THE BROADCAST BRIDGE
Connecting IT to Broadcast

Audio For Broadcast

Part 3 - Broadcast Audio Processing Tools

A Themed Content Collection from The Broadcast Bridge

Supported by:

LAWO

Themed Content Collection

Contents

Series Overview

By Tony Orme. The Broadcast Bridge.

Audio For Broadcast is a Themed Content Collection which serves as an audio course and reference resource for broadcast technologists.

It covers the science and practical applications of all aspects of audio in broadcast. It is not aimed at audio A1's, it is intended as a resource for the 'all-rounder' engineers and operators who encounter and must deal with audio on a day-to-day basis but who are not audio specialists... and everyone who wants to broaden their knowledge of how audio for broadcast works.

Sometimes you don't need to know everything about something. In our frenetic and challenging working lives more and more jobs are multi-skilled and adaptive, and we're often expected to cover more functions than we are comfortable with. We can't all be experts. Sometimes we just need enough knowledge to get the job done.

Broadcast Audio introduces some basic concepts for people who aren't audio people. While the audio signal chain is very different to what it was a generation ago, some fundamentals remain. Things like signal processing haven't changed and still need to be done well, but the series will also look at how we got to where we are.

It will look at how consumer technology and broadcast intent has influenced the production of content, and what that means for the people who work to put everything together. Because you never know when you might be asked to lend a hand.

Audio For Broadcast is a five part series:

Available now:

Part 1. Theory, The Console, Monitoring & Metering

Part 2. Broadcast Audio Capture

Part 3. Audio Processing Tools

Future parts due in 2023:

Part 4. Routing, Sync' & Latency
- Routing & Asset Sharing
- Synchronization
- Latency & Delay Compensation

Part 5. Broadcast Audio Workflows
- Traditional Studio Signal Flow
- Outside Broadcast & Distributed Teams
- Cloud Based Audio

Dynamics Processors

By Kevin Emmott. The Broadcast Bridge.

Keeping audio levels under control is the foundation of audio mixing, and Dynamics Processors give us tools to automate level control in various ways.

In modern broadcast environments it's good to know a little bit about everything. Knowing a little bit about what other people do and how they do it not only provides a better appreciation of what it takes to knit together a live production, but it also promotes empathy and builds confidence when we're asked to pitch in.

In broadcast audio, where preset processing chains are often already in place, if everything is set up properly, there should be little to do outside of a tweak here and there. But sound is annoyingly unpredictable, so it's important to understand some of the techniques sound engineers use to help manage sound as it reaches the control room.

Audio processors, whether they are standalone or built into a mixing console, are essential tools in helping to tame unruly signals and provide the ability to manipulate those signals in a multitude of ways and Dynamics Processors are a fundamental part of that tool kit.

Dynamics Processors help sound engineers protect against unpredictable spikes; they are used to adhere to loudness regulations, to ensure broadcasts are intelligible, to tailor broadcast output to fit a specific genre, or a specific medium, and are not as frightening as they first seem.

In fact, Dynamics can be summed up in ten words, and if you are pushed for time, you can skip the rest after this next line:

Compressors make big sounds quieter; expanders make quiet sounds bigger.

Home On The Range

Of course, in reality Dynamics aren't just about volume; they provide operators with the ability to manage the dynamic range, which is the difference between the loudest and the quietest parts of an audio signal.

Compressor/limiters and expander/gates provide this functionality, and you can hear real-world examples of how these tools can affect broadcast output every time you turn on the radio in the car.

From the quietest whispers of a clarinet to the intensity of the brass section, classical radio has a very wide dynamic range; in a symphony, the range is integral to the storytelling. To maintain that, classical music broadcasters apply very little compression to their output.

But it's no good for the school run when it's competing with traffic, road noise and chattering children. To hear the full dynamic range, you really need to crank up the volume to hear the quiet bits, which makes the loud bits deafening.

To resolve this, drive-time commercial music stations will compress the dynamic range of their output to make it more audible. Compressing the peaks to create a narrower range means the content can be broadcast at a higher volume without the risk of encroaching on loudness regulations or distortion/clipping. In this kind of environment it makes the whole programme easier to listen to.

Putting On The Squeeze

A compressor reduces the dynamic range of an audio signal by attenuating (reducing) the loudest parts of the signal – usually the transient at the start of the waveform – while leaving the quieter parts untouched. This smooths out the signal, reducing the peaks and making the overall signal quieter.

Compression is controlled with a combination of four settings: threshold, ratio, attack and release.

Threshold and ratio settings work as a team; without each other they are useless.

