Immersive Wave Experimentation and *The Matrix*

Dirk-Jan van Manen and Johan O. A. Robertsson

In the 1999 blockbuster motion picture *The Matrix* (see imdb.com/title/tt0133093), in perhaps one of the most memorable exchanges in recent movie history, Morpheus, played by American actor Laurence Fishburne, tells Neo, played by Canadian actor Keanu Reeves, "Unfortunately, no one can be told what The Matrix is. You'll have to see it for yourself."

As Neo and the viewers soon learn, The Matrix turns out to be an elaborate simulation, generated by extremely powerful and sentient computers, that acts as a virtual reality for all the senses and constantly tricks the human mind and body into believing that everything it is seeing, hearing, touching, tasting, and smelling is real. Designed to keep humans enslaved mentally and physically, while exploiting them as batteries, Neo could therefore not be told what The Matrix is because the illusory reality created and maintained by it was so compelling that it could no longer be distinguished from reality itself by anyone trapped inside.

Now, scientists from the Swiss Federal Institute of Technology, Zurich, Switzerland (see ethz.ch/en.html) are making headway on an ambitious and much less nefarious project to build a "Matrix" for acoustic waves. Their goal? To trick acoustic and elastic waves propagating in finite-size laboratories into behaving exactly like they would if they were propagating in much larger, extended media. By doing this, the scientists can study wave physics unimpeded by laboratory boundaries and under fully controlled conditions. To achieve this, the scientists propose lining the boundaries of experimental domains with hundreds of active sources and receivers and using technologies like those applied in noise-canceling headphones to compute and control the wavefield on those boundaries in real time.

The task at hand could hardly be any more daunting; not only do the scientists want to suppress the reflections from the physical boundaries, but they also want to include reflections from user-defined virtual scatterers in the extended medium. Furthermore, they want to make it possible for such virtual reflections to interact with the virtual environment again after interacting with scatterers in the experimental domain, that is, facilitating multiple scattering between real and virtual objects.

The Problem of Scale in Laboratory Wave Propagation Experimentation

To understand the scientists' motivation for trying to construct a Matrix for acoustic waves, it helps to know a bit more about laboratory wave experimentation and seismology. Earth materials support compressional (P) wave speeds of up to approximately 14 km/s. Granted, this maximum is only reached for P waves at the core-mantle boundary (and at pressures beyond most, if not all, laboratories), but because earthquakes produce waves with periods of several tens to hundreds of seconds, wavelengths in the Earth can easily reach hundreds of kilometers.

This gives rise to a genuine chicken and egg situation. To research wave physics at those low frequencies and long wavelengths, scientists need to do laboratory experimentation. But to do laboratory experimentation, they need to scale down the experiment, often by factors of 10,000 or more. To scale down the experiment, however, they need to know the physics over those four orders (or more) of scale. Clearly, an impossible task!

Such problems would be drastically reduced, however, if scientists had access to a Matrix for acoustic waves. Although they would still need to know the physics of the embedding medium, the ability to truncate the medium without boundary reflections opens the door to truly low-frequency, long-wavelength experiments that could reduce the required scale factors by, for example, two orders of magnitude or more. This would therefore

https://doi.org/10.1121/AT.2022.18.4.40

enable laboratory wave physics experiments at frequencies and wavelengths much closer to the physics that scientists are really interested in.

A Matrix for Numerical Modeling: Immersive Boundary Conditions

To further their ambitious goals, scientists looked toward a numerical modeling method that they had developed previously, called immersive boundary conditions (IBCs). IBCs were initially conceived as nonreflecting boundaries in numerical simulations.

Without special treatment, the (free) edge of any computational domain on which the wave equation (WE) is discretized (e.g., gridded to enable the computations) acts as a reflecting boundary for the simulated waves, essentially trapping the waves on the numerical domain and masking reflections of real interest. Note that simply fixing the edge of the computational domain does nothing to solve the problem because rigid boundaries reflect waves equally (albeit with different polarity). IBCs thus prevent boundary reflections by treating the edge specially during such computations.

Because directly solving discretizations of the WE yields some of the most complete numerical solution methods but also the most computationally expensive, scientists have been looking for ways to make the computational domain as small as possible. Thus, it is a major computational advantage to be able to truncate the computational domain without generating boundary reflections.

