

**Wiley Series on Information
and Communication Technology**

Russell T. Hsing and Vincent K. N. Lau, Series Editor



SMART DATA PRICING



Editors

**Soumya Sen • Carlee Joe-Wong
Sangtae Ha • Mung Chiang**

WILEY

SMART DATA PRICING

**WILEY SERIES ON INFORMATION
AND COMMUNICATION TECHNOLOGY**

Series Editors: T. Russell Hsing and Vincent K. N. Lau

A complete list of the titles in this series appears at the end of this volume.

SMART DATA PRICING

Edited by

**Soumya Sen, Carlee Joe-Wong, Sangtae Ha,
and Mung Chiang**

WILEY

Copyright © 2014 by John Wiley & Sons, Inc. All rights reserved

Published by John Wiley & Sons, Inc., Hoboken, New Jersey
Published simultaneously in Canada

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning, or otherwise, except as permitted under Section 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 750-4470, or on the web at www.copyright.com. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permission>.

Limit of Liability/Disclaimer of Warranty: While the publisher and author have used their best efforts in preparing this book, they make no representations or warranties with respect to the accuracy or completeness of the contents of this book and specifically disclaim any implied warranties of merchantability or fitness for a particular purpose. No warranty may be created or extended by sales representatives or written sales materials. The advice and strategies contained herein may not be suitable for your situation. You should consult with a professional where appropriate. Neither the publisher nor author shall be liable for any loss of profit or any other commercial damages, including but not limited to special, incidental, consequential, or other damages.

For general information on our other products and services or for technical support, please contact our Customer Care Department within the United States at (800) 762-2974, outside the United States at (317) 572-3993 or fax (317) 572-4002.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic formats. For more information about Wiley products, visit our web site at www.wiley.com.

Library of Congress Cataloging-in-Publication Data:

Smart data pricing / edited by Soumya Sen, Carlee Joe-Wong, Sangtae Ha, and Mung Chiang
pages cm

Includes index.

ISBN 978-1-118-61166-1 (hardback)

1. Telecommunication—Pricing. 2. Resource allocation. I. Sen, Soumya, 1982-
HE7631.S57 2014
384'.043—dc23 2013051204

Printed in the United States of America

ISBN: 9781118611661

10 9 8 7 6 5 4 3 2 1

CONTENTS

FOREWORD	xv
PREFACE	xvi
CONTRIBUTORS	xx
I SMART DATA PRICING IN TODAY’S ECOSYSTEM	1
1 Will Smart Pricing Finally Take Off?	3
<i>Andrew Odlyzko</i>	
1.1 Introduction	3
1.2 Telecom Mistakes	7
1.3 Voice and Other Missed Opportunities in Telecom	10
1.4 The Telecom Industry and Innovation	12
1.5 The Large Telecommunications Revenues	12
1.6 The High Potential for Profits in Telecommunications	13
1.7 Telco (R)evolutions	14
1.8 Capital Intensity	16
1.9 Mysteries of Investment, Costs, Profits, and Prices	18
1.10 A Historical Vignette: Bridger Mitchell and Flat Rates	20
1.11 Another Historical Vignette: Flat Rates for Data	24
1.12 Directions for Smart Pricing Research and Deployment	25
1.13 Growth in Demand	26
1.14 Technology Trends	27
1.15 Conclusions	28
Acknowledgments	29
References	29
2 Customer Price Sensitivity to Broadband Service Speed: What Are the Implications for Public Policy?	35
<i>Victor Glass, Stela Stefanova, and Ron Dibelka</i>	
2.1 Introduction	35
2.2 Model	38
2.3 Data	39

2.4	Variable Descriptions	39	
2.5	Results	41	
2.6	Conclusions	44	
	References	45	
3	Network Neutrality with Content Caching and Its Effect on Access Pricing		47
	<i>Fatih Kocak, George Kesidis, and Serge Fdida</i>		
3.1	Introduction	47	
3.2	Background	49	
3.3	Two Different Eyeball ISPs	51	
3.4	Three Different Congestion Points Per ISP, Fixed Caching Factors	52	
3.5	One Congestion Point Per ISP, Fixed Caching Factors	55	
3.6	Three Different Congestion Points Per ISP, Fixed Caching Factors, Multiple Providers of One of the Types	56	
3.7	Numerical Experiments	57	
3.8	Future Work	62	
	References	64	
II	TECHNOLOGIES FOR SMART DATA PRICING		67
4	Pricing under Demand Flexibility and Predictability		69
	<i>Ozgur Dalkilic, John Tadrous, Atilla Eryilmaz, and Hesham El-Gamal</i>		
4.1	Introduction	69	
4.2	Pricing Under Demand Flexibilities	71	
4.2.1	The Day-Ahead Electricity Market with Flexible Consumers	72	
4.2.2	Optimal Time-Dependent Pricing under Convexity Assumptions	77	
4.2.3	Optimal Bundle Pricing under Discreteness Assumptions	78	
4.2.4	Numerical Examples and Insights	79	
4.3	Pricing Under Predictable Demand	80	
4.3.1	Pricing for Demand Shaping and Proactive Download in Data Networks	83	
4.3.2	Cost Minimization via Proactive Data Service and Demand Shaping	86	
4.3.3	Pricing Policies Attaining Modified Profiles	89	
	References	92	
5	Dual Pricing Algorithms by Wireless Network Duality for Utility Maximization		97
	<i>Chee Wei Tan and Liang Zheng</i>		

