



# Presentation from the VLDB 2015 PC Chairs



Chen Li  
UC Irvine



Volker Markl  
TU Berlin



33

Research Sessions

6

Tutorials

7

Industrial Sessions

4

Keynotes

3

Demo Sessions

2

Panels

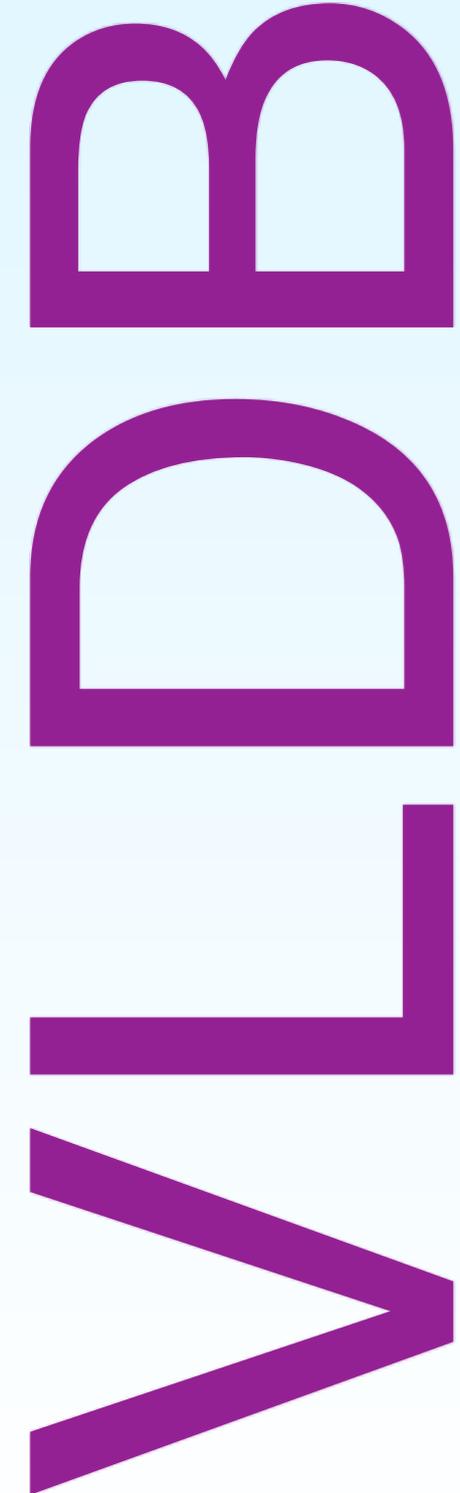
2

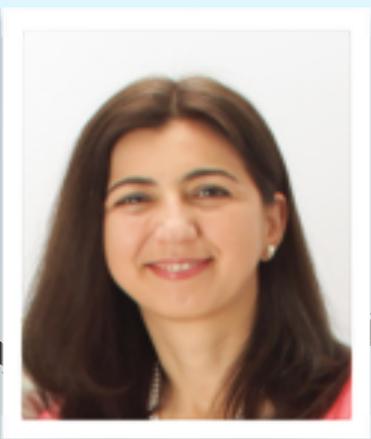
Poster Sessions

9

Workshops

2015 Turing Award Lecture!





Fatma Özcan

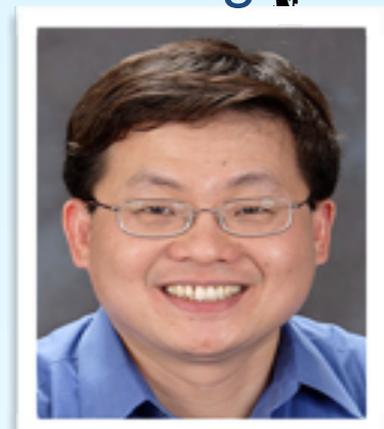


Felix Naumann

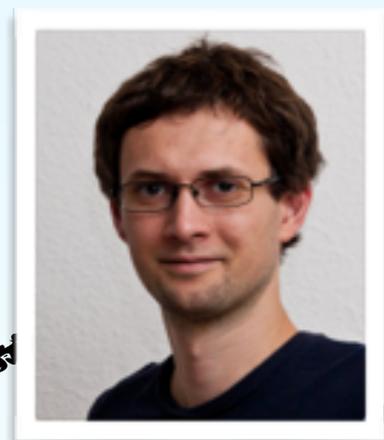
Jignesh M. Patel



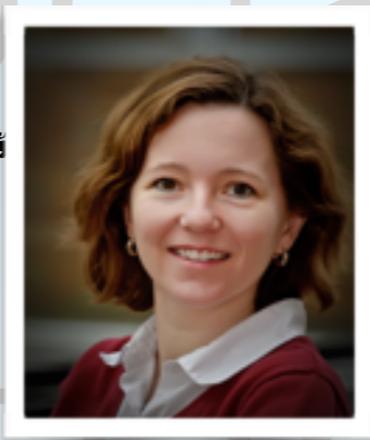
Kevin C. Chang



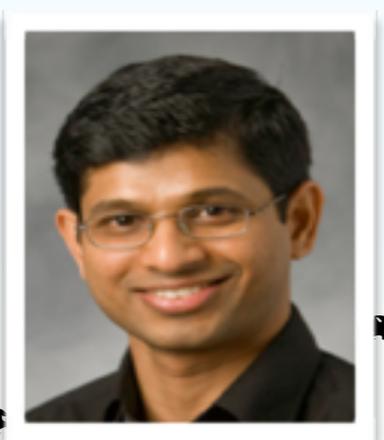
# Associate Editors



Rainer Gemulla



Magdalena Balazinska



Shivnath Babu

Stefan Manegold



Yi Chen









# Committee

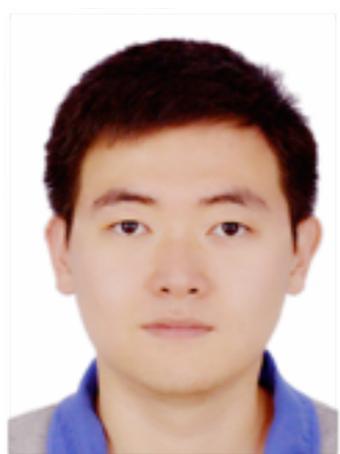
- Sihem Amer-Yahia (Chair)
- Beng Chin Ooi
- Walid Aref
- Patrick Valduriez



# Best Paper Award

# Constructing an Interactive Natural Language Interface for Relational Databases

Tuesday, 15:30-17:00 (Research Session 7)



Fei Li

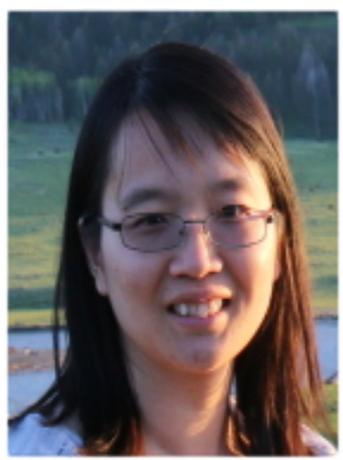


H. V. Jagadish

The paper proposes an interactive natural language interface for relational databases, which enables novice users to construct complex queries. It improves the usability of an RDBMS, as it enables anyone to use to ask questions to a database system. This paper is likely to start a new line of research as well as products. For a query expressed in natural language, the interface interacts with the user in several steps (as we do in real life to make our questions more precise) in determining the query semantics and subsequently generating the corresponding SQL. At each step, the system interactively presents to the user its own understanding of the query through alternatives, as opposed to just final answers. The authors rely on a query tree structure to represent the interpretation of an NLP query from the database's perspective, which facilitates verification by users, and translation into SQL. The system (NLIDB) was implemented following the component-based approach, where each component can be independently constructed, optimized or substituted. The experiments involve real users and verify the feasibility of the approach and illustrate the strengths of the system/ approach.



