
By Reza Rejaie

ANYONE CAN BROADCAST VIDEO OVER THE INTERNET

A mesh-based P2P streaming mechanism provides TV-like distribution service to average computer users over the Internet with no special support from the network.

Being able to stream live video from a desktop to a large number of end systems over the Internet allows even average users to host their own TV shows at home. The availability of broadband connections, along with users' increasing ability to capture and view multimedia streams (video and audio), enables average users to send (and receive) decent-quality video streams to (and from) the Internet. The missing component for providing such TV-like service has been a scalable one-to-many distribution mechanism that accommodates live streaming content. Such a mechanism should ensure in-time delivery of streaming content from a source to each interested user throughout the Internet.

Since the early 1990s, it's been expected that such one-to-many service would be offered through IP multicast. However, its limited deployment by Internet service providers has motivated a new approach based on peer-to-peer (P2P) overlays. In this paradigm, participating end systems, or peers, establish pairwise connections to form an application-level overlay. The P2P overlay is used to deliver streaming content from the source to all participating peers with no special support from the network or the need for high-bandwidth connections. This class of one-to-many streaming

over P2P overlays—usually referred to as P2P streaming [1, 3, 4, 6]—involves two key features:

Scalability of available resources. Individual peers contribute their outgoing bandwidth by forwarding their available content to other connected peers. The aggregate available bandwidth for serving content (system capacity) scales organically with the number of participating peers; and

Ease of deployment. Groups of participating peers easily deploy the system, running

The missing component for TV-like service over the Internet is a SCALABLE ONE-TO-MANY DISTRIBUTION MECHANISM that accommodates live streaming content.

the application on their desktop computers with no special support from the network.

Here, I review the key design challenges and explore mesh-based P2P streaming as a promising approach for providing this service. The key role of end systems in content delivery introduces several challenges in designing P2P streaming mechanisms:

Bandwidth heterogeneity and asymmetry. Participating peers in a session are likely to have heterogeneous-access-link bandwidth. The access-link bandwidth of individual peers is often asymmetrical, that is, the incoming bandwidth is greater than the outgoing bandwidth. The heterogeneity and asymmetry of access-link bandwidth imply that various connections of the overlay are likely to have different bandwidth. Therefore, a high-bandwidth peer may need to simultaneously receive streaming content from multiple peers in order to obtain the desired streaming quality.

Dynamics of peer participation (churn). Participating users in a P2P system (and thus their computers) may join or leave a session at any arbitrary point in time. These user-driven dynamics result in ongoing changes of the overlay topology that could disrupt the streaming-content delivery through the overlay.

Utilization of available resources. While the aggregate available resources scale with the user population, the system's ability to utilize available resources may not scale proportionally for two reasons: The structure of the overlay may prevent some peers from contributing their outgoing bandwidth; for example, leaf peers in a tree-shaped overlay do not forward content to other peers. And participating peers may lack sufficient useful content to fully utilize the available bandwidth to their connected peers.

MESH-BASED P2P STREAMING

The recent success of BitTorrent, a P2P file distribution protocol, has inspired a new and promising approach to P2P streaming called mesh-based, or data-driven, P2P streaming. In mesh-based P2P streaming, participating peers form a randomly

connected mesh-shaped overlay and incorporate swarm-like content delivery similar to BitTorrent [2]. This approach overcomes all three challenges—bandwidth heterogeneity, peer dynamics, and resource utilization—outlined earlier, accommodating delivery of good-quality streams to a large number of users.

Having multiple neighbors enables each peer to gracefully cope with the departure of any of them due to churn. Accommodating bandwidth heterogeneity implies that the quality of the delivered stream to each peer should be proportional to its incoming access-link bandwidth. Each peer could determine the number of its neighbors proportional to its incoming bandwidth to achieve proper bandwidth connectivity to the overlay.

An elegant approach to facilitating bandwidth heterogeneity involves encoding a video stream with Multiple Description Coding (MDC). An MDC encoder organizes a video stream into multiple substreams in which each of them can be independently decoded to produce a low-quality version of the video [5]. Decoding several unique descriptions leads to progressively better quality. With MDC-encoded content, a low-bandwidth user may receive only one description of the video and view only a low-quality version of the video, while a high-bandwidth peer receives all descriptions of the video and views the maximum quality. While MDC provides the flexibility to accommodate bandwidth heterogeneity among peers, the content-delivery mechanism must still ensure that each peer receives a proper number of descriptions.

