Datenbanksysteme II: Implementation of Database Systems

Ulf Leser

- Slides in English, Vortrag auf Deutsch
- Much input from
  - Prof. J-C Freytag, HU Berlin
  - Prof. K-U Sattler, TU Illmenau
  - Prof. A Kemper, Dr. Eickler, TU München
- AGNES
- Prof Freytag / Prof Leser
- What you should (must) know to follow this course
- What do you study?

Some Motivation

A classical first contact with database

- Company: “I need to track my customers”
  - Name, address, profession, prior contracts
- Naïve engineer: “No problem”
  - ~1984: Turbo Pascal 4.0, Schneider CPC646, 512 KB main memory
  - Each customer one record / line in file
  - Load customers from disk into memory
  - Repeat until “Q”
    - (S)earch and list customers
    - (E)dit customer
    - (D)elete customer
    - (I)nsert customer
    - (Q)uit
  - Write customers to disk
- Invoice: ... DM

Story part 2

- Company: “I need to track my offers”
  - Customers have projects and call for bids, company makes offers
  - Many customers have many projects over time
- Naïve engineer: “No problem”
  - Reuse existing architecture
  - Load offers from disk into memory
  - Repeat until “Q”
    - (S)earch and list offers
    - (E)dit offer
    - (D)elete offer
    - (I)nsert offer
    - (Q)uit
  - Write offers to disk
- Invoice: more DM

Story part 3

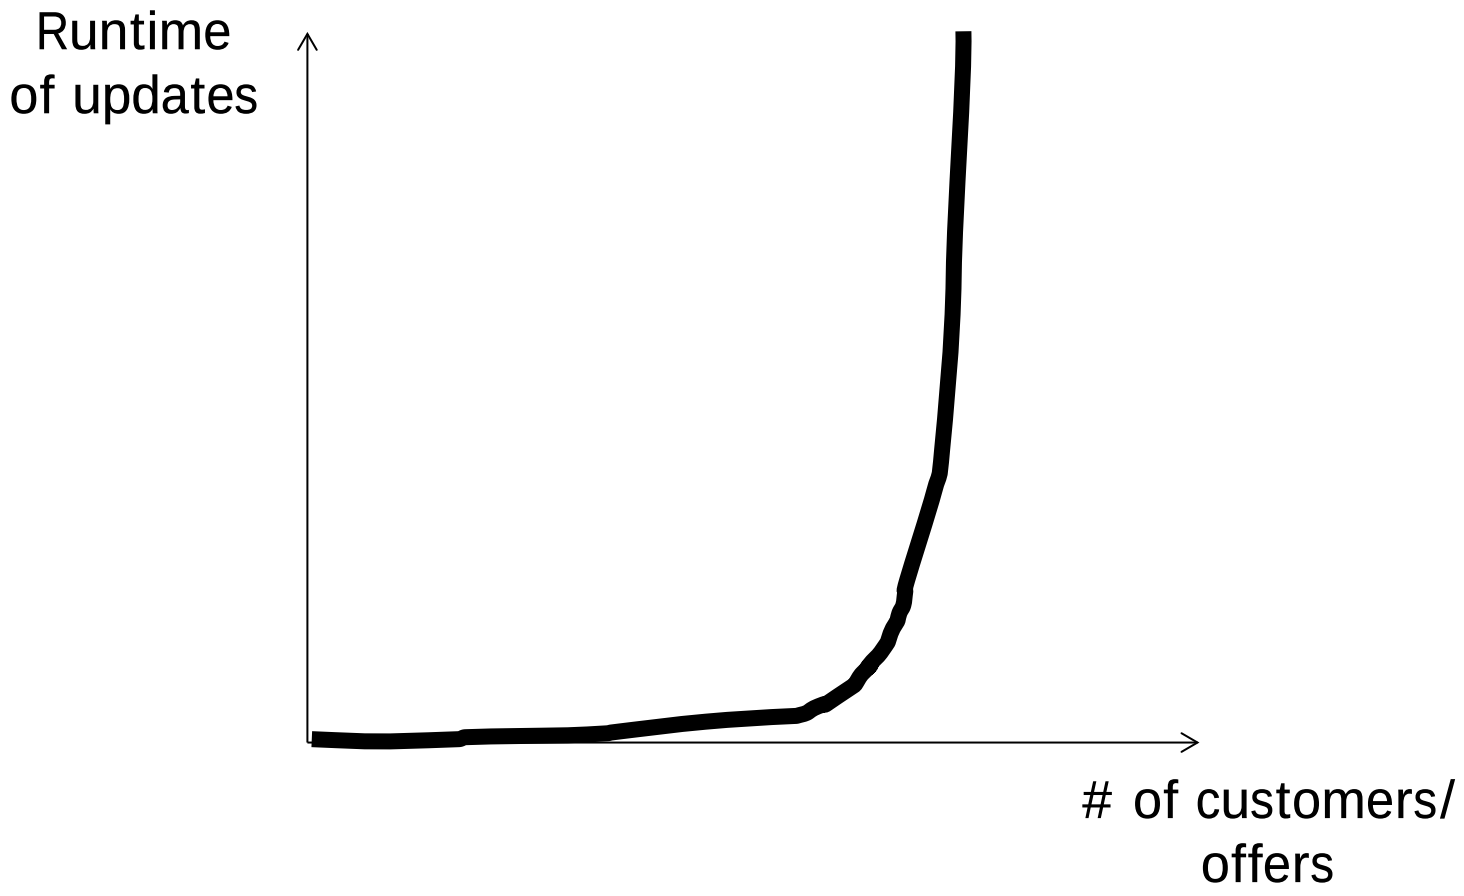

Part 3

- Disaster!
  - ~ 500 customers
  - ~ 40 offers per customer
  - ~ 2KB per offer
  - Gives 500*40*2.000 = 40 MEGABYTE
    - This was the size of a hard disc at that time
  - No way to load and hold all data at startup
- Wrong architecture
- Solution: Buy a RDBMS

Lessons learned

- Scalability in data management is an issue
  - Scalable: Graceful runtime degradation with increasing data volumes
  - Not scalable: Works fine for small datasets, breaks down for large datasets
- Data is business-critical
  - If offers-file corrupted – company goes out of business
    - Which will affect engineer too!
- Think before you start programming
  - Project 100% over budget (database license)
  - Project 300% over time (6 months instead of 2)

Second Example: Linguistic Databases

(Victor Rosenfeld, 2013)

Wunder	gibt	es	immer	wieder	!
Wunder	geben	es	immer	wieder	!
Acc.Pl.Neut	3.Sg.Pres.Ind	3.Nom.Sg.Neut	--	--	--
NN	VVFIN	PPER	ADV	ADV	$.

