


FROM THE SOUND UP: REVERSE-ENGINEERING ROOM SHAPES FROM ACOUSTIC SIGNATURES

Willem Boning Arup Acoustics, New York, NY, USA willem.boning@arup.com

Alban Bassuet Tippet Rise Art Center, Fishtail, MT, USA alban.bassuet@tippetrise.org



1 INTRODUCTION

The process of designing a room for music often begins with a rough shape known to perform well acoustically and functionally. The shape is then refined through analysis, simulation and optimization to meet acoustical targets. But what if it were possible to invert the process, starting by composing the room's sound and then generating a matching shape? In this paper, we propose a method for acoustical reverse-engineering. We describe techniques that allow the designer to construct a sound signature characterizing the room's acoustics and use that signature to generate a solution space of corresponding room geometries.

Acoustical reverse-engineering can be considered a design-oriented inverse problem. When the parameters that define a model or system are unknown (like the shape of a room), an inverse problem can be used to reconstruct those parameters from the system's output data (like acoustical measurements). Perhaps the best-known example of an inverse problem related to acoustics is the drum shape problem posed by Kac, who claimed that it is possible to infer the shape of a drumhead from the sound it makes when struck.¹ Gordon and Webb later found that a single sound impulse could in fact describe multiple, isospectral drumhead shapes (fig. 2).² Their finding reflects a structural difference between forward and inverse problems: a forward problem has a single solution, whereas an inverse problem may consist of a model space of multiple solutions that are either equivalent or probabilistically distributed.³ When a drumhead is struck, a single, unambiguous impulse is produced. That same impulse, however, could have been produced by different drumhead shapes.

A simple inverse problem can be used in the early stages of design to determine an acoustically appropriate room volume. The volume can be back-calculated by defining the seating capacity and desired reverberation time and rearranging the variables in Sabine's formula.⁴ If reverberation time were considered the main driver of acoustic quality, solving this inverse problem would suffice as a design template. But early reflections, which are essential for strength, clarity, intimacy and envelopment, are just as important to the listening experience.^{5,6,7} A room reverse-engineered from reverberation time and seat count will not produce a particular shape—any convex, non-coupled geometry will do as long as it conforms to the required volume. Early reflections, however, are not



Figure 1: Gordon and Webb pose with paper models of isospectral drums.

shape-agnostic. Two rooms with the same reverberation time, one tall and narrow and the other short and wide, will sound different because early reflections will arrive at listeners' ears from different directions, at different delays and at different levels of attenuation. Acknowledging the importance of early sound requires a more comprehensive approach to acoustical reverse-engineering, one with early reflections at the heart of the process.

Our method for acoustical reverse-engineering draws from two lines of room acoustics research. The first, acoustic scene reconstruction, involves using inverse methods to infer deterministic room shapes from early reflections. Gunel describes a technique for estimating a physical room's shape by analyzing a directional impulse response.⁸ Antonacci et al. describe a method for inferring a two-dimensional room boundary by analyzing a series of mono-channel impulse response recorded at different locations.⁹ Dokmanić, Lu and Vetterli reconstruct a two-dimensional room shape from an impulse recorded on a single mono microphone¹⁰ and Dokmanić et al. extend the method to a three-dimensional room by analyzing an impulse response recorded by an array of arbitrarily placed mono microphones.¹¹ This line of research is forensic in its pursuit of a single solution, with advances made as increasingly simple data is used to infer increasingly complex room shapes. Rather than generating a single shape from an impulse response, our aim is to create a solution space offering acoustically equivalent but geometrically diverse outcomes to the designer.

We also draw from work done by Bassuet and Woodger to patent a process for designing rooms based on virtual acoustic signatures.¹² The authors synthesized a number of ideal signatures based on impulse responses recorded in a survey of historical rooms for music. They then developed a room in pursuit of each signature, starting with a basic shape and refining it iteratively until acoustic analysis software confirmed a match. The reverse-engineering method we propose in this paper begins with a virtual sound signature similar to that described by Bassuet and Woodger but presents a new technique for translating the signature into a room shape. Rather than using intuition and iteration as the mechanisms for design, and rather than requiring the designer to hone in on one solution, our method automatically generates a geometric solution space that allows the designer to freely explore a range of unforeseen shapes and forms.

