PREMIER REFERENCE SOURCE

Streaming Media with Peer-to-Peer Networks

Wireless Perspectives

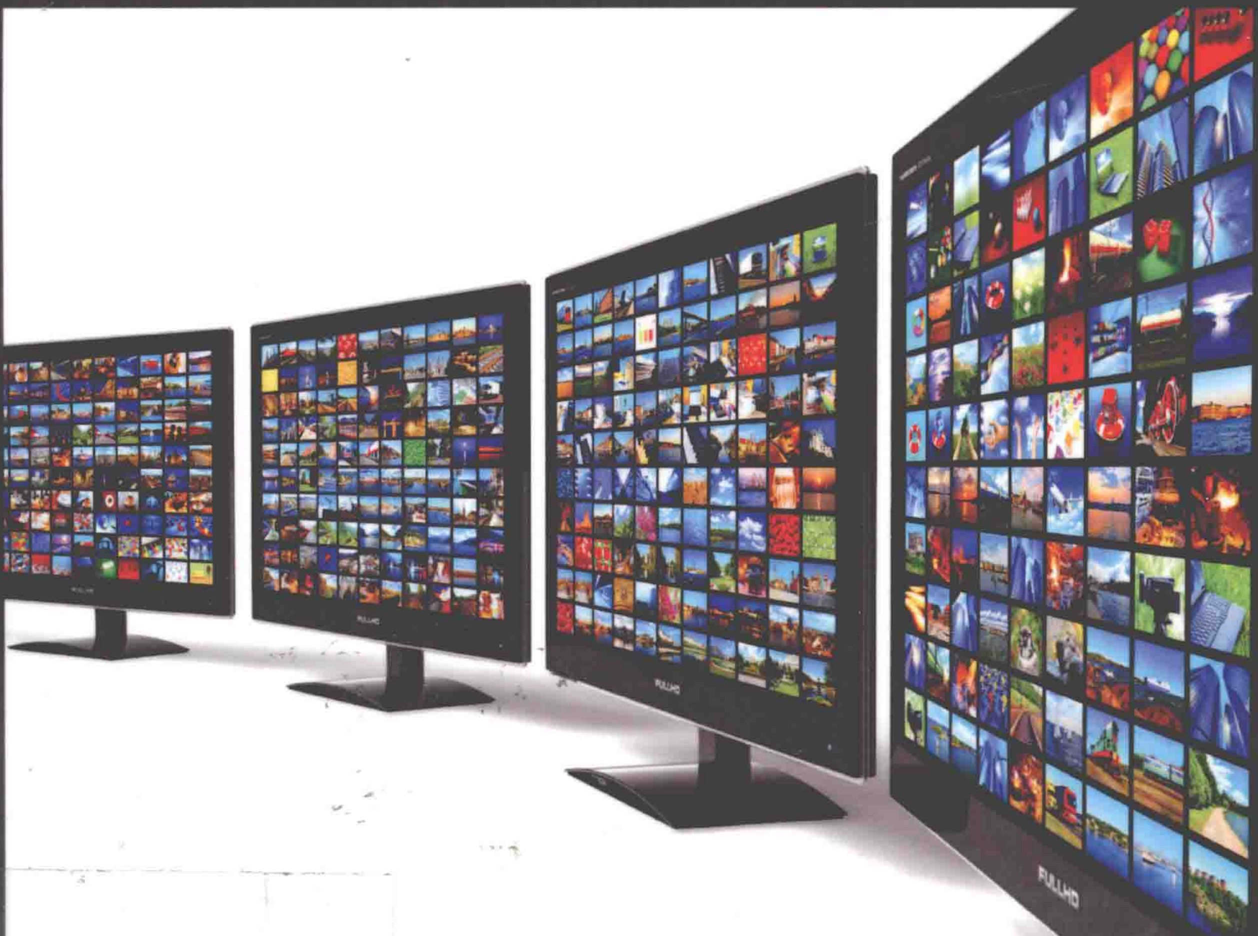

Martin Fleury & Nadia Qadri

Streaming Media with Peer-to-Peer Networks:

