The Sound Designer’s Phenomenological Field Guide

A collection of observable sound phenomenon
for sound design applications.

By Joyce Ciesil

Special Thanks

Áine Palmer
Jill BC Du Boff
Justin Ellington
Mike Backhaus

Table of Contents

Introduction

As a human being, I have found that my attention is very rarely something I can control. Referred to by my partner as my “special interests”, I will routinely devote nearly all of my energy and time to a slew of random subjects. Whether it’s Birding, Detroit style pizza, or the events leading up to the 1979 Disco Demolition, I often find myself deeply invested in a niche topic. For as long as I can remember this has been the case, and only one topic has ever been able to hold my attention for a meaningful amount of time; that topic is sound. While my relationship to it has changed over the course of my life, I have been obsessed with sound since I was a teenager, and I have spent the last decade working professionally as a Sound Designer and Composer.

When I interviewed for the Sound Design program at the David Geffen School of Drama at Yale, I was asked the same question that I’d been asked almost a decade earlier when I applied to the Sound Design program at DePaul University: why sound? My answer was the same each time: “because I think it’s magic”.

No matter how much I learn about sound physics, music theory, or how we as human beings hear, sound remains a beautiful and engrossing part of my life.

This field guide was born out of that fascination and a desire to share it with others.

The descriptions of each sound characteristic and phenomena listed within are simplified for ease of comprehension; in that sense, this is intended as an introductory text with which you can begin to familiarise yourself with how sound works and how it can be designed. If any idea manages to capture your attention, let this be a starting point- there are more comprehensive books on sound design and engineering and I highly encourage you to seek them out.

Self Portrait by Joyce Ciesil

What is Sound? What is Noise?

While the difference between sound and noise can be debated, for our purposes, and in the context of this field guide, we will be using the following definitions:

Sound is a vibration that travels through a medium (typically air) that is then heard and interpreted when it reaches a person's ear. Sound is a subjective interpretation of a vibration.[1]

Noise is any form of vibration or disturbance that is either unwanted or unable to be interpreted.[2] Noise can be sound, however noise is objective- it exists even without someone to interpret it.

EXAMPLE: If a tree falls in the forest it creates a noise. If someone is around to hear that noise, it becomes a sound.

If you hear something, and are unable to determine what it is- it's a noise. You may try to make it a sound through description.

EXAMPLE: What was that?

It sounded like a duck.

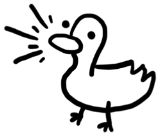

At this point I feel that it is important to note that this definition relies upon the sound being heard and interpreted by a listener. In this field guide we'll be moving forward with the assumption that our listener has two ears, one on each side of their head, and that they are able to hear the given human hearing range of 20hz to 20,000hz.[3]

In reality there is no "standard" human being and I have never met anyone with perfect hearing (nor would I like to). When we hear sound, we interpret it, and because each and every human on this planet has their own perspective, no sound will ever sound exactly the same to two different people. Its beauty isn't found in perfection or harmony, it's in gaining a better understanding of the world around us.

Being able to hear music and the words of our loved ones is just the icing on top.

Characteristics of Sound

Using the definition of sound we've establisted, we can identify some of its physical characteristics. This is not an exhaustive list. Our aim is not an advanced understanding of each of them, but rather a baseline with which we can understand the sound phenomena that will follow.

Frequency and Pitch

Frequency, typically measured in Hertz (hz), measures how often an event repeats itself in a set unit of time. In our context, frequency refers to the number of oscillations a sound wave makes within a second. Sounds with lower frequency vibrations are associated with low pitches, and sounds with higher frequency vibrations are associated with higher pitches. **Pitch** is a subjective way to qualify (describe) and quantify (count) Frequency. Oftentimes Pitch is associated with musical notes, and in a western musical context lower pitches are associated with lower note values while higher pitches are associated with higher note values.[4]

Speed or Velocity

The **speed of sound**, sometimes referred to as **velocity**, is not static. However, it is typically assumed that the speed of sound through air is 1,125 feet per second (at 68° Fahrenheit) or 343 meters per second (at 20° Celsius).[5] Though it seems counterintuitive, sound travels faster through a high density medium and slower through a low density medium. This means that high humidity and high temperatures cause sound to move faster, and low humidity and low temperatures cause sound to move slower.

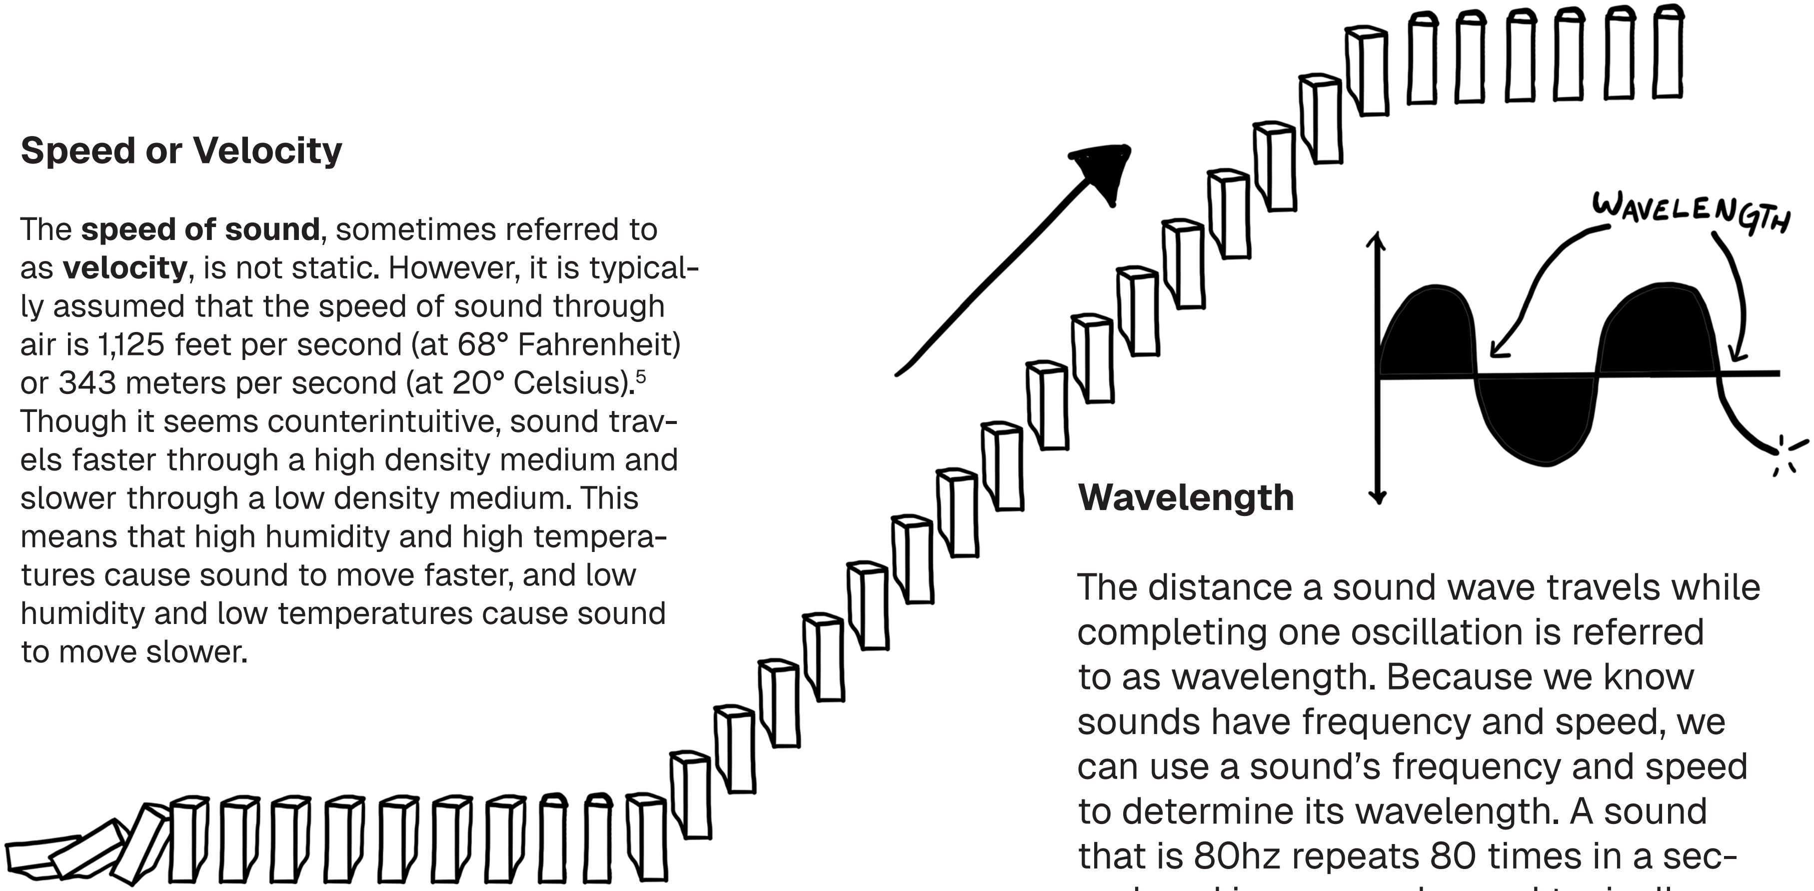

EXAMPLE: Dominoes arranged closer to each other topple over faster than dominoes arranged farther apart, because it takes less time for each domino to collide with the next. Sound waves travel through air in a similar fashion. As Air particles collide with each other, they transfer the vibration of sound. Densely packed air particles are able to collide with each other faster and allow the vibration to travel faster as a result.

Wavelength

The distance a sound wave travels while completing one oscillation is referred to as wavelength. Because we know sounds have frequency and speed, we can use a sound's frequency and speed to determine its wavelength. A sound that is 80hz repeats 80 times in a second, and in a second sound typically travels 1,125 feet. Therefore an 80hz sound wave is 14.0625 feet in length. This equation is written mathematically as:

$$14.0625=1125/80 \text{ or } \lambda=V/F$$

Where the Greek symbol Lambda (λ) denotes wavelength, V denotes the speed, and F denotes the frequency.[6]

Loudness, Quiet, and Volume

Often measured in decibels (dB), Sound Pressure Levels (SPL), or Volume units (vu), **Loudness**, **Quiet**, and **Volume** are how we typically describe the power of a sound wave. Loud, high volume sounds require a great deal of energy and result in higher dB, SPL, or vu, where quiet, low volume sounds require less energy and result in lower dB, SPL, or vu. It is to be noted that loud, high-energy-level sounds can cause damage to ears and in extreme cases other physical objects like wine glasses, lighting fixtures, or windows. In contrast, low-energy-level sounds are sometimes too quiet to be heard. As sound moves through a space and expands its volume it follows what is called the inverse square law. This means that the farther a sound travels, the quieter it becomes.[7]

EXAMPLE: The wings of a mosquito don't generate much energy, so the sound of a mosquito flying is incredibly quiet, and can only be heard when it's right by your ear.

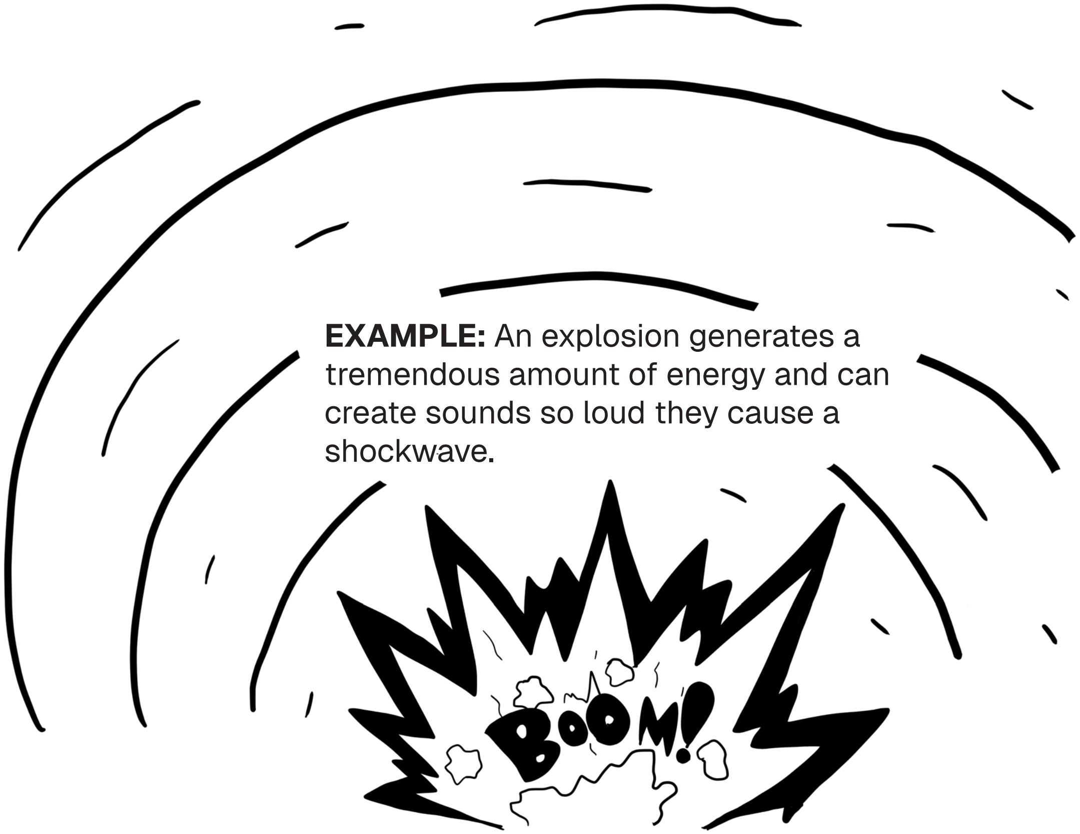

EXAMPLE: An explosion generates a tremendous amount of energy and can create sounds so loud they cause a shockwave.

