

Pricing Communication Networks

Economics, Technology and Modelling

Costas Courcoubetis

Athens University of Economics and Business, Greece

Richard Weber

University of Cambridge, UK



WILEY

Pricing Communication Networks

WILEY-INTERSCIENCE SERIES IN SYSTEMS AND OPTIMIZATION

Advisory Editors

Sheldon Ross

Department of Industrial Engineering and Operations Research, University of California,
Berkeley, CA 94720, USA

Richard Weber

Statistical Laboratory, Centre for Mathematical Sciences, Cambridge University,
Wilberforce Road, Cambridge, CB3 0WB

BATHER–Decision Theory: An Introduction to Dynamic Programming and Sequential
Decisions

CHAO/MIYAZAWA/PINEDO–Queueing Networks: Customers, Signals and Product Form
Solutions

COURCOUBETIS/WEBER–Pricing Communication Networks: Economics, Technology
and Modelling

DEB–Multi-Objective Optimization using Evolutionary Algorithms

GERMAN–Performance Analysis of Communication Systems: Modeling with
Non-Markovian Stochastic Petri Nets

KALL/WALLACE–Stochastic Programming

KAMP/HASLER–Recursive Neural Networks for Associative Memory

KIBZUN/KAN–Stochastic Programming Problems with Probability and Quantile
Functions

RUSTEM–Algorithms for Nonlinear Programming and Multiple-Objective Decisions

WHITTLE–Optimal Control: Basics and Beyond

WHITTLE–Neural Nets and Chaotic Carriers

The concept of a system as an entity in its own right has emerged with increasing force in the past few decades in, for example, the areas of electrical and control engineering, economics, ecology, urban structures, automation theory, operational research and industry. The more definite concept of a large-scale system is implicit in these applications, but is particularly evident in such fields as the study of communication networks, computer networks, and neural networks. The *Wiley-Interscience Series in Systems and Optimization* has been established to serve the needs and researchers in these rapidly developing fields. It is intended for works concerned with the developments in quantitative systems theory, applications of such theory in areas of interest, or associated methodology.

Pricing Communication Networks

Economics, Technology and Modelling

Costas Courcoubetis

Athens University of Economics and Business, Greece

Richard Weber

University of Cambridge, UK



WILEY

Copyright © 2003

John Wiley & Sons Ltd, The Atrium, Southern Gate, Chichester,
West Sussex PO19 8SQ, England

Telephone (+44) 1243 779777

Email (for orders and customer service enquiries): cs-books@wiley.co.uk

Visit our Home Page on www.wileyeurope.com or www.wiley.com

All Rights Reserved. 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 under the terms of the Copyright, Designs and Patents Act 1988 or under the terms of a licence issued by the Copyright Licensing Agency Ltd, 90 Tottenham Court Road, London W1T 4LP, UK, without the permission in writing of the Publisher. Requests to the Publisher should be addressed to the Permissions Department, John Wiley & Sons Ltd, The Atrium, Southern Gate, Chichester, West Sussex PO19 8SQ, England, or emailed to permreq@wiley.co.uk, or faxed to (+44) 1243 770620.

This publication is designed to provide accurate and authoritative information in regard to the subject matter covered. It is sold on the understanding that the Publisher is not engaged in rendering professional services. If professional advice or other expert assistance is required, the services of a competent professional should be sought.

Other Wiley Editorial Offices

John Wiley & Sons Inc., 111 River Street, Hoboken, NJ 07030, USA

Jossey-Bass, 989 Market Street, San Francisco, CA 94103-1741, USA

Wiley-VCH Verlag GmbH, Boschstr. 12, D-69469 Weinheim, Germany

John Wiley & Sons Australia Ltd, 33 Park Road, Milton, Queensland 4064, Australia

John Wiley & Sons (Asia) Pte Ltd, 2 Clementi Loop #02-01, Jin Xing Distripark, Singapore 129809

John Wiley & Sons Canada Ltd, 22 Worcester Road, Etobicoke, Ontario, Canada M9W 1L1

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic books.

Library of Congress Cataloging-in-Publication Data

Courcoubetis, Costas.

Pricing communication networks : economics, technology, and modelling / Costas

Courcoubetis, Richard Weber.

p. cm.—(Wiley-Interscience series in systems and optimization)

Includes bibliographical references and index.

ISBN 0-470-85130-9 (alk. Paper)

1. Information technology—Finance. 2. Computer networks—Mathematical models. 3. Digital communications—Mathematical models. I. Weber, Richard. II. Title. III. Series.

HD30.2 .C68 2003

384'.043—dc21

2002191081

British Library Cataloguing in Publication Data

A catalogue record for this book is available from the British Library

ISBN 0-470-85130-9

Typeset in 10/12pt Times by Laserwords Private Limited, Chennai, India

Printed and bound in Great Britain by Biddles Ltd, Guildford, Surrey

This book is printed on acid-free paper responsibly manufactured from sustainable forestry in which at least two trees are planted for each one used for paper production.

We dedicate this book to Dora and Persefoni, the muses of my life (C. Courcoubetis),
and to Richard, my father (R. Weber).

