

RESEARCH GETTING INTERNET VIDEO READY FOR PRIME TIME



If you've struggled to watch poor quality Web video, then you

might not believe that someday soon your TV will display crisp-looking movies and sitcoms sent via the Internet. But electrical engineering Professor Bernd Girod won't be surprised. In fact, he thinks that Internet video will soon look better than the broadcast television you are used to. His research is all about getting video on the Internet ready for prime time.

The Internet was not designed for video, which explains why video often appears small, blurry, sluggish, and choppy. Communications companies nevertheless want to use the Internet for television and movies because it is nearly ubiquitous and handles digital data with incredible versatility. The Internet would be ideal for delivering services like video on demand, if only the picture and sound were enjoyable.

Girod is helping that happen by combating the network delays, or "latency" that dogs Internet video. The Internet chops up data into little "packets." Each packet may contain a

single frame of video. The problem with the Internet is that packets often get lost between the server and the end user's "client" computer. For data like a Web page, this is no big deal. The computers can keep communicating, taking many seconds if necessary, to re-send lost data until it is all accounted for and assembled properly. But this kind of latency and disorder are intolerable in video.

Traditionally, streaming video software programs have worked around this problem by downloading a chunk of the video before starting playback. Once this buffer is downloaded and the video starts playing, the software scrambles behind the scenes to gather all the frames that are coming up next. But buffering doesn't always buy enough time (resulting in more buffering), forces the viewer to endure a big upfront delay, and doesn't allow for two-way communication such as live videoconferencing, or interactivity such as fast-forwarding and changing channels.

Rather than fixing the Internet, Girod and his students have slashed delays and improved quality by working around the Internet's video delivery flaws. They have, for example, improved a technique called "motion compensation" which is essential to compress video to manageable data transmission rates. With motion compensation, the server and client predict how the next frame will look by recognizing the motion in the scene. This trick avoids sending the same scene contents over and over again. Only what's truly new is transmitted.

Through their work, Girod and students have made motion compensation more robust by enabling "multiframe prediction" based on several past pictures instead of just the most recent frame. Their idea is now a key component of the latest video compression standard, known as H.264 or MPEG-4 AVC.

Girod and former PhD student Yi Liang engineered a system that is robust to packet losses by having the server and the client compare notes about what packets have been lost. In this way, the server can avoid basing motion compensation on frames that never made it from the server to the client. With multiframe prediction, there are usually several good alternative frames to choose from. This extra resilience to packet losses results in so much better video quality, that retransmissions of lost packets can be omitted altogether.

Girod and Liang also developed "adaptive media playout," which buys time by subtly slowing down the speed at which video (and audio) plays, based on network congestion and delays. Originally, they conceived the idea for telephony over the Internet. By applying sophisticated signal processing techniques, Girod and Liang are able to buy a few hundred extra milliseconds without

creating a perceptible difference in playback. The extra time is often all the client needs to capture late-arriving packets. In essence, by smartly slowing down playback, computers can often avoid stopping to rebuild a buffer.



Girod appears in Internet video via SCPD

Continued next page

Continued from page 2, Girod

Yet another important research topic in Girod's group, which includes PhD students Mark Kalman and Eric Setton, is "packet scheduling" in which the server, aware of network delays, tries to make intelligent choices about when to send packets and which ones don't have to be sent at all.

All of these advances, some of which have already been incorporated into practical systems, can be combined to provide high-quality video with latency of less than a second over a congested, loss-prone network. This is important for video transmission because, for all its benefits, the Internet is just such a network.

Girod's Web site: www.stanford.edu/~bgirod/