Wireless Perspectives

Martin Fleury
University of Essex, UK

Nadia Qadri
COMSATS Institute of Information Technology, Pakistan

Information Science
REFERENCE

Managing Director:	Lindsay Johnston
Senior Editorial Director:	Heather A. Probst
Book Production Manager:	Sean Woznicki
Development Manager:	Joel Gamon
Acquisitions Editor:	Erika Gallagher
Typesetter:	Lisandro Gonzalez
Cover Design:	Nick Newcomer

Published in the United States of America by
Information Science Reference (an imprint of IGI Global)
701 E. Chocolate Avenue
Hershey PA 17033
Tel: 717-533-8845
Fax: 717-533-8661
E-mail: cust@igi-global.com
Web site: http://www.igi-global.com

Library of Congress Cataloging-in-Publication Data

Streaming media with peer-to-peer networks: wireless perspectives / Martin Fleury and Nadia Qadri, editors.
p. cm.
Includes bibliographical references and index.
Summary: "This book offers insights into current and future communication technologies for a converged Internet that promises soon to be dominated by multimedia applications, at least in terms of bandwidth consumption"--Provided by publisher.
ISBN 978-1-4666-1613-4 (hardcover) -- ISBN 978-1-4666-1614-1 (ebook) -- ISBN 978-1-4666-1615-8 (print & perpetual access) 1. Streaming technology (Telecommunications) 2. Internet. I. Fleury, Martin, 1951- II. Qadri, Nadia, 1980-
TK5105.386.S738 2012
004.67'8--dc23

2012002468

British Cataloguing in Publication Data
A Cataloguing in Publication record for this book is available from the British Library.

All work contributed to this book is new, previously-unpublished material. The views expressed in this book are those of the authors, but not necessarily of the publisher.

Editorial Advisory Board

List of Reviewers

Preface

This book presents a set of research-based investigations of a number of paradigms, which work together in various ways. From the book title we have: "peer-to-peer network," multimedia "streaming," and "wireless perspectives," in other words, mobile applications of peer-to-peer streaming. It turns out that there is or there is suspected to be a synergy between one type of wireless network and peer-to-peer networks. This is possible because peer-to-peer networks are overlay networks that are superimposed upon an existing physical network. The issue of the compatibility of the two paradigms is explored in Chapter 5. The wireless networks in question are ad hoc or infrastructureless, meaning that they can exist without the presence of a centralized access point to redirect traffic between access points over a wired infrastructure network. As this is a book about paradigms, it seemed fitting to include a final chapter which considered information-centric networks. These, too, are overlays, and may present the next step in locating multimedia streaming sources without the use of router-based addressing at the network layer of the protocol stack. In fact, service oriented architectures, XML routers, deep packet inspection, and content delivery networks (see later) together with peer-to-peer overlay technologies are a clear effort from the top to move towards a networking which is focused on information rather than router-based addressing. However, before considering the contents of the book's chapters further, it is worth asking the question "Where do these paradigms stand in the commercial world at the time of writing?"

Streaming is generally used for pre-encoded video streaming in order to avoid hefty storage requirements, particularly on mobile battery-powered devices. This requirement will grow as the era of high-definition laptops is entered. Streaming can also help protect content from copying, as there is no stored copy at the end-point. For live encoded video and two-way interactive video (such as video phone and video conferencing), streaming is certainly necessary. Chapter 4 has more on advances in adaptive streaming. Streaming seems to have been introduced commercially by Real Networks with a pioneering audio broadcast of a baseball event in 1995. The company went on to introduce video streaming in 1997 so that by 2000, about 85% of streaming content in the Internet was apparently in the Real format. However, the Real Player, which now streams using a proprietary codec rather than originally an H.263 standard codec, appears to have suffered a market decline, at least in respect to video streaming. One cause of this may have been the obtrusiveness of Real Player upgrades and auxiliary programs, or it could be the cost of the servers, or indeed competition with Adobe Flash Player (except on Apple devices). It is worth noting that though Real did prevent the writing of their streams to store, there now exists software to do this (de-streaming), thus allowing viewing times to be shifted. There are also filters to allow different players, other than the Real Player, to be used.

