

Relational Algebra

CSE462 Database Concepts

Demian Lessa

Department of Computer Science and Engineering
State University of New York, Buffalo

January 26–28, 2011

Relational Algebra

Relational Algebra (RA) is an algebra of relations that provides simple yet powerful ways to construct new relations from existing ones. It is related both to **first-order logic** and **set algebra**.

- Modern systems do not use RA as their query language.
- Instead, they use a concrete language such as SQL.
- It is important to note, however, that RA is at the core of SQL.
- DBMSs translate queries into RA (or variant) during query processing.

Relational Algebra (cont.)

Why RA?

- I can do anything with <your favorite PL>!
 - Yes, but only in principle. In practice...
 - How do you represent tuples in <your favorite PL>?
 - How do multiple users share, query, and updated their data?
 - How do you achieve this while keeping, e.g., physical data independence?
- Practical importance.
 - **Strictly less powerful** than <your favorite PL>.
 - Easy to use, e.g., fewer and simpler syntactic constructs.
 - Simpler allows the DBMS to search for efficient query evaluations.
 - Still, expressible enough that it is practically useful.
- Limitations.
 - Finiteness of relations. Not all operations are closed.
 - Aggregates (e.g., MIN, MAX, AVG);
 - Recursion (e.g., transitive closure);
 - Ordering.
