

Commenced Publication in 1973

Founding and Former Series Editors:

Gerhard Goos, Juris Hartmanis, and Jan van Leeuwen

Editorial Board

David Hutchison

Lancaster University, UK

Takeo Kanade

Carnegie Mellon University, Pittsburgh, PA, USA

Josef Kittler

University of Surrey, Guildford, UK

Jon M. Kleinberg

Cornell University, Ithaca, NY, USA

Alfred Kobsa

University of California, Irvine, CA, USA

Friedemann Mattern

ETH Zurich, Switzerland

John C. Mitchell

Stanford University, CA, USA

Moni Naor

Weizmann Institute of Science, Rehovot, Israel

Oscar Nierstrasz

University of Bern, Switzerland

C. Pandu Rangan

Indian Institute of Technology, Madras, India

Bernhard Steffen

TU Dortmund University, Germany

Madhu Sudan

Microsoft Research, Cambridge, MA, USA

Demetri Terzopoulos

University of California, Los Angeles, CA, USA

Doug Tygar

University of California, Berkeley, CA, USA

Gerhard Weikum

Max Planck Institute for Informatics, Saarbruecken, Germany

Raghunath Nambiar Meikel Poess (Eds.)

Topics in Performance Evaluation, Measurement and Characterization

Third TPC Technology Conference, TPCTC 2011
Seattle, WA, USA, August 29–September 3, 2011
Revised Selected Papers

Volume Editors

Raghunath Nambiar
Data Center Group
Cisco Systems, Inc.
3800 Zankar Road
San Jose, CA 95134, USA
E-mail: rnambiar@cisco.com

Meikel Poess
Server Technologies
Oracle Corporation
500 Oracle Parkway
Redwood Shores, CA 94065, USA
E-mail: meikel.poess@oracle.com

ISSN 0302-9743 e-ISSN 1611-3349
ISBN 978-3-642-32626-4 e-ISBN 978-3-642-32627-1
DOI 10.1007/978-3-642-32627-1
Springer Heidelberg Dordrecht London New York

Library of Congress Control Number: 2012944423

CR Subject Classification (1998): C.4, D.2.8, J.1, H.2.7-8, K.6, H.4, K.1

LNCS Sublibrary: SL 2 – Programming and Software Engineering

© Springer-Verlag Berlin Heidelberg 2012

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.

The use of general descriptive names, registered names, trademarks, 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.

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

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, and one vehicle for achieving this objective is the sponsorship of the Technology Conference Series on Performance Evaluation and Benchmarking (TPCTC). With this conference, 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.

This book contains the proceedings of the Third TPC Technology Conference on Performance Evaluation and Benchmarking (TPCTC 2011), held in conjunction with the 37th International Conference on Very Large Data Bases (VLDB 2011) in Seattle, Washington, from August 29 to September 3, 2011, including 12 selected papers and two keynote papers.

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 2011 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.

Raghunath Nambiar
Meikel Poess

TPCTC 2011 Organization

General Chairs

Raghunath Nambiar (Cisco)
Meikel Poess (Oracle)

Publicity Committee Chair

Nicholas Wakou (Dell)

Program Committee

Alfredo Cuzzocrea	University of Calabria, Italy
Alain Crolotte	Teradata, USA
Neoklis Polyzotis	University of California, Santa Cruz, USA
Badridine Khessib	Microsoft, USA
Brian Cooper	Google, USA
Harumi Kuno	HP Labs, USA
Kai Sachs	SAP, Germany, USA
Marco Vieira	University of Coimbra, Portugal
Masaru Kitsuregawa	University of Tokyo, Japan
Michael Brey	Oracle, USA
Michael Molloy	Dell, USA
Patrick O'Neil	University of Massachusetts at Boston, USA
Peter Thawley	Sybase, USA
Tilman Rabl	University of Toronto, Canada

Publicity Committee

Andrew Masland (NEC)
Matthew Lanken (Oracle)
Peter Thawley (Sybase)
Michael Majdalany (L&M Management Group)
Forrest Carman (Owen Media)

