

Immanuel Trummer

Gates Hall, Cornell University
Ithaca (NY), 14850
USA

itrummer@cornell.edu
www.itrummer.org
Revised 3/2017

INTERESTS

My research focuses on optimization problems that arise in the context of big data analytics. In particular, I have studied various generalizations of the classical Query Optimization problem. Those generalizations are necessary in order to accurately model the capabilities of modern query execution platforms. I am also exploring the potential of Quantum Computing for solving analytics related optimization problems. This research branch is based on a grant giving me access to a D-Wave 2X adiabatic quantum annealer. Beyond optimization, I am interested in Text Mining and machine learning.

EDUCATION

Ecole Polytechnique Fédérale de Lausanne (EPFL) 2010-2016
PhD in Computer Science
Advisor: Christoph Koch

University of Stuttgart & Ecole Centrale de Nantes 2003-2010
Double Diploma in Computer Science & Engineering
Obtained with Distinction - Ranked among top five students

AWARDS & HONORS

- Google Faculty Research Award in Database Management
- Publication selected by CACM as CACM Research Highlight
- Selected for ACM SIGMOD Research Highlight Award 2015
- Invitation to publish in “Best of VLDB 2015” (VLDB Journal)
- Google European PhD Fellowship in structured data analysis
- USRA grant for accessing a quantum annealer (machine price: \$15 million)
- EPFL IC Teaching Assistant Award 2015
- Scholarship of the German National Academic Foundation
- First graduation prize by the Computer Science Forum Stuttgart
- Scholarship for Academic Excellence by the University of Stuttgart
- Scholarship for the TIME double degree program
- 2nd prize at German national music competition

RESEARCH EXPERIENCE

Assistant Professor Start: 2016
Cornell University, Computer Science Department Ithaca, USA

Graduate Student Researcher 2010-2016
EPFL, DATA Lab (until 2013: AI Lab) Lausanne, Switzerland

Traditional query optimization, multi-objective query optimization, multi-objective parametric query optimization, and probably approximately optimal query optimization:

- I unified the research branches of parametric query optimization and of multi-objective query optimization by introducing the MULTI-OBJECTIVE PARAMETRIC QUERY OPTIMIZATION problem. The corresponding paper was invited for publication as ACM SIGMOD Research Highlight.
- I developed a decomposition method for the query optimization problem that allows to solve it by MASSIVE PARALLELIZATION using large clusters with hundreds of nodes.
- I developed APPROXIMATION SCHEMES for multi-objective query optimization allowing to gradually trade optimization time for query plan optimality guarantees.
- I developed an INCREMENTAL ALGORITHM for multi-objective query optimization that enables users to find their preferred cost tradeoff in an interactive process.

- I developed a REDUCTION from query optimization to mixed integer linear programming allowing to leverage integer programming solver implementations for query optimization.
- I developed a RANDOMIZED ALGORITHM for multi-objective query optimization that is tailored to this problem and handles significantly larger queries than prior approaches.
- I introduced PROBABLY APPROXIMATELY OPTIMAL query optimization which models situations in which query optimizers need to estimate predicate selectivity via sampling.

Machine learning and text mining:

- In collaboration with researchers from Google Mountain View, I have developed a system that mines SUBJECTIVE ENTITY-PROPERTY ASSOCIATIONS from Web text at a very large scale. This system learns entity type and property specific user behavior models in an unsupervised manner and exploits them to interpret collected text fragments more reliably.

Quantum computing:

- I am experimenting with a D-Wave 2X adiabatic quantum annealer, located at NASA Ames Research Center in California. I evaluate the long-term potential of QUANTUM COMPUTING for solving analytics related optimization problems.
- I have shown that the MULTIPLE QUERY OPTIMIZATION problem can be solved on the D-Wave quantum annealer with speedups of up to three orders of magnitude compared to traditional approaches.

Undergraduate Researcher

University of Stuttgart

2003-2010

Stuttgart, Germany

Cloud computing:

- I developed a method for finding OPTIMAL PROVISIONING strategies for Cloud applications.

PUBLICATIONS

Journal Articles

- Immanuel Trummer, Christoph Koch.
Multi-objective parametric query optimization.
Published as *ACM SIGMOD Research Highlight 2015*.
- Immanuel Trummer, Christoph Koch.
Multi-objective parametric query optimization.
“Best of VLDB 2015” (*VLDB Journal*).
- Immanuel Trummer, Boi Faltings, Walter Binder.
Multi-objective quality-driven service selection –
A fully polynomial time approximation scheme.
TSE, 2014,40(2):167-191.

