



Sounds from *The Garden*

Butt Music from Hell

Al Biles



R·I·T

Overview

- Audio in Interactive Media
- Case Study: *The Garden*, including...
- *The Butt Music from Hell*
- Advice for doing audio in your project

Audio

- Critical to immersive user experience
 - No longer an enhancement, now an expectation
 - Bar approaching quality of movie sound
- Last set of assets added to most projects
 - Get leftover resources (memory & machine cycles)
 - Absorb overflows from earlier phases (money & time)
 - Have to react to what's already done
- Audio gets squeezed

Linear vs. Interactive Media

- Linear (movies, videos, etc.)
 - Only one path through content, fixed timing
 - Can fine tune audio to fit visuals
- Interactive (simulations, games, VR, installations)
 - Multiple, semi-unpredictable paths through content
 - Timing depends on user actions
 - Audio has to be agile, more reactive

Conventional Audio Assets

- Sound effects / Foley art
 - Visible source in the narrative space
- Ambient sound
 - Can't see source, but can hear it & expect it
- Voice work
 - Dialog, narration
- Music (cues)
 - Theme, background, source
- Interface (UI) sounds
 - Aural feedback for user actions

Diegesis

- Diegetic: Heard by characters in narrative space
 - Sound effects, ambient, dialog, exclamations
 - Source music: Musicians playing in scene, radio, etc.
 - Place in narrative space: L-R pan, level, reverb, etc.
- Non-diegetic: Heard only by user/player
 - Narration, voiceover
 - Theme, background, incidental music
 - Interface sounds
 - Place in user/player's head

Point of Perception (PoP)

- PoP = Point of View + Point of Audition
- 3rd-person (Objective camera)
 - PoP is observer of narrative world through window
 - Diegetic sounds all come from within the window
- 1st-person (Subjective camera)
 - PoP is user/player's position *inside* narrative world
 - Sounds can come from anywhere, even off screen
- Almost all movies are 3rd-person
- Lots of games are 1st-person

Sound Effects / Foley Art

- Trigger
 - User/player actions, game events
 - Simple: Just call play()
- Localization
 - Diegetic => place sound at visible source
 - Use LR pan, gain, EQ, reverb, Doppler (if moving)
 - Dev environment might provide real-time 3D
 - Place sound at source location in virtual world
 - Place mic at camera location
 - Physically model what mic hears

Ambient Sounds

- Trigger
 - When entering new setting, changing game state
 - Loop textures (croaking frogs, wind, street noise)
 - Randomly play discrete sounds (bird calls, siren)
- Localization
 - Diegetic => place them somewhere in the world
 - Looped textures placed “everywhere”
 - Discrete sounds placed “anywhere”

Voice Work – Narration

- Trigger
 - Usually at a game state change
 - Play speech to completion
- Localization
 - Non-diegetic => Usually monaural, minimal effects
 - Supposed to sound inside user's head

Voice Work – Dialog

- Between NPCs (Non Playing Characters)
 - Trigger a line when player speaks
 - Localize to NPC location, with appropriate effects
 - Often synched to animation (especially in cut scenes)
- Between NPC and player
 - NPC addressing player no problem (same as above)
 - Player's speech is a problem
 - Speech recognition not ready for prime time
 - Most dialogs are text-based, at least for player input
 - Messy, to say the least

Music

- Theme music
 - Foreground: Make memorable first impression
 - Usually a stand-alone “tune”
- Background/underscoring
 - Establish setting, enhance mood
 - Not a “tune”: Shouldn’t draw attention

Music

- Incidental music
 - Literally accompanies an incident/event
 - Can “magnify” the event (stinger)
 - Maybe foreshadow the event?
- Source music
 - Diegetic
 - Musicians in narrative space, music on radio
 - Usually localized to origin of music

Interactive Music

- Music should change with in-game events
 - Triggered by narrative or by user/player actions
 - Can't just play cue once and be done
- Usually requires deconstructing the cue
 - Horizontally: into loopable phrases
 - Vertically: into layerable parts
- Events trigger transitions *within* the cue
- Different compositional process

Case Study: Garden of Earthly Delights



The Triptych





The Game

- Game set in the Hell panel
 - “Trapped in a world to which they do not belong, players must find a way to escape. Hell is a strange place, with lost souls locked into hells of their own creation.”
- Player Goal
 - Become the new overlord of the region
- Mechanics: 2 Modes
 - Spirit: Move freely but can't interact with NPCs
 - Flesh: Possess NPC bodies, interact, use stuff
- Gameplay: 1st person PoP
 - Explore and solve puzzles to unlock NPC bodies and possess them

The Project

- 2 years: Most work in spring 2014 & spring 2015
- 12+ faculty, 30+ students from 5+ academic units
- Lots of teams
 - Game design
 - Concept art
 - 3D modelling & animation
 - Textures & UI
 - Development (programmers)
 - **Audio**
 - Research

Quick Project Critique

- Great learning experience...
- Goal morphed
 - From: serious/educational game aimed at the Prado
 - To: an adult-ish game that embraced the hell part
- Scope: Too ambitious (What else is new...)
- Semester: Artificial timeline constraint
 - Students doing final projects for a class
 - All came in at end of semester: No time to integrate
- Bottom line: Needs more polish



Audio Assets

- 140+ individual assets used, all original
 - Lots of SFX, voice work
 - Some ambient, interface sounds
 - Only one piece of music
- Simple interactivity => No middleware needed
- Used Unity localization for some sounds

Audio Team

- Faculty lead
- Students from Interactive & Game Audio class
 - 2 students from spring 2014, 4 from spring 2015
 - Their final projects
- I defined initial list of needed assets, served as consultant, basically stayed out of the way
- Students did all production work & delivered assets to programmers, except...
- ...I did the tune






Audio Students!