2 CONSTRUCTING A SOUND SIGNATURE

Our method for acoustical reverse-engineering begins with a sound signature that characterizes the acoustical environment of an audience area in a future room. The signature will be used to infer the paths sound waves must take to produce the same acoustic impression, and those paths will in turn be used to generate a series of sound-reflecting surfaces that define the room's overall shape.

The sound signature we adopt for the purpose of reverse-engineering is a "hybrid model" signature consisting of the sound source, a series of image sources and statistically-defined reverberation (fig. 2). The hybrid model was first described by Vorländer¹³ has served as the basis of a number of room acoustic simulation softwares, including Odeon,¹⁴ CATT,¹⁵ Spat¹⁶ and RAVEN.¹⁷ The model cannot represent all of the information that would be captured in a recorded impulse response and it is poor at characterizing rooms in which the sound impression is dominated by properties of diffraction, modal behavior, high absorption, and (if a statistical tail is used) non-ergodic reverberation.^{18,19} But it has proven effective for characterizing medium to large Sabinian rooms with low absorption and strong, relatively broadband early reflections, all qualities typically considered to be desirable in rooms for unamplified music. And crucially, the hybrid model contains enough spatial information to construct a corresponding room geometry.

Source #	Distance (m)	Atten. (dB)			
1	18	-25			
Image Source #	Az. angle (deg)	El. angle (deg)	Delay (ms)	Dist. atten. (dB)	Add. atten. (dB)
1	37.6	2.3	11.80	-26.9	0.0
2	324.6	2.2	15.30	-27.3	0.0
3	46.5	38.7	42.40	-30.2	0.0
4	316.6	37.2	45.70	-30.5	0.0
5	0	60.7	57.90	-31.6	0.0
6	168.4	1.1	96.20	-34.1	0.0
7	200	1.1	97.80	-34.2	0.0
8	37.6	55.7	64.10	-32.0	0.0
9	324.6	54.4	66.10	-32.2	0.0
10	163.3	24.4	111.00	-35.0	0.0
11	204.7	21.2	113.30	-35.1	0.0
Reverberation #	Reverb. time (s)				
1	2.4				

Figure 2: Sound signature displayed in table format.

For purposes of acoustical simulation, the image sources in the sound signature are defined in terms of incidence angle, delay time, attenuation due to distance and attenuation due to absorption or scattering. For purposes of inferring their reflection paths in space, delay is converted into distance (according to the speed of sound in air) and the images are located as points in space relative to a single receiver. To extend an acoustical signature beyond one "originating" receiver location, the image sources must be made audible to multiple receivers across an audience area. But do the images translate relative to each receiver or do they remain fixed in place with reference to the originating receiver? If we were to take the first approach, the incidence, delay and attenuation of each image would remain constant, guaranteeing uniform sound at every seat in the house. While this scenario may seem like the democratic ideal, it is nevertheless impossible to achieve physically, as each set of images for each receiver would require its own set of reflection surfaces—six images multiplied by 500 audience members would yield 3,000 surfaces, many of which would occlude each other (fig. 3). We take the second approach, in which one fixed set of image sources, and one set of corresponding surfaces, serves the entire audience area. All of the images remain audible from seat to seat but gradually change in incidence angle, delay and attenuation moving away from the originating receiver. The range of perceived incidence for each image source can be found by translating the audience perimeter along a vector from the originating receiver to the image source as shown in figure 4. The acoustic variability can thus be predicted and controlled to ensure that the acoustic character remains consistent across the audience area.




Figure 3: Image sources referenced to one originating receiver (l) and translated to give the same early sound impression to multiple receivers (r).




Figure 4: Image sources fixed in space for all receivers (l) and the resulting variability of perceived image source incidence across the receiver area.

2.1 Sound Signature Design Interface

A sound signature may be abstracted from an impulse response as shown by Antonacci²⁰ or derived from a model of a room by image source method (ISM) as described by Allen and Berkley²¹ and Borish.²² Alternately, a sound signature can be designed from scratch, offering the designer the opportunity to invent an entirely new acoustical environment. To facilitate the design of new acoustic signatures, we created a software tool that allows the user to construct, modify and listen to sound signatures in real time. In the tool's user interface (UI), constructed in Processing, the designer can set his or her distance from the virtual sound source, position and reposition image sources, and adjust the reverberation time (fig. 6).

