


# Indexing

Introduction to Databases  
CompSci 316 Fall 2016




---

---

---

---

---

---

---

---

## Announcements (Thu., Nov. 10)

- Project milestone #2 due today
- Homework #3 sample solution to be posted on Sakai by this weekend
- Homework #4 to be assigned next Tuesday

---

---

---

---

---

---

---

---

## What are indexes for?

- Given a value, locate the record(s) with this value
  - `SELECT * FROM R WHERE A = value;`
  - `SELECT * FROM R, S WHERE R.A = S.B;`
- Find data by other search criteria, e.g.
  - Range search
    - `SELECT * FROM R WHERE A > value;`
  - Keyword search

} Focus of this lecture

database indexing

Search

---

---

---

---

---

---

---

---

### Dense and sparse indexes

- **Dense:** one index entry for each search key value
  - One entry may "point" to multiple records (e.g., two users named Jessica)
- **Sparse:** one index entry for each block
  - Records must be **clustered** according to the search key

123	Milhouse	10	0.2
142	Bart	10	0.9
279	Jessica	10	0.9
345	Martin	8	2.3
456	Ralph	8	0.3
512	Nelson	10	0.4
679	Sherri	10	0.6
697	Terri	10	0.6
857	Lisa	8	0.7
912	Windel	8	0.5
997	Jessica	8	0.5

---

---

---

---

---

---

---

---

---

---

---

---

### Dense versus sparse indexes

- Index size
  - Sparse index is smaller
- Requirement on records
  - Records must be clustered for sparse index
- Lookup
  - Sparse index is smaller and may fit in memory
  - Dense index can directly tell if a record exists
- Update
  - Easier for sparse index

---

---

---

---

---

---

---

---

---

---

---

---

### Primary and secondary indexes

- **Primary index**
  - Created for the **primary key** of a table
  - Records are usually clustered by the primary key
  - Can be sparse
- **Secondary index**
  - Usually dense
- SQL
  - PRIMARY KEY declaration automatically creates a primary index, UNIQUE key automatically creates a secondary index
  - Additional secondary index can be created on non-key attribute(s):  
`CREATE INDEX UserPopIndex ON User (pop);`

---

---

---

---

---

---

---

---

---

---

---

---

### ISAM

- What if an index is still too big?
  - Put a another (sparse) index on top of that!
  - ISAM (Index Sequential Access Method), more or less

Example: look up 197

---

---

---

---

---

---

---

---

### Updates with ISAM

Example: insert 107  
Example: delete 129

---

---

---

---

---

---

---

---

### B<sup>+</sup>-tree

- A hierarchy of nodes with intervals
- Balanced (more or less): good performance guarantee
- Disk-based: one node per block; large fan-out