Commercial streaming tends to employ simulcast for pre-encoded content, in which different versions of the same pre-encoded stream are stored on a server, about five versions on the Akamai *content delivery network (CDN)*. Of course, this involves considerable storage and organization and management of that storage. The pixel/frame resolution is often varied resulting in different bitrates to suit the bandwidth bottleneck that commonly exists over an *Asymmetric Digital Subscriber Connection (ADSL* —marketed as "broadband"). Commercial streaming also tends to use pseudo-streaming, which is also known as progressive download. Pre-encoded bitstreams are broken into chunks, each of which is downloaded, while a previous chunk is played. Because browser plug-in compatibility is seen as important, connection-oriented *Transmission Control Protocol (TCP)* is used as the transport protocol for *HyperText Transfer Protocol (HTTP)*, or in Adobe's case as the basis of *Real Time Messaging Protocol (RTMP)*, which in turn may be encapsulated in HTTP. (RTMP is a proprietary protocol, which reportedly appears in one variant to use an anonymous Diffie-Hellman mechanism for its authentication, making content vulnerable to snooping.) The widespread adoption of the Adobe Flash Player (present on about 95% of browsers in July 2011) acts as a non-technical constraint on the form of streaming. Microsoft Silverlight has about 53% penetration of browsers. Both of these players use proprietary software.

It is sometimes said that TCP is necessary to penetrate firewalls and *network address translation (NAT)*. However, it is the negotiation of a connection that is needed for this purpose and a protocol such as *Datagram Congestion Control Protocol (DCCP)* can set up a connection before reverting to *User Datagram Protocol (UDP)* transport. In fact, Skype is thought to use the STUN protocol for NAT and firewall traversal by UDP streams. As is often observed, TCP, because of its built-in reliability, is unsuitable for a real-time service, as congestion can lead to frequent resends of the data and a slowing down of the transmission rate. The problem is worse in a wireless network as lost packets may be due to channel conditions, in which case, staunching the flow will not make the adverse channel conditions go away. Buffers can smooth out the delays. However, large buffers militate against *click-and-view* type streaming, as they lead to long start-up delays. This is only partially mitigated by placing the content in caches nearer the user. Apart from this, pseudo-streaming is unsuitable for the streaming of two-way video. However, it should be remarked that *peer-to-peer (P2P)* streaming also uses chunks rather than bitstreams, though the difficulty of lost chunks is now solved by means of having multiple sources for the same chunk. Chapter 8 considers the chunk method of streaming for wireless mobile ad hoc networks and their vehicle-bound related networks (VANETs).

In recognition of these problems with pseudo-streaming, commercial operators have introduced adaptive pseudo-download (one variant is known as HTTP Live Streaming) in which the target download rate is not fixed at start-up time but can be adapted according to measured or estimated conditions on the network. Chunk sizes tend to be between 2 and 15 s(econds) in duration. Unfortunately, this is unlikely to be suitable for wireless networks, as the channel can change over a time interval much smaller than a chunk's download time. On the other hand, if a chunk's size is reduced then compression coding gain is reduced as well. The complexity of managing selection of chunks in the context of stateless HTTP is also a considerable burden. Instead, there are a number of alternative ways to implement adaptation. If simulcast is to be used then an H.264/*AVC (Advanced Video Coding)* codec supports switching frames, which can remove the overhead of spatially coded I-frames as switching points and provide smoother transitions at those switching points. Bitrate transcoder banks can also serve to dynamically adapt to changes in congestion. However, transcoders represent a hardware investment that can be removed by scalable codecs. Such codecs as the *Scalable Video Coding (SVC)* extension to H.264 are said to have almost the same coding efficiency as H.264/AVC. The research perspective for scalable streaming is considered in Chapter 7. However, at the time of writing hardware H.264/SVC codecs do not appear to be available, which implies that mobile devices with comparatively limited processing ability are not able to decode SVC streams. A commercial solution by Vidyo for video conferencing appears to only send the base layer to 3G/4G devices, whereas other SVC layers are sent to static devices.