SFX / Foley

- No library sounds in final game
- Recorded either in sound lab or in field (literally)
- Sounds for all occasions
 - Footsteps (walk cycles)
 - Cards, coins, knife sounds, mugs clanking
 - Opening doors, crates; picking things up, dropping
 - Heartbeats, dismembering, soul sucking, possession
 - Drinking, pouring, potion sounds
 - Interface sounds





For example: Walk Cycles

- Record lots of footsteps individually or together
- Edit them together into a loop
 - Maybe have alternative loops
 - Maybe have alternative footsteps within a loop
- Different surfaces
 - Leaves 
 - Mud  
 - Water 
 - Wood 

More Foley Examples

- Scroll Opening 
- Fire on Spawn 
- Happiness-Sucking Potion 
- Death Gargle 
- Swan Death 

Ambient Examples

- Ghost whispering 
- Wind 
- Demon horde skittering 
- Scary Voice 

Voice Work

- Dialogs handled with text
 - User interaction uses dialog boxes
 - Decided not to pronounce NPC text
- Did lots of voice acting for several mechanics
 - Reassembling/disassembling wife's body
 - Boss battle with the rabbit
- Lots of ambient speech: Voices in a mob

Voice Work Examples

- Intro Announcement 

- Wife (dis)assembly



- Crowd heckling at boss battle




- Oooohhh!



Music

- No theme or background music
- Only one tune: Linear & diegetic
- Used in mechanic where player solves audio puzzle to unlock tune & bestow new ability
- Leveraged irresistible feature of painting...

The Band!



The Butt Music



Butt Music Aligned



- Plainchant Notation
- Four-line staff for note pitches
- No note lengths
- No measure lines
- No time signature
- Supposed to have lyrics
- Gregorian Chant

Real Plainchant

Climacus:
a Neume

C clef:
This line is C

Vertical Line: Pause or rest
Longer Shorter

III.

Punctum

XVII. s.

5.

C Rédo in únum Dé- um, Pátrem omnipoténtem,

factó-rem cæ-li et térræ, vi-sibí-li-um ómni-um, et

invi-sibí- li-um. Et in únum Dóminum Jé-sum Chrí-

TABLE 1: Plainchant notation of the fifteenth century

	Roman	Ambrosian	Gothic	Hungarian*	Modern Transcription
Punctum					
Virga					
Podatus					
Clivis					
Climacus					
Scandicus					
Torculus					
Porrectus					
Cephalicus					none
Epiphonus					none
Bivirga					
Virga cum orisco**					with ornament
C clef					substitute G clef
F clef					
B rotundus					
B naturalis					
B quadratus					
Custos, Direct					

- Neumes tell how to sing a syllable of the text
- Pitch [sequences] for a single syllable
- Several different systems
- Replaced by our familiar 5-line staff
- Still used in some churches today

* Hungarian notation, gothicized by Austrian influence into a form called Messine-German, spread over an area that coincides with the borders of medieval Hungary. See the introduction to *Missale Notatum Strigoniense ante 1431 in Posonio*, ed. Janka Szendrei and Richard Rybaric (Budapest, 1982), pp. 46–67.

** A name coined here to describe the combination of virga and oriscus that commonly appeared in late fifteenth-century manuscripts and printed books.

Bosch Was No Musician



- No Lyrics: Blasphemy!
- Clef painted out?
- Neumes don't match any known system (just a brush gesture?)
- Neumes don't line up well with lines, so pitches are ambiguous
- Bosch likely using music as "decoration," not as actual music

Music as Decoration



What the Hell...

- Transcribe it anyway!
- Others have done this
- Recently in February, 2014
 - Amelia Hamrick
 - Student at Oklahoma Christian U.
 - Straightforward interpretation
 - Impose 4/4 time by default
 - Initial blog post went viral
 - Got recorded with choral “lyrics”



Amelia's Transcription

Amelia posted a piano version of the torture-victim's backside hymnal

Adapted from Hieronymus Bosch's Garden of Earthly Delights



Hieronymus Bosch



Choral Recording

- Gregorian chant in a really, really big cathedral
- On repeat, imposed 3/4 time with guitar & bass
- Alleged lyrics:

Butt song from hell,

This is the butt song from hell.

We sing from our asses while burning in purgatory

The butt song from hell,

The butt song from hell.

Butts!



AI's Interpretation



- Assume C clef
- Assume all neumes are punctums
- Interpret notes as literally as possible
- Handle ambiguities as musically as possible
- Use butt crack as rests in bottom two staves
- Use instruments from the painting

Butt Music

from Garden of Earthly Delights

Hieronymus Bosch
as told to Al Biles

AI's Score

Musical notation for measures 1-4. The score is in 4/4 time. The upper staff contains a melody of quarter notes: C4, D4, E4, F4, G4, A4, B4, C5, B4, A4, G4, F4, E4, D4, C4. The lower staff contains a bass line with quarter notes and rests: C3, G2, C3, F2, G2, C3, F2, G2, C3, F2, G2, C3, F2, G2, C3.

5

Butt Crack

Musical notation for measures 5-8. The upper staff continues the melody: C4, D4, E4, F4, G4, A4, B4, C5, B4, A4, G4, F4, E4, D4, C4. The lower staff continues the bass line: C3, G2, C3, F2, G2, C3, F2, G2, C3, F2, G2, C3, F2, G2, C3. A 'Butt Crack' rest is indicated above the final note of the upper staff.