Max fan-out: 4

---

---

---

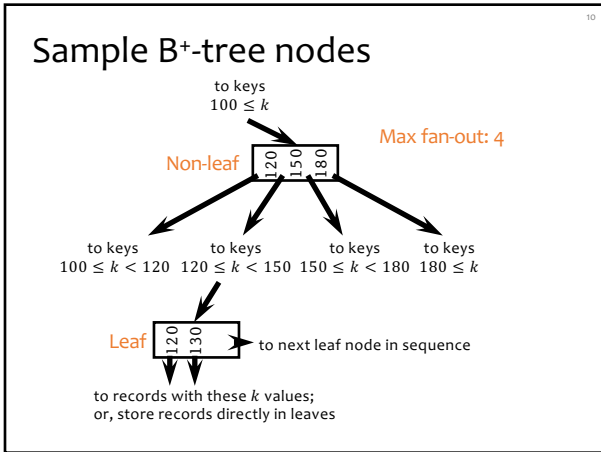
---

---

---

---

---




---

---

---

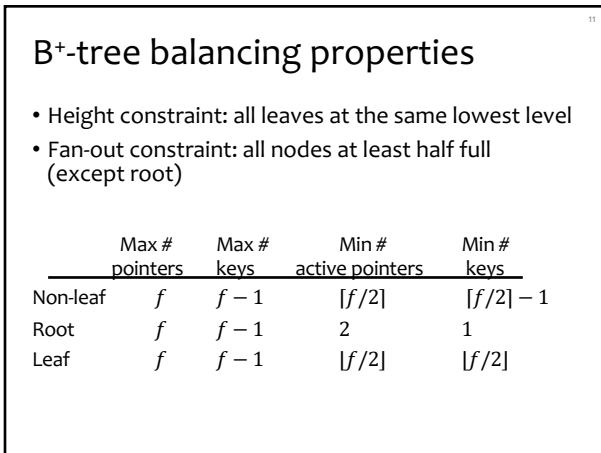
---

---

---

---

---




---

---

---

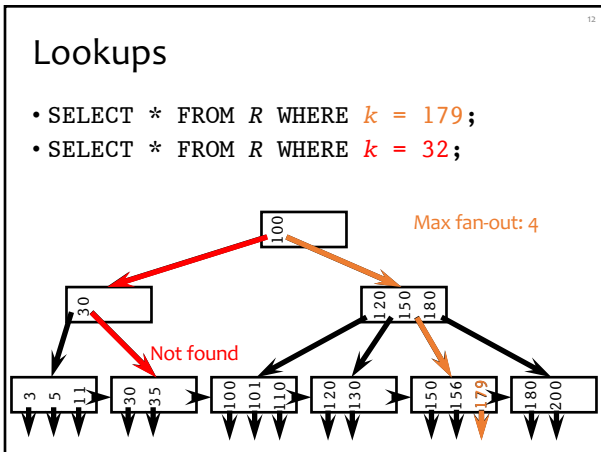
---

---

---

---

---




---

---

---

---

---

---

---

---

### Range query

- SELECT \* FROM R WHERE  $k > 32$  AND  $k < 179$ ;

Max fan-out: 4

Look up 32...

And follow next-leaf pointers until you hit upper bound

---

---

---

---

---

---

---

---

### Insertion

- Insert a record with search key value 32

Max fan-out: 4

Look up where the inserted key should go...

And insert it right there

---

---

---

---

---

---

---

---

### Another insertion example

- Insert a record with search key value 152

Max fan-out: 4

Oops, node is already full!