The threshold sets the level at which the dynamics processing kicks in, and only signals which are higher than the threshold will be affected; by definition, any signal lower than the threshold is left alone. But in order to have any effect at all, the compressor needs to know what to do with those signals.

This is set by the ratio. When a signal exceeds the set threshold its input to output ratio is changed by a set ratio. A ratio of 1:1 will have no effect at all – it simply means that for every 1dB of input signal beyond the threshold setting, the output will also be 1dB. However, a ratio setting of 3:1 would mean that for every input signal 3dB above the threshold it would reduce to 1dB, bringing down the level of those troublesome transients by 2dB. A ratio of 5:1 reduces 5dB of input above the threshold to 1dB of output, resulting in a reduction of 4dB.

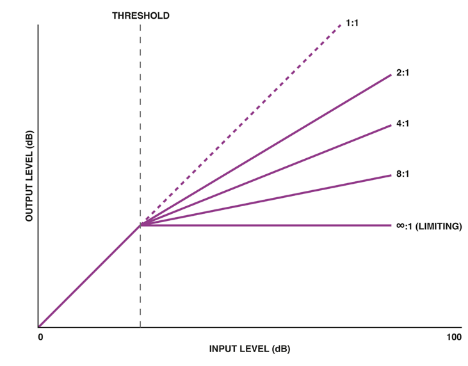

Basic compressor

Ratios of 10:1 or more changes the game significantly, but we'll get to that later.

On The Attack

The attack parameter defines how quickly the compressor kicks into action once the signal has exceeded the threshold. A shorter attack time will reduce the signal as soon as it hits the threshold, while a longer attack will allow some of the transient to pass through before any

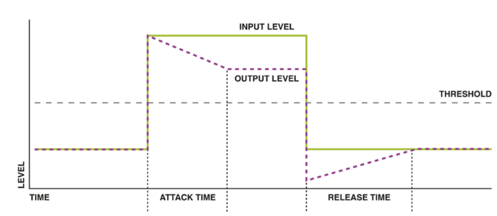

Attack & Release

gain reduction is applied. Conversely, the release time defines when the compressor stops processing the input signal and how quickly it returns to applying no processing after hitting the threshold.

These settings can significantly shape the sound: for example, as a faster attack and release time lets less of the transient through, it can result in a more aggressive sound and boost the perceived loudness.

In some outboard gear, some of these parameters are already set, but the dynamics section of a broadcast console will all be user definable. There is no right or wrong way to set these up; like all things in live broadcast, the sound engineer's primary consideration is always to meet the needs of the broadcast, and so compression levels and how aggressive they are will depend on the content, the output, and the broadcast medium.

If You Kneed It

Another element to dynamics settings is the knee setting. The knee determines the range around the threshold where gain reduction will be applied. A hard knee means that gain reduction is applied at the assigned ratio as soon as the input signal hits the threshold. A softer knee starts to apply gain reduction before the signal hits the threshold, from a lower point in the input signal, and will apply full reduction when the signal reaches a certain level above the threshold.

Knee settings can be used to produce a more transparent gain reduction and broadcast consoles will often have a graphical representation of the setting to help visualise what it will do to the signal. You'll know it when you see it; it looks like a leg bending at the knee.

One consequence of using compression is that it makes the overall signal quieter, and some parts of its output will be at a lower level than its input. If the application of compression makes the overall signal too quiet, make-up gain can be applied to the signal just prior to output. Make-up gain shifts whole signal equally to bring it back up to the required level.

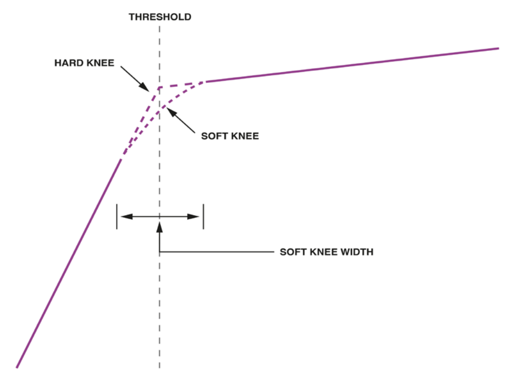

Compressor knee

In this way, a signal can be processed with the right output level while avoiding distortion (or clipping, in a digital workflow) caused by rogue peaks which exceed the maximum level the equipment can handle.

Limiting

A compressor becomes a limiter when a minimum ratio of 10:1 (remember that?) is applied. This high ratio effectively compresses the signal so hard that it's as if it stops it dead. In broadcast a limiter is often applied at the last stage of the signal processing chain to help ensure that the signal is within loudness regulations prior to broadcast, and to prevent it from being squashed by an automatic broadcast limiter prior to going to air.

