SIGMOD 2017
Teaser Talks

Wednesday - May 17th 2017

Controlling False Discoveries During Interactive Data Exploration

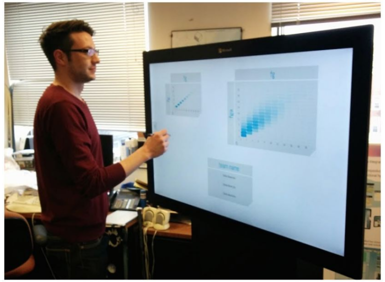

Problem:

Visual analytics & data exploration are prone to false discovery.

Challenge:

Infer hypotheses from visualizations automatically.

Control false discovery interactively with theoretical guarantee.

System:

QUDE: Quantifying Uncertainties in Data Exploration.

Demo: *Safe Visual Data Exploration*

Session 11, Continental B, 11am.

Z Zhao, L De Stefani, E Zgraggen, C Binnig, E Upfal and T Kraska

Computer Science Department, Brown Univesrity

MacroBase: Prioritizing Attention in Fast Data

Peter Bailis; Edward Gan; Samuel Madden; Deepak Narayanan;
Kexin Rong; Sahaana Suri

Problem:

Massive fast data streams :: i) relational analytics not enough
ii) no ML systems for automatically reducing streams at scale

Solution:

New engine combining streaming classification + explanation
This paper: architecture, unsupervised estimation, sketching
Exciting results from production and open source use

This is the next major challenge for dataflow-based analytics

Data Canopy: Accelerating Exploratory Statistical Analysis

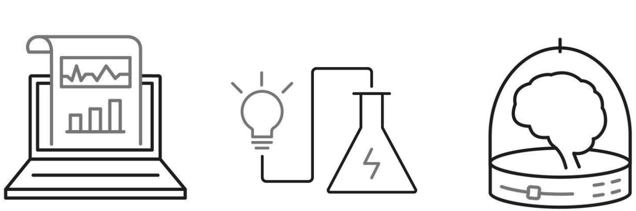

Statistics are everywhere!

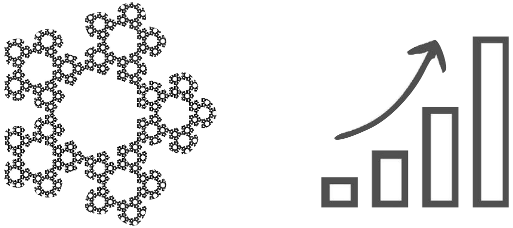

Repetitive statistics and data access

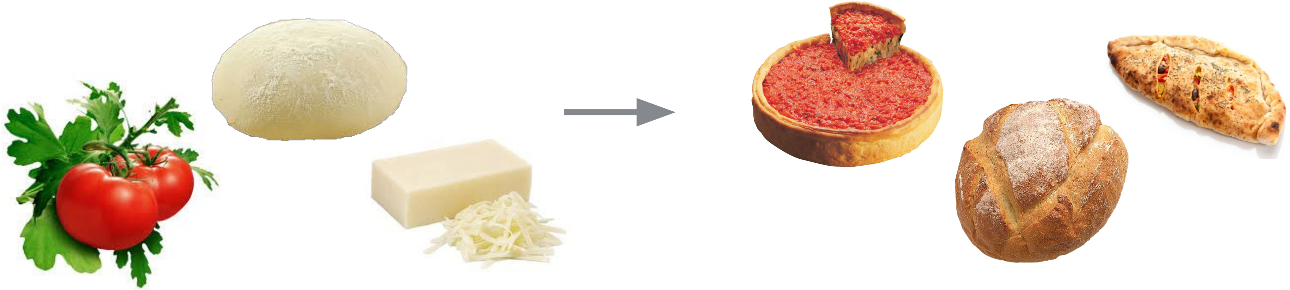

Data Canopy synthesizes statistics from basic ingredients

Abdul Wasay | Xinding Wei | Niv Dayan | Stratos Idreos

Beliefs...

Beta Probabilistic Databases

A Scalable Approach to Belief Updating and Parameter Learning

Niccolo' Meneghetti   Oliver Kennedy   Wolfgang Gatterbauer

evidence

Do Ada and Bob work for the same company?
yes yes yes no yes no no …
Does Ada have a job?
no yes yes no no yes …
Does Bob work for HP?
no no no no yes no no …
…

model

name	emp	*tid*	a	b	$\langle\boldsymbol{\theta}\rangle$	$h[\boldsymbol{\theta}]$
Ada	HP	x_1	6	4	.6	−0.507
Ada	IBM	x_2	18	12	.6	−1.014
Bob	HP	x_3	2	2	.5	−0.125

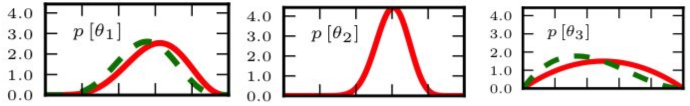

Problem: training a probabilistic model from sampled answers to Boolean queries

Main Contribution: a novel, tuple-independent probabilistic DB with Beta priors

Database Learning: Toward a Database that Becomes Smarter Every Time

Beliefs...

Yongjoo Park, Shahab Tajik, Michael Cafarella, Barzan Mozafari U.Michigan, Ann Arbor

11AM | Continental C

Previous

40 Years of DB research: Repetition of "Query → Answer"

Our Work

DB Learning: More queries → Fast query processing
(**ML:** More observation → More Accurate)

The first work that becomes faster every time by exploiting the answers to past queries

Beliefs...

Staging User Feedback toward Rapid Conflict Resolution in Data Fusion

Romila Pradhan Siarhei Bykau Sunil Prabhakar

Data Item	S_1	S_2	S_3	S_4
Zootopia		Howard	Spencer	Spencer
Kung Fu Panda	Stevenson		Nelson	
Inside Out		leFauve	Docter	
Finding Dory				Stanton
Minions	Coffin	Renaud		
Rio	Jones		Saldanha	

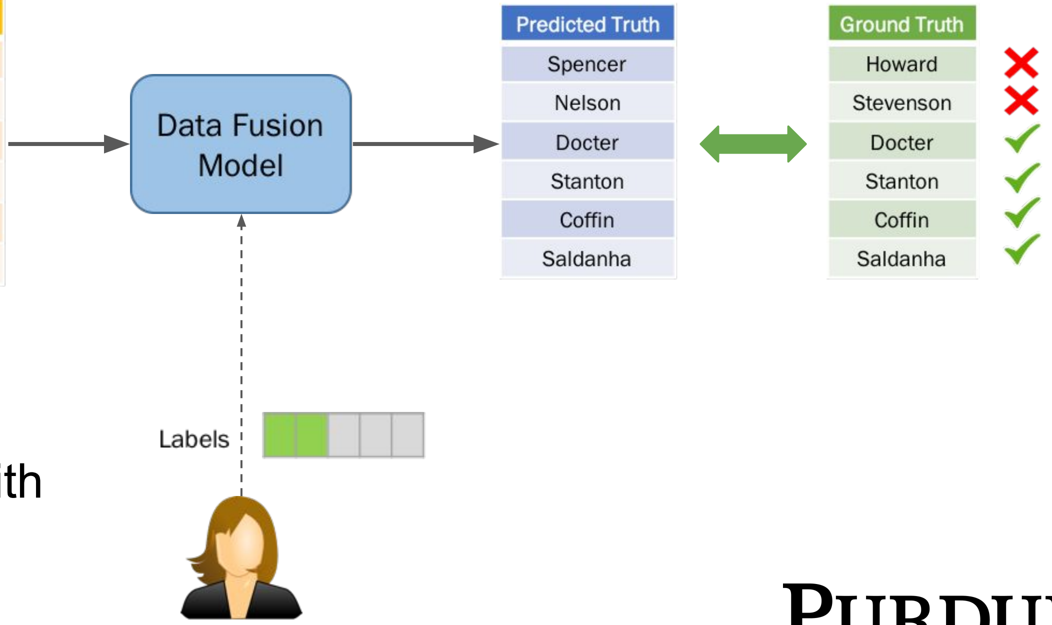

How to be most effective with user feedback?

