

A system to keep cloud-based gamers in sync

August 31 2023, by Adam Zewe



Scientists from MIT and Microsoft Research have developed a unique technique that can synchronize media streams being streamed over different networks to multiple devices with less than 10 milliseconds of inter-stream delay. They used this technique to synchronize audio and video streams in cloud gaming, but it could also be more broadly applied in AR/VR applications. Credits: Credit: Jose-Luis Olivares/MIT

Cloud gaming, which involves playing a video game remotely from the cloud, witnessed unprecedented growth during the lockdowns and gaming hardware shortages that occurred during the heart of the COVID-19 pandemic. Today, the burgeoning industry encompasses a \$6 billion global market and more than 23 million players worldwide.

However, interdevice synchronization remains a persistent problem in cloud gaming and the broader field of networking. In cloud gaming, video, audio, and haptic feedback are streamed from one central source to multiple devices, such as a player's screen and controller, which typically operate on separate networks. These networks aren't synchronized, leading to a lag between these two separate streams. A player might see something happen on the screen and then hear it on their controller a half second later.

Inspired by this problem, scientists from MIT and Microsoft Research took a unique approach to synchronizing streams transmitted to two devices. Their system, called Ekho, adds inaudible white noise sequences to the game audio streamed from the cloud server. Then it listens for those sequences in the audio recorded by the player's controller.

Ekho uses the mismatch between these noise sequences to continuously measure and compensate for the interstream delay.

In real cloud gaming sessions, the researchers showed that Ekho is highly reliable. The system can keep streams synchronized to within less than 10 milliseconds of each other, most of the time. Other synchronization methods resulted in consistent delays of more than 50 milliseconds.

And while Ekho was designed for cloud gaming, this technique could be used more broadly to synchronize media streams traveling to different devices, such as in training situations that utilize multiple augmented or virtual reality headsets.

"Sometimes, all it takes for a good solution to come out is to think outside what has been defined for you. The entire community has been fixed on how to solve this problem by synchronizing through the network. Synchronizing two streams by listening to the audio in the room sounded crazy, but it turned out to be a very good solution," says Pouya Hamadani, an [electrical engineering](#) and [computer science](#) (EECS) graduate student and lead author of a paper describing Ekho.

Hamadani is joined on the paper by Doug Gallatin, a [software developer](#) at Microsoft; Mohammad Alizadeh, an associate professor of electrical engineering and computer science and a member of the Computer Science and Artificial Intelligence Laboratory (CSAIL); and senior author Krishna Chintalapudi, a principal researcher at Microsoft Research. The paper will be presented at the ACM SIGCOMM conference.

Off the clock

At the heart of interstream delay in cloud gaming is a fundamental problem in networking known as clock synchronization.

"If the controller and the screen could look at their watches and at the same time see the same thing, then we could synchronize everything to the clock. But a lot of theoretical work on clock synchronization shows that there are certain bounds you can never overcome," Hamadani says.

Many approaches attempt clock synchronization by ping-pong messaging, where a device sends a ping message to the server, which sends a pong message back. The device counts how long it takes the message to return, and cuts that value in half to calculate the network delay.

But the path over the network is likely asymmetric, so it may take more time for the message to reach the server than it does for the return message. Therefore, this method is unreliable and can introduce hundreds of milliseconds of error. Humans can typically perceive interstream delay once it reaches 10 milliseconds.

"So if something happens on the screen, we want it to happen within 10 milliseconds on the controller, as well," Hamadianian explains.

He and his collaborators decided to try listening to game audio to synchronize these separate streams.

In cloud gaming, the microphone on the player's controller records audio in the room, including game audio played by the speakers on the screen, which it sends back to the server. But using this for synchronization is unreliable because the room audio contains background noise.

So they designed Ekho to add identical sequences of extremely low-volume white noise, known as pseudo noise, to the game audio before it is streamed to the player's screen. It uses these pseudo-noise segments for synchronization.

Before building Ekho, the researchers conducted a user study to prove that players could not hear the pseudo noise in the game audio. These noise sequences are also resilient to compression, which is important because audio sent from the controller is highly compressed to speed the data transfer.

Pseudo noise, real success

The Ekho-Estimator module adds pseudo-noise sequences to the game audio. When it receives the recorded game audio from the controller, it listens for those markers and tries to line up the streams. This enables it

to precisely calculate the inter-stream delay.

The Ekho-Estimator sends that information to the Ekho-Compensator module, which either skips a few milliseconds of sound or adds a few milliseconds of silence to the game audio being sent by the server, which synchronizes the streams.

They tested Ekho on real cloud streaming sessions and found that it was superior to other synchronization methods, even when the microphone quality was poor or [background noise](#) was picked up by the recording.

Ekho limited interstream delay to less than 10 milliseconds for nearly 87 percent of the time during streams. No other method the team tested was able to cut that delay to less than 50 milliseconds.

"The traditional way of doing this, which involves trying to measure the synchronization error using the underlying network, the errors are significantly larger. When we started this project, we weren't sure whether this could even be done. But the accuracy we can get down to with Ekho, at sub-millisecond levels, it is unheard of," says Chintalapudi.

Impressed by these results, the researchers want to see how well Ekho performs in more complex situations, such as synchronizing five controllers to the same screen device. Also, since Ekho was targeted for cloud gaming, it has range limitations. Future work could seek to enhance Ekho so it can synchronize devices at either end of a very large room, like a concert hall.

"Using inaudible [white noise](#) as a sort of 'timekeeper' is a great example of how out-of-the-box thinking can produce unexpected results," says Alizadeh. "The technique could improve user experience, not just in cloud gaming but potentially in any multidevice streaming scenario."

More information: Ekho: Synchronizing Cloud Gaming Media Across Multiple Endpoints: [www.microsoft.com/en-us/research ... gcomm23-final146.pdf](https://www.microsoft.com/en-us/research/publications/gcomm23-final146.pdf)

This story is republished courtesy of MIT News (web.mit.edu/newsoffice/), a popular site that covers news about MIT research, innovation and teaching.

Provided by Massachusetts Institute of Technology

Citation: A system to keep cloud-based gamers in sync (2023, August 31) retrieved 28 April 2024 from <https://techxplore.com/news/2023-08-cloud-based-gamers-sync.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.