Amplitude

Equilibrium is a state of balance in which opposing forces cancel out, when the air around us is at equilibrium we typically associate this with silence. As a sound travels through the air, the vibrations create areas of pressure higher than equilibrium **(compression)** and areas of pressure lower than equilibrium **(rarefaction)**. The difference between the compression and rarefraction of each sound wave is referred to as **Amplitude**, and larger differences are associated with increased loudness.[8]

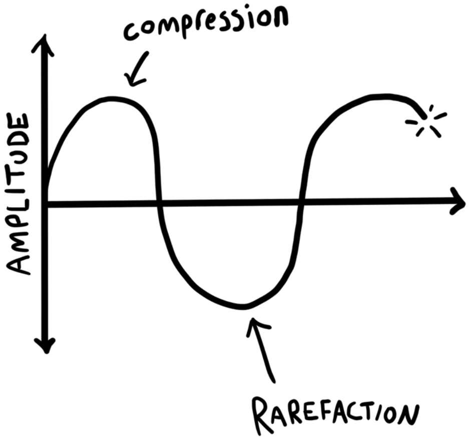

EXAMPLE: Imagine a guitar string. At rest, the string is at equilibrium. When it is struck with a guitar pick the string vibrates back and forth. If we view the string on a flat plane, its movements up represent compression, and its movement down represents rarefaction.

Reflection, Absorption, and Transmission

Put simply, reflection is when a sound collides with something and bounces off of it, absorption is when a sound collides with something and is soaked up by it, and transmission is when a sound collides with something and passes through it. Whether a sound reflects, absorbs, or transmits when it collides with a object is dependent on the frequency of the sound, and the quality of the object it is colliding with. As a general rule, high frequencies are commonly reflected, mid frequencies are commonly absorbed, and low frequencies tend to be transmitted.[9]

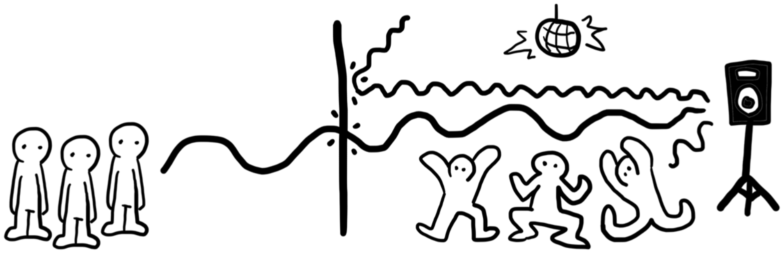

EXAMPLE: Imagine you're dancing in a nightclub to some funky fresh beats. The music is bumping out of the clubs sound system: the high frequencies are reflecting off of the walls of the club, the mid frequencies are being absorbed by the walls and bodies of those dancing, and the low frequencies transmit through the walls so that the bassline can be heard outside the club by the people waiting in line.

Timbre or Tone Quality

When two different voices or instruments make sounds that are similar, or even virtually identical in frequency and loudness, timbre is the word we use to describe the tonal qualities of each that make them different.[10] Words used to describe timbre are highly subjective, and it's recommended to clarify your vocabulary with your collaborators, but words often used to describe timbre include the following: bright, dark, warm, cold, mellow, sharp, tinny, boomy, nasally, gentle, and harsh.

EXAMPLE: A saxophone and trumpet could each play the same note, however due to the method each instrument is played the resulting sounds have a different timbre.

EXAMPLE: Wooden doors and metal doors have a different timbre when they are shut.

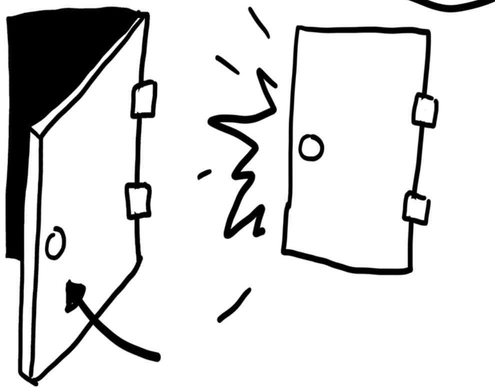

EXAMPLE: Individual human voices have different timbres, but even the same human voice can have a different timbre based on how it is mic'd.

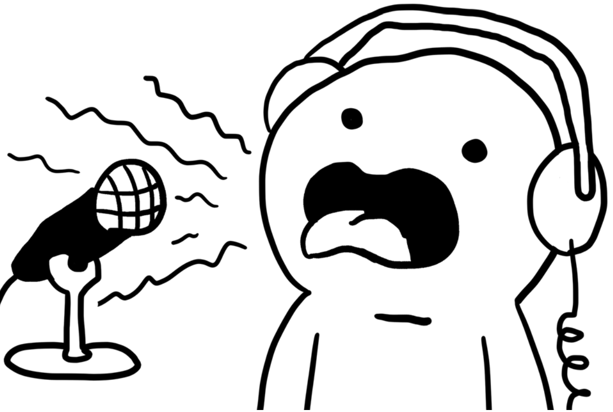

Diffraction and Refraction

When a sound is traveling through a single medium and the speed of sound is constant, diffraction describes the way that the sound spreads out and moves around obstacles.[11] However, when sound is traveling through a medium that has varying density, sound experiences refraction which causes the sound to change directions due to the fact that the speed of sound is not consistent across the medium.[12]

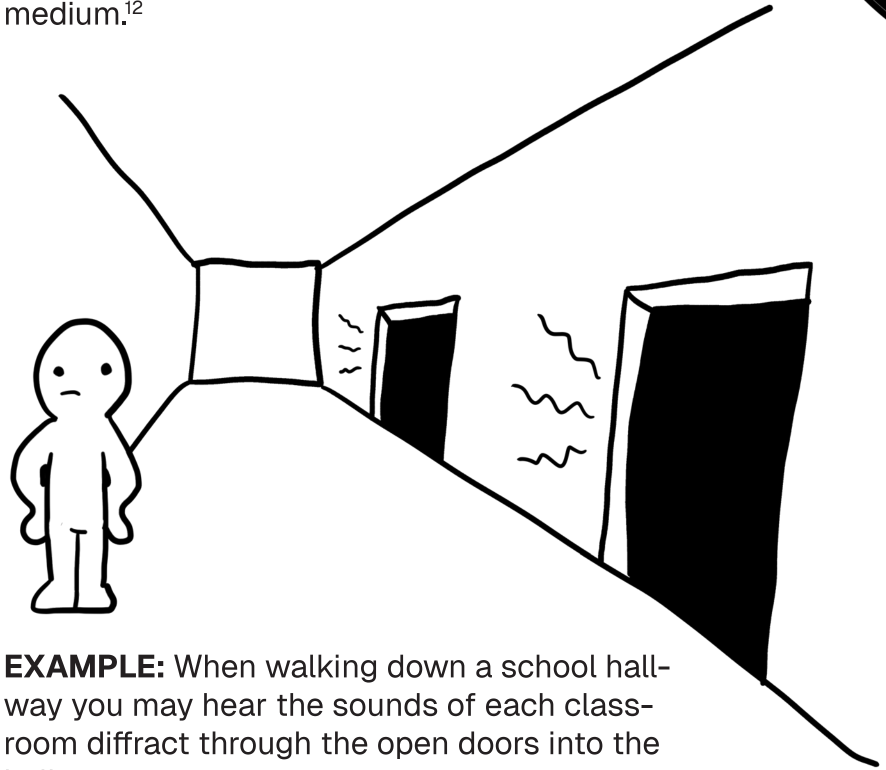

EXAMPLE: When walking down a school hallway you may hear the sounds of each classroom diffract through the open doors into the hallway.

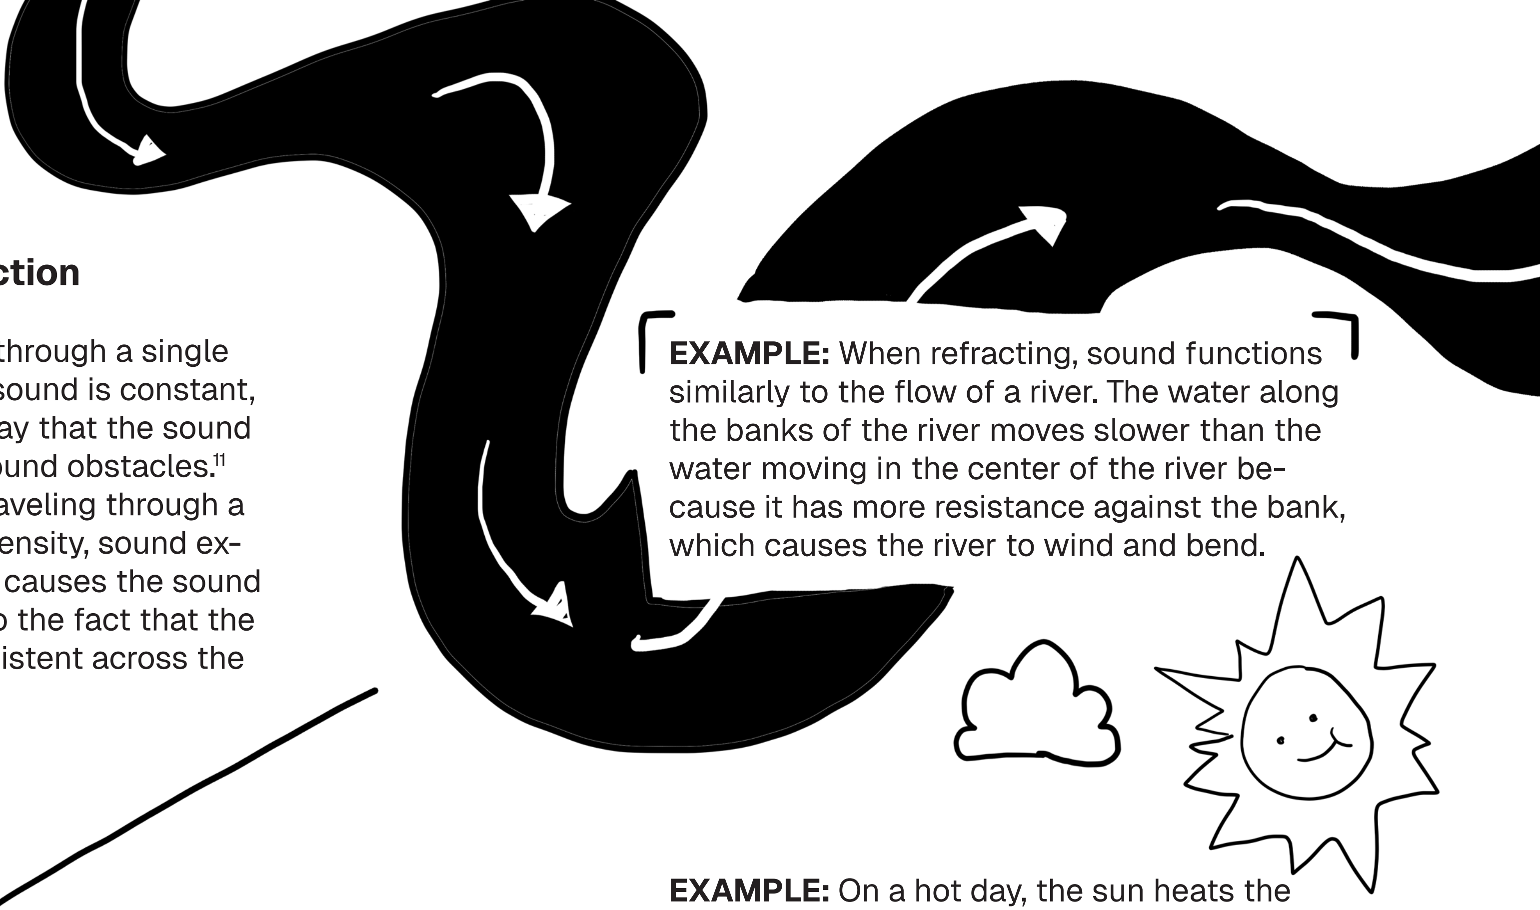

EXAMPLE: When refracting, sound functions similarly to the flow of a river. The water along the banks of the river moves slower than the water moving in the center of the river because it has more resistance against the bank, which causes the river to wind and bend.

EXAMPLE: On a hot day, the sun heats the earth, which then heats the surrounding air. As the air rises into the atmosphere it cools down. This creates a gradient of temperature, and causes the sound waves to refract upwards, because the sound waves nearest the ground are travelling faster than the sound waves higher in the air.

