Welcome to CO 572: Advanced Databases

Holger Pirk

Purpose of this Lecture

Figuring stuff out

- What you know
  - This should mostly be revision (tell me if it is not)
- What we're trying to achieve
- What I will expect
- What you can expect

Disclaimer

He's new at this

Disclaimer

He's new at this

- Not an excuse for gross incompetence

An example

What Resource bounds this query's performance

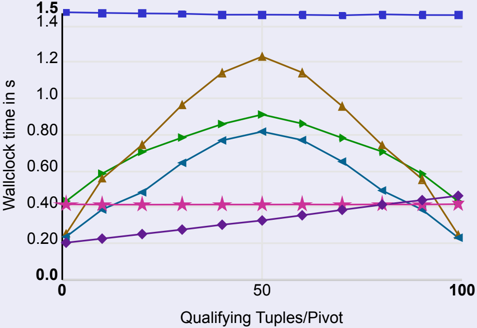

An example

Option A

CPU!

Option B

Memory!

Option C

Disk!

Option D

What?

Another example

What is this Workload's Isolation Level

Effect of long read-only transactions

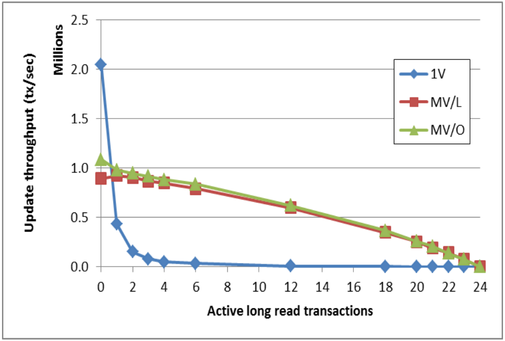

Workload:

- Short txns 10R+ 2W
- Long txns: R 10% of rows

24 threads in total

- X threads running short txns
- 24-X threads running long txns

Another example

Option A

Repeatable Read!

Option B

Serializable!

Option C

Read Committed!

Option D

I Have no idea!

Takeaways

Bottom line

- You (probably) don't know everything
- I certainly don't
- But I know a lot of people who do – ASK!!!
- If noone knows, it is research (see me!)

What is a Database Management System?

Well,

as it says, a

- Database
- Management
- System

What is a Database?

Merriam Webster

Database A usually large collection of organized data

Data Factual information

In Context

- Pretty much any structured collection of data points/data items
  - A relational table
  - A set in your favorite programming language
  - A vector in your favorite programming language
  - A graph
  - A stack of index cards

What is Management?

Merriam Webster

- the conducting or supervising of something
- to handle or direct with a degree of skill
- to work upon or try to alter for a purpose
- judicious use of means to accomplish an end

In Context

- Provide everything you need to manage your data
  - while taking advantage of degrees of freedom
- Usually prescribe an external interface: data model, protocol & semantic guarantees
  - Relations, Documents/Trees, Graphs, Arrays
- Results matter, internal organization is a degree of freedom
  - Storage, Processing, Resilience, etc. ← This is our focus

What is a System?

Merriam Webster

- a regularly interacting or interdependent group of items forming a unified whole

In Context

- Often made up from components,
  - that interact
  - to achieve a greater goal
- Usually applicable to many situations (i.e., generic)
- How is that different from a well-designed application?
  - The goal is domain-agnostic
  - Components serve a data-management purpose, not a domain-purpose

So, then what is a Database Management System?

The Greater Goal

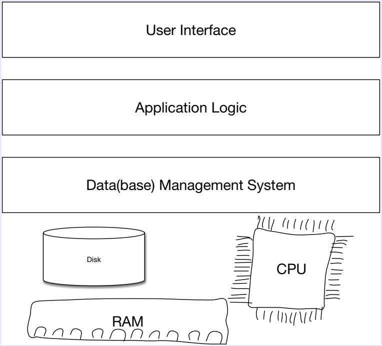

Figure: The role of a DBMS in a software application

So, then what is a Database Management System?

Functionality of a DBMS

- Storage
- Query Processing
- Transaction Processing
- External Interface
- (Access Control)

So, then what is a Database Management System?

Transaction Processing

- Getting data into the database

Storage

- Organizing data for the purpose

Query Processing

- Retrieving data from storage
- Perform a modest amount of computation

External Interface

- The best are standardized
- Often using some kind of query language (like SQL)

What a Database Management System is not?

- A runtime for your applications (though people have tried)
  - some support user-defined functions
  - some even have embedded webservers, middleware, . . .
    - a horrible idea
- A place to store intermediate state
- A filesystem

Why DBMSs do not make good filesystems

Option A

DBMSs are not good at direct-lookups

Option B

If you limit degrees of freedom, performance suffers

Option C

DBMSs do not scale

Option D

They do not?

What is a data management system?

Components

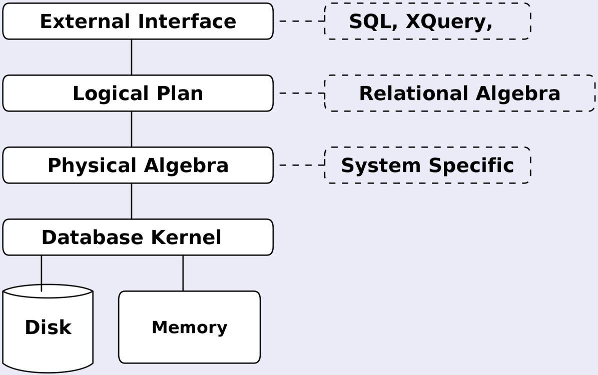

External Interface

Logical Plan

Physical Algebra

Database Kernel

Disk

Memory

SQL, XQuery,

Relational Algebra

System Specific

What Database Management Systems exist?

