Advances in Database Technology — EDBT 2020

23rd International Conference
on Extending Database Technology
Copenhagen, Denkmark, March 30–April 2, 2020
Proceedings

Editors

Angela Bonifati
Yongluan Zhou
Marcos Antonio Vaz Salles
Alexander Böhm
Dan Olteanu
George Fletcher
Arijit Khan
Bin Yang

open proceedings

Advances in Database Technology — EDBT 2020
Proceedings of the 23rd International Conference
on Extending Database Technology
Copenhagen, Denkmark, March 30–April 2, 2020

Series ISSN: 2367-2005

Editors

Angela Bonifati, Lyon 1 University, France
Yongluan Zhou, University of Copenhagen, Denmark
Marcos Antonio Vaz Salles, University of Copenhagen, Denmark
Alexander Böhm, SAP Research, Germany
Dan Olteanu, University of Oxford, United Kingdom
George Fletcher, Eindhoven University of Technology, The Netherlands
Arijit Khan, Nanyang Technnological University, Singapore
Bin Yang, Aalborg University, Denmark

OpenProceedings.org
University of Konstanz
University Library
78457 Konstanz, Germany

OpenProceedings ISBN: 978-3-89318-083-7

DOI of this front matter: 10.5441/002/edbt.2020.01

Foreword

The International Conference on Extending Database Technology (EDBT) is a leading international forum for database researchers, developers, and users to present and discuss novel and cutting-edge ideas and techniques, and to showcase tools and experiences related to data management. Data management is an essential enabling technology that has applications in several scientific, business and social communities, and runs on diverse technical platforms associated with the web, enterprises, clouds and mobile devices. The database community has a continuing tradition of contributing with models, algorithms and architectures to the set of tools and applications that enable day-to-day functioning of our societies. Faced with the broad challenges of today's applications, data management technology constantly broadens its reach, exploiting new hardware and software to achieve innovative results and embracing new challenges in the years to come.

EDBT 2020 solicited submissions of original research contributions, descriptions of industrial solutions and applications, and proposals for tutorials and software demonstrations. We encouraged submissions of research papers related to all aspects of data management. We also encouraged submissions of visionary papers as well as innovative system papers and experimental analyses papers. In addition to long research paper submissions, EDBT 2020 again encouraged the submission of short research papers, which provide an excellent opportunity to describe significant work or research in progress that can foster the discussions at the conference. Short papers are presented as posters at plenary poster sessions of the conference. This year, they will also be communicated at a plenary lightning talks session.

The program committees of EDBT accepted 30 out of 151 submitted regular research papers, resulting in an acceptance rate of 20% for the research track; 26 out of 85 submitted short research papers, resulting in an acceptance rate of 31% for short research papers; 16 out of 37 demos, resulting in an acceptance rate of 43% for the demonstration track; and 10 out of 37 industrial and application papers, resulting in an acceptance rate of 27% for industrial and application papers. The papers will be presented in eight research paper sessions, three industrial and application sessions, as well as two plenary poster and demonstration sessions.

The program additionally features five workshops, one of which is the well-established DOLAP workshop that has successfully been co-located with EDBT since many years. Finally, the conference program includes four tutorials and an EDBT and ICDT joint session on climate change.

I would like to thank all authors for their contributions: a successful conference crucially depends on high-quality submissions. I also would like to thank all senior reviewers and reviewers for serving on the EDBT 2020 program committee, in particular for the high quality and timely handling of all reviews and discussions. This community service requires a lot of work on a tight schedule, and is what makes our research community function and ensures the sustained impact of our research. Thanks to their valuable effort we can look forward to an exciting program and attractive EDBT conference in Copenhagen from March 30–April 2, 2020.

