

Data Structures: Indexing

Indexing and Hashing

People(name, age, address, title)


SELECT *

FROM People

WHERE age = 35

How to answer the query fast?

age=35



record



Outline

- Dense/Sparse Indexes
 - Indexes on Sorted (Sequential) Files
 - Indexes on Unsorted Files (Secondary Indexes)
 - Applications of Secondary Indexes and Buckets
- B Trees
- Hashing
- Multi-dimensional Indexes

Age
(search key)

Sorted File



10	
20	

30	
40	


50	
60	

70	
80	

90	
100	

Dense Index

Sorted File




Pointer for **each**
key value (age)

Can do binary search ..

Sparse Index


Sorted File



Pointer for the
lowest key in
each block

Sparse 2nd level

Sorted File



Dense/Sparse Index Summary


- Dense: Store pointer to each record/key value.
 - Can tell if a record exists without accessing the file
 - Efficient access
 - Needed for secondary indexes (as we'll see later)
- Sparse: Store pointer to the lowest key in each block.
 - Less index space. May fit in memory.
 - Better for insertions/deletions

Insertions/Deletions in Sorted Files

- Many options:
 - Leave extra space in data and/or index blocks to accommodate expansion
 - Overflow blocks (next 2 slides)
 - Localized reorganization (moving records within or among blocks)
 - Complete reorganization (done once in a while, in conjunction with above techniques)

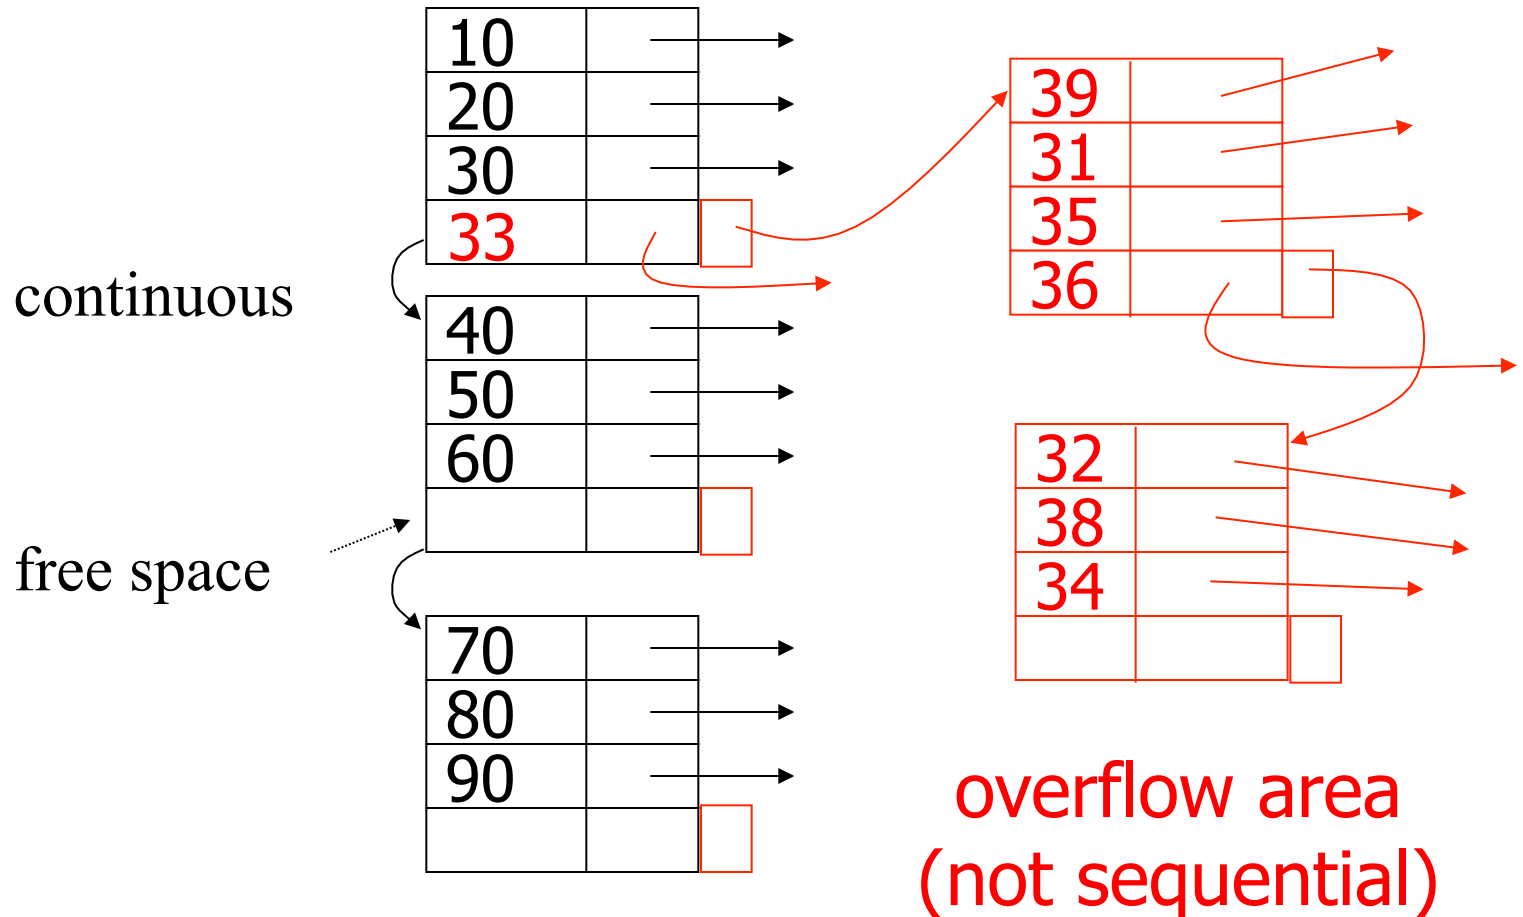
Overflow Data Blocks

– insert record 25



Overflow Index Blocks

Index (sequential)




Outline

- Dense/Sparse Indexes
 - Indexes on Sorted (Sequential) Files
 - Indexes on Unsorted Files (Secondary Indexes)
 - Applications of Secondary Indexes and Buckets
- B Trees
- Hashing
- Multi-dimensional Indexes