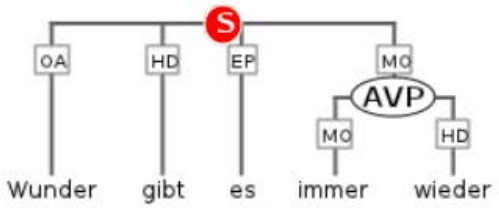

Semantic types

Steilpass Wunder gibt es immer wieder ! Erst spielen die Dallgower Gemeindevertreter so statisch und verzagt wie die deutsche Abwehrreihe der Fußballkicker . Und dann kommt aus der Tiefe solch ein fulminanter Steilpass , von dem man hofft , dass die Seeburger oder Groß-Glienicker Mitspieler ihn aufnehmen können . Ein Befreiungsschlag ist es allerdings nicht , weil es vorerst keine Gefahr fürs Dallgower Tor gab . Die Seeburger und einige Groß-Glienicker haben den Ball erst zurückgespielt und dann um so drängender wieder gefordert . Nun sollen sie zeigen , wie sie die Chance verwerten . Eine Diskussion , wo künftig die Trainerkabine stehen soll , wäre in der jetzigen Spielsituation verheerend . Und eine Parallele zu den deutschen Grotten-Kickern gibt es immer noch . Auch wenn die Spieler aus den verschiedenen Vereinen zusammengewürfelt sind , sie müssen sich daran gewöhnen , dass sie nun in einer Mannschaft " Döberitzer Heide " spielen . Und das heißt gemeinsam und nicht gegeneinander . Ermahnungen von der Seitenlinie , miteinander fair umzugehen und sich nicht beim kleinsten Schubser gegenseitig zu zerfleischen , sind normalerweise überflüssig . Vorerst allerdings

Cross-sentence links

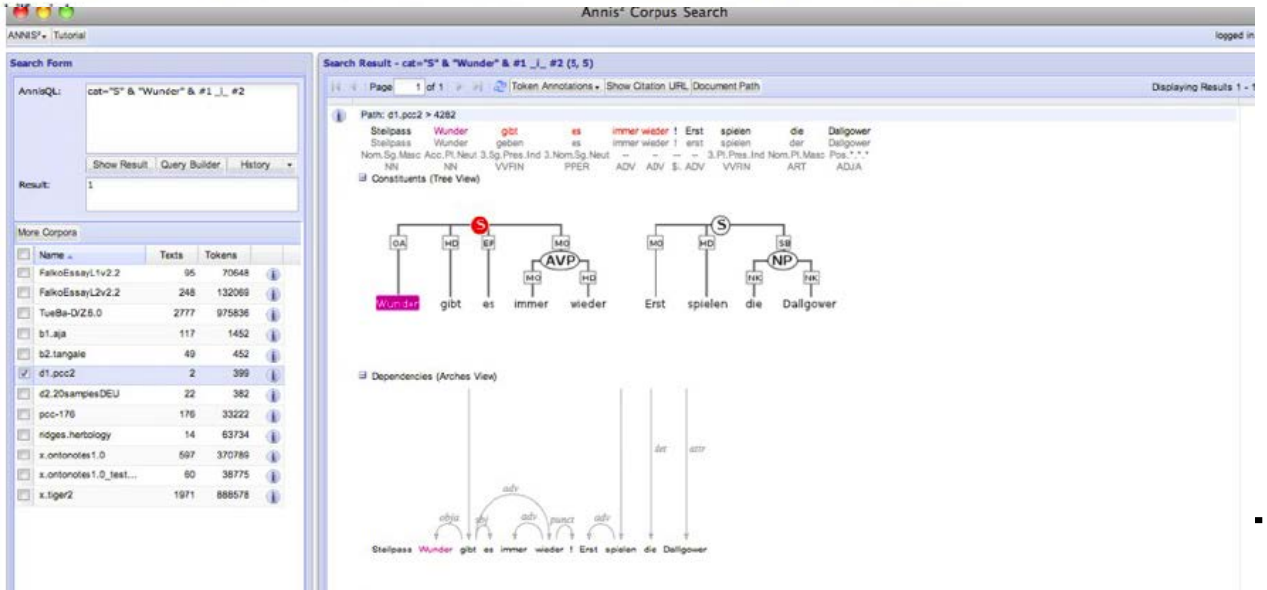

AQL Queries

Wunder	gibt	es	immer	wieder	!
Wunder	geben	es	immer	wieder	!
Acc.Pl.Neut	3.Sg.Pres.Ind	3.Nom.Sg.Neut	--	--	--
NN	VVFIN	PPER	ADV	ADV	$.

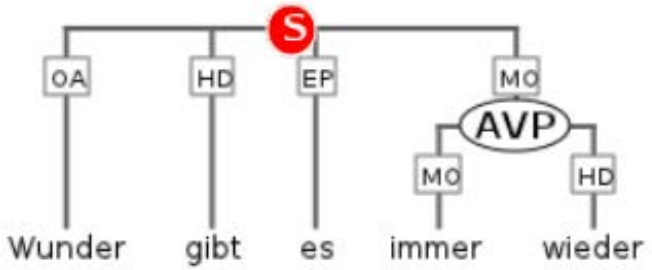

cat=„S“ &	Find all sentences; bind to variable #1
„wunder“ &	Find all token „Wunder“; bind to variable #2
#1 _i_ #2	Join: remove #1 which do not include #2

Let‘s do it right - PostGreSQL

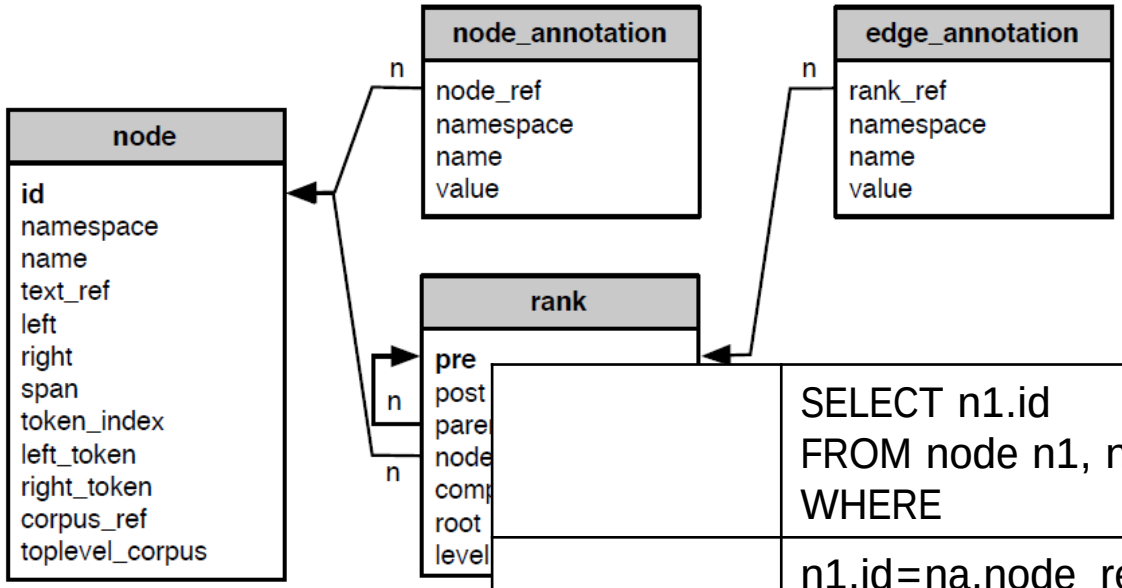

	SELECT n1.id FROM node n1, node n2, node_annotation na WHERE
cat=„S“ &	n1.id=na.node_ref AND na.name=„cat“ AND na.value=„S“ AND
„wunder“ &	n2.span=„Wunder“ AND
#1 _i_ #2	n1.text_ref=n2.text_ref AND n2.right<=n1.right AND n2.left>=n1.left

