

*Commenced Publication in 1973*

Founding and Former Series Editors:

Gerhard Goos, Juris Hartmanis, and Jan van Leeuwen

## Editorial Board

David Hutchison

*Lancaster University, Lancaster, UK*

Takeo Kanade

*Carnegie Mellon University, Pittsburgh, PA, USA*

Josef Kittler

*University of Surrey, Guildford, UK*

Jon M. Kleinberg

*Cornell University, Ithaca, NY, USA*

Friedemann Mattern

*ETH Zurich, Zurich, Switzerland*

John C. Mitchell

*Stanford University, Stanford, CA, USA*

Moni Naor

*Weizmann Institute of Science, Rehovot, Israel*

C. Pandu Rangan

*Indian Institute of Technology Madras, Chennai, India*

Bernhard Steffen

*TU Dortmund University, Dortmund, Germany*

Demetri Terzopoulos

*University of California, Los Angeles, CA, USA*

Doug Tygar

*University of California, Berkeley, CA, USA*

More information about this series at <http://www.springer.com/series/7408>

Raghunath Nambiar · Meikel Poess (Eds.)

# Performance Evaluation and Benchmarking for the Era of Artificial Intelligence

10th TPC Technology Conference, TPCTC 2018  
Rio de Janeiro, Brazil, August 27–31, 2018  
Revised Selected Papers

*Editors*

Raghunath Nambiar  
Advanced Micro Systems, Inc.  
Santa Clara, CA, USA

Meikel Poess  
Oracle Corporation  
Redwood Shores, CA, USA

ISSN 0302-9743                      ISSN 1611-3349 (electronic)  
Lecture Notes in Computer Science  
ISBN 978-3-030-11403-9              ISBN 978-3-030-11404-6 (eBook)  
<https://doi.org/10.1007/978-3-030-11404-6>

Library of Congress Control Number: 2018967046

LNCS Sublibrary: SL2 – Programming and Software Engineering

© Springer Nature Switzerland AG 2019

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, express or implied, with respect to the material contained herein or for any errors or omissions that may have been made. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

This Springer imprint is published by the registered company Springer Nature Switzerland AG  
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

# Preface

The Transaction Processing Performance Council (TPC) is a non-profit organization established in August 1988. Over the years, the TPC has had a significant impact on the computing industry's use of industry-standard benchmarks. Vendors use TPC benchmarks to illustrate performance competitiveness for their existing products, and to improve and monitor the performance of their products under development. Many buyers use TPC benchmark results as points of comparison when purchasing new computing systems.

The information technology landscape is evolving at a rapid pace, challenging industry experts and researchers to develop innovative techniques for evaluation, measurement and characterization of complex systems. The TPC remains committed to developing new benchmark standards to keep pace with these rapid changes in technology. One vehicle for achieving this objective is the TPC's sponsorship of the Technology Conference Series on Performance Evaluation and Benchmarking (TPCTC) established in 2009. With this conference series, the TPC encourages researchers and industry experts to present and debate novel ideas and methodologies in performance evaluation, measurement, and characterization.

The First TPC Technology Conference on Performance Evaluation and Benchmarking (TPCTC 2009) was held in conjunction with the 35th International Conference on Very Large Data Bases (VLDB 2009) in Lyon, France, during August 24–28, 2009.

The Second TPC Technology Conference on Performance Evaluation and Benchmarking (TPCTC 2010) was held in conjunction with the 36th International Conference on Very Large Data Bases (VLDB 2010) in Singapore during September 13–17, 2010.

The Third TPC Technology Conference on Performance Evaluation and Benchmarking (TPCTC 2011) was held in conjunction with the 37th International Conference on Very Large Data Bases (VLDB 2011) in Seattle, Washington, USA, during August 29 – September 3, 2011.

The 4th TPC Technology Conference on Performance Evaluation and Benchmarking (TPCTC 2012) was held in conjunction with the 38th International Conference on Very Large Data Bases (VLDB 2012) in Istanbul, Turkey, during August 27–31, 2012.

The 5th TPC Technology Conference on Performance Evaluation and Benchmarking (TPCTC 2013) was held in conjunction with the 39th International Conference on Very Large Data Bases (VLDB 2013) in Riva del Garda, Italy, during August 26–30, 2013.

The 6th TPC Technology Conference on Performance Evaluation and Benchmarking (TPCTC 2014) was held in conjunction with the 40th International Conference on Very Large Data Bases (VLDB 2014) in Hangzhou, China, during September 1–5, 2014.

