

Illuminating Acoustic Reflection Orders Using SketchUp + Light Rendering

by J. Parkman Carter



© J. Parkman Carter 2016

This is a free e-book intended for educational and research purposes. You are free to share it (in unmodified form) with whomever you wish.

This project was partially funded by a grant from the Robert Bradford Newman Student Award Fund in honor of Theodore John Schultz. The Schultz Grants provide partial funding for teachers or researchers interested in developing improvements in architectural education, and acknowledge the contributions of the late Dr. Schultz for his years of distinguished research contributions to the field of acoustics and for his dedicated service as a founding member of the Newman Student Award Fund Advisory Committee. The Newman Fund is sponsored by the Technical Committee in Architectural Acoustics in the Acoustical Society of America.

All product names identified throughout this e-book are trademarks or registered trademarks of their respective companies, and their use here is not intended to convey endorsement or other affiliation with this e-book. While reasonable care has been exercised in the preparation of this e-book, the author and the Acoustical Society of America assume no responsibility for errors or omissions, or for damages resulting from the use of the information contained herein.

Copyright Disclaimer Under Section 107 of the Copyright Act 1976, allowance is made for "fair use" for purposes such as criticism, comment, news reporting, teaching, scholarship, and research. Fair use is a use permitted by copyright statute that might otherwise be infringing. Non-profit, educational or personal use tips the balance in favor of fair use.

FAIR USE NOTICE: This e-book contains copyrighted material the use of which has not always been specifically authorized by the copyright owner. Such material is made available to advance understanding of ecological, political, human rights, economic, democracy, scientific, moral, ethical, and social justice issues, etc. It is believed that this constitutes a "fair use" of any such copyrighted material as provided for in section 107 of the US Copyright Law. In accordance with Title 17 U.S.C. Section 107, this material is distributed without profit to those who have expressed a prior general interest in receiving similar information for research and educational purposes.

Contents

Abstract	v
Acknowledgments	vii
How to view this interactive PDF	ix
1. Introduction	1
1.1 Early Reflections	4
1.2 Image Source Method	7
1.3 Early Engineering Precedents	19
1.4 Art Precedents: “Infinite Rooms”	21
1.5 Caveats: Light vs. Sound	28
2. Method	31
2.1 SketchUp Modeling	32
2.2 Light Rendering with V-Ray	35
2.2.1 Reflective Mirror Material	36
2.2.2 Emissive Source Material	48
2.2.3 Receiver Position Camera	51
2.2.4 Render Quality Settings	53
2.2.5 Saving Renderings	55
2.2.6 Controlling Reflection Orders	56
2.2.7 Color-Coding Multiple Sources	59
2.3 360° Panoramic Renderings	63
2.3.1 Equirectangular Projection	63
2.3.2 “Hacking” V-Ray’s Camera Optics	66
2.3.3 QuickTime Virtual Reality (QTVR)	72
3. Conclusion	77
Image Credits	78
References	79
About the Author	80

Abstract

The conventional architecture workflow tends to—quite literally—“overlook” matters of sound, given that the modeling tools of architectural design are almost exclusively visual in nature. The modeling tools used by architectural acousticians, however, produce visual representations which are, frankly, less than inspirational for the design process.

This project develops a straightforward scheme to visualize acoustic reflection orders using light rendering with the freely available and widely popular Trimble SketchUp 3D modeling software. In addition to allowing architectural designers to visualize acoustic reflections in a familiar modeling environment, this scheme also gives the acoustician a platform to invite designers into a deeper discussion about sound and architecture. A historical background of the method and specific techniques will be presented.

Acknowledgments

This project has been generously supported by the Theodore John Schultz Grant for the Advancement of Acoustical Education (2014-2015 award). I also wish to thank students and faculty at Rensselaer Polytechnic Institute and Parsons School of Design for indulging this metaphor in sonic illumination.

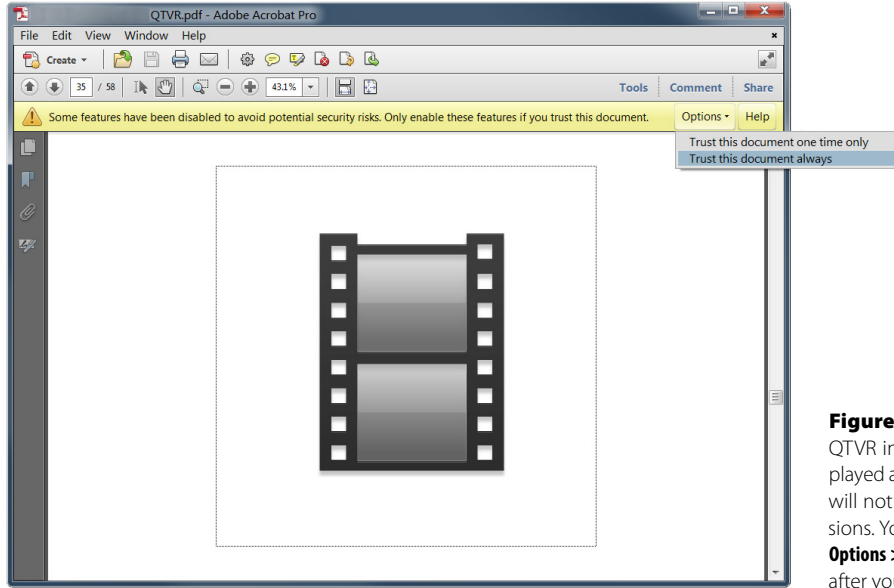


Figure 1: The first time you encounter a QTVR in this PDF, you will see the icon displayed at left, and the interactive image likely will not open due to Adobe Reader permissions. You will see the warning above. Click **Options > Trust this document always**, and thereafter you will see the QTVR.

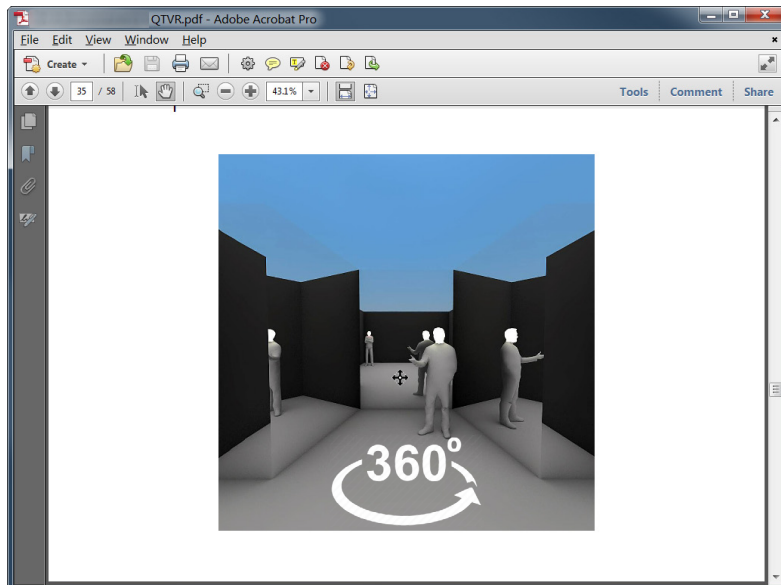


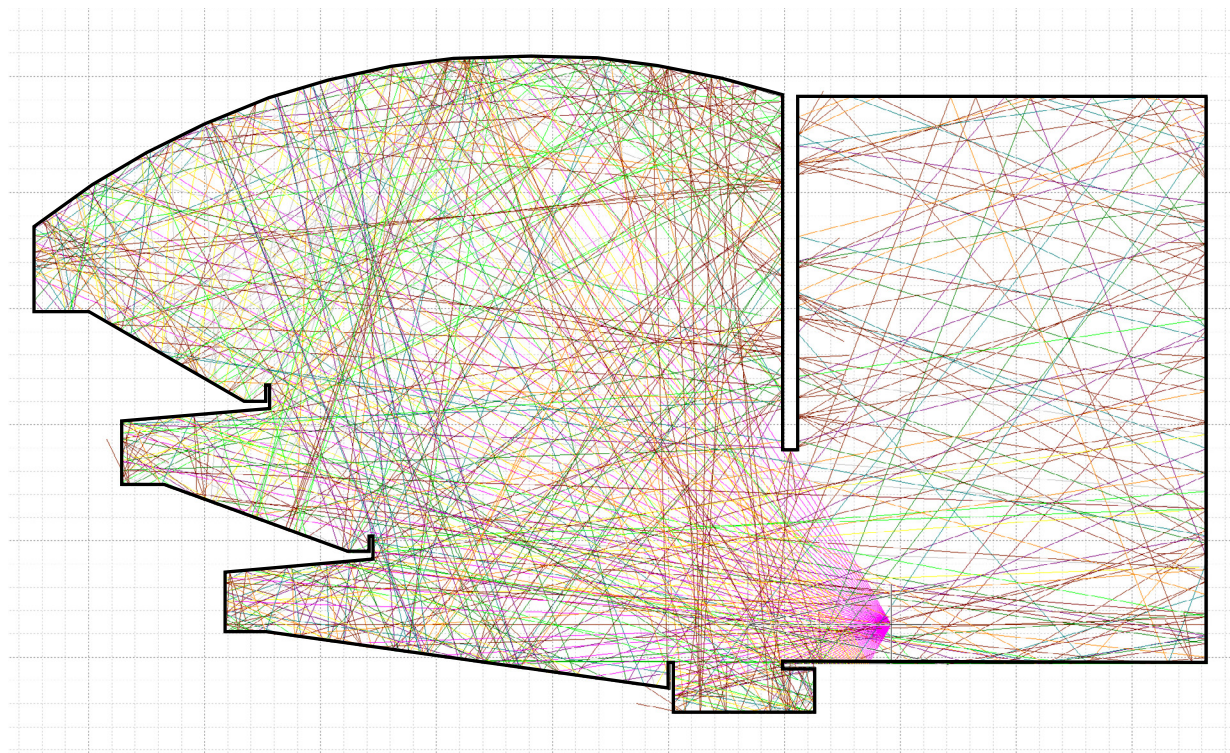
Figure 2: To interact with a QTVR, click and drag the left mouse button in the direction you wish to "steer" your view. You can look all around the 360° image, including above and below your point of view.

How to view this interactive PDF

This document is an interactive PDF that includes some features which may be unfamiliar to some readers. Embedded hyperlinks will appear in the familiar [blue underlined](#) format, and will take you to outside web references via your default internet browser. Figure 1-2 (pg. 5) is an embedded video file – simply click the image to play the video.

The interactive images whose format may be unfamiliar to readers are the embedded QuickTime Virtual Reality (QTVR) files on pages 73 and 75. These are 360° spherical panoramic images which you can “steer” around in to see the entire scene from a fixed point of view. Interacting with QTVR files is similar to using Google Maps in Street View mode. The QTVR format will allow us to create immersive renderings of our spaces under study.

The first attempt to use the interactive figures in this PDF may trigger Adobe Reader’s permissions warning. Please refer to Figures 1 and 2 on the opposite page to see how to apply the correct permissions and interact with the interactive 360° view.



1 Introduction

Visualizing architectural acoustics is tricky. Numerous strategies exist for portraying sound's propagation in space and its interaction with boundary surfaces, but many of the tools to do so are highly specialized and the results are typically illegible to non-acousticians. This means that designers, clients, and end-users have no clear way to participate in the discussion of a project's acoustics.

There is an entire profession of architectural acousticians and for good reason, but as someone who collaborates with designers and works with design students who are eager to address sound in their projects, my motivation is to bring more ears to the design table. All hearing people – even highly visual folks like architects! – have a lifetime's experience of sound in buildings, yet the arcane language of acoustics leaves them believing the topic is beyond their purview. Today the practice of acoustical design remains specialized, intimidating, non-intuitive, and ultimately opaque to many stakeholders, so my goal is to find ways to make the topic more approachable.

One key barrier to a larger discussion of architectural acoustics is the fact that architectural designers and acoustic engineers use tools and workflows which are considerably different. Admittedly, architects rely on drafting, modeling, rendering, and image-editing tools which are, to a fault, entirely visual. But it is worth noting that those tools are commonly used by many other disciplines as well, which widens the conversation around design documents. Part of why the discipline of acoustics is inaccessible

Figure 1-1 (opposite): While visually compelling, the ray-tracing method of portraying sound paths as vectors in an architectural plan or section drawing (shown) is difficult to comprehend, and certainly doesn't leave the designer confident about what acoustical actions to take.

2 ILLUMINATING ACOUSTIC REFLECTION ORDERS

by comparison is due to the fact that the most common tools do not cross over into other disciplines very well, if at all. Acousticians use several different mathematical calculations, some esoteric visualization methods (such as that seen in Figure 1-1), and auralization. Arguably the most powerful computational technique currently used in acoustics modeling, auralization uses virtual models to generate acoustic profiles from a particular space, and then allows the acoustician to test different sounds in the modeled space in a sophisticated and yet direct way: “this is what the room would sound like...” The resulting simulations can be very effective for communicating to a broad audience the acoustical effects created by various design and material choices.

However, while the results can be quite accessible, the auralization workflow is highly specialized and beyond the mettle of most designers who would otherwise welcome the inclusion of sound into their workflow. I have had occasion to teach auralization to both graduate design students and graduate acoustics students and, though I continue to be excited by its potential for design and research, I am always struck by how inaccessible the niche software is, even to students who already understand the acoustic principles involved. And these are not folks who are averse to learning software – indeed architecture and acoustics students today have to learn a dizzying array of software packages. But it seems that auralization software remains too advanced, non-intuitive, and/or cumbersome to integrate meaningfully into the schematic design phases of a project when basic architectural geometries and materialities are being considered and tested.

In the “illuminating acoustic reflections” method which follows, I am proposing a sort of *visualization bridge* between architectural design modeling and auralization. Of course, auralization packages can produce several kinds of graphic visualizations in addition to audio simulations. But even disregarding the challenges already mentioned, the wireframe model views, animated vectors flying about, and color-coded matrices are more suited to engineering analysis than design inspiration. The visualization technique I will show is criminally simple. I believe it is more intuitive and accessible for designers and design students given that it is based on the most common

design software, and works readily with design modeling efforts that are already underway. This allows acoustical thinking to be integrated into the design modeling workflow during the generative stages, rather than being brought in near the end to “fix” a design’s acoustics.

Like many acousticians, I don’t believe that acoustical design should be a fringe topic in architecture, reserved only for music performance spaces. A building’s performance must be evaluated against myriad criteria – energy footprint, historical and cultural aesthetics, social and economic responsibility – but given how significantly a building both determines the soundscape of its inhabitants and impacts the larger soundscape of its contextual environment, it is critical that architectural acoustics join the progressive design discussion alongside daylighting and natural ventilation. It is my hope that designers and design students alike can begin to more confidently grasp how acoustic behavior is related to design decisions. The method which follows is not intended to replace auralization, but rather to use existing software common to the design workflow to produce more intuitive visualization. As such, this method can help shepherd designers into deeper discussions about acoustics and make auralization more relevant to more architectural projects, as designers embrace sound thinking as a vital part of their process.

Although the method presented here is a procedure for creating a visual analogue of acoustic reflections, it is actually based upon some of the earliest modeling techniques used by acoustical engineers, and is directly relevant to many acoustic algorithms in use today. In the early 20th century, engineers would demonstrate the relationship of sound propagation and early reflections in a hall by constructing scaled physical models with shiny surfaces and using a light as the sound source. Today we can take this concept much further – and faster, and more easily – using digital modeling and rendering tools. But first, let’s establish what early acoustic reflections are and how they are important to the sound character of a space.

1.1 Early Reflections

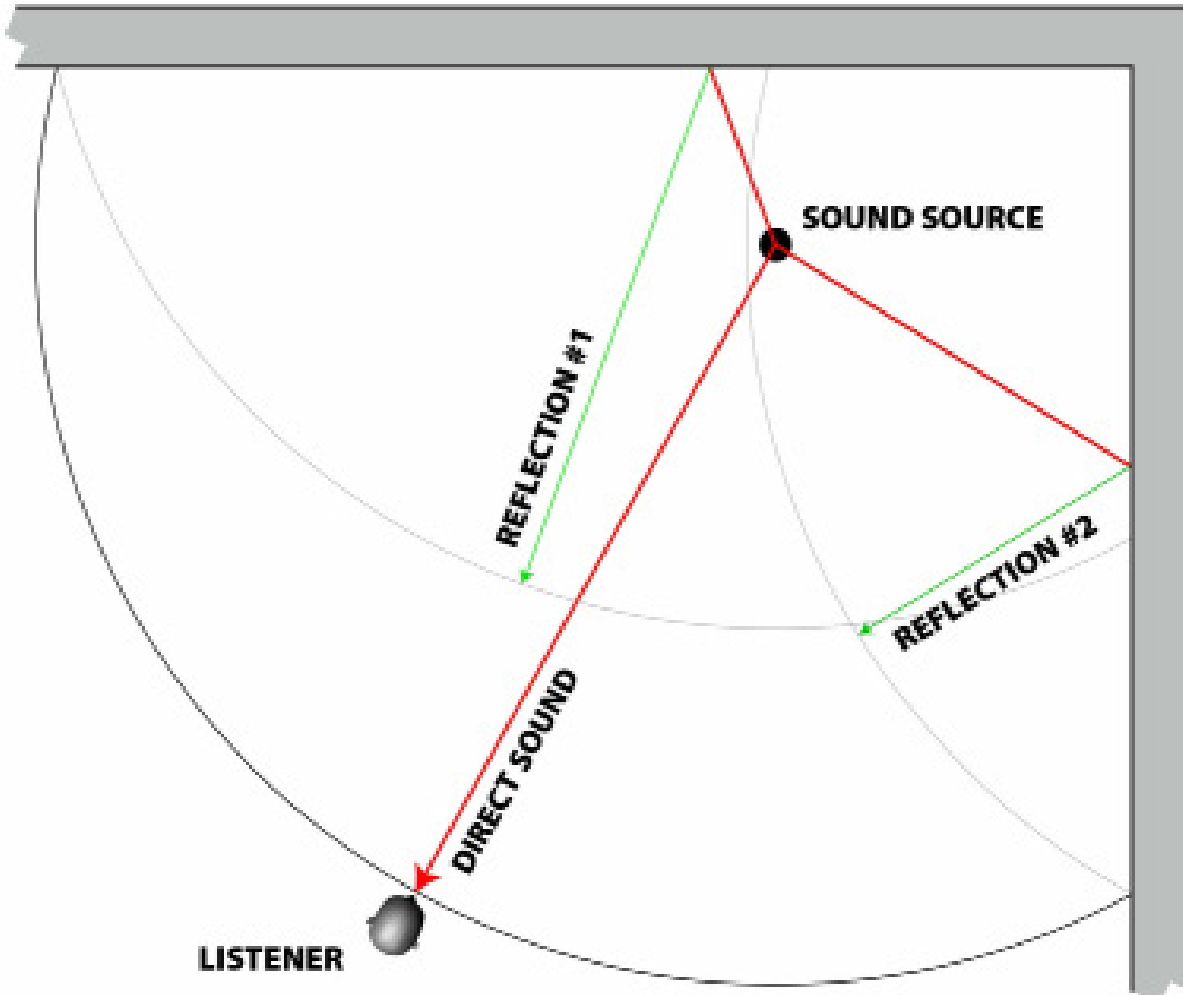
As an audio consultant during the CD era, I had many occasions to configure hi-fi systems in a variety of listening rooms. My colleagues and I would often use a simple method to locate the placement of acoustic treatments, one that required no calculations or drafting of layouts. Since many listening rooms doubled as functional living rooms, it was desirable to minimize the amount of absorption paneling on the walls, and thus effective placement of the panels was critical. Our trick was to grab a CD and have the client sit in the “sweet spot” listening position, then place the disc somewhere on a side wall with the label side against the wall and the shiny, mirror-like side reflecting into the room. We would then slide the disc along the wall until the client could see the reflection of a loudspeaker’s tweeter. Marking this location and repeating the process of locating both speakers’ reflections on each of the walls allowed us to precisely position a minimum number of panels. This practical technique also allowed us to check the positioning of rugs on the floor between the listening position and the loudspeakers.

The goal was to absorb the critical early reflections from surfaces in the listening room relevant to the listening position. The CD provided a simple and handy mirror to identify those first reflections. With those surface locations addressed, the rest of the room could be left untreated and the hi-fi system would still remain capable of a high degree of soundstaging and imaging performance. This is because absorbing the early reflections prevents substantial room “coloration” from interfering with the loudspeakers’ direct sound and “smearing” the system’s capabilities otherwise. Let’s take a closer look at how these early sonic reflections contribute to the perception of architectural acoustics.

When a sound event occurs in some bounded space, the first and loudest component of the sound to reach a listener some distance away is the portion of the sound wave known as the “direct” sound. Any nearby surfaces which are acoustically reflective (i.e. not constructed of materials which absorb or transmit sound elsewhere) will redirect the sound’s energy, creating new portions of wave front which will arrive at the listener slightly

Figure 1-2 (interactive): Click the animation below to see how reflections are delayed relative to the direct sound. Also, note the two different representations of the sound: as a wave front and as a directional vector.

delayed in time and, because of the longer distance traveled and any absorptive qualities of the wall, somewhat lower in level relative to the direct sound. The animation in Figure 1-2 demonstrates this principle (at a speed much slower than the actual speed of sound in air, of course).



6 ILLUMINATING ACOUSTIC REFLECTION ORDERS

The spatial relationship of the sound source and listener locations relative to the various reflective surfaces of a particular space comprise the experience of architectural acoustics. That is to say, the sound source is experienced very differently within architectural spaces than it would be if the source and the listener were high up in open air with no reflective surfaces nearby. In an architectural space, if we were to make an impulsive sound that is as brief as possible (such as a balloon burst) and, at the listener's position, plot the decreasing levels of all the discrete sound reflections over time on a graph, we would have what is known as an "impulse response" which describes the architectural acoustics of that particular space for a particular spatial relationship between the sound source and listener. In practice, acousticians identify three components within the impulse response:

1. The **direct sound** itself, at the front of the impulse response;
2. The **early reflections**, which often occur so close in time to the direct sound that they meld with our perception of the direct sound; and
3. **Reverberation**, which is the late energy in a space's response when several orders of acoustic reflections have merged into a combined wash of dissipating sound energy (see Figure 1-3).

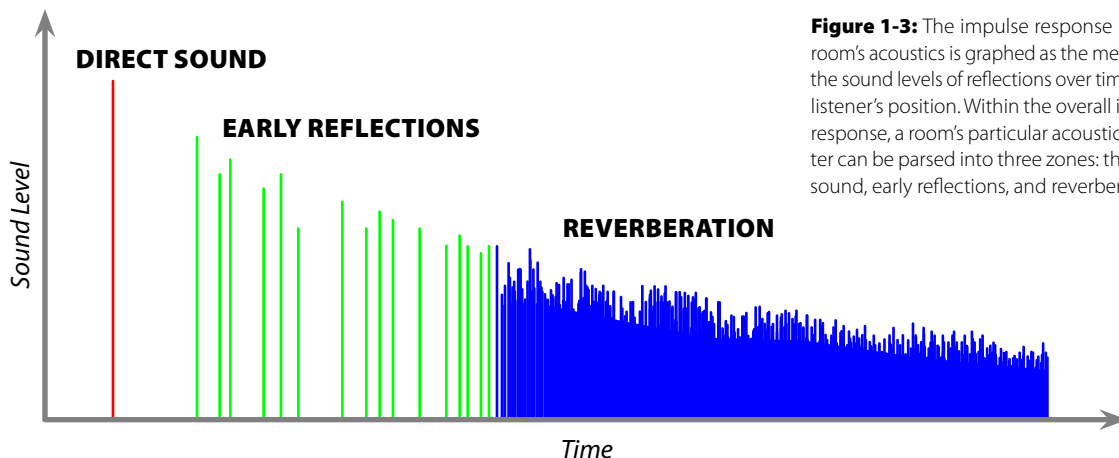


Figure 1-3: The impulse response (IR) of a room's acoustics is graphed as the measure of the sound levels of reflections over time at the listener's position. Within the overall impulse response, a room's particular acoustic character can be parsed into three zones: the direct sound, early reflections, and reverberation.

The perceived balance of direct sound, early reflections, and reverberation combine to produce the distinct experience of a particular space's acoustic character. The early reflections have a pronounced impact on the listener's sense of the direct sound. In the earlier example of the hi-fi listening room, the desire is to minimize or remove entirely the early reflections so that the acoustics which are reproduced in the music recording are obscured as little as possible by one's local room acoustics. But in music performance halls the strategies for addressing early reflections are quite different.

With acoustic (i.e. unamplified) instruments in larger halls, the early reflections can be critical in reinforcing the sounds on stage and helping them to reach the back of the audience. Furthermore, the melding of the early reflections with the perception of the direct sound can also enlarge the apparent size of the sound source (acousticians use a metric known as "apparent source width" or ASW) as well as make the sound of the hall more enveloping or "intimate" to the listener. In other cases, such as with very large halls or with specific geometries, the reflections may occur delayed enough in time from the direct sound that perceivably distinct echoes result, which can be distracting or unpleasant.

Whether trying to remove them or strategically reinforce them, early reflections play a significant role in acoustics, and thus we need ways to locate these primary reflections relative to the sound source and listener position(s). Once we understand the layout in space of early reflections, architectural decisions about geometry and materiality (absorptive vs. reflective) can be made with confidence.

1.2 Image Source Method

Several analogies exist to describe the propagation of sound in space, with ripples in a pond being perhaps the most popular. Wave behavior is quite complex, so the directionality of sound propagation in space is often reduced to the vector corresponding to the wave front's direction of travel. In Figure 1-2, we saw both the wave fronts and the associated vectors represented. But how did we know exactly where the reflection points on the

8 ILLUMINATING ACOUSTIC REFLECTION ORDERS

boundary surfaces would occur? Conveniently, the law of specular reflection for light holds true for acoustic vectors as well: *the angle of incidence equals the angle of reflection*. The easiest way to plot the direction of the reflected vector is to draw a straight line from the receiver/listener position, **not** to the sound source position (that would be the direct sound), but to a *virtual reflection of the sound source*, as if the reflecting surface were an actual mirror. In Figure 1-4, we see how these virtual reflections – also known as “image sources” – determine the directionality of the acoustic vectors.

