

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, 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

Gerhard Weikum

Max Planck Institute for Informatics, Saarbrücken, Germany

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

Raghunath Nambiar · Meikel Poess (Eds.)

Performance Evaluation and Benchmarking

Traditional - Big Data -
Internet of Things

8th TPC Technology Conference, TPCTC 2016
New Delhi, India, September 5–9, 2016
Revised Selected Papers

Editors

Raghunath Nambiar
Cisco Systems, Inc.
San Jose, CA
USA

Meikel Poess
Oracle Corporation
Redwood City, CA
USA

ISSN 0302-9743 ISSN 1611-3349 (electronic)
Lecture Notes in Computer Science
ISBN 978-3-319-54333-8 ISBN 978-3-319-54334-5 (eBook)
DOI 10.1007/978-3-319-54334-5

Library of Congress Control Number: 2017932124

LNCS Sublibrary: SL2 – Programming and Software Engineering

© Springer International Publishing AG 2017

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.

Printed on acid-free paper

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

The original version of the cover and title page were revised. The title was corrected from “Performance Evaluation and Benchmarking. Traditional - Big Data - Interest of Things” to “Performance Evaluation and Benchmarking. Traditional - Big Data - Internet of Things”. An erratum to cover and frontmatter can be found at DOI: [10.1007/978-3-319-54334-5_11](https://doi.org/10.1007/978-3-319-54334-5_11)

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.

This book contains the proceedings of the 8th TPC Technology Conference on Performance Evaluation and Benchmarking (TPCTC 2016), held in conjunction with the 41st International Conference on Very Large Data Bases (VLDB 2016) in New Delhi, India, during September 5–9, 2016, including selected peer-reviewed papers as well as an invited paper and a keynote paper.

The hard work and close cooperation of a number of people contributed to the success of this conference. We would like to thank the members of TPC and the organizers of VLDB 2016 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.

January 2017

Raghunath Nambiar
Meikel Poess

TPCTC 2016 Organization

General Chairs

Raghunath Nambiar	Cisco, USA
Meikel Poess	Oracle, USA

Program Committee

Alain Crolotte	Teradata, USA
Akon Dey	University of Sydney, Australia
Berni Schiefer	IBM, Canada
Chaitanya Baru	SDSC, USA
Daniel Bowers	Gartner, USA
Dhabaleswar Panda	The Ohio State University, USA
Francois Raab	Infosizing, USA
Harumi Kuno	HP Labs, USA
Marco Vieira	University of Coimbra, Portugal
Michael Brey	Oracle, USA
Paul Cao	HP, USA
Reza Taheri	VMWare, USA
Tilman Rabl	University of Toronto, Canada
Yanpei Chen	Splunk, USA

Publicity Committee

Raghunath Nambiar	Cisco, USA
Andrew Bond	Red Hat, USA
Miso Cilimdžić	Microsoft, USA
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 that defines transaction processing and database benchmarks and distributes vendor-neutral performance data to the 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 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
USA
Voice: 415-561-6272
Fax: 415-561-6120
E-mail: 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.

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 2016 Organization

Full Members

Actian
Cisco
Cloudera
Dell
DataCore
Fujitsu
HP Enterprise
Hitachi
Huawei
IBM
Inspur
Intel
Lenovo
Microsoft
Oracle
Pivotal
Red Hat
SAP
Teradata
VMware

Associate Members

IDEAS International
San Diego Super Computing Center
Telecommunications Technology Association
University of Coimbra, Portugal
CAICT

Steering Committee

Andrew Bond (Red Hat)
Michael Brey (Oracle)
Matthew Emmerton (HP)
Raghunath Nambiar (Cisco)
Jamie Reding (Microsoft)

Public Relations Committee

Andrew Bond (Red Hat)
Raghunath Nambiar (Cisco), Chair
Miso Cilimdzc, Microsoft, USA
Meikel Poess (Oracle)
Reza Taheri (VMware)

Technical Advisory Board

Andrew Bond (Red Hat)
Paul Cao (HP)
Matthew Emmerton (IBM)
John Fowler (Oracle)
Jamie Reding (Microsoft), Chair
Nicholas Wakou (Dell)

Technical Subcommittees and Chairs

TPC-ACID-AR: John Fowler
TPC-Pricing: Jamie Reding
TPC-C: Jamie Reding
TPC-DI: Meikel Poess
TPC-DS: Meikel Poess
TPC-E: Matthew Emmerton
TPC-H: Miso Cilimdzc
TPCx-V: Reza Taheri
TPC-VMS: Reza Taheri
TPCx-BB: Bhaskar Gowda
TPCx-HS: Tariq Magdon-Ismail

Working Groups and Chairs

TPC-IoT: Raghunath Nambiar

Contents

Industry Standards for the Analytics Era: TPC Roadmap	1
<i>Raghunath Nambiar and Meikel Poess</i>	
TPCx-HS on the Cloud!	7
<i>Nicholas Wakou, Michael Woodside, Arkady Kanevsky, Fazal E Rehman Khan, and Mofassir ul Islam Arif</i>	
From BigBench to TPCx-BB: Standardization of a Big Data Benchmark	24
<i>Paul Cao, Bhaskar Gowda, Seetha Lakshmi, Chinmayi Narasimhadevara, Patrick Nguyen, John Poelman, Meikel Poess, and Tilmann Rabl</i>	
Benchmarking Spark Machine Learning Using BigBench.	45
<i>Sweta Singh</i>	
Benchmarking Exploratory OLAP.	61
<i>Mahfoud Djedaini, Pedro Furtado, Nicolas Labroche, Patrick Marcel, and Verónica Peralta</i>	
Lessons from OLTP Workload on Multi-socket HPE Integrity Superdome X System	78
<i>Srinivasan Varadarajan Sahasranamam, Paul Cao, Rajesh Tadakamadla, and Scott Norton</i>	
Benchmarking Distributed Stream Processing Platforms for IoT Applications	90
<i>Anshu Shukla and Yogesh Simmhan</i>	
AdBench: A Complete Benchmark for Modern Data Pipelines	107
<i>Milind Bhandarkar</i>	
Lessons Learned: Performance Tuning for Hadoop Systems	121
<i>Manan Trivedi and Raghunath Nambiar</i>	
Work-Energy Profiles: General Approach and In-Memory Database Application	142
<i>Annett Ungethüm, Thomas Kissinger, Dirk Habich, and Wolfgang Lehner</i>	
Erratum to: Performance Evaluation and Benchmarking	E1
<i>Raghunath Nambiar and Meikel Poess</i>	
Performance and Energy Analysis Using Transactional Workloads	159
<i>Anastasia Ailamaki, Danica Porobic, and Utku Sirin</i>	
Author Index	161