

Lecture Notes in Electrical Engineering

Volume 92

For further volume
<http://www.springer.com/series/7818>

Antonio Liotta · George Exarchakos

Networks for Pervasive Services

Six ways to upgrade the internet

Antonio Liotta
Department of Electrical Engineering
Eindhoven University of Technology
P.O. Box 513
5600 MB Eindhoven
The Netherlands
e-mail: a.liotta@tue.nl

George Exarchakos
Department of Electrical Engineering
Eindhoven University of Technology
P.O. Box 513
5600 MB Eindhoven
The Netherlands
e-mail: g.exarchakos@tue.nl

ISSN 1876-1100

e-ISSN 1876-1119

ISBN 978-94-007-1472-4

e-ISBN 978-94-007-1473-1

DOI 10.1007/978-94-007-1473-1

Springer Dordrecht Heidelberg London New York

© Springer Science+Business Media B.V. 2011

No part of this work may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, microfilming, recording or otherwise, without written permission from the Publisher, with the exception of any material supplied specifically for the purpose of being entered and executed on a computer system, for exclusive use by the purchaser of the work.

Cover design: eStudio Calamar, Berlin/Figueres

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

To Maria, Dikeos and Marina

Foreword

When I began working for Peter Kirstein's group at the University College London (UCL) in 1980 the department had already been part of the embryo Internet for seven years. The first "local area networks" were being deployed, and it was becoming clear that the future would consist of many networks, using a variety of technologies, all of which would need to interwork. Accordingly the Internet community adopted an architecture (designed by Robert Kahn and Vint Cerf) in which all networks must implement a common "internet protocol" (IP) to carry packets of data across and between networks. *Networks* were connected to their neighbours by computers termed "gateways" (we now call them "routers"). IP did not attempt to correct any errors that might arise; that was left to the "hosts"—the computers attached to the networks that were the sources and sinks of data. Hosts implemented a protocol called the "transmission control protocol" (TCP), which arranged for the re-transmission of any packets that did not arrive intact.

Back in 1980 routers were based on refrigerator-sized "mini-computers" which cost tens of thousands of pounds (hundreds of thousands of euros in today's terms). Connecting a computer to a LAN cost more than £1,000. However, by the time Antonio Liotta joined us at UCL in 1995 things were very different; costs had plummeted, personal computers were widespread and the Internet now comprised thousands of networks and millions of hosts. The applications that generated the bulk of the traffic though—file transfer, email and the burgeoning world wide web—still matched the requirements that had inspired the development of the TCP all those years before. Underlying TCP is the assumption of a client–server model; the server computer has something the client wants and the TCP delivers it complete and error free with high probability. TCP achieves this by trading timeliness for reliability. That is fine for applications like email—no one cares much if an email message is delayed for a few seconds provided it arrives intact. However, it was clear to Antonio and others researching in the late 1990s that new applications were on the horizon, many of which would not fit the TCP client–server model at all well. Some, such as streaming audio and video, would not tolerate TCP-induced delays. Others were abandoning the asymmetric client–server model in favour of a more egalitarian "peer-to-peer" approach typified by

file-sharing applications such as Napster. Yet another development, exemplified by Antonio's own research, turned the client–server model on its head by moving the servers ("agents") around the network in order to complete a task in the most efficient way.

Today those anticipated developments have arrived with a vengeance. Traditional TCP-based applications now form just a small minority of Internet traffic, perhaps no more than 15%, and most of that is world wide web. Streaming applications comprise around 10% and much of the rest is peer-to-peer file sharing—mostly of videos. (Things change so rapidly that in 12 months time these estimates will all likely be wrong!) Just as the applications have changed, so have the devices on which they run. The number of mobile phones in the world already vastly exceeds the number of networked computers and, increasingly, these phones themselves are on the Internet. Not only will there be hundreds of millions of them but they will move about! Mobility brings its own set of problems: wireless connections are subject to rapid changes in transmission speeds and error rates; an IP address is no longer a reliable clue as to where in the world a host is located.

Plainly the Internet has already adapted somewhat to support today's applications and host mobility. However, the adaptations have often been piece-meal, and stresses and strains sometimes appear. Researchers today must not only look at how better to adapt the Internet to today's applications but must also anticipate the huge changes that are, inevitably, around the corner. The authors of this book have, between them, accumulated many years researching novel techniques for optimising novel technologies within the Internet. They are well placed to understand the problems that must be solved and what solutions might be feasible. They begin by describing the key features of the Internet as it has evolved and the problems that must be addressed if it is to become flexible enough to support today's applications and mobility. They then look at what further adaptations may be needed within the Internet of the future. They do not make the mistake of claiming to know precisely what will be needed. Rather, they have used their knowledge to identify, as their subtitle makes clear, *Six ways to upgrade the Internet*. They explain these upgrades with the aid of carefully chosen examples and illustrations. The result is a book that will be of great benefit to students who wish to progress from an understanding of what the Internet is now towards an understanding of the motivations and techniques that will drive its future.