5.1	Introduction	97
5.2	Utility Maximization	99
5.3	The Wireless Network Duality	103
5.3.1	Wireless Network Duality and Algorithms	105
5.3.2	Smooth and Nonsmooth Utility	105
5.3.3	Nonsmooth Special Case: $U(\boldsymbol{\gamma}) = \min_{l=1, \dots, L} \frac{\gamma_l}{\beta_l}$	106
5.3.4	Wireless Network Duality	108
5.3.5	Interference Load Minimization	112
5.3.6	Utility Maximization Algorithm	113
5.3.7	A Software Implementation	116
5.3.8	Connection between Dual Algorithm and Pricing Function in Game Theory	117
5.4	Numerical Examples	119
5.5	Conclusion	122
	References	123

6 Human Factors in Smart Data Pricing

127

Soumya Sen, Carlee Joe-Wong, Sangtae Ha, and Mung Chiang

6.1	Introduction	127
6.2	Methodology	128
6.2.1	Designing Systems with Users in Mind	128
6.2.2	Expert Evaluations	132
6.2.3	Conducting a Field Trial	133
6.2.4	Choosing an Evaluation Method	135
6.3	HCI Lessons from the Energy Market	135
6.4	User Psychology in Home Networks	136
6.4.1	Network Management and QoS Control	136
6.4.2	Implications of Throttling	138
6.4.3	Response to Capping	139
6.5	User Psychology in Bandwidth Pricing	140
6.5.1	Effects of Variable Pricing	140
6.5.2	Effects of Speed-Tier Pricing	141
6.5.3	Effects of Dynamic Time-Dependent Pricing	142
6.6	Day-Ahead Dynamic TDP	143
6.7	Perspectives of Internet Ecosystem Stakeholders	144
6.7.1	Operator Perspectives	144
6.7.2	Consumer Viewpoints	145
6.7.3	Content Provider Considerations	146
6.7.4	Application Developer Concerns	147
6.7.5	Policy Evolution	147
6.8	Lessons from Day-Ahead Dynamic TDP Field Trials	148
6.8.1	Trial Objectives	148
6.8.2	Trial Structure	148

6.8.3 Application User Interface 152
6.8.4 Trial Results 155
6.9 Discussions and Conclusions 162
Acknowledgments 164
References 164

III USAGE-BASED PRICING 167

7 Quantifying the Costs of Customers for Usage-Based Pricing 169

László Gyarmati, Rade Stanojevic, Michael Sirivianos, and Nikolaos Laoutaris

7.1 Introduction 169
7.2 The Cost of a Customer in a Network 170
7.2.1 Datasets Used in the Case Studies 171
7.3 Discrepancy, the Metric of Comparing Different Cost-Sharing Policies 172
7.4 How Do We Compute the Costs of the Customers? 173
7.4.1 Case Study: F-Discrepancy in Backbone Networks 175
7.5 Where Do We Meter the Traffic? 180
7.5.1 Case Study: M-Discrepancy in Backbone Networks 181
7.6 What Is the Impact of the Diverse Costs of the Devices? 183
7.6.1 Case Study: TCO Discrepancy in Backbone Networks 184
7.7 Who is Liable for the Incurred Costs? 185
7.7.1 Case Study: L-Discrepancy in Backbone Networks 188
7.8 Related Work 190
7.9 Conclusions 191
References 191

8 Usage-Based Pricing Differentiation for Communication Networks: Incomplete Information and Limited Pricing Choices 195

Shuqin Li and Jianwei Huang

8.1 Introduction 195
8.1.1 Related Work 197
8.2 System Model 198
8.3 Complete Price Differentiation Under Complete Information 200
8.3.1 User's Surplus Maximization Problem in Stage 2 200
8.3.2 Service Provider's Pricing and Admission Control Problem in Stage 1 200

8.3.3	Properties	204	
8.4	Single Pricing Scheme	205	
8.4.1	Problem Formulation and Solution	205	
8.4.2	Properties	206	
8.5	Partial Price Differentiation Under Complete Information	209	
8.5.1	Three-Level Decomposition	210	
8.5.2	Solving Level 2 and Level 3	212	
8.5.3	Solving Level 1	214	
8.6	Price Differentiation Under Incomplete Information	217	
8.6.1	Extensions to Partial Price Differentiation under Incomplete Information	220	
8.7	Connections with the Classical Price Differentiation Taxonomy	220	
8.8	Numerical Results	221	
8.8.1	When is Price Differentiation Most Beneficial?	221	
8.8.2	What is the Best Trade-Off of Partial Price Differentiation?	226	
8.9	Conclusion	227	
	Appendix 8.A	228	
8.A.1	Complete Price Differentiation Under Complete Information with General Utility Functions	228	
8.A.2	Proof of Proposition 8.1	231	
8.A.3	Proof of Lemma 8.2	232	
8.A.4	Proof of Theorem 8.4	233	
8.A.5	Proof of Theorem 8.6	238	
	References	238	
9	Telecommunication Pricing: Smart Versus Dumb Pipes		241
	<i>Atanu Lahiri</i>		
9.1	Introduction	241	
9.2	Uniform Ordering	243	
9.2.1	Dumb Pipe	244	
9.2.2	Smart Pipe	247	
9.2.3	Smart Pipe Versus Dumb Pipe	249	
9.3	Nonuniform Ordering	255	
9.3.1	Smart Pipe Versus Dumb Pipe Revisited	255	
9.4	Conclusion	264	
	References	266	
IV	CONTENT-BASED PRICING		267
10	Economic Models of Sponsored Content in Wireless Networks with Uncertain Demand		269
	<i>Matthew Andrews, Ulas Ozen, Martin I. Reiman, and Qiong Wang</i>		