A warm thanks to Anastasia Alaimaki and Tamer Özsu for serving on the Test-of-Time Award committee to select the paper from EDBT 2010 that has had the most lasting influence. Wook-shin Han, Erhard Rahm and Nesime Tatbul generously accepted to serve on the Best Paper committee. The EDBT 2020 program is the result of the joint effort of many people who shared their experience and time to contribute to the EDBT 2020 program and make the conference a great success. Alexander Boehm served as PC chair for industrial and application papers; George Fletcher as PC chair for the demonstration track; Dan Olteanu as tutorial chair; Alexandra Poulovassilis as workshop chair. My warmest thanks to all these people. The general chairs, Yongluan Zhou and Marcos Vaz Salles and the other local organizers worked hard to make all necessary arrangements for a successful event. Special thanks to Arijt Khan, the EDBT proceedings chair; Davide Mottin, the publicity chair; Bin Yang, the local executive chair; Boris Düdder, the sponsorship chair; Nanna Højholt, the finance chair and Yiwen Wang, the website chair, for tirelessly finding solutions for all our requests. Norman Paton was most helpful in advising and coordinating with the EDBT Executive Board. Last but not least, I would like to thank Marc H. Scholl for assembling the EDBT proceedings on openproceedings.org I hope that you find EDBT 2020 inspiring, enriching, and enjoyable and look forward to meeting you in Copenhagen.

Angela Bonifati
EDBT 2020 Program Chair

Program Committee Members

Research Program Committee

Angela Bonifati (Lyon 1 U, France) – Chair

Senior Program Committee Members

Karl Aberer (EPFL Lausanne, Switzerland)
Walid Aref (Purdue U, USA)
Michael Benedikt (U Oxford, UK)
Michael Böhlen (U Zurich, Switzerland)
K. Selcuk Candan (Arizona State U, USA)
Kevin Chang (U Illinois at Urbana-Champaign, USA)
Vassilis Christophides (INRIA, France)
Daniel Deutch (Tel Aviv U, Israel)
Floris Geerts (U Antwerp, Belgium)
Jan Hidders (VU Brussel, Belgium)
Katja Hose (Aalborg U, Denmark)
Christoph Koch (EPFL Lausanne, Switzerland)
Georgia Koutrika (Athena Research Center, Greece)
Ulf Leser (Humboldt-U Berlin, Germany)
Guoliang Li (Tsinghua U, China)
Chengkai Li (U Texas at Arlington, USA)
Eric Lo (Chinese U Hong Kong, China)
Evaggelia Pitoura (U Ioannina, Greece)
Louiqa Raschid (U Maryland, USA)
Sherif Sakr (U Tartu, Estonia)
Semih Salihoglu (U Waterloo, Canada)
Kai-Uwe Sattler (TU Ilmenau, Germany)
Arash Termehchy (Oregon State U, USA)
Riccardo Torlone (Roma Tre U, Italy)
Peter Triantafillou (U Warwick, UK)
Yannis Velegrakis (Utrecht U, The Netherlands)