In swarm-like content delivery, the source provides various segments of that content to at least one peer in the session; participating peers can thus exchange their available segments until each one has the segments it requires. Swarming couples push-content reporting with pull-content requesting. Each peer periodically reports its newly available content segments to all its neighbors while requesting specific new segments from each of them. The segments requested by each peer from a neighbor are determined by a packet-scheduling algorithm based on the available content and bandwidth from

its neighbor. As the key component of content delivery, the packet-scheduling algorithm to each peer aims to utilize available bandwidth from individual neighbors in order to maximize its received quality, or the number of received descriptions.

Swarm-like content delivery utilizes available resources in the system in a scalable fashion if available content among peers is diverse during the session. A file-swarmling mechanism (such as BitTorrent) can accommodate this goal, as all content is available at the source. However, in a live streaming session content is progressively generated over time and must be delivered in a timely fashion. To effectively incorporate swarming into live P2P streaming, participating peers can maintain the same playout time, which is τ seconds behind the source's playout time; each peer views the content with a τ -second delay. This delay continuously provides τ -seconds worth of content that can be used by peers for swarming. In essence, τ determines the maximum available time for delivery of each segment, as well as the minimum required buffering at individual peers. The greater the value of τ , the more effective the swarm-like content is able to utilize the outgoing bandwidth of a larger group of participating peers; the cost is a longer playout delay between the source and the participating peers [3].

The P2P streaming approach has become increasingly popular among developers. P2P streaming software (and applications) is available for downloading on the Web, including wwtv (www.wwtv.com) and sopcast (www.sopcast.com), and is being used to broadcast popular events (such as the World Cup 2006) to thousands of concurrent viewers over the Internet. While no technical information is publicly available about them, they appear to use a modified version of BitTorrent to swarm a recent window of content, as discussed earlier.

These systems illustrate the feasibility of using mesh-based P2P streaming for scalable broadcasting of video over the Internet. However, several fundamental design issues about this approach are not well understood and require further investigation:

Overlay properties. The effect of overlay properties (such as peer degree, or number of neighbors, and directed vs. undirected pairwise connections) on the performance of content delivery is important but has not yet been fully studied. For example, increasing the peer degree improves the availability (and potential diversity) of content among an individual peer's neighbors but reduces bandwidth connectivity from each of these neighbors; *Pattern of content delivery.* The collective effect of

packet scheduling at individual peers determines the overall pattern of content delivery from the source to individual peers through a randomly connected overlay mesh. How the similarity (or difference) of the packet-scheduling algorithm among participating peers affects the performance of content delivery is an open question; *Coding issues.* MDC codecs are still only research prototypes able to generate a limited number of descriptions (often two or three). The codecs are neither widely available nor flexible enough to use in P2P streaming systems; and

Security and incentives. The open nature of the P2P communication model raises concerns about security, as well as about the incentives used to motivate individual peers to contribute resources. These issues must be addressed to limit the potential damage that might be caused by misbehaving users on the system.

All these issues are being explored by researchers and developers at the University of Oregon (mirage.cs.uoregon.edu/PRIME) and Microsoft Research (research.microsoft.com/projects/CoopNet), as well as at other research organizations.

CONCLUSION

Mesh-based P2P streaming is a promising approach for enabling computer users to broadcast video from their desktops and devices to a large number of users across the Internet, hosting their own shows from home. Even though P2P streaming systems are available today, further research is required to deepen our understanding of the performance, robustness, and limitations of this approach in practice. **C**

REFERENCES

1. Castro, M., Druschel, P., Kermarrec, A.-M., Nandi, A., and Singh, A. SplitStream: High-bandwidth content distribution in a cooperative environment. In *Proceedings of ACM SOSP* (Lake Bolton, NY, Oct. 19–22, 2003), 298–313.
2. Cohen, B. BitTorrent; www.bittorrent.com.
3. Magharei, N. and Rejaie, R. Understanding mesh-based peer-to-peer streaming. In *Proceedings of the ACM International Workshop on Network and Operating Systems Support for Digital Audio and Video* (Newport, RI, May 22–23, 2006), 56–61.
4. Padmanabhan, V., Wang, H., and Chou, P. Resilient peer-to-peer streaming. In *Proceedings of the IEEE International Conference on Network Protocols* (Atlanta, Nov. 4–7, 2003).
5. Wang, Y., Reibman, A., and Lin, S. Multiple description coding for video delivery. *Proceedings of the IEEE* 93, 1 (Jan. 2005), 57–70.
6. Zhang, X., Liu, J., Li, B., and Yum, T.-S. CoolStreaming/DONet: A data-driven overlay network for peer-to-peer live media streaming. In *Proceedings of IEEE INFOCOM* (Miami, Mar. 13–17, 2005).

REZA REJAE (reza@cs.uoregon.edu) is an assistant professor in the Department of Computer and Information Science at the University of Oregon, Eugene, OR.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright of Communications of the ACM is the property of Association for Computing Machinery and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.