However, the interactions of the waves with structures in the larger background medium are still often important. Thus, it would be ideal to be able to truncate the domain without generating reflections *and* retain, through the application of the special boundary conditions, the interactions of the waves with the background medium. IBCs were specifically developed for this purpose. They immerse a truncated modeling domain in a background medium, allowing waves to propagate seamlessly between the truncated domain and the background. They thus enable arbitrary-order scattering interactions (van Manen et al., 2007).

Briefly, IBCs work as follows (see **Figure 1**). To start the computation, a source (**Figure 1,** *black star*) generates "incident waves" (**Figure 1,** *curved arrow*) that are injected onto a truncated numerical domain D_{sct} enclosed by injection boundary B. At every time step of the simulation, wavefield values observed a finite distance away from the edge, on the extrapolation surface S, are extrapolated to the edge of the computational domain (i.e., to B) and through the background model D'_{sct}. This extrapolation is achieved by convolving (i.e., filtering) the observed wavefield values with a large number of precomputed impulse responses that fully characterize the background model (including, e.g., the "background scatterer"). These impulse responses model both the direct propagation between the extrapolation surface S and the injection boundary B (**Figure 1,** *short arrow*; denoting "extrapolated waves") as well as the interactions with the background medium (**Figure 1,** *curved arrow*; denoting "long-range interactions").

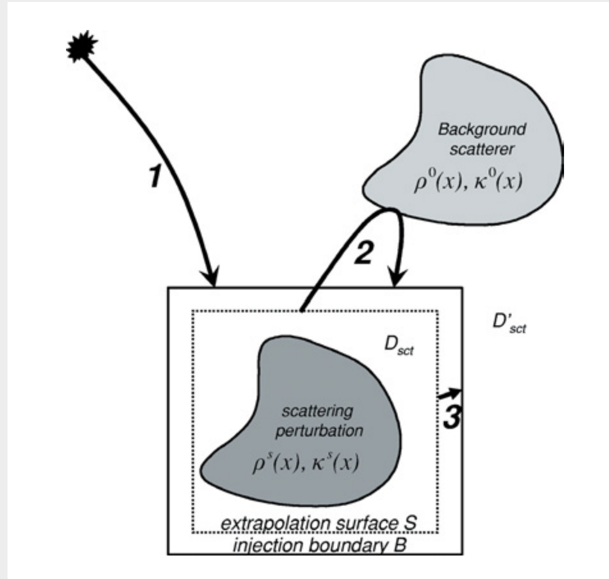

Figure 1. *Immersive boundary conditions (IBCs). At every time step of a numerical simulation, wavefield values recorded a finite distance away from the edge (i.e., on extrapolation surface S) are extrapolated to the edge of the computational domain (i.e., to injection boundary B) as well as through the background model. This allows predicting the outgoing and ingoing waves at the edge very accurately and therefore enables canceling, respectively, injecting those waves. D_{sc}, truncated numerical domain; D'_{sct}, background model. Reproduced from van Manen et al. (2007), with permission of Acoustical Society of America. Copyright 2007, Acoustical Society of America. See text for details.*

This extrapolation allows predicting the waves that arrive at the boundary very accurately. The predicted outgoing waves are then used to cancel the actual outgoing waves at the edge of the computational domain, like the way anti-sound is used in noise-canceling headphones. The predicted ingoing waves, from the interaction with the larger background model, are instead injected and radiate onto the numerical domain, and this is, in fact, what allows reflections from the background model to be retained.

The extrapolation is done for every time step of the simulation. Similarly, the cancellation and injection take place at every time step. Because the computed ingoing waves radiated back onto the truncated numerical domain can be extrapolated again (and again) after a second (third, fourth, ...) interaction with the scatterers, the method is fully recursive, generating all orders of interactions with the background medium. The resulting wavefields match reference wavefields computed on the full model to within numerical precision. Thus, the scattering on the immersed truncated domain cannot be distinguished from the scattering on the full model, essentially constituting a Matrix for the numerical modeling of acoustic waves.