London, January 2011

Graham Knight
Department of Computer Science
University College London

Preface

Through the eyes of billions of Internet users, we have learned how the ease of communication can ignite phenomenal innovation. It is fascinating to witness the new habits and social phenomena created by the Web. However, what happens behind the scenes of our digital ecosystem? It is the network that moves our data around, handles the peak-hour traffic and strives to smoothly deliver the audio-video streams. Networks play a vital role in sustaining the unrelenting evolution of the most demanding Web systems.

Networks have to keep up with unprecedented data volumes while adapting to new communication patterns or, rather, new kinds of traffic. Most applications are now *pervasive*. We expect them to be accessible everywhere, without compromise. We expect the same “look and feel,” and the same quality and functionality, irrespective of any other technological constraints. Hence, many fear that the emerging breed of *pervasive applications* will soon render the Internet obsolete. As a matter of fact, a worldwide effort to reinvent the Internet is well underway by the “Future Internet” research community.

Through our active involvement in the investigation and teaching of network protocols, we have come to realize how difficult it is to grasp networking concepts that exceed the horizon of TCP/IP (i.e., the Internet protocol). When it comes to *advanced network protocols*, specialist literature abounds with creative proposals. Yet, very few protocols manage to step out of the laboratory and into the commercial world.

Perhaps our most ambitious task in writing this book was to extract a selection of remarkable ideas from the scientific literature and make them accessible to the non-specialist reader. Our book does not have the objective of embracing the *Future Internet*, though it does introduce a series of *network mechanisms* that will certainly find a place in the next-generation network. We propose six ways to upgrade the Internet and make it more *ubiquitous, reactive, proactive, information-driven, distribution-efficient* and *searchable*. In the final chapter, we offer some considerations about the Future Internet, though we have resisted the temptation to give any specific technical solutions.

This book is self-contained and is meant for anybody with an interest in the *post-Internet* era. We use the book to teach *ad hoc networks* and *P2P networks* in our *Communicating Systems* course at the Technical University Eindhoven (The Netherlands). You do not need to have a background in *computer networks* because all necessary concepts are summarized in the first two chapters. We have had to face the challenge of teaching *networking* to students who are not keen mathematicians: our efforts are reflected in this book which does not contain equations or mathematical formulations, but is enriched by examples and illustrations.

Yet, this is not another book for “dummies.” Whoever has taken a classic course in *computer networks* will find our book to be a useful tool for gaining a deeper understanding of more *advanced network mechanisms*.

We hope that scholars in the field will find inspirational ideas within these pages for their research.

January 2011

Antonio Liotta
George Exarchakos

Acknowledgments

This book has been in our mind for several years, but it started as a concrete project only after we moved to the *Technical University of Eindhoven* (The Netherlands). At TU/e we found the inspirational energy that sustained our efforts. Several improvements came after discussions with our colleagues at *Electrical Engineering and Computer Science* and thanks to the interactions with our *Communicating Systems* students.

Springer was instrumental in bringing the book to the light of the day. We were lucky to work with an enthusiastic publishing editor who not only championed the book but also taught us a lot of things about writing for non-specialist readers. Many thanks to Rachel Hopkins who patiently copyedited the book in astounding detail.

We are particularly grateful to Lisandro Granville for injecting great ideas into the book. Graham Knight and Raouf Boutaba were the first to provide scientific feedback on the manuscript. Alessandro Liotta helped making [Chaps. 1, 3 and 10](#) readable to the non-technologist.

On a more personal note, we are immensely grateful to our respective partners for supporting us throughout the writing process.

Contents

1	On the Way to the Pervasive Web	1
1.1	The Net, a Tool for Everyone	1
1.2	The Inexorable Transformation of Internet Applications	3
1.3	The Application's Mutiny	5
1.4	Everything on the Move	9
1.5	New Interaction Paradigms Emerge	10
1.6	The Scent of Pervasive Applications	12
1.7	The Billion Dollar Question	13
	References	14
2	The Network, As We Know It	15
2.1	The Multiple Facets of Networks	15
2.2	Networks from the Eyes of an Ordinary User	16
2.3	Invite a Programmer to Understand What's in the Cloud	18
2.4	A Network Engineer to Turn a Switch into a Router	20
2.5	The Computer Science of a Router	23
2.6	Simple Math to Stabilize the Net	27
2.7	Life of a Commuter	33
2.8	The Three Fundamental Principles	35
	References	38
3	Six Problems for the Service Provider	39
3.1	The Net has Ossified	39
3.2	Problem 1: Not Truly Ubiquitous	42
3.3	Problem 2: The Unresponsive Net	44
3.4	Problem 3: Too Much, Too Stale Signaling	44
3.5	Problem 4: Lack of Parallelism	46
3.6	Problem 5: Data Agnosticism	48
3.7	Problem 6: Inadequate Net-Search Engine	49
3.8	Concluding Remarks	50
	References	50