Secondary indexes

Sequence
field



30	
50	

20	
70	

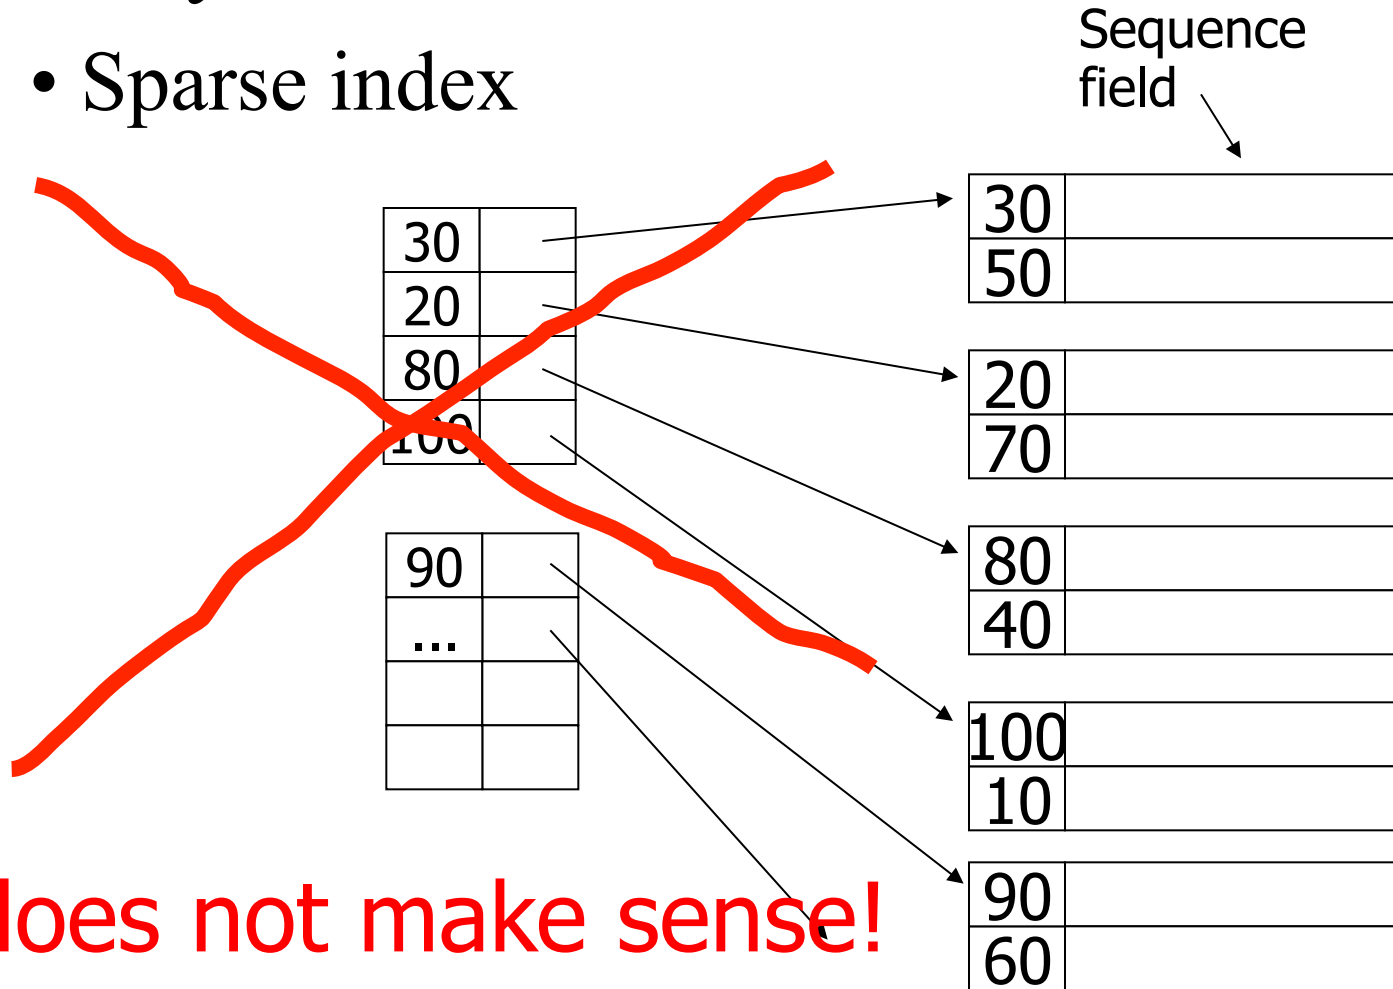
80	
40	

100	
10	

90	
60	

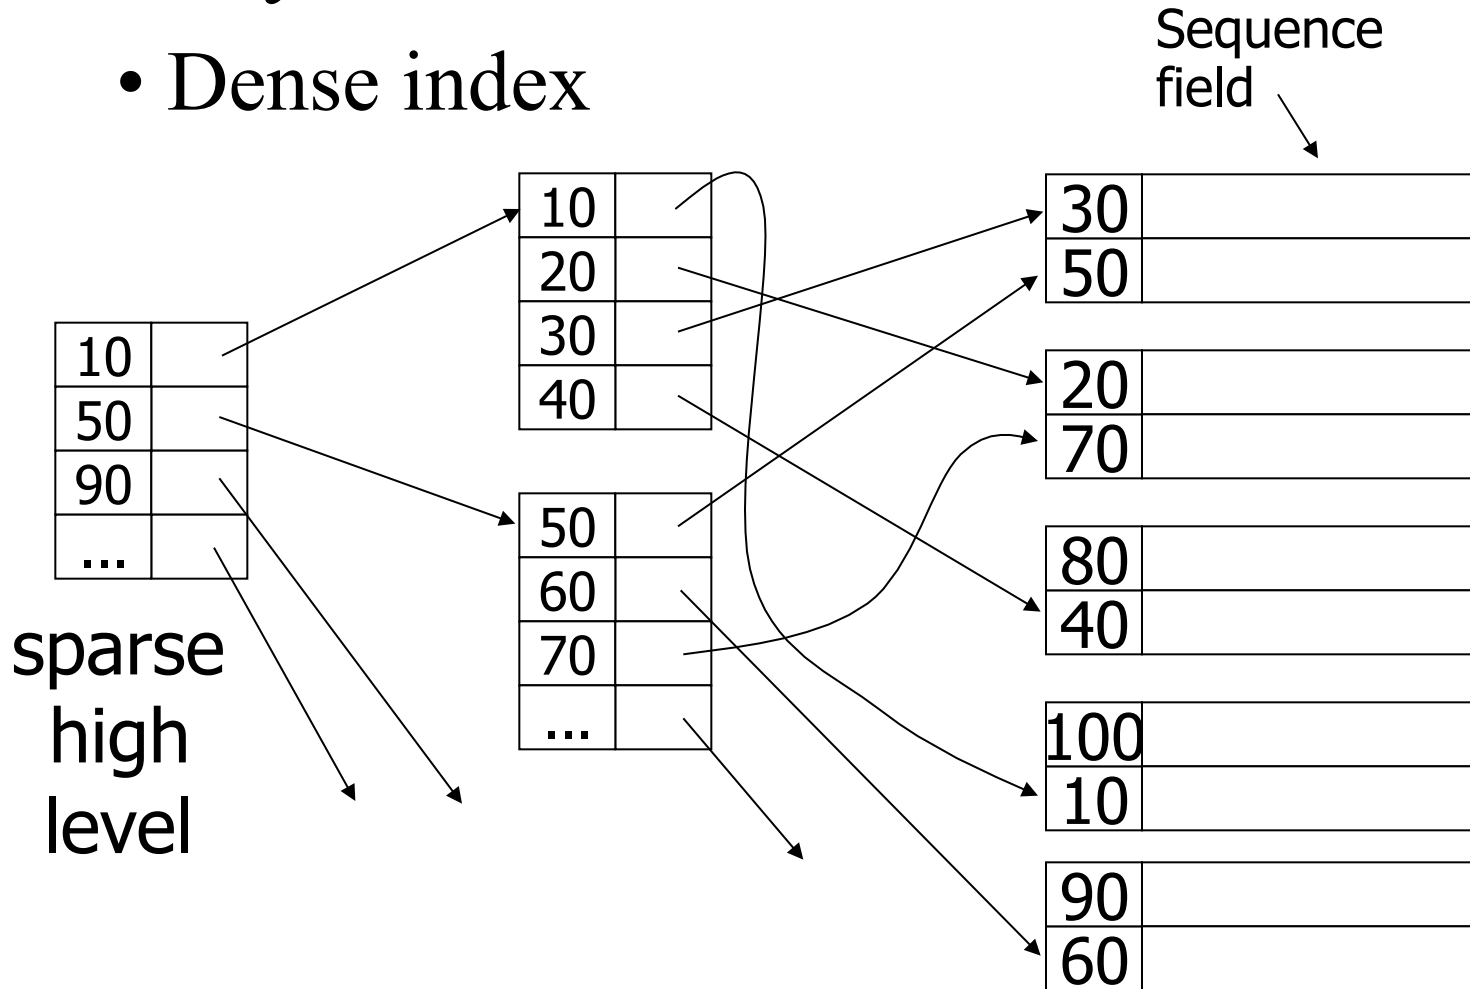
Secondary indexes

- Sparse index



Secondary indexes

- Dense index



With secondary indexes:

- Lowest level **has to be** dense
- Other levels are sparse

Duplicate values?

20	
10	


20	
40	

10	
40	

10	
40	

30	
40	

Duplicate values: Use “buckets”



buckets

Outline

- Dense/Sparse Indexes
 - Indexes on Sorted (Sequential) Files
 - Indexes on Unsorted Files (Secondary Indexes)
 - Applications of Secondary Indexes and Buckets
- B Trees
- Hashing
- Multi-dimensional Indexes

Use-case for Secondary Indexes: Clustered Files

Movie(title, year, studioName, ..., ...)

Studio(name, address, president)

Typical Query

Select title, year

From Movie JOIN Studio

Where president = “xxx”

Clustered Files

- Previous example:
 - Efficient to store movies with its studio
- Answering the Query:
 - Find studio record using an index on ‘president’
 - Find movies stored **with** the studio record found
 - However, **index on movies must be secondary.**

Studio A		Studio B		Studio C	
		Movies of studio A		Movies of studio B	

Applications of Buckets

Indexes

Name: primary


Dept: secondary

Floor: secondary

Table

EMP (name,dept,floor,...)

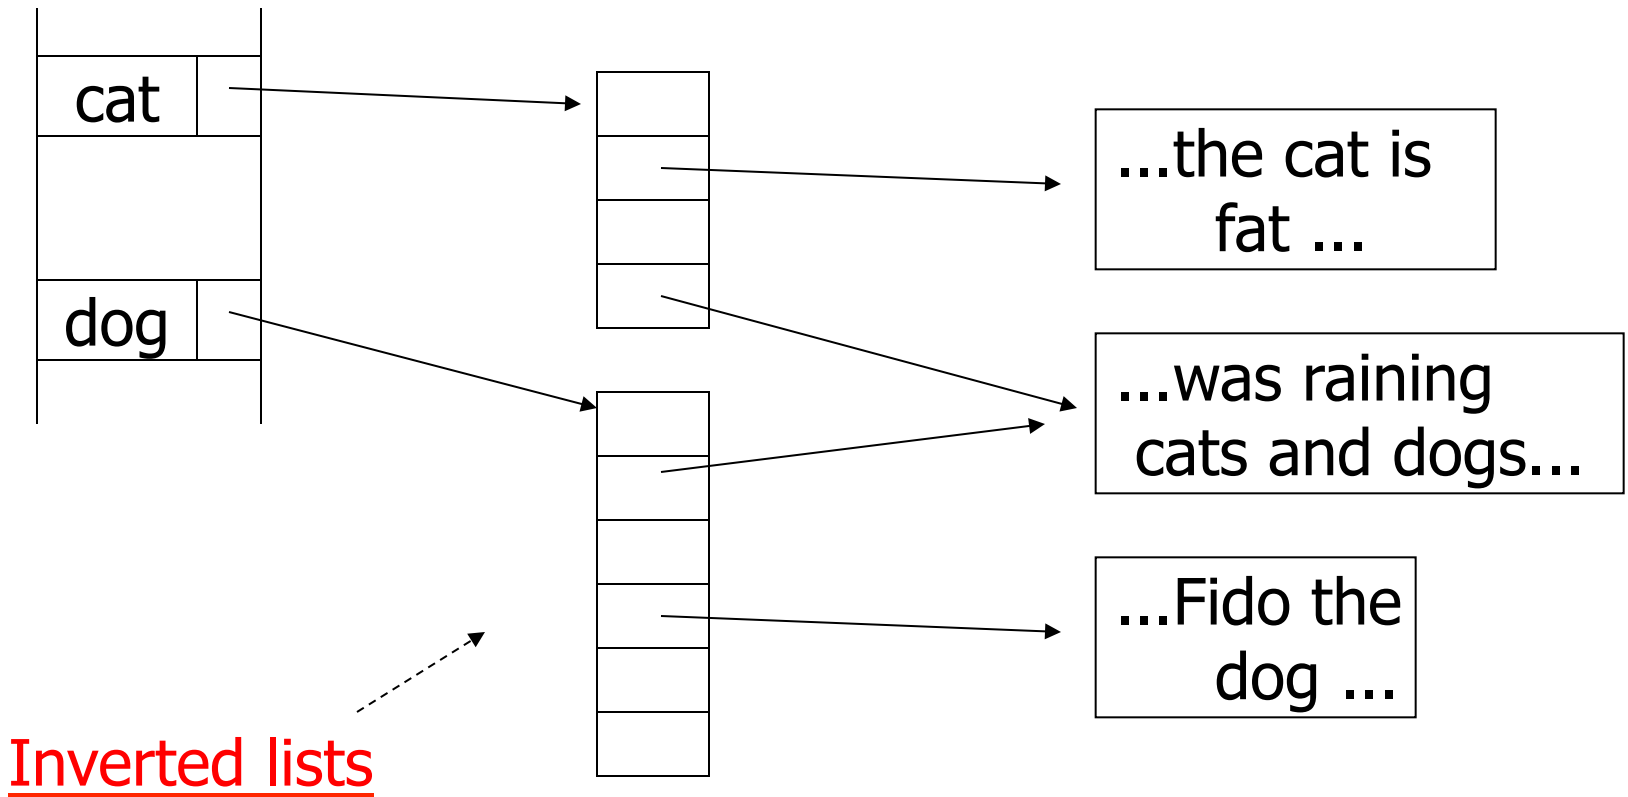
Query: Get employees in (Toy Dept) \wedge (2nd floor)