Close The Gates

Now we understand the compressor it's easier to understand how an expander works because it's technologically similar to a compressor (while being completely the opposite).

Rather than reducing the dynamic range by reducing louder sounds, an expander increases the dynamic range of the signal. There are two types of expansion: upwards expansion - where signals above the threshold are increased (making loud sounds louder), and downwards expansion - where signals below the threshold are reduced (making quiet sounds quieter). Most expanders found in broadcast systems are downwards expanders.

A gate is like an extreme downwards expander because it filters out quiet parts of the signal and allows just the louder sections to pass through.

A threshold is set in the same way as with a compressor, but now signals which are above the threshold will be sent to the output of the gate. Signals below the threshold level will be attenuated (reduced) by a set amount of dB.

Attack and release settings work the same as a compressor. Gates also have a hold (or delay) setting to determine how long the gate stays open. Some gates have a hysteresis control which sets two thresholds; one for the opening level of the gate and the other for the closing level. This is usually controlled using an offset from the opening threshold to the closing threshold – which is generally set a few dB lower and is relative to the open threshold – this can help prevent oscillation if a gate is opening and closing rapidly.

An expander does a similar job to a gate, but instead of muting the signal below the threshold, it uses ratio settings in the same way as a compressor. It works in the same way, so a ratio of 2:1 means that for every 1dB below the threshold, the signal will be reduced to 2dB, and so on.

Expander/gates are super useful for cleaning up signals by filtering out unwanted low-level background noise from the signal such as crowd hubbub or machinery, or for removing any background spill which is caught on a nearby microphone, which can often happen with using multiple mics on a drumkit, for example. Reducing rather than completely removing this noise with a gate can produce a more natural sounding result.

Getting It Right From The Start

Getting the dynamic range right at the input stage helps protect against signal overloads later in the signal chain and is an essential element of the mixers' tool kit. In our next section we'll look at equalization (EQ) which is used to manipulate specific frequency components to shape the sound for aesthetic or correctional applications.

Equalizers (EQ)

By Kevin Emmott. The Broadcast Bridge.

EQ is one of the central tools of the audio production process and with a modest amount of knowledge and practice, a little can go a very long way to improving the subjective quality of a broadcast.

In addition to everything else a sound operator has to deal with, critical listening is a constant. Critical listening is not just about knowing when an audio source sounds bad – we can all do that – but is about identifying why it sounds bad, which part of the signal is causing the problem and how it can be fixed.

Equalisation (EQ) is one of the tools that sound engineers use to help fix these issues, and while modern mixing equipment can show where a signal is misbehaving with beautifully precise EQ curves and detailed visual representations, most audio operators don't mix with their eyes. That's why they have ears.

As we have seen , dynamics help keep signals confined within a dynamic range, but EQ can enhance those signals to compensate for external factors such as external noise and challenging room acoustics, and to provide more context to the output for specific content. In essence, it corrects nasty artefacts and aesthetically shapes the audio to the broadcast output.

Used in combination with dynamics, these two tools can fix the vast majority of challenging input sources. In the good old analogue days, an input signal would work its way down a channel strip through EQ and into dynamics. These days, modern digital consoles or workstations provide lots of ways to combine EQ and dynamics, with plug-ins or external hardware inserts providing even more options.

Unsurprisingly, changes to EQ will affect a signal's dynamics and vice versa, and there is a great deal of spirited debate over which should come first (spoiler: both sides are correct).

But to understand how they affect each other, let us start with what EQ is, and how it is applied.

Everybody Hz

Frequency is measured in Hertz (Hz). As adults, every sound we hear sits on a frequency range between 20Hz (low) to 20kHz (high).

EQ is used to adjust the character and tone of a sound by boosting or attenuating (reducing) specific frequencies in this range. Expensive-looking hi-fi's in the 1980's often had graphic equalisers which did this. Like supercharged bass and treble controls but with more slices. Graphic equalisers provide a series of controls (usually sliders) with each one affecting a narrow band of the frequency spectrum. The design is intuitive, controls on the left affect low frequencies and they work their way over to those on the right controlling high frequencies.

These controls don't add anything to the actual signal – they simply allow a listener to change the emphasis of frequency ranges within what already exists. Broadcast equipment goes beyond what a graphic equaliser can achieve by throwing even more controls at it. With a parametric EQ (PEQ) instead of having one control per band we are given a small set of adjustable controls.