4 Spontaneous Networks	51
4.1 The Gift of Ubiquity.	51
4.2 Spontaneous Connectivity	53
4.3 The Hidden-Terminal Problem.	54
4.4 The Exposed-Terminal Problem.	55
4.5 Preventive Measures to Avoid Collision	55
4.6 Path Discovery in a Volatile Networks	58
4.7 The KISS Approach	59
References	62
5 Reactive Networks	65
5.1 Why Networks on Demand?	65
5.2 A Traffic-Free Network	66
5.3 Our First Path	66
5.4 Path Management.	69
5.5 Our Second Path	73
5.6 Global Synchronization.	73
5.7 Error Management	75
5.8 Remarks on Reactive Networks	77
References	77
6 Proactive Networks	79
6.1 From Reactive to Responsive	79
6.2 Keep the Network Ready	80
6.3 How do I Find My Multipoint Relay?	81
6.4 Life of an OLSR Node	82
6.5 The Node's Information Repository	84
6.6 Shortest Path over the MPR Sub-topology.	84
6.7 A Complete Example	85
6.8 How Proactive Can You Be?.	87
6.9 The Power of Hybrid Protocols	90
References	93
7 Content-Aware Networks	95
7.1 Routers Should Read the Content.	95
7.2 A Network on Top of the Physical Network	96
7.3 Centralized Assignment of Node Identifiers.	99
7.4 Centralized Entry Point Discovery	99
7.5 Multiple Bootstrap Servers	102
7.6 Decentralized Assignment of Node Identifiers	104
7.7 Entry Point Discovery via Underlying Links	104
7.8 Content is an Asset at the Edges	107
References	108

8	Distribution-Efficient Networks	111
8.1	Publishing Goes Beyond Bootstrapping.	111
8.2	The Two Flavors of Virtual Networking	112
8.3	Creating Unstructured Neighborhoods.	113
8.4	Making Yourself Known in Unstructured Neighborhoods	116
8.5	Unstructured Resource Publishing	117
8.6	Secure a Role in Structure Worlds	121
8.7	Build Strict Formations.	122
8.8	Place Links and Resources into a Structured Ring	126
8.9	Data-Awareness via Protocol-Agnosticism.	129
	References	130
9	Discovering Virtual Resources	133
9.1	Four Ways to Reach a Resource	133
9.2	Assessment of Discovery Mechanisms	134
9.3	Containing the Proliferation of Discovery Messages	134
9.4	Blind Discovery for Unstructured Networks	135
9.5	Informed Discovery in Unstructured Networks	138
9.6	Discovery in Loosely-Structured Networks	139
9.7	Deterministic Discovery in Structured Networks	142
	References	144
10	A Peek at the Future Internet.	145
10.1	The Fourth Networking Principle: Beyond Mere Connectivity	145
10.2	Internet of Things: Sense and Influence Your Environment	146
10.3	Small, Large Networks	147
10.4	Manage the Autonomics	150
10.5	Dependable Networks	150
10.6	The Fine Line Between Freedom, Security and Privacy	151
10.7	Energy-Efficient Networks	152
10.8	No Matter What, the Network will Remain Generative.	153
	References	154
Index	157	

Acronyms

AODV	Ad-hoc on-demand distance vector
APS	Adaptive probabilistic search
BFS	Breadth-first search
CIDR	Classless inter-domain routing
CPU	Control processing unit
CS	Client–server
CTS	Clear-to-send
DFS	Depth-first search
DHCP	Dynamic host configuration protocol
DHT	Distributed hash table
DiffServ	Differentiated services
DNS	Domain name system
DSL	Digital subscriber line
DV	Distance vector
FTTH	Fiber-to-the-home
HTL	Hops-to-live
IARP	Intrazone routing protocol
ICT	Information and communication technologies
ID	Identifier
IDS	Iterative deepening search
IERP	Interzone routing protocol
IntServ	Integrated services
IP	Internet protocol
IPTV	Internet protocol television
IS	Intelligent search
ISP	Internet service provider
KISS	Keep it simple and stupid
LAN	Local area network
LS	Link state routing
LSP	Link state packet
MACA	Multiple access with collision avoidance

MANET	Multiple ad hoc network
mBFS	Modified breadth-first search
MPR	Multipoint relays
NDP	Neighbor discovery protocol
NL	Neighbor list
NO	Network operator
OLSR	Optimized link state routing
P2P	Peer-to-peer
PC	Personal computer
PKI	Public key infrastructure
QoE	Quality of experience
QoS	Quality of service
RERR	Route error
RFIDs	Radio-frequency identification
RREP	Route reply
RREQ	Route request
RTS	Request-to-send
RW	Random walks
TC	Topology control
TCP	Transmission control protocol
TIB	Topology information base
TTL	Time-to-live
TV	Television
VoD	Video on-demand
VoIP	Voice over IP