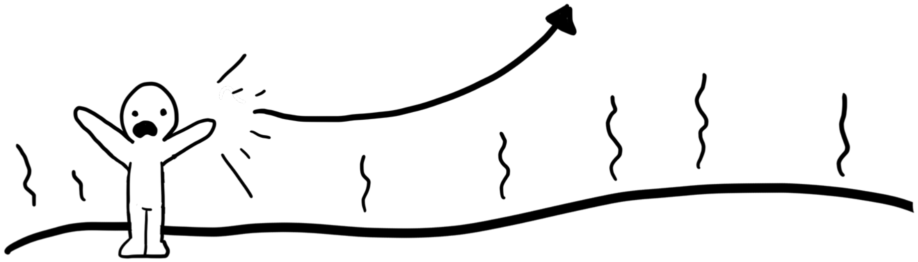

What is Sound Design?

While the definition of sound design varies from industry to industry, I like to define sound design as the art of using sound and sound technology to influence a listener's perspective and experience.

Sound can allow us to transport the audience through space and into the mind and body of characters; it can allow us to slow down the passage of time, and even assign weight and significance to everyday objects. The music we choose can impact a listeners mood and at the most fundamental level sound is what allows our listeners to hear the story or experience that they are engaging in.

Sound helps us gather information about our surroundings, a sound designer curates that information for the audience.

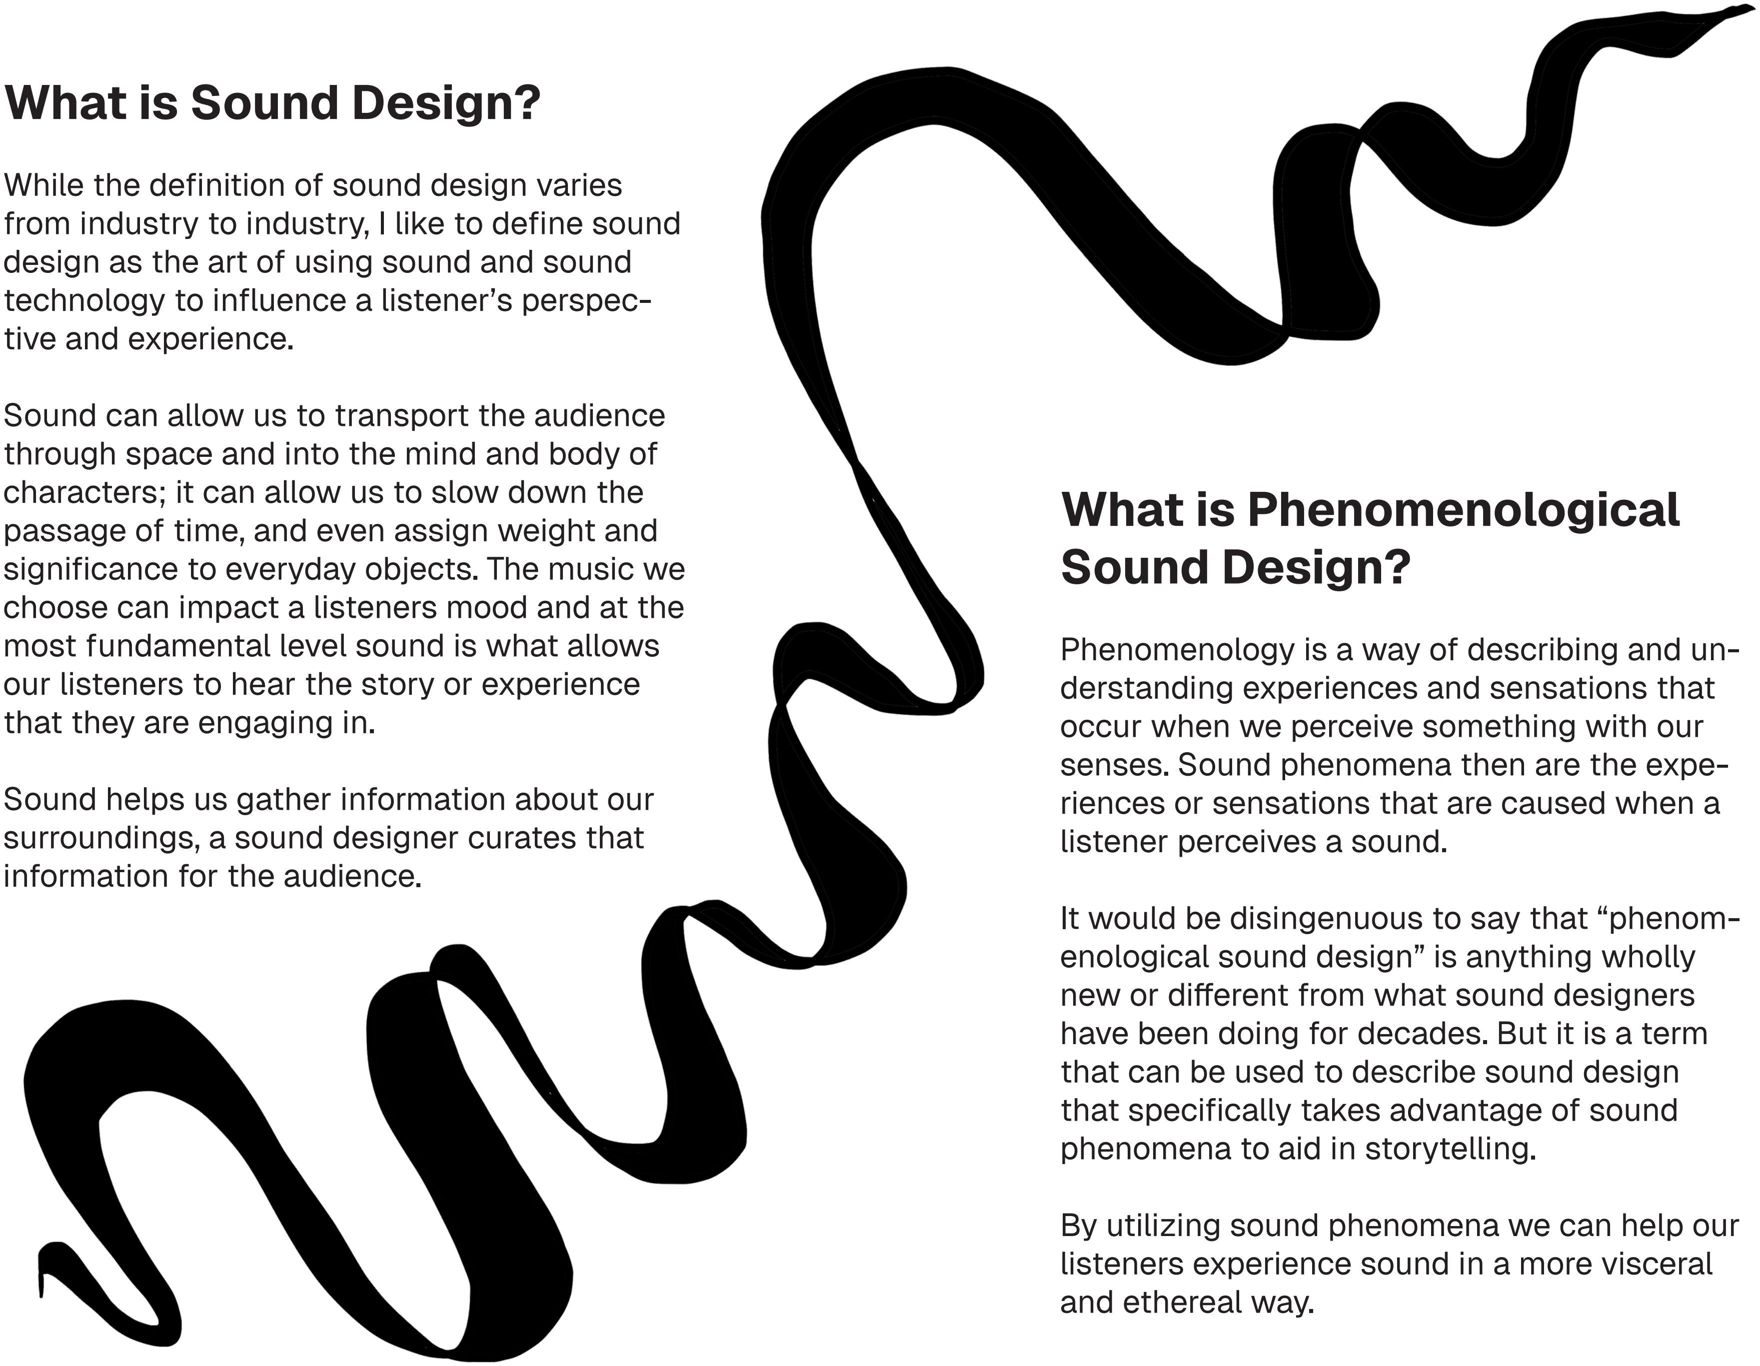

What is Phenomenological Sound Design?

Phenomenology is a way of describing and understanding experiences and sensations that occur when we perceive something with our senses. Sound phenomena then are the experiences or sensations that are caused when a listener perceives a sound.

It would be disingenuous to say that "phenomenological sound design" is anything wholly new or different from what sound designers have been doing for decades. But it is a term that can be used to describe sound design that specifically takes advantage of sound phenomena to aid in storytelling.

By utilizing sound phenomena we can help our listeners experience sound in a more visceral and ethereal way.

Observable Sound Phenomena

Precedence Effect

A.K.A. the Haas Effect and the Law of the First Wavefront

When two identical sounds originating from two different sources arrive at our listener's ears shortly after each other, our listener combines them in their brain into one singular sound. Our listener's brain then attributes the source of the sound to the direction they first heard it.[13]

How does it work?

Because our listener has two ears, one on each side of their head, when sound comes from a singular source both ears hear it. Their brain analyzes which ear heard the sound first, and how much later their second ear heard the sound. The brain then determines which direction the sound came from with this information.

EXAMPLE: Imagine our listeners' ears are 1.125 feet apart, and they're listening to a sound that's coming from a speaker 11.25 feet to their left. The sound from the speaker will hit their left ear 10 millisecond after it leaves the speaker, before hitting their right ear 1 millisecond later. Their brain would then conclude that the sound originated from their left, because their left ear heard the sound first.

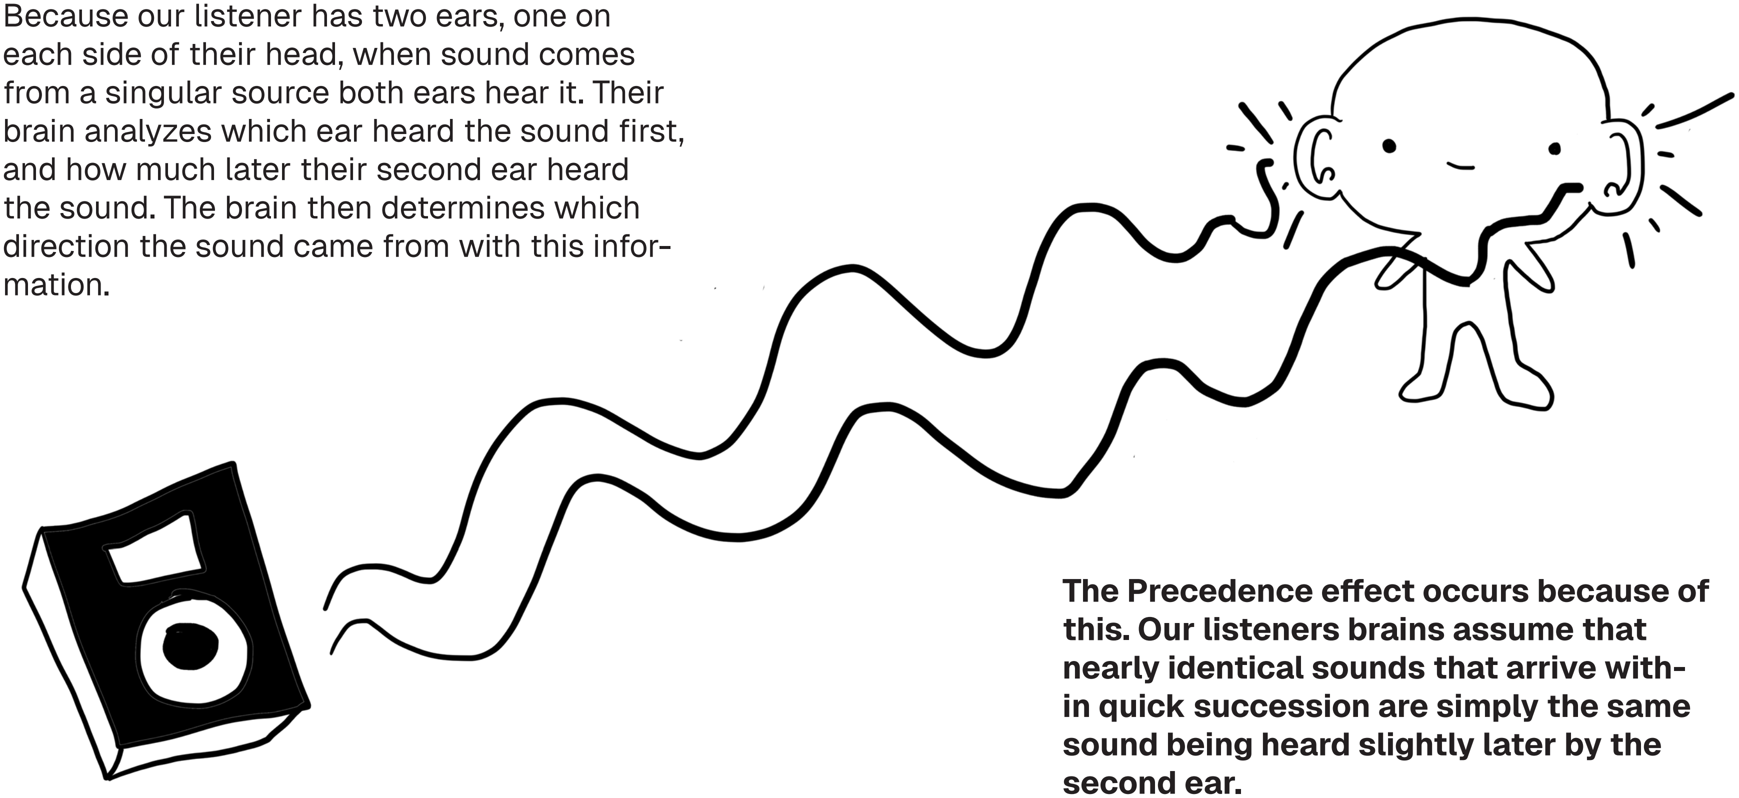

The Precedence effect occurs because of this. Our listeners brains assume that nearly identical sounds that arrive within quick succession are simply the same sound being heard slightly later by the second ear.