Figure 2. *IBCs.* ***A:*** *snapshot of a reference simulation computed on a geophysical subsurface model for a source at the surface (**red star**).* ***B:*** *snapshot of a simulation computed on a much smaller part of the same model (i.e., **black square**) but truncated using IBCs. Wavefields computed on the full and the truncated model match each other to within numerical precision. Just from observations of the scattering within the **black square**, it is impossible to tell whether the simulation was run on the full or the truncated model. After Broggini et al. (2017), with permission. See text for details.*

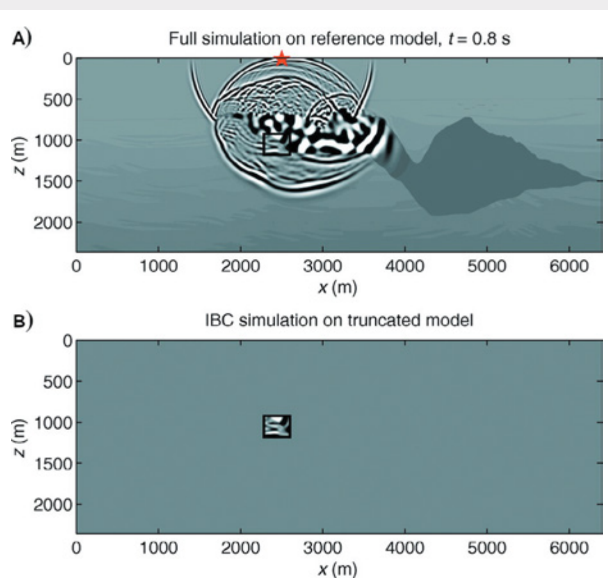

An example of an IBC simulation on a severely truncated geophysical model is shown in **Figure 2**. The truncated model is approximately one-hundredth of the size of the full model, indicating that expensive discretizations of the WE on the full model can be almost completely replaced by convolutions (i.e., filtering operations) with precomputed impulse responses in simulations that employ IBCs.

Beyond Numerical Modeling: Immersive Wave Experimentation

A breakthrough came in 2013 with the realization that IBCs can also serve to truncate *physical* wave propagation experiments without boundary reflections and to immerse *physical* experiments in virtual numerical environments in real time (Vasmel et al., 2013). In **Figure 1**, this entails replacing the truncated numerical domain D_{sct} with a *physical* experimentation domain. Furthermore, the extrapolation surface S and the injection boundary B now need physical implementations. These are needed both to measure and cancel, respectively, the physical outgoing waves and to inject and radiate, respectively, the virtual ingoing waves back onto the experimental domain.

With these changes, using IBCs to immerse physical experiments in real time, so-called immersive wave experimentation (IWE), is, in principle, feasible. However, a major challenge still remains, namely the extrapolation of the wavefield values to the boundary of the experimental domain as well as through the numerical environment, including any analog-to-digital and digital-to-analog conversions, must be completed *before* the waves arrive there. That is, the measurement, extrapolation, and control must be *faster* than the waves propagate in the physical experiment. They must be completed faster than the speed of sound!

The requirement for these substantial and real-time computations poses severe latency requirements to any control system implementing such an extrapolation. For example, if the recording surface and the physical boundary are separated by 30 cm and the background medium is water (with a sound speed of 1,500 m/s), then the total system latency cannot exceed 200 µs (i.e., one-fifth of a thousandth of a second). Furthermore, because the wavefield at any point on the recording surface can affect the wavefield at any point on the physical boundary and at any time in the future, the extrapolation constitutes a very computationally intensive task, requiring significant computational resources.

A key enabler for immersive wave experimentation therefore has been a low-latency, acquisition, compute, and control system implementing the IBCs developed in close collaboration with NI (formerly National Instruments) Switzerland. Utilizing so-called field-programmable gate arrays (FPGAs) for the extrapolation computation, the 500-FPGA strong system connects 800 simultaneous inputs to 800 simultaneous outputs through 640,000 individual impulse responses. It does so with a fixed, ultralow system latency of 200 μs, realizing the minimum extrapolation distance of 30 cm for a background medium of water.

Two-Dimensional Immersive Wave Experimentation: Turning a Circular into a Square Domain

Using this ultralow latency system, the first successful two-dimensional (2D) immersive experiments were performed by Becker et al. (2020). This involved immersing a circular physical domain with rigid boundaries into a slightly larger virtual square numerical domain in real time. For arbitrary sound sources in the circular physical experiment, the resulting wavefields, including the tails of multiple-scattered energy in the recordings (i.e., the coda) were virtually indistinguishable from corresponding reference fields modeled for the square domain directly (see **Figure 3**).

Figure 3, A and **B,** shows *experimentally* measured wavefield values (in mV) without and with, respectively, active IBCs recorded in the *circular* physical domain bounded by rigid boundaries. **Figure 3, C** and **D**, shows *numerically* modeled wavefield values (in mV), modeled for the *circular* physical domain and for the slightly larger *square* domain, respectively, bounded by rigid boundaries. **Figure 3, C** and **D**, serves as references for the results in **Figure 3, A** and **B**, respectively. Thus, if the immersion is successful, then the wavefield in **Figure 3B** should look like the wavefield modeled for the square domain in **Figure 3D**, even though it is recorded in the circular physical domain!