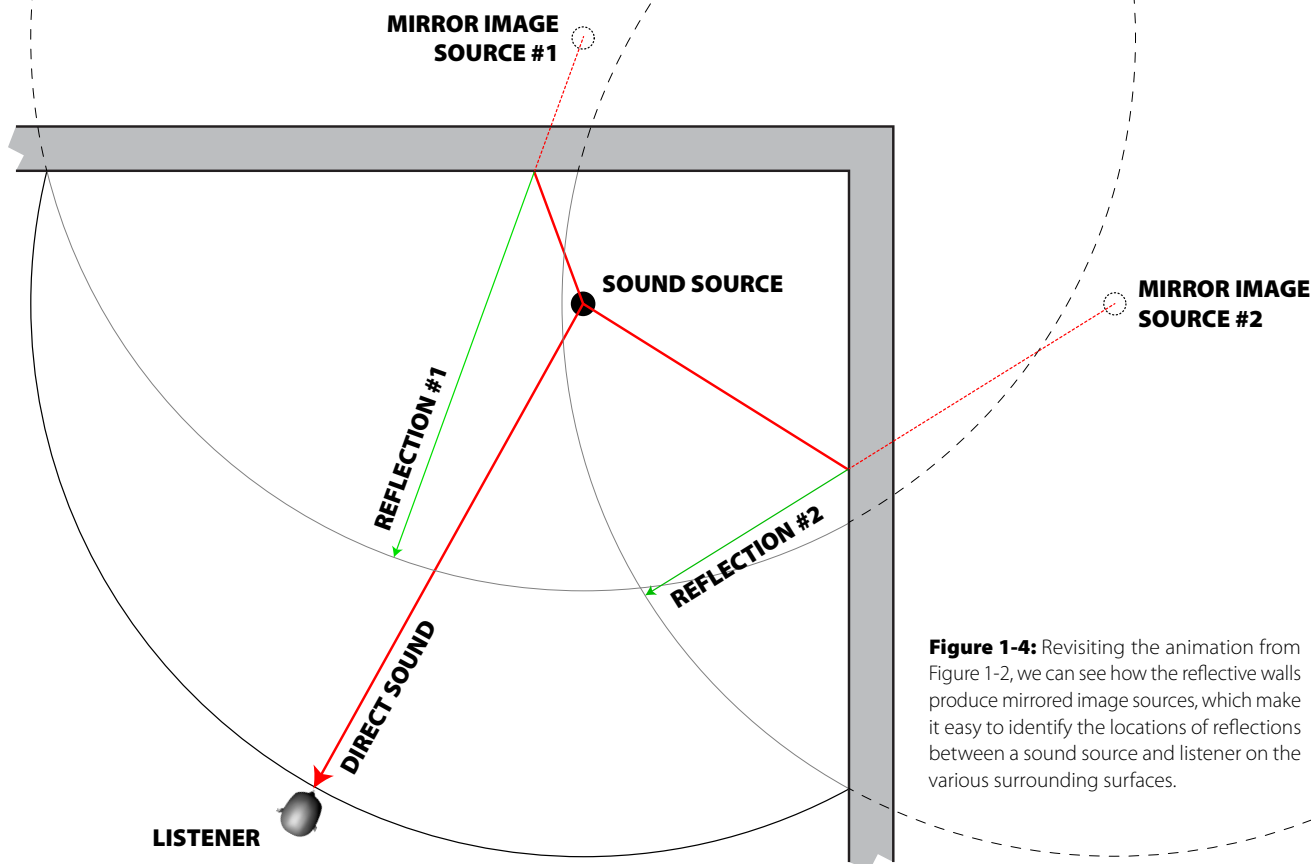
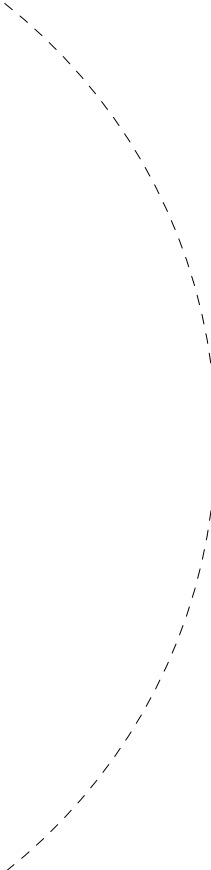


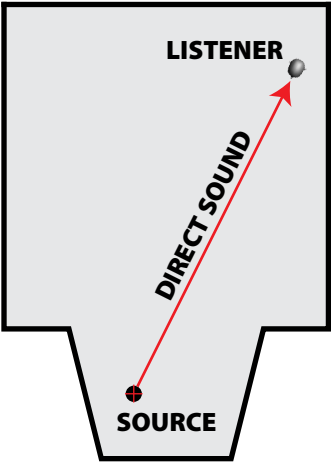
Figure 1-4: Revisiting the animation from Figure 1-2, we can see how the reflective walls produce mirrored image sources, which make it easy to identify the locations of reflections between a sound source and listener on the various surrounding surfaces.



In the above example we see two different “first-order” reflections – first-order because they have bounced off of one surface to reach the listener. But when we are dealing with a fully contained space with reflective surfaces on all sides of the source and listener, we encounter several orders of reflections as virtual images themselves are reflected. This corresponds to sound bouncing continuously around a space until the sound energy dissipates and dies out. The spatial relationship between the sound source and listener, the space’s particular geometry, and the space’s particular balance of reflective vs. absorptive surfaces together determine the presence and distribution of these acoustic reflections.

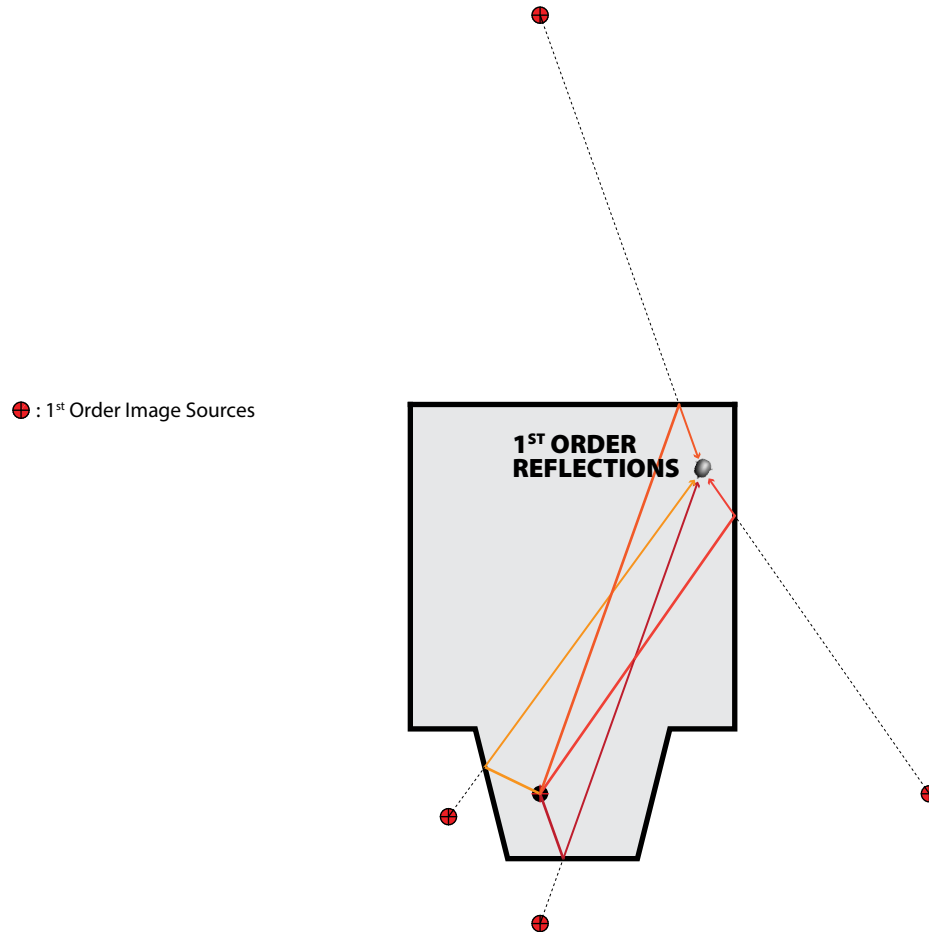
The Image Source Method (ISM) is a procedure for geometrically locating virtual image sources and their paths from the source back to the listener. On the following pages, we see a 2D plan view of a simple room with a sound source located in an alcove and a listener located near an opposite corner. Advancing through Figures 1-5a through 1-5f illustrates the direct sound path followed by the first, second, and third order reflections as derived from their virtual image sources. There are also a couple examples of invalid reflections which cannot exist due to the alcove’s geometry and the positions of the sound source and listener.

Figure 1-5a: We begin with the 2D floor plan of a room with the source located in an alcove and the listener located near an opposite corner, with the direct sound indicated.



Direct Sound • 1st Order • 2nd Order • 3rd Order • All Reflections • Invalid Reflections

Figure 1-5b: The first-order reflections are determined by virtual image sources that would be reflected by the walls surrounding the source and listener, as if those walls were mirrors.



Direct Sound • 1st Order • 2nd Order • 3rd Order • All Reflections • Invalid Reflections

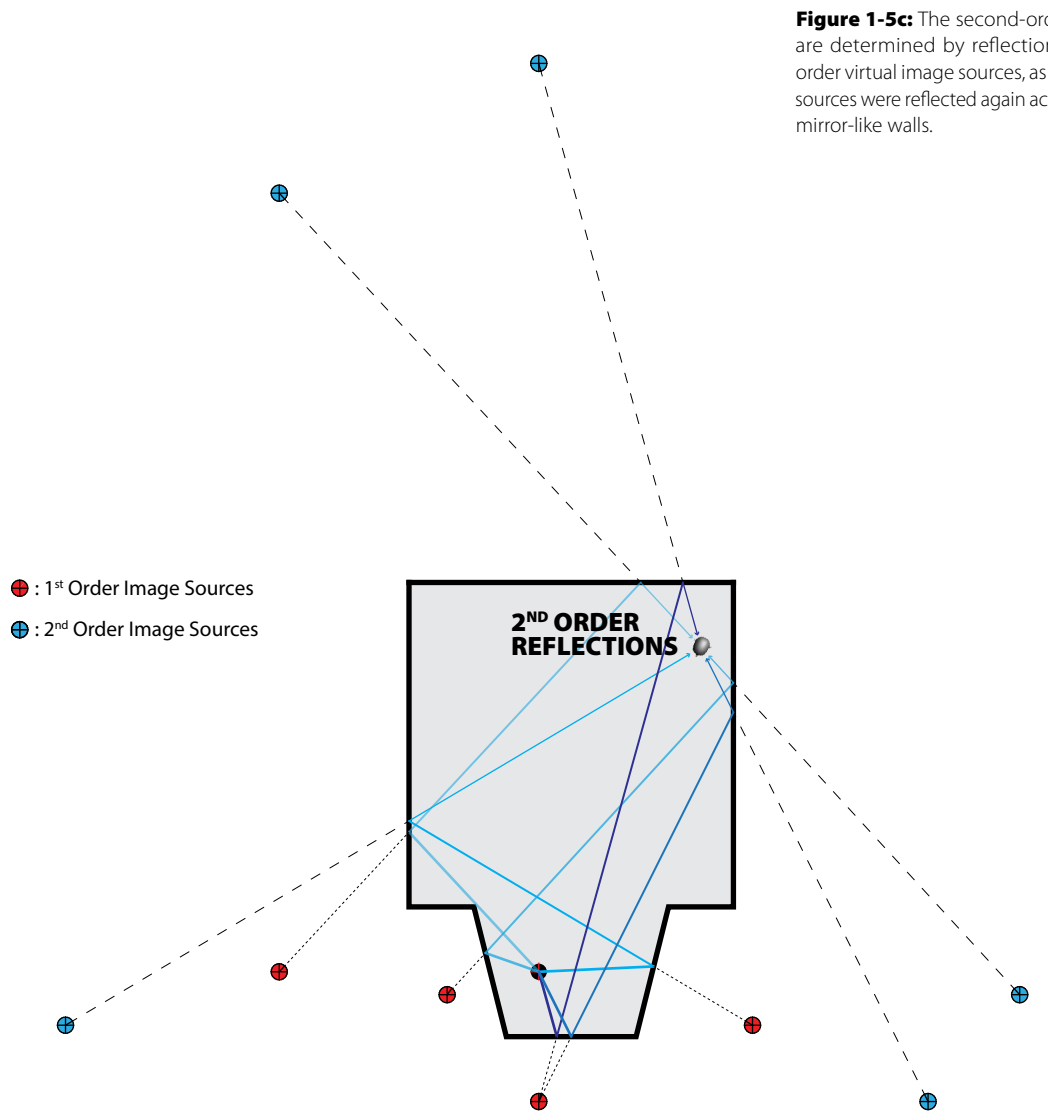


Figure 1-5c: The second-order reflections are determined by reflections of the first-order virtual image sources, as if those image sources were reflected again across the room's mirror-like walls.

Direct Sound • 1st Order • 2nd Order • 3rd Order • All Reflections • Invalid Reflections

Figure 1-5d: The second-order reflections determine the location of the third order reflection sources and thus the paths between the source and listener.

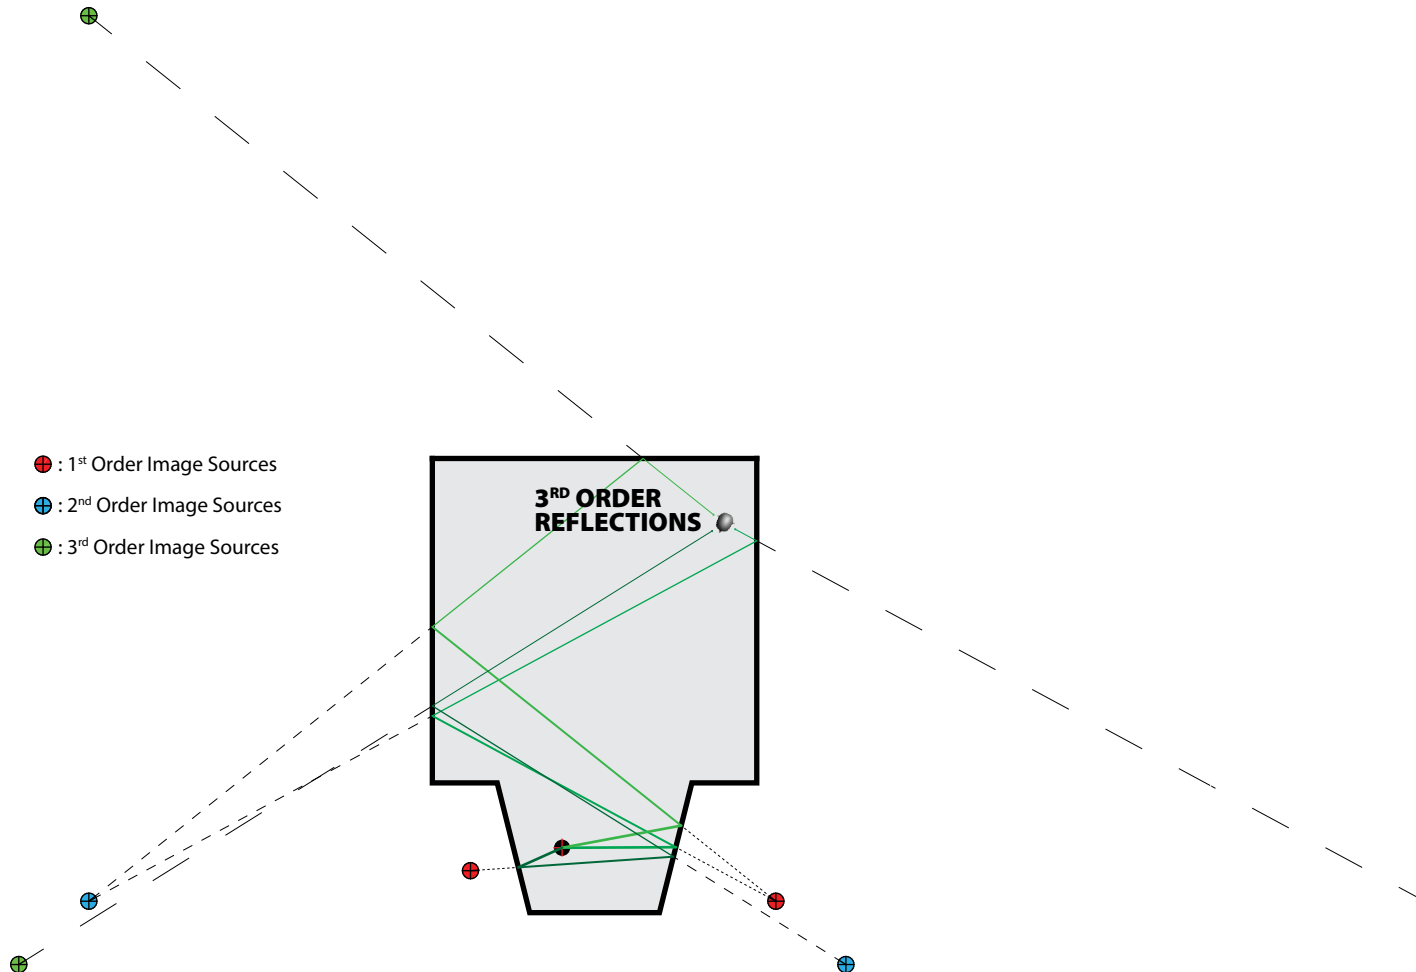


Figure 1-5e: This plot shows all the virtual image sources and associated paths between the source and listener for the first-, second-, and third-order reflections

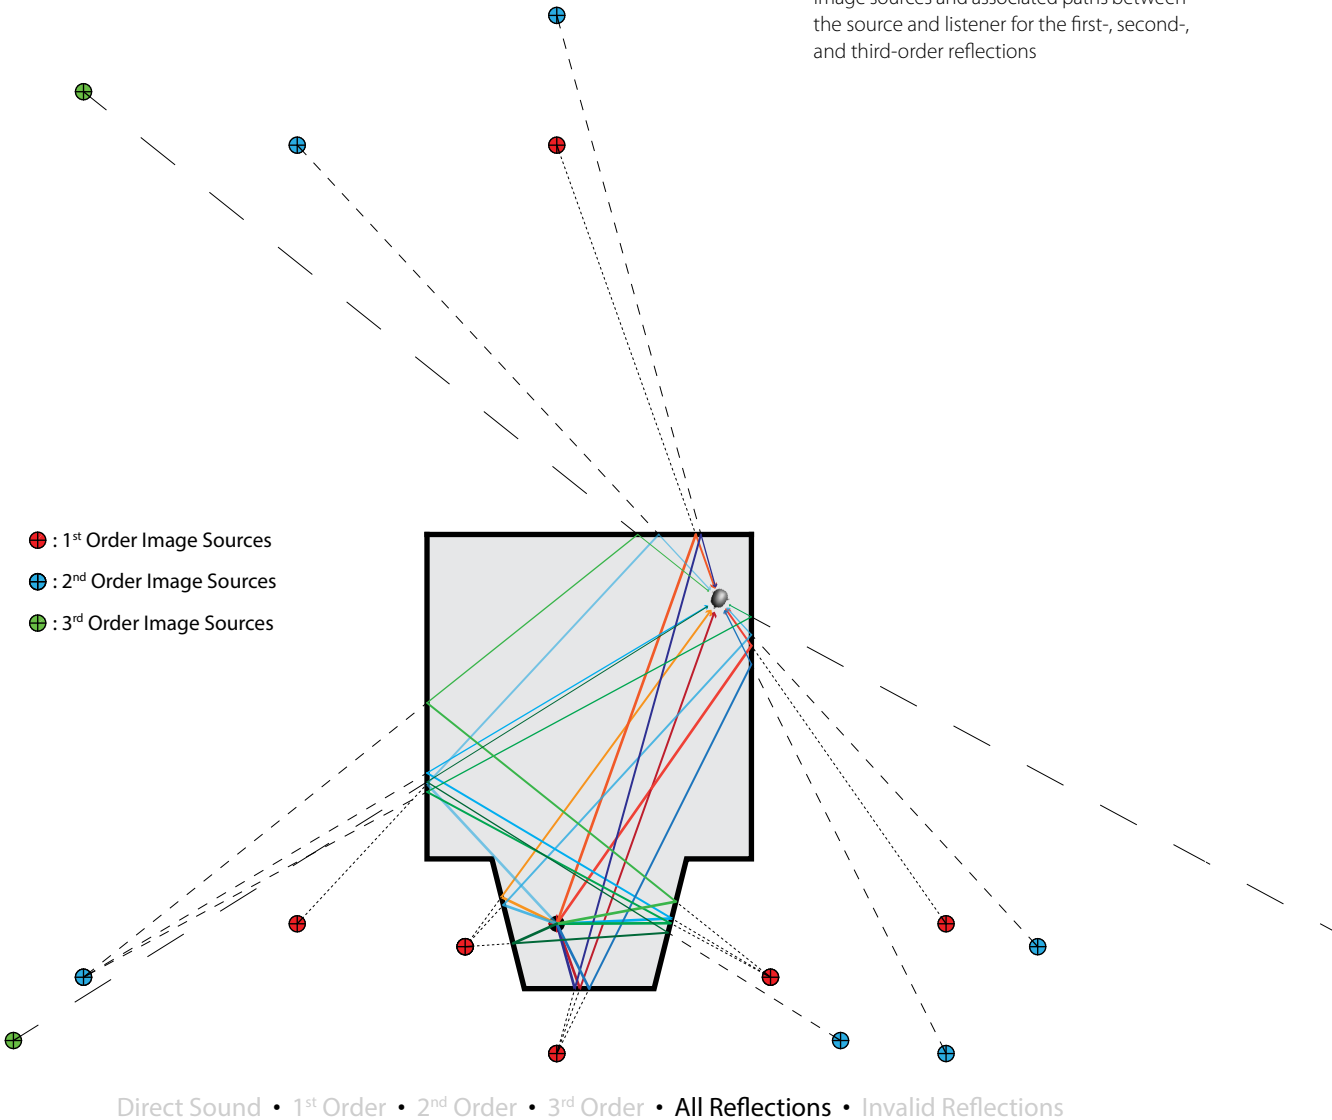
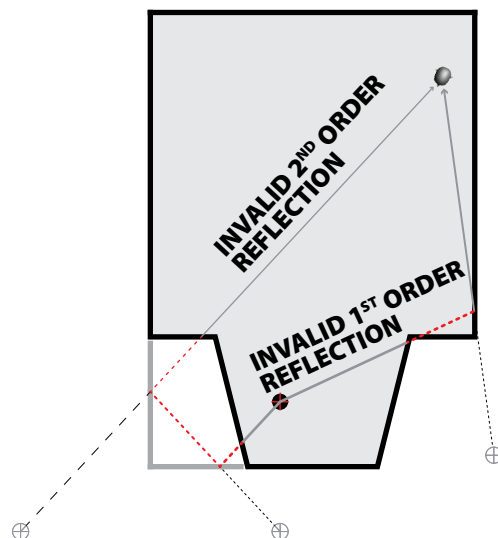
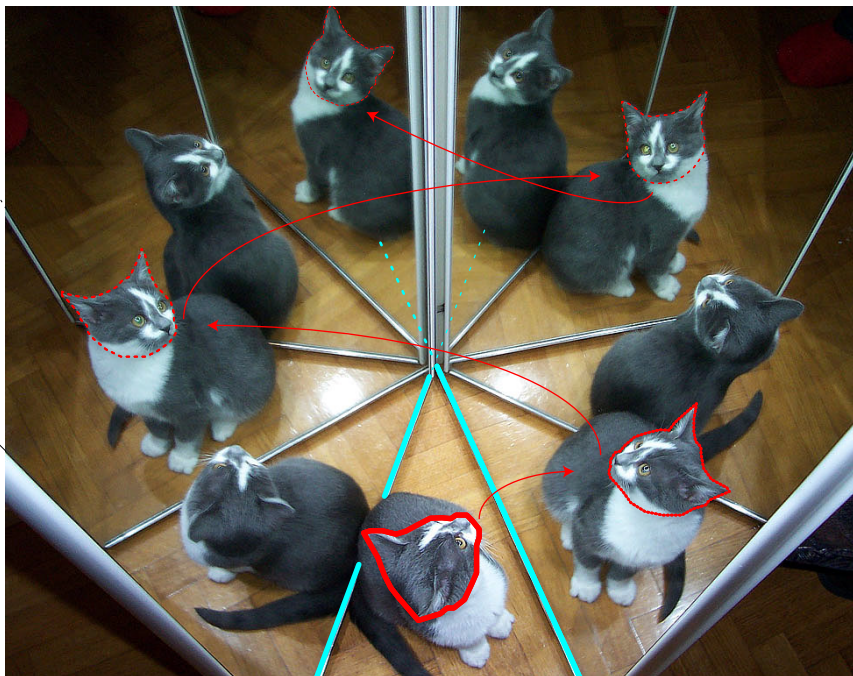


Figure 1-5f: Here we see examples of invalid reflections, as determined by the room's geometry and the spatial relationship between sound source and listener locations.



Direct Sound • 1st Order • 2nd Order • 3rd Order • All Reflections • Invalid Reflections



While simple in concept, it should be clear that calculating several orders of virtual image reflections can become both tedious and complex quite quickly. And this example only shows 2D relationships! Obviously, adding a third dimension to the reflections creates further complexity. The plan view demonstrated is useful for measuring out the virtual image source locations. But let's look at another simple example that demonstrates the experience of 3D reflections from the point of view of the listener. In Figure 1-6, let's designate the kitty to be the sound source and the point of view of the photographer to be the listener's position. We will take the two mirrors on either side of the kitty to be acoustically reflective surfaces, and we will pretend the floor material is absorptive. Click through the gallery to understand the distribution of the different reflection orders, and their spatial relationship to the reflecting surfaces and to the listener.

Figure 1-6a: If we take the kitty to be the sound source and the photographer to be the listener, this example shows how the two reflective walls create four orders of reflections, as reflections themselves are reflected in turn. This figure shows the contribution starting with the right side mirror's first-order reflection.

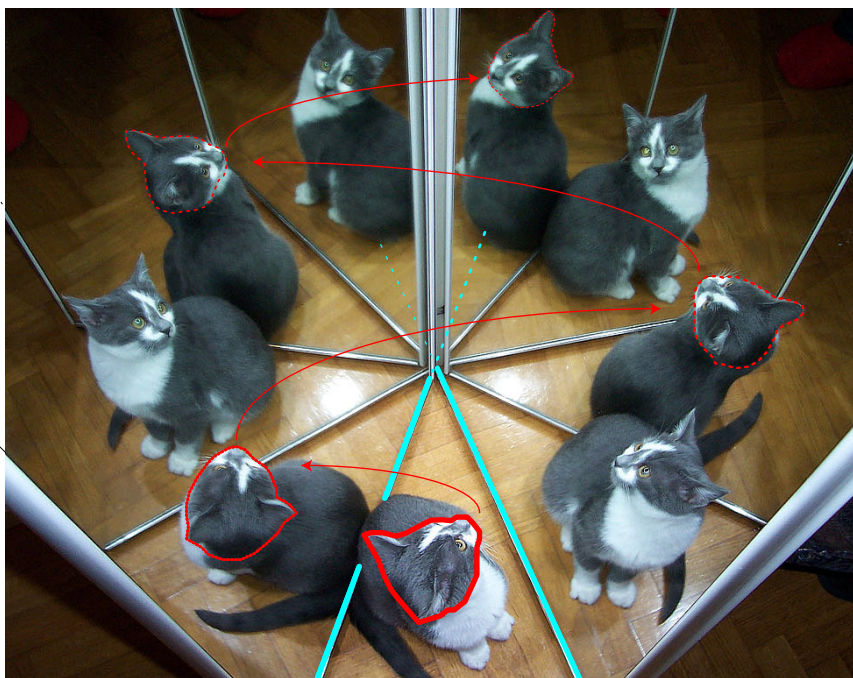
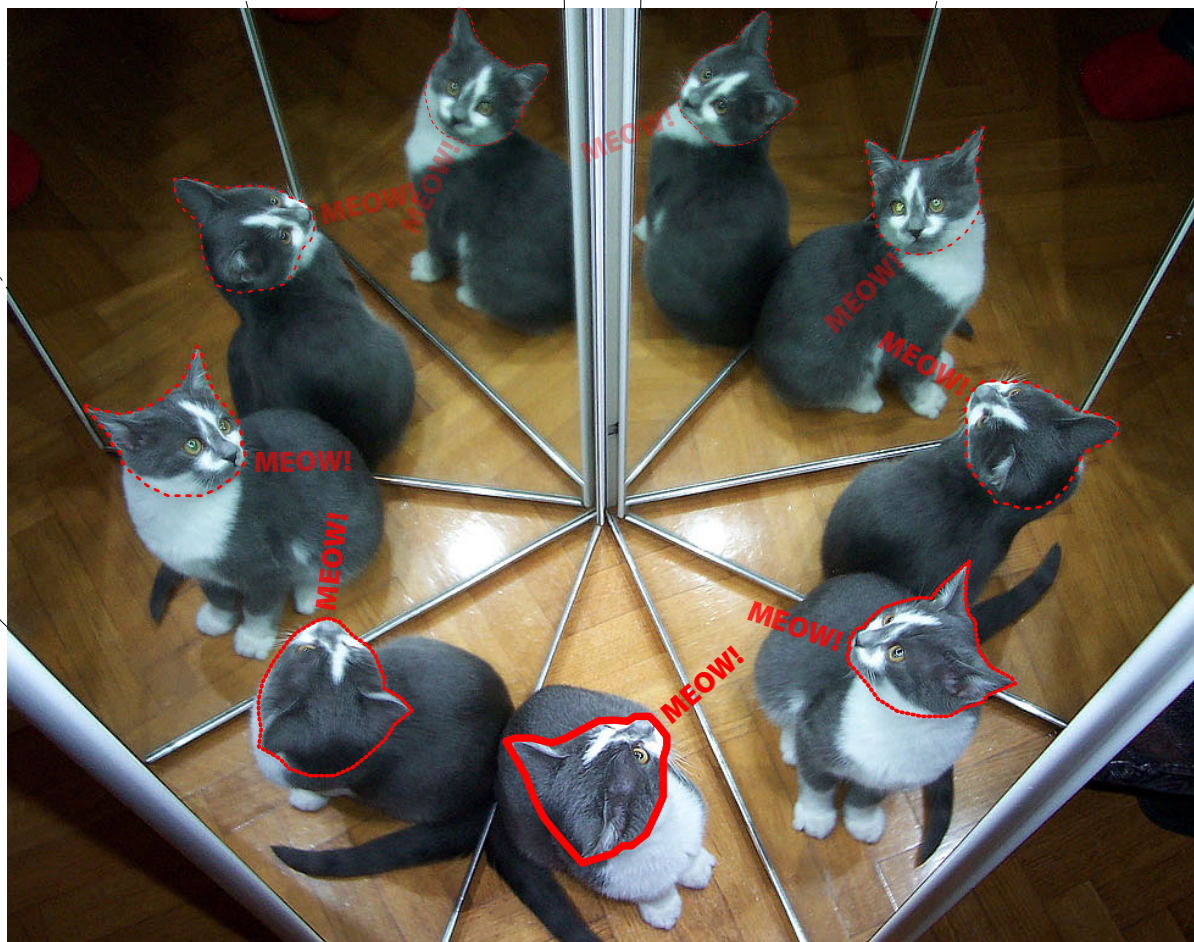


Figure 1-6b: Similar to the previous figure, this figure shows the four orders of reflections contributed by the left side mirror's first order reflection.

