

Performance Evaluation in Database Research: Principles and Experience

Stefan Manegold
CWI
Amsterdam, The Netherlands
stefan.manegold@cwi.nl

Ioana Manolescu
INRIA Saclay—Île-de-France
Orsay, France
ioana.manolescu@inria.fr

ABSTRACT

A significant part of today's database research focuses on improving performance of a specific system. Quantitative experiments are the best way to validate such results. However, performing experiments is not always easy. Besides the complexity of the system under test, designing an experiment, choosing the right environment and parameter values, analyzing the data which is gathered, and reporting it to a

third party in an expressive and intelligible way is hard.

In this tutorial, we present a general road-map to the above steps, including tips and tricks on how to organize and present code that performs experiments, so that an outsider can repeat them.

The tutorial is primarily aimed at MS and PhD students seeking to improve their experiment practices, but more senior attendants may also find it interesting.

ABOUT THE AUTHORS

Stefan Manegold (CWI, Amsterdam) Stefan Manegold is a researcher in the database group at CWI in Amsterdam, The Netherlands. His research work comprises database architectures and data management on modern hardware as well as database-supported XML processing, with a particular interest in optimization, performance and benchmarking. He is co-founder of the DaMoN workshop series (co-located with SIGMOD since 2005) and co-chair of ExpDB 2007.



Ioana Manolescu (INRIA Saclay—Île-de-Paris) Ioana Manolescu is a researcher in the Gemo group of INRIA Saclay, in France. Her research work is centered around XML data management, distributed data management systems, and Web application modeling. She is a founder of two SIGMOD-affiliated workshops, XIMEP on XQuery processing, and ExpDB on experimental evaluation in database research.



Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the ACM. To copy otherwise, or to republish, to post on servers or to redistribute to lists, requires a fee and/or special permissions from the publisher, ACM.
EDBT 2009, March 24–26, 2009, Saint Petersburg, Russia.
Copyright 2009 ACM 978-1-60558-422-5/09/0003 ...\$5.00