PURDUE UNIVERSITY

Social Network

Discovering Your Selling Points: Personalized Social Influential Tag Exploration

Yuchen Li; Ju Fan; Dongxiang Zhang; Kian-Lee Tan

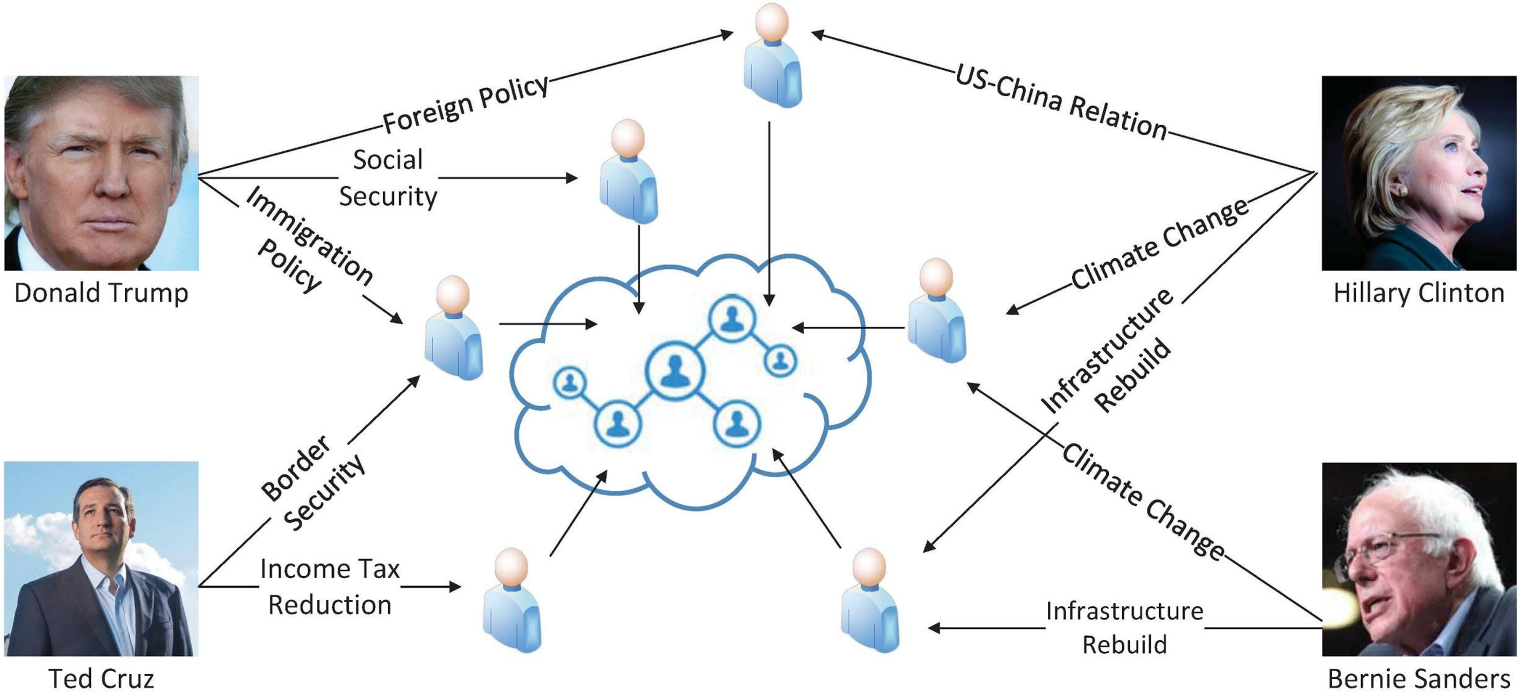

Social Network

Coarsening Massive Influence Networks

for Scalable Diffusion Analysis

Naoto Ohsaka, Tomohiro Sonobe, Sumio Fujita, Ken-ichi Kawarabayashi

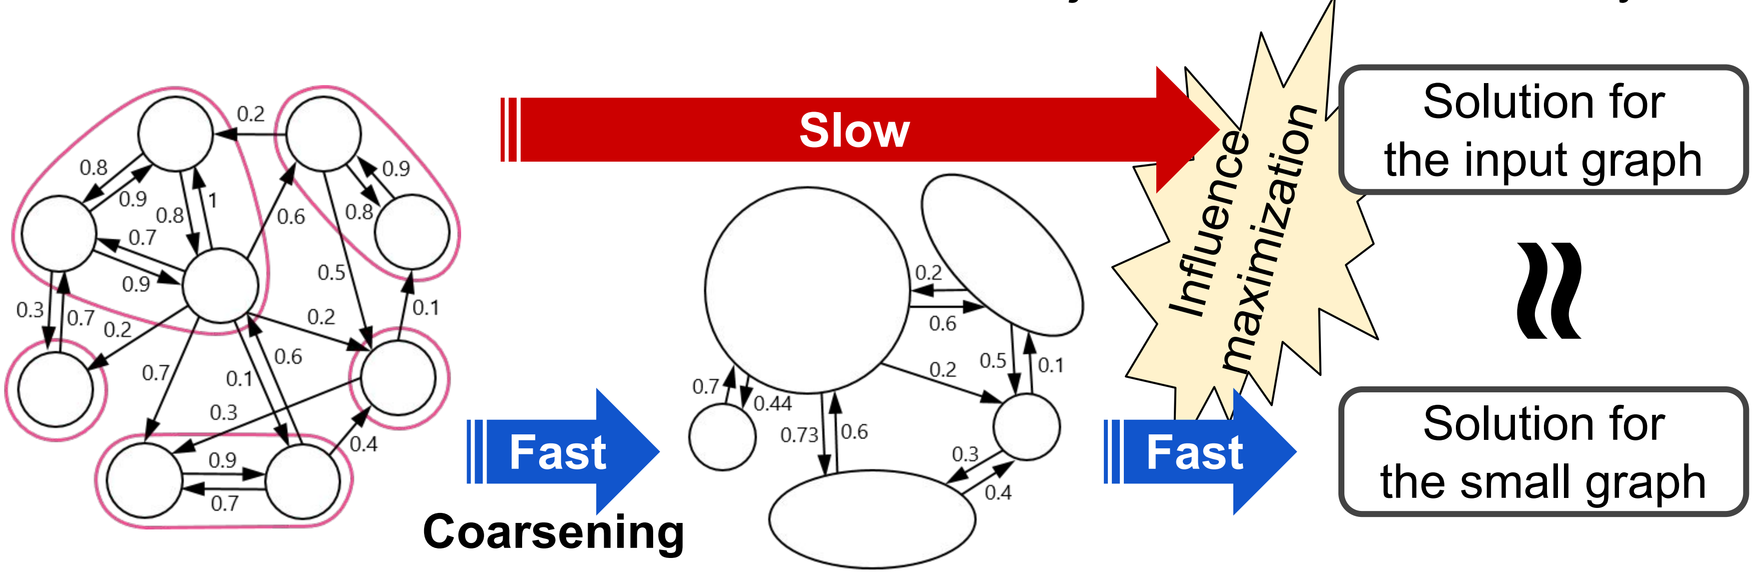

We propose

reduction strategy, scalable algorithm, analysis framework

Accuracy guarantee 1 hour for billion edges 2—30× faster

Debunking the *Myths* of Influence Maximization

Social Network

Reveal and Rectify

- *Myths* and *Mis-claims*
- *Bugs* and *undocumented assumptions*

Decision Tree

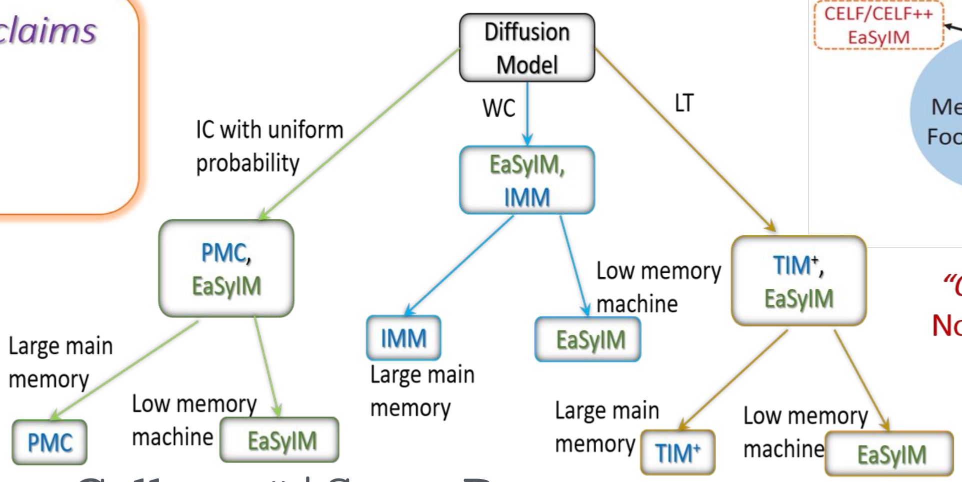

Categorization

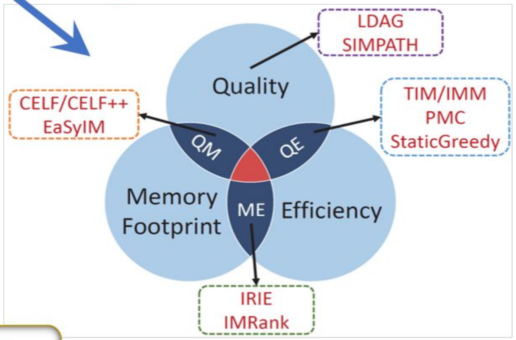

"One Size Doesn't Fit All"
No Single State of the Art!

Akhil Arora* | Sainyam Galhotra* | Sayan Ranu

Interactive Mapping Specification with Exemplar Tuples

Give me a few tuples, I'll get you a mapping

Source

Company

IdCompany	*Name*	*Town*
'C1'	'AA'	'Paris'
'C2'	'Ev'	'Lyon'

Flight

Departure	*Arrival*	*IdCompany*
'Lyon'	'Paris'	'C1'
'Paris'	'Lyon'	'C2'

Travel Agency

IdAgency	*Name*	*Town*
'A1'	'TC'	'L.A.'

Target

Carrier

Id	*Name*	*Town*
'Id1'	'AA'	'Paris'
'Id2'	'Ev'	'Lyon'
'Id3'	'TC'	'L.A.'

Departure

Town	*IdCarrier*
'Lyon'	'Id1'
'Paris'	'Id2'

Arrival

Town	*IdCarrier*
'Paris'	'Id1'
'Lyon'	'Id2'

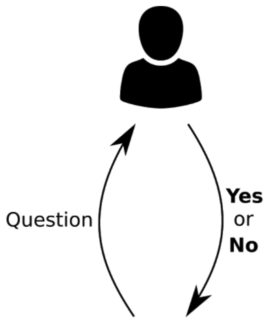

Simple Boolean Interactions

Final mapping