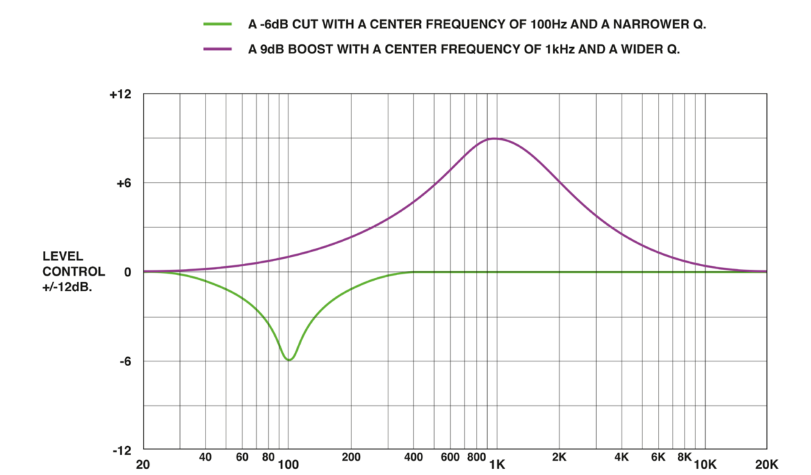

FREQUENCY SCALE FROM 20Hz TO 20kHz. THE SCALE IS LOGARITHMIC, NOT LINEAR.

The frequency control allows us to select the centre of the frequency band to be adjusted - it can usually be set anywhere between 10Hz to 20kHz.

The gain control determines the amount that the frequency will be boosted or attenuated. This is measured in decibels (db).

Because these adjustments will also affect some of the frequencies either side of the chosen value, the bandwidth (also referred to as "Q" for quality) control determines how wide this band will be. Rather than measuring this as a frequency range, the Q setting is measured in octaves; this is because the frequency range we hear is logarithmic rather than linear, so moving the bandwidth up and down the frequency range as an octave measurement adjusts the filter logarithmically and maintains the musicality of the filter across the whole spectrum.

Using a lower Q setting makes the band wider, influencing a wider range of frequencies either side of the centre value and will produce a more natural, subtle sound. A higher Q setting produces a narrower band, which influences a much narrower range of frequencies either side of the centre value.

Ding Dong

Even though sound engineers use their ears to mix, there are some fabulously descriptive names to help operators visualise what is going on. The EQ filter we have been describing is known as a bell curve, because it looks exactly like a bell. How lucky is that?

With a very narrow Q setting, the bell becomes very slender to become a notch filter. A notch filter affects fewer of its surrounding frequencies, and although this sounds less natural it is incredibly useful to sound operators.

Notch filters are used to remove noises which occur at a particular frequency. When combined with a high attenuation, its narrow bandwidth can pinpoint and entirely remove problematic frequencies while having minimal effect on surrounding frequencies. A popular use of a notch filter is to remove microphone feedback, or to filter out noise from a mains supply at around 50Hz/60Hz. They can also be used to reduce sibilance within vocals on those troublesome S's which tend to occur around 5kHz to 8kHz.

High Pass

As well as controlling troublesome frequencies, EQ is also used to make sonic improvements to a signal. The explosion in podcasting and the affordability of consumer recording technology means that the internet is at bursting point with tutorials on how to EQ the voice, and vocals can be hugely improved with very minor adjustments.

A high pass filter is used to reduce or remove low frequency noises on a signal, and is often applied to vocal tracks to simply remove information which is not required by attenuating low frequencies while allowing high frequencies to pass though. As the frequency range of a human voice starts at around 100Hz, setting a filter at 100Hz can remove a lot of unnecessary low-end noise such as air conditioning machines, traffic and generators.

At the top end of the scale a low pass filter attenuates high frequencies while allowing low frequencies to pass through. In music production an engineer may keep more of the high end to provide some brightness, but in a broadcast environment, with a reporter on location, it is often unnecessary and the application of a low pass filter can clean up the signal by once again removing information which is not adding anything to the broadcast.

For both pass filters, the Q setting switches to a slope to determine the cutoff rate. In the same way as a bell curve, the gradient is measured in dB (for gain) and in octaves (for frequency).

Enjoy Your Shelf

A shelf is similar to a pass filter, although shelves can also be used to boost a frequency signal as well as reduce it. A low shelf will affect frequencies below the curve at the lower end of the spectrum, while a high shelf affects frequencies at the higher end of the frequency range. They are more subtle than a pass filter by attenuating the sound below the curve rather than removing it entirely.

Which Came First, The Chicken Or The EQ?

While EQ and dynamics work well together, there are no rules as to what that relationship should be.

Like everything in broadcast, treatment will depend on both the audio source and the motivation for the output. Every live signal is unique and unpredictable, and there are multiple factors which will determine how they are treated.

