Geometric Sound Propagation

Sung-eui Yoon

Slides are from Carl Schissler, Anish Chandak, and Dinesh Manocha

Class Objectives are:

- Binaural audio and Head Related Transfer Function (HRTF)
- Sound Propagation Phenomena
- Geometric Methods for Sound Propagation

Sound Propagation

The process by which sound is emitted from a source, interacts with the environment, and is received by a listener.

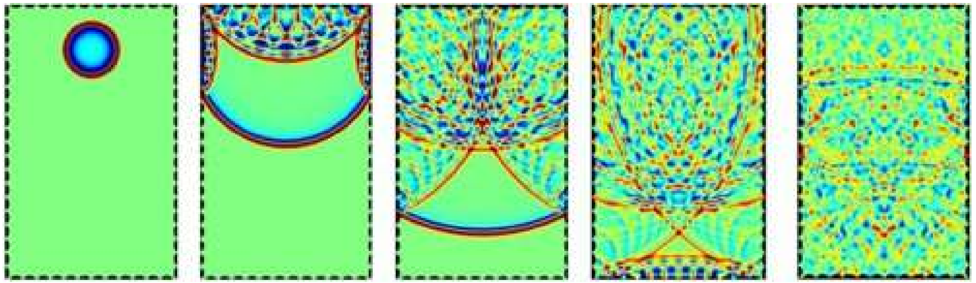

In our daily lives

Spatial Sound Is Everywhere

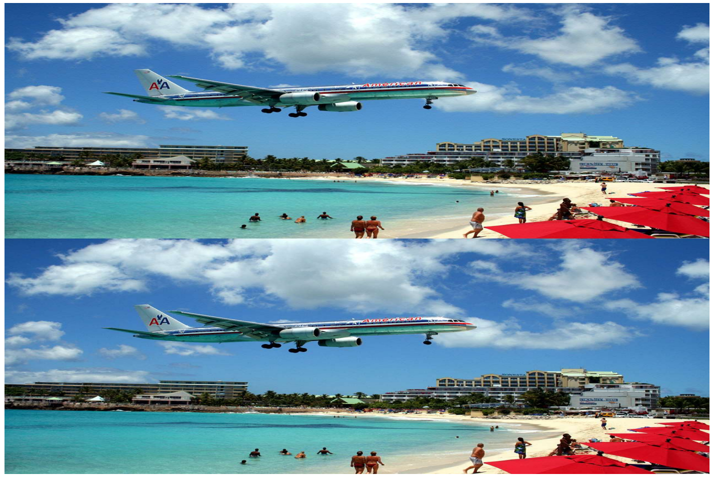

Princess Juliana International Airport

Emergency Vehicles

Spatial sound is crucial to human hearing

Spatial Sound Is Everywhere

In our daily lives

Basketball stadium

Grace Cathedral, San Francisco

Spatial sound is crucial to human hearing

Applications

Games

MAG (PS3): *up to 256 players*

Real Racing 2 (iPhone)

