Places in the Note

The Haas Effect
(the Precedence Effect)
in Rehearsal, Performance, Sound Reinforcement and Mixing

How how mere milliseconds can
enhance and clarify a music performance.

How timing and delay can affect
our sense of sound location and depth.

by John Blasquez

The Haas Effect

Leading and following in Musicianship

I'm about to describe some simple but rarely perceived temporal truths. We'll be talking timing, timing, timing.

In private and group music lessons I devote a considerable amount of time explaining the concept of "leading and following." The topic arrises in lessons where I accompany or play with students, in group rehearsals, and in music recording mixing sessions.

When musicians really connect, and when they create their richest performances and recordings, they do not play their notes exactly at the same time.

On hearing this many people think, "What?! Don't we all painstakingly strive for exactitude in musical timing?" Yes, indeed, timing is of utmost importance. But that doesn't mean that our notes should all occur simultaneously, exactly on downbeats and upbeats.

Notes played exactly at the precise same moment result in a temporal masking. And the overall effect is a lack of clarity and timbral brilliance. So the great players add a bit of space between instruments, ordering them, so that featured instruments and/or delicate instruments (or voices) lead slightly ahead.

A subtle temporal spread greatly enhances and clarifies the presence of each instrument, much like the use of a panned "stereo image" in music recording—a good stereo spread offers much greater clarity than a mono mix. Generally speaking, the melody leads slightly ahead of the accompaniment.

The leading instrument naturally shines in detail, the listener relaxes as they needn't strain to hear because the lead instrument is easily heard. An added benefit: the performer sounds more confident, just from "stepping in front".

I was once told that Charles Mingus used the phrase "places in the note." I've never found a reference to this quote, but it certainly suggests that there are times to lead, and times to follow, and times to place your notes "between" various instruments.

Indeed we need highly refined and steady timing to achieve the goals I'll propose, but we're definitely overriding the simplistic concept that all notes should land "bang on the beat."

The Haas Effect (aka the Precedence Effect) describes these temporal phenomena.

NOTE: When you listen to examples of the Haas Effect, for instance in tutorials on youTube, or if you're the adventurous kind and plan to construct your own experiments, be sure to listen over headphones or good quality stereo speakers. Here's a commendable overview of the Haas Effect.

Ideal instrument ordering

To my ear, the instrument ordering that promotes the greatest clarity is:

- melody first (if shared, delicate instruments go first)
- harmony second
- then chords
- followed by bass
- and is percussion backmost (in most styles)

... although some styles and artists prefer the percussion or bass in the lead.

Chords, bass and percussion are generally referred to as "back up." And there you have it— backup is in back. The temporal position is expressed in the label. However the backside placement is slight enough that it never sounds like "dragging" but actually adds excitement and drive to the blend.

Soloist timing

There's a type of leading that I call "soloist timing." When and artist steps ahead into soloist timing—managing a proper amount of lead—it doesn't sound like rushing. The main effect is that the artist's instrument is better heard. It sounds louder, you can hear all the detail of the note onset, and you get a richer quality of tone. And if your fellow musicians understand that you're intentionally surfing the leading edge, they won't close the gap.

Rushing and Dragging

Well timed leading and following will not equate to acceleration ... *if* the other members allow you the space. When a soloist leads excessively the sound becomes strident and pushy. When a soloist follows behind it sounds like they're dragging. A single rushing musician can cause a performance, rehearsal, or jam to speed up ... as the other musicians try to close the gap. Similarly a single dragging musician can cause a performance to slow down as the other musician try to close the gap.

Delicate instruments

Delicate instruments are those with low volume, short sustain (short note durations), lower pitch, and simple tone (few harmonics.) Consider staying behind them.

Hard concepts to grasp

It may seem absurd that we'll be attempting to time the distances between instruments to within roughly 40 to 60 milliseconds. Like, what's a millisecond? And how am I supposed to hear or measure that?? Truthfully people hear and respond at this level ... with surprising precision. The trick is doing so consciously, intentionally, and to produce the desired effect. (Beginners have a uncanny tendency to play behind, to hide out, and it works—but that's only a good place if it enhances the music.)