# Best Paper Award



Jiexing Li

# Resource Bricolage for Parallel Database Systems

Wednesday, 13:30–15:00 (Research Session 15)



Jeffrey Naughton

This is a core database systems paper that addresses a real problem, mainly how to deal with the heterogeneity in the machines composing a cluster-based database system. The paper proposes a mechanism ("Resource Bricolage") to make efficient use of heterogeneous hardware when processing a workload in a parallel database system. It addresses a very relevant problem (clusters don't grow homogeneously) and is the first paper on this subject. The approach is relatively simple and practical, using linear programming to optimize data distribution - and thus resource consumption - in a cluster. The techniques were implemented on top of Microsoft SQL server parallel data warehouse. Overall, this is an excellent and impactful paper. One can envision many extensions that can follow up from this research.



Rimma V. Nehme

Best Paper Award  
(Runner-up)





# Committee

- Alfons Kemper (Chair)
- Viktor Leis
- Justin Levandoski
- Uwe Röhm
- Pinar Tözün



# Best Demo Award



# Vizdom: Interactive Analytics through Pen and Touch

Andrew Crotty (Brown University)

Alex Galakatos (Brown University)

Emanuel Zraggen (Brown University)

Carsten Binnig (Brown University)

Tim Kraska (Brown University)

Thursday, 15:30-17:00 (Demo Session 3)

Best Demo Award





## Papers

## Talks

## Posters

**Research Track Papers**  
(including Rollover Papers  
from VLDB 2014)

**33 Research Sessions**  
18 min/paper,  
1 min intro, 15 min  
presentation, 2 min Q&A

**Tuesday**  
Sessions 1-12  
**Thursday**  
Sessions 13-33

**Industrial Track Papers**

**7 Industrial Sessions**  
30 min/paper  
1 min intro, 25 min  
presentation, 4 min Q&A

**Tuesday**

**VLDB Journal Papers**

**Tuesday**

Mapping Talks to Posters



151

The total number of accepted papers in the Research Track out of **710** submissions: **139** will be presented this week and **12** will rollover to **VLDB 2016**.

21

The total number of rollover papers from VLDB 2014.

20

The total number of accepted papers in the Industrial Track out of **68** submissions.

