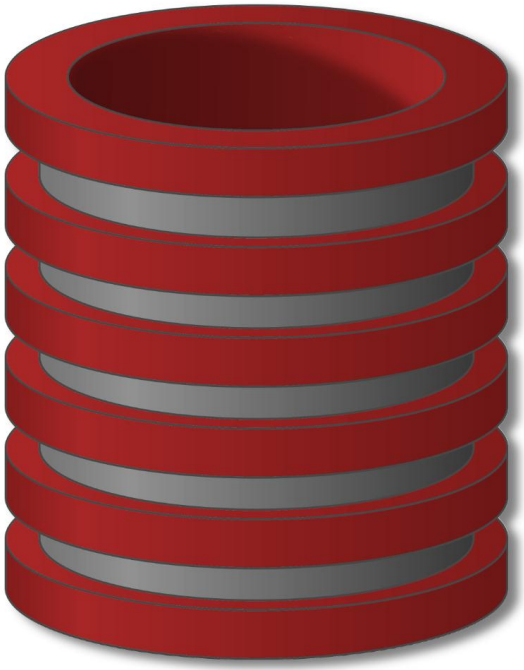

Relational Databases

Querying Relational Databases

Jennifer Widom

Steps in creating and using a (relational) database

1. Design schema; create using DDL
2. "Bulk load" initial data
3. Repeat: execute queries and modifications

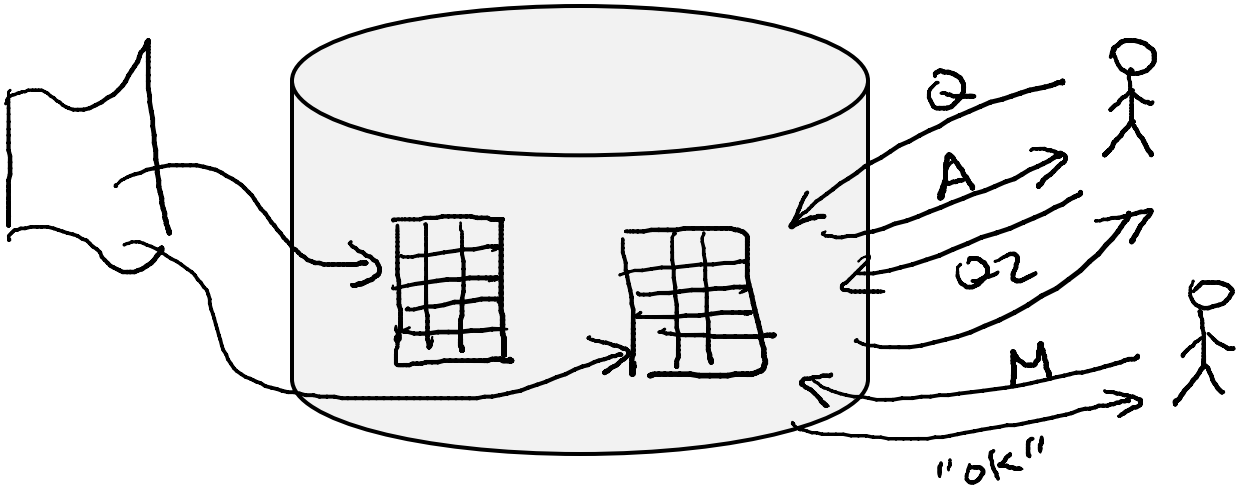

Ad-hoc queries in high-level language

- *All students with GPA > 3.7 applying to Stanford and MIT only*
- *All engineering departments in CA with < 500 applicants*
- *College with highest average accept rate over last 5 years*

- Some easy to pose; some a bit harder
- Some easy for DBMS to execute efficiently; some harder
- "Query language" also used to modify data

Not correlated

DML

Queries return relations (“compositional”, “closed”)

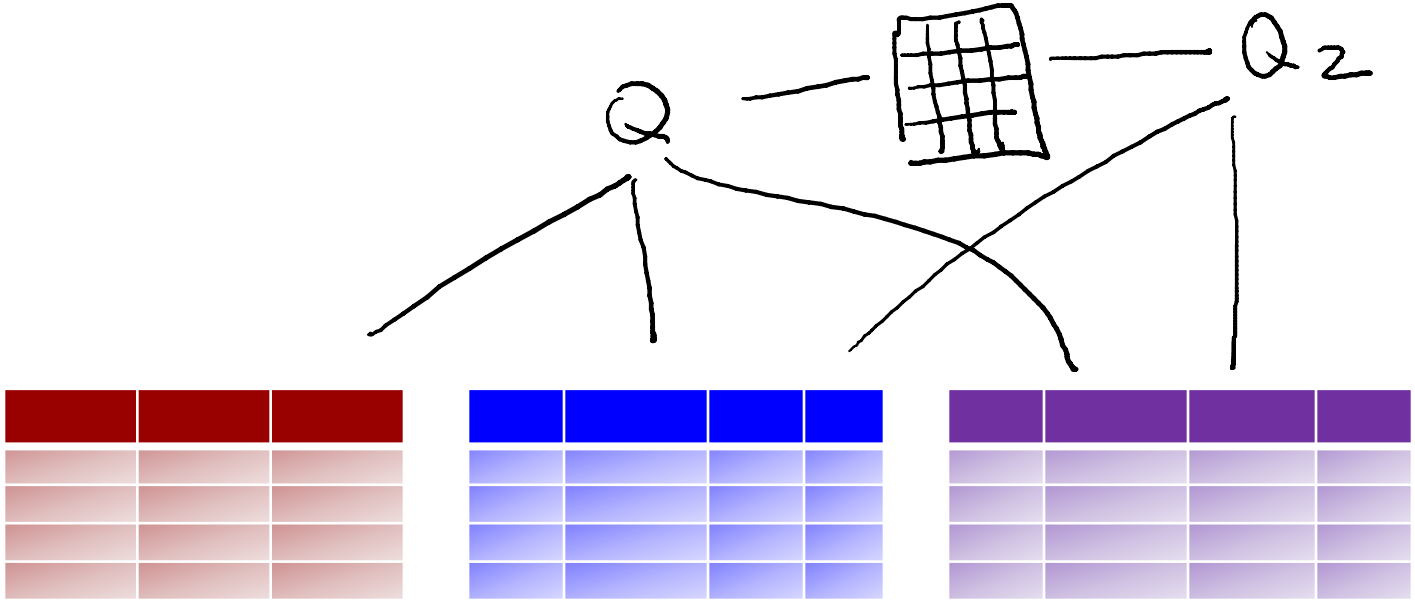

Query Languages

- Relational Algebra — formal

$\Pi_{ID}\, \sigma_{GPA>3.7 \wedge cName='Stanford'}(Student \bowtie Apply)$

- SQL — actual / implemented

```
Select Student.ID
From Student, Apply
Where Student.ID=Apply.ID
And GPA>3.7 and college='Stanford'
```

IDs of students with GPA > 3.7 applying to Stanford