Places in the Beat

So what's the reason for this exploration and information?

To illuminate the simple truth that people with unimpaired hearing her and unconsciously respond to a surprising level of auditory precision.

And to emphasize the significance of the concept of leading and following in music. And to explain this in quantifiable, technical terms. To show that we may indeed want to aim for various placed "in the beat."

Digital Control vs. Analog

Matters of timing and individual note volume became apparent,

After I initially began working with MIDI (via MIDI recording software, commonly called a sequencer or Digital Audio Workstation (DAW)—I immediately had precise editing control over note volume and timing at a level never realistically and consistently possible in the analog world of tape recording.

Getting the Rough cut

Prior to any editing I had to capture a performance. The piano keyboard is the most common "MIDI controller" (input device) and indeed I've always used a piano keyboard as my MIDI controller.

My instrumental expertise is with stringed instruments and my piano skills definitely lag *far* behind, so my piano performances generally need significant clean up.

When I initially recorded via the piano keyboard the result was good—but sometimes it took multiple passes (and/or punch-ins) to get a reasonable "performance." And even then, not unexpectedly, it was marred by many small timing and note volume issues.

> NOTE: In MIDI applications note volume is called velocity (as it is essentially equivalent to the force with which the key is struck ... and the term volume was already used in other areas of MIDI. ***)

Quantization

I used the application's Quantize function to polish the captured performance and "snap notes onto the grid." Fortunately the note onsets in my edited performance were close enough to the nearest eighth note or sixteenth, so this worked magically.

When I first used 100% quantization (i.e. all note onsets moved *exactly* to the start of the designated beat or sub-beat. I expected the result to sound robotic, rigid or lacking warmth. But to my surprise, a monophonic melody like a Bach Partita for Solo Violin sounded darn good, like the precision we strive for. And I also found that quantizing 96 to 97% gave a slightly more pleasant result, a little more natural with the slightest ebb and flow.

Quantizing polyphonic music

Light quantization produced good results on monophonic music. So what about quantizing polyphonic music?

First I tried a guitar piece written in two voices. Then piano in three voices. Again quantization greatly improved the overall timing in my humble keyboard-entered recording.

But now there was something noticeably disturbing about the playback. It sounded heavy. And the bass notes seemed predominant even if I lowered the "volume" of the bass notes beyond a desirable threshold.

Next I tried a multitrack project with piano, clarinet, bass and drums. Again, when quantized to 97% the individual instrument timing was improved on each track, and sounded great when soloed. But the overall mix of instruments was muddied when I played all tracks together.

The mystery slowly unraveled as I remembered how I would ask percussionists to layback in rehearsals and performances. So I moved the onset of all percussion slightly behind by two to four ticks, and voila. Then I set the piano and bass in between the leading clarinet and the drums. Now it the quartet sounded lighter, clearer and far more realistic. And the drums were really "in the pocket."

> NOTE: There are styles like (Cuban music and Disco) where the drums may lead so they're focused for maximum "beat." There are good example of Cuban drumming on Gloria Estefan's Mi Terra albumBut this is rarely the effect I'm after when mixing.)
>
> In contrast, listen to Gene Kruppa on a good day with Benny Goodman or other big bands. He has an uncanny ability to strongly drive the band while simultaneously "disappearing—and his powerful but feather light presence is not achieved by playing softer, but simply by following, slightly behind.
>
> Other good examples:
>
> - the performance by Bone Pony's drummer on "Blue Blue Blue"—it's hard to believe there's a full set of traps playing the whole way with this talented an acoustic garage band. That's usually a recipe for disaster!
> - And on a recording of Scatter the Mud -- Eileen Eivers fiddles behind the pipes, and so you mainly hear to pipe's timbre and little fleck of the fiddle's attacks and the sustains of ringing open strings, but otherwise not much of the fiddle's tone is apparent, because it let's the fiddle lead.
> - The interplay of Joni Mitchell's dulcimer and James Taylor's guitar starting at 0:523 on "A Case of You," where the leading position gently drifts back and forth between the accompaniment instruments.