As the user adjusts the acoustical signature, its parameters are ported from the UI to an audio simulation patch constructed in MaxMSP with Spat objects. The patch handles the direct sound and image sources by panning, delaying and attenuating an anechoic music signal through individual taps, and simulates reverberation using Spat's cluster and reverberation objects, with the room's mixing time estimated based on the image source delays. All of the output taps are then combined and decoded for second-order Ambisonics or binaural playback. Because updates to the sound signature are processed and auralized in real time, the user has the choice of listening to changes while making adjustments or can save and retrieve sound signatures to listen to A-B comparisons. The user is also free to switch between different anechoic samples to test how well the sound signature supports different types of music.

The sound signature UI is currently limited by the fact that it is reality-agnostic, in that the user can set any image sources and reverberation he or she wants regardless of whether they are physically realizable. Using the interface thus requires some forethought about the future room's overall envelope and audience capacity. For example, two 15 ms-delayed images arriving from either side




Figure 5: Sound signature design interface.

of the room would likely be very difficult to achieve in a 2,500-capacity concert hall (unless it were extremely narrow and extremely long). Likewise, an image source located straight down would be hard to realize in any room. To increase the likelihood of a realizable sound signature, the user can derive image sources from the future room's maximum envelope and use those to get a rough sense of scale.

As the sound signature produced in the UI will not remain constant from seat to seat in the eventual room, it is the designer's task to craft a signature that can be successfully extended over a larger audience area. To experience how the signature will vary, the interface allows the user to translate his or her listening location away from the originating receiver position to hear the source and image sources from different vantage points. While a single sound signature can be applied to the entire audience area, it may also be desirable to construct multiple signatures and assign them to different audience areas within the same room. The designer can use multiple signatures to create a room with distinctly different acoustical environments or, alternately, to head off acoustical defects. For example, an image source that is well-integrated at the originating receiver position may arrive outside the window of integration at another receiver position some distance away, requiring the signature to be modified or replaced.

3 TRANSLATING A SOUND SIGNATURE INTO A ROOM SHAPE

The image sources in a sound signature are made audible in a room by deducing the location, orientation and extents of planar, sound-reflecting surfaces. A single order of reflection yields one deterministic solution. The first-order reflection plane and reflection point for the originating receiver can be found by inverting the image source method as shown in figure 6. At two orders of reflection, however, a large solution space opens up. The first reflection point (with respect to the receiver) may be located anywhere along the incidence vector up to the plane of first-order reflection. The second reflection point can then be located at any point on the surface of an ellipsoid with the first reflection point and source as foci. Figure 7 shows a partial second-order solution space.




Figure 6: First-order reflection solution.

A reflection surface or pair of surfaces must be extended to make an image source audible to a larger audience area. This is done by inverting the ISM visibility test for a sample of receivers on the audience area boundary, as shown in figure 8. Surfaces can also be extended by the same inverse test to make images of a larger source area audible. To ensure broadband reflections reach the receivers at the boundary of the audience area, the surface edges must be further offset. The offset distance can be approximated by applying Rindel's equation for sizing reflectors, which factors in reflection angle and characteristic distance between source, reflection point and receiver.²³

The materiality of a surface is determined by the originating image source's attenuation. If its attenuation is due to distance alone, a rigid, massive and smooth material will be required. If additional attenuation is specified, the material must be assigned scattering or absorbing properties.




Figure 7: Second-order partial solution space with sample paths drawn (above) and selected path (below).




Figure 8: Second-order solution extended to multiple receivers and sources.

3.1 Solution Space UI

We constructed a user interface in Grasshopper, a parametric modeling plugin for Rhino, that allows the designer to navigate the solution space and create reflection geometry for each image source in a sound signature (fig. 9). The user can toggle between first and second-order solution modes. In first-order mode, only one solution is possible and no further input is required from the user: what you see is what you get. In second-order mode, the user is given controls to explore the larger solution space: a one-dimensional slider sets the position of the first reflection point along the incidence vector and a two-dimensional slider sets the second reflection point's UV coordinates on the ellipsoid. Given the two reflection points, the patch calculates the angles of the reflecting surfaces and their extents for the audience and source areas in the 3D model. As the user adjusts the points, he or she can observe the reflecting surfaces shift and morph in real time, and is free to pick an outcome based on aesthetic, functional or other considerations with the confidence that every solution will be acoustically equivalent.