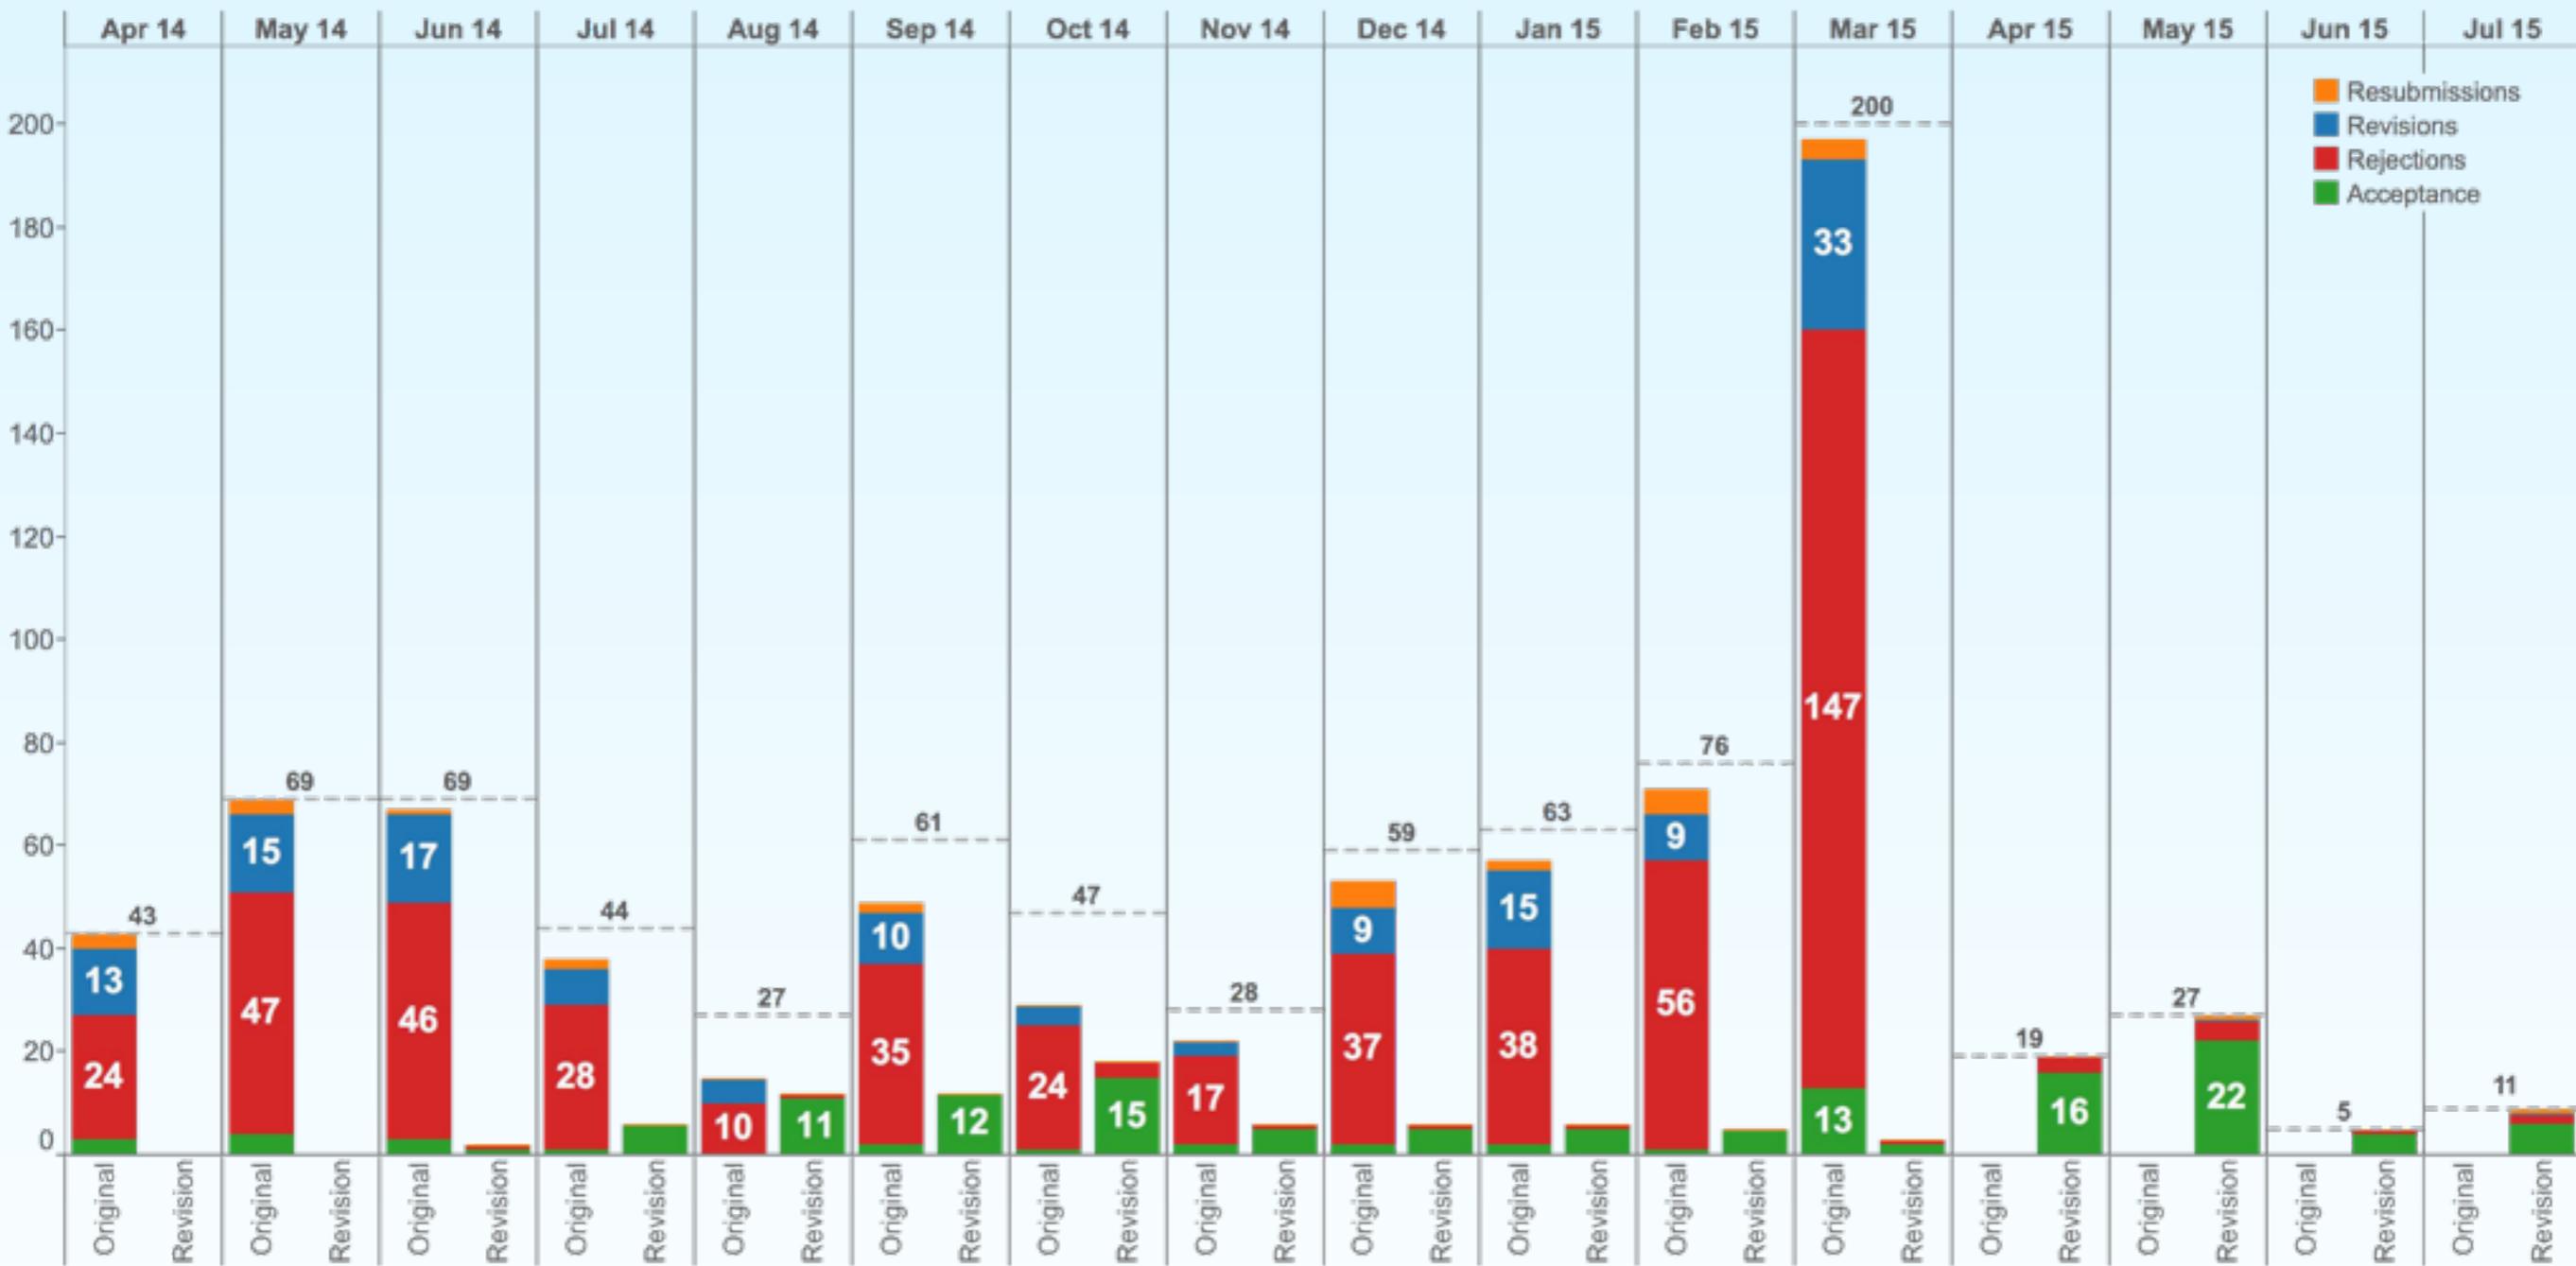
49

The total number of accepted papers in the Demonstration Track out of **148** submissions.