Program Committee Members

Bernd Amann (Sorbonne U – LIP6, France)
Akhil Arora (EPFL Lausanne, Switzerland)
Elena Baralis (Politecnico di Torino, Italy)
Denilson Barbosa (U Alberta, Canada)
Senjuti Basu Roy (New Jersey Inst. of Techn., USA)
Luigi Bellomarini (Banca d'Italia, Italy)
Sonia Bergamaschi (U Modena Reggio Emilia, Italy)
Laure Berti-Equille (IRD, France)
Arnab Bhattacharya (IIT Kanpur, India)
Luc Bouganim (INRIA-UVSQ, France)
Andrea Calì (U London, Birkbeck College, UK)
Bogdan Cautis (Paris Sud U, France)
Lei Chen (Hong Kong U Sc. & Techn., China)
Dario Colazzo (Paris Dauphine U, France)
Bin Cui (Peking U, China)
Alfredo Cuzzocrea (U Calabria, Italy)
Sabrina De Capitani di Vimercati (U Milan, Italy)
Stefania Dumbrava (ENSIIE Rennes, France)
Donatella Firmani (Roma Tre U, Italy)
Rainer Gemulla (U Mannheim, Germany)
Paolo Guagliardo (U Edinburgh, UK)
Xi He (U Waterloo, Canada)
Xin Huang (Hong Kong Baptist U, China)
Zsolt Istvan (IMDEA Software Institute, Spain)
Panos Kalnis (King Abdullah UST, Saudi Arabia)
Vana Kalogeraki (Athens U Ec. & Busin., Greece)
Verena Kantere (U Ottawa, Canada)
Yaron Kanza (AT&T Labs – Research, USA)
Asterios Katsifodimos (Delft UT, The Netherlands)
Xiang Lian (Kent State U, USA)
Ping Lu (Beihang U, China)
Paolo Missier (Newcastle U, USA)
Davide Mottin (Aarhus U, Denmark)
Behrooz Omidvar-Tehrani (Grenoble Alpes, France)
Eric Peukert (Leipzig U, Germany)
Holger Pirk (Imperial College London, UK)
Dimitris Plexousakis (Institute of CS, FORTH, Greece)
Giuseppe Polese (U Salerno, Italy)
Arnau Prat (U Politècn. Catalunya, Spain)
Mohammad Sadoghi (UC Davis)
Carlo Sartiani (U della Basilicata, Italy)
Stefanie Scherzinger (OTH Regensburg, Germany)
Petra Selmer (Neo4j, UK)
Juan F. Sequeda (Capsenta Labs, USA)
Hala Skaf-Molli (U Nantes, France)
Kostas Stefanidis (U Tampere, Finland)
Gábor Szárnyas (Budapest U Tech. & Eco., Hungary)
Ernest Teniente (U Politècn. Catalunya, Spain)
Jens Teubner (TU Dortmund, Germany)
Farouk Toumani (U Clermont Auvergne, France)
Anthony K. H. Tung (National U Singapore)
Wendy Hui Wang (Stevens Inst. of Techn., USA)
Nikolay Yakovets (TU Eindhoven, The Netherlands)
Demetrios Zeinalipour (U Cyprus, Greece)

Industrial/Applications Program Committee

Roy (Sudipto) Chowdhuri (Salesforce)
Colin Florendo (Google)
Prasanta Ghosh (Microsoft)
Fisnik Kastrati (Huawei)
Martin Kaufmann (Teradata)
Justin Levandoski (Amazon)
Stefan Mandl (Exasol)
Norman May (SAP SE)
Ismail Oukid (TU Dresden)
Fatma Ozcan (IBM)
Orestis Polychroniou (Amazon)
Danica Porobic (Oracle)
Harald Schöning (Software AG)
Peter Thawley (Amazon)
Adrian Vogelsgesang (Tableau)
Hannes Voigt (Neo4j)
Till Westmann (Couchbase)
Panfeng Zhou (Alibaba)

Demonstration Program Committee

Ashvin Agrawal (Microsoft Research, USA)
Toshiyuki Amagasa (U Tsukuba, Japan)
Sihem Amer-Yahia (CNRS LIG, France)
Sourav Bhowmick (Nanyang Techn. U, Singapore)
Carsten Binnig (TU Darmstadt, Germany)
Chee-Yong Chan (National U Singapore)
Adriane Chapman (U Southampton, UK)
Dong-Wan Choi (Inha U, South Korea)
Rick Cole (Tableau Software, Inc., USA)
Vasiliki Kalavri (ETH Zurich, Switzerland)
Romans Kasperovics (SAP, Germany)
Wolfgang Lehner (TU Dresden, Germany)
Hannes Mühleisen (CWI, Netherlands)
Hubert Naacke (Sorbonne U, France)
Keisuke Nakano (U Tohoku, Japan)
Jorge Pérez (U Chile)
Antonella Poggi (Sapienza U Roma, Italy)
Cristian Riveros (Pontificia U Catolica Chile)
Jagan Sankaranarayanan (Google, USA)
Yuya Sasaki (Osaka U, Japan)
Michael Schmidt (Amazon, USA)
Sławek Staworko (U Lille 3, France)
Dimitri Surinx (Hasselt U, Belgium)
Alex Thomo (U Victoria, Canada)
Xiaolan Wang (Megagon Labs, USA)
Raymond Chi-Wing Wong (Hong Kong UST)
Meihui Zhang (Beijing Inst. Techn., China)

Conference Organization

General Chairs
Yongluan Zhou, University of Copenhagen, Denmark
Marcos Antonio Vaz Salles, University of Copenhagen, Denmark

