University of Washington Database Group

Newsletter - Summer 2017

UW Database Group

We perform research in various aspects of data management systems, theory, and applications. Be sure to visit our website for news and upcoming events!

PAUL G. ALLEN SCHOOL
OF COMPUTER SCIENCE & ENGINEERING

We are excited to announce that we are now part of the Paul G. Allen School of Computer Science and Engineering

EVENTS

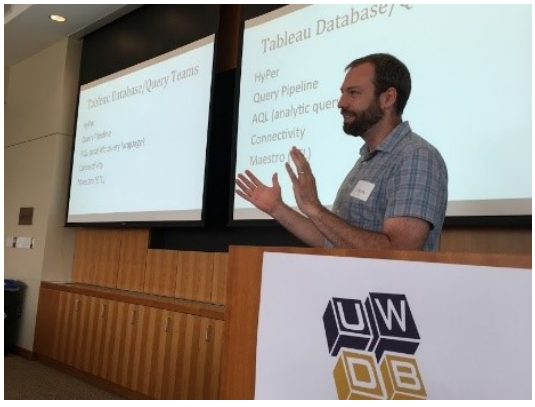

UWDB organized its first workshop for the affiliate members this June on database usability. The workshop was a great success, with presentations and open exchanges from member companies, students from the group, and other invited speakers around campus. Shown in the picture: top row - Justin Talbot, Senior Research Scientist, Tableau and Bolin Ding, Researcher, Data Management, Exploration and Mining, Microsoft Research ; bottom row: Sean Mooney, Chief Research Information Officer, SOM, Biomedical Informatics and Medical Education , Doug Brown, Database Software Architect, Teradata.

AWARDS, HONORS & RECOGNITION

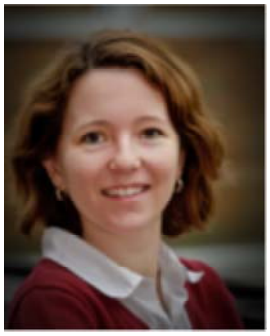

- Magdalena Balazinska was awarded the ACM SIGMOD Test of Time Award : Magdalena Balazinska, Hari Balakrishnan, Samuel Madden, Michael Stonebraker. Fault-tolerance in the Borealis distributed stream processing system. SIGMOD Conference 2005: 13-24.
- Congratulations also for her appointment as a member of the NSF CISE Advisory Council and recently, as Director of the University of Washington's eScience Institute.

eScience Institute
ADVANCING DATA-INTENSIVE DISCOVERY IN ALL FIELDS

Alvin Cheung won a National Science Foundation CAREER Award for "Generating Application-Specific Database Management Systems" in which he aims to automate the process of domain specialization in database management systems (DBMSs). In this project, Alvin will leverage recent advances in programming systems and data management to build tools that automatically understand database application semantics — streamlining what is currently a complex and error-prone process to improve a variety of applications that rely on DBMSs, from banking, to social media, to scientific discovery.

CONGRATULATIONS TO OUR STUDENTS FOR WINNING AWARDS AT SIGMOD 2017 !!!

- Shumo Chu, Daniel Li, and Chenglong Wang won the Best Demonstration Award for Demonstration of the Cosette Automated SQL Prover
- Maaz Ahmad earned a Demonstration Award Honorable Mention for Optimizing Data-Intensive Applications Automatically By Leveraging Parallel Data Processing Frameworks;

- Brandon Haynes and Artem Minyaylov earned a Demonstration Award Honorable Mention for VisualCloud Demonstration: A DBMS for Virtual Reality.
- Jennifer Ortiz won second place at the SIGMOD Student Reseach Competition.

Please see below for details about the work on the Research Highlights below.

RESEARCH HIGHLIGHTS

I. CONFERENCE PAPERS

Bas Ketsman, Dan Suciu**: A Worst-Case Optimal Multi-Round Algorithm for Parallel Computation of Conjunctive Queries. PODS 2017: 417-428**

Abstract:

We study the optimal communication cost for computing a full conjunctive query Q over p distributed servers. Two prior results were known. First, for one-round algorithms over skew-free data the optimal communication cost per server is $m/p1/\tau * (Q)$ where m is the size of the largest input relation, and τ is the fractional vertex covering number of the query hypergraph. Second, for multi-round algorithms and unrestricted database instances, it was shown that any algorithm requires at least $m/p1/ p * (Q)$ communication cost per server, where $\rho*(Q)$ is the fractional edge covering number of the query hypergraph; but no matching algorithms were known for this case (except for two restricted queries: chains and cycles).