→ **Intersect** toy bucket and 2nd Floor buckets to get set of matching EMP's

Inverted Lists in Information Retrieval

Documents



IR QUERIES

- Find articles with “cat” and “dog”
- Find articles with “cat” or “dog”
- Find articles with “cat” and not “dog”


- Find articles with “cat” in title
- Find articles with “cat” and “dog”
within 5 words

Outline


- Dense/Sparse Indexes
 - Indexes on Sorted (Sequential) Files
 - Indexes on Unsorted Files (Secondary Indexes)
 - Applications of Secondary Indexes and Buckets
- B Trees
- Hashing
- Multi-dimensional Indexes

B Tree Example


Root



Sample non-leaf



Sample leaf node



Main Challenge in B Trees

- To ensure good search cost, the tree should be kept of **minimum and uniform height**. I.e., keep the tree:
 - full (packed), and
 - balanced (almost uniform height).
- This can be difficult -- in face of insertions and deletions.

Solution: **Keep the tree “semi-full” (i.e., each tree node half-full).**

- Makes it easy to keep the tree balanced and semi-optimal.

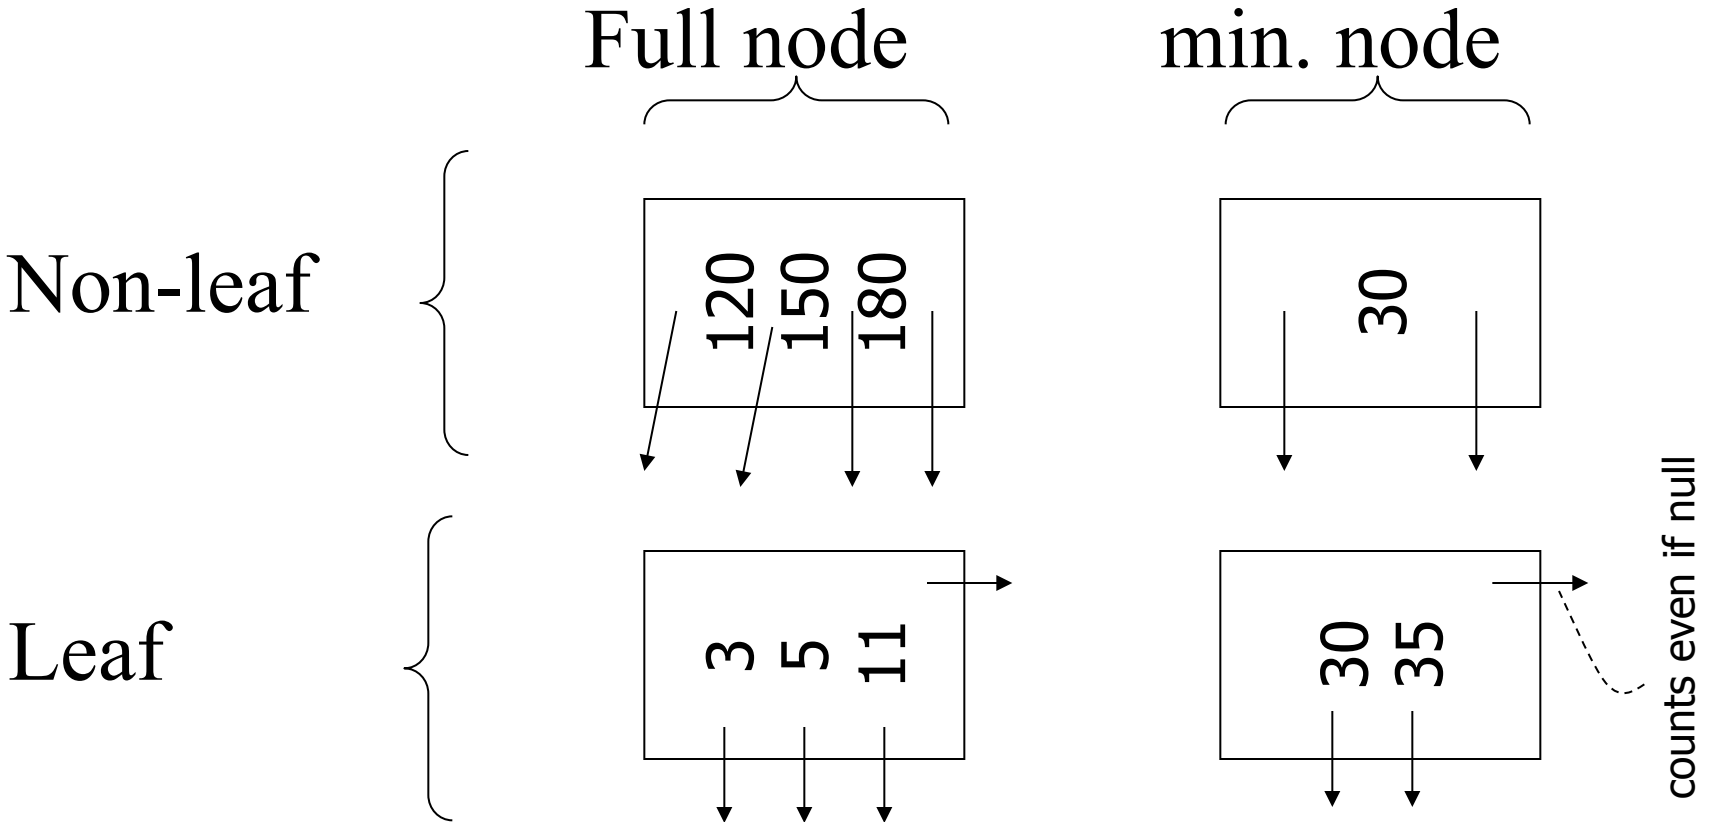
Size of nodes

$$\left\{ \begin{array}{l} n+1 \text{ pointers} \\ n \text{ keys} \end{array} \right. \quad \text{(fixed)}$$

Use at least (to ensure a **balanced** tree)

- Non-leaf: $\lceil (n+1)/2 \rceil$ pointers (to nodes)
- Leaf: $\lfloor (n+1)/2 \rfloor$ pointers (to data)

n=3




Insertions into B tree

- (a) simple case
 - space available in leaf
- (b) leaf overflow
- (c) non-leaf overflow
- (d) new root


