

Peer-to-Peer Video Streaming

Peer-to-Peer Video Streaming

Eric Setton

Streaming Media Systems Group
Hewlett-Packard Laboratories
Palo Alto, CA

Bernd Girod

Information Systems Laboratory
Department of Electrical Engineering
Stanford University
Stanford, CA

Eric Setton
Streaming Media Systems Group
Hewlett-Packard Laboratories
Palo Alto, CA 94304
eric.setton@hp.com

Bernd Girod
Information Systems Laboratory
Department of Electrical Engineering
Stanford, CA 94305
bgirod@stanford.edu

ISBN-13: 978-0-387-74114-7

e-ISBN-13: 978-0-387-74115-4

Library of Congress Control Number: 2007932753

© 2007 Springer Science+Business Media, LLC

All rights reserved. This work may not be translated or copied in whole or in part without the written permission of the publisher (Springer Science+Business Media, LLC, 233 Spring Street, New York, NY 10013, USA), except for brief excerpts in connection with reviews or scholarly analysis. Use in connection with any form of information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed is forbidden.

The use in this publication of trade names, trademarks, service marks and similar terms, even if they are not identified as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

Printed on acid-free paper.

9 8 7 6 5 4 3 2 1

springer.com

To my parents
Eric Setton

Preface

Live video broadcasting over the Internet requires an infrastructure capable of supporting a large number of simultaneous unicast connections. Since the costs of providing this service grow with the number of viewers, television networks have been reluctant to offer it to their customers on a large scale. Peer-to-peer architectures are an alternative where viewers contribute their resources to the network to act as relays, hence overcoming the need for a dedicated content delivery infrastructure.

Peer-to-peer video streaming systems offer the same advantages as peer-to-peer file transfer networks but face additional challenges since data transfer needs to occur continuously to avoid playout interruptions. This is particularly difficult since the peers are connected to the Internet by links which may have different capacity and reliability. Moreover, data delivery paths may simply disappear without prior notice, e.g., when a peer leaves the broadcast. This challenging environment is a perfect field of application for recent advances in compression, streaming, and networking and a catalyst for new progress. Remarkably, functioning solutions have emerged and the research community now expects that in the future peer-to-peer video streaming system will be used for large-scale live television distribution over the Internet.

One of the goals of this book is to provide an overview of today's state-of-the-art peer-to-peer video streaming technology and to show how it can be improved in terms of video quality, robustness, and latency. We present adaptive video coding and streaming techniques which enhance the performance of conventional client-server systems and extend them to peer-to-peer multicast. We focus on throughput-limited environments where congestion often hampers interactivity and fast response times. We analyze the benefits of scheduling packet transmissions and retransmissions in a way which adapts to the particular properties of video streams and to the changing topology of peer-to-peer networks. The performance of the solutions we propose is assessed by analyzing the results of experiments carried out over simulated networks with large numbers of peers.

The contents of this book is the result of research carried out over the course of the last few years at Stanford University. The text itself is based on the 2006 Ph.D. dissertation of the first author. Compared to this original work, the presentation has been revised substantially and additional material of interest has been added.

The authors would like to thank the members of the Image, Video and Multimedia Systems group at Stanford and the members of the Streaming Media Systems group at HP Labs for many insightful discussions. We are grateful to John Apostolopoulos for his continuous support and his many helpful comments and to Jeonghun Noh for his work on the peer-to-peer control protocol. Amy Brais and Jennifer Evans, at Springer, were very supportive and provided us with a lot of assistance.

Palo Alto, CA, USA

*Eric Setton
Bernd Girod*

Contents

1	Introduction	1
2	Background	5
2.1	Video Compression	5
2.1.1	H.264 Video Coding	5
2.1.2	Distortion Models	7
2.2	Video Streaming	8
2.2.1	Error Resilience	8
2.2.2	Congestion Control	11
2.2.3	Path Diversity	12
2.3	Multicast Architectures	13
2.3.1	IP Multicast	13
2.3.2	Content Delivery Networks	13
2.4	Peer-to-Peer Systems	14
2.4.1	Peer-to-Peer File Transfer, the Example of BitTorrent ..	14
2.4.2	Peer-to-Peer Streaming	17
3	Streaming over Throughput-Limited Paths	19
3.1	Video Encoding for Throughput-Limited Paths	19
3.1.1	End-to-End Rate-Distortion Performance Model	20
3.1.2	Experimental Results	28
3.2	Congestion-Distortion Optimized Scheduling	33
3.2.1	Channel Model	34
3.2.2	Evaluating a Schedule	35
3.2.3	Randomized Schedule Search	41
3.2.4	CoDiO Light	45
3.2.5	Experimental Results	47
3.3	Chapter Summary	61

4	Peer-to-Peer Control Protocol	63
4.1	Protocol Description	64
4.1.1	Different Peer States	64
4.1.2	Different Tree Connection States	69
4.1.3	Multicast Source	74
4.1.4	Protocol Settings	75
4.2	Experimental Protocol Evaluation	77
4.2.1	Experimental Setup	77
4.2.2	Control Protocol Traffic Distribution	79
4.2.3	Join and Rejoin Latency	79
4.2.4	Scalability	80
4.2.5	Limiting Throughput	82
4.3	Chapter Summary	86
5	Video Streaming over a Peer-to-Peer Network	87
5.1	Video Streaming Protocol	88
5.1.1	Video Packet Transmission	88
5.1.2	Retransmissions	89
5.2	Peer-to-Peer CoDiO Scheduling	89
5.2.1	Sender-Driven Prioritization	91
5.2.2	Distortion-Optimized Retransmission Scheduling	93
5.2.3	Scheduler Evaluation	95
5.3	Experimental Results	102
5.3.1	Video Sessions	102
5.3.2	Diversity	102
5.3.3	CoDiO P2P	103
5.4	Chapter Summary	114
6	Conclusions and Future Work	115
6.1	Conclusions	115
6.2	Future Work	116
A	Video Experiments	117
A.1	Video Streaming	117
A.1.1	Encoding Structures	117
A.1.2	Latency-Constrained Video Streaming	118
A.1.3	Error-Resilient Decoding	118
A.1.4	Quality Metric	119
A.2	Video Sequences	121
A.2.1	Container	121
A.2.2	Foreman	123
A.2.3	Mobile	125
A.2.4	Mother & Daughter	127
A.2.5	News	129
A.2.6	Salesman	131

References 133

Index 147