Keynote

Umeshwar Dayal (HP Labs)

Invited Talk

Karl Huppler (IBM)

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 directions. 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

Voice: 415-561-6272

Fax: 415-561-6120

Email: 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 email 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 2011 Organization

Full Members

AMD
Bull
Cisco
Dell
Fujitsu
HP
Hitachi
IBM
Ingres
Intel
Microsoft
NEC
Oracle
ParAccel
Sybase
Syncsort
Teradata
Unisys
VMware

Associate Members

Ideas International
ITOM International Co
Telecommunications Technology Association
University of Coimbra, Portugal

Steering Committee

Karl Huppler (IBM), Chair
Charles Levine (Microsoft)
Mike Nikolaiev (HP)
Raghunath Nambiar (Cisco)
Wayne Smith (Intel)

Public Relations Committee

Nicholas Wakou (Dell), Chair
Andrew Masland (NEC)
Meikel Poess (Oracle)
Peter Thawley (Sybase)
Raghunath Nambiar (Cisco)

Technical Advisory Board

Jamie Reding (Microsoft), Chair
Peter Thawley (Sybase)
Dave Steinhoff (ParAccel)
Matthew Emmerton (IBM)
Mike Nikolaiev (HP)
Nicholas Wakou (Dell)
Wayne Smith (Intel)

Table of Contents

Shaping the Landscape of Industry Standard Benchmarks: Contributions of the Transaction Processing Performance Council (TPC)	1
<i>Raghunath Nambiar, Nicholas Wakou, Andrew Masland, Peter Thawley, Matthew Lanken, Forrest Carman, and Michael Majdalany</i>	
Metrics for Measuring the Performance of the Mixed Workload CH-benCHmark	10
<i>Florian Funke, Alfons Kemper, Stefan Krompass, Harumi Kuno, Raghunath Nambiar, Thomas Neumann, Anisoara Nica, Meikel Poess, and Michael Seibold</i>	
Towards an Enhanced Benchmark Advocating Energy-Efficient Systems	31
<i>Daniel Schall, Volker Hoefner, and Manuel Kern</i>	
Optimization of Analytic Data Flows for Next Generation Business Intelligence Applications	46
<i>Umeshwar Dayal, Kevin Wilkinson, Alkis Simitsis, Malu Castellanos, and Lupita Paz</i>	
Normalization in a Mixed OLTP and OLAP Workload Scenario.....	67
<i>Anja Bog, Kai Sachs, Alexander Zeier, and Hasso Plattner</i>	
Measuring Performance of Complex Event Processing Systems	83
<i>Torsten Grabs and Ming Lu</i>	
Benchmarking with Your Head in the Cloud	97
<i>Karl Huppler</i>	
Extending TPC-E to Measure Availability in Database Systems.....	111
<i>Yantao Li and Charles Levine</i>	
SI-CV: Snapshot Isolation with Co-located Versions	123
<i>Robert Gottstein, Ilia Petrov, and Alejandro Buchmann</i>	
Introducing Skew into the TPC-H Benchmark	137
<i>Alain Crolotte and Ahmad Ghazal</i>	
Time and Cost-Efficient Modeling and Generation of Large-Scale TPCC/TPCE/TPCH Workloads.....	146
<i>Christina Delimitrou, Sriram Sankar, Badriddine Khessib, Kushagra Vaid, and Christos Kozyrakis</i>	

When Free Is Not Really Free: What Does It Cost to Run a Database Workload in the Cloud?	163
<i>Avrilia Floratou, Jignesh M. Patel, Willis Lang, and Alan Halverson</i>	
A Fine-Grained Performance-Based Decision Model for Virtualization Application Solution	180
<i>Jianhai Chen, Dawei Huang, Bei Wang, Deshi Ye, Qinming He, and Wenzhi Chen</i>	
A PDGF Implementation for TPC-H	196
<i>Meikel Poess, Tilmann Rabl, Michael Frank, and Manuel Danisch</i>	
Author Index	213