More Complicated Queries

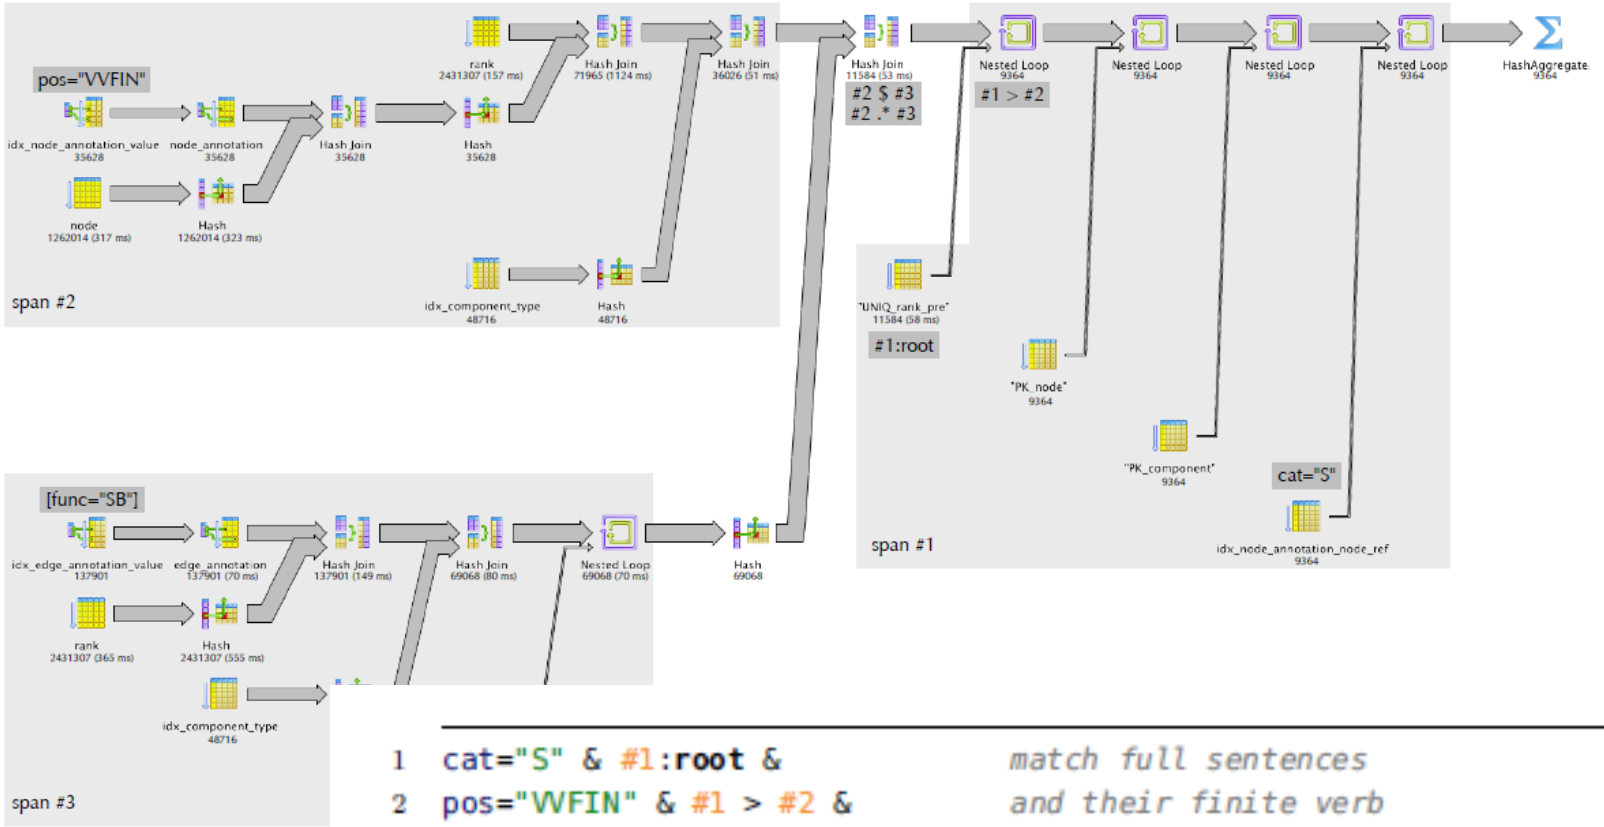


```
1  cat="S" & #1:root &          match full sentences
2  pos="VVFIN" & #1 > #2 &      and their finite verb
3  node & #1 >[func="SB"] #3 &  and their subject
4  #2 .* #3 &                   where the verb precedes the subject
5  meta::Genre="Politik"        only consider documents of the genre Politik
```

Listing 1: Annis query matching sentences in which the subject follows the verb.

Did we do it right?

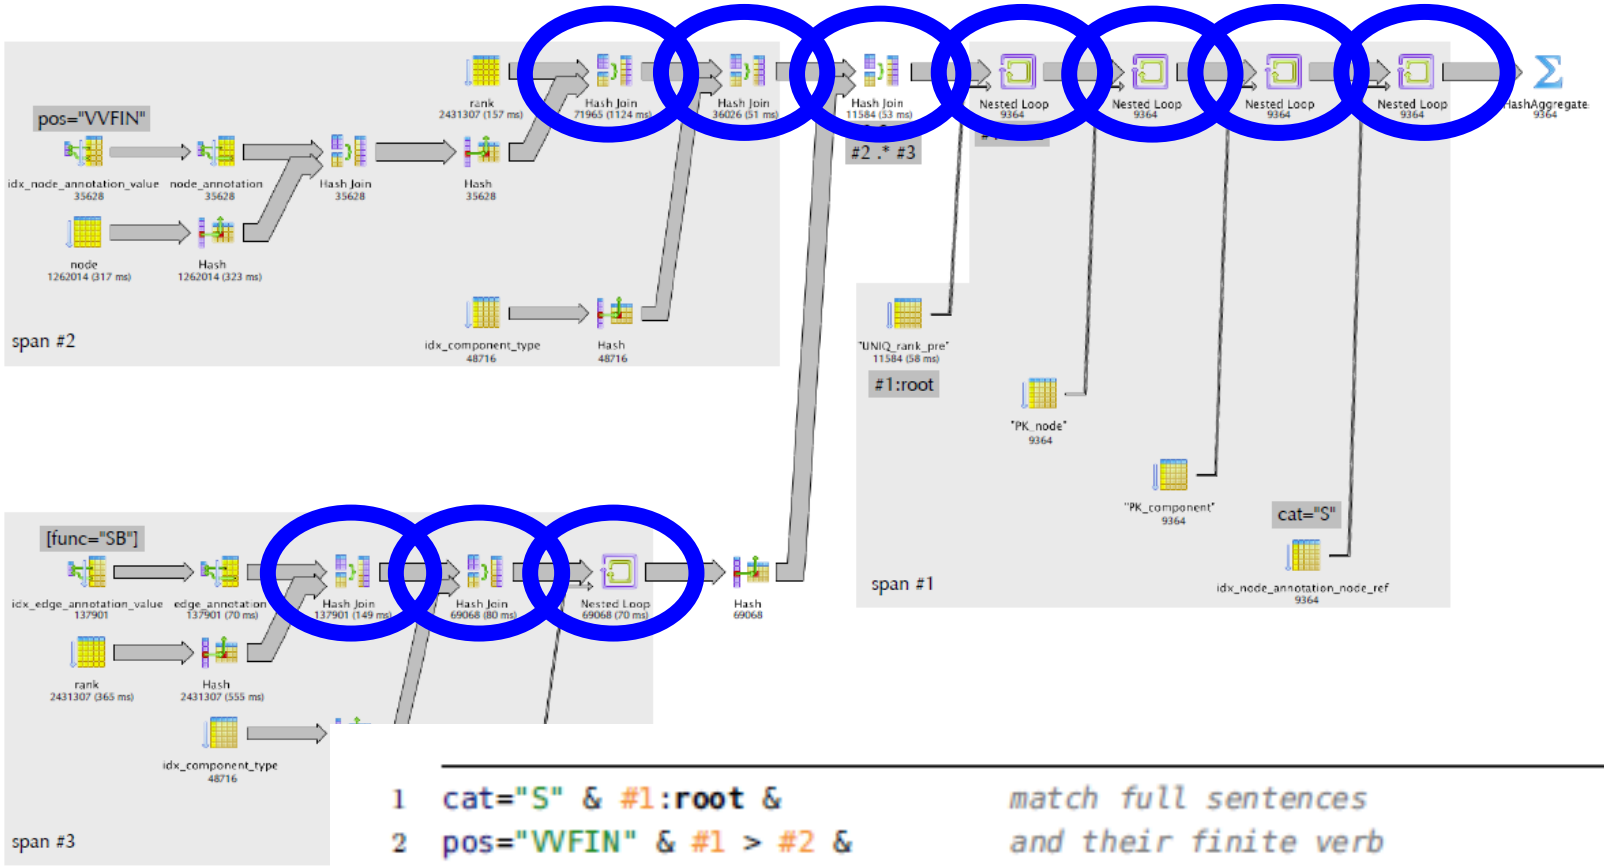


```
1  cat="S" & #1:root &          match full sentences
2  pos="VVFIN" & #1 > #2 &      and their finite verb
3  node & #1 >[func="SB"] #3 &  and their subject
4  #2 .* #3 &                   where the verb precedes the subject
5  meta::Genre="Politik"        only consider documents of the genre Politik
```