It would be difficult to accurately localize the acoustic reflection orders through listening tests alone. This is why such visualization can be instrumental in determining the role of different architectural surfaces which contribute to our experience of sound and acoustics. In the example with the kitty, it is much easier to explain to someone that, in such an architectural configuration of reflective surfaces, we wouldn't simply hear a "meow" confined to the kitty's mouth. We would hear a larger "ring-shaped" meow, with more sonic emphasis coming from the front of that ring where the direct sound and earliest reflections are situated. Thus, a simple visualization using mirrors and the natural behavior of light (which corresponds more or less with the behavior of sound propagation vectors) allows us to demonstrate spatial acoustics in a way that other more abstract visualizations and auralization do not. But this approach is not new!



1.3 Early Engineering Precedents

For nearly as long as architectural acoustics has been a professional discipline, engineers have used light and optics to “illuminate” acoustic behavior. Over a century ago, W. C. Sabine used the Schlieren technique to create spark photographs revealing the propagation of waves in small scale architectural models. More recently, M. Nagata has used optically reflective materials such as glass or Plexiglas to position orchestra shells (Long, 2006). The correspondence in directionality between optical and acoustical reflectivity means that when an audience member can see a visual reflection of the stage musicians in an orchestra shell’s panels, the acoustical reinforcement is also in alignment.

The luminous acoustic rendering technique we will cover here is essentially a recreation – and powerful extension – of another optical modeling technique devised by Japanese architectural engineer Takeo Satow, which was originally reported in the 1929 paper “Acoustics of Auditorium Ascertained by Optical Treatment in Models” (World Engineering Congress, Paper No. 118). Satow built simple models by extruding the 2D plan and section views of an auditorium, and then covering the interior surfaces with a reflective mirror material. To create directional beams emanating from a “sound” source, he placed a light bulb inside a cylinder with slits and filled the model with smoke so that the reflecting beams of light would be visible (see Figures 1-7 and 1-8). In 1940, R. Vermeulen used the tactic to analyze the acoustical geometries of the large assembly hall in the United Nations Building in Geneva (see Figure 1-9).

These simple models directly embody the Image Source Method, and are illustrative indeed for depicting fundamental acoustic behavior in architectural spaces. Given that the principles of specular reflection are conveniently shared between optics and acoustics, these early 2D model experiments (2D because they only depict a *plane* of reflections in either the section or plan view) inspire further questions: could the technique be used to illuminate more complex 3D spaces? How could we both reveal the complex series of reflections and do so without the smoke (for making

Figure 1-6c (opposite): This example shows how the “ring” of four orders of reflections enlarges the sound of the kitty, and demonstrates that such visualization can make acoustic behavior more apparent.

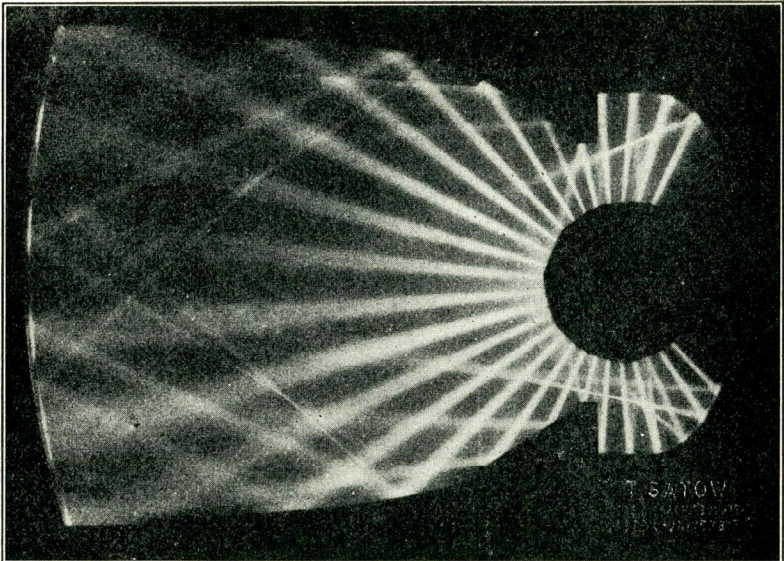


Figure 1-7: Plan view example of Satow's optical model of an auditorium, using mirror material on the walls and a source comprised of a light bulb in an enclosure with slits.

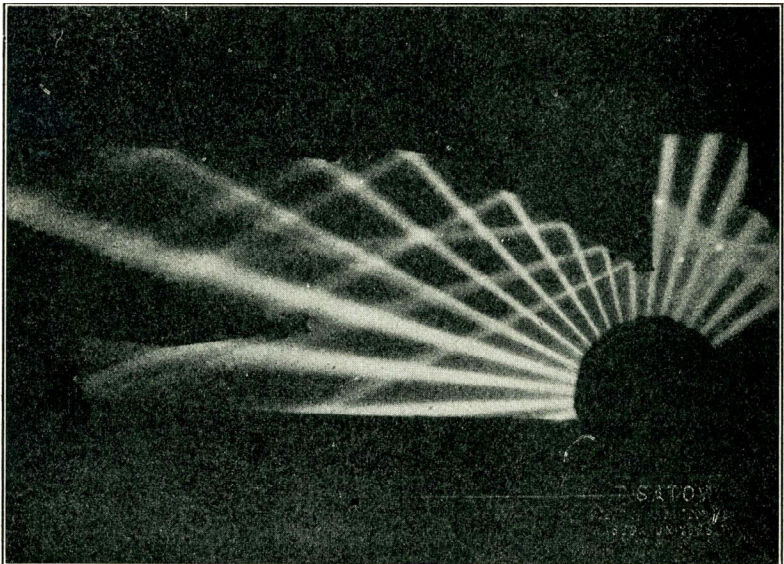
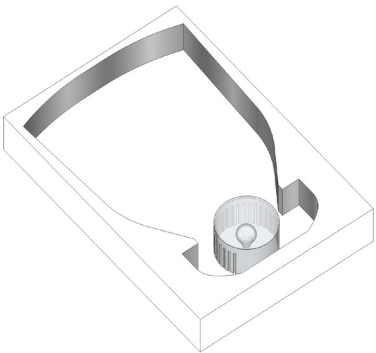


Figure 1-8: Similar to the above example, this optical model by Satow extruded a 2D section view of an auditorium to show source rays reflecting off the ceiling and distributing into the hall, as well as collecting in the fly tower space above the stage.

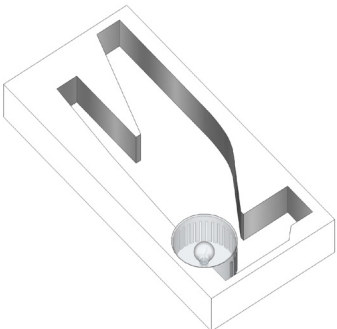
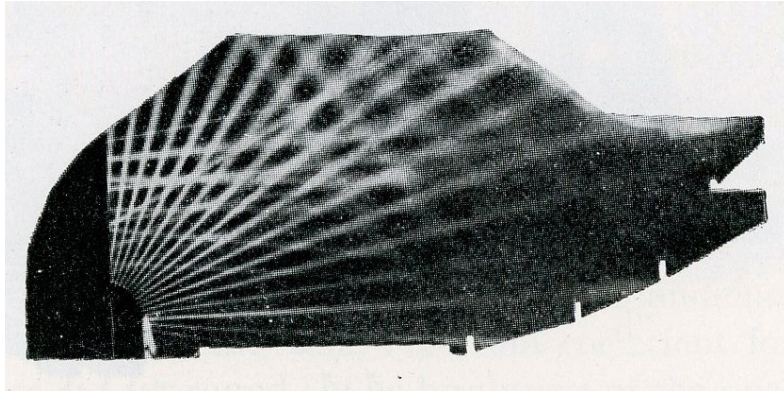


Figure 1-9: In this section view optical model by Vermeulen, we see how the ceiling and proscenium (above stage, left) contribute to the distribution of sound into the hall of the United Nations Building in Geneva.



the light legible)? Could the point of view be shifted from 2D architectural plan and section views to more experiential perspectival views from the listener's position? Disregarding construction complexity and costs for a moment (anticipating the advantages of digital virtual modeling to come later), couldn't we just build full scale models so that we could see the reflection orders mapped right onto the interior surfaces themselves, thereby visualizing how acoustic reflections are spatialized?

Enter the artists!

1.4 Art Precedents: “Infinity Rooms”

While acousticians have typically confined their light-based reflection studies to small scale models, a number of artists have created “infinity room” installations large enough to walk around in, covering a room's surfaces entirely with mirrors. The experience can be breathtaking – stepping into what is often a modest enclosure, and then suddenly finding yourself floating in what appears to be boundless space, with countless reflections spread out into the distance. The effect is expansive and mesmerizing, and invites contemplation about open space and our increasingly confined relationship to the interior. But when we consider that sound distribution works on similar reflection principles, such installations also become powerful demonstrations of acoustic space.

One of the most well-known and prolific mirror room artists is Yayoi Kusama, whose installations often include myriad small elements such as polka-dotted pieces on the floor or tiny suspended point lights which dramatize the infinite space by seeming to fill it. See Figures 1-10 through 1-13. [This video](#) demonstrates what it was like to step into Kusama's 2002 installation *Fireflies on the Water* at the Whitney Museum in NYC, which incorporated a water surface to add a shimmering effect to the infinite reflections. In [this interview](#), performance artist Adele discusses the inspiration she derived from Kusama's recent mirror room installation *The Souls of Millions of Light Years Away* at The Broad museum in LA, which she used as a visual accompaniment to her live performance at the 2016 BRIT Awards.

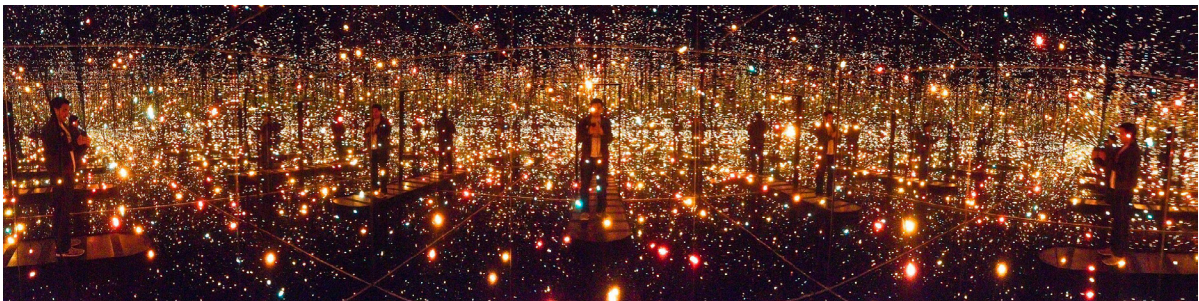


Figure 1-13: Yayoi Kusama, *I Who Have Arrived in Heaven* (2013), David Zwirner Gallery, NYC.



Figure 1-10 (opposite left): Yayoi Kusama, *Infinity Mirror Room – Phalli's Field* (1965), R. Castellane Gallery, NYC.

Figure 1-11 (opposite): Yayoi Kusama, *Infinity Mirror Room – Phalli's Field* (1998), Museum Boij-mans Van Beuningen, Rotterdam.

Figure 1-12 (opposite bottom): Yayoi Kusama, *Fireflies on the Water* (2002), Whitney Museum of Art, NYC.

Thilo Frank's installation *The Phoenix is closer than it appears* (2011) is visually simpler than Kusama's mirror rooms, but includes a playful (if destabilizing) element activated by the visitor: a swing! See Figures 1-14 and 1-15 to get a sense of how a small room can open up into infinite space through reflections of reflections. In [this video](#), if we take the videographer's point of view to be the listener's position and the swinging woman to be the sound source, we can see clearly how the acoustic reflections shift as the source's spatial positioning changes. This dynamic behavior would be quite difficult indeed to convey through auralizations or calculations!

Perhaps most relevant to the method we will cover, Ursula Mayer's *Acoustic Mirror* (2004) installation staged music performances in a mirror room, directly linking the principles of optical and acoustical reflections (see Figure 1-16). Such mirror rooms are, in essence, an enactment of the Image Source Method, if we understand all mirror surfaces to represent acoustically reflective materials in the architectural space we are considering. Such mirror rooms provide a powerful and immediately legible 3D demonstration of how acoustic reflections are distributed depending on the spatial relationships between the source, receiver, and surrounding surfaces.



Figure 1-14: Thilo Frank, *The Phoenix is closer than it appears* (2011), Museum of Modern Art, Denmark. Exterior and entrance.

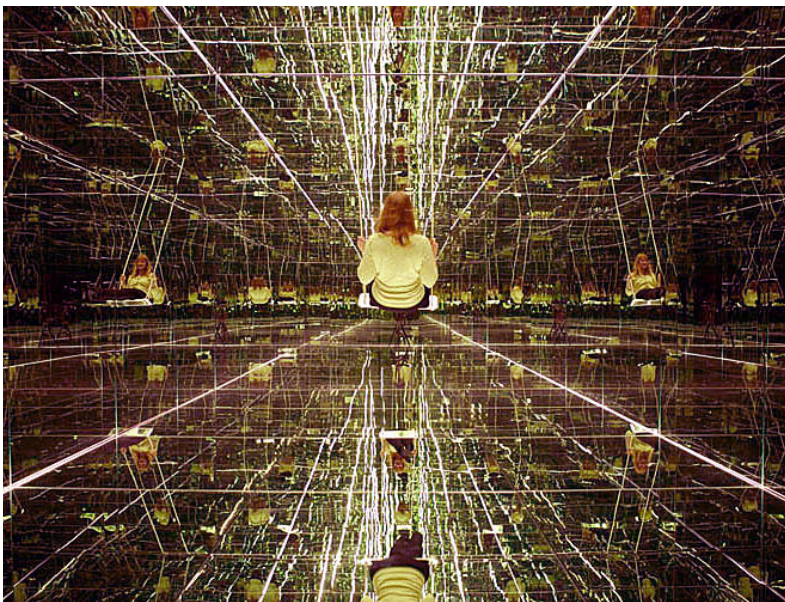


Figure 1-15: Thilo Frank, *The Phoenix is closer than it appears* (2011), Museum of Modern Art, Denmark. Interior.

Figure 1-16: Ursula Mayer, *Acoustic Mirror* (2004), Shedhalle, Zürich. Note that if the camera position corresponds to the listener's position, the optical reflections of the musicians (sound sources) also correspond to acoustic reflections.



Figure 1-17: Artist Doug Aitken covered the walls of his stairwell with mirrors. If we think of him as a sound source, the optical reflections can also tell us about the acoustics of the space.



Artist Doug Aitken covered the walls of his stairwell in mirrors, and in Figure 1-17 it is straightforward to understand three primary points about using the Image Source Method in 3D space:

1. The perceived size of the sound source – in this case, let's pretend Aitken in the blue shirt is speaking – is noticeably expanded by the proximity of reflecting surfaces;
2. Other prominent acoustic reflections can be easily mapped to the room's surfaces by locating the virtual image reflections (note the cluster to the far right in the image); and
3. The photograph's point of view directly corresponds to the receiver/listener's position relative to the source and reflections.

Full-scale mirror rooms are clearly more experiential and revealing of complex 3D reflections than the 2D scale models used by acousticians, but it should be obvious that they are also more demanding and complicated to construct. The question becomes: is it possible to create such reflective behavior – including the spatially complex deeper orders of reflections – within more flexible virtual models? You bet! Some CGI (computer-generated imagery) film artists have even been able to integrate actual actors from a green screen studio into virtual mirror rooms created after the fact in post-production. Have a look at the examples from the 2014 film *Divergent* (© Red Wagon Entertainment, Summit Entertainment) in Figure 1-18 – an impressive proof of concept! An interview with the visual effects supervisor Jim Berney is available [here](#).

For our purposes, the method of rendering acoustic reflections is considerably easier and more straightforward than that CGI example, but before we dive into the funhouse of creating virtual mirror rooms from scratch, some caveats are in order.

Figure 1-18: The CGI artists led by Jim Berney for the 2014 film *Divergent* (© Red Wagon Entertainment, Summit Entertainment) proved that virtual mirror rooms can be created digitally, and even integrated with live actors. Top images: green screen studio stills with live actor. Bottom images: final composites in virtual mirror room.



1.5 Caveats: Light vs. Sound

The Image Source Method allows us to use light as a stand-in for sound. Although both forms of energy are vibrational in nature, of course *light and sound are different*. For one thing, our experience of sound is more intricately tied to perception *over time*, since sound waves travel at a fraction of the speed of light – 1,130 feet per second in air, as opposed to light which is a million times faster. Due to the way light rendering works, in the visualization method which follows we are rendering the reflection orders as they would occur *instantaneously* rather than over time (as with impulse responses and auralization). While doing this departs from the temporal behavior of sound propagation, it allows us to focus on the totality of the geometrical and material relationships which impact the distribution of reflections in a particular space – precisely that domain which is of concern to designers (as well as acousticians).

Acousticians will be quick to point out that the Image Source Method is based solely on specular reflection, and does not account for diffusion and diffraction effects. Furthermore, this method does not allow us to analyze the *spectral character* of reflective surfaces – in practice, materials may be absorptive for some frequencies and reflective for others. For example, thin drywall on studs can be fairly absorptive to low frequencies but reflective to high frequencies. Finally, due to the extremely long wavelengths of low frequencies (from 6-60 feet for deep bass sounds) their propagation behavior in and around architectural spaces produce certain phase and “modal” effects which cannot be represented by the Image Source Method (to be fair, most auralization engines rely at least partially on ISM algorithms and likewise cannot be relied upon for low frequency response below ~250 Hz).

The “illuminating acoustic reflections” rendering method is essentially a visualization of the Image Source Method, *literally*, albeit with considerably more power and flexibility. However, by designating certain architectural surfaces as either optically/acoustically reflective or not, we are of course simplifying the description of acoustics compared to more detailed calculations and auralization. Given the photorealistic capabilities of light rendering

engines, it would be possible to assign diffusely reflective materials (such as brushed metal) to certain surfaces as a stand-in for partially reflective acoustic materials. But at that point I would suggest our light-for-sound analogy would become overwrought. The entire point of this visualization bridge between the designer's 3D modeling experiments and the acoustician's auralizations is that we make *fundamental acoustic behavior accessible* – as determined by spatial relationships, surface geometries, and reflective/absorptive material choices – so that designers can confidently consider acoustics in the early stages of a project. As a design becomes more detailed and issues like low frequency room modes and spectral balance become more important, that is the perfect time to begin dialogue with an acoustician who can perform higher level calculations and complement the visualizations with in-depth auralization.

The hope here is twofold: 1) designers discover a way to integrate sound thinking into their process, and 2) with that heightened sound awareness, more ears – and acousticians! – are at the design table. All architectural spaces stand to benefit from acoustical thinking, and if this method motivates more architects to seek out auralization – perhaps, over time, even developing an aural intuition to match their visual intuition – we will have succeeded indeed. While the following visualization method cannot replace proper training in acoustics (the resources listed in the *References* section are a good start), understanding the principles and techniques involved should make anyone comfortable with 3D modeling and rendering much more comfortable with understanding acoustic behavior. As one gains experience visualizing acoustics while also paying attention to the sounds encountered in architectural spaces, one will develop the ability to interpret spatial reflections and anticipate acoustic behavior from visualizations.

With that said, let's fire up the rendering engines.



+



2 Method

The method demonstrated here consists of five general steps:

1. *Create a 3D virtual model of the architectural space under study.*
2. *Apply a mirror material to acoustically reflective surfaces.*
3. *Apply a light-emitting material to sound source objects.*
4. *Set the camera to correspond to the listener position.*
5. *Render and process the image.*

How is this anything more than just using contemporary modeling software to recreate the physical optical models that acousticians experimented with nearly a century ago? First, 3D virtual models are much faster and cheaper to build, and are far more efficient and flexible when creating myriad variations for testing purposes. There is also the obvious advantage that most architectural projects will already have a working 3D virtual model which can easily be adapted for these acoustical imaging studies.

As we will soon see, light rendering in virtual models also allows us to do things that physics forbids, such as dictating how deep the light reflections can penetrate a mirror. We can also “paint” any sound source object with a light-emitting material, making it simple to set up models with one or more sources which remain distinct within reflections.

Finally, scaled physical models make it rather difficult to occupy the listener’s position. With virtual modeling and rendering, we have complete

control over the camera's position and orientation. Again defying physics, we can also dictate that the "photographer" doesn't populate and obscure the mirror reflections, and we can also create 360° optics which are simply impossible with actual camera lenses.

All of these advantages amount to a simple but powerful, intuitive and highly flexible method for identifying architectural and material effects on acoustics. These effects will shape the important early reflections and thus the "voice" of the space under study.

To be most approachable and effective, this method has been devised with the designer's typical toolset in mind. This toolset consists of a 3D virtual modeler, a light renderer, and an image editor. Though the general steps could be applied with a wide variety of software packages found in architectural practice, for the demonstrations here we will use the most popular 3D virtual modeler, light renderer, and image editor: Trimble's SketchUp, Chaos Group's V-Ray, and Adobe's Photoshop, respectively.

2.1 SketchUp Modeling

It can be argued that SketchUp is the "iPod" of software 3D modelers. It is far from the most advanced modeler, but it has disrupted the industry by being incredibly easy to learn and fast to use. Not unlike the way low-quality MP3s stumped the audiophile scene, SketchUp caught its more sophisticated (and expensive) brethren off-guard by pursuing a "sketchy" rather than photorealistic look. But don't let the fun, sketchy styles fool you: SketchUp supports many third party rendering packages which allow you to generate highly detailed and photorealistic imagery from SketchUp.

The real power of SketchUp lies in its incredibly wide adoption by designers and 3D folks of all stripes: from hobbyists and woodworkers to visual artists and game developers, from theater set designers and filmmakers to archaeologists and crime scene analysts, from interior designers to architects to urban planners. This broad base of users numbers in the *millions*, unlike other specialized 3D modelers whose user bases perhaps number in the thousands. This enormous community of users coupled with

SketchUp's open scripting platform (anyone can write their own custom tools and plugins using Ruby Script and share them freely with others) results in an active and supportive community of users at all skill levels. The quantity of free tutorial resources and add-on tools is astounding. Furthermore, the searchable 3D Warehouse is a treasure trove of models and components built and shared by others which can be freely loaded into your models. Need a grand piano source in your study but don't feel like modeling one from scratch? [3D Warehouse](#) has several options for you.

For those who are new to SketchUp, follow the links in the *References* section at the end of this document to find plenty of help and instruction. One can expect to be modeling confidently in SketchUp within a few hours of opening the program for the first time. For the demonstrations here, I will be using SketchUp 2016 Pro, however the free version SketchUp Make will also work just fine, as will older versions of SketchUp (this project began way back with SketchUp 8).

While we will not be covering the basics of SketchUp modeling here, there are a few pointers worth mentioning to make your modeling habits more precise and efficient. Follow the links here for specific tutorials:

- Use [Groups](#) and the [Outliner](#) to separate and organize your geometry and material variations.
- Use [Layers](#) to control the visibility of different light sources (i.e. sound sources, in this method) so that different source positions and combinations can be tested easily.
- Use the [Position Camera](#) tool to precisely locate and orient the point of view (i.e. listening position).
- Use [Scenes](#) to save specific camera/listener locations for controlled comparisons between different scenarios.

Later we will look at a variety of models for demonstrating the possibilities of this method. For now, to cover the relevant modeling and rendering settings, I will work with a 3D model of the 2D plan drawings used in the *Introduction* section to illustrate the Image Source Method.



2.2 Light Rendering with V-Ray

Chaos Group's V-Ray for SketchUp plugin is not the least expensive rendering engine available for SketchUp, but it is widely regarded as the best and most popular one due to its power, flexibility, and quality. Even if you have never heard of V-Ray, you have likely seen plenty of imagery created by it. V-Ray's photorealism and animation capabilities have made it popular for everything from real estate brochures to international design competition visualizations, from car commercials to CGI effects in blockbuster films.

V-Ray was chosen for this method primarily because it is already familiar to many designers. For those who are new to V-Ray, it is available at significantly reduced rates with education licenses, and a [free trial period](#) is also available. Another reason why V-Ray is so popular is that it is available as a plugin for several of the most commonly used modelers, which means skills transfer easily between platforms. The free rendering engine [Kerkythea](#) is popular with SketchUp users on a limited budget. This method should translate to renderers other than V-Ray just fine, the principles are the same. For the demonstrations in this document, I will be using V-Ray 2.0 for SketchUp.

Similar to auralization packages, V-Ray's interface and operation are admittedly complex, with some rather arcane settings. Rest assured you do not have to possess a complete command of every button and checkbox in V-Ray before using it effectively. We will cover precisely those steps and settings which are relevant for this method. While V-Ray may seem overwhelming to newcomers at first, unlike the limited number of auralization packages that integrate with SketchUp, V-Ray is a much larger enterprise and thus has undergone far more development and smoother integration with SketchUp.

[V-Ray's online manual and tutorials](#) are very useful references. There are several other independent tutorials and forum communities easily found online which are also valuable resources for learning and troubleshooting V-Ray. That said, it is the goal of this document to take you through the specific steps of this method in a comprehensive and straightforward way, such that everything you need to know to get started will be covered here.

Figure 2-1 (opposite): To demonstrate the rendering steps, we will use a 3D model of the 2D plans used in the *Introduction* to illustrate the Image Source Method.



Figure 2-2: On this toolbar, we will only need the three leftmost buttons for editing materials and rendering options. The blue “R” button is the rendering button itself.

Figure 2-3: The only lights from this toolbar that we will need are the Omni Light (left) for general lighting and the Sphere light (right) for designating sound sources.

Before we proceed through the specific steps and settings, let’s take a look at the two V-Ray for SketchUp (VfS) toolbars. We will not require much of V-Ray’s extensive capabilities, and will only need the buttons outlined in red in the above figures. On the **VfS: Main Toolbar**, we will use (from left to right) the **Material Editor**, the **Option Editor**, and certainly the blue “R” render button. On the **VfS: Lights** toolbar, we will only need the **Omni Light** (left) and the **Sphere Light** (right). The differences between these two seemingly similar light sources will be explained soon.

We will now go through the steps and settings for creating the materials, lighting, and camera optics to turn our SketchUp models into visualizations of acoustical reflections.

2.2.1 Reflective Mirror Material

After creating a SketchUp model for studying acoustic reflections, the first rendering step is to decide which surfaces will be designated as reflective, then create and apply a mirror material to those surfaces. Acousticians will need to forgive this rather binary approach of designating surfaces as either reflective or not. Recall that the point here is to reveal how different architectural surfaces are contributing generally to the strength of early reflections. For our purposes, we only need to consider those common building materials which are hard and inert enough to reflect much of the middle range of the audible spectrum (the range where human voices and most musical instruments reside): concrete, steel, glass, masonry, tile, hardwood floors, plaster and gypsum walls. Materials which are primarily absorptive or

diffusive we will designate as non-reflective: carpet, soft furniture, acoustical ceiling tiles, drapery, and other geometrically complex textures.

