Hearing Tests

Dirk Noy & Gabriel Hauser: The WSDG Acoustic Lab

SAM INGLIS

The advice of professional acousticians can be hard to evaluate on paper — but what if you could actually hear the results of their work before building starts?

"If you have a recording space that is 100 square metres and four metres high, and you tell the client 'You're gonna have a 0.7 second reverberation time, is that OK for you?', the client usually says 'I don't know. How does it sound?' Because it's hard to relate these kinds of numbers to a sound, even for some professionals. They can say 'I like that room, that sounds great.' But is it 0.6? Is it 1 second? It makes it a lot easier to actually listen to it."

Gabriel Hauser is explaining the motivation behind the Acoustic Lab, an innovative tool developed by the acousticians and studio design team at the Walters-Storyk Design Group. The idea is that clients no longer have to make choices on the basis of visual mock-ups and dry data. Instead, they can actually hear what WSDG's work will achieve, and audition different choices that are better made at the design stage than after the fact.

"Let's say you have a client here and you want to discuss materials for the ceiling treatment, for example," continues WSDG's Dirk Noy. "You have a choice between the expensive ceiling cloud, which is super reflective or super absorptive, a medium-priced version, and a cheap version that the carpenter from around the corner could make. You can have the client make a decision based on numbers and graphs, but you can also have them listen to it. We can model all the three cases, and we can have an intelligent conversation about it. The thinking behind it is that we can have the dialogue with people that have no acoustical background, and don't know the terminology, but they still know how to listen."

Real & Virtual Models

At present, there are three Acoustic Labs, located in New York, Berlin and Basel. I visited Gabriel and Dirk at the Swiss site, to experience the lab for myself and find out more about how it works.

Acoustical simulatons of one sort or another have actually been around for a surprisingly long time. From the 1930s onwards, the acoustic properties of a building would be tested by constructing a miniature version of that building in scales of 20:1, 10:1 or even 8:1. High-frequency test signals would then be fed into these models. "People built little scale models and took measurements in them," says Dirk. "It had a tiny source, usually a spark generator; you had little receivers with microphones and then you could measure and then scale the measurement result and learn something about the absorption characteristics, reverb times *et cetera* of that space. And it's still being done for bigger projects."

There are, however, obvious drawbacks to this approach. For one thing, the need to scale the results down by a factor of 10 or 12 means that you need to generate and record test signals that extend to 30 to 40 kHz, and even then, you won't learn much about how the real building will behave at frequencies higher than 3 or 4 kHz. For another, it's expensive and time consuming, as Dirk relates: "Unfortunately, when you have to change the front of the balcony railing or something, you have to actually bring in real stuff to make the change, or perforate it for real. Then you find, 'Oh, no, I made too many holes and I have to redo the whole thing,' so another three days are lost! So this is a nice approach, but very few projects actually support this type of procedure."

On all but the largest projects, physical scale models have now been replaced by virtual counterparts. Of course, architects too use CAD tools to model their structures; in theory, these architectural models could

WSDG | ARCHITECTURAL ACOUSTIC CONSULTING | MEDIA SYSTEMS ENGINEERING

WSDG use a program called EASE to generate 3D models of building interiors.

Dirk Noy (top) is the founder of WSDG's Swiss operation; Gabriel Hauser (above) joined soon after it was set up, and both have been involved in hundreds of projects since.

be used as a basis for acoustic calculations, but WSDG almost always prefer to create their own. "You can take an AutoCAD drawing and put it into the acoustical simulation environment, but the level of detail is too high in the architectural drawing," says Dirk. "There might be a door handle that you don't really want to have in your acoustical model because it will increase the calculation time tenfold without giving any additional information. So we usually make the models ourselves."

"Creating the model is an acoustician's work, really," agrees Gabriel. "For example, there could be a screen that I don't need to include in an acoustic model because it's acoustically transparent, it's just a visual screen. Or if a column is really small, say 10 centimetres in diameter, I just leave it out, because for the wavelengths that we're looking at, it's too small. So creating the model is the first step that already needs acoustic knowledge. That's why we don't use architectural models, except as a baseline."

Conversely, an acoustic model can also include data that might not be part of a standard architectural model, most obviously the reflective and absorptive properties of the materials involved.