9

Butt Crack

Musical notation for measures 9-12. The upper staff continues the melody: C4, D4, E4, F4, G4, A4, B4, C5, B4, A4, G4, F4, E4, D4, C4. The lower staff continues the bass line: C3, G2, C3, F2, G2, C3, F2, G2, C3, F2, G2, C3, F2, G2, C3. A 'Butt Crack' rest is indicated above the final note of the upper staff.

13

Musical notation for measures 13-16. The upper staff continues the melody: C4, D4, E4, F4, G4, A4, B4, C5, B4, A4, G4, F4, E4, D4, C4. The lower staff continues the bass line: C3, G2, C3, F2, G2, C3, F2, G2, C3, F2, G2, C3, F2, G2, C3. A slur is placed over the final note of the upper staff.

Assume all quarter notes in C major (white keys on piano)

Set drones to C

Two butt-crack rests

Add a pitched drum line to provide an uneven rhythm

Hold the last note


Rendered with MIDI

Hurdy Gurdy



- Decided to play the melodic line on a Hurdy Gurdy because:
 - It's a really cool period instrument
 - This is the first depiction of a **buzzing drone string** (called a dog)
- You can allegedly buy replicas of this one for \$5000 (!)

Anatomy of a Hurdy Gurdy



- Rosined wheel excites strings as it's cranked
- Keys shorten strings like an autoharp
- Keys don't affect drone strings
- Buzzer bridge adds different timbre to one drone string
- Basically a string version of a bagpipe (also existed)

Faking a Hurdy Gurdy

- Synthesized with MIDI instruments
- Used a Yamaha MU-128 tone generator (synth)
- Melodic line doubled on two instruments:
 - A weird, physically modeled bowed instrument
 - A detuned sitar
- Drone is two octaves using a Chinese Jing Hu
- Drum is a Japanese Taiko Drum (pitched)
- Recorded a MIDI “performance”

Butt Music

from Garden of Earthly Delights

Hieronymus Bosch
as told to AI Biles



AI's Recording

Monaural mix with a stereo reverb

Not localized in game, just played back

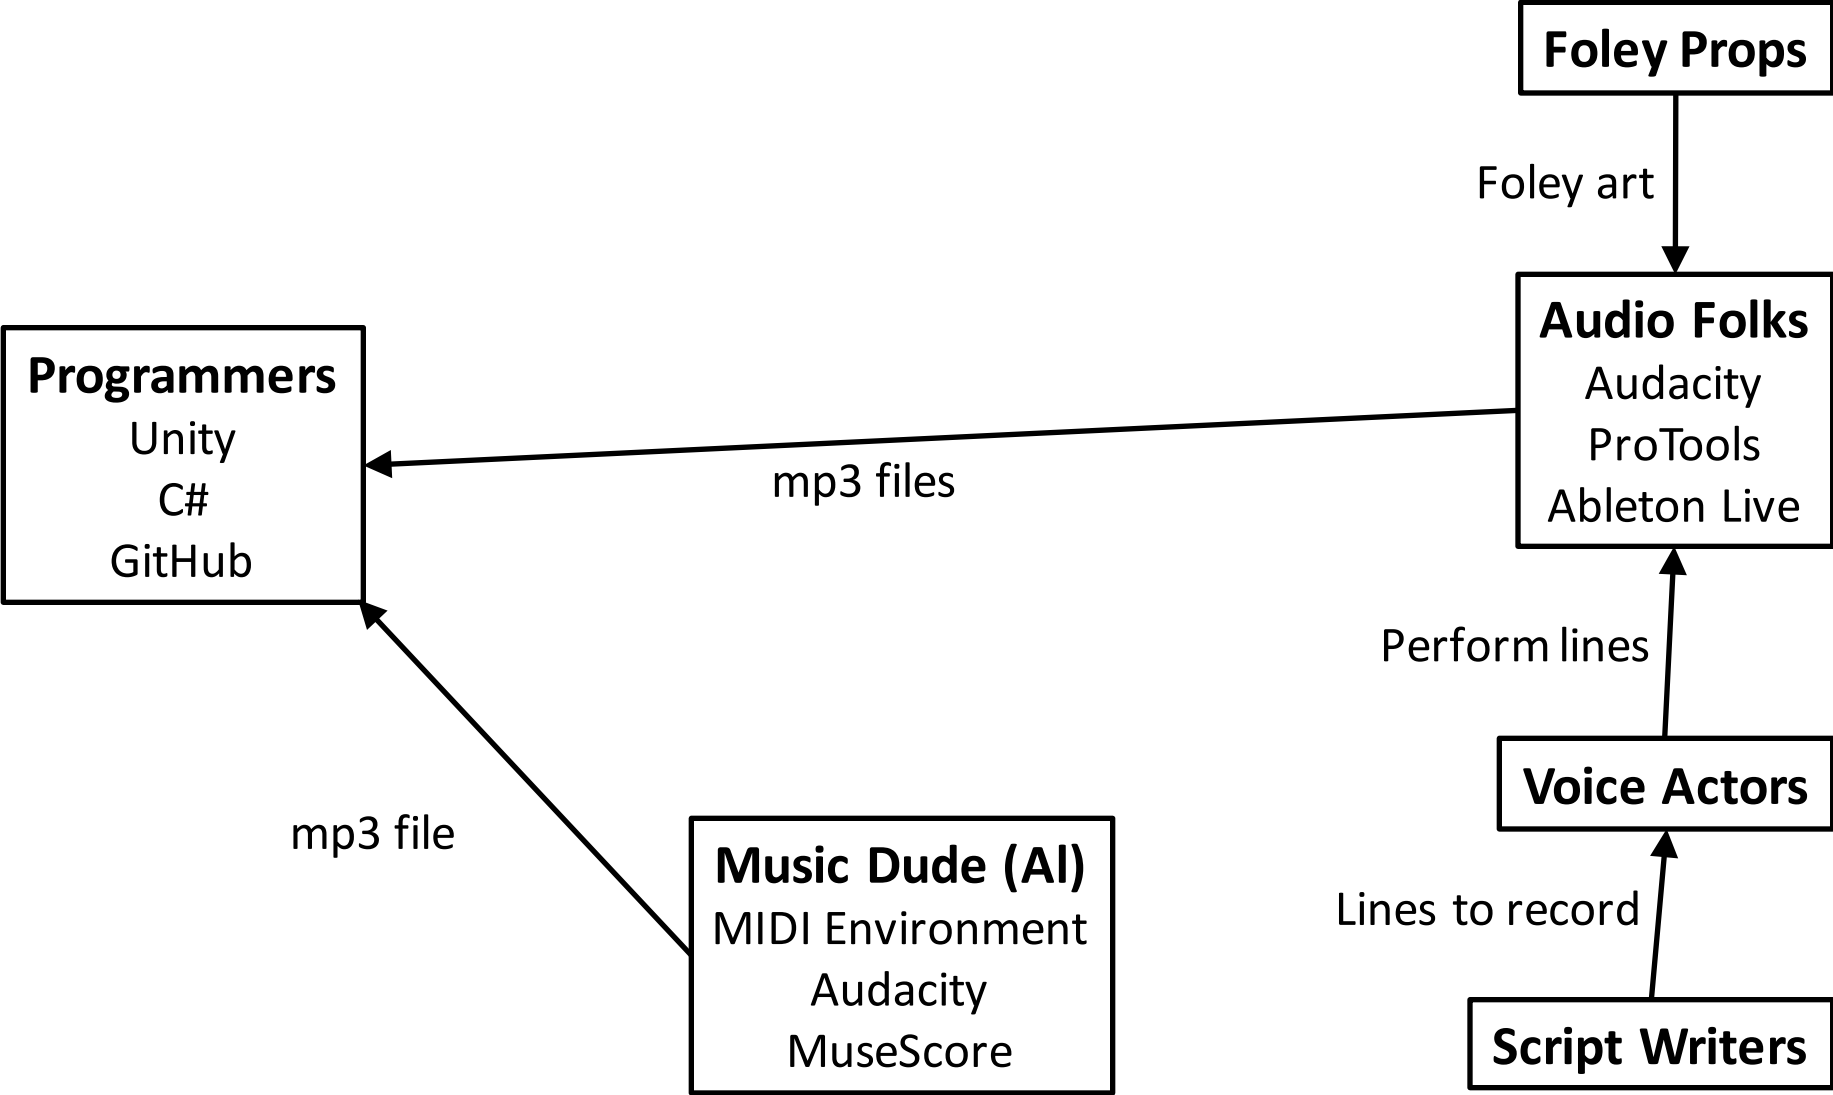
Game mechanic: Player discovers first few notes of the tune, which triggers entire tune & unlocks new ability