There's much more to say and many more examples to offer regarding quantization and leading/following.

Anecdotally, I was teaching leading/following at a music camp. Daily we would sit in a group for 50 minutes, all members of the group playing the melody at once except for a couple playing light accompaniment.

The assignment was to allow a single person to take the lead by edging their timing forward *without* increasing their volume (for the duration of the full AABB structure of the melody).

I rarely had to throw a flag on the play. And after the six of seven members got a couple of tries "out front" the inevitable conclusion of the group members was that the soloist playing stood out, was more toneful, and sounded louder as they simply took the lead.

We did the same thing again and once in front the soloist experimented with lowering their volume. Again their presence and lead was heard even when they played softly.

Then we practiced the opposite: laying back and playing louder to make our instrument heard; the effect was insignificant resulted led to overplaying of instruments and forced tone without much gain in presence or clarity.

After dinner that evening a student approached me. He said a person across the table asked what workshops he was attending. My student described our experiments and the apparent result. The person across the table turned out to be none other than Dr. Gareth Loy. He smiled and said, of course, the Haas Effect. And explained it to my student. a talented copy writer and editor, who then enlightened me.

After camp I engaged in a bit of research it became apparent the Haas Effect was indeed the quantifiable phenomenon.

I

The effect of "Volume" on musicality (MIDI velocity)

I previously explained how increasing the precision of note onset timing (via MIDI quantization) improved the overall sound of my humble performances. After completing a couple of satisfying quantization experiments I turned my attention to MIDI velocity —"velocity" is MIDI parlance for what we normally think of as note volume.

I was excited about this because, due to my piano skill level, my biggest and most frequent errors were not timing, but notes played too loud or too soft. So I was eager to correct those flaws.

All MIDI Velocities = x

In my first experiment I set all notes to one specific value. The result was horrid. It sounded like the relentless banging of a piano student not wanting to practice.

And so it was immediately apparent. Smoothness in music is not due to a consistent amount of note volume, but rather due to fluctuations in volume.

Later experiments indeed proved that simple accent patterns help tremendously—and those patterns vary with the style of music.)

Human variation or "randomness" further improved the sound within the framework a particular accent pattern,. But when applied to a musical phrase or score where all velocities = x accent patterns can bring the piece a much improved and more lyrical sound.

Accent Patterns

There is much to say about accent pattens and how to edit them in a MIDI application.

Suffice it to say:

- In classical music there's often the pattern in 4/4 time of accenting beat 1 and 3. Often emphasizing beat 1 over beat 3. Notes falling on the "ands" between the beat are generally lighter, as are most short notes. Triplets are tricky.
- In Jazz, Swing, Blues and most other styles of popular music the accents tend to fall on beat 3 and 4. Often emphasizing beat 4 over beat 2. Notes falling on the "ands" between the beat are generally lighter, as are most short notes.
- There are many exceptions!

Again, there is MUCH more to say about accent patterns, but I'll move on to other interesting affects related to the Haas Effect.

Aural illusions

We are going to look at some very unusual and counterintuitive properties of sound ... or more precisely, the way that our brain interprets sound waves that contact our listening hardware (ears) and audio neuro-circuitry.

Some of these properties and effects clearly fall into the realm of aural illusions. To gain some perspective, and for analogy, let's look at some common optical illusions.

