

PETE STIRLING: So there's a few different approaches to this kind of a problem, but the biggest challenge here is the computational complexity, so one way to overcome that is to precompute everything, so you can lay out a grid of points in the world and then compute the audio pressure at all of those points, and also the relationship between the points, the visibility and all of those factors, to work out how the sound would propagate through the space, and then just sort of use the nearest one or interpolate. And there's a few drawbacks to going this direction. So one of the main ones is the large offline compute time. It's going to take a long time to bake out that data, and all of that data is going to take up a lot of memory, so you're going to have to load it all up and keep it resident in memory, so that's another drawback here, and I think one of the biggest ones – and this comes from the large compute time – that makes it a really long iteration time, so when you're trying to dial it in and get all that finetuning, you don't want to wait half-an-hour or an hour, however long it takes, to bake it all out, to hear it again. You're going to forget what you were listening for, and you're not going to have a smooth workflow working that way. And then the technical limitation here is its limit to static environments, so if you've got a door in the space, and you've baked out the reverb with that door closed, in game when that door opens, the sound is not actually going to pass through that door, or if it does, you're going to have to work around some limitations of the pre-baking. So we wanted to find a better way to approach this, so we set up some requirements for the solution we were going to develop. So firstly, the geometry needs to be able to move dynamically. You need to be able to open doors, and that changes how the reverb, the propagation, and the inclusion of the whole scene sounds. It needs to be high quality. We know sound designers have very particular tastes, they're not going to settle for a compromise on quality to get a little bit of accuracy. We need to be maintaining a really high quality bar for the audio, and we need to give them control over how it sounds. We need to give them the right knobs and dials, so they can dial it in to sound how they want the space to sound, and give them control, and you have to be able to play a lot of sounds. It's computationally complex, but people aren't going to settle for just having half-a-dozen sounds. You still need to be able to play 64 sounds or more in your scene at the same time. And so Facebook Reality Labs took these pretty difficult constraints that we put on them and worked on the problem, and they came up with something really special. I've got a demo of what they put together right here. So this scene, you can see it's a concrete space, it's open at the top, and there's 21 sources playing simultaneously in this. They're all stems of a different song, and they're all moving around dynamically, and throughout the video, the geometry of the space also changes, so I'll play that for you now. [video] You have a small room. It is quite small. And then it moves. [video] You can hear how the space opens up, as that wall disappears off into the distance, the reverb gets bigger, and the whole space sounds like you're in that space, and this is really, really amazing. When you're in there and you hear it, it really reinforces the visuals, and you really get to appreciate how important getting the reverb correct is, and so we've taken the technology they've developed and we're working on integrating it into the audio SDK right now. We've got a little progress video to show you, this far from final, but this is that same system in Oculus Home. [video] And excuse the mess there. And so this is a really interesting environment to model because it's not entirely an indoor environment, because it's got that large open wall, but it's not exactly an outdoor environment either, so trying to do the reverb for this correctly with impulse responses or a parametric reverb can be very difficult, and it's highly asymmetric.

[00:05:00]

You've got this big wall behind you, and so modeling that correctly with existing technology would be very, very difficult, if possibly impossible, so we really thought this was a great place to dog-food this tech, and dog-fooding this tech was really important for us because we wanted to make sure that what we were delivering to you would meet Tom's high quality standards, as well as the performance targets of hitting frame rate and all of that, and not compromising in any way.

TOM SMURDON: One thing Pete's not telling you is how fantastic this sounds in headphones. It's hard to tell here through speakers, just how spatialized everything is, so when you're standing in this room, if you've been in Home before, you know like right there, we're standing back against the wall, as things are bouncing around the room, you hear the thing slap back off the wall behind you. It sounds kind of odd in here, but when it's in the headphones, everything sounds like it's right. You hear the slap back behind you. You hear all the long trails, like off in front of you, just like the reverb feels more in front of you. As you turn to the left, then you still hear the wall, like you will know when you're next to a wall without even seeing anything. You can feel everything, it's amazing. This video, Pete was just up in Seattle from Australia last week, and he was over at Facebook Reality Labs, where they've been working on this, he brought the video back, he was like, "That has to go in the talk." As soon as I heard it, I sat there and listened to it like over and over again in the studio, just trying to get like, oh, this is fantastic, and like I've just seen the video, I know it needs some more work before it goes in, but I'm super enthused and happy about where they're at right now.

PETE STIRLING: Yeah, and this is really interesting, and I think one thing we've learned from all of this is how it changes how we deal with reverb. Reverb itself has been worked on for a very long time, we've come a long way with reverb technology, and making a pleasing-sounding reverb is relatively straightforward and it's a well-explored area, but making an accurate reverb is completely new territory, and trying to get an accurate reverb that's also pleasing is a whole new challenge.