June 30–July 5, 2019
Amsterdam, Netherlands

Association for Computing Machinery

Advancing Computing as a Science & Profession

SIGMOD/PODS 2019
AMSTERDAM

SIGMOD '19

Proceedings of the 2019 International Conference on Management of Data

General Chairs:
Peter Boncz & Stefan Manegold (CWI, Netherlands)

Program Chair:
Anastasia Ailamaki (EPFL, Switzerland)

Program Vice-Chairs:
Amol Deshpande (University of Maryland, USA)
Tim Kraska (MIT, USA)

Association for
Computing Machinery

Advancing Computing as a Science & Profession

The Association for Computing Machinery
1601 Broadway, 10th Floor
New York, NY 10019-7434

ISBN: 978-1-4503-5643-5

Additional copies may be ordered prepaid from:

ACM Order Department
PO Box 30777
New York, NY 10087-0777, USA

Phone: 1-800-342-6626 (USA and Canada)
+1-212-626-0500 (Global)
Fax: +1-212-944-1318
E-mail: acmhelp@acm.org
Hours of Operation: 8:30 am – 4:30 pm ET

Printed in the USA

Welcome to SIGMOD 2019 - The 2019 ACM SIGMOD International Conference on the Management of Data!

This year, the conference is held in the city center of Amsterdam, capital of The Netherlands. Amsterdam is an internationally oriented city, home to people with origins from all over the world. This used to be already the case even back in the 16th and 17th century, when Amsterdam was the world's biggest trading and financial center; establishing the world's first stock exchange in 1602.

SIGMOD/PODS 2019 is held in the original Amsterdam Stock and Commodities Exchange, constructed by Dutch architect Berlage between 1896 and 1903, which now serves as the well-equipped Amsterdam Conference Center. This architect and his apprentices (the school of Berlage) left an important mark on the city, being responsible for a major expansion of the city in the early 20th century. The sculptures and drawings in the Exchange building refer to the people behind the commodities traded in the various rooms ("Effecten" - stock; "Graan" - grains), e.g., depicting farmers in the grain exchange room; as a reminder that trading affects society. Amsterdam is a city that offers many cultural activities, including the world-famous classical Concertgebouw Orchestra, as well as many museums (Amsterdam Museum, Rijksmuseum, Rembrandthuis, Anne Frank Huis). In a slight break with SIGMOD tradition, the SIGMOD opening reception will be held one day later, on Tuesday night, when the SIGMOD/PODS attendees will have exclusive access to the Van Gogh museum. The Wednesday conference dinner is organized across the water in Amsterdam North, in Noorderlicht Cafe in a festival-like environment. This used to be harbour area and was less-populated and industrial, but in the recent decade has become a hotspot for nightlife activities.

Amsterdam is also increasingly a hub for data science companies and services, with multiple universities and CWI in the vicinity; which all participate in the organization of SIGMOD/PODS 2019. On Thursday night, after the SIGMOD program finishes, there will be a meetup of Amsterdam Data Science, where the local data science community will be able to mingle with our data management research community.

The SIGMOD 2019 Research Program Committee consists of the Program Chair, two Program Vice Chairs, a core committee with 37 members, and a regular committee with 98 members. During the reviewing period, we solicited additional reviews from 16 external reviewers and occasional input from 10 assistant reviewers. The committee received 430 submissions, out of which 12 were desk-rejected (i.e., without review). There was no bidding; instead, reviewer assignments were made using input from Microsoft's Conference Management System, the Toronto Paper Matching System, and the reviewers' background (the detailed assignment procedure is described in a paper which has been submitted for publication to SIGMOD Record). The core committee members had (roughly) double the reviewing load of the regular committee members, and in addition acted as discussion leaders and meta-reviewers for their assigned papers. There were two rounds of submissions, with deadlines in July and November, respectively. Initially, each paper received three reviews. At this point authors could read the reviews and provide feedback about potential factual errors (disclosed to the reviewers) or sensitive issues about potential mishandling (confidentially to the chair). Two additional reviews were solicited for a paper if (a) the reviewers' expertise level was suboptimal, or (b) if there was significant score discrepancy in the first three reviews, or (c) if it was heading for rejection but

had received a weak accept (or higher) by at least one reviewer. Papers were discussed extensively online; 10 were accepted based on the first round of reviews, while 311 were rejected. The authors of the remaining 97 papers were asked to revise their papers to address reviewers' criticisms; 78 revisions were ultimately accepted for a total of 88 papers which are presented in the research track. Finally, 12 papers were shepherded after acceptance to guarantee that the camera-ready version addresses all of the reviewers' comments.