The UI allows the user to visualize solutions for all of the image in the signature simultaneously. Viewing the individual solutions together helps the designer position surfaces so that they do not collide with each other or occlude each others' reflection paths. To verify that no surfaces are obstructing the image sources' audibility, the user can toggle on a forward ISM visibility check to identify any "deaf spots" in the audience. The same check can be used to identify any unintended early reflection paths, for example a first-order reflection off one of the two surfaces designed to realize a second-order solution.




Figure 9: Solution space explorer UI.

The group of reflection surfaces that emerge from the solution space UI will be discontinuous, so the designer must complete the room's shape by adding infill geometry. Once the room is closed, the patch calculates the room's volume and calculates the Sabins required to meet the sound signature's reverberation time, which the user may then apply to the audience area and/or the infill surfaces. (If any image sources are attenuated more than due to distance alone, some of the Sabins will be automatically applied to those images' reflecting surfaces.)

4 VALIDATION

Once a room geometry is completed, it must be analyzed to determine whether its acoustics match those constructed in the sound signature design interface. Strictly speaking, there should be no audible difference between the virtual sound signature and an impulse response recorded at the originating receiver position in the resulting physical room. The same expectation of similarity applies to displacements within the virtual sound signature and their corresponding positions in the real room.

Because we have not yet completed a project with reverse-engineering as the basis of design (two are currently in development), we present a partial validation as proof of concept. We derived a sound signature from an existing room shape and used it to reverse-engineering two radically different room shapes. Would the three rooms sound the same? We simulated their acoustics in CATT v9 and analyzed the results by comparing numerical measures, visualizing 3D Impulse Responses (3DIRs), and conducting A-B listening tests. In addition to checking for similarity between comparable receiver positions between the three rooms, we also wanted to assess the variance in sound impression within each room. Would generating the room from a single set of fixed image sources produce an acoustical pattern that would hold over a large audience area?

4.1 Setup

We derived an initial sound signature from a design for a long and narrow hall with a cross-shaped section. The room had been designed to have reverberant yet intimate and enveloping sound characteristics, supported by an array of strong early reflections from multiple directions. We estimated the reverberation time by Sabine formula and identified key early reflections by ISM for a source on stage and a receiver placed a little over halfway back in the audience area. The resulting sound signature is shown in figure 3 above. Using the solution space UI, we then constructed




Figure 10: Three room shapes created from one sound signature.

surfaces corresponding to each image source in the signature, sizing them to reflect sound from the source to the entire audience area. Where surfaces unavoidably overlapped or intersected, we trimmed them to ensure a continuous, occlusion-free envelope. We then added infill geometry to complete the room shape. We carried out this process three times, creating three different room shapes (fig. 10-11). Room A is a recreation of the originating room's cross-shaped geometry. Room B is a product of first-order solutions only, resulting in a symmetrical but non-orthogonal form. We constructed Room C from a combination of first and second-order solutions, enabling us to craft a more organic, asymmetrical shape.



Figure 11: Interior views of rooms A, B and C

4.2 Computer simulation

To simulate the acoustics of each room, we imported their geometries into CATT v9. We assigned the reflecting, infill and stage surfaces a mid-frequency absorption coefficient of 0.02 and a scattering coefficient 0.3 and assigned the audience plane a mid-frequency absorption coefficient of 0.98 and a scattering coefficient of 0.9. (We assigned greater-than-typical values for absorption in order to produce a clearer comparison of the effects the rooms' reflecting surfaces). We specified a single, omnidirectional source at the same location as in our initial ISM derivation and in addition to the originating receiver, added 13 other receiver positions spread across the audience area. Each room configuration was rendered in CATT-TUCT v1.1 using algorithm 2 with 1,000,000 rays emitted and diffraction disabled.