The 7th TPC Technology Conference on Performance Evaluation and Benchmarking (TPCTC 2015) was held in conjunction with the 41st International Conference on Very Large Data Bases (VLDB 2015) in Kohala Coast, USA, during August 31 – September 4, 2015.

The 8th TPC Technology Conference on Performance Evaluation and Benchmarking (TPCTC 2016) was held in conjunction with the 42nd International Conference on Very Large Data Bases (VLDB 2016) in New Delhi, India, during September 5–9, 2016.

The 9th TPC Technology Conference on Performance Evaluation and Benchmarking (TPCTC 2017) was held in conjunction with the 43rd International Conference on Very Large Data Bases (VLDB 2017) in Munich, Germany, during August 28 – September 1, 2017.

This book contains the proceedings of the 10th TPC Technology Conference on Performance Evaluation and Benchmarking (TPCTC 2018), held in conjunction with the 44th International Conference on Very Large Data Bases (VLDB 2018) in Rio de Janeiro, Brazil, from August 27th to August 31st, 2018.

The hard work and close cooperation of a number of people have contributed to the success of this conference. We would like to thank the members of TPC and the organizers of VLDB 2018 for their sponsorship; the members of the Program Committee and Publicity Committee for their support; and the authors and the participants who are the primary reason for the success of this conference.

December 2018

Raghunath Nambiar  
Meikel Poess

# **TPCTC 2018 Organization**

## **General Chairs**

Raghunath Nambiar, Cisco, USA

Meikel Poess, Oracle, USA

## **Program Committee**

Daniel Bowers, Gartner, USA

Michael Brey, Oracle, USA

Paul Cao, HPE, USA

Alain Crolotte, Teradata Corporation, USA

Ajay Dholakia, Lenovo, USA

Karthik Kulkarni, Cisco, USA

Dhableswar Panda, The Ohio State University, USA

Tilmann Rabl, TU Berlin, Germany

Reza Taheri, VMware, USA

## **Publicity Committee**

Raghunath Nambiar, Cisco, USA

Andrew Bond, Red Hat, USA

Paul Cao, HPE, USA

Vo Ngoc Phu, Duy Tan University, Vietnam

Meikel Poess, Oracle, USA

Reza Taheri, VMware, USA

Michael Majdalany, L&M Management Group, USA

Forrest Carman, Owen Media, USA

Andreas Hotea, Hotea Solutions, USA

# About the TPC

## Introduction to the TPC

The Transaction Processing Performance Council (TPC) is a non-profit organization focused on developing industry standards for data-centric workloads and disseminating vendor-neutral performance data to industry. Additional information is available at <http://www.tpc.org/>.

## TPC Memberships

### Full Members

Full Members of the TPC participate in all aspects of the TPC's work, including development of benchmark standards and setting strategic direction. The Full Member application can be found at <http://www.tpc.org/information/about/app-member.asp>.

### Associate Members

Certain organizations may join the TPC as Associate Members. Associate Members may attend TPC meetings, but are not eligible to vote or hold office. Associate membership is available to non-profit organizations, educational institutions, market researchers, publishers, consultants, governments, and businesses that do not create, market, or sell computer products or services. The Associate Member application can be found at <http://www.tpc.org/information/about/app-assoc.asp>.

### Academic and Government Institutions

Academic and government institutions are invited to join the TPC and a special invitation can be found at <http://www.tpc.org/information/specialinvitation.asp>.

## Contact the TPC

TPC

Presidio of San Francisco

Building 572B (surface)

P.O. Box 29920 (mail)

San Francisco, CA 94129-0920

Voice: (+1)415-561-6272

Fax: (+1)415-561-6120

E-mail: [info@tpc.org](mailto:info@tpc.org)



## **How to Order TPC Materials**

All of our materials are now posted free of charge on our website. If you have any questions, please feel free to contact our office directly or by e-mail at [info@tpc.org](mailto:info@tpc.org).

## **Benchmark Status Report**

The TPC Benchmark Status Report is a digest of the activities of the TPC and its technical subcommittees. Sign-up information can be found at the following URL: <http://www.tpc.org/information/about/email.asp>.