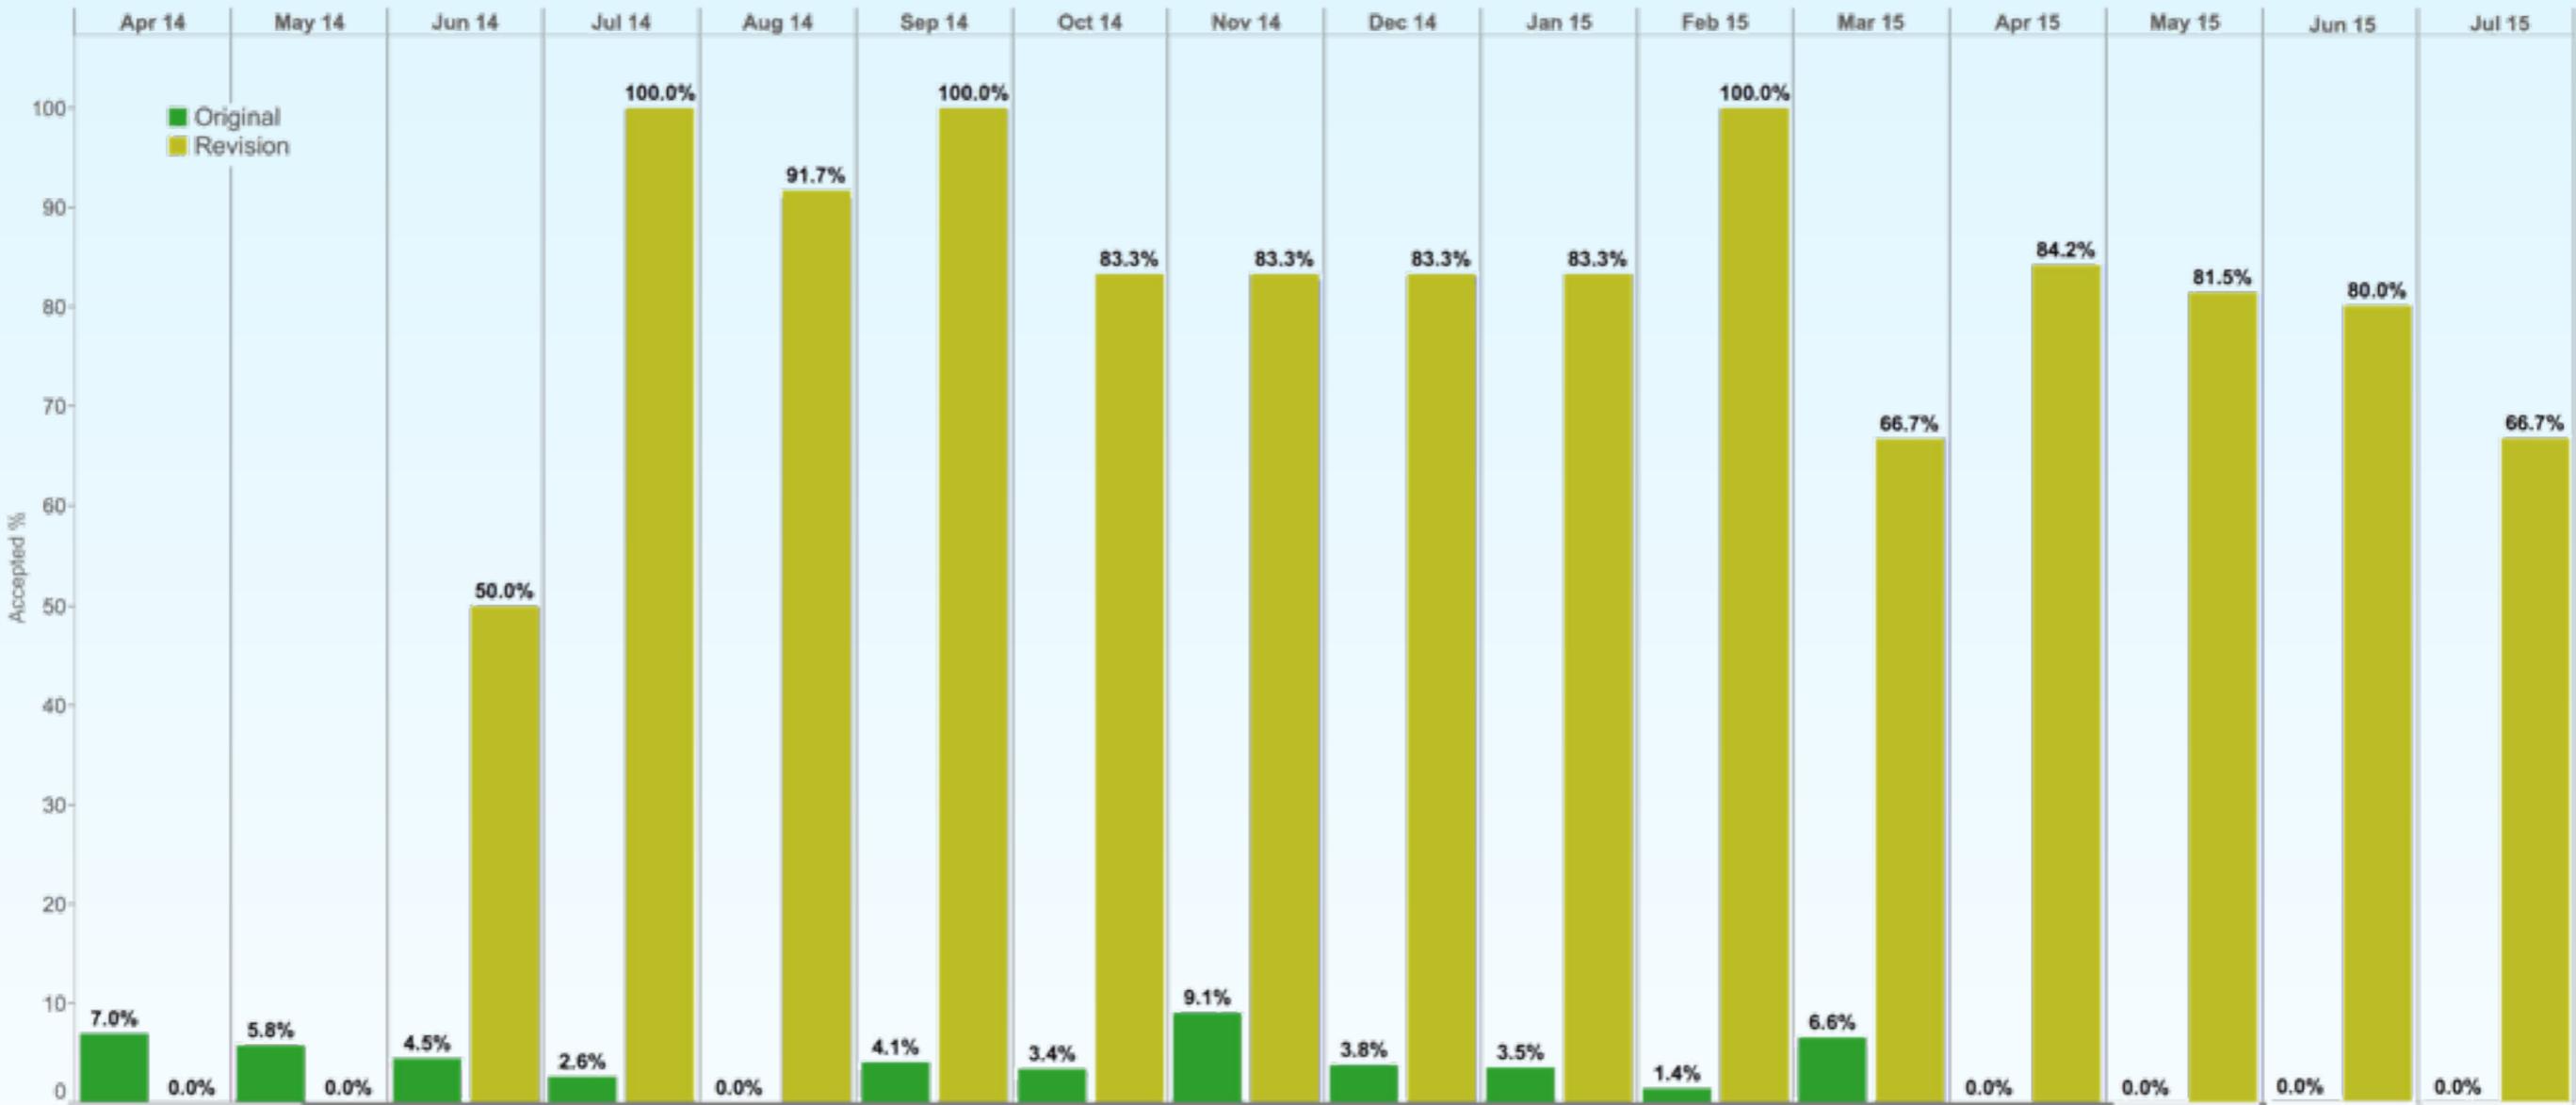
# VLDB '15 Summary







# Research Track

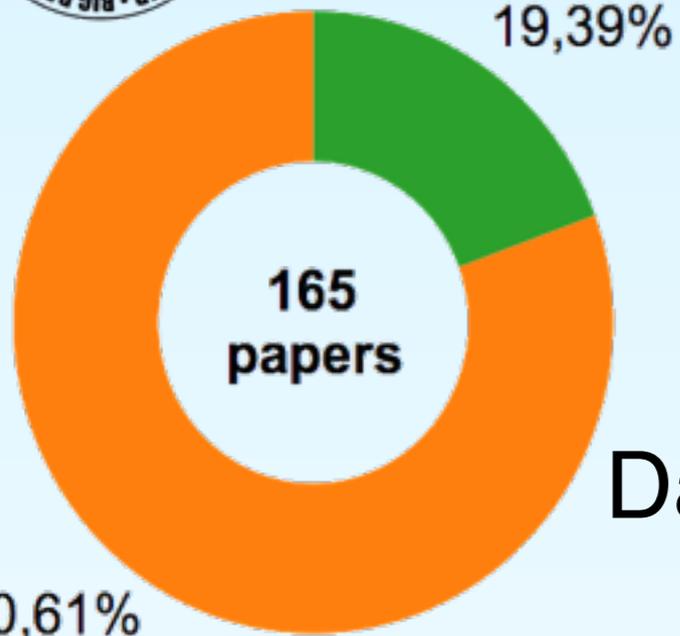


# Acceptance Ratio