$$\mathbf{m_1} : Company(c1, aa, paris_1) \wedge Flight(lyon, paris_2, c1) \rightarrow \exists id1, Firm(id1, aa, paris_1) \wedge Departure(lyon, id1) \wedge Arrival(paris_2, id1)$$

$$\mathbf{m_2} : TravelAgency(a1, tc, la) \rightarrow \exists id3, Firm(id3, tc, la)$$

A. Bonifati, U. Comignani, E. Coquery, R. Thion

Foofah: Transforming Data By Example

Mappings...

"Mark" Zhongjun Jin, Michael R. Anderson, Michael Cafarella, H. V. Jagadish

Motivation

Current data transformation methods:

1) consume too much user effort

2) require high expertise

3) programming-skills

Results

1. Save about 60% of user effort than state-of-art Programming By Demonstration data transformation tool; requires little expertise.
2. Can handle about 90% real-world data transformation tasks

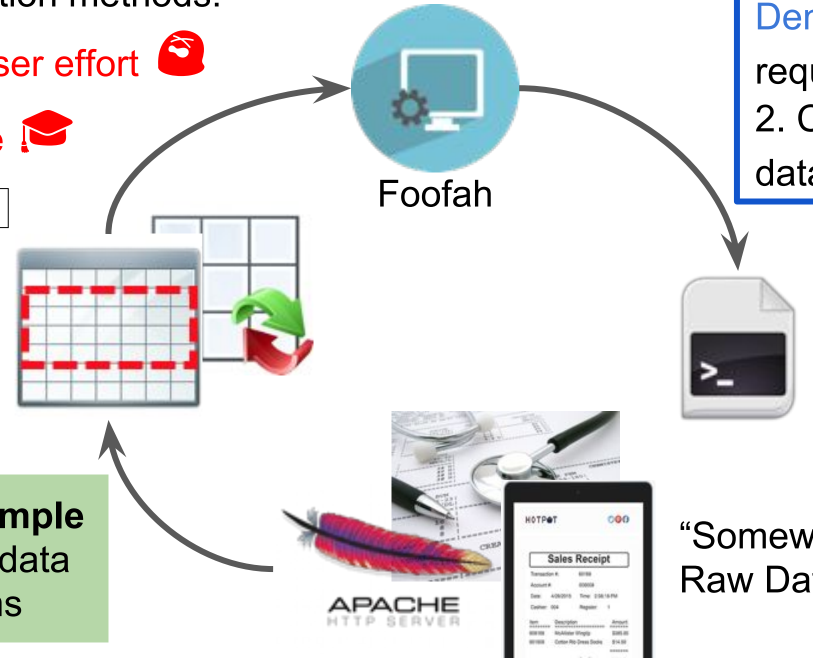

Input-output Example

Synthesized Program

"Somewhat" Structured Raw Data

Programming-By-Example approach synthesizing data transformation programs

UNIVERSITY OF MICHIGAN

Mappings...

QIRANA: A Framework for Scalable Query Pricing

Shaleen Deep, Paris Koutris

University of Wisconsin-Madison

Problem : Design scalable pricing framework

- increased demand for data markets
- need for pricing systems with formal guarantees

Proposed Solution : QIRANA

- scalable
- arbitrage-free data pricing framework
- allows customizability for seller
- history-aware pricing

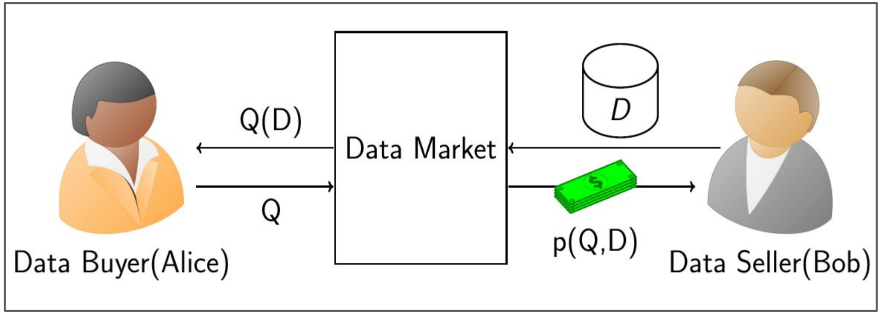

Talk@Session 14, Lake Michigan

Access Path Selection in Main-Memory Optimized Data Systems:

Should I Scan or Should I Probe?