In this paper we describe a multi-round algorithm that

computes any query with load $m/p^{1/\rho^*(Q)}$ per server, in the case when all input relations are binary. Thus, we prove this to be the optimal load for all queries over binary input relations. Our algorithm represents a non-trivial extension of previous algorithms for chains and cycles, and exploits some unique properties of graphs, which no longer hold for hyper-graphs.

Project website: https://homes.cs.washington.edu/~suciu/bas-pods-2017.pdf

---

Definition A Probabilistic Database is (**W**, P), where **W** is an incomplete database, and P: **W** → [0,1] a probability distribution: $\sum_{i=1,n} P(W_i) = 1$

Laurel Orr, Dan Suciu, Magdalena Balazinska: **Probabilistic Database Summarization for Interactive Data Exploration. PVLDB 10(10): 1154-1165 (2017)**

Abstract:

We present a probabilistic approach to generate a small, query-able summary of a dataset for interactive data exploration. Departing from traditional summarization techniques, we use the Principle of Maximum Entropy to generate a probabilistic representation of the data that can be used to give approximate query answers. We develop the theoretical framework and formulation of our probabilistic representation and show how to use it to answer queries. We then present solving techniques and give three critical optimizations to improve preprocessing time and query accuracy. Lastly, we experimentally evaluate our work using a 5 GB dataset of flights within the United States and a 210 GB dataset from an astronomy particle simulation. While our current work only supports linear queries, we show that our technique can successfully answer queries faster than sampling while introducing, on average, no more error than sampling and can better distinguish between rare and nonexistent values.

Project website: http://db.cs.washington.edu/projects/entropydb/

Parmita Mehta, Sven Dorkenwald, Dongfang Zhao, Tmer Kaftan, Alvin Cheung, Magdalena Balazinska, Ariel Rokem, Andrew Connolly, Jacob Vanderplas, and Yusra AlSayyad: **Comparative Evaluation of Big-Data Systems on Scientific Image Analytics Workload. PVLDB 2017 Volume 10 Issue 11.**

Abstract:

Scientic discoveries are increasingly driven by analyzing large volumes of image data. Many new libraries and specialized database management systems (DBMSs) have emerged to support such tasks. It is unclear how well these systems support real-world image analysis use cases, and how performant the image analytics tasks implemented on top of such systems are. In this paper, we present the first comprehensive evaluation of large-scale image analysis systems using two real-world scientic image data processing use cases. We evaluate five representative systems (SciDB, Myria, Spark, Dask, and TensorFlow) and find that each of them has shortcomings that complicate implementation or hurt performance. Such shortcomings lead to new research opportunities in making large-scale image analysis both efficient and easy to use.

Project website: http://homes.cs.washington.edu/%7Emagda/papers/mehta-vldb17.pdf

---

Chenglong Wang, Alvin Cheung, Rastislav Bodík: **Synthesizing highly expressive SQL queries from input-output examples. PLDI 2017: 452-466**

Abstract

SQL is the de facto language for manipulating relational data. Though powerful, SQL queries can be difficult to write due to their highly expressive constructs. Using the programming-by-example paradigm to help users write SQL queries presents an attractive proposition, as evidenced by online help forums such as Stack Overflow. However, developing techniques to synthesize SQL queries from input-output (I/O) examples has been difficult due to SQL's rich set of operators. In this paper, we present a new scalable and efficient algorithm to synthesize SQL queries from I/O examples. Our key innovation is the development of a language for abstract queries, i.e., queries with uninstantiated operators, that can express a large space of

SQL queries efficiently. Using abstract queries to represent the search space nicely decomposes the synthesis problem into two tasks: (1) searching for abstract queries that can potentially satisfy the given I/O examples, and (2) instantiating the found abstract queries and ranking the results. We implemented the algorithm in a new tool, called SCYTHE, and evaluated it on 193 benchmarks collected from Stack Overflow. Our results showed that SCYTHE efficiently solved 74% of the benchmarks, most in just a few seconds. Queries synthesized by SCYTHE range from simple ones involving a single selection to complex ones with six levels of nested queries.
Project website: https://scythe.cs.washington.edu/