Applications

Virtual Reality

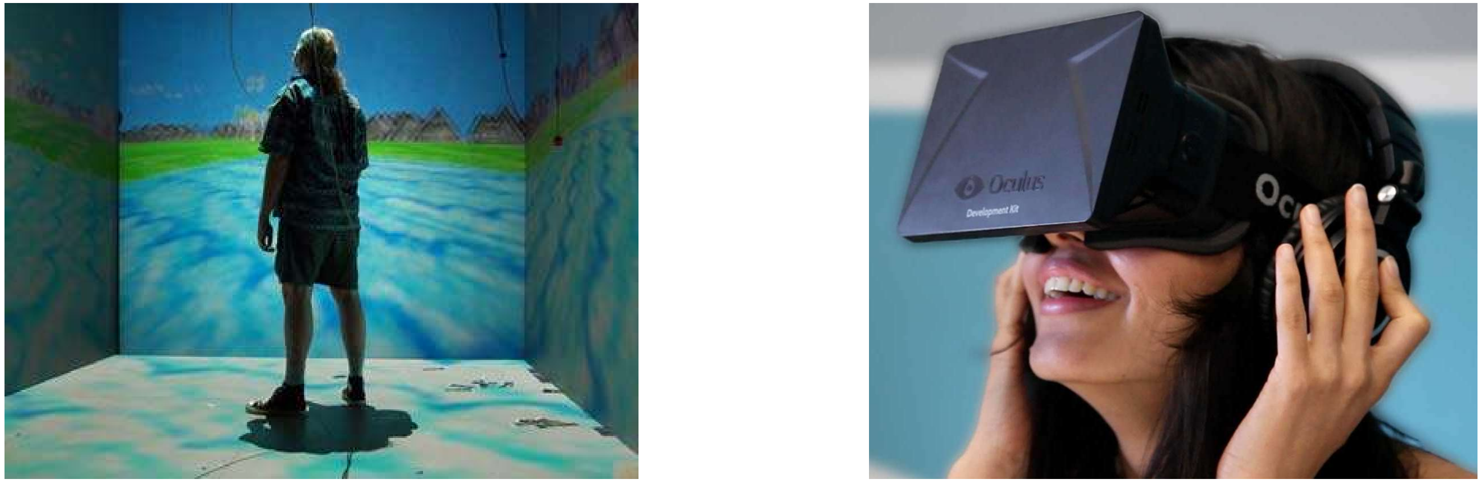

Applications

Training Simulations

Applications

Architectural Acoustics

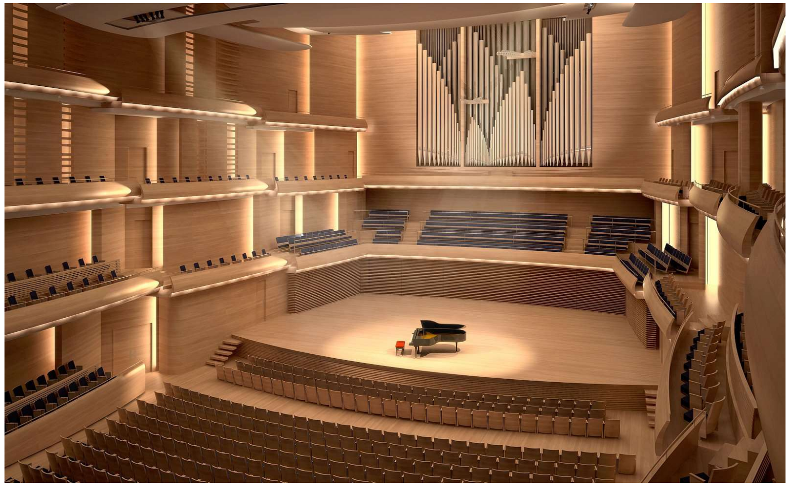

Audio-Visual Coherence

Direction

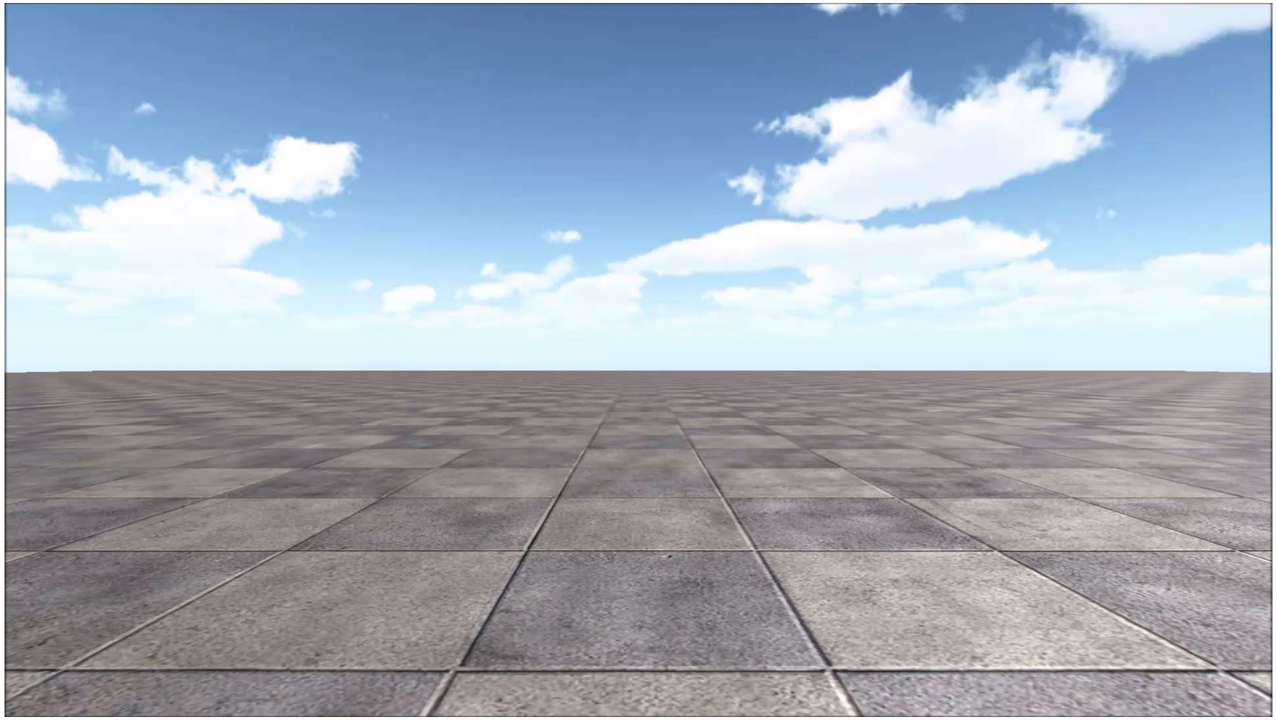

Occlusion

Audio-Visual Coherence

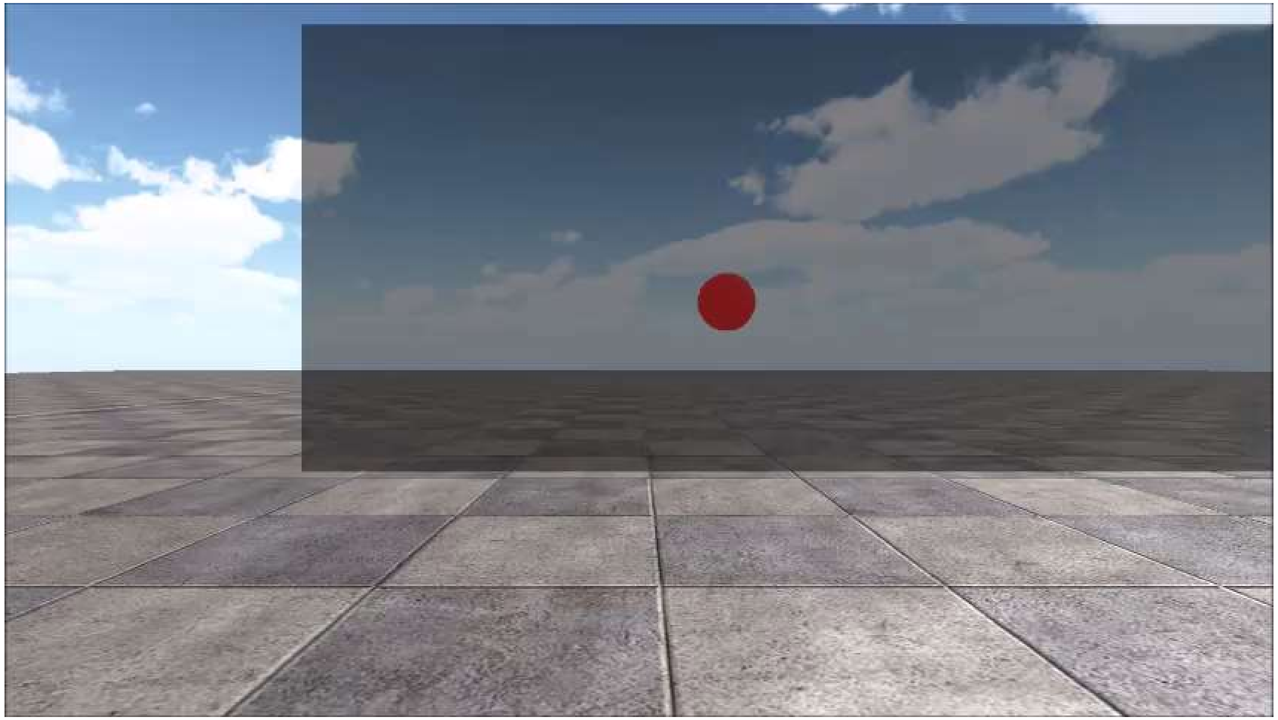

Reverberation

Audio-Visual Coherence

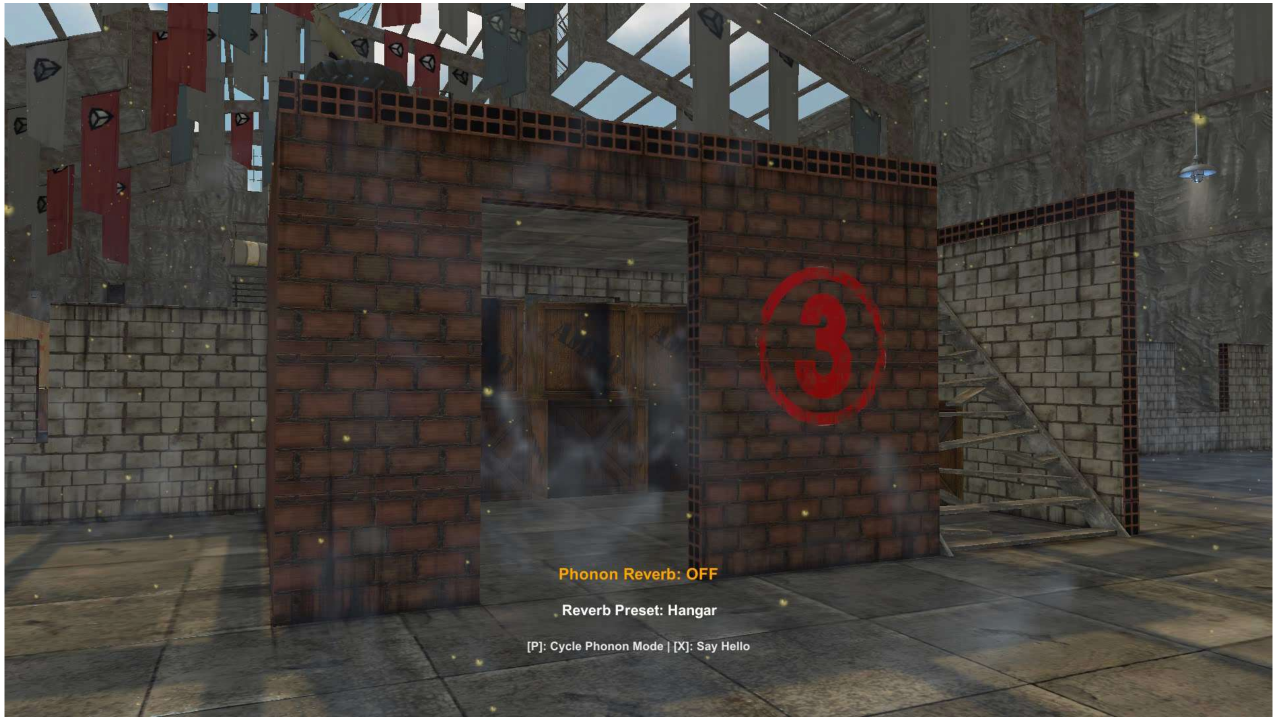

Phonon Reverb: OFF

Reverb Preset: Hangar

[P]: Cycle Phonon Mode | [X]: Say Hello

Audio-Visual Coherence

HMD Capabilities

Orientation Tracking

Position Tracking

Audio-Visual Coherence

HMD Capabilities – Orientation Tracking

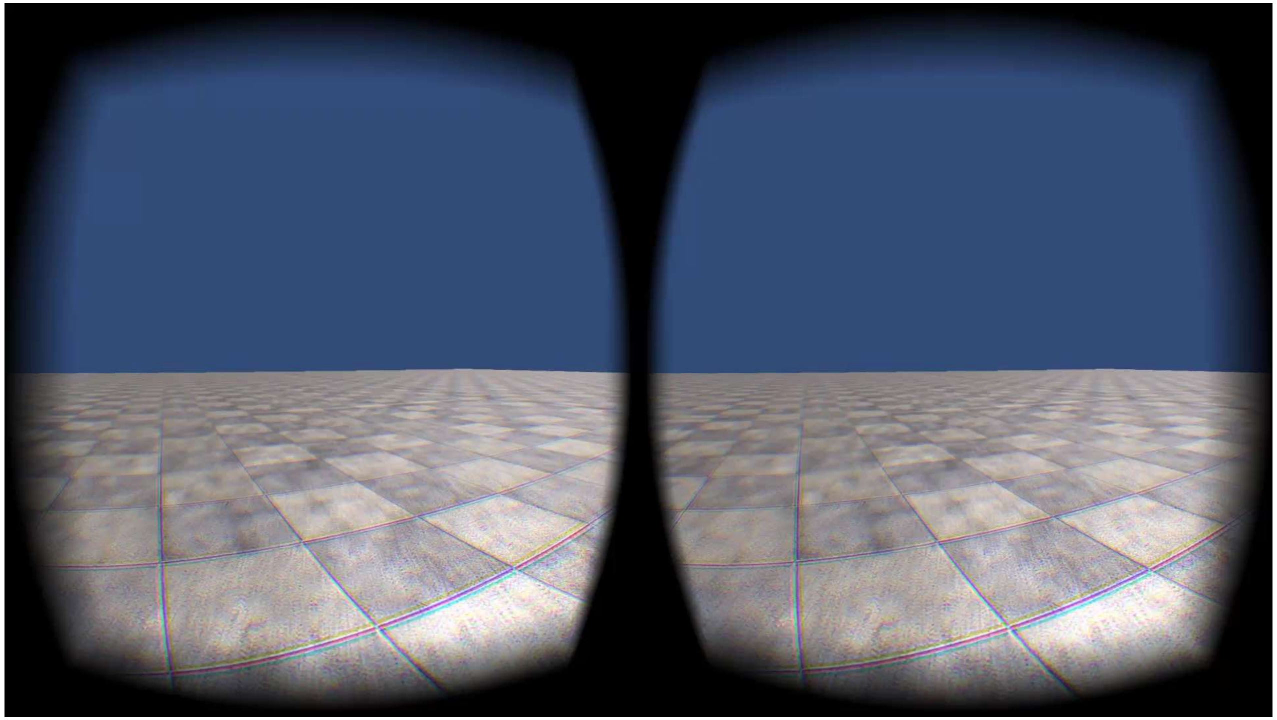

Components of Spatial Sound

The big picture

Sound is emitted

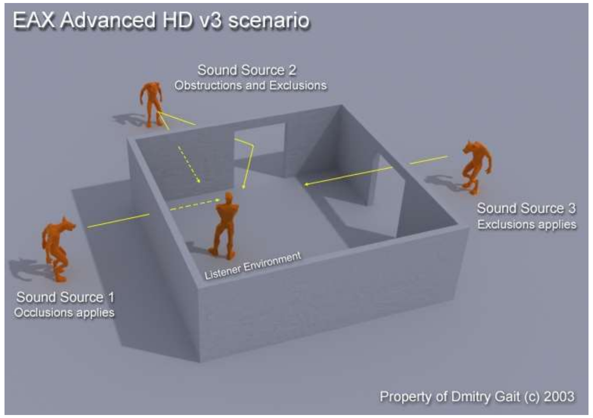

Interacts with environment

Environmental Acoustics
reverb, occlusion, reflection

Reaches listener's ears

Binaural Audio
HRTFs

Components of Spatial Sound

The big picture

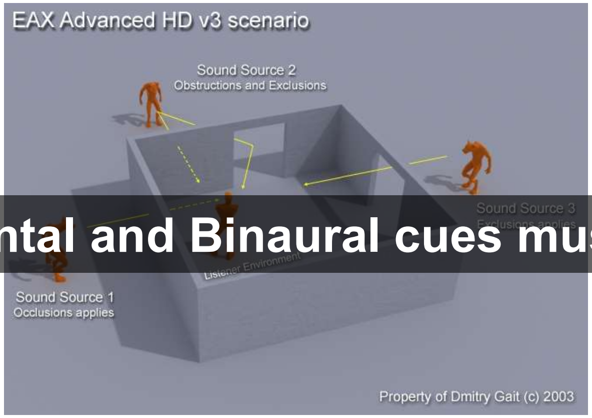

Environmental and Binaural cues must be combined

Sound is emitted

Interacts with environment

Environmental Acoustics
reverb, occlusion, reflection

Reaches listener’s ears

Binaural Audio
HRTFs

Binaural Audio

Listening with both ears

- We use both ears when locating a sound source
- Each ear receives slightly different sound
  - Time difference
  - Intensity difference
  - Spectral cues

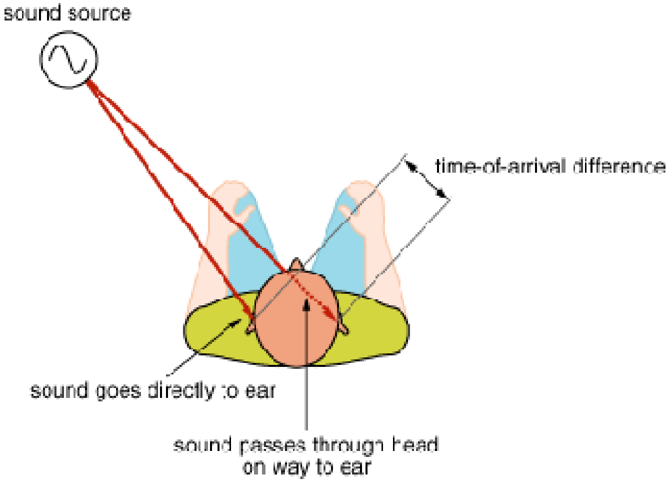

Inter-aural Time Difference

Sound arrives quicker at one ear than the other

- Source is usually closer to one ear than the other
- Sound arrives at different times at left and right ears

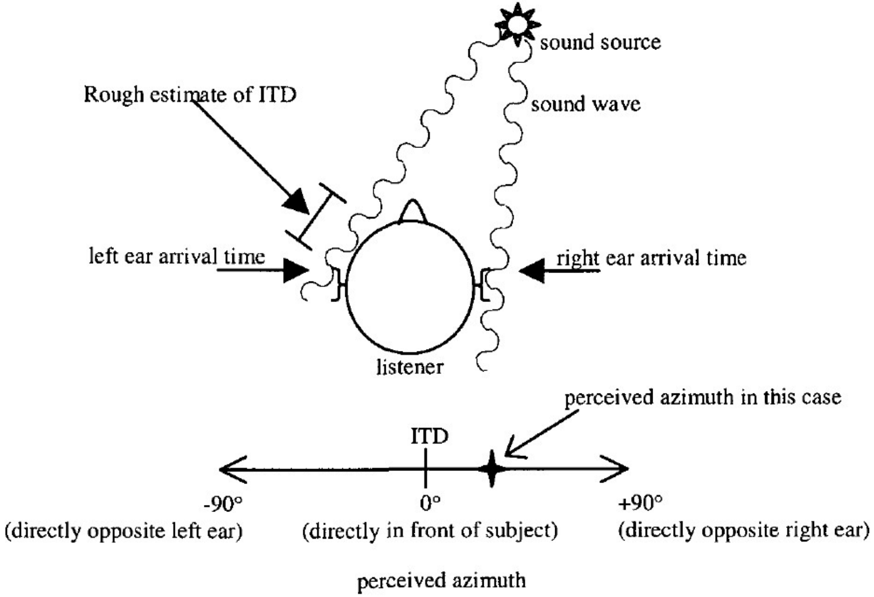

Inter-aural Intensity Difference

Sound is louder in one ear than the other

- Sound is usually closer to one ear than the other

- Sound arrives with different intensities at left and right ears

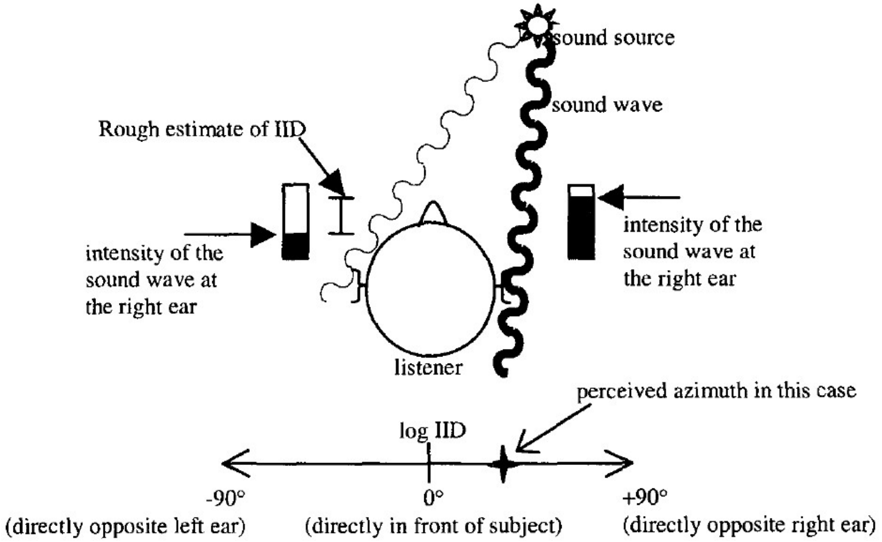

Cone of Confusion

Why ITD and IID are not enough

- Multiple source directions have the same ITD and IID
- Such directions lie on a **cone of confusion**
- ITD and IID are not enough, we need more directional cues for front-back, up-down disambiguation

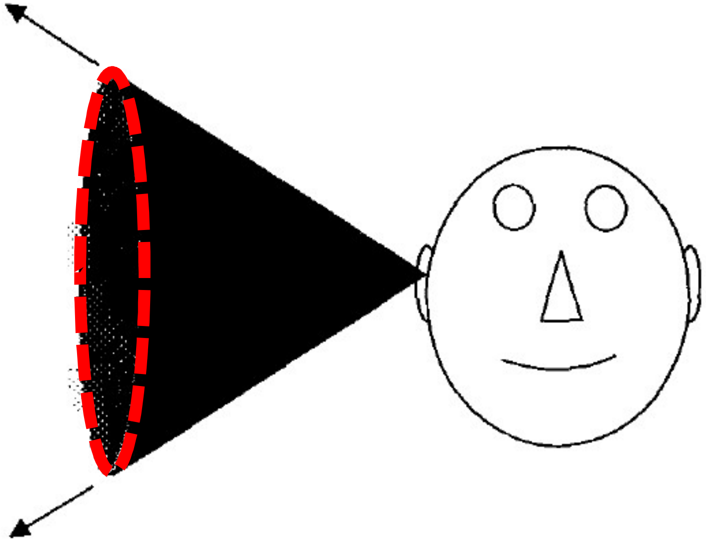

Spectral Cues

Frequency and phase differences

- Sound is affected by head, outer ear, shoulders, torso
- Outer ear performs frequency-dependent filtering
- Head shape causes multiple scattering
- These are cues we use every day!