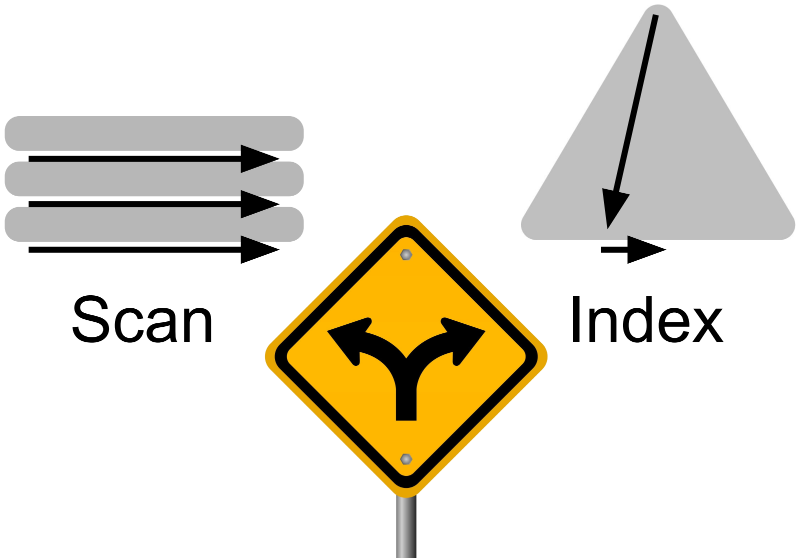

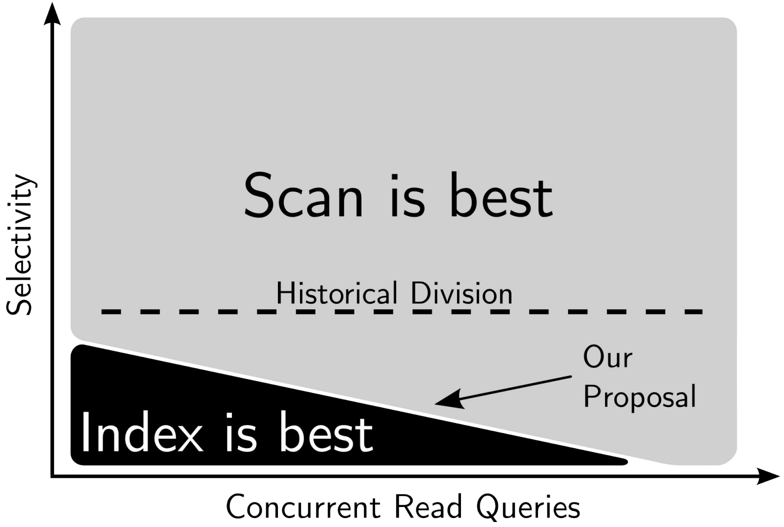

DASlab @ Harvard SEAS

Michael S. Kester Manos Athanassoulis Stratos Idreos

Optimization of Disjunctive Predicates for Main Memory Column Stores

Fisnik Kastrati; Guido Moerkotte

UNIVERSITY OF MANNHEIM

- **Problem: Optimization of disjunctive predicates**
  $(p_{1,1} \wedge \ldots \wedge p_{1,n}) \vee \ldots \vee (p_{m,1} \wedge \ldots \wedge p_{m,n})$
- Current optimization schemes for disjunctive predicates are based on heuristics which produce poor plans
  - Query is transformed into either **DNF** or **CNF** and then optimized
- When optimizing disjunctive predicates, the true optimization potential cannot be achieved by means of traditional plans
  - We can fill this gap by means **Bypass Processing**
- We propose a new algorithm which exploits Bypass processing
  - Our experiments show an improvement in plan quality by an average factor of over **2000** vs. heuristics used in RDBMSs

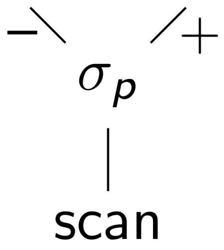

11AM
Lake Erie

A Top-Down Approach to Achieving **Performance Predictability** in Database Systems

Jiamin Huang; Barzan Mozafari; Grant Schoenebeck; Thomas F. Wenisch

Q1. How to automatically **identify** root causes of performance variance in a complex codebase?

Q2. How to make database systems **more predictable but also faster**?

Q3. How our techniques **improved MySQL** and are deployed on **2M+ servers** today?

Teaser Talks (Second Part)

AQP (2)

Two-Level Sampling for Join Size Estimation

Yu Chen; Ke Yi

Hong Kong University of Science and Technology, Hong Kong SAR, China

- **Problem: Estimate join size**
  - **Selection predicates given at query time**
- **Our solution: Sampling**

$$\text{sample}(\sigma_c(R)) = \sigma_c(\text{sample}(R))$$

  - **Two-Level Sampling**
  - **One Pass, Unbiased, Smaller Error**
  - **Beats previous sampling methods**
- **Applications**
  - **Query optimization (join order)**
  - **Approximate query answering (COUNT, SUM, AVG)**

$$\text{Customer} \bowtie \text{Order}$$

$$\Downarrow$$

$$\sigma_{\text{Age}<=35}(Customer) \bowtie \sigma_{\text{TotalPrice}>200}(Order)$$

Customer

Cust key	Name	...	Age
1	Lizabeth	...	41
2	Elliott	...	65
3	Helga	...	20
4	Parker	...	47
5	Wilford	...	22

Order

Ord key	Cust key	Total Price
1	2	322
2	4	553
3	5	420
4	3	82
5	3	120
6	1	604
7	2	418

By the end of my talk I hope to convince you to replace all your Bloom filters with a new data structure, the counting quotient filter (CQF).

- Problem: Bloom filters lack features that many applications need.
  - Can’t count, merge, resize, delete, scale out of RAM, etc.
  - Applications are forced to work around the limitations.
- Solution: Counting quotient filter (CQF).
  - Supports counting, merging, scaling out of RAM, etc.
  - Counts skewed input distributions efficiently.
  - Faster and smaller than a Bloom filter.
- Several computational biology and streaming applications teams are already replacing Bloom filters with CQFs in their code.

Prashant Pandey Michael A. Bender Rob Johnson Rob Patro

BePI: Fast and Memory-Efficient Method for Billion-Scale Random Walk with Restart

[1]Jinhong Jung; [1]Namyong Park; [2]Sael Lee; [1]U Kang

[1]Seoul National Univ. and [2]The State Univ. of New York Korea

- **Problem: Random Walk with Restart**
  - RWR measures relevance scores between nodes in graphs
  - How can we compute RWR scores quickly in very large graphs?
- **Proposed method: BePI**
  - Fast and scalable by taking the advantages of both preprocessing and iterative approaches
- **Experimental Results**
  - Process **100x** larger graphs and requires **130x** less memory space than existing preprocessing methods
  - Compute RWR scores up to **9x** faster than its competitors

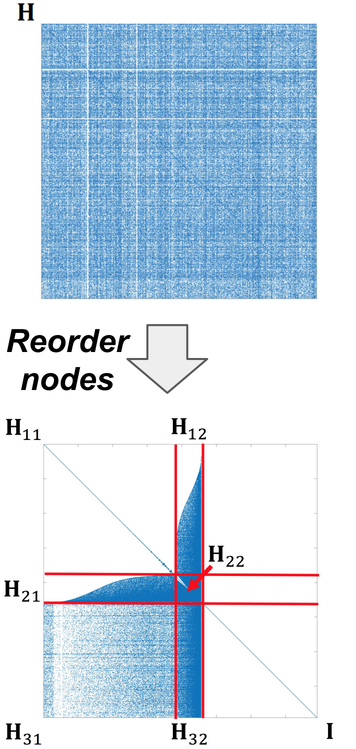

Session 16, Continental B

Determining the Impact Regions of Competing Options in Preference Space

User Pref.

