Peer-to-Peer Television: Ready for Prime Time?

Bernd Girod (Fellow, IEEE), Stanford University, Stanford, CA, USA
bgirod@stanford.edu

Peer-to-peer (P2P) live video multicast transmits a stream to a large population of Internet nodes, utilizing the uplink bandwidth of participating peers. Media delivery is accomplished via a distributed protocol that lets peers self-organize into distribution trees or meshes. Different from P2P file sharing, such as BitTorrent, audio and video are streamed in real-time to provide all connected peers with a viewing experience akin to broadcast television. Unlike content delivery networks (CDNs), P2P multicast does not require an expensive dedicated infrastructure. It can, in principle, scale to an unlimited number of users as its resources grow with each new peer.

Trees vs. Meshes. It is common to contrast tree-based and mesh-based overlay networks for P2P live multicast [1,2]. Tree-based systems build at least one, but more typically several complementary distribution trees and split the original stream among them. Packets flow along the same paths, until peers depart, new peers join, or changing network conditions require a re-arrangement. In mesh-based systems, peers trade chunks of data not yet received within a sliding window, resulting in a random "gossip" dissemination of the stream. Mesh-based protocols are often an extension of the BitTorrent protocol, which itself is not suitable for streaming media delivery. The distinction trees vs. meshes is somewhat arbitrary, as multiple trees form a mesh, while a particular chunk in a mesh-based system traverses a spanning tree (if no duplicate delivery is made).

Push vs. Pull. The more important distinction is a consequence of the route persistence of tree-based vs. mesh-based systems. Tree-based systems with their infrequent route changes lend themselves to push streams to the peers. When a packet arrives at a peer, it is forwarded instantly to the next peer without further ado. This is not the case for the data-driven, mesh-based systems, which usually wait until another peer "pulls" a chunk it needs, resulting in an additional delay. Because of its simplicity and conceptual elegance, mesh-pull is used by a number of popular systems, such as CoolStreaming [3], PPLive, or PPStream. Alas, these systems test the user's patience with start-up delays ranging from 10s of seconds to minutes [4]. Large differences in playout delay of different peers can also be annoying in some settings. In contrast, tree-push systems, like ESM (End-System Multicast) [2,5] or our own SPPM (Stanford Peer-to-Peer Multicast) [6,7] can bring startup delays down to a few seconds.

Minimizing Transport Latency. Building on our research on low-latency media streaming, we have developed a P2P live video multicast system designed for short transport latency. This work resulted in SPPM [6,7], which has been optimized and tested extensively and has recently been commercially deployed. The SPPM protocol builds multiple complementary multicast trees to minimize end-to-end delay and pushes an H.264/AVC video stream (and accompanying audio) down these trees. Error control is achieved by retransmissions requested from alternative parents, and congestion is avoided by video/network-aware packet scheduling using the CoDiO (Congestion-Distortion Optimized) framework.

How Many Trees? With SPPM, eight trees are typically used, providing multipath robustness against packet losses and peer churn. More importantly, eight trees also provide sufficiently fine-grained bandwidth aggregation. With a single tree, peers that cannot upload at the full bit-rate would not be able to contribute at all. In our experience, a 500-600 kbps stream would force over half of the peers to be "leeches" in today's Internet. With N trees, the smallest contribution is 1/Nth of this rate, and N=8 permits almost all peers to contribute to at least one tree. Note that the distribution is rather skewed; about 20% of the peers provide 70% of the total uplink bandwidth.

Peer Accessibility. Firewalls and Network Address Translation (NAT) restrict the accessibility of most Internet nodes and thus pose a major hurdle for P2P systems. In data we collected for 33,000 users during an SPPM multicast of the 2008 Electronic Sports World Cup [8], about 2/3 of the peers do not accept a direct connection request. With the help from an external STUN server, the situation is much improved. Now, about 15% of randomly drawn peer pairs cannot connect. This is still a problem for small peer groups (consider the extreme case of one sender, one receiver). An easy solution routes the video stream through a universally accessible media gateway, which acts as an intermediary. As the number of peers grows, it becomes increasingly unlikely that a peer cannot find another peer that it can connect to. We measured for the event above that external server assistance increased the available peer uplink bandwidth by 2.3X, bringing it within 80% of the theoretical maximum. Alas, external assistance to traverse NATs and firewalls can add seconds to the pre-roll delay.