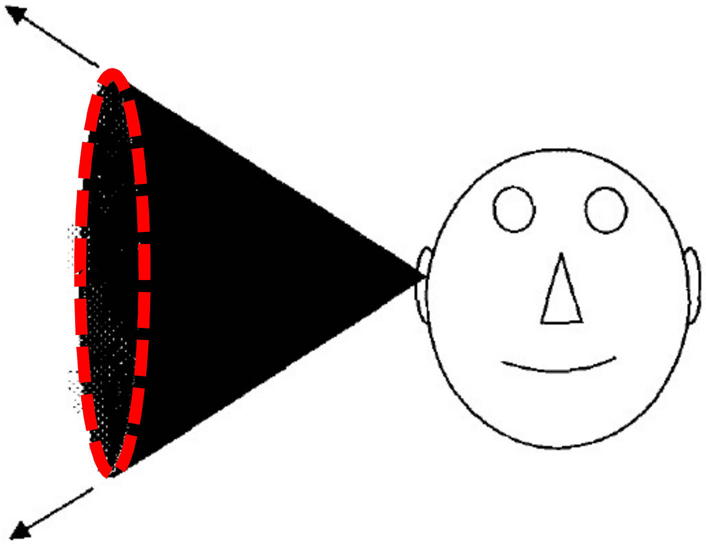

Head Related Transfer Function (HRTF or HRIR)

Unifies all binaural cues

- **HRTF = ITD + IID + Spectral cues**
- Special signal **recorded** for different source positions around listener
- One HRTF for each ear
- Models effect of head, shoulders, outer ear, torso, …

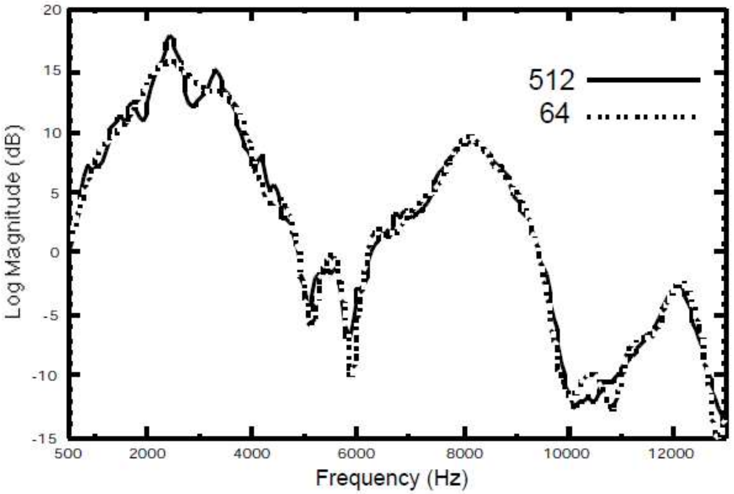

HRTF as a function of frequency

The gritty origin story

History of Binaural Audio

- Research began: **1939** (!)
- “Modern” algorithms developed: **Early 1990s**
- **Are we reinventing the wheel?**

New Developments in Binaural Audio

And some unanswered questions

- **Head Tracking**
  Head movements are an important location cue
- **Real-Time Environmental Acoustics**
  Audio-visual coherence is an equally important cue
- **Individualized HRTFs**
  Do we need to record each user’s HRTF?

Binaural Audio: Summary

The bottom line

- **Binaural audio** = use of **HRTFs**
- For each incoming direction, there is a pair of HRTFs (one per ear)
- The HRTFs transform the sound to what is heard in each ear
  - Allows front/back, up/down, left/right localization

Sound Propagation

Given: source(s), listener(s), obstacle(s), propagation medium

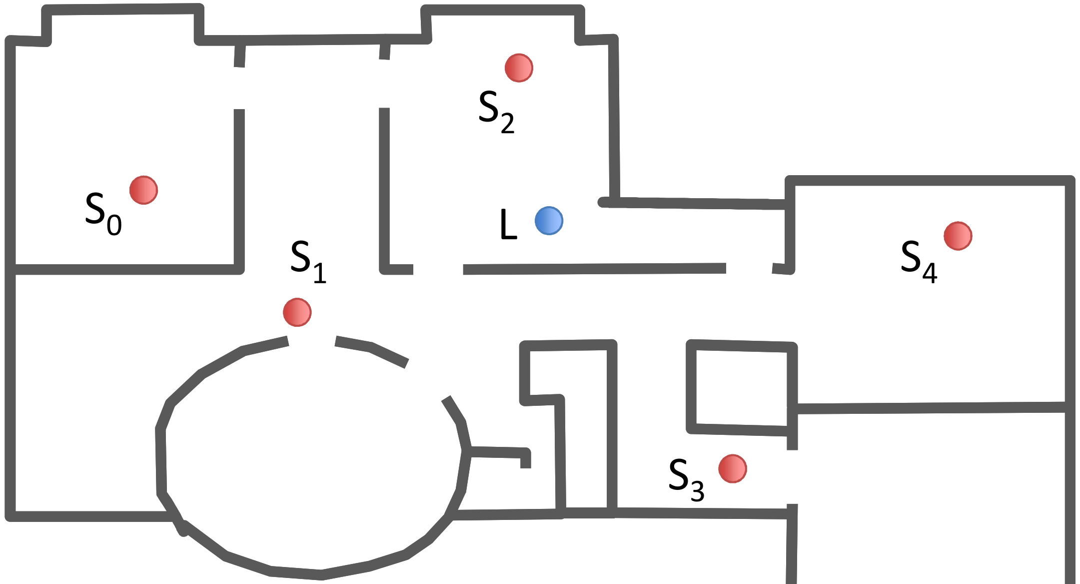

Sound Propagation

Given: source(s), listener(s), obstacle(s), propagation medium

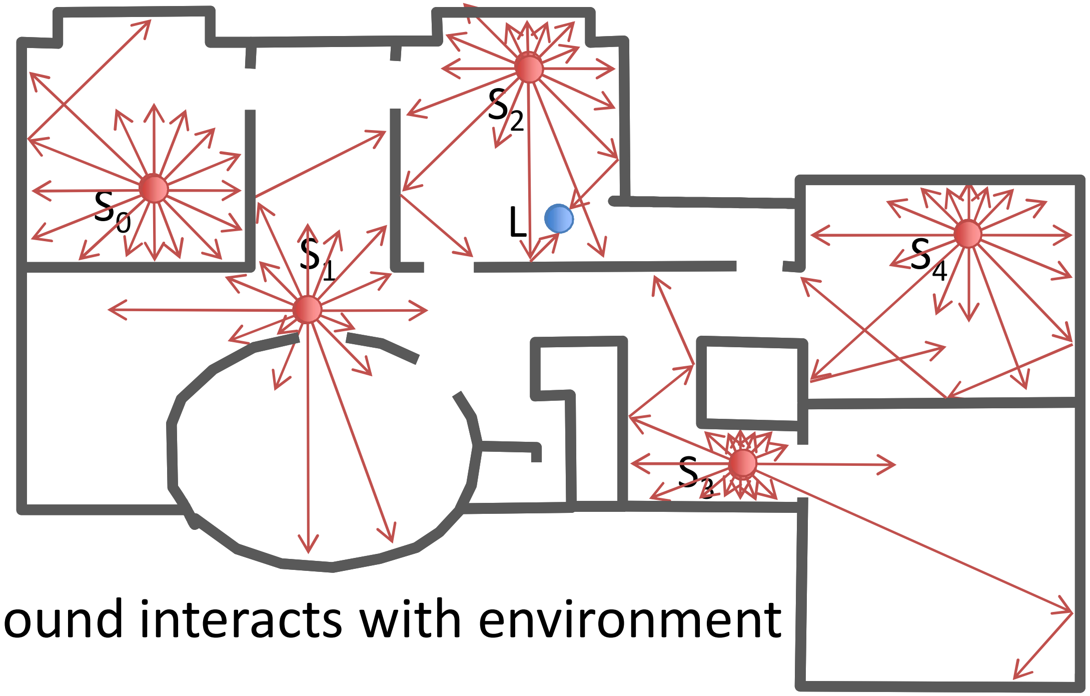

Determine how sound interacts with environment

Sound Propagation Phenomena

Specular reflections
Diffuse reflections
Diffraction
Transmission
Early Reflections
Late Reverberation
Source Modeling
Spatial Sound

There are many different types of complex interactions that must be modeled.

For interactive applications, propagation must be updated at >10Hz

Sound propagation is a challenging task!

Sound Propagation Phenomena

Specular reflections
Diffuse reflections
Diffraction
Transmission
Early Reflections
Late Reverberation
Source Modeling
Spatial Sound

Sound is reflected about the normal at the same angle it arrived.

Good for 'Mirror'-like surfaces, perfect reflectors

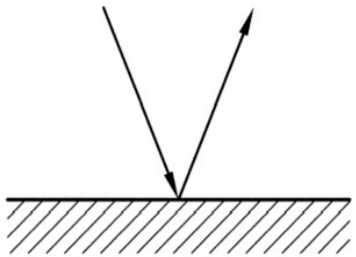

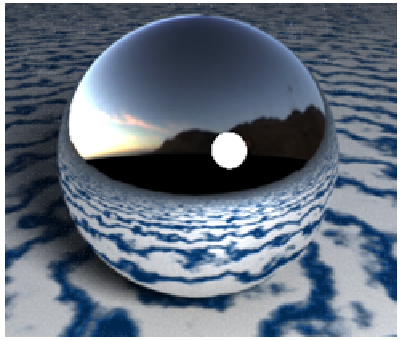

Sound Propagation Phenomena

Specular reflections
Diffuse reflections
Diffraction
Transmission
Early Reflections
Late Reverberation
Source Modeling
Spatial Sound

Sound is scattered when reflected due to small surface variations

Frequency-dependent scattering, BRDF

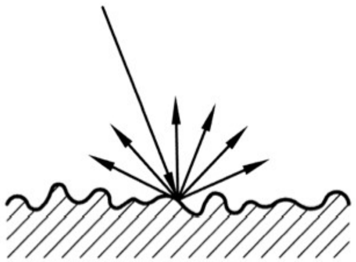

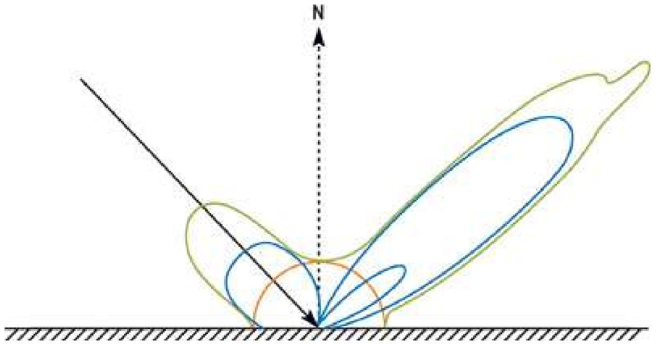

Sound Propagation Phenomena

Specular reflections
Diffuse reflections
Diffraction
Transmission
Early Reflections
Late Reverberation
Source Modeling
Spatial Sound

Low-frequency sound is scattered by objects or features of similar size to the wavelength.

Hear sources without being able to see them

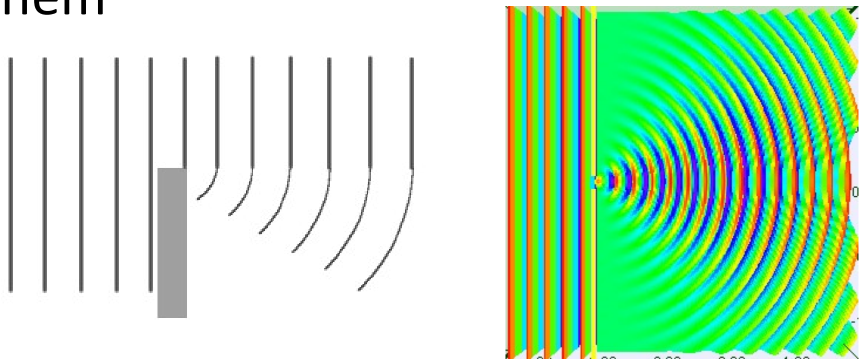

Diffraction

Listen to low-pass diffraction effect

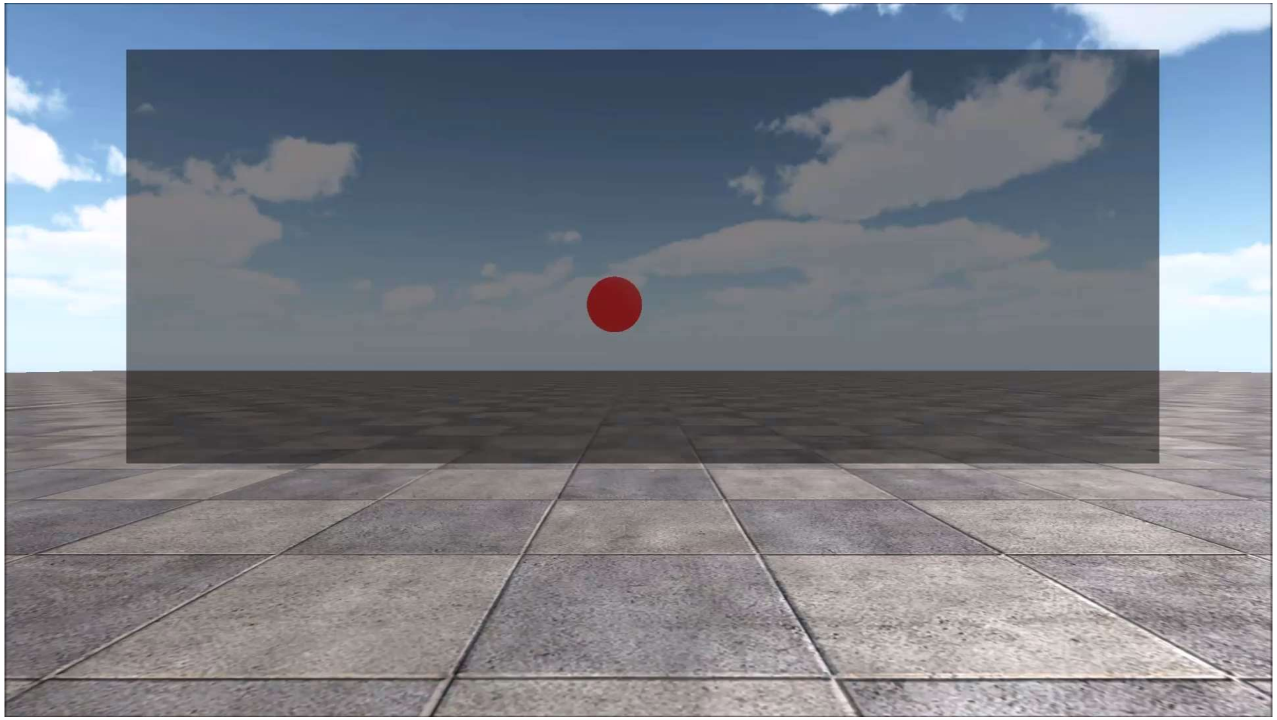

Sound Propagation Phenomena

Specular reflections
Diffuse reflections
Diffraction
Transmission
Early Reflections
Late Reverberation
Source Modeling
Spatial Sound

