

# Database Management Systems

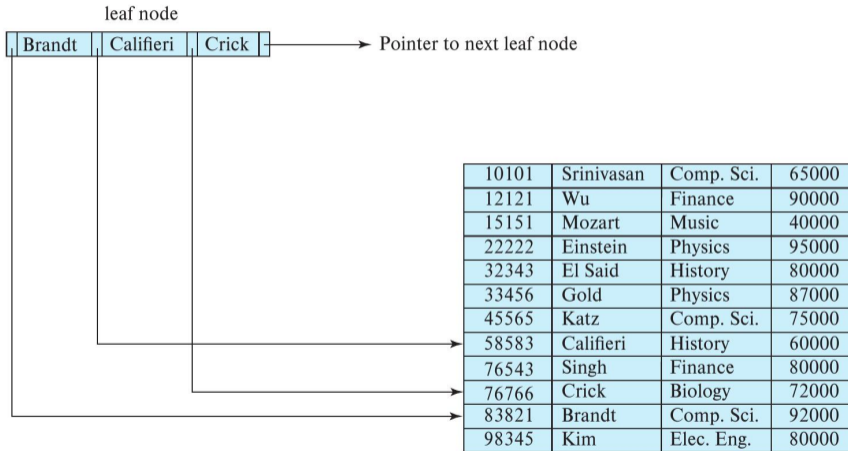
Madhavan Mukund

<https://www.cmi.ac.in/~madhavan>

Sai University

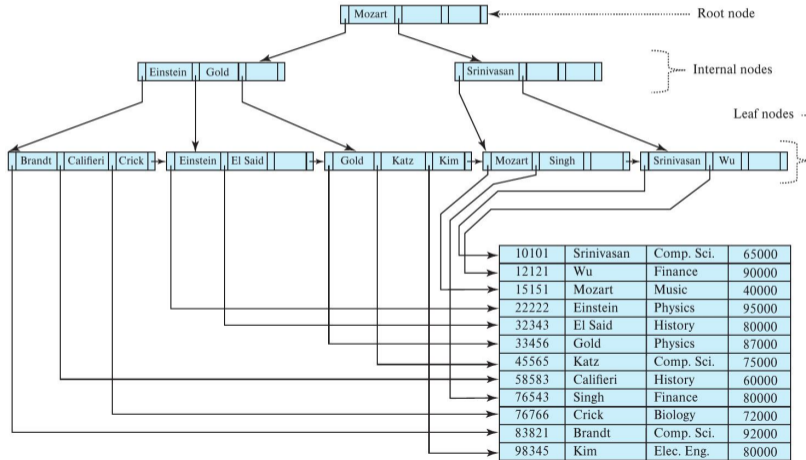
Lecture 16, 20 October 2023

- Leaf nodes form a dense index — linked list of leaves, each one block



# B+ trees

- Leaf nodes form a dense index — linked list of leaves
- Non-Leaf nodes form a sparse index



- Leaf nodes form a dense index — linked list of leaves
- Non-leaf nodes form a sparse index
- Constraints — assume  $n$  keys and pointers can fit in a block
  - Each leaf has at least  $\lceil (n-1)/2 \rceil$  key values
  - Each non-leaf has at least  $\lceil n/2 \rceil$  pointers
  - Height of the tree is proportional to  $\log_{n/2}(N)$