- 10.1 Introduction 269
 - 10.1.1 Research Questions 270
 - 10.1.2 Previous Work 271
 - 10.1.3 Designing Contracts Under Uncertain Demand 272
 - 10.1.4 The Models 273
- 10.2 Analyzing Sponsored Content When EUs Pay Per Byte 276
 - 10.2.1 Content Provider’s Problem 276
 - 10.2.2 Service Provider’s Problem 277
 - 10.2.3 A Pareto Analysis of the Two-Parameter Contract 279
 - 10.2.4 Summary of the Analysis with a Contract Price c and Additional Revenue from End Users 280
 - 10.2.5 Numerical Example 281
- 10.3 Analyzing Sponsored Content in the Case of EU Quotas 282
 - 10.3.1 Case 1: Sponsorship-Insensitive Transition Probabilities 284
 - 10.3.2 Case 2: Sponsorship-Sensitive Transition Probabilities 285
- 10.4 Summary 287
 - References 287

11 CDN Pricing and Investment Strategies under Competition 289

Yang Song, Lixin Gao, and Arun Venkataramani

- 11.1 Introduction 289
- 11.2 Related Works 291
 - 11.2.1 The Pricing of a Monopoly CDN 291
 - 11.2.2 CDNs in Content Delivery Supply Chain 292
 - 11.2.3 Compare CDN and Other Multiple-Choice Markets 293
- 11.3 Background 294
 - 11.3.1 Static Analysis 294
 - 11.3.2 Predictive Analysis 295
 - 11.3.3 Dynamic Analysis 296
 - 11.3.4 Summary 300
- 11.4 Content Producers’ CDN Selection Problem 300
 - 11.4.1 Precise-Coverage Model 300
 - 11.4.2 Approximate-Coverage Model 301
- 11.5 CDN Pricing Game Under Competition 302
 - 11.5.1 Two-CDN Pricing Games 302
 - 11.5.2 The n -CDN Pricing Games 307
- 11.6 CDN Competition Under Market Structure Change 308
 - 11.6.1 Assumptions 309

11.6.2	Market State Change Through CDN Federation	309
11.6.3	The Dynamic CDN Game	311
11.7	Conclusion	317
	Acknowledgments	318
	References	318
12	Smart Pricing and Market Formation in Hybrid Networks	321
	<i>Aris M. Ouksel, Doug Lundquist, and Sid Bhattacharyya</i>	
12.1	Spectrum Shortage	321
12.2	Peer-To-Peer Networking	323
12.3	Commercial Viability	325
12.4	Self-Balancing Supply/Demand	328
12.5	Hybrid Network Model Overview	330
12.5.1	Organization	330
12.5.2	Algorithms	331
12.5.3	Hardware	331
12.5.4	Distributed Accounting	331
12.5.5	Network Security	332
12.6	Incentive Modeling	332
12.7	Flow Model	333
12.8	Prioritization Model	336
12.8.1	Divisible Incentives	337
12.8.2	Indivisible Incentives	338
12.9	Conclusion	338
	References	339
13	To Tax or To Subsidize: The Economics of User-Generated Content Platforms	341
	<i>Shaolei Ren and Mihaela van der Schaar</i>	
13.1	Introduction	341
13.2	Model	343
13.2.1	Intermediary	344
13.2.2	Content Producers	345
13.2.3	Content Viewers	346
13.3	Profit Maximization on User-Generated Content Platforms	346
13.3.1	Definition of Equilibrium	346
13.3.2	Optimal Content Viewing	347
13.3.3	Equilibrium Content Production	350
13.3.4	Optimal Payment Rate	352
13.3.5	Overjustification Effects	356
13.4	Extension to Heterogeneous Production Costs	356
13.5	Conclusion	361
	References	361

V	MANAGING CONTENT DELIVERY	363
14	Spare Capacity Monetization by Opportunistic Content Scheduling	365
	<i>Bell Labs and Alcatel-Lucent</i>	
14.1	Summary	365
14.2	Background	367
14.3	The Plutus Approach	368
14.3.1	Pricing Model	371
14.4	Architecture and Design	375
14.4.1	Components	376
14.4.2	Client-Side Monitoring of Available Capacity	382
14.5	Performance Evaluation	383
14.5.1	Network Utilization	383
14.5.2	Delay	384
14.5.3	User Experience	386
14.6	Conclusions and Future Work	387
	Acknowledgments	387
	References	388
15	Asynchronous Content Delivery and Pricing in Cellular Data Networks	391
	<i>Vijay Gabale, Umamaheswari Devi, Ravi Kokku, and Shivkumar Kalyanraman</i>	
15.1	Introduction	391
15.1.1	Surging Mobile Data Traffic and Declining Operator Profits	391
15.1.2	Traffic Variations and Peak-Time Congestion	392
15.1.3	Yield Management through Smart Pricing	392
15.2	User Survey	393
15.2.1	Setup and Goals	393
15.2.2	State of User QoE	394
15.2.3	Delay Tolerance	394
15.2.4	Delay Elasticity by Traffic Type	395
15.2.5	Price Sensitivity	396
15.2.6	Adoption	396
15.2.7	Pricing Interface	397
15.3	Time-Shifting Traffic	398
15.3.1	Time-Shifting Taxonomy	398
15.3.2	Comparison of the Time-Shifting Alternatives	400
15.4	Pricing to Enable Delivery-Shifting	402
15.4.1	Computing (Price, EDT) Options	402
15.4.2	Integration with an MNO's Infrastructure	404