In the demonstration model here, I will designate the walls as reflective and the floor as non-reflective. I will also designate the ceiling as non-reflective so that I can omit it and let the sky provide general illumination. (Later we will look at how to create general illumination in a completely sealed interior.)

- In the **VfS: Main Toolbar**, click on the leftmost “**M**” button to open the **Material Editor**.
- Under the Materials List, right click on **Scene Materials** and select **Create Material > Standard**. See Fig. 2-4.
- Right click on the newly added **DefaultMaterial**, select **Rename Material**, and call it “**Mirror_walls**”. See Fig. 2-5.
- Right click on **Mirror_walls**, and select **Create Layer > Reflection**. Notice the addition of a **Reflection** settings tab on the right side. See Fig. 2-6.
- Click the triangle to the left of **Mirror_walls** to expand the layers, right click on the **Diffuse** layer and select **Remove Layer**. See Fig. 2-7.

If you click the **Preview** button on the upper left of the **Material Editor**, you will now see what appears to be a sphere made of clear glass (see Fig. 2-8). Whenever you add a new reflection layer to a material in V-Ray, it automatically loads a Fresnel texture map, which is useful for various glass and metal materials when you want the reflection behavior to follow Fresnel laws. You are already familiar with Fresnel behavior: if you stand in front of a clear window with equal amount of light on both sides, you will see through the window. The further you move away from perpendicular to the glass, the more reflective and less transparent the window becomes. At very shallow angles to the glass, the window will approach the reflectivity of a mirror. This variable reflective behavior which depends on viewing angle is described by an Index of Refraction (IOR) number. Because the automatically loaded Fresnel map has IOR 1.55 which is typical for glass, and because we removed the diffuse layer underneath the reflection layer, our material looks like clear glass for the moment.

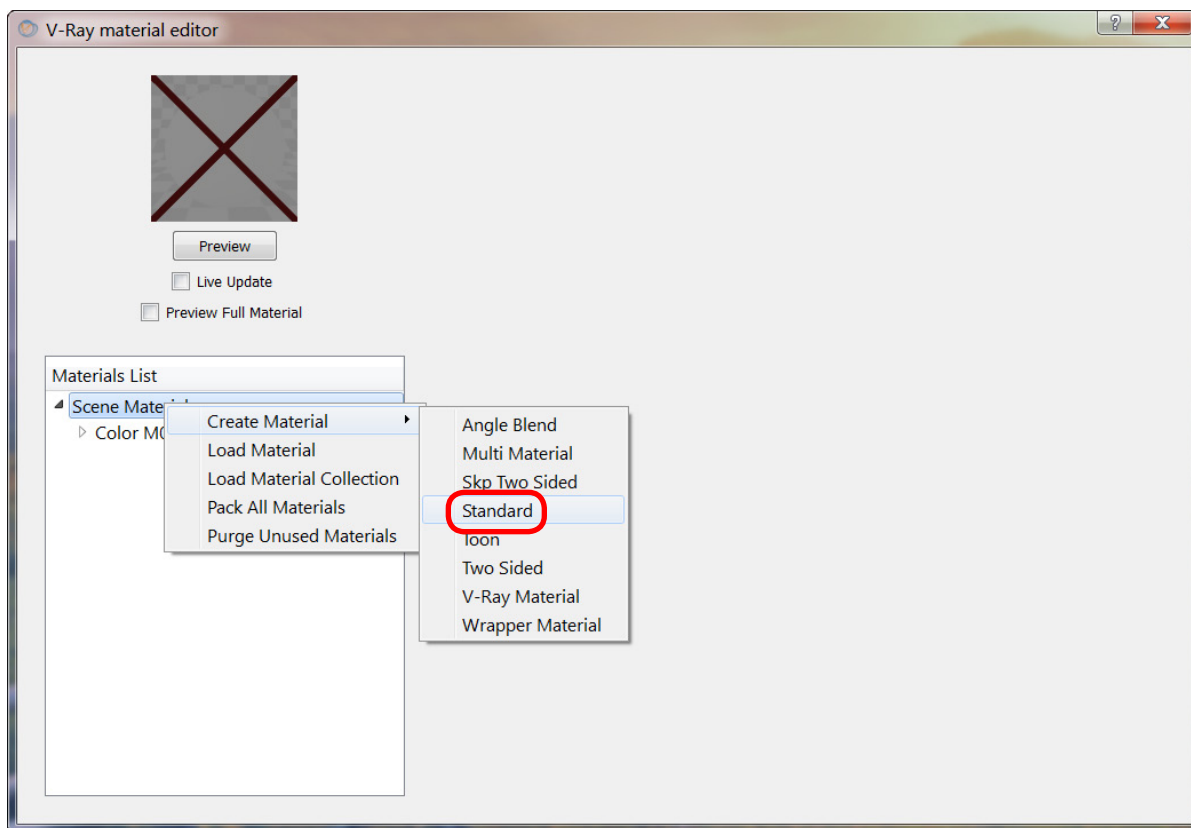


Figure 2-4: Create a new material in V-Ray's **Material Editor** by right clicking on **Scene Materials** and selecting **Create Material > Standard**. This will add a new **DefaultMaterial** to the list of materials in the current model.

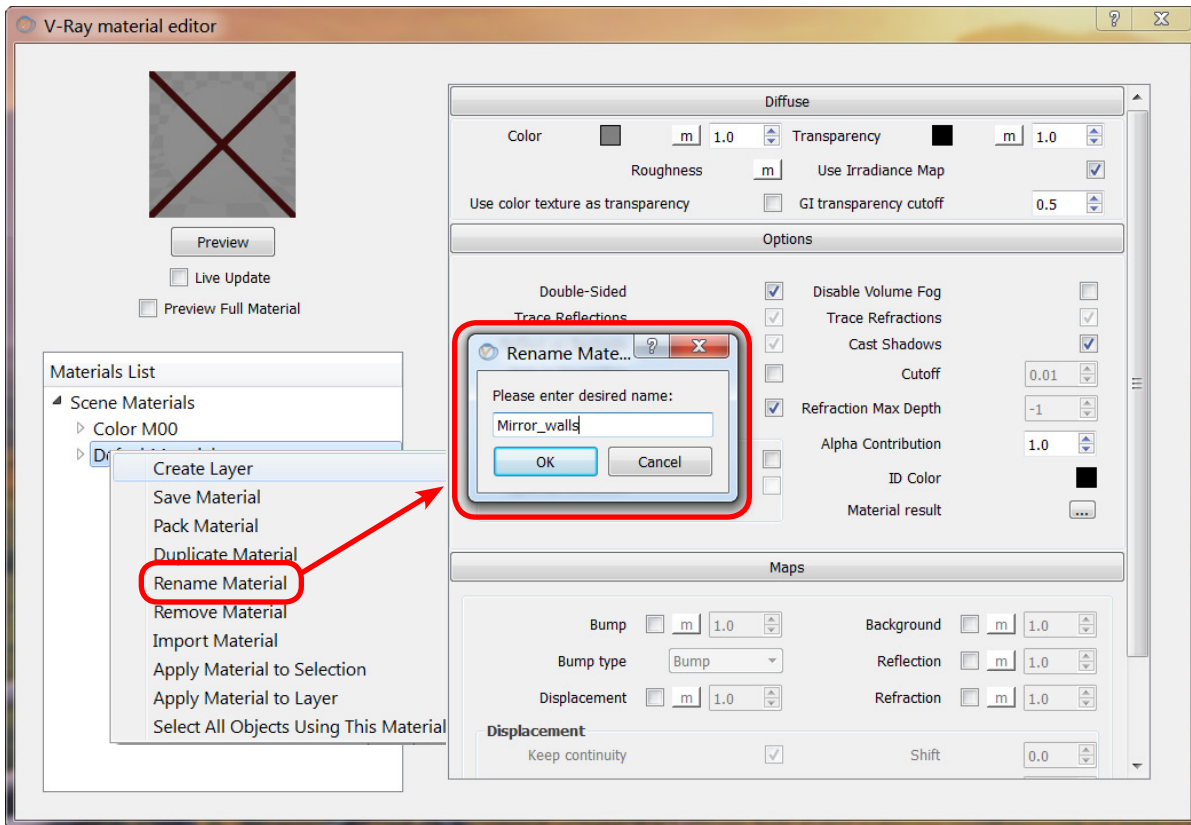


Figure 2-5: Rename the new material by right clicking on **DefaultMaterial** and selecting **Rename Material**. In the pop-up window, change the name to **Mirror_walls** and click **OK**.

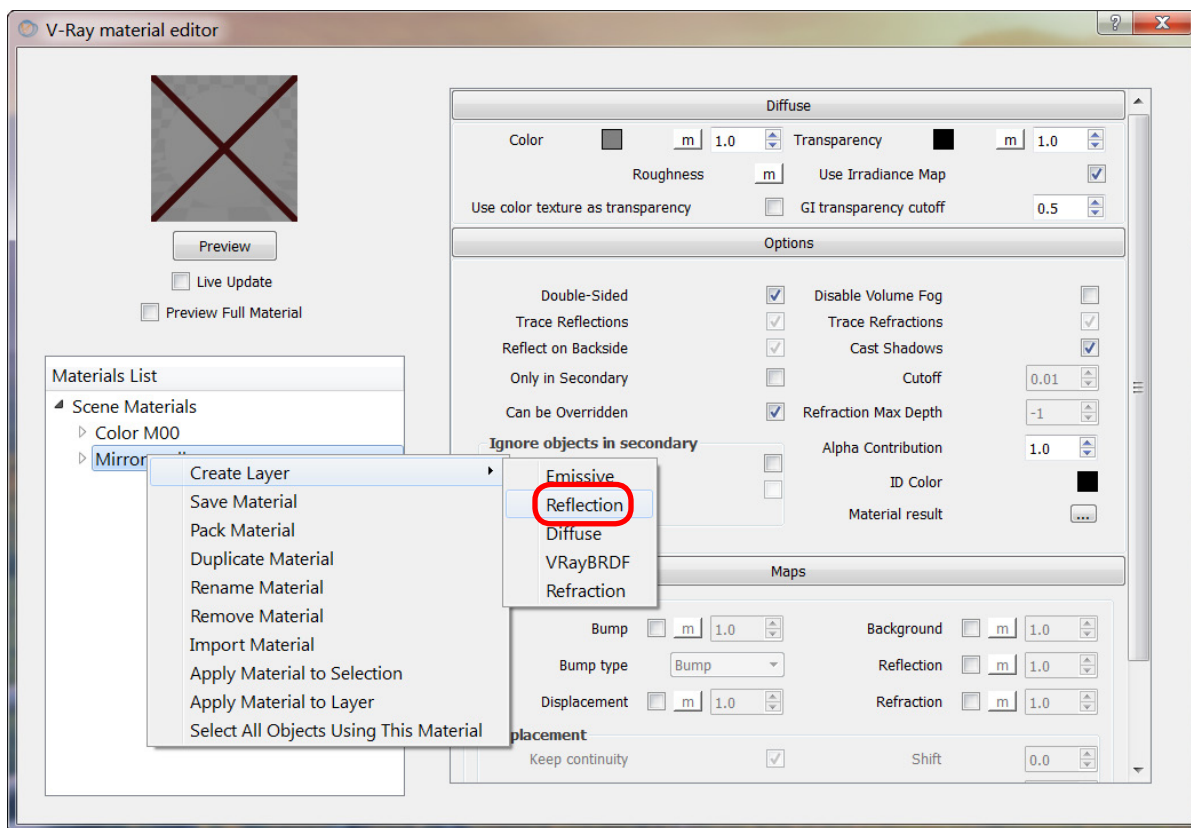


Figure 2-6: Add a reflection layer by right clicking on **Mirror_walls** and selecting **Create Layer > Reflection**.

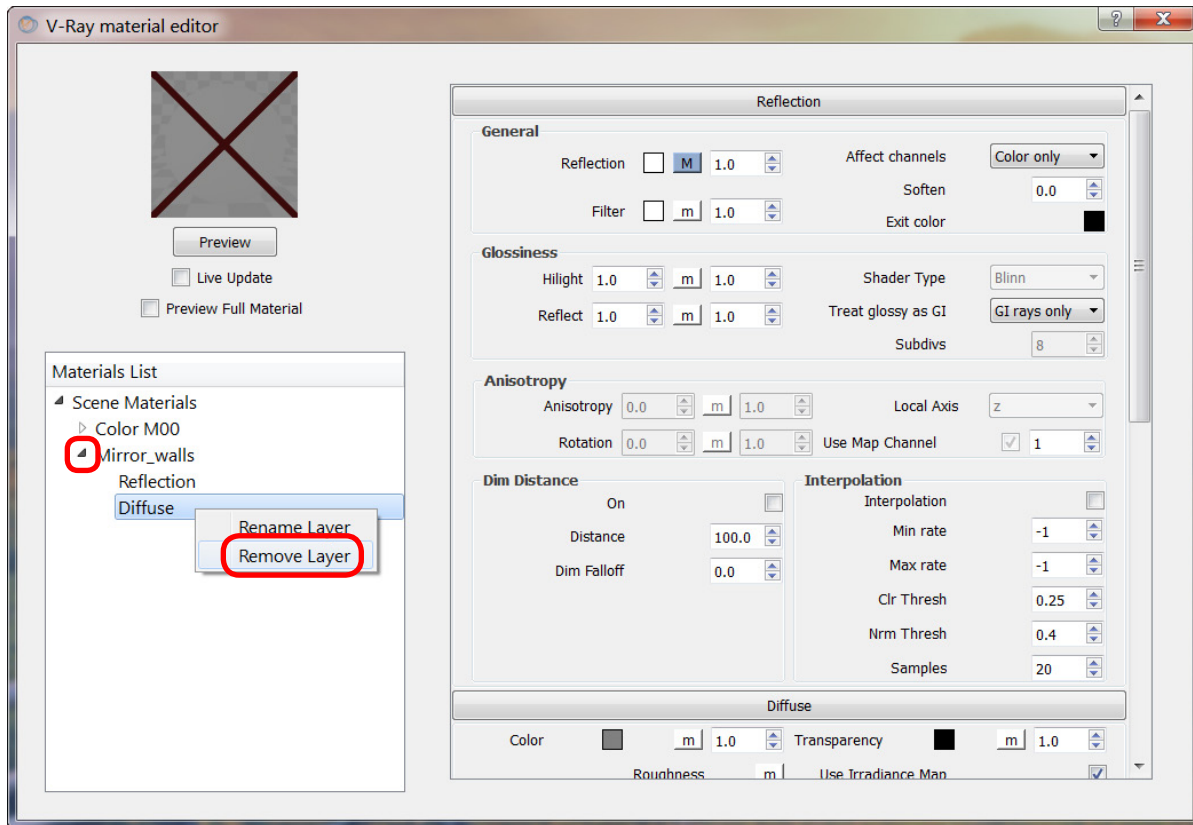
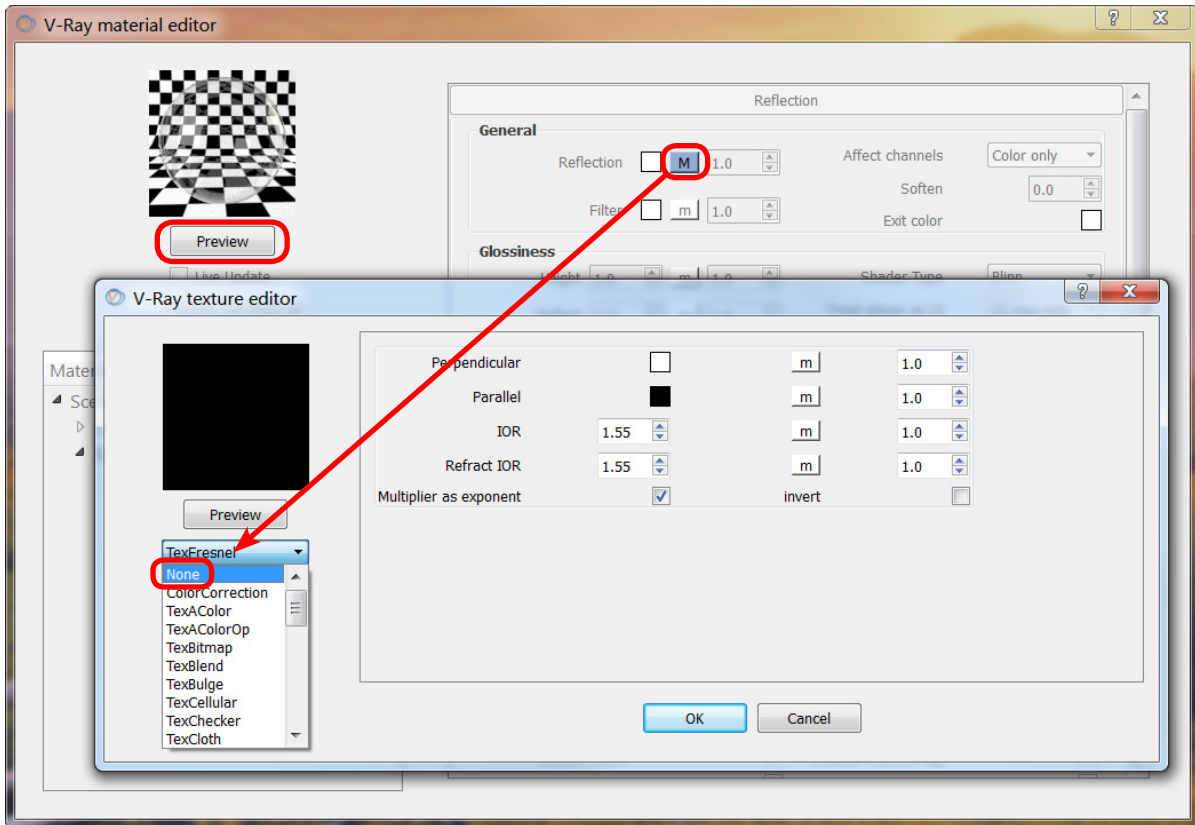


Figure 2-7: Remove the **Diffuse** layer by clicking the triangle to the left of **Mirror_walls** to expand the material's layers, right clicking on **Diffuse**, and selecting **Remove Layer**.



There are two ways to make our material behave more like a mirror. The first and simplest method would be to remove the Fresnel map altogether. In V-Ray, material maps are loaded into slots behind buttons designated “m” with a white background (no map applied) or “M” with a blue background (Map applied). To remove the Fresnel map, in the **Material Editor** under the **Reflection** tab’s **General** subsection, click the blue “M” next to **Reflection**. This brings up the **Texture Editor** window, where the dropdown menu on the left will let you select **None** instead of the **TexFresnel** texture map. Click **OK** on the **Texture Editor** window, and then click **Preview** again on the **Material Editor**, and you will see the sphere update to a mirror material. Without the Fresnel

Figure 2-8: Clicking the **Preview** button reveals that our material currently looks like clear glass. By clicking the “M” button to bring up the **Reflection** layer’s **Texture Editor**, we could choose **None** instead of the **TexFresnel** map. This would give us a perfectly reflective material, however we can make our mirror material appear more realistic...

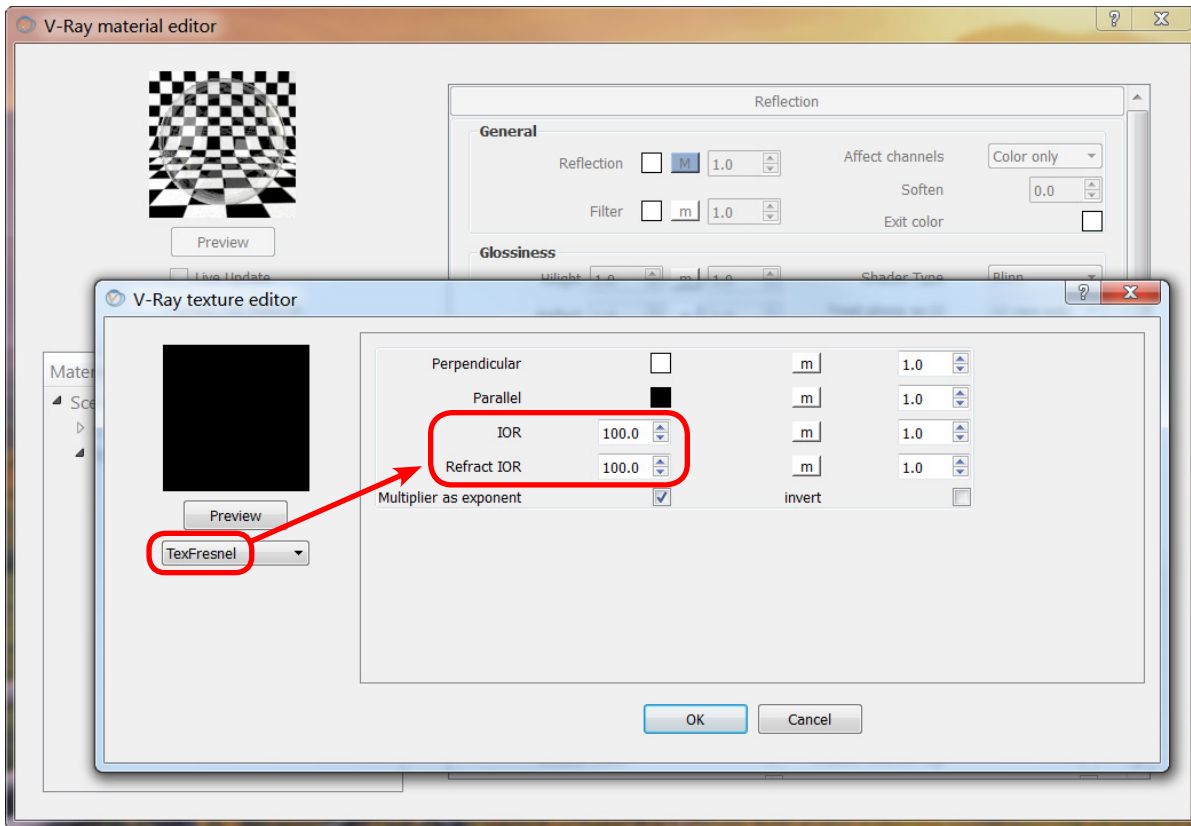


Figure 2-9: ...by keeping the **TexFresnel** texture map and adjusting the Index of Refraction settings (both **IOR** and **Refract IOR**) to the maximum value of **100.0**. Click **OK**. This creates a highly reflective material that exhibits a slight amount of reflectivity fall-off, just as in real-world mirrors. This will make the spatiality of the reflections more legible.

texture map, the material is completely reflective regardless of viewing angle: an ideal mirror.

The only catch here is that we never see an ideal mirror in the real world! Using a perfectly reflective material in our model would indeed exhibit geometrically correct reflections, but would appear somewhat strange because our experience tells us to expect some slight imperfection and reflectivity fall-off with real-world mirrors.

The second method of arriving at our mirror material makes a slight but appreciable adjustment to the realism of the mirror material's reflective behavior, and thus makes the spatiality of the reflections more legible. If you

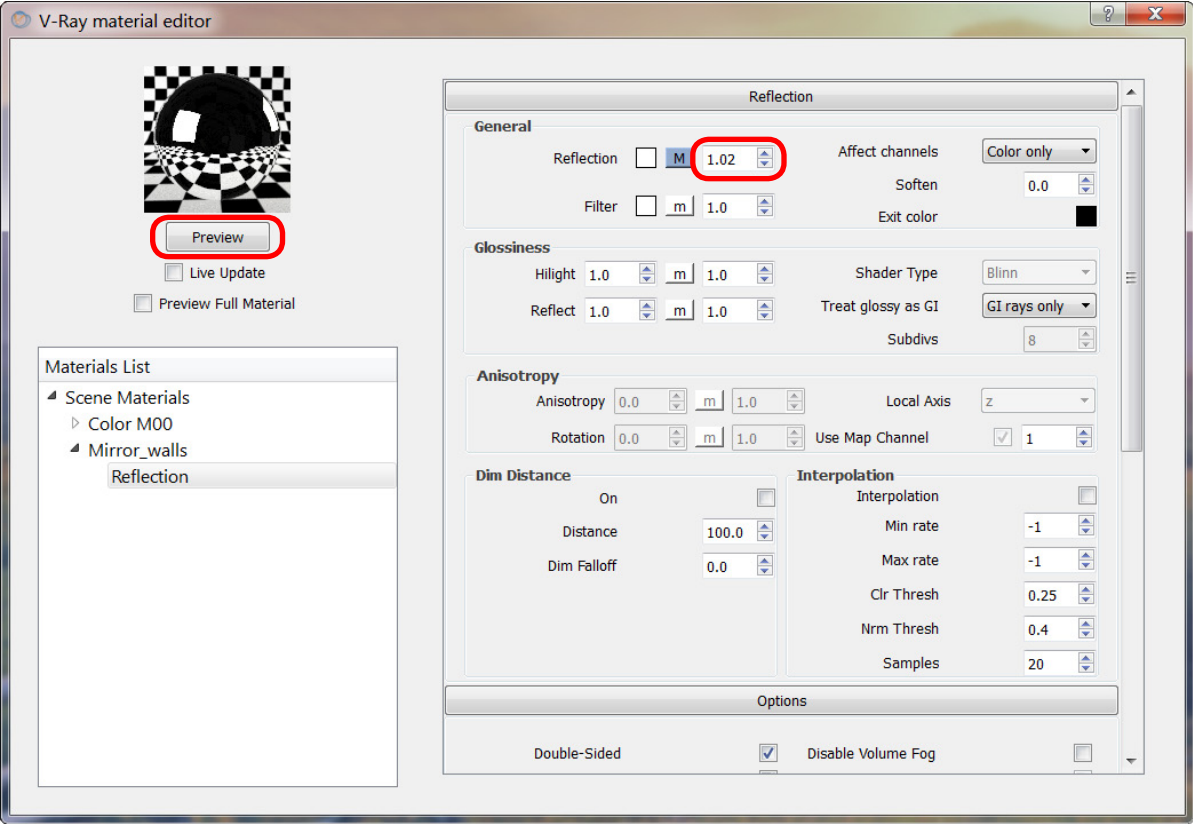


Figure 2-10: Adjust the **Reflection** strength in .01 increments to “clean” the clarity of reflection depth in the mirror material. Click the **Preview** button again and the sample sphere now shows our final mirror material.

removed the Fresnel texture map previously, restore it by clicking the “m” next to **Reflection** on the **Material Editor’s Reflection** tab, then select **TexFresnel** in the **Texture Editor’s** dropdown menu and click **OK**. Then:

- On the **Texture Editor**, set both the TexFresnel map’s **IOR** and **Refract IOR** to the maximum value of **100.0**. Click **OK**. See Fig. 2-9.
- On the **Material Editor**, the numeric value to the right of the blue “M” of the **Reflection** map controls the strength of the applied map. If you find that the mirror material’s reflectivity falls off too much and you would prefer the reflections’ clarity to penetrate deeper, adjusting this setting in .01 increments will have the effect of “cleaning” the reflection depth. In my examples here, I am using a setting of **1.02**. Click **Preview** to see our final mirror material applied to the sample sphere. See Fig. 2-10.

There are two ways to apply a material created in V-Ray to the surfaces in your SketchUp model. The first way is to use SketchUp’s selection tool to select all the faces to which you want to apply the V-Ray mirror material, and then in V-Ray’s **Material Editor** right click **Mirror_walls** and choose **Apply Material to Selection**. However, for whatever reason, this is not as fast as applying the material directly in SketchUp using the paint bucket tool:

- In V-Ray’s **Material Editor**, left click **Mirror_walls**, and notice that in SketchUp’s **Materials** panel the corresponding swatch has also been selected. Click the **Mirror_walls** swatch in SketchUp’s **Materials** panel to activate the paint bucket tool.
- Using SketchUp’s paint bucket tool, click all the faces in the model to which you want to have the V-Ray mirror material applied.

To create a test render with the mirror material applied, I click the “R” render button on the **VfS Main Toolbar**. Figures 2-11 and 2-12 show the bare SketchUp model and the V-Ray test render respectively (other V-Ray settings are involved, which we will get to shortly).