Contents

Preface	xv
List of Acronyms	xix
A Networks	1
1 Pricing and Communications Networks	3
1.1 The Market for Communications Services	3
1.1.1 The Communications Revolution	3
1.1.2 Communications Services	3
1.1.3 Information Goods	4
1.1.4 Special Features of the Communications Market	5
1.2 Developments in the Marketplace	6
1.3 The Role of Economics	9
1.3.1 Overprovision or Control?	10
1.3.2 Using Pricing for Control and Signalling	12
1.3.3 Who Should Pay the Bill?	13
1.3.4 Interconnection and Regulation	14
1.4 Preliminary Modelling	16
1.4.1 Definitions of Charge, Price and Tariff	16
1.4.2 Flat Rate versus Usage Charging	17
1.4.3 Dynamic Pricing in an Internet Cafe	18
1.4.4 A Model for Pricing a Single Link	19
1.5 A Guide to Subsequent Chapters	21
1.6 Further Reading	22
2 Network Services and Contracts	23
2.1 A Classification of Network Services	24
2.1.1 Layering	24
2.1.2 A Simple Technology Primer	25
2.1.3 Value-added Services and Bundling	28
2.1.4 Connection-oriented and Connectionless Services	30
2.1.5 Guaranteed and Best-effort Services	32
2.2 Service Contracts for Transport Services	33
2.2.1 The Structure of a Service Contract	33
2.2.2 Policing Service Contracts	36

2.2.3	Static and Dynamic Contract Parameters	37
2.3	Further Reading	39
3	Network Technology	41
3.1	Network Control	41
3.1.1	Entities on which Network Control Acts	42
3.1.2	Timescales	43
3.1.3	Handling Packets and Cells	43
3.1.4	Virtual Circuits and Label Switching	44
3.1.5	Call Admission Control	45
3.1.6	Routing	46
3.1.7	Flow Control	48
3.1.8	Network Management	50
3.2	Tariffs, Dynamic Prices and Charging Mechanisms	50
3.3	Service Technologies	51
3.3.1	A Technology Summary	51
3.3.2	Optical Networks	53
3.3.3	Ethernet	54
3.3.4	Synchronous Services	56
3.3.5	ATM Services	57
3.3.6	Frame Relay	59
3.3.7	Internet Services	60
3.4	Other Types of Services	71
3.4.1	Private and Virtual Networks	71
3.4.2	Access Services	73
3.5	Charging Requirements	76
3.6	A Model of Business Relations for the Internet	77
3.7	Further Reading	82
4	Network Constraints and Effective Bandwidths	83
4.1	The Technology Set	84
4.2	Statistical Multiplexing	85
4.3	Accepting Calls	86
4.4	An Elevator Analogy	87
4.5	Effective Bandwidths	90
4.6	Effective Bandwidths for Traffic Streams	91
4.6.1	The Acceptance Region	94
4.7	Some Examples	95
4.8	Multiple QoS Constraints	99
4.9	Traffic Shaping	100
4.10	Effective Bandwidths for Traffic Contracts	102
4.11	Bounds for Effective Bandwidths	103
4.12	Deterministic Multiplexing	105
4.13	Extension to Networks	107
4.14	Call Blocking	108
4.15	Further Reading	109

B	Economics	111
5	Basic Concepts	113
5.1	Charging for Services	113
5.1.1	Demand, Supply and Market Mechanisms	113
5.1.2	Contexts for Deriving Prices	114
5.2	The Consumer's Problem	116
5.2.1	Maximization of Consumer Surplus	116
5.2.2	Elasticity	118
5.2.3	Cross Elasticities, Substitutes and Complements	118
5.3	The Supplier's Problem	119
5.4	Welfare Maximization	120
5.4.1	The Case of Producer and Consumers	120
5.4.2	The Case of Consumers and Finite Capacity Constraints	123
5.4.3	Discussion of Assumptions	124
5.4.4	Peak-load Pricing	125
5.4.5	Walrasian Equilibrium	126
5.4.6	Pareto Efficiency	127
5.4.7	Discussion of Marginal Cost Pricing	130
5.5	Cost Recovery	131
5.5.1	Ramsey Prices	131
5.5.2	Two-part Tariffs	133
5.5.3	Other Nonlinear Tariffs	135
5.6	Finite Capacity Constraints	137
5.7	Network Externalities	138
5.8	Further Reading	140
6	Competition Models	141
6.1	Types of Competition	141
6.2	Monopoly	143
6.2.1	Profit Maximization	143
6.2.2	Price Discrimination	144
6.2.3	Bundling	148
6.2.4	Service Differentiation and Market Segmentation	149
6.3	Perfect Competition	151
6.3.1	Competitive Markets	152
6.3.2	Lock-in	152
6.4	Oligopoly	154
6.4.1	Games	154
6.4.2	Cournot, Bertrand and Stackelberg Games	157
6.5	A Unifying Social Surplus Formulation	160
6.6	Further Reading	160
C	Pricing	161
7	Cost-based Pricing	163
7.1	Foundations of Cost-based Pricing	163
7.1.1	Fair Charges	164

