VLDB2011
37th International Conference on Very Large Data Bases, Seattle, WA

Proceedings of the VLDB Endowment

Volume 4, No. 12 – Industrial, Applications and Experience Track, Demonstrations, Challenges & Vision Track, Tutorials and Panels
Proceedings of the 37th International Conference on Very Large Data Bases, Seattle, WA

Editor-in-Chief:
H. V. Jagadish

Guest Editors:
José Blakeley, Joseph M. Hellerstein, Nick Koudas, Wolfgang Lehner, Sunita Sarawagi, Uwe Röhm

PVLDB – Proceedings of the VLDB Endowment

Volume 4, No. 12, August 2011.

The 37th International Conference on Very Large Data Bases, Seattle, WA.

Copyright 2011 VLDB Endowment

Volume 4, Number 12:
VLDB2011 Industrial Track, Demonstrations, Challenges and Vision Track, Tutorials and Panels

Pages ii – viii and 1213 – 1515

ISSN 2150-8097, August 2011.

Additional copies only online at: portal.acm.org, arxiv.org/corr, and www.vldb.org

TABLE OF CONTENTS

Front Matter

Industrial, Applications, and Experience Track Papers

Demonstration Track Papers

Challenges and Vision Track Papers

Tutorials

Panels

VLDB 2011 PROGRAM COMMITTEES

VLDB 2011 General PC Co-Chairs

José Blakeley, Microsoft

Joseph Hellerstein, University of California – Berkeley

Industrial and Applications Track Co-Chairs

Berthold Reinwald, IBM Almaden Research Center

Phil Bohannon, Yahoo! Research

Industrial and Applications Track Program Committee

Roger Barga (Microsoft Research)

Christof Bornhoevd (SAP Palo Alto)

Brian Cooper (Google)

Anish Das Sarma (Yahoo! Research)

Mario Inchiosa (Netezza)

Avinash Lakshman (Facebook)

Zhen Hua Liu (Oracle)

Jun Rao (LinkedIn)

Mehul Shah (HP Labs)

Adam Silberstein (Yahoo! Research)

Alkis Simitsis (HP Labs)

Garret Swart (Oracle)

Yuanyuan Tian (IBM Research)

Yu Xu (Teradata)

Peter Zabback (Microsoft)

Calisto Zuzarte (IBM)

Demonstration Track Co-Chairs

Jignesh Patel, University of Wisconsin – Madison

Masatoshi Yoshikawa, University of Kyoto

Demonstration Track Program Committee

Sourav Bhowmick (Nanyang Technological University, Singapore)

Dhruba Borthakur (Facebook)

Sarah Boulakia (LRI Orsay, France)

Sang Cha (Seoul National University, Korea)

Venkatesh Ganti (Google)

Shahram Ghandeharizadeh (U. Sourthern California)

Yoshiharu Ishikawa (Nagoya University, Japan)

Volker Markl (TU Berlin, Germany)

Jun Miyazaki (Nara Institute of Science and Technology)

Bongki Moon (University of Arizona)

Atsuyuki Morishima (Tsukuba University, Japan)

Rimma Nehme (Microsoft Research)

Jun Rao (LinkedIn)

Rajeev Rastogi (Yahoo! Research India)

Keishi Tajima (Kyoto University, Japan)

Xiaofang Zhou (University of Queensland, Australia)

Challenges and Vision Track Chair

Gerhard Weikum, Max Planck Institute for Informatics

Challenges and Vision Track Program Committee

Anastassia Ailamaki (EPF Lausanne)
Sihem Amer-Yahia (Yahoo! Research)
Philip Bernstein (Microsoft Research)
Michael Cafarella (University of Michigan)
Susan Davidson (University of Pennsylvania)
Laura Haas (IBM Almaden Research Center)
Meichun Hsu (HP Labs)
Hank Korth (Lehigh University)

VLDB 2011 Tutorial Program Co-Chairs

Qiong Luo, Hong Kong University of Science & Techn.
Gerome Miklau, University of Massachusetts – Amherst

PVLDB Information Director

Gerald Weber (University of Auckland)

VLDB 2011 Proceedings Chair

Uwe Röhm (University of Sydney)

Steering Committee

Serge Abiteboul, Peter Apers, Philip Bernstein, Elisa Bertino, Peter Buneman, Martin Kersten, Z. Meral Ozsoyuglu

LETTER FROM THE INDUSTRIAL TRACK CHAIRS

As data volumes increase dramatically, and grid- and cloud-processing techniques are adopted by more and more industrial and government entities, the challenges for processing, migrating and searching this data are multiplying. As a result, a variety of platforms, development environments and management tools are required, each with a unique set of technical challenges. This year's Industrial Track of VLDB is a snapshot of this development from several areas of industry. In particular, these proceedings include several papers on managing workflow, programming environments, debugging and analysis task on grid frameworks. In addition, there is interesting work on replaying workloads for testing, data warehousing, clinical data management, email processing and data cleaning. The papers were drawn from a strong body of submissions, and on behalf of the program committee, we would like to thank the authors for the obvious effort put into their submissions.

Berthold Reinwald, IBM Almaden Research Center
Phil Bohannon, Yahoo! Research
VLDB 2011 Industrial, Applications, and Experience Track Chairs