

Performance Evaluation for Network Services, Systems and Protocols

Stênio Fernandes

Performance Evaluation for Network Services, Systems and Protocols

Foreword by Professor Antonio Pescapè

 Springer

Stênio Fernandes
Centro de Informática
Universidade Federal de Pernambuco
Recife, Pernambuco, Brazil

ISBN 978-3-319-54519-6 ISBN 978-3-319-54521-9 (eBook)
DOI 10.1007/978-3-319-54521-9

Library of Congress Control Number: 2017932749

© Springer International Publishing AG 2017

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. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Printed on acid-free paper

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

Foreword

The advancement of computer networks has been remarkable in these recent times. The processes that were commonly used to facilitate the operations of these networks have quickly become obsolete, giving way to quicker and better technology and ultimately better computer networks. The increased use of better virtualization technologies has brought about a collaborative effort to improve the operations of the computer network systems, the services rendered, and other important protocols. The result is very impressive, meeting the expectation of consumers in terms of reliability and speed. The changes experienced in network virtualization (NV), software-defined networking (SDN), network functions virtualization (NFV), and other similar fields have made it important to focus more on the means through which the actual performance can be determined and evaluated for newer innovations. These performance evaluations are crucial to the service providers so they can design and plan their future networks.

I can confirm the importance of this book to the entire computer networking research and industrial community. I have played an active role in the research community of network monitoring and measurement since 2000. I have made useful contributions in the fields of traffic analysis and modeling, traffic classification, traffic generation, performance monitoring, network security, and also cloud and SDN monitoring. I am a full professor at the University of Naples Federico II, where I lecture students on computer networks and analysis of Internet performance. I have coauthored over 180 research papers published in international journals (e.g., *IEEE/ACM Transactions on Networking*, *Communications of the ACM*, *IEEE TPDS*, *IEEE TNSM*, *Computer Networks*, etc.) and conferences (e.g., SIGCOMM, NSDI, Infocom, IMC, PAM, Globecom, ICC, etc.). I have been honored with a Google Faculty Award, several best paper awards, Microsoft and Amazon research grants, and two IRTF (Internet Research Task Force) ANRP (Applied Networking Research Prize) awards.

Prof. Stênio Fernandes has solid record of research publications in the field of computer communications and networks. He has published over 120 research papers in a number of international peer-reviewed conferences and journals. His research interests cover the crucial aspects of performance evaluation of network

and communication systems and Internet traffic measurement, modeling projects, and analysis. I can affirm that this book is a reflection of his experiences with academic and industrial research projects related to performance evaluation of computer networks. In this book, Prof. Stênio Fernandes gives a comprehensive perspective of the methods that are used to accurately determine performance evaluation of modern computer networks. In this book, the crucial and advanced features of performance evaluation techniques are clearly explained in a way that the reader will understand the methods of conducting the right evaluation plans. Taking excerpts from the scientific literature, this book addresses the most relevant aspects of experimentation, simulation, and analytical modeling of modern networks. The readers will have a better understanding of applied statistics in computer networking and how the functions on theory and the best practices in the field intersect. This book will identify the current challenges that industrial and academic researchers face in their work and also the potentials for better innovations in this field.

University of Naples Federico II
Naples, Italy

Antonio Pescapè

Acknowledgments

I have nursed the dream of writing this book for a very long time. My position as a member of technical program committees serving for a large number of important scientific conferences in the computer networking field has given me the opportunity to witness in wonder the number of excellently written papers that have been rejected due to lapses and lack of rigor in their performance evaluation and analysis. It is common to see authors come up with brilliant ideas, but they fail to scientifically prove the validity of these ideas. A poor performance evaluation will cast aspersions on any papers' claims about its contributions and relevance to the field. The case is the same for scientific journals; I have been privileged to act as a referee for many important journals in the field of computer networks and communications. Going through the exhibition area during a scientific conference, I met Susan Lagerstrom-Fife, an editor (Computer Science) at Springer, USA. After the usual pleasantries, I asked her about the requirements needed to write a book for Springer. I got some useful information and took action, and I can happily say that this book is the result of that productive conversation. I would like to thank Susan and her assistant Jennifer Malat for guiding me along this long road.

It was an interesting and difficult experience writing this book. I experienced what writers call “the writer’s block” often. Now I know how real it is, and I can confirm that it is not a very happy experience. I was able to overcome this challenge by reading good books on focus and productivity. I owe a lot of my success in overcoming this challenge to Barbara Oakley whose course “Learning How to Learn” on Coursera played a vital role in helping me develop my mind and sharpen my skills at a higher level. I was very happy to have the opportunity to thank her in person when she came to give a talk at Carleton University, Ottawa in Canada, in May 2016. I will not stop expressing my sincere gratitude to her for putting out all that useful information for free.

Communicating your ideas to a diverse audience is not a very easy task. Writing a book chapter that entails the reviews of essential concepts of statistics was difficult to organize and deliver. I would like to thank Alexey Medvedev, who has a PhD in mathematics (2016) from Central European University, for assessing all the equations and mathematical concepts in that chapter.