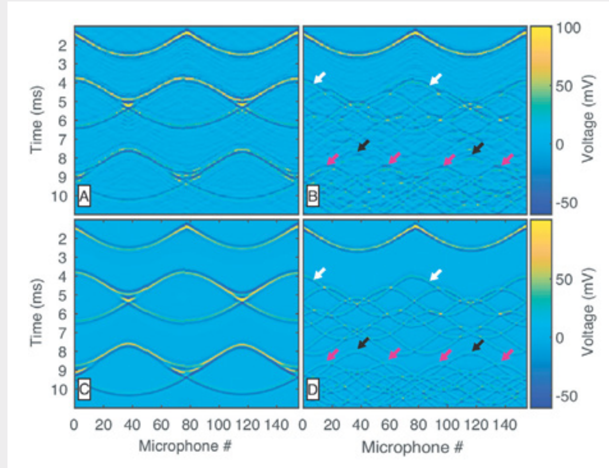

Figure 3. *Two-dimensional (2D) acoustic immersive wave experimentation (IWE). A circular experimentation domain with rigid boundaries is immersed in a slightly larger square numerical domain in real time.* ***A:*** *wavefield recorded in the circular physical domain when the IBC is off.* ***B:*** *wavefield recorded in the circular domain when the IBC is on.* ***C:*** *modeled reference wavefield for the circular domain.* ***D:*** *modeled reference wavefield for the square domain. The close match between the fields in* ***B*** *and* ***D*** *confirms the successful immersion of the circular domain. From Becker (2020), with permission. See text for details.*

Comparing **Figure 3B** with **Figure 3D**, it is clear that the immersed response indeed closely matches the modeled response for the virtual square domain and that the imprint of the underlying circular physical domain has been successfully removed. In **Figure 3, B** and **D,** the *white arrows* denote the first arriving virtual reflection from the bottom boundary of the virtual square domain (introduced onto the physical domain by the IBC). In **Figure 3, B** and **D,** *other arrows* denote waves that have interacted with the virtual domain at least twice, with the *black arrows* denoting the bottom and top and the *pink arrows* denoting the left and right sides. Finally, taking advantage of the fact that the virtual medium can be any desired medium, including a nonphysical one, Becker et al. (2020) showed that it is possible to immerse a (circular) physical domain into a (square) region of nonphysical energy gain materials.

More recently, it has been shown that IBCs can also be used to make objects appear and disappear in real time acoustically. Following up on the theoretical and numerical work by van Manen et al. (2015), Becker et al. (2021)

demonstrated how to cloak a rigid object experimentally for arbitrary, unknown incident wavefields (including for broadband wavefields generated by moving sources).

Next Generation Virtual Acoustics

Encouraged by the successes of 2D IWE, recent studies have started exploring three-dimensional applications of IWE, including on the human scale and in the audible frequency range. Virtual acoustics (VA) is the acoustic counterpart of virtual reality (VR) (see Vorländer, 2020, for an introduction). Compared with existing VA approaches such as wavefield synthesis (Berkhout, 1988; Ahrens, 2012) and higher order ambisonics (Zotter and Frank, 2019), IBC-based VA environments hold out the promise of true, two-way shared acoustic experiences. At least three different types of IBC-based VA environments are envisaged.

Immersed Acoustic Spaces

Here, a single, physical space is equipped with IBCs. The aim is the highest possible realism in terms of placing one or more people and sound sources in that space jointly in a modeled or measured virtual acoustic environment. A number of applications for such immersed acoustic spaces are being considered. One example is use a "show" room where clients can listen to different architectural designs and hear the acoustic impact of using different materials before a building is constructed. Another example are concert halls with reconfigurable acoustic responses.

With IBCs, a simple rectangular room can, in principle, have the perfect response for any type of music. Finally, such spaces could also be used for mobility training of visually impaired people. Some visually impaired people use clicking sounds to navigate acoustically (see bit.ly/3RhpkzY for an example). Immersed acoustic spaces could provide safe, yet acoustically realistic, environments where visually impaired people can learn to navigate.

