Lecture Notes in Computer Science 11178

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

Rena Bakhshi · Paolo Ballarini
Benoît Barbot · Hind Castel-Taleb
Anne Remke (Eds.)

Computer Performance Engineering

15th European Workshop, EPEW 2018
Paris, France, October 29–30, 2018
Proceedings

Springer

Editors
Rena Bakhshi
eScience Center
Amsterdam
The Netherlands

Paolo Ballarini
University Paris-Saclay
Gif-sur-Yvette
France

Benoît Barbot
Université Paris-Est Créteil
Créteil
France

Hind Castel-Taleb
Telecom SudParis
Evry
France

Anne Remke
Westfälische Wilhelms Universität Münster
Münster
Germany

ISSN 0302-9743 ISSN 1611-3349 (electronic)
Lecture Notes in Computer Science
ISBN 978-3-030-02226-6 ISBN 978-3-030-02227-3 (eBook)
https://doi.org/10.1007/978-3-030-02227-3

Library of Congress Control Number: 2018957280

LNCS Sublibrary: SL2 – Programming and Software Engineering

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

Preface

This LNCS volume contains the proceedings of the 15th European Performance Engineering Workshop (EPEW), held in Paris, France, during October 29–30, 2018. At its 15th edition the annual EPEW workshop series is a well-established workshop aimed at providing researchers from both academia and industry with a forum for debate/networking over a broad range of topics across the performance engineering realm, including, dependability and security modeling, performance-oriented model verification and testing, hardware and software systems case studies, applications/extensions of queuing theory, and network design. Following the tradition of EPEW, the papers presented at the 2018 edition reflect the diversity of modern performance engineering, where theoretical aspects (involving formalisms such as graphs, trees, queueing networks, stochastic automata as well as mathematical methods such as simulation, product form solutions, game theory, optimization, model checking) are often combined with specific applications (e.g., delay-tolerant networks, mobile cloud computing, smart buildings, fault-tolerant systems, distributed databases).

EPEW 2018 received 27 submissions from nine countries around the world, including Asia, North America, and Europe. Each paper was peer reviewed by an average of three Program Committee (PC) members and assessed on the basis of its relevance to the workshop community, its novelty, and its technical quality. The review outcome led to the selection of 17 high-quality contributions for publication in the proceedings and presentation at the workshop.

We would like to thanks the keynote speakers we were honored to host at EPEW 2018, namely, Dr. Gerardo Rubino, a senior researcher at Inria/IRISA Rennes (France) with a strong background both in quality of services for network applications as well as in quantitative analysis of probabilistic models, and Dr. Benny Van Houdt, a senior lecturer at the computer science department of the University of Antwerp (Belgium) with strong diverse expertise in performance evaluation of computer systems and networks.

We would also like to warmly thank all PC members and external reviewers for their quality work in the review process. Furthermore we would like to thank both the SAMOVAR laboratory of Télécom SudParis, the LACL laboratory of the University of Paris-Est Créteil, and the MICS laboratory of CentraleSupélec for their support in making the organization of EPEW 2018 possible. We are very grateful both to the EasyChair team, for their useful conference system that we used for managing papers submission/reviewing process, and to Springer for their editorial support. Above all, we

would like to thank the authors of the papers for their contribution to this volume. We hope these contributions will be as useful and inspiring to the readers as they were to us.