Even simple things like Microphone placement can make a huge difference to the level and frequency content of the signal, while whispering, shouting, and singing all change the level and frequency spectrum of the input. Different microphone types play a big part, such as ribbon mics with their high frequency roll off and more pronounced lower frequencies if the subject is closer to the mic.

All this means that processing ordering must fit the workflow and what the broadcast demands of it.

With an 'EQ before compression' model, boosts or cuts to the EQ will change the level of the signal going into the compressor, and an engineer may have to keep adjusting the threshold of the compressor. If a source is going to need EQ tweaking over time due to changes in the surrounding environment, this might not be a good option. Tweaking the EQ after the compressor gets rid of this issue, but means that artifacts present in the incoming signal can drive the compressor in ways that aren't helpful.

These relationships are complex and sound operators may often choose to correct issues with EQ pre-compression, and sculpt the sound for aesthetic reasons post-compression. In live broadcast the EQ will often be used to give some gravitas to a voice, or to notch out an annoying low-end rumble or mains feedback.

It is easy to fall into the trap of assuming that EQ is about creatively boosting the frequencies which sound good. It can be better to approach EQ as a subtractive, corrective tool - removing any unwanted or harsh frequencies to reveal and therefore emphasize the desired, pleasing frequencies. Once unwanted artifacts have been sculpted away, sparing use of frequency boosting usually works well.

Ultimately, while complimenting each other beautifully, EQ and compression do very different jobs. Perhaps a better way to look at them is not as such a linear process, but to assess whether an EQ-ed signal needs any compression, or whether a compressed signal needs any EQ.

You know, by critically listening to it. Either way, the constant is that the sound engineer is always listening.

When Is A Sound Good?

By Christian Scheck. Head Of Marketing Content at Lawo.

When asked what "good sound" means to them, each audio engineer will give you their take on what really counts.

Unsurprisingly, opinions tend to differ quite a bit as there is no one-size-fits-all. And this is precisely what makes it so interesting to listen to mixes by respected audio engineers with different aesthetics and angles.

Live Performances

While they all agree that a mix should first and foremost serve the song and sound like all sonic elements it contains are there for a reason, the way to get there tends to vary. Obviously, studio mixes can be tweaked down to the last dB of EQ correction and revisited in case this change does not produce the intended result. A live context, on the other hand, mainly focuses on sound reinforcement that requires a slightly different approach and gives you far less time to perfect your mix.

Also, any tweak you make has an immediate impact on the sound the audience hears. Unless, of course, the desk has a function called "Listen Sense" (or something similar) that allows operators to adjust the channel parameters "offline", using headphones in PFL mode, without the audience noticing what is going on behind the scenes. Once the sound engineer is happy with the result, they can apply the changes to the live mix the audience hears.

But we are digressing here. It turns out that it is much easier to agree on what a good mix should *not* contain: the low end should not be boomy, the high end must not be ear-piercing, the sonic image should not be "hollow", and all parts need to sit well in the mix without acrobatic gain riding that may ruin the overall balance in a variety of respects. Also, the frequency spectrum needs to be adjusted in such a way as to avoid that the venue resonates back at you in an unpleasant way. This is far from easy, because an empty venue does not behave in the same way when it is packed.

All of the above require signal processing: using EQs to focus on the important frequencies of each signal source, dynamics to ensure that all signals remain audible without dominating the sonic image when they shouldn't, panning to improve signal separation and widen the sound image, and effects for sweetening. A pinch of punch may also be highly welcome.

What furthermore counts for experienced sound engineers is a console's basic sound, i.e., the one with no processing at all. What happens when you throw all relevant faders open at sound check? Is the sound muddy, or muffled, at first? Do you automatically reach for the EQ and boost the 2~4kHz range, because you know from experience that this is the only way to get a relatively "natural" sound?

Some consoles on the market already "sound great" right off the bat, allowing you to focus on finetuning critical aspects within seconds rather than minutes. This would be something to look out for before deciding on a new desk.

Another aspect is dynamics: what is on offer and, more importantly, how does it (alter the) sound? Should the compressor be almost inaudible at high ratio settings? Can you get it to not pump out of control? To what extent does the limiter color the sound, which may require additional EQ'ing? How accurately does the noise gate respond? And what about the expander, the de-esser and the possibility to easily implement parallel compression?

As always, quality comes at a price, which is usually worth paying, because the console has a service life of at least ten years. Besides, a mixing desk worth its salt offers a host of other features that are equally indispensable. Even though some audio engineers have taken to mixing even live events in the box, i.e., on a computer, a physical mixing console cannot easily be replaced with an inexpensive fader panel that can be purchased for a song. The ability to perform certain tweaks simultaneously in different sections (VCA grouping, bussing, dynamics, EQ, panning, etc.) is still an exclusive feature of "real" consoles.