A 2D example of an immersed acoustic space is shown in a top view in **Figure 4**. In this and the following examples, a floor plan of Amiens Cathedral, Amiens, France, has been used to model 2D acoustic impulse responses that completely define the virtual environment.

Figure 4, *top left*, shows a snapshot of a wavefield propagating in a physical space of 3 × 4.5 m while that space is immersed in the larger virtual environment of 28 × 12 m using IBCs. **Figure 4,** *bottom*, shows a snapshot of a wavefield directly computed for the union of the physical and the virtual environments (i.e., the full model) for reference. **Figure 4,** *top right*, shows a zoom in of the reference wavefield. Note the absence of reflections from the rigid walls of the physical space (**Figure 4,** *top left*) because they are canceled by the IBC. Furthermore, it is clear that the field in the immersed space is identical to the reference field, as expected. In **Figure 4,** the *black shapes* denote the real and virtual scatterers (e.g., internal walls and pillars). In **Figure 4,** the *red rectangles* denote the locations of the rigid walls of the physical space and the *blue rectangles* denote the surfaces from which waves are extrapolated.

Shared Acoustic Spaces

Here, two physical spaces with reflecting boundaries are equipped with IBCs, but the IBCs only cover one wall in

Figure 4. *Numerical example of an "immersed acoustic space." **Top left:** a physical space of 3 × 4.5 m (**red rectangle**) is immersed in a virtual space of 28 × 12 m (i.e., a 2D acoustic model of Amiens Cathedral, Amiens, France) using IBCs. Note the absence of boundary reflections in the resulting immersed acoustic space. **Bottom:** reference field directly computed for a full model consisting of the union of the physical and the virtual space. **Top right:** zoom in of the reference field. **Solid black shapes:** scatterers; **white dot:** a (human) source. Acoustic model created from a floor plan by Gothika. See text for details.*

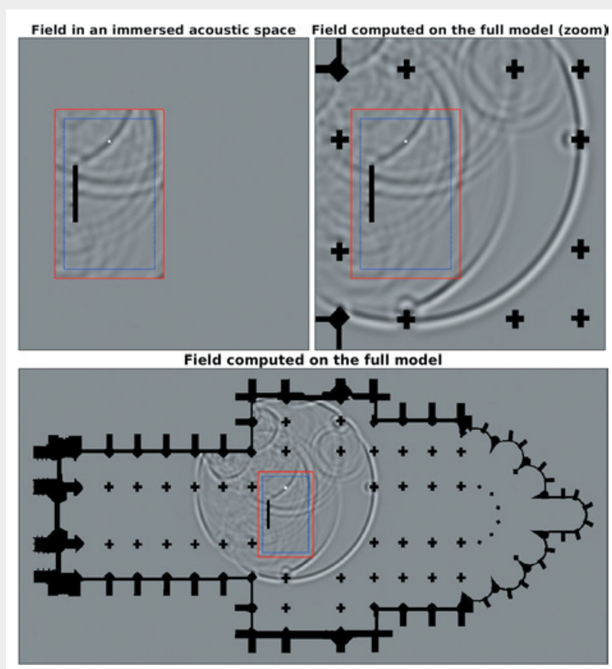

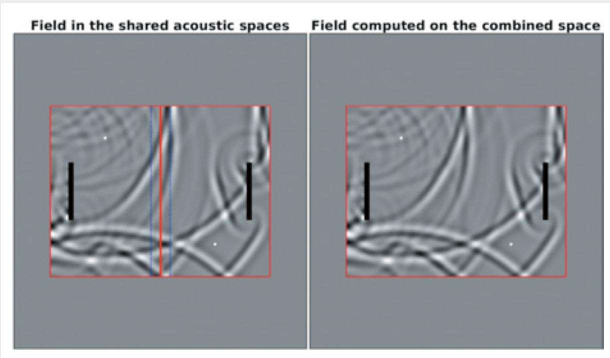

Figure 5. *Numerical example of a "shared acoustic space." **Left:** a physical space of 3 × 4.5 m (**left red rectangle**) is connected to another physical space of 3 × 4.5 m (**right red rectangle**) using IBCs along the adjacent faces to create a shared acoustic space of 6 × 4.5 m. Note the absence of boundary reflections along the adjacent faces. **Right:** reference field directly computed for the combined space of 6 × 4.5 m. The union of the fields in the shared acoustic spaces is identical to the reference field. **Solid black shapes:** scatterers; **white dots:** a (human) source and receiver; **thin blue lines:** recording surfaces from which the field is extrapolated. See text for details.*