Listing 1: Annis query matching sentences in which the subject follows the verb.

RDBMS Feature: Indexes, Materialized Views

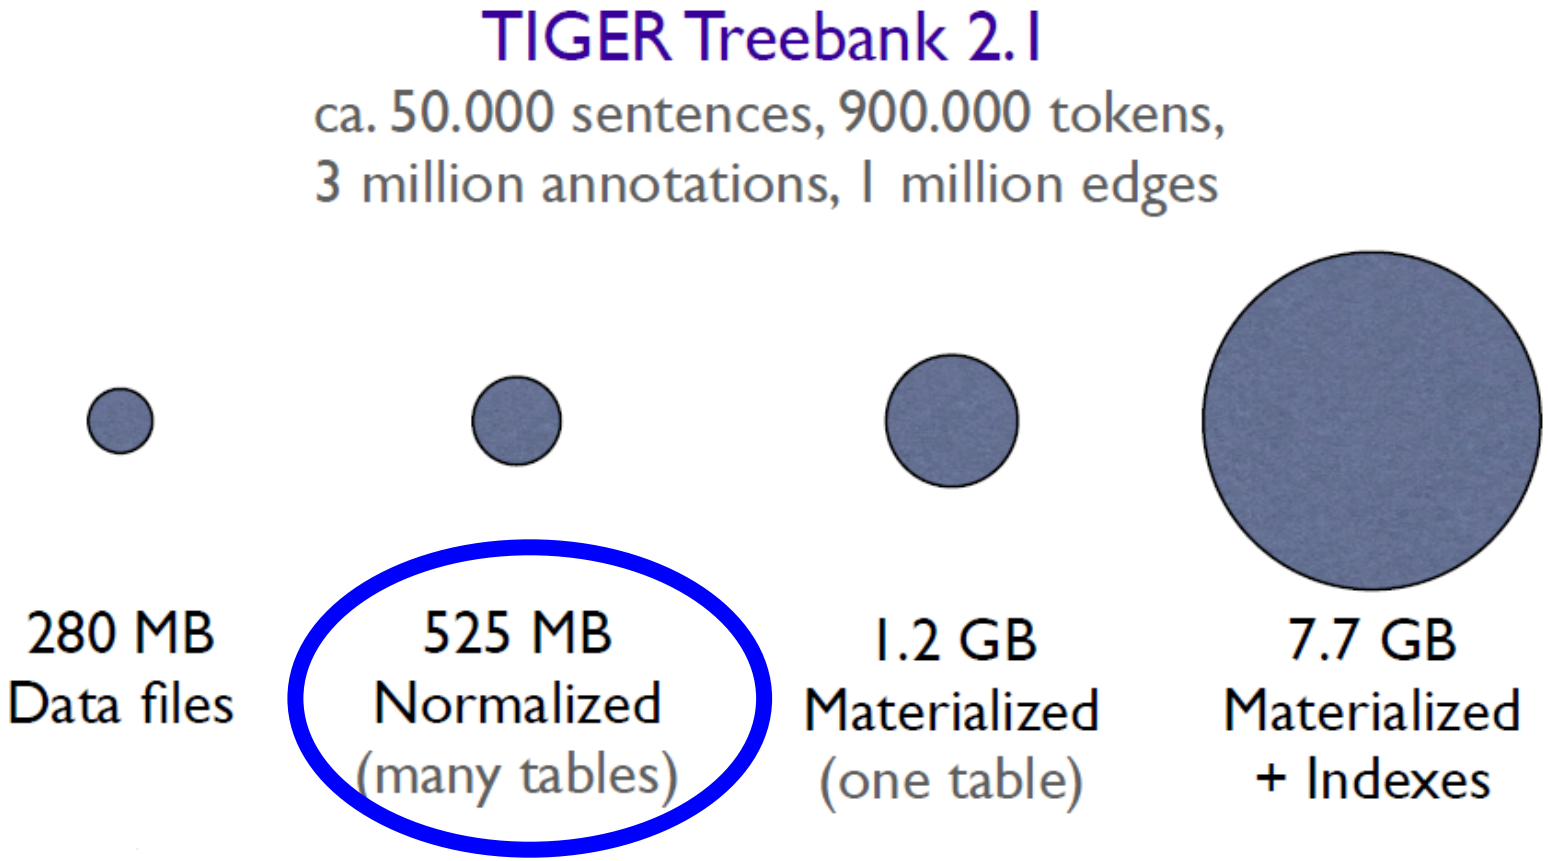

It is 2016 – we can keep this in memory!