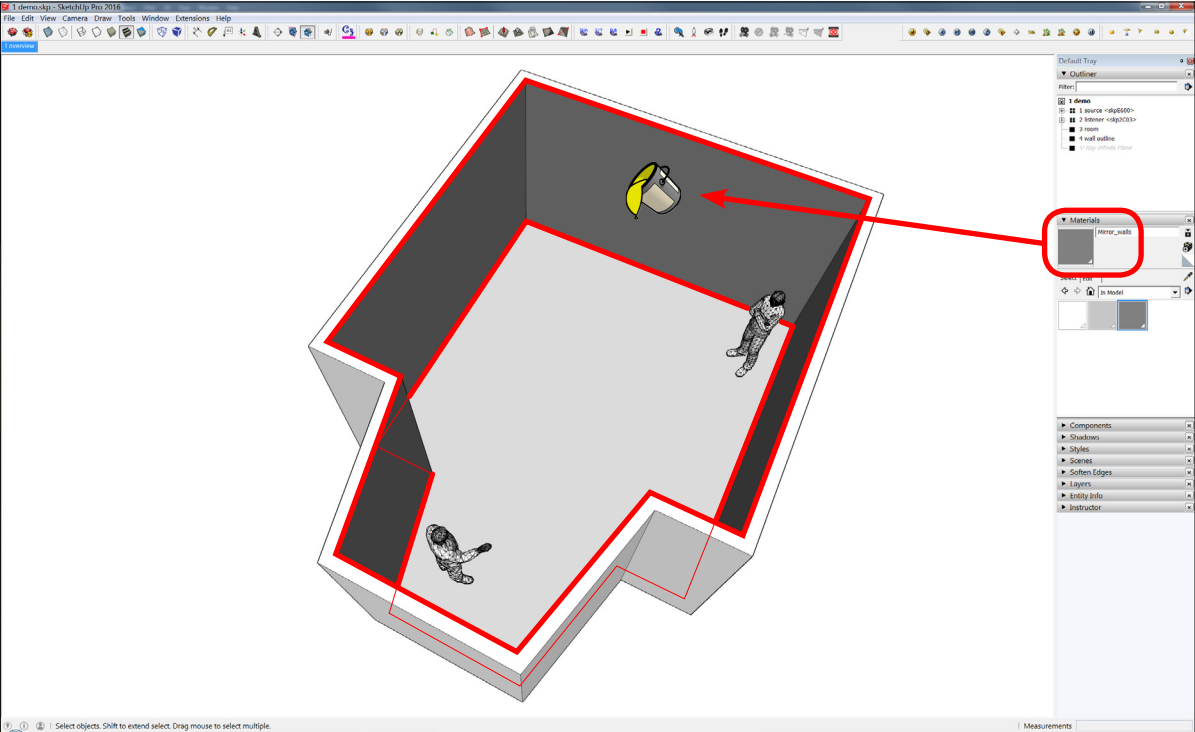


Figure 2-11: The fastest way to apply materials created in V-Ray to surfaces in the SketchUp model is to select the **Mirror_walls** material in the V-Ray **Material Editor**, which will automatically update the material swatch in SketchUp's **Materials** panel. Then click the **Mirror_walls** swatch in the SketchUp **Materials** panel to activate the paint bucket tool and click on surfaces to apply the material. For this demo model, I have painted just the walls with the mirror material, outlined in red.



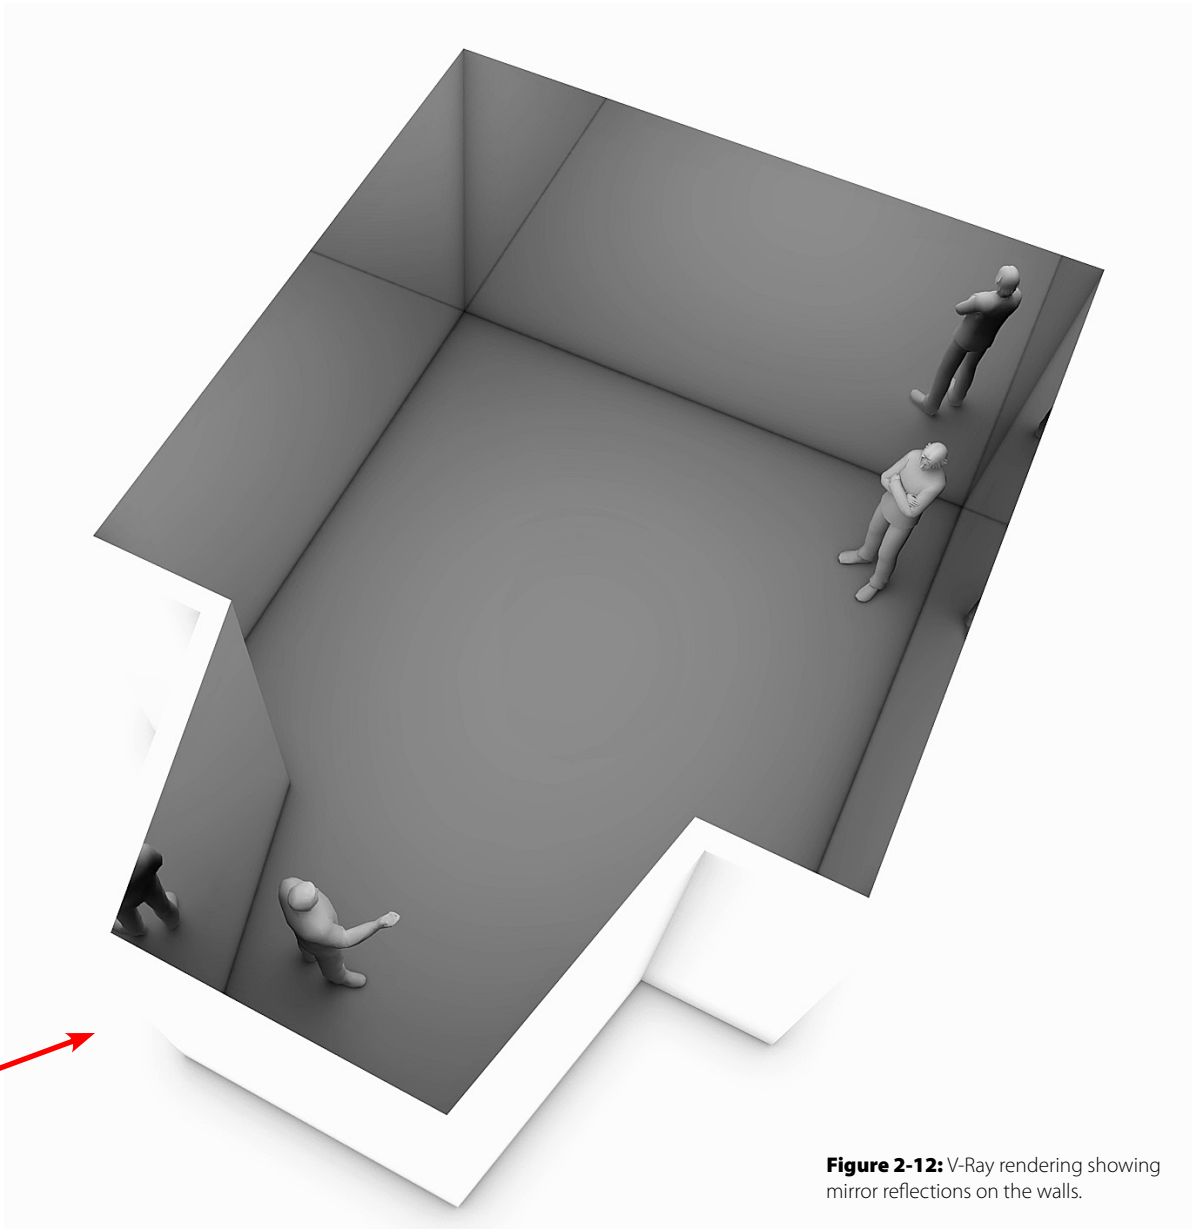
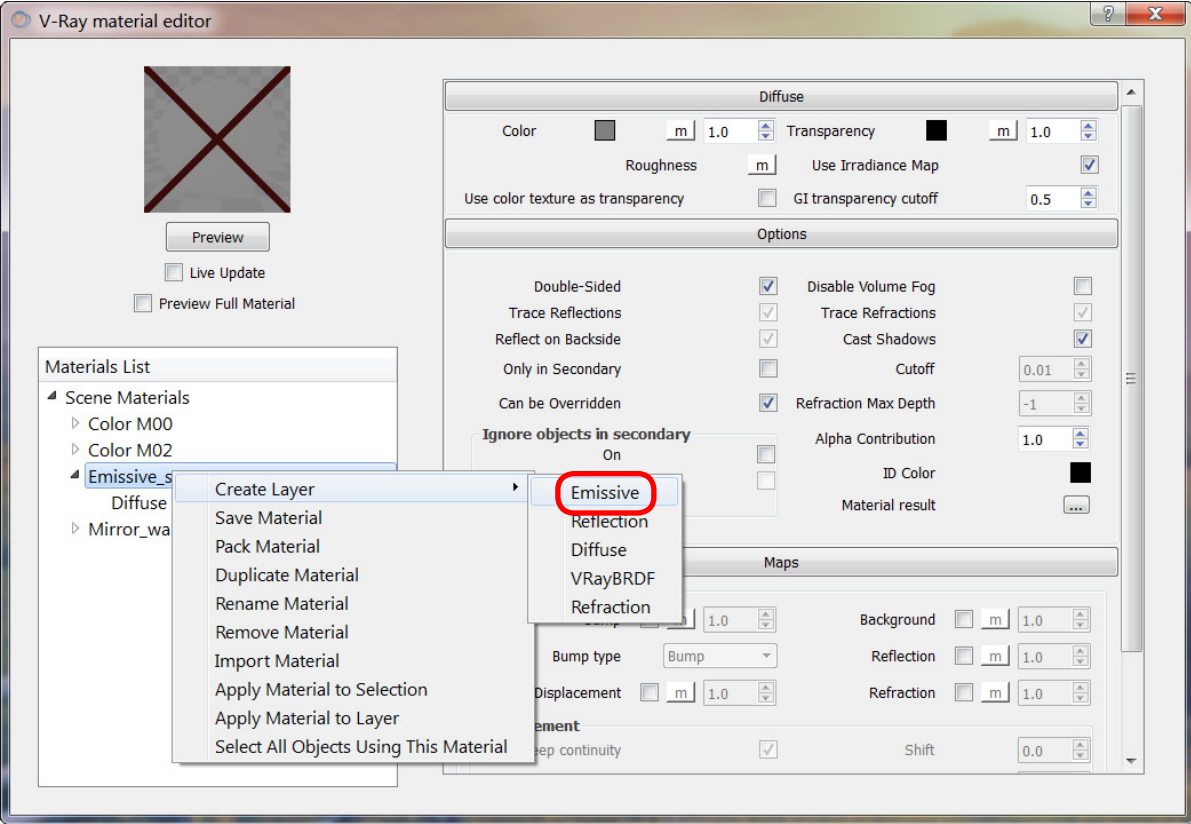


Figure 2-12: V-Ray rendering showing mirror reflections on the walls.



2.2.2 Emissive Source Material

For the source located in the demo model's alcove, I imported a 3D model of a speaking man from www.archive3d.net. I have created a subgroup of his head geometry so that I can more easily treat it separately from his body. We will be making his head a light source by creating a luminous "emissive" material and painting his head with it. We will then have control over the luminous properties of our chosen sound source just like we would have over any other light object. V-Ray's emissive material makes it easy for us to paint any object (human head, piano, loudspeaker, vehicle, etc.) as a light/sound source which will be easy to see in our model's reflective surfaces.

Figure 2-13: After creating and renaming a new material, add an emissive layer by right clicking on **Emissive_source1** and selecting **Create Layer > Emissive**.

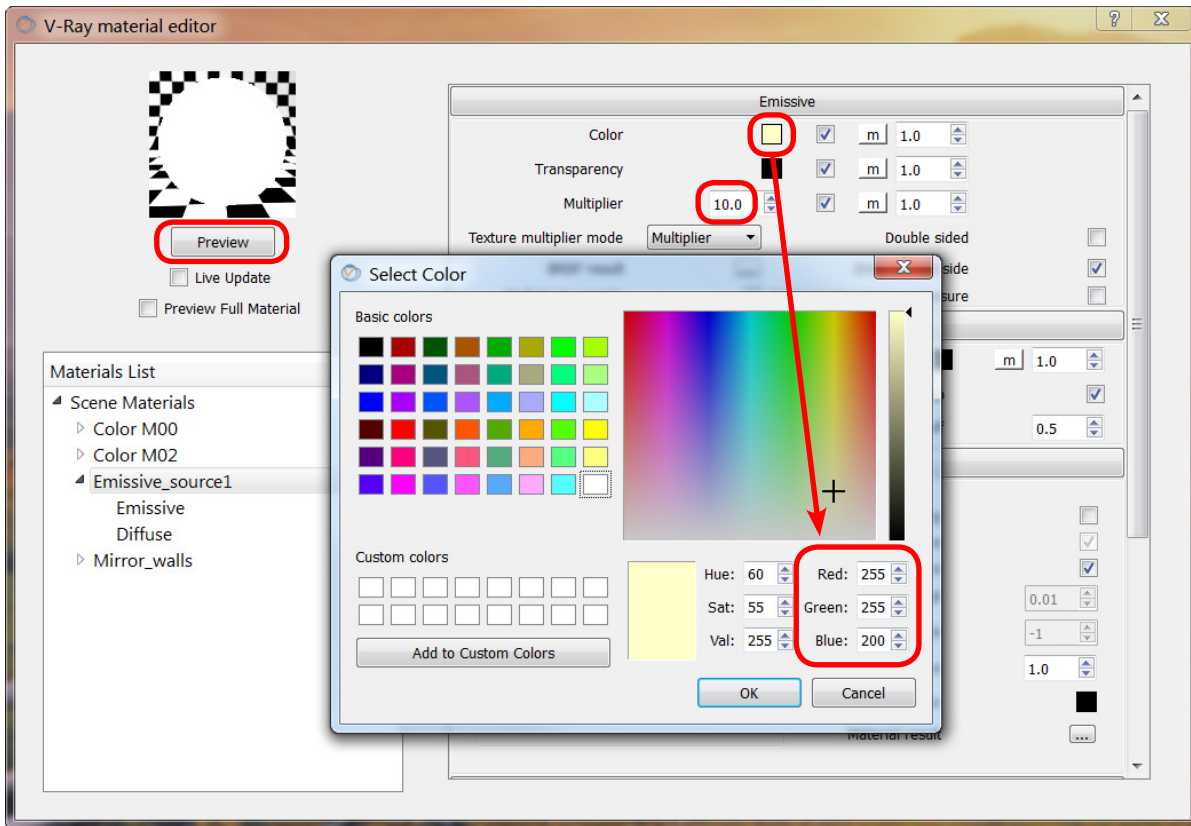


Figure 2-14: You can set the color of the emissive material's light by clicking the swatch to the right of the **Color** setting on the **Emissive** tab, and adjusting the RGB values in the **Select Color** pop-up window. The **Multiplier** controls the intensity of the light. Clicking the **Preview** button shows the emissive material applied to the sample sphere, however the luminous behavior will be more apparent in the context of a rendering.

- Under the Materials List, right click on **Scene Materials** and select **Create Material > Standard**. (Refer to Fig. 2-4.)
- Right click on the newly added **DefaultMaterial**, select **Rename Material**, and call it "Emissive_source1". (Refer to Fig. 2-5.)
- Right click on **Emissive_source1**, and select **Create Layer > Emissive**. See Fig. 2-13.
- **Optional:** give the light a specific color. On the **Emissive** tab, click the **Color** swatch to open the **Select Color** window. To give my emissive material a slightly yellowish incandescent appearance, I have reduced the amount of blue in the RGB settings. See Fig. 2-14.

- For now, I have increased the **Emissive** layer's **Multiplier** to **10.0**. This directly controls the brightness of our **Emissive_source1** material. See Fig. 2-14. We will return to this setting in a following section about exposure.
- *Tip:* set the **Diffuse** layer's color to be the same as the **Emissive** layer's color (in Fig. 2-14, this setting is hidden behind the **Select Color** pop-up window). Since the **Diffuse** color is what appears in the SketchUp model, this makes it easy to determine which objects have the emissive material applied in SketchUp before rendering. See Fig. 2-15.
- Clicking the **Preview** button shows our emissive material applied to the sample sphere, and indicates that it is indeed bright enough to turn it white. It will appear more like incandescent light in a test render. See Fig. 2-16.

Figure 2-15: The **Diffuse** layer color is apparent in the SketchUp model before rendering.

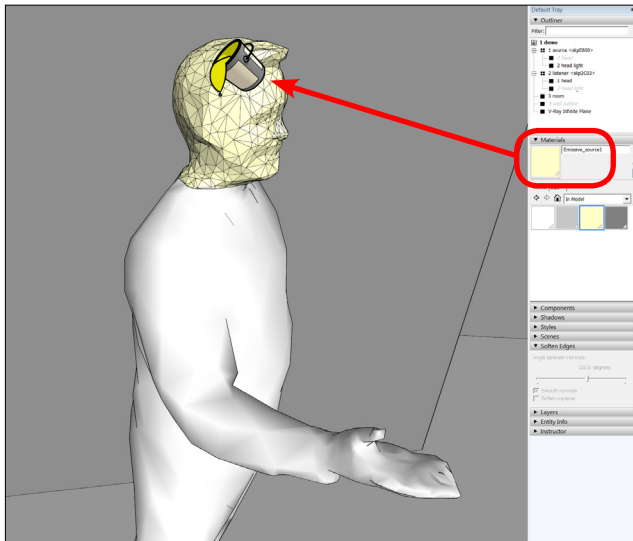
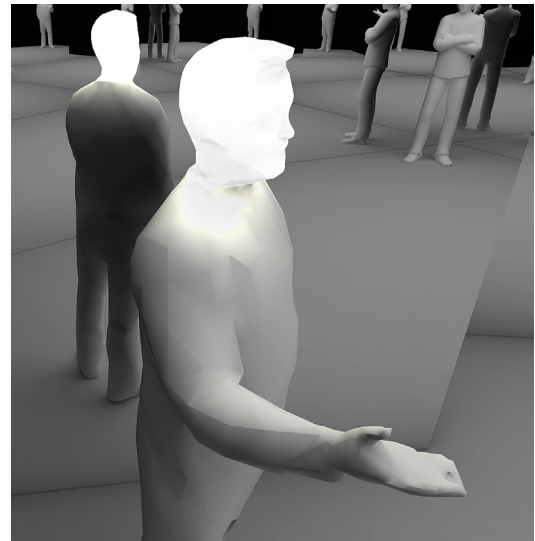


Figure 2-16: Test render of the light/sound source with our **Emissive_source1** material.



2.2.3 Receiver Position Camera

As you can see in the demo model, I have also imported another 3D model of a person to occupy the listener's position. He will be useful in demonstrating some other techniques later, but for now I am only using him to mark the location of the receiver's position (this can also be done using guides in SketchUp). For the first round of renderings, I only want to see the light/sound source and the room. Here we can defy real-world physics by removing the camera/listener from all the reflections. Try doing that in a mirror room! This is useful when you want to focus on the source reflections without them being cluttered by reflections of the receiver. To do this, simply hide the SketchUp group containing any objects you don't want to see in the reflections.

Figure 2-17: Setting a field of view (FOV) which corresponds (more or less) to human vision is done by selecting **Camera > Field of View** in SketchUp, and then entering “35 mm” in the Value Control Box (VCB) in the bottom right corner of the program window.

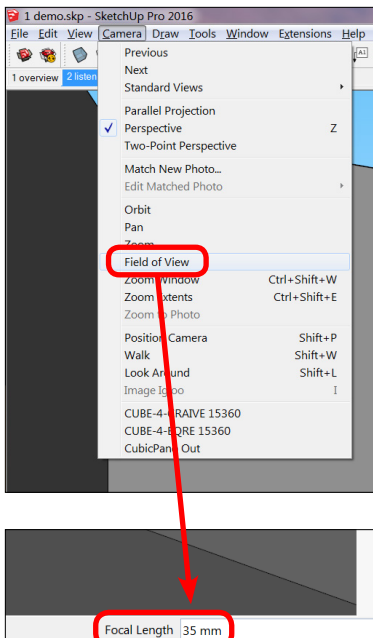
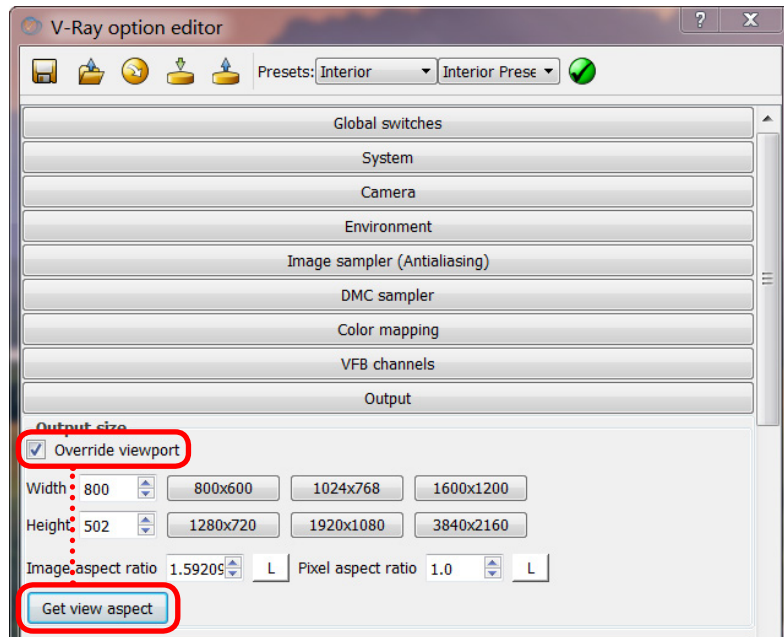


Figure 2-18: Use V-Ray's **Output** tab to get the image aspect ratio from SketchUp.



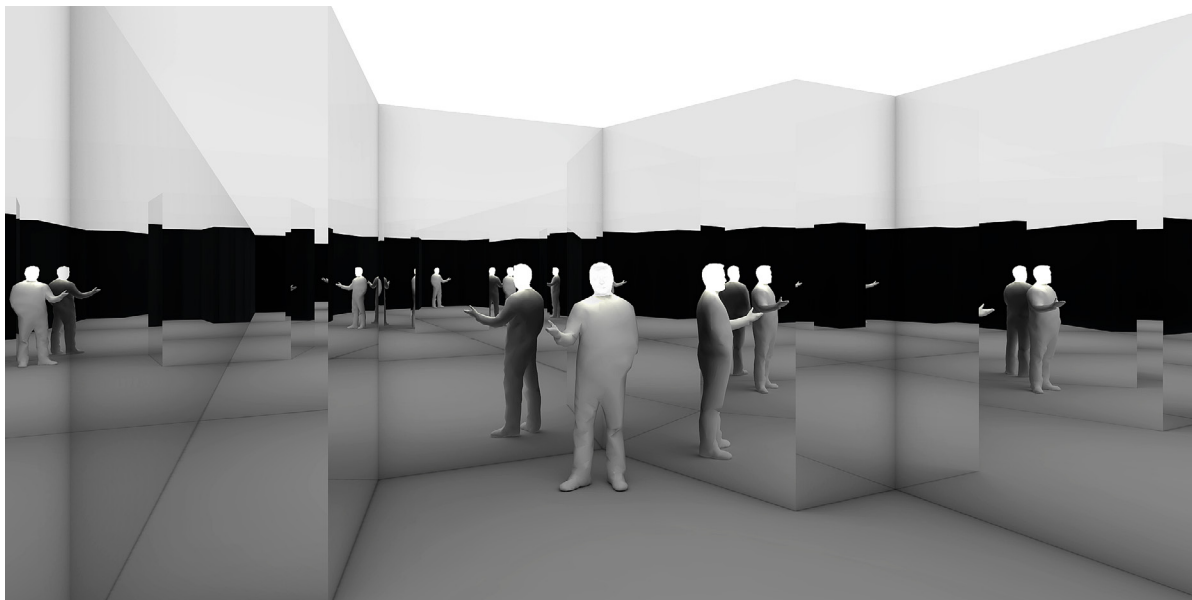
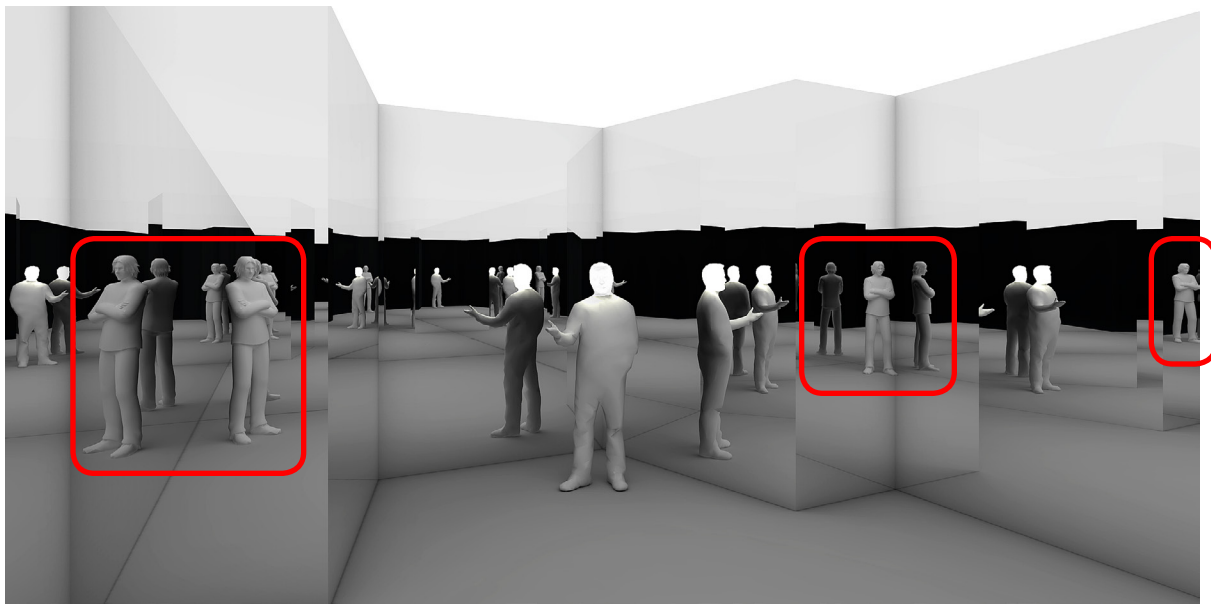


Figure 2-19 (opposite): Here we have used the 3D model of the listener to position and orient the camera's field of view relative to the sound source. But notice how our own reflections (outlined in red) populate the reflections as well...

Figure 2-20 (opposite): ...In this rendering we see how easy – and clarifying – it is to hide the listener model (i.e. *ourselves*) from the reflections, allowing us to focus on the distribution of the sound source reflections.

- The first step is to position the camera in the receiver's location at the appropriate height, and aim its view toward the source (if you do not know how to do this, refer to the camera positioning tutorial link in the *SketchUp Modeling* section).
- The next step is to set the desired field of view (FOV) in SketchUp: **Camera > Field of View** using the Value Control Box (VCB) at the bottom right of the modeling window to enter in the desired horizontal field of view. This can be in degrees by entering "deg" after the number, or in focal lengths corresponding to typical lenses by entering "mm" after the number. To match the typical field of human vision, I have entered "35 mm" in Fig. 2-17.
- To make V-Ray's output match the FOV of the SketchUp window, on the **V-Ray Option Editor's Output** tab, either 1) uncheck **Override viewport** (which will default to the native resolution of the SketchUp window on the particular monitor screen), or 2) if you want to customize the resolution (smaller for faster test renders, or larger for higher quality final renders), be sure to check **Override viewport** to set the desired resolution and then click the **Get view aspect** button. See Fig. 2-18.
- To remove unwanted reflections, simply hide those groups in SketchUp by right-clicking on them in the **Outliner** and selecting **Hide**. The difference between Figs. 2-19 and 2-20 shows what happens when we hide ourselves (i.e. the listener) from the reflections, allowing us to focus on just the sound source's reflections. The power of virtual modeling to bypass real physics (when it is convenient and/or clarifying) should be apparent!

