



## Tutorial 06 Noise Reduction Understanding and avoiding unwanted sound material

written by Stephan Schütze

### Overview

As a term, “noise” can be used to describe any audible sound. Noise is also described in various dictionaries as “an unpleasant or undesirable sound.” It is this second definition that has become the more common use for the word. Noise is often used to refer to any aspect of a signal that is undesirable. Video signals can contain noise in the form of static or lines or a blurry signal and when referring to audio it usually refers to a high frequency hiss or a low frequency hum, but any sound that is an unwanted part of a project can easily find itself stuck with the noise label.

There are many processes and plug-ins that describe themselves as “Noise reduction” that function in different ways with varying levels of success depending exactly what you are trying to achieve, but there is no one magic bullet for removing unwanted sound material in a mix. This tutorial will go through a few of the basics of minimising unwanted noise and suggest a few approaches that might be useful.

### Understanding noise

First it’s important to understand what noise is and how it works. As previously mentioned some noise can be anything that’s undesirable, but there are sounds that are specifically labeled noise. These include white, pink and brown noise. White noise is the most commonly recognised and is often characterised by the hissing sound of a TV or radio when not tuned in properly. Technically, white noise consists of a mix of frequencies all of equal power output. This basically equates to every frequency existing all at the same level. Brown noise (also called red noise) is different in that it has a relationship of frequency to power. As the frequencies get lower they get stronger in power and so are more noticeable. Pink noise kind of exists in the space between white and brown noise. It does not have equal power across all frequencies, but it also does not have the linear increase that brown noise has.



Figure 1. Sound file in regular display mode

The above image illustrates a wave file loaded into Adobe Audition. In the standard display mode the horizontal axis represents time and the vertical axis represents amplitude. The sound above is an example of typical human speech and as such has an irregular pattern of peaks and troughs that reflect the constantly varying amplitude of spoken words.

All sounds are made up of a range of different frequencies combined to create the sounds we hear. Some editing tools can display a spectral view of a sound file. This spectral view shows the frequencies that exist in the sound file and colour codes them depending on how strong they are. Figure 1. shows a regular sound file of human speech. In a normal editor, the wave pattern is displayed. The horizontal direction represents time and the vertical direction represents amplitude. Figure 2. shows the same sound file in spectral view. This time the vertical direction indicates frequency. The higher up the vertical line the higher the frequency. The colors range from black, meaning no signal at all, through dark purple and blue for very weak power for those frequencies, up through red to yellow and then white for extreme power levels of those frequencies. Figure 2. is a spectral view of the frequencies and their power levels in human speech.