MonetDB: A Main-Memory Column Store

- Workload: 330 real-life queries
- MonetDB is a RDBMS, but
  - All data kept in main memory
  - No indexes – all scans
  - Column store: Keep column values together, not tuples
- Advantages
  - No IO, buffering, caching, ...
  - Much better cache utilization for scans (outweighs missing indexes)
  - Column compression (memory, faster scans)
- Still relational: Many joins

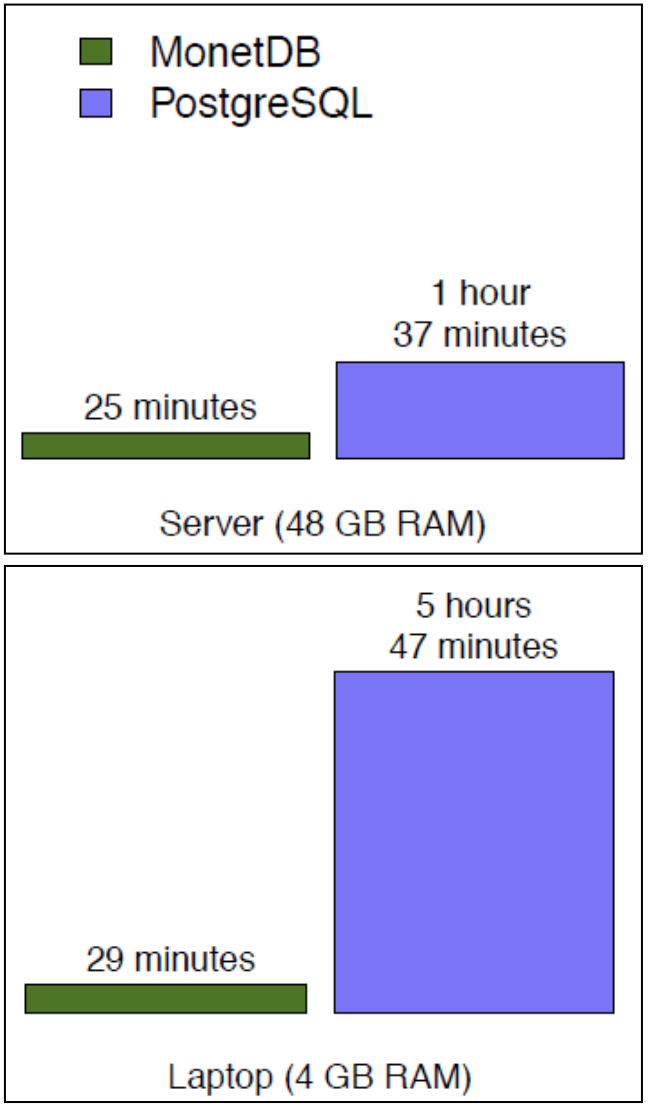

Query Optimization is Difficult to Predict

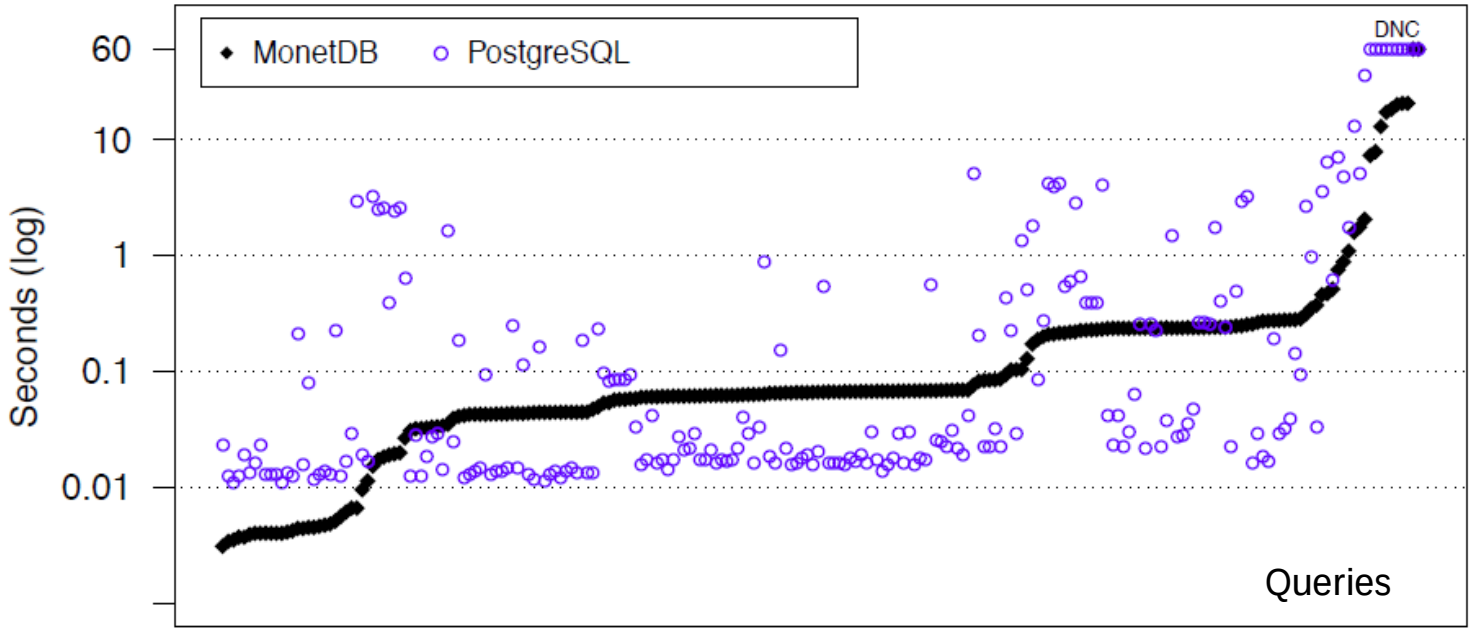

- PostGreSQL is faster for many queries despite IO
- But if it is slower, it is much slower (log scale)

Even Better: A Graph Store

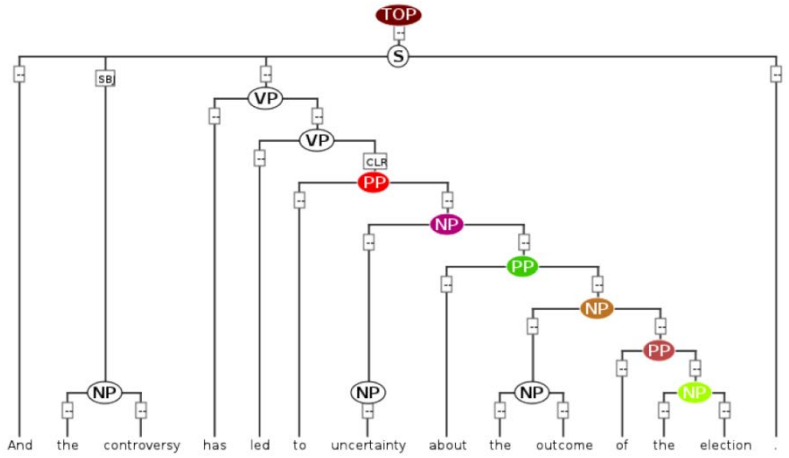

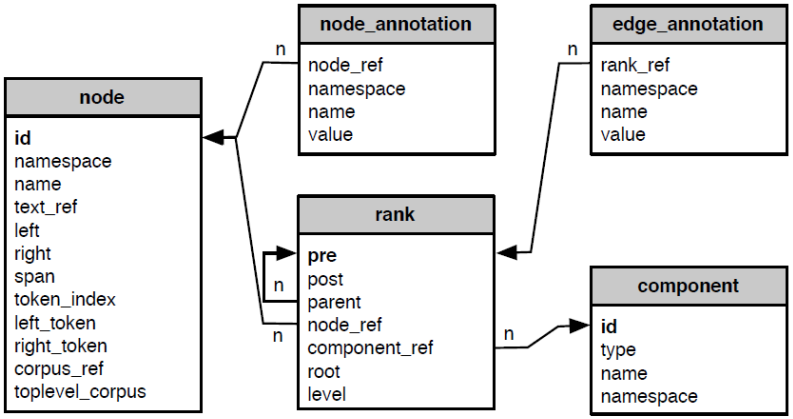

- AQL queries navigate through graphs
- Relational: One join for (almost) every edge traversal
- graphANNIS: AQL on a main-memory graph data model
  - No joins, but following pointers
    - Implemented as indexes into arrays
  - Indexes to find the right nodes quickly (to start traversals)