4.3 Results and analysis

Figure 12 shows values for ISO-3328:3–2012 parameters G, T-30, EDT, C-80, and LF, extracted directly from the CATT-TUCT echogram at the 500 Hz octave band, with the range of just-noticeable-difference (JND) for each parameter indicated by error bars. The results show very strong acoustical similarity between the three rooms, with nearly all parameter values falling within JND ranges of each other from position to position. The results also show evidence of a relatively consistent sound impression across the audience area (save for receivers at the very front, where direct sound inevitably dominates, and receiver 14 at the very back). The most notable discrepancies are that room B's EDT values are consistently lower and room B's LF values are consistently higher than those of room A. The differences are slight, however, with almost all values falling inside the JND range.

Figure 13 shows 3DIRs for receiver positions 6, 12 and 14. The 3DIRs were created by spatializing 2nd-order B-format impulse responses generated in CATT according to a method developed by Bassuet,⁷ and show the directionality, intensity and arrival time of sound reflections relative to the direct sound. Like the numerical parameters, the 3DIRs indicate strong similarity between comparable receiver positions. The delay time window and directionality of reflections is consistent between rooms with only slight differences in energy level. The only notable discrepancy is at receiver positions 12 and 14, where a reflection arriving from the upper left is visible in rooms B and C but not in room A. Within each room, the reflections change in intensity, delay and incidence angle from position to position but importantly, they all remain present in the signature, again confirming the extension of the originating signature into a consistent pattern across the audience area.

As a final subjective test of whether the rooms would sound the same, we carried out a blind A-B listening survey in the Arup SoundLab, focusing on receiver positions 6, 9, 12 and 14. 12 participants, all acoustic and audiovisual consultants, were asked to compare 11 pairs of 2nd-order B-format impulse responses convolved with an anechoic sample of Handel's *Water Music* and decoded for Ambisonics playback over the SoundLab's speaker array in terms of loudness, clarity, early-to-late energy balance, image width, envelopment and reverberation. For each parameter, participants could respond that sample B was greater than A (+1), equal to A (0), or less than A (-1). Participants were also asked to note the most significant difference they heard between the two IRs. For each comparison, the sample was first played twice, switching from A to B and then from B to A at the halfway points. Participants were then allowed to request as many repetitions of the sample in A-B and/or B-A order as they needed to complete the comparison. Each comparison included one IR from room A and one IR from either rooms B or C, with the playback order varied. The one exception was one pair of identical IRs, included as a control for differences between the two halves of the anechoic sample.




Figure 12: Comparison of acoustical parameters for the three rooms. Position 8 is the originating receiver.

Figure 14 shows averages of perceived differences in acoustical parameters in rooms B and C compared to room A for receiver positions 9, 6, 12 and 14, for an average of the positions and for the control. A score of +1 would indicate a unanimous perception of a parameter being greater in room B or C than in room A. A score of -1 would indicate unanimous perception of the parameter being less in room B or C than in room A. A score of 0 indicates either unanimous perception that the parameter was the same in both rooms or evenly-balanced disagreement. The scores are generally low, and their absolute values are for the most part less than those of the same-IR control, indicating that the rooms sounded the same or at least very similar depending on the participant and the receiver position. Written comments support the same conclusion. For the comparison between room C and A at receiver position 6, for example, participants wrote comments including, "Very subtle," "Hard to compare," "Didn't hear much difference" and "Any difference too subtle to




Figure 13: 3D Impulse responses at receiver positions 6, 12 and 14.




Figure 14: Averages of perceived differences in acoustical parameters in rooms B and C compared to room A.

distinguish.” Averaging the perceived differences for each parameter over the four receiver positions, the only distinction between room B and room A seems to be that it is slightly more enveloping when measured against the control. None of the averaged parameter differences in room C exceed those of the control by more than 0.1, making it difficult to claim any overall difference.