Sound is transmitted into and through a material, exits and continue propagation

Effects like refraction, attenuation, different speed of sound

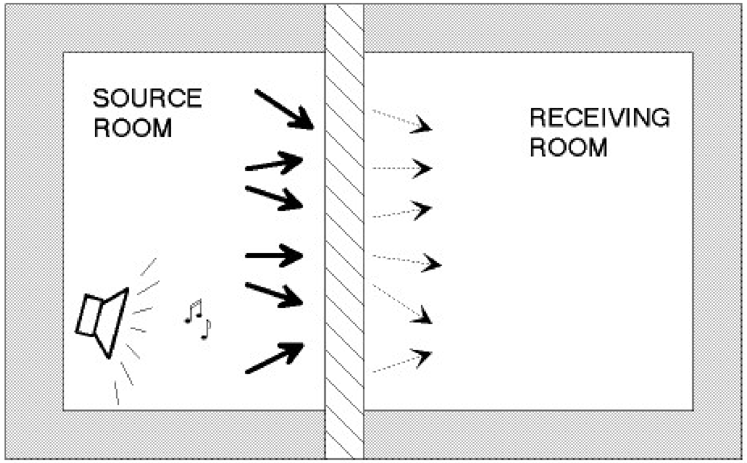

Sound Propagation Phenomena

Specular reflections
Diffuse reflections
Diffraction
Transmission
Early Reflections
Late Reverberation
Source Modeling
Spatial Sound

The first sound paths that arrive at the listener

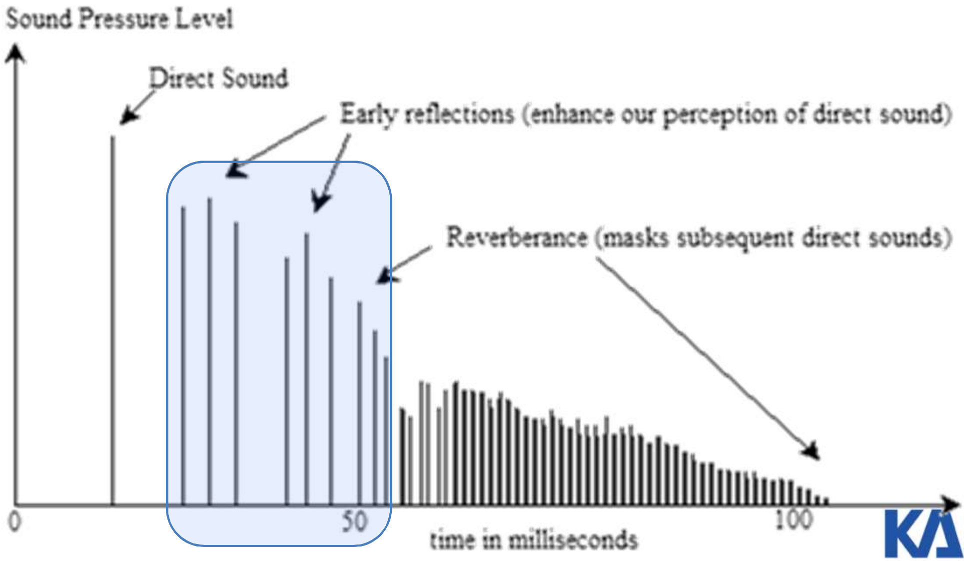

Sound Propagation Phenomena

Specular reflections
Diffuse reflections
Diffraction
Transmission
Early Reflections
Late Reverberation
Source Modeling
Spatial Sound

Many many paths arrive at the listener after 50-100 ms

- Distinct echos transition to smooth reverberation

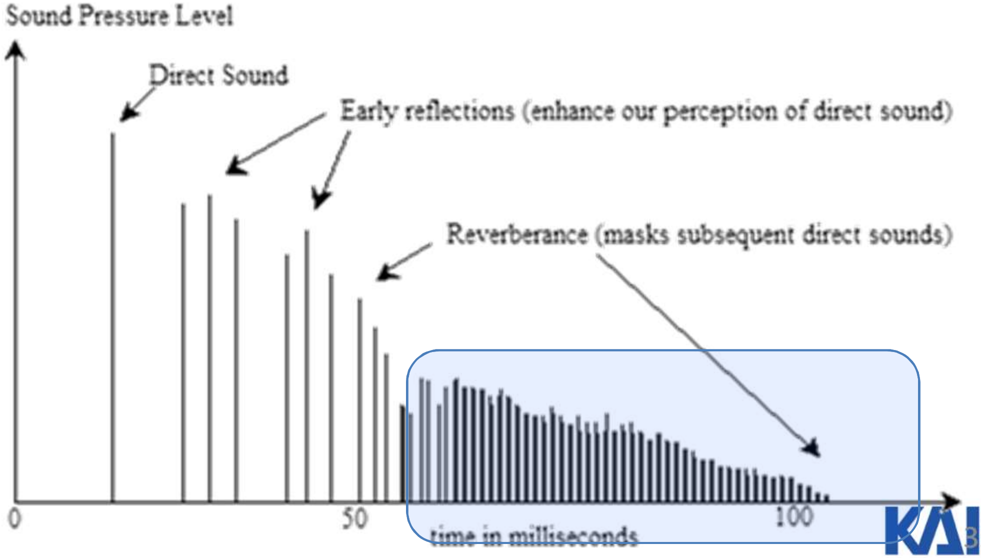

Reflections and Reverberation

Distribution of reflections over time

- **Early reflections**
  Heard as distinct echoes, give sense of individual objects (walls, buildings)

Early Reflection Audio

- **Late reflections (Reverberation)**
  Heard as smooth, long echoes, give sense of size, shape of space

Late Reflection Audio

Reverberation Characteristics

Varies with Shape

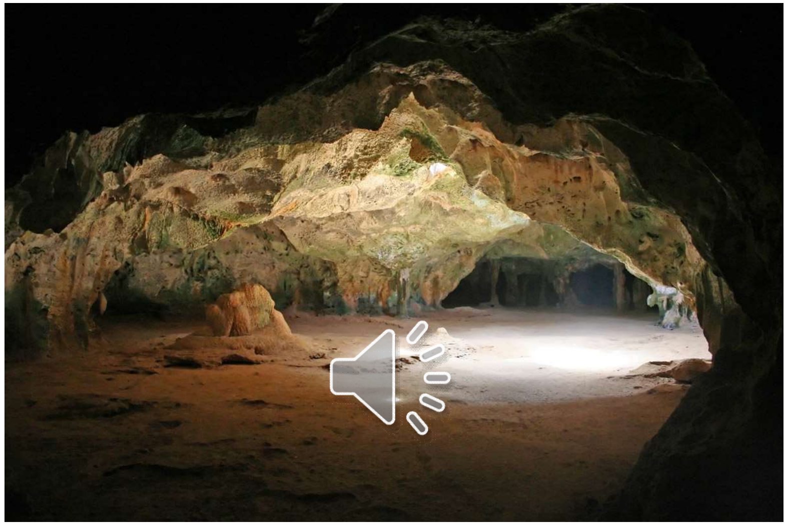

Reverberation Characteristics

Varies with Acoustic Material

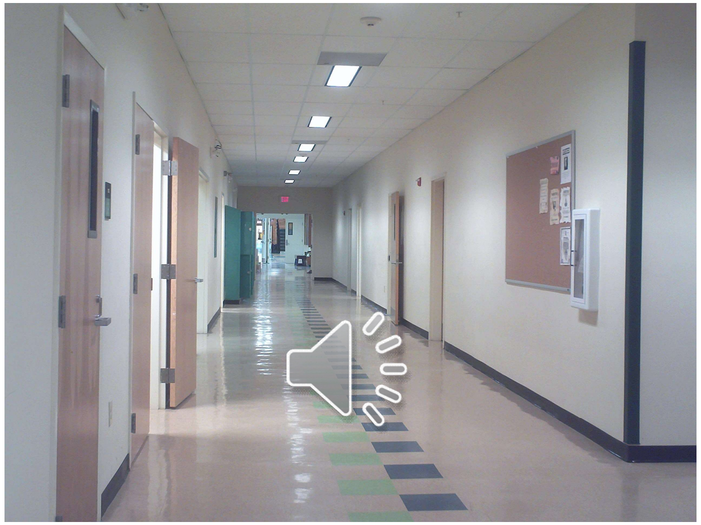

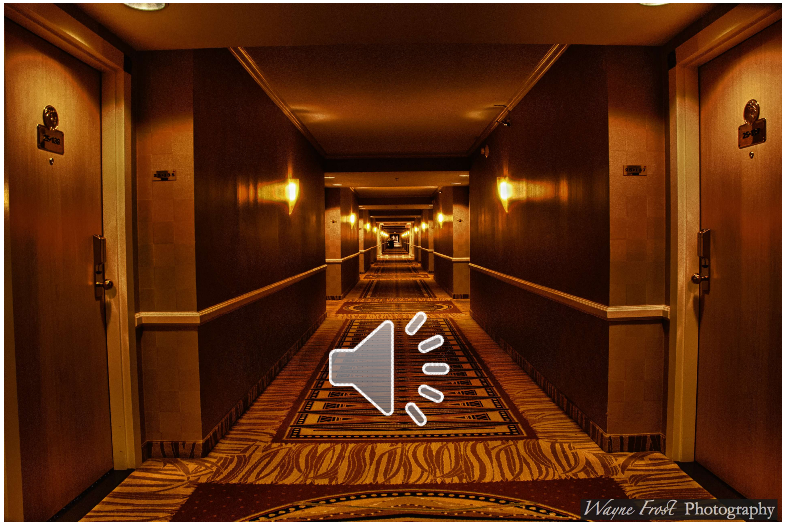

Reverberation Characteristics

Varies with Scene

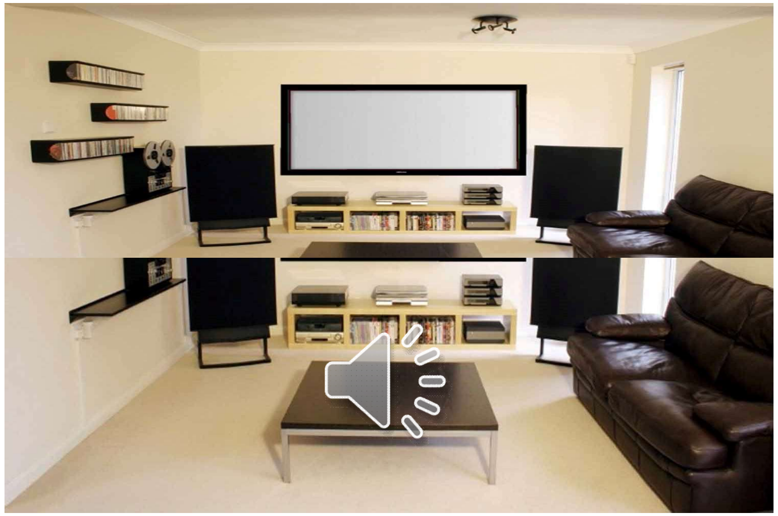

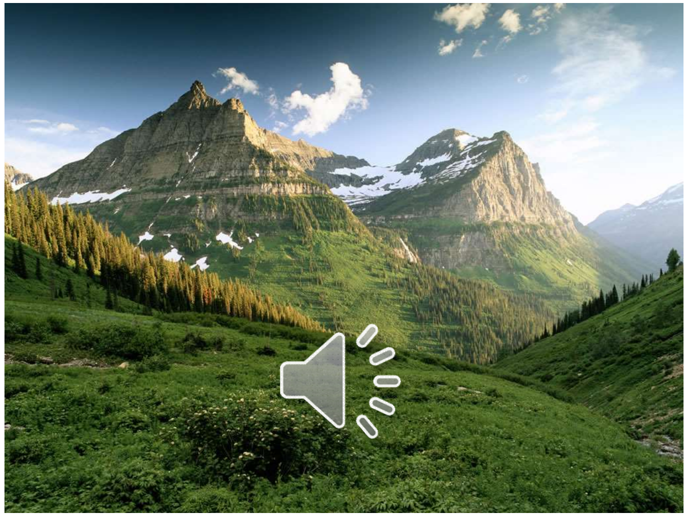

Sound Propagation Phenomena

Specular reflections
Diffuse reflections
Diffraction
Transmission
Early Reflections
Late Reverberation
Source Modeling
Spatial Sound

Sound sources can have different representations that must be handled.

