


CALREC AUDIO PRIMER



Basic introduction to Calrec functionality and applications

Published in association with



Keeping audio professionals up to speed in the fast-changing world of sound and image.

Calrec Audio Ltd

Nutclough Mill
Hebden Bridge
West Yorkshire
England UK
HX7 8EZ

Tel: +44 (0)1422 842159
Fax: +44 (0)1422 845244
Email: enquiries@calrec.com

calrec.com

No part of this manual may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and scanning, for any purpose, without the prior written consent of Calrec Audio Ltd.

Whilst the Company ensures that all details in this document are correct at the time of publication, we reserve the right to alter specifications and equipment without notice. Any changes we make will be reflected in subsequent issues of this document. The latest version will be available upon request. This publication is for International usage.

Calrec Audio Ltd reserve the right to change specifications without notice. E & O.E.

The established policy of Calrec Audio Ltd. is to seek improvements to the design, specifications and manufacture of all products. It is not always possible to provide notice outside the company of the alterations that take place continually.

Despite considerable effort to produce up to date information, no literature published by the company nor any other material that may be provided should be regarded as an infallible guide to the specifications available nor does it constitute an offer for sale of any particular product.

Alpha, Sigma, Omega, Zeta, Hydra Audio Networking and Bluefin High Density Signal Processing (HDSP) are registered trade marks of Calrec Audio Ltd. **Dolby®E** is a registered trade mark of Dolby Laboratories, Inc. All other trade marks are acknowledged.

© 2018 Calrec Audio Ltd. All Rights Reserved.

AUDIO PRIMER

CONTENTS

Introduction	5		
Introduction	7	Surround panning	39
		Precise Positioning with LCR	40
		Using a Surround Microphone	41
		Downmixing	42
Digital Audio Consoles	9		
Inputs, Busses & Outputs	10	Mix-Minus	43
Assignability	11	Mix Minus	44
Layers	12		
Running Levels and headroom	13	GPIO - Control Logic	45
VCA-Style Groups	14	GPIO (Relays and Optos)	46
Metering	16		
Redundancy and Hot-Plugging	17	Appendix A: Basic Audio Theory	47
		Audio Waves	48
Equalization	19	Simple Sine waves	49
Concepts of Equalization	20	Phase and Polarity	50
Filter Types	21	Decibels	52
Shelf and Bell Curves	22	Equal Loudness Contours	53
Q	23	Octaves	54
		Frequency Response	55
Dynamics	25	Polar Plots	56
Compressor	26	MS Processing	57
Limiter	28	Delay	59
Gate	29	Basics of Digital Sampling	60
Expander	31	Digital Clocking	62
		Binary numbers	64
Signal Flow	33		
Signal Flow	34		
Processing Order	35		
Surround, Pans & Downmixing	37		
Working in Surround	38		

AUDIO PRIMER INTRODUCTION

INTRODUCTION

This guide is intended to be a basic overview of audio theory, digital processing, applications and techniques relating to the broadcast industry, tailored specifically for the current range of Calrec consoles.

Explanations of basic audio concepts are covered only to provide a grounding in order to understand more advanced concepts and applications. There are a multitude of high quality textbooks available on the subject of audio theory and this document is not intended to be a replacement for any of these.

Explanations of certain Calrec features are described in this document in order to separate them from operational content which can be found in the Operator Manual relevant to your product. The Operator Manuals describe the methods to access and control these features, whereas this manual describes them in a results and application oriented manner.

This document uses a certain level of abstraction in order to make the descriptions independent of any one product and as such you should consult the relevant operator manual to confirm which features or processes apply to your product.

AUDIO PRIMER

DIGITAL AUDIO CONSOLES

INPUTS, BUSSES & OUTPUTS

Audio mixing consoles are made up of a number of input channels, busses and outputs. A bus is a term used to describe a path containing multiple signals, ie multiple input channels mixed together. A console may have many busses with different names such as Mains, Tracks, Auxes. These are all of equivalent audio quality and any can be outputted as program content, however the features available on each bus type may vary. For example a Calrec Main output may have EQ and dynamics whereas an Aux may not.

Audio group busses are used as an intermediary stage between inputs and outputs. Multiple inputs can be routed to a group, which then allows them all to be controlled together, eg by adjusting the group fader, all the inputs routed through that group will be adjusted. Groups also offer features similar to channels so amongst other things, EQ and dynamics can be adjusted to affect the audio from multiple channels. The group itself is then routed to the output rather than the individual channels which are routed to the group.

Technically, groups are busses, but not outputs, although if required a Calrec group can be outputted using its direct output. All input channels and groups have direct outputs which are not busses as they output only the audio from that channel or group (even though the group direct output contains audio derived from a bus). Direct outputs have independent level controls, so a return of a single input can be sent and adjusted without affecting the program mix. If using group direct outputs as a substitute for a bus output, care needs to be taken to ensure the direct output level is set correctly.

An important operational concept of digital audio consoles is assignability.

Traditional analog console design features “channel strips” where controls for each channel are laid out in a column above each channel fader. Controls are repeated across the surface for each channel.

A digital control surface does not need to have an individual control for each parameter. Surface controls can change their function to reflect the current mode, for example rather than needing a full set of EQ controls on each fader, you can have one EQ section capable of making adjustments to any fader.

As controls can adjust parameters on any fader, they need to be focused on the correct path before adjusting. On a Calrec console, each fader has an assign button, selecting this makes that fader the focus of the assignable control panels. This is referred to as the “currently assigned fader”. Selecting an assign button on a different fader deselects the first one and will refresh all the assignable panels to display and allow changes of settings on the newly selected fader.


Benefits

Without needing a set of controls for each fader, more space is available to display information about controls and settings.

Various assign controls can be positioned on the surface to allow the operator to change parameters from a central location. The obvious benefit of this is being able to remain in the sweet spot of the monitoring setup while retaining total control of the whole surface.

The flexible layouts of the Apollo and Artemis systems allow each operator (there could be up to three on each surface) to configure where their assign

FIGURE 1 - COMMON CONTROLS



have “Wild” controls dedicated to each fader. The function of the wild control can be set to control almost any parameter, so for example if you want an input gain (or any other control) above each fader that can be adjusted without assigning the fader, this is still possible.

If a knob or button breaks, it affects all channels / faders - a good design allows that control can be adjusted elsewhere, such as from the PC Application. The Apollo / Artemis design uses standard assign panels that can be in any mode, so if there is a problem with the panel normally used