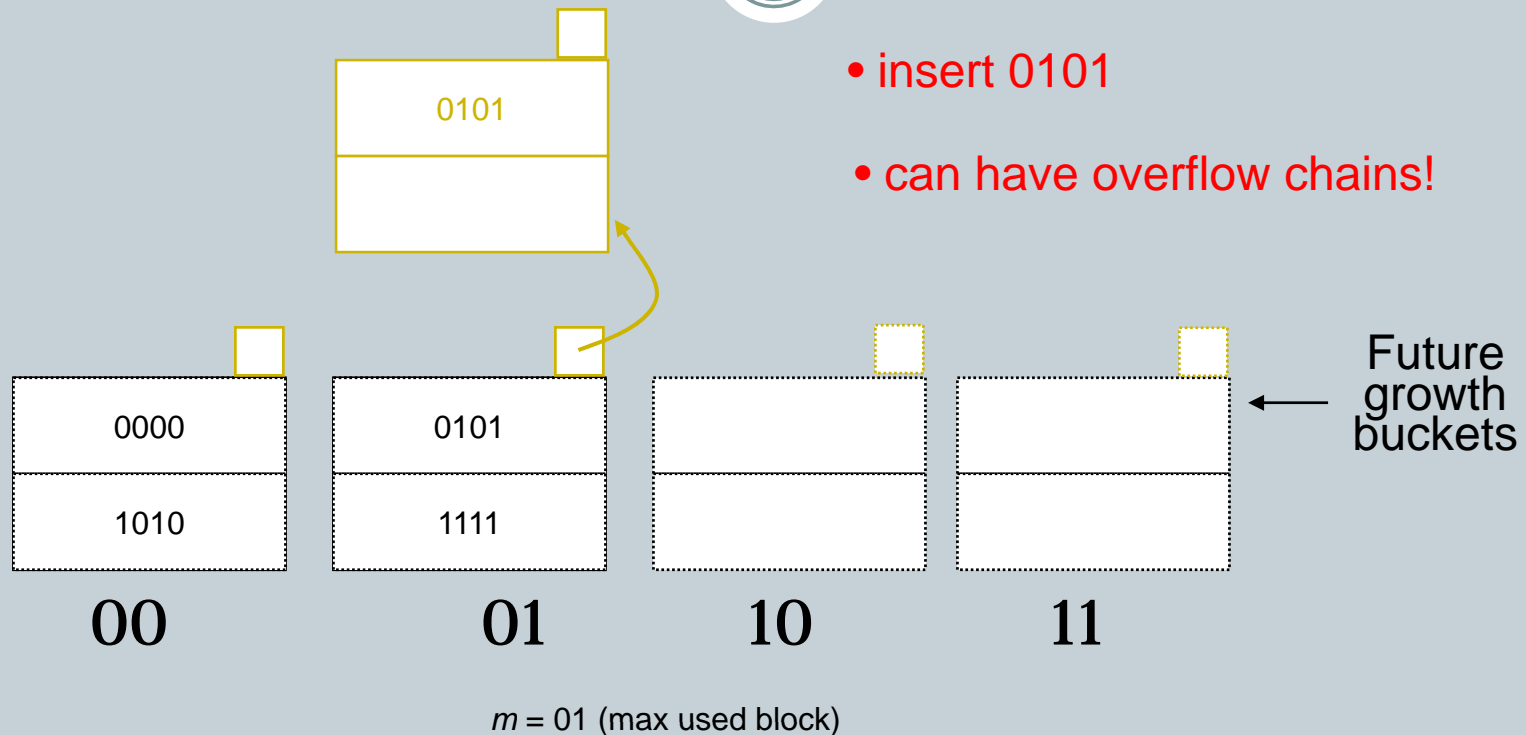


(b) File grows linearly



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

24



Rule

If $h(k)[i] \leq m$, then

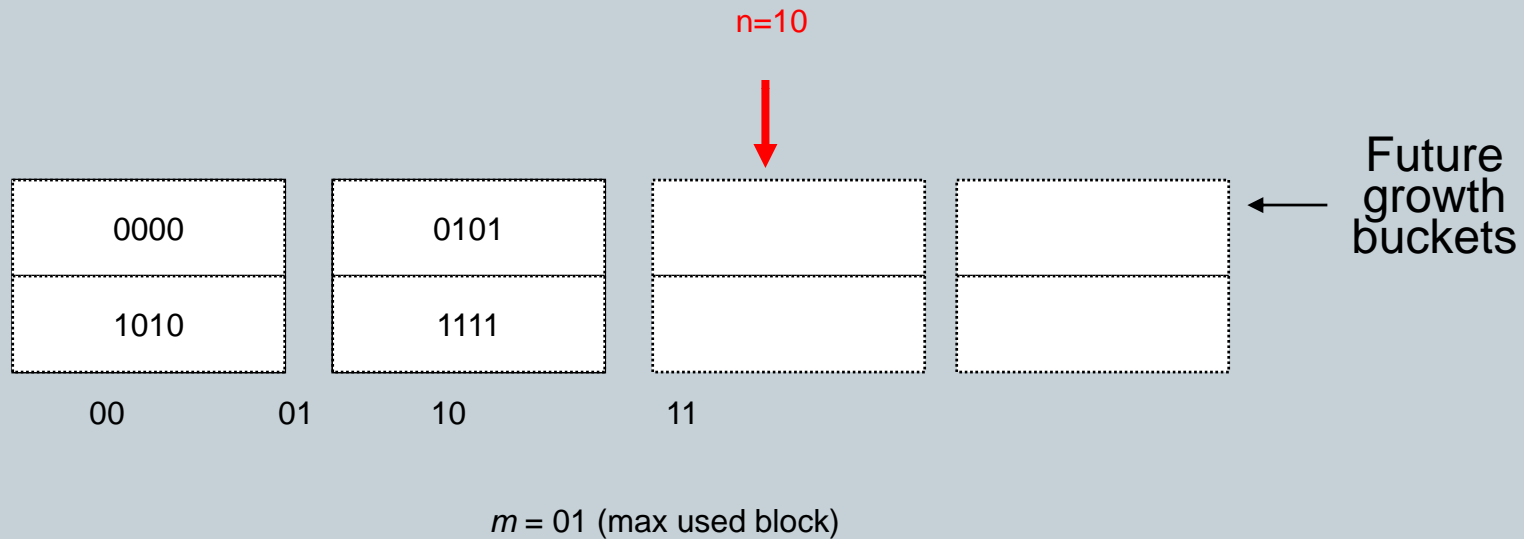
look at bucket $h(k)[i]$

else, look at bucket $h(k)[i] - 2^{i-1}$

Note

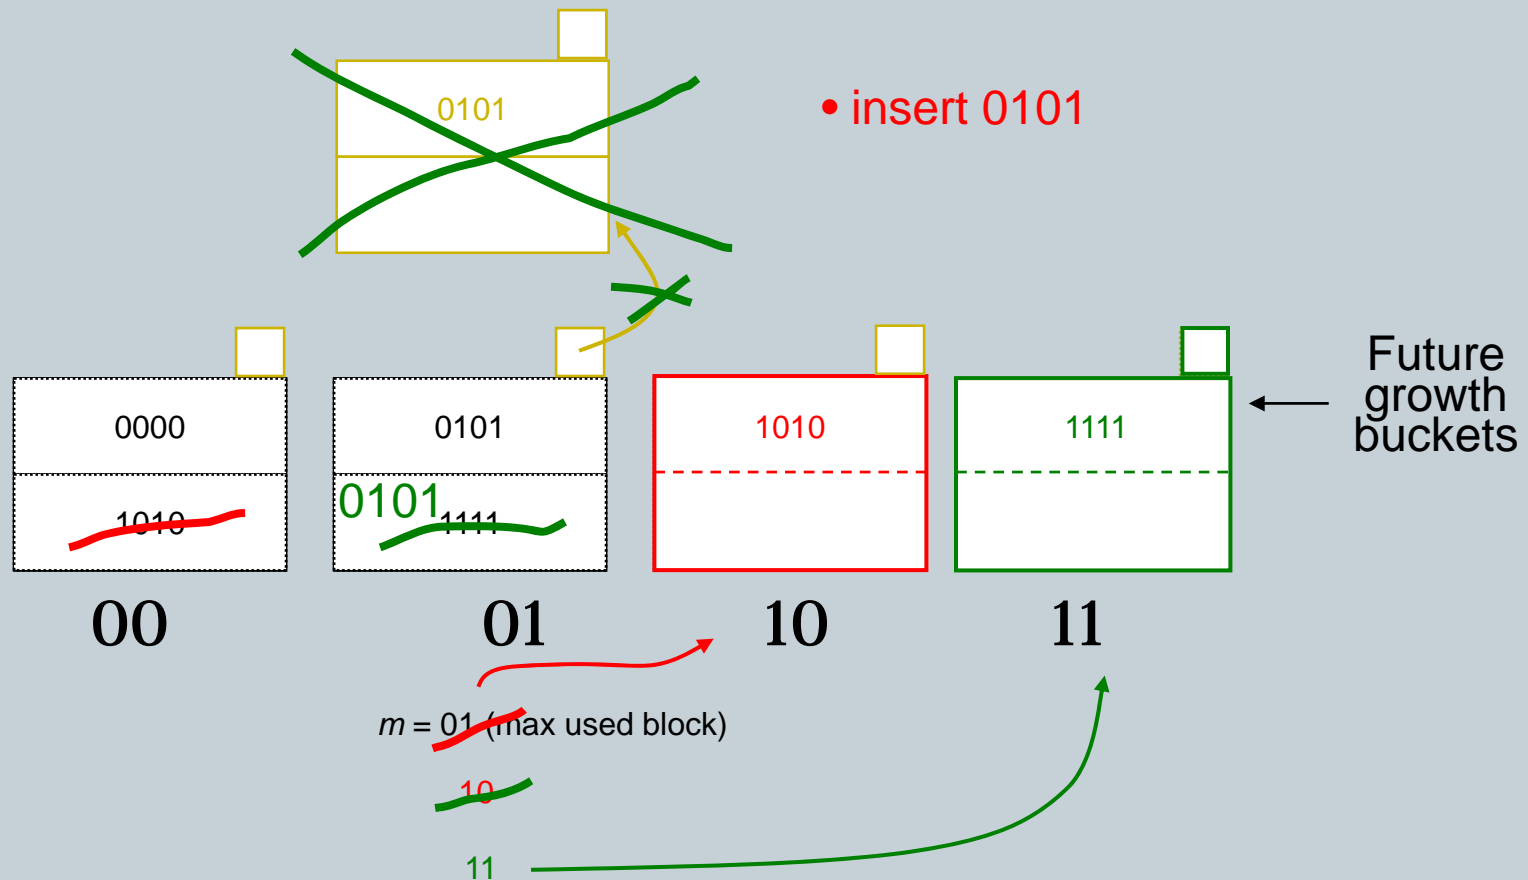
25

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



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

26



Example Continued: How to grow beyond this?

27

$i = 2$ ~~3~~

