

SpringerBriefs in Computer Science

Series editors

Stan Zdonik, Brown University, Providence, Rhode Island, USA

Shashi Shekhar, University of Minnesota, Minneapolis, Minnesota, USA

Xindong Wu, University of Vermont, Burlington, Vermont, USA

Lakhmi C. Jain, University of South Australia, Adelaide, South Australia, Australia

David Padua, University of Illinois Urbana-Champaign, Urbana, Illinois, USA

Xuemin (Sherman) Shen, University of Waterloo, Waterloo, Ontario, Canada

Borko Furht, Florida Atlantic University, Boca Raton, Florida, USA

V.S. Subrahmanian, University of Maryland, College Park, Maryland, USA

Martial Hebert, Carnegie Mellon University, Pittsburgh, Pennsylvania, USA

Katsushi Ikeuchi, University of Tokyo, Tokyo, Japan

Bruno Siciliano, Università di Napoli Federico II, Napoli, Italy

Sushil Jajodia, George Mason University, Fairfax, Virginia, USA

Newton Lee, Newton Lee Laboratories, LLC, Tujunga, California, USA

More information about this series at <http://www.springer.com/series/10028>

Ye Tian · Min Zhao · Xinming Zhang

Internet Video Data Streaming

Energy-Saving and Cost-Aware Methods

Ye Tian
School of Computer Science
and Technology
University of Science
and Technology of China
Hefei, Anhui
China

Xinming Zhang
School of Computer Science
and Technology
University of Science
and Technology of China
Hefei, Anhui
China

Min Zhao
School of Computer Science
and Technology
University of Science
and Technology of China
Hefei, Anhui
China

ISSN 2191-5768 ISSN 2191-5776 (electronic)
SpringerBriefs in Computer Science
ISBN 978-981-10-6522-4 ISBN 978-981-10-6523-1 (eBook)
<https://doi.org/10.1007/978-981-10-6523-1>

Library of Congress Control Number: 2017954878

© The Author(s) 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 Nature Singapore Pte Ltd.
The registered company address is: 152 Beach Road, #21-01/04 Gateway East, Singapore 189721, Singapore

Preface

Video has become more and more prevalent on the Internet. Studies show that YouTube has already accounted for 19.27 and 13.19% of the Internet traffics in Europe and in North America, respectively, in 2014, and it is estimated that by the year 2020, 82% of the global IP traffic will be video content. Most of Internet videos are distributed through Content Delivery Networks (CDNs). A CDN is a large distributed system consisting hundreds of thousands of dedicated content servers. Unlike peer-to-peer video streaming networks, running a CDN for holding a large-scale video service is expensive, as there are considerable costs on bandwidth usage as well as the energy consumed by the network servers.

In this book, we first review the key issues and design choices of several representative CDNs, for instance, Akamai and Google, and discuss the energy-saving techniques for server clusters, data centers, and CDNs. We then tackle the problem of saving a video streaming CDN's operating expense, including both its energy cost and the traffic cost. From our measurement study on the CDN infrastructure of Youku, which is the largest video service site in China, we find that there exists an inherent conflict between improving a video streaming CDN's energy efficiency for power saving and maintaining the CDN's ISP-friendly server selection policy. To solve this conflict, we propose a cost-aware capacity provisioning algorithm, which dynamically plans the service capacities of a CDN's server clusters in numerous ISPs and optimizes its overall operating cost regarding both the energy consumptions and the cross-ISP traffics. By using the workload derived from real-world measurement and applying actual power and bandwidth price parameters, we show with simulation experiments that our approach can significantly reduce a video streaming CDN's overall operating cost and avoid frequent server switches effectively. Finally, we discuss the directions for the future research.

For understanding the book, readers should have general knowledge on computer networks and algorithms. We hope the book can be helpful for the audience, in particular for the researchers in the Internet and networking area.

Hefei, Anhui, China
August 2017

Ye Tian
Min Zhao
Xinming Zhang

Acknowledgements

We would like to thank our colleagues, who provide valuable insights and advices during the book preparation. We also want to thank the editors and staff at Springer for their assistance. This book is not possible without their enormous help. Finally, we owe our gratitude to our families, who support us throughout this book.

We would like to acknowledge the supports from the National Natural Science Foundation of China (Project Nos. 61672486, 61379130, and 61672485), Key Project of the New Generation Mobile Wireless Broadband Communication Networks from MIIT of China (Project No. 2017ZX03001019-004), and the Anhui Provincial Natural Science Foundation (Project No. 1608085MF126).

Contents

1	Introduction	1
1.1	Background	1
1.2	CDN Objectives	2
1.2.1	Meeting Service Level Agreement	2
1.2.2	Reducing Traffic Expense	3
1.2.3	Reducing Energy Expense	4
1.3	Challenges in CDN	4
1.3.1	System Architecture	4
1.3.2	Server Selection	5
1.3.3	Capacity Provisioning	6
1.4	Overview of the Book	6
	References	7
2	Content Delivery Networks and Its Interplay with ISPs	9
2.1	CDNs in Real World	9
2.1.1	Akamai	9
2.1.2	Limelight	11
2.1.3	Amazon CloudFront	12
2.1.4	Google	12
2.1.5	Bing	13
2.1.6	YouTube	14
2.1.7	Netflix and Hulu	15
2.1.8	CDNs in China	15
2.2	Interplay Between CDN and ISP	16
2.2.1	ISP Prospective	16
2.2.2	CDN Prospective	16
	References	17

3	Energy Management	19
3.1	Energy Saving for Data Center	19
3.1.1	DVFS-Based Energy Saving	19
3.1.2	Energy Saving with Virtualization	20
3.1.3	Energy Saving with Capacity Right-Sizing	21
3.2	Energy Saving for CDN	22
3.2.1	Energy-Aware Load Balancing	22
3.2.2	Battery-Based Power Saving	22
	References	23
4	Cost Measurement for Internet Video Streaming	25
4.1	Measurement Methodology and CDN Architecture	25
4.2	Server Selection Policy Analysis	28
4.2.1	Server Selection Characteristics	28
4.2.2	Understanding Server Selection Dynamics	29
4.2.3	ISP-Friendliness	31
4.2.4	An ISP View of ISP-Friendliness Violation	32
4.2.5	Discussion	33
4.3	Energy-Aware Capacity Provisioning	33
4.4	Implication and Motivation	34
	References	34
5	Capacity Provisioning for Video Content Delivery	35
5.1	Problem Statement	35
5.1.1	The Network Model	35
5.1.2	Cost Function	36
5.1.3	Problem Formulation	36
5.2	Capacity Provisioning Algorithm	37
5.2.1	Characteristics of CDN Workload	38
5.2.2	Estimating Global Capacity Lower Bound	39
5.2.3	Optimal ISP Capacity Allocation for Cost Saving	40
5.2.4	Allocating ISP Capacities with Reduced Server Switches	41
	Appendix	42
	References	43
6	Performance Evaluation	45
6.1	Experiment Setup	45
6.2	Evaluation and Comparison	47
6.2.1	Overall Performance	47
6.2.2	Performance from ISP Perspective	50
6.2.3	Influence of Erroneous Workload Predictions	52
6.2.4	Influence of View Session Lengths	52
6.2.5	Influence of Power and Bandwidth Prices	54
	References	55

7 Concluding Remarks	57
7.1 Conclusions	57
7.2 Future Research Directions	57
References	59