7.1.2	Subsidy-free, Support and Sustainable Prices	165
7.1.3	Shapley Value	170
7.1.4	The Nucleolus	172
7.1.5	The Second-best Core	172
7.2	Bargaining Games	174
7.2.1	Nash's Bargaining Game	174
7.2.2	Kalai and Smorodinsky's Bargaining Game	176
7.3	Pricing in Practice	177
7.3.1	Overview	177
7.3.2	Definitions Related to the Cost Function	179
7.3.3	The Fully Distributed Cost Approach	181
7.3.4	Activity-based Costing	184
7.3.5	LRIC+	187
7.3.6	The Efficient Component Pricing Rule	188
7.4	Comparing the Various Models	190
7.5	Flat Rate Pricing	191
7.6	Further Reading	194
8	Charging Guaranteed Services	195
8.1	Pricing and Effective Bandwidths	196
8.1.1	The Network Case	201
8.2	Incentive Issues in Pricing Service Contracts	202
8.3	Constructing Incentive Compatible Tariffs from Effective Bandwidths	204
8.3.1	The Time-volume Charging Scheme	205
8.3.2	Using General Measurements	207
8.3.3	An Example of an Actual Tariff Construction	208
8.3.4	Competition	210
8.3.5	Discouraging Arbitrage and Splitting	211
8.4	Some Simple Pricing Models	212
8.4.1	Time-of-day Pricing	212
8.4.2	Combining Guaranteed with Best-effort	213
8.4.3	Contracts with Minimum Guarantees and Uncertainty	214
8.5	Long-term Interaction of Tariffs and Network Load	216
8.6	Further Reading	218
9	Congestion	219
9.1	Defining a Congestion Price	220
9.1.1	A Condition for Capacity Expansion	222
9.1.2	Incentive Compatibility	222
9.1.3	Extensions	222
9.2	Connection with Finite Capacity Constraints	223
9.3	Models in which Users Share Congested Resources	224
9.3.1	A Delay Model for a $M/M/1$ Queue	224
9.3.2	Services Differentiated by Congestion Level	225
9.3.3	A Blocking Model	225
9.4	Congestion Prices Computed on Sample Paths	227
9.4.1	A Loss Model	228
9.4.2	A Congestion Model with Delay	229

9.4.3	Bidding for Priority	230
9.4.4	Smart Markets	230
9.5	An Incentive Compatible Model for Congestion Pricing	231
9.6	Further Reading	232
10	Charging Flexible Contracts	235
10.1	Notions of Fairness	237
10.2	The Proportional Fairness Model	239
10.2.1	A Primal Algorithm	241
10.2.2	A Dual Algorithm	243
10.2.3	User Adaptation	243
10.2.4	Stochastic Effects and Time Lags	244
10.2.5	Proportional Fairness with a Congestion Cost	244
10.3	An Internet Pricing Proposal	245
10.4	A Model of TCP	247
10.5	Allocating Flows by Effective Bandwidth	249
10.6	User Agents	250
10.7	Pricing Uncertainty	254
10.8	The Differentiated Services Approach	256
10.8.1	Paris Metro Pricing	257
10.9	Towards a Market-Managed Network	259
10.10	Further Reading	260
D	Special Topics	261
11	Multicasting	263
11.1	The Requirements of Multicasting	264
11.2	Multicasting Mechanisms at the Network Layer	265
11.3	Quality of Service Issues	267
11.3.1	Multicast Application Requirements	267
11.3.2	Network Mechanisms	268
11.4	Flow Control Mechanisms	269
11.5	The Economic Perspective	271
11.5.1	A Model for Allocating Multicast Bandwidth	271
11.5.2	The Problem of Sharing Common Cost	272
11.5.3	Formation of the Optimal Tree	275
11.5.4	Cost Sharing and Multicast Trees	275
11.6	Settlement	277
11.7	Further Reading	278
12	Interconnection	279
12.1	The Market Structure	279
12.1.1	Peering Agreements	279
12.1.2	Interconnection Mechanisms and Incentives	281
12.1.3	Interconnection Pricing	283
12.2	Competition and Service Differentiation	284
12.3	Incentives for Peering	285

12.4 Incentive Contract Issues	286
12.5 Modelling Moral Hazard	287
12.6 Further Reading	290
13 Regulation	291
13.1 Information Issues in Regulation	292
13.1.1 A Principal-agent Problem	292
13.1.2 An Adverse Selection Problem	296
13.2 Methods of Regulation	297
13.2.1 Rate of Return Regulation	297
13.2.2 Subsidy Mechanisms	297
13.2.3 Price Regulation Mechanisms	300
13.3 Regulation and Competition	301
13.4 Regulation in Practice	302
13.4.1 Regulation in the US	302
13.4.2 Current Trends	305
13.5 Further Reading	306
14 Auctions	309
14.1 Single Item Auctions	311
14.1.1 Take-it-or-leave-it Pricing	311
14.1.2 Types of Auction	312
14.1.3 Revenue Equivalence	313
14.1.4 Optimal Auctions	315
14.1.5 Risk Aversion	317
14.1.6 Collusion	318
14.1.7 The Winner's Curse	318
14.1.8 Other Issues	319
14.2 Multi-object Auctions	320
14.2.1 Multi-unit Auctions	320
14.2.2 Combinatorial Bidding	321
14.2.3 Double Auctions	322
14.2.4 The Simultaneous Ascending Auction	323
14.2.5 Some Issues for Multi-object Auctions	324
14.3 Auctioning a Bandwidth Pipeline	327
14.4 Further Reading	330
Appendix A Lagrangian Methods for Constrained Optimization	333
A.1 Regional and Functional Constraints	333
A.2 The Lagrangian Method	333
A.3 When Does the Method Work?	335
A.4 Shadow Prices	336
A.5 The Dual Problem	337
A.6 Further Reading	338