$n$  ptrs  
 $n-1$  keys

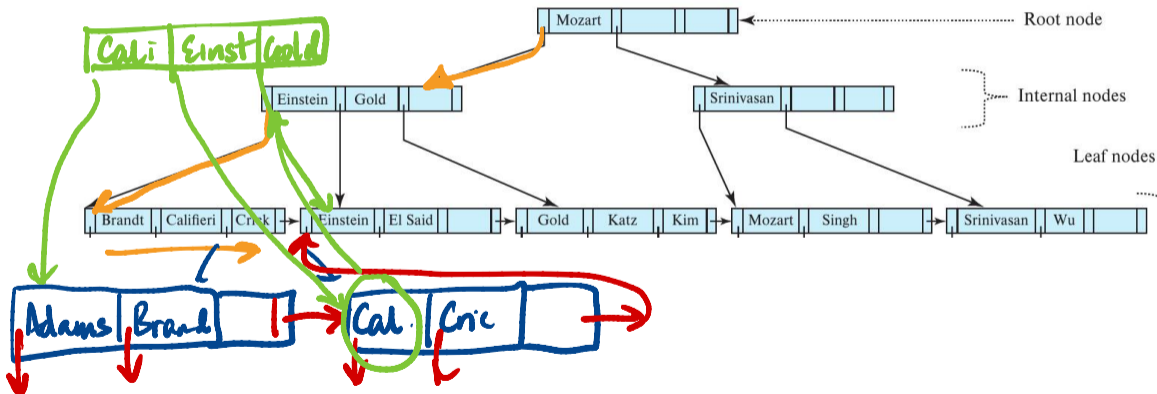
$P_1, K_1, \dots, P_{n-1}, K_{n-1}, P_n$

# B+ trees — insertion

- Insert Adams


$$\geq \left\lceil \frac{n-1}{2} \right\rceil$$

$$\frac{(n-1)+1}{2} = \frac{n}{2}$$



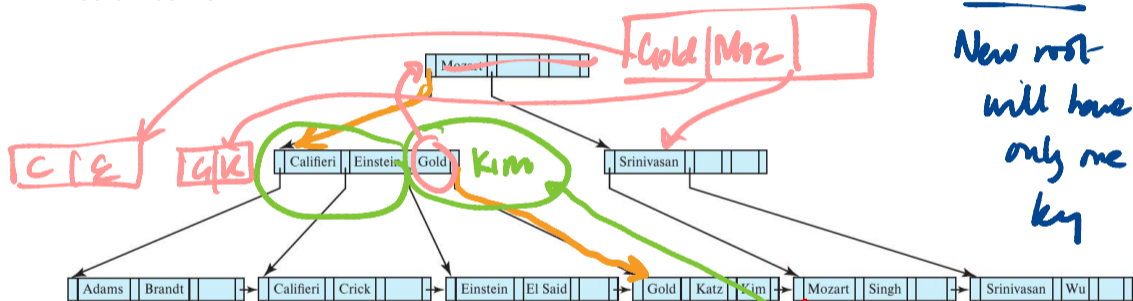
# B+ trees — insertion

## ■ Insert Adams



# B+ trees — insertion

■ Insert Adams



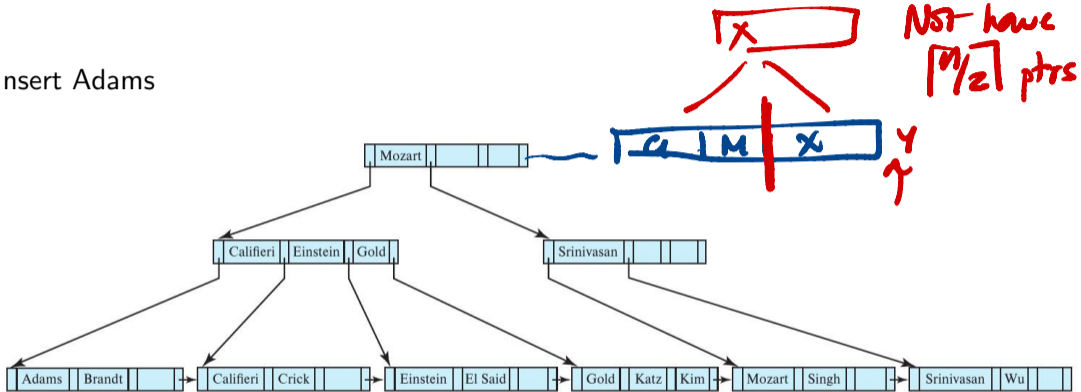
Split root  
New root will have only one key