There is a third way where a console manufacturer decides not to include any modulation, reverb, delay, etc., effects at all, because what sounds perfect to one audio engineer may not be what the next favors. Leaving room in the user interface for convenient, touchscreen-based, sweet-spot control of whichever plug-ins an audio engineer decides to use stands the best chance of satisfying all users. This approach furthermore has a positive effect on the desk's purchase price.

Signal processing is too delicate a matter to skimp on bread-and-butter signal processing features every operator immediately falls in love with, or simply expects. Let's therefore get the basics spot-on first. Allowing users to choose their preferred brand of effects plug-ins—and to change their minds without having to replace the console—looks like a very mindful and sustainable approach. For only you know what good sound means to you.

Divide And Rule

Much of the above also applies to the broadcast world, of course. The workflow is slightly different, however, because standards-based SMPTE ST2110 IP networking plays a much bigger part. How so?

First of all, a broadcast operation often has several consoles for as many audio control rooms. While for global events and other high-profile productions the channel count occasionally goes wild, with in excess of 200 sources that need to be mixed, most projects involve between 30 and 128 channels. Lawo's A__UHD Core, the DSP engine for the mc² console family, accommodates up to 1024 fully-featured DSP channels, which may look like overkill.

Yet, users do not need to activate all channels the engine can muster. A flexible licensing system allows them to start out with 256 channels and then add more slices of 256 as their needs evolve. The underlying philosophy of this licensing system furthermore sparked a second approach: one A__UHD Core can be shared by 4, 8, 16 or even 32 mixers. Not all of them need to be mixing consoles, by the way. It is perfectly possible to control one or several "slices", each of which has its own routing matrix and mixing console peripherals and is operationally completely independent, from a software UI (a so-called headless mixer) or using Ember+ commands. In such a setting, 1U of processing power is enough for up to 32 audio and/or master control rooms, which is good news for your energy consumption and rack space requirements.

IP networking also allows for workflows where one or several consoles are in different parts of the world and still access a single A__UHD Core, which can sit in yet another location. This is frequently used in live settings, often involving 5.1.4 Dolby Atmos immersive audio productions. The console is operated in city A and controls the DSP engine in city B anywhere in the world.

Another consideration for consoles—whether used for live performances or in broadcast—is how much can be automated, and how. AutoMix is most helpful for situations where several speakers debate on-set or co-host a show. Keeping their levels in check and allowing prioritizing one over all others makes the life of a sound engineer much easier. A welcome feature of this functionality is that inactive microphones are muted or attenuated in such a way that the overall ambience remains natural.

AutoMix can be used for any signal, from mono and stereo to multiple immersive channels, to minimize background noise and crosstalk with reduced sound coloration. Truncated sentences and late fade-ins are things of the past, enabling the sound engineer to focus on overall balance and sound quality.

A Downmix function for the creation of stereo renditions from immersive mixes guarantees perfect conversions into amazingly authentic sound images using just a few parameters.

In combination with the KICK software, which has been mandatory in the Bundesliga for several years, and a compatible tracking system, Lawo's AutoMix function is even able to create larger-than-life ball-kick, referee and moaning noises by mixing the signals of as many as 24 microphones placed around the pitch. With a view to a clean immersive audio result, this function needs to work in such a way as to keep the ambience level constant and phase-aligned to avoid nasty coloration, while always emphasizing the important noises on and off the pitch.

Almost all TV productions involve several cameras, and directors are known to request frequent cuts among them. How does the sound engineer cope with that, as the audio that goes with the footage is very likely contributed by different microphones assigned to other audio channels than the ones they had been working on until the cut? The answer is called Audio-follows-Video. It provides automated transitions and the perfect coupling of image and sound. It is based on assigning each camera's tally to an event that can be selected for one or more channels, with a total of 128 available events. Parameters such as Rise Time, On Time, Hold Time, Max Time and Fall Time can be used to set the processing envelope in order to create amazingly smooth and natural-sounding transitions from camera to camera.

It's All About the Sound

You may have noticed that all the bells and whistles that make the audio workflow smoother in a busy broadcast production environment focus on achieving a natural sound with a constant ambience level free from phasing and unpleasant level jumps.

With a Lawo console, you are always firmly in control. And don't be surprised if you are complimented on your sound.

Noise & Signal Repair

By Kevin Emmott. The Broadcast Bridge.

Understanding where noise creeps in and how to minimize it are key audio skills but sometimes, inevitably, modern noise reduction tools are a lifesaver.

Shhhhh!

As we have already learned, noise reduction is a lot of work for just one pair of ears.