0000	0101	1010	1111		0101
	0101				0101

0 00

0 01

0 10

0 11

100

101

...

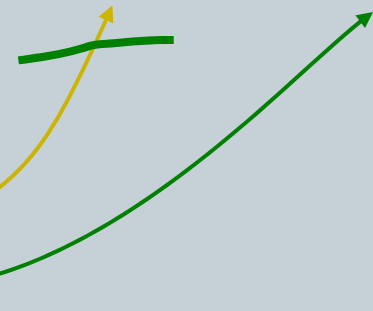
~~100~~ 101

~~110~~ 111

$m = 11$ (max used block)

100

101



When do we expand file?

28

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

Linear Hashing

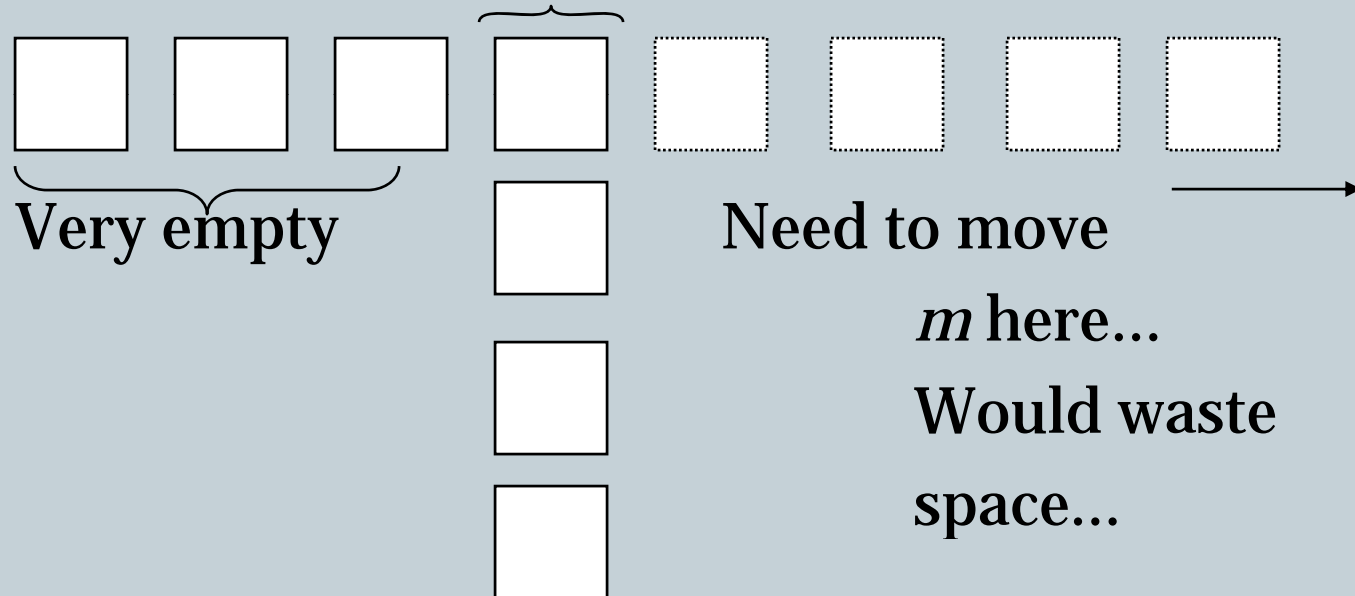
29

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

Example: BAD CASE

30

Very full



Summary

31

Hashing

- How it works
- Dynamic hashing
 - Extensible
 - Linear

Next:

32

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

Indexing vs. Hashing

33

- Hashing good for probes given key

e.g., `SELECT ...`

`FROM R`

`WHERE R.A = 5`

Indexing vs. Hashing

34

- INDEXING (Including B Trees) good for Range Searches:
e.g.,

```
SELECT  
FROM R  
WHERE R.A > 5
```

Index definition in SQL

35

- Create index name on rel (attr)
- Create unique index name on rel (attr)

└──────────→ defines candidate key

- Drop INDEX
name

SQL

36

CANNOT SPECIFY TYPE OF INDEX

(e.g. B-tree, Hashing, ...)