September 2018

Rena Bakhshi
Paolo Ballarini
Benoît Barbot
Hind Castel
Anne Remke

Organization

Program Committee

Elvio Gilberto Amparore	University of Turin, Italy
Urtzi Ayesta	LAAS-CNRS, France
Rena Bakhshi	Netherlands eScience Center, The Netherlands
Paolo Ballarini	CentraleSupelec, France
Simonetta Balsamo	Università Ca' Foscari di Venezia, Italy
Benoit Barbot	Paris-Est Creteil University, France
Marco Beccuti	University of Turin, Italy
Marco Bernardo	University of Urbino, Italy
Marco Biagi	University of Florence, Italy
Olivier Brun	LAAS-CNRS, France
Laura Carnevali	University of Florence, Italy
Hind Castel	Telecom SudParis, France
Tadeusz Czachórski	Institute of Theoretical and Applied Informatics, Polish Academy of Sciences, Poland
Dieter Fiems	Ghent University, Belgium
Jean-Michel Fourneau	DAVID, Université de Versailles St. Quentin, France
Stephen Gilmore	University of Edinburgh, UK
Boudewijn Haverkort	University of Twente, The Netherlands
András Horváth	University of Turin, Italy
Gábor Horváth	Budapest University of Technology and Economics, Hungary
Emmanuel Hyon	Paris Nanterre University, France
Alain Jean-Marie	Inria, France
William Knottenbelt	Imperial College London, UK
Samuel Kounev	University of Würzburg, Germany
Lasse Leskelä	Aalto University, Finland
Andrea Marin	University of Venice, Italy
Marco Paolieri	University of Southern California, USA
Nihal Pekergin N.	Paris-Est Creteil University, France
Agapios Platis	University of the Aegean, Greece
Philipp Reinecke	Cardiff University, UK
Anne Remke	WWU Münster, Germany
Markus Siegle	University of the Bundeswehr Munich, Germany
Miklos Telek	Budapest University of Technology and Economics, Hungary
Nigel Thomas	Newcastle University, UK
Petr Tuma	Charles University, Czech Republic
Enrico Vicario	University of Florence, Italy

Joris Walraevens	Ghent University, Belgium
Qiushi Wang	Nanyang Technological University, Singapore
Katinka Wolter	Freie Universitüt zu Berlin, Germany
Huaming Wu	Tianjin University, China

Additional Reviewers

Hüls, Jannik
Pilch, Carina
Santoni, Francesco
Walter, Jürgen

Keynotes

Performance Evaluation Targeting Quality of Experience

Gerardo Rubino

Inria, France

When we must evaluate the performance of a computing facility, a communication network, a Web service, we typically build a model and, then, we use it to analyze one or several metrics that we know are important for capturing the performance aspect of interest of the considered system (mean response time, mean backlog of jobs/packets/requests/…, loss probabilities, etc.). Typical tools for analyzing the model are queuing theory results, Markov chain algorithms, discrete event simulation, etc. If we specifically consider the case of applications or services operating on the Internet and focusing on video, or audio, or voice content (IP telephony, video streaming, video-conferences, …), in most cases the ultimate target is the perception the user has about the delivered service, how satisfied she is with its quality, and this perception concentrates in that of the quality of the content (how good was the reception of the transmitted voice over the IP channel, or of the play of the movie requested to the VoD server, etc.). We call it Perceptual Quality (PQ), and it is the main component of the user-centered Quality of Experience for these fundamental classes of applications. In theory, PQ is the mandatory criteria to take care of the user when designing the system. Needless to say, these classes of apps and/or services are responsible for a very important component of today’s and tomorrow’s Internet traffic, and they represent a large fraction of total traffic in volume. The PQ is usually evaluated using subjective tests, that is, by means of panels of human observers. In this area many standards exist, depending on the type of media, the type of usage, etc. A subjective testing session provides, at the end, a number measuring the PQ, that is, quantifying this PQ. When quantifying (when measuring), we typically refer to the Mean Opinion Score (MOS) of the video or voice sequence, and a standard range for MOS is the real interval [1, 5], ‘1’ the worst, ‘5’ the best. In this presentation, we will argue that using an appropriate approach to measure this PQ, we can rely on our classic tools in performance evaluation (queuing models, low level stochastic processes, etc.) while focusing our effort in analyzing directly this PQ central aspect of our systems. Instead of saying “if the offered traffic and the system service rate satisfy relation R, then the throughput of the system is high, which is good, but the delay is also a little bit high, which is not very nice,…”, we can say “if the offered traffic and the system service rate satisfy relation R, then the PQ is high enough”. That is, instead of showing how the throughput, the mean backlog, the mean response time, … evolve with some parameters that can be controlled to tune the system’s performance, we can work directly with the PQ, the ultimate target, and still use our M/M/* queues, Jackson networks, or whatever model is relevant in our study. This allows obtaining results concerning our final goal, that is, keep the user happy when looking at the video stream, or when using her telephone, at a reasonable cost. Detailed examples will be given using the author own proposal for the automatic measure of PQ, called PSQA (for Pseudo Subjective Quality Assessment) and based on Machine Learning tools., that provides a rational function of several parameters returning the current PQ, parameters that may include those low level metrics.