15.5	Simulation Results	406	
15.5.1	Performance Measures	406	
15.5.2	Simulation Setup	406	
15.5.3	Results	408	
15.6	Conclusion	411	
	References	412	
16	Mechanisms for Quota Aware Video Adaptation		415
	<i>Jiasi Chen, Amitabha Ghosh, and Mung Chiang</i>		
16.1	Introduction	415	
16.1.1	Two Conflicting Trends	415	
16.1.2	Current Approaches in Practice	416	
16.2	Related Work	417	
16.2.1	Video Adaptation	417	
16.2.2	Video Streaming Protocols	417	
16.2.3	Quota Aware Video Adaptation	418	
16.3	A Potential Solution: QAVA	418	
16.3.1	Trading off Quality Versus Cost Versus Volume	418	
16.3.2	Incentives for Players in QAVA Ecosystem	419	
16.3.3	Design Considerations	420	
16.4	QAVA System Design	421	
16.4.1	A Modular Architecture Design	421	
16.4.2	Module Placement	423	
16.4.3	QAVA Operational Example	424	
16.5	Stream Selection	425	
16.5.1	Video Request, Utility, and Cost Model	425	
16.5.2	Stream Selection as Knapsack Problems	427	
16.5.3	Solving Finite-Horizon Markov Decision Process	429	
16.6	User and Video Profilers	430	
16.6.1	Profiling User Behavior	430	
16.6.2	Profiling Video Cost and Utility	432	
16.7	Performance Evaluation	433	
16.7.1	Experimental Setup	433	
16.7.2	Comparing Stream Selection Algorithms	434	
16.7.3	Single-User Examples	434	
16.7.4	Multiuser Stream Selection	434	
16.7.5	Sensitivity to Prediction Error	437	
16.8	Conclusions	438	
	References	438	
17	The Role of Multicast in Congestion Alleviation		441
	<i>Alan D. Young</i>		
17.1	Congestion in Cellular Networks	441	

- 17.2 Video, The Application 442
- 17.3 Why is Unicast not Ideal for All Video? 444
- 17.4 Why is Multicast Better for Video in Some Circumstances? 445
- 17.5 Broadcast, Multicast, and Unicast Architectures for the Delivery of Video 447
- 17.6 Future Potential Architectures Mixing Broadcast, Multicast and Unicast 449
- 17.7 Conclusions 450
Reference 451

VI PRICING IN THE CLOUD 453

18 Smart Pricing of Cloud Resources 455

Yu Xiang and Tian Lan

- 18.1 Data Center VM Instance Pricing 457
 - 18.1.1 Dynamic Scheduling and Server Consolidation for Fixed Pricing Scheme 457
 - 18.1.2 Price Estimation for the Uniform Pricing Scheme 458
- 18.2 Data Center SLA-Based Pricing 461
- 18.3 Data Center Time-Dependent Pricing 466
 - 18.3.1 Electricity Cost 467
 - 18.3.2 Workload Constraints 468
- 18.4 Conclusion and Future Work 474
References 474

19 Allocating and Pricing Data Center Resources with Power-Aware Combinatorial Auctions 477

Benjamin Lubin and David C. Parkes

- 19.1 Introduction 477
 - 19.1.1 Related Work 478
- 19.2 A Market Model of Data Center Allocation 480
 - 19.2.1 Buyer Valuation Model 482
 - 19.2.2 Defining The Goods in the Market 486
 - 19.2.3 Seller Cost Model 487
- 19.3 Experimental Results 489
 - 19.3.1 Scalability and Complexity 492
- 19.4 Going Beyond Processing and Power 493
- 19.5 Pricing 495
- 19.6 Conclusions 497
Acknowledgments 498
References 498

FOREWORD

Smart phones, tablets, and other video and music streaming devices fuel an exploding demand for network, cloud, and content services. Providers find it difficult to increase revenue to match the investments required to address this demand. The wireless networks are getting stressed and the quality of service suffers. The experience of other industries suggests that smarter pricing mechanisms might improve the matching of resources and users and the revenue of providers, thereby increasing user welfare both in the short term and long term. Researchers are exploring this possibility and a number of recent workshops on this topic attest to the perceived urgency of developing effective approaches.

This collection of papers presents the analysis of the pricing of network services and content conducted by leading researchers from industry and academia. The topics include the following: the tension between the users' preference for simple tariffs and potential benefits of more complex schemes; the users' sensitivity to quality of service and their willingness to shift demand; economic incentives for efficient caching and infrastructure improvements; and pricing schemes for content and for cloud resources.

Researchers will welcome this timely and broad coverage of Smart Data Pricing (SDP).

JEAN WALRAND
University of California, Berkeley, CA

PREFACE

As the demand for data in both wired and wireless broadband networks continues to grow every year, Internet Service Providers (ISPs) are increasingly turning to pricing both as a congestion management tool and as a revenue generation model. This evolution in the pricing regime is evidenced by the elimination of flat-rate plans in favor of \$10/GB or higher usage based overage fees in the United States and various other countries in Asia and Europe. This rapid transition from unlimited data plans to a reign of penalty-based mechanisms, including throttling, capping, and usage-based fees, all within a span of just 4 years as witnessed in the United States is shown in Figure 1. Consequently, Smart Data Pricing (SDP) will play a major role in the future of mobile, broadband, and content. SDP refers to a departure from the traditional flat-rate or byte-counting models to considering pricing as a network management solution. Thus, SDP will impact not only end users and network operators, but will also engage content providers, policy makers, mobile advertisers, vendors, and device suppliers. SDP incorporates the following principles:

1. ***Pricing for end-user Quality of Experience (QoE) and not just linear byte-counting:*** Simple policies like usage-based pricing (byte-counting) (i) force users to pay the same amount per unit of bandwidth consumed irrespective of the congestion levels on the network, and (ii) fail to account for the fact that different applications have different bandwidth requirements to attain a certain QoE for the user. SDP should try to match the price for delivering application-specific desired QoE requirements of the user to the ISP's congestion cost at the time of delivery.
2. ***Application layer control to impact physical layer resource management:*** Today's smart devices, with their easy-to-use graphical user interfaces, can potentially enable consumer-specified choice for access quality. Whether done manually or in an automated mode, users' specifications of their willingness to pay for their desired QoE of different applications can be taken in as inputs at the APP layer and used to control PHY layer resource allocation and media selection (e.g., WiFi offloading versus 3G). But enabling this interaction requires consumer trials to understand how to design incentives and create interfaces that can be effective in modifying end-user behavior.