2.2.4 Rendering Quality Settings

Many of the settings buried within the V-Ray Options Editor's different tabs are concerned with the controlling the quality of light, shadow, and material appearance in the final render. For our purposes, we only need to concern ourselves with controlling two settings: 1) selecting a global quality preset, and 2) adjusting the output resolution.

- At the top of the **V-Ray Options Editor**, click the drop-down menu next to “Presets” and select **Interior**. From the next drop-down menu to the right, select a preset based on your needs: the “Test” and “Low” settings are perfect for conducting test renders at the fastest rendering speeds, while the slower “High” and “Very High” settings should be used for final renderings. After making your preset selection, be sure to click the **green check button** to activate the preset.
- On the **Output** tab, you can affect the rendering speed and quality by adjusting the output resolution **Width** and **Height**. Obviously, lower resolution settings will have less detail but will render faster than higher resolution settings. It is common to set these values quite low for test renders, then change to a much higher resolution for the final rendering. If you want to match the viewport in the SketchUp window, remember to click the **Get view aspect** button. See Fig. 2-21.

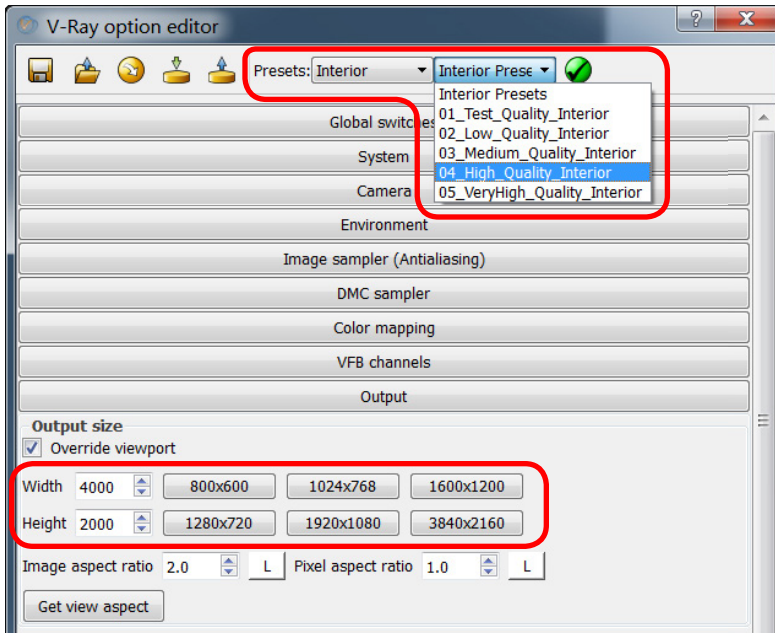


Figure 2-21: The simplest way to control rendering quality is through the **V-Ray Option Editor**’s presets: use the left drop-down menu to select the **Interior** category, then use the right drop-down menu to select a preset depending on the need for speed vs. quality. Be sure to click the **green check button** to apply the preset selection. The other variable which affects the speed and quality of the rendering is the resolution, which is set on the **Output** tab: the higher the resolution, the finer (but slower) the final rendering will be.

2.2.5 Saving Renderings



Figure 2-22: Clicking the blue “R” render button on the **VFS: Main Toolbar** begins the rendering procedure.

Once you click the **blue “R” render button** on the **VFS: Main Toolbar** (Fig. 2-21), V-Ray automatically opens two new windows. The smaller **V-Ray Progress Window** (Fig. 2-22) shows you a handy progress bar and notes about where in the process the rendering engine is. The **V-Ray Frame Buffer** (Fig. 2-23) is where you watch your rendering cook. Depending on the speed and memory of your computer, you will see a number of rendering “buckets” working their way around the image as the ray-tracing algorithms do their work.

- Once the **V-Ray Progress Window** progress bar is showing 100% complete, in the **V-Ray Frame Buffer** click the **single disk icon** to save the rendering. See Fig. 2-23.
- Name the file accordingly. A good practice is to indicate in the filename pertinent information about source and receiver locations, and the number of reflection orders (more on that next).

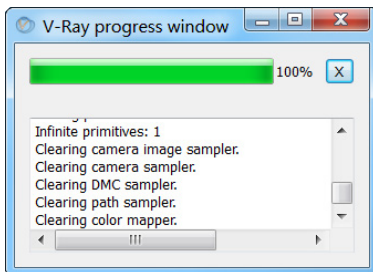
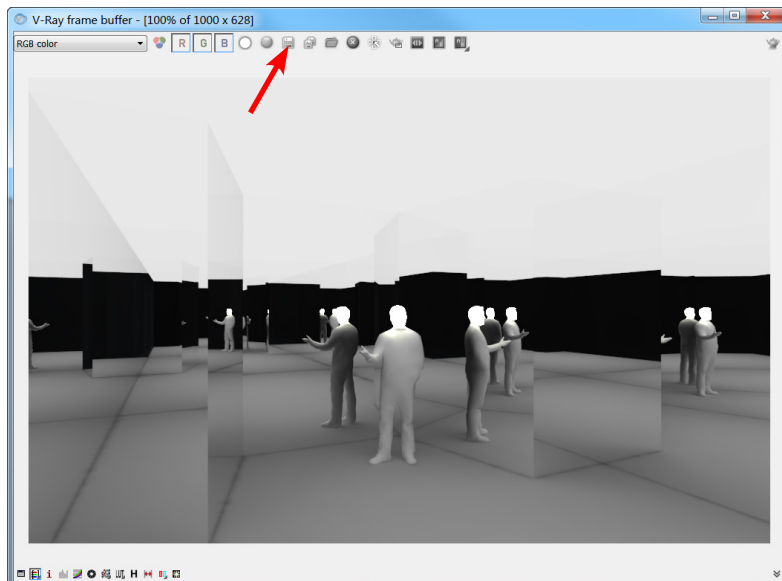


Figure 2-23 (above): The **V-Ray Progress Window** automatically opens when you begin a new rendering, and lets you know how far along in the process the renderer currently is.

Figure 2-24 (right): The **V-Ray Frame Buffer** also automatically opens upon starting a new render and allows you to monitor the rendering’s “cooking” process. When complete, simply click the disk icon to save the image.



Rendering the image is the time-consuming part. Saving out the completed rendering shouldn't take more than a few seconds. Now that we know how to set up and generate renderings, let's look at some advanced uses of V-Ray to visualize acoustic reflections.

2.2.6 Controlling Reflection Orders

Did you notice in Figures 2-19 and 2-20 that our reflections don't seem to penetrate as deeply as those in the artists' infinity rooms? There appears to be a black wall in the distance, where our mirror surfaces arbitrarily decide to stop behaving like mirrors... Exactly! This is one of the most powerful "hacks" of this visualization method, and it certainly cannot be achieved in real-world physics.

Not unlike auralization software which uses the Image Source Method in its algorithms to calculate the spatial distribution of reflections, V-Ray needs a way to manage computational efficiency when rendering reflective materials. Under the **V-Ray Option Editor's Global Switches** tab, there is a **Max depth** setting – with this option checked, the user can dictate how many orders of reflections the render will process. For typical visualization work, you would only use as many orders as necessary to achieve acceptable realism. For our purposes, this parameter allows us to precisely visualize the spatial contribution of different acoustic reflection orders. See Figs. 2-26 through 2-29.

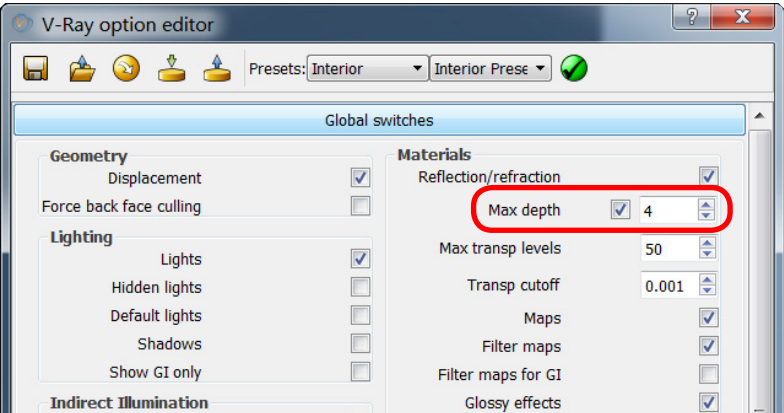
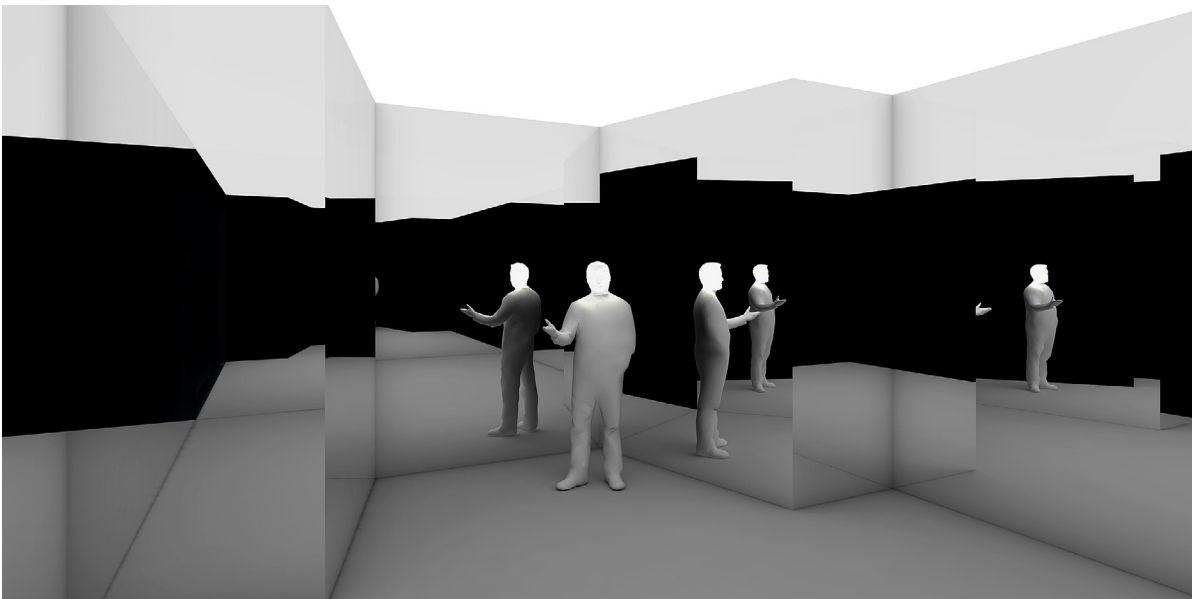
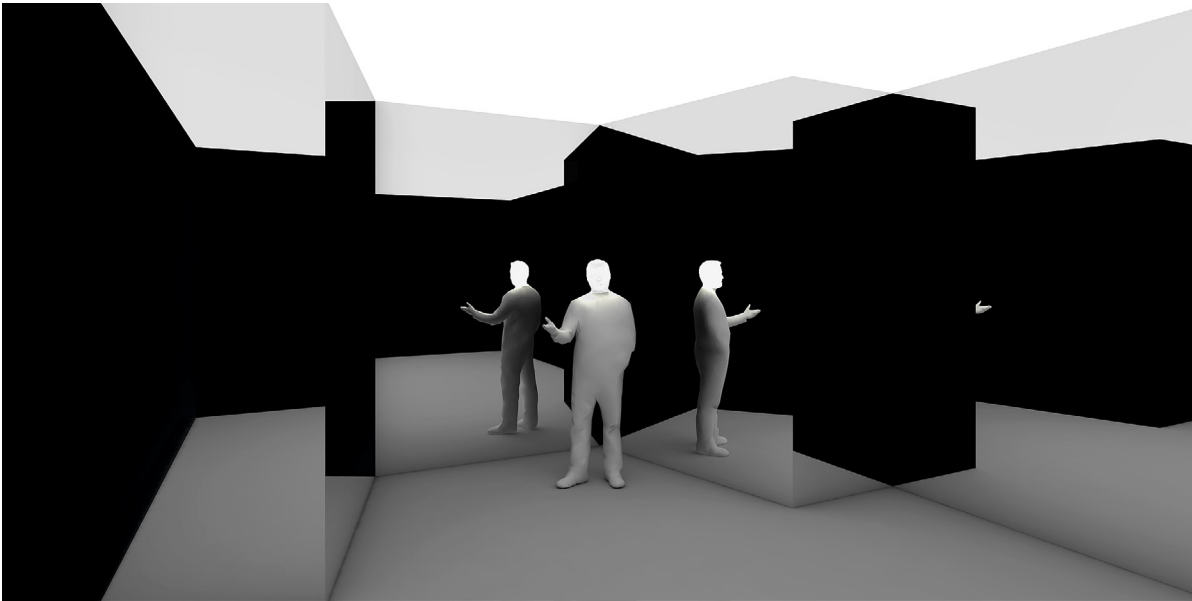
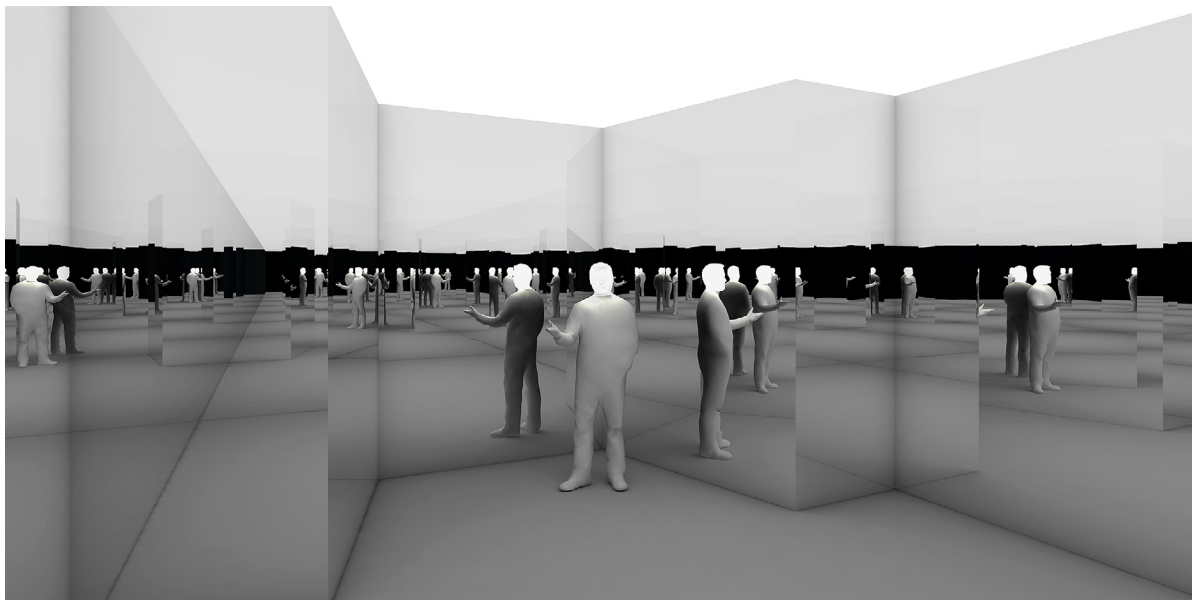
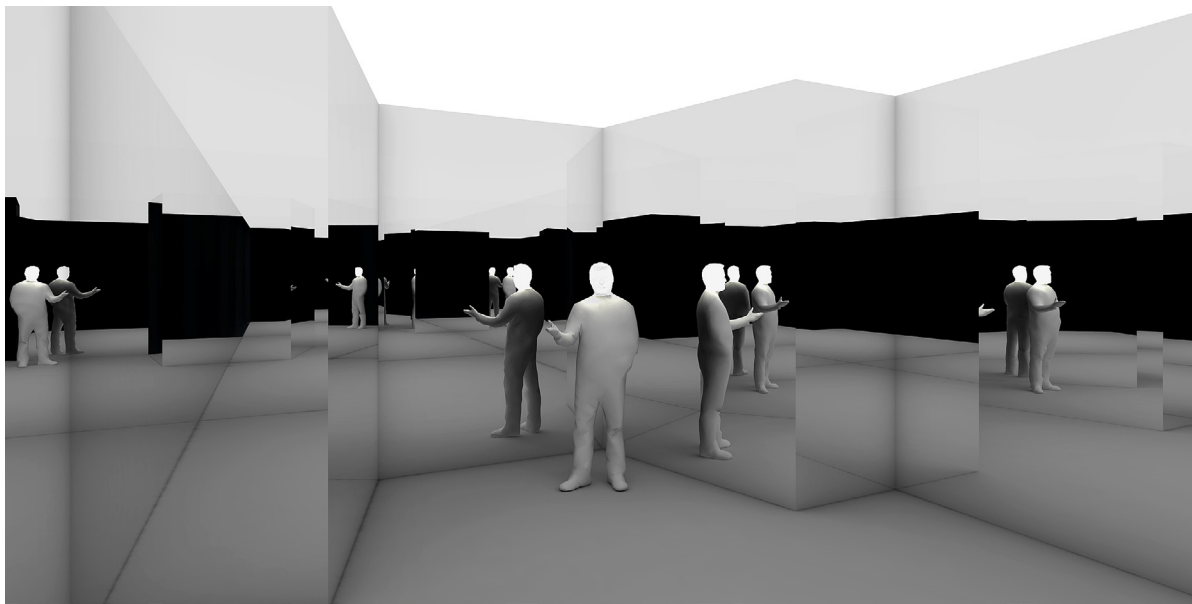


Figure 2-25 (left): The **V-Ray Option Editor's Global Switches** tab has a **Max depth** setting which allows us to precisely control the number of reflections orders. Simply check the box and select a suitable number depending on what you are trying to discern.

Figure 2-26 (opposite above): One order of reflections.

Figure 2-27 (opposite below): Two orders of reflections.





2.2.7 Color-Coding Multiple Sources

Another unique capability of this visualization method entails creating differently colored emissive materials for different sources, so that we can clearly see how the distinct reflections are distributed per source. For this example, I will turn our second 3D model of a person – the one we were previously using as the listener position – and paint his head with a new red emissive material. A new camera/listener is positioned so that we can see both sources in the space. Creating the second emissive source material is fast and easy using the previous source material as a starting point.

Figure 2-28 (opposite above): Four orders of reflections.

Figure 2-29 (opposite below): Ten orders of reflections.

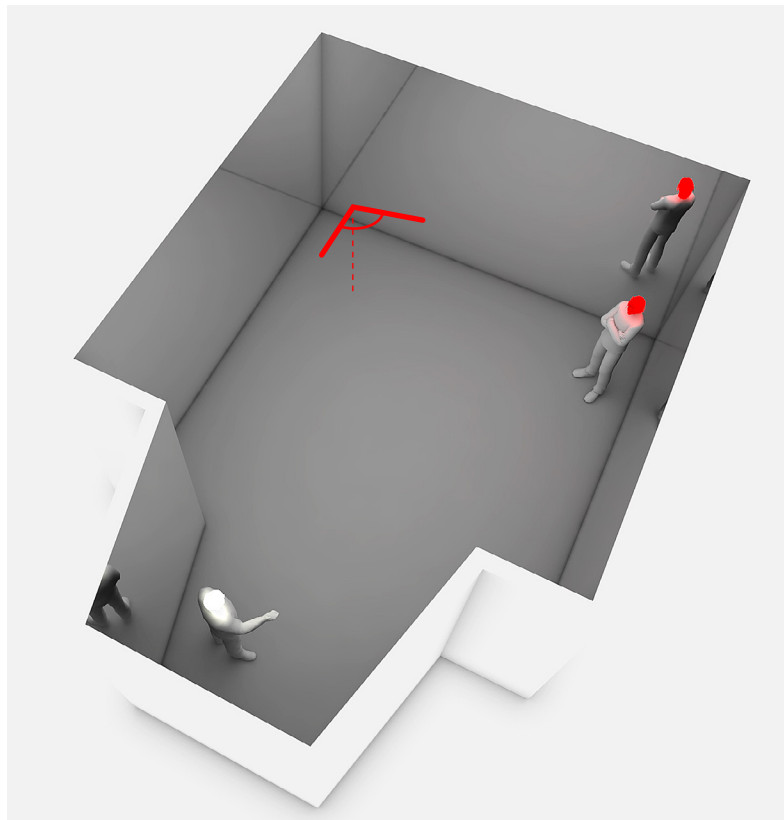
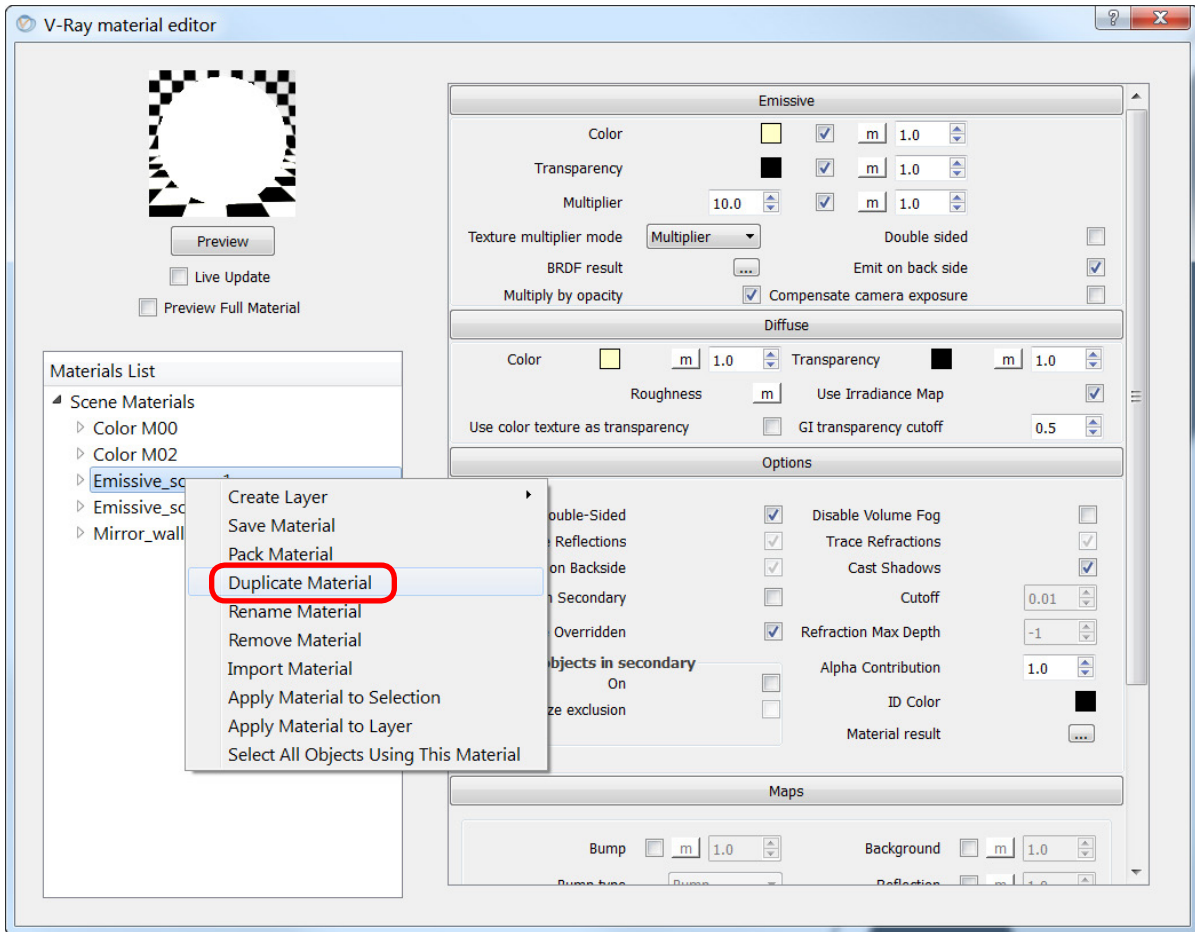


Figure 2-30: We will convert our listener into a second sound source in our space and give this speaker a different emissive color to distinguish his reflections. Note that we have also changed the camera/listener position to the viewing angle represented in red.



- Open the **V-Ray Material Editor** and find our **Emissive_source1** material in the **Materials List**. Right click on this material and choose **Duplicate Material**.
- This conveniently creates a new material named **Emissive_source2** with luminous properties which are identical to the material we already configured in the earlier section on creating emissive source materials.

Figure 2-31: Duplicate the Emissive_source1 material to quickly create an Emissive_source2 material with the same luminous properties.

