Lecture Notes in Computer Science 9629

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, Zürich, 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

Anne Remke · Boudewijn R. Haverkort (Eds.)

Measurement, Modelling and Evaluation of Dependable Computer and Communication Systems

18th International GI/ITG Conference, MMB & DFT 2016
Münster, Germany, April 4–6, 2016
Proceedings

Springer

Editors
Anne Remke
Universität Münster
Münster
Germany

Boudewijn R. Haverkort
University of Twente
Enschede
The Netherlands

ISSN 0302-9743 ISSN 1611-3349 (electronic)
Lecture Notes in Computer Science
ISBN 978-3-319-31558-4 ISBN 978-3-319-31559-1 (eBook)
DOI 10.1007/978-3-319-31559-1

Library of Congress Control Number: 2016933469

LNCS Sublibrary: SL2 – Programming and Software Engineering

© Springer International Publishing Switzerland 2016
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.

Printed on acid-free paper

This Springer imprint is published by Springer Nature
The registered company is Springer International Publishing AG Switzerland

Preface

Welcome to the proceedings of MMB and DFT 2016! We are very pleased to present this LNCS volume with its contributions on performance and dependability evaluation techniques for distributed and embedded systems, computer and software architectures, and communication networks.

This volume contains the papers that were presented at the 18th International GI/ITG Conference on Measurement, Modelling and Evaluation of Computing Systems and Dependability and Fault Tolerance (MMB and DFT 2016) held during April 4–6, 2016, in Münster, Germany.

Following a thorough review procedure with at least three reviews per submission and a careful selection process, the Program Committee of MMB and DFT 2016 compiled an interesting scientific program comprising 12 regular papers and three tool presentations.

Since the start of the biennial MMB conference series in the early 1980's, we have seen substantial changes in the field of performance evaluation, dependability, and fault-tolerance of computer and communication systems. This is, for example, reflected in the relatively large number of submissions that deal with a very interesting and highly relevant field of research, namely, smart grids. We believe that for this community, it is very important to address new and exciting applications and investigate how the knowledge that is available in our community can be applied to these.

Besides the main program, the conference hosted three satellite workshops covering related research topics:

- The 8th International Workshop on Practical Applications of Stochastic Modelling (PASM)
- The Third Workshop on Network Calculus (WoNeCa)
- The Workshop on E-mobility and Smart Grids: Challenges and Opportunities (E-mobility)

By hosting these workshop, we hope to foster interaction between strongly related communities.

We were very fortunate to include two very interesting and relevant keynote presentation in the conference program:

- "DDoS 3.0: How Terrorists Bring Down the Internet" by Prof. Dr. Ir. Aiko Pras, University of Twente, The Netherlands
- "From Transient Analysis to Probabilistic Model Checking of Markov Regenerative Processes" by Prof. Dr. Enrico Vicario, University of Florence, Italy

To enable cross-fertilization between the conference and the satellite workshops, we included two additional invited talks that covered the research areas of the satellite workshops and that fit, from our perspective, very well in the scope of the main conference:

- "Critical Machine-to-Machine Communications: Performance Models vs. Reality in the 10^{-10} Regime," by Prof. Dr. James Gross, Royal Institute of Technology, Stockholm, Sweden
- "Open Analysis of Crowdsourced Car Sensor Data: The enviroCar Project," by Dr. Christoph Stasch, 52 North, Münster, Germany

As conference chairs, we express our gratitude to all members of the Program Committee and all external reviewers for their dedicated service, maintaining the quality objectives of the conference, and for the timely provision of their valuable reviews.

We thank all the authors for their submissions, all the speakers for their lively presentations, and all the participants for their contributions to interesting discussions. We acknowledge the support of the EasyChair conference system and express our gratitude to its management team for their commitment to serve the scientific community. Further, we thank Springer for unceasing support and excellent management of the LNCS publishing process.

Finally, it is our hope that readers will find these MMB and DFT 2016 proceedings informative and useful for their future research on measurement, modelling, analysis, and performance evaluation of advanced computer and communication systems.

