Project Triton

Pre-computed Environmental Wave Acoustics

Nikunj Raghuvanshi

Microsoft Research

John Tennant

THE COALITION

Acoustics improves immersion

- We understand our surroundings audio-visually
- Can hear behind head, around corners: persistent awareness
- Need good acoustics in games

- VR/AR –

Positional Audio =
Directional Audio + Environmental Acoustics

Game acoustics today

- Compared to lighting, CPU/RAM are limited…

…but sound's behavior is more complex: wave effects crucial

- Mix of approximate methods w/ pitfalls and manual tagging
- Designer: Tedious work, limited time for creative expression
- Why not follow the path of game lighting?
  Reliable physics + design

The problem

Part I: Accurate acoustic simulation =
robust results for dynamic sources & listener. But expensive!

@GDC 2011: Real-time wave acoustics: full core, GBs of RAM, custom audio engine.
Raghuvanshi, et. al., SIGGRAPH 2010, "Precomputed wave simulation for real-time sound propagation of dynamic sources in complex scenes"

Part II: How to design acoustic data & implement expressive audio?

Triton is the first shipped system to meet both challenges

Triton's solution

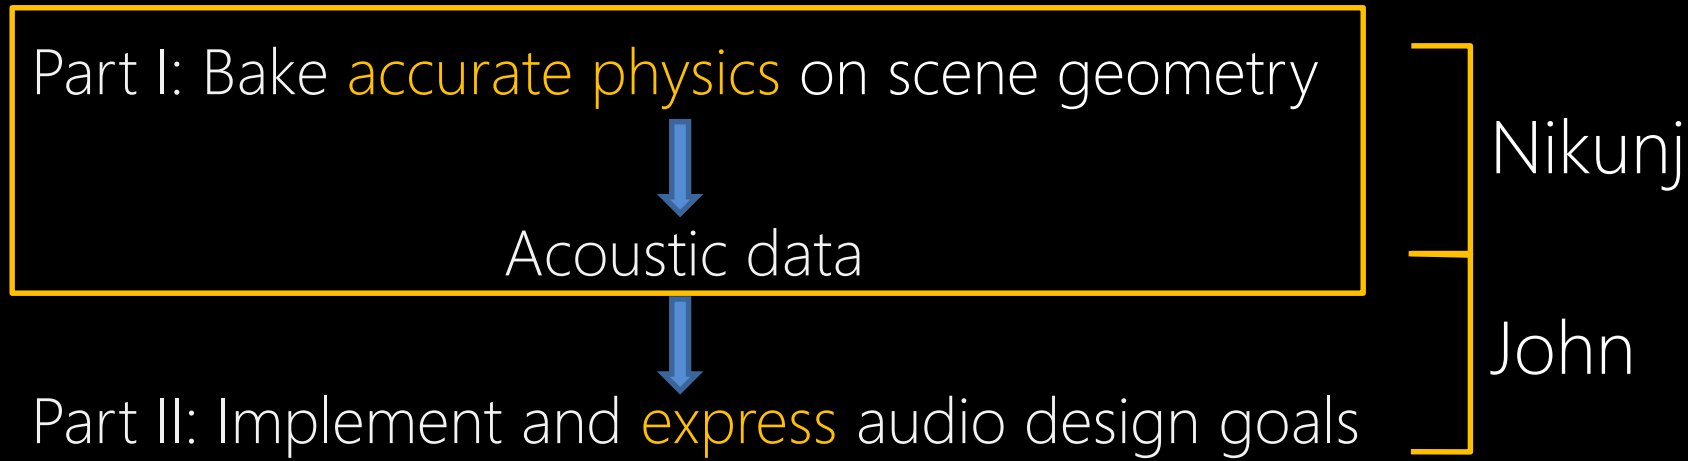

Goal: *Combine fast wave acoustics with expressive audio design to bring game spaces to life*

200× slomo, horizontal slice of 3D simulation

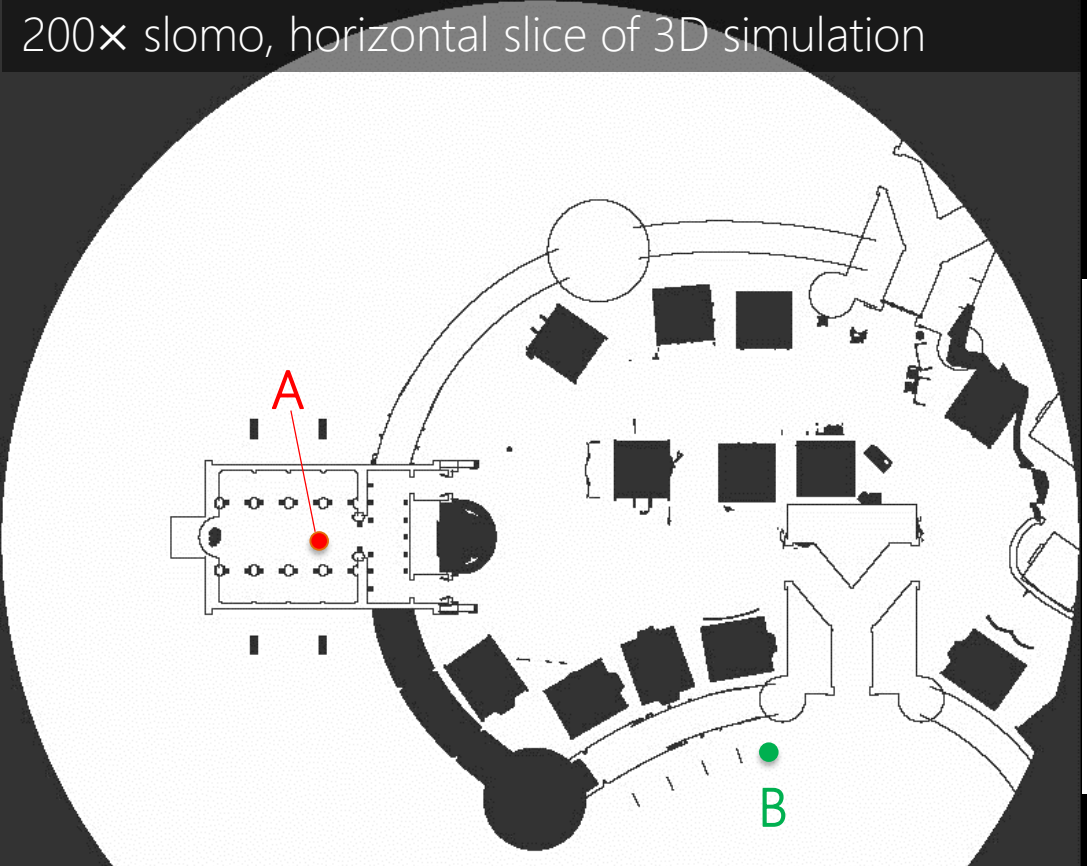

Wave acoustics

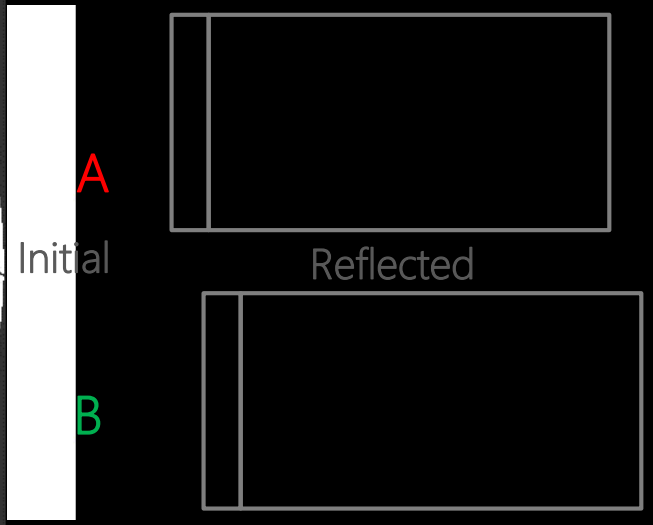

Impulse Response

- Diffraction: Sound waves bend, we hear around corners.
- Reverberation: We hear sound bouncing around

100 m

Microsoft Research

Important acoustic effects

1. Obstruction
2. Occlusion
3. Wetness ratio
4. Reverb decay rate

…

Effects 1 & 2: Obstruction & occlusion

Obstruction ~ initial (direct) sound energy
Occlusion ~ initial (direct) + reflected (indirect)

Strong obstruction, little occlusion

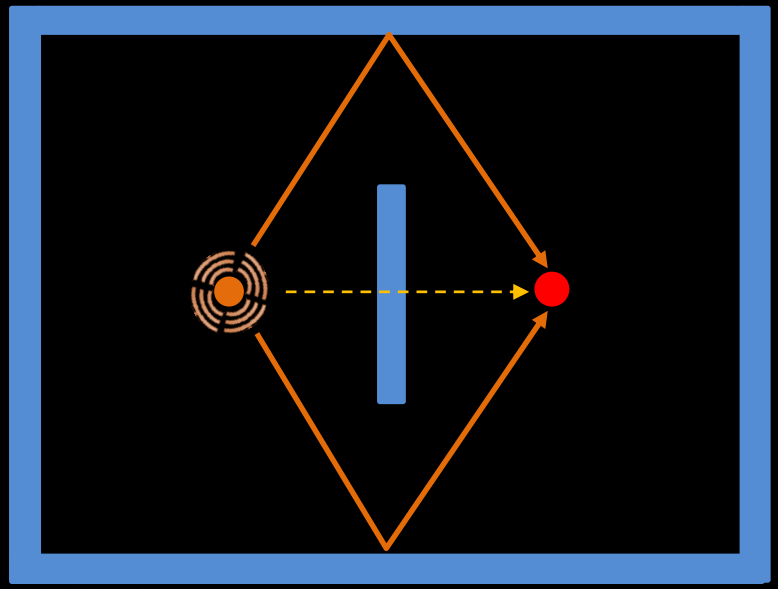

Strong obstruction, strong occlusion

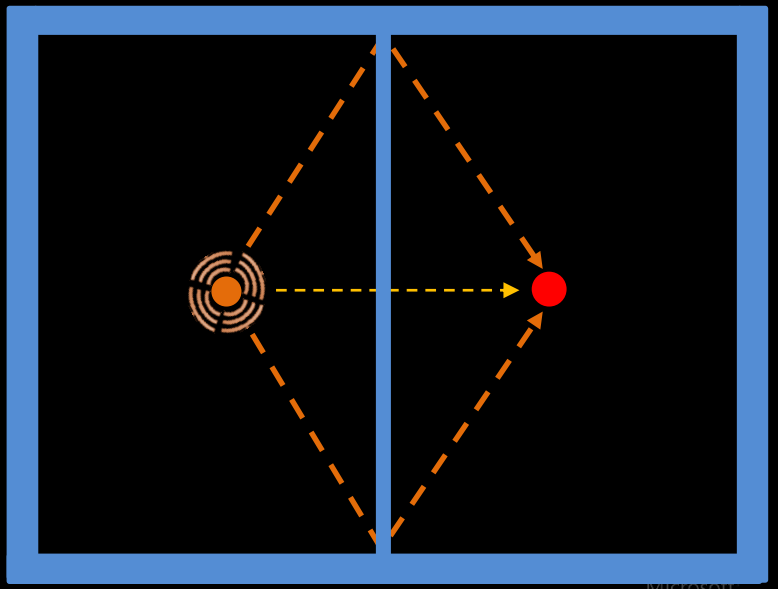

Effect 1: obstruction (initial energy)

Existing approximate method: Shortest path length

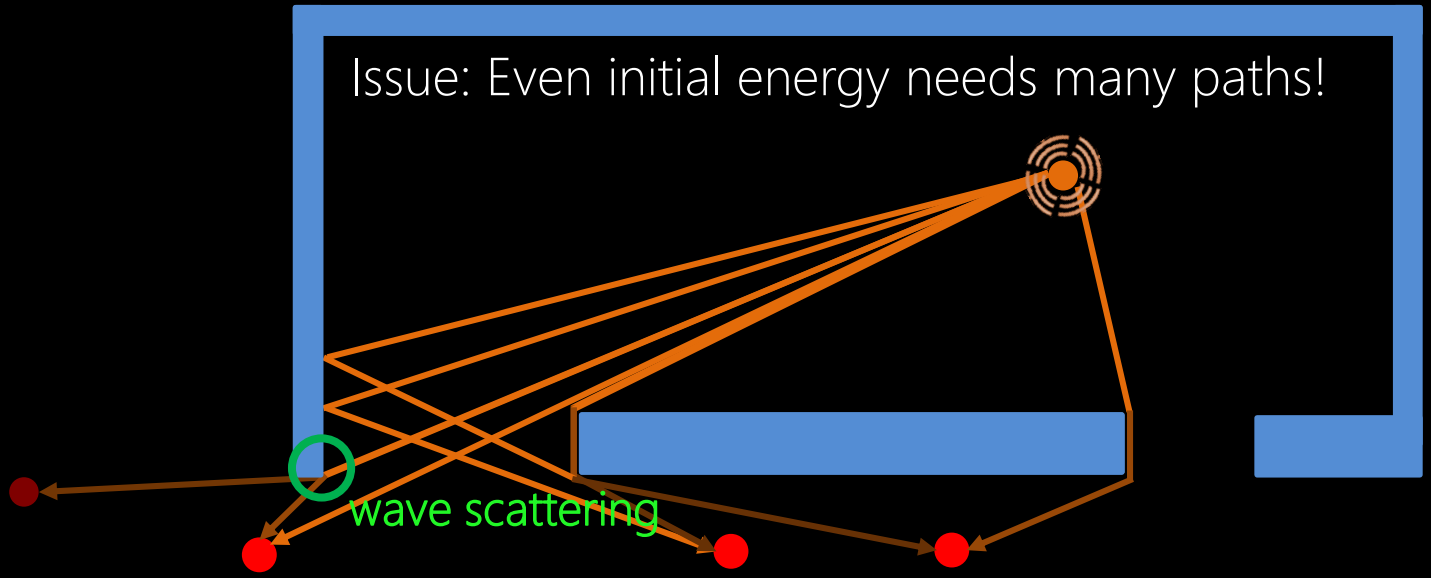

Upshot: Risk of making sound too weak

Effect 2: Occlusion?

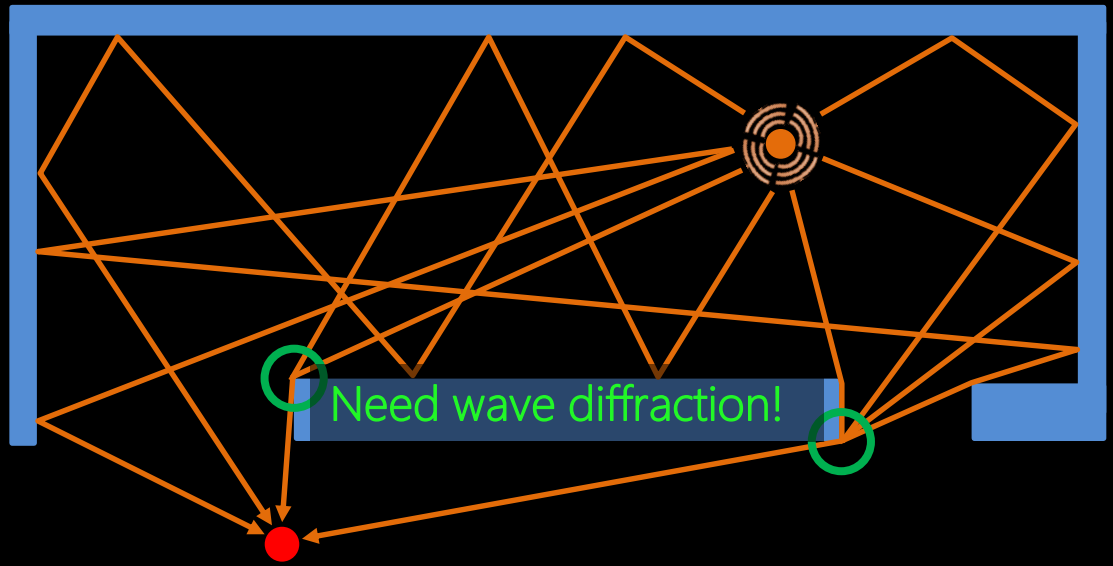

Ignore diffraction → sound too weak

Missing paths → glitches

Fast approximate method does NOT exist!

Effect 3: Wetness ratio (sense of distance)

Wetness ~ ratio of reflected and initial energy

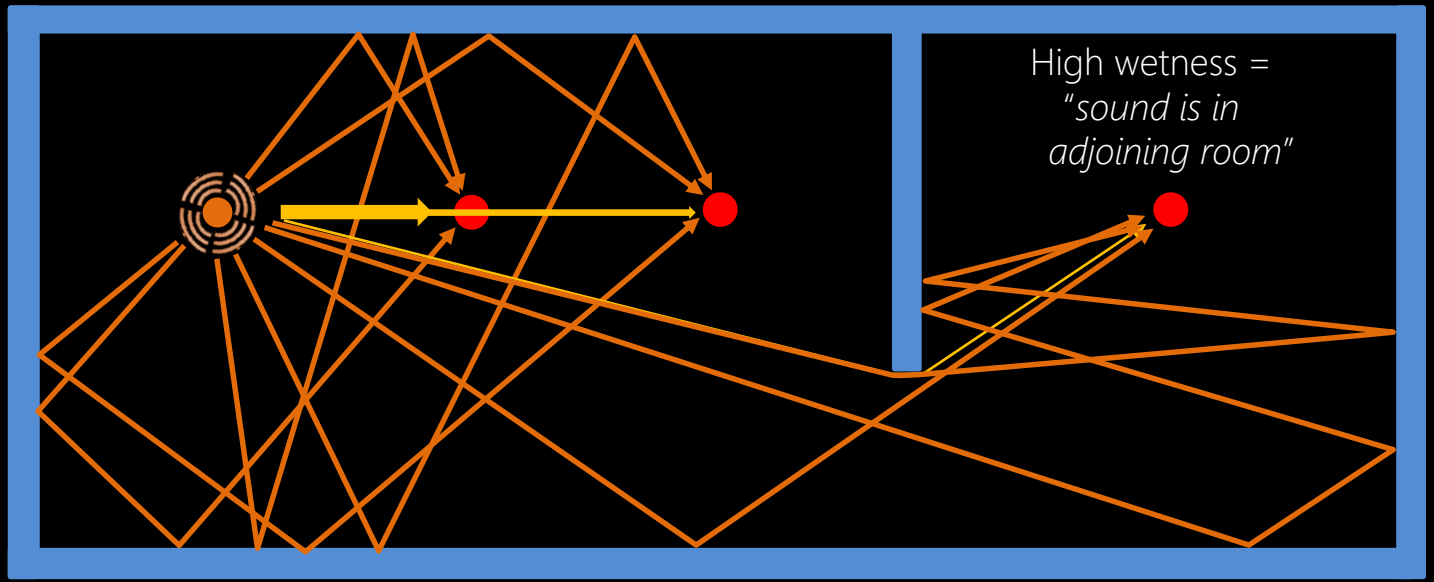

Common method: Designer specifies reverb distance roll-off curves

Room size and effect of geometry are ignored!


