

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

University of Dortmund, 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 of Computer Science, Saarbruecken, Germany

Raghunath Nambiar Meikel Poess (Eds.)

Performance Evaluation and Benchmarking

First TPC Technology Conference, TPCTC 2009
Lyon, France, August 24-28, 2009
Revised Selected Papers



Springer

Volume Editors

Raghunath Nambiar
Hewlett-Packard Company
11445 Compaq Center Dr W
Houston, TX 77070, USA
E-mail: raghu.nambiar@hp.com

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

Library of Congress Control Number: 2009938720

CR Subject Classification (1998): C.4, D.2.8, D.2, D.4.8, H.3.4, K.6.2

LNCS Sublibrary: SL 2 – Programming and Software Engineering

ISSN	0302-9743
ISBN-10	3-642-10423-1 Springer Berlin Heidelberg New York
ISBN-13	978-3-642-10423-7 Springer Berlin Heidelberg New York

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.

springer.com

© Springer-Verlag Berlin Heidelberg 2009
Printed in Germany

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India
Printed on acid-free paper SPIN: 12792528 06/3180 5 4 3 2 1 0

Preface

First established in August 1988, the Transaction Processing Performance Council (TPC) has shaped the landscape of modern transaction processing and database benchmarks over two decades. Now, the world is in the midst of an extraordinary information explosion led by rapid growth in the use of the Internet and connected devices. Both user-generated data and enterprise data levels continue to grow exponentially. With substantial technological breakthroughs, Moore's law will continue for at least a decade, and the data storage capacities and data transfer speeds will continue to increase exponentially. These have challenged industry experts and researchers to develop innovative techniques to evaluate and benchmark both hardware and software technologies.

As a result, the TPC held its First Conference on Performance Evaluation and Benchmarking (TPCTC 2009) on August 24 in Lyon, France in conjunction with the 35th International Conference on Very Large Data Bases (VLDB 2009). TPCTC 2009 provided industry experts and researchers with a forum to present and debate novel ideas and methodologies in performance evaluation, measurement and characterization for 2010 and beyond.

This book contains the proceedings of this conference, including 16 papers and keynote papers from Michael Stonebraker and Karl Huppler.

A number of people have contributed to the success of this conference. I would like to thank the members of TPC and the organizers of the VLDB 2009 conference for their support and sponsorship. I would also like to thank members of the Program Committee, Publicity Committee, the authors and the conference participants for their contributions in making this conference a big success.

November 2009

Raghunath Nambiar

TPCTC 2009 Organization

General Chair

Raghunath Nambiar (Hewlett-Packard)

Program Committee Chair

Meikel Poess (Oracle)

Publicity Committee Chair

Nicholas Wakou (Dell)

Program Committee

Alain Crolotte (Teradata)

Berni Schiefer (IBM)

Krithi Ramamritham (IIT Bombay)

Masaru Kitsuregawa (University of Tokyo)

Mike Molloy (Dell)

Murali Krishna (Hewlett-Packard)

Omer Trajman (Vertica)

Paul Larson (Microsoft)

Yicheng Tu (University of South Florida)

Publicity Committee

Jerrold Buggert (Unisys)

Matthew Lanken (Oracle)

Peter Thawley (Sybase)

Forrest Carman (Owen Media)

Michael Majdalany (L&M Management Group)

Keynote Speakers

Michael Stonebraker (MIT)

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. Over the past two decades, the TPC has had a significant impact on industry and on expectations around benchmarks. Vendors and end users rely on TPC benchmarks to provide real-world data that are backed by a stringent and independent review process.

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 form 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 form can be found at: <http://www.tpc.org/information/about/app-assoc.asp>.

Academic and Government Institutions

Academic and government institutions are invited to participate in the TPC's bimonthly meetings, 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: <http://www.tpc.org/information/about/email.asp>.