CONTENTS xiii

Appendix B Convergence of Tatonnement 339

 B.1 The Case of Producers and Consumers 339

 B.2 Consumers with Network Constraints 340

References 341

Index 353

Preface

This book is about pricing issues in modern communications networks. Recent technology advances, combined with the deregulation of the communication market and the proliferation of the Internet, have created a new and highly competitive environment for communication service providers. Both technology and economics play a major role in this new environment. As recent events in the marketplace make clear, the success of a communication services business is not guaranteed by new technology alone. An important part of any business plan for selling communications services is pricing and competition issues. These should be taken into account from the start. Traditionally, engineers have devised communication services without reference to how they should be priced. This is because communication services have been provided by large monopolies, with guaranteed incomes. The bundling and pricing aspects of individual services have been secondary. However, services are now sold in competitive markets and an important part of the service definition is how it should be priced. Technology can place severe restrictions on how this can be done. The following are some reasons why the pricing of communications services is now exciting to study:

1. Pricing affects the way services are used, and how resources are consumed. The value that customers obtain from services depends on congestion and on the way services are priced.
2. Communication service contracts provide for substantial flexibility. Pricing plays an important role as an incentive mechanism to control performance and increase stability.
3. Modern networking technology provides new possibilities for producers and the consumers to exchange economic signals on fast time scales. This allows for the creation of new flexible services that customers can control and by which they can better express their needs for quality. This was not possible until a few years ago, since previously services were statically defined and the network operator was in total in control.
4. There is no unique way to price. Issues such as 'flat' versus 'usage-based' charging have important effects on the short and long term network operation and its competitive position. These must be understood by people designing pricing policies.
5. Competition can be greatly influenced by the architecture of a networks and the ability of few players to control bottleneck resources in parts of the network, such as the access. New networks should be designed so that they provide an open competition environment in all parts of the supply chain for services. Competition and regulation issues are important in today's communication market.

6. Communication services are economic goods and must be priced accordingly. There are generic service models that capture aspects such as quality and performance and can be used to derive optimal prices in a services market. They can be used to propose tariffs with the desired incentive properties by pricing the appropriate service contract parameters.

We began this book after five years of research focused in pricing the rich family of ATM services and the newly emerging Internet. We believe there is a need for a book that can explain the provision of new services, the relation of pricing and resource allocation in networks, and the proliferation of the Internet and the debate on how to price it. We have had in mind as readers graduate students and faculty in departments of Electrical Engineering, Computer Science, Economics and Operation Research, telecoms engineers, researchers and engineers who work in research and industrial laboratories, and marketing staff in telecoms companies who need to understand better the technology issues and their relation to pricing. Our experience is that most of these people have only part of the background needed to follow such important subjects. Readers with engineering and OR background usually lack the economics background. Economists usually know little about communications technology and usually underestimate its importance. We have sought to write in a way that all readers will find stimulating. The book should interest anyone with some technology and mathematics background who wishes to understand the close relation of communication networks and economics. Of course, economists may skip the chapters on basic economics.

When we started this book, ATM technology was already declining in importance as an alternative to the Internet. However, there continues to be a practical demand for services such as ATM and Frame Relay. These can be put into the same generic model as the provision of WAN connectivity services. Similar concepts will apply in future extensions of Internet services that provide quality guarantees, such as differentiated services and integrated services. Consequently, we not only deal with the Internet, but also with effective bandwidths and statistical multiplexing.

The scope of this book is broad. It covers most of the concepts that are needed to understand the relation of economics and communications. We do not claim to provide a complete unifying framework, but explain many concepts that are generic to the problem of pricing. This is not a 'how to price' recipe book. Rather, it explores relevant subjects. It provides the basic models and terminology needed for a non-specialist reader to understand subtle topics where technology, information and economics meet. It explains the architecture of the communications market and provides a simple and intuitive introduction to network services at all levels, from the infrastructure to transport. We have tried to make the book technology independent, emphasizing generic service aspects and concepts.

The reader does not have to be an expert in communications or read several books on networking technology numbering hundreds of pages in order to understand these basic concepts. This may be of great benefit to a reader with an economics or operation research background. The same holds for readers with no economics background. We explain relevant microeconomic concepts in enough detail that the reader can follow many issues in network economics, without having to study advanced economic textbooks. However, we are not economists and do not claim to cover all topics in network economics. We hope that we do provide the reader with a useful summary of many key issues and definitions in basic economics. Those who wish to study these ideas in more depth can turn to economics textbooks. For instance, our section on game theory should remind those readers who have

previously studied it of those concepts from the subject that we use in other parts of the book. Readers who have not studied game theory before should find that the section provides a readable and concise overview of key concepts, but they will need to look elsewhere for details, proofs and further examples.