Triton OFF
Pause(No) LoadScreen(No)
You do NOT have debug mem enabled
1. GearMainMenuEntry_P
2. SP_Farm_B_P
3. GearMainMenuEntry_P
4. SP_Farm_B_P
5. GearMainMenuEntry_P
6. SP_Farm_B_P
7. GearMainMenuEntry_P
8. SP_Farm_B_P
2 checkpoint reloads
SERVER
Warnings: 20686
1st Player Difficulty: Casual


Triton ON
Pause(No) LoadScreen(No)
You do NOT have debug mem enabled
1. GearMainMenuEntry_P
2. SP_Farm_B_P
1 checkpoint reloads
Play_VEH_DB_Kestrel_Loop | Occ: -5.4, Obs: 0.0
Dry: -5.42
Wet: -36.57
Refl. decay: 0.93
Reverb decay: 0.94
Warnings: 20803
Outdoorness: 0.07
Overriding? No
[Tweaks] Wetness: -5.58 dB, Decays: [0.5, 1.0, 2.0] -> [0.5, 1.0, 2.0] s
1st Player Difficulty: Casual

Initial Energy only → Obstruction

+

Reflected Energy only → Wetness ratio

=

Total energy → Occlusion

Effect 4: Decay rate

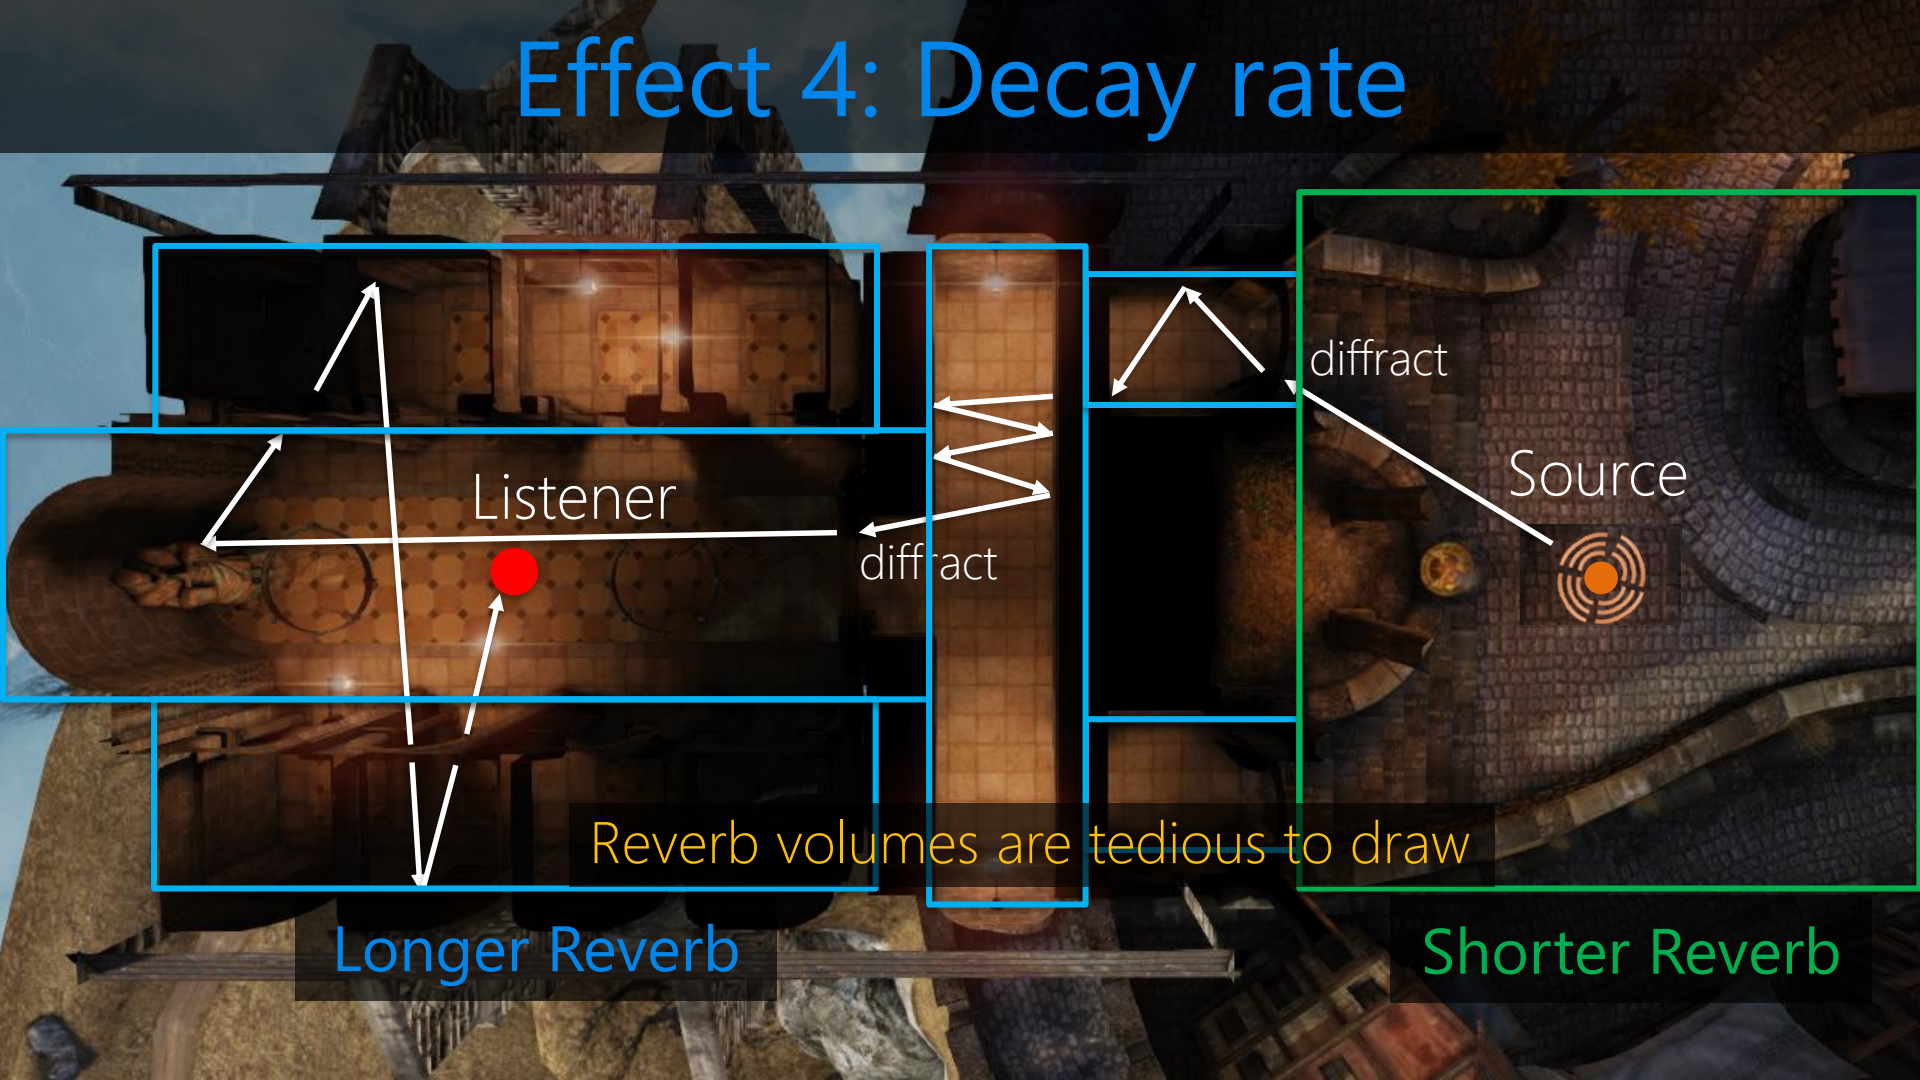

Effect 4: Decay rate

Using Listener location → short reverb!

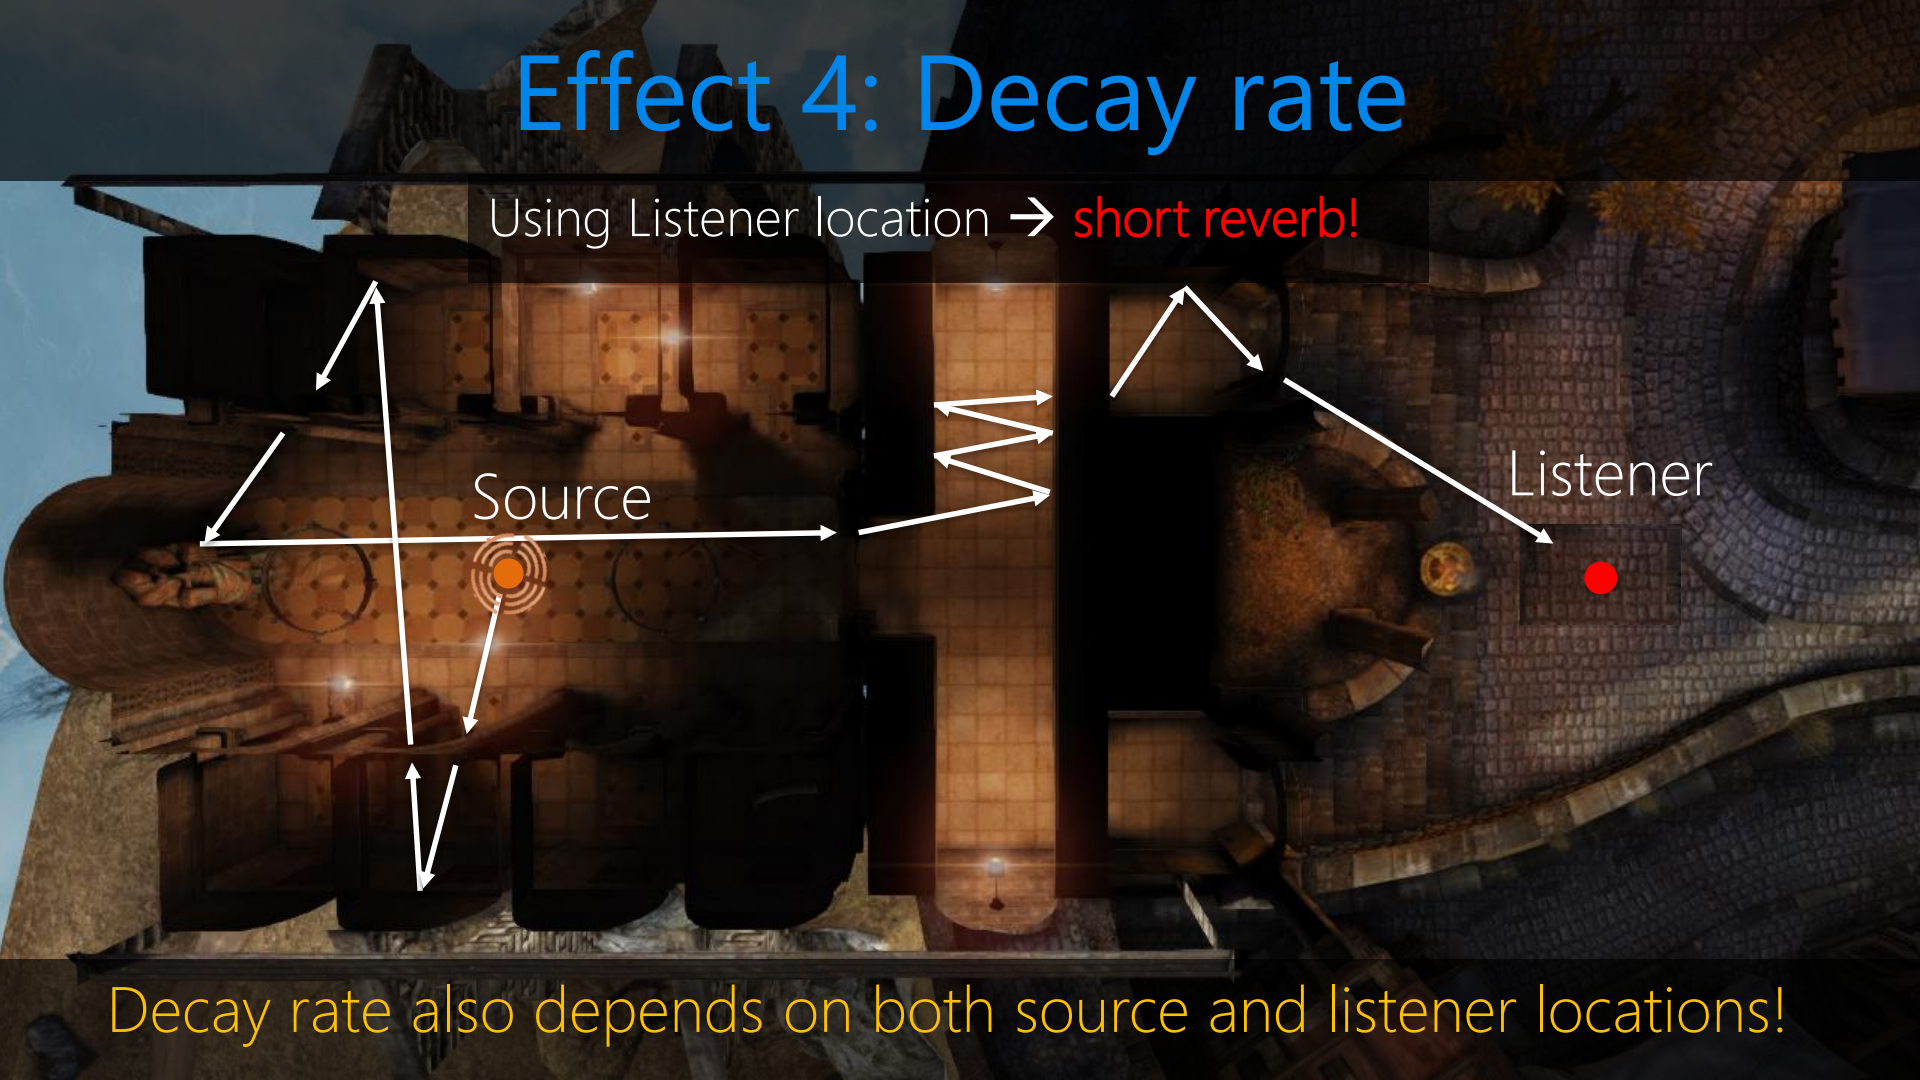

Decay rate also depends on both source and listener locations!


Triton ON
You do NOT have debug mem enabled
1. GearMainMenuEntry_P
2. SP_Farm_B_P
3. GearMainMenuEntry_P
4. SP_Farm_B_P
5. GearMainMenuEntry_P
6. SP_Farm_B_P
7. GearMainMenuEntry_P
8. SP_Farm_B_P
9. GearMainMenuEntry_P
10. SP_Rescue_B_P
11. GearMainMenuEntry_P
12. SP_Rescue_B_P
2 debug jumppoints used
9 checkpoint reloads
SERVER
Warnings: 1217
1st Player Difficulty: Casual
Consistent reverb: source *and* listener


Triton ON
252
You do NOT have debug mem enabled
1. GearMainMenuEntry_P
2. SP_Farm_B_P
3. GearMainMenuEntry_P
4. SP_Farm_B_P
5. GearMainMenuEntry_P
6. SP_Farm_B_P
7. GearMainMenuEntry_P
8. SP_Farm_B_P
9. GearMainMenuEntry_P
10. SP_Rescue_B_P
11. GearMainMenuEntry_P
12. SP_Rescue_B_P
2 debug jumppoints used
16 checkpoint reloads
SERVER
Warnings: 2285
1st Player Difficulty: Casual
Acoustics conveys *location*

Important acoustic effects

1. Obstruction
2. Occlusion
3. Wetness ratio
4. Reverb decay rate

- Depend on *both* (moving) source and listener location
- Depend on scene geometry
- Triton can model them

Triton = baked wave simulation

- Wave simulation: Accurate & reliable results on complex scenes
- Runtime = lookup + interpolation. Light on CPU.
- Need *dynamic* source & listener: large RAM!
- Baking is restricted to static geometry
  - Feasible first step
  - Dynamic scenes (doors/destruction): could layer heuristics on top, like lighting