■ Insert Lamport



# B+ trees — insertion

## ■ Insert Adams




## ■ Insert Lampport



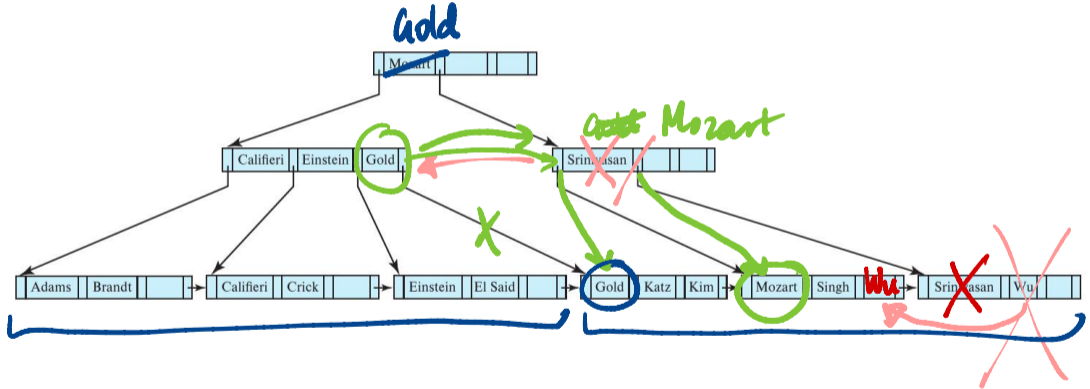
# B+ trees — insertion

- Insert Adams
- Insert Lamport

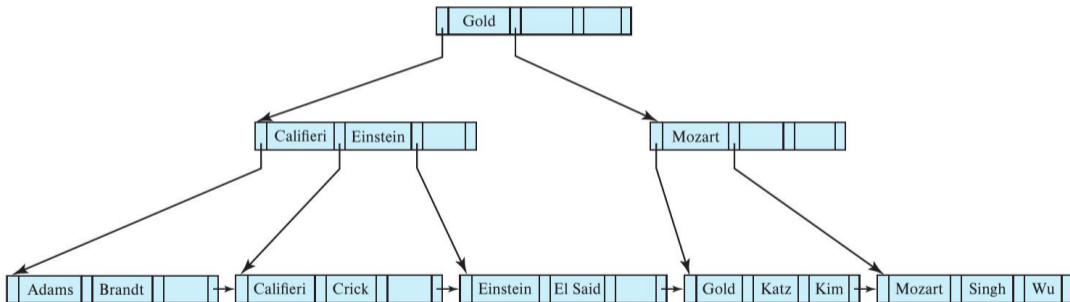


# B+ trees — deletion

## ■ Delete Srinivasan




- Delete Srinivasan



# B+ trees — deletion

## ■ Delete Srinivasan



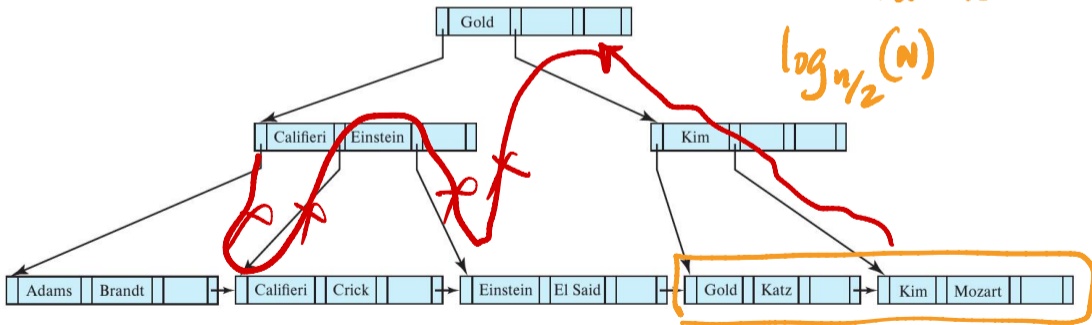
## ■ Delete Singh and Wu

# B+ trees — deletion

- Delete Srinivasan
- Delete Singh and Wu

All ops preserve  
"half full"  
condition

$$\log_{n/2}(N)$$



# Query processing

- Translate the query from SQL into relational algebra
- Evaluate the relational algebra expression

Like a compiler

Query  $\rightarrow$  Machine code

Query  $\rightarrow$  Algorithm

# Query processing

- Translate the query from SQL into relational algebra
- Evaluate the relational algebra expression
- Challenges