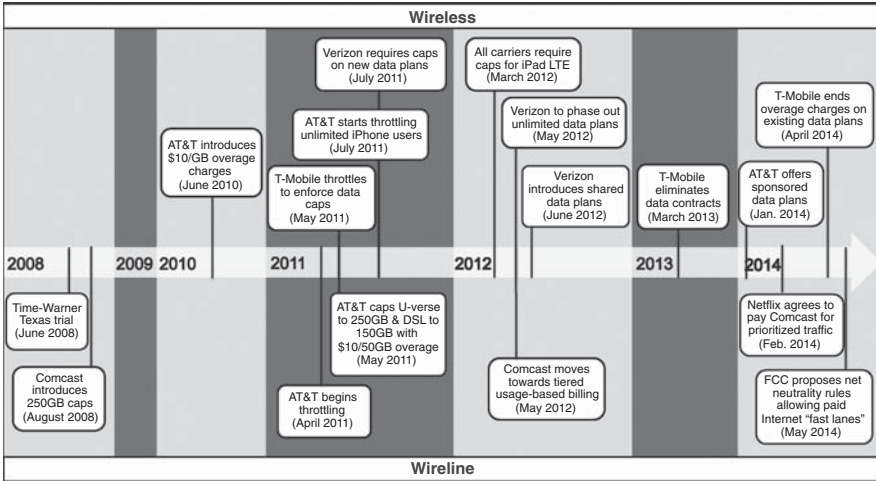


Figure 1 Timeline of the evolution in pricing plans in the United States.

3. *Incorporating edge devices as a part of the network management system:*

Instead of only managing traffic in the network core, SDP explores ways to make edge devices (e.g., smart mobile devices and customer-premise equipments like gateways) a part of the network resource allocation and management system. For example, instead of throttling traffic in the network core using the policy charging and rules function (PCRF), the edge devices (e.g., home gateways) themselves can locally regulate demand based on a user’s budget, QoE requirements, and network load or available prices. Such measures to push control from the network core out to end users, while preserving the end-to-end principles of the Internet, have been gaining attention among networking research groups (for example, the M3I¹ collaboration in Europe).

SDP can refer to (a) time/location/app/congestion dependent dynamic pricing, (b) usage-based pricing with throttling/booster, (c) WiFi offloading/proactive caching, (d) two-sided pricing/reverse billing/sponsored content, (e) quota-aware content distribution, (f) shared data pricing, and any combination or extension of the above. For instance, two-sided pricing can include QoE enhancements, or it may simply refer to content providers partially subsidizing data. SDP can benefit end users, network operators, and content providers by improving users’ Quality of Experience; lowering ISP congestion and CapEx/OpEx, thus increasing their revenue/profit margin and decreasing churn, and encouraging more consumption and ad revenue for content/app providers. But to realize these benefits, SDP requires pricing models that capture the interplay between technical and economic factors, as well as interfaces between network providers and content & application providers; effective user interface designs;

¹<http://www.m3i.org/>

field trials; and a combination of smart ideas, systematic execution, and informed policy.

This volume of collected essays on SDP has immensely benefitted from the annual SDP Forum, which organizes workshops to bring together industry experts, academics, and regulators for in-depth discussions on the topic. SDP 2012 was held in Princeton, New Jersey, and the SDP 2013 and 2014 Workshops were held in conjunction with IEEE INFOCOM in Turin, Italy and Toronto, Canada. The workshops have been attended by professionals from AT&T, Verizon, Comcast, NECA, Alcatel-Lucent, Cisco Systems, Qualcomm, Microsoft, ACS, and many other leading networking companies. It therefore comes with little surprise that several of the chapters in this volume have been contributed by industry researchers and showcase some cutting-edge research in this area.

The first three chapters of this book discuss SDP's feasibility in the current Internet ecosystem. The first chapter looks back on previous efforts to promote SDP and asks whether the current market climate will be more receptive. The next chapter approaches SDP's feasibility from a customer perspective, using empirical data to examine their price sensitivity. Finally, the third chapter incorporates regulatory concerns by examining network neutrality in the context of content caching.

The next three chapters address SDP's technical feasibility. The first chapter in this section develops a pricing model that accounts for the flexibility and predictability of customer demand. The second chapter focuses on wireless networks, showing how pricing can be used to make wireless resource allocation more efficient. The last chapter focuses on SDP's interface between ISPs and users, examining how the ISP can communicate prices to users through interfaces on their devices.

The next three chapters of the book shift to variants on usage-based pricing, a particular form of SDP. The first chapter examines whether usage-based pricing can in fact help ISPs by quantifying the distribution of infrastructure costs among ISP customers. The next two chapters then turn to differentiated pricing: the first of these develops a model for differentiated usage-based pricing, while the second examines the benefits of non-differentiated and differentiated pricing for ISPs and end users.

Another form of SDP, content-based pricing, is discussed in the next four chapters. The first chapter discusses a variant of usage-based or capped pricing, in which content providers subsidize the delivery of their content to end users, sponsoring users' Internet access. The second chapter shifts the focus to content delivery networks and the impact of competition on their pricing and investment, while the third chapter discusses the economics of a hybrid model in which content delivery can be offloaded to a secondary P2P network during congested times. The last chapter considers the economics of content providers, focusing on how the owners of user-generated content platforms, e.g., social networking websites, can best monetize this content.

The next four chapters discuss technical aspects of realizing economically efficient models of content delivery. The first chapter investigates the idea of opportunistic content transfer, offloading traffic to times of lower congestion with a monetary discount given during times of lower congestion. The next chapter considers a similar idea, in which sessions like content transfers can be spread over time, but with prices determined by the deadline of each session's completion. The third chapter focuses

on video content, and shifts the focus away from ISPs to consider how a user might distribute a budget for consuming videos over time. Finally, the last chapter considers multicast technology and how it can alleviate network congestion.