Illustrated pipeline: Scene

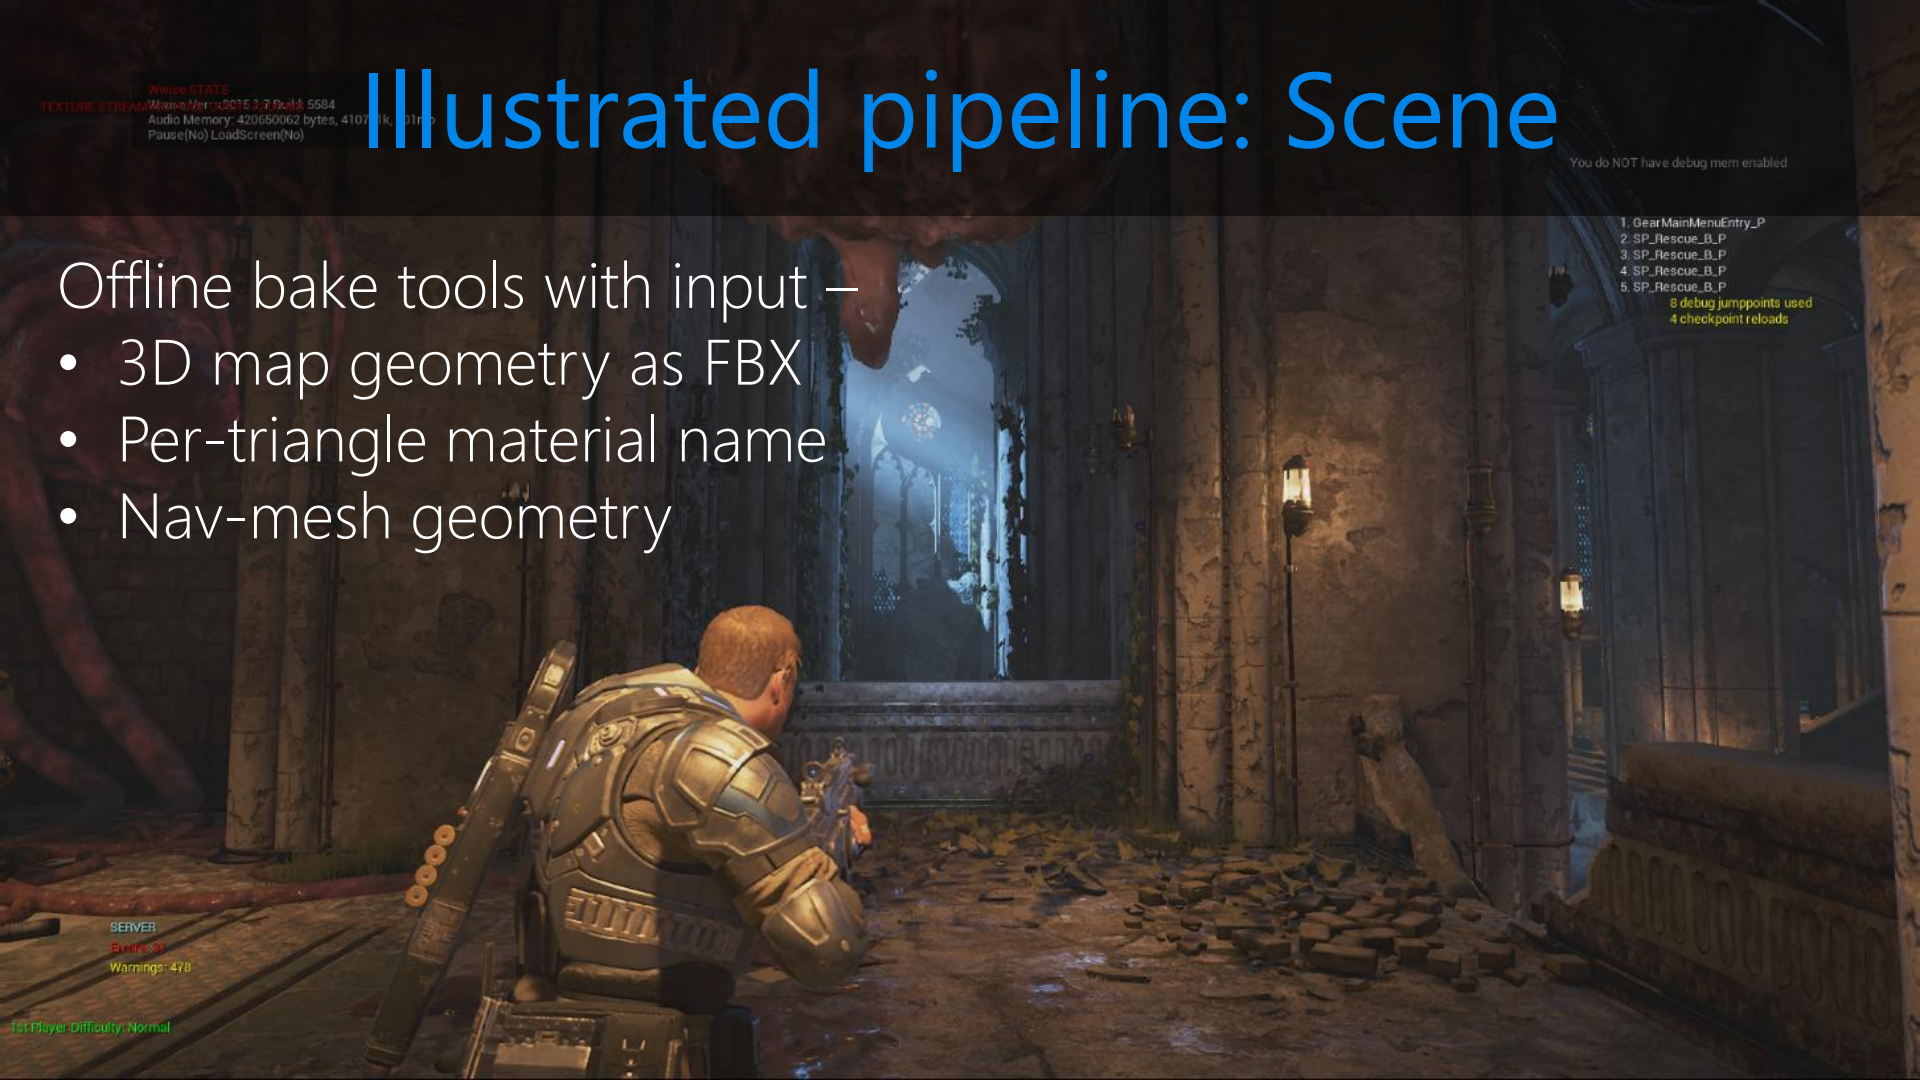

Offline bake tools with input –

- 3D map geometry as FBX
- Per-triangle material name
- Nav-mesh geometry

Start baking: Voxelize!


How to guide player sampling? Boxes?

Nav mesh to the rescue!

Auto-layout **adaptively**-sampled *player probes*

Triton map bake =
For each player probe,
3D wave simulation

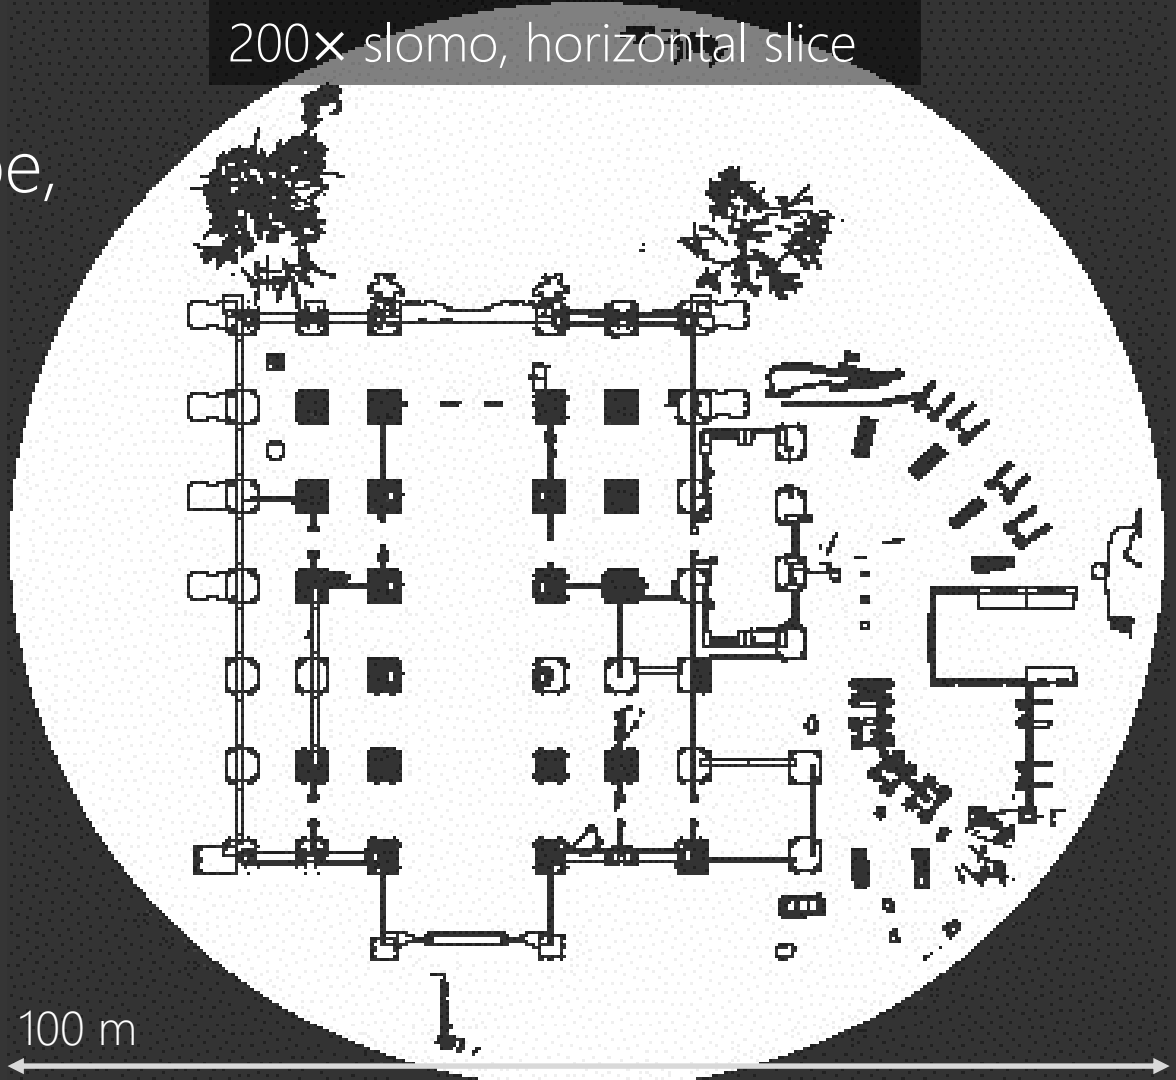

Main Idea: Response → 4 perceptual parameters [for each source & listener location pair]

- Initial ("direct") Energy
- Reflections Energy
- Reflections decay rate
- Reverb decay rate

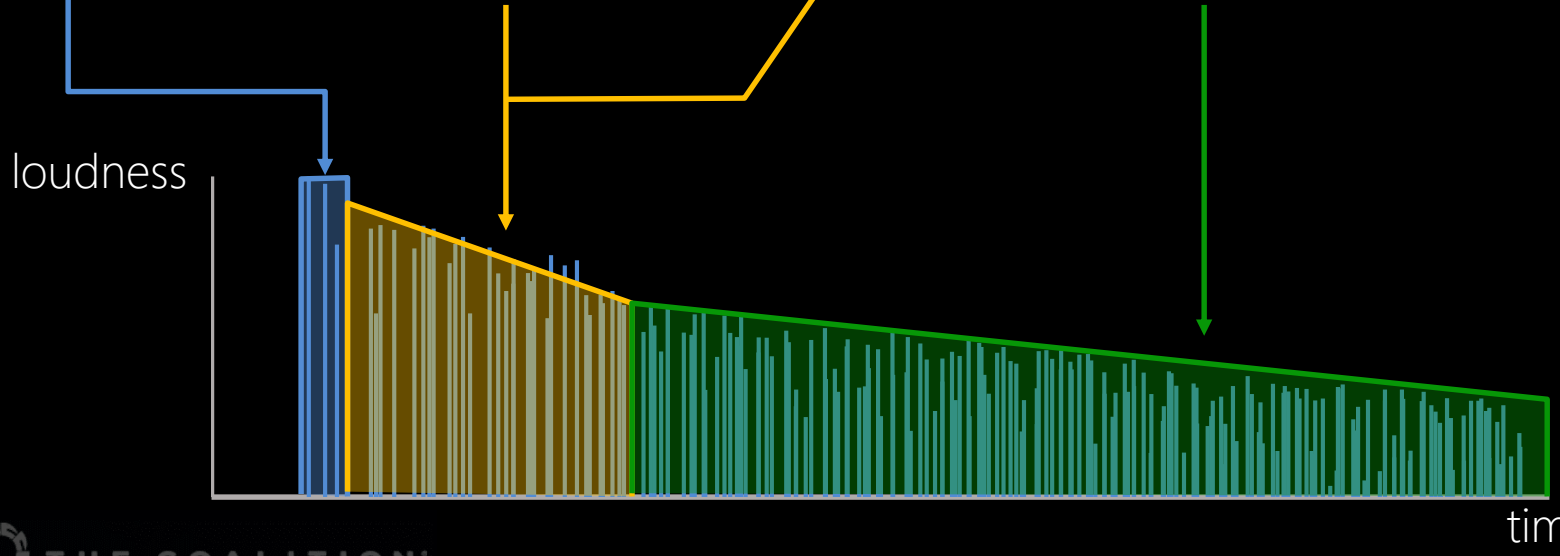

Per-player-probe data: loudness

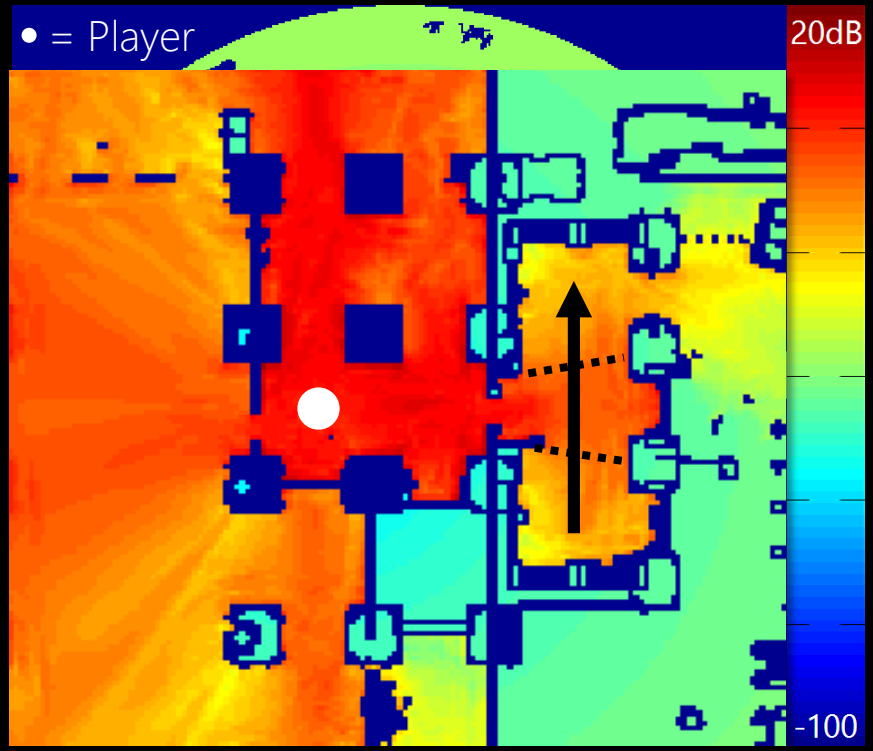

Initial energy (0-10ms)

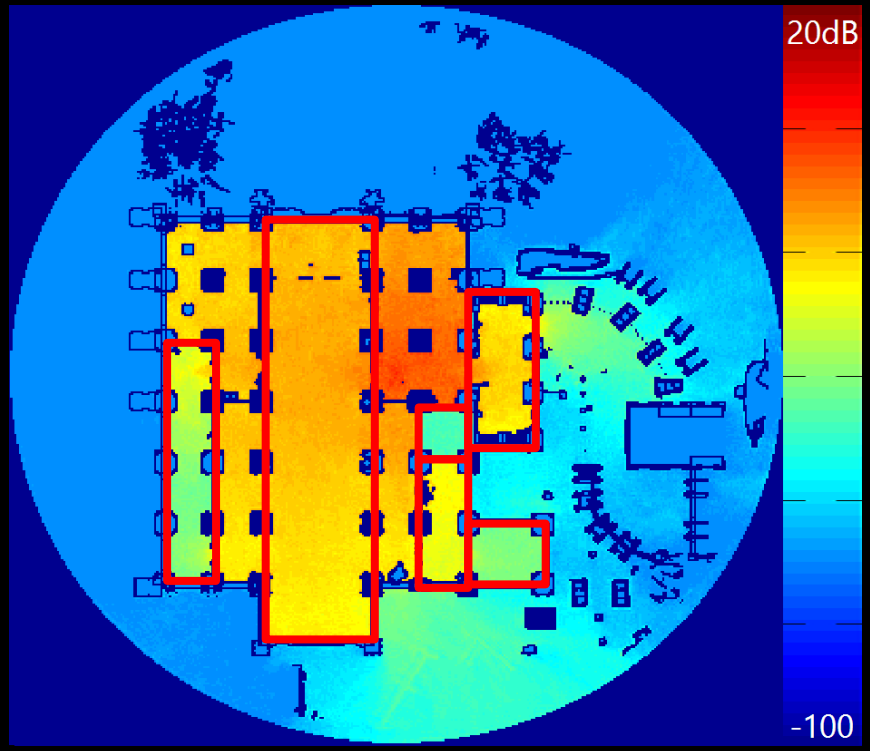

Reflected energy (10-210ms)