Understanding the dynamic range of a signal and recognizing the frequencies within it is a critical part of shaping its output. We've already talked about how to refine the effects of noise interference by applying dynamics and EQ, techniques which can be developed and honed over time and which are useful tools to control noise and ensure compliance.

Here's some more good news; there are other real-time tools that you can employ which will do a lot of the heavy lifting for you, straight away, with no messing around.

It's A Kind Of Magic

You don't have to be an engineer to understand the benefits of automatic noise reduction. Anyone who has tried noise-cancelling headphones knows what a huge difference it makes. Turning on that switch on a tube or a plane and hearing all the extraneous noise automatically sucked out of the air is like listening to a magic trick.

Automatic Noise Cancelling (ANC) on headphones is pretty simple – tiny microphones sample the ambient noise and your headphones invert the phase to cancel out any outside noise. Voila. Instant immersion.

While automatic noise reduction for live broadcast isn't that simple, there are a variety of hardware and software plug-ins which take the edge off, but it is probably best practice to give yourself a head start.

Adding Noise

As we know, lots of things add noise. Environmental artifacts like aircon, traffic, wind and lighting; electrical interference from power lines; even the equipment adds noise, like mic self-noise and cable interference.

Some of these things are unavoidable, and efficient planning and filtering can help mitigate some of the repeat offenders, but they all add to what is known as the noise floor. The noise floor is the accumulation of unwanted noise which is inherent in the signal; the higher the noise floor, the more difficult it will be to distinguish the quieter elements of the audio we actually want to hear.

Analogue equipment introduces far more noise into the signal chain, from its own electrical components or from artifacts which are added as it travels down the input cables. Digital processing is far more forgiving, but the truth is that all electronic equipment produces some noise and every piece of equipment and additional process in the signal chain will add more and more of it.

The noise floor is always going to be part of the signal and sensible audio engineers aim to keep the noise floor as low as possible from the start. So now is probably a good time to talk about gain staging.

Gain staging

Gain staging is about achieving the highest possible signal-to-noise ratio on any given input. In other words, minimising unwanted noise while maximizing dynamic range and headroom.

It involves setting the input and output levels of each device in the signal chain. This means making sure there is enough headroom at each stage to ensure it doesn't distort further down the line. Combining multiple signals together can also push the limitations, so leaving enough headroom on each individual one is an important consideration.

Due to its nature, gain staging for analogue equipment is more about hiding the noise floor by boosting the input to maximise the signal, and the tendency has always been to set the input levels on analogue equipment higher to minimise the amount of inherent noise which can be heard.

In a digital system, engineers tend to use a different approach. Although a digital signal will likely have a lower noise floor, digital audio also has an absolute level of 0dBFS (decibels relative to Full Scale – we covered this in more detail when we looked at metering in part one of the series). While pushing the limits of an analogue signal can be a stylistic choice, at 0dBFS digital audio runs out of quantizing levels, distorts quite horribly and should be avoided at all costs.

Noise Reduction & Signal Repair

Even with the best start, signals will still degrade; mic clipping, wind noise, lossy compression, poor microphone technique, analogue-to-digital conversion, crowd noise, electrical interference, the relentless pace of modern life…even in the most controlled environments we are never actually in control.

Real-time noise reduction and signal repair products give engineers an opportunity to actively reduce or eliminate such noises during live broadcasts at the touch of a button. Each will depend on the characteristics of the noise and the desired output, so we're back to asking what the broadcaster is trying to achieve and using the right tools for the job.

They may also affect the overall tonal balance, but while there is always a trade-off between how effective noise reduction is and the effect on the overall signal, this can often be mitigated by using the right equipment in the right part of the signal chain.

Available as both hardware and software plug-ins, these products can solve specific issues. Dedicated hardware processors tend to be more common for real-time noise reduction and signal enhancement in live environments as they offer very low-latency processing.

Meanwhile, software plug-ins for DAWs are used mainly in post-production to clean up for TV or film.

Let's start with some simple fixes.

All About The De

Many signal repair functions have features which have helpful names which describe exactly what they do, which is nice. De-essers, de-poppers, de-noise, de-rustle and de-wind functions all do exactly what you think they might do.

Software like this works by learning where in the frequency spectrum the problem exists and isolating those frequencies with the goal of providing greater intelligibility to voices.

De-essing simply reduces excessive sibilant sounds found in vocals by attenuating frequencies in the sibilant range. It does exactly the same thing as a notch filter to reduce sibilance around 5kHz to 8kHz, but it does so at the push of a button.

