

CS347: Data Management

Introduction addendum

Professor Daniel P. Miranker

Objectives:

- Provide an overview of the course
- Set expectations

1: Intro. Addendum

Data Management

Formality

Relation schema*: the name of a table, its attributes (column names)

- Example 1: Name(First, Last)
- Example 2: T1(a1, ..., an)

Database schema: a set of relation schema

- Example 1:
{Name(student-id, First, Last), Class(class-name, student-id)}
- Example 2:
{R(a1, a2, a3), S(b1,b2)}

* *First use of a formal term will be underlined. Terminology is important*


1: Intro. Addendum

Data Management

DBMS Architecture


Storage Manager

- Exploit memory hierarchy to compensate for slow disks.
 - working sets (from OS)
 - search algorithms
- Specifics
 - manage a *heap* of disk pages
 - allocation of main memory (buffer management)
 - index methods, e.g. B+ tree (access paths)



1: Intro. Addendum

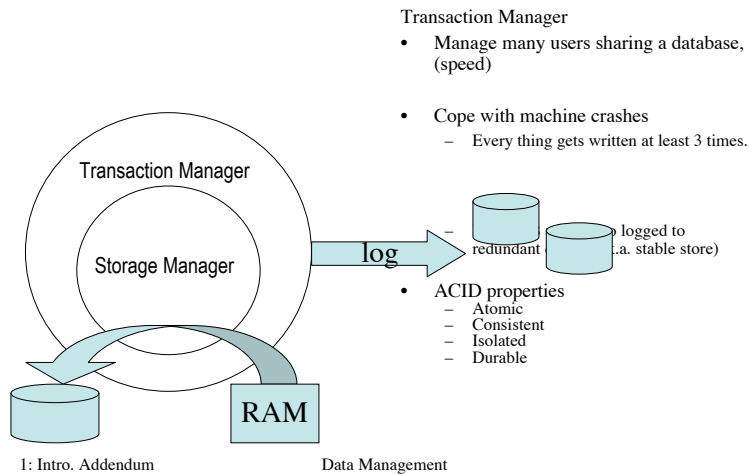
Data Management



1: Intro. Addendum

Data Management

DBMS Architecture, 2



ACID Properties

Atomicity: **all** actions of a transaction happen, or **none** happen.

Consistency: if a transaction is consistent, and the database starts from a consistent state, then it will end in a consistent state.

Isolation: the execution of one transaction is isolated from other transactions.

Durability: if a transaction commits, its effects persist in the database.

I: Intro. Addendum


Data Management

Transaction:

i. collection of actions

that preserve consistency

ii. a process that manifests the ACID properties



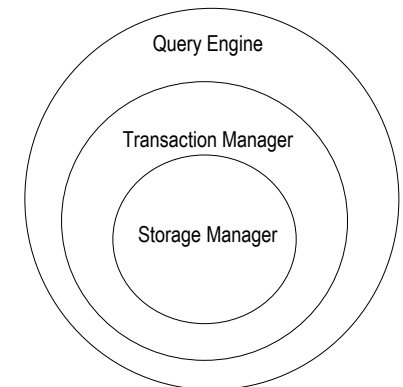
I: Intro. Addendum

Data Management

DBMS Architecture 3

Query Engine

- SQL execution environment
 - parse
 - compile to logical operators
 - optimize: Choose a good set of access paths and sequence of database operators (a.k.a. a physical plan)



I: Intro. Addendum

Data Management

DBMS Architecture 3

Query Engine

- SQL execution environment
 - parse
 - compile to logical operators
 - optimize: Choose a good set of access paths and sequence of database operators (a.k.a. a physical plan)