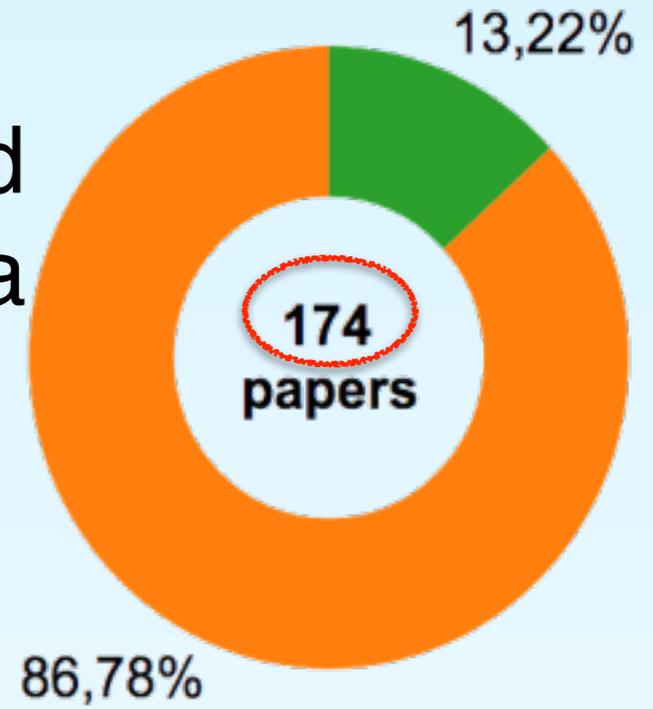


Accepted  
Submission

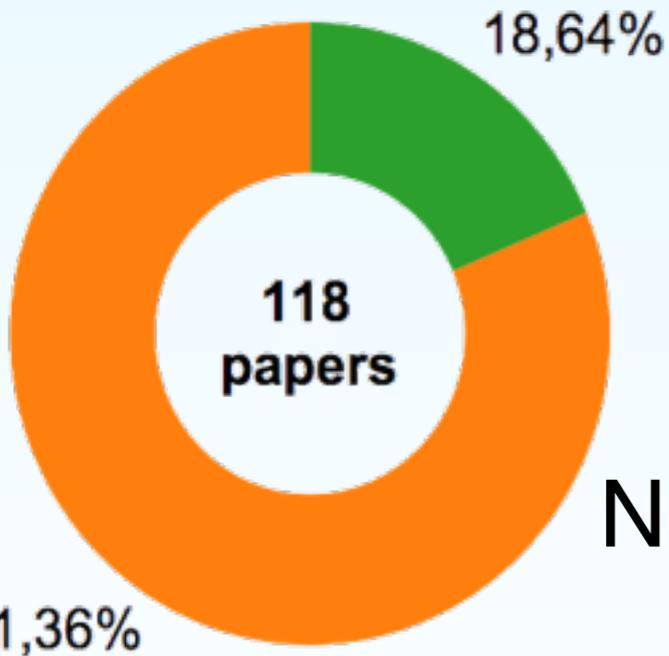
# Text, Data Types, and Semi-structured Data