CDNs are another related area of commercial activity. In brief, a CDN distributes content from centrally placed servers to edge servers that act as caches for that content. The content may be load-balanced between these *point-of-presences (PoPs)*. When (say) a streamed video clip is selected from a central server, the request is diverted to an edge server or PoP and the originating requester application is also informed of the download source. Selection of a suitable edge server is by some criterion such as cost or quality of service. The problem of occasional high demand for (say) a popular TV program (flash crowds that can overwhelm client server systems) is reduced by a CDN but still presents a weakness.

A hybrid P2P CDN can help in this situation, as popular content is downloaded by the peers themselves to the network edge. Thus P2P assisted systems are more scalable than pure CDNs. Peers act as alternative content servers. In effect, this removes the cost of distributing data from a central server to the edge that a CDN provider must meet. However, there are some disadvantages of P2P CDNs. Due to the asymmetric nature of twisted-pair telephone line transmission under *Asymmetric Digital Subscriber Line (ADSL)*, there may not be enough bandwidth to conveniently support the uploading of content from a peer. A P2P CDN may also generate more traffic than a traditional CDN, as material is downloaded to multiple peers and then can be uploaded. It was probably the additional traffic generated and in particular that caused the original version of the BBC iPlayer (a UK catch-up TV service operated by a public broadcaster) to be converted from P2P delivery or rather a hybrid P2P delivery service from Kontiki. In doing that the BBC avoided the risk of *Internet Service Provider (ISP)* traffic blocking or traffic shaping. Another issue is the cost to the peer of acting as a P2P host in terms of reduction in compute power and occupation of upload bandwidth. The risk of downloading malware and the need to protect content providers' intellectual property must also be addressed in a commercial setting (as is promised for example by Kontiki, which operates a closed public-key infrastructure).

A clean-slate hybrid P2P solution to the delivery of streaming content is illustrated by China Telecom's *Media Telecom Network (MTN)*. This is intended to provide live and on-demand TV through cache servers and by exchange between the users' terminals. These can consist of PCs, set-top-boxes, and mobile phones. In contrast to research, which tends to concentrate on the P2P streaming mechanisms and architecture, a system such as MTN must manage subscribers, provide billing (including authentication, authorization, and accounting), and collect viewing statistics. In such a system terminal, churn is a constant issue for a P2P system but randomized selection of suitable peers is a way to counter that. Content exchange management is responsible for making the initial connection and the establishment of the initial streaming path, after which suitable peers are also identifies. There is also a need to maintain an electronic program guide. All content is subject to digital rights management, which requires authentication as part of the key exchange process. Thus, in such a large scale system, P2P exchange is one part of a much large system.

Streaming of TV programs is one of the commercial success stories of P2P (along, of course, with Skype, from the originators of the Kazaa file-sharing application — acquired by Microsoft in October 2011). The background to these streaming systems is considered in Chapter 3. These applications appear particularly popular in Asia and tend to concentrate on live programs rather than pre-recorded and stored TV programs. For example, QQLive is only available in Chinese but apparently attracts 3 million web site visitors per day. PPS.tv is available in a variety of Asian languages, with reportedly several hundred thousand users viewing at around 200-500 kbps per stream, i.e. low resolution TV, hosted by an originating server with a link no more bandwidth capacity than that of domestic ADSL. Other Asian P2PTV applications are TVUPlayer, PPLive, and PPStream. TVUplayer is one of the few that also carry U.S. channels as well. It should be noted that some re-broadcasting apparently comes without a license to do so.