each space. Furthermore, the impulse responses used in the extrapolation do not contain the acoustic imprint of Amiens Cathedral (or any other virtual medium) but contain only the direct waves required to predict the outgoing waves on the one wall in each space. In this configuration, the focus is *only* on the level of realism achieved in terms of acoustic interactions between people in the different physical spaces and *not* on changing the acoustic characteristics of the physical spaces themselves. Sound propagates correctly between people in the resulting shared space. The resulting room characteristics are the characteristics of the two physical spaces combined. This gives the correct sense of "where people are" and ensures that people know "where to look" in response to acoustic cues (i.e., when complementing sound by matching images).

Example applications of such shared acoustic spaces include virtual meeting rooms and conference venues where a closer interaction and sensory experience between multiple participants in different physical venues are needed. Shared acoustic spaces provide this because they place each participant in the acoustically correct position in the shared space. Another application envisaged are virtual musical performances, bringing remote musicians together acoustically. In this case, the physical rooms used should have reasonable acoustics to begin with because they are not modified but only combined in this configuration.

An example of a shared acoustic space is shown in **Figure 5**. In **Figure 5,** *left*, snapshots of the wavefields propagating in two separate physical spaces of 3 × 4.5 m are shown adjacent to each other while those spaces are acoustically connected to each other using IBCs. The IBCs are only implemented along the right wall of the left physical space and the left wall of the right physical space. Note the absence of reflections from the walls along which the IBCs are implemented. That is, wave fronts intersecting the walls along which the IBCs are implemented are continuous. In **Figure 5,** *right*, a snapshot of the wavefield directly computed for an acoustic space of 6 × 4.5 m (i.e., the size of the combined space) is shown. Once again, the resulting wavefields in the physical spaces match the reference wavefield perfectly.

Shared Immersed Acoustic Spaces

Here, the focus is both on getting the sound to propagate correctly between people in different physical spaces and on immersing the two (or more) physical spaces in a single larger virtual environment. This requires IBCs that completely surround the two physical spaces. As a result, the existing room characteristics of the two spaces are completely suppressed in favor of bringing the two parties together in one larger virtual acoustic environment (i.e., Amiens Cathedral in these examples). This is the most technologically advanced as well as the costliest setup. Example applications of shared immersed acoustic spaces include virtual recording studios for music professionals and concert halls spanning multiple physical venues.

Shared immersed acoustic spaces are illustrated in **Figure 6**. In **Figure 6,** *top left*, the resulting wavefields in the two immersed physical spaces are shown adjacent to each other. In **Figure 6,** *bottom*, a directly computed reference wavefield is shown. In **Figure 6,** *top right*, a zoom in of the reference wavefield is shown. As can be seen, the wavefields, including all higher order interactions between the physical scatterers (e.g., **Figure 6,** *black shapes*, in the physical spaces) and the virtual scatterers (e.g., **Figure 6,** *black shapes*, in the model of Amiens Cathedral), match perfectly.

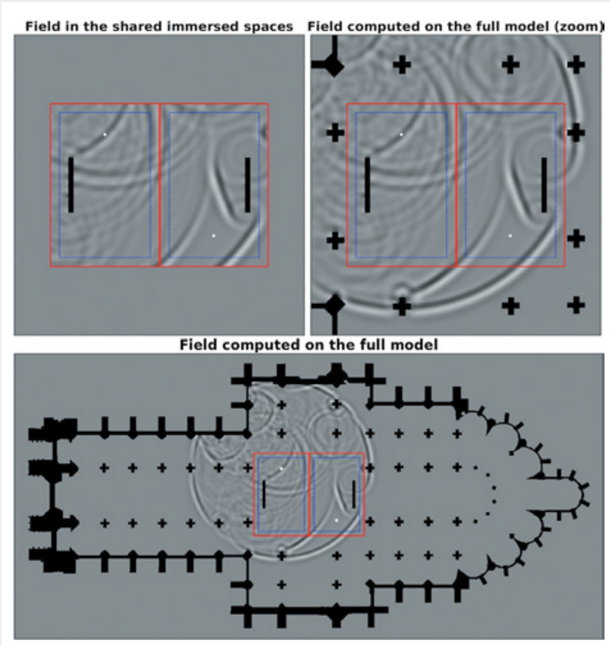

Figure 6. *Numerical example of a "shared immersed acoustic space." **Top left:** two physical spaces of 3 × 4.5 m (**left** and **right red rectangles**) are simultaneously immersed in a virtual space of 28 × 12 m (i.e., a 2D acoustic model of Amiens Cathedral) using IBCs. Note the absence of boundary reflections in the shared immersed spaces. **Bottom:** reference field directly computed for the full model consisting of the union of the physical and virtual spaces. **Top right:** zoom in of the reference field. The field in the shared immersed spaces is identical to the reference field computed for the full model directly. See text for details.*

Can One Hear the Shape of a Room?