OR PARAMETERS

(e.g. Load Factor, Size of Hash,...)

ATTRIBUTE LIST \Rightarrow MULTIKEY INDEX

(next)

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

Multi-key Index

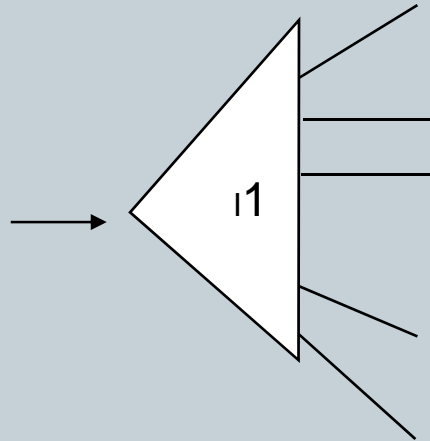
38

**Motivation: Find records where
DEPT = “Toy” AND SAL > 50k**

Strategy I:

39

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

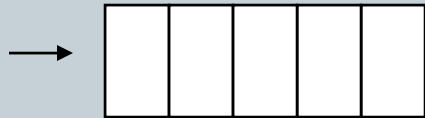


Strategy II:

40

- Use 2 Indexes; Manipulate Pointers

Toy



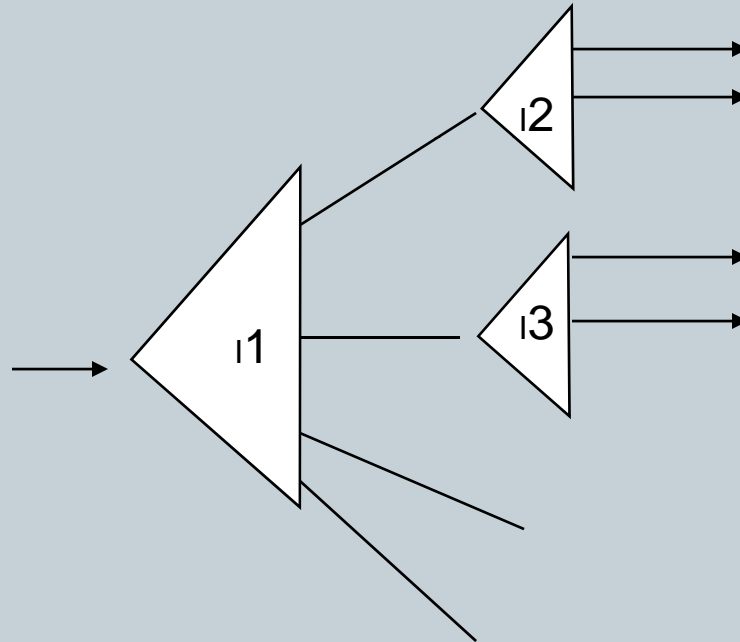
Sal
> 50k

Strategy III:

41

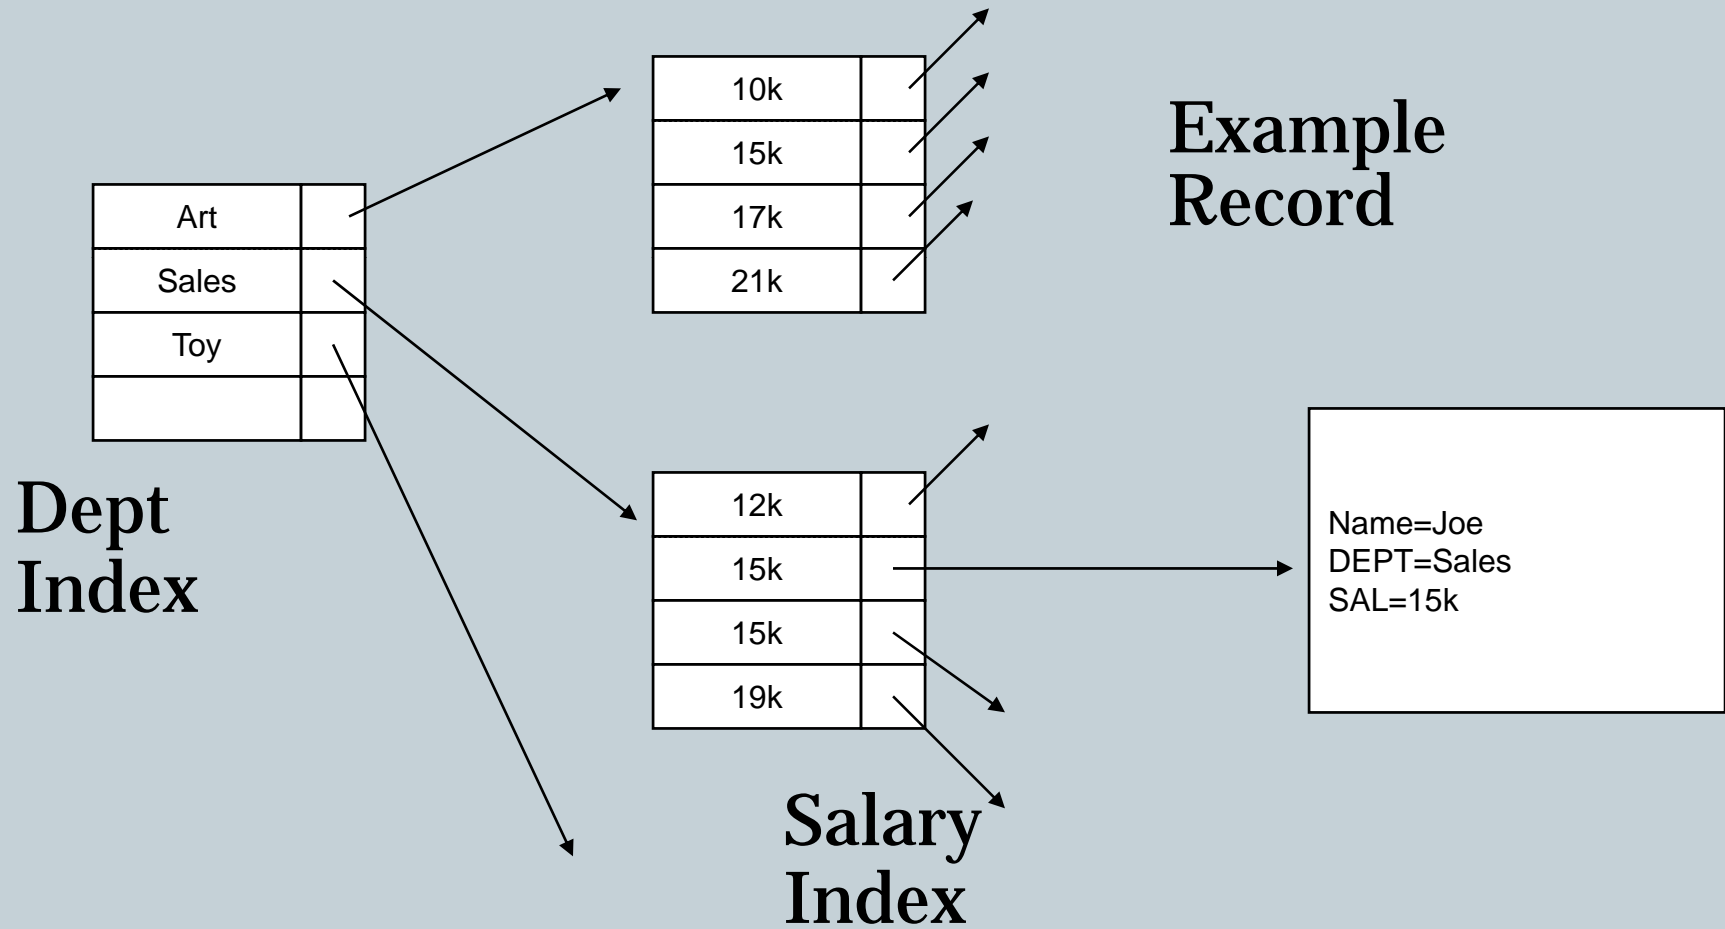
- **Multiple Key Index**

One idea:



Example

42



For which queries is this index good?

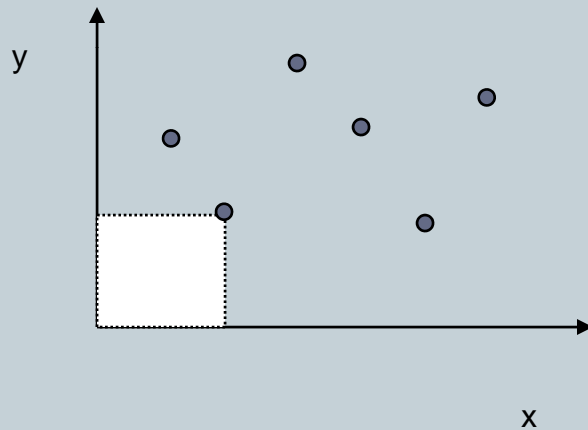
43

- ☐ 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:

44

- **Geographic Data**



DATA:

$\langle X1, Y1, \text{Attributes} \rangle$

$\langle X2, Y2, \text{Attributes} \rangle$

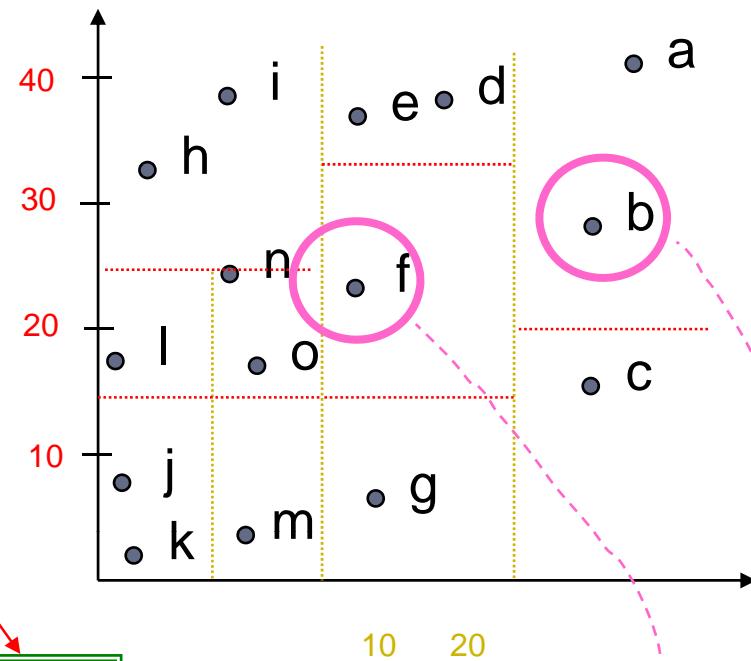
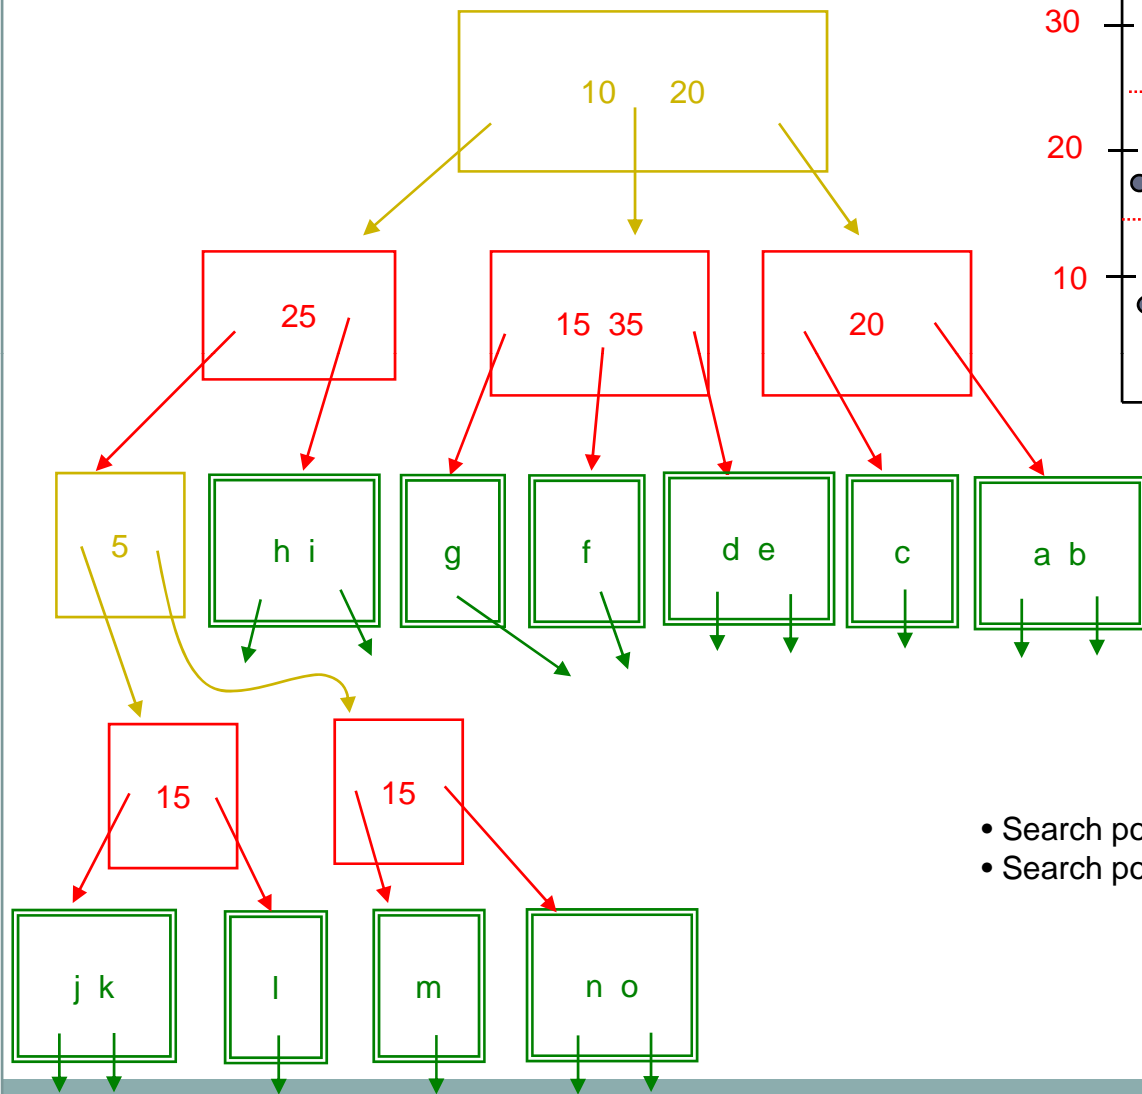
\vdots