February 2016

Anne Remke
Boudewijn R. Haverkort

Organization

Program Committee

Lothar Breuer	University of Kent, UK
Peter Buchholz	TU Dortmund, Germany
Hans Daduna	University of Hamburg, Germany
Hermann De Meer	University of Passau, Germany
Johannes Dr. Riedl	ITG MMB
Klaus Echtle	University of Duisburg-Essen, Germany
Markus Fidler	Leibniz Universität Hannover, Germany
Reinhard German	University of Erlangen, Germany
Gerhard Hasslinger	T-Systems ENPS Darmstadt, Germany
Boudewijn Haverkort	University of Twente, The Netherlands
Holger Hermanns	Saarland University, Germany
Joost-Pieter Katoen	RWTH Aachen University, Germany
Peter Kemper	College of William and Mary, USA
Udo Krieger	Otto Friedrich University, Germany
Kai Lampka	Uppsala University, Sweden
Wolfram Lautenschlaeger	Alcatel-Lucent Deutschland AG, Germany
Axel Lehmann	Universität der Bundeswehr München, Germany
Ralf Lehnert	Technical University of Dresden, Germany
Michael Menth	University of Tübingen, Germany
Peter Reichl	University of Vienna, Austria
Anne Remke	University of Twente, The Netherlands
Ramin Sadre	Université catholique de Louvain, Belgium
Francesca Saglietti	University of Erlangen-Nuremberg, Germany
Jens Schmitt	TU Kaiserslautern, Germany
Markus Siegle	Bundeswehr University Munich, Germany
Dietmar Tutsch	University of Wuppertal, Germany
Kurt Tutschku	Blekinge Institute of Technology (BTH), Sweden
Oliver Waldhorst	Daimler AG, Germany
Verena Wolf	Saarland University, Germany
Bernd Wolfinger	Universität Hamburg, Germany
Katinka Wolter	Freie Universität zu Berlin, Germany
Armin Zimmermann	Technische Universität Ilmenau, Germany

Additional Reviewers

Beck, Michael
Berger, Daniel
Heidtmann, Klaus
Heimgaertner, Florian
Jongerden, Marijn
Junges, Sebastian
Krcal, Jan
Krüger, Thilo
Lück, Alexander
Mandarawi, Waseem

Abstracts of Invited Talks

DDoS 3.0 - How Terrorists Bring Down the Internet

Aiko Pras, José Jair Santanna, Jessica Steinberger, and Anna Sperotto

University of Twente
Enschede, The Netherlands
{a.pras, j.j.santanna, a.sperotto}@utwente.nl,
jessica.steinberger@h-da.de

Abstract. Dependable operation of the Internet is of crucial importance for our society. In recent years Distributed Denial of Service (DDoS) attacks have quickly become a major problem for the Internet. Most of these attacks are initiated by kids that target schools, ISPs, banks and web-shops; the Dutch NREN (SURFNet), for example, sees around 10 of such attacks per day. Performing attacks is extremely simple, since many websites offer "DDoS as a Service"; in fact it is easier to order a DDoS attack than to book a hotel! The websites that offer such DDoS attacks are called "Booters" or "Stressers", and are able to perform attacks with a strength of many Gbps. Although current attempts to mitigate attacks seem promising, analysis of recent attacks learns that it is quite easy to build next generation attack tools that are able to generate DDoS attacks with a strength thousand to one million times higher than the ones we see today. If such tools are used by nation-states or, more likely, terrorists, it should be possible to completely stop the Internet. This paper argues that we should prepare for such novel attacks.

Open Analysis of Crowdsourced Car Sensor Data - The enviroCar Project

Christoph Stasch, Albert Remke, Arne de Wall, and Matthes Rieke

52°North Open Source Software GmbH, Münster, Germany
{c.stasch,a.remke,a.dewall,m.rieke}@52north.org

Cars are equipped with various sensors used to monitor the engine and its environment. By using the so-called On-Board-Diagnostics II (OBD-II) interface, these sensors can be assessed by external devices. The enviroCar project[1] consists of an open infrastructure that utilizes this technology in order to enable drivers to collect, analyze, share and discuss car sensor data [1]. As shown in Fig. 1, the enviroCar infrastructure consists of an app, a server component, various analysis tools, and a community portal. The enviroCar app allows car drivers to connect their Android mobile phones to car sensors using an OBD bluetooth adapter. The app provides feedback while driving and allows uploading recorded tracks to the enviroCar server, where the data is publicly and anonymized accessible as open data. Thereby, the user still has full control on all of his tracks and can view them or delete them, in case he does not want a specific track to be shared. The enviroCar server is implemented as a RESTful Web Service with a MongoDB at the backend. It receives new tracks as JSON and provides several additional formats like CSV for download.