# Query processing


- Translate the query from SQL into relational algebra
- Evaluate the relational algebra expression
- Challenges
  - Many equivalent relational algebra expressions

$\sigma_{salary < 75000}(\pi_{salary}(instructor))$  vs  $\pi_{salary}(\sigma_{salary < 75000}(instructor))$



# Query processing

- Translate the query from SQL into relational algebra
- Evaluate the relational algebra expression
- Challenges
  - Many equivalent relational algebra expressions  
 $\sigma_{salary < 75000}(\pi_{salary}(instructor))$  vs  $\pi_{salary}(\sigma_{salary < 75000}(instructor))$
  - Many ways to evaluate a given expression



# Query processing

- Translate the query from SQL into relational algebra
- Evaluate the relational algebra expression

- Challenges

- Many equivalent relational algebra expressions

$\sigma_{salary < 75000}(\pi_{salary}(instructor))$  vs  $\pi_{salary}(\sigma_{salary < 75000}(instructor))$


- Many ways to evaluate a given expression

- Query plan

- Annotate the expression with a detailed evaluation strategy ~~key values~~

# Query processing

- Translate the query from SQL into relational algebra
- Evaluate the relational algebra expression
- Challenges
  - Many equivalent relational algebra expressions  
 $\sigma_{salary < 75000}(\pi_{salary}(instructor))$  vs  $\pi_{salary}(\sigma_{salary < 75000}(instructor))$
  - Many ways to evaluate a given expression
- Query plan
  - Annotate the expression with a detailed evaluation strategy ~~key values~~
    - Use index on *salary* to find instructors with *salary* < 75000



# Query processing

- Translate the query from SQL into relational algebra
- Evaluate the relational algebra expression
- Challenges
  - Many equivalent relational algebra expressions  
 $\sigma_{salary < 75000}(\pi_{salary}(instructor))$  vs  $\pi_{salary}(\sigma_{salary < 75000}(instructor))$
  - Many ways to evaluate a given expression
- Query plan
  - Annotate the expression with a detailed evaluation strategy key values
    - Use index on *salary* to find instructors with *salary* < 75000
    - Or, scan entire relation, discard rows with *salary*  $\geq$  75000

# Query optimization

- Choose plan with lowest cost

We know what the tables look like

# Query optimization

- Choose plan with lowest cost
- Maintain **database catalogue** — number of tuples in each relationn, size of tuples,  
...

# Query optimization

- Choose plan with lowest cost
- Maintain **database catalogue** — number of tuples in each relation, size of tuples, ...
- Assess cost in terms of disk access and transfer, CPU time, ...

$10^{-3}$

$10^{-9}$

SSDs  
 $10^{-6}$

———— ignore

# Query optimization

- Choose plan with lowest cost
- Maintain **database catalogue** — number of tuples in each relation, size of tuples, ...
- Assess cost in terms of disk access and transfer, CPU time, ...
- For simplicity, ignore in-memory costs (CPU time), restrict to disk access



# Query optimization

- Choose plan with lowest cost
- Maintain **database catalogue** — number of tuples in each relation, size of tuples, ...
- Assess cost in terms of disk access and transfer, CPU time, ...
- For simplicity, ignore in-memory costs (CPU time), restrict to disk access
- Disk accesses
  - Relation  $r$  occupies  $b_r$  blocks
  - **Disk seeks** — time  $t_S$  per seek
  - **Block transfers** — time  $t_T$  per transfer

Read vs write

# Query optimization

- Choose plan with lowest cost
- Maintain **database catalogue** — number of tuples in each relation, size of tuples, ...
- Assess cost in terms of disk access and transfer, CPU time, ...
- For simplicity, ignore in-memory costs (CPU time), restrict to disk access
- Disk accesses
  - Relation  $r$  occupies  $b_r$  blocks
  - **Disk seeks** — time  $t_S$  per seek
  - **Block transfers** — time  $t_T$  per transfer
- Other factors — buffer management etc