I would also like to thank all my colleagues from universities around the world, most especially from the Universidade Federal de Pernambuco (Brazil), the University of Ottawa (Canada), and Carleton University (Canada) for the encouragements and kind support that helped me finish this book. I send special thanks to my former supervisor Professor Ahmed Karmouch (University of Ottawa) and my colleague Professor Gabriel Wainer (Carleton University). I also would like to extend my sincere gratitude to many network engineers I met at the Internet Engineering Task Force meetings between 2014 and 2017; I would like to thank them for the support and tips they offered me while I was writing this book.

Finally, I would like to thank my family and friends for showing their concerns with the regular question – “How’s the book writing going?” Many of the challenges I faced while writing this book made me to sometimes be unavailable and impatient. I promise to catch up with you all over coffee, wine, music concerts, and physical activities. This book would not have been possible without the love, support, and appreciation of my work expressed by my wife Nina and my children Victor and Alice. I also wish to dedicate this book to my mother Penha and my father (in memoriam) Fernando.

Contents

- 1 Principles of Performance Evaluation of Computer Networks 1**
 - 1.1 Motivation: Why Do We Need to Assess the Performance of Computer Networks? 1
 - 1.2 Classical and Modern Scenarios: Examples from Research Papers. 2
 - 1.2.1 Performance Evaluation in Classical Scenarios. 3
 - 1.2.2 Performance Evaluation in Modern Scenarios. 16
 - 1.3 The Pillars of Performance Evaluation of Networking and Communication Systems. 30
 - 1.3.1 Experimentation/Prototyping, Simulation/Emulation, and Modeling 30
 - 1.3.2 Supporting Strategies: Measurements 39
 - References. 40
- 2 Methods and Techniques for Measurements in the Internet 45**
 - 2.1 Passive vs. Active vs. Hybrid Measurements. 47
 - 2.2 Traffic Measurements: Packets, Flow Records, and Aggregated Data 51
 - 2.3 Sampling Techniques for Network Management. 54
 - 2.4 Internet Topology: Measurements, Modeling, and Analysis 57
 - 2.4.1 Guidelines for creation, selection, and registration of an Autonomous System (AS) 59
 - 2.4.2 Internet Topology Discovery: Tools, Techniques, and Datasets 61
 - 2.5 Challenges for Traffic Measurements and Analyses in Virtual Environments 64
 - 2.5.1 Cloud Computing Environments 64
 - 2.5.2 Virtualization at Network Level. 67
 - 2.6 Bandwidth Estimation Methods. 69
 - References. 70

3	A Primer on Applied Statistics in Computer Networking	75
3.1	Statistics and Computational Statistics	75
3.2	I'm All About That Data	78
3.3	Essential Concepts and Terminology	80
3.4	Descriptive Statistics	82
3.4.1	I Mean It (Or Measures of Centrality)	83
3.4.2	This Is Dull (Or Measures of Dispersion)	83
3.4.3	Is It Paranormally Distributed? (Or Measures of Asymmetry and Tailedness)	86
3.5	Inferential Statistics	89
3.5.1	Parameter Estimation: Point vs. Interval	90
3.5.2	Estimators and Estimation Methods	95
3.6	The Heavy-Tailed Phenomenon	99
3.6.1	Outlier Detection	100
3.6.2	Heavy-Tailed Distributions and Its Variations (Subclasses)	102
3.6.3	Evidence of Heavy-Tailedness in Computer Networks	107
	References	111
4	Internet Traffic Profiling	113
4.1	Traffic Analysis	114
4.1.1	Identification and Classification	114
4.1.2	Techniques, Tools, and Systems for Traffic Profiling	118
4.2	Industrial Approach for Traffic Profiling: Products and Services	130
4.3	Traffic Models in Practice	132
4.3.1	Workload Generators	133
4.4	Simulation and Emulation	137
4.4.1	Discrete-Event Simulation and Network Simulation Environments	140
4.4.2	Practical Use of Network Simulators and Traffic Profiles	147
	References	148
5	Designing and Executing Experimental Plans	153
5.1	Designing Performance Evaluation Plans: Fundamentals	153
5.2	Design of Experiments (DoE)	156
5.2.1	The DoE Jargon	159
5.2.2	To Replicate or to Slice?	161
5.3	DOE Options: Choosing a Proper Design	164
5.3.1	Classification of DOE Methods	165
5.3.2	Notation	166

5.4	Experimental Designs	166
5.4.1	2^k Factorial Designs (a.k.a. Coarse Grids)	166
5.4.2	2^{k-p} Fractional Factorial Designs	167
5.4.3	m^k Factorial Designs (a.k.a. Finer Grids)	167
5.5	Test, Validation, Analysis, and Interpretation of DOE Results	167
5.6	DOEs: Some Pitfalls and Caveats	168
5.7	DOE in Computer Networking Problems	169
5.7.1	General Guidelines	169
	References	174
 Erratum to: Performance Evaluation for Network Services, Systems and Protocols		 E1