Shumo Chu, Konstantin Weitz, Alvin Cheung, Dan Suciu: **HoTTSQL: proving query rewrites with univalent SQL semantics. PLDI 2017**

Abstract:

Every database system contains a query optimizer that performs query rewrites. Unfortunately, developing query optimizers remains a highly challenging task. Part of the challenges comes from the intricacies and rich features of query languages, which makes reasoning about rewrite rules difficult. In this paper, we propose a machine-checkable denotational semantics for SQL, the de facto language for relational database, for rigorously validating rewrite rules. Unlike previously proposed semantics that are either non-mechanized or only cover a small amount of SQL language features, our semantics covers all major features of SQL, including bags, correlated subqueries, aggregation, and indexes. Our mechanized semantics, called HoTT SQL, is based on K-Relations and homotopy type theory, where we denote relations as mathematical functions from tuples to univalent types. We have implemented HoTTSQL in Coq, which takes only fewer than 300 lines of code and have proved a wide range of SQL rewrite rules, including those from database

research literature (e.g., magic set rewrites) and real-world query optimizers (e.g., subquery elimination). Several of these rewrite rules have never been previously proven correct. In addition, while query equivalence is generally undecidable, we have implemented an automated decision procedure using HoTTSQL for conjunctive queries: a well-studied decidable fragment of SQL that encompasses many real-world queries.
Project website: http://cosette.cs.washington.edu/

Shumo Chu, Chenglong Wang, Konstantin Weitz, Alvin Cheung: **Cosette: An Automated Prover for SQL. CIDR 2017**

Abstract:

Deciding query equivalence is an important problem in data management with many practical applications. Solving the problem, however, is not an easy task. While there has been a lot of workdone in the database research community in reasoning about the semantic equivalence of SQL queries, prior work mainly focuses on theoretical limitations. In this paper, we present Cosette, a fully automated prover that can determine the equivalence of SQL queries. Cosette leverages recent advances in both automated constraint solving and interactive theorem proving, and returns a counterexample (in terms of input relations) if two queries are not equivalent, or a proof of equivalence otherwise. Although the problem of determining equivalence for arbitrary SQL queries is undecidable, our experiments show that Cosette can determine the equivalences of a wide range of queries that arise in practice, including conjunctive queries, correlated queries, queries with outer joins, and queries with aggregates. Using Cosette, we have also proved the validity of magic set rewrites, and confirmed various real-world query rewrite errors, including the famous COUNT bug. We are unaware of any prior tool that can automatically determine the equivalences of a broad range of queries as Cosette, and believe that our tool represents a major step towards building provably-correct query optimizers for real-

world database systems

Project website: http://cosette.cs.washington.edu/

Jingjing Wang, Tobin Baker**,** Magdalena Balazinska, Daniel Halperin, Brandon Haynes, Bill Howe, Dylan Hutchison, Shrainik Jain, Ryan Maas, Parmita Mehta, Dominik Moritz, Brandon Myers, Jennifer Ortiz, Dan Suciu, Andrew Whitaker, Shengliang Xu: **The Myria Big Data Management and Analytics System and Cloud Services. CIDR 2017.**

Abstract:

In this paper, we present an overview of the Myria stack for big data management and analytics that we developed in the database group at the University of Washington and that we have been operating as a cloud service aimed at domain scientists around the UW campus. We highlight Myria's key design choices and innovations and report on our experience with using Myria for various data science use-cases.

Project website: http://myria.cs.washington.edu/

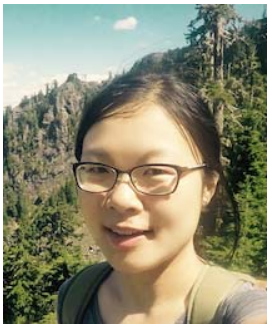

Jingjing Wang, Magdalena Balazinska: **Elastic Memory Management for Cloud Data Analytics. USENIX Annual Technical Conference 2017: 745-758**

Abstract:

We develop an approach for the automatic and elastic management of memory in shared clusters executing data analytics applications. Our approach, called ElasticMem, comprises a technique for dynamically changing memory limits in Java virtual machines, models to predict memory usage and garbage collection cost, and a scheduling algorithm that dynamically reallocates memory between applications. Experiments with our prototype implementation show that our approach outperforms static memory allocation leading to fewer

query failures when memory is scarce, up to 80% lower garbage collection overheads, and up to 30% lower query times when memory is abundant.