[1] Bo Tang; [2] Kyriakos Mouratidis; [1] Man Lung Yiu
[1] The Hong Kong Polytechnic Univ., [2] Singapore Management Univ.

Who are our target customer?

Hilton HOTELS & RESORTS

How about our market share?

BMW

Where am I most competitive?

kSPR queries

Look-ahead techniques

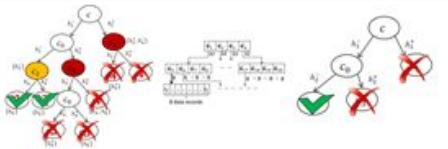

Progressively processing

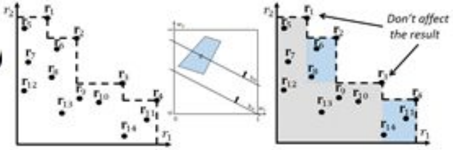

CellTree

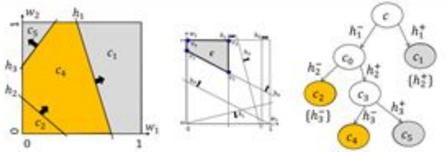

Datasets

Hotels.com MySTATS STATS House Stats CarStats

Impact regions

Customer Profiling

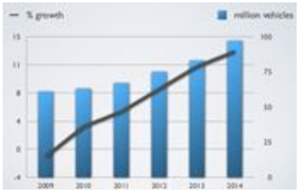

Market Analysis

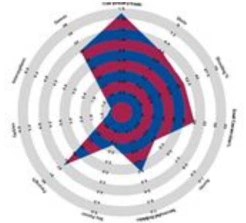

Competitive Analysis

(Session 17, 14:00-14:25, Continental C)

User Pref.

Efficient Computation of Regret-ratio Minimizing Set: A Compact Maxima Representative

Abolfazl Asudeh; Azade Nazi; Nan Zhang; Gautam Das

Regret-ratio minimizing set can serve for (approximately) answering maxima queries when convex hull is large.

We make *several fundamental theoretical as well as practical advances in developing such a compact set*.

- In 2D: we develop an optimal linearithmic time algorithm by leveraging the ordering of skyline tuples.
- In HD: we develop an *approximation algorithm* that runs in linearithmic time and guarantees a regret ratio, within *any arbitrarily small user-controllable distance* from the optimal regret ratio.

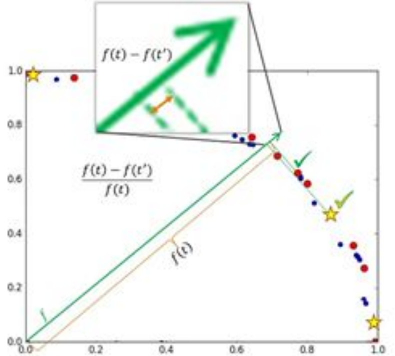

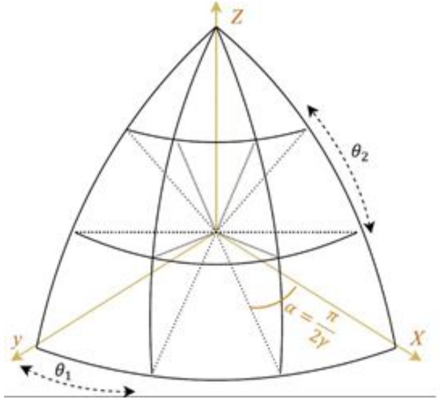

User Pref.

FEXIPRO: Fast and Exact Inner Product Retrieval in Recommender Systems

[1] Hui Li; [2] Tsz Nam Chan; [2] Man Lung Yiu; [1] Nikos Mamoulis

[1] The University of Hong Kong and [2] Hong Kong Polytechnic University

- **Problem: top-k inner product retrieval**
  - Matrices **Q** and **P** come from matrix factorization
  - Large $\mathbf{q}^T\mathbf{p}$ indicates a possible recommendation
- **Our Framework FEXIPRO**
  - Use Thin SVD, Integer Approximation and Monotonicity Reduction to manipulate data
  - Orthogonal to existing systems
- **Experiments**
  - At least **one order of magnitude** faster than existing methods
  - Single-thread FEXIPRO is faster than multi-thread Intel MKL and requires much less memory

$$\begin{pmatrix} 1.6 & 1.3 & 0.7 & 1.0 & 0.4 \\ 0.6 & 0.8 & 2.7 & 2.8 & 2.2 \end{pmatrix} \quad \mathbf{P}$$

$$\underbrace{\begin{pmatrix} 3.2 & -0.4 \\ 3.1 & -0.2 \\ 0 & 1.8 \\ -0.4 & 1.9 \end{pmatrix}}_{\mathbf{Q}^T} \begin{pmatrix} \mathbf{4.9} & \mathbf{3.8} & 1.2 & 2.1 & 0.4 \\ \mathbf{4.8} & \mathbf{3.9} & 1.6 & 2.5 & 0.8 \\ 1.0 & 1.4 & \mathbf{4.9} & \mathbf{5.0} & 4.0 \\ 0.5 & 1.0 & \mathbf{4.9} & \mathbf{4.9} & 4.0 \end{pmatrix} \quad \mathbf{Q}^T\mathbf{P}$$

Large Values = Good Recommendation

User Pref.

Feedback-Aware Social Event-Participant Arrangement

Jieying She; Yongxin Tong; Lei Chen; Tianshu Song

- Background & Motivation
  - Event arrangement in EBSN
  - Existing studies
    - The satisfaction scores are hard to learn
    - Users may not accept the arrangements

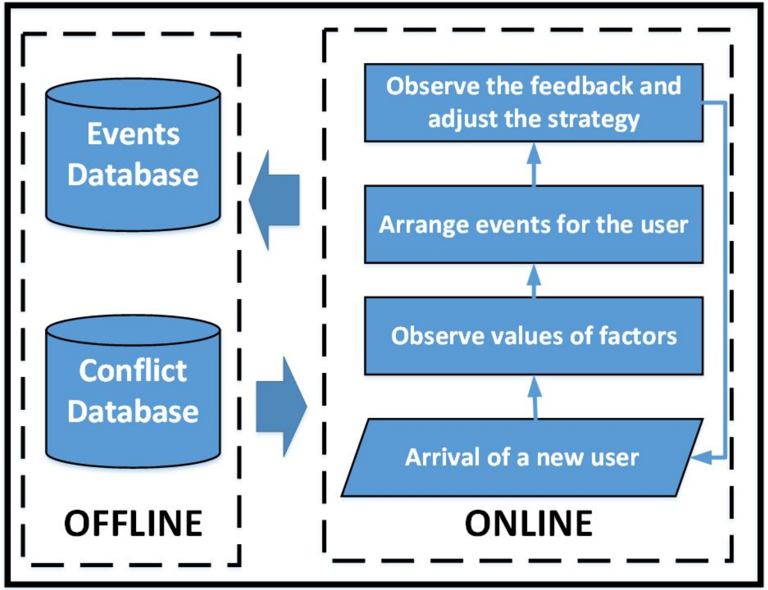

- Solutions
  - Multi-arm Bandit (MAB) based framework
    - Thompson Sampling based solution
    - Upper Confidence Bound (UCB) based solution
- Experiments
  - The ***Thompson Sampling*** based solution ***does not perform well*** under FASEA
  - The ***UCB*** based solution ***is the best in overall*** by extensive experiments on both real and synthetic datasets