Rays, Cones & Pyramids

Once all the acoustically significant surfaces have been incorporated into the model, it can be used to generate not only a visual but also an auditory simulation of the space. Key to this is the use of a technique derived from ray-tracing. The user specifies where in the room a hypothetical source (such as a musician or a PA speaker) should be placed, along with the location and directivity of a virtual multi-channel microphone array capturing that source. WSDG's software will then project paths extending outwards in all directions from the source. If these do not reach the receiver either directly or after a given number of reflections from walls and other surfaces, they are disregarded. If they do reach the receiver, they are adjusted for the absorption coefficients of the materials they've encountered along the way, and an impulse response is generated. Eventually, all of these responses are summed to create a single impulse response that represents the complete sound captured by that virtual microphone for the chosen source.

"The ray-tracing or cone-tracing or pyramid-tracing algorithms radiate energy from the source in a certain room area and see what happens with it, if it bounces back towards the receiver or not," explains Gabriel. "So they actually follow cones of sound, so to speak, and if one of those hits the receiver, you start to narrow it down and look at that cone a bit more precisely."

When suitable impulse responses have been obtained, these can then be convolved onto test tones for measurement — but they can also be applied to anechoic recordings, to create an auralisation that recreates the experience of standing in the as-yet-unbuilt concert »