Domain	Closed Source	Open Source
Relational	Oracle DBMS, IBM DB2, Microsoft SQL Server	PostgreSQL, MySQL, SQLite
OLAP	VectorWise, Vertica, Snowflake, Impala, Redshift	MonetDB
OLTP	VoltDB	Silo, H-Store
Graphs	Virtuoso	Neo4J
Trees/Documents	DocumentDB, IMS (1966!)	CouchDB

What is a database management application?

Not a system

- The boundary is blurry
- Not generic
  - Domain-specific
  - Hard to generalize
  - Often contains domain-specific tricks

Here is a proposed spectrum

- Yelp
- A mobile app for geo-services
- A library to manage unordered collections of tagged coordinates
- A spatial data management library
- A relational database
- A block storage system

(Relational) Database Management Systems

Relations are sets of tuples

Places

Name (PK)	Lonitude	Latitude	Stars
Yummy Burgers	37.9	21.1	4.1
Great Coffee Inc.	36.1	18.9	*NULL*
The Ice Cream Shop	*NULL*	*NULL*	4

Reviews

User	Place Name (FK)	Comment
Holger	Yummy Burgers	Just awful, never again
Peter	The Ice Cream Shop	Quite alright

Normalization

Denormalization leads to redundancy

- Why is redundancy bad
  - ...
  - ...
  - ...

The goal of normalization

- Eliminate redundancy

Normalization

Denormalization leads to redundancy

- Why is redundancy bad
  - ...
  - ...
  - ...

The goal of normalization

- Eliminate redundancy

1NF No complex attributes

2NF No functional dependencies within tables

3NF All primary keys are defined

Schemas

The definition of the attributes of the tuples in your relations (duh)

```
create table review (stars int,
comment varchar(1024), user int);
```

But also integrity contraints (uniqueness, keys, foreign-keys)

```
alter table review add
foreign key(user) references user(name);
```

Some people distinguish external and internal schema

- The idea of internal schemas is misleadingly simplistic
  - there may not even be a schema

Internal Storage

Many degrees of freedom

- Data could be stored
  - Column-wise, Row-wise, Hybrid
  - In trees, graphs, etc.
- Directly in indices
- Normalized or denormalized
- Compressed
- On disk, in memory or on GPUs
- On a remote machine or even in the cloud

Transactions

Isolated Run like you were alone on the system

Atomic Run completely or not at all

Consistent The interesting thing here is that there may be inconsistency in between

Durable After the transaction commits, even power outage won't undo the transaction

Question

Why are databases a bad place for intermediate state

Why does any of this matter?

Why does any of this matter?

DBMSs make lots of money

- We are talking about it 50 billion dollar market (in 2017)
  - That is only the pure sales volume of relational DBMSs
  - add administration, tuning, application development, . . .

they also are an important part of our life

- Online shops
- Banks
- Online content
- Your phone (SQLite probably has billions of running instances)
- . . .

Why would I care?

Why do I need to know?

- If you ever happen to work on a DBMS (not very likely)
- To make you a kick-ass DBA (somewhat more likely)
- To apply data management techniques outside the field (extremely likely)
  - Bragging rights if you're the guy to implement a radix-partitioned in-memory hash-join
  - Some of this actually comes up if you interview at Google, Facebook, Microsoft, etc.

What does this course look like

Some Admin stuff

- Course is taught in two halves
  - First halve on single-node databases (taught by your's truly)
  - Second halve on distributed databases (taught by Peter McBrien)
- Register for this course (at level 2)!!!
  - Bad things will happen if you don't
  - Ask your cohort admin person how to do it

Things I will expect

- Basic, knowledge about algorithms & data structures
  - Arrays, linked lists, trees, heaps, . . .
  - Sorting, binary searching, graph/tree-traversal
- Some basic computer architecture knowledge
  - Main memory, disks, caches, multicore, . . .
- Programming C++
  - Programming for one of the two coursework assignments
- Honesty when preparing coursework
  - Don't try me!

Coursework

- There will be three assignments
  - Using a database (with me)
  - Database internals (also with me)
  - Distributed data processing (with Peter McBrien)
  - This is feedback
- Exam in the end
  - Relevant is what is discussed in class

Books

Fundamentals of Database Systems

Ramez Elmasri, Shamkant Navathe
Sixth edition, Pearson new international edition., Pearson,

Database Systems: The Complete Book

Hector Garcia-Molina, Jeffrey D. Ullman, Jennifer Widom
2nd ed., Pearson Education

Database Systems: Practical Approach to Design, Implementation, and Management

Thomas M. Connolly, Carolyn E. Begg
Sixth edition, Global edition., Pearson Education Limited

Webpage

`https://www.doc.ic.ac.uk/~pjm/CO572/`

Outline

- Object-Relational mapping
- Data Storage
- Querying relational data
- Join formulation and evaluation
- Query planing and optimization
- Processing models
- Secondary storage

This was it

This was it

See you on friday for a class on Object Relational Mapping (thank me later!)

Remember

- Register for this course
    - at Level 2!!!