Queries:

45

- 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 $\langle X_i, Y_i \rangle$?

Example



- Search points near f
- Search points near b

Queries

47

- 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 $b = \langle 7, 24 \rangle$

Geographic index

48

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

Two more types of multi key indexes

49

- **Grid**
- **Partitioned hash**

Grid Index



Key 2

X1 X2 Xn

Key 1

V1
V2
.....
Vn

To records with key1=V3, key2=X2

CLAIM

51

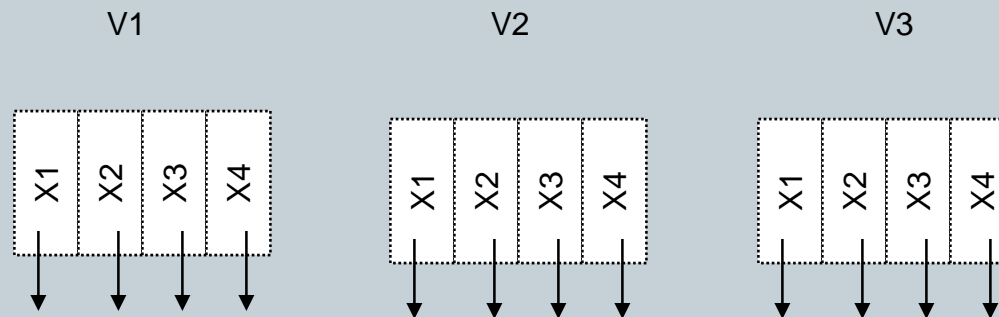
- Can quickly find records with
 - key 1 = $V_i \wedge$ Key 2 = X_j
 - key 1 = V_i
 - key 2 = X_j
- And also ranges....
 - E.g., key 1 $\geq V_i \wedge$ key 2 $< X_j$

But there is a catch with Grid Indexes!

52

- How is Grid Index stored on disk?

Like
Array...