Optical Illusions

When most people look at the following objects they think the lower one is larger:

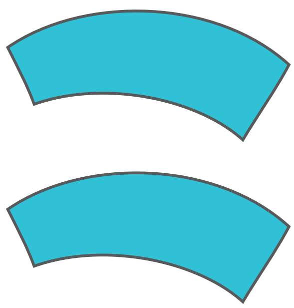

But as we can see here, thanks to the outline, they are exactly the same size:

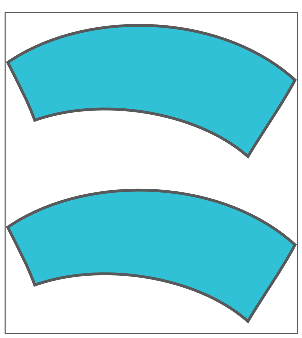

In this classic illusion the bottom line looks shorter than the top one:

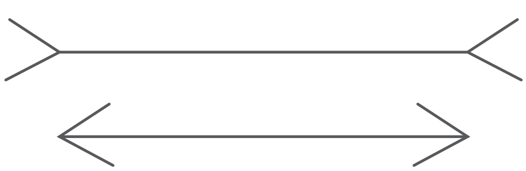

And here our eye is convinced that the center circle on the right is larger than the center circle on the left.

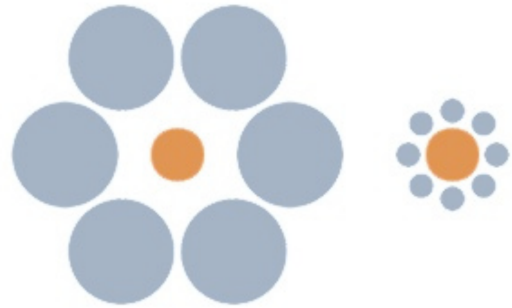

And in the following image we're absolutely certain that the only straight lines are the top and bottom ones:

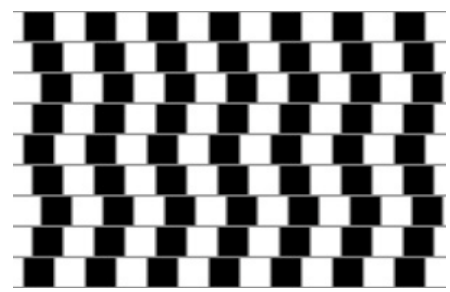

And in this dizzyingly image it's difficult to perceive that there are no spirals or overlaps, just four circles:

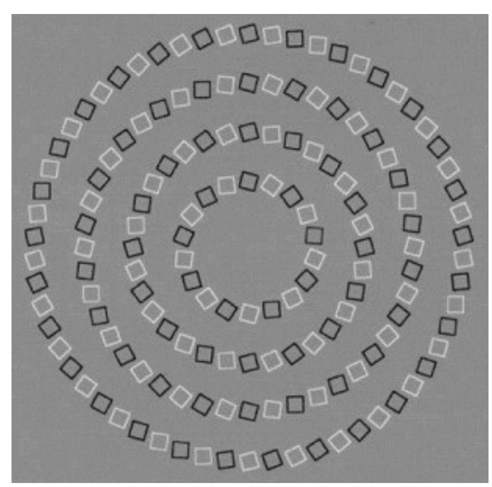

We can measure and verify that reality differs from what we "see." Clearly, our senses cannot be trusted 100%. This is true in the fields of audio and music. So let's look at some very useful aural illusions.