EDBT Program Chair
Angela Bonifati, Lyon 1 University, France

ICDT Program Chair
Carsten Lutz, University of Bremen, Germany

EDBT Industrial/Application Chair
Alexander Böhm, SAP Research, Germany

EDBT Demonstrations Chair
George Fletcher, Eindhoven University of Technology, The Netherlands

Tutorial Chair
Dan Olteanu, University of Oxford, United Kingdom

Workshops Chair
Alex Poulovassilis, Birbeck University of London, United Kingdom

EDBT Proceedings Chair
Arijt Khan, Nanyang Technnological University, Singapore

ICDT Proceedings Chair
Jean Christoph Jung, University of Bremen, Germany

Local Executive Chair
Bin Yang, Aalborg University, Denmark

Sponsorship Chairs
Boris Düdder, University of Copenhagen, Denmark

Publicity Chair
Davide Mottin, Aarhus University, Denmark

Finance Chair
Nanna Højholt, University of Copenhagen, Denmark

Website Chair
Yiwen Wang, University of Copenhagen, Denmark

Test-of-Time Award

Established in 2014, the Test-of-Time Award awarded by the Extended Database Technology (EDBT) Conference recognizes papers presented at EDBT Conferences that have had the most impact in terms of research, methodology, conceptual contribution, or transfer to practice over the past ten years.
The EDBT 2020 Test of Time Award committee was formed by Anastasia Ailamaki (Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland), Tamer Oszu (University of Waterloo, Canada), and Angela Bonifati (Lyon 1 University, France).

After careful consideration, the committee has decided to select the following paper from the EDBT 2010 conference as the EDBT ToT Award winner for 2020:

Optimizing joins in a map-reduce environment

by Foto Afrati and Jeff Ullman

published in EDBT 2010 Proceedings, pp. 99–110, DOI: 10.1145/1739041.1739056.

This paper presented optimization strategies for executing multi-way joins in a map-reduce environment. It focused on large-scale data and provided algorithms to choose the number of map-keys and shares in order to minimize the communication cost among the map and reduce processes.
The committee members agreed that this paper clearly pioneered the field of join processing in map-reduce environments. It has triggered substantial follow-up research and impact on big data processing in parallel and distributed architectures.

The EDBT Test-of-Time award for 2020 will be presented during the EDBT/ICDT 2020 Conference in Copenhagen, Denmark, as part of the Awards session on March 31, 2020.

Best Paper Award

The EDBT 2020 Best Paper Award committee was formed by Wook-shin Han (Postech, Korea), Erhard Rahm (University of Leipzig, Germany), Nesime Tatbul (Intel & MIT, USA), and Angela Bonifati (Lyon 1 University, France). After careful consideration, the committee has decided to select the following paper as the EDBT Best Paper for 2020:

Provenance for Probabilistic Logic Programs

by Shaobo Wang, Hui Lyu, Jaichi Zhang, Chenyuan Wu, Xinyi,Chen, Wenchao,Zhou, Boon Thau Loo, Susan B. Davidson, Chen Chen.

DOI: 10.5441/002/edbt.2020.14

Abstract: Despite the emergence of probabilistic logic programming (PLP) languages for data driven applications, there are currently no debugging tools based on provenance for PLP programs. In this paper, we propose a novel provenance model and system, called P3 (Provenance for Probabilistic logic Programs) for analyzing PLP programs. P3 enables four types of provenance queries: tra- ditional explanation queries, queries for finding the set of most important derivations within an approximate error, top-K most influential queries, and modification queries that enable us to modify tuple probabilities with fewest modifications to program or input data. We apply these queries into real-world scenar- ios and present theoretical analysis and practical algorithms for such queries. We have developed a prototype of P3, and our evaluation on real-world data demonstrates that the system can support a wide-range of provenance queries with explainable results. Moreover, the system maintains provenance and execute queries efficiently with low overhead.

The EDBT Best Paper Awards for 2020 will be presented during the EDBT/ICDT 2020 Conference in Copenhagen, Denmark, as part of the Awards session on March 31, 2020.

Table of Contents

Short papers

Industry and Applications Papers

Demonstrations

Tutorials