Similarly, de-popping can reduce the intensity of “p” and “b” sounds caused by bursts of air hitting the microphone, while algorithms to remove clicks and crackles can remove noises caused by electrical or digital errors in the audio signal.

In audio post environments, many of these tools are bundled together in software packages which can be accessed in a DAW and remove noise, clicks, hums, sibilance and other artifacts automatically by analysing an audio file and applying a combination of fixes.

Keeping It Live

Noise reduction for live production has the same motivation to preserve dialogue, but it usually occupies a different place in the signal chain. It works by using a combination of algorithms and hardware components to detect noise and apply appropriate filters.

Rack-mounted units are physically located in the ACR or in an outside broadcast truck, although 12v portable units can reduce ambient noise at source on location with a reporter.

Either way, real-time noise reduction hardware is usually placed before the dynamics/EQ processing to minimise the amount of manual remedial work. If there is any automix functionality being applied to the channels – such as might be used to automatically duck microphones on a panel/discussion show – it would also be applied beforehand so that each channel is as clean as possible prior to joining the automix bus; it is far easier to clean the channels before this process rather than it having to react to frequently changing output levels.

It also needs to be lightning fast; not only can it be used to clean up signals for the mix, but it might also be used for in-ear monitoring (IEM); for example, to remove mic spill on an IEM feed for backing vocalists. For these reasons, near-zero latency is key.

Sometimes you might want some background noise; a cheering crowd may add some colour and atmosphere to the broadcast. To enable this, most noise reduction units will have separate attenuation and bias controls to provide more control. An attenuation control allows the engineer to control how much the noise is attenuated, which allows them to keep some noise in if it adds to the broadcast.

Conversely, a bias control affects how much influence the tool has on the signal, and many manufacturers will also provide the opportunity to apply these controls across specific frequency bands to allow fine tuning around each voice.

Adapting To The Environment

This technology is improving over time, with algorithms able to calculate changing levels in background noise and apply noise attenuations at different frequencies to optimise suppression. The ability to adapt to changing environments and apply changes in real time is beneficial wherever it is not possible to control the wider environment (so, you know, everywhere).

Live dialogue noise suppression is not more important than it used to be, but it is more common. With more OTT and OTA channels, podcasts, digital channels and other on-demand services we've never had as much content to choose from. More content in more locations with more people and more distractions.

We can't control our environments, but with some knowledge about how audio works, and some help from clever tech, we can at least exert some control over our content.

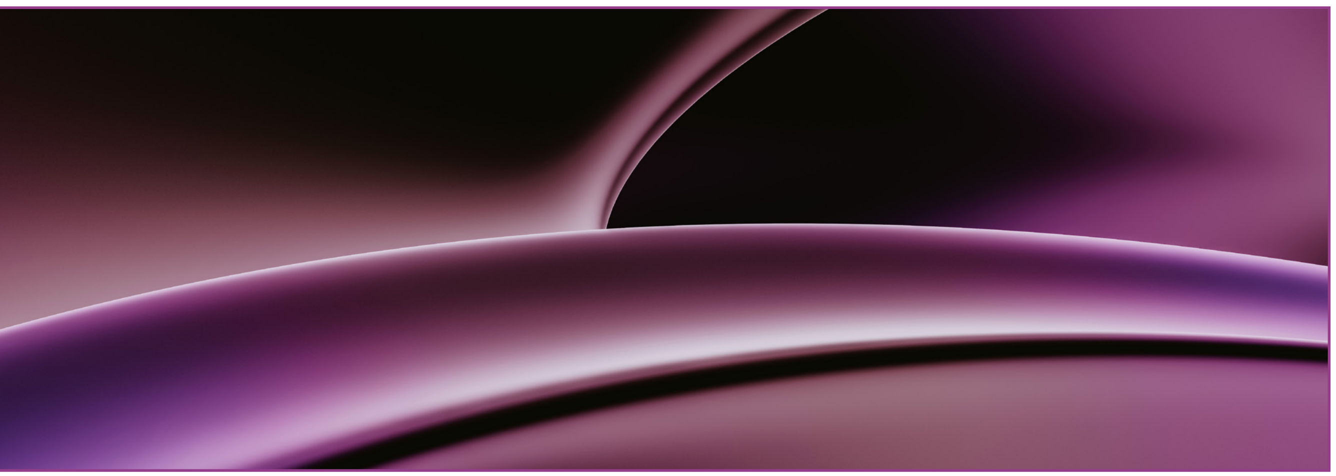

Themed Content Collection

EG ESSENTIAL GUIDES

ARTICLES

MEDIA

The Broadcast Bridge publishes hundreds of free articles like this every year.
Subscribe online to stay informed via email or our social channels.

www.thebroadcastbridge.com