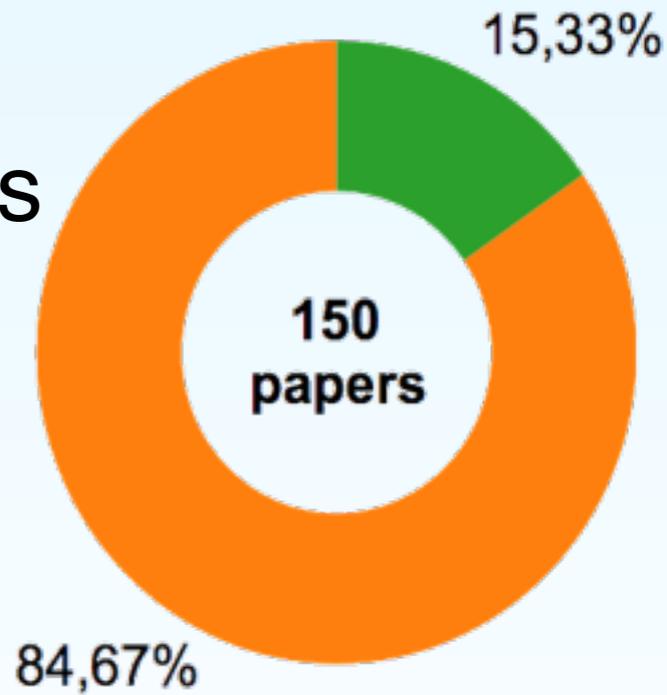
Database Engines



Applications



Novel Database Architectures



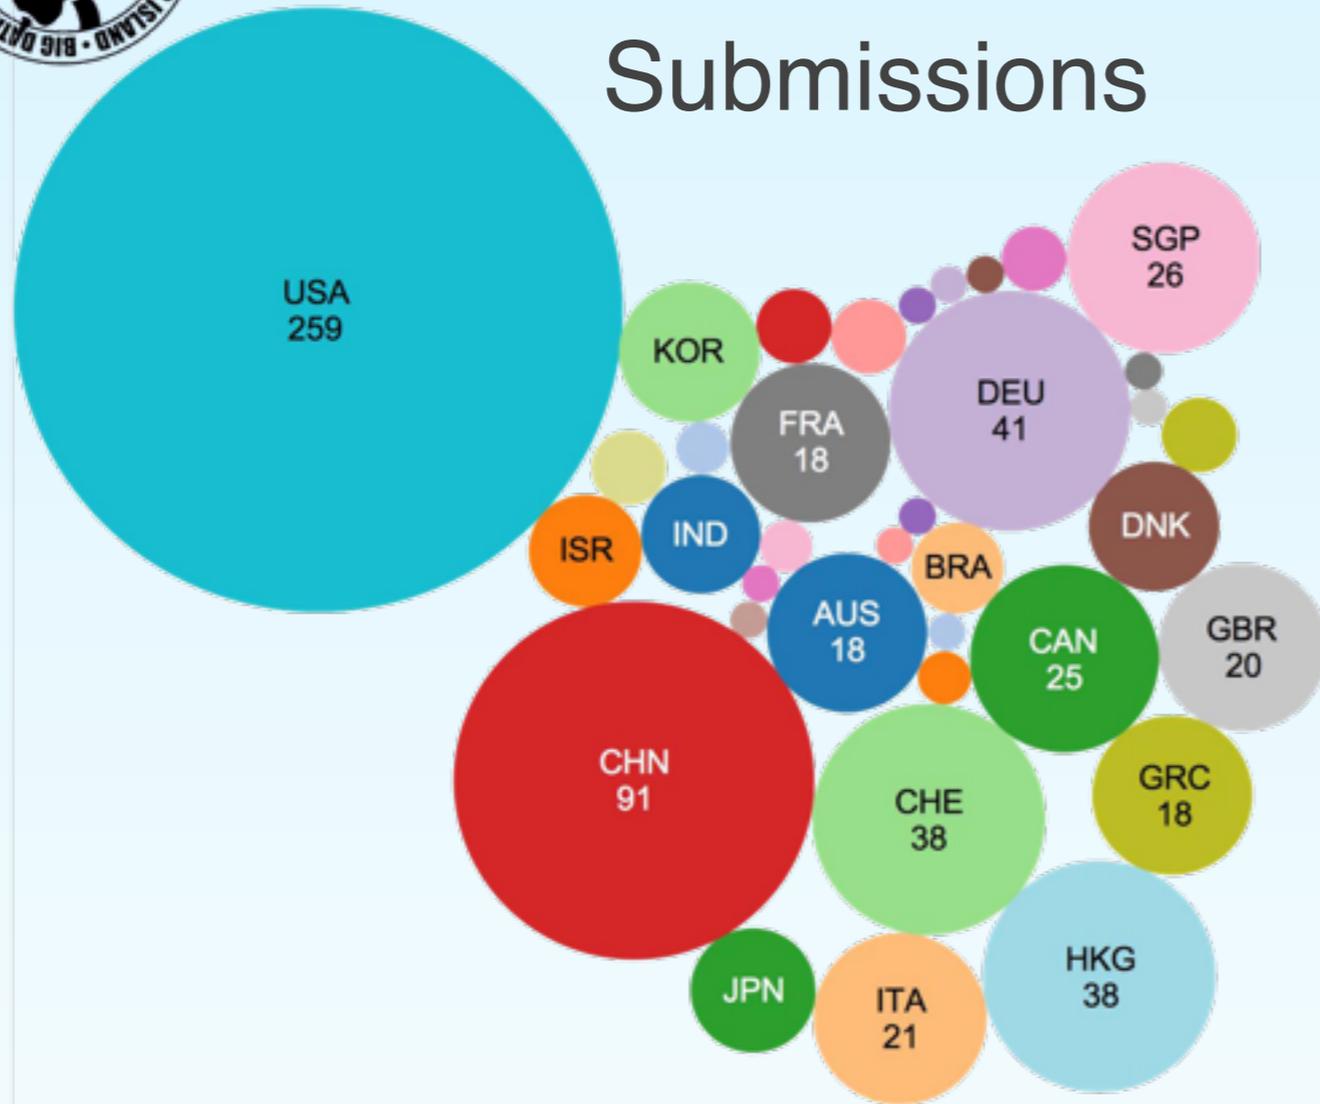