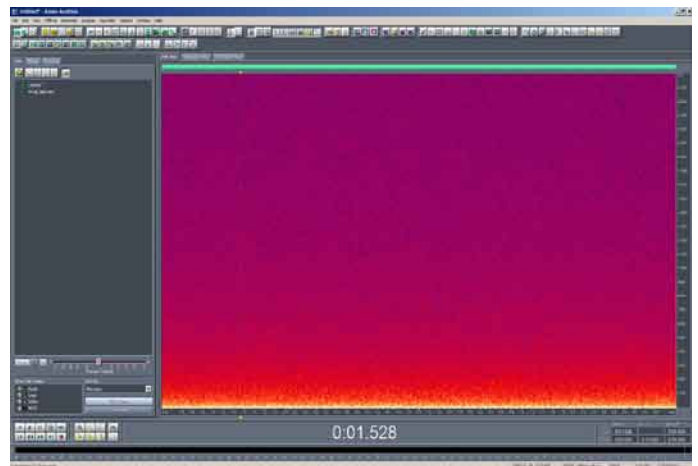


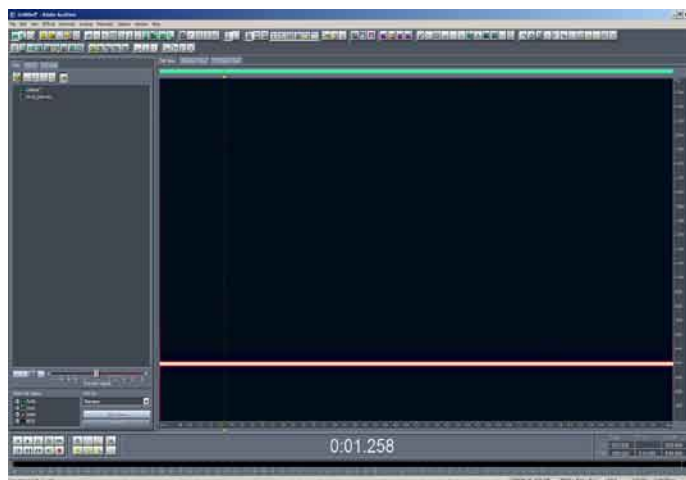
Figure 2. Sound file in spectral display mode

Stronger frequency material is displayed in brighter colors and the full range of frequencies that make up a sound can be quickly identified. Areas of silence can still be identified, but it becomes apparent that there is far more sound information present in the wave file than the standard view may have indicated.

Human speech tends to fall in the range from about 100-500Hz and it's in this range that the strongest frequencies are displayed in the spectral view of figure 2. There is also frequency material in both higher and lower frequencies at much weaker signals. Most sounds will include material across a range of frequencies that may not be obvious when listening to the sound. This frequency information can include harmonics from the actual sound, other sounds that might also have been recorded with the primary sound, and of course elements of hiss, vibration or hum all of which are referred to as noise. Noise sometimes exists at the location of the original sound. Objects like fans, machinery and natural sounds can all add to unwanted material, but noise can also sometime appear through the process of recording. Equipment, cables and transferring of data can often generate noise into a signal.



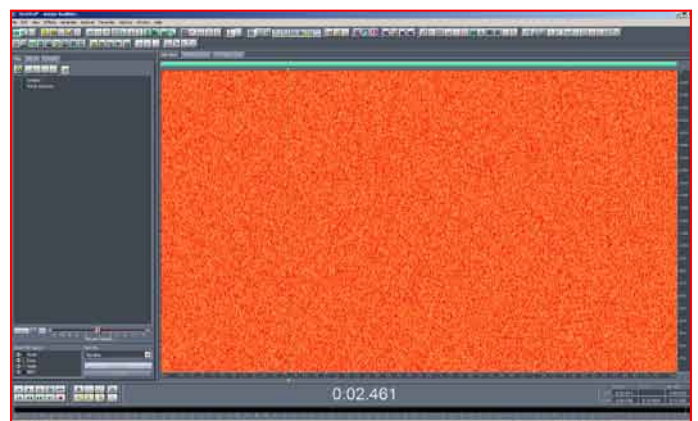
**Figure 4. Brown Noise spectral view**  
 Because the frequency material at the lower end of brown noise is stronger than at the high end, the spectral display shows bright colors at the bottom of the screen gradually decreasing in intensity as they progress towards the top of the screen.



**Figure 3. 4000Hz Sinewave Spectral view**  
 Under spectral view a sine wave is displayed as a single line as its frequency information is extremely limited. Only a pure tone will display this behavior. Most sounds are made up of a great many frequencies mixed together to create the final sound.

Figure 3. illustrates the spectral view of a sine wave created at 4000 Hz. Because it is a generated sound and because sine waves are pure tones, the only frequency information is a line at the 4000 Hz point in the spectrum. This is very rare to see as most sounds are a combination of many frequencies, but it is a good demonstration of how the frequencies are displayed. If you were to increase or decrease the pitch of the sine wave, the line in spectral view would adjust up or down accordingly. Figure 4. is an example of brown noise. This time there is far more frequency information, but because brown noise consists of stronger frequencies in the lower frequency range the bottom of the spectral view is where the brighter colors are displayed. Spectral displays like this can be useful for sound editors as they quickly display the region where the strongest frequencies are present. This can be useful for mixing and equalising sounds.