# TPC 2018 Organization

## Full Members

Actian  
Alibaba  
AMD  
Cisco  
Dell EMC  
DataCore  
Fujitsu  
Hewlett Packard Enterprise  
Hitachi  
Huawei  
IBM  
Inspur  
Intel  
Lenovo  
Microsoft  
Nutanix  
Oracle  
Pivotal  
Red Hat  
SAP  
Teradata  
Transwarp  
TTA  
VMware

## Associate Members

IDEAS International  
University of Coimbra, Portugal  
China Academy of Information and Communications Technology

## Steering Committee

Andrew Bond, Red Hat, USA  
Michael Brey (Chair), Oracle, USA  
Matthew Emmerton, IBM, USA  
Raghunath Nambiar, AMD, USA  
Jamie Reding, Microsoft, USA

## **Public Relations Committee**

Andrew Bond, Red Hat, USA  
Paul Cao, HPE, USA  
Gary Little (Chair), Nutanix  
Raghunath Nambiar, AMD, USA  
Meikel Poess, Oracle, USA  
Reza Taheri, VMware, USA

## **Technical Advisory Board**

Paul Cao, HPE, USA  
Matthew Emmerton, IBM, USA  
Gary Little, Nutanix  
Jamie Reding (Chair), Microsoft, USA  
Da-Qi Ren, Huawei, USA  
Ken Rule, Intel, USA  
Nicholas Wakou, Dell, USA

## **Technical Subcommittees and Chairs**

TPC-C: Jamie Reding, Microsoft, USA  
TPC-H: Meikel Poess, Oracle, USA  
TPC-E: Matthew Emmerton, IBM, USA  
TPC-DS: Meikel Poess, Oracle, USA  
TPC-DI: Meikel Poess, Oracle, USA  
TPCx-HS: Tariq Magdon-Ismail, VMware, USA  
TPCx-BB: Bhaskar Gowda, Intel, USA  
TPCx-V: Reza Taheri, VMware, USA  
TPCx-HCI: Reza Taheri, VMware, USA  
TPCx-IoT: Karthik Kulkarni, Cisco, USA  
TPC-Pricing: Jamie Reding, Microsoft, USA  
TPC-Energy: Paul Cao, HPE, USA

## **Working Group and Chair**

TPC-AI: Raghunath Nambiar, AMD, USA

# Contents

Industry Panel on Defining Industry Standards for Benchmarking Artificial Intelligence . . . . .	1
<i>Raghnath Nambiar, Shahram Ghandeharizadeh, Gary Little, Christoph Boden, and Ajay Dholakia</i>	
UniBench: A Benchmark for Multi-model Database Management Systems . . .	7
<i>Chao Zhang, Jiaheng Lu, Pengfei Xu, and Yuxing Chen</i>	
PolyBench: The First Benchmark for Polystores . . . . .	24
<i>Jeyhun Karimov, Tilmann Rabl, and Volker Markl</i>	
Benchmarking Distributed Data Processing Systems for Machine Learning Workloads . . . . .	42
<i>Christoph Boden, Tilmann Rabl, Sebastian Schelter, and Volker Markl</i>	
Characterizing the Performance and Resilience of HCI Clusters with the TPCx-HCI Benchmark . . . . .	58
<i>H. Reza Taheri, Gary Little, Bhavik Desai, Andrew Bond, Doug Johnson, and Greg Kopczynski</i>	
Requirements for an Enterprise AI Benchmark . . . . .	71
<i>Cedric Bourrasset, France Boillod-Cerneux, Ludovic Sauge, Myrtille Deldossi, Francois Wellenreiter, Rajesh Bordawekar, Susan Malaika, Jean-Armand Broyelle, Marc West, and Brian Belgodere</i>	
Towards Evaluation of Tensorflow Performance in a Distributed Compute Environment . . . . .	82
<i>Miro Hodak and Ajay Dholakia</i>	
A Comparison of Two Cache Augmented SQL Architectures . . . . .	94
<i>Shahram Ghandeharizadeh and Hieu Nguyen</i>	
Benchmarking and Performance Analysis of Event Sequence Queries on Relational Database . . . . .	110
<i>Yuto Hayamizu, Ryoji Kawamichi, Kazuo Goda, and Masaru Kitsuregawa</i>	
Data Consistency Properties of Document Store as a Service (DSaaS): Using MongoDB Atlas as an Example. . . . .	126
<i>Chenhao Huang, Michael Cahill, Alan Fekete, and Uwe Röhm</i>	
Lessons Learned from the Industry’s First TPC Benchmark DS (TPC-DS) . . .	140
<i>Manan Trivedi and Zhenqiang Chen</i>	
<b>Author Index . . . . .</b>	<b>155</b>