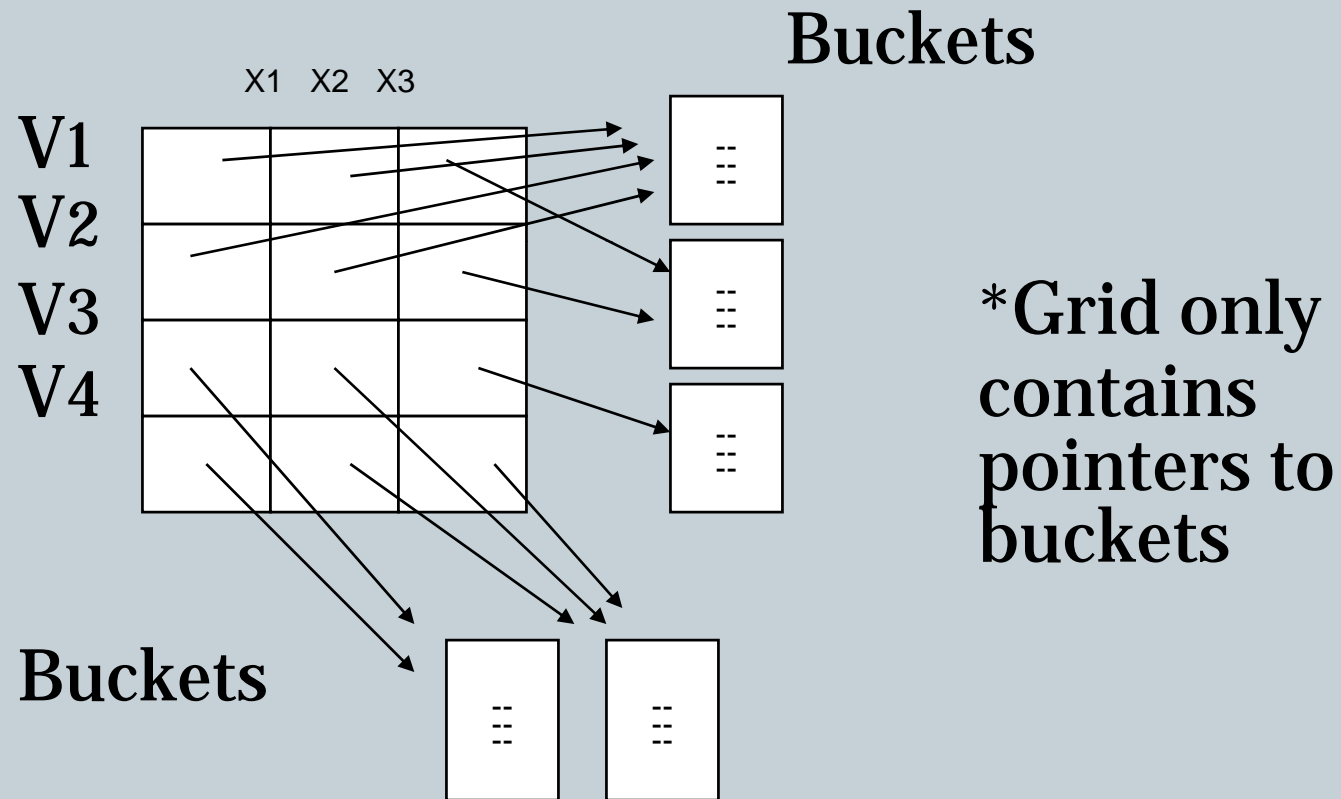


Problem:

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

Solution: Use Indirection

53



With indirection:

54

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

Can also index grid on value ranges

55

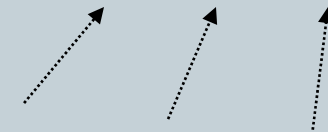
Salary

0-20K		1
20K-50K		2
50K- ∞		3

Grid

1	2	3
Toy	Sales	Personnel

Linear Scale



Grid files

56

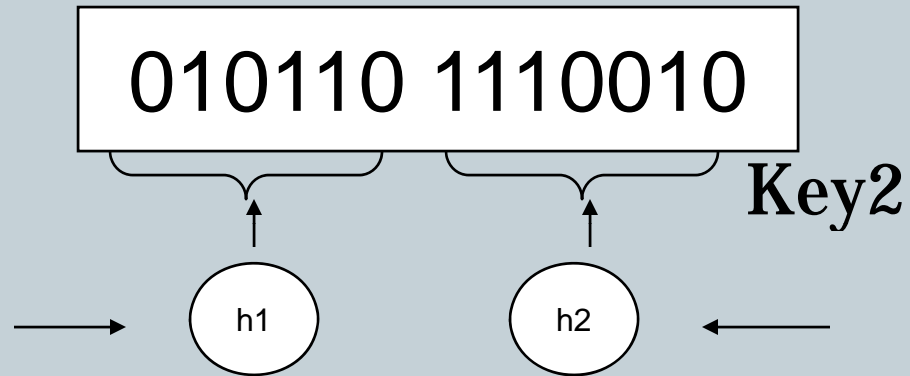
- ⊕ Good for multiple-key search
- ⊖ Space, management overhead (nothing is free)
- ⊖ Need partitioning ranges that evenly split keys

Partitioned hash function

57

Idea:

Key1

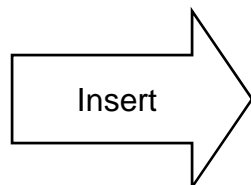


Key2

h1(toy)	=0	000
h1(sales)	=1	001
h1(art)	=1	010
:		011

h2(10k)	=01	100
h2(20k)	=11	101
h2(30k)	=01	110
h2(40k)	=00	111
:		

<Fred>
<Joe><Sally>



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

h1(toy) =0

h1(sales) =1

h1(art) =1

.

h2(10k) =01

h2(20k) =11

h2(30k) =01

h2(40k) =00

.

.

000

001

010

011

100

101

110

111

<Fred>
<Joe><Jan>
<Mary>
<Sally>
<Tom><Bill>
<Andy>

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

h1(toy) =0

h1(sales) =1

h1(art) =1

.

h2(10k) =01

h2(20k) =11

h2(30k) =01

h2(40k) =00

.

- Find Emp. with Sal=30k

000	<Fred>
001	<Joe><Jan>
010	<Mary>
011	
100	<Sally>
101	
110	<Tom><Bill>
111	<Andy>

look here

h1(toy) =0

h1(sales) =1

h1(art) =1

.

h2(10k) =01

h2(20k) =11

h2(30k) =01

h2(40k) =00

.

- Find Emp. with Dept. = Sales

000

001

010

011

100

101

110

111

<Fred>
<Joe><Jan>
<Mary>
<Sally>
<Tom><Bill>
<Andy>

look here

Summary

62

Post hashing discussion:

- Indexing vs. Hashing
- SQL Index Definition
- Multiple Key Access
 - Multi Key Index

Variations: Grid, Geo Data

- Partitioned Hash