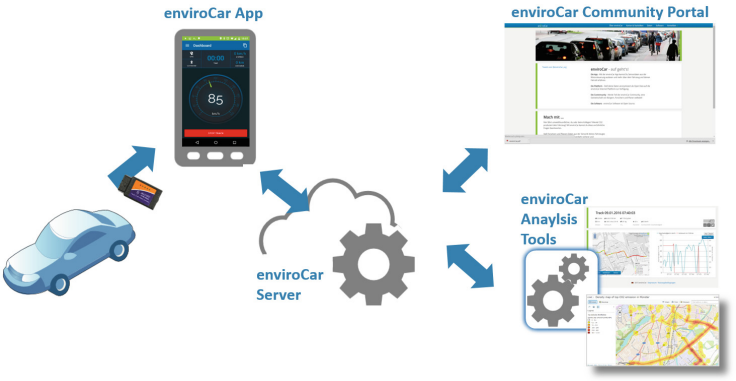

Fig. 1. Overview on the enviroCar infrastructure

The enviroCar community portal serves as the main entry point for enviroCar members. Members can explore and analyze their own tracks, compare their driving statistics to other members, and share tracks via social media platforms like Facebook

[1] General information about the project can be found at http://www.envirocar.org.

or Twitter. Several additional analysis tools using the open enviroCar data set (or subsets of it) are currently developed. These include, for instance, an R package[2] allowing to load enviroCar tracks into R and to apply further statistical analysis or interpolations. Based upon this, a fuzzy-based map matching algorithm following Quddus [2] has been implemented in R to match the track measurements to street segments in OSM[3]. In addition, several online maps that aggregate the tracks are available, e.g. for showing emission hotspots or aggregated speed measurements[4].

Current research of the enviroCar project focuses on improving and automating the map matching of tracks, on developing common interfaces and tools to collaborate on and exchange analysis functionality and discuss analysis results. Other topics include statistical analysis of trajectories. As fuel consumption and emissions are not directly measured, they need to be estimated from other parameters like mass air flow and lambda voltage sensors which measure the proportion of oxygen exhaust. Improving this estimation and accessing the uncertainty in the estimates is a further topic of current research.

Future potential applications of the enviroCar infrastructure are manifold: Urban and traffic planners can use the platform for discussing traffic measures and monitoring the measures' effects with the public. Scientists may utilize the data for developing and evaluating novel analysis methods and algorithms. As an example, first attempts for using the data for consumption-based routing have resulted in promising results. However, for this purpose the data base still needs to be enlarged and problems like, for example, selection bias need to be considered.

While the current data is gathered from fuel-powered cars, we also consider the enviroCar infrastructure as a basis for monitoring the future deployment of e-cars. The approach for consumption-based routing may also be applied to e-cars. Furthermore, information about power consumption of individual drivers may be used to derive individual ranges that drivers may reach without re-charging the battery.

References

1. Bröring, A., Remke, A., Stasch, C., Autermann, C., Rieke, M., Möllers, J.: enviroCar: a citizen science platform for analyzing and mapping crowd-sourced car sensor data. Trans. GIS **19**(3), 362–376 (2015)
2. Quddus, M.A.: High integrity map matching algorithms for advanced transport telematics applications. PhD thesis, Imperial College London, UK (2006)

[2] More information about the enviroCaR package can be found at https://github.com/enviroCar/enviroCaR.

[3] The package can be downloaded from https://cran.r-project.org/src/contrib/Archive/fuzzyMM/.

[4] See the Maps & Statistics section on http://envirocar.org for more information.

From Transient Analysis to Probabilistic Model Checking of Markov Regenerative Processes

Enrico Vicario

Department of Information Engineering, University of Florence, Florence, Italy
enrico.vicario@unifi.it

Keywords: non-Markovian models • Stochastic Petri nets • Numerical solution • Markov regenerative processes • Markov renewal theory • Probabilistic model checking

1 Talk Outline

In the engineering of systems exposed to the intertwined effects of concurrency and uncertainty, verification of quantitative properties of stochastic models enables early assessment of design choices and provides model driven guidance for implementation and integration stages. To this end, probabilistic model checking enables a systematic practice through which the same model can be verified against multiple probabilistic properties specified in some well defined language, able to analyze the impact on quality of different patterns of behavior, and open to automated regression verification when the model evolves.