Exploiting Common Patterns for Tree-Structured Data

Zhiyi Wang Shimin Chen

(Institute of Computing Technology, Chinese Academy of Sciences)

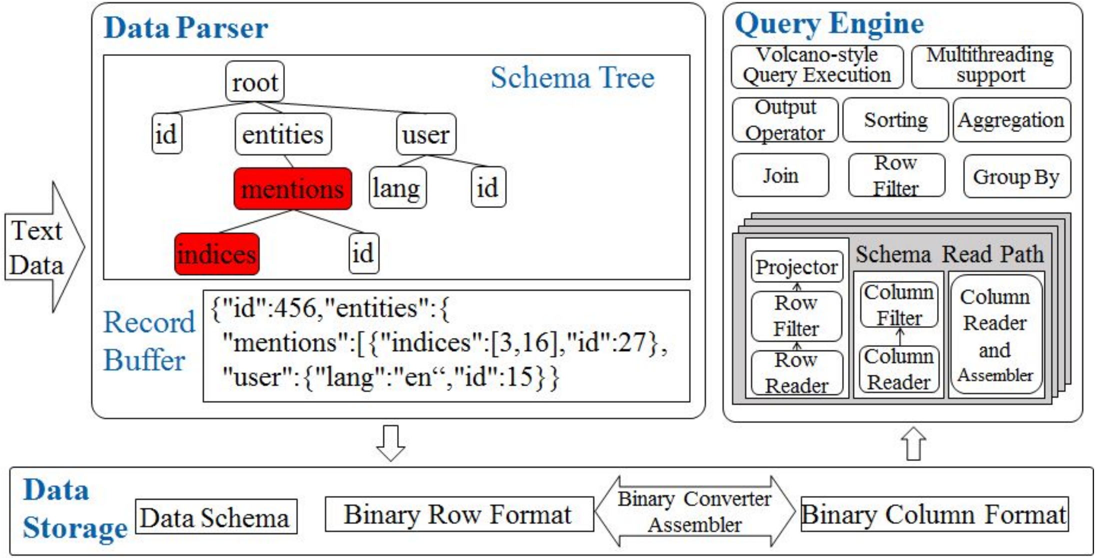

STEED:

System for TrEE structured Data

- **Supports tree-structured data:** e.g., JSON, Protocol Buffers, etc.
- **Exploits real-world data pattern: simple path optimization**
- **Achieves 10~1000X speedup compared to state-of-the-art systems**

Extracting & Analyzing *Hidden* Graphs from RDBMSs

Tree & Graph (2)

<u>Konstantinos Xirogiannopoulos</u>, Amol Deshpande | University of Maryland, College Park

Graph-structured **data** can enable **analyses** ***impossible*** using SQL analytics

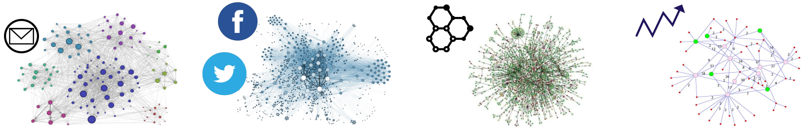

Email Networks *Social Networks* *Protein Interaction Networks* *Stock Trading Networks*

But first...***where*** is your data ***stored***??

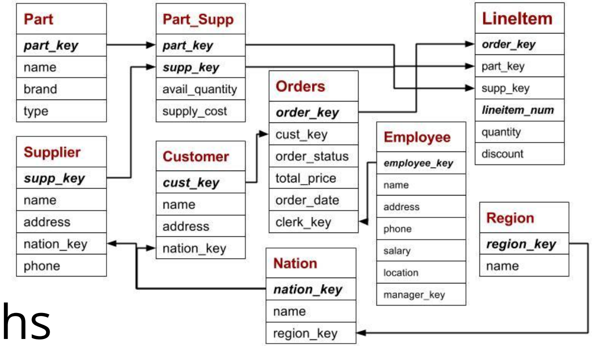

Problem: When ***extracting*** graphs from RDBMSs, graphs often ***orders-of-magnitude*** *larger* than the **input** tables. Graph may ***not fit*** in memory

Solution: A software layer (GraphGen) over the database that ***efficiently*** loads in a ***condensed-representation***, and enables efficient processing through various ***APIs***.

TrillionG: A Trillion-scale Synthetic Graph Generator using a Recursive Vector Model

Himchan Park; Min-Soo Kim

- Motivation
  - Large-scale realistic graphs for benchmarks, e.g., RMAT, Kronecker
  - Low-level core techniques for rich graph generation, e.g., LDBC, gMark
- Solutions
  - A Vertex Scope approach (AVS)
  - Recursive Vector Model (*RecVec*)
- Experiments
  - More realistic graphs by adding noises
  - Schema-driven rich graph
  - **A trillion edges graph within 2 hours using just 10 PCs**

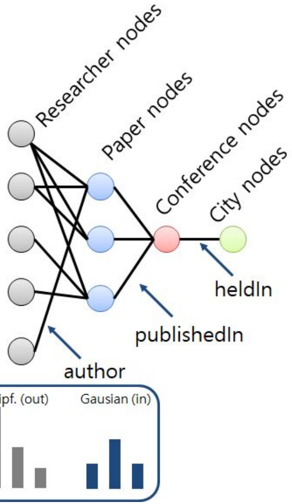

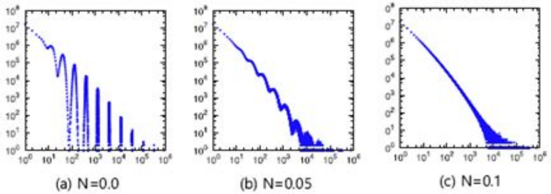

Schema Independent Relational Learning

ML

Jose Picado; Arash Termehchy; Alan Fern; Parisa Ataei

Oregon State University

Mission: Find any sign of **life** on Earth.

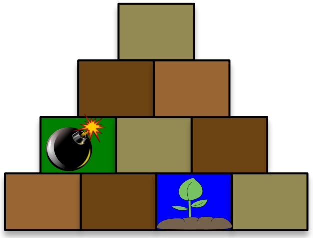

boxes			
box	item	color	desc
1	bomb	green	solid
2	plant	blue	wet

contains	
box	color
1	bomb
2	plant

color	
box	item
1	green
2	blue

description	
box	desc
1	solid
2	wet

The result of **current** learning algorithms depend on the schema.

life(x) :- boxes(y,x,z,wet).

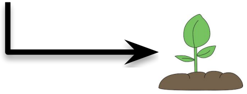

life(x) :- contains(y,x), color(y,green).

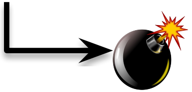

People represent same data using different schemas.

We want to learn same accurate answers over all possible schemas for the same data.

life(x) :- boxes(y,x,z,wet).

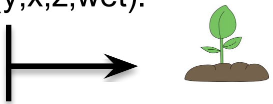

life(x) :- contains(y,x), color(y,z), description(y,wet).

Castor: schema independent, accurate and efficient. It leverages concepts of schema design.