Thomas Krause (submitted)

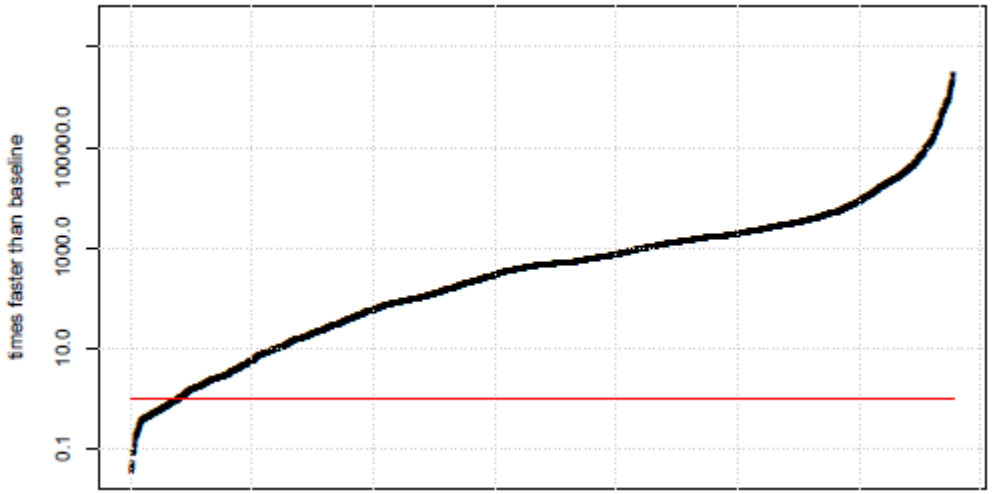

- Workload: ~3300 real-life queries against ~20 corpora
- graphANNIS versus PostGreSQL
  - ~40 times faster on the entire workload
  - Faster for 97% of all queries
  - Not much slower for the remaining 3%

Databases are Infrastructure

Classical Three-Tier Architecture

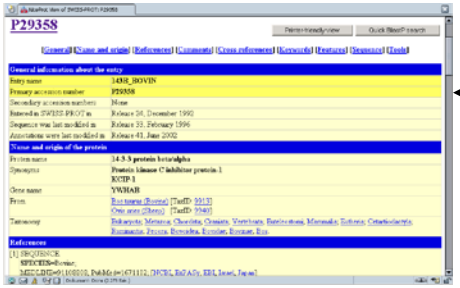

Servlets/
EJB

Presentation

Upwards: OO Interface
Application Server
Downwards: SQL

Application logic
“Business processes”

DBMS

Database 1

Database 2

“State”
Storage, Search
Recovery, Consistency

Today's Database Systems

- RDBMS are essential parts of enterprise infrastructures
  - More important than OS
  - Long-running, expensive, essential investment
  - Holds the most important business assets: Data, information
- Database administrator is a well-paid profession
  - Developers write SQL & business logic
  - Admins make SQL run fast
  - Many programmers, fewer DB developers, very few DB admins
  - A skills much demanded in industry
- RDMS became an often "invisible" piece of software
  - "So nützlich wie fließendes Wasser" (G. Weikum, MPI Saarbrücken)

Main Features

- Data needs to be stored
  - Disk access (or cache utilization) is the main bottleneck
  - Hence: Minimize access time -> minimize disk access
- Data is manipulated by many clients
  - Concurrent access quickly screws up data
  - Hence: Synchronize access
- Data is used by many apps with different requirements
  - No good to design application specific “optimal” data structures
  - Hence: Use application independent languages and models
- Systems crash
  - Crashes cannot be avoided
  - Hence: Protect consistency by logging, constraint enforcement, ...

DBS2: Implementation of Database Systems

- Lecture 4 SWS
  - Wednesday, 11 – 13 , 3.113
  - Thursday, 11 – 13 , 3.113

- Contact

Ulf Leser
Room: IV.401
Tel: (030) 2093 – 3902
Mail: leser (at) informatik.hu-berlin.de

Exercises and Examination

- Exercises run by Jörg Bachmann (DBIS)
  - Presence & commitment are necessary
  - Implementation of file-/ buffer-/ index manager in C++
  - Wednesday, 13-15, 3.113
  - Starts today (probably not!)
- Examination
  - Oral or written?
  - Oral exam dates will be set mid-January
  - Admission: Passing the exercises

Slides

- Slides are available shortly after the lecture
- Please send me
  any errors
- Slides are
  - not a script
  - no substitution for
    listening to the lecture
  - no substitution for
    reading a book

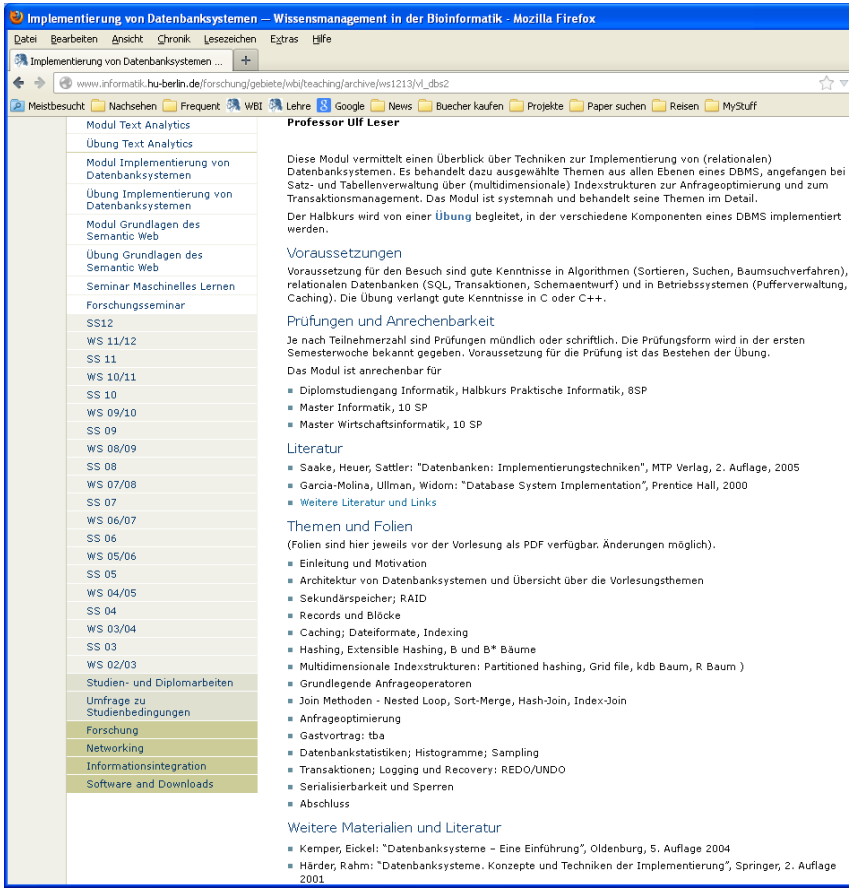

Literature

- Primary
  - Saake, Heuer, Sattler “Datenbanken: Implementierungstechniken”, mitp Verlag, 2005 (2. Auflage)
  - Garcia-Molina, Ullman, Widom: “Database System Implementation”, Prentice Hall, 2000
- Other
  - Kemper, Eickel: “Datenbanksysteme – Eine Einführung”, Oldenburg, 5. Auflage 2004
  - Härder, Rahm: “Datenbanksysteme. Konzepte und Techniken der Implementierung”, Springer, 2. Auflage 2001
  - R. Elmasri und S.B. Navathe: Fundamentals of Database Systems, Benjamin Cummings
    - Deutsche Übersetzung: „Grundlagen von Datenbanksystemen“, Pearson, 2002