In 1966, Polish American mathematician Mark Kac (1966) wrote an article with the now famous title *Can One Hear the Shape of a Drum?* (see also Scott and Morrison, 2022). The frequencies at which a drumhead vibrates depend on its shape. In the article, Kac (1966) asks the question whether the shape can be uniquely determined if the frequencies are known. It took until the early 1990s for three mathematicians, Gordon, Webb, and Wolpert, to show that there indeed exist multiple shapes, nontrivially related to one another, that vibrate with exactly the same frequencies (thus finally answering Kac's question with a clear no) (Gordon et al., 1992).

More recently, Dokmanic et al. (2011), in an article titled *Can One Hear the Shape of a Room: The 2D Polygonal Case*, considered the somewhat related problem of estimating the 2D room geometry from a single acoustic room impulse response (RIR). They discussed the uniqueness of the mapping between the geometry of a planar polygonal room and a single RIR and proposed an algorithm that performs such room shape estimation "blindfolded."

The successful application of 2D acoustic IWE by Becker et al. (2020) by turning a circular wave propagation domain into a square wave propagation domain in real time and the ongoing efforts to apply such technologies to the next generation virtual acoustics force us to reconsider and give new answers to such old questions. Although it may be possible to estimate the shape of the room from a single acoustic RIR, how do we actually know that this RIR was obtained in a physical room with that exact shape and not in a physical room with a different shape that was immersed in another virtual environment using IBCs or IWE? Indeed, when immersing a physical space in a virtual space using IBCs, not just a single RIR but all RIRs are consistent with the virtual space. And thus, even more sophisticated techniques than considered by Dokmanic et al. (2011) that we haven't thought of are all bound to fail. Like Neo, in the famous movie, one would have to find a way to unplug oneself from such a Matrix for acoustic waves to know for sure.

Combining Matrices for Sight and Sound

It is no exaggeration to say that we are at the beginning of another digital transformation of society. The speed and complexity of cost-effective computers, sensors, and connectivity have reached a level that enables immersing individuals in virtual environments and is already being pioneered by companies such as Google and Meta. When exactly we will reach the level of realism portrayed in the movie *The Matrix* is hard to say because the illusion has to work for all the senses (not only sight and hearing but also touch, smell, and taste). The elegant solution by the writers of *The Matrix* was to just bypass the corresponding sensory organs and instead connect the simulation directly to the human information superhighway that is the brain. Although controversial experiments on animals are being carried out by some groups, it will likely still take several decades before humans would be willing to experiment on a large scale with such direct neural connections. Until that time, combining matrices for sight (VR) and sound (VA) seems like an obvious next step. Then touch and, eventually, smell and taste can be added. The question is probably not whether such a combined Matrix can be good enough to fool all humans all the time (as was the goal in the movie) but whether we may soon reach a point where they are good enough that the mind drops its resistance from time to time and we start forgetting, at least temporarily, that what we are experiencing is not "the real thing." Until that time, look out for exciting applications of immersive wave technology in acoustics!

References

Ahrens, J. (2012). *Analytic Methods of Sound Field Synthesis*. Springer Science & Business Media, Heidelberg, Germany.

Becker, T. S. (2020). *Immersive Wave Experimentation: Linking Physical Laboratories and Virtual Simulations in Real-Time*. Doctoral Dissertation, ETH Zurich, Zurich, Switzerland. https://doi.org/10.3929/ethz-b-000410579.

Becker, T. S., Börsing, N., Haag, T., Bärlocher, C., Donahue, C. M., Curtis, A., Robertsson, J. O., and van Manen, D.-J. (2020). Real-time immersion of physical experiments in virtual wave-physics domains. *Physical Review Applied* 13(6), 064061.

Becker, T. S., van Manen, D.-J., Haag, T., Bärlocher, C., Li, X., Börsing, N., Curtis, A., Serra-Garcia, M., and Robertsson, J. O. (2021). Broadband acoustic invisibility and illusions. *Science Advances* 7(37), eabi9627.