Mean Field Models for (Large-Scale) Load Balancing Systems

Benny Van Houdt

University of Antwerpen, Belgium

This talk focuses on the behavior of load balancing systems and is composed of three parts.

In the first part we revisit a classic mean field result on load balancing in large distributed systems. More specifically, we focus on the celebrated power-of-two choices paradigm and its mean field limit. We subsequently introduce a theorem for the class of density dependent population processes established in the 1970s by Kurtz and discuss some of the technical issues involved to extend this result to the stationary regime. We end the first part by illustrating the accuracy of the mean field limit using simulation.

In the second part we introduce the refined mean field approximation, which is a technique to compute a 1/N correction term to improve the accuracy of classic mean field limits. This technique can be used to more accurately approximate the performance of small systems, e.g., consisting of N = 10 servers, and can be applied to any density dependent population processes with limited effort. We focus on the different computational steps involved to compute this correction term and illustrate its accuracy on various numerical examples.

In the final part of the talk we discuss some recent results on load balancing schemes that select servers based on workload information (as opposed to queue length information). Such systems are motivated by load balancing systems that use late binding or redundancy. We present explicit results for the workload and response time distribution when the job sizes follow an exponential distribution and indicate how to compute these distributions for non-exponential jobs sizes.

This talk is based on joint work with Nicolas Gast (Inria) and Tim Hellemans (University of Antwerp).

Contents

On the Degradation of Distributed Graph Databases with Eventual Consistency 1
Paul Ezhilchelvan, Isi Mitrani, and Jim Webber

To What Extent Does Performance Awareness Support Developers in Fixing Performance Bugs? 14
Alexandru Danciu and Helmut Krcmar

Deriving Symbolic Ordinary Differential Equations from Stochastic Symmetric Nets Without Unfolding 30
Marco Beccuti, Lorenzo Capra, Massimiliano De Pierro, Giuliana Franceschinis, and Simone Pernice

Mean Value Analysis of Closed G-Networks with Signals 46
Jean-Michel Fourneau

Extending the Steady State Analysis of Hierarchical Semi-Markov Processes with Parallel Regions 62
Marco Biagi, Enrico Vicario, and Reinhard German

Formal Parameter Synthesis for Energy-Utility-Optimal Fault Tolerance 78
Linda Herrmann, Christel Baier, Christof Fetzer, Sascha Klüppelholz, and Markus Napierkowski

Performance Model of Apache Cassandra Under Heterogeneous Workload Using the Quantitative Verification Approach 94
Al Amjad Tawfiq Isstaif and Nizar Alhafez

Modelling Smart Buildings Using Fault Maintenance Trees 110
Alessandro Abate, Carlos E. Budde, Nathalie Cauchi, Arnaud van Harmelen, Khaza Anuarul Hoque, and Mariëlle Stoelinga

Performance Impact of Misbehaving Voters 126
Mohammed Alotaibi and Nigel Thomas

Information Flow Security for Stochastic Processes 142
Jane Hillston, Andrea Marin, Carla Piazza, and Sabina Rossi

Towards Probabilistic Modeling and Analysis of Real-Time Systems 157
Laura Carnevali, Luca Santinelli, and Giuseppe Lipari

An Ontology Framework for Generating Discrete-Event Stochastic Models 173
Ken Keefe, Brett Feddersen, Michael Rausch, Ronald Wright, and William H. Sanders

A Mixed Strategy for a Competitive Game in Delay Tolerant Networks. 190
Thi Thu Hang Nguyen, Olivier Brun, and Balakrishna Prabhu

Second Order Fluid Performance Evaluation Models for Interactive 3D Multimedia Streaming 205
Enrico Barbierato, Marco Gribaudo, Mauro Iacono, and Pietro Piazzolla

Modeling the Effect of Parallel Execution on Multi-site Computation Offloading in Mobile Cloud Computing. 219
Ismail Sheikh and Olivia Das

An OpenFlow Controller Performance Evaluation Tool 235
Zhihao Shang, Han Wu, and Katinka Wolter

Product-Form Queueing Networks with Batches 250
P. G. Harrison

Author Index 265