(a) Insert key = 32

n=3




(b) Leaf Overflow: Insert key = 7

n=3



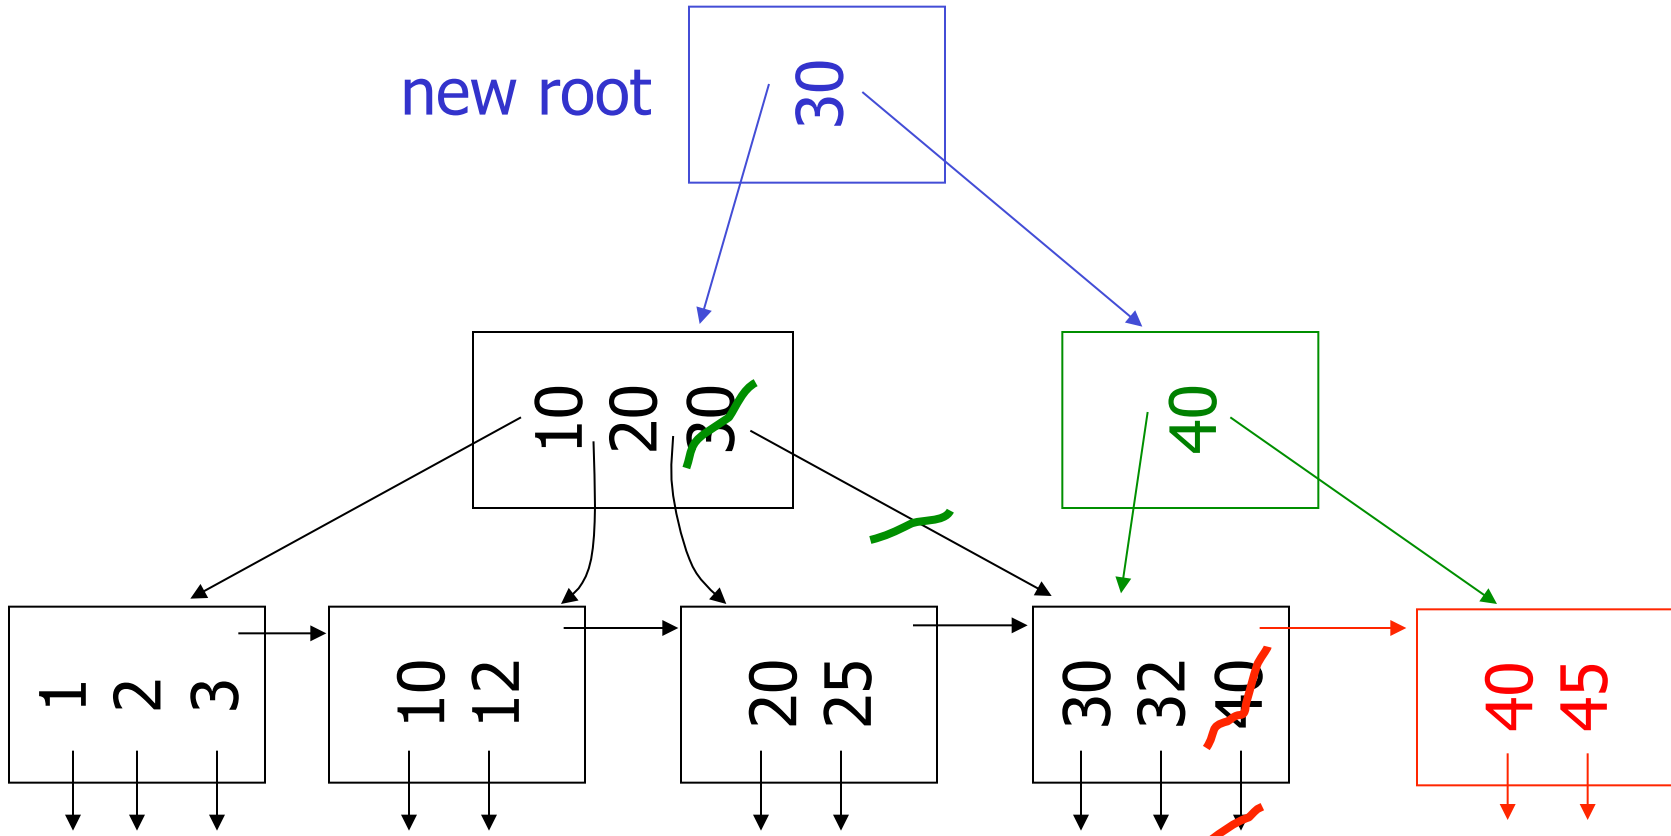
(c) Non-leaf Overflow: Insert key = 160

n=3



(d) New root: Insert 45

n=3




Deletion from B tree

- (a) Simple case - no example
- (b) Coalesce with neighbor (sibling)
- (c) Re-distribute keys
- (d) Cases (b) or (c) at non-leaf