Conference Publications

- Immanuel Trummer, Christoph Koch.
A fast randomized algorithm for multi-objective query optimization.
SIGMOD 2016.
- Immanuel Trummer, Christoph Koch.
Multiple query optimization on the D-Wave 2X adiabatic quantum computer.
VLDB 2016.
- Immanuel Trummer, Christoph Koch.
Parallelizing query optimization on shared-nothing architectures.
VLDB 2016.
- Immanuel Trummer, Christoph Koch.
An incremental anytime algorithm for multi-objective query optimization.
Talk Recording: <https://www.youtube.com/watch?v=J54gVI9UAo>
SIGMOD 2015.

- Immanuel Trummer, Alon Halevy, Hongrae Lee, Sunita Sarawagi, Rahul Gupta.
Mining subjective properties on the Web.
Talk Recording: <https://www.youtube.com/watch?v=a9RYBydQRXA>
SIGMOD 2015.
- Immanuel Trummer, Christoph Koch.
Multi-objective parametric query optimization.
Talk Recording: <https://www.youtube.com/watch?v=h03IaSfFtJY>
VLDB 2015.
INVITED TO “BEST OF VLDB 2015” VLDBJ ISSUE.
SELECTED AS ACM SIGMOD RESEARCH HIGHLIGHT.
- Immanuel Trummer, Christoph Koch.
Approximation schemes for many-objective query optimization.
SIGMOD 2014.
- Mehdi Riahi, Thanasis Papaioannou, Karl Aberer, Immanuel Trummer.
Utility-driven data acquisition in participatory sensing.
EDBT 2013.
- Immanuel Trummer, Boi Faltings.
Optimizing the tradeoff between discovery, composition,
and execution cost in service composition.
ICWS 2011.
- Immanuel Trummer, Boi Faltings.
Dynamically selecting composition algorithms
for economical composition as a service.
ICSOC 2011.
- Immanuel Trummer, Frank Leymann, Ralph Mietzner, Walter Binder.
Cost-optimal outsourcing of applications into the clouds.
CloudCom 2010.

Patents

- Immanuel Trummer, Boi Faltings.
A method for multi-objective quality-driven service selection.
US Patent Application 13/670,864 (US) 2012.

Theses

- Immanuel Trummer.
From massive parallelization to quantum computing:
seven novel approaches to query optimization.
PhD Thesis, 2016.
- Immanuel Trummer.
Cost-optimal provisioning of cloud applications.
Diploma Thesis, 2010.

Technical Reports

- Immanuel Trummer, Christoph Koch.
Solving the join ordering problem via mixed integer linear programming.
<http://arxiv.org/pdf/1511.02071v1.pdf>, 2015.
- Immanuel Trummer, Christoph Koch.
Probably approximately optimal query optimization.
<http://arxiv.org/pdf/1511.01782v1.pdf>, 2015.

TEACHING & MENTORING

Lecturer, Introduction to Database Systems Spring 2017
This course introduces Cornell students to database management systems, covering topics such as SQL and relational algebra, efficient processing of relational operators, query optimization, concurrency control, recovery, and NoSQL databases.

Lecturer, Advanced Database Systems Fall 2016
This course covers advanced topics around database management systems. It involves extensive paper reading and discussions.

Teaching Assistant, Big Data Fall 2014
I mentored four teams of students (eight to ten students per group) throughout their semester course projects. The projects I supervised ranged from the design and development of a system for crowdsourced SQL processing to the computer-based analysis of language shift in a corpus containing newspaper articles spanning nearly 200 years. One of my teams won the course-internal competition for the best project.

I received the EPFL Teaching Assistant Award 2015 in appreciation of my work as a teaching assistant for the Big Data course.

INDUSTRIAL EMPLOYMENT

Google 5/2014-9/2014
Intern, Web Answers & Web Tables Mountain View, USA
I designed and implemented a system for mining subjective property associations from Web text. The resulting system was successfully used to infer billions of entity-property associations from a Web snapshot and led to a publication at SIGMOD 2015. I was supervised by Alon Halevy and Hongrae Lee.

IBM 7/2007-10/2007
Intern, Extreme Blue Program Böblingen, Germany
Within a team of four students, I took part in the design and implementation of a system for RFID-based temperature tracking during DHL transports.

Agricultural Ministry of Mali 4/2006-6/2007
Software Developer Nantes, France
I designed and implemented an IT system (database, Web interface and standalone client) for the storage and treatment of agricultural data. This project was commissioned by the agricultural ministry of Mali to the Junior Enterprise at Ecole Centrale de Nantes.

Alcatel 7/2006-8/2006
Intern, R&D Division Paris, France
I designed and implemented a Web interface for network surveillance.

LANGUAGES

German Native speaker.
English Published and taught in English for six years.
French Eight years of studies in francophone countries; taught courses in French.

References are available upon request.