There is no one unifying model for network services. We provide models for several services and leave others of them out. These models allow network services to be priced similarly to traditional economic goods. These models can be used by network engineers as a framework to derive prices for complex transport services such as ATM, Frame Relay, IP VPNs, etc. We model the Internet and its transport services and discuss certain issues of fairness and resource allocation based on pricing for congestion. This provides a deeper understanding of the feedback aspects of the Internet technology, and of the recent proposals to provide for a richer set of bandwidth sharing mechanisms. We also provide the theoretical framework to price contracts in which parameters can be dynamically renegotiated by the users and the network. Finally, we give the reader a simple but thorough introduction to some current active research topics, such as pricing multicasting services, incentive issues in interconnection agreements between providers, and the theory of price regulation. For completeness, we also provide a simple introduction to auction mechanisms which are currently used to allocate scarce resources such as spectrum.

We hope to introduce non-specialists to concepts and problems that have only been accessible to specialists. These can provide both a practical guideline for pricing communication services and a stimulation for theoretical research. We do not review in extreme detail the existing literature, although we provide basic pointers. A guide to references appears at the end of each chapter. We seek to unify and simplify the existing state-of-the-art by focusing on the key concepts. We use mathematics to make the ideas rigorous, but we hope without being unnecessary detailed. About 80% of the results in the book have been published elsewhere and 20% are new. The level of the mathematics is at that of first year university student's knowledge of calculus and probability, and should be accessible to students and engineers in the field. Appendix A covers some important ideas of solving constrained optimization problems using Lagrange multipliers. The book has parts which are more technology specific and other parts that are more theoretical. Readers can take their pick.

We have found it convenient to divide the book in four parts. An overview of their contents can be found at the end of Chapter 1. Possible course that could be taught using this book are as follows:

1. An introductory course on pricing: Sections 1.4, 2.1, 3.2–3.3, 4.1–4.5, 4.10, 5.2–5.4.3, 5.4.7, 6.1–6.3, 7.3, 7.5, 8.1–8.4, 9.1–9.4, and Chapter 10.
2. An advanced course on mathematical modelling and pricing: Section 1.4, Chapter 2, Sections 3.1–3.3 and 3.5, Chapter 4, Sections 5.1–5.4, 5.6, 6.1–6.3, Chapters 8, 9 and 10.
3. A course on telecoms policy issues and regulation: Chapter 1, Sections 2.1, 3.2–3.6, Chapters 5 and 6, Sections 7.1–7.1.2, 7.3–7.5, Chapters 12 and 13, Sections 14.1–14.1.3, 14.2 and 14.3.
4. A course on game-theoretic aspects of pricing: Sections 5.1–5.4, 6.1, 6.4, 7.1–7.2, Chapters 9, 10, 11, Sections 12.4–12.5, 13.1, and Chapter 14.
5. An introductory network services and technology course: Sections 1.1–1.2, 2.1 and Chapter 3.

Acknowledgment

There are many people with whom we have enjoyed stimulating discussions while working on this book. These include especially Frank Kelly and Pravin Varaiya, who have done so much to inspire research work on pricing communications. They include also our partners in the Ca\$hman and M3i projects, and Panos Antoniadis, Gareth Birdsall, Bob Briscoe, John Crowcroft, Manos Dramitinos, Ioanna Constantiou, Richard Gibbens, Sandra Koen, Robin Mason, Georges Polyzos, Stelios Sartzetakis, Vassilis Siris, Georges Stamoulis and Jean Walrand.

List of Acronyms

ATM	Asynchronous Transfer Mode
ABR	Available Bit Rate
BGP	Border Gate Protocol
BSP	Backbone Service Provider
CAC	Connection Acceptance Control
CBR	Constant Bit Rate
CDVT	Cell Delay Variation Tolerance
CLP	Cell Loss Probability
CPNP	Calling Party Network Pays
CS	Consumer Surplus
DS	Differentiated Services
DWDM	Dense Wavelength Division Multiplexing
ECPR	Efficient Component Pricing Rule
ERP	Enterprise Resource Planning
FCFS	First Come First Serve
FDC	Fully Distributed Cost
IBP	Internet Backbone Provider
IGMP	Internet Group Management Protocol
IETF	Internet Engineering Task Force
ILEC	Incumbent Local Exchange Carrier
IS	Integrated Services
ISDN	Integrated Services Digital Network
ISP	Internet Service Provider
LAN	Local Area Network
LMDS	Local Multipoint Distribution Service
LRIC	Long Run Incremental Cost
MAN	Metropolitan Area Network
MC	Marginal Cost
MPEG	Moving Picture Experts Group
MPLS	MultiProtocol Label Switching
M-ECPR	Market determined Efficient Component Pricing Rule
PCR	Peak Cell Rate
NAP	Network Access Provider
PHB	Per Hop Behaviour
POP	Point of Presence
QoS	Quality of Service
RBOC	Regional Bell Operating Company

RFC	Request for Comments
RSVP	Resource Reservation Protocol
SCR	Sustainable Cell Rate
SDH	Synchronous Digital Hierarchy
SLA	Service Level Agreement
SMG	Statistical Multiplexing Gain
SONET	Synchronous Optical NETWORK
SW	Social Welfare
TELRIC	Total Element LRIC
TCP/IP	Transmission Control Protocol/Internet Protocol
TCA	Traffic Conditioning Agreement
UNE	Unbundled Network Element
UBR	Unspecified Bit Rate
UDP	User Datagram Protocol
VBR	Variable Bit Rate
VC	Variable Cost
VC	Virtual Circuit
VPN	Virtual Private Network
WAN	Wide Area Network
WWW	World Wide Web
XSP	Access Service Provider

Part A

Networks

1

Pricing and Communications Networks

This chapter describes current trends in the communications industry. It looks at factors that influence pricing decisions in this industry, and some differing and conflicting approaches to pricing. Section 1.1 is about the market for communications services. Section 1.2 is about present developments in the marketplace. Section 1.3 is about issues that pricing must address. Section 1.4 presents some introductory modelling.