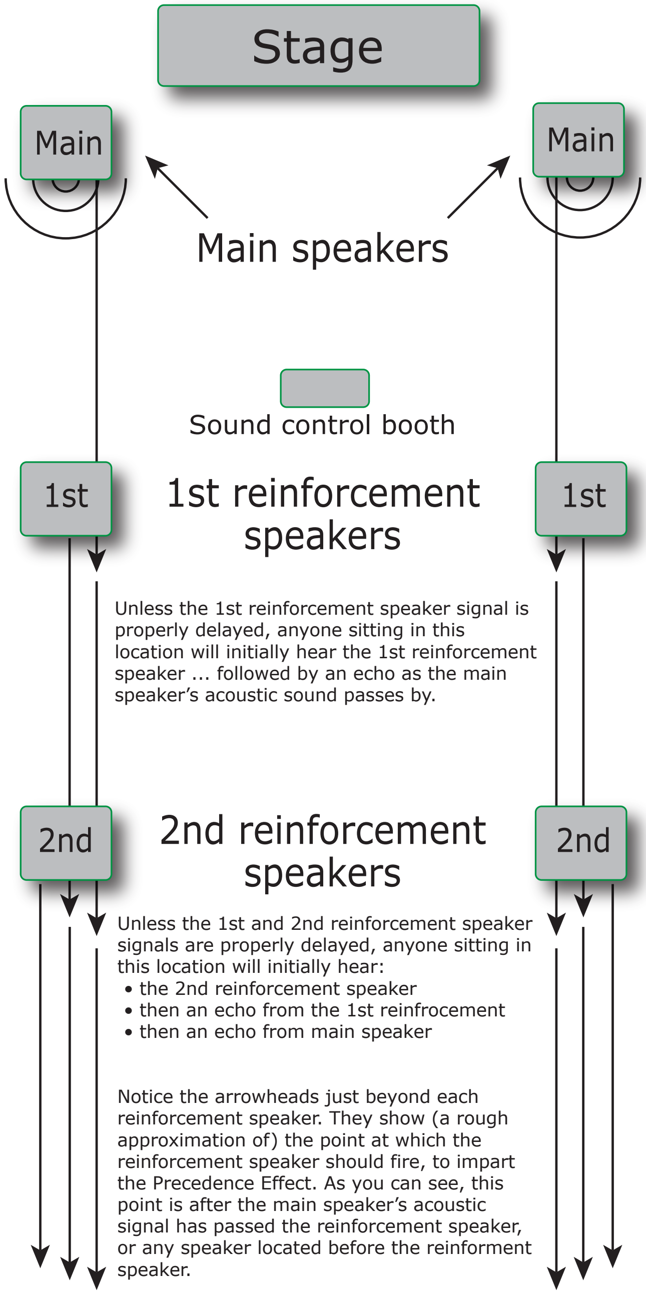

The following discussion illustrates a common sound amplification scenario. I present it here to underscore the significance of the Precedence Effect, and that it has a large impact on our perceptions. Minute and seemingly untenable timings can make the difference between clear, full, open sound vs. a sound cloudy and crowded. Timing can clarify the presence of a particular instrument or voice by intentionally focusing the listener's ear on it.

Managing sound volume in large settings

Large venues want amplified sound to reach everyone "in the house" roughly at the same volume level. A typical amplification system includes two main speakers positioned between the stage and audience—we'll refer to these as main speakers, and the signal sent to the as the main signal.

Setting the main speaker volume to suit those far from the stage may generate volume too loud for those listening near the stage. To distribute sound as evenly as possible, sound engineers often add one or more sets of reinforcement speakers, as shown here. But it's not as simple as plugging in more speakers.*

Synchronizing with Delay

Because sound travels rather slowly, there are some real factors to address when using reinforcement speakers. If the main speakers and reinforcement speakers emit the performance signal simultaneously, anyone hearing a reinforcement speaker will also hear an echo as the main speaker's acoustic signal passes by.

To eliminate the echo sound engineers delay the performance signal—so it emanates from the reinforcement speakers as the main speaker's acoustic sound reaches them.

But there's a bit more to managing good sound. Next we'll look at an interesting technique that leverages the way the brain generates perceptions ... based on the timing of sound waves reaching the ears.

The Precedence Effect & Haas Effect

It feels more natural for the listener when he or she focuses their attention toward the performance ... rather than consciously noticing that the sound is coming from nearby reinforcement speakers.

To create the illusion that the amplified sound emanates from the main speakers, sound engineers employ the Precedence Effect, which is related to the Haas Effect. They delay the signal sent to the reinforcement speakers an additional 2/100 to 10/100 of a second, so the main speaker sound reaches the ear before the sound from the reinforcement speakers. This creates a psychoacoustic aural illusion. The "ear" naturally focuses on the sound it hears first, so it "listens" deep ... perceiving the sound source as located toward the stage and main speakers, and thus focuses the listeners attention to the stage.