Point sources
Area sources
Directional sources

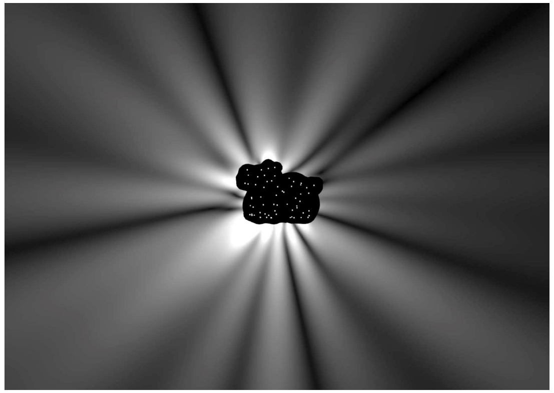

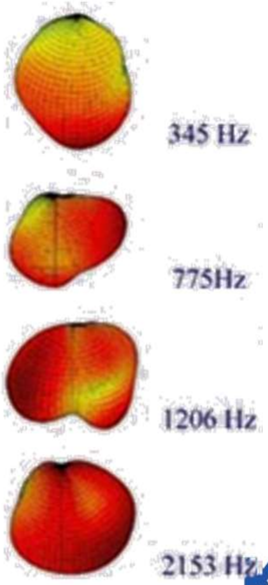

Sound Propagation Phenomena

Specular reflections
Diffuse reflections
Diffraction
Transmission
Early Reflections
Late Reverberation
Source Modeling
Spatial Sound

To produce effects like localization, directional sound effects must be modeled.

Different sound arrives at each ear

Brain interprets difference between signals to determine direction.

Head-Related
Transfer Function (HRTF)

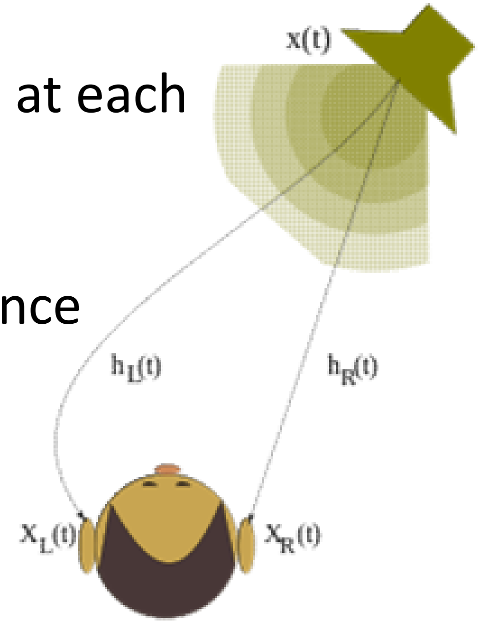

Impulse Responses

Compute: impulse response (IR) for each source-listener pair

IR: time-domain 1D filter
- captures response of linear system (sound propagation)
- can have directional component
- multiple frequency bands
- convolve with source audio to get propagated sound

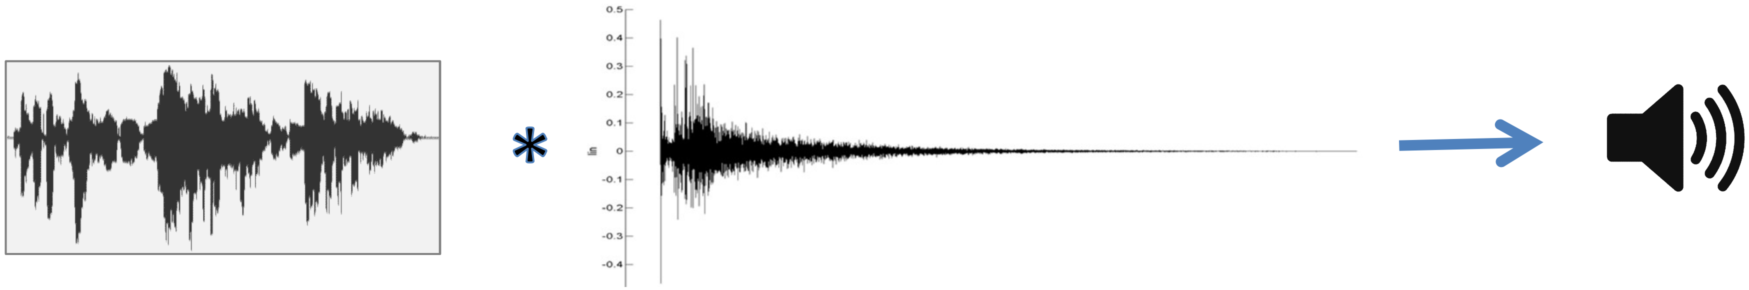

Definition: Sound Path

A path through the scene from a source to a listener.

Consists of:

- Series of multiple interactions with scene:
  - reflections, diffractions, transmissions, change of media
- Delay time (distance)
- Attenuation factor (frequency-dependent)
- Directions:
  - from source
  - from listener

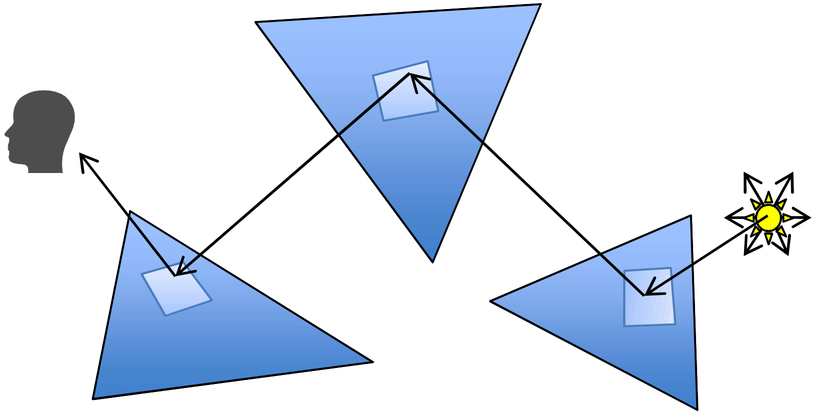

Frequency-Dependent Propagation

Model attenuation factors as a vector of frequency coefficients.

- e.g. – (63Hz, 125Hz, 250Hz, 500Hz, 1000Hz, 2000Hz, 4000Hz, 8000Hz)

Apply coefficients to filtered frequency bands during audio rendering

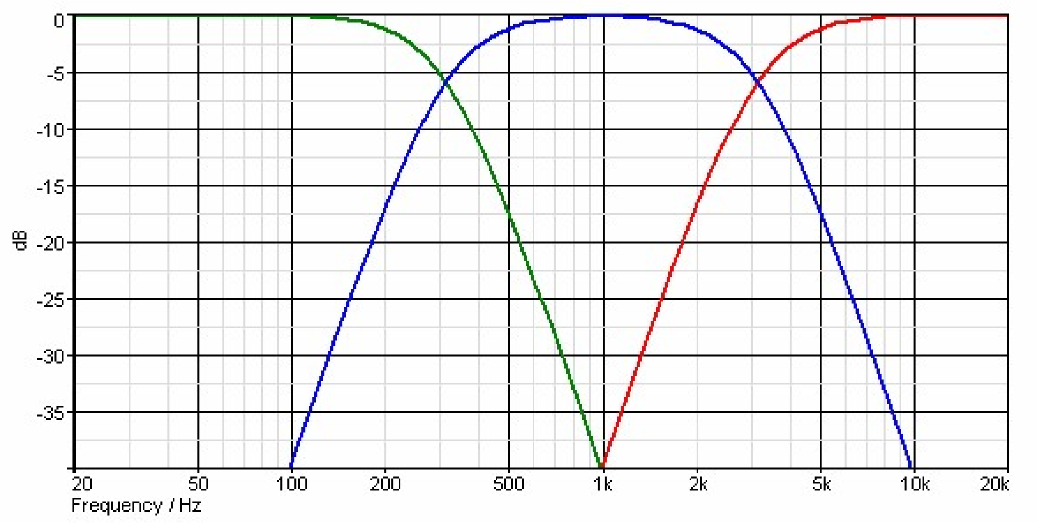

Sound Materials

A model of how sound interacts with a surface/object:

- Reflection ***R***: how much of incident sound is reflected?
- Scattering ***S***: how much of reflected sound is diffusely scattered?
- Transmission ***T***: how much of incident sound is transmitted through material
- $R + T \leq 1$
- All parameters can vary for different frequencies!

Discovering Environmental Interactions

How to figure out what an environment sounds like

- Follow **sound waves** as they **propagate**

- Most popular approach: **ray tracing**
  - Algorithms exist for reflections, scattering
  - New techniques developed for diffraction

- Output of ray tracing: **room impulse response**
  - Similar to HRTF, but models effect of environment

History of Environmental Audio

The even grittier origin story

- **1930**
  Early research began

- **Late 1990s**
  Advances in environmental audio for games

- **2005 – Present**
  Renewed research interest

New Developments in Environmental Audio

What is different this time around?

- **Real-Time Ray Tracing**
  Lots of research in early 2000s
- **More Data-Parallel Compute**
  SIMD, GPU computing, …
- **New, Advanced Algorithms**
  Precomputation algorithms: like lightmapping for sound

Background: Ray Tracing

Given:

- obstacles represented by geometric primitives (e.g. triangles)
- ray start position, ray direction, max distance

Find:

- nearest intersection point
- surface normal
- object ID
- primitive ID
- material ID

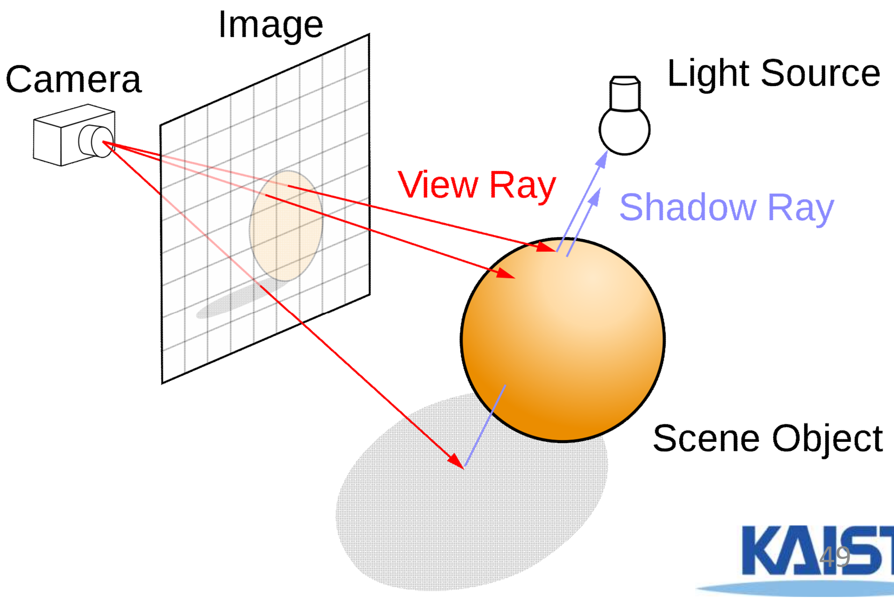

Background: Ray Tracing

Commonly used for offline rendering in graphics

Lots of work on acceleration structures:
- KD-tree
- Bounding volume hierarchy: BVH
- SIMD

Background: Binary Space Partition (BSP)

Recursively partition space by planes
- Cells are convex.
- What happens when object is cut by plane?
    - need to split object

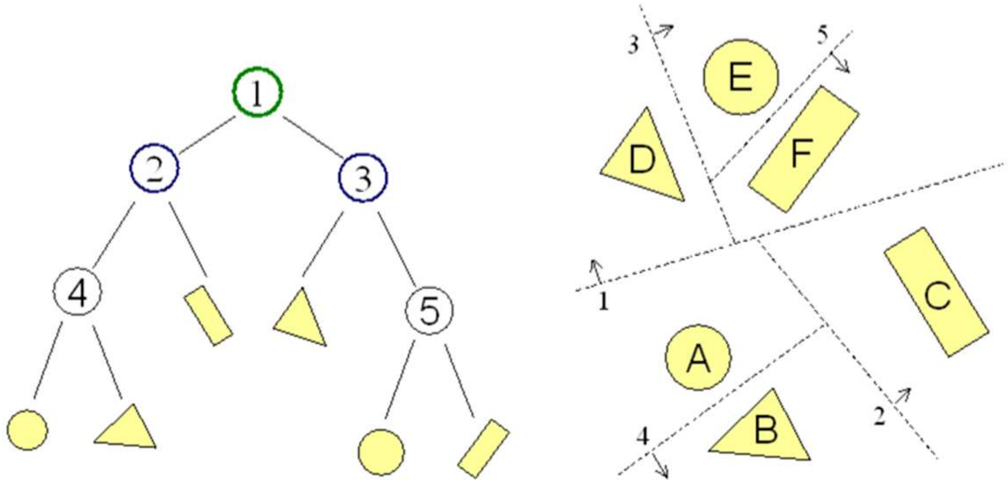

Background: KD-Tree

Special case of BSPs with axis-aligned split planes

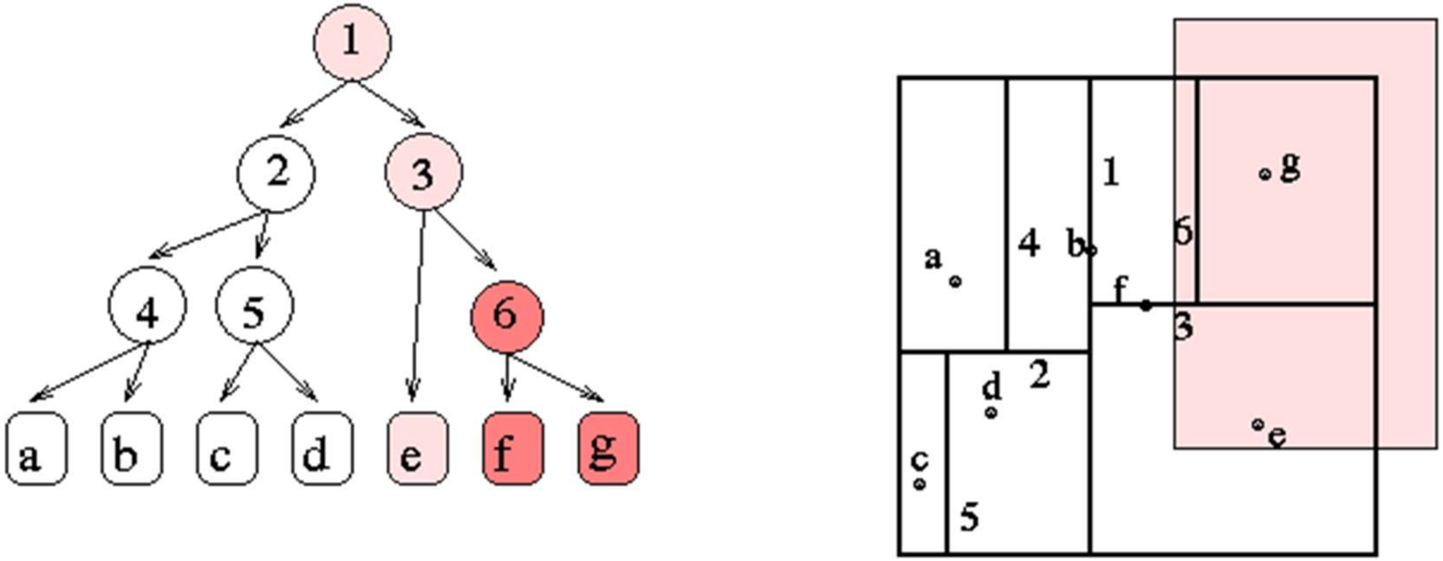

Background: Bounding Volume Hierarchy (BVH)

Hierarchy of bounding volumes: AABB, OBB, spheres, etc.