# Top Primary Areas



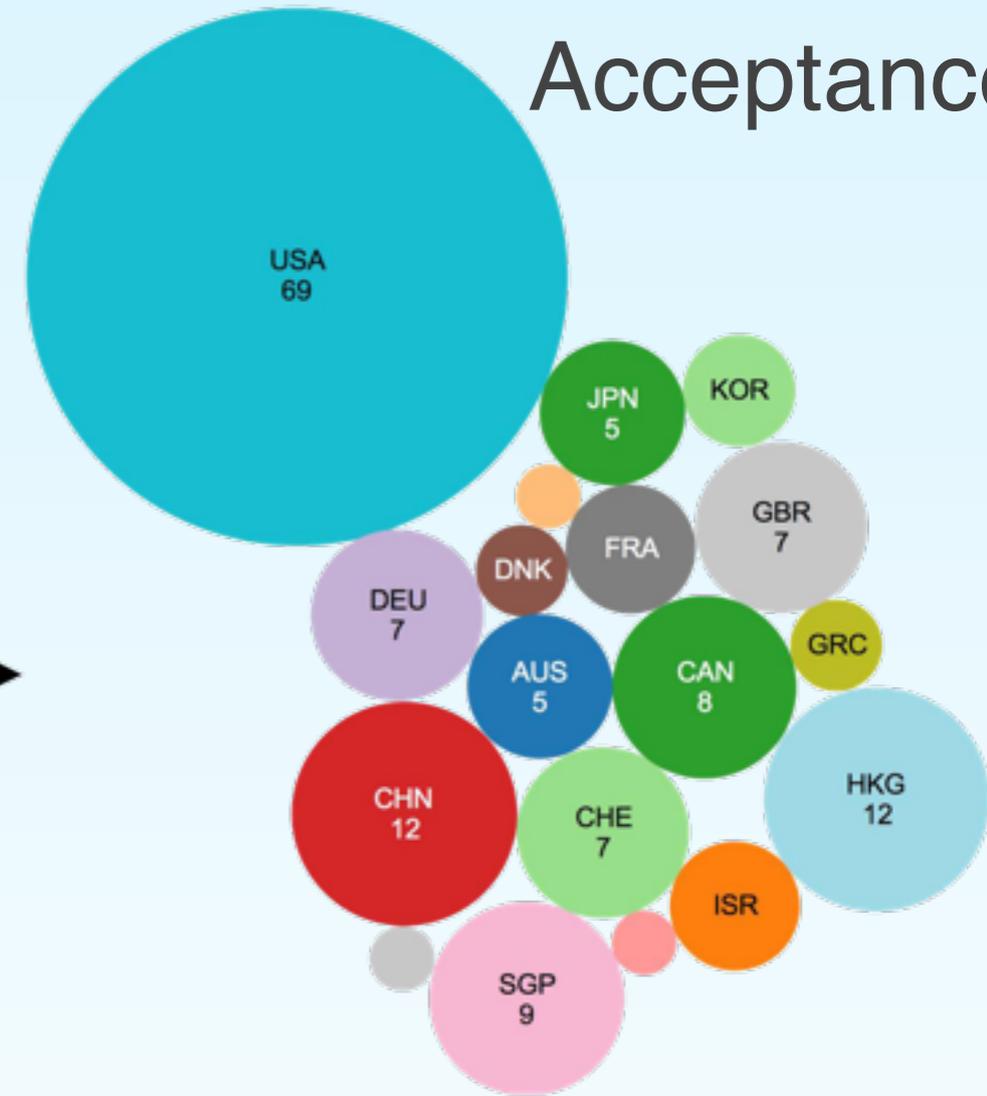




# Submissions

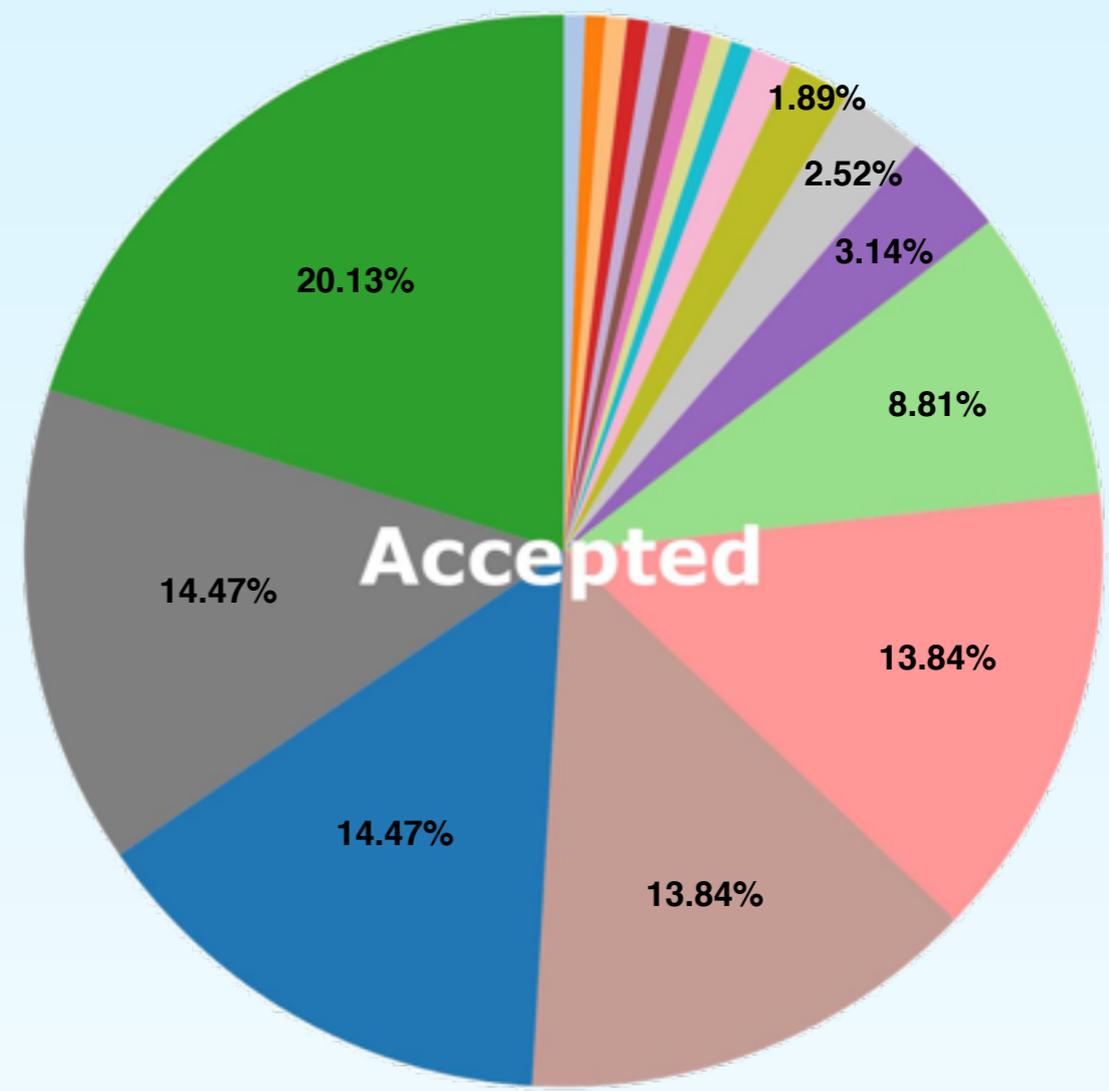
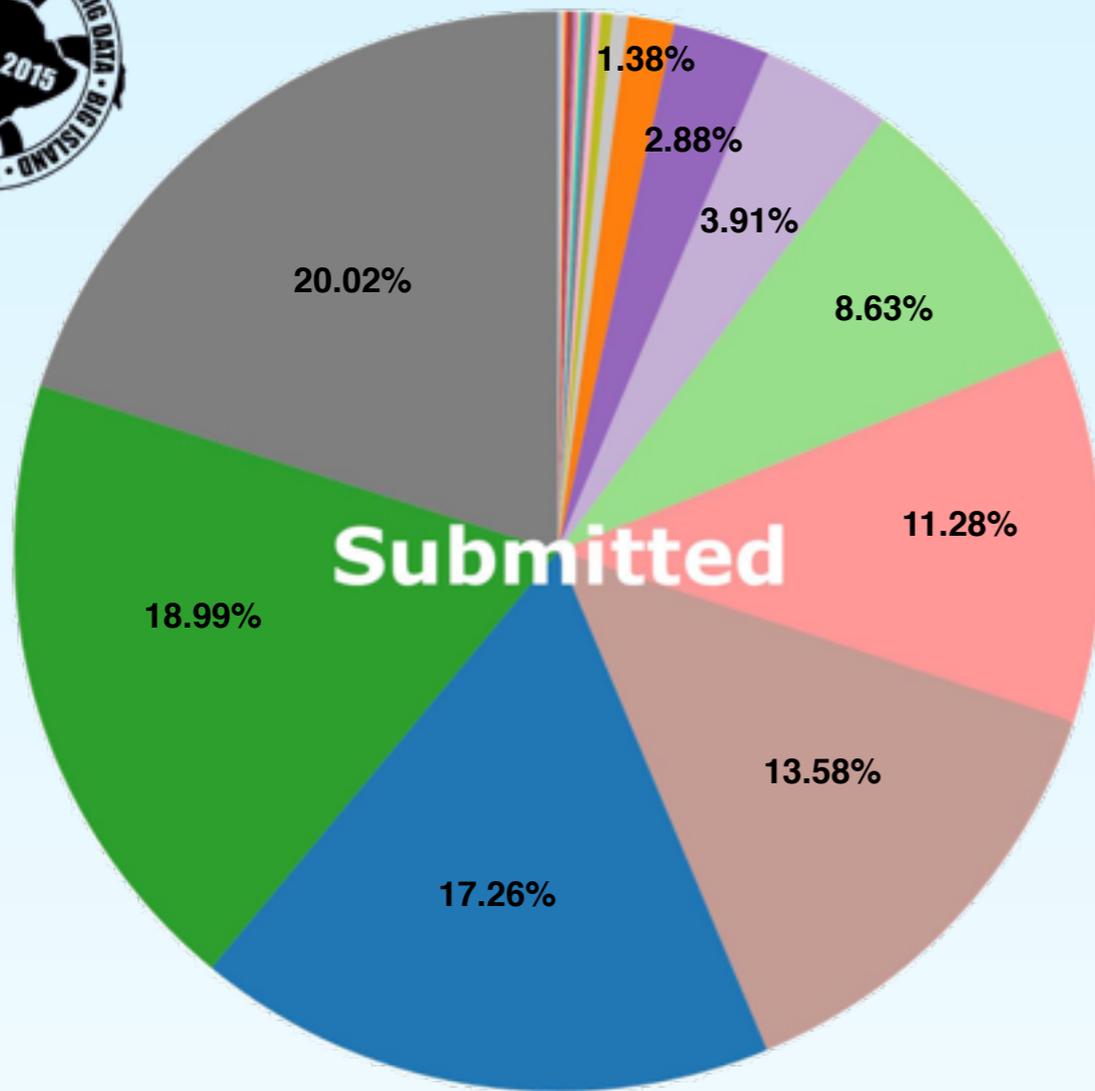