EXAMPLE: Imagine the same situation as before, but we've added another speaker 11.25 feet away from our listener on the right. If we play the same sound from both speakers the sound arrives into each ear at the same time, and our listener will perceive it as coming from both directions. But if we add 1 or 2 milliseconds of delay to our right speaker, it will help trick our listener's brain into hearing the sound originating from the left.

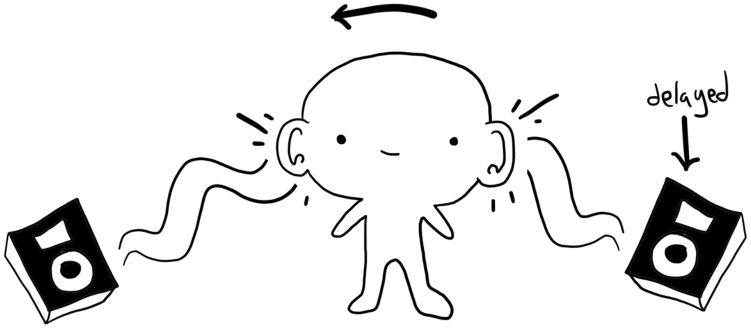

NOTE: The Precedence Effects effectiveness is also related to the type of content the speakers are playing. Typically shorter delay times like 1- 4ms work for content that is a short click, but delay times up to 20 ms or even 40ms can work for sustained tones. If the delay between the arrival times of the two sound sources is too great, the listener will no longer hear the sounds as one, but instead as two distinct sounds.[14]

Design Applications

The Precedence effect can be used in recording, but it is primarily used when tuning sound systems. When working in large theaters that require amplification, Sound Designers can utilize the Precedence Effect to trick the audience into believing that the sound they are hearing is coming from the stage, when in reality it is coming from their speakers.

EXAMPLE: An actor is speaking into a microphone, and the microphone is going to a speaker. The Sound Designer can delay the speaker such that the sound of the actor's voice reaches the audience slightly before the sound from the speaker. When done with a sufficiently small delay, the audience will experience the Precedence Effect. It will feel like the sound is coming from the actor on stage, even though a large percentage of the volume is coming from a nearby speaker.

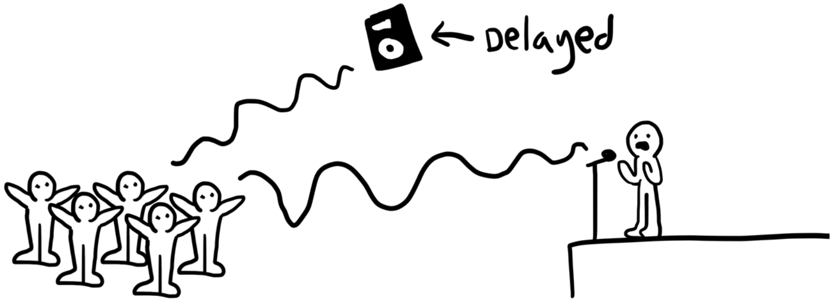

The Doppler Effect

When the distance between our listener and a sound source changes while the source is emitting sound, our listener will hear a change in the sound's frequency (i.e. pitch) based on the speed the sound source is moving and the change in distance between it and our listener. In other words, the pitch of the sound will increase as the sound source moves closer and decrease as it moves away.

How does it work?

The speed of sound is relatively slow. We know that the speed of sound relates directly to wavelength, and wavelength relates directly to frequency.

If a sound source is moving towards a listener at sufficient speed while it is making a sound, each cycle of the sound wave is emitted from a position closer than the previous one. Conversely, as it moves away from a listener, each sound wave is emitted from a position further away than the previous one. Functionally this shortens each wavelength as it approaches the listener, and lengthens each wavelength as it moves away. Because shorter wavelengths mean higher frequencies, and longer wavelengths mean lower frequencies, the Doppler Effect changes the perceived pitch of the sound.[15]

EXAMPLE: If an ambulance is driving by our listener at sufficient speed while its siren is blaring, our listener observes the siren increase in pitch as it approaches and decrease in pitch as it drives away.[16]

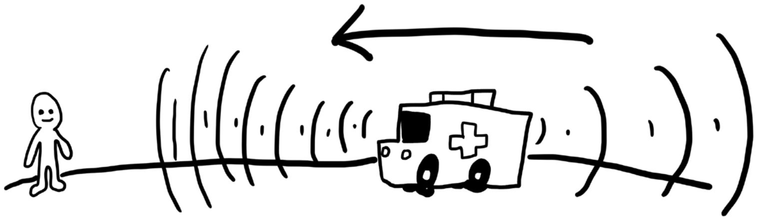

Design Applications

When working with recorded media, whether it's podcasting, audio books, film, or recorded sound effects for live theater, the Doppler Effect can be used to add realism to the sound effect. A sound effect of a car driving past a microphone may already have the Doppler Effect built in the recording, but if a Sound Designer wanted to add the sound of music playing from the car's stereo, the sound designer could manually add the Doppler Effect to the music by manipulating its rate of playback. Additionally, if a sound designer wants to give the impression that a sound effect that was initially static is now rapidly approaching the listener, they could add a small increase in rate in addition to an increase in volume, to indicate its proximity to the listener is changing.

Reverb and Echo

When a sound reflects off of a surface and is heard by a listener, it is referred to as Reverb (short for reverberation) or Echo. Reverb describes the subtle effect of a sound reflecting in a large space (like someone singing in a church), while Echo usually describes the more dramatic effect of a sound being clearly reflected back at the listener, giving the illusion that the sound is being repeated back after a short delay (like when someone shouts into a canyon).

How Does it Work?

Reverb and Echo utilize the speed of sound and the delay that is incurred when sound reflects off of surfaces. The distinction between Reverb and Echo relates directly to the amount of time between the source audio and the reflection, as well as the clarity of the reflection. Reverb is typically associated with shorter delay times, and less clear reflections, while Echo is associated with large delay times and clear reflections.

EXAMPLE: If someone is singing in a large hall, and our listener is standing near the singer, our listener hears the singer but they also hear the reflections of the singer's voice off of the walls. In this instance the reflections are coming from a variety of angles and distances away from our listener, and the time it takes the sound to reflect off the wall and return to our listener is too great for the Precedence Effect (see page 22) to take place. Our listener then hears the singer's voice and the natural reverb of the room.[17]

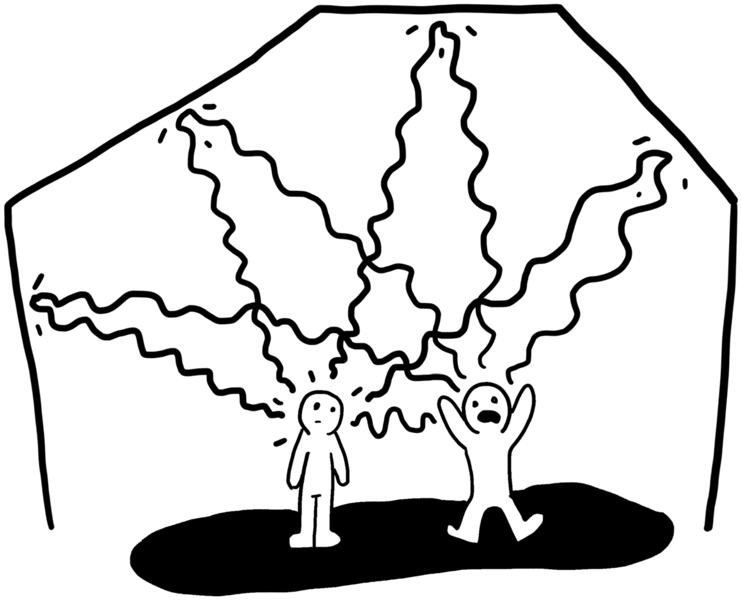

EXAMPLE: If our listener is standing next to a cave explorer, and they shout into a long cave. Our listener will hear the voice of the explorer and then after a brief pause they will hear the reflection of their voice off of the back wall of the cave. This is called Echo. The Echo that our listener hears will be far more clear than the reflections of the reverb, because as opposed to the many angles of incidence in the large hall, the back of the cave is flat and more evenly reflects the sound. The delay between the initial sound of the spelunker and the echo is also greater, as the back of the cave is further away from our listener than the walls of the large hall.

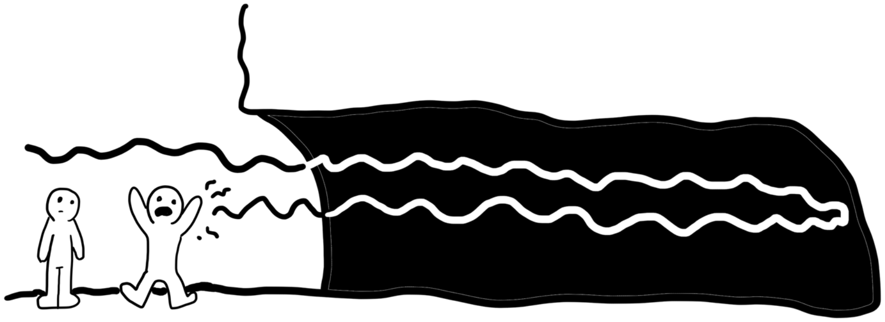

NOTE: Typically in the context of Reverb and Echo the original sound is referred to as the Dry Signal, and the reflections are referred to as the Wet Signal. The relationship between the two is called the Dry/Wet Mix.

Design Applications

If we know that Reverb and Echo rely on sound reflecting off of surfaces, Sound Designers can use Reverb and Echo to give listeners an impression of the size of a room. Sounds can be recorded in a studio setting without reverb, then played in a space that has natural reverb or echo and re-recorded. Spring and Plate Reverb units make use of sound transmission to generate a Reverb- like effect by transmitting sound through a physical medium. And digital Echo and Reverb plugins use software to generate a Reverb- or Echo-like effect by manipulating the original sound and adding delay.

EXAMPLE: When creating a sound design for a podcast you may want it to sound like the characters are in a large cavernous space. In order to do this, you could add a Reverb or Echo plugin to give listeners the impression that the walls of the space are quite far away even without listeners being able to see them.

EXAMPLE: When recording an album, you may want to add a room mic to capture the natural reverb in the room that the band is playing in. You can then control the amount of natural reverb on the track when you mix the album.

NOTE: Reverb and Echo are not mutually exclusive, a sound can have both naturally.

Shepard Tone

A Shepard Tone is a tone that is created by combining a number of overlapping sustained sounds that are separated by octaves and have varying volume levels. A series of Shepard Tones played in a sequence is called a Shepard Scale and can give the illusion of an infinitely ascending or descending scale.[18] A popular variation of this, known as the Shepard-Risset glissando, utilizes sliding pitches that appear to slide upwards or downwards infinitely.[19]

How Does it Work?

In a Shepard Tone the lower and higher octave sounds are played at a lower volume than the middle octaves, with the lowest and highest of those octaves being barely audible. When a Shepard Scale is ascending, the notes of the scale give the impression that they are infinitely ascending because the notes most audible to our listener are ascending. As notes reach the point where they would no longer be able to ascend they become inaudible and are replaced by inaudible notes at lower octaves. As the lower octave notes in each tone are constantly ascending in pitch and volume as well, our listener experiences a continual increase in pitch for the notes they are able to hear.

NOTE: This sound phenomenon functions similarly to a barber's pole illusion.

Design Applications

When composing or adding ambient drones to a piece, whether it is in podcasting, video games, film, or live theater, a Shepard Scale or Shepard-Risset glissando can create a dramatic impression.

EXAMPLE: In theater, a sound designer may employ the use of an ascending Shepard scale in order to convey the rising tension of a scene over a long period of time. An infinitely ascending pitch can evoke a sense of anxiety in the audience in expectation of a dramatic explosion of emotions.