- Start at root, traverse nodes that ray intersects
- Test vs. primitives in intersected leaf nodes

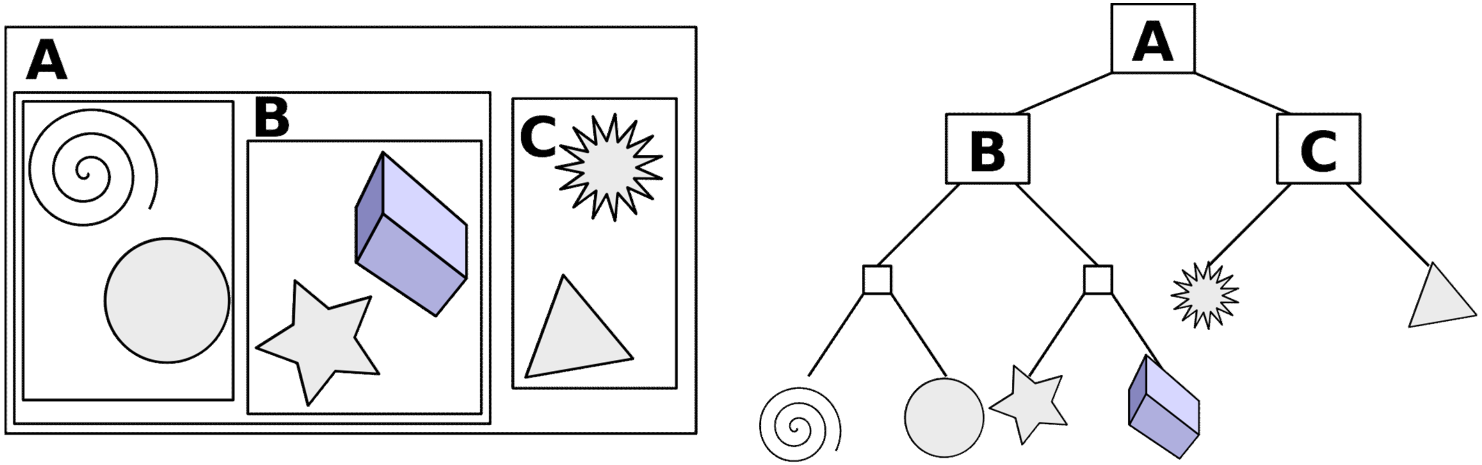

Sound Propagation – Wave vs. Geometric

Wave:	Geometric:
Sound = wave	Sound ≈ particles, acoustic energy
Limited to low frequencies	Better for high frequencies
Complexity: O(volume), O(freq^{4})	O(log(# primitives)) per ray
Pre-computed	Interactive
Static scenes	Dynamic scenes

Geometric Methods

Monte Carlo Methods	[Allred and Newhouse 1958; Haviland and Thanedar 1973]
Image Source Method	[Allen and Berkley 1979; Borish 1984]
Beam Tracing	[Funkhouser et al. 1998; Tsingos et al. 2001]
Frustum Tracing	[Taylor et al. 2009; Chandak et al. 2009]
Ray Tracing	[Krokstad et al. 1968; Vorländer 1989; Lentz et al. 2007; Taylor et al. 2012]

Image Source Method

- Compute set of virtual 'image' sources for a sound source S.

- Each image source corresponds to a **specular** path through the scene.

- Complexity increases with reflection order and # reflecting planes in scene

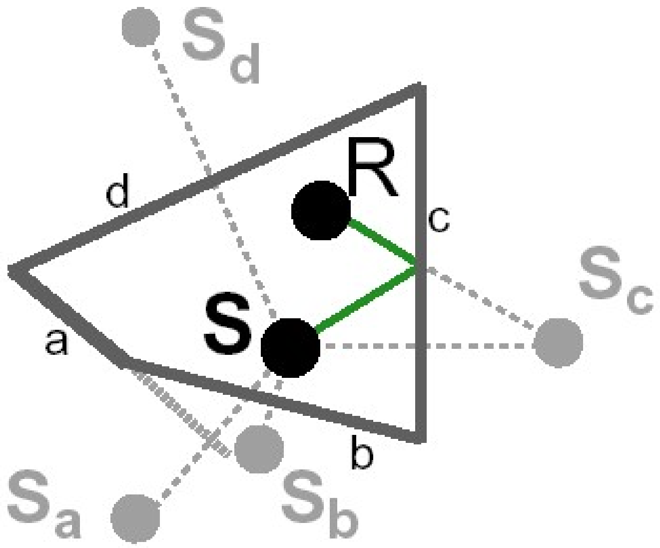

[Allen and Berkley 1979; Borish 1984]

Beam Tracing

Treat sound as 'beams'

Advantages:

- no aliasing issues, good for dynamic listeners
- handles diffraction

Disadvantages:

- expensive preprocessing step
- can't handle dynamic sources or geometry
- can't handle diffuse reflections

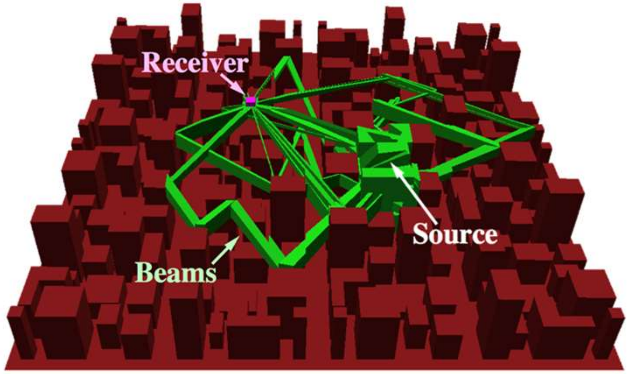

[Funkhouser et al. 1998;
Tsingos et al. 2001]

Beam Tracing

Example: input scene

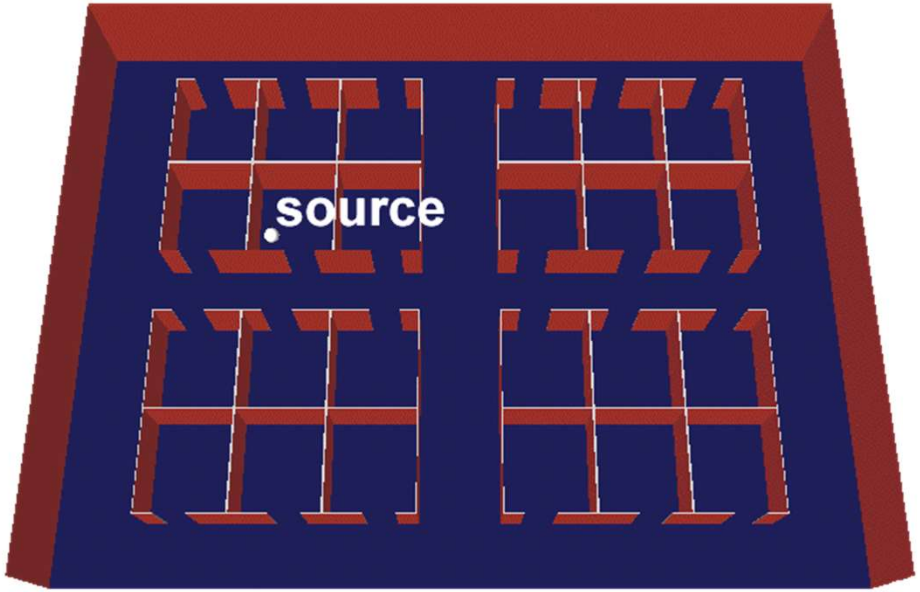

Beam Tracing - Preprocessing

- Partition 3D space into convex regions (BSP tree)
- Build adjacency graph

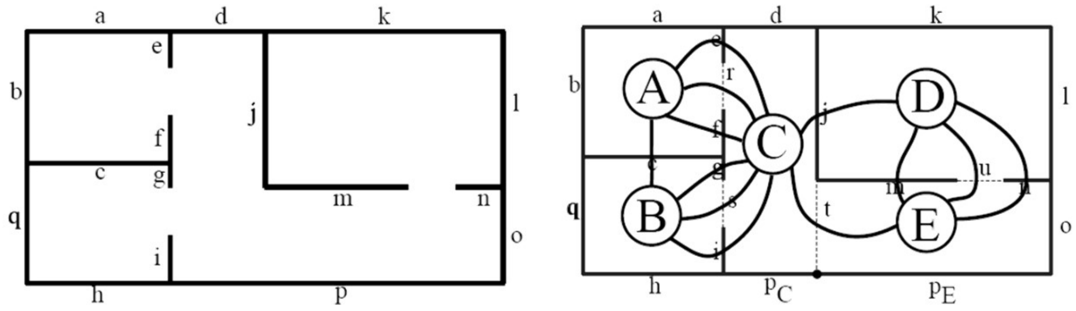

(a) Input model. (b) Cell adjacency graph.

Figure 7: Example spatial subdivision.

Beam Tracing - Preprocessing

- Partition 3D space into convex regions (BSP tree)

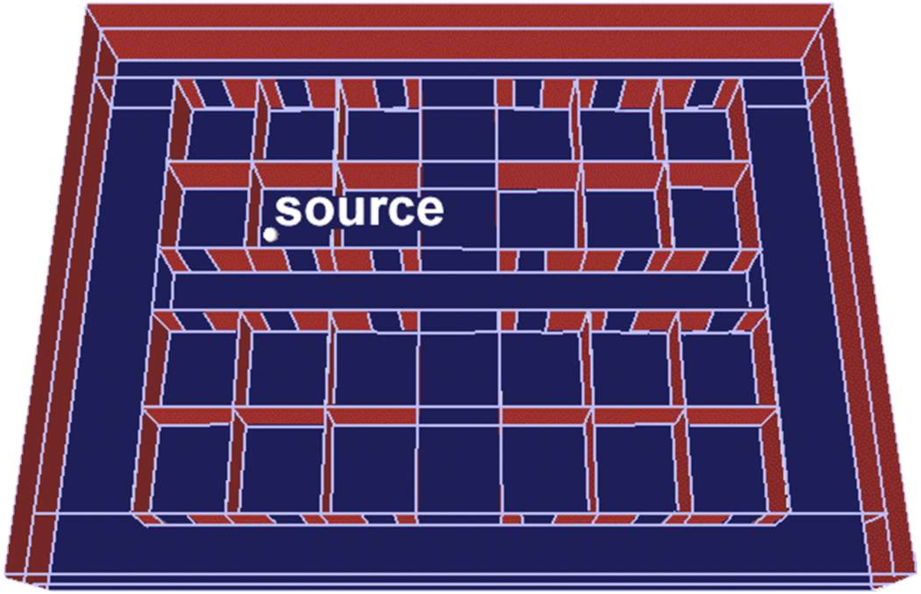

Beam Tracing - Preprocessing

- Compute beam tree
- Node information:
  - cell ID
  - beam & beam apex
  - cell boundary
  - parent node ID
  - attenuation