The last two chapters of the book consider pricing in the cloud. The first chapter investigates and compares three different schemes for pricing data center resources, namely real-time instance pricing, deadline-based service level agreements, and time-dependent pricing. The last chapter proposes using combinatorial auctions to price and allocate resources in a data center while taking into account its electricity constraints.

The diversity of topics explored in these book chapters reflects SDP's broad potential impact. Indeed, SDP brings together ideas from such diverse fields as network engineering, economics, human-computer interaction, data science, and technology policy to answer fundamental questions about broadband pricing. Yet there remain significant emerging themes which this book does not cover. For instance, little rigorous analysis has been done on shared data plans, which have recently become mainstream in the U.S. Perhaps more significantly, "network neutrality" is emerging as a fundamental issue, with new regulations from the FCC and Netflix's agreement with Comcast to pay for a separate "fast lane" for its streaming traffic. And as more and more devices become connected to the Internet, pricing for the Internet of Things is becoming an important question. The emergence of these and other topics will ensure that SDP remains an exciting and relevant research topic in the years to come.

SOUMYA SEN, CARLEE JOE-WONG, SANGTAE HA, AND MUNG CHIANG

ACKNOWLEDGMENTS

We would like to thank all of the participants of the first, second, and third Smart Data Pricing Workshops, held respectively in Princeton, New Jersey on July 30 and 31, 2012; Turin, Italy on April 19, 2013; and Toronto, Canada on May 2, 2014. We are also grateful to all of the contributing authors for their time and effort, as well as our colleagues who served as reviewers for the contributions.

CONTRIBUTORS

MATTHEW ANDREWS, Bell Labs, Alcatel-Lucent, Murray Hill, NJ
RANDEEP BHATIA, Bell Labs, Alcatel-Lucent, Murray Hill, NJ
SID BHATTACHARYYA, Department of Information and Decision Sciences,
University of Illinois at Chicago, Chicago, IL
JIASI CHEN, Princeton University, Princeton, NJ
MUNG CHIANG, Princeton University, Princeton, NJ
OZGUR DALKILIC, The Ohio State University, Columbus, OH
UMAMAHESWARI DEVI, IBM Research, Bangalore, India
RON DIBELKA, National Exchange Carrier Association, Inc., Whippany, NJ
HESHAM EL-GAMAL, The Ohio State University, Columbus, OH
ATILLA ERYILMAZ, The Ohio State University, Columbus, OH
SERGE FDIDA, UPMC, Paris, France
VIJAY GABALE, IBM Research, Bangalore, India
LIXIN GAO, University of Massachusetts, Amherst, MA
AMITABHA GHOSH, UtopiaCompression Corporation, Los Angeles, CA
VICTOR GLASS, National Exchange Carrier Association, Inc. Whippany, NJ
BHAWNA GUPTA, Bell Labs, Alcatel-Lucent, Murray Hill, NJ
LÁSZLÓ GYARMATI, Telefonica Research, Barcelona, Spain
SANGTAE HA, University of Colorado, Boulder, CO
JIANWEI HUANG, The Chinese University of Hong Kong, Hong Kong, China
CARLEE JOE-WONG, Princeton University, Princeton, NJ
SHIVKUMAR KALYANRAMAN, IBM Research, Bangalore, India
GEORGE KESIDIS, The Pennsylvania State University, University Park, State
College, PA
FATIH KOCAK, The Pennsylvania State University, University Park, State
College, PA

- RAVI KOKKU, IBM Research, Bangalore, India
- ATANU LAHIRI, University of Washington, Seattle, WA
- TIAN LAN, George Washington University, Washington, DC
- NIKOLAOS LAOUTARIS, Telefonica Research, Barcelona, Spain
- SHUQIN LI, Alcatel-Lucent Shanghai
- BENJAMIN LUBIN, Boston University Boston, MA; Harvard University, Cambridge, MA
- DOUG LUNDQUIST, Department of Information and Decision Sciences, University of Illinois at Chicago, Chicago, IL
- ANDREW ODLYZKO, University of Minnesota, Minneapolis, MN
- ARIS M. OUKSEL, Department of Information and Decision Sciences, University of Illinois at Chicago, Chicago, IL
- ULAS OZEN, Ozyegin University, Istanbul, Turkey
- DAVID C. PARKES, Harvard University, Cambridge, MA
- MARTIN I. REIMAN, Alcatel-Lucent Bell Labs
- SHAOLEI REN, Florida International University, Miami, FL
- MIHAELA VAN DER SCHAAR, University of California, Los Angeles, Los Angeles, CA
- SOUMYA SEN, Carlson School of Management, University of Minnesota, Minneapolis, MN
- MICHAEL SIRIVIANOS, Cyprus University of Technology, Limassol, Cyprus
- YANG SONG, University of Massachusetts, Amherst, MA
- RADE STANOJEVIC, Telefonica Research, Barcelona, Spain
- STELA STEFANOVA, National Exchange Carrier Association, Inc., Whippany, NJ
- JOHN TADROUS, The Ohio State University, Columbus, OH
- CHEE WEI TAN, City University of Hong Kong, Hong Kong, China
- ARUN VENKATARAMANI, University of Massachusetts, Amherst, MA
- QIONG WANG, University of Illinois Urbana-Champaign, Champaign, IL
- YU XIANG, George Washington University, Washington, DC
- ALAN D. YOUNG, P & Y Associates, LLC
- LIANG ZHENG, City University of Hong Kong, Hong Kong, China

PART I
Smart Data Pricing in Today's
Ecosystem

1 Will Smart Pricing Finally Take Off?

ANDREW ODLYZKO

1.1 INTRODUCTION

Will smart pricing dominate telecommunications? We certainly do see growth in sophisticated pricing in many areas of the economy. Congestion charges for cars entering central business districts and “smart” electric meter deployments are spreading. Airlines are even beginning to auction seat upgrades [1]. And there is no shortage of desire for smart pricing in telecommunications. For a survey of recent developments, see Reference 2. Many new technological developments, such as software-defined networking (SDN), are touted as facilitating differentiated services and differentiated pricing. The overwhelming consensus of the industry, as well as of the research community, and of regulators, is that flat rates are irrational. Thus, for example, in 2011, Jon Leibowitz, the then-Chairman of the US Federal Trade Commission could not “quite understand why something like metering hasn’t taken off yet.” (See Reference 3 for references to this and similar recent quotes, as well as for a summary of the arguments in favor of flat rates.)