Empirical evidence [3] shows that most quantitative requirements encountered in the construction of software intensive systems can be effectively expressed through a set of probabilistic specification patterns, where the most prominent role is played by the probabilistic until operator $\mathcal{P}_{\geq p}\{\phi_1\ Unt^{[\alpha,\beta]}\phi_2\}$ which specifies that: with probability not lower than p, some property ϕ_2 will be eventually satisfied within the time bound $[\alpha, \beta]$ and property ϕ_1 is satisfied in all the states visited until that time.

A number of techniques and tools have been proposed, relying on statistical discrete event simulation or numerical solution. In particular, numerical solution approaches aim at computing results with high accuracy and confidence through exhaustive state-space analysis, often relying on some restriction on the class of models amenable to verification. In the most notable case, if all model durations are exponentially distributed (EXP), the model always satisfies the Markov condition, and an efficient numerical solution can be attained by composition of behaviors according to a renewal argument referred to the time point α [1, 2].

However, the construction of a valid model may require that some durations break the EXP memoryless property and be generally distributed (GEN), as occurring for instance in aging processes accumulating memory over time, or in real-time systems or network protocols where correctness depends on firm time bounds. In a more

philosophical perspective, since the properties that are being verified capture a firm requirement on the time interval $[\alpha, \beta]$ in which ϕ_2 must be satisfied, it is much likely that the system under verification will rely on structural mechanisms enforcing firmly bounded response times.

Unfortunately, when the model includes GEN durations, the state of the system will depend on time elapsed between past events, and the Markov condition can be satisfied only at some special regeneration points. In this case, probabilistic model checking becomes much harder, combining together the complexities of non-Markovian analysis with the additional constraints posed by the model checking formulation. In a structural perspective, much of this depends on the overlapping memories contributed by durations in the model and by the time constraints in the property specification.

In this talk, we recall the salient traits of the method of stochastic state classes [4] implemented in the Oris tool (www.oris-tool.org) for transient analysis of models with multiple concurrent GEN durations. We specifically focus on the class of models that always encounter a regeneration within a bounded number of steps, and we report on recent results [5] that exploit stochastic state classes as a measure of probability over sets of runs and apply the principles of Markov regenerative analysis to enable efficient evaluation of a probabilistic until operator. The outlined solution, also provides the basis for a reflection about hurdles and structural limits that arise when Markov regenerative analysis is cast in the shape of probabilistic model checking.

References

1. Baier, C., Haverkort, B., Hermanns, H., Katoen, J.-P.: Model-checking algorithms for continuous-time Markov chains. IEEE Trans. Softw. Eng. **29**(6), 524–541 (2003)
2. Donatelli, S., Haddad, S., Sproston, J.: Model checking timed and stochastic properties with CSL^{TA}. IEEE Trans. Softw. Eng. **35**(2), 224–240, (2009)
3. Grunske, L.: Specification patterns for probabilistic quality properties. In: ICSE 2008, pp. 31–40. ACM, May 2008
4. Horváth, A., Paolieri, M., Ridi, M., Vicario, E.: Transient analysis of non-Markovian models using stochastic state classes. Perform. Eval. **69**(7–8), 315–335 (2012)
5. Paolieri, M., Horváth, A., Vicario, E.: Probabilistic model checking of regenerative concurrent systems. IEEE Trans. Softw. Eng. Accepted August 2015 (to appear)

Critical Machine-to-Machine Communications: Performance Models vs. Reality in the 10^{-10} Regime

James Gross

KTH, Stockholm, Sweden
james.gross@ee.kth.se, www.jamesgross.org

Abstract. Over the last few years, so called critical machine-to-machine communications has received more and more research attention. Spurred by flexibility and cost constraints in various industries, this area refers to wireless communication systems that can guarantee extremely high reliabilities at rather low latencies. Envisioned requirements reach down to maximum application layer packet error rates of 10^{-10} over latencies of a few milliseconds. While such systems potentially have a big relevance for safety-critical applications in industry, it is open how such systems should be designed.

In this talk, we will address selected design issues of such systems from a practical and theoretical perspective by employing communication-theoretic arguments, stochastic network calculus and probabilistic model checking. We will show that the area consists of a rich set of mostly open performance evaluation questions: Under which conditions is such communication possible at all? Which system components play a key role for the performance? Do model-based findings carry over to practical settings? Which methods can be employed to develop such systems in practise?

Contents