Per-player-probe data : early sound decay

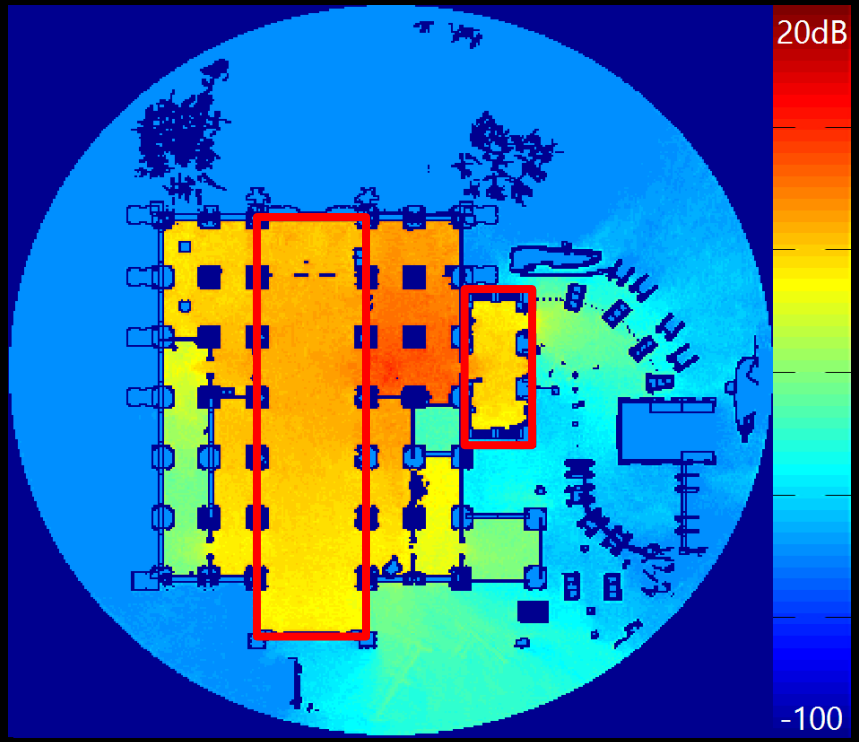

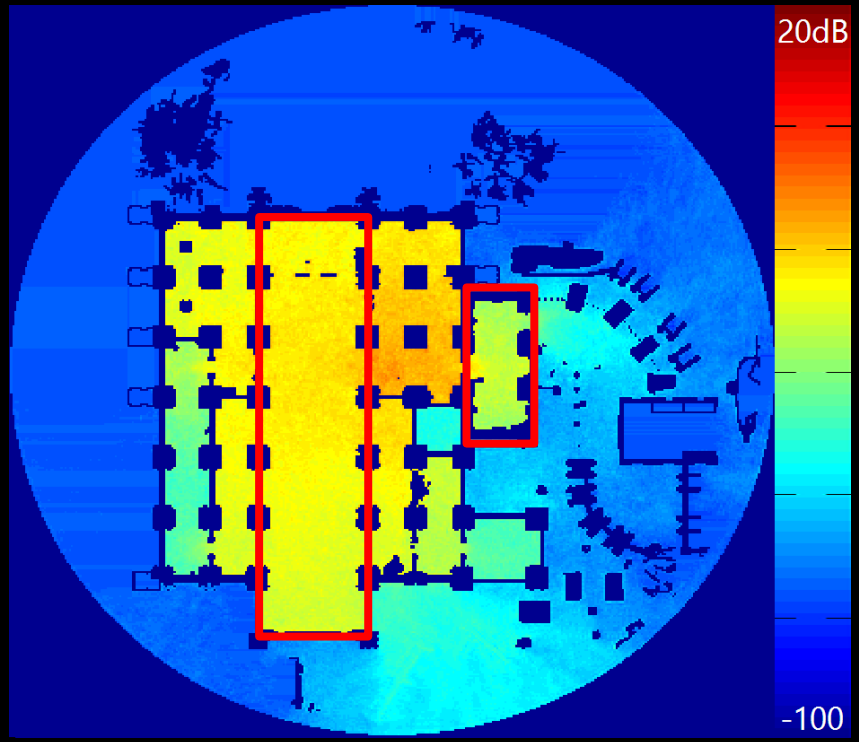

after 250ms

50 TB → 100 MB

Baking for moving sources & listener is costly

- 100 machines → ~4 hours
  - ~10-20 minutes per player probe
  - ~1000-1500 player probes per Gears campaign map
- Trivially parallel: double machines, half time
- Bake tool runs on PC/Xbox (latter as “bake game”)

Runtime: Interpolated lookup

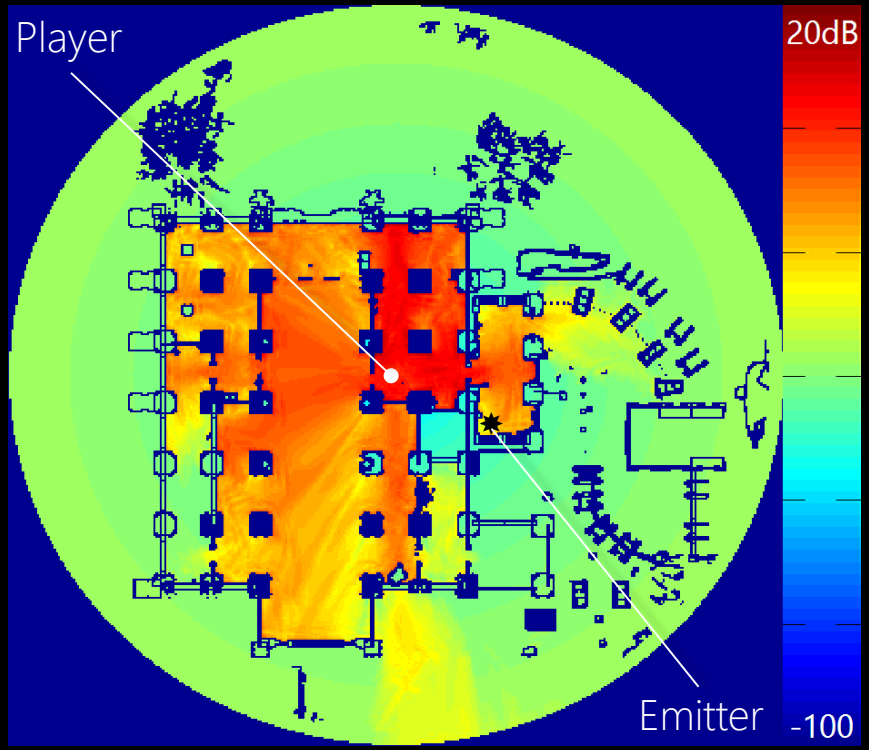

- For each parameter: Decompress slice (Zlib) & do interpolation
- Fixed caches, 20MB overhead
- Worker thread on shared core
- Updated at visual frame rates [30/60FPS in single/multi-player]
- Averages ~100μs per query, spikes of ~200μs on cache miss
- 32 sources served per-frame ~10-20% of shared Xbox core

Game integration

Source pos.

Player pos.

Game

Triton lib

QueryAcoustics()

Baked acoustic data

Wwise reverb sends, occlusion/obstruction

Interpret/ design

Perceptual Data

- initial energy
- reflections energy
- reflections decay rate
- reverb decay rate

Wave baking + perceptual compression = practical acoustics

Technical	Efficient	CPU: ~100μs for acoustic lookup RAM: ~100MB for Campaign, ~20MB for Multiplayer
	Robust	Smooth results on complex level geometry
Design	Automation	No geometry to clean, no volumes to draw
	Expressivity	Interpret Triton data via Unreal blueprint and Wwise

Paper reference: *Raghuvanshi & Snyder, SIGGRAPH 2014,*
"Parametric wave field coding for precomputed sound propagation"

Triton's design

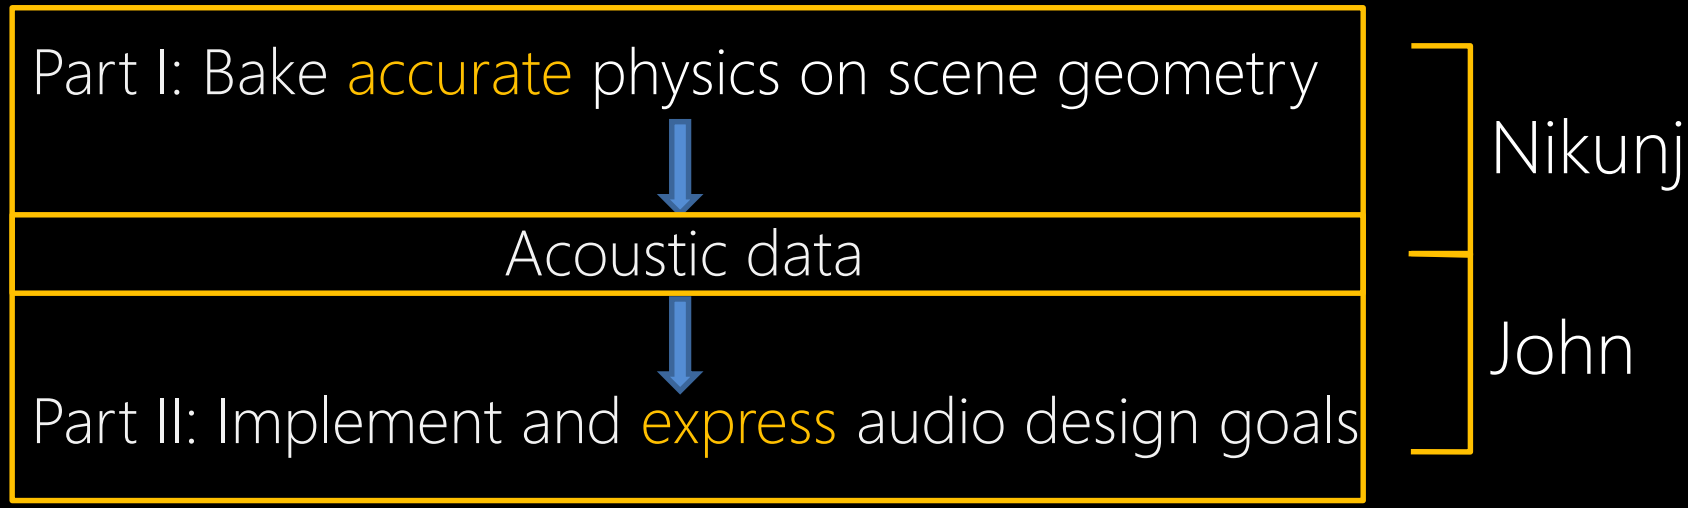

Combine fast wave acoustics with expressive audio design to bring game spaces to life

How does a Sound Designer work with this?

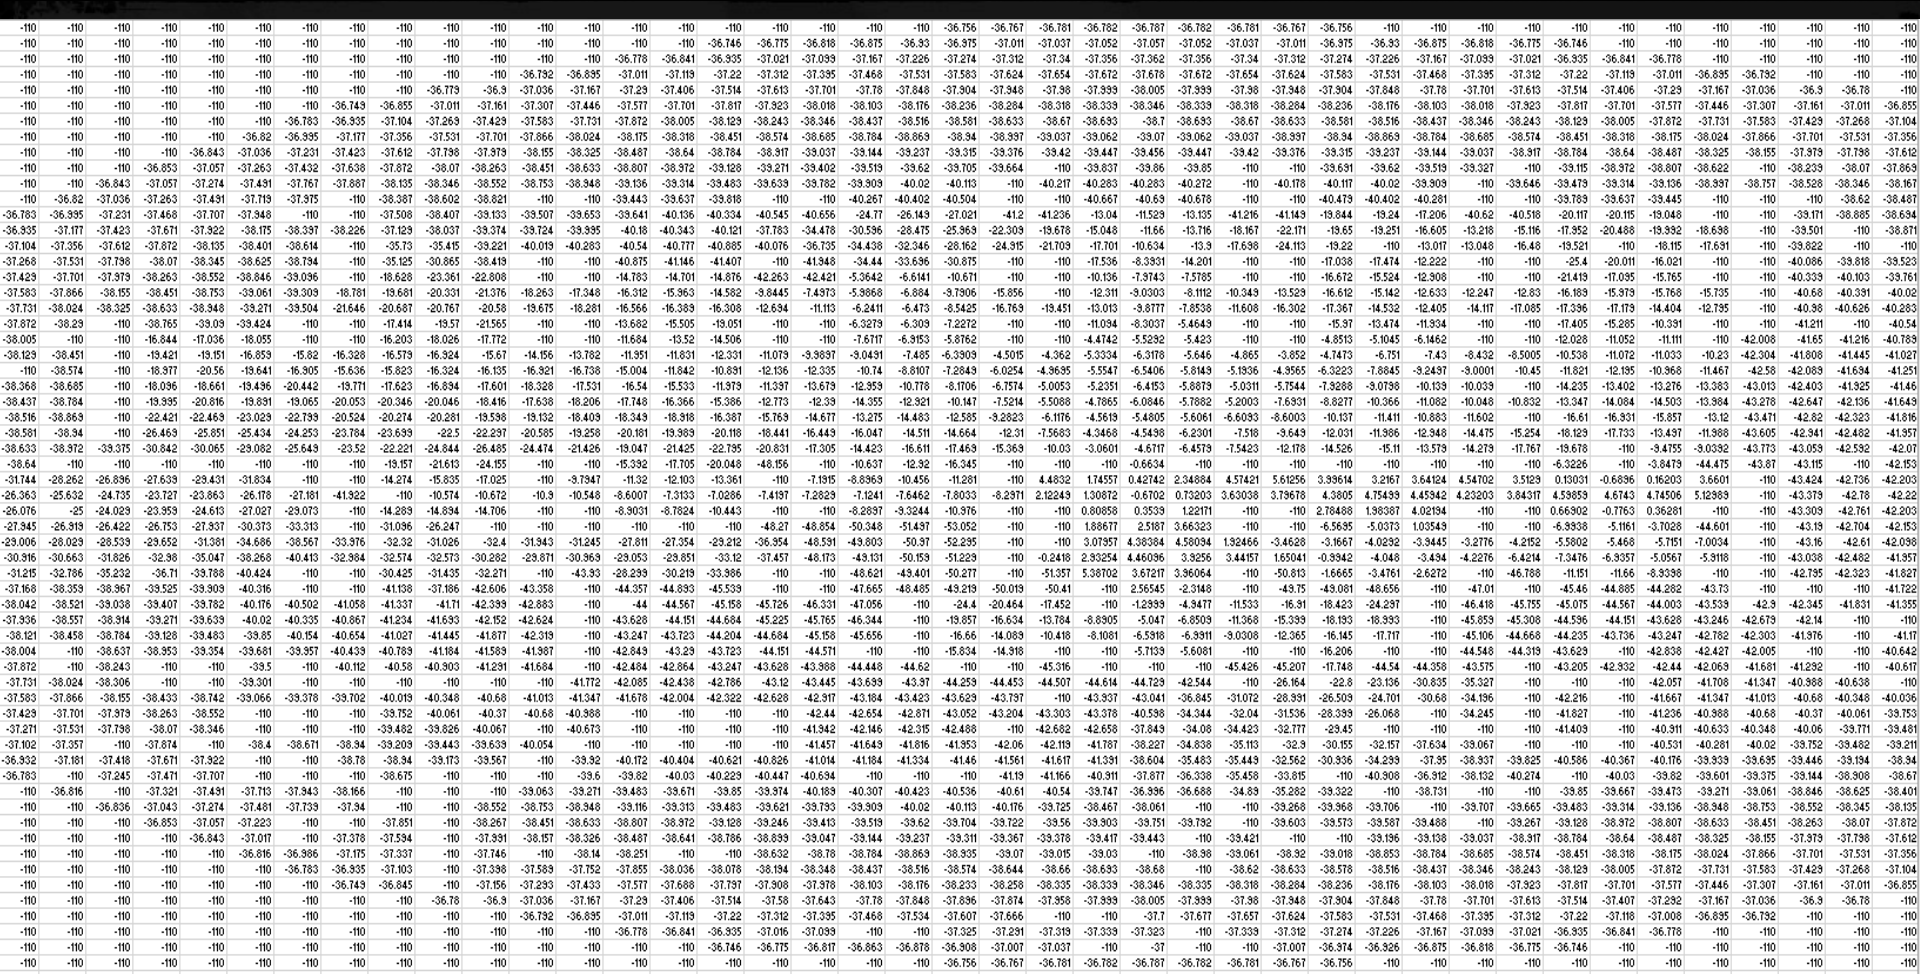

What Triton Is

- Listener and Source Positions → Occlusion/Obstruction, Reverb Wetness, and Reverb Decay Rate values.
- This data is inert and must be *interpreted*.
- The RAM requirements are next-gen.

SP_Rescue_B.ace	11/30/2015 6:51 PM	ACE File	100,468 KB
SP_Dam_A.ace	11/19/2015 5:44 AM	ACE File	117,872 KB

What Triton Isn't

Triton is a Data-Set //not a 'Reverb Plug-in'

Triton Implementation Challenge

The ideal Triton implementation with no CPU/RAM limitations

- Per-Sound Impulse Responses Updated Every Frame
  no middleware support, over RAM budget [10-100GB] ☹
- Per-Sound Convolution Reverb
  ~100 instances; way over CPU budget ☹

Not practical for video game audio…

V1.0

Implementation Progression

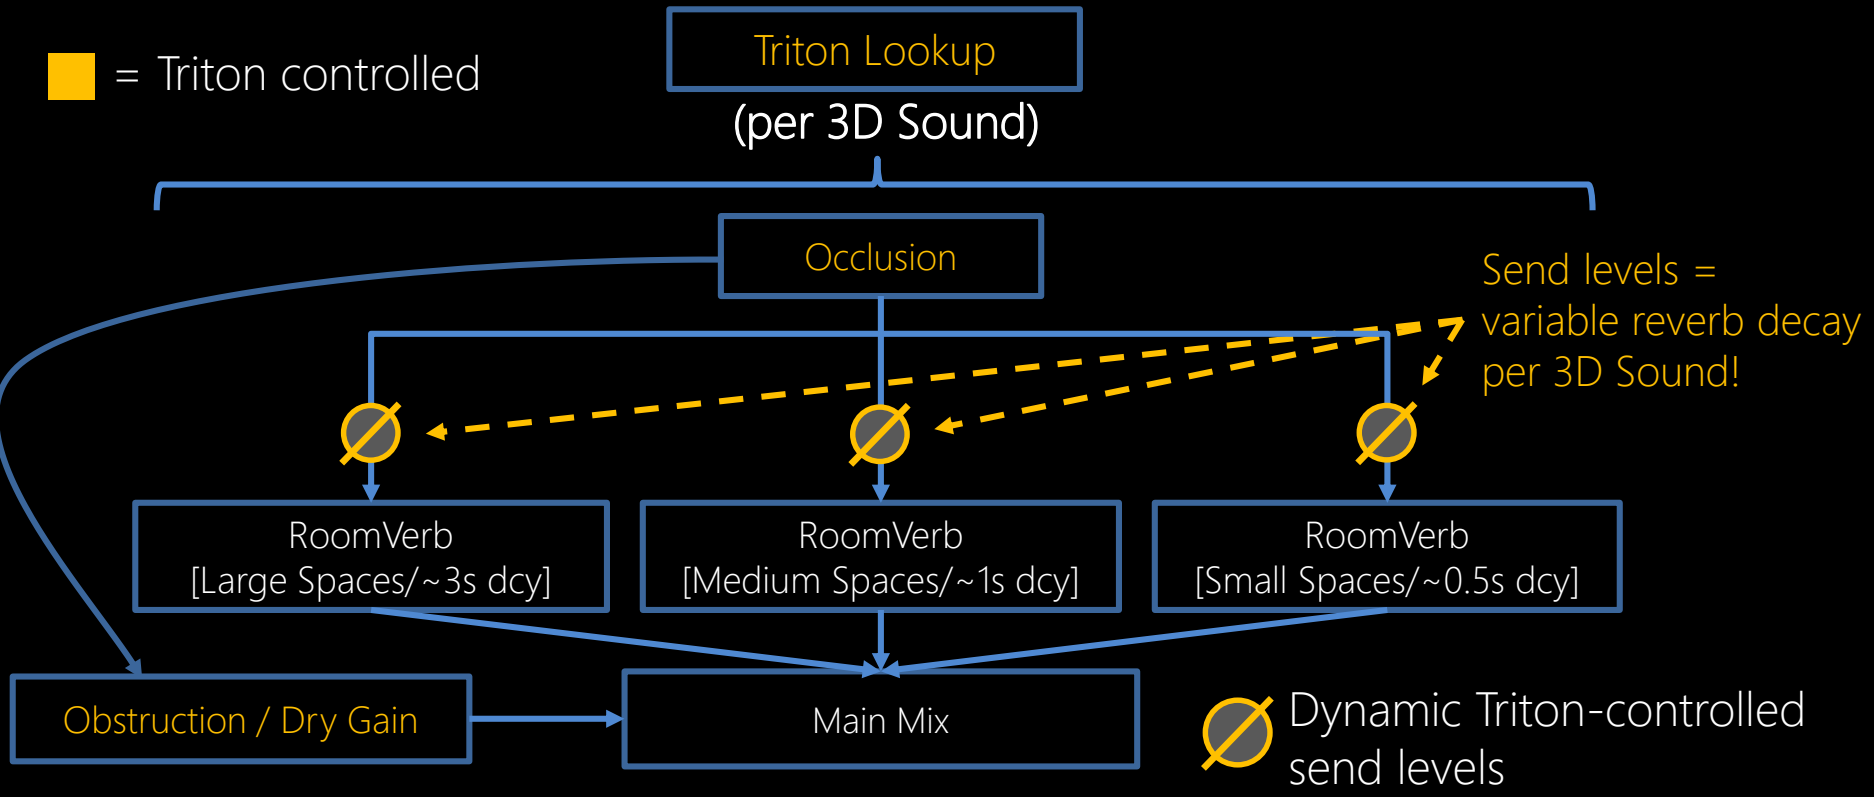

V1.0

Implementation Progression

V1.0 – Summary:

- Triton Reverb Data controls levels sent 3 instances of Wwise RoomVerb.
- Occlusion/Obstruction uses Wwise built in architecture.
  - Strengths: Occlusion/Obstruction is excellent. (Video Example) Dynamic Space Transitions, Responsive, Revelatory moments, Successful Proof-of-Concept.
  - Weaknesses: No solution for outdoor reverb, RoomVerb sounds crowded when stacked, unpredictable behaviour ~25% of the time. Not yet shippable.