Überblick

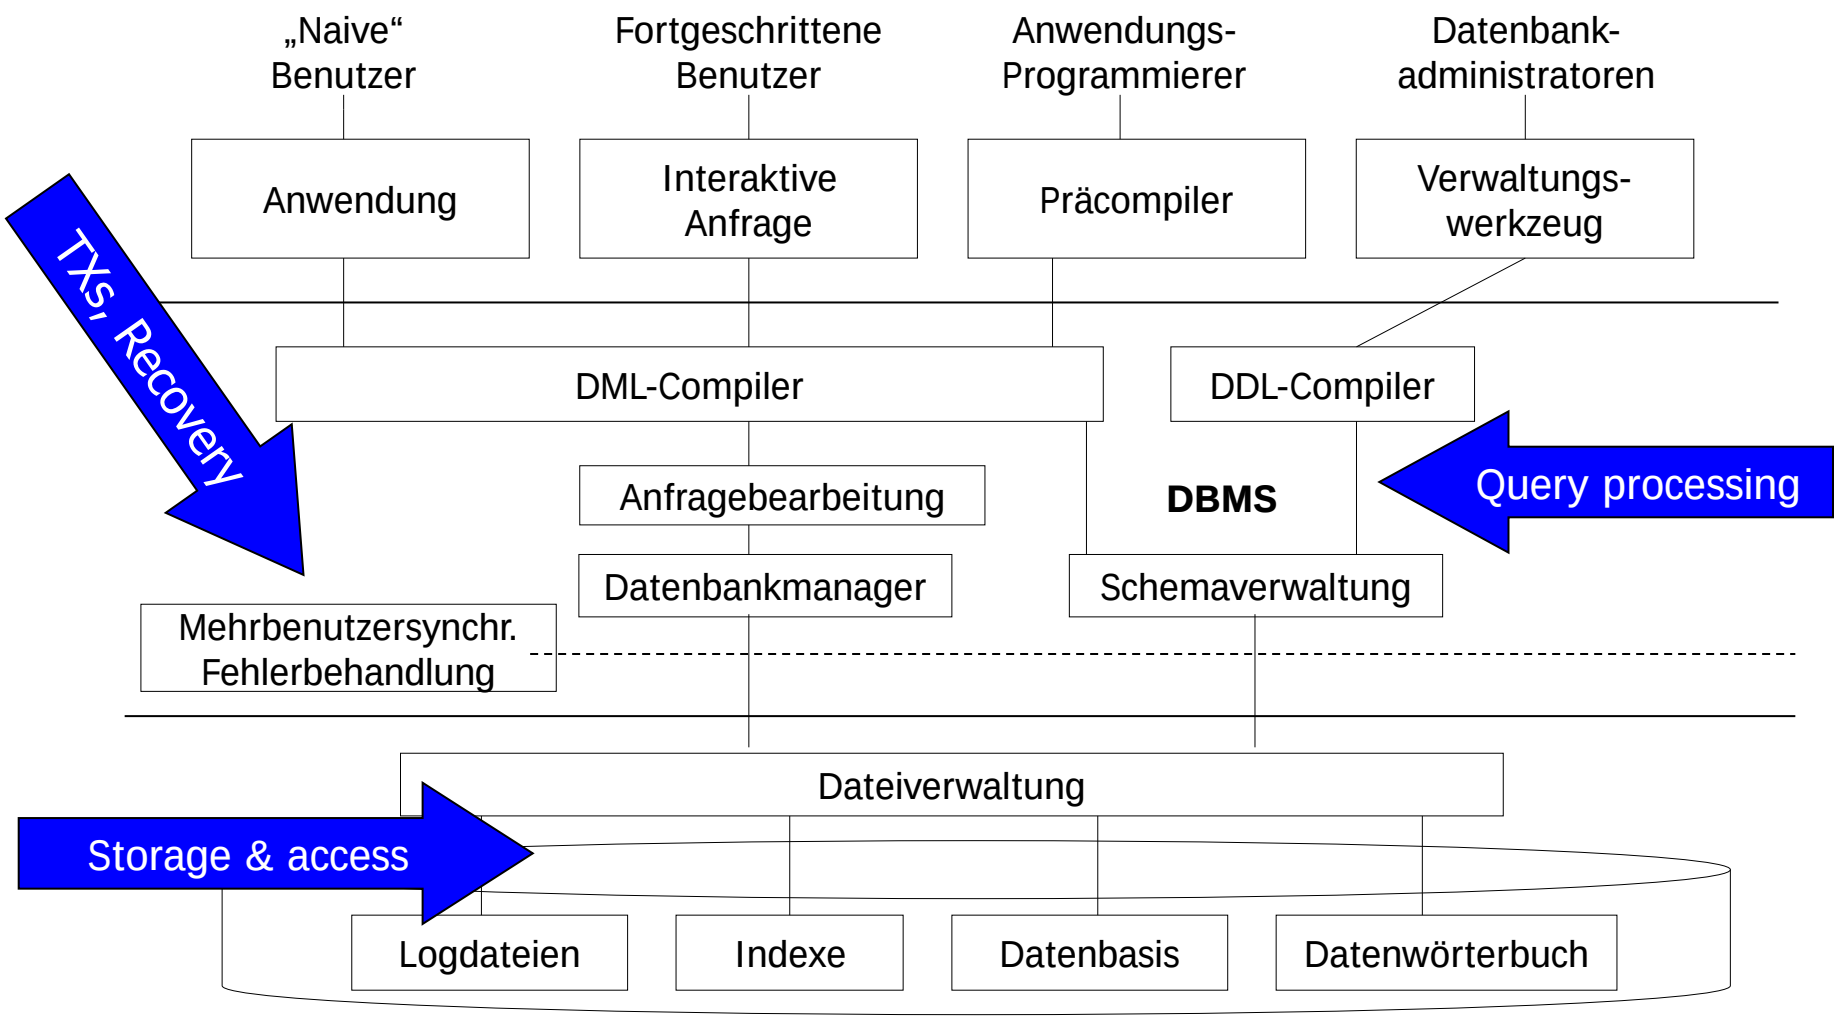
„Naive“ Benutzer
Fortgeschrittene Benutzer
Anwendungs-Programmierer
Datenbank-administratoren
Anwendung
Interaktive Anfrage
Präcompiler
Verwaltungs-werkzeug
TXs, Recovery
DML-Compiler
DDL-Compiler
Query processing
Anfragebearbeitung
DBMS
Datenbankmanager
Schemaverwaltung
Mehrbenutzersynchr.
Fehlerbehandlung
Dateiverwaltung
Storage & access
Logdateien
Indexe
Datenbasis
Datenwörterbuch

Contents

- Introduction
- Overview and architecture
- Storage and access methods
  - B*-Trees, Extensible hashing, index-sequential files ...
  - Multidimensional indexing: Grid-files, kd-Trees, R-Trees ...
- Query processing and optimization
  - Physical relational operators
  - Cost-based optimization
- Recovery
- Transactions and concurrency control