(b) Coalesce with sibling

– Delete 50


n=4



(c) Redistribute keys

– Delete 50

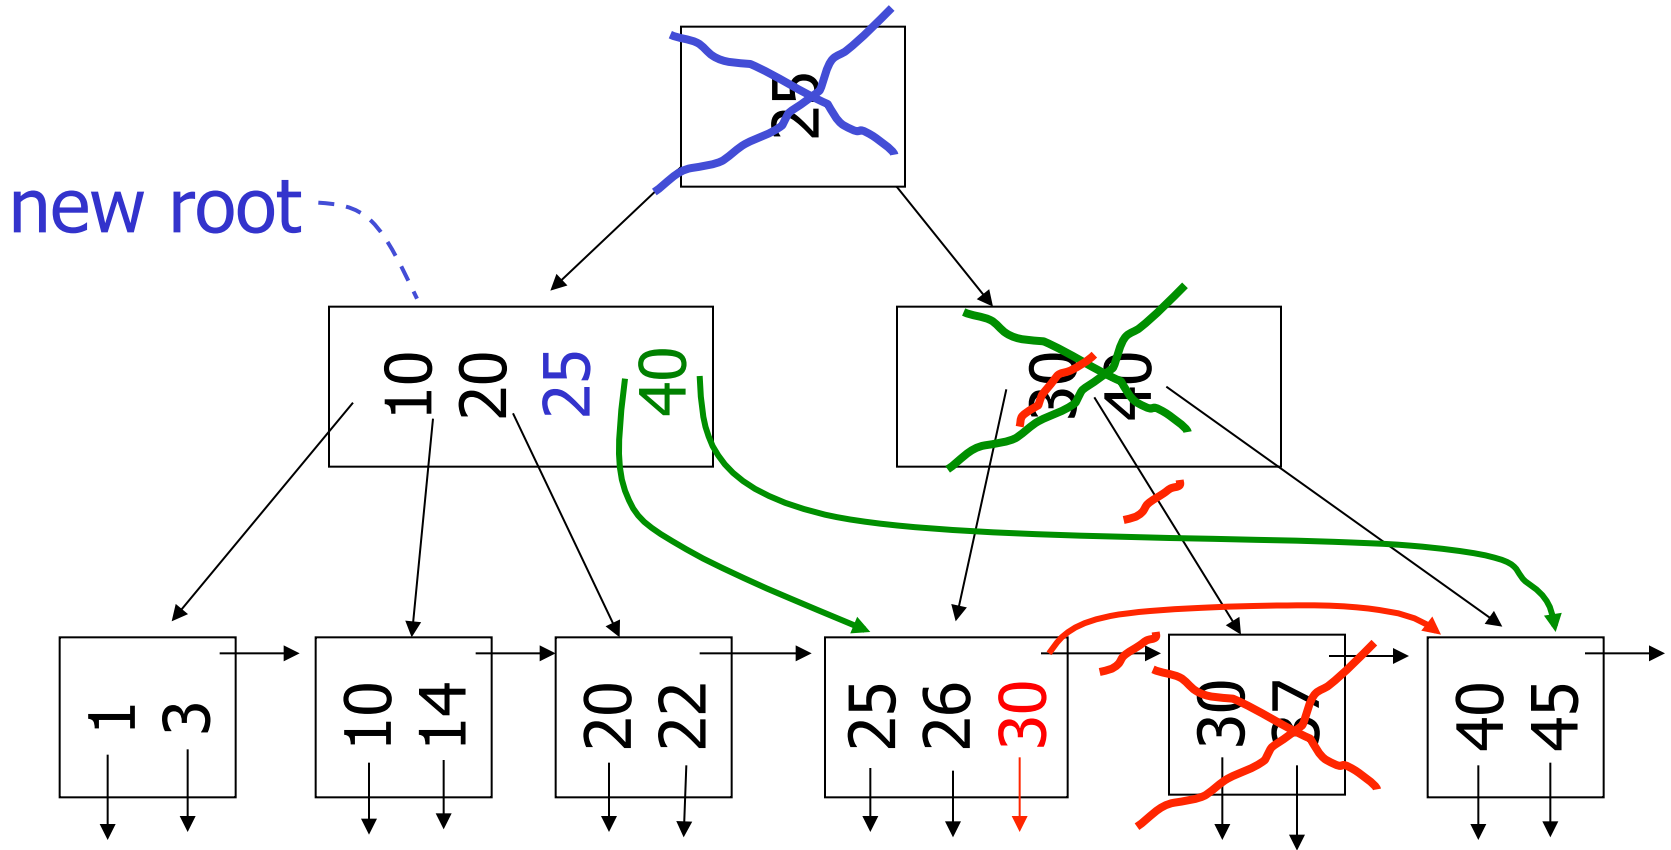
n=4



(d) Non-leaf coalesce

– Delete 37

n=4



B-tree deletions in practice

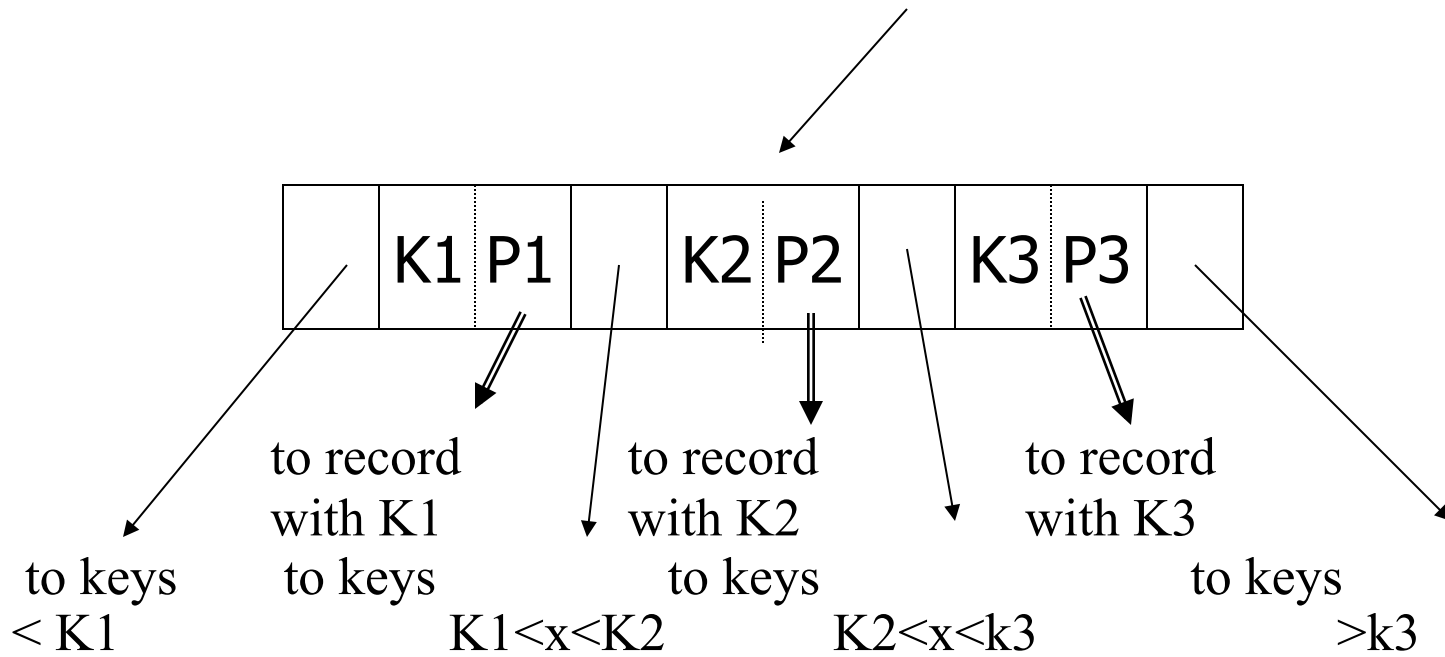
- Often, coalescing is not implemented
 - Too hard and not worth it!

B-Trees vs Indexes on Sorted Files

- B Tree Advantages:
 - Fast inserts and deletes (splits/merges are rare) at the cost of sub-optimal number of levels.
 - But, in most practical cases – the number of levels is still only 3-4.


Another Variant of B-Trees

- What we have discussed till now are actually called “B+ trees”.
- The original B-Tree (without the +):
 - Avoids key values being stored multiple times
 - Has pointers to records in non-leaf nodes also



B-tree Variant Example (n=2)

- sequence pointers
not useful now!
(but keep space for simplicity)



Pros and Cons of the Variant

Pros

- Has faster lookup

Cons

- Non-leaf and leaf has different sizes
- For fixed block size, less “bushier”

➡ Not preferred

Outline

- Dense/Sparse Indexes
- B Trees
- Hashing
 - Overall Idea
 - Handling growth: (i) Extensible Hashing, (ii) Linear Hashing
- Multi-dimensional Indexes

Hashing

key \rightarrow h(key)







Buckets
(typically 1
disk block)

Two Alternatives


(i) Buckets contain the tuples

key \rightarrow h(key)



Two Alternatives

(ii) Buckets contain **pointers** to tuples



- Above needed if tuples are already stored (e.g., for secondary indexes)

Example hash function

- Key = 'x₁ x₂ ... x_n' *n* byte character string
- Have *b* buckets
- h: add x₁ + x₂ + x_n
 - compute sum modulo *b*

EXAMPLE 2 records/bucket

INSERT:

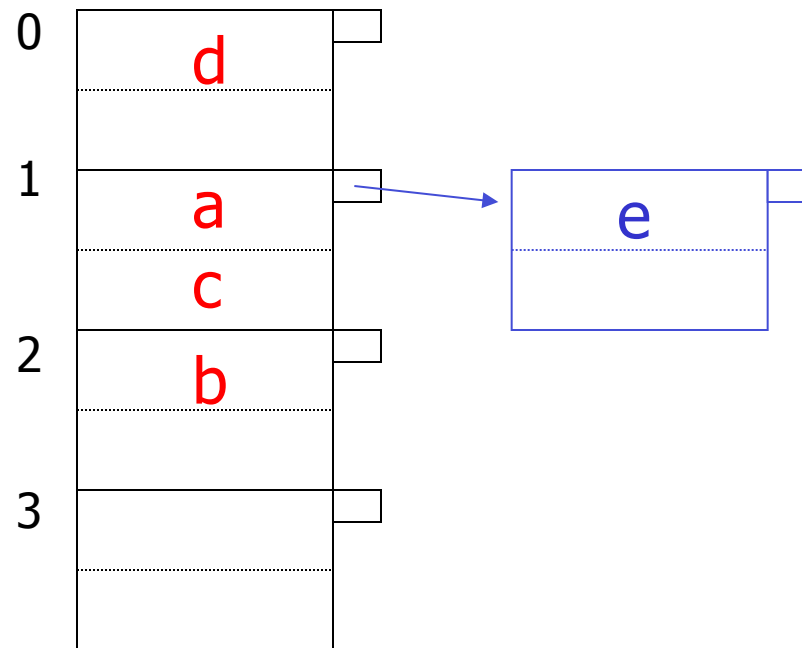
$$h(a) = 1$$

$$h(b) = 2$$

$$h(c) = 1$$

$$h(d) = 0$$

$$h(e) = 1$$




EXAMPLE: deletion

Delete:

e

f

c



Good hash
function:

Expected # of keys/bucket
is the same for all buckets

Within a bucket:

- Do we keep keys sorted?
- Yes, if CPU time critical and inserts/deletes are not too frequent

Rule of thumb:


Space Utilization =

$$\frac{\text{\# keys in the system}}{\text{total \# keys that fit (without overflows)}}$$

- Keep space utilization between 50% to 80%
- If $< 50\%$, wasting space
- If $> 80\%$, overflows significant

How do we cope with growth?

- Overflows and reorganizations
- Dynamic hashing




- Extensible
- Linear

Outline

- Dense/Sparse Indexes
- B Trees
- Hashing
 - Overall Idea
 - Handling growth: (i) Extensible Hashing, (ii) Linear Hashing
- Multi-dimensional Indexes