EXAMPLE: In a videogame the use of a descending Shepard-Risset glissando could be used to add to the feeling that the player character is falling. This effect is particularly useful if the amount of time the player character is falling is not predetermined.

The Levitin Effect

Listeners, even those without any musical training, are typically able to remember sounds in the correct key.[20]

How does it works?

While there is not necessarily a uniformly agreed upon cause of the Levitin Effect, Levitin offers the possible explanation that it is because a listener's brain has the ability to recognize patterns.

Design Applications

When establishing a recurring musical theme or even when using a recurring tonal sound effect, the Levitin Effect tells us that our audience will typically remember the original key, which allows us to change the key in order to imply emotional shifts.

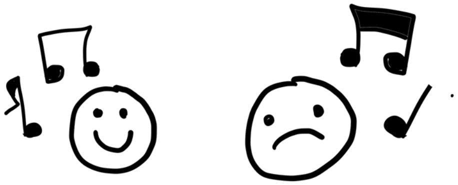

EXAMPLE: A composer can create a recurring musical phrase to be played throughout the course of a film. If the motif shifts into a minor key once the characters are confronted by their nemesis, the music might imply that things have taken a turn for the worse.

EXAMPLE: A sound designer could use an elevator ding in a minor key for a play. As the main character has their soul sucked out of them by their office job each day, they return to work and hear the sad elevator ding. At the end of the play, when the main character quits their job, the elevator ding could shift to a major key to imply optimism.

The Cocktail Party Effect

When exposed to a large amount of stimuli like a noisy environment, listeners are typically able to hear sounds that their brains have assigned importance to while filtering out unnecessary noise.

How does it work?

This effect is called the Cocktail Party Effect because it is readily observable at cocktail parties, where listeners are able to hear the conversation they are a part of while filtering out the noise of other conversations nearby. Interestingly, while our listener may not be listening to the other conversations in the room; if someone else nearby were to say our listeners name they would typically hear it amongst the noise.[21] This implies that our brains are constantly sifting through incoming stimuli and assigning different degrees of importance to them, without us being consciously aware of it.[22]

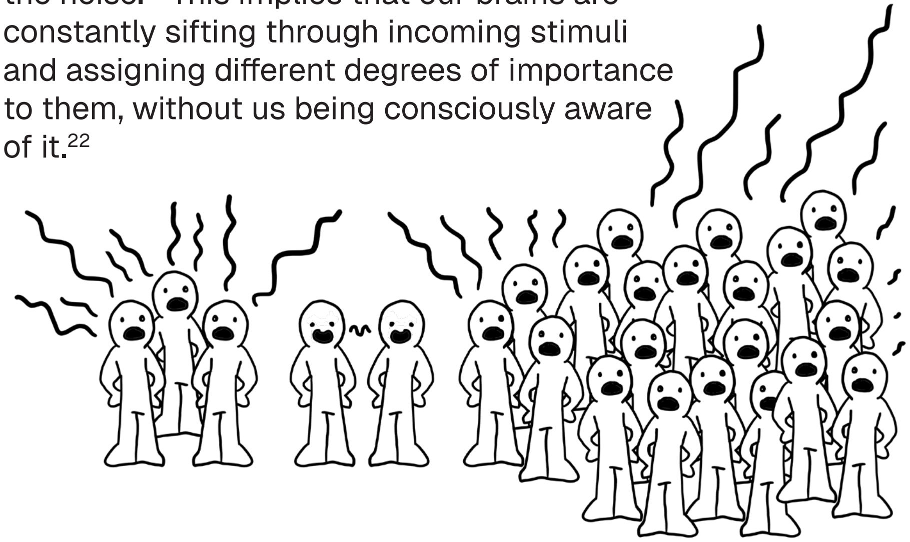

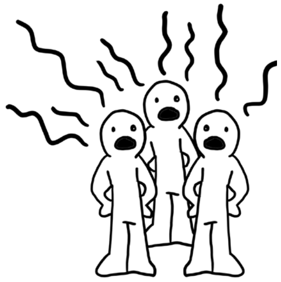

Design Applications

Our brains will naturally and unconsciously filter out unnecessary noise; sound designers can manipulate what sounds are filtered out, and the existence of the Cocktail Party Effect can make this gesture feel natural.

EXAMPLE: In a film setting, a Sound Designer could establish the sounds of a pirate ship during a storm at a high volume, and decrease them subtly as the characters begin shouting to each other. The audience will be aware of the impact that the rain, waves, and lightning are having on the ship, but they'll still be able to accurately hear the dialogue, mirroring how our listeners' brains would function in a similar setting.

EXAMPLE: In a live theatrical setting, a Sound Designer could artificially raise the noise floor by playing white noise or a room tone through the sound system. Over time, especially if the room tone is introduced slowly, the audience will begin to filter it out. At a dramatic moment in the play the Sound Designer could then stop the room tone and dramatically lower the noise floor. Giving an impression of the air being sucked out of the room, and making the audience incredibly aware of the silence.

Proximity Effect

AKA a Bass Tip-Up

The Proximity Effect is how we describe an increase of low frequencies when a sound source is moved closer to a directional microphone.[23]

How does it work?

When two sounds are in sync, meaning the compression and rarefaction of their sound waves line up, we describe the sounds as being in phase. Directional microphones utilize ports, or sound holes, on the sides of their diaphragm in order to create a difference in phase between the sound in front of and behind the diaphragm. This phase difference allows the diaphragm to move more easily as the pressure in front of the diaphragm and behind the diaphragm are not equal.

When you speak into the microphone off axis, the sound enters through the ports at the same time and the soundwaves in front of and behind the diaphragm are in phase. This makes the pressure in front of and behind the diaphragm equal, and it makes it harder for the diaphragm to move. This causes sounds that are off axis to sound quieter.[24]

When you're speaking and you move a directional microphone closer to your mouth, the compression and rarefaction of your voice largely stay the same while their proximity to the microphone changes. The microphone perceives an increase of amplitude as you get closer. This results in the pressure difference on the diaphragm rising. As a result, we hear an increase in low frequencies.

Design Applications

Sound Designers can utilize the proximity effect to change the timbre of a voice or instrument when micing them.

EXAMPLE: Radio broadcasters make use of the proximity effect to add depth and a sultry quality to their voices.

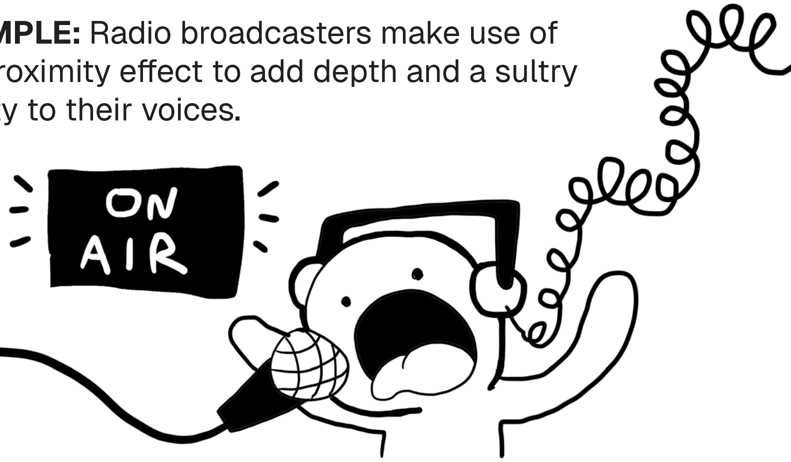

EXAMPLE: A rap artist might use the proximity effect to add more low frequencies and intensity to their voice.

Illusory Continuity of Tones

An auditory illusion that occurs when a listener hears a tone interrupted for a short period of time by a noise. In this instance the listener will hear the burst of noise, but perceive that the initial tone never stopped playing.[25]

How Does it Work?
There is no verified explanation for how this phenomenon occurs; generally, however, our brain intentionally fills in gaps for us.

In order for the illusion to occur, the volume of the noise needs to be as loud as (if not louder than) the tone, and the tone's volume needs to remain relatively static before and after the burst occurs. The pitch can remain constant, ascend, or descend, and the listener's brain will still fill in the gap. This phenomenon can also extend to other non-continuous sounds like human speech, and drum patterns, though results may vary.

Design Applications

When composing intricate rhythms for a drum kit, it is often impossible to hit a high hat on every sixteenth note. However, by having your performer hit a crash or ride symbol on the sixteenth, the listeners' brains fill in the information and believe that the high hat pattern was continuous.[26]

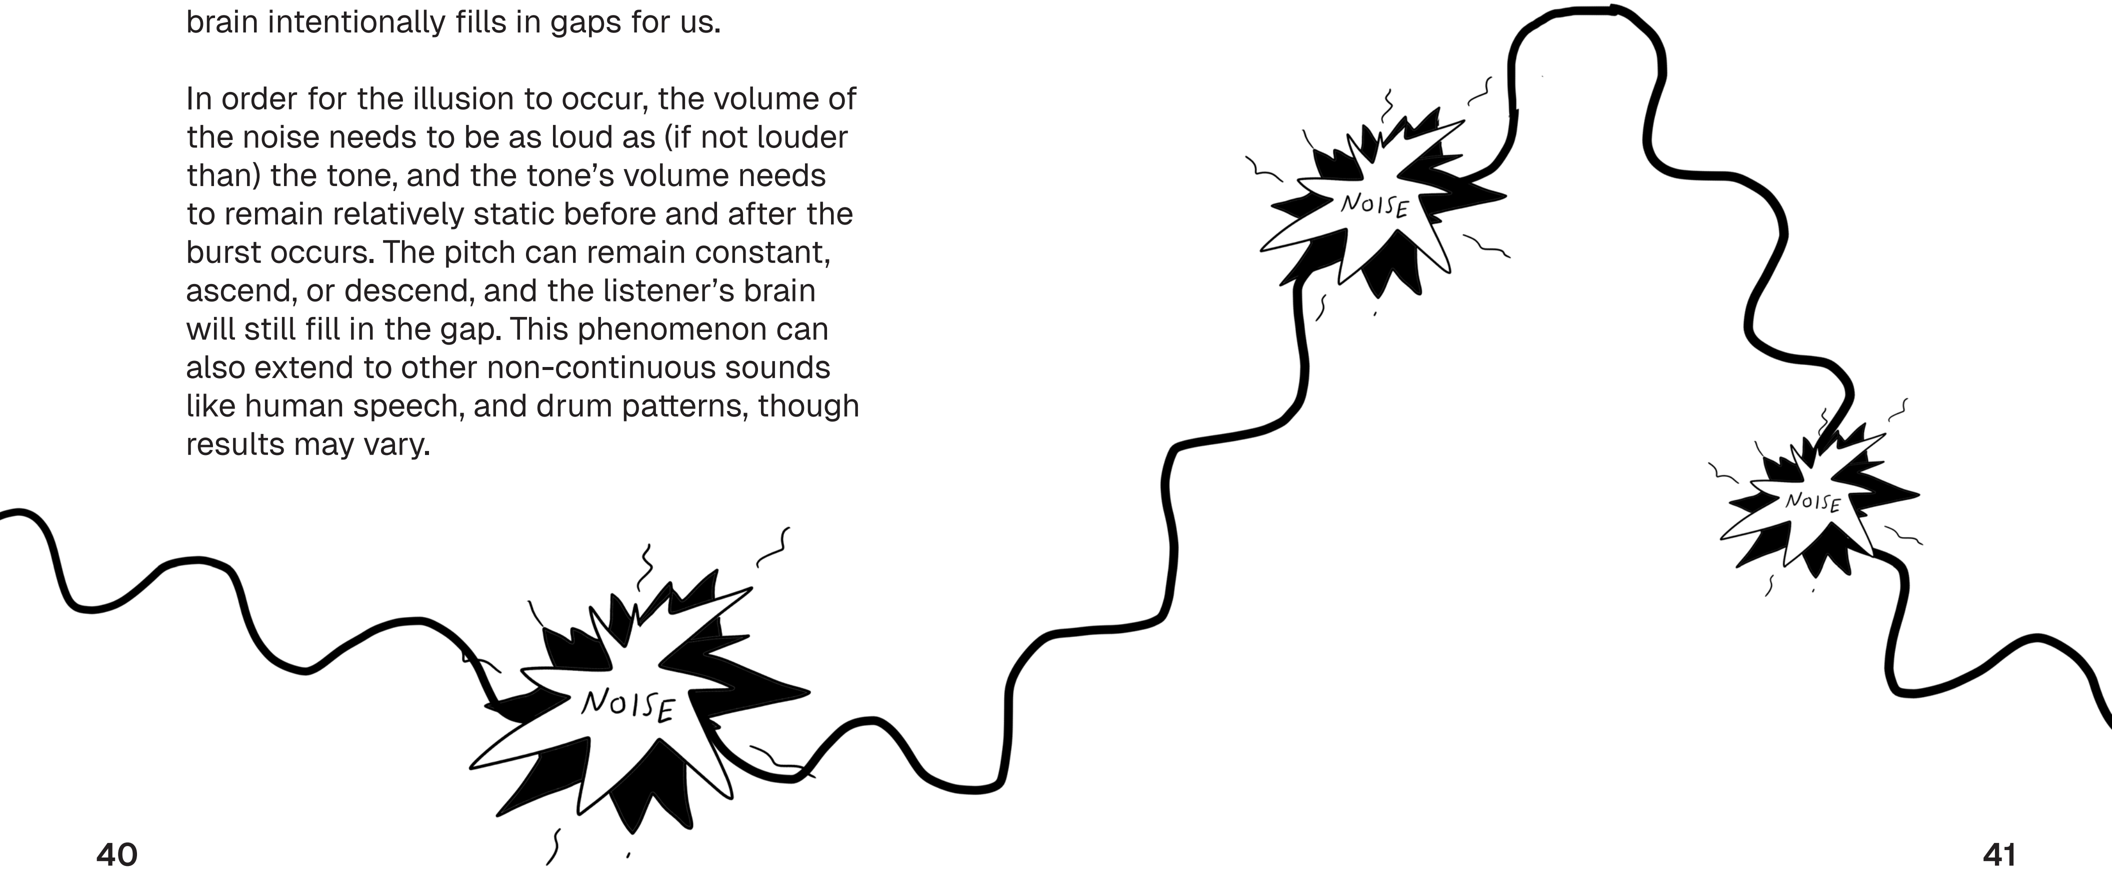

The Acoustic Startle Reflex

Sounds higher than 80db that are not anticipated can cause listeners' eyes to blink and their muscles to involuntarily tense. Similarly, sounds that listeners are not expecting are typically perceived as being louder than they actually are.[27]

How does it Work?
The Acoustic Startle Reflex exists in our brains to help us respond quickly to stimuli that may indicate we are in danger. Sounds that startle us might be perceived as louder because our brain is attributing greater importance to them. This allows us to react quicker.