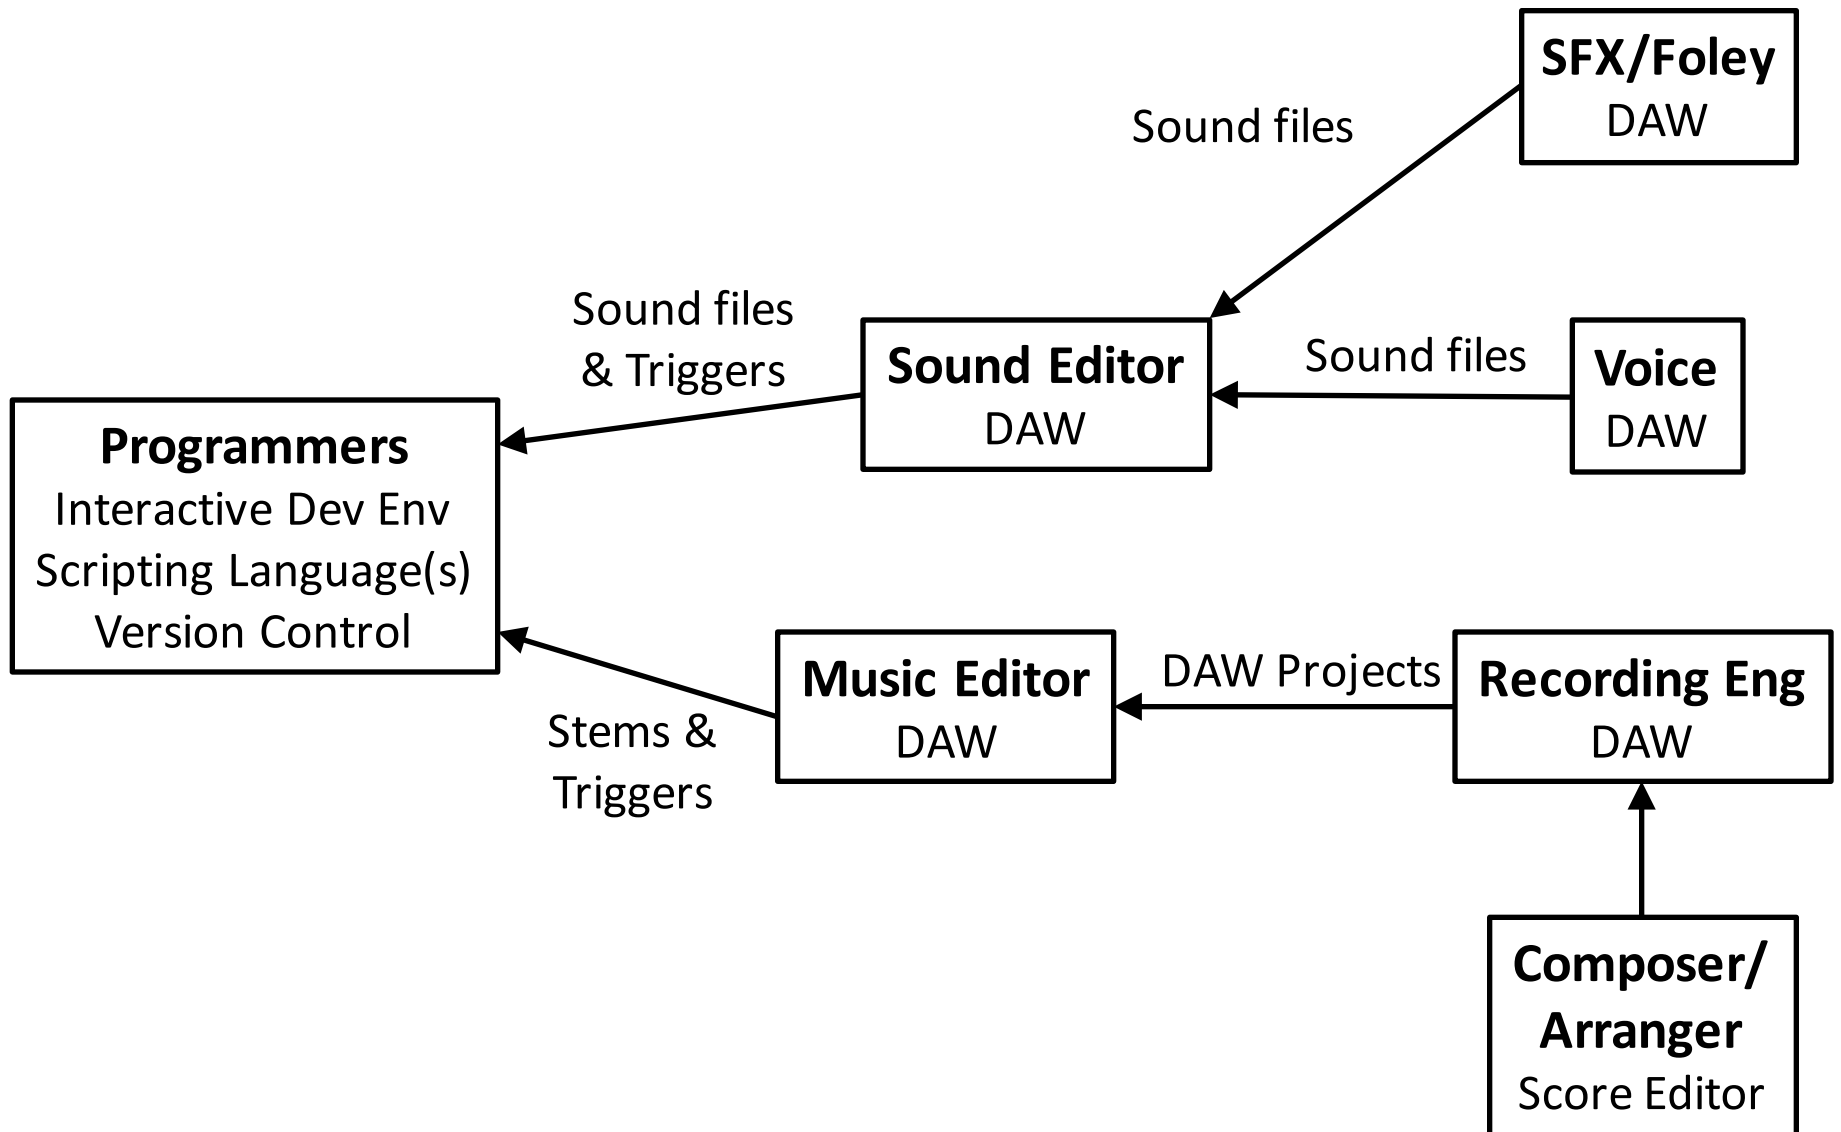
Segue to Advice: Audio Workflow

- Workflow is central to productive production
 - Especially with multiple people on a team
 - Especially especially with multiple teams
 - Especially³ with big, complex, multi-faceted projects
- Helps get a handle on software for different tasks
- Connections among tasks