Extensible Hashing: Two ideas

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

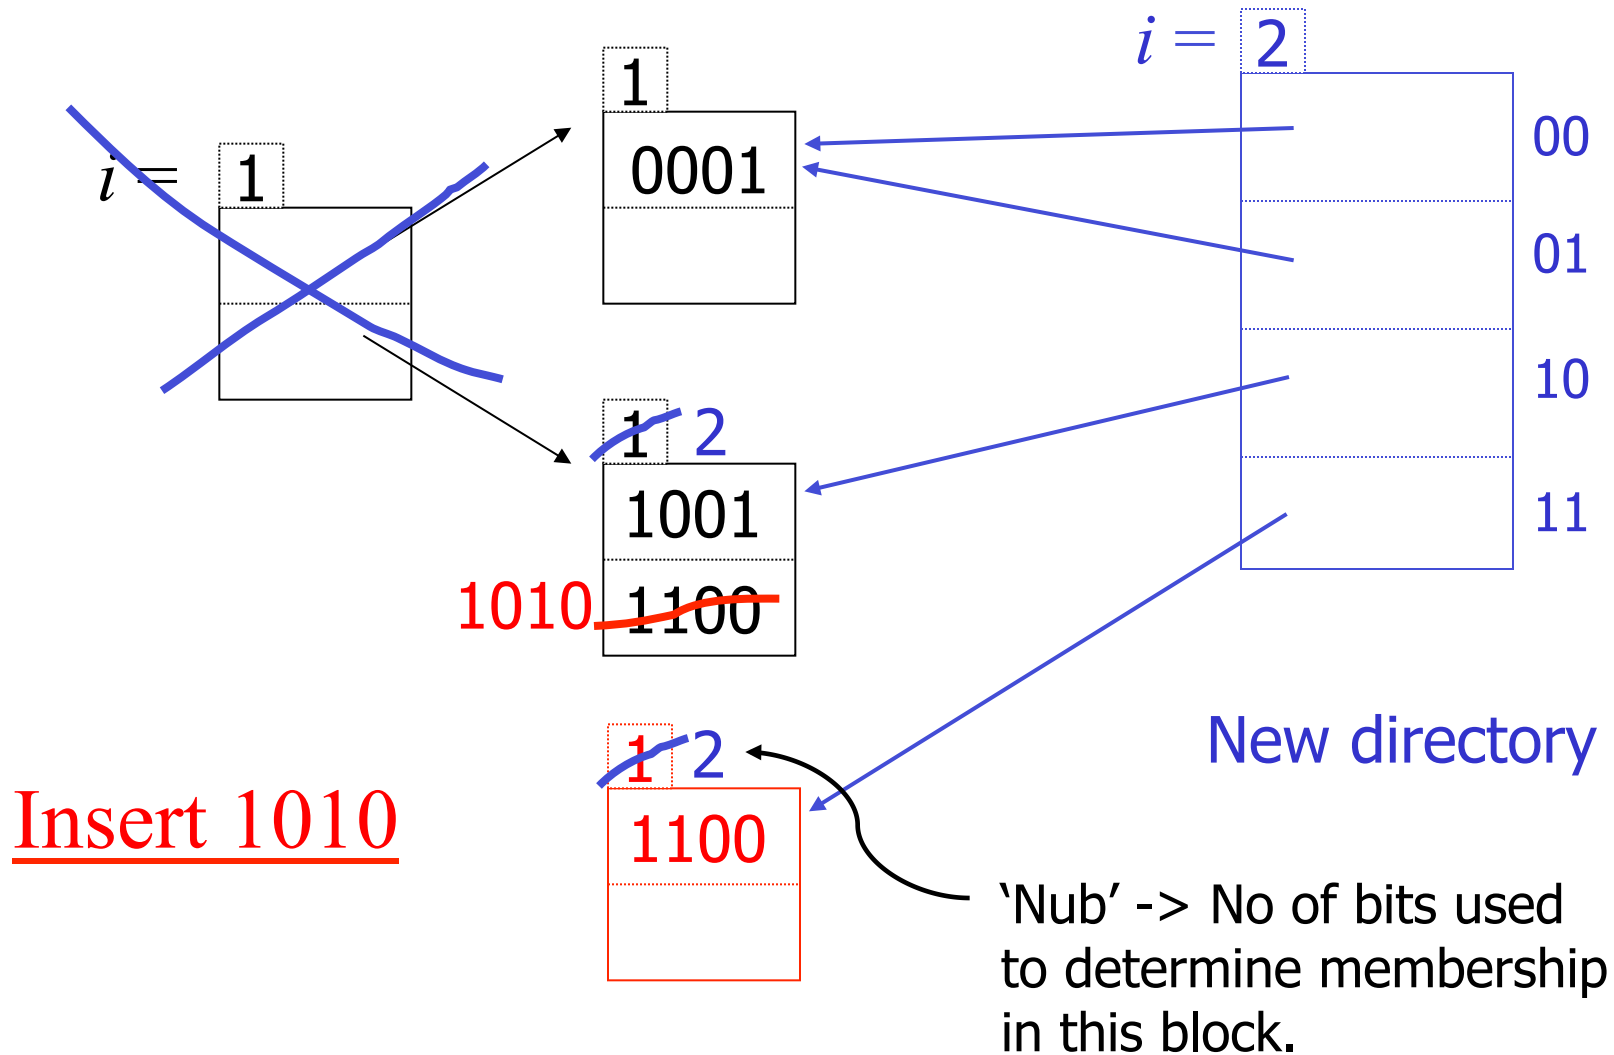


use $i \rightarrow$ grows over time....


(b) Use directory



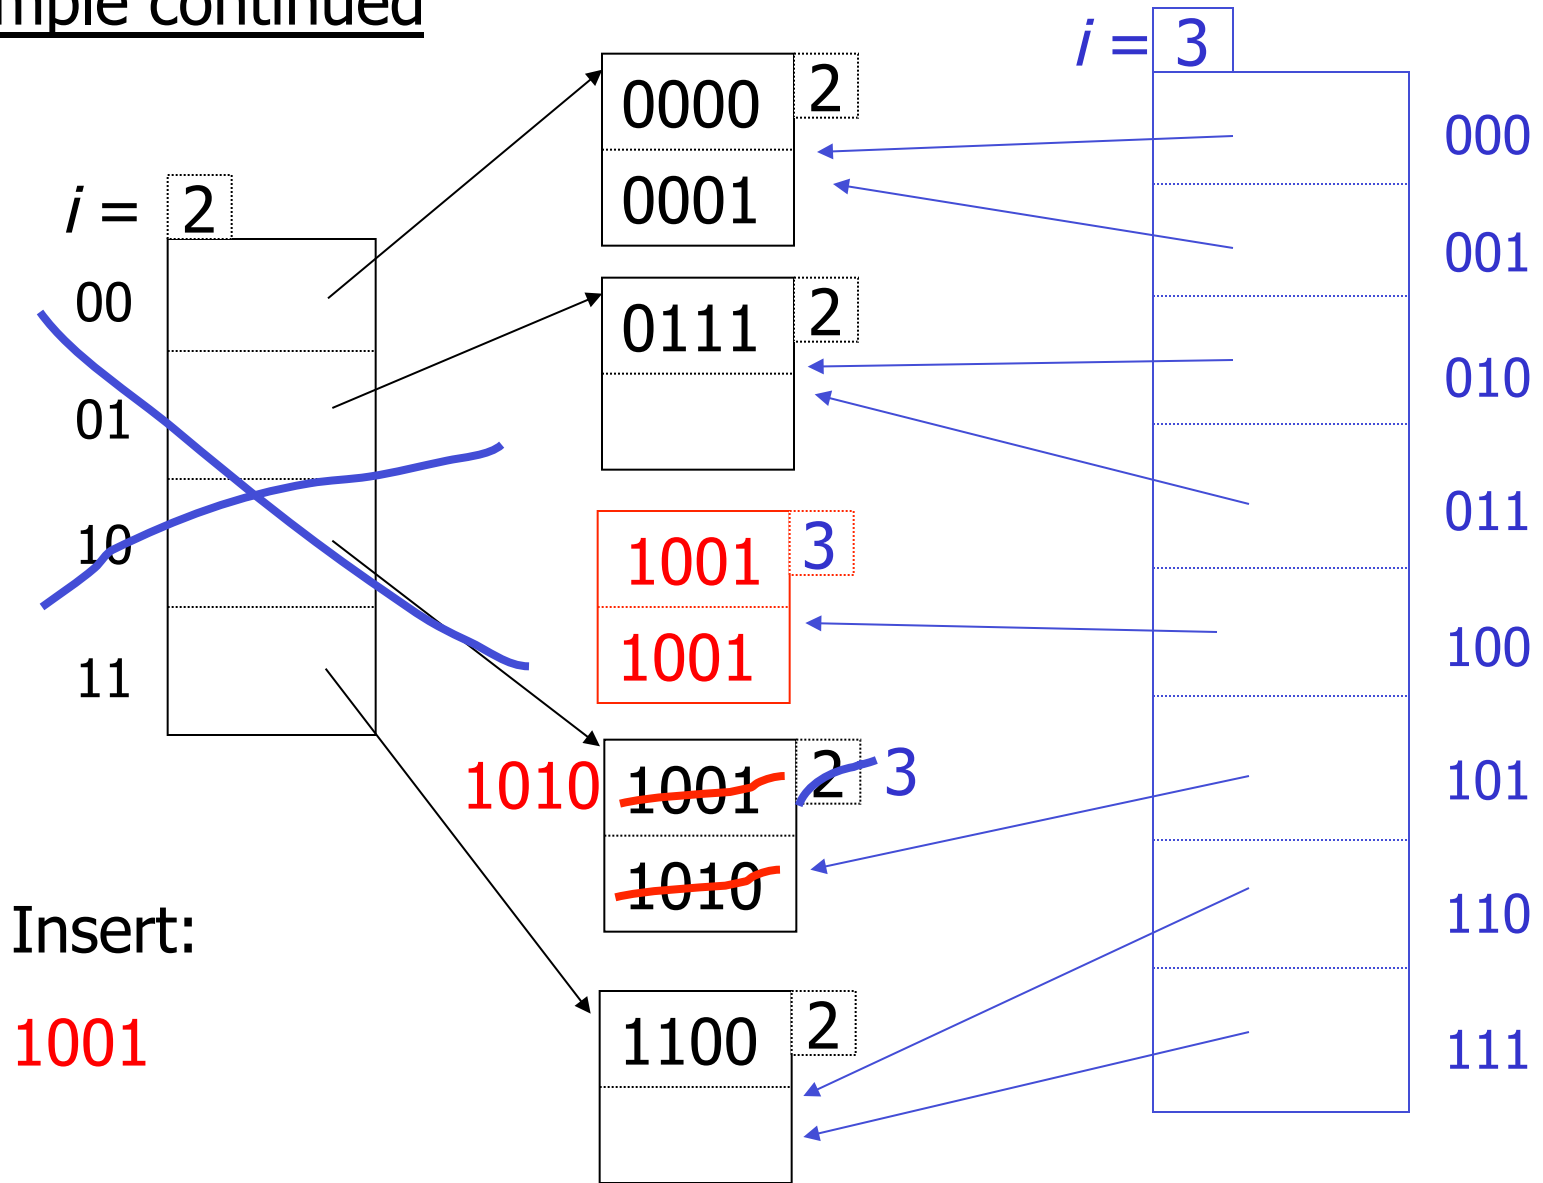
Example: $h(k)$ is 4 bits; 2 keys/bucket



Example continued



Example continued



Extensible hashing: deletion

- Reverse insert procedure

Extensible Hashing Summary

- If the directory fit in memory, then total I/Os is always ONE.

Negatives:

- Increase in i requires a lot of work.
- Directory doubles in size – so, sudden impact.
- Bad case: Assume blocks store 2 records. If there are 3 keys are “almost same”, then i will end up being very large.