From the results of our numerical, visual and aural analyses, it is clear that the three rooms we generated from a single sound signature sound nearly the same, and from some receiver positions are acoustically indistinguishable. The only consistent differences we observed across the three analyses involve spatial impression. We attribute these differences not to a fault in our reverse-engineering method but to two aspects of our validation procedure that could have been better controlled. First, the relatively high scattering coefficient we applied to the reflecting surfaces likely attenuated second-order reflections more than first-order reflections, which may explain why room B, reverse-engineered to return first-order reflections only, was perceived to more enveloping. Some variance in spatial impression between the halls may also be due to small “deaf spots” caused by trimmed reflection surfaces to make them fit together into a continuous envelope. Room A required substantially more trimming than rooms B and C in order to match the originating room’s shape. A forward ISM analysis of selected surfaces in the three rooms demonstrates that in room A, upper right side reflections are indeed inaudible from the right edge of the audience area (fig. 15).




Figure 15: Upper right side reflections in rooms A, B and C. Note the lack of reflections reaching receivers next to the right side wall in room A.

5 CONCLUSION

In this paper we presented a new method for acoustical reverse-engineering. We illustrated how a virtual sound signature can be used to generate a solution space of geometries to achieve the same acoustic impression physically. We also showed that a single signature can be used to define an acoustical environment with an acceptable range of variation for a larger audience. In our validation experiment, we demonstrated that the solution space for one sound signature can produce room shapes that look different but sound the same, not just from one listening position but across the entire audience area.

By linking a sound signature to all of its possible geometric solutions, reverse-engineering frees the designer to experiment outside the aural, formal and typological conventions of concert hall design. The designer can sculpt any sound knowing it will be physically realizable and can craft any shape knowing it will not compromise the acoustics.

From an aural perspective, reverse-engineering offers an opportunity to explore new sound aesthetics that might be unlike those of any existing room for music. Conversely, it offers an




Figure 16: The interior of room B showing basic geometry and added architectural flourishes.

opportunity to replicate the sound from an existing hall without needing to recreate its shape. A room can sound like a shoebox without looking like a shoebox, for example. Our method also gives the designer control over how uniform or disparate the sound impressions are within a single space. A reverse-engineered hall could offer listeners the choice between two (or more) distinct but equally engaging acoustical environments. Control over the consistency of a room's sound sets reverse-engineering apart from real-time auralization techniques, which allow the designer to modify a room shape and hear the resulting sound in real-time, but only from one listening position at a time, and with no guarantee that the sound impression will extend to a larger audience area.

While reverse-engineering gives the designer a great amount of sonic flexibility, it currently limits control over the shape of the late sound to setting reverberation time. Bradley and Soulodre, Bassuet and Lokki have found that variations in late energy can affect perceptions of envelopment and spaciousness,^{24,7,25} but the relationship between geometry and late sound is still relatively unexplored, so more research must be done before it can be integrated into the reverse-engineering process.

From a visual and formal perspective, reverse-engineering offers an opportunity to create new spatial environments that are not simply variations on the traditional concert hall models. In our validation exercise, we took advantage of this opportunity to create three very different looking room




Figure 17: Diagram for a new performing arts venue in Montana featuring a reverse-engineered inner shell.

shapes from the same initial sound signature. The freedom comes at the relatively small price of refraining from napkin sketches at the start of a project in favor of developing a flexible design language capable of transforming the raw geometric output of the solution space. It goes without saying that the final rooms do not need to be hard-edged and gray-colored (fig. 16). Above, we mentioned the possibility for a sound signature from a historical room to be reverse-engineered into a radical new shape. By the same token an architect should also be able to take a radical new sound signature and reverse-engineer it into a conservative, orthogonal room shape. We also note that surfaces generated through reverse-engineering do not need to be connected into a continuous room shape. For a new performing arts center in Montana, we are proposing an inner array of reverse-engineered panels open to a larger acoustical volume (fig. 17). In this “deconstructed” concert hall, the panels will realize the images in our sound signature by returning early reflections to the audience and the outer box will provide the reverberation.

Acoustical reverse-engineering cannot realize all types of room sounds, nor is it capable of generating all types of room shape. A sound signature dominated by uneven late sound, for example, would be hard to realize, and the solution space will never automatically return a curved surface. But despite its exceptions, reverse-engineering remains a powerful tool for exploring new sonic and formal environments. By freeing the designer from preconceptions about how a room for music should look or sound, we hope the method will stimulate designers to invent a great diversity of engaging, surprising and stimulating spaces for performing and listening to music.