Design Applications
Sound Designers can use the Startle Reflex to scare their audience.

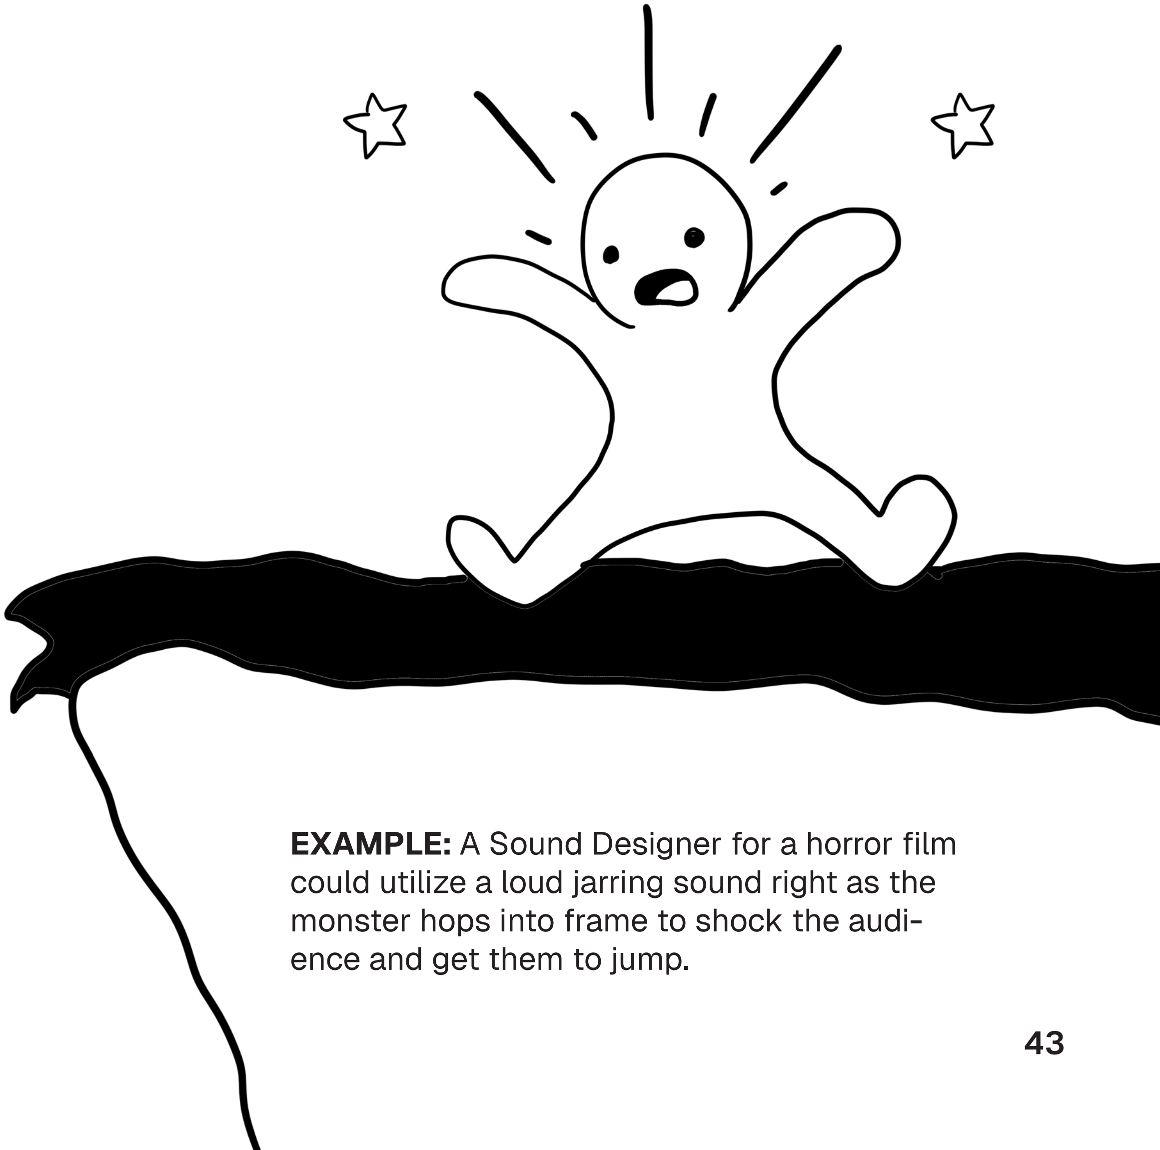

EXAMPLE: A Sound Designer for a horror film could utilize a loud jarring sound right as the monster hops into frame to shock the audience and get them to jump.

Speech-to-song Illusion

When a short spoken phrase is looped and played on repeat for an extended period of time listeners will begin to hear rhythms and pitch within the phrase and it will begin to sound like a song.[28]

How Does it Work?

Across almost every language, even in non-tonal languages like English, Irish and German, people tend to speak with fluctuating pitch and rhythm. Over the course of a sentence or conversation, the varying tone and inflections don't read musically because they are not necessarily repeating or taking place in any structured meter. But when a section is taken and repeated on a loop, the listener's brain begins to notice the repeating rhythm and tones, and they begin to hear the music in it.

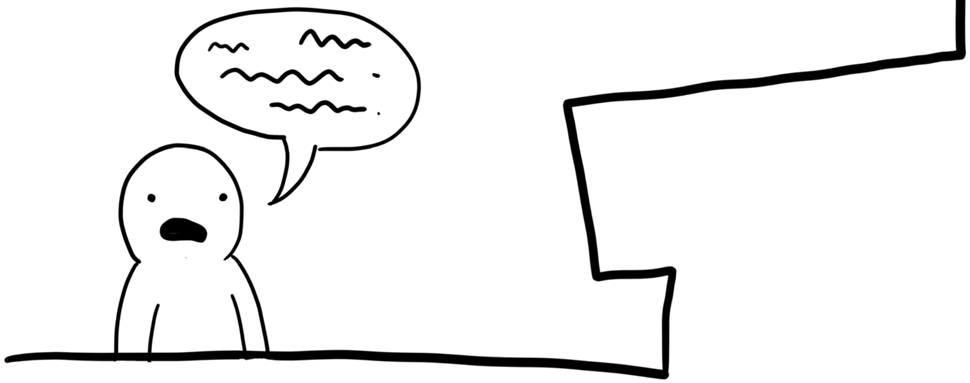

Design Applications

When composing electronic dance music, utilizing a looped vocal sample can add a fun and interesting base to work on top of.

In a live performance setting, looping the last spoken phrase of a scene in order for it to build into a transition where the looped vocals are a central part of the song can add a haunting element to the transition.

Sympathetic Resonance

When a vibratory body, like a guitar string, tuning fork, or wine glass is exposed to an external vibration that matches or harmonizes with the frequency it makes when it vibrates, the vibratory body will begin to vibrate sympathetically.[29]

How does it work?

When a vibratory body is struck it begins to-move rapidly. As the object vibrates physically, it also vibrates the air around it generating sound. If that soundwave collides with another vibratory object, specifically one that would vibrate naturally at that same frequency or a harmonic frequency, that object will absorb the sound wave and transfer it back into a physical vibration. Take, for example, a guitar string: when the string is struck, it vibrates the air around it and other guitar strings tuned to the same frequency will begin to sympathetically vibrate. This extends beyond instruments though to almost all elements of the physical world.

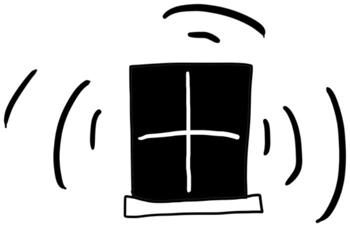

EXAMPLE: You place a subwoofer in the middle of a room filled with odds and ends. When it begins to play tones, you hear different objects in the room vibrate. Depending on how loud the subwoofer is playing you may even begin to hear the windows rattle.

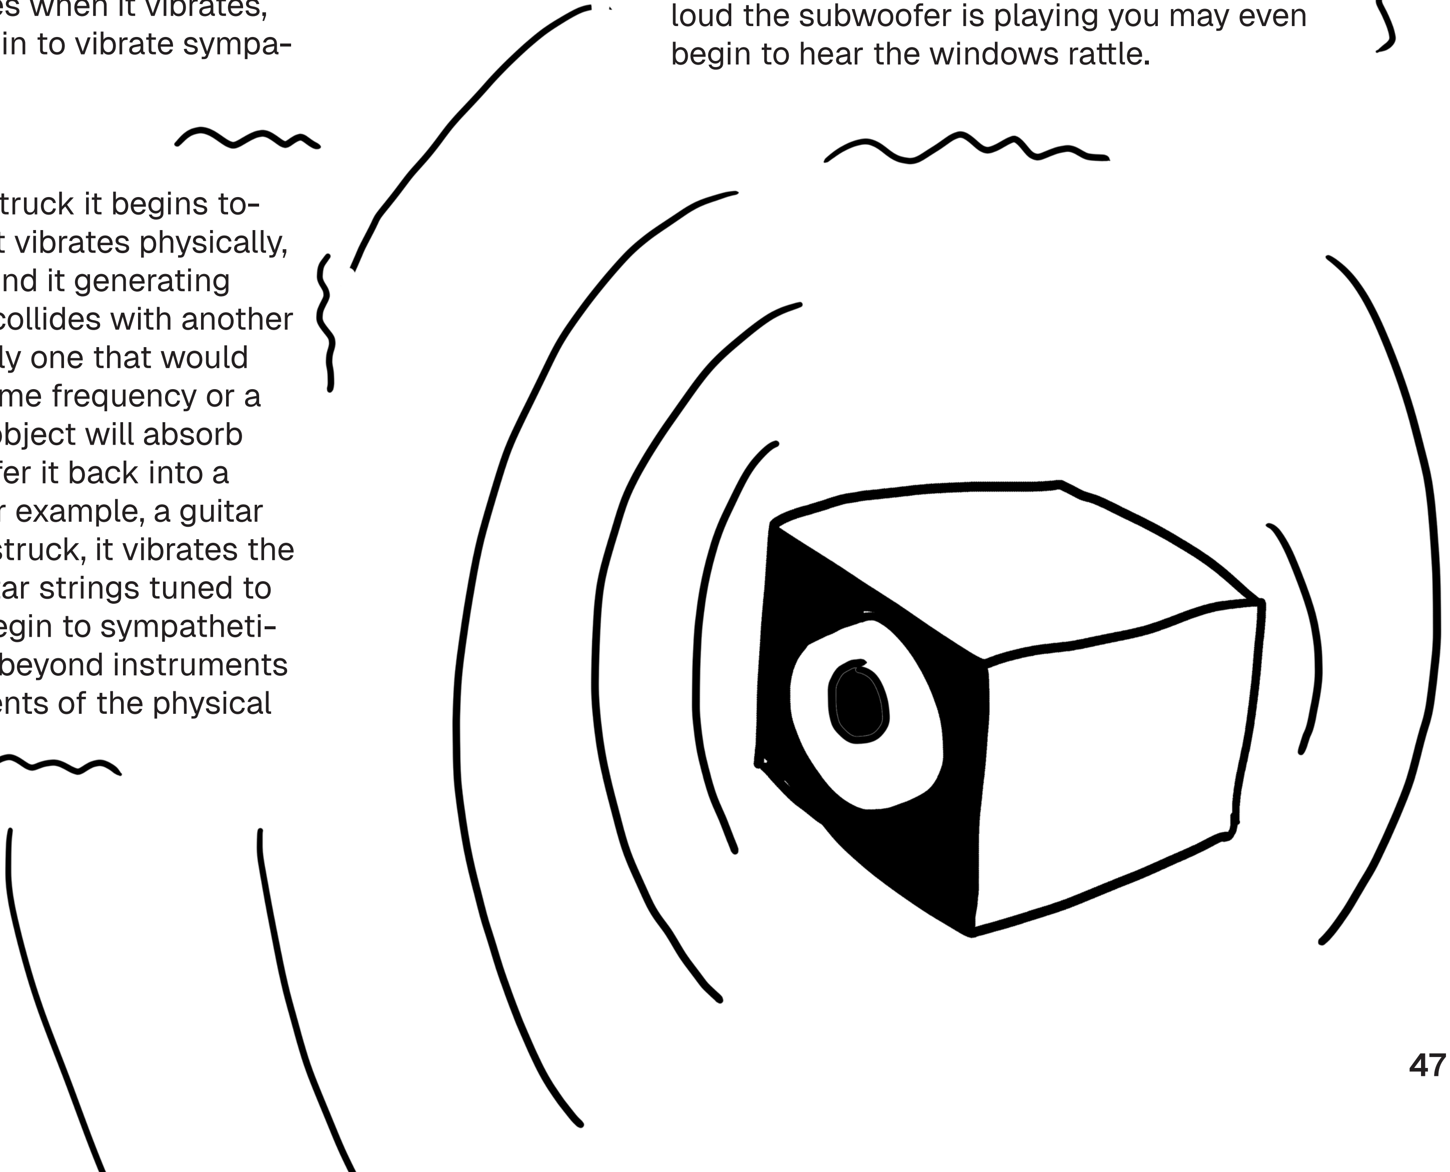

Design Applications

When tuning a sound system your subwoofers will typically let you know the sympathetic resonance of the room. As you play a very bassy song you may begin to hear parts of the room rattle, it can sometimes be very loud and annoying. If you'd like to remove this rattle you have two option:

Option 1: Use an equalizer to reduce the harmonic frequencies that cause certain parts of the room to vibrate. (You can determine which frequencies are causing the vibrations by playing different frequency sine waves through the sub until you've identified which frequencies cause the room to rattle).