Echoless repeated sound fused together!

The Hass Effect states that if two or more same or similar sounds occur within a tiny window of time, the nervous system fuses them into a perception of a single sound. This fusing occurs because the auditory nervous system has a "threshold frame rate." also referred to as the listener's "echo threshold." Sounds become fused when they occur within 6 to 100 milliseconds of each other. However fusing depends on the complexity of the sound. Simple sounds like clicks fuse at a very short distances, and become clearly heard as individual sounds even at 6 to 12 milliseconds. Voice and music are more likely to "unfuse" between 40 and 100 milliseconds. (A milliseconds is 1/1000th of a second. So 40 milliseconds is 4/100ths of a second.)

As explained, when sounds occur below the listener's "echo threshold" no echo is heard, because the sounds are perceived as fused. However there are significant fused effects that occur below the echo threshold. And the following is of primary significance.

The first emitted sound will "take precedence" and be the most prominently perceived sound,—yet we hear a single fused sound ... and no echo. (Some phasing and tone coloring may occur, but that is not of particular relevance to this discussion.)

We have familiar and obvious visual analogies for this sort of phenomenon. At approximately 30 frames per second movies, video, and animation fool "the eye" into perceiving smooth continuous motion. When the visual input exceeds the frame rate of the visual nervous system, we perceive something quite different from the reality of the visual input.

And for another example, generally we don't see the raster refreshing on computer displays and TVs. it happens too quickly. Of particular interest old analog TVs and computer displays used interlaced video. This meant they showed only half the image at a time, quickly flashing alternating lines; equivalent to showing every other line of this paragraph for a slit second.

The Speeds of Sound

To set the delay properly we need to know the speed of sound. The speed of sound is very slow when compared to the speed of light. Unlike the speed of light, the speed of sound it is not a constant, and it's better measured in feet per second than miles per hour. It's speed varies dramatically depending on the medium in which it travels, such as air, various gases, water, other fluids, and solid materials like metals, rock or glass.

When no specification is expressed, the "speed of sound" refers to the speed of sound "in air."

The air speed of sound is affected by altitude, heat, humidity, air pressure and medium density. These factors must be accounted for when calculating delays for reinforcement speakers ... however, as I understand it, air pressure and medium density roughly cancel each other.

Altitude has considerable bearing, because in nature we find lower temperature and lower air pressure at higher altitudes, and because increased altitude ultimately lowers humidity.

The speeds of sound in feet per second:

- 1115 ft./second at sea level
- 1050 ft./second at 5000 feet

Calculating the Delay

Say we're at sea level, and the distance between the main speaker and the 1st reinforcement speaker is 60 yards (180 feet.) We divide 180 feet by 1115 feet per second. So it takes .16 seconds for the sound of the main speaker to reach the 1st reinforcement speaker.

Adding the Precedence Effect

Add a .01 to .10 delay (i.e. by setting your total delay between .17 and .26 seconds) and you've got your depth-producing, attention-focusing Precedence Effect. In other words, the sound appears to be originating from the stage area.

References

Precedence Effect:

http://en.wikipedia.org/wiki/Precedence_effect

Speed of Sound:

http://www.grc.nasa.gov/WWW/k-12/Missions/Jim/Project1ans.htm

http://en.wikipedia.org/wiki/Speed_of_sound

Haas Effect google search:

Haas Effect in sound reinforcement

* **NOTE**: There additional benefits to using reinforcement speakers. A loss of clarity occurs as sound travels through long distances of air, through random sound wave reflections and standing waves. By delivering a fresh signal near the listener, reinforcement speakers partially resolve these degradations.