Outline


- Dense/Sparse Indexes
- B Trees
- Hashing
 - Overall Idea
 - Handling growth: (i) Extensible Hashing, (ii) [Linear Hashing](#)
- Multi-dimensional Indexes

Linear hashing

- Another dynamic hashing scheme

Two ideas:

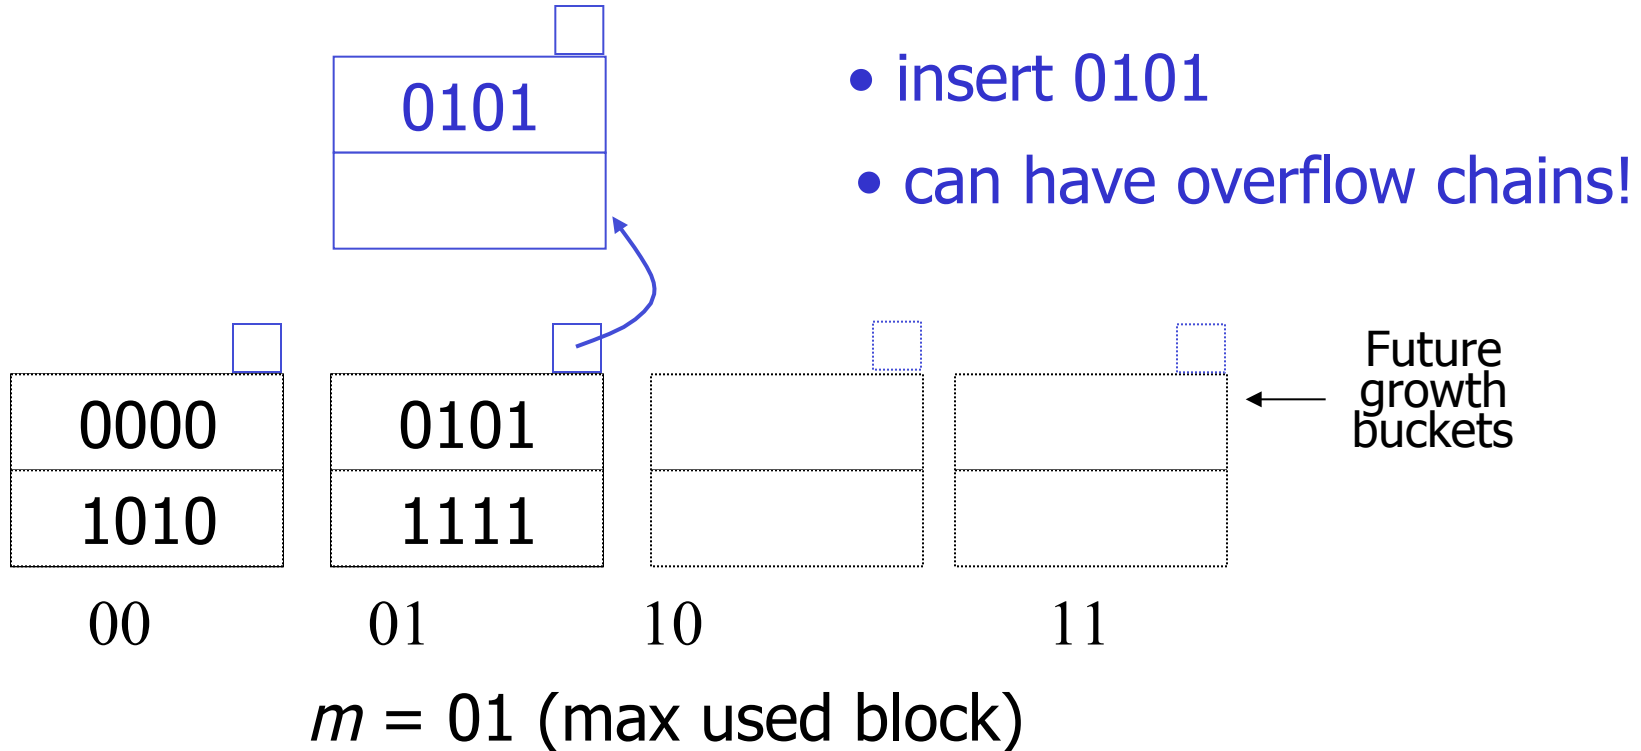
(a) Use i low order bits of hash



(b) File grows linearly



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



Rule

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

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

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


When do we expand file?

Space Utilization =

$$\frac{\text{\# keys in the system}}{\text{total \# keys that fit (without overflows)}}$$


- If $>$ threshold then increase m
(and maybe i)

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




Example Continued: How to grow beyond this?

$i =$ ~~2~~ 3



BAD CASE for Linear Hashing



Indexing vs Hashing

- Hashing good for probes given key
e.g.,
SELECT ...
FROM R
WHERE R.A = 5
- INDEXING (Including B Trees) good for
Range Searches:
e.g.,
SELECT
FROM R
WHERE R.A > 5

Outline

- Dense/Sparse Indexes
- B Trees
- Hashing (Extensible, Linear)
- Multi-dimensional Indexes

Multi-dimensional Indexes

Motivation

Find records where

DEPT = “Toy” AND SAL > 50k

General Multi-dimensional Queries

Assume tuples with $\langle w, x, y, z \rangle$

1. Partial-Match Queries:

- Find tuples where “ $x=5$ and $y=6$ ”

2. Range Queries

- Find tuples where “ $5 < x < 10$ and $y > 5$ ”

3. Nearest Neighbor Queries

- Find tuples “closest” to $\langle 4, 3, 5, 7 \rangle$.

4. Where-am-I Queries (over object shapes)

- Find objects containing $\langle 2, 3, 6, 7 \rangle$

Multi-dimensional Indexes


Motivation

Find records where

DEPT = “Toy” AND SAL > 50k


Strategy I: Use one-dimensional index

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




Strategy II: Manipulate Pointers

- Use 2 Indexes; manipulate pointers



Strategy III: Multi-dimensional Indexes

One idea:



Example

Art	
Sales	
Toy	

Dept
Index

10k	
15k	
17k	
21k	

12k	
15k	
15k	
19k	

Salary
Index

Example
Record

Name=Joe
DEPT=Sales
SAL=15k


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


Types of multidimensional indexes

- Grid Files
- Partitioned Hash
- Tree like structures
 - Multi-key indexes (previous slide)
 - kd-tree
 - Quad tress
 - R trees

Grid Index

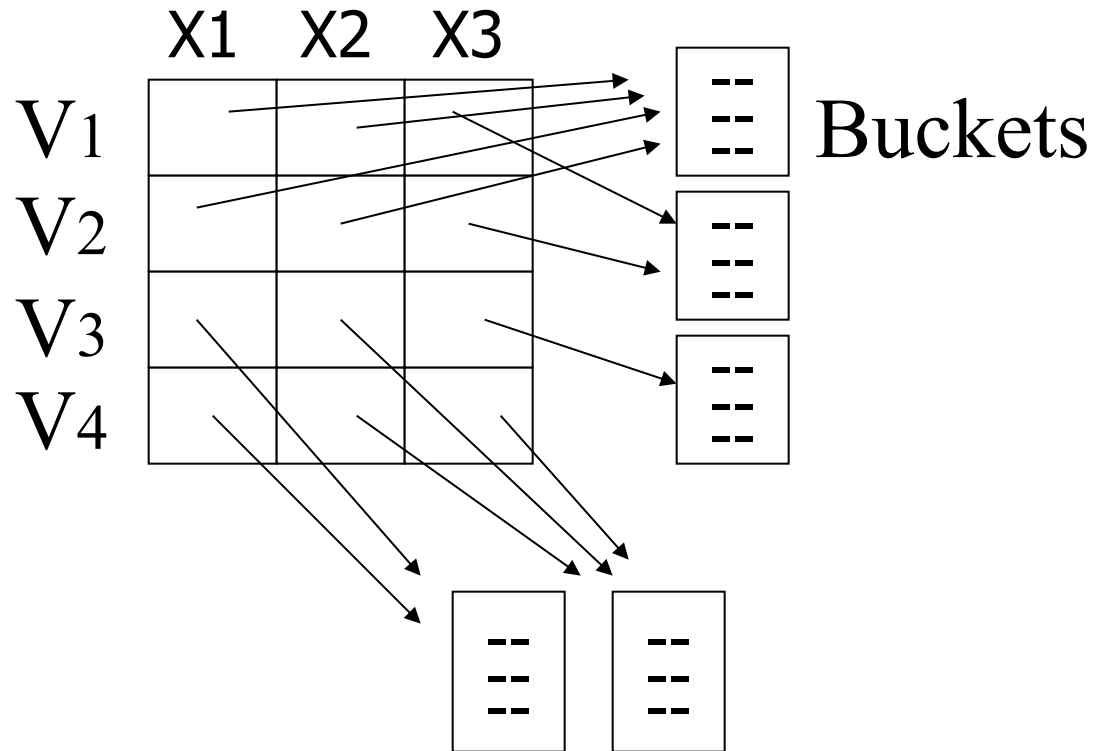


Storing Grid Index



- Each $\langle V_i, X_j \rangle$ entry should be of uniform size (array elements must be of same size).
➔ Store a single pointer to a bucket (of tuples/pointers).

I.e., use indirection




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 $< X_j$

Grid Files with Value-Ranges

- Visualize a grid file as “stripes” in each dimension.



Grid Files with Value-Ranges



Salary

0-40K	1
40K-100K	2
100K-500k	3

Grid

Age

1	2	3
0-30	30-50	50-100



Grid files


- Need to carefully split the ranges.
- May need extra indexes to determine row and/or column numbers from a key value
 - These indexes may be small enough to reside in main memory.
 - Most cases, these indexes could be replaced by “binary search”.
- Good for partial-match and range queries.

Types of multidimensional indexes

- Grid Files
- Partitioned Hash
- Tree like structures
 - Multi-key indexes (previous slide)
 - kd-tree
 - Quad tress
 - R trees


Partitioned Hash Function

Idea:



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	
:			



<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


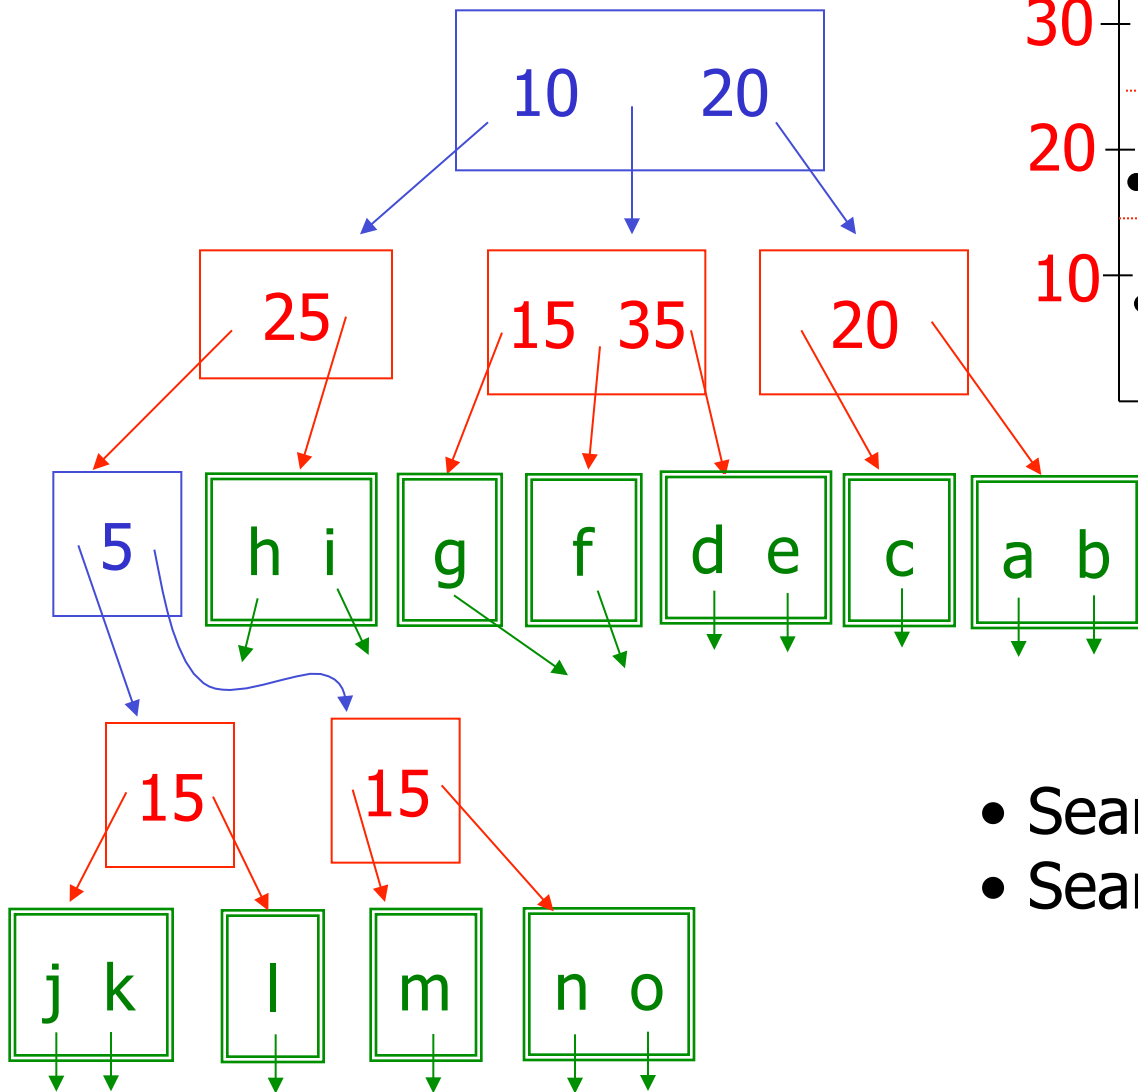
Grid Files vs. Partitioned Hash

- Grid Files are good for range-queries, partial-matches. Can be used for nearest-neighbor queries (by considering bigger and bigger “ranges”)
- Partitioned Hash are good only for partial-matches.

Types of multidimensional indexes

- Grid Files
- Partitioned Hash
- Tree like structures
 - Multi-key indexes (previous slide)
 - kd-tree
 - Quad tress
 - R trees

kd-Tree



- Search points near f
- Search points near b

Trees-like structures


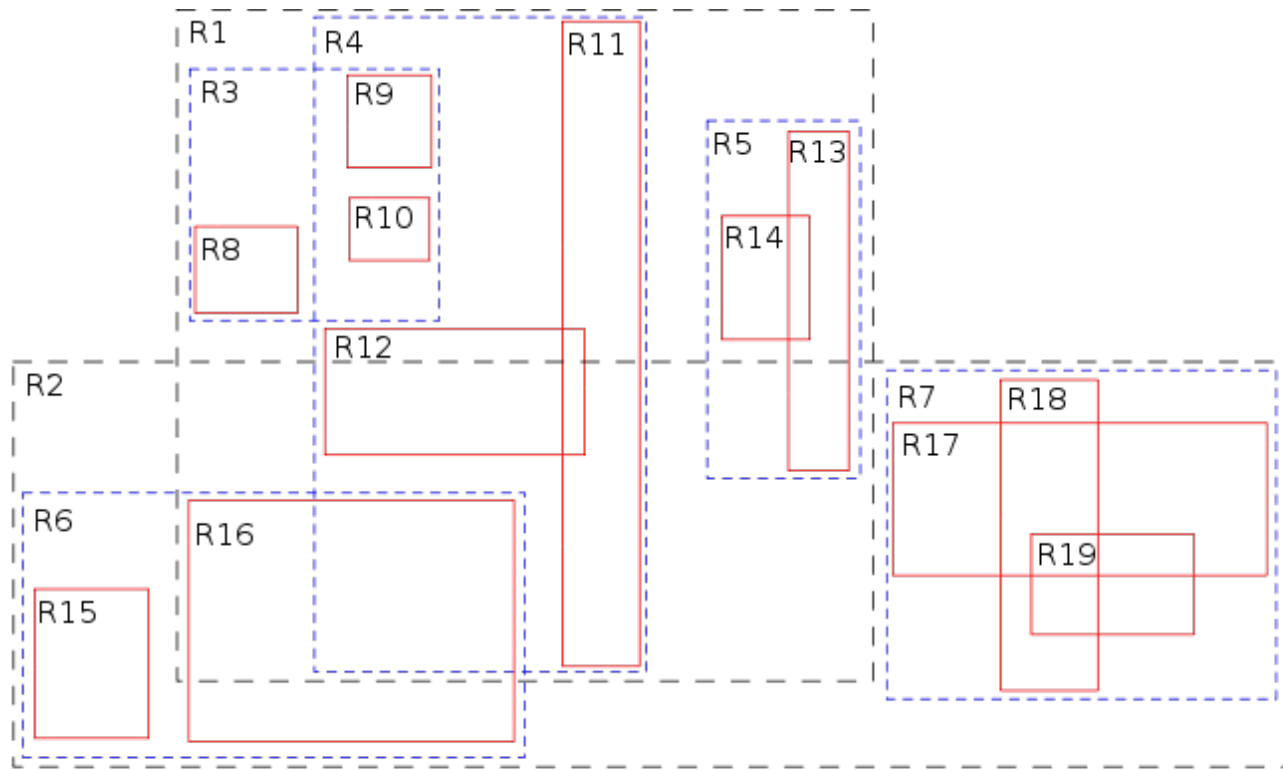
- kd-Trees
 - Interior node (A, V) has two children
 - $(A\text{-value} < V)$ and $(A\text{-value} > V)$
- Quad Trees
 - Each node corresponds to a rectangle
 - Partitions its region into sub-regions using a point (thus, each level divides across **all** dimensions)
- R-tree (NEXT ..) *used to store regions.*
 - Similar to Quad tree, except partitions may not be disjoint; may not cover the entire region.

R-trees

- Suppose we want to store “regions” instead of points.
- E.g., to store objects of geographic maps, we must store points, roads, buildings, spatial-regions, etc. A typical query could be:
 - Find all universities within 100km of the given point?
 - Find all roads that intersect with a given region?

R-trees

- R-trees are similar to kd-trees, wherein each node represents rectangle but:
 - The “subrectangles” corresponding to the children may not be disjoint, and together they may not completely cover the parent’s rectangle.
- R-tree can also be looked upon as built “bottom-up” – where given objects/regions are grouped into bigger rectangles and so on.



R-trees

- **Insertion:** Given a rectangle to insert.
 - Start from the root, and pick **one of the children** that contains the given rectangle. Continue till you reach the leaf. Store.
- **Where-am-I Query:** Given a rectangle/point, find intersection/containing regions:
 - Start from the root, and traverse **all** of the children that intersect/contain the given rectangle. Continue till you reach the leaf.

Outline

- Dense/Sparse Indexes
- B Trees
- Hashing (Extensible, Linear)
- Multi-dimensional Indexes (Grid, Partitioned Hash, kd-tree, quad-trees, R-trees, **BitMaps**)

BitMap Indexes

- Store **binary vectors for each attribute value**.
The vector represents the record numbers that contain that value. E.g., a vector 010100001 for “dog” signifies that the 2nd, 4th, and 9th records contain the value “dog”.
- Need to store such vectors for each value and each attribute.
- Records must have permanent numbers.

BitMaps: Challenges/Issues

- Good for range, partial-match (bucket intersections) queries.
- Need to create other indexes (secondary or binary) for overall operation: to find a bit-vector for a value, and then to find the i -th record of the table.
- Vectors can be compressed for efficiency. Note that a vector is expected to have few 1s.
- **Record Insertions:** New record number.
- **Record Deletions:** “Retire” the number, thus, not effecting the numbering.