Occlusion: firefight

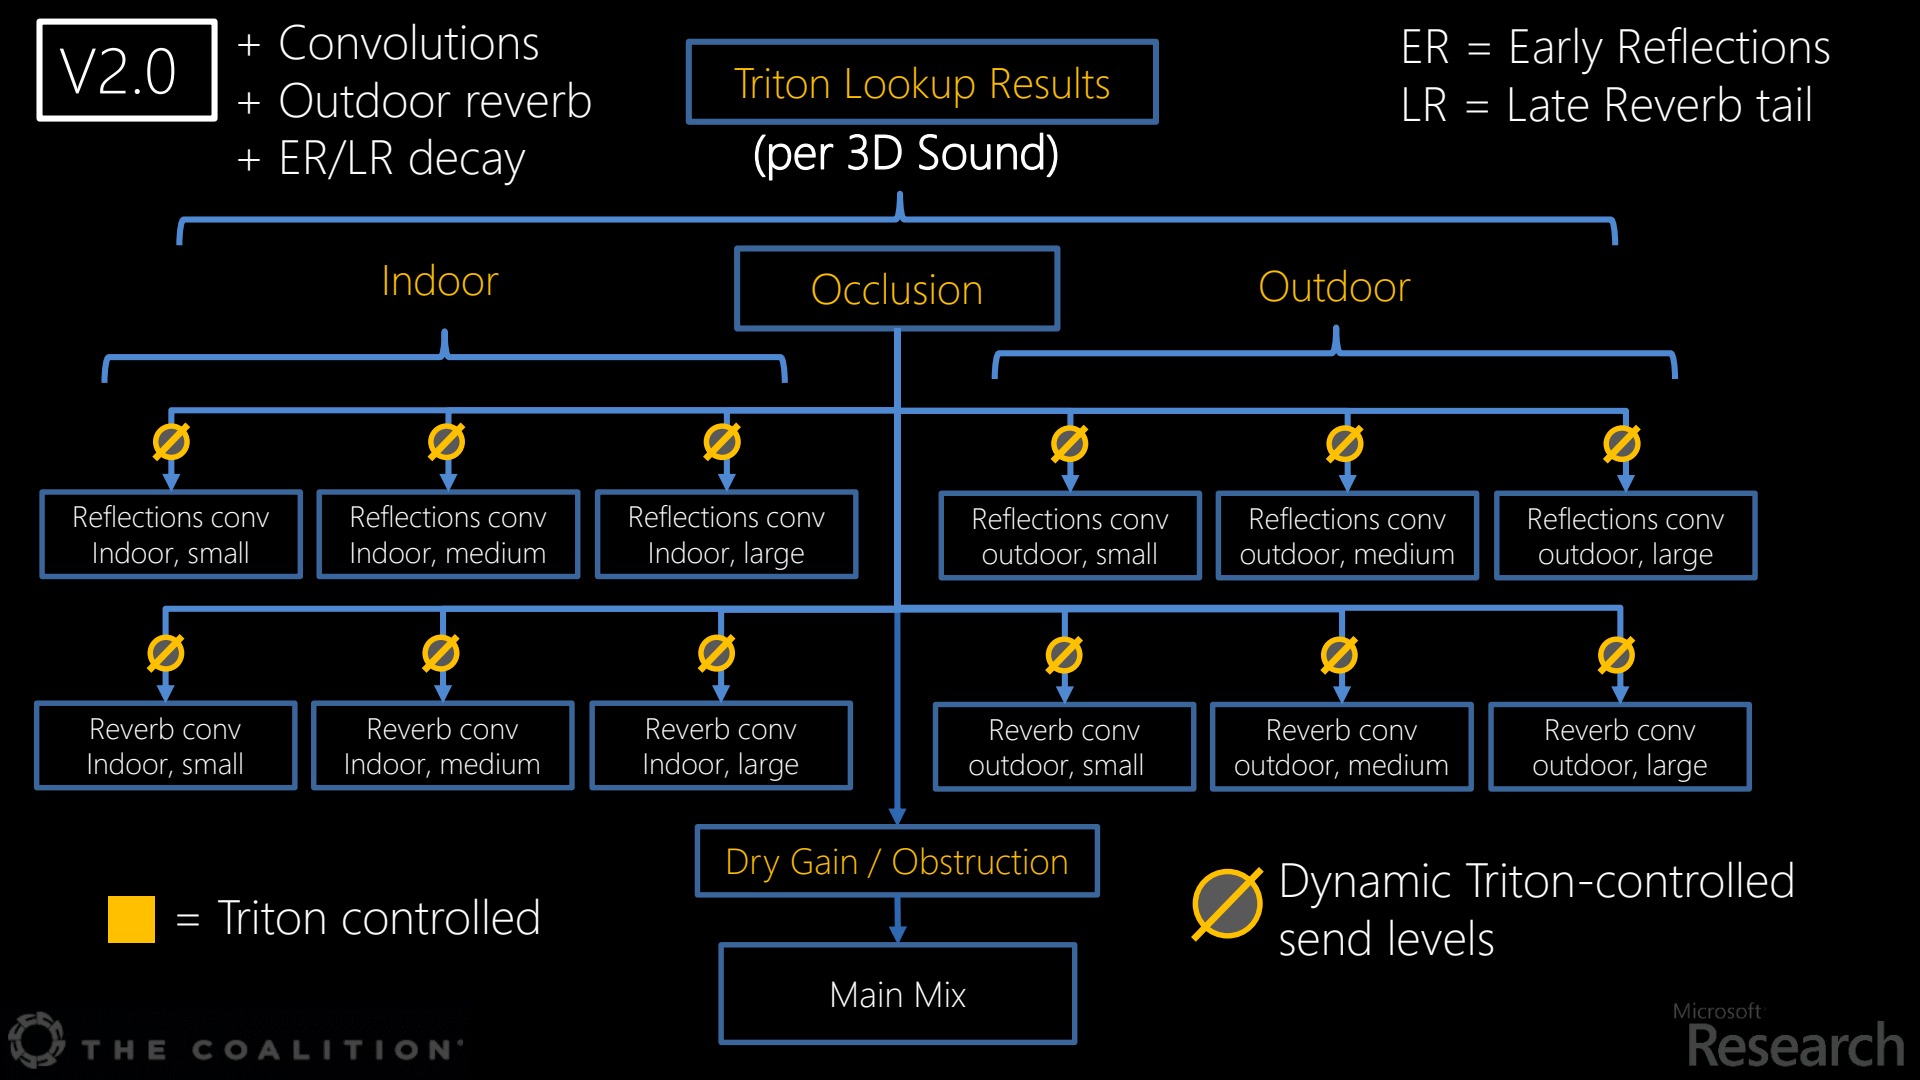
V2.0
+ Convolutions
+ Outdoor reverb
+ ER/LR decay
Triton Lookup Results
(per 3D Sound)
ER = Early Reflections
LR = Late Reverb tail
Indoor
Occlusion
Outdoor
Reflections conv Indoor, small
Reflections conv Indoor, medium
Reflections conv Indoor, large
Reflections conv outdoor, small
Reflections conv outdoor, medium
Reflections conv outdoor, large
Reverb conv Indoor, small
Reverb conv Indoor, medium
Reverb conv Indoor, large
Reverb conv outdoor, small
Reverb conv outdoor, medium
Reverb conv outdoor, large
Dry Gain / Obstruction
Main Mix
= Triton controlled
Dynamic Triton-controlled send levels
THE COALITION
Microsoft Research

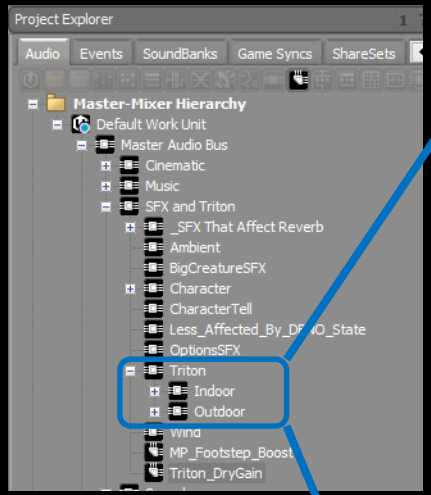
Project Explorer
Audio
Events
SoundBanks
Game Syncs
ShareSets
Master-Mixer Hierarchy
Default Work Unit
Master Audio Bus
Cinematic
Music
SFX and Triton
_SFX That Affect Reverb
Ambient
BigCreatureSFX
Character
CharacterTell
Less_Affected_By_DF_IO_State
OptionsSFX
Triton
Indoor
Outdoor
Wind
MP_Footstep_Boost
Triton_DryGain

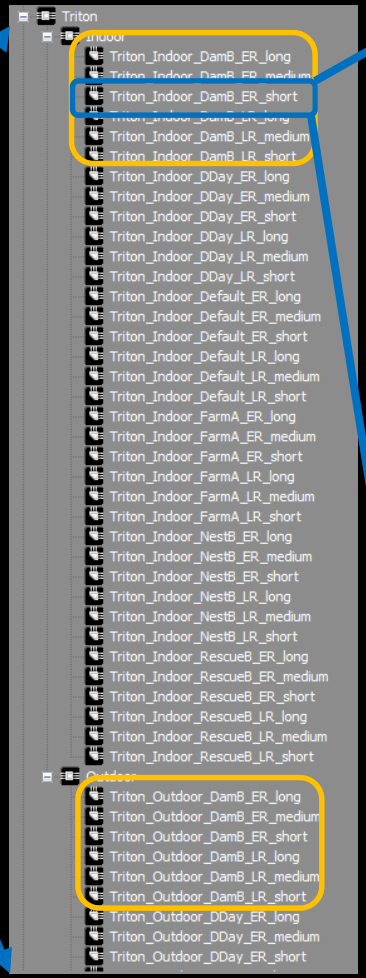
Triton
Indoor
Triton_Indoor_DamB_ER_long
Triton_Indoor_DamB_ER_short
Triton_Indoor_DamB_LR_medium
Triton_Indoor_DamB_LR_short
Triton_Indoor_DDay_ER_long
Triton_Indoor_DDay_ER_medium
Triton_Indoor_DDay_ER_short
Triton_Indoor_DDay_LR_long
Triton_Indoor_DDay_LR_medium
Triton_Indoor_DDay_LR_short
Triton_Indoor_Default_ER_long
Triton_Indoor_Default_ER_medium
Triton_Indoor_Default_ER_short
Triton_Indoor_Default_LR_long
Triton_Indoor_Default_LR_medium
Triton_Indoor_Default_LR_short
Triton_Indoor_FarmA_ER_long
Triton_Indoor_FarmA_ER_medium
Triton_Indoor_FarmA_ER_short
Triton_Indoor_FarmA_LR_long
Triton_Indoor_FarmA_LR_medium
Triton_Indoor_FarmA_LR_short
Triton_Indoor_NestB_ER_long
Triton_Indoor_NestB_ER_medium
Triton_Indoor_NestB_ER_short
Triton_Indoor_NestB_LR_long
Triton_Indoor_NestB_LR_medium
Triton_Indoor_NestB_LR_short
Triton_Indoor_RescueB_ER_long
Triton_Indoor_RescueB_ER_medium
Triton_Indoor_RescueB_ER_short
Triton_Indoor_RescueB_LR_long
Triton_Indoor_RescueB_LR_medium
Triton_Indoor_RescueB_LR_short
Outdoor
Triton_Outdoor_DamB_ER_long
Triton_Outdoor_DamB_ER_medium
Triton_Outdoor_DamB_ER_short
Triton_Outdoor_DamB_LR_long
Triton_Outdoor_DamB_LR_medium
Triton_Outdoor_DamB_LR_short
Triton_Outdoor_DDay_ER_long
Triton_Outdoor_DDay_ER_medium
Triton_Outdoor_DDay_ER_short

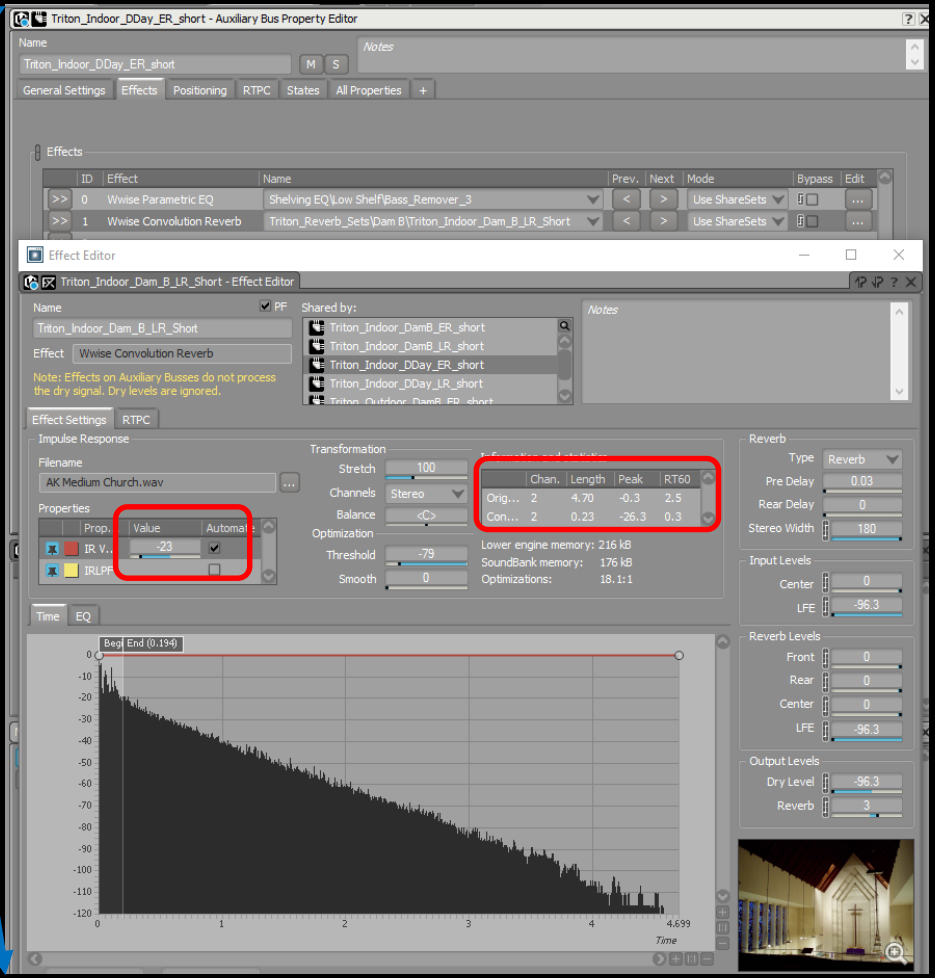
Triton_Indoor_DDay_ER_short - Auxiliary Bus Property Editor
Name
Triton_Indoor_DDay_ER_short
Notes
General Settings
Effects
Positioning
RTPC
States
All Properties
Effects
ID
Effect
Name
Prev.
Next
Mode
Bypass
Edit
0
Wwise Parametric EQ
Shelving EQ\Low Shelf\Bass_Remover_3
Use ShareSets
1
Wwise Convolution Reverb
Triton_Reverb_Sets\Dam B\Triton_Indoor_Dam_B_LR_Short
Use ShareSets
Effect Editor
Triton_Indoor_Dam_B_LR_Short - Effect Editor
Name
Triton_Indoor_Dam_B_LR_Short
PF
Shared by:
Triton_Indoor_DamB_ER_short
Triton_Indoor_DamB_LR_short
Triton_Indoor_DDay_ER_short
Triton_Indoor_DDay_LR_short
Notes
Effect
Wwise Convolution Reverb
Note: Effects on Auxiliary Busses do not process the dry signal. Dry levels are ignored.
Effect Settings
RTPC
Impulse Response
Filename
AK Medium Church.wav
Properties
Prop.
Value
Automat
IR V...
-23
IRLPF
Transformation
Stretch
100
Channels
Stereo
Balance
<C>
Optimization
Threshold
-79
Smooth
0
Chan.
Length
Peak
RT60
Orig...
2
4.70
-0.3
2.5
Con...
2
0.23
-26.3
0.3
Lower engine memory: 216 kB
SoundBank memory: 176 kB
Optimizations: 18.1:1
Reverb
Type
Reverb
Pre Delay
0.03
Rear Delay
0
Stereo Width
180
Input Levels
Center
0
LFE
-96.3
Reverb Levels
Front
0
Rear
0
Center
0
LFE
-96.3
Output Levels
Dry Level
-96.3
Reverb
3
Time
EQ
Beg
End (0.194)
0
-10
-20
-30
-40
-50
-60
-70
-80
-90
-100
-110
-120
0
1
2
3
4
4.699
Time