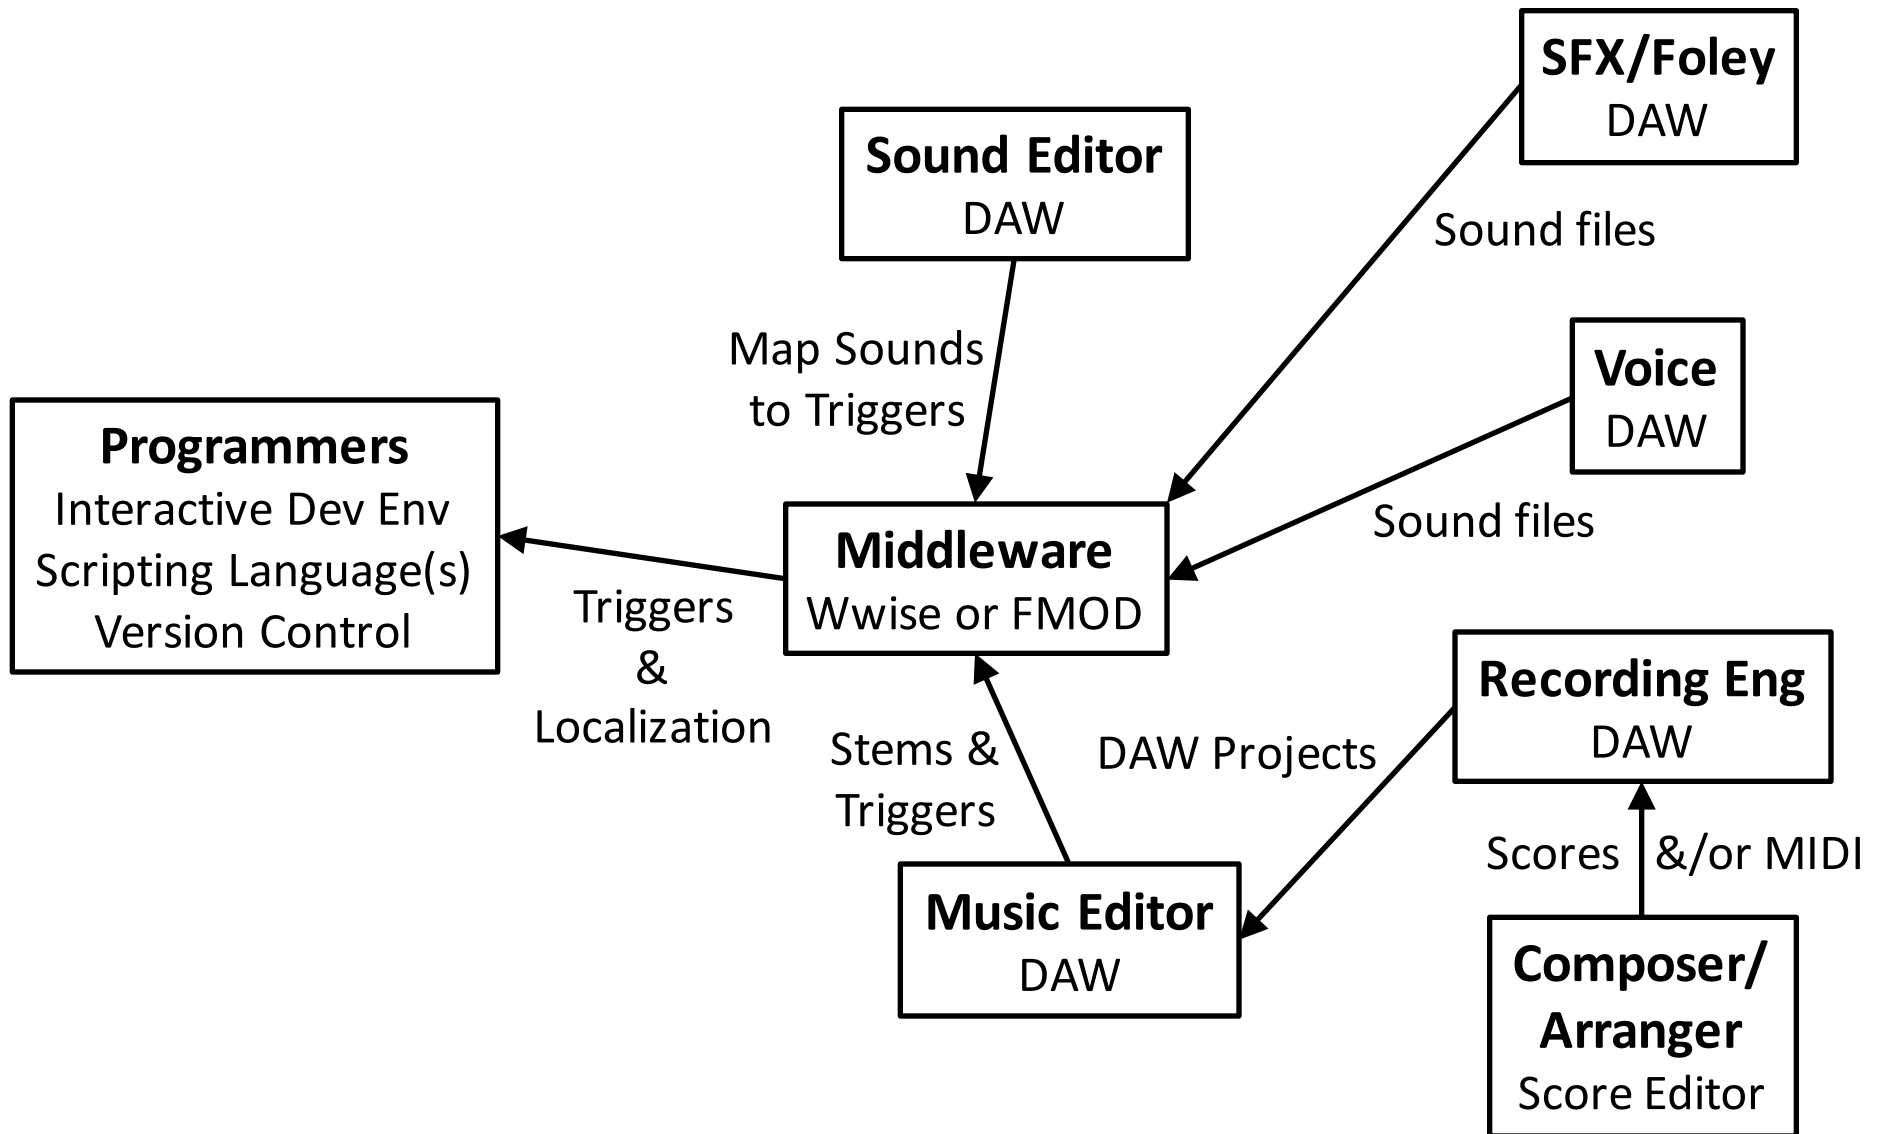
Garden Audio Workflow: Simple



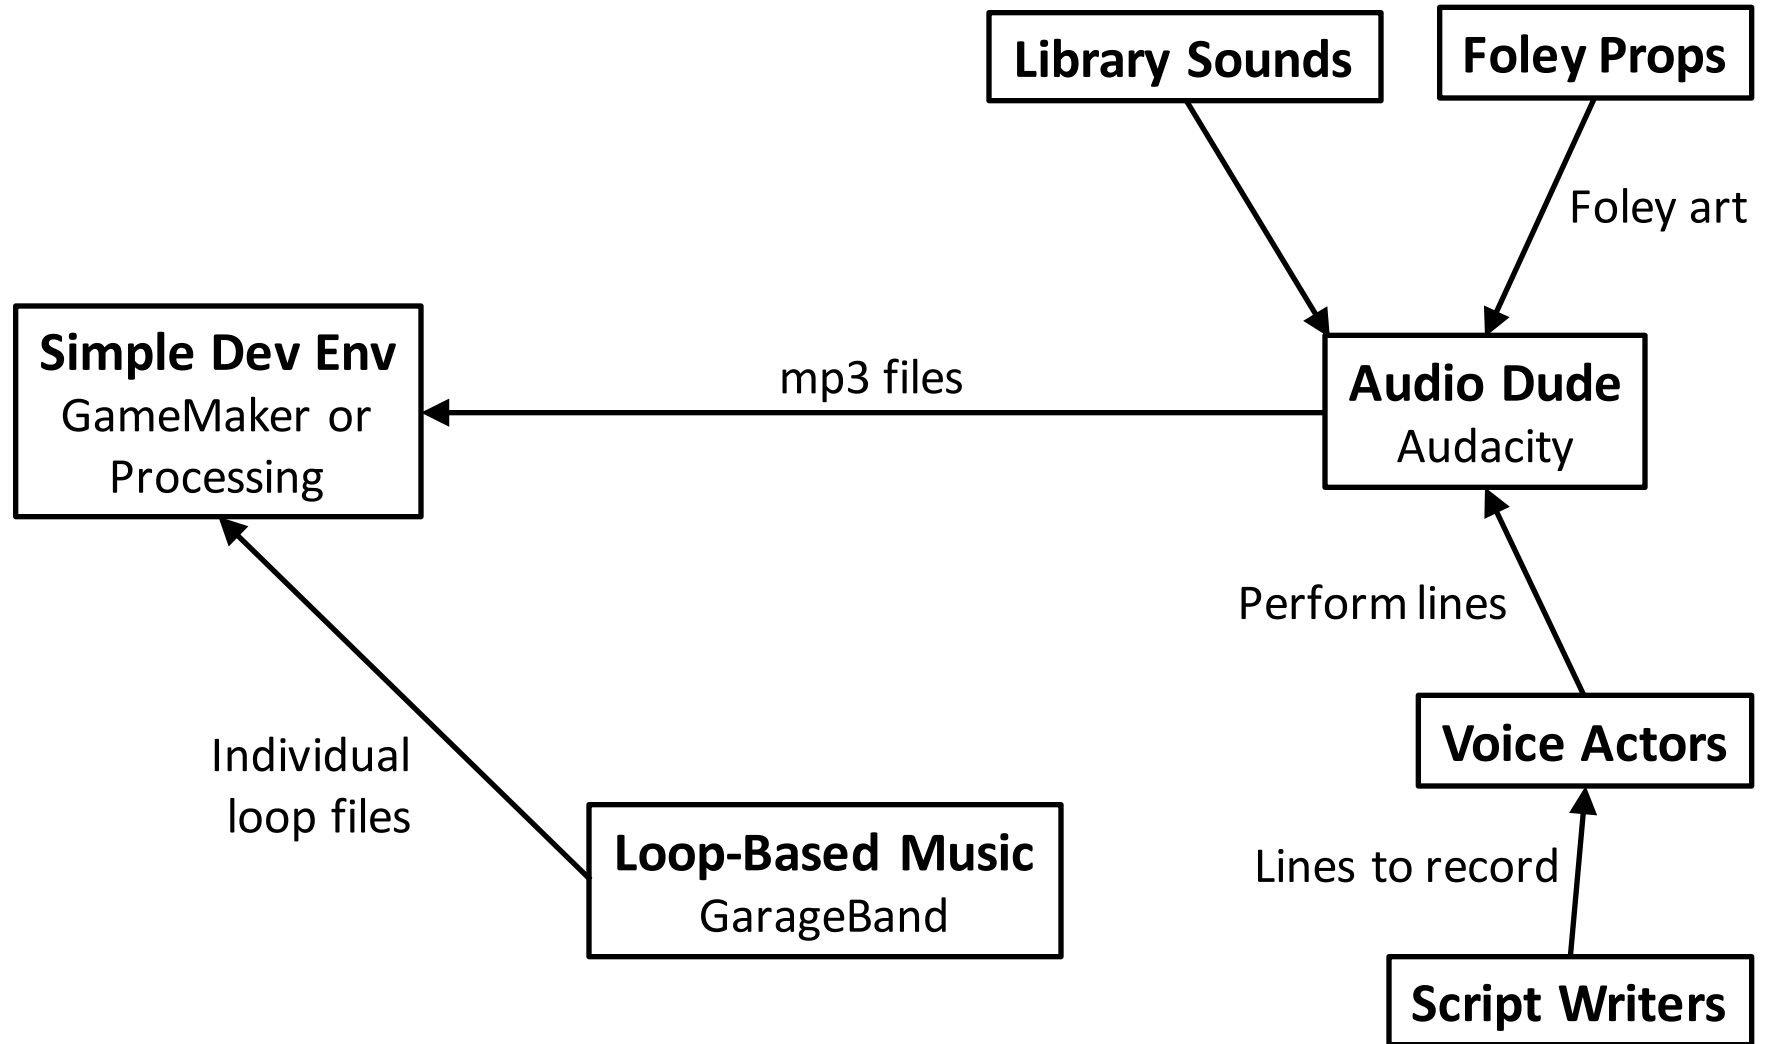
A Generalized Audio Workflow



Audio Workflow with Middleware



Low-Road Audio Workflow



Lots of Audio-Related Software!