An illustration of white noise looks like a constant pattern. This is because of the flat frequency response of white noise. All the frequencies present are of equal power, and so the spectral view is like a constant color pattern. Pink noise falls somewhere between white noise and brown noise. Pink noise does contain frequency material of differing signal levels, but not in the linear manner that brown noise does. Using a spectral display and your ears, you can isolate material that you want to keep and track down sounds that you want to classify as noise. Once you have identified noisy material then you need to remove it. There are various approaches to reducing noise that will be more or less effective depending on the circumstances.



**Figure 5. White Noise spectral display**  
 White noise consists of frequency material all of equal power levels. This results in a spectral display that is essentially a constant color pattern.




Figure 5. Pink Noise spectral display

While not constant in level like white noise, Pink noise does not scale frequency to power levels in the same way as brown noise. It falls somewhere between the two in characteristics.

## Starting clean

Always try to capture the best signal possible. I know it's probably become a cliché these days but it really is the best practice to get into. Recording the cleanest clearest source material in the first instance is the best option for a good end result. The reason why this might sound clichéd is because lots of professionals keep saying it, and the reason why lots of professionals keep saying it is because a lot of people don't listen. I have talked with a scary number of people who believe the "noise reduction" button in the audio software will fix all issues and they can quite happily go ahead and record a soft melodic piano line in the same room as a washing machine or an industrial drill press. These same scary folk then find themselves devastated when their final mix sounds more like the insides of their car than a piece of music. Noise reduction is as much about prevention as it is about "curing" the sound once it is recorded.

There are two useful rules to keep in mind when recording.

1. Set the record level high when recording. The sound that you want to capture should be a good strong signal on your recording device without peaking. If there are other sounds in the area try waiting for them to stop, or reposition yourself to minimize how much of the other sound you capture. See the basic recording tutorial for more information.
2. Where possible, use an appropriate microphone. If all you want to capture is your voice, then use a good close in mic that has a very limited range of pickup. That way background noise is not an issue because the mic shouldn't even detect it (these are the types of mics that location journalists use). If however you are trying to record a bird high up in a tree, that same microphone will be completely inappropriate. You need something with far more range, and also a microphone that has a more direct recording pattern would be better (such as a shotgun mic).

Of course this might be difficult as audio equipment and especially microphones can be very expensive. If it's the case that you simply can't afford more than one mic then you will need to adjust your expectations when you are recording. Do some tests when you first purchase a mic so you have a good understanding of its limitations, then you can plan to work around what you can and cannot achieve.

The absolute top end gear provides the user with the ability to record what is present without adding extra noise to a sample. Line noise, hum from a mixer or amp or signal degradation can all be issues in lower end equipment, but the main advantage of not having these issues is being able to clearly hear what else may be a problem. If you know your equipment is clean then you can really listen to your environment to ensure nothing else is generating noise. Air-conditioning, fluorescent lighting and electronic equipment will all generate sounds of some level; do not underestimate how well a microphone can pick up these sounds. Common computer fans are a terrible source of noise and can often be picked up by a mic even from another room. Recording out of doors adds a whole range of other potential contamination issues.

## Keeping it clean

What some people may not realise is that every step of the recording process can potentially add more noise to the mix. This is far less of an issue these days with digital recording gear, but any analogue stage of the recording process needs to be dealt with carefully. A microphone is essentially an analogue device in that it takes the movement of sound waves and converts them into an electrical signal via the movement of a diaphragm much like that in the human ear. This is an analogue process even if everything else is entirely digital. In most cases there are further analogue stages as the signal from most microphones is transferred down a standard mic cable into the recording device. The quality and serviceability of the cables and connectors are all places where noise can sneak into a recording. This is the best reason to maintain your audio cables in good condition. Phantom power boxes, line mixers, pre amps and anything else the signal passes through can also add potential noise. Keep cables in good condition and use as short a cable as possible between the mic and the recorder.

## Once it's gotten noisy what then?

Assuming you have tried your hardest to sample clean material and there is still some unwanted signal, (this can happen for various reasons no matter how careful you are). What can be done to remove unwanted sound material?

Most noise reduction tends to work through manipulating frequencies. By either boosting the strength of wanted frequencies or reducing the strength of unwanted frequencies the overall effect of noise can be reduced. The actual processes can vary, but overall frequency manipulation often cannot completely remove noise from a mix. This is because sounds in general cross a large range of frequencies especially when harmonics are taken into account. So even if the sound you want to keep is a low frequency sound, reducing or removing high frequencies can alter the original sound. Sometimes it becomes a trade off between removing noise and altering the material you want to keep.

## Low hum or vibrations

With outdoor recording, especially recording natural sounds anywhere near or in a large city you will almost always record a low frequency rumble or hum. In general the sounds of people and traffic generate a fairly audible low end vibration. If the sounds you want to record are fairly high in the frequency range then it is usually safe to entirely remove the low end frequencies. Either use a spectrum view that allows you to select all the frequencies below a certain range and delete them or an EQ tool that can reduce frequency signals can to remove the low end rumbles. The beauty of a spectral view is that you can actually see the sound material you want to keep and what frequencies it covers and then select the frequencies below that range and delete them entirely. Not all audio software contains a spectral view however so an EQ process is more often used. Because audio harmonics are of a higher frequency than the source sound it means that removing frequencies below the source is far less likely to alter the wanted sound. Just be careful that the sound you want to keep is not a complex sound that covers a wide frequency range and accidentally remove some of its lower end. The beauty of modern software based systems is the ability to undo changes made to your material, so feel free to play around with removing different frequencies and seeing just how far you can push things. This is a great way to learn safely in the knowledge that at any time you can revert to the original file if you stuff things up.

## Hi-end hiss

This is where things can get tricky. A lot of noise reduction programs allow you to select a region of hiss or unwanted sound and then use that sample as an indicator of the frequencies to suppress. The software will analyse the sample you have selected and use those frequencies as a guide as to which frequencies to suppress. In practice it is an easy step of selecting a silent area of a recording where there is nothing but hiss and the software will go through and remove these frequencies and thus remove the hiss. The problem with this is that it does not allow for the fact that those same frequencies might be present in the material you want to keep. In all instances of working with these tools the trick is to try various levels and listen very carefully to the results. Make sure that the process of removing the hiss is not also taking the life out of your original material. In some cases if you have recorded a signal too soft and there is a lot of hiss present, it really can be better to return to the studio or location and record the material again. Should this not be possible because the recording was a rare opportunity or just too expensive to try and record again then you will need to go through the material slowly and carefully. Applying a light treatment of reduced higher frequencies through EQing can remove some of the hiss, but you need to monitor what you are doing. Unfortunately there is no magic bullet to remove noise.



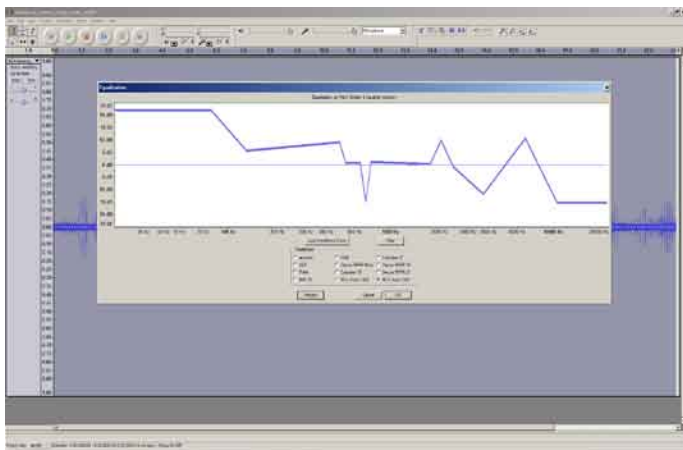
Figure 6. Noise Reduction tool

This tool within Adobe Audition allows the user to select a region of a sound file as a sample of the material to be removed as noise. The program will then analyze the sound file looking for material that matches the set of frequencies and amplitudes of the selected sample and suppresses them. The level of how much the tool suppresses this material can be varied to improve the results.

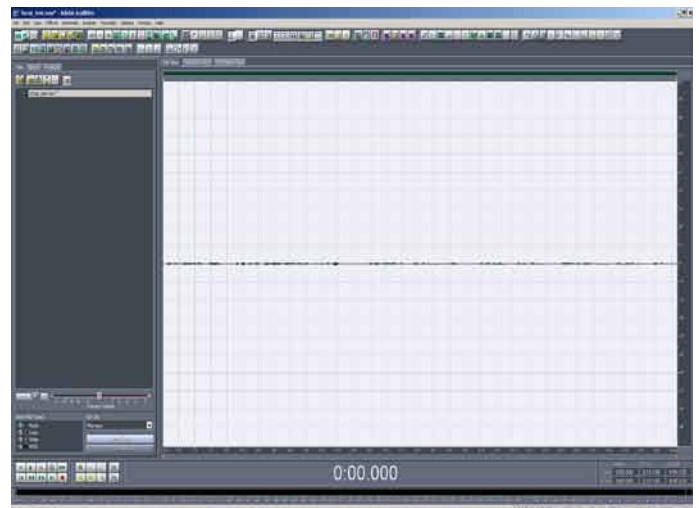
Removing unwanted high end noise can be more perilous than low frequency noise because of harmonics. Musical sounds have more obvious harmonics than sound effects, but sounds may contain material of higher frequencies than may initially be apparent. Also a sound will often contain frequency material that you may not be able to directly hear or recognize, but that can be obvious if it is removed. Sounds can quickly become muddy or muffled if you reduce or remove too much high frequency material. It can take considerable training to be able to listen to a sound and instantly recognise the frequency material that is a part of that sound. As an experiment, choose any sound file and in an audio editor use the EQ function to change the levels of various frequency groups. Taking a particular frequency group to either extreme is a quick way of hearing how significant that group is as part of the overall sound. It can be quite revealing to find out how fundamental to a sound a particular range of frequencies are when it may not be obvious when the sound is first heard.

## Amplitude processing and noise floors

Another way in which some editors deal with noise reduction is by analysing amplitude levels. Usually the primary sound in a recording is the loudest material present. Background sounds and unwanted noise tend to be at a lower level than the primary material. Amplitude processing works by selecting all sound material below a certain amplitude threshold or noise floor and suppressing it. In programs that allow real time altering of the threshold range, you can actually hear various elements of a sound cease as the threshold is increased. Ultimately the entire sound file can be reduced if the noise floor is raised too high. Again, this method works on the assumption that the primary sound will never reduce in amplitude below the set threshold, but in the same way that the primary sound may share frequencies with unwanted material it is possible for a primary sound to fade in or fade out from below the threshold level. The best rule in both situations is to not simply trust the process to a computer to fix in one hit. Rather to carefully monitor the sound as you vary the levels of both amplitude and frequency noise reduction to achieve the best balance of reducing noise while maintaining the integrity of the primary sound material.



**Figure 7. Equalizing tool in Audacity**  
By understanding how equalization can effect a sound you can learn to use it to reduce the effects of unwanted noise material in a sound file. Radically changing the EQ values of a sound file will quickly demonstrate how a sound can be manipulated.



**Figure 8. Overprocessed Noise reduction**  
This is the same wave file as figure 1. but it has been processed with a noise reduction tool with the cutoff threshold set to maximum. The result is the almost total suppression of the sound file. This is an extreme example but illustrates that the tools can only work with the information provided to them by the user.

Noise has been a serious issue in sound and music production for many years and it is not an issue that will ever be easily and consistently solved. Understanding how it works is a more useful tool for avoiding noise issues than relying on automatic software solutions.