Map =
Dam B

V2.0

Implementation Progression

V2.0 – Summary:

- Triton Data controls levels to 12 Wwise Convolution Reverbs.
  - Strengths: Indoor/Outdoor threshold blending is excellent (video example), dynamic reverb decay response
  - Weaknesses: Setup and calibration, bugs were difficult to stomp

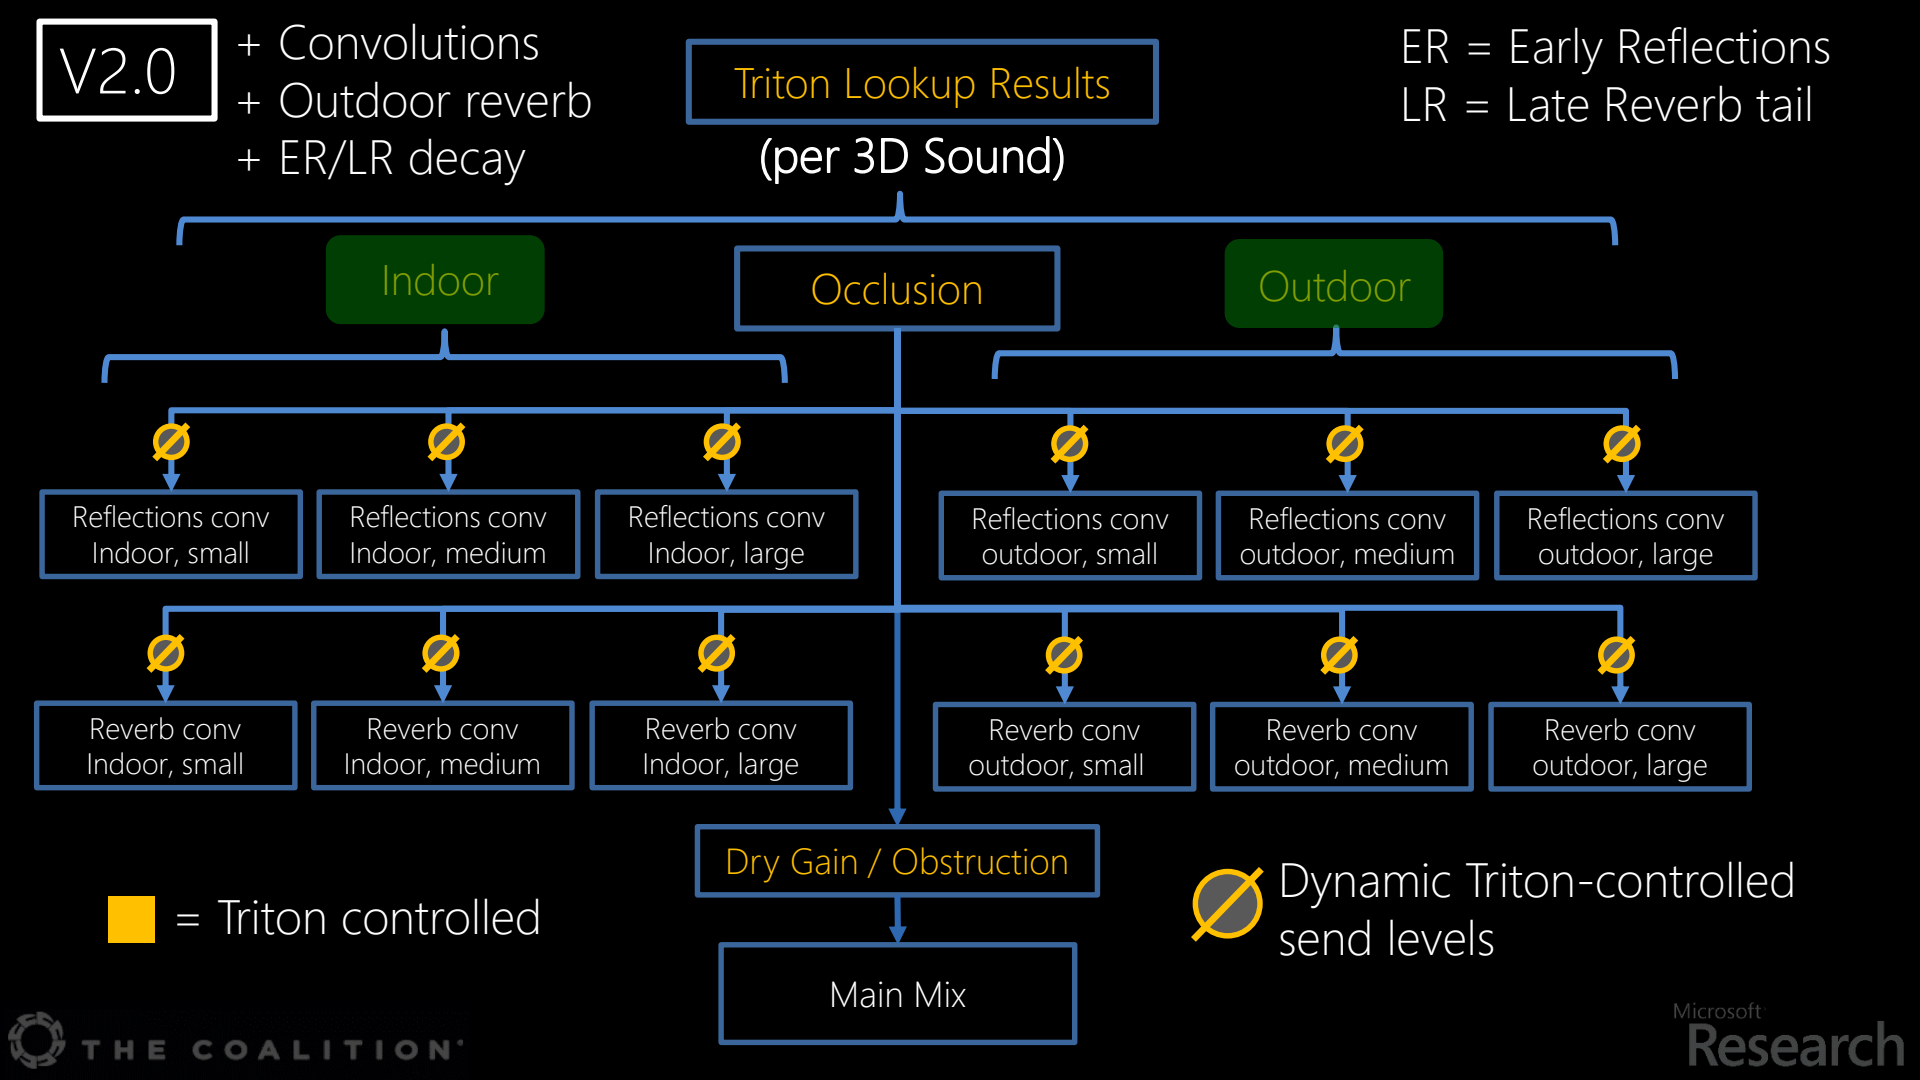
V2.0
+ Convolutions
+ Outdoor reverb
+ ER/LR decay
Triton Lookup Results
(per 3D Sound)
ER = Early Reflections
LR = Late Reverb tail
Indoor
Occlusion
Outdoor
Reflections conv Indoor, small
Reflections conv Indoor, medium
Reflections conv Indoor, large
Reflections conv outdoor, small
Reflections conv outdoor, medium
Reflections conv outdoor, large
Reverb conv Indoor, small
Reverb conv Indoor, medium
Reverb conv Indoor, large
Reverb conv outdoor, small
Reverb conv outdoor, medium
Reverb conv outdoor, large
Dry Gain / Obstruction
Main Mix
= Triton controlled
Dynamic Triton-controlled send levels
THE COALITION
Microsoft Research

Mixing Indoor/Outdoor Reverbs Through Thresholds

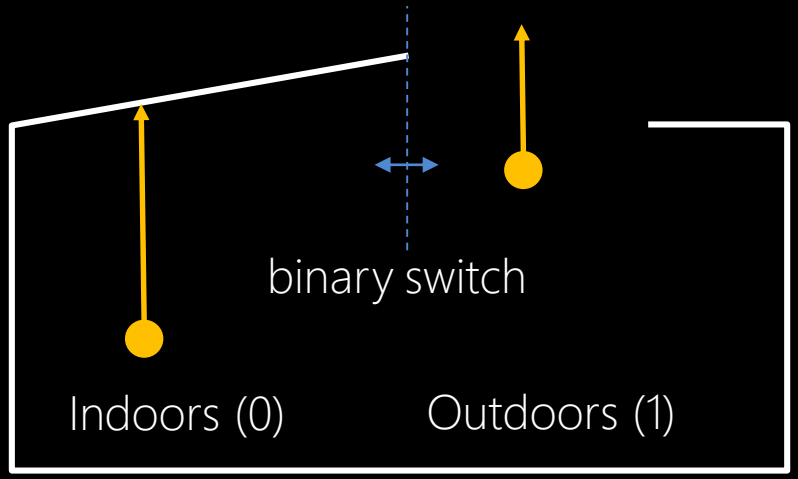

Problem: **This concept is not binary.**

Solution: Triton pre-computes a smooth “Outdoorness”

Solution: $\text{Outdoorness} = \frac{\text{Energy reaching sky}}{\text{Total Energy shot}}$

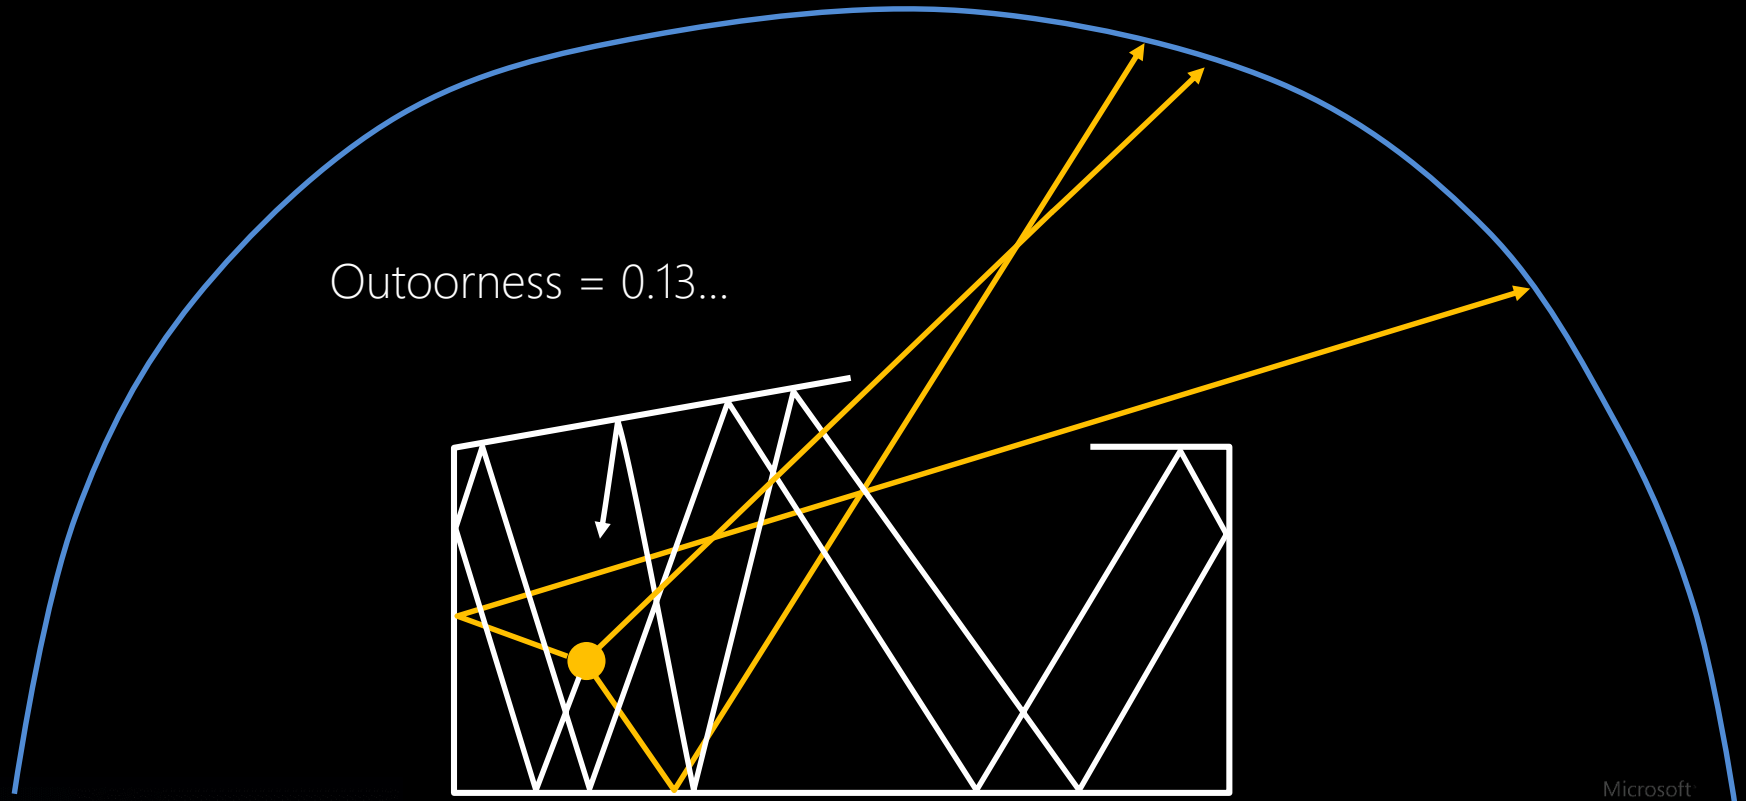

Outdoorness changes **smoothly** as player approaches transition

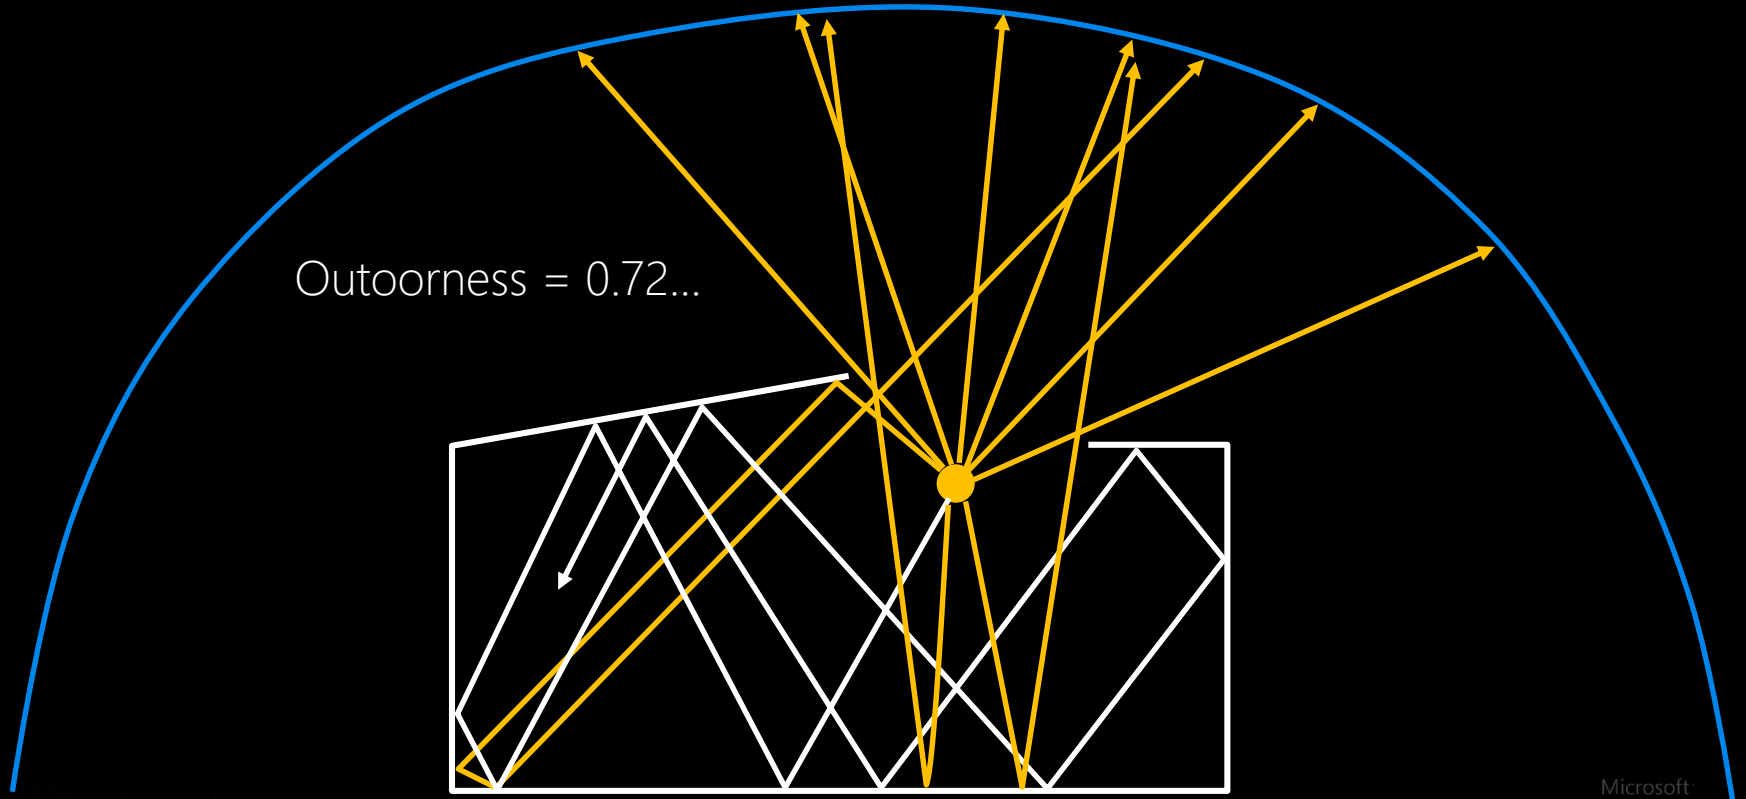

Unforeseen Uses: Threshold Blending for Rain and Wind

- Triton uses this for determining sends for inside/outside reverb blending
- Exposed as RTPC: Driving sound of Wind, Rain, etc.