1.1 The market for communications services

1.1.1 The Communications Revolution

We are in the midst of a revolution in communications services. Phenomenal advances in fibre optics and other network technology, enhanced by the flexible and imaginative software glue of the World Wide Web have given network users a technology platform that supports many useful and exciting new services. The usefulness of these services is magnified because of *network externality*. This is the notion that a network's value to its users increases with its size, since each of its users has access to more and more other users and services. This is one of the facts that spurs the drive towards worldwide network connectivity and today's Internet revolution — a revolution which is changing the way we engage in politics, social life and business. It is said that the electronic-economy, based as it is upon communications networks that provide businesses with new ways to access their customers, is destined to be much more than a simple sector of the economy. It will someday be *the economy*.

In a world that is so thoroughly changing because of the impact of communications services, the pricing of these services must play an important role. Of course a price must be charged for something if service providers are to recover their costs and remain in business. But this is only one of the many important reasons for pricing. To understand pricing's other roles we must consider what type of product are communications services and the characteristics of the industry in which they are sold.

1.1.2 Communications Services

The number of connections that can be made between n users of a network is $\frac{1}{2}n(n-1)$. This gives us *Metcalfe's Law* (named after the inventor of Ethernet), which says that the value of

a network increases as the square of the number of users. It relates to the idea of network externality and the fact that a larger network has a competitive advantage over a smaller one, because each of the larger network's users can communicate with a greater number of other users. It makes the growth of a large customer base especially important. With this in mind, a network operator must price services attractively. In this respect, communications services are like any economic good and fundamental ideas of the marketplace apply. One of these is that decreasing price increases demand. Indeed, it is common for providers to give away network access and simple versions of network goods for free, so as to stimulate demand for other goods, build their customer base and further magnify network externality effects.

The above remarks apply both to modern networks for data communication services and to the traditional telecommunications networks for voice services, in which the former have their roots. Throughout this book we use the term 'telecommunications' when referring specifically to telephony companies, services, etc., and use the broader and encompassing term 'communications' when referring both to telephony, data and Internet. It is interesting to compare the markets for these networks. For many years the telecommunications market has been supplied by large regulated and protected monopolies, who have provided users with the benefits of economy of scale, provision of universal service, consistency and compatibility of technology, stable service provision and guaranteed availability. Services have developed slowly; demand has been predictable and networks have been relatively easy to dimension. Prices have usually been based upon potential, rather than actual, competition. In comparison, the market for modern communications services is very competitive and is developing quite differently. However, the markets are alike in some ways. We have already mentioned that both types of network are sensitive to network externality effects.

The markets are also alike in that in that network topology restricts the population of customers to whom the operator can sell and network capacity limits the types and quantities of services he can offer. Both topology and capacity must be part of the operator's competitive strategy. It is helpful to think of a communications network as a factory which can produce various combinations of network services, subject to technological constraints on the quantities of these services that can be supported simultaneously. Severe congestion can take place if demand is uncontrolled. A central theme of this book is the role of pricing as a mechanism to regulate access to network resources and restrict congestion to an acceptable level.

Traditional telecoms and modern data communications are also alike in that, once a network of either type is built, the construction cost is largely a *fixed cost*, and the variable operating costs can be extremely small. If there is no congestion, the *marginal cost* of providing a unit of communications service can be almost zero. It is a rule of the marketplace that competition drives prices towards marginal cost. Thus, a danger for the communications industry is that the prices at which it can sell communications services may be driven close to zero.

In summary, we have above made three elementary points about pricing: lowering price increases demand; pricing can be used to control congestion; competition can drive prices to marginal cost.

1.1.3 Information Goods

It is interesting to compare communications services with *information goods*, such as CDs, videos or software. These share with communications services the characteristic of being costly to produce but cheap to reproduce. The first copy of a software product bears all the production cost. It is a sunk cost, mainly of labour. Many further copies can be produced

at almost no marginal cost, and if the software can be distributed on the Internet then its potential market is the whole Internet and its distribution cost is practically zero. Similarly, once a network is built, it costs little to provide a network service, at least while there is no congestion. This also shows that information goods and network services can sometimes be viewed as public goods, like highways. Assuming that the installed network capacity is very large (which is nearly true given today's fibre overprovisioning), the same information good or network service can be consumed by an arbitrary number of customers, increasing its value to its users (due to externalities) and the value to society. This is in contrast to traditional goods like oranges and power; a given orange or kilowatt-hour can be consumed by a single customer and there is a cost for producing each such additional unit.

The similarity cannot be pushed too far. We must not forget that a network has a continuing running cost that is additional to the one-time cost of installation. This includes network management operations, amongst which accounting and billing are particularly costly. The cost of selling a single copy of a piece of software is small compared to the cost of maintaining, monitoring and billing a network service. It is not surprising that cost, among many other economic factors, influences the evolution of networking technology. One reason for the acceptance of Internet technology and the Internet Protocol (IP) is that there it is less costly to manage a network that is based on a single unifying technology, than one that uses layers of many different technologies.