Yet there are reasons for caution in the rush to smart pricing. After all, the modern consensus about its desirability is not new. It goes back centuries, to the days of snail mail. Furthermore, industry has often either stumbled onto flat or almost flat rates, or been forced into them, all against its will, and ended up benefiting. Thus, for example, US wireless service providers have been boasting of the low per-minute voice call revenues that reign in United States, much lower than in most of the world. What they universally neglect to mention is that these low prices are the result of the success of the block-pricing plan introduced by AT&T Wireless in 1998, which also eliminated roaming and long-distance charges. This plan, the result not of a careful study of historical precedents or the economics of communications but rather the fruit of a desperate carrier looking for a way to gain customers, was widely derided but proved unexpectedly popular. It forced the rest of the industry to follow suit with

similar plans and led to large increases in voice usage (see, e.g., the chart in Reference 4). The end result is that the United States has the world's highest per-subscriber voice usage, yielding those low average per-minute prices that the industry boasts of. Probably not coincidentally, US wireless service providers are among the world's most profitable. This story, and others similar to it, should make one cautious about rushing to follow the industry consensus. This is true even when such a consensus is fortified by scholarly studies, because those tend to be even more biased toward fine-grained pricing. The telecom industry and telecom researchers have historically been notorious for not understanding what is in the industry's own interests.

The traditional preoccupation with smart pricing is likely to be reinforced by the economics of telecom. Contrary to common opinion, it is not all that capital intensive. As is demonstrated in Section 1.8, telecom is simply not in the same category as such large and important industries as electricity or roads when it comes to the ratio of capital investment to revenues. Telecom is primarily about service, customer inertia, and territorial strategic plays (where the territories may be physical or virtual).

Although the telecom industry is not very capital intensive, communications is extremely valuable and any society is willing to pay astonishing amounts for it. As an example, by some measures, the United States spends almost 50% more on telecom services than it does for electricity. (See Section 1.5 for more data and references.) Furthermore, in spite of all the complaints from the industry about its supposedly impoverished state, there appears to be very large profits in many parts of it. As this passage is being written in the summer of 2013, Verizon is in the process of buying out Vodafone's 45% stake in the Verizon Wireless unit for \$130 billion. This means that the whole of Verizon Wireless is being valued at almost \$300 billion. As will be shown in Section 1.9, that is about four times the cost of replacing all the tangible assets of that enterprise. It is also almost enough to replace the entire US telecom infrastructure, both wireless and wired, with the latter redone in fiber. This is anomalous by traditional standards, but then, as will be discussed in Section 1.9, the entire economy is behaving anomalously, with very high corporate profits, low interest rates, and low capital investment. Whether this is a temporary aberration, or whether we are in a new economic era, remains to be seen. However, telecom is very much in the mainstream of this historically unusual behavior, and so many traditional yardsticks of financial performance may not apply.

While the telecom industry has often been blind to profitable opportunities, it has always been aware that high profits are possible. However, it has usually faced difficulties in using their favorite methods for profit extraction because of various combinations of legal and regulatory constraints and the peculiar nature of demand for communication services. Table 1.1 shows an approximation of current prices paid by users for varying amounts of data from various services.

This table demonstrates the main problem faced by telecom. The most valuable information can often be conveyed in just a few bits. Thus, for example, in the early days of postal services, when receivers paid on delivery, information would often be transmitted in the form of small modifications in the address. The addressee would then scan the envelope, figure out what the message was, and refuse to accept (and pay for) the letter.

TABLE 1.1 Price per Megabyte

SMS	\$1000.00
Cellular voice	1.00
Wireline voice	0.10
Residential Internet	0.01
Backbone Internet	0.0001

Practices from two centuries ago may seem irrelevant, but in fact they are very instructive, as the basic economic issues have always been the same, even as technology has changed drastically, cf. [5]. Thus, for example, today, we have the telecom industry investing heavily in deep packet inspection. In the past, post offices had employees hold letters up against burning candles to make sure that there were no enclosures that were subject to extra fees. The basic incentive is to extract as much value as possible, and that usually requires fine-grained pricing to achieve successful price discrimination. But usually, in communication as well as in transportation, limits are placed on what service providers are allowed to do. The net neutrality debate is just another instance of the ancient conflict between economic efficiency and fairness in markets [6]. Giving unfettered control of any critical service to any provider, or an oligopoly of providers, either *de jure* or *de facto* (by allowing natural monopoly mechanisms to operate), is equivalent to abolishing property rights with the usual negative impacts on innovation and efficiency. Hence, we have almost always had constraints, such as those of common carriage. The real question is about the appropriate level of constraints.

Public talk of capacity limits is often just a public relations measure, designed to overcome opposition to service provider strategies. Thus, for example, in early 2013, Michael Powell, the President of the US cable industry association [and former Chairman of the Federal Communications Commission (FCC)] admitted, contradicting many earlier declarations by a variety of executives and experts, that “cable’s interest in usage-based pricing was not principally about network congestion, but instead about pricing fairness” [7]. Whenever business leaders talk of “fairness,” it is generally safe to assume that they are really after extracting more revenues through differential pricing. This is neither a novel nor is it nefarious. In fact, differential pricing was and is at the core of regulatory economics, as it can be used to promote social welfare, and has been frequently mandated by governments. However, historically, the degree of price discrimination that was allowed varied depending on economics, with more discrimination being allowed when the costs of providing those services have been large [8]. The question for the near future is whether modern telecom should be allowed more power to discriminate. Further, even if it is given that power, one should consider whether it would be wise to use it. The right answer depends on the balance between growth in demand and improvements in technology.