Peter Boncz, Stefan Manegold
SIGMOD'19 General Chairs
CWI, Netherlands

Anastasia Ailamaki
SIGMOD'19 Program Chair
EPFL, Switzerland

Table of Contents

Research 3: Information Extraction

Session Chair: Guoliang Li

Industry 1: Data Applications

Session Chair: Marco Serafini

SIGMOD Keynote 2

Session Chair: Stefan Manegold

Panel

Research 4: Distributed Data Management

Session Chair: Holger Pirk

Research 5: Provenance

Session Chair: Alexandra Meliou

Research 6: Streams

Session Chair: Jonathan Goldstein

Industry 2: Storage & Indexing

Session Chair: Alexander Shraer

Research 7: Modern Hardware

Session Chair: Justin Lewandowski

Research 8: Data Integration/Cleaning

Session Chair: Paolo Papotti

Research 9: Query Processing & Optimization 2

Session Chair: Jun Yang

Research 10: Graphs 1

Session Chair: Angela Bonifati

Award Talks

Research 11: Systems & Machine Learning

Session Chair: Matthias Boehm

Research 12: Indexing

Session Chair: Stratos Idreos

Research 13: Fairness, Uncertainty

Session Chair: Ke Yi

Research 14: Graphs 2

Session Chair: Sourav S Bhowmick

Research 15: Graphs 3

Session Chair: Xuemin Lin

Research 16: Machine Learning

Session Chair: Theodoros Rekatsinas

Research 17: Scalability

Session Chair: Norman May

Industry 3: Data Platforms

Session Chair: Ying Zhang

Student Abstracts

Demonstrations

Tutorials

Workshop Summaries

SIGMOD 2019 Organization

SIGMOD Program Chair:
Anastasia Ailamaki (EPFL, Switzerland)

SIGMOD Program Vice-Chairs:
Amol Deshpande (University of Maryland, USA)
Tim Kraska (MIT, USA)

SIGMOD General Chairs:
Peter Boncz (CWI & Vrije Universiteit Amsterdam, The Netherlands)
Stefan Manegold (CWI & Universiteit Leiden, The Netherlands)

SIGMOD Honorary Chair:
Martin Kersten (CWI & Universiteit van Amsterdam, The Netherlands)

SIGMOD Local Arrangements/Organization Chairs:
George Fletcher (Eindhoven University of Technology, The Netherlands)
Asterios Katsifodimos (Delft University of Technology, The Netherlands)

SIGMOD Sponsorship Chairs:
Semih Salihoglu (University of Waterloo, Canada)
Tilmann Rabl (HPI Potsdam, Germany)

SIGMOD Finance Chair:
Hannes Mühleisen (CWI, Amsterdam & Vrije Universiteit Amsterdam, The Netherlands)

SIGMOD Registration Chair:
Maurice van Keulen (University of Twente, Enschede, The Netherlands)

SIGMOD Publicity/Social Media Chairs:
Jan Hidders (Vrije Universiteit Brussel, Belgium)
Torsten Grust (Universität Tübingen, Germany)

SIGMOD Web/Information Chair:
Holger Pirk (Imperial College London, UK)

SIGMOD Mentorship Chairs:
Parth Nagarkar (New Mexico State University, USA),
Qiong Luo (Hong Kong University of Science & Technology, Hong Kong, China)

SIGMOD Proceedings Chair:
Ziawasch Abedjan (TU Berlin, Germany)

SIGMOD Tutorials Chair:
Ioana Manolescu (INRIA)
Hakan Hacigumus (Google, USA)

SIGMOD New Researcher Symposium Chairs:
Katja Hose (Aalborg University, Denmark)
Spyros Blanas (Ohio State University, USA)

SIGMOD Workshops Chairs:
Ihab Ilyas (University of Waterloo, CA)
Angela Bonifati (Lyon 1 University, FR)
Benny Kimelfeld (Technion, IL)

SIGMOD Undergrad Research Contest:
Jana Giceva (Imperial College London, UK)
Eugene Wu (Columbia University, USA)