Option 2: Go around the room and either remove or secure any of the objects that are vibrating so that they no longer make a noise.

Option 2 is almost always the preferred option, as Option 1 leaves you with less Bass frequencies to play with, and can affect how your content sounds.

Alternately, you may decide you want to shake the room, or better yet a specific object in the room. The resonant frequency of an object can be crudely found by striking it and seeing what sort of tone it generates. Tapping a wine glass with a fork will tell you what frequency it vibrates at, and then playing that frequency back at the wine glass will cause it to sympathetically resonate. Sufficient volume can even cause the wine glass to vibrate so intensely it will shatter.

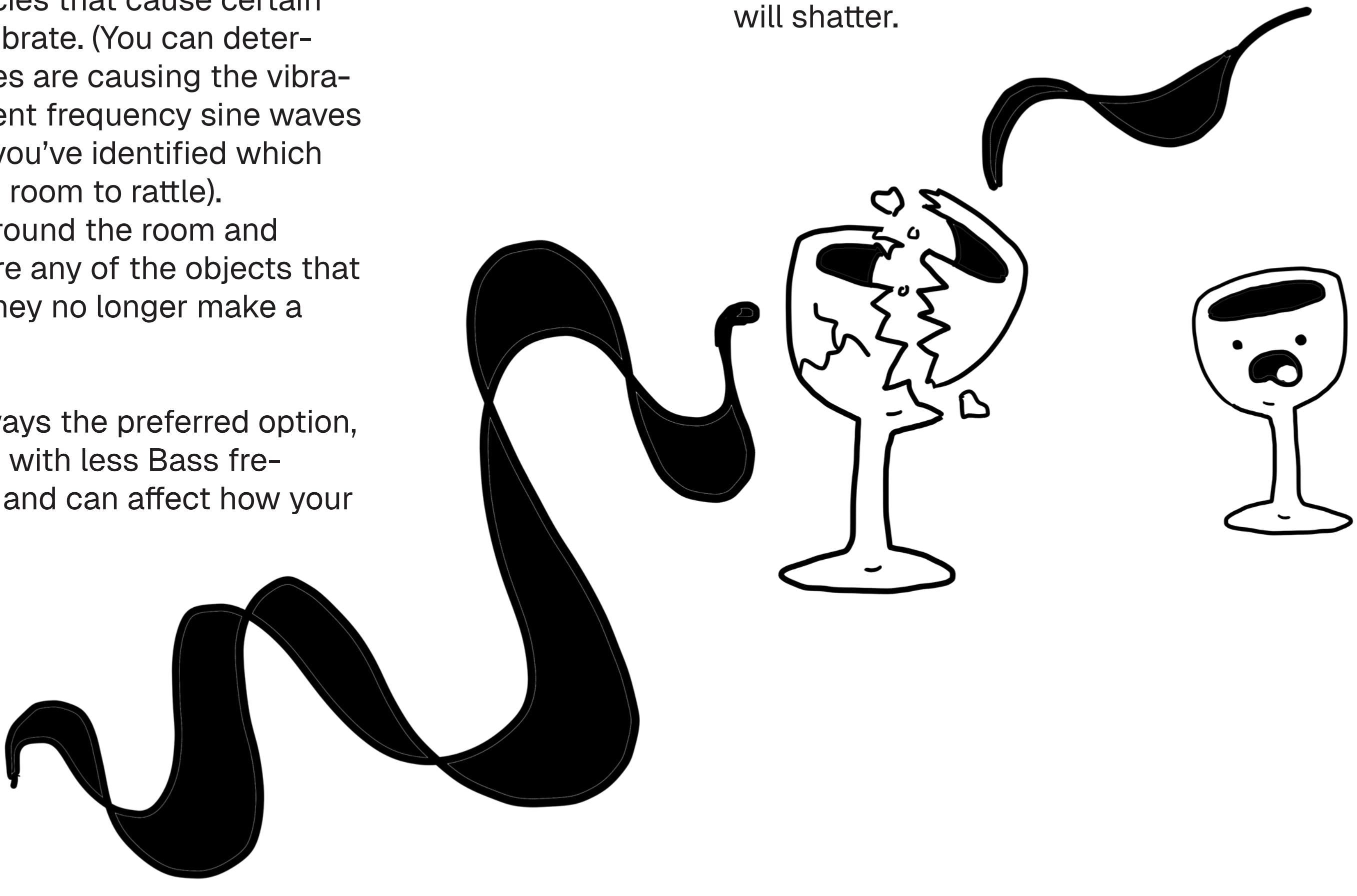

Cymatics

Similar to Sympathetic Resonance, Cymatics is a vibrational phenomena that occurs when soundwaves are absorbed and converted into physical vibrations. As sound waves (or sometimes just specific vibrational frequencies applied physically) interact with a surface -such as a metal plate or a membrane- the surface will begin to vibrate. Depending on the composition of the surface and the frequency applied, different areas of the surface will vibrate more intensely than others. A medium placed on top of the surface (like sand) will then reveal which areas are vibrating more intensely via displacement, sometimes revealing intricate patterns.

How does it work?

Functionally, Cymatics works through sympathetic vibration. As the surface vibrates, because of its composition different areas will tend to vibrate more intensely than others.[30] The artistry then comes from the selection of a surface and the selection of the medium placed on top of the surface. Common pairings are metal plates and sand as well as bowls filled with a water and cornstarch mixture.

Design Applications

Cymatics are a physical manifestation of sound waves and their design applications tend to be visual. When working in a recorded medium Cymatics can be used to create patterns and visuals that can then be matched in post production to different sounds. It's important to remember that Sympathetic Resonance is caused by sound waves being turned into a physical vibration. When being used for Cymatics, the tones typically need to be quite loud, and they typically need to be tones that are pure and steady, like a sine wave. In live settings featuring an audience the tones needed to produce certain patterns may be unpleasant for the listeners.

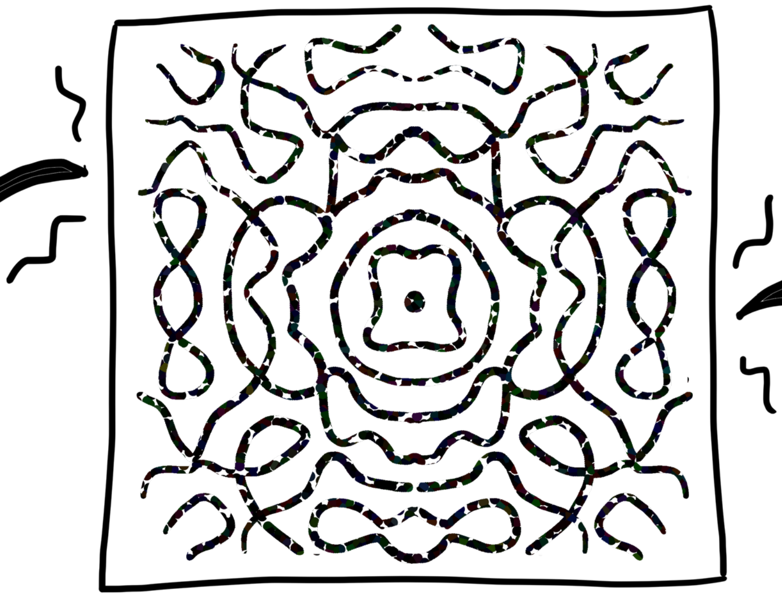

Risset's Rhythm

It's French, the T is silent

The rhythmic equivalent of the Shepard-Risset Glissando, a Risset rhythm is a rhythm that appears to be increasing or decreasing in tempo infinitely.

How Does it Work?

A Risset Rhythm applies the same concept and design of a Shepard Tone to tempo. Imagine a simple rhythm that consists of just whole notes, half notes, and quarter notes. The whole notes and quarter notes are quieter, while the half notes are the loudest. As the rhythm increases in tempo to double time the half notes become our new quarter notes and decrease in volume, and as the whole notes become our new half notes they increase in volume. Each time the slowest note goes from barely audible to being the most present note, it is replaced by a barely audible note playing as slow as it was originally, which in this example, is half time the current dominant tempo.[31]