Ready for Prime Time? Early deployments have demonstrated that P2P multicast can scale to hundreds of thousands of simultaneous users. NATs/firewalls are a nuisance, but surmountable. With today's broadband access, P2P live video multicast with a few hundred kbps is possible. This is enough for good-quality CIF (352x288) video, or bad VGA (640x480). HD video will become feasible, as bandwidth grows, possibly in less than 5 years. Alas, outside of the technical realm, it is questionable whether ISPs will continue to foot the bill for "free" uplink bandwidth, as P2P live multicast gets ready for prime time. A number of researchers have noted that many limitations of P2P video multicast can be overcome today by wedding P2P with a CDN. Such marriage is likely to be happy, as P2P protocols are well suited to operate a CDN. The CDN can dynamically inject bandwidth when and where it is needed, while P2P acts as an outer layer that boosts the bandwidth of the CDN many-fold.

Beyond Prime Time. In our current research at Stanford, we are experimenting with various extensions of P2P video multicast [9]. One is the extension of SPPM for heterogeneous mobile clients, based on distributed transcoding in the overlay. Another is the seamless integration of live streaming and time-shifted streaming. While it is nearly impossible to achieve low latencies by grafting streaming capabilities onto a non-real-time P2P protocol such as BitTorrent, low-latency streaming systems such as SPPM can be easily extended to support distributed storage for time-shifting. The spatial counterpart of time-shifted streaming is Interactive Region-of-Interest (IRoI) streaming. We allow different users to watch different portions of the high-spatial-resolution video with arbitrary zoom factors and distribute the shared portions of the video stream using P2P live multicast. This approach is promising also for efficient distribution of free-viewpoint TV, where each viewer only needs one view at a time (or two for stereo), and the continuous broadcast of tens or hundreds of views is prohibitive.

References

[1] N. Magharei, R. Rejaie, and Y. Guo, "Mesh or Multiple-Tree: A Comparative Study of Live P2P Streaming Approaches," in Proc. IEEE INFOCOM'07, Anchorage, AL, May 2007

[2] J. Liu, S. Rao, B. Li, and H. Zhang, "Opportunities and Challenges of P2P Internet Broadcast," Proc. IEEE, vol. 96, no. 1, pp. 11–24, Jan. 2008.

[3] X. Zhang, J. Liu, B. Li, and T.-S. P. Yum, "BCoolStreaming/DONet: A Data-driven Overlay Network for Peer-to-Peer Live Media Streaming," in Proc. INFOCOM'05, Miami, FL, Mar. 2005.

[4] X. Hei, C. Liang, J. Liang, Y. Liu, K. W. Ross, "A Measurement Study of a Large-Scale P2P IPTV System," IEEE Trans. Multimedia, vol. 9, no. 8, pp. 1672-1686, Dec. 2007.

[5] Y. Chu, S. Rao, S. Seshan, and H. Zhang, "A Case for End System Multicast," IEEE J. Sel. Areas Commun., vol. 20, pp. 1456–1471, Oct. 2002.

[6] P. Baccichet, J.Noh, E. Setton, and B. Girod, "Content-Aware P2P Video Streaming with Low Latency," *Proc. IEEE International Conference on Multimedia and Expo, ICME 2007*, Beijing, China, Jul. 2007.

[7] E. Setton, P. Baccichet, and B. Girod, "Peer-to-Peer Live Multcast: A Video Perspective," Proc. IEEE, vol. 96, no. 1, pp. 25-38, Jan. 2008.