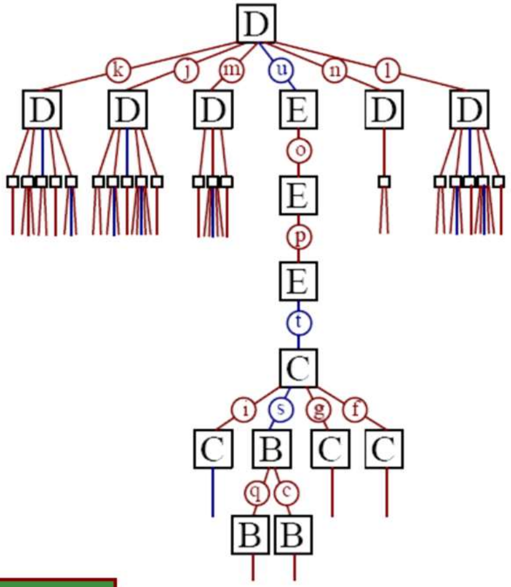

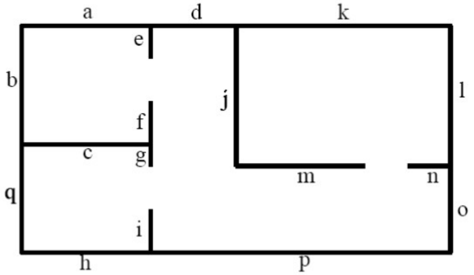

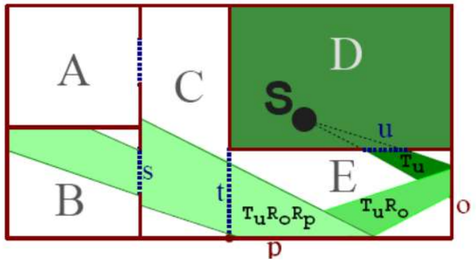

Beam Tracing - Runtime

Find cell containing source position

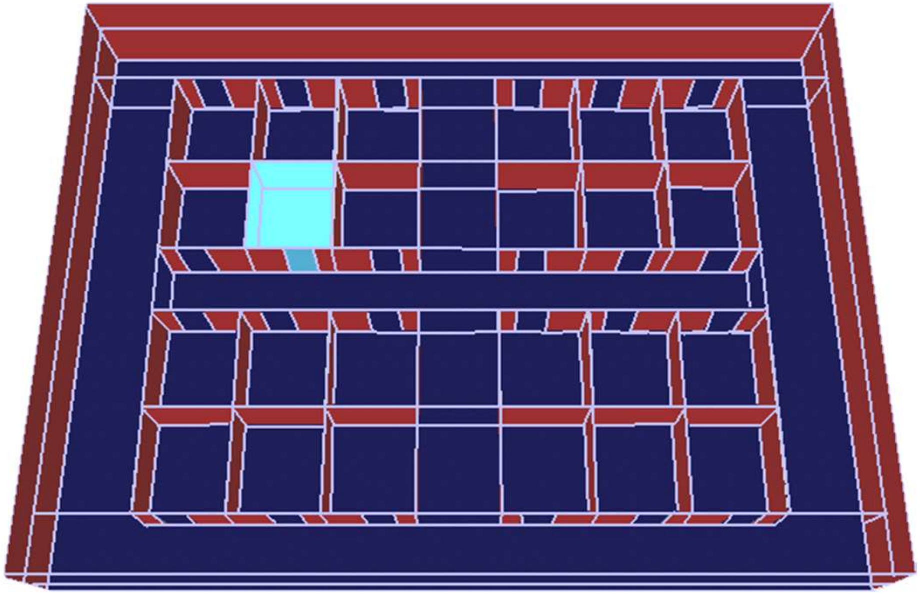

Beam Tracing - Runtime

Example path

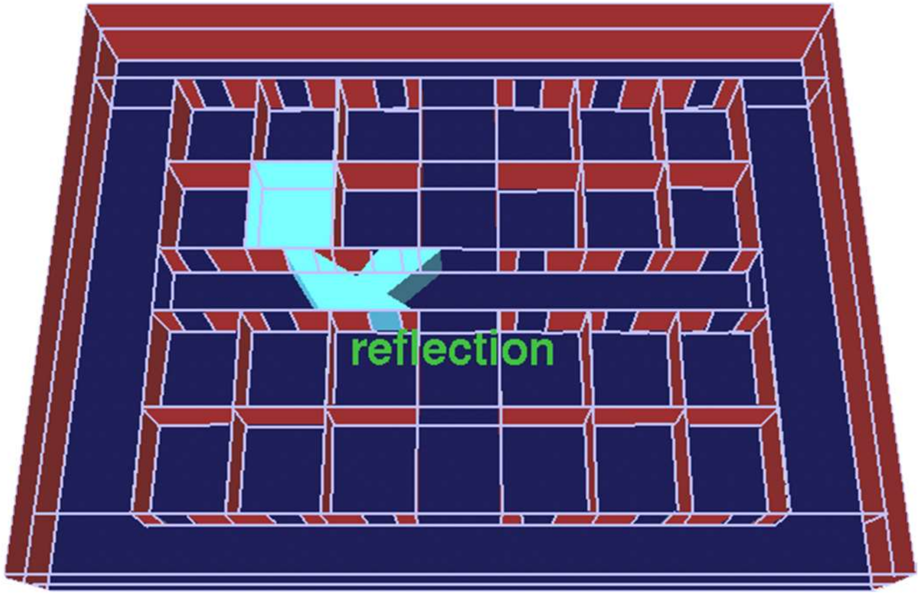

Beam Tracing - Runtime

Example path

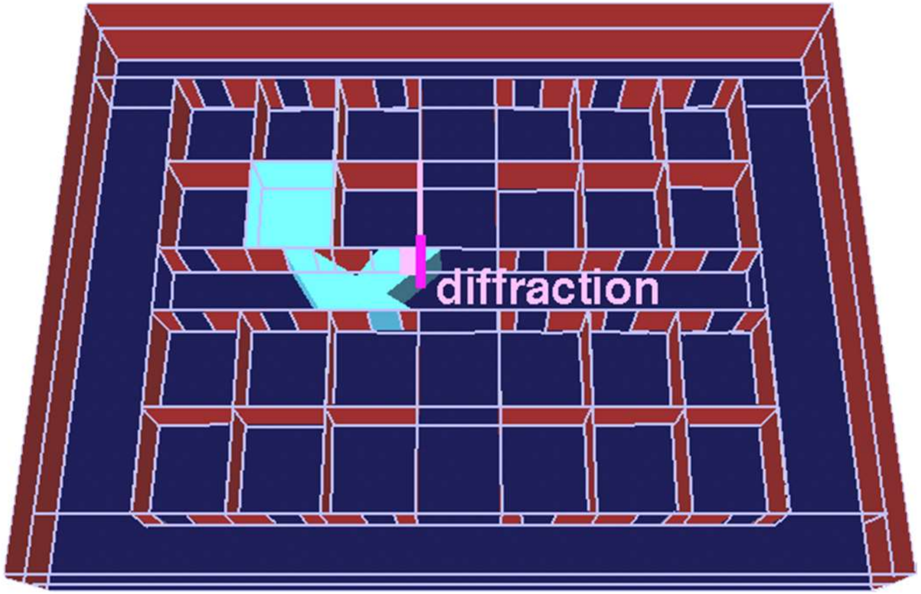

Beam Tracing - Runtime

Example path

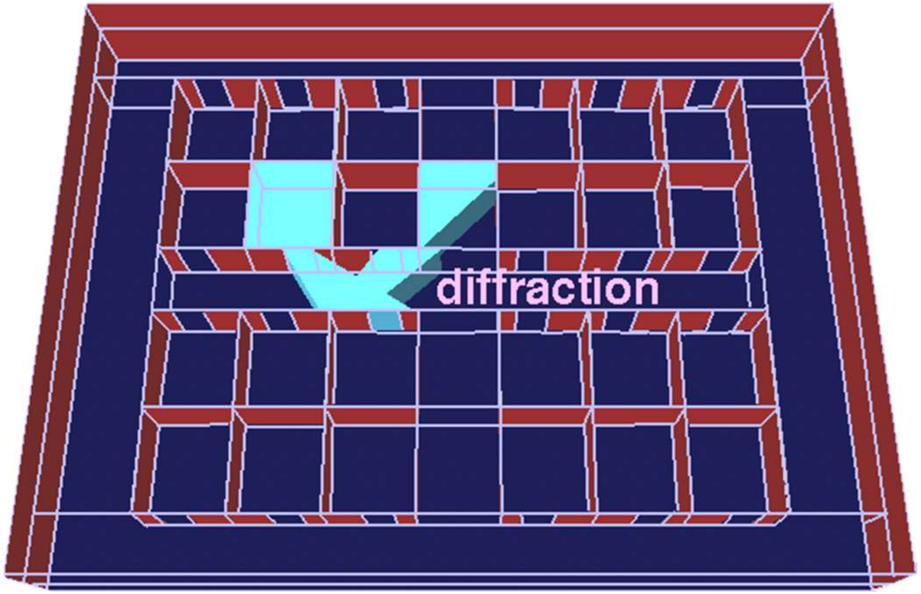

Beam Tracing - Runtime

Example path

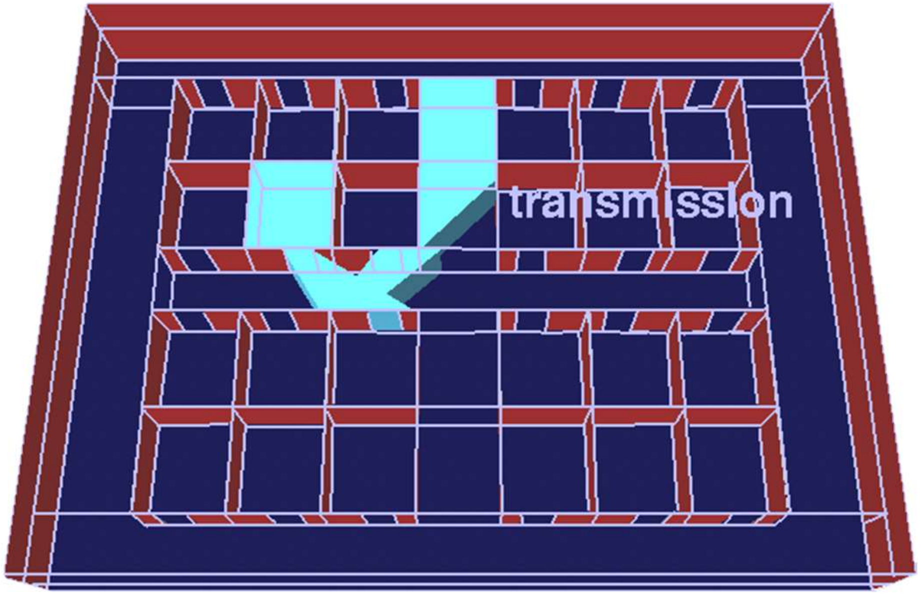

Beam Tracing - Runtime

Example path

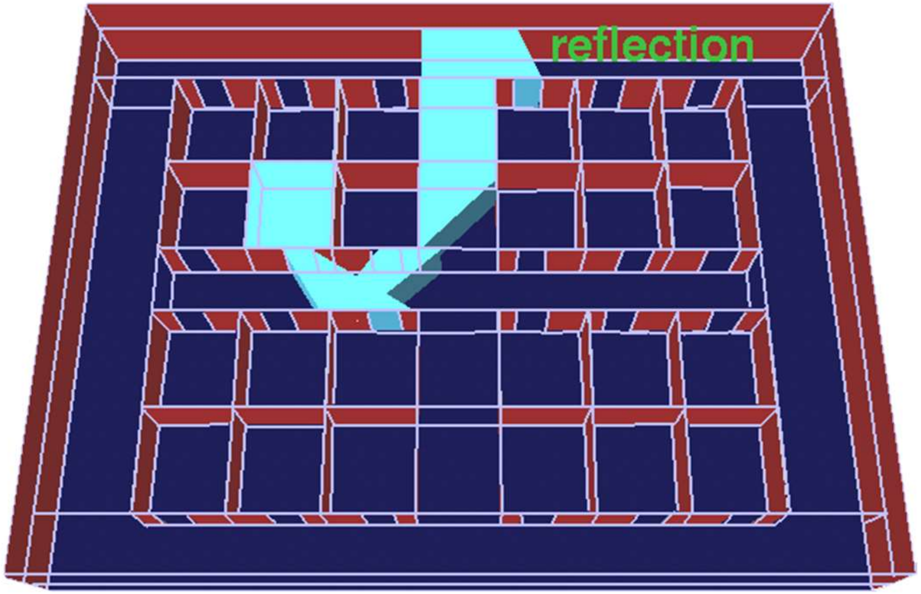

Beam Tracing - Runtime

Example path

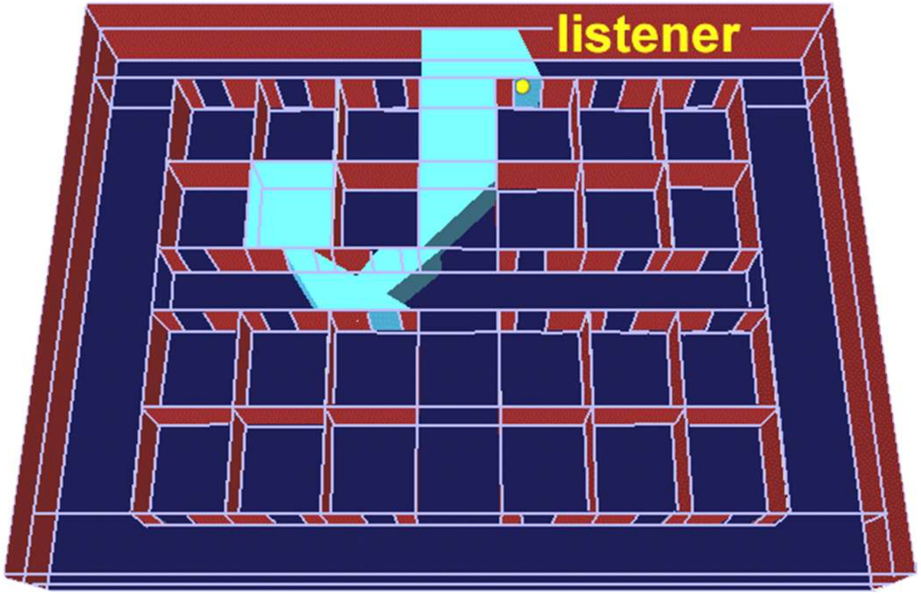

Beam Tracing - Runtime

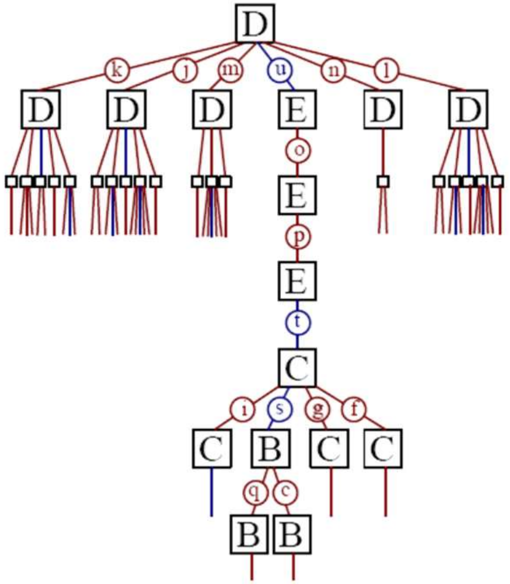

1. Find cell C containing listener (log N)
2. For each beam in C, check if listener is inside
3. If yes, there is a path.
4. Compute attenuation, path length, direction.
5. Construct path by traversing beam tree.
6. Compute IR.