TPC 2009 Organization

Full Members

AMD, Bull, Dell, Fujitsu, Fusion IO, HP, Hitachi, IBM, Ingres, Intel, Kickfire, Microsoft, NEC, Netezza, Oracle, ParAccel, Sun, Sybase, Syncsort, Teradata, Unisys, Vertica, VMware and XSPRADA

Associate Members

Ideas International, ITOM International Co and TTA

Steering Committee

Karl Huppler (IBM), Chair
Charles Levine (Microsoft)
Jerrold Buggert (Unisys)
Mike Molloy (Dell)
Raghunath Nambiar (Hewlett-Packard)

Public Relations Committee

Nicholas Wakou (Dell), Chair
Jerrold Buggert (Unisys)
Matthew Lanken (Oracle)
Peter Thawley (Sybase)
Raghunath Nambiar (Hewlett-Packard)

Technical Advisory Board

Mike Brey (Oracle), Chair
Jamie Reding (Microsoft)
Matthew Emmerton (IBM)
Mike Molloy (Dell)
Omer Trajman (Vertica)
Rick Freeman (Unisys)
Wayne Smith (Intel)

Table of Contents

Transaction Processing Performance Council (TPC): Twenty Years Later – A Look Back, a Look Ahead	1
<i>Raghunath Othayoth Nambiar, Matthew Lanken, Nicholas Wakou, Forrest Carman, and Michael Majdalany</i>	
A New Direction for TPC?	11
<i>Michael Stonebraker</i>	
The Art of Building a Good Benchmark	18
<i>Karl Huppler</i>	
Database Are Not Toasters: A Framework for Comparing Data Warehouse Appliances	31
<i>Omer Trajman, Alain Crolotte, David Steinhoff, Raghunath Othayoth Nambiar, and Meikel Poess</i>	
The State of Energy and Performance Benchmarking for Enterprise Servers	52
<i>Andrew Fanara, Evan Haines, and Arthur Howard</i>	
From Performance to Dependability Benchmarking: A Mandatory Path	67
<i>Marco Vieira and Henrique Madeira</i>	
Overview of TPC Benchmark E: The Next Generation of OLTP Benchmarks	84
<i>Trish Hogan</i>	
Converting TPC-H Query Templates to Use DSQGEN for Easy Extensibility	99
<i>John M. Stephens Jr. and Meikel Poess</i>	
Generating Shifting Workloads to Benchmark Adaptability in Relational Database Systems	116
<i>Tilmann Rabl, Andreas Lang, Thomas Hackl, Bernhard Sick, and Harald Kosch</i>	
Measuring Database Performance in Online Services: A Trace-Based Approach	132
<i>Swaroop Kavalanekar, Dushyanth Narayanan, Sriram Sankar, Eno Thereska, Kushagra Vaid, and Bruce Worthington</i>	
Issues in Benchmark Metric Selection	146
<i>Alain Crolotte</i>	

Benchmarking Query Execution Robustness	153
<i>Janet L. Wiener, Harumi Kuno, and Goetz Graefe</i>	
Benchmarking Database Performance in a Virtual Environment	167
<i>Sharada Bose, Priti Mishra, Priya Sethuraman, and Reza Taheri</i>	
Principles for an ETL Benchmark	183
<i>Len Wyatt, Brian Caulfield, and Daniel Pol</i>	
Benchmarking ETL Workflows	199
<i>Alkis Simitsis, Panos Vassiliadis, Umeshwar Dayal, Anastasios Karagiannis, and Vasiliki Tziouvara</i>	
A Performance Study of Event Processing Systems	221
<i>Marcelo R.N. Mendes, Pedro Bizarro, and Paulo Marques</i>	
The Star Schema Benchmark and Augmented Fact Table Indexing	237
<i>Patrick O'Neil, Elizabeth O'Neil, Xuedong Chen, and Stephen Revilak</i>	
An Approach of Performance Evaluation in Authentic Database Applications	253
<i>Xiaojun Ye, Jingmin Xie, Jianmin Wang, Hao Tang, and Naiqiao Du</i>	
Author Index	267