Quick Review of Apr 17 material

- Multiple-Key Access
  - There are good and bad ways to run queries on multiple single keys
- Indices on Multiple Attributes
  - Combining two keys into a single concatenated attribute
  - Grid files
    - crude array with one dimension and linear scale for each attribute
    - more than one cell may point to a given bucket of values
    - array grid may be dynamically resized during use
  - Other alternatives are spatial databases: R-tree, quad-trees, k-d tree
- Bitmap Indices
  - linear array of bits: bit j is set if tuple j has the attribute that this bitmap tracks (e.g., “Moonroof”: bit j is 1 if record j is a car with moonroof)
  - Queries are answered by combining several bitmaps using **and, or, not**

Today

- HW #4: due Thursday April 24 (next class)
  - Questions: 12.11, 12.12, 12.13, 12.16
- No HW for next week
- Today:
  - Start Chapter 13: Query Processing

Query Processing

- SQL is good for humans, but not as an internal (machine) representation of how to calculate a result
- Processing an SQL (or other) query requires these steps:
  - parsing and translation
    - turning the query into a useful internal representation in the extended relational algebra
  - optimization
    - manipulating the relational algebra query into the most efficient form (one that gets results the fastest)
  - evaluation
    - actually computing the results of the query