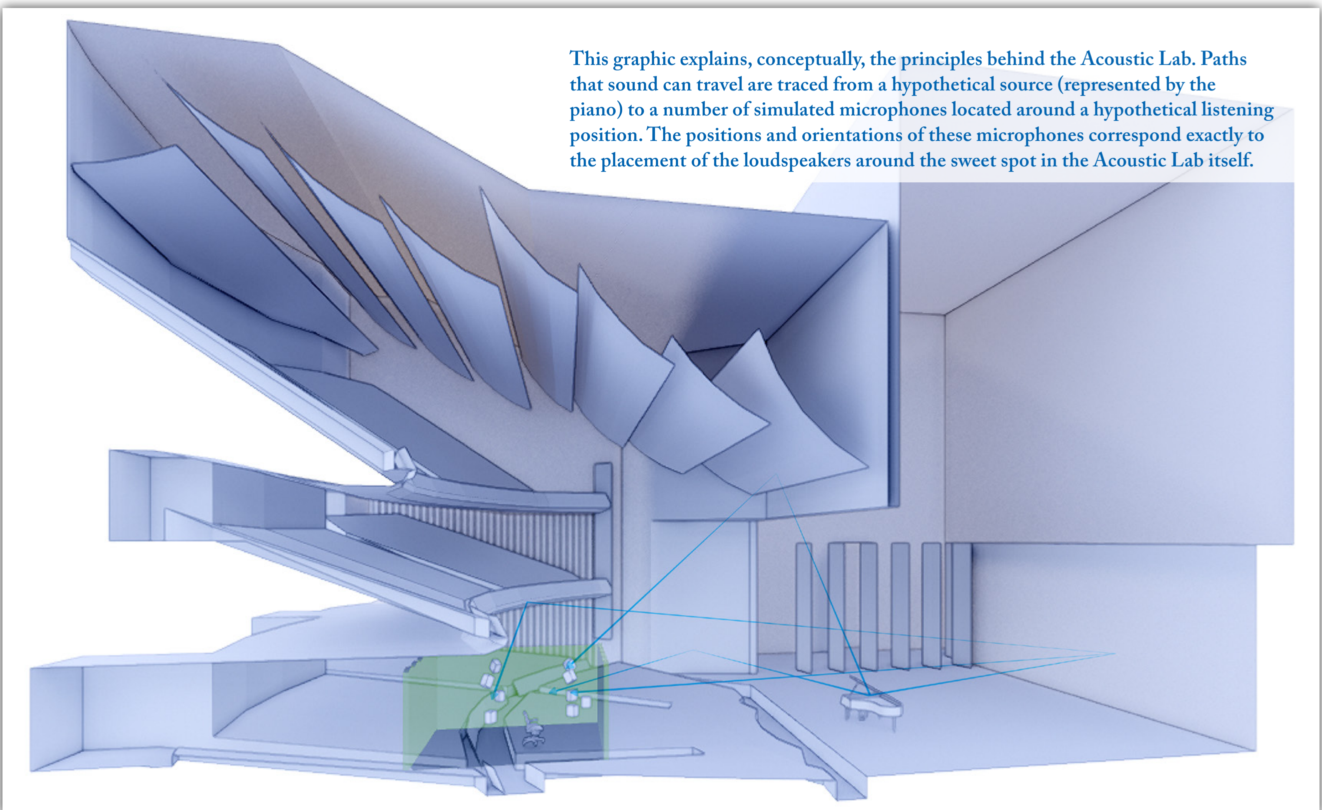

This graphic explains, conceptually, the principles behind the Acoustic Lab. Paths that sound can travel are traced from a hypothetical source (represented by the piano) to a number of simulated microphones located around a hypothetical listening position. The positions and orientations of these microphones correspond exactly to the placement of the loudspeakers around the sweet spot in the Acoustic Lab itself.

» hall, live room or railway station. The directivity, position and number of the virtual microphones is directly related to the intended playback system. For example, the Acoustic Lab in WSDG's Berlin office uses an Ambisonics speaker array to recreate the soundfield at the virtual listening position. Modelling a space for playback on a first-order Ambisonics rig would involve creating impulse responses for three orthogonal figure-8 microphones and one omni, all positioned at the same point in space. In theory, you could also create a binaural auralisation for headphone listening, by mimicking what would be captured on a pair of in-ear mics.

At the Basel Acoustic Lab, however, auralisations are played back over a spaced array of nine loudspeakers: five are arranged in a conventional surround setup, with the remaining four above them to provide an additional height dimension. "There's this format called Auro-3D, where you have the standard 5.1 setup and then you have a second layer that is higher up," explains Gabriel. "We installed that in the Vienna Symphonic Library Synchrostage control room in Vienna and I listened to some recordings there that they did. We switched on and off this height channel, and it was like the same kind of experience as switching off the surrounds in the 5.1 setup. It was like everything was getting flat. And it's the same kind of dimensional loss that you hear. These four channels really added to the spaciousness of what you perceived. So this experience, and the fact that our acoustic model can provide 5.1 impulse response sets led to the decision to try this kind of setup in our lab."

The position and directivity of each speaker in the lab needs to correspond exactly to that of the virtual microphones in the model, so it's necessary to generate a separate impulse response for each. "The impulse response that is calculated is not only weighted with the absorption coefficient of the surface but also with the microphone directivity," explains Dirk. "So that means if sound is being reflected at the rear of the auditorium and comes back from here, it has full level for the surround microphone [*corresponding to, say, the left surround speaker in a 5.0 array*]. But if it hits this microphone here [*he indicates the front left loudspeaker in the array*], it gets attenuated via microphone directivity pattern, so that's kind of an additional attenuation for each of those five channels."

Compared with the Basel setup, Dirk and Gabriel describe the Berlin Ambisonics system as "more precise, but less tangible". "With B-format, you have an auralisation that is basically at one point," says Gabriel. "The microphone is at a single point, so there's no time difference between left, centre and right channels as there is in 5.1. This time difference gives you a bit more added spaciousness, I would say."

Translating these impulse responses into something the client can hear is straightforward: they are simply loaded into the convolver plug-in in Magix's Samplitude and applied to the source material. Dirk: "Every channel has its own room simulator with the impulse response that corresponds to whatever direction the speaker's at."

Sight & Sound

Ideally, perhaps, the Acoustic Lab itself would be an anechoic chamber, to ensure that real room reflections don't get intermingled with the virtual ones. In practice, however, it's located in a well-treated and reasonably spacious room on the ground floor of the WSDG offices. Listening to auralisations in the Acoustic Lab can be quite an uncanny experience: the relatively subdued acoustic signature of the lab itself disappears, and the virtual space takes over. It also reinforces the extent to which our senses work together. One of WSDG's example projects is a remodelled concert hall at the local music conservatoire. This is quite a lively space, and without any visual cues to go on, an anechoic piano »

» recording heard from the perspective of a virtual listener sounds reverberant and splashy rather than truly immersive. It's hard to sustain a sense of 'being there' when you have only to open your eyes and see that the far wall is not 15 metres away!