CoolStreaming was a pioneering live P2PTV streaming solution from about 2000 onwards. At one point in time, CoolStreaming had as many as 80,000 concurrent viewers by 2008. CoolStreaming used a BitTorrent form of content distribution but targeting the chunk-based players of the Real Player, Windows Media Player variety (refer back to the earlier discussion). At the same time it carefully matched the resolution of streams to the download capability of access links. A gossip-based broadcast algorithm distributed TV program availability. The early version of CoolStream (according to its originators) suffered from a slow start-up time and possible reliance for download on a single source, which in turn could lead to rejection from the download of very popular content (called "flash crowds"). These problems were subsequently resolved firstly by servers acting as push agents (after the first pull operation by a client) rather than, as previously, clients pulling all video chunks from a server. Multipath and multi-stream delivery was deployed to make the service more reliable. However, notice that for live streaming it may be necessary to send redundant streams to ensure continuity, thus causing an increase in bandwidth usage, though probably no more than if multiple independent streams are distributed from the same server. Copyright issues seem to have closed CoolStreaming in 2005. However, Roxbeam Corporation, operating from Beijing, now has a P2P multicast TV solution, along with CDN support (a strengthened form of P2P using pre-positioned servers within China), VoD, and a service called RayCAST, which is a content production kit for P2P broadcasting. Roxbeam also now targets set-top boxes acting as P2P clients.

There is an issue over the quality of the viewing experience for such a service and Chapter 9 shows how mobile viewing experiences should best be assessed. While quality-of-service (QoS) assessment both on the network and in the quality of the video is sufficient for static networks, for mobile networks the variety of device displays, the expectations of viewers and the type of content are issues that need to be considered. These are generally grouped under a quality-of-experience (QoE) heading, which also takes into account the video viewing quality at the end device. Thus, QoE is not the same as QoS.

In Western Europe, a company like Zattoo has not apparently made as much progress as Asian P2PTV has. In part, this is due to licensing restrictions which are an obstacle to the re-broadcast of main-stream TV programs. There is also a problem with providing revenue, which in Switzerland at one point came from adding advertisements to re-broadcast German TV, which was then challenged in the courts by film content owners. Interestingly, from the point-of-view of this book is that there is now an iPad application for Zattoo, which implies that mobile streaming may provide a target for such P2P TV in Western Europe. Joost is another English-language P2P TV streaming service, with revenue clearly provided from advertisements, as these interrupt streams at regular intervals. Licensing issues restrict content viewing outside the USA.

Another area in which (hybrid) P2P TV may find a use is in narrowcasting, with corporate broadcasting being a case in point. The advantage of P2P in a corporate setting is that it reduces the cost of server deployment, while at the same time avoids the problems of ISP blocking or traffic shaping that occur in public networks. For example, Kontiki Inc. now offers such a service, after its previous foray into public Internet TV. Tele-education across university and college networks could also benefit from peer-assisted video streaming. Both corporate networks and educational establishments (at least in Anglo-

Saxon commercially minded universities) have a need for the delightfully-named ego-casting, in which the CEO or the principal of the university broadcast their views to the workers/academic and office staff.

Some of the chapters have been mentioned in the previous exploration of the commercial impact of P2P streaming. Chapter 1 is a short chapter that acts as an extension to this Preface, in that it considers the practicalities of extending the P2P paradigm to the commercial world. In particular, it considers Fonera wireless routers, which allow others to share some of their WLAN signal with other Fon users, thus creating a Fon spot (by analogy with wireless hotspot). One type of user exchanges free roaming for free access to others' Fon spots. It is also possible to receive revenue from the purchase of passes when these passes result in access to one's Fon spot. However, this business model is not widely implemented. In the UK, where this preface is being written, one of the largest ISP's had joined with Fon to create a network of Fon spots, and this editor was able to use such a hotspot when Internet access using ADSL was somehow discontinued by his erstwhile ISP. Thus, Fon spots can also act as backup networks. Chapter 1 also considers fixed-mobile convergence, which P2P can assist in. Chapter 2 explores this in more detail from the point-of-view of he convergence of the Internet with cellular phone networks, which requires an integration mechanism, namely the *IP Multimedia Subsystem (IMS)*. Chapter details a plan to integrate P2P with an overall system, which should be compared to the MTN of China Telecom.