5 Schichten Architektur

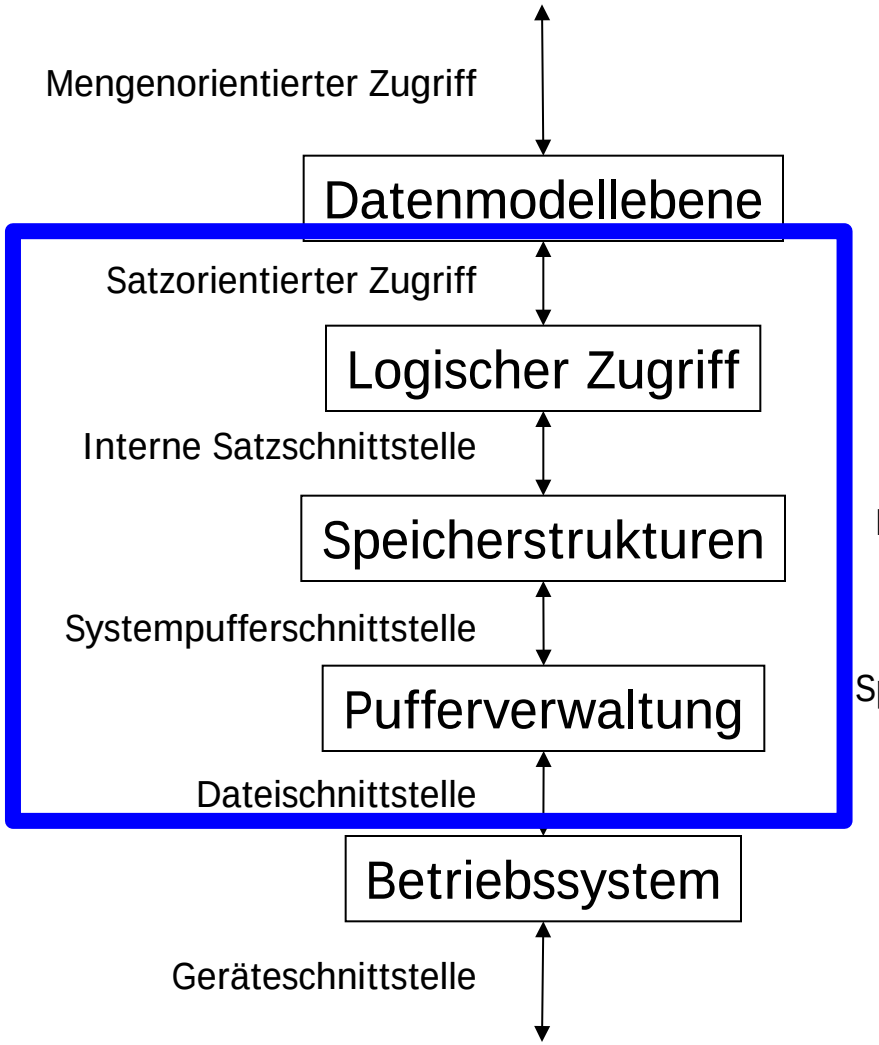

Anfrageübersetzung, Zugriffspfadwahl, Zugriffskontrolle, Integritätskontrolle

Sortierung, Transaktionsverwaltung, Cursorverwaltung

Record Manager, Sperrverwaltung, Log / Recovery

Speichermanagement, Puffermanagement, Caching-Strategien

Externspeicherverwaltung

Guests

- TBA ...

Feedback 2012 / 2013

Alter	Geschlecht	Gefehlt	Teilnehmerzahl	Warum kommen?	Studiengang	Fachsemester	Freundlich	Fragen	Sprache	Präsentation	Beispiele	Konzeption	Überblick	Viel neues	Kritische Auseinandersetzung	Nützlich	Lernziele	Materialien	Tempo	Schwierigkeit	Arbeitsaufwand	Dozent	Vorlesung	Abweichung vom Optimum
21	M	1	3	2	BA	5	6	5	6	5	6	5	5	6	4	6	5	5	3	3	3			**8**
28	M	0	3	2	?	5	6	6	4	6	6	6	6	6	5			5	3	3	3	1	1	**4**
35	M	4	3	2	DI	8	6	6	6	4	6	6	6	6	5	6	6		3	3	4	1	1	**4**
25	M	5	3	2	DI	11	6	6	6	4	6	5	5	6	4	5	5	5	3	3	3	1	1	**9**
35	M	2	3	2	?	6	6	6	6	6	6	6	6	5	6	5	5		2	2	3	1	1	**5**
24	M	0	3	1,2,3	DI	9	6	6	6	6	6	6	6	6	6	6	6	6	3	3	3	1	1	**0**
25	M	1	3	2,3	MW	3	6	6		5	6	5	6	6	6	6	6	5	3	3	3	1	1	**3**
22	M	4	3	2	MI	1	6	5	6	5	5	6	6	4	4	5			3	4	3	1	2	**10**
24		5	2	2	MI	1	6	6	6	5	5	5	6	6	5	5	6		3	4	3	1	1	**6**
22	M	0	2	2	MI	1	6	6	5	6	6	6	6	5		6	6	4	3	4	3	1	1	**5**
24	M	0	3	2	MI	1	6	5	6	5	5	6	6	5	4	5	6	4	2	3	3	1	2	**11**
25,9		**2,0**	**2,9**			**4,6**	**6,00**	**5,73**	**5,70**	**5,18**	**5,67**	**5,67**	**5,82**	**5,55**	**5,14**	**5,44**	**5,63**	**4,86**	**2,82**	**3,18**	**3,09**	**1,00**	**1,20**	
			3,0				**6,00**	**6,00**	**6,00**	**6,00**	**6,00**	**6,00**	**6,00**	**6,00**	**6,00**	**6,00**	**6,00**	**6,00**	**3,00**	**3,00**	**3,00**	**1,00**	**1,00**	
			0,10				**0,00**	**0,27**	**0,30**	**0,82**	**0,33**	**0,33**	**0,18**	**0,45**	**0,86**	**0,56**	**0,38**	**1,14**	**0,18**	**-0,18**	**-0,09**	**0,00**	**-0,20**	
							0,00	0,27	0,30	0,82	0,33	0,33	0,18	0,45	0,86	0,56	0,38	1,14	0,18	0,18	0,09	0,00	0,20	6,28

Free Text

- **Besonders gut**
- 5 Auftreten Dozent
- 4 Small Talk am Anfang
- 5 Beispiele
- Deutsche Vorlesung
- Engagement Dozent
- 2 Gastvorträge
- 2 Konkreter Anwendungsbezug
- Übung
- Tempo

- **Verbesserung**
- Struktur Folien
- Tafelbild zu unklar
- Warum nur Oracle/MySQL
- Liegt zeitlich schlecht
- Fehlt: PL/SQL, UDF, ...
- Entweder nur vormittags oder nur nachmittags
- Praktikum nicht schwer, aber aufwändig
- Übung enger mit Vorlesung verbinden

Datenbanken@Informatik

- A predefined focus area in our Master

- Datenbanken 1: Grundlagen (BA)
- Information Retrieval (BA)
- Datenbanken 2: Implementierung
- Data Warehousing und Data Mining
- Informationsintegration (inkl. verteilter Anfrageoptimierung)
- Neue Konzepte und Techniken für Datenbanksysteme
- Techniken und Konzepte zum Schutz der Privatsphäre
- Datenbanktheorie
- Frequent seminars