However, the same auralisation produces a completely different psychological reaction when you listen while viewing a 3D model of the space on an Oculus Rift VR headset. The newfound integration of auditory and visual data suddenly places you in the room with utter believability; as long as you remain in the sweet spot, you can turn your head, look up, look down; and what had previously sounded like a bad recording of a piano now convincingly comes across as a real acoustical event taking place in the hall where you're seated. Even a relatively crude 3D mock-up of a railway station has the same transformatory effect on the listening experience.

This audio-visual experience makes very clear the value of the Acoustic Lab as a tool for bringing projects to life, and presenting choices to clients in a way that allows them to make subjective judgements. "If the clients say 'That's a neat gimmick. What do we learn from it?' then it failed," says Gabriel. "But if the clients say 'Ah, OK, now I understand what you mean when you say 0.56 STI is better than 0.51,' that will be the perfect scenario."

Size Isn't Everything

Client feedback so far has been positive, but it's early days, and there are still some limitations on the realism with which WSDG can reproduce an acoustic environment. Partly this is down to simplifications in the models themselves. "For example, in the acoustic program that we are using, you enter absorption data in octave bands," says Gabriel. "This, of course, is very rough, because between 1kHz and 2kHz a lot can happen. Just having two different kinds of absorption coefficients for these octave frequencies is a rough approximation." Likewise, models currently assume that absorption is constant regardless of angle of incidence, which is not the case with all real materials.

To generate the auralisations heard in the Acoustic Lab, a single anechoic recording is separately convolved with the impulse response for each virtual microphone in the model.

From the point of view of studio design, however, the most significant limitation is probably to do with room size. In a football stadium or railway station, room modes are so low in frequency as to fall well out of the audible range, so there is no need for them to be simulated. In a small control room, by contrast, managing room modes is perhaps the most important and challenging part of the acoustician's job; but WSDG's impulse responses cannot capture this aspect of a room's behaviour. "This kind of algorithm is purely ray-based, using geometrical acoustics," says Gabriel. "It's not a wave-based acoustics simulation program, which means that as soon as we come into the region of eigenmodes instead of a statistical soundfield, the model is not accurate any more. In smaller rooms with a Schroeder frequency of 150Hz, this program is only useful at 250 or 500 Hz upwards, and the low frequencies are not very accurate."

Wave Behaviour

WSDG do, of course, model the low-frequency properties of control rooms, but they use different software to do so. One challenge for the future is to integrate these separate tools, as Gabriel acknowledges. "The quality of a control room is, I would say, 60 to 80 percent about the low end. And that's exactly what we cannot, at this moment, simulate with this program, because it uses geometric acoustics and not wave-based acoustics. But we are also using a wave-based program exactly for calculation of control rooms, low-frequency behaviour, modal response, to know where we have to put what kind of treatment, to make this room sound great. And this program is actually capable of outputting an impulse response, but only sensibly up to the Schroeder frequency — so we could have some kind of crossover frequency where we use the simulation from program A and then switch over to the simulation of program B, and then get a comprehensive picture of what's going on."

"The challenge of the tools, really, is combining geometrical acoustics and wave acoustics," agrees Dirk. "Geometrical acoustics are good for mid and high frequencies, bad for low frequencies. And wave acoustics is good for low frequencies but too much calculation — like really much — for high frequencies. So people need to kind of combine those two in one simulation world and then have a kind of a filter or crossover to manage the results for it."

Recreating a small control room within the Acoustic Lab would also require its own acoustics to be further optimised, as Gabriel explains. "If the reverberation time of the project is low, we run into problems with our demo room, which is not anechoic. It's just pretty much controlled. So I would say the reverberation time of the project needs to be at least twice what we have here, which means 0.4 seconds [*RT60*]: so that's medium-sized or large-sized control rooms, medium-sized recording rooms, and higher. But I think we're moving into the right direction with the tools that are available."

In time, most of these challenges could be overcome, and WSDG are already working on some of them. Others are possible in theory, but don't offer enough commercial benefit to justify the enormous investment in time and research. "We are not primarily a research lab!" insists Dirk.


The Acoustic Lab in Basel uses an Auro-3D spaced speaker setup, but other configurations are possible.

"We are not IBM or whatever. We do projects, mainly. This is kind of a support tool to facilitate dialogue between different stakeholders on a project. Sometimes the struggle as an acoustician is to explain to the world what you are doing, because it's very abstract, and the goal really is to create a dialogue enhancer for talking to people about acoustics who have no idea about acoustics."