Scalable Kernel Density Classification using Threshold-Based Pruning

ML

Stanford University

Edward Gan; Peter Bailis

Complex Distributions

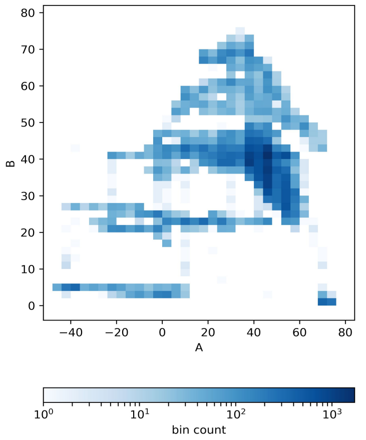

Expensive Estimators

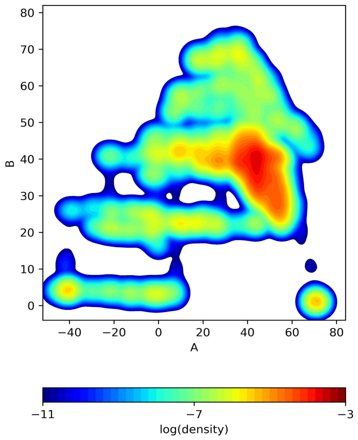

Classification Results

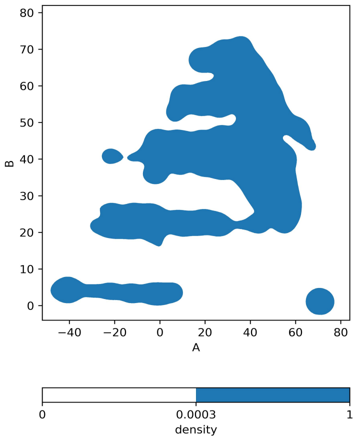

Kernel Density Estimation

+

Classification Predicate

Kernel Density Classification

K-d tree Index

ML + Predicate Pushdown: Asymptotic (1000x) Speedups

ML

The BUDS Language for Distributed Machine Learning

Zekai "Jacob" Gao; Shangyu Luo; Luis Perez; Chris Jermaine

RICE

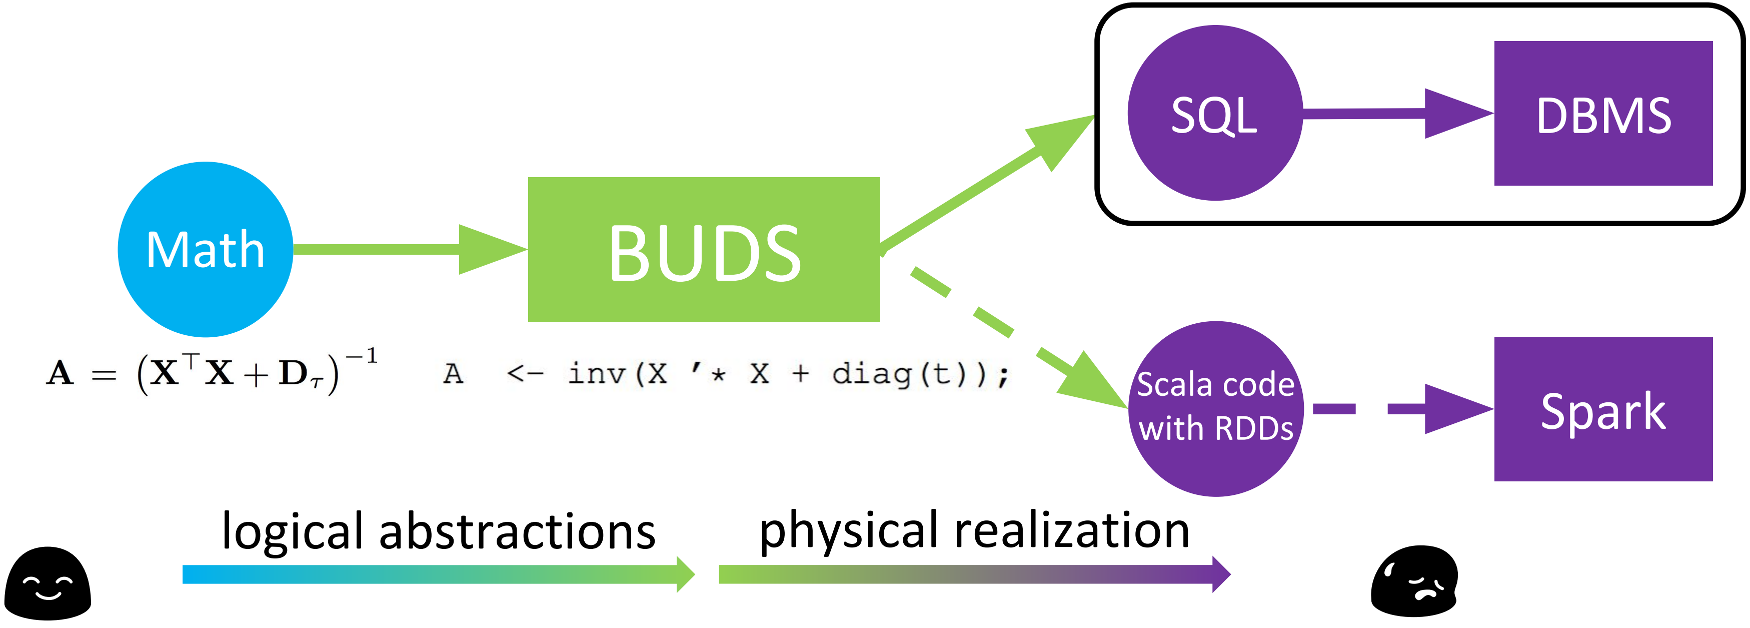

$$\mathbf{A} = (\mathbf{X}^\top \mathbf{X} + \mathbf{D}_\tau)^{-1}$$

```
A  <- inv(X '* X + diag(t));
```

Our results: the BUDS optimized implementations have competitive performance compared to the hand-coded SQL implementations

Improving the Life of a Data Scientist

ML

Zoi Kaoudi; Jorge Quiane; Sara Thirumuruganathan; Sanjay Chawla; Divy Agrawal

Data scientist

What people think he does

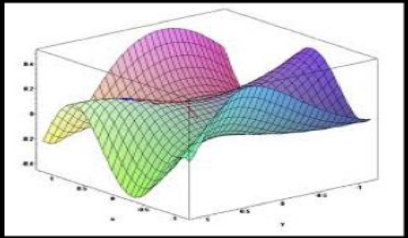

What he thinks he does

What he actually does

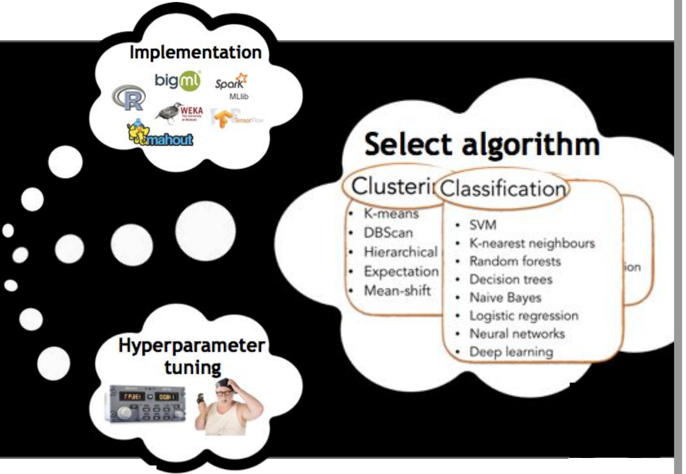

Best performance without sacrificing accuracy

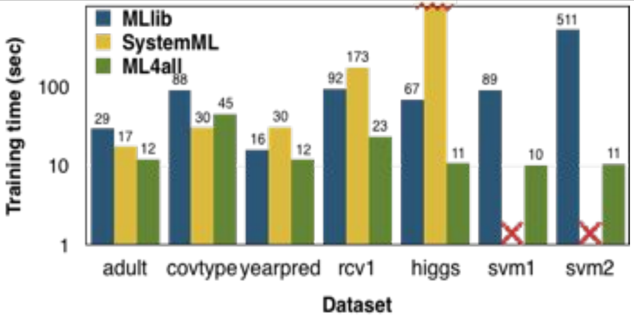

An Experimental Study of Bitmap Compression and Inverted List Compression

Jianguo Wang; Chunbin Lin; Yannis Papakonstantinou; Steven Swanson

15	28	45	64	80	168

The integer array compression problem

bitmap compression

inverted list compression

Which is better?