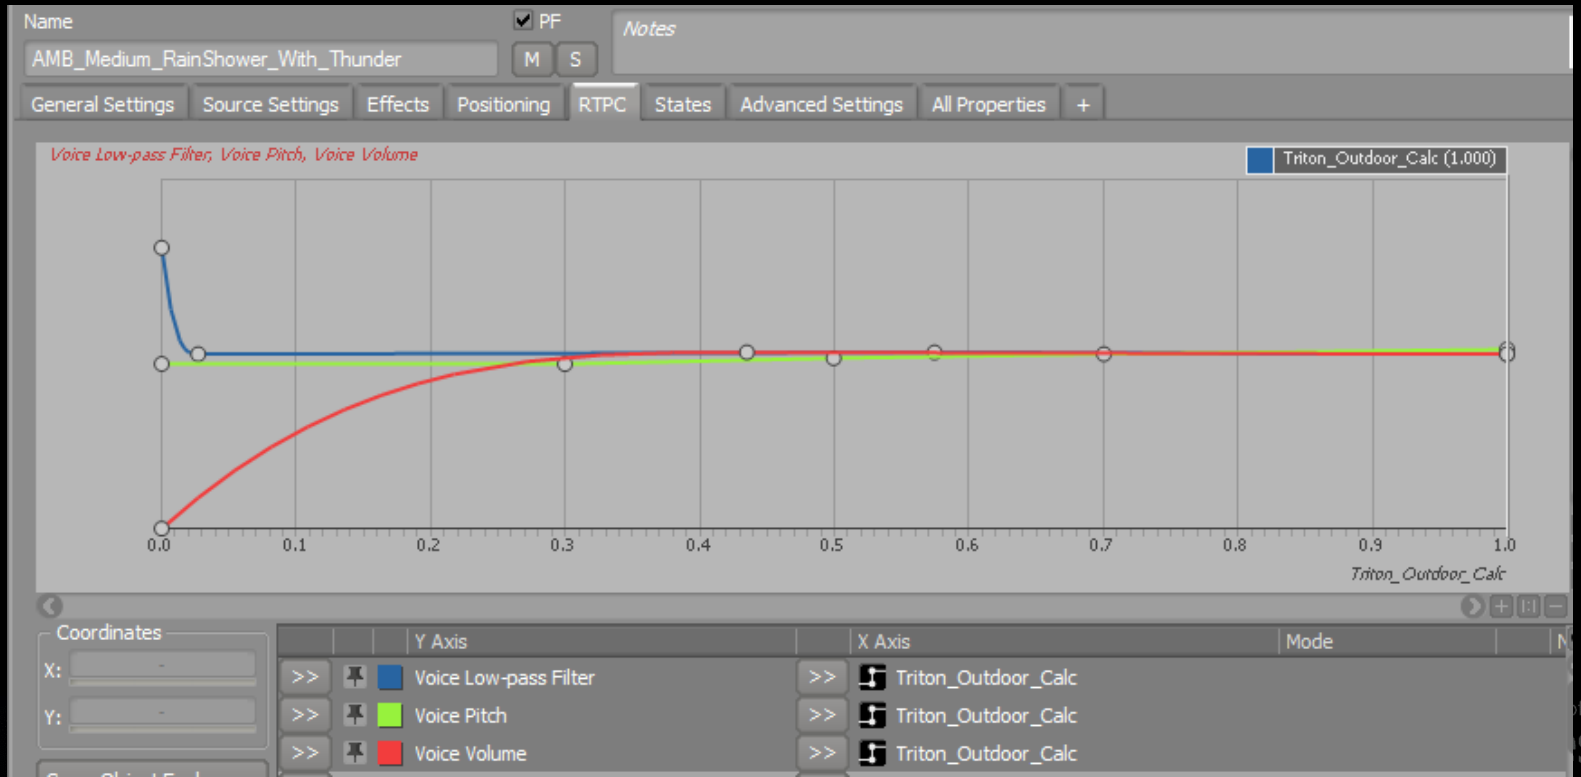

Outdoorness: Rain

V2.0 Generated some uncomfortable questions:

Triton Reverb worked well *sometimes*—why?

Time investment for shippable quality?

Storytime…

One day, we were given a deadline by
Rod Fergusson (Gears Co-Creator/Studio Head) and
Chuck Osiega (Creative Director):

"Get Triton Reverb sounding shippable on one map in 2 weeks or we switch to a proven conventional reverb solution."

Change of Strategy

1. Reduce number of sounds going into Triton Reverb.
2. Reduce the number of reverb outputs.

Simplification revealed **root causes** of the problems.

Triton-calculated Occlusion/Obstruction remained enabled for almost all 3D sounds.

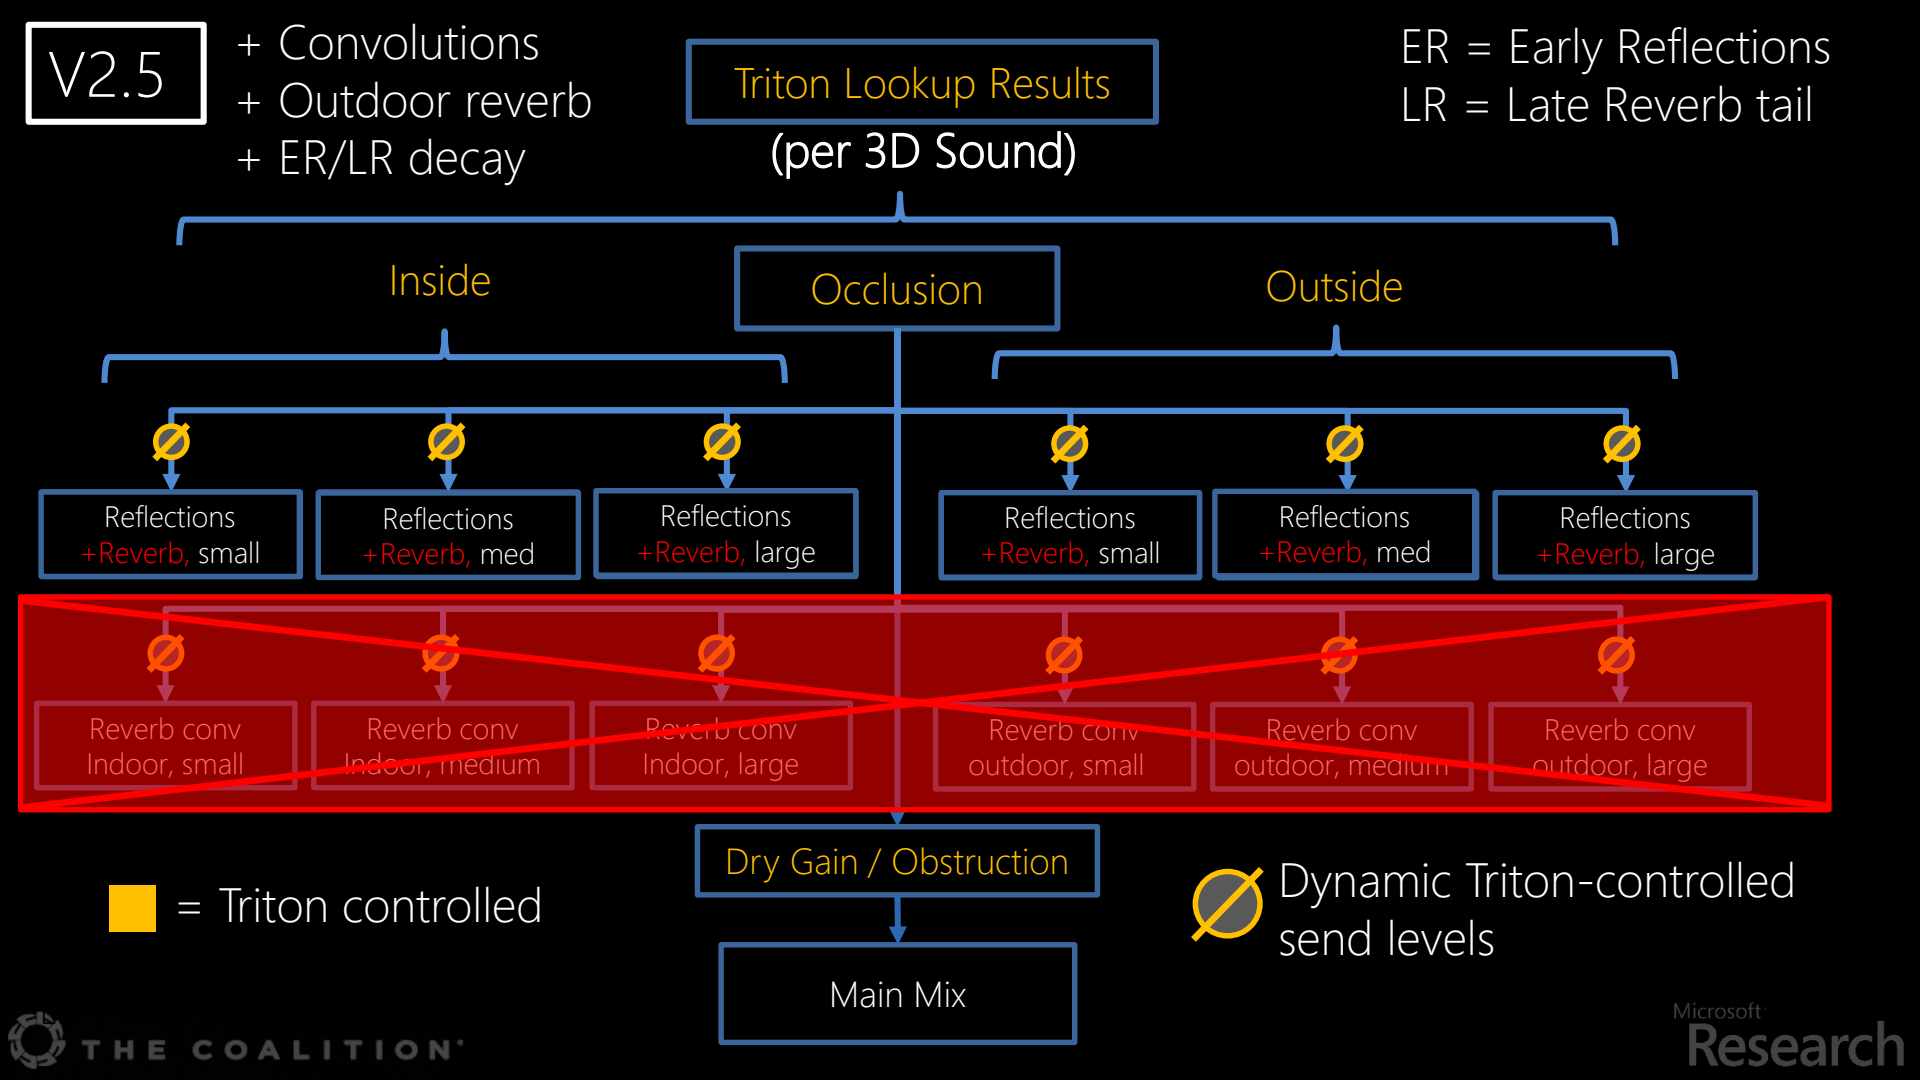
V2.5
+ Convolutions
+ Outdoor reverb
+ ER/LR decay
Triton Lookup Results
(per 3D Sound)
ER = Early Reflections
LR = Late Reverb tail
Inside
Occlusion
Outside
Reflections +Reverb, small
Reflections +Reverb, med
Reflections +Reverb, large
Reflections +Reverb, small
Reflections +Reverb, med
Reflections +Reverb, large
Reverb conv Indoor, small
Reverb conv Indoor, medium
Reverb conv Indoor, large
Reverb conv outdoor, small
Reverb conv outdoor, medium
Reverb conv outdoor, large
Dry Gain / Obstruction
Main Mix
= Triton controlled
Dynamic Triton-controlled send levels
THE COALITION
Microsoft Research

Root Causes of our Triton Challenges

1. Physical dynamic range vs. Aesthetic dynamic range
2. Dry Gain added to point source sounds
3. Physics-based vs. Aesthetics-based reverb wetness/decay times: Emotionally Interpreted Spaces

A Shift in Thinking:
How "Real" is "Too Real" for Game Acoustics?

A Shift in Thinking: How “Real” is “Too Real” for Game Acoustics?

Triton datasets are formed with reality based calculations. BUT:
Not all physical aural phenomena are creatively desirable all the time!

Lesson 1: Dynamic Range

Triton’s ‘real world’ dynamic range needs interpretation into a smaller gameplay range.

Sound designers in film and games have ALWAYS made this interpretation to create a mix.
Realizing that we needed to apply this same interpretive process to
Triton-Reverb inputs was crucial to shipping.

Physical dB Range vs. Game dB range

Decibel/SPL Range Chart		
	Real World Sounds	Gears 4 Sounds in "real world dB range"
0dB	Threshold of human hearing	
10dB	Breathing	
20dB	Whispering	
30dB	Room tone/Light ambient tones	Quiet Level Ambience
40dB		Typical World Sound (Fire, Trees in Light Wind… etc.)
50dB	Conversation	→ Squad Speech ←
60dB	Busy street	
70dB	Vacuum, noisey restaurant	
80dB	Kitchen Blender	Loud World Sound (Large Machinery… etc.)
90dB	Lawn Mower	
100dB	Helicopter 30m above	Kestrel
110dB	Loud Rock Concert	
120dB	Thunderclap (close range)	
130dB	Gunshot	Lancer
140dB		
150dB		
160dB		
170dB		
180dB		
190dB		Grenade (close range)
200dB	Eardrum Rupture	

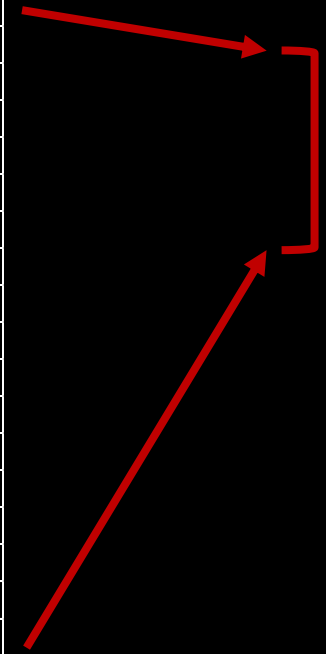

~35dB of usable Dynamic Range

Final adjustment of occlusion/obstruction curves to account for game-relative dynamic range

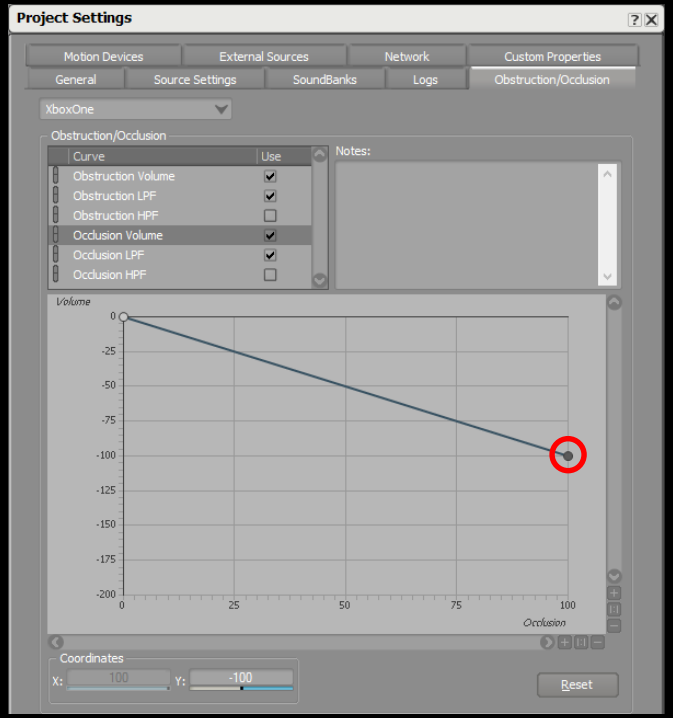

Physics based dynamic range: 100dB

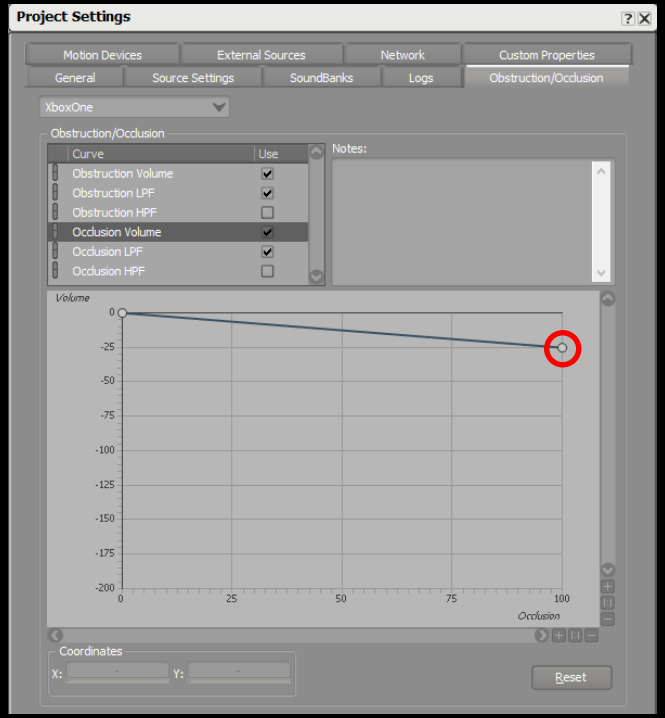

Game based dynamic range: 25dB

A Shift in Thinking: How “Real” is “Too Real” for Game Acoustics?

