Seminar

Optimizing data management on new hardware (OpDaMNeHa)

Summer Term 2014

Lehrgebiet Informationssysteme

Weiping Qu
qu@cs.uni-kl.de

AG Datenbanken und Informationssysteme

AG Heterogene Informationssysteme

Goals

a) Familiarize yourself with a scientific topic independently
b) Find scientific literature on web or in the library, cite correctly
c) Prepare a written composition, presentation, and discussion
d) Time Management, Deadline!
e) Warm-up for your bachelor/master thesis

Plan I

Task 0: Kick-off meeting (today, April 23rd)

- Assign topics

Task 1: Searching for Literature (2 weeks later, May 7th)

- Read the basic papers provided by your supervisor
- Understand your topic
- You have to search for further papers[1]
- Collect a list of related papers and send it to your supervisor

Task 2: Annotated Table of Contents (TOC) (10 days later, May 16th)

- Prepare an annotated TOC for your supervisor (approx. 1.5 pages)
- Including chapter names, section names and comments

Links

1. For finding related scientific papers
(*download free only in university network*):

Google scholar

- scholar.google.de

ACM Digital Library

- http://dl.acm.org/

DBLP

- http://www.informatik.uni-trier.de/~ley/db/

CiteSeerX

- http://citeseerx.ist.psu.edu/index

NEW REGISTRATION RULE:

Each student is allowed to quit the seminar in the first week.

If one quits, another student from the waiting list can continue his/her work before the first deadline (literature list).

If one quits later than the first week, his/her seminar registration in next semester will be directly moved to the waiting list.

Plan II

Task 3: Written Composition (about 7 weeks later, July 7th)

- Construct fine-grained sub-deadlines with your supervisor
- LNCS[1] Layout, PDF Format required
- Length: 6000–8000 Words (net.) ≅ 15–20 Pages
- Correct and complete bibliography
- Don't copy and paste directly from e.g. papers or web articles! (we will check!)

Task 4: Presentation (July 11th,14th,18th, 21st)

1. LNCS (Lecture Notes in Computer Science) Word/Latex template
   - http://www.springer.com/computer/lncs?SGWID=0-164-6-793341-0

Presentation

a) Length: 60 Minutes
   - ❑ 45 minutes for the presentation
   - ❑ 15 minutes for the discussion

b) Presentation:
   - ❖ Projector (private or one of our notebooks)

c) You must submit your presentation electronically, one week after you gave your talk

How you successfully pass the seminar

- **All certificates will be graded!**
- **"Basic Requirements (i.e. 4.0)":**
  - Meet the deadlines!
  - Don't copy and paste!
  - Decent presentation
  - Be present when others give their talks (must explain reasons in case one cannot attend certain talks)
  - Get involved in the discussions (at least one question)
- **"Criteria for grading"**
  - Quality of your written composition
  - Quality of your presentation (Animate slides! Not text-reading!)
  - Quality of your answers to the questions raised by others in the discussion part of your presentation
  - Be active in discussion
  - Overall impression of your supervisor

Schedule

- April 23th: Kick-off meeting
- May 7th: Deadline for literature list
- May 16th: Deadline for annotated TOC
- July 7th: Deadline for revised written composition
- Final presentations:
  - Friday, July 11th
  - Monday, July 14th
  - Friday, July 18th
  - Monday, July 21st

All deadlines are strict!

Outline

Database	New hardware	Writing

- Transaction, ACID props
- Buffer management
- Recovery
- Indexing
- Query processing:
  - Selection
  - Join
  - Aggregation
  - Data loading
	- Multi-core CPUs
	- GPUs
	- FPGAs
	- Flash memory, SSD
	- A lot of memories
		- Read
		- Understand
		- Write
		- Present
		- Show you‘ve understood and you‘ve written

Thank your for your attention!

Questions before introducing topics

<<Durable transactions for memory using Flash>>

Architecture:

cheap, fast solid-state storage (SSD), inexpensive DRAM, and multi-core CPUs

Database technique covered:

ACID semantics, logging

Proposal:

a new system architecture called Hathi that provides an in-memory transactional heap made persistent using high-speed flash drives.