# Acceptance



- |  |  |  |  |
|--|--|--|--|
| <span style="color: blue;">■</span> AUS        | <span style="color: purple;">■</span> DEU      | <span style="color: blue;">■</span> IND        | <span style="color: purple;">■</span> PAK      |
| <span style="color: lightblue;">■</span> BEL   | <span style="color: brown;">■</span> DNK       | <span style="color: lightblue;">■</span> IOT   | <span style="color: brown;">■</span> QAT       |
| <span style="color: orange;">■</span> BGD      | <span style="color: brown;">■</span> EGY       | <span style="color: orange;">■</span> ISR      | <span style="color: pink;">■</span> SAU        |
| <span style="color: lightorange;">■</span> BRA | <span style="color: pink;">■</span> ESP        | <span style="color: lightorange;">■</span> ITA | <span style="color: pink;">■</span> SGP        |
| <span style="color: green;">■</span> CAN       | <span style="color: pink;">■</span> FIN        | <span style="color: green;">■</span> JPN       | <span style="color: grey;">■</span> SWE        |
| <span style="color: lightgreen;">■</span> CHE  | <span style="color: grey;">■</span> FRA        | <span style="color: lightgreen;">■</span> KOR  | <span style="color: grey;">■</span> THA        |
| <span style="color: red;">■</span> CHN         | <span style="color: grey;">■</span> GBR        | <span style="color: red;">■</span> NLD         | <span style="color: yellowgreen;">■</span> TUR |
| <span style="color: lightcoral;">■</span> CYP  | <span style="color: yellowgreen;">■</span> GRC | <span style="color: lightcoral;">■</span> NZL  | <span style="color: yellowgreen;">■</span> TWN |
| <span style="color: purple;">■</span> CZE      | <span style="color: lightblue;">■</span> HKG   | <span style="color: purple;">■</span> OMN      | <span style="color: cyan;">■</span> USA        |

Contributors



# Subject Areas

- |   |  |
|---|--|
| <span style="color: blue;">■</span> Benchmarking and Performance Measurement              | <span style="color: gray;">■</span> Tree, Graph, and Semi-structured Data            |
| <span style="color: orange;">■</span> Data Cleansing and Data Profiling                   | <span style="color: orange;">■</span> Benchmarks and Administration                  |
| <span style="color: red;">■</span> Fuzzy, Probabilistic, and Approximate Data             | <span style="color: purple;">■</span> Innovative Systems                             |
| <span style="color: brown;">■</span> Multi-core, Main memory, and other emerging hardware | <span style="color: lightpurple;">■</span> Languages, User interfaces, and Usability |
| <span style="color: pink;">■</span> Query Processing                                      | <span style="color: lightgreen;">■</span> Experiments and Analysis                   |
| <span style="color: yellowgreen;">■</span> Vision Track                                   | <span style="color: lightcoral;">■</span> Information Integration                    |
| <span style="color: cyan;">■</span> Web Data Management                                   | <span style="color: brown;">■</span> Novel DB Architectures                          |
| <span style="color: gray;">■</span> None of the above                                     | <span style="color: blue;">■</span> Applications                                     |
| <span style="color: pink;">■</span> Spatial Databases and GIS                             | <span style="color: green;">■</span> Database Engines                                |
| <span style="color: yellowgreen;">■</span> User Interfaces                                | <span style="color: gray;">■</span> Text, Semi-structured data, and Data Types       |



**FRIDAY**

To be held on Friday



Bring your laptops!

**x2**

Each tutorial will be presented twice.

# 8 Tutorials on Publicly Available Open Source Big Data Systems



**REEF**



**Flink**



BOSS Workshop



# Engineering Database Hardware and Software Together



## Juan Loaiza, Oracle

As Senior Vice President of Systems Technology at Oracle, Juan Loaiza is in charge of developing the mission-critical capabilities of Oracle Database, including data and transaction management, high availability, performance, in-memory processing, enterprise replication, and Oracle Exadata.

Mr. Loaiza joined the Oracle Database development organization in 1988. Mr. Loaiza holds BS and MS degrees in computer science from the Massachusetts Institute of Technology.

Keynote





# Big Plateaus of Big Data on the Big Island

## Todd Walter, Teradata



Todd Walter is the Chief Technologist for Teradata across the Americas region. With substantive expertise in big data, database engineering and systems architecture, he works closely with Teradata customers, colleagues, and alliance partners to evaluate and prioritize initiatives — and implement data strategy and analytics. As a pragmatic visionary, Walter helps customer business analysts as well as technologists better understand all of the astonishing possibilities of big data and analytics in view of emerging as well as existing capabilities of information infrastructures.

Todd works with organizations of all sizes and levels of experience, from start-ups to Fortune 100 companies at the leading edge of adopting big data, data warehouse and analytics technologies. Walter has been with Teradata for nearly 28 years, contributing significantly to Teradata's unique design features and functionality. He holds more than a dozen Teradata patents and is a Teradata Fellow, the highest technical award granted by the company. Todd served for more than ten years as Chief Technical Officer of Teradata Labs, responsible for vision, strategy and technical leadership of the Teradata product line before taking on his current strategic consulting role.

Keynote



# Big Data Research: Will Industry Solve all the Problems?



## Magdalena Balazinska, University of Washington

Magdalena Balazinska is an Associate Professor in the department of Computer Science and Engineering at the University of Washington and the Jean Loup Baer Professor of Computer Science and Engineering. She's the director of the IGERT PhD Program in Big Data and Data Science. She's also a Senior Data Science Fellow of the University of Washington eScience Institute. Magdalena's research interests are in the field of database management systems. Her current research focuses on big data management, scientific data management, and cloud computing. Magdalena holds a Ph.D. from the Massachusetts Institute of Technology (2006). She is a Microsoft Research New Faculty Fellow (2007), received an NSF CAREER Award (2009), a 10-year most influential paper award (2010), an HP Labs Research Innovation Award (2009 and 2010), a Rogel Faculty Support Award (2006), a Microsoft Research Graduate Fellowship (2003-2005), and multiple best-paper awards.

Keynote



**Many thanks to everyone who devoted their time & energy to make VLDB 2015 possible.**

**Enjoy the conference!**



We thank Reda Al Azmeh for his support in preparing figures and this PC Chair presentation!

Acknowledgement