SIGMOD Programming Contest:
Ravi Rajwar (Intel, USA)
Pınar Tözün (IT University of Copenhagen, Denmark)

SIGMOD Core Program Committee:
Ashraf Aboulnaga (Qatar Computing Research Institute, Qatar)
Azza Abouzied (New York University Abu Dhabi, United Arab Emirates)
Gustavo Alonso (ETH Zürich, Switzerland)
Peter Alvaro (University of California, Santa Cruz, USA)
Sourav Bhowmick (Nanyang Technological University, Singapore)
Carsten Binnig (Brown University, USA)
Matthias Boehm (IBM, USA)
Angela Bonifati (University of Lyon, France)
K. Selcuk Candan (Arizona State University, USA)
Gautam Das (University of Texas, Arlington, USA)

Daniel Deutch (Tel Aviv University, Israel)
Johannes Gehrke (Microsoft, USA)
Jonathan Goldstein (Microsoft, USA)
Bill Howe (University of Washington, USA)
Stratos Idreos (Harvard University, USA)
Ihab Ilyas (University of Waterloo, Canada)
Zachary Ives (University of Pennsylvania, USA)
Chris Jermaine (Rice University, USA)
Raghav Kaushik (Microsoft, USA)
Martin Kersten (CWI, Netherlands)
Wolfgang Lehner (Technical University of Dresden, Germany)
Justin Levandoski (Microsoft, USA)
Feifei Li (University of Utah, USA)
Xuemin Lin (University of New South Wales, Australia)
Alexandra Meliou (University of Massachusetts Amherst, USA)
Paolo Papotti (Eurecom, France)
Srinivasan Parthasarathy (Ohio State University, USA)
Holger Pirk (Imperial College London, UK)
Evaggelia Pitoura (University of Ioannina, Greece)
Theodoros Rekatsinas (Stanford University, USA)
Dan Suciu (University of Washington, USA)
Jens Teubner (Technical University of Dortmund, Germany)
Marcos Antonio Vaz Salles (University of Copenhagen, Denmark)
Jiannan Wang (Simon Fraser University, Canada)
Jun Yang (Duke University, USA)
Ke Yi (Hong Kong University of Science & Technology, Hong Kong)
Nan Zhang (Pennsylvania State University, USA

SIGMOD Regular Program Committee:

Ioannis Alagiannis (Microsoft, USA)
Foteini Alvanaki (EPFL, Switzerland)
Nicolas Anciaux (French Institute for Research in Computer Science and Automation, France)
Raja Appuswamy (Eurecom, France)
Manos Athanassoulis (Harvard University, USA)
Peter D Bailis (Stanford University, USA)
Alex Beutel (Google, USA)
Spyros Blanas (Ohio State University, USA)
Alexander Boehm (SAP SE, Germany)
Renata Borovica (University of Melbourne, Australia)
Lei Cao (MIT, USA)
Raul Castro Fernandez (MIT, USA)
Badrish Chandramouli (Microsoft, USA)
Lei Chen (University of Santo Tomas, Philippines)
Shimin Chen (University of Chinese Academy of Sciences, China)
Reynold Cheng (Univ. of Hong Kong, Hong Kong)
James Cheng (Chinese Univ. of Hong Kong, Hong Kong)
Xu Chu (Georgia Institute of Technology, USA)
Carlo Curino (Microsoft, USA)
Mahashweta Das (Hewlett-Packard, USA)
Sudipto Das (Microsoft, USA)
Chris De Sa (Cornell University, USA)
Dong Deng (Massachusetts Institute of Technology, USA)
AnHai Doan (University of Washington, USA)
Eduard Dragut (Temple University College of Education, USA)
Georgios Fakas (Uppsala University, Sweden)
Wenfei Fan (University of Edinburgh, UK)
Alan Fekete (University of Sydney, Australia)
George Fletcher (Eindhoven University of Technology, Netherlands)
Wolfgang Gatterbauer (Northeastern Univ., USA)
Jana Giceva (Imperial College London, UK)
Boris Glavic (Illinois Institute of Technology, USA)
Torsten Grust (University of Tübingen, Germany)
Alon Halevy (Megagon Labs, USA)
Michael Hay (Colgate University, USA)
Melanie Herschel (University of Stuttgart, Germany)
Vagelis Hristidis (UC Riverside, USA)
H. V. Jagadish (University of Michigan, USA)
Alekh Jindal (Microsoft, USA)
Ryan Johnson (Amazon Web Services)
Eser Kandogan (IBM, USA)
Manos Karpathiotakis (EPFL, Switzerland)
Alfons Kemper (Technical University of Munich, Germany)