*Which blocks to keep in memory?*

(A1) Linear search

Read  $b_r$  blocks

"Seek" only first block

$$b_r \cdot b_T + 1 \cdot b_S$$

$$T_{\theta}(r)$$

# Selection

(A1) Linear search

(A2) Clustering index, equality on key — index height  $h_i$

single value

# of index queries to reach specific record

$h_i$  index lookups


$t_s + t_r$

Actual record

$t_s + t_r$

$$(h_i + 1) \cdot (t_s + t_r)$$

$$\sigma_{\theta}(r)$$



# Selection

- (A1) Linear search
- (A2) Clustering index, equality on key — index height  $h_i$
- (A3) Clustering index, equality on nonkey




table is sequential w.r.t key

$h_i$  index lookups  $t_s + t_T$

$b$  blocks with matching records  $+ t_s + b \cdot t_T$

$$h_i(t_s + t_T) + t_s + b \cdot t_T$$

# Selection

- (A1) Linear search
- (A2) Clustering index, equality on key — index height  $h_i$
- (A3) Clustering index, equality on nonkey
- (A4) Secondary index (key, non-key)



$h_i (t_s + t_r) + 1$  block lookup

$(h_i + 1) (t_s + t_r)$ , as before

Search key order  $\neq$   
table order

# Selection

- (A1) Linear search
- (A2) Clustering index, equality on key — index height  $h_i$
- (A3) Clustering index, equality on nonkey
- (A4) Secondary index (key, non-key)

$h_i(t_s + t_T) + n(t_s + t_T)$

↳  $n$  records match this search key  
each is in a different block


Expensive

# Selection

- (A1) Linear search
- (A2) Clustering index, equality on key — index height  $h_i$
- (A3) Clustering index, equality on nonkey
- (A4) Secondary index (key, non-key)
- (A5) Clustering index, comparison — sorted on  $A$

Salary < 70000

$$h_i (t_s + t_T) + (t_s + b t_T)$$




$x < \text{Salary} < y$

1 seek, b transfers



# Selection

- (A1) Linear search
- (A2) Clustering index, equality on key — index height  $h_i$
- (A3) Clustering index, equality on nonkey
- (A4) Secondary index (key, non-key)
- (A5) Clustering index, comparison — sorted on  $A$
- (A6) Clustering index, comparison — not sorted on  $A$



Linear Scan

$$\ln(t_{st} + t_T) \text{ index}$$
$$n \cdot (b_s + b_T)$$

# Complex selections

Conjunctions, disjunctions and negations

(A7) Conjunctive selection using one index

$$\sigma_{\theta_1 \wedge \theta_2 \wedge \dots \wedge \theta_k}(r)$$

↑  
Index

$$\sigma_{\theta_2}(r)$$

→ check  $\theta_1 \wedge \theta_3 \wedge \dots \wedge \theta_k$

$$\sigma_{\theta_1 \vee \theta_2}(r)$$

∨

¬θ

and

or

not

# Complex selections

Conjunctions, disjunctions and negations

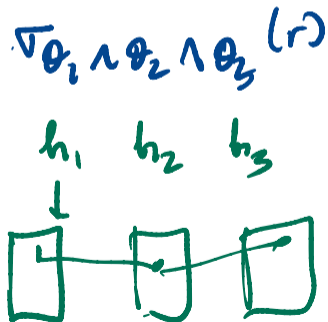
(A7) Conjunctive selection using one index

(A8) Conjunctive selection using composite index

# Complex selections

Conjunctions, disjunctions and negations

- (A7) Conjunctive selection using one index
- (A8) Conjunctive selection using composite index
- (A9) Conjunctive selection using intersection of pointers



# Complex selections

Conjunctions, disjunctions and negations

- (A7) Conjunctive selection using one index
- (A8) Conjunctive selection using composite index
- (A9) Conjunctive selection using intersection of pointers
- (A10) Disjunctive selection by union of pointers

Linear scan

$$\bigvee_{h_1} \bigvee_{h_2} \bigvee_{h_3} - (r)$$

# Complex selections

## Conjunctions, disjunctions and negations

- (A7) Conjunctive selection using one index
- (A8) Conjunctive selection using composite index
- (A9) Conjunctive selection using intersection of pointers
- (A10) Disjunctive selection by union of pointers
- (Neg) Negation

Linear scan



Join

$$\sigma_{\theta} (r_1 \times r_2)$$

In memory sort

vs External (disk based) sort