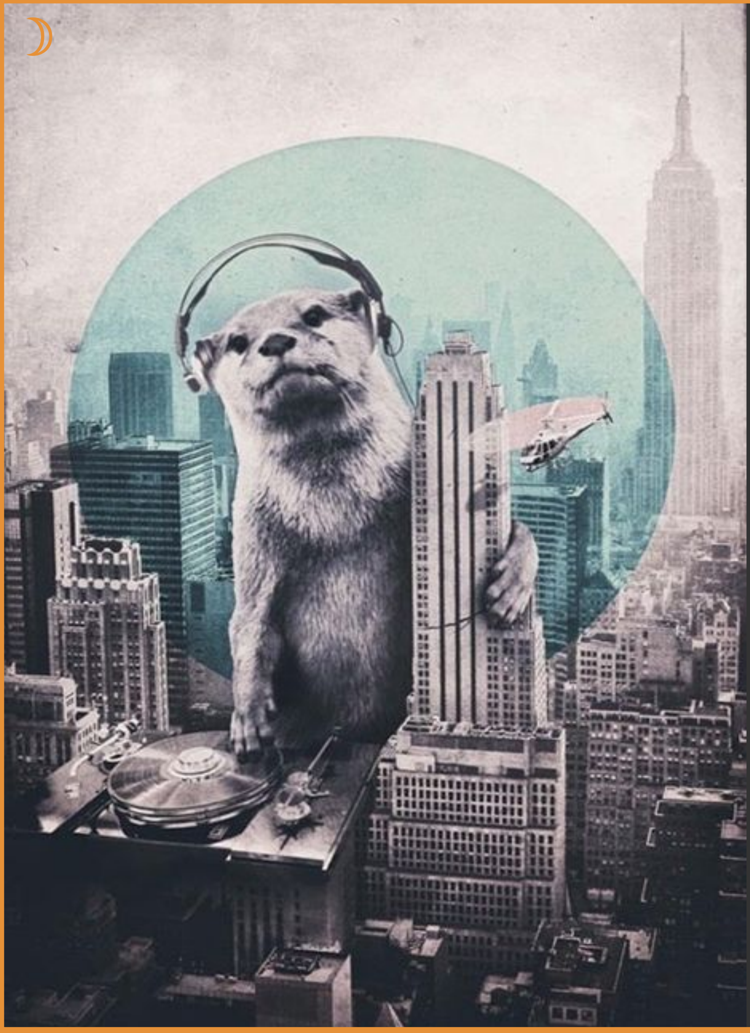

OtterTune

Automatic Database Management System Tuning Through Large-scale Machine Learning

OtterTune leverages past experiences to tune new DBMS configurations.

OtterTune outperforms tuning scripts and Amazon RDS.

Can OtterTune match a DBA?

Dana Van Aken; Andrew Pavlo; Geoffrey J. Gordon; Bohan Zhang

Solving the Join Ordering Problem via Mixed Integer Linear Programming

Immanuel Trummer; Christoph Koch

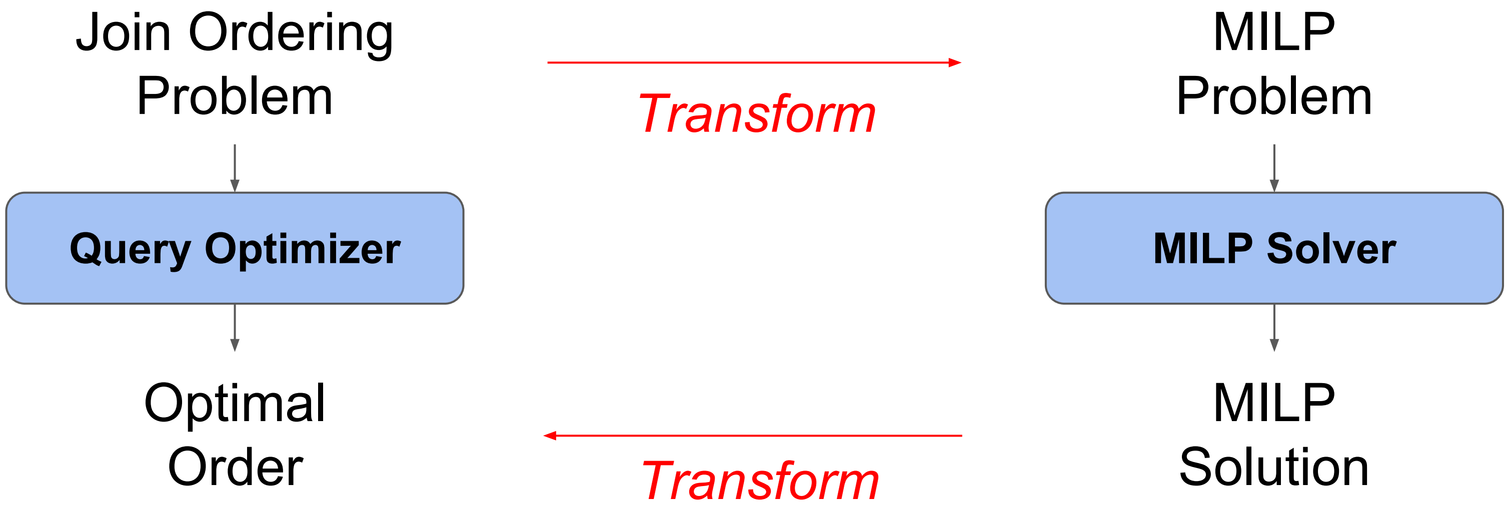

Optimizing Queries with Up to 40 Tables ...

Amazon Aurora: Design Considerations for High Throughput Cloud-Native Relational Databases

<u>Problem:</u> Once storage and compute are decoupled the I/O bottleneck moves to the network

- Offload redo processing from compute
- Purpose-built scale-out multi-tenant log-structured distributed storage service designed for databases
- Storage volume striped across hundreds of nodes over 3 availability zones (AZ)
- Six copies of data, two in each AZ to protect against correlated AZ+1 failures

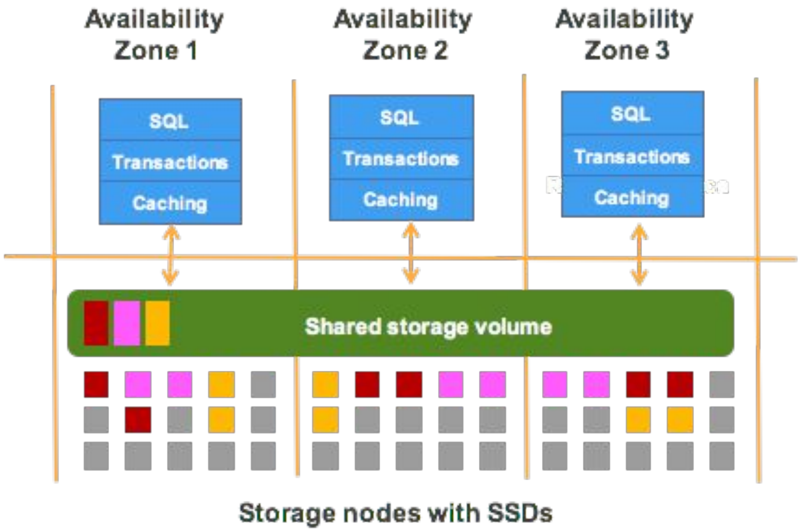

<u>Result:</u> Better durability, availability, & jitter …

… which are all the same thing (on different time scales)