Supervisor:

Prof. Härder

Topic 2

<<OS paging and buffer managment>>

Architecture:

inexpensive DRAM, traditional HDD

Database technique covered:

data locality

Proposal:

using virtual memory paging mechanism to balance hot/cold data in memory and HDD.

Supervisor:

Prof. Härder

Topic 3

<<Enhancing recovery using an SSD buffer pool extension>>

Architecture:

SSD

Database technique covered:

recovery

Proposal:

exploit the SSD persistence for recovery and normal restart with shorter recovery times.

Supervisor:

Caetano Sauer

Topic 4

<<Energy-proportional query execution using a cluster of wimpy nodes>>

Architecture:

a distributed shared-nothing DBMS

Database technique covered:

ennergy efficiency, load schedule

Proposal:

dynamically adjust clusters for energy and performance

Supervisor:

Daniel Shall

Topic 5

<<Vectorization vs. Compilation in Query Execution>>

Architecture:

modern CPUs

Database technique covered:

query compiler, parallel exeuction using SIMD(single instruction multiple data)/vectorizd processing instructions

Proposal:

compiling SQL queries (Project, Select, Join) directly to executable programs running with SIMD instructions

Supervisor:

Daniel Shall

Topic 6

<<Sort/hash join on modern multi-core CPUs>>

Architecture:

multi-core CPU, Intel Core i7

Database technique covered:

sort-merge/hash joins

Proposal:

optimizing join on multi-core CPUs

Supervisor:

Daniel Shall

Topic 7

<<Scalable aggregation on multi-core CPUs>>

Architecture:

memory

Database technique covered:

in memory parallel aggregation

Proposal:

hybrid independent/shared parallel processing

Supervisor:

Yong Hu

Topic 8

<<Optimizing select conditions on GPUs>>

Architecture:

GPU

Database technique covered:

selection queries

Proposal:

optimizing selection queries on highly parallel GPUs by taking branches into consideration

Supervisor:

Prof. Dessloch

Topic 9

<<Join processing on GPUs >>

Architecture:

GPU

Database technique covered:

Join processing

Proposal:

optimizing joins on GPUs with limited memories (gigabytes)

Supervisor:

Prof. Dessloch

<<Generalized index scan on multi-core CPU/GPU>>

Architecture:

multi-core CPU, GPU

Database technique covered:

Index-based scan

Proposal:

optimizing search using CPUs and GPUs

Supervisor:

Yong Hu

Topic 11

<<Peak performance- remote memory revisited>>

Architecture:

memory

Database technique covered:

storage scalability

Proposal:

utilize memory provided by remote machines for database performance

Supervisor:

Yong Hu

Topic 12

<<In-memory indexing on modern architectures>>

Architecture:

memory

Database technique covered:

Index-based scan

Proposal:

in-memory index different from block-optimized B+-trees

Supervisor:

Weiping Qu

<<Instant loading for main memory databases>>

Architecture:

memory

Database technique covered:

Data loading

Proposal:

fast loading of large data in&out of memory database

Supervisor:

Weiping Qu

<<Data processing on FPGAs>>

Architecture:

many-core FPGAs

Database technique covered:

database operations, e.g. sort

Proposal:

accelerates data processing using FPGA, integrate FPGA with existing architecture

Supervisor:

Weiping Qu

Select your topic now!

1) Durable transactions for memory using Flash
2) OS paging and buffer management
3) Enhancing recovery using an SSD buffer pool extension
4) Energy-proportional Query Execution using a Cluster of Wimpy Nodes
5) Vectorization vs. compilation in query execution
6) Sort/hash join on modern multi-core CPUs
7) Scalable aggregation on multi-core CPUs
8) Optimizing select conditions on GPUs
9) Join processing on GPUs
10) Generalized search tree index scan on multi-core CPU/GPU
11) Peak performance – remote memory revisited
12) In-memory indexing on modern architectures
13) Instant loading for main memory databases
14) Data processing on FPGAs

NEW REGISTRATION RULE:

Each student can decide whether to quit the seminar in the first week.

If one quits, another student from the waiting list can continue his/her work before the first deadline (literature list).

If one quits later than the first week, his/her seminar registration in next semester will be directly moved to the waiting list.