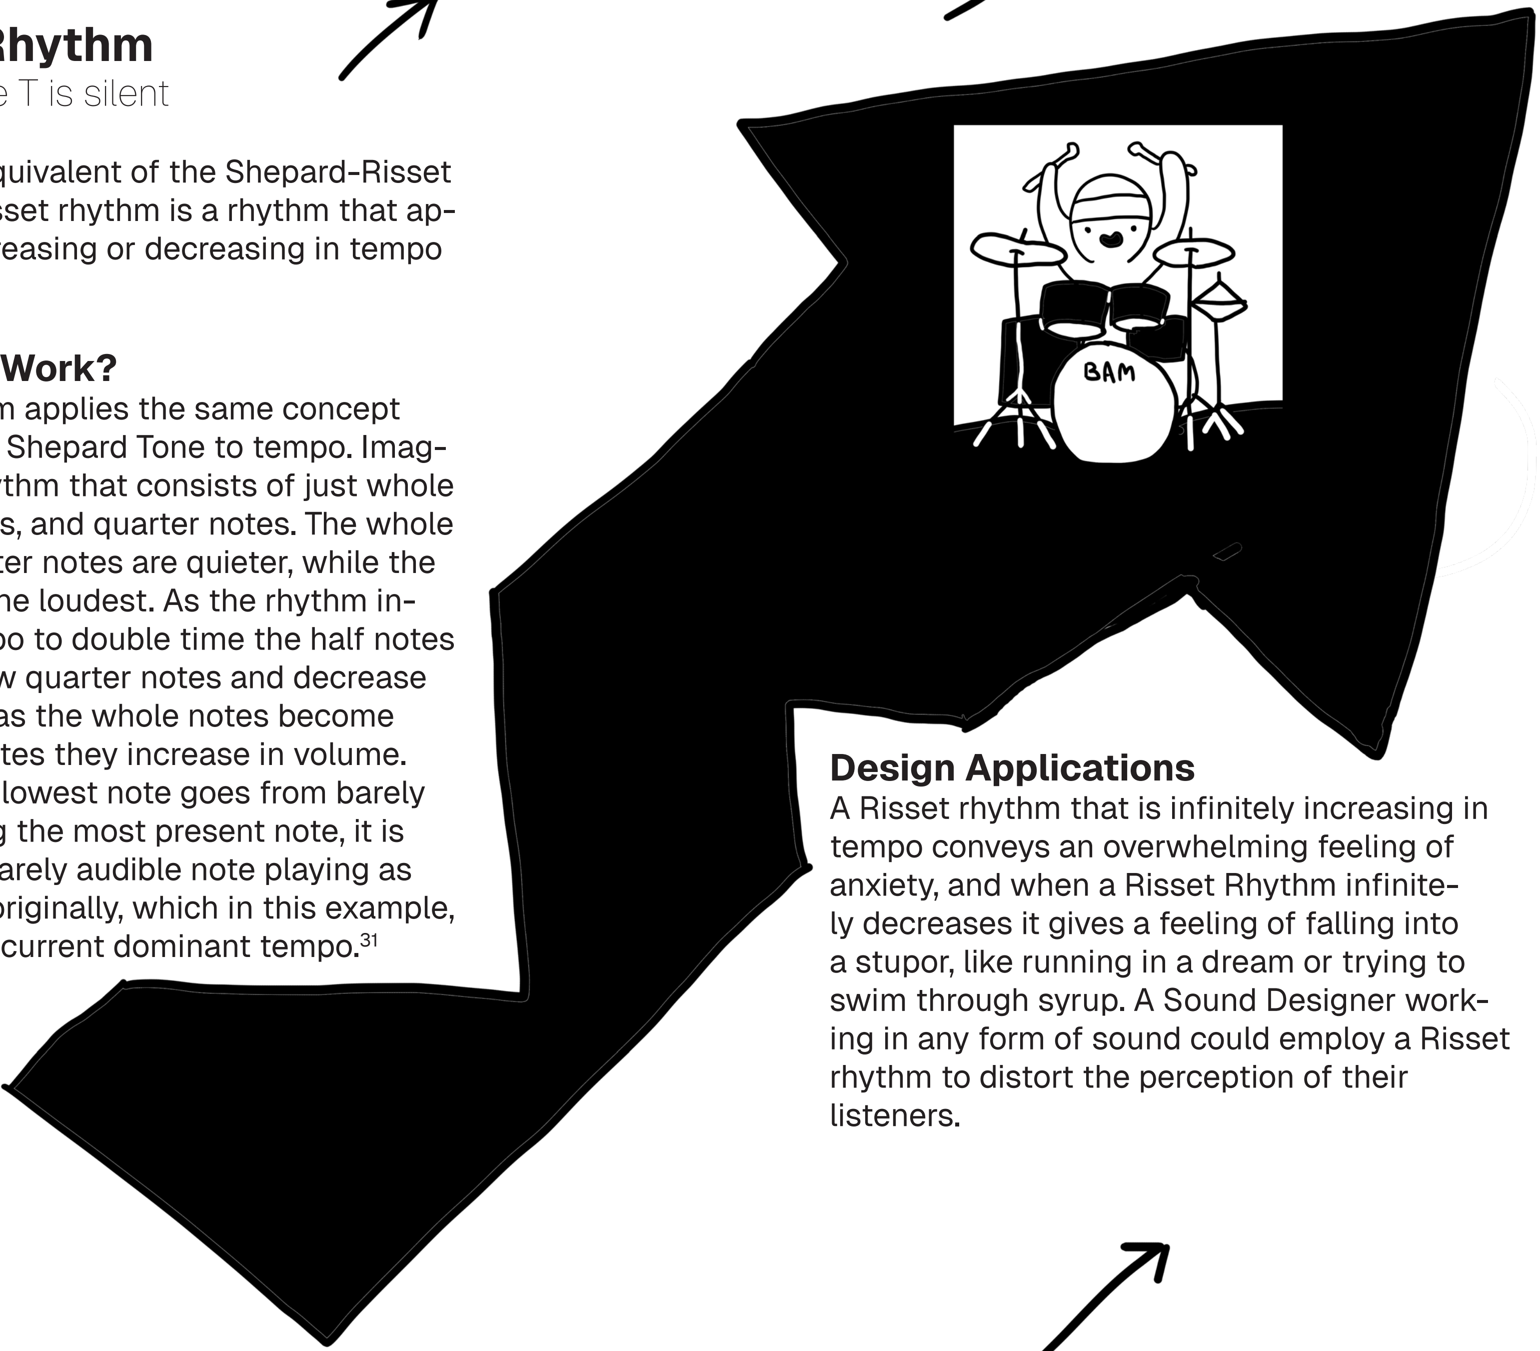

Design Applications

A Risset rhythm that is infinitely increasing in tempo conveys an overwhelming feeling of anxiety, and when a Risset Rhythm infinitely decreases it gives a feeling of falling into a stupor, like running in a dream or trying to swim through syrup. A Sound Designer working in any form of sound could employ a Risset rhythm to distort the perception of their listeners.

Ciesil's Perspective

The sound something makes is typically tied to the size of the object, whether it is in terms of loudness or frequency. Larger objects tend to make louder low frequency sounds, whereas smaller objects tend to make quieter high frequency sounds. In design, Ciesil's Perspective states that Large objects make large sounds, small objects make small sounds, but small objects can make large sounds when viewed closely or given importance.

How it works?

Admittedly this is a self insert, but its inclusion in this field guide comes from a desire to put a name to a sound design principle that is as pervasive as it is effective. Ciesil's Perspective is an acknowledgement that to design a sound for an object or experience is to design the listener's perception of that object or experience. Just like when we imagine the weight of an object when we see someone struggle to lift it, the sound that object makes when it is dropped allows us to imagine its weight. Sound as a design tool teaches us about an object, and Ciesil's Perspective acknowledges that we can depart from the reality of the world around us in order to make a sound fit the reality of the designed world.

Design Applications

Sound Designers in any medium can use Ciesil's Perspective to create sounds that do not already exist, or add importance to objects by changing their sound.

EXAMPLE: An ant walking along a tile floor makes almost no sound, but if we were sound designing a video game where players shrink down and view that ant from the perspective of a dust mite, the ant's footsteps would be thunderous.

EXAMPLE: A ring landing on a table makes a bright chiming sound, but if a Sound Designer was working on a fantasy movie and that ring happened to be bound with dark magic to cloud the hearts and minds of men, when the ring hit's the table it may make a loud booming thud.

Closing Remarks

Thank you for making it this far. I hope that the contents of this field guide helped to spark inspiration or curiousity within you.

When I used to work as the Audio Manager at the University of Chicago I would often be asked to come in and guest lecture a class on sound design as part of a section of their intro to design class. My closing remarks at each of those classes is what I'll leave you with here:

"If you can do anything else with your life, I encourage you to do it. If you're at all interested in money, fame, or stability, you will most likely not find it in sound design. But, if you find yourself obsessed with sound, to the point where you can't do anything else: we'd love to have you."

Endnotes

1 This definition is my own, but based partially on the Merriam-Webster definition: ‘sound’, The Merriam-Webster Dictionary (2025) [online], https://www.merriam-webster.com/dictionary/sound [accessed 3 May 2025].

2 Likewise, see ‘noise’, The Merriam-Webster Dictionary (2025) [online], https://www.merriam-webster.com/dictionary/noise [accessed 3 May 2025].

3 Gary D. Davis and Ralph Jones, The Sound Reinforcement Handbook, 2nd ed (Hal Leonard, 1989), 1.

4 John A. Leonard, Theatre Sound (Routledge, 2001), 17-22.

5 The speed of sound can be calculated at different temperatures using this online resource: Tim Brice and Todd Hall, ‘The Speed of Sound Calculation’, https://www.weather.gov/epz/wxcalc_speedofsound [accessed 3 May 2025].

6 Leonard, Theatre Sound, 19.

7 Davis and Jones, The Sound Reinforcement Handbook, 43.

8 Davis and Jones, The Sound Reinforcement Handbook, 1.

9 Davis and Jones, The Sound Reinforcement Handbook, 53-56.

10 Deena Kaye and James LeBrecht, Sound and Music for the Theatre, 2nd edition (Focal Press, 2000), 33.

11 Philip McCord Morse and K. Uno Ingard, Theoretical Acoustics, International Series in Pure and Applied Physics (McGraw-Hill, 1968), 449-450.

12 Morse and Ingard, Theoretical Acoustics, 708-710.

13 Helmut Haas, ‘Influence of a Single Echo on Audibility of Speech’, JOURNAL OF THE AUDIO ENGINEERING SOCIETY, 20.2 (1972), pp. 146-159.

14 Jens Blauert, Spatial Hearing: The Psychophysics of Human Sound Localization, Rev. ed. (MIT Press, 1997), 203-222.

15 As first described in Christian Doppler’s treatise, ‘Über das farbige Licht der Dopppelsterne und einiger anderer Gestirne des Himmels’. See the text anx translation in Alec Eden, The Search for Christian Doppler (Springer Vienna : Imprint: Springer, 1992), 93-104.

16 Victoria Deiorio, The Art of Theatrical Sound Design: A Practical Guide (Methuen Drama, 2019), 26.

17 Leonard, Theatre Sound, 22.

18 Roger N. Shepard, ‘Circularity in Judgments of Relative Pitch’, The Journal of the Acoustical Society of America, 36.12 (1964), pp. 2346–53.

19 Jean-Claude Risset, ‘Pitch and Rhythm Paradoxes: Comments on “‘Auditory Paradox Based on Fractal Waveform’” [J. Acoust. Soc. Am. 79, 186–189 (1986)]’, The Journal of the Acoustical Society of America, 80.3 (1986), pp. 961–62.

20 Daniel J. Levitin, ‘Absolute Memory for Musical Pitch: Evidence from the Production of Learned Melodies’, Perception & Psychophysics, 56.4 (1994), pp. 414–23.

21 E. Colin Cherry, ‘Some Experiments on the Recognition of Speech, with One and with Two Ears’, The Journal of the Acoustical Society of America, 25.5 (1953), pp. 975–79. For a review of recent research that considers the cocktail party phenomenon, see Simon Haykin and Zhe Chen, ‘The Cocktail Party Problem’, Neural Computation, 17.9 (2005), pp. 1875–902.

22 Adelbert W. Bronkhorst, ‘The Cocktail Party Phenomenon: A Review of Research on Speech Intelligibility in Multiple-Talker Conditions’, Acustica United with Acta Acustica, 86.1 (2000), pp. 117–28.

23 Leonard, Theatre Sound, 47.

24 See Leonard, Theatre Sound, 47; Davis and Jones, The Sound Reinforcement Handbook, 120; Shure Service and Repair, 'Why Does Proximity Effect Occur?', https://service.shure.com/s/article/why-does-proximity-effect-occur?language=en_US&region=en-US, [accessed 3 May 2025].

25 Lars Riecke, A. John van Opstal, and Elia Formisano, 'The Auditory Continuity Illusion: A Parametric Investigation and Filter Model', Perception & Psychophysics, 70.1 (2008), pp. 1–12.

26 Thanks to Minjae Kim for providing this example.

27 Peter J. Lang, Margaret M. Bradley, and Bruce N. Cuthbert, 'Emotion, Attention, and the Startle Reflex', Psychological Review, 97.3 (1990), pp. 377–95.

28 Diana Deutsch, liner notes to Phantom Words and Other Curiosities (2003, Philomel Records), 16-17.

29 Hermann von Helmholtz and Alexander John Ellis, On the Sensations of Tone as a Physiological Basis for the Theory of Music, 5th ed (Longmans, Green, 1930), 36.

30 Hans Jenny, Cymatics: Wave Phenomena, Vibrational Effects, Harmonic Oscillations, with Their Structure, Kinetics and Dynamics (Basilius Presse, 1974).

31 Risset, 'Pitch and Rhythm Paradoxes'.

Bibliography

Blauert, Jens, Spatial Hearing: The Psychophysics of Human Sound Localization, Rev. ed. (MIT Press, 1997)

Cherry, E. Colin, 'Some Experiments on the Recog nition of Speech, with One and with Two Ears', The Journal of the Acoustical Society of America, 25.5 (1953), pp. 975–79,

Davis, Gary D., and Ralph Jones, The Sound Reinforcement Handbook, 2nd ed (Hal Leonard, 1989)

Deiorio, Victoria, The Art of Theatrical Sound Design: A Practical Guide (Methuen Drama, 2019)

Deutsch, Diane, Phantom Words and Other Curiosities (2003, Philomel Records)

Eden, Alec, The Search for Christian Doppler (Springer Vienna : Imprint: Springer, 1992)

Haas, H., 'Influence of a Single Echo on Audibility of Speech', Journal of the Audio Engineering Society, 20.2 (1972), pp. 146–59

Haykin, Simon, and Zhe Chen, 'The Cocktail Party Problem', Neural Computation, 17.9 (2005), pp. 1875–902

Helmholtz, Hermann von, and Alexander John Ellis, On the Sensations of Tone as a Physiological Basis for the Theory of Music, 5th ed (Longmans, Green, 1930)

Jenny, Hans, Cymatics: Wave Phenomena, Vibrational Effects, Harmonic Oscillations, with Their Structure, Kinetics and Dynamics (Basilius Presse, 1974)

Kaye, Deena, and James LeBrecht, Sound and Music for the Theatre, 2nd edition (Focal Press, 2000)

Lang, Peter J., Margaret M. Bradley, and Bruce N. Cuthbert, 'Emotion, Attention, and the Startle Reflex', Psychological Review, 97.3 (1990), pp. 377–95

Leonard, John A., Theatre Sound (Routledge, 2001)

Levitin, Daniel J., 'Absolute Memory for Musical Pitch: Evidence from the Production of Learned Melodies', Perception & Psychophysics, 56.4 (1994), pp. 414–23

Morse, Philip McCord, and K. Uno Ingard, Theoretical Acoustics, International Series in Pure and Applied Physics (McGraw-Hill, 1968)

Riecke, Lars, A. John van Opstal, and Elia Formisano, 'The Auditory Continuity Illusion: A Parametric Investigation and Filter Model',

Risset, Jeane Claude, 'Pitch and Rhythm Paradoxes: Comments on "'Auditory Paradox Based on Fractal Waveform'" [J. Acoust. Soc. Am. 79, 186–189 (1986)]', The Journal of the Acoustical Society of America, 80.3 (1986), pp. 961–62

Shepard, Roger N., 'Circularity in Judgments of Relative Pitch', The Journal of the Acoustical Society of America, 36.12 (1964), pp. 2346–53