The main problem, past, present, and future, that is faced by telecom is that the most valuable information usually requires just a few bits to convey. The second main problem is that because of technological progress, transmission capacity is growing. Thus the industry is faced with the challenge of persuading users to pay for big pipes

when the additional value that enlarging those pipes provides is not all that high. (There are arguments that the value of transmission capacity, as well as that of computing power and storage, should be measured on a logarithmic scale, so that going from what is now a slow 1 Mbps link to a 1 Gbps one corresponds only to an increase in value from 6 to 9, cf. [9].) At the moment, that additional capacity is consumed largely by video. But the value is still dominated by the low bandwidth voice and texting.

The general conclusion of this work, based on the study of trends in demand and supply, is that in wireline communication, the critical issue faced by the telecom industry is not handling overpowering exafloods of traffic, as has often been claimed, cf. [10–12], but stimulating demand to fill the growing capacity of transmission systems [13]. The most effective way to do that is to offer flat rates and open access to encourage innovation. To the extent that any market segmentation is needed, it is best handled by offering flat rate services with different peak speeds. Pricing by volume of traffic (whether using caps or other schemes) may be attractive at the moment to service providers preoccupied with trying to protect their traditional subscription video service revenues. However, it is an ineffective instrument that does not address any of the issues well and, in the long run, is likely to damage not only the economy as a whole but also the profits of service providers. Any truly “smart pricing” measures, such as congestion charges, are likely to be detrimental to the industry.

These general conclusions for wired communications apply directly mainly to the richer and more industrialized countries. Even in those, there is likely to be exceptional situations where the cost structure forces some “smart pricing” approaches. For poor countries, the best choices along the frontier of feasible technological and business models is likely to lean further toward “smart pricing.” This would be consistent with the general observation, cf. [5], that at the consumer level, sophisticated pricing is most appropriate for large and relatively infrequent transactions, and simple pricing for small and frequent ones. This is also what we observe in the market today, with the greatest proliferation of “smart pricing” in less-developed countries, where the relative burden of telecommunications charges is higher.

In wireless communication, the optimal choice even in rich countries appears to be different than that for wireline, because of a different balance between feasible supply and potential demand. There have been widespread projections that wireless data traffic would continue to double each year, as it had done for several years. Those are now being disproved, as growth rates are declining (see Section 1.13). Still, those rates are high, and there is far more traffic that are likely to use the radio path if that were feasible, as wireless data traffic is under 5% of wireline. Coupled with the low value of most of this data, and the resulting low likelihood of service providers being able to extract large new revenues, it appears probable that the incentives for the industry will be to constrain usage and to implement differentiated quality of service to protect the most valuable low bandwidth applications. So somewhat finer-grained pricing is likely to prevail in this domain than in wireline. Still, the need to limit what Nick Szabo [14] has aptly called the *mental transaction costs* involved in fine-grained pricing, and related concerns, is likely to restrict the complexity of schemes that succeed. The sophisticated pricing plans so beloved of researchers are likely to be confined to

areas such as business-to-business dealings and may be of limited applicability even there.

However, the strong prejudice in favor of “smart pricing” among both industry leaders and academic researchers guarantees that many schemes will be developed, and quite a few will be deployed. Chances are that, as was true of many sophisticated prioritization schemes developed for voice private branch exchanges (PBXs) or early data switches, they will not see much use. But for those cases where they might be used, it appears that most of current research, as well as academic instruction, is missing some important ingredients. As is discussed in Section 1.12, it will likely be important to explore the most effective ways to introduce noise and other impairments into communication systems to provide differential quality of service. (On the other hand, there will likely also be demand for methods to detect such actions.)

The next section reviews briefly some of the main fallacies that invigorate the push for “smart pricing.” This is followed by a section on some missed opportunities in telecommunications, demonstrating how this industry tends to “stumble to success,” pursuing mistaken goals, and prospering by accident. Section 1.4 has a very brief discussion of the reasons telecom has been so poor at innovating in services and is likely to remain poor in the future. Section 1.5 discusses this industry’s place in the entire economy. Section 1.6 points out that high profits have not infrequently been obtained in this sector. Section 1.7 sketches the main changes that have taken place in the money flows in telecommunications in recent decades. Section 1.8 demonstrates that, contrary to general opinion, this industry is not all that capital intensive. Section 1.9 discusses some of the puzzles of the modern economy and the degree to which the telecom industry exhibits similar behavior. Section 1.12, cited earlier, discusses some missing ingredients in modern research and education, should “smart pricing” become widespread. Sections 1.10 and 1.11 take a historical look at some earlier work on telecom pricing and the degree to which it reflected the prejudices we observe today. Sections 1.13 and 1.14 then discuss the growth in the demand for data traffic and improvements in transmission technologies and what the contrasts are with those that for optimal pricing strategies. Finally, Section 1.15 provides the conclusions.

1.2 TELECOM MISTAKES

Many of the basic but general issues that have a bearing on the possible adoption of smart pricing have already been explored in the literature (see, e.g., [5, 15–17]) and so will be touched on very lightly here. However, they do need to be mentioned, because there are many misapprehensions about the nature of telecom and these issues often have an important bearing on the optimal choices of pricing policies. For example, we are constantly told that content is king. (Content is taken here to mean material prepared by professionals for wide distribution and not, as some use it, to denote anything in digital form.) But

Content is not king.