Changing the perceived location of a sound with delays

To follow this discussion you'll need to understand the concept of pan. Audio circuitry imitates binaural hearing by using a separate channel for each of our ears. During playback that signal is sent to left and right speakers, such as stereophonic headphones.

Audio mixers commonly have 8 or 16 separate channels of mono or stereo signals. Each channel strip has a pan knob that determines the amount of the signal sent to the left and right stereo channels.

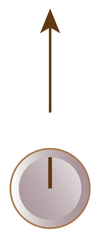

Center pan

When a mono audio signal is “center panned” the signal is sent to both ears (to the right and left stereo channels) at equal volumes. As a result we perceive as “straight ahead”

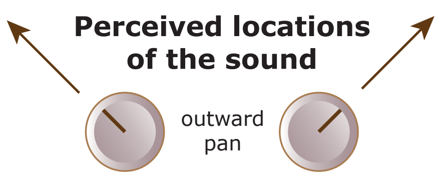

Outward pan

When two different mono audio signal are panned outward, as shown, we hear one signal located to the left, and the other located to the right.

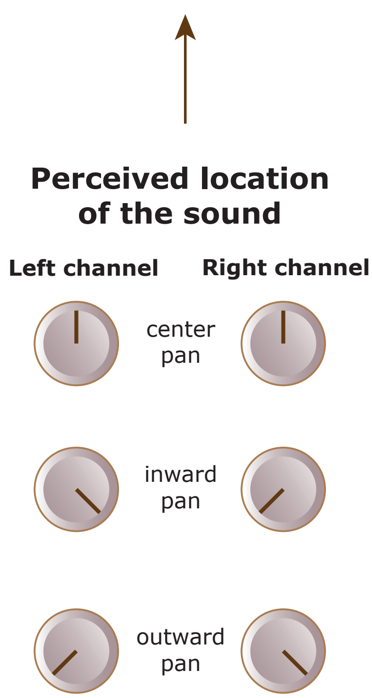

Balanced panning of matching sounds

Regardless of the various pan settings, as shown here, when matching signal levels are sent to both channels, we perceive the location of the sound is straight ahead.

This happens whenever right/left pan settings mirror each other —as long as the channel signals and volume match. Some phase cancellation may occur due to the pan, but the perceived location of the sound is always straight ahead.

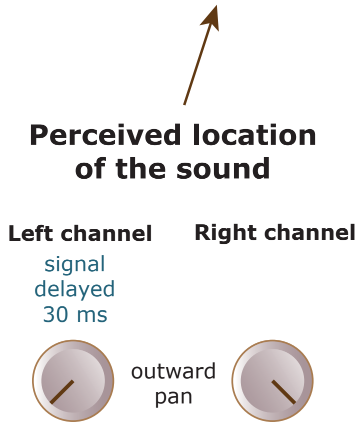

Panning with the Precedence Effect

In accord with the Haas Effect and the Precedence Effect, the perceived location shifts when we insert a delay on either channel. The diagram to the left shown a delay applied to the left channel.

Using the a hard “outward pan” (the Left channel panned hard left; the Right cannel panned hard right) an interesting result occurs when we insert a delay on the left channel, as indicated in the image to the left.

The sound of the delayed is temporally masked. This means the time delay lessens it’s perceived presence. It sounds quieter and it’s definition and detail are less clear.

This makes the right channel “sound” slightly louder. And and it shifts the perceived location rightward.

The effect occurs unless the level of the delayed channel is increased by roughly 15dB (decibels) or when the level of the non delayed channel is decreased by roughly 15dB. Either such change in increased amplitude would override (or cancel out) the precedence effect.

So as we can see the impact of a slight delay is significant.

NOTE: In order for the two sounds to fuse into a single heard event the delay must be lower than the perceived echo threshold of the listener—this is the core tenet of the Haas Effect.

This threshold can vary from about 3 to 100 milliseconds (3 to 10 hundredths of a second.) Simple sounds, such as clicks, have a lower fuse threshold. Voice and audio have slightly higher thresholds.