Chapters 3 and 4 are related in that Chapter 3 reviews the development of P2P streaming from P2P file sharing while Chapter 4 continues the story by reviewing the need for streaming adaptation to varying network conditions. One place where such adaptation is certainly needed is in wireless networks, as the time granularity of channel change can be very rapid. In fact, as reported in Chapter 4, existing commercial solutions will be insufficient during periods of device mobility. Chapter 5 is the first of a number of chapters in this book that examine ad hoc networks. As previously mentioned, the chapter takes a hard look at the prospects of mapping the P2P paradigm onto ad hoc networks. Such networks have applications in emergencies and disasters when the existing network infrastructure has been removed. They are also attractive to security forces. In the cellular world, they have a role through multi-homing in providing backup capacity when cells are overcrowded or they can extend the coverage of a cellular system. However, the business model for capacity sharing for extensions of cellular systems remains unclear. In one area, that of wireless vehicular networks, there probably is a business model justifying the 'infotainment' as well as a case for public safety, reporting traffic congestion and accidents. Chapter 6 shows, particularly for vehicular networks, that more detailed modeling is required, particularly as it is not easy to set up wireless networks test-beds for automobiles or provide analytic solutions. Chapter 7 goes on to consider the feasibility of peer-to-peer video streaming in such environments.

This book is essentially about research, and from chapter 7 onwards, possible research directions are considered. Currently, as previously mentioned though scalable video coding has the potential to simplify video streaming within mobile networks, development of such systems is at the cutting edge of research, especially for peer-to-peer streaming. As Chapter 9 demonstrates, scalable video coding brings the prospect of adaptation to peer-to-peer networks. In that case, streaming should be aware both of the P2P network overlay and also the network underlay upon which the P2P network is mapped. If mobile P2P streaming is to become a reality then some way of judging the users' quality of experience is needed, if only to try to improve it in situations in which it is compromised. As previously mentioned Chapter 9 examines this issue as well as giving a guide to video streaming quality of experience assessment. Peer-to-peer streaming is a form of application layer multicast, occasioned by the failure of IP multicasting to be effectively and universally deployed. Chapter 10 places P2P multicast in the context of reliable application-layer multicasting and reviews the research that is still ongoing in this field.

Finally, the last section of the book considers related network development areas. Chapter 11 continues the theme of ad hoc networks but now instead of the familiar WLAN networks, the topic is extended into wireless mesh networks and in particular to broadband wireless technologies such as WiMAX. It is sometimes forgotten by European technologists, who believe that Long Term Evolution(-advanced) broadband wireless will predominate that IEEE 802.11m, is an at least if not better technology and that it is now widely deployed in much of the world, as WiMAX forum statistics confirm. Chapter 12 continues the theme of wireless network exploration with a look at multi-homed devices (ones that can communicate over more than one network type). This field presents a fertile research topic and represents a possible target of P2P streaming, once routing problems have been addressed. Chapter 12 represents a very detailed investigation of how to optimize routing in the mobile multi-homed environment. As was mentioned, at the beginning of this preface, Chapter 13 represents a future view of networks in which just as in P2P networking, an overlay is used. However, in this case routing is based on information rather than node virtual topology. In fact, looking back to chapter one, one can see that this chapter represents the latest turn in fixed-mobile convergence, of which P2P communication has a strategic part (see also Chapter 2). It is hoped that the books' content will be stimulating, giving a foundation in the first section, moving on to more specific topics in the second section, with research contributions on future networks in the final section.