Oliver Kennedy (University at Buffalo, USA)
Hideaki Kimura (Oracle, USA)
Donald Kossmann (Microsoft, USA)
Paraschos Koutris (University of Washington, USA)
Rajasekar Krishnamurthy (IBM, USA)
Arun Kumar (UC San Diego, USA)
Viktor Leis (Technical University of Munich, Germany)
Guoliang Li (Tsinghua University, China)
Chengkai Li (University of Texas, Arlington, USA)
Ashwin Machanavajjhala (Duke University, USA)
Sam Madden (MIT, USA)
Nikos Mamoulis (University of Ioannina, Greece)
Norman May (SAP SE, Germany)
Gerome Miklau (University of Massachusetts Amherst, USA)
Ingo Müller (ETH Zürich, Switzerland)
Thomas Neumann (Technical University of Munich, Germany)
Milos Nikolis (University of Oxford, UK)
Ismail Oukid (SAP SE, Germany)
Tamer Ozsu (University of Waterloo, Canada)
Themis Palpanas (Paris Descartes, France)
Ippokratis Pandis (Amazon, USA)
Charalampos Papamanthou (University of Maryland, USA)
Andy Pavlo (Carnegie Mellon University, USA)
Danica Porobic (Oracle, USA)
Dan Ports (University of Washington, USA)
Iraklis Psaroudakis (Oracle, Switzerland)
Ravi Ramamurthy (Microsoft, USA)
Mary Roth (IBM, USA)
Sudeepa Roy (Duke University, USA)
Sudip Roy (Google, USA)
Florin Rusu (UC Merced, USA)
Mohammad Sadoghi (UC Davis, USA)
Yasushi Sakurai (Kumamoto University, Japan)
Ken Salem (University of Waterloo, Canada)
Semih Salihoglu (University of Waterloo, Canada)
Mohamed Sarwat (Arizona State University, USA)

SIGMOD Regular Program Committee (continued):

Sebastian Schelter (New York University, USA)
Prithviraj Sen (IBM, USA)
Divesh Srivastava (AT&T, USA)
Radu Stoica (IBM, Switzerland)
Julia Stoyanovich (Drexel University, USA)
Wang-Chiew Tan (Megagon Labs, USA)
Nan Tang (Qatar Computing Research Institute, Qatar)
Yufei Tao (Chinese University of Hong Kong, Hong Kong)
Nesime Tatbul (MIT, USA)
Farhan Tauheed (Oracle, USA)
Arash Termehchy (Oregon State University, USA)
Martin Theobald (University of Luxembourg, Luxembourg)
Immanuel Trummer (Cornell University, USA)
Eugene Wu (Columbia University, USA)
Da Yan (University of Alabama, Birmingham, USA)
Cong Yu (Google, USA)
Ce Zhang (ETH Zürich, Switzerland)
Yongluan Zhou (University of Copenhagen, Denmark)
Wenchao Zhou (Georgetown University, USA)

SIGMOD External Program Committee:

Panagiotis Antonopoulos (Microsoft, USA)
Akhil Arora (EPFL, Switzerland)
Magdalena Balazinska (University of Washington, USA)
Stella Giannakopoulou (EPFL, Switzerland)
Georgia Koutrika (Athena Research Center, Greece)
Barzan Mozafari (University of Michigan, USA)
Rachel Pottinger (University of British Columbia, Canada)
Babak Salimi (University of Washington, USA)
Hanghang Tong (Arizona State University, USA)
Charalampos Tsourakakis (Boston University, USA)
Nalini Venkatasubramanian (University of California, Irvine, USA)
Yinghui Wu (Washington State University, USA)
Yaoliang Yu (University of Waterloo, Canada)
Matei Zaharia (Stanford University, USA)
Emanuel Zgraggen (Massachusetts Institute of Technology, USA)
Kostas Zoumpatianos (Harvard University, USA)

SIGMOD Assisting Program Committee:

Peng Cheng (Hong Kong University of Science & Technology, Hong Kong)
Stefania Dumbrava (INRIA, France)
Cibele Freire (Wellesley College, USA)

Sainyam Galhotra (University of Massachusetts, Amherst, USA)
Zeynep Korkmaz (University of Waterloo, USA)
Anil Pacaci (University of Waterloo, USA)
Aida Sheshbolouki (University of Waterloo, USA)
Yue Wang (Hong Kong University of Science & Technology, Hong Kong)
Xiaofei Zhang (University of Memphis, USA)
Libin Zheng (Hong Kong University of Science & Technology, Hong Kong)

SIGMOD Industrial Track PC Chairs:
Lyublena Antova (Datometry)
Jignesh Patel (UW Madison, USA)

SIGMOD Industrial Track PC Members:
Alexander Shraer (Apple, USA)
Allison Holloway (Oracle, USA)
Avrilia Floratou (Microsoft, USA)
Danica Porobic (Oracle, USA)
Entong Shen (Amazon Web Services, USA)
Jianjun Chen (Huawei US Research Center, USA)
Josiane Xavier Parreira (Siemens AG Austria)
Marco Serafini (University of Massachusetts Amherst, USA)
Mark Callaghan (Facebook, USA)
Markus Weimer (Microsoft)
Mohamed Soliman (Datometry, USA)
Per-Ake Larson (University of Waterloo, Canada)
Vaishnavi Sashikanth (Google, USA)
Ying Zhang (MonetDB Solutions, Netherlands)

SIGMOD Demo Track PC Chairs:
Thomas Heinis (Imperial College London, UK)
Fatma Ozcan (IBM Almaden, USA)

SIGMOD Demo Track PC Members:
Alexandros Koliousis (Imperial College London, UK)
Anisoara Nica (SAP SE, Waterloo)
Bertram Ludaescher (University of Illinois, USA)
Boris Glavic (Illinois Institute of Technology, USA)
Danica Porobic (Oracle)
Dong-Wan Choi (Inha University)
Eric Lo (Chinese University of Hong Kong, China)
Fabian Suchanek (Telecom Paris Tech Univ., France)
Fabio Porto (LNCC)
Florin Rusu (UC Merced, USA)
Georgia Koutrika (Athena Research Center)
Guoliang Li (Tsinghua University)
Hannes M√ºhleisen (CWI)
Huiping Cao (New Mexico State University, USA)
Iman Elghandour (University of Copenhagen)
Jaydeep Sen (IBM Research AI, USA)
Johann Gamper (Free Univ. of Bozen-Bolzano, Italy)
Kai Zheng (University of Electronic Science & Technology of China)
Katja Hose (Aalborg University)
Kyriakos Mouratidis (Singapore Management University)
Lei Chen (Hong Kong University of Science & Technology)
Manos Karpathiotakis (Facebook)
Maria Luisa Sapino (U. Torino)
Maya Ramanath (IIT Delhi)
Mitesh Vasa (IBM Watson - Almaden, USA)
Mohamed Eltabakh (Teradata)
Mohammad Sadoghi (UC Davis, USA)
Mustafa Canim (IBM Research)
Nan Tang (Qatar Computing Research Institute, HBKU)
Niketan Pansare (IBM Almaden, USA)
Odysseas Papapetrou (EPFL)
Olga Papaemmanouil (Brandeis University)
Parth Nagarkar (NMSU)
Pelin Angin (METU, Turkey)
Peter Fischer (University of Augsburg)
Renata Borovica-Gajic (University of Melbourne)
Ricardo Torres (IC-Unicamp)
Semih Salihoglu (University of Waterloo, Canada)
Senjuti Basu Roy (NJ Institute of Technology, USA)
Stefanie Scherzinger (OTH Regensburg)
Sudip Roy (Google)
Till Westmann (Couchbase)
Vasilis Efthymiou (IBM Research - Almaden, USA)
Verena Kantere (University of Ottawa, Canada)
Xiaochun Yang (Northeastern University, USA)
Xu Chu (GATECH, USA)
Yongxin Tong (Beihang University)

SIGMOD 2019 Sponsors & Supporters

Sponsors:

Diamond Supporter:

Platinum Supporters:

Gold Supporters:

Silver Supporters:

Startup Supporters:

Platinum Publisher:

Gold Publisher:

Silver Publishers:

Student Support:

Institutional Support: