



## Tutorial 08

### Positional Sound

#### Mono, stereo and surround sound explained

written by Stephan Schütze

#### Overview

This tutorial will endeavour to explain some of the ways in which sound works relative to its position and how we hear it in the real world as compared to through speakers. Understanding how sound behaves can assist a sound designer in working with positional sound. Whether using mono, stereo or any of the various surround sound formats there are some aspects of sound behaviour that can be used to create effective mixes.

#### Sound all around

The way humans hear sound in the real world through their ears is very different to how sound is represented through technology. When listening to a band on a CD through either speakers or headphones, each of the instruments will sound from only two sources, the left or right speakers (or headphone speakers). In the real world the sound from a live band will travel to the listener's ears from almost every direction. Firstly there is the direct source of the instruments themselves; each drum emits sound directly from the drum skin, acoustic instruments like saxophone and trumpet produce sound directly from the end of the instrument and a singer emits sound from their mouth. An electronic instrument's sound comes from any speaker that the signal has been sent to. Speakers will also produce signals from vocal and instrument microphones attached to them via an amplifier.




Figure 1. Sound generated in a live space

Each instrument on stage produces a direct source of sound waves indicated in red. As the sound waves travel through the performance space they will reflect off of surfaces such as walls indicated by blue. As sound expands in every direction like a ripple in a pond the reflections from the walls will also reflect in every direction. As a result the entire room is filled with sound waves travelling in every direction.

Once a sound leaves its initial source it will travel within the environment in which it was generated. Soundwaves travel through air and are reflected or absorbed depending on what type of surface they come into contact with. In a typical event space, the walls will reflect some sound, while carpet, seats and people will absorb some sound (This all depends on what materials the surfaces are made from). Regardless of what and how much sound is reflected or absorbed the end result is that soundwaves from the instruments will reach the listeners ears from nearly every direction, some directly and some as reflected sound. (See the reverb tutorial for more information on reflected sound) Every day humans hear true surround sound everywhere they go. Sound waves propagate in every direction as well as reflecting off surfaces, so at any time our ears receive sound waves coming from every possible direction.

In a natural environment this true surround sound is easier to observe. In a forest the sounds of birds, insects and wind are each individually generated from many different directions simultaneously; left, right, up, down, in front and behind. Human ears constantly receive direct sound waves from every direction. It is easier to hear the directional nature of these sounds away from a city as there is not the same level of background noise as in a city or the same level of reflected sounds. This level of surround sound would be almost impossible to recreate with technology. It would need a sphere with every inch of its interior surface housing a tiny speaker capable of transmitting a sound to the ears of someone standing inside it. While this might be an interesting scientific or artistic installation it is not very practical for watching movies or listening to music.

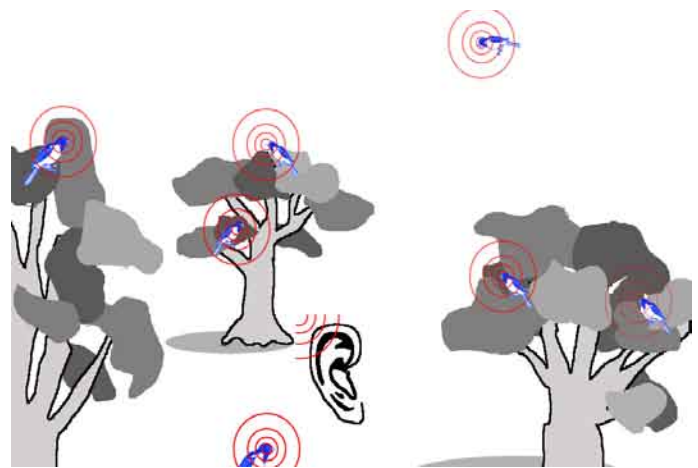


Figure 2. Sounds generated in a natural environment

While it is still possible to have reflected sound outdoors (of trees, the ground or any object) in general there is far less reflective sound and more direct sound. Without the distraction of the many sounds that are present in a city environment it is easy to identify sound sources and track the direction of origin. Individual sounds are more distinct in a quiet outdoor setting whereas in a busy city setting each sound often just becomes part of the overall noise

What is referred to as a mono sound is how most sounds are generated in the real world. Mono means one; in reference to audio it means a single signal or track. There can be unlimited sounds mixed together, but they form a single signal when transmitted. So a radio with one speaker can only ever produce a mono signal. In the real world most things that generate sound generate a mono sound, i.e. the source of the sound is a single location or point in space. Even a complex object like a car engine is made up of lots and lots of mono sounds. Each component in the engine that emits a sound will do so at a single location. It is the combination of many sound events that make up the overall engine sound. With a single sound source a person can listen and judge where that sound is coming from. The shape of the human ear provides information to the brain to allow us to judge which direction a sound is coming from, past experience with similar sounds helps us to determine how far away the sound source is.

The various curves and contours of the ear reflect sound waves as they approach the ear in a particular way. The brain interprets the way a sound wave is reflected in the outer ear and uses this information to tell us which direction the sound originated from. The reflections of sounds on each ear, and the difference between the signal received in the right ear and the signal received in the left ear pinpoint the location of the sound. This is how when a human closes their eyes they tell the direction of a sound source.

**Example:** *The sound of constantly dripping water from a single source.*

With eyes closed, humans can still tell which direction the source of a drip is coming from. A person may need to turn their head and concentrate for a few seconds, but usually most people will be able to accurately indicate the direction of the drip. This is an instinctual ability in humans. As far as how far away the sound is, this is more based on experience. Once we have heard the sound of water dripping and understand what is occurring that information is stored in the brain as a memory. One important factor that is remembered is that a dripping tap is not a very loud sound. Even placing an ear right next to the water as it falls does not produce very powerful sound waves. Because of this experience we can guess that if we can hear dripping water even with eyes closed that the source cannot be too far away (For a sound to be heard from a long distance the original sound waves would have to be quite powerful to travel over the distance). The knowledge of generally how loud a drip of water is allows the brain to calculate how far away the drip is occurring. Multiplying the single drip into the hundreds of impacts of a rain storm the brain is still dealing with the same information, simply multiplied. For each drip the brain experiences a sensation of position in relation to the ears as well as an amount of distance from the ears. A rain storm is one of the most perfect examples of true surround sound as thousands of sound sources are generated all around at different distances.



Figure 3. Rain drops surrounding the listener

In the same way as raindrops cause ripples on water their sound generates sound waves that travel out from each drop. The ear hears these drops and the brain can calculate the position and distance depending on the relative volume. In the diagram the weaker red lines represent sound that is further away from the ear and as a result quieter. Even with eyes closed a listener is aware that they are surrounded by rain falling in many different locations at varying distances.

It is easy to distinguish a mono recording of rainfall from a stereo recording. The mono signal is identical through both speakers and so can sound very flat and lacking in depth. By comparison a stereo recording contains unique material in each output track. This far more realistic to listen to as it is a lot closer to what would be heard in the real world. Add two further channels to produce a basic surround sound environment and it would appear even more realistic. The listener is now surrounded by sound sources. This is the motivation behind creating film TV and games with surround sound environments to increase the realism of the audio material. Surround sound can be very effective if used well in film, TV and games, but like anything it needs to be used well or else it can be a wasted resource.

### Creating a basic stereo sound

#### Quick Stereo Mix

It is possible to create a basic stereo sound file in the following way.

-Select a prerecorded mono sound file of birds in a forest. The length is not important but at least a few seconds will be needed to get a good effect. The longer the file the more flexibility it provides.

-Alternatively go out and record a sound file somewhere appropriate. Ensure the sound file is mono, or for the purposes of this tutorial make it a mono sound file in a sound editing program.

I am using Audacity for this tutorial as it has the features I need and being free means its perfect for anyone who wants to follow the tute step by step.

1. Open Audacity and either drag and drop the sound file into Audacity or go through the file menu and open it.
2. Use the cursor to click and drag an area of the sound file so it is high lit. I would recommend selecting about the last 35-40% of the sound file. There should now be about 40% highlighted and 60% not highlighted.




Figure 4. Selecting 40% of track to cut/detach

3. Use the edit / cut command (ctrl x) to remove the highlighted area. The sound file should now only display the non highlighted area (the original 60%)
4. Use the project /new audio track command to create a new blank audio track. This should appear below your existing track.




Figure 5. Adding a new empty track to Audacity

5. Click anywhere in the empty audio track and use the edit/paste command. This should paste the end 40% from the initial sound file into the new empty track below it.

There are now two tracks that can be played simultaneously, but the project still does not sound stereo. Stereo is far more than just having sound coming out of two speakers or each headphone. The same signal coming out of two speakers is simply a mono mix to both channels; true stereo consists of different signals for each channel. There are two tracks in Audacity with different material they are both panned to the center so what has been created is a new mono sound by mixing the two tracks into the center position. It is a different sound than the original, but its still mono. To make it stereo there is one last step.




Figure 6. 2 tracks but still mono

The red circles on the panning sliders show that both tracks are panned at the center so even though there are two different tracks they will be mixed together to create a mono track. The pan positions need to be changed

6. Slide the pan slider for track one to the extreme left and then slide the pan slider for track two to the extreme right (the opposite for both is also fine as long as track one and two are panned differently).



Figure 7. Altering the pan position of each track

The red circles indicate the pan sliders for each track. Note that they have been moved to the absolute extremes of left and right. This is the most obvious way to produce a stereo effect. Varying the positions will create different sounding effects.

This has created a stereo sound. When listened to there should be different sounds coming out of each speaker at different times which is what defines a stereo sound. Two channels with different signals being played together.

Now, why didn't I just put a full copy of the original sound file in both track one and two and do the pan step? Would that have created a stereo file? No. By panning the two copies of the same files extreme left and right it would have created the stereo effect of spatial distance, but there would still have been only one signal coming from both channels. So it would have created a wide panned mono sound

Sometimes in music it can be effective to wide pan the same signal to both sides. With drum tracks this can give the effect of a two drums playing the same line and can be really effective. But there will also usually be other instruments mixed in as well so the overall mix will consist of different material between left and right.

Use Audacity and experiment with sliding the pan for each track into different positions. Extreme left and right is the most obvious way of creating stereo and the result is the most dramatic stereo effect, but sometimes it can be too much. Try panning only a slight way from center and see how different it sounds. As the entry in the F2A glossary states, Panning is the positioning of a mono signal within the stereo field, this is positioning where we want the sound to be. Extreme right panning is placing the signal distant from the ear on the right side, while close right panning is placing the signal close to the ear. This tutorial will cover depth of field later.

I will now explain why not to cut the original track exactly in half and use 50% in each track. One of the most common reasons to want to create a stereo track is to create a background or ambient track to create a general feel for the location. In this case we have created a forest with birds all around. If this was a scene from a film or TV then we might need to fill considerable time with the ambience track so we need the background ambience to play for quite some time. If simply cut in half and placed in each track the file would have created a stereo sound that would only run for half the time of the original. A looping file could be made out of this stereo file, but if the original file was not very long then a short loop what becomes too obvious too quickly would be produced (see how to make looping files).

The following method will produce a longer stereo sound.

1. Select all track one and highlight it.
2. Use the edit/copy command to copy all of track one's material.
3. Click anywhere in track one to the right hand side of the wave file in blank space. (if the wave file fills the entire Audacity window area then use the view/zoom out command (ctrl 3) this should give you some empty space at the end of the wav file.
4. Use the edit/paste function (ctrl v)




Figure 8. Copy/Pasting to create a longer sound

If each track is able to loop seamlessly then a longer file can be created by copying and pasting the entire track over and over. The difference in lengths between the two tracks will result in an offset of loop points which will create a much longer period of time without repeated material

The paste function will automatically paste the material directly onto the end of the current wav file making a new longer wav file. (Not all audio programs will do this, many will paste directly on the cursor's position creating a gap between the existing and new material if the cursor is positioned that way. This is ultimately more flexible and I find Audacity's behavior in this kind of strange. It's very helpful for people new to audio editing but it is very limiting for more experienced users)

Repeat steps 1-4 for track two. The two tracks are of different lengths, either continue to paste onto the end of the shorter track until they are roughly the same length or select any excess material from the longer track and delete it to match track lengths. By copying and pasting different lengths between track one and track two it has effectively offset the loop points where they start to repeat (copying a file and pasting it onto its own end is the most basic way of creating a loop). By offsetting the loop points between the left and right this is effectively creating a loop that will be much longer.

For example. A source file that is 7 seconds long, broken it into two pieces of 3 and 4 seconds, then used as above will not get a repeating loop until the 12 second mark. The new loop point will be where the two tracks both finish simultaneously. So using this method it is possible to turn a 7 second mono file into a 12 second stereo file.

Finally highlight both tracks (or use ctrl a) and use the project/quick mix command to quickly join the two tracks together in their current state. This allows the file to be saved out as a new file of a stereo mix down of the two mono tracks. In Audacity you need to use the file/export command and you can choose either wav, mp3 or ogg vorbis as your final file format.

This method can construct stereo files out of mono sounds very quickly and easily.

Be aware that if the mono sounds do not cleanly loop then the extending of tracks by copying and pasting can result in unwanted click sounds. See the tutorial on looping sounds to avoid this and other issues with creating loops.

There is one major disadvantage to using this method for creating loops and that is that it creates a stereo file of greater intensity than the original mono files. What I mean is this.

In the original mono recording lets imagine the forest had two birds in it talking to each other. When the original file is cut up the material used is taken from further down the time line and added to the beginning, splitting them between left and right sides. This becomes is a mono sound with two birds in the right channel and a mono file with two birds in the left channel, The number of birds in the new stereo forest has doubled. For creating basic ambience effects this might not be a problem, in fact adding more birds to the forest might be a good result, but it is not always appropriate.

Using a different example of the forest can more clearly illustrate this issue.

Example. A seven second mono recording of Tom and Jane talking to each other. Cut the recording into 4 seconds and 3 seconds and paste them into tracks 1 and 2 as before with the birds. The first obvious issue here is that it cannot be looped as this will result in the conversation being repeated over and over, but this is a different issue that is covered in the loops tutorial. The main issue with quickly stereo panning this material is that it will create Tom and Jane talking out of the right side and simultaneously talking out of the left side. Not only are each of them talking over themselves but they are moving unnaturally around our stereo environment. This undesirable effect can happen with any kind of material and is why often it is better to use a slower and more exact method of creating a stereo sound.




Figure 9. Looping sounds or varying lengths  
By adding different length loops together it is possible to create a segment of sound that will play for longer without repetition. Tracks 2 and 3 are both under 5 seconds long, but together they will play for 14 seconds without repeating material.

The following method can be used to create a quick stereo mix from a mono dialogue recording.

1. Open Audacity and either drag your sound file and drop it into Audacity or go through the file menu and open it.
  2. Select all of track one and edit/copy (ctrl c) it.
  3. Use the project/ new audio track command to create a second track
  4. Edit/paste (ctrl v) a second copy of the initial track into track 2
  5. Slide the pan slider for track one to the extreme left and then slide the pan slider for track two to the extreme right (the opposite for both is also fine as long as track one and two are panned differently).
- There is now the same material in both track 1 and track 2 with the two tracks panned to the extremes. Decide which person is on which side (doesn't matter which, assign Tom on the right side and Jane on the left).
6. Activate the solo function on track one.
  7. Go through track one until you find a section where Jane is speaking. Highlight this section of Jane's speech.

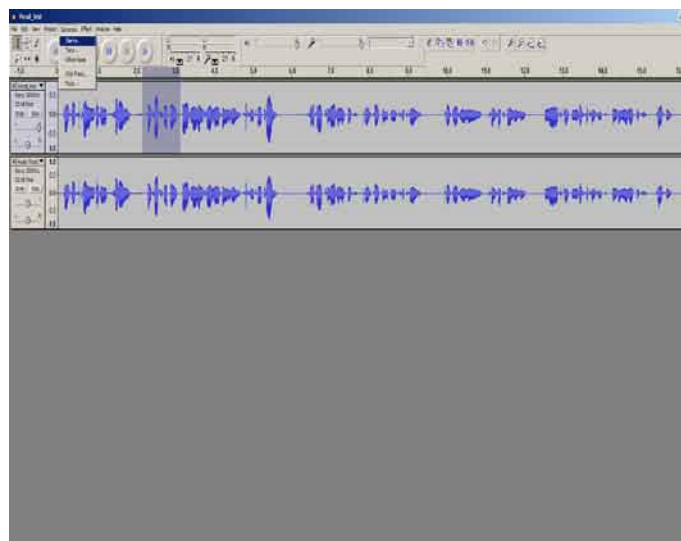


Figure 9. Selectively removing a voice from each track  
By removing all of one voice from track 1 and all of the other voice from track 2 it is possible to create a stereo conversation.

8. Use the Generate / silence command. A pop up window should appear. Normally this window will ask to specify the length of silence to create in a particular track. In this case because there is an area highlighted beforehand it will default to that exact value. Click ok  
The selected area of Jane's speech should now be completely silent.
9. Go through the remainder of track one and highlight and replace all of Jane's speech with silence.
10. Once this is complete unsolo track 1 and solo track 2. then repeat the entire process but this time highlight and silence Tom's speech sections.

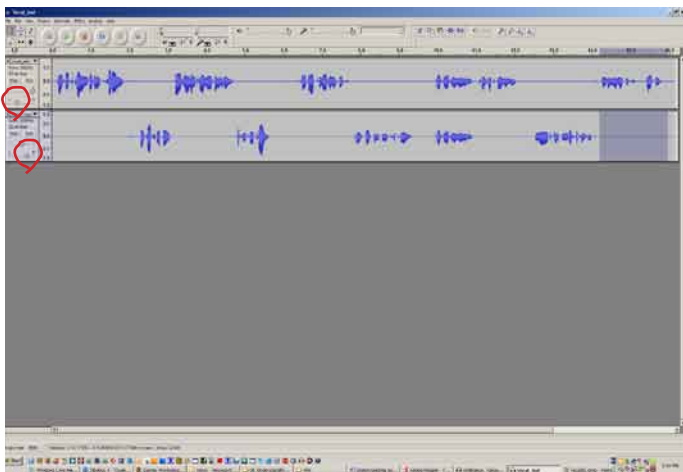


Figure 10. The two isolated voice tracks  
Note the pan positions

Once this has been done playing the two tracks together should sound as if Tom and Jane are having a conversation with the listener standing in the middle. By adjusting the panning it is possible to position the two speakers as close or as far away from the listener as required. The same process can be used for any number of different speakers. Use a new track for each speaker, and delete all dialog except for the speaker assigned to that track. Pan each speaker slightly differently and it can create the illusion of a group of people standing around the listener.

With a more complex sound environment, such as a film sound track, each sound needs to be individually placed into the sound environment and made to "fit". In the same way as sounds need to be made louder or softer to better suit what occurs on screen, sounds are also panned various degrees to the left or right to suit the visuals. In a complex sound environment this can take considerable time, but when done properly the results can be fantastic. Exactly the same process applies to surround sound there are just more positional locations in which to place a sound. Be careful not to overuse sound positioning. Having every sound positioned to the extremes will not always sound good and can take the impact out of special events. One of the best uses of surround sound in a film I have heard was when almost the entire film used only the front speakers, one particular event utilized the rear speakers as well. It was such a drastic change to suddenly have sound all around that it was a real highlight to that point in the film.

### Complex stereo placement

When creating a more complex stereo mix each sound added to the project is assigned a position to achieve the end result. In some cases several sounds may share the exact same position in the mix in other cases every single sound will have its own location. For example: creating the sound of a police car driving by from left to right. The sounds used include the sound of a car driving and the sound of a police siren. As the siren would be attached to the car in the real world the position of the two elements need to match as they move. Both sounds need to move from the left to the right at the same rate otherwise it will create the effect that the siren is moving independently to the car. With the sound of an explosion each individual element would blast away from the initial point at different speeds and angles so each element could be panned individually. The chaotic nature of the explosive event is such that it does not matter where in the mix the elements are positioned.


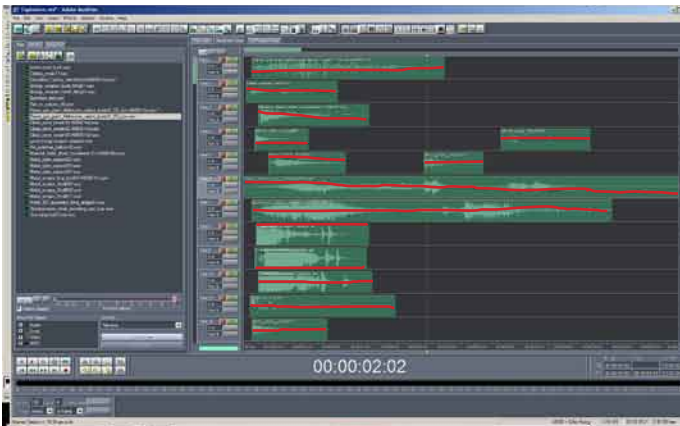


Figure 11. Matching pan compared to random pan  
The positions and movement rate of both the car and siren sound must match to create a realistic sound effect, whereas the various elements of an explosion are chaotic and can be set completely randomly without creating any issues.

Like any aspect of sound design positional panning is often a matter of trial and error. Unless a sound needs to be specifically positioned to match an event, try moving each sound element around in the mix to get a feel for where it might work best. Separating sound of similar frequencies can be a good approach as it can prevent those sounds from clashing. Panning should be done before level balancing because moving a sound further out positionally will lower its effective level in the mix, so try and establish a basic position for the sounds before you finalize all the volume levels. Also think of each sound as a tiny narrative, if there is a large object that wooshes past the listener traveling on the right hand side, there should probably be an appropriate impact sound for that object on that side of the mix at a later time. Perhaps no one will notice this level of detail in amongst the mayhem of a large explosion sound effect, but it will be worth the effort for anyone who does.

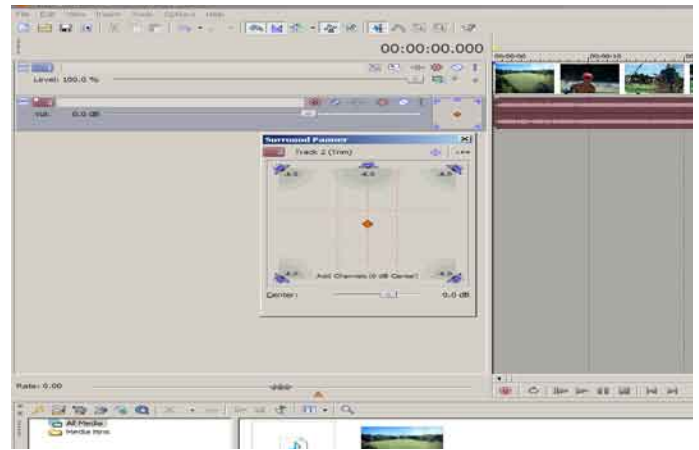


**Figure 12. Varying pan positions creates depth of field**  
The many sound files used to create the explosion sound have a great variety of panning positions applied to them. Some are slightly offset from centre while others are set to the extremes. Some sound also shift from one side of centre to the other to simulate an explosion element in motion. The overall effect of this variety is to create a much greater depth of field in the final sound effect.

### Surround sound

As many sound effects libraries are mono and many methods of recordings sounds produce mono sounds, creating a sound environment will usually require the designer to position each sound within the mix. With a stereo project this is done by panning sounds either to the left or the right to various degrees. In a surround sound project the sound needs to be placed somewhere within the available speaker positions. Many surround systems also allow for sounds to be directly allocated to the centre front speaker or to the sub channel. The centre front channel is usually used for dialogue to keep speech clear and distinct from other sounds, and the sub channel is reserved for sounds with low frequency content. This setup emulates a situation where most sounds in the environment are further back almost like a backdrop, while speech comes from a direct forward source, much as if someone was standing directly in front of you talking. The sub channel is non directional; this is because low frequency sounds have very long sound waves that travel further, but are difficult to isolate from a directional point of view.

Positioning sounds in a surround environment is essentially the same as a stereo environment in that you select a direction and then position the sound to the appropriate degree in that direction. With stereo it is a single line with left and right at either extreme, a surround sound environment allows sounds to be positioned within a broader area, usually a circle. The position of the sound defines how much of its signal is sent to each speaker. A sound can be entirely channeled to one speaker or positioned in such a way that it is audible through two or more speakers. If a sound contain low frequency information then that sound is also routed out of the sub channel as well, in this way a sound will play through the appropriate speakers to indicate its special location as well as the sub channel to boost its low end frequencies.



**Figure 13. Surround sound panner on Sony Vegas**  
The surround panner in Vegas allow the user to alter each tracks positional information by dragging an icon within the pan area. This can be automated to alter over time allowing for complex surround environments to be created.

Mixing in surround sound like any creative process is up to the tastes of the creative team and the requirements of the project but there are various guidelines that can help to avoid a bad mix.

- Too much movement between channels can be very disorientation for the listener. This might sometimes be exactly the effect required, but often even if accompanying visuals with lots of movement exactly tracking the movement in 3D space may not produce the most desirable effect.
- Sound for film, TV and game is already altered greatly from how it functions in reality so modifying it further to maintain a solid continuity that support the narrative is often more important than trying to match every single shake and tilt of the camera.
- Overuse of the surround speakers can mitigate their usefulness as a narrative tool. Although almost every sound we hear in real life comes to us from every direction, in media productions surround speakers or often used more for effect. In games the rear speakers can be an essential tool to support game-play and narrative, in film and TV they are often used for effect instead of trying to simulate reality.
- Allow your ears to decide. Regardless of design or desire if something just doesn't sound right when it is positioned in a 3D space then change it. There are no rules saying surround sound must be used, and must be used in these circumstances. Work with what sounds suitable for the project.

Sound effects don't need to be complex to be good or work effectively. Sometimes it will be necessary to construct a complex sound from many different sources mixed and positioned carefully to suit a project but it does not mean every sound needs to be created in this way. A subtle use of positioning can be very effective if used creatively. Like all artistic processes it depends on what you are trying to achieve.