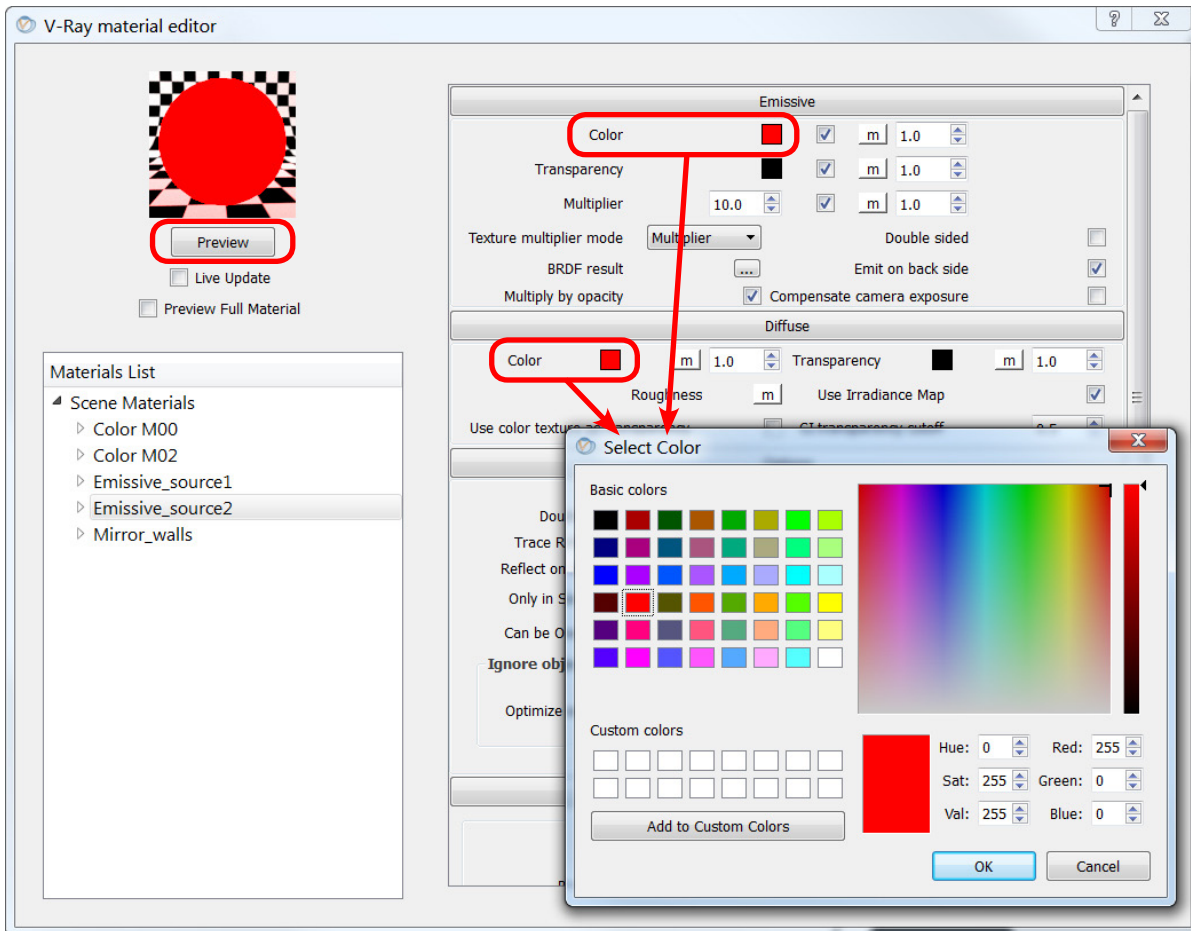
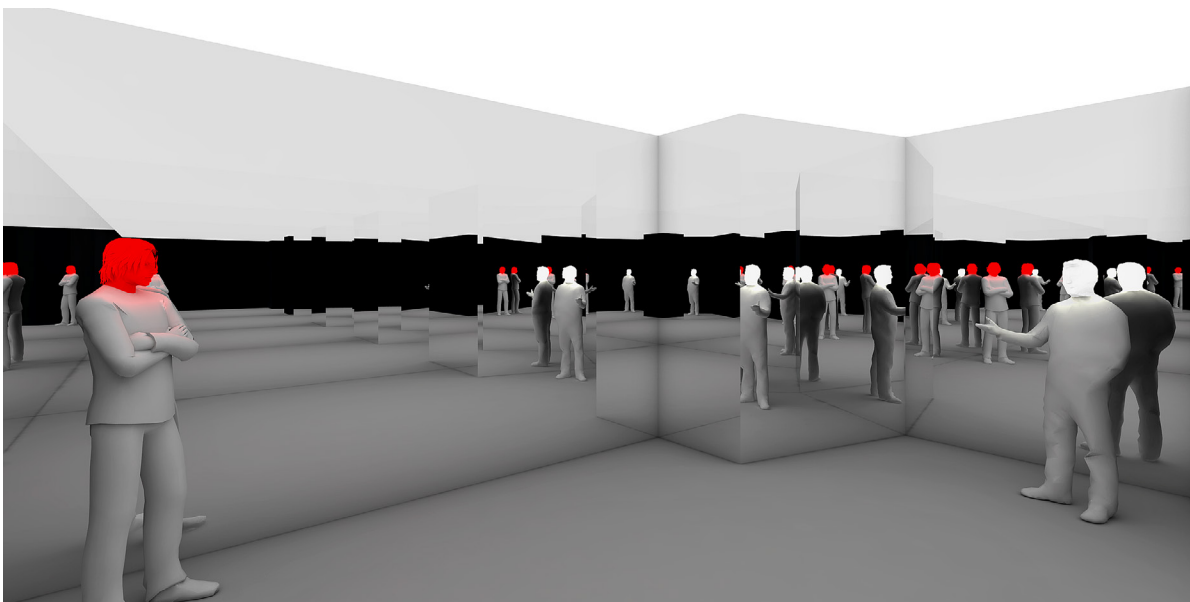
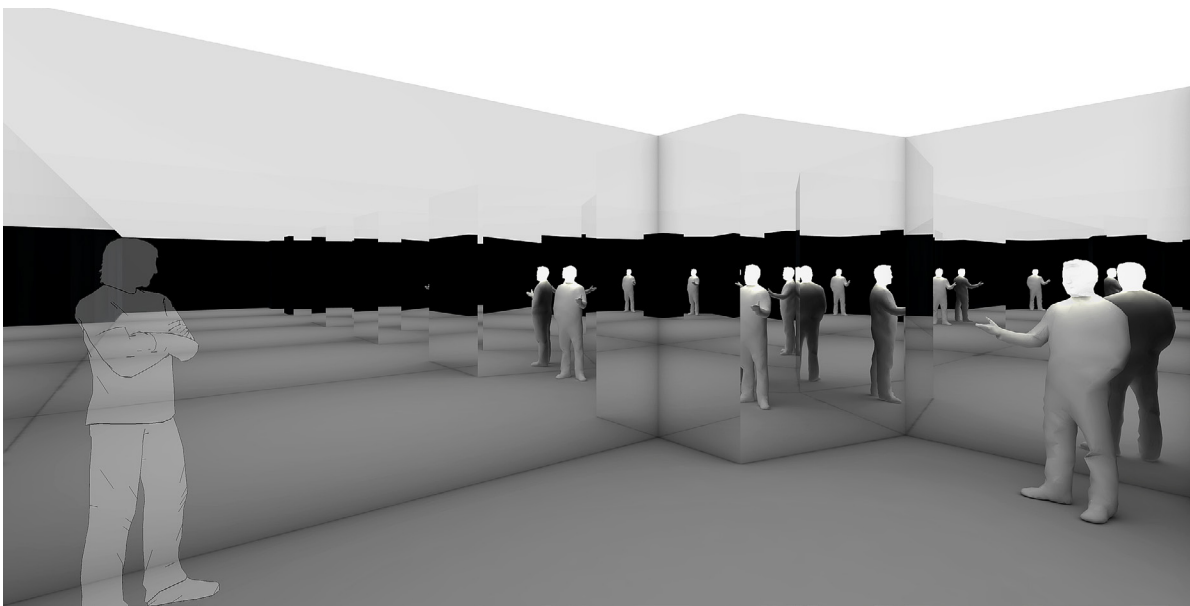


Figure 2-32: Distribution of reflections for two different sound sources, with the second source given a red color to distinguish the reflections relative to the sources.

- All we need to do now is alter the new **Emissive_source2** material's light color on the **Emissive** tab by clicking the color swatch and selecting a color. Here I have chosen pure red.
- We will also alter the **Diffuse** color to match so that we can distinguish it when painting objects with the material in SketchUp.
- Click the **Preview** button to verify the color properties, and then paint the second speaker's head with the new emissive material.



To demonstrate the impact of the second source, let's create two renderings from the new camera/listener position: one without the second source present in the reflections and one with it. See Figs. 2-33 and 2-34. Notice how, perhaps surprisingly, the red source has a significant amount of acoustic reflection energy coming from the alcove (on the right side of the image). Also counterintuitive: notice how little the side wall contributes to the acoustic reflections of either source!

2.3 360° Panoramic Renderings

So far we have been representing the listener's field of view in the traditional format of a cropped photographic image. Of course, hearing is an omnidirectional sense, so we need more comprehensive optics if we want to create renderings which convey the full spatial spectrum around the listening position. There are a few ways to capture and display the complete field of view, but first we must become familiar with a particular geometrical mapping projection known as the "equirectangular" projection.

2.3.1 Equirectangular Projection

Cartography has given us many different projections for unwrapping a sphere, each with its own distortion and emphasis. The most useful format for our purposes will be the equirectangular projection, which is presented as a 2:1 aspect ratio rectangle covering 360° horizontally and 180° vertically. This means that the far left (-180°) and far right (+180°) edges of the image are adjacent on the sphere. The top edge corresponds to looking straight up (+90°), also known as the zenith. The bottom edge corresponds to looking straight down (-90°), or the nadir. With this format, an orthogonal reference grid can be overlaid onto the image to denote spherical angles around the camera/listener. Also useful, this projection format minimizes distortion around the horizon, which corresponds well to our lateral experience, reserving the most distortion for the poles. In Figures 2-35 through 2-37 we see examples of how a sphere – whether a globe or spherical panorama – is unwrapped and flattened into an easy to read equirectangular image.

Figure 2-33 (opposite above): Distribution of reflections for a single sound source from the new camera/listener position.

Figure 2-34 (opposite below): Distribution of reflections for two different sound sources, with the second source given a red color to distinguish the reflections relative to the other source. Notice how the visualization method reveals substantial acoustic reflection activity in the alcove even for the red source, who is located nowhere near the alcove!

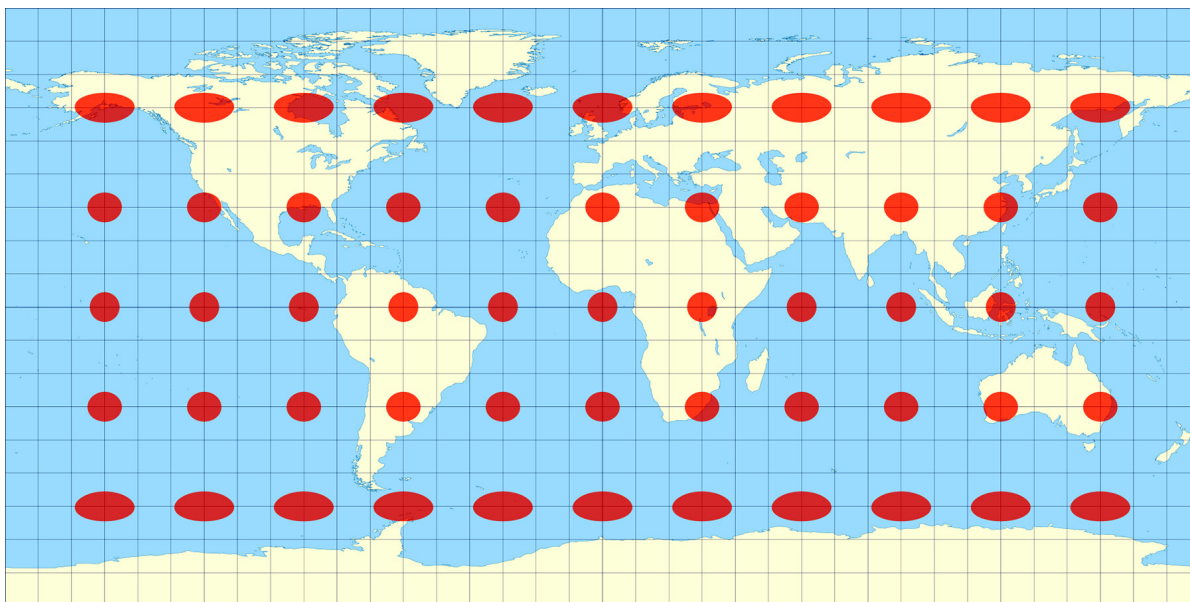
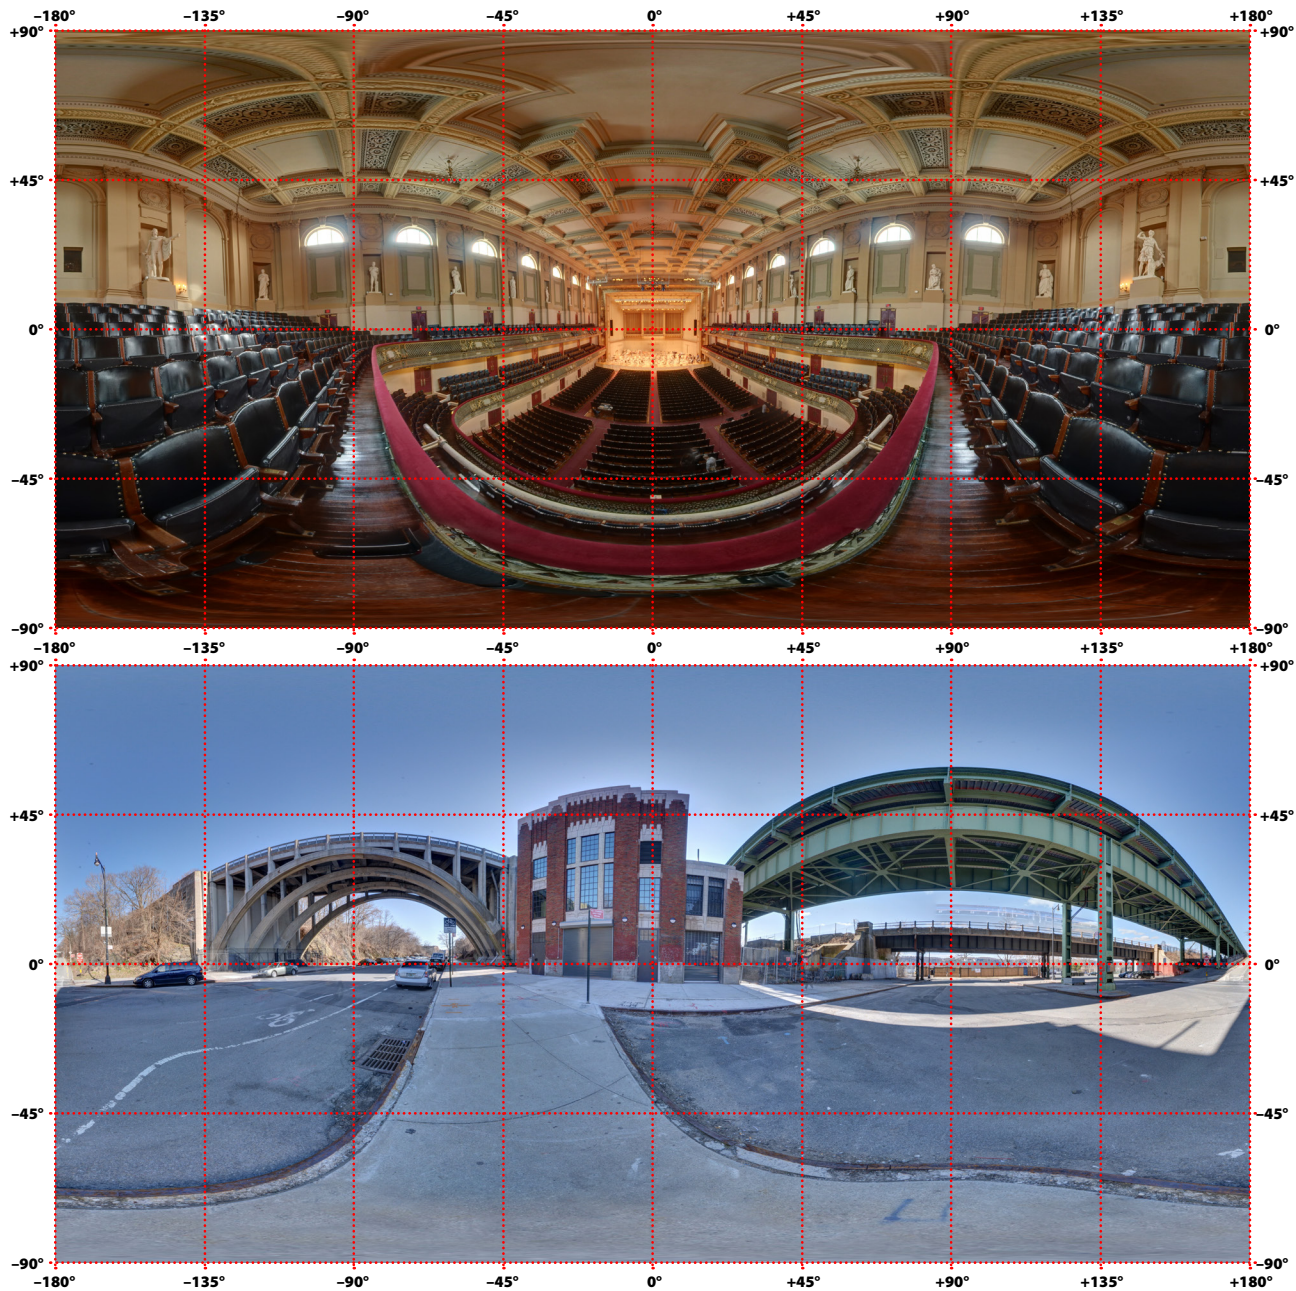


Figure 2-35 (above): The Tissot Indicatrix reveals the gradual distortions undergone when a globe is flattened to an equirectangular projection.

Figure 2-36 (opposite above): An equirectangular projection of a spherical panorama taken from the second balcony of Boston Symphony Hall. The orthogonal grid of angular measurements is overlaid to illustrate how the equirectangular projection maps a $360^\circ \times 180^\circ$ sphere to a 2:1 flat plane.

Figure 2-37 (opposite below): An equirectangular projection of a spherical panorama taken in north Manhattan.



2.3.2 “Hacking” V-Ray’s Camera Optics

We could use an image editor to laboriously stitch together multiple renderings into one image with a larger field of view, but there is no need for that. Once again, V-Ray allows us to defy real-world physics: we can set the V-Ray virtual camera “lens” to see the entire 360° field of view at once – not possible with actual lens optics! First, position the camera in SketchUp at the listener’s position and orient the view to what will be the center of the panorama, making sure that the line of sight is level to the ground. Then:

- In the **V-Ray Option Editor**, expand the **Camera** tab.
- Under the **Type** drop-down menu, select **Spherical**.
- Check the **Override FOV** box, and set the value to **360.0**. See Fig. 2-38.

Next we need to configure V-Ray to output a proper equirectangular format, so we must override the SketchUp viewport and make sure the

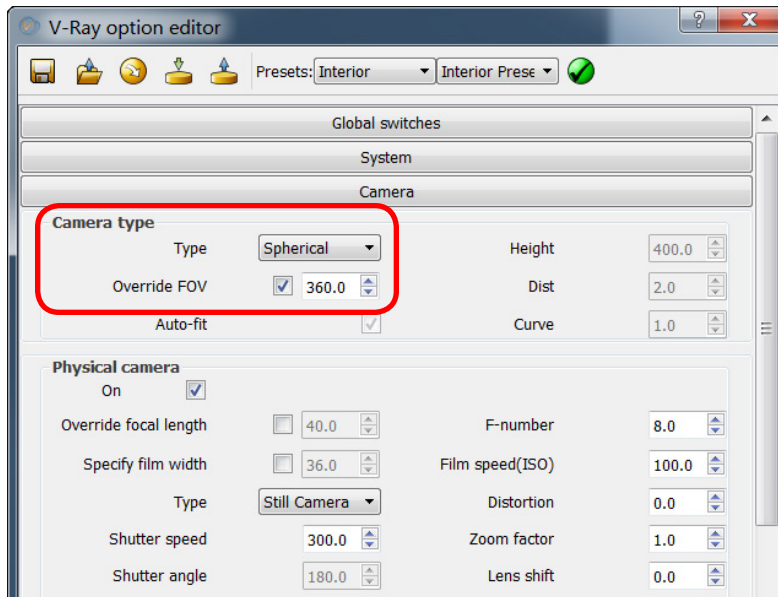


Figure 2-38: To generate equirectangular projections directly in the renderings, we must force V-Ray’s virtual camera to achieve what no real-world optics can: a full 360° lens. In the **V-Ray Option Editor**, under the **Camera** tab, use the **Type** drop-down menu to select the **Spherical** camera, then check the **Override ROV** box and change the angle to **360.0** degrees.

aspect ratio is 2:1 with whatever resolution we are using:

- In the **V-Ray Option Editor**, expand the **Output** tab.
- Check **Override Viewport**.
- Set the **Width** to be twice the **Height**, e.g. 4000 and 2000 respectively.
- Make sure the **Image Aspect Ratio** updates to **2.0**. See Fig. 2-21.

Now when you render, the V-Ray Frame Buffer will not correspond to the SketchUp viewport, but instead will generate a perfect equirectangular projection. We will demonstrate with two examples as positioned in Fig. 2-39.

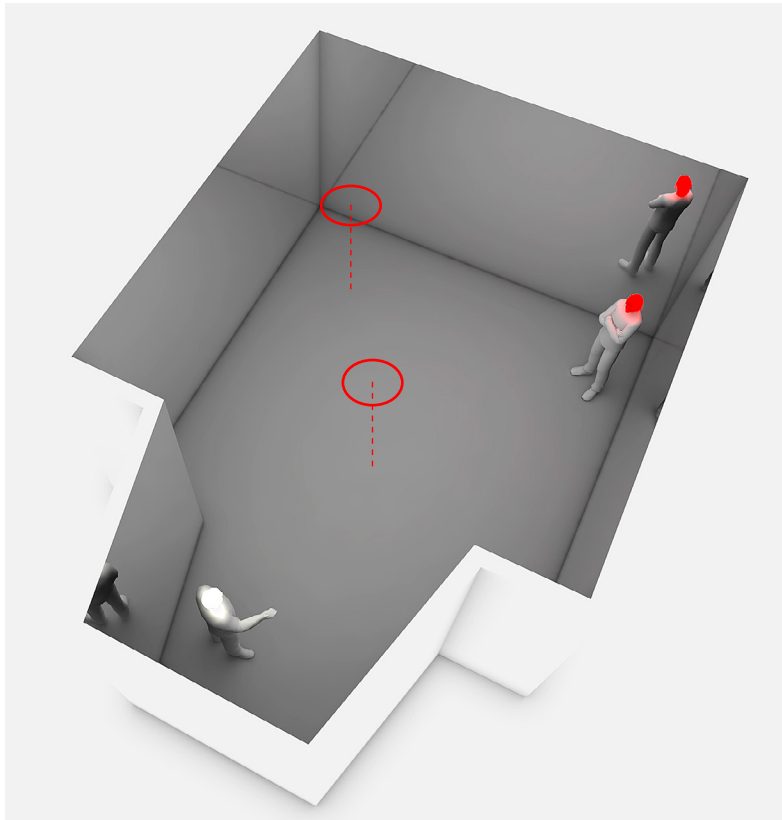
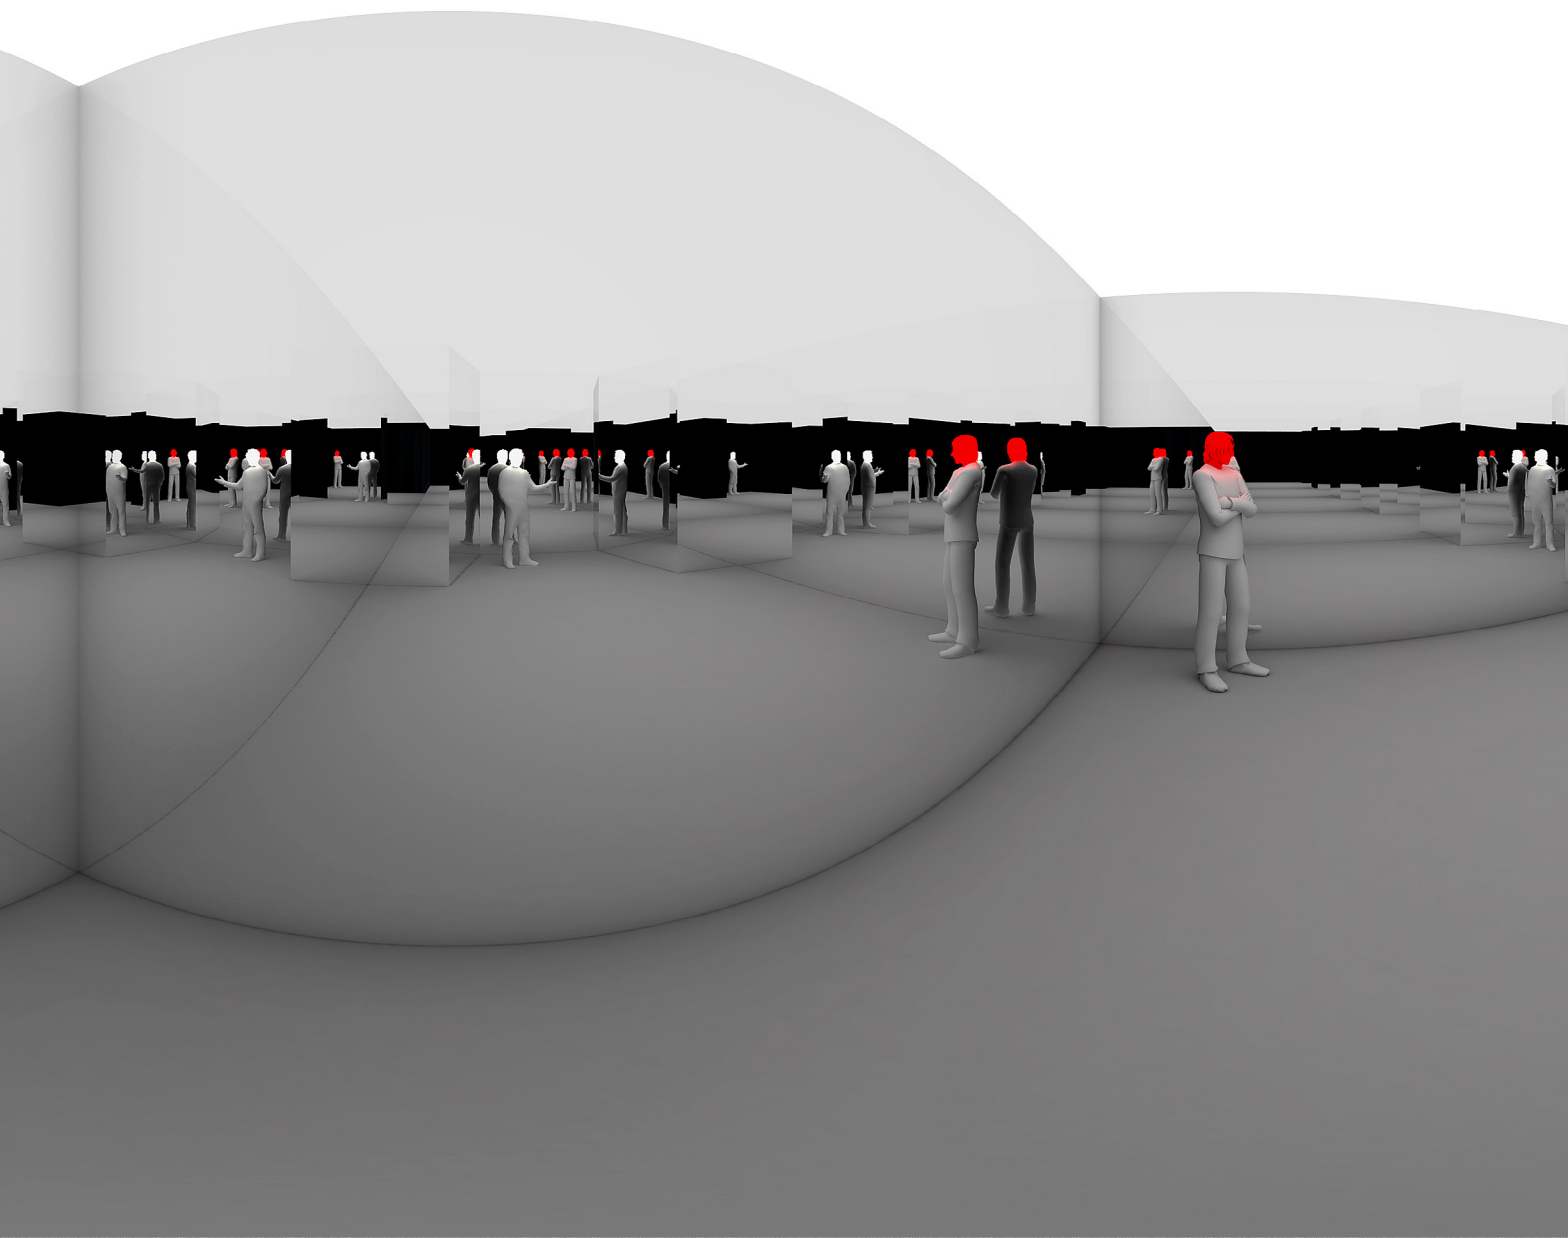


Figure 2-39: For the two equirectangular spreads on the following pages, we will look at the differences between two camera/listener positions: the first in the far left corner opposite the alcove, and the second in the center of the room.

Figure 2-40: The equirectangular projection from the corner camera/listener position. Notice where the clusters of reflections occur for each sound source.