6 ACKNOWLEDGMENTS

The work presented in this paper was initiated and developed in partnership with Matthew Berstch. The authors would also like to acknowledge the generous technical assistance of Yshai Yudekowitz, Charles Avis, Han Dong and Terence Caulkins.

7 REFERENCES

1. M. Kac. 'Can one hear the shape of a drum?', *The American Mathematical Monthly* 73(4), 1-23. (1966).
2. Gordon and D. Webb. 'You can't hear the shape of a drum', *American Scientist* 84(1) 46-55. (1996).
3. A. Tarantola. *Inverse Problem Theory and Methods for Model Parameter Estimation*. Philadelphia: Society for Industrial and Applied Mathematics. (2005).
4. L. Beranek. *Concert Halls and Opera Houses*, 2nd ed. New York: Springer. (2004), p. 541.
5. A.H. Marshall. 'Acoustical determinants for the architectural design of concert halls', *Architectural Science Review* 11(3), 81-87. (1968).
6. Y. Jurkiewicz and E. Kahle. 'Early reflection surfaces in Concert Halls - a new quantitative criterion', *Proc. Acoustics '08*. Paris (2008).
7. A. Bassuet. 'New acoustical parameters and visualization techniques to analyze the spatial distribution of sound in music spaces.' *Proc. International Symposium on Room Acoustics*. Melbourne (2010).
8. B. Gunel. 'Room shape and size estimation using directional impulse response measurements', *Proc. Forum Acusticum*. Seville (2002).
9. F. Antonacci, et al. 'Inference of Room Geometry From Acoustic Impulse Responses', *IEEE Transactions on Audio, Speech and Language Processing* 20(10). (2012).

10. I. Dokmanić, Y.M. Lu, and M. Vetterli. 'Can one hear the shape of a room: the 2-D polygonal case', Proc. ICASSP. Prague (2011).
11. I. Dokmanić, et al. 'Acoustic echoes reveal room shape', Proc. National Academy of Sciences of the United States of America. (2013).
12. G.A. KnickKrehm, A. Bassuet, G. Ellerington, and A. N. Woodger. Methods and systems for improved acoustic environment characterization, U.S. Patent No. 8,396,226. (2013), sheet 42.
13. M. Vorländer. 'Simulation of the transient and steady-state sound propagation in rooms using a new combined ray-tracing/image-source algorithm', J. Acoust. Soc. Am. 86 (1), 172-178. (1989).
14. G. Naylor. 'Treatment of Early and Late Reflections in a Hybrid Computer Model for Room Acoustics', Proc. 124th ASA meeting. New Orleans (1992).
15. B-I. Dalenbäck. 'A New Model for Room Acoustic Prediction and Auralization'. Doctoral Thesis. Chalmers University of Technology. (1995).
16. J-M. Jot. 'Synthesizing Three-Dimensional Sound Scenes in Audio or Multimedia Production and Interactive Human-Computer Interfaces.' Proc. 5th Int. Conf. Interface to Real & Virtual Worlds. (1996).
17. S. Pelzer et al. 'Interactive real-time simulation and auralization for modifiable rooms', Proc. International Symposium on Room Acoustics. Toronto (2013).
18. B-I. Dalenbäck. 'Whitepaper regarding diffraction.' Report. (2012).
19. M. Vorländer. Auralization: Fundamentals of Acoustics, Modelling, Simulation, Algorithms and Acoustic Virtual Reality. Berlin: Springer Verlag. (2008), p. 213.
20. F. Antonacci, et al. 'Inference of Room Geometry'.
21. J.B. Allen and D.A. Berkley. 'Image method for efficiently simulating small-room acoustics', J. Acoust. Soc. Am. 65, 943-950. (1979).
22. J. Borish. 'Extension of the image model to arbitrary polyhedra', Journal of the Acoustical Society of America 75(6), 1827-1836. (1984).
23. J. H. Rindel. 'Attenuation of Sound Reflections due to Diffraction', Proc. Nordic Acoustical Meeting. Aalborg (1986).
24. J.S. Bradley and G.A. Soulodre. 'The influence of late-arriving energy on spatial impression', J. Acoust. Soc. Am. 97, 2263-2271. (1995).
25. T. Lokki. 'Throw away that standard and listen: your two ears work better', Proc. International Symposium on Room Acoustics. Toronto (2013).