



ISAM: Indexed-Sequential- Access-Method

Adapted from Prof Joe Hellerstein's notes
<http://db.cs.berkeley.edu/dbcourse/notes.html>





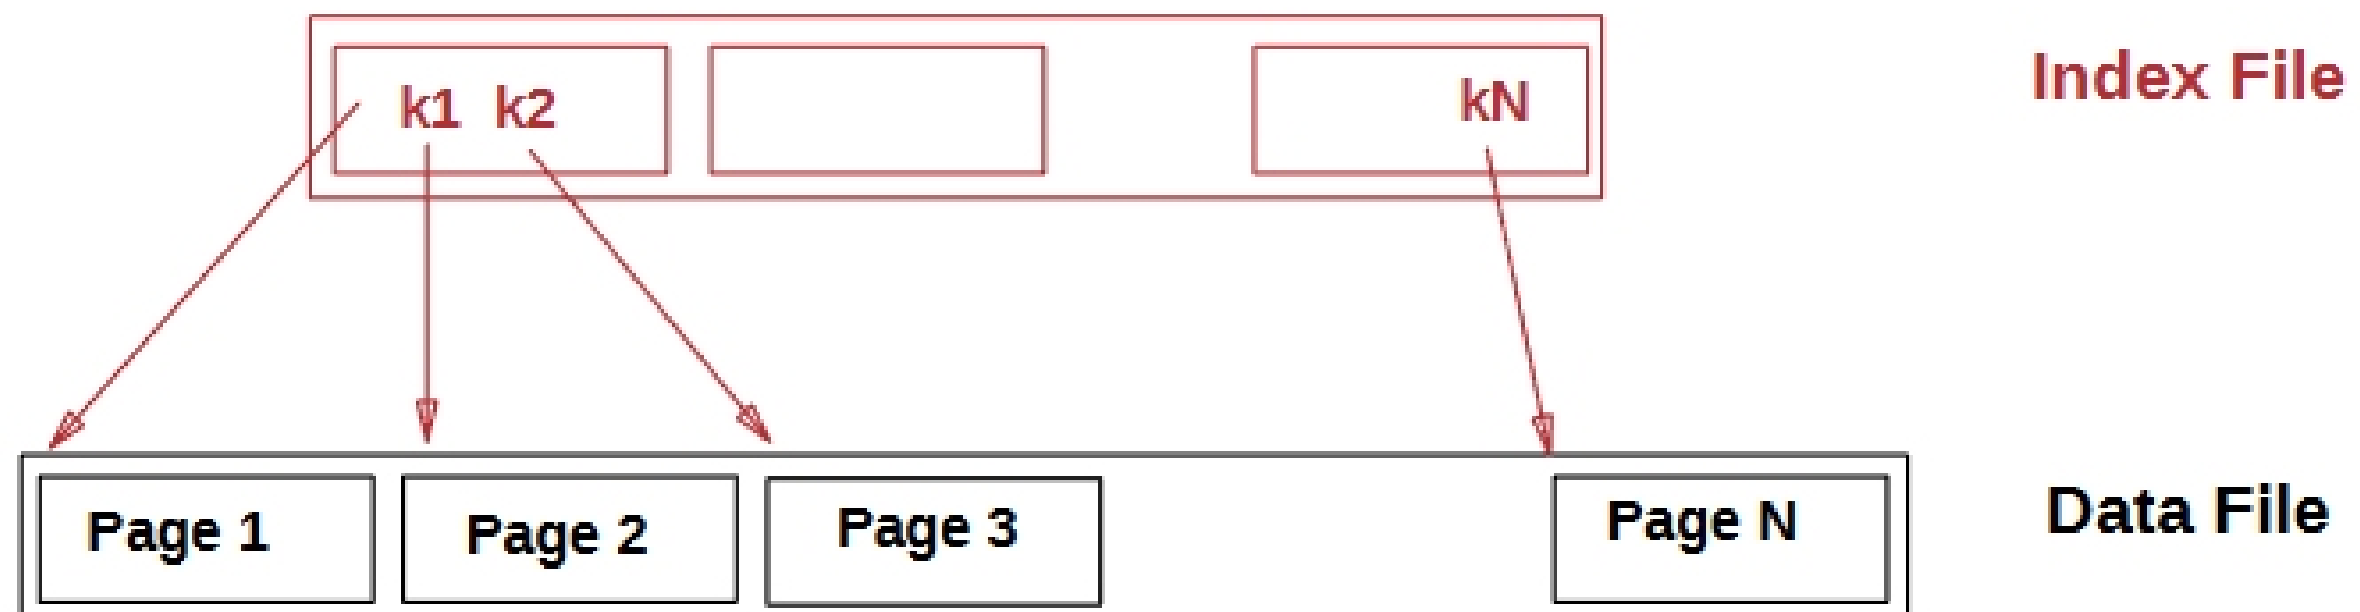
A Note of Caution

- **ISAM is an old-fashioned idea**
 - B+-trees are usually better, as we'll see
 - Though not *always*
- **But, it's a good place to start**
 - Simpler than B+-tree, but many of the same ideas
- **Upshot**
 - **Don't** brag about being an ISAM expert on your resume
 - **Do** understand how they work, and tradeoffs with B+-trees



Range Searches

- **“Find all students with $gpa > 3.0$ ”**
 - If data is in sorted file, do binary search to find first such student, then scan to find others.
 - Cost of binary search can be quite high.
- **Simple idea: Create an ‘index’ file.**
 - Level of indirection again!



- *Can do binary search on (smaller) index file!*