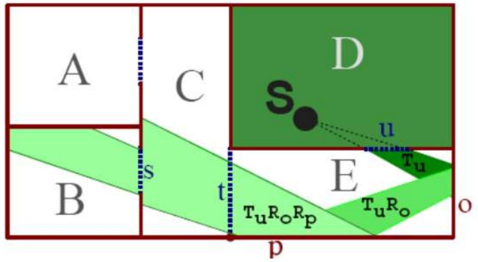

Beam Tracing - Runtime

Example path

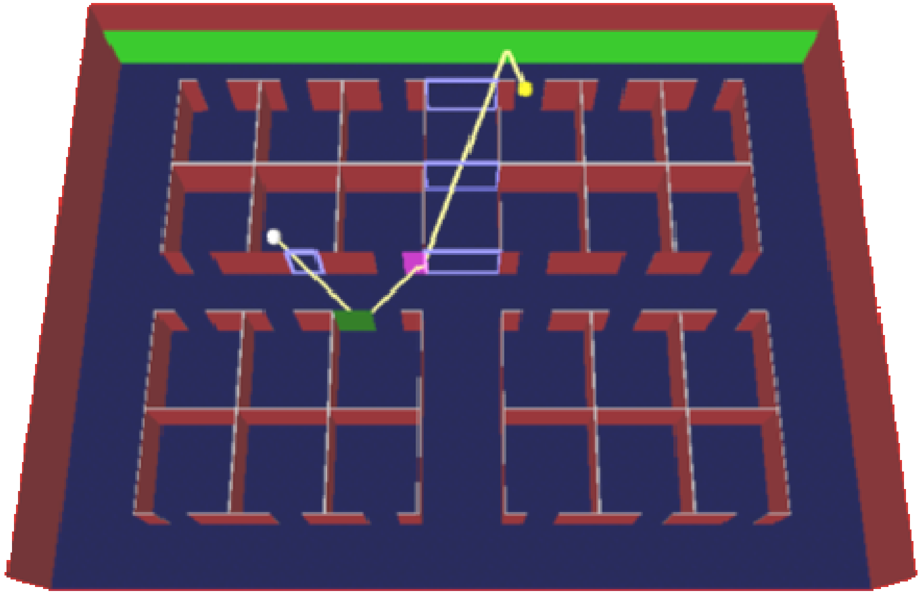

Beam Tracing - Runtime

Example path

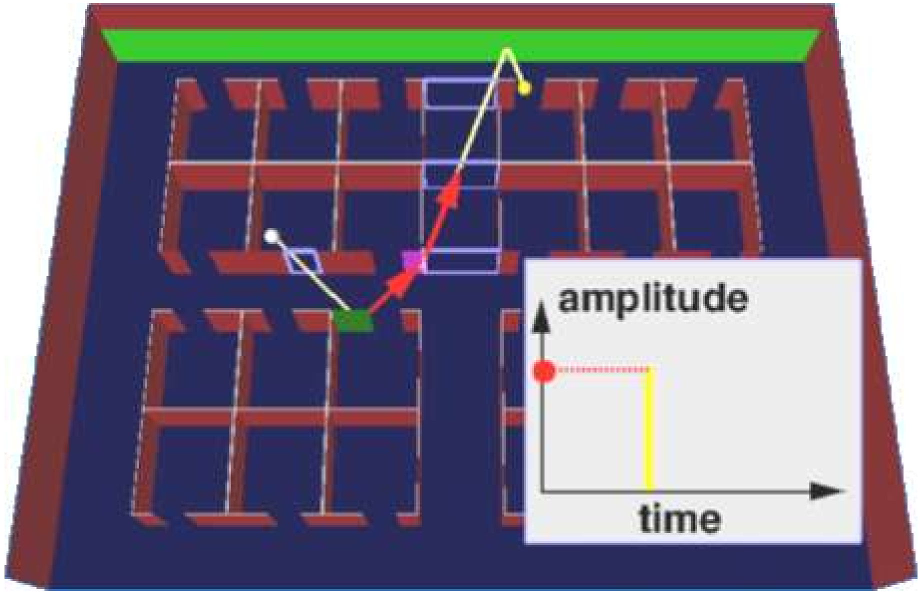

Beam Tracing - Runtime

All paths

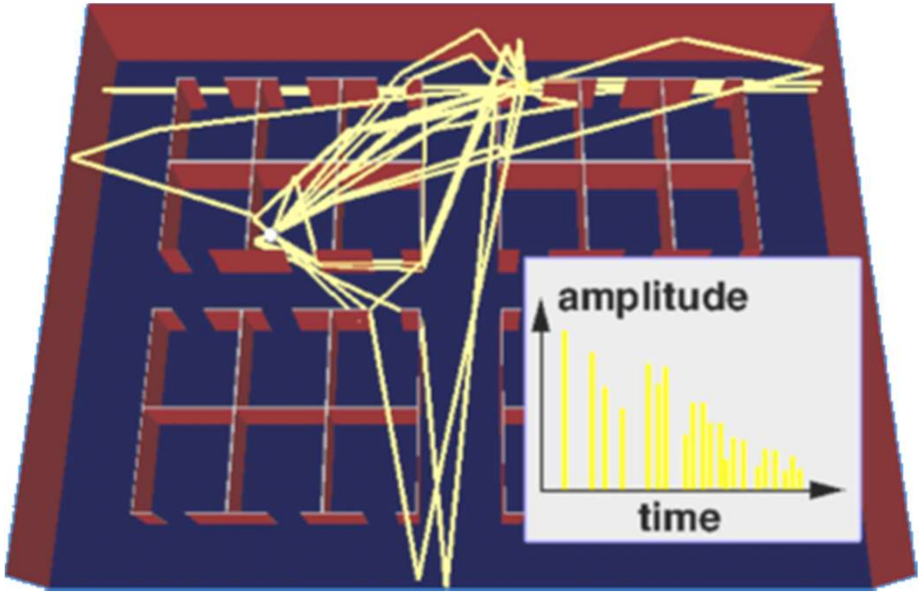

Frustum Tracing

Trace frustums through the scene.

- Handles diffraction, complex geometry, dynamic source/listener.

Disadvantages:

- can miss some paths
- no diffuse reflections

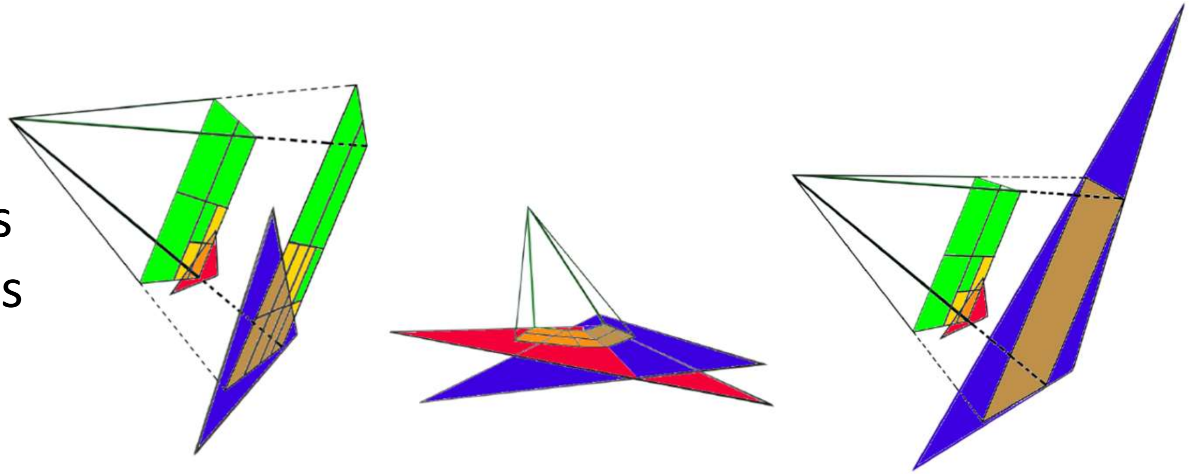

[Lauterbach 2007; Taylor et al. 2009; Chandak et al. 2009]

Frustum Tracing – Pipeline Overview

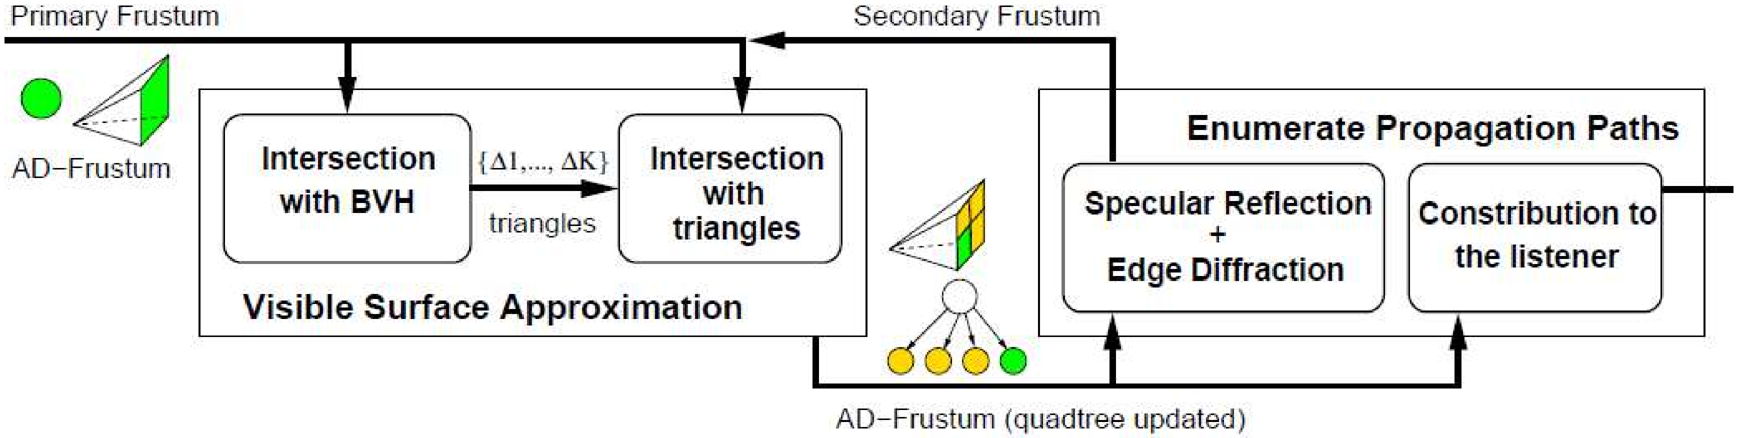

Primary Frustum

AD−Frustum

Secondary Frustum

Visible Surface Approximation

Intersection with BVH

{ΔI,..., ΔK} triangles

Intersection with triangles

Enumerate Propagation Paths

Specular Reflection + Edge Diffraction

Constribution to the listener

AD−Frustum (quadtree updated)

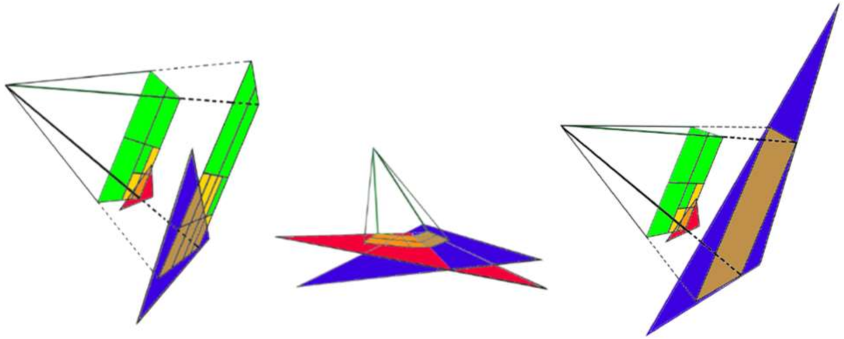

Frustum-Triangle Intersection, Subdivision

Class Objectives were:

- Binaural audio and Head Related Transfer Function (HRTF)
- Sound Propagation Phenomena
- Geometric Methods for Sound Propagation

Next Time…

- **Advanced propagation methods**

Homework

- Go over the next lecture slides before the class
- Watch 2 SIG/I3D/HPG videos and submit your summaries every Tue. class
  - Just one paragraph for each summary

Example:

Title: XXX XXXX XXXX

Abstract: this video is about accelerating the performance of ray tracing. To achieve its goal, they design a new technique for reordering rays, since by doing so, they can improve the ray coherence and thus improve the overall performance.

Any Questions?

- **Submit four times in Sep./Oct.**
- **Come up with one question on what we have discussed in the class and submit at the end of the class**
  - **1 for typical questions**
  - **2 for questions that have some thoughts or surprise me**