[8] J. Noh, P. Baccichet, and B. Girod, "Experiences with a Large-Scale Deployment of the Stanford Peer-to-Peer Multicast," *Proc. International Packet Video Workshop - PV2009*, Seattle, WA, May 2009.

[9] J. Noh, P. Baccichet, F. Hartung, A. Mavlankar, and B. Girod, "Stanford Peer-to-Peer Multicast (SPPM) – Overview and Recent Extensions," *Proc. International Picture Coding Symposium, PCS 2009*, Chicago, IL, May 2009.

Bernd Girod is Professor of Electrical Engineering at Stanford University, California. He also holds a courtesy appointment with the Stanford Department of Computer Science and serves as Director of the Stanford Center for Image Systems Engineering (SCIEN). His current research interests include image and video coding, networked media systems, and image-based retrieval.

He received his M. S. degree in Electrical Engineering from Georgia Institute of Technology, in 1980 and his Doctoral degree from University of Hannover, Germany, in 1987. He joined Massachusetts Institute of Technology, Cambridge, MA, USA, and was an Assistant Professor of Media Technology at the Media Laboratory there until 1990. From 1990 to 1993, he was Professor of Computer Graphics and Technical Director of the Academy of Media Arts in Cologne, Germany, jointly appointed with the Computer Science Section of Cologne University. From 1993 until 1999, he held the Telecommunications Chair at University of Erlangen-Nuremberg, Germany, and was the Head of the Telecommunications Institute I and director of the Telecommunications Laboratory. He served as the Chairman of the Electrical Engineering Department from 1995 to 1997.

As an entrepreneur, Professor Girod has worked successfully with several start-up ventures as founder, investor, director, or advisor. Most notably, he has been a co-founder and Chief Scientist of Vivo Software, Inc., Waltham, MA (1993-98); after Vivo's aquisition, 1998-2002, Chief Scientist of RealNetworks, Inc. (Nasdaq: RNWK). He has served on the Board of Directors for 8x8, Inc., Santa Clara, CA, (Nasdaq: EGHT) 1996-2004, and for GeoVantage, Inc., Swampscott, MA, 2000-2005. In 2007, he co-founded Dyyno, Inc. Palo Alto, CA. From 2004 to 2007, he also served as Chairman of the Steering Committee of the new Deutsche Telekom Laboratories at the Technical University of Berlin.

Professor Girod has authored or co-authored one major text-book (printed in 3 languages), four monographs, and over 400 book chapters, journal articles and conference papers, and is a named inventor of over 20 US patents. He has been a member of the IEEE Image and Multidimensional Signal Processing Technical Committee from 1989 to 1997 and has served on the Editorial Boards for several journals in his field, among them as founding Associate Editor for the IEEE Transactions on Image Processing and Area Editor for Speech, Image, Video & Signal Processing of the IEEE Transactions on Communications. He has served on numerous conference committees, e.g., as Tutorial Chair of ICASSP-97 in Munich and again for ICIP-2000 in Vancouver, as General Chair of the 1998 IEEE Image and Multidimensional Signal Processing Workshop in Alpbach, Austria, as General Chair of the Visual Communication and Image Processing Conference (VCIP) in San Jose, CA, in 2001, and General Chair of Vision, Modeling, and Visualization (VMV) at Stanford, CA, in 2004, and General Co-Chair of ICIP-2008 in San Diego.

Professor Girod was elected Fellow of the IEEE in 1998 'for his contributions to the theory and practice of video communications.' He has been named 'Distinguished Lecturer' for the year 2002 by the IEEE Signal Processing Society. He received the 2002 EURASIP Best Paper Award (with J. Eggers) and the 2004 EURASIP Techical Achievement Award, the IEEE Multimedia Communication Best Paper Award in 2007, and the EURASIP Image Communication Best Paper Award 2008. He was elected a member of the German National Academy of Sciences (Leopoldina) in 2007 and a Fellow of EURASIP in 2008.