Berkhout, A. J. (1988). A holographic approach to acoustic control. *Journal of the Audio Engineering Society* 36, 977-995.

Broggini, F., Vasmel, M., Robertsson, J. O., and van Manen, D.-J. (2017). Immersive boundary conditions: Theory, implementation, and examples. *Geophysics* 82(3), T97-T110.

Dokmanić, I., Lu, Y. M., and Vetterli, M. (2011). Can one hear the shape of a room: The 2-D polygonal case. In *Proceedings of the 2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Prague, Czech Republic, May 22-27, 2011, pp. 321-324.

Gordon, C., Webb, D. L., and Wolpert, S. (1992). One cannot hear the shape of a drum. *Bulletin of the American Mathematical Society* 27(1), 134-138.

Kac, M. (1966). Can one hear the shape of a drum? *The American Mathematical Monthly* 73(4P2), 1-23.

Scott, E. K. E., and Morrison, A. (2022). The acoustics of the modern jazz drum kit. *Acoustics Today* 18(4), 31-39. https://doi.org/10.1121/AT.2022.18.4.31.

van Manen, D.-J., Robertsson, J. O. A., and Curtis, A. (2007). Exact wave field simulation for finite-volume scattering problems. *The Journal of the Acoustical Society of America* 122(4), EL115-EL121.

van Manen, D.-J., Vasmel, M., Greenhalgh, S., and Robertsson, J. O. A. (2015). Broadband cloaking and holography with exact boundary conditions. *The Journal of the Acoustical Society of America* 137(6), EL415-EL421.

Vasmel, M., Robertsson, J. O., van Manen, D. J., and Curtis, A. (2013). Immersive experimentation in a wave propagation laboratory. *The Journal of the Acoustical Society of America* 134(6), EL492-EL498.

Vorländer, M. (2020). Are virtual sounds real? *Acoustics Today* 16(1), 46-54. https://doi.org/10.1121/AT.2020.16.1.46.

Zotter, F., and Frank, M. (2019). *Ambisonics: A Practical 3D Audio Theory for Recording, Studio Production, Sound Reinforcement, and Virtual Reality*. Springer Nature, Cham, Switzerland. https://doi.org/10.1007/978-3-030-17207-7.

About the Authors

Dirk-Jan van Manen
dirkjan.vanmanen@erdw.ethz.ch

Department of Earth Sciences
Institute of Geophysics
Swiss Federal Institute of Technology
(ETH Zürich)
8092 Zürich, Switzerland

Dirk-Jan van Manen has an MSc in geophysics from Utrecht University, Utrecht, The Netherlands, and a PhD in geophysics from the University of Edinburgh, Edinburgh, United Kingdom. After his Marie Curie Industrial Fellowship at Schlumberger, he worked in their R&D department for eight years as a geophysicist and research manager. Since 2014, he has been a senior researcher and lecturer at ETH Zürich, Zürich, Switzerland. He is a recipient of the Society of Exploration Geophysicists (SEG) J. Clarence Karcher Award (2008) and the European Association of Geoscientists and Engineers (EAGE) Arie van Weelden Award (2007) and a corecipient of the ENI New Frontiers of Hydrocarbons Award (2015). He currently leads the research at the Centre for Immersive Wave Experimentation, ETH Zürich.

Johan O. A. Robertsson
johan.robertsson@erdw.ethz.ch

Department of Earth Sciences
Institute of Geophysics
Swiss Federal Institute of Technology
(ETH Zürich)
8092 Zürich, Switzerland

Johan O. A. Robertsson has an MSc from Uppsala University, Uppsala, Sweden, and a PhD from Rice University, Houston, Texas. After two years as a postdoc at the Swiss Federal Institute of Technology (ETH Zürich), Zürich, Switzerland, he joined Schlumberger in 1996. In his last assignment, he was the research director and scientific advisor at Schlumberger Cambridge Research. Since 2012, he has been a professor at ETH Zürich where he is the chair of applied geophysics. He is a corecipient of an ENI award (2015), a European Research Council (ERC) advanced grant (2016), the European Association of Geoscientists and Engineers (EAGE) Conrad Schlumberger Award (2016), and the EAGE Guido Bonarelli Award (2020). In 2020, he was a Society of Exploration Geophysicists (SEG) honorary lecturer. He is currently head of the Department of Earth Sciences at ETH Zürich.