finale®



Audio Production Tools

- Audio Editors: Audacity (Open source)
 - Good news
 - Easy to use
 - Lots of effects
 - Large user community
 - Fine for simple SFX, voice acting
 - Bad news
 - No effects chains: Very cumbersome to mix
 - No loop feature: Have to copy/paste
 - No MIDI: only audio for music
 - Very limited for music

Audio Production Tools

- Digital Audio Workstations: ProTools, Logic
 - Good news
 - Industrial strength
 - Effects chains
 - MIDI
 - Almost required for non-trivial music
 - Bad news
 - Licensed (not free)
 - Learning curve
 - Overkill for simple things

Audio Production Tools

- Loop-based DAWs: FL Studio, Ableton Live
 - Good news
 - Loop-based composition/production
 - Full-featured DAW
 - Can “compose” with minimal musical experience
 - Bad news
 - Steep learning curve
 - License issues
 - GarageBand: Good for beginners
 - Flatter learning curve
 - Limited as a DAW

Audio Tools

- Score Editors (Finale, MuseScore)
 - Good news
 - Generate sheet music for performance groups
 - Generate MIDI versions of the score
 - Great for traditional composers
 - Bad news
 - Music only, traditional music at that
 - No audio
 - Not very useful for non-musicians

Audio Tools

- Audio Middleware (Wwise, FMOD)
 - Sits between DAW & game/audio engine
 - Goal: Connect composers, audio folks, programmers
 - Composers can build interactive music w/o programming
 - Audio folks can edit/mix audio in game w/o programming
 - Programmers can place trigger points in code
 - Powerful way to streamline audio workflow
 - Great on projects with complex interactive audio
 - Overkill on small projects with simple audio

Advice for Your Projects

- Do spotting sessions as early as possible
 - From story boards, prototypes, early versions
 - List all the sounds you need
 - Determine trigger points & sequences
 - Time the visuals & synch points within
- Preproduction work
 - Decide which sounds you'll buy vs. make
 - Experiment with Foley props
 - Write scripts for voice work
 - Find voice actors (among your team, most likely)

Advice – Recording

- Recording session
 - At least two people
 - At least one to make sounds
 - At least one to twiddle knobs & push buttons
 - Quiet place (turn off fans, appliances, etc.)
 - Decent mic (USB mic or stereo recorder)
 - Set levels correctly (use headroom, but **never** clip)
 - Save everything **uncompressed** (wav, aiff)
 - Archive originals and write-protect them

Advice – Editing

- Start with copies of archived original recordings
- Work uncompressed (at least 44.1 kHz, 16 bit)
- Don't be stingy with tracks
- Save the multi-track project, not just mixed result
- Save frequently, make alternative versions
- Deconstruct complex sounds when possible
 - May need to deliver stems, not fully mixed sound
 - Keep ambient textures and discreet sounds separate
 - Interactive music requires individual parts, phrases

Advice – Voice Work

- Always have a script, at least to start!
- Record dialogs in sequence in real time
- Record monaurally, preferably one mic per actor
- Directing the actors
 - Leave space between lines unless dialog must overlap
 - Go for good enunciation with appropriate feeling
 - Get different emotional takes
 - Record multiple takes, but remember...
 - You are ***not*** Stanley Kubrick!!!

Advice – Effects

- When to apply?
 - Baked into asset at production time OR
 - Applied in-game by audio engine (active localization)
- Order in which to apply effects (guideline)
 - Compression (squeeze dynamic range)
 - Level (really could be any time)
 - EQ (Treble/bass on steroids)
 - Chorus/phase/flange/distortion
 - Panning (L-R)
 - Reverb

Advice – Putting It All Together

- Lots of audio assets at same time is a problem
 - Can mask each other
 - Can be distracting, overwhelming, break immersion
- Goal: Hear most important sounds clearly
- Solutions
 - Duck less important sounds (dial them down or out)
 - Use audio dimensions to spread sounds out

Audio Priorities

1. Sounds needed for user/player actions
2. Sounds that match visuals
3. Sounds that you'd expect to hear
4. Sounds that enhance events
5. Sounds that convey emotions
6. Sounds that set a mood
7. Sounds that fill gaps

Audio Dimensions

- Location: L-R pan
 - Pan competing sounds to different locations
 - Not always feasible if tied to on-screen sources
- Frequency: Center pitch, bandwidth
 - Control frequency range of sounds with EQ
 - Important sound only one using a particular band
 - Avoid loud “noise” sounds (they eat all frequencies)
- Timbre/envelope/origin
 - Sounds with different attributes can coexist better

Development Process

- Should be spiral model, not waterfall
 - Fixing one issue often creates other issues
 - Whack-a-Mole problem
 - Keep DAW projects around to remix
- Middleware can help
 - Can add Wwise to Unity dev interface
 - Tweak levels, other parameters of specific assets in real time during gameplay

Overall

- Listen with fresh ears
 - Take lots of breaks & do other things
 - Get other folks to listen, playtest
- Document Everything!
 - Each asset should have a dev trail
 - Blog/Wiki can be useful (if folks use it)
 - Content management very useful: Learn how to use it
 - Librarian: Very useful roll for someone to play
- Have fun!!!!

Questions?