Lesson 2: Dry Gain

Small enclosed spaces can amplify dry sound in reality.

When simulated in gameplay, this can sound unexpected.

Solution: The ceiling of this value was originally +12db - we clamped this gain to +3dB

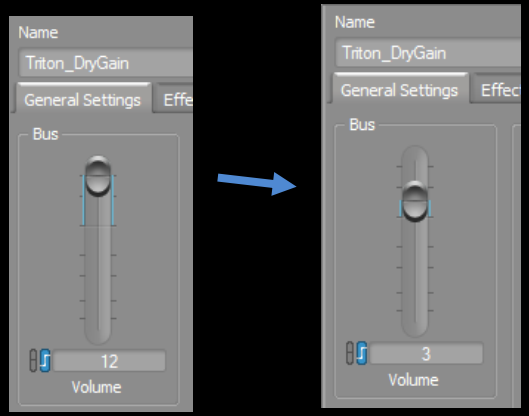

A Shift in Thinking: How “Real” is “Too Real” for Game Acoustics?

Lesson 3:

Triton simulations can violate expectations:

- small spaces can generate **very** long decay times.
- large spaces can generate short decay times.

In games, we have emotional requirements for reverb to inform storytelling.

Solution: Scriptable designer-based interpretations of the Triton data. AKA The “Triton Tweak Acoustics node”.

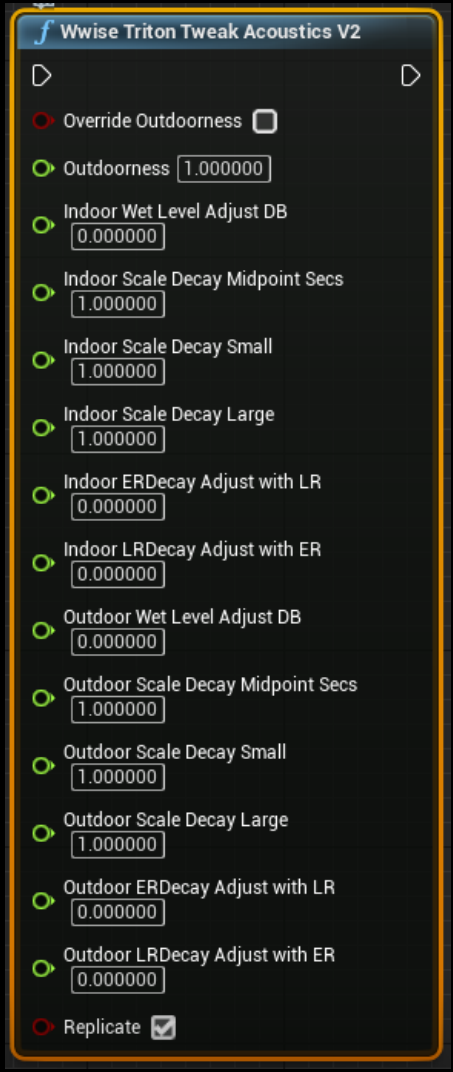

Designer crafted spaces: reverb/rain transition

Designer crafted spaces: reverb/rain transition

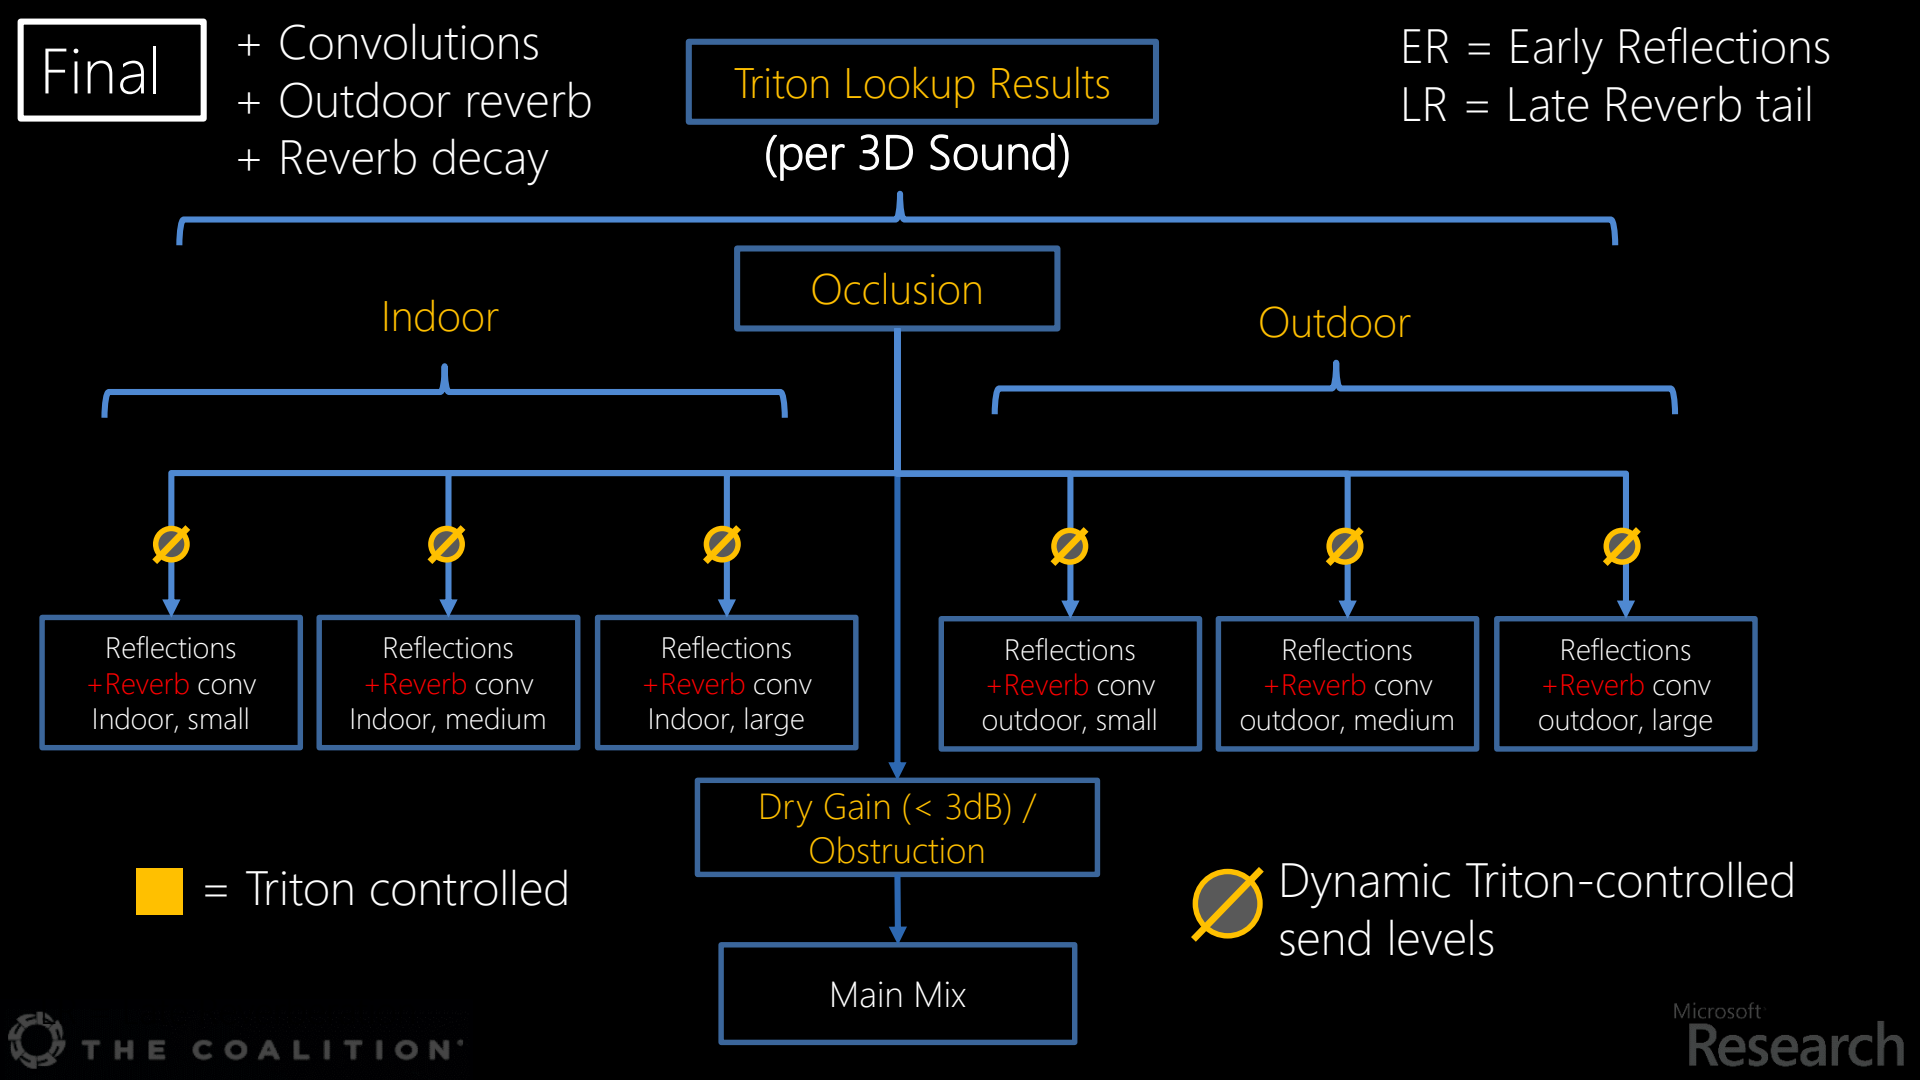
Final
+ Convolutions
+ Outdoor reverb
+ Reverb decay
Triton Lookup Results
(per 3D Sound)
ER = Early Reflections
LR = Late Reverb tail
Occlusion
Indoor
Outdoor
Reflections +Reverb conv Indoor, small
Reflections +Reverb conv Indoor, medium
Reflections +Reverb conv Indoor, large
Reflections +Reverb conv outdoor, small
Reflections +Reverb conv outdoor, medium
Reflections +Reverb conv outdoor, large
Dry Gain (< 3dB) / Obstruction
Main Mix
= Triton controlled
Dynamic Triton-controlled send levels
THE COALITION
Microsoft Research

Gears 4 Triton Implementation: Final

Lessons Learned; Summary

- Tuning and Testing of this unprecedented system was our biggest risk.
  - Reduce number of Wwise parameters Triton is driving. Simplify! Limit Triton Reverb to a sensible dB range.
- Focus on the ***emotionally*** motivated "Cinematic Experience" instead of the reality-inspired "Realistic Experience".

Our Trip to The Uncanny Valley…

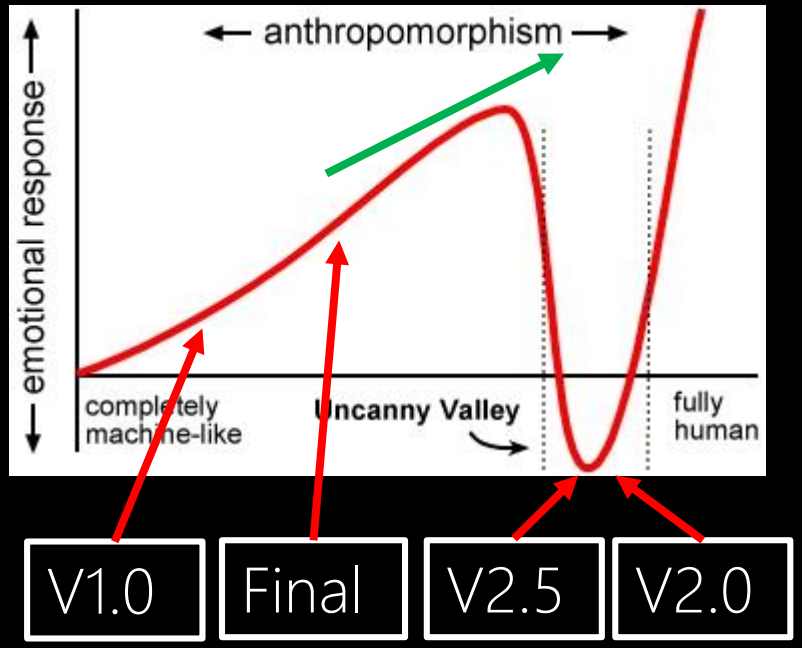

We've only scraped the surface of what's possible with Triton.


Triton ON
300
1. GearMainMenuEntry_P
2. SP_Rescue_B_P
3. GearMainMenuEntry_P
4. SP_Dam_A_P
2 debug jumppoints used
ACT IV - CHAPTER 2
NO DETOURS
SERVER
1st Player Difficulty: Casual

Conclusion

- Baked wave acoustics for dynamic sources & listener is now practical
- Designed occlusion, obstruction, reverb wetness, and decay times in production game levels
- Key idea: Do accurate physics → perceptual data → game
- Future
  - Next: directional initial sound and early reflections, outdoor echoes, streaming 100→25MB RAM, faster bakes
  - Portability
  - Longer term: dynamic geometry e.g. doors/destruction

Special Thanks…

John Morgan (Audio Director),
Mike Rayner (Technical Lead),
Rod Fergusson (Studio Head),
Chuck Osieja (Creative Director)
…for believing we could pull it off.

Jimmy Smith (Lead Audio Programmer),
Clarence Chu (Audio Programmer),
Jess Kube (Tools Programmer),
Michael Oliver (Intern Programmer)
…for their indispensable work.

John Snyder (Principal Researcher),
…for his steadfast collaboration on Triton research for the last 9 years

Chakravarty Reddy Alla Chaitanya (Research Intern)
…for helping invent adaptive sampler

Kristofor Mellroth (Head of Audio),
…for bridging Research and Games

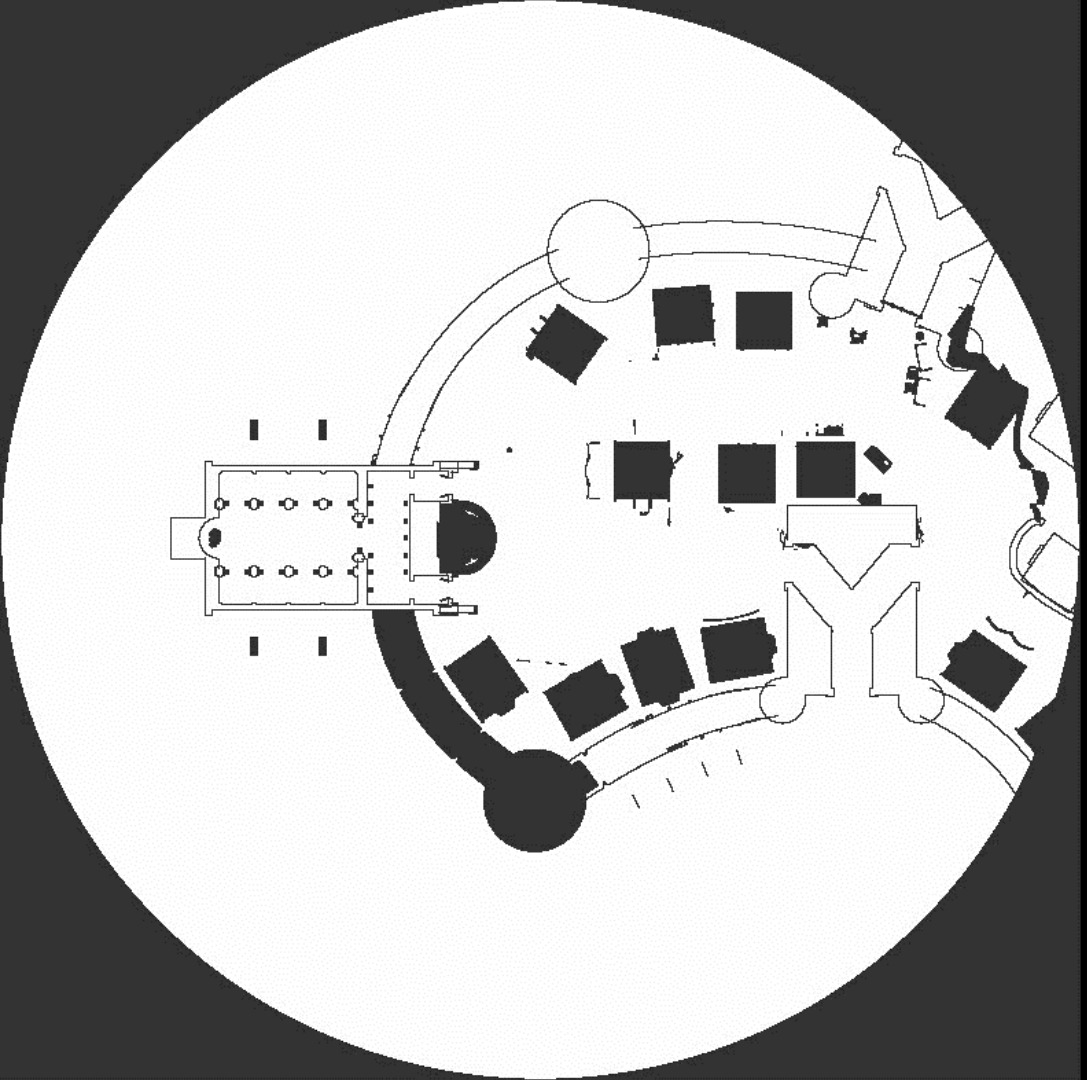

Questions?

Slide deck with videos and papers available @www.nikunjr.com

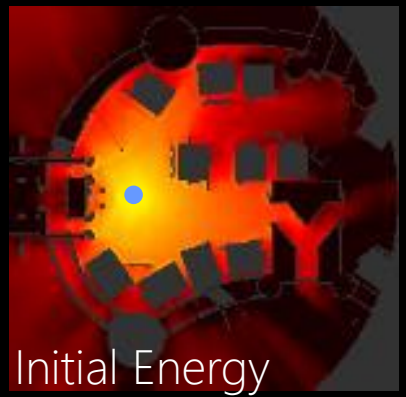

Microsoft Research

WE'RE HIRING!

BEERS OF WAR

- Email – andrewgl@Microsoft.com
- Visit – http://thecoalitionstudio.com/#Join-Us

THE COALITION™