There are some lessons to be learned from the fact that information goods can sell at both low and high prices. Consider, for example, the fact that there are hundreds of newspaper web sites, where entertaining or useful information can be read for free. It seems that publishers cannot easily charge readers, because there are many nearly equivalent sites. We say that the product is 'commoditized'. They may find it more profitable to concentrate on differentiating their sites by quality of readership and use this in selling advertising. In contrast, a copy of a specialist software package like AutoCad can sell for thousand of dollars. The difference is that its customer base is committed and would have difficulty changing to a competing product because the learning curve for this type of software is very steep. Similarly, Microsoft Word commands a good price because of a network externality effect: the number of people who can exchange documents in Word increases as the square of the number who use it. These examples demonstrate another important rule of the marketplace: if a good is not a commodity, and especially if it has committed customers, then it can sell at a price that reflects its value to customers rather than its production cost.

We have noted that both traditional telecoms and modern communications services are sensitive to network topology and congestion. This is not so for an information good. The performance of a piece of software running on a personal computer is not decreased simply because it is installed on other computers; indeed, as the example of Microsoft Word shows, there may be added value if many computers install the same software.

1.1.4 Special Features of the Communications Market

One special feature of the market for communications services, that has no analogy in the market for information goods (and only a little in the market for telecommunications), is that in their most basic form all data transport services are simply means of transporting data bits at a given quality level. That quality level can be expressed such terms as the probability of faithful transmission, delay and jitter. A user can buy a service that the operator intended for one purpose and then use it for another purpose, provided the quality

level is adequate. Or a user can buy a service, create from it two services, and thereby pay less than he would if he purchased them separately. We say more about the impact of such *substitutability*, *arbitrage* and *splitting* upon the relative pricing of services in Section 8.3.5.

Another thing that makes communication transport services special is their reliance on *statistical multiplexing*. This allows an operator to take advantage of the fact that data traffic is often bursty and sporadic, and so that he can indulge in some amount of overbooking. He need not reserve for each customer a bandwidth equal to that customer's maximum sending rate. Statistical multiplexing produces economy of scale effects: the larger the size of the network, the more overbooking that can take place, and thus the size of the customer base that can be supported increases more than proportionally to the raw quantity of network resources. It is intuitive that a network service that is easier to multiplex should incur a lesser charge than one which is more difficult to multiplex. There are many multiplexing technologies and each is optimized for a particular type of data traffic. For instance, SONET (Synchronous Optical NETwork) is a multiplexing technology that is optimized for voice traffic (which is predictable and smooth), whereas the Internet technology is optimized for data traffic (which is stochastic and bursty).

Simple economic goods are often specified by a single parameter, such as number of copies, weight, or length of a lease. In contrast, contracts for data communications services are specified by many parameters, such as peak rate, maximum throughput and information loss rate. Contracts for services that support multimedia applications are specified by additional parameters, such as ability to sustain bursty activity, and ability and responsibility to react to changing network conditions. Since service contracts can be specified in terms of so many parameters, their potential number is huge. This complicates pricing. How are we to price services in a consistent and economically rational way? Moreover, contracts are more than simple pricing agreements. For example, a contract might give a user the incentive to smooth his traffic. Customers also benefit because the quality of the service can be better and lower priced. This poses questions of how we can reasonably quantify a customer's network usage and price contracts in a way that makes pricing a mechanism for controlling usage.

1.2 Developments in the marketplace

In the next two sections, we look at some important factors that affect the present market for communications services. We make some further arguments in favour of the importance of pricing. We describe the context in which pricing decisions occur, their complexity and consequences. Some of these issues are subject to debate, and will make most sense to readers who are familiar with present trends in the Internet. Some readers may wish to skip the present section on first reading.

There have been two major developments in the marketplace for telecoms services: the development of cost-effective optical network technologies, allowing many light beams to be packed in a single fibre; and the widespread acceptance of the Internet protocols as the common technology for transporting any kind of digitized information. Simultaneously, the Internet bubble of late 1990s has seen an overestimation of future demand for bandwidth and overinvestment in fibre infrastructure. Together, these factors have created a new technology of such very low cost that it threatens to disrupt completely the market of the traditional telephone network operators, whose transport technologies are optimized for voice rather than data. It has also commoditized the market for transport services to such an extent that companies in that business may not be able to recover costs and effectively compete.

One reason for this is that the Internet is a ‘stupid’ network, which is optimized for the simple task of moving bits at a single quality level, irrespective of the application or service that generates them. This makes the network simple and cheap. Indeed, the Internet is optimized to be as efficient as possible and to obey the ‘end-to-end principle’. To understand this principle, consider the function: ‘recovery from information loss’. This means something different for file transfer and Internet radio. The end-to-end principle says that if such a function is invoked rarely, and is not common to all data traffic, then it is better to install it at the edge of the network, rather than in each link of the network separately. Complexity and service differentiation is pushed to the edges of the network. The reduction in redundancy results in a simpler network core. Customer devices at the edges of the network must provide whatever extra functionality is needed to support the quality requirements of a given application.