---

---

---

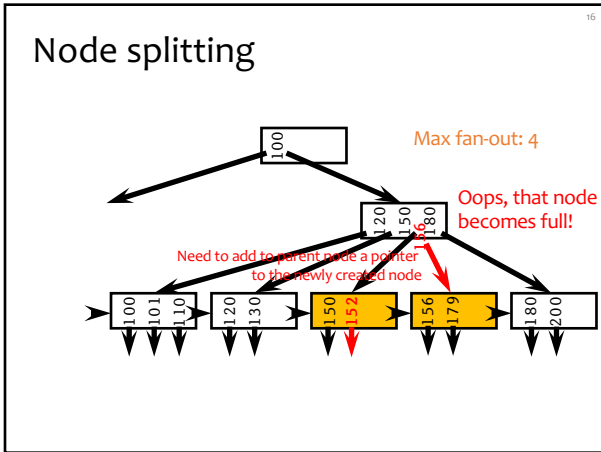
---

---

---

---

---




---

---

---

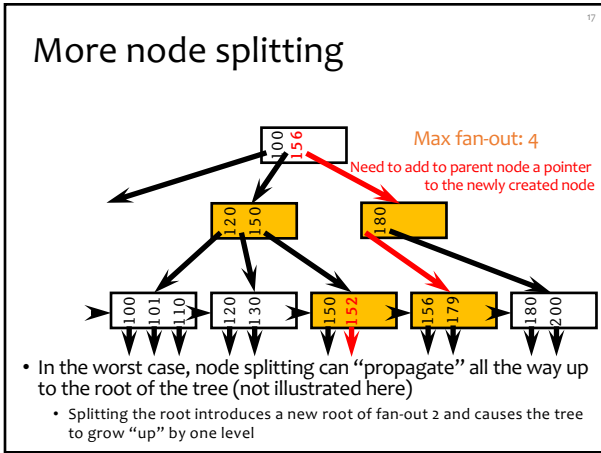
---

---

---

---

---




---

---

---

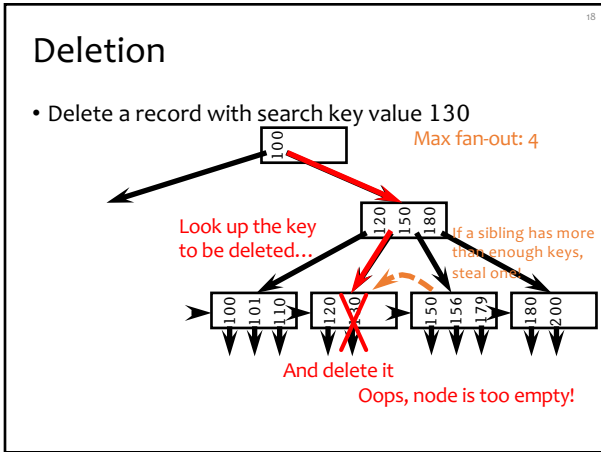
---

---

---

---

---




---

---

---

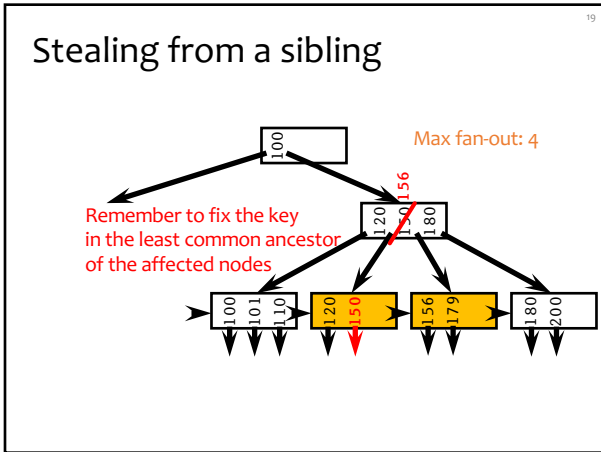
---

---

---

---

---




---

---

---

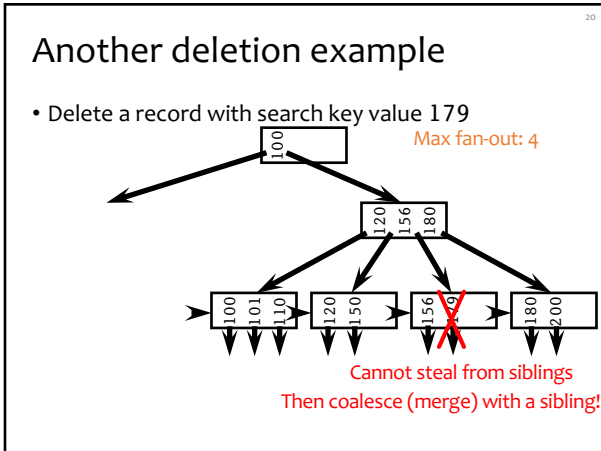
---

---

---

---

---




---

---

---

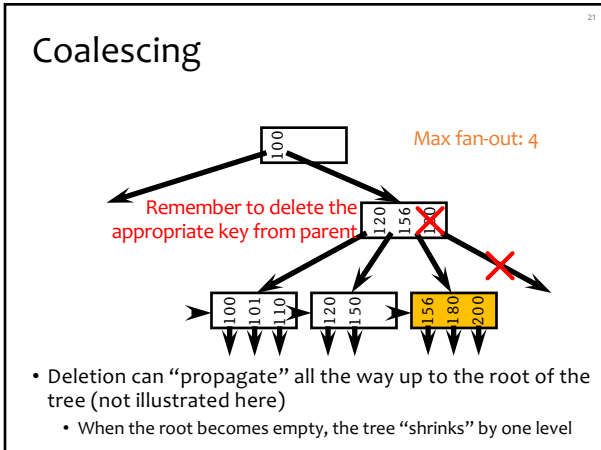
---

---

---

---

---




---

---

---

---

---

---

---

---

## Performance analysis

- How many I/O's are required for each operation?
  - $h$ , the height of the tree (more or less)
  - Plus one or two to manipulate actual records
  - Plus  $O(h)$  for reorganization (rare if  $f$  is large)
  - Minus one if we cache the root in memory
- How big is  $h$ ?
  - Roughly  $\log_{\text{fanout}} N$ , where  $N$  is the number of records
  - B<sup>+</sup>-tree properties guarantee that fan-out is least  $f/2$  for all non-root nodes
  - Fan-out is typically large (in hundreds)—many keys and pointers can fit into one block
  - A 4-level B<sup>+</sup>-tree is enough for “typical” tables

---

---

---

---

---

---

---

---

## B<sup>+</sup>-tree in practice

- Complex reorganization for deletion often is not implemented (e.g., Oracle)
  - Leave nodes less than half full and periodically reorganize
- Most commercial DBMS use B<sup>+</sup>-tree instead of hashing-based indexes because B<sup>+</sup>-tree handles range queries

---

---

---

---

---

---

---

---

## The Halloween Problem

- Story from the early days of System R...

```
UPDATE Payroll
SET salary = salary * 1.1
WHERE salary >= 100000;
```

- There is a B<sup>+</sup>-tree index on Payroll(salary)
- The update never stopped (why?)
- Solutions?

---

---

---

---

---

---

---

---



25

### B<sup>+</sup>-tree versus ISAM

- ISAM is more **static**; B<sup>+</sup>-tree is more **dynamic**
- ISAM can be more compact (at least initially)
  - Fewer levels and I/O's than B<sup>+</sup>-tree
- Overtime, ISAM may not be balanced
  - Cannot provide guaranteed performance as B<sup>+</sup>-tree does

---

---

---

---

---

---

---

---

26

### B<sup>+</sup>-tree versus B-tree

- B-tree: why not store records (or record pointers) in non-leaf nodes?
  - These records can be accessed with fewer I/O's
- Problems?

---

---

---

---

---

---


---

---

27

### Beyond ISAM, B-, and B<sup>+</sup>-trees

- Other tree-based indexes: R-trees and variants, GiST, etc.
  - How about binary tree?



- Hashing-based indexes: extensible hashing, linear hashing, etc.
- Text indexes: inverted-list index, suffix arrays, etc.
- Other tricks: bitmap index, bit-sliced index, etc.

---

---

---

---

---

---

---

---