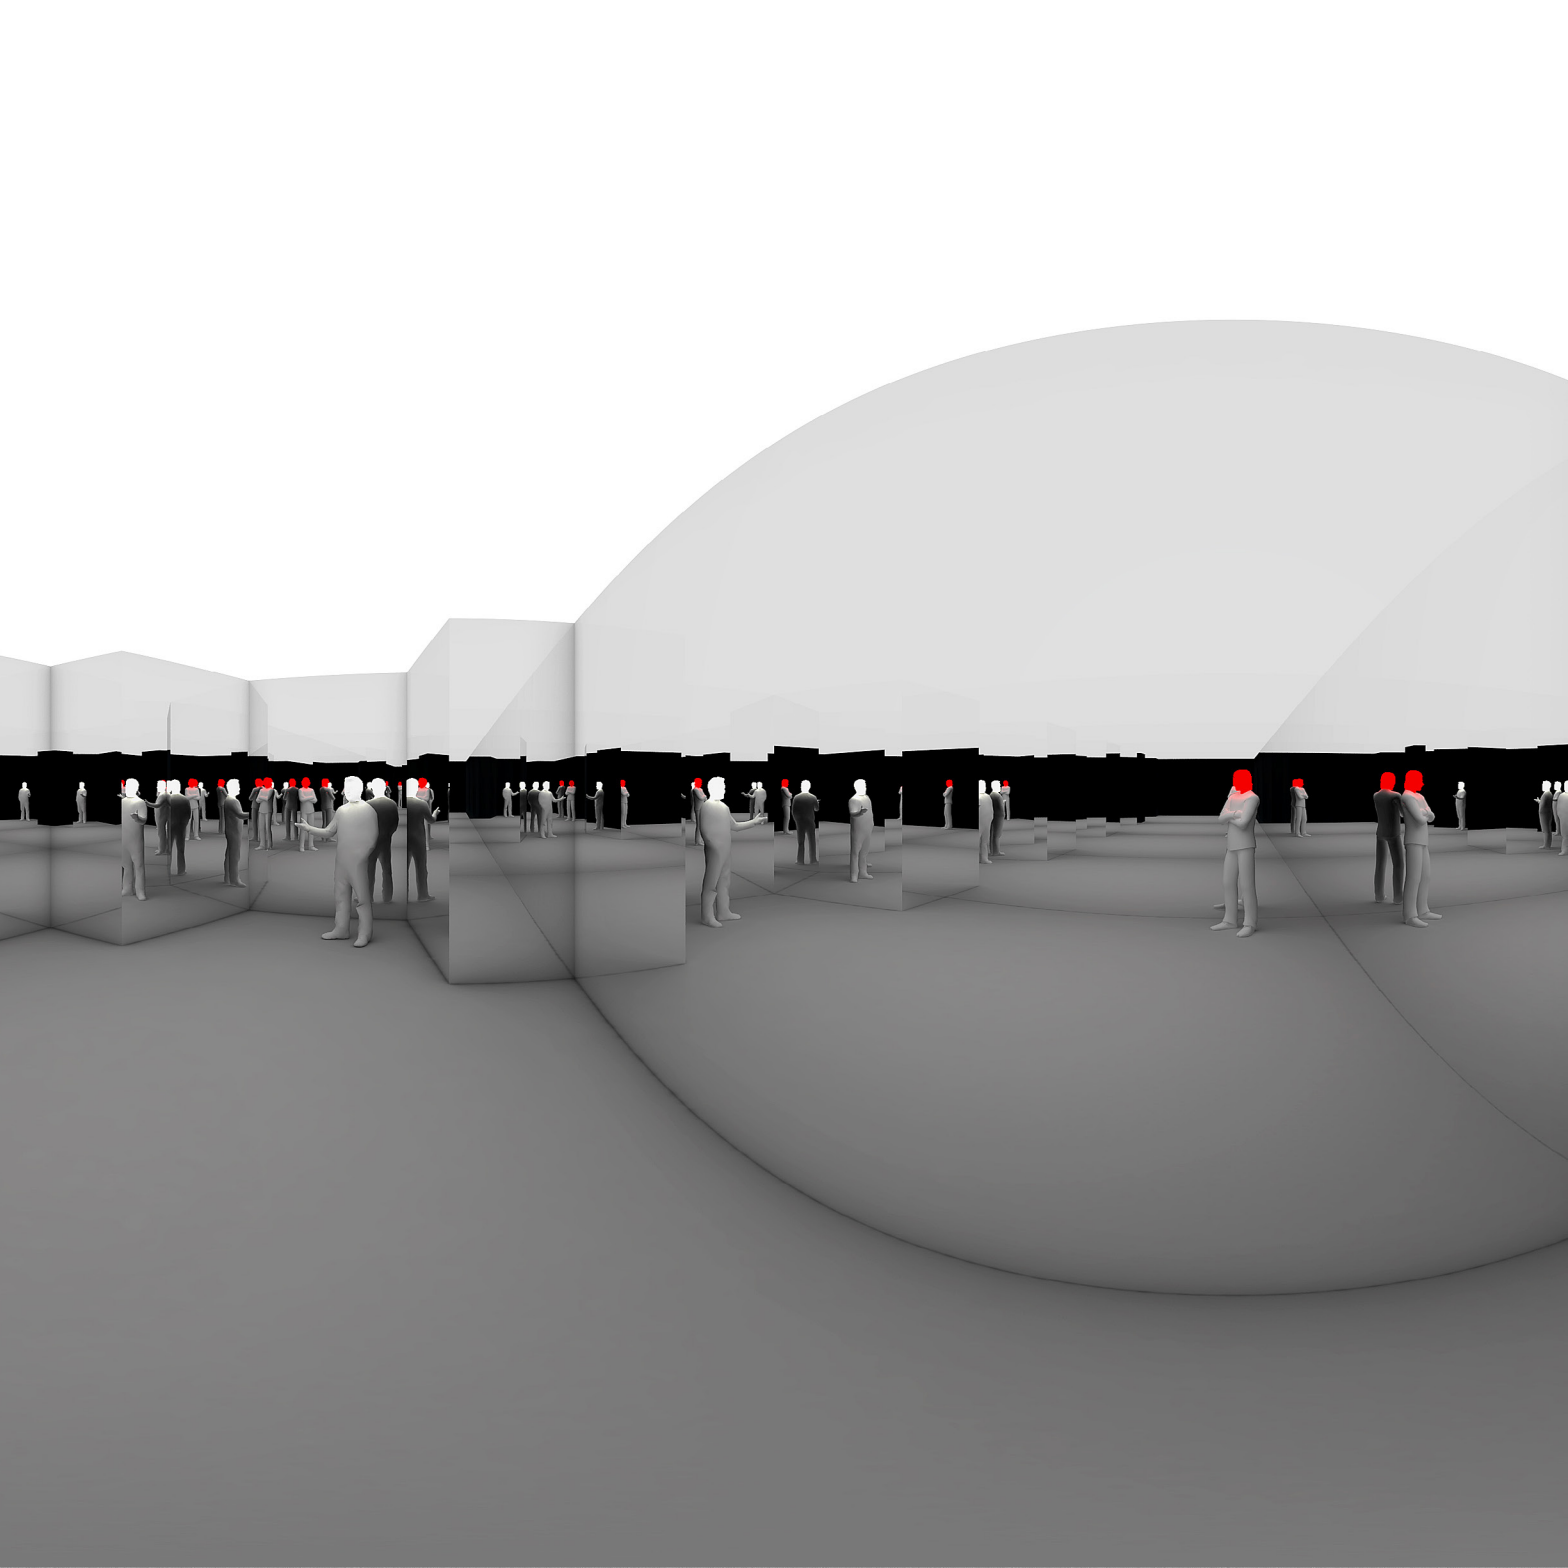
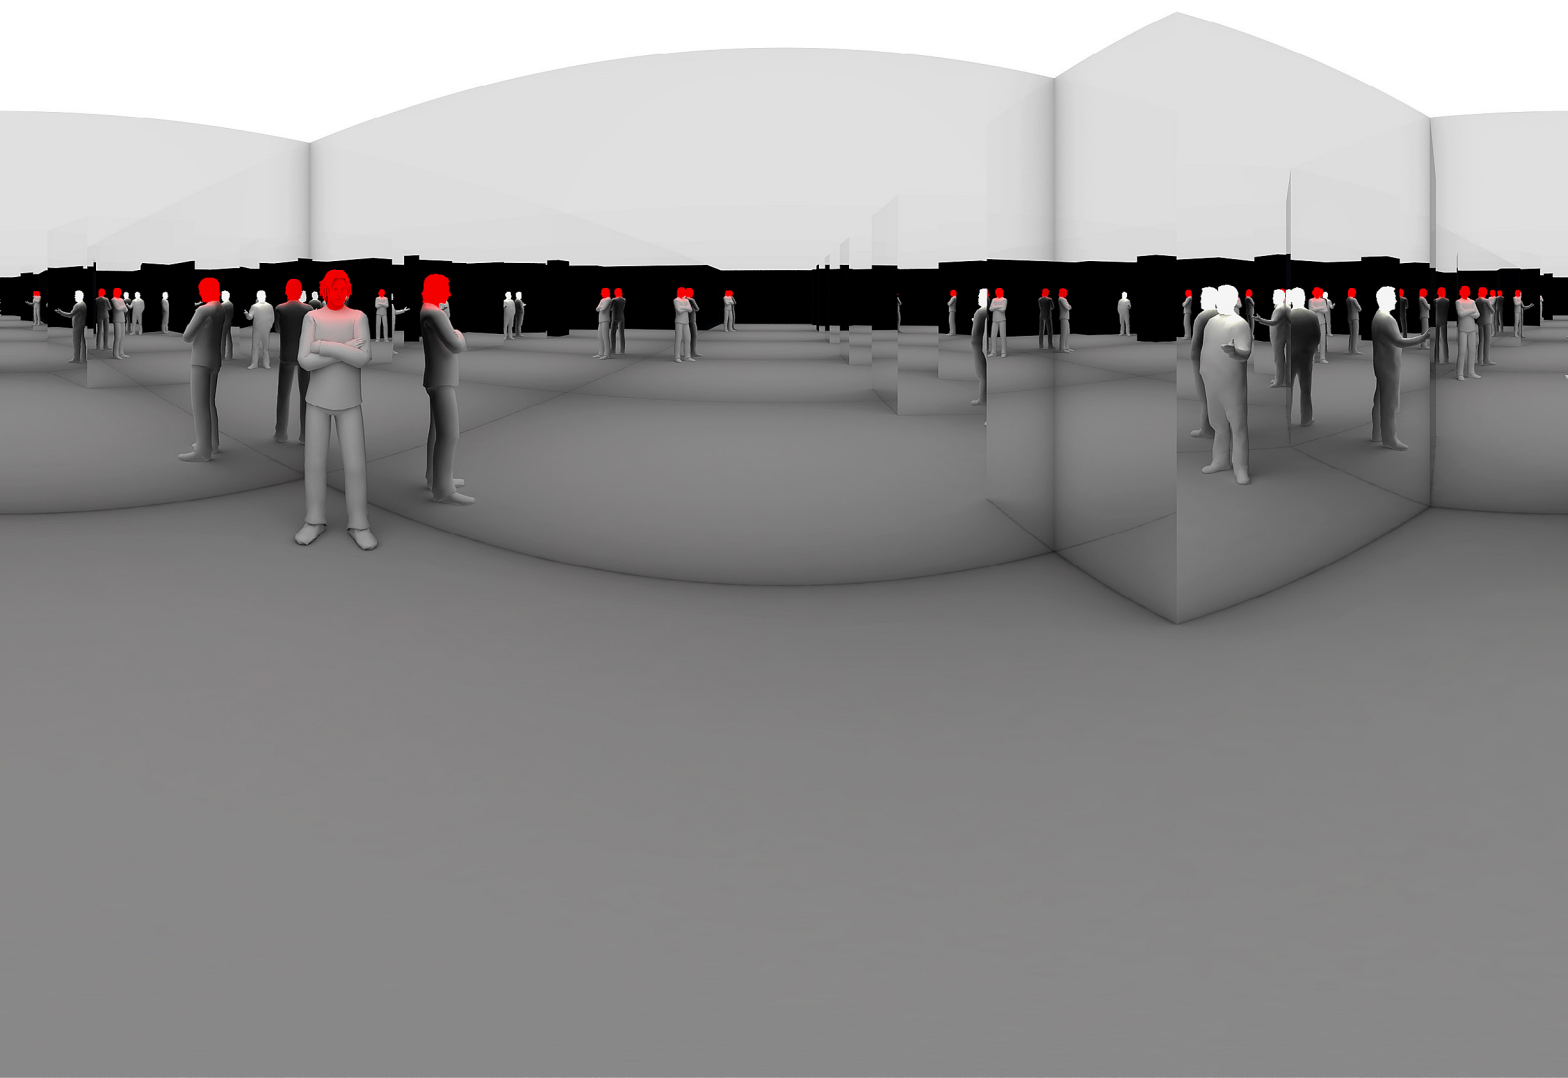
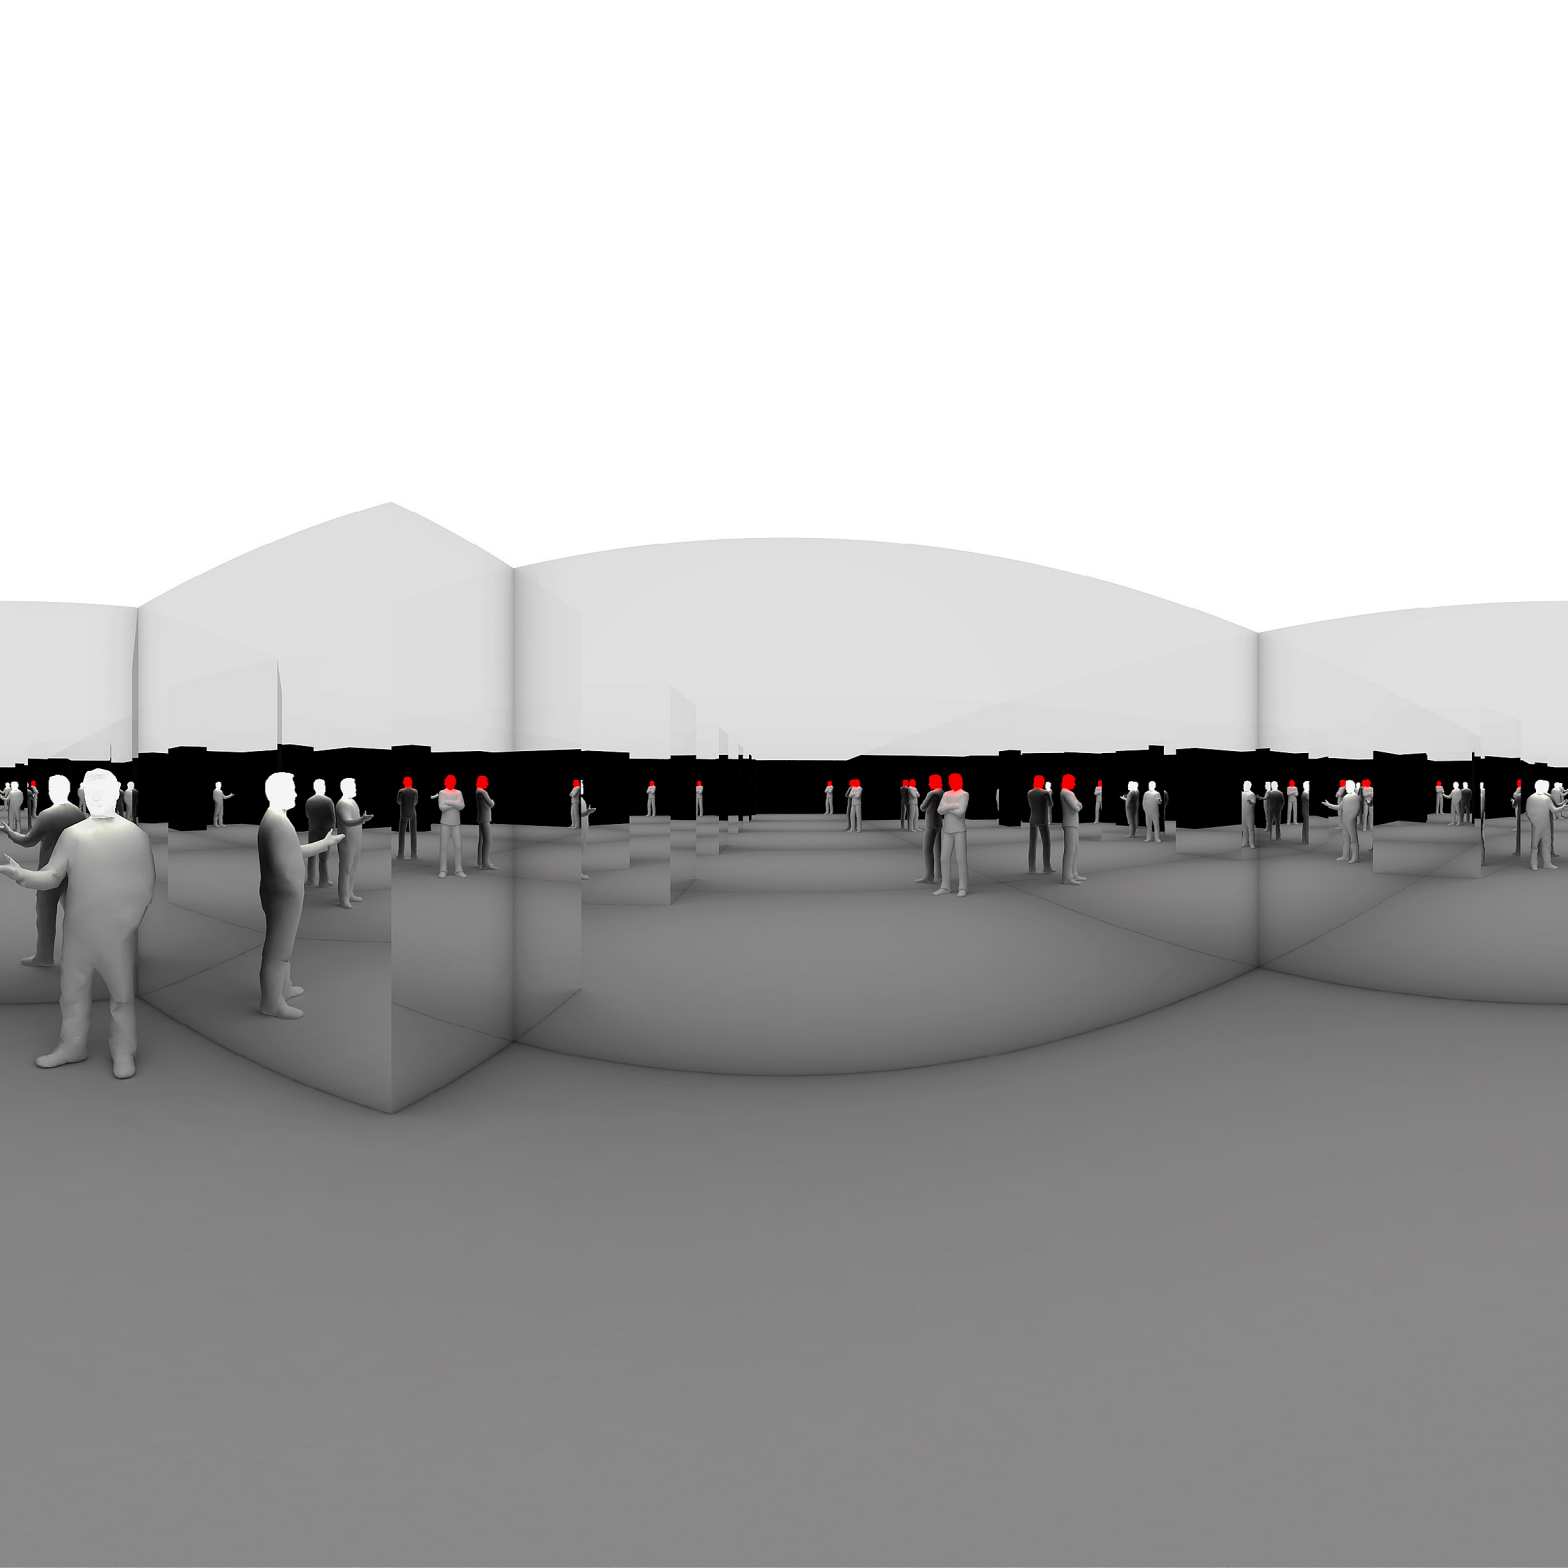


Figure 2-41: The equirectangular projection from the middle camera/listener position. What are the differences compared to the previous listener position? Notice how confined the white source is to the alcove now, and how much more distributed around the space the red source is.





2.3.3 QuickTime Virtual Reality (QTVR)

The flattened equirectangular projection is extremely useful for surveying patterns across the entire field of view in a single image. But what if the increased geometrical distortion above or below the horizon are distracting, and you'd like to study the spatial distribution of reflections in a more intuitive way? Another advantage of the equirectangular projection is that it is a common format for interactive viewers.

You are probably already familiar with navigating spherical panoramas: interactive 360° photography platforms such as Google Maps' Street View allow you to "steer" around a spherical projection on a flat screen. While you cannot see the entire field of view at once, wherever you look exhibits only modest geometrical distortion compared to the equirectangular projection. This makes looking up and down appear more natural, and is useful for immersing into a particular point of view from a particular listening position.

There are a few ways to view your equirectangular renderings as interactive files. The popular spherical panography stitching software PTGui includes an excellent application, PTGui Viewer, which can open the equirectangular images directly, giving the user complete control over the size and aspect of the viewer window. The navigation is extremely stable and smooth.

Another popular option is to convert the equirectangular images to QuickTime Virtual Reality (QTVR) .mov files. When these are opened with any standard QuickTime player, the user can interactively steer around the panorama. Furthermore, QTVR files can be placed inside interactive PDF files and Adobe Reader will allow the user to steer the panorama right within the document. Popular software for converting equirectangular images to QTVR .mov files include [PTGui](#) and Garden Gnome's [Pano2VR](#).

To wrap up this overview of the illuminated acoustics method, take a spin in the following two embedded QTVR files. Have a closer look at the acoustic reflection orders of our demonstration model from the two different vantage points.

Figure 2-42 (opposite, interactive): Click and drag to swivel around in the QTVR example from the corner listening position.

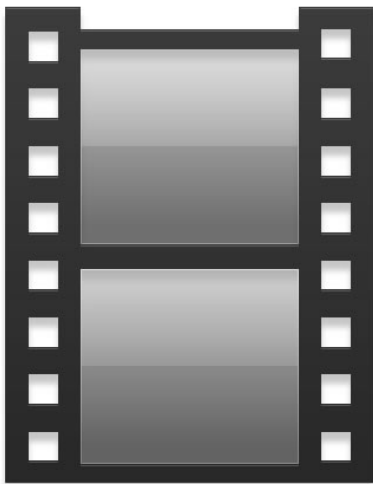
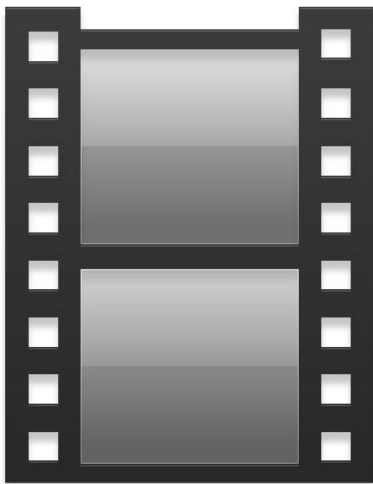
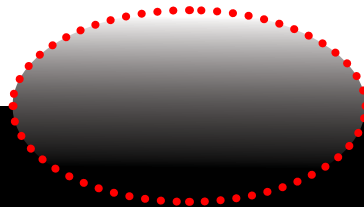


Figure 2-43 (opposite, interactive): Click and drag to swivel around in the QTVR example from the middle listening position.



**DEAF
ARCHITECTS**



**BLIND
ACOUSTICIANS**

3 Conclusion

The method demonstrated here revives an early acoustical modeling technique based on luminous sources and reflective surfaces to optically reveal sound paths within architectural models. Using contemporary software tools commonly found in the designer's workflow, we have seen how rendering virtual 3D models with mirror surfaces and emissive source materials gives us visualization capabilities not possible with physical models. We can dictate the number of reflection orders, and precisely control the brightness and color of sources and reflections. We can generate 360° immersive points of view, and even render the point of view as a sound source itself.

As a purely visual embodiment of the Image Source Method, this approach is not intended to replace auralization, but rather to provide a path for designers to include acoustical thinking in their regular modeling and visualization methods. Ultimately, the aim is to encourage more practitioners to join the vital discussion of sound in architecture.

We all have a lifetime's experience of sound in architecture, and yet divided disciplines and specialized tools have prevented larger discussions from flourishing between architects and acousticians. This project acknowledges the stark divide represented on the cover of Robert E. Apfel's 1998 book *Deaf Architects & Blind Acousticians: A Guide to the Principles of Sound Design*. It is my sincere hope that the techniques shared here might soften that disciplinary edge somewhat.

Image Credits

All images, drawings, models, renderings, animations, and screenshots were generated by the author except for those noted here. The plan drawings and models used for the examples in the Method chapter were derived from a project in an Architectural Acoustics class taught by Todd Brooks at RPI. Software icons are from Trimble SketchUp and Chaos Group V-Ray.

Covers: www.wallpaperest.com

Figure 1-6: www.webneel.com

Figure 1-7: Knudsen, V. O. (1947). *Architectural Acoustics*, p. 70. New York: John Wiley & Sons, Inc.

Figure 1-8: Ibid.

Figure 1-9: Knudsen, V. O., & Harris, C. M. (1950). *Acoustical Designing in Architecture*, p. 202. New York: John Wiley & Sons, Inc.

Figure 1-10: <http://arttattler.com/archiveyayoikusama.html>

Figure 1-11: Ibid. (Photo by Bob Goedewaagen)

Figure 1-12: <http://imgur.com/gallery/l4ZUI> (Panorama by DrImgur)

Figure 1-13: <http://hauteliving.com/2015/11/5-la-art-installations-you-need-to-see-this-fall/591708/?market=san-francisco> (Photo by Andrew Toth/Getty Images)

Figure 1-14: <http://www.thilofrank.net/the-phoenix-is-closer>

Figure 1-15: Ibid.

Figure 1-16: http://archiv.shedhalle.ch/eng/archiv/2005/karnevaleske/the-matische_reihe/projekte_3/mayer/index.shtml

Figure 1-17: <http://www.smithsonianmag.com/innovation/doug-aitken-is-redefining-how-we-experience-art-180947643/> (Photo by Timothy Archibald)

Figure 1-18: <http://www.studiodaily.com/2014/03/how-they-did-it-divergents-mirror-room-sequence-and-more/>

References

[Trimble SketchUp online training tutorials and resources.](#)

[SketchUp video tutorials on YouTube.](#)

[V-Ray for SketchUp online manual and tutorials.](#)

Allen, J. B., & Berkley, D. A. (1979, Apr). Image Method for Efficiently Simulating Small-Room Acoustics. *J. Acoust. Soc. Am.*, 65, 943-950.

Blauert, J. (1997). *Spatial Hearing: The Psychophysics of Human Sound Localization*. Cambridge: The MIT Press.

Borish, J. (1984, Jun). Extension of the Image Model to Arbitrary Polyhedra. *J. Acoust. Soc. Am.*, 75, 1827-1836.

Cremer, L., & Müller, H. A. (1982). *Principles and Applications of Room Acoustics* (Vol. 1). (T. J. Schultz, Trans.) New York: Applied Science Publishers.

Everest, F. A. (2001). *Master Handbook of Acoustics, 4th Edition*. New York: McGraw-Hill.

Kleiner, M. (2012). *Acoustics and Audio Technology, 3rd Edition*. Fort Lauderdale: J. Ross Publishing.

Knudsen, V. O. (1947). *Architectural Acoustics*. New York: John Wiley & Sons, Inc.

Knudsen, V. O., & Harris, C. M. (1950). *Acoustical Designing in Architecture*. New York: John Wiley & Sons, Inc.

Long, M. (2006). *Architectural Acoustics*. Burlington: Elsevier Academic Press.

Rindel, J. H. (2000). The Use of Computer Modeling in Room Acoustics. *Journal of Vibroengineering*, 3(4), 219-224.

Rindel, J. H. (2002). Modelling in Auditorium Acoustics: From Ripple Tank and Scale Models to Computer Simulations. *Forum Acusticum Sevilla 2002*, 33, pp. 31-35. Seville: OAI.

Robertson, B. (2014, March 28). *How They Did It: Divergent's Mirror Room Sequence (and More)*. Retrieved from www.studiodaily.com: <http://www.studiodaily.com/2014/03/how-they-did-it-divergents-mirror-room-sequence-and-more/>

Stephenson, U. (1990). Comparison of the Mirror Image Source Method and the Sound Particle Simulation Method. *Applied Acoustics*, 29, 35-72.

Vermeulen, R. (1940, Nov). Optical Experiments with Models on the Directional Distribution of Sound in Halls. *Philips Technical Review*, 5, 321-323.

About the Author

J. Parkman Carter was trained as a luthier at the Totnes School of Guitarmaking in Devon, UK. He received a BA in Integrated Arts from Bard College, and then worked for several years building custom instruments and loudspeakers, while also working as a consultant in high-end audio and video systems.

Both alarmed and inspired by the Architecture 2030 Challenge, which showed that the built environment is responsible for half of the planet's carbon footprint, Carter returned to school in 2008 to study architecture in earnest. After completing three Masters Degrees in Architecture, Architectural Lighting Design (both at Parsons School of Design), and Architectural Acoustics (at Rensselaer Polytechnic Institute), Carter's current research focuses on the limitations of conventional architectural representation to convey salient qualities of natural light and soundscape in the built environment.

Carter is currently a PhD Candidate in Architectural Sciences at RPI in Troy, NY, and also teaches in the School of Constructed Environments at Parsons School of Design in New York, NY.

A longer version of this e-book is forthcoming, which includes advanced rendering techniques and several application examples. To be notified when it is available, send your email address to cartej@rpi.edu.