Project website:

https://homes.cs.washington.edu/~jwang/publications/elasticmem-atc.pdf

II. DEMOS

Shumo Chu, Daniel Li, Chenglong Wang, Alvin Cheung, Dan Suciu: Demonstration of the Cosette Automated SQL Prover. SIGMOD Conference 2017 (Winner of the Best Demo Award)

Abstract:

In this demonstration, we showcase Cosette, the first automated prover for determining the equivalences of SQL queries. Despite theoretical limitations, Cosette leverages recent advances in both automated constraint solving and interactive theorem proving to decide the equivalences of a wide range of real world queries, including complex rewrite rules from the database literature. Cosette can also validate the inequality of queries by finding counter examples, i.e., database instances which, when executed on the two queries, will return different results. Cosette can find counter examples of many real world inequivalent queries Including a number of real-world optimizer bugs. We showcase three representative applications of Cosette: proving a query rewrite rule from magic set rewrite, finding counter examples from the infamous optimizer bug, and an interactive visualization of automated

grading results powered by Cosette, where Cosette is used to check the equivalence of students' answers to the standard solution. For the demo, the audience can experience through the three applications, and explore the Cosette by interacting with the tool using an easy-to-use web interface.

Project website http://cosette.cs.washington.edu/

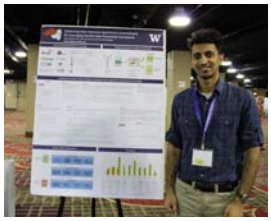

Maaz Bin Safeer Ahmad, Alvin Cheung:
Optimizing Data-Intensive Applications Automatically By Leveraging Parallel Data Processing Frameworks. SIGMOD Conference 2017: 1675-1678
(Honorable mention for the Best Demo Award)

Abstract:

In this demonstration we will showcase CASPER, a novel tool that enables sequential data-intensive programs to automatically leverage the optimizations provided by parallel data processing frameworks. The goal of CASPER is to reduce the inertia against adaptation of new data processing frameworks---particularly for non-expert users---by automatically re-writing sequential programs written in general purpose languages to the high-level DSLs or APIs of these frameworks. Through CASPER's browser-based interface, users can enter the source code of their Java applications and have it automatically retargeted to execute on Apache Spark. In our interactive presentation, we will use CASPER to optimize sequential implementations of data visualization programs as well as image processing kernels. The optimized Spark implementations along with the original sequential implementations will then be executed simultaneously on the cloud to allow the demo to audience compare the runtime performances and outputs in real-time.

Project website: http://casper.uwplse.org/

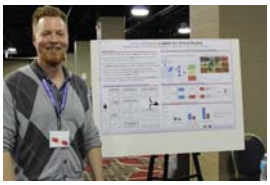

Brandon Haynes, Artem Minyaylov, Magdalena Balazinska, Luis Ceze, Alvin Cheung:

VisualCloud Demonstration: A DBMS for Virtual Reality. SIGMOD Conference 2017: 1615-1618

(Honorable mention for the Best Demo Award)

Abstract:

We demonstrate VisualCloud, a database management system designed to efficiently ingest, store, and deliver virtual reality (VR) content at scale. VisualCloud targets both live and prerecorded spherical panoramic (a.k.a.360) VR videos. It persists content as a multidimensional array that utilizes both dense (e.g., space and time) and sparse (e.g., bit and frame rate) dimensions. VisualCloud uses orientation prediction to reduce data transfer by degrading out-of-view portions of the video. Content delivered through VisualCloud requires 40% less bandwidth than existing methods and scales to many concurrent connections. This demonstration will allow attendees to view both live and prerecorded VR video content served through VisualCloud. Viewers will be able to dynamically adjust tuning parameters (e.g., bitrates and path prediction) and observe changes in visual fidelity.

Project website: http://db.cs.washington.edu/projects/visualcloud/

Helga Gudmundsdottir**,** Babak Salimi, Magdalena Balazinska, Dan R. K. Ports, Dan Suciu:

A Demonstration of Interactive Analysis of Performance Measurements with Viska. SIGMOD Conference 2017: 1707-1710

Abstract:

The ultimate goal of system performance analysis is to identify the underlying causes for performance differences between different systems and different workloads. We make this goal easier to achieve with Viska, a new tool for generating and interpreting

performance measurement results. Viska leverages cutting-edge techniques from big data analytics and data visualization to aid and automate this analysis, and helps users derive meaningful and statistically sound conclusions using state-of-the-art causal inference and hypothesis testing techniques.