In terms of the overall importance of understanding trends in video streaming, including P2P steaming, and the contribution of mobile video watching in particular, a number of salient facts can be pointed to. The number of users watching multimedia content over the Internet, both live and on-demand, has dramatically increased in recent years, and is now the leading source of Internet traffic, surpassing even peer-to-peer file sharing (before peer-to-peer streaming the leading contribution of the peer-to-peer paradigm). This trend has been reported in several surveys, and is confirmed by network operators. Over one third of the top 50 sites by volume are based on video distribution (Source: Cisco Visual Networking Index: Usage, 2010-2015). Netflix, a provider of on-demand Internet streaming video, at the moment accounts by itself for almost 30% of peak period downstream traffic (Source: Sandvine's Global Internet Phenomena Report - Spring 2011). The sum of all forms of video is expected to exceed 90% of global user traffic by 2014. The web is becoming one of the main broadcasting platforms; this is due to the recent availability of the so called *three screens* (digital TV, PC, and smartphones) capable of accessing it, and the increase of the number of faster connections to the network. The most recent sport events, such as the FIFA World Cup and the Vancouver 2010 Winter Games, have been watched through streaming by several million unique simultaneous viewers. Mobile Internet data traffic stands at 237 Peta Bytes (PB) per month in 2010 but is set to rise to 6,254 PB in 2015 and could exceed that of the wired Internet by 2015 according to Cisco's published estimates. Mobile Internet is by far the fastest growing new service with a forecast increase in data traffic over the next five years of 92%. In fact some markets such as South America according to Sandvine have largely replaced fixed or wired access with wireless or mobile access. Overall according to Cisco, unmanaged IP consumer traffic is predicted to be 53.3 Exabytes per month by the end of 2015, as opposed to 11.8 for managed networks, and 4.9 over mobile networks. This book is a modest contribution to understanding these developments.

The book will be of interest to industry managers, as it offers insights into current and future communication technologies for a converged Internet that promises soon to be dominated by multimedia applications (at least in terms of bandwidth consumption). Chapter 1, together with this preface, marks out that relevance, as well as indicating the challenges of selecting a viable business plan. Graduate researchers will find that individual chapters identify research areas and provide pointers to and a review of the research literature in the area. Some chapters in the final section also act as detailed research reports. Some of the chapters are also suitable as a basis of Masters' dissertation topics and likewise provide a kit to start such a project off. As such, a library copy will act as a valuable resource to such students and researchers.

Martin Fleury
University of Essex, UK

Table of Contents

Detailed Table of Contents

Section 1
Foundations

Peer- to-peer technology has dramatically transformed the landscape of the Internet's traffic in recent years. In this introductory chapter, the authors highlight how the technology relates to the convergence of fixed and mobile networks with features that work irrespective of location, access technology, or user-interactive devices.

The growth of the Internet and its popular services are forcing telecom operators to provide advanced services to their subscribers, as traditional voice services are no longer enough to attract more customers. To enable more innovative and value-added IP services and take advantage of the services that the Internet provides, the IP Multimedia Subsystem (IMS) is introduced. The IMS provides a complete access-agnostic architecture and framework that facilitates the convergence of the mobile network, removing the gap between the two most successful communication networks: cellular and Internet network. The harmonized All-IP platform has the potential to provide all Internet services with a more cost-effective and more efficient architecture than the circuit-switched networks do. However, by merging two of the most successful networks, the integration of two network models with different concerns and motivations is not without its problems, among which, the scalability issue is the most essential when supporting content delivery services. The purpose of this chapter is to study and design a new content delivery network infrastructure, PeerMob, merging the Peer-to-Peer technology with the IMS framework, which benefits IMS with scalability, reliability, and efficiency features coming with decentralized P2P architecture. The chapter also puts this P2P IMS paradigm under realistic network conditions and strenuous simulation to evaluate the performance of the P2P IMS system.