The fact that the Internet is stupid is one of the major reasons for its success. However, it also means that a provider of Internet backbone services (the ‘long-haul’ part of the network, national and international) is in a weak bargaining position if he tries to claim any substantial share of what a customer is prepared to pay for an end-to-end transport service, of which the long-haul service is only a part. That service has been commoditized, and so in a competitive market will be offered at near cost. However, as noted previously, the cost of building the network is a sunk cost. There is only a very small variable cost to offering services over an existing network infrastructure. The market prices for network services will be almost zero, thus making it very difficult for the companies that have invested in the new technologies to recover their investments and pay their debts. As some have said, *the best network is the hardest one to make money running* (Isenberg and Weinberger, 2001).

This ‘paradox of the best network’ does not surprise economists. As we have already noted, there is little profit to be made in selling a commodity. The telephone network is quite different. Customers use only simple edge devices (telephones). All value-added services are provided by the network. Network services are constructed within the network, rather than at the edges, and so operators can make money by being in control. Similarly, video and television distribution use service-specific networks and make good profits. Telephone networks are optimized for voice and not for data. Voice streams are predictable in their rates, while data is inherently bursty. Due to the overspecified requirements (for reliability and voice quality), the technologies for voice networks (SONET and SDH) are an order of magnitude more expensive than the technology for providing simple bit moving services of comparable bandwidth, as provided by the Internet using the new optical transmission technologies. The extra quality per bit offered by telephone network infrastructures does not justify their substantially greater costs. Moreover, the large network capacity available may let the quality of the bits provided by the new Internet technology networks approach that provided by the telephone network. Unfortunately, these voice-centred technologies are not so old as to be easily written-off. Existing operators invested heavily in them during the late 1980s and mid 1990s, encouraged by regulators who allowed them a ‘return on assets’, that is, a profit proportional to the assets under their control. This makes it hard for operators to abandon their voice-centred infrastructures and build new networks from scratch.

The above arguments suggest that network operators deploying the new Internet over fibre technologies should be able to carry voice at substantially less cost than traditional network operators, and so drive them out of business. They will also be able to offer a rich set of high bandwidth data services, which are again cheaper for them to provide.

However, things are not entirely rosy for these new network operators. They have their own problem: namely, a bandwidth glut. During the Internet bubble of the late

1990s investors overestimated the growth in the demand for data services. They believed there would be an unlimited demand for bandwidth. Many companies invested heavily in building new fibre infrastructures, at both the metropolitan and backbone level. DWDM (Dense Wavelength Division Multiplexing) made it possible to transport and sell up to 80 multiple light waves (using present technology) on a single strand of fibre. Gigabit Ethernet technologies combined with the Internet protocols allowed connectivity services to be provided very inexpensively over these fibre infrastructures. Using present technologies each light wave can carry up to 10 Gbps of information, so that a single fibre can carry 800 Gbps. Although DWDM is presently uneconomic in the metropolitan area, it makes sense in the long-haul part of the network. It has been estimated that there are now over a million route-miles of fibre installed worldwide, of which only about 5% is lit, and that to only about 8% of the capacity of the attached DWDM equipment. Thus there is potential for vastly more bandwidth than is needed. Some experts believe that fibre is overprovisioned by a factor of ten in the long-haul part of the networks. Further bad news is that demand for data traffic appears to be increasing by only 50% per year, rather than doubling as some had expected.

The result is that the long-haul bandwidth market has become a commodity market, in which demand is an order of magnitude less than expected. A possible reason is miscalculation of the importance of complementary services. High-capacity backbones have been built without thinking of how such 'bandwidth freeways' will be filled. The business plans of the operators did not include the 'bandwidth ramps' needed, i.e. the high-bandwidth access part that connects customers to the networks. The absence of such low priced high-bandwidth network access services kept backbone traffic from growing as predicted. Besides that, transport services have improved to such an extent that technology innovation is no longer enough of a differentiating factor to provide competitive advantage. Prices for bandwidth are so low that it is now very hard for new network operators to be profitable, to repay the money borrowed for installing the expensive fibre infrastructure, or to buy expensive spectrum licenses.

Existing operators of voice-optimized networks are also affected. Their income from highly priced voice calls has reduced, as voice customers have migrated to the Internet technology of voice-over-IP networks, while the demand for voice remains essentially constant. They have not seen a compensating increase in demand for data services, which in any case are priced extremely low because of competition in that commoditized market. Some local service providers are even selling data services at below cost because of their expensive legacy network technology, while simultaneously installing the new IP over fibre technology in parts of their networks to reduce their costs. Of course infrastructure is not the only cost of providing traditional access and voice services. A larger part of the cost is for orders, repairs, customer service and support. This cost will always be reflected in customers' bills. Thus local operators, who have traditionally been in a monopoly position, do live in a somewhat protected environment because they have a steady income from their large and loyal base of telephone customers. Competition is fiercest in the long-haul part of the network, where new technologies can be easily deployed, economies of scale are great, and many operators compete.

It may seem paradoxical to have such severe sustainability problems in a growth industry such as telecommunications. Although the pie is growing, the business models seem to have some serious flaws. This is due to miscalculations, and because companies have tried to become simultaneously both retail and wholesale service providers, with the result that they have been competing with their own customers. Some experts envisage extreme scenarios. In one such scenario, the regulator acquires and controls the complete fibre infrastructure in