

Introduction

- ♦ As for any index, 3 alternatives for data entries k*:
 ③ Data record with key value k
 - ③ <k, rid of data record with search key value k>
 - \odot <k, list of rids of data records with search key k>
- Choice is orthogonal to the *indexing technique* used to locate data entries k*.
- Tree-structured indexing techniques support both *range searches* and *equality searches*.
- ISAM: static structure; <u>B+ tree</u>: dynamic, adjusts gracefully under inserts and deletes.

Database Management Systems, R. Ramakrishnan and J. Gehrke

























B+-tree Search Performance

- * Assume leaf pages can hold L data entries
- $\boldsymbol{\diamond}$ Assume B+-tree has order \boldsymbol{d}

atabase Management Systems, R. Ramakrishnan and J. Gehrke

- $\boldsymbol{\ast}$ Assume the tree has to index N data entries
- What is the best-case search performance (measured in number of I/Os)?
- * What is the worst-case search performance

A Note on `Order'

atabase Management Systems, R. Ramakrishnan and J. Gehrke

Database Management Systems, R. Ramakrishnan and J. Gehrk

- Order (d) concept replaced by physical space criterion in practice (`at least half-full').
 - Index pages can typically hold many more entries than leaf pages.
 - Variable sized records and search keys mean different nodes will contain different numbers of entries.
 - Even with fixed length fields, multiple records with the same search key value (*duplicates*) can lead to variable-sized data entries (if we use Alternative (3)).

Bulk Loading of a B+ Tree
 If we have a large collection of records, and we want to create a B+ tree on some field, doing so by repeatedly inserting records is very slow.
 <u>Bulk Loading</u> can be done much more efficiently.
 <u>Bulk Loading</u> can be done much more efficiently.
 Initialization: Sort all data entries, insert pointer to first (leaf) page in a new (root) page.

 3*
 4*
 6*
 9*
 10*
 11*
 12*
 13*
 20*
 23*
 31*
 35*
 36*
 38*
 41*
 44*

Summary of Bulk Loading

* Option 1: multiple inserts.

- Slow.

- Does not give sequential storage of leaves.
- * Option 2: <u>Bulk Loading</u>
 - Has advantages for concurrency control.
 - Fewer I/Os during build.
 - Leaves will be stored sequentially (and linked, of course).
 - Can control "fill factor" on pages.

Database Management Systems, R. Ramakrishnan and J. Gehrke