Project website:

http://homes.cs.washington.edu/~magda/papers/helga-sigmod17-demo.pdf

III. WORKSHOP PAPERS

Dylan Hutchison, Bill Howe, Dan Suciu: **LaraDB: A Minimalist Kernel for Linear and Relational Algebra Computation. BeyondMR@SIGMOD 2017: 2:1-2:10**

Abstract:

Analytics tasks manipulate structured data with variants of relational algebra (RA) and quantitative data with variants of linear algebra (LA). The two computational models have overlapping expressiveness, motivating a common programming model that affords unified reasoning and algorithm design. At the logical level we propose LARA, a lean algebra of three operators that expresses RA and LA as well as relevant optimization rules. We show a series of proofs that position LARA at just the right level of expressiveness for a middleware algebra: more explicit than MapReduce but more general than RA or LA. At the physical level we find that the LARA operators afford efficient implementations using a single primitive that is available in a variety of backend engines: range scans over partitioned sorted maps.

To evaluate these ideas, we implemented the LARA operators as range iterators in Apache Accumulo, a popular implementation of Google's BigTable. First we show how LARA expresses a sensor quality control task, and we measure the

performance impact of optimizations LARA admits on this task. Second we show that the LARADB implementation outperforms Accumulo's native MapReduce integration on a core task involving join and aggregation in the form of matrix multiply, especially at smaller scales that are typically a poor fit for scale-out approaches. We find that LARADB offers a conceptually lean framework for optimizing mixed-abstraction analytics tasks, without giving up fast record-level updates and scans.
Project website:https://arxiv.org/abs/1703.07342

IV. INVITED TALKS

Alvin Cheung:

- Cosette: A Solver for SQL Equivalences
  In January 2017, the first RDP (Reasoning about Declarative Programs) Workshop was held as part of the POPL 17 conference. The workshop was attended by researchers from both programming languages and database research. Alvin gave an invited talk on Cosette, and discussed how programming languages techniques can be applied to solve problems in data management research, along with new challenges in both research areas**.**

Magda Balazinska:

- Intel Tech Talk (online). Title: **"Comparative Evaluation of Big-Data Systems on Scientific Image Analytics Workloads**." June 2017. Presentation slides here.
- Teradata. Title: **"Performance SLAs for Cloud Data Analytics**". April 2017. Presentation slides here.
- Keynote talk at the ICDE 2017 PhD Symposium. Title: "Research with Real Users". April 2017. Presentation slides here.

- Georgia Tech. Title: "The Myria Big Data Management System and Cloud Service". Mary Jean Harrold Memorial Lecture. **Distinguished Lecture** in the School of Computer Science. November 2016. Presentation slides here.
- University of Buffalo. Title: "The Myria Big Data Management System and Cloud Service" Computer Science **Department Colloquium**. November 2016. Presentation slides here.

Dan Suciu:

- Dan Suciu: **Communication Cost in Parallel Query Evaluation: A Tutorial. PODS 2017: 319.** The slides are available here and the video recording is available here.

  Project website: http://homes.cs.washington.edu/~suciu/project-aitf.html

- An invited talk at the 2017 Alberto Mendelzon Workshop, where he talked about Shannon-type inequalities for optimal query processing. The talk explained among other things how the recent concept of submodular width of a query can be used for optimal query processing

UWDB GROUP FUN

The group celebrated the end of the school year with a party at Magda's place in Hawthorne Hills and the 4th of July fireworks at Dan's place in South Lake Union. Warm wishes from UWDB and see you in Fall Quarter!

Twitter Website Facebook

Computer Science & Engineering at the University of Washington is consistently ranked among the top programs in the nation. We educate tomorrow's innovators, conduct cutting-edge research in the principal areas of the field, lead a broad range of multi-disciplinary initiatives that demonstrate the transformative power of computer science and computer engineering, and are widely recognized for our success in promoting diversity.

The Paul Allen Center for Computer Science & Engineering

Our home page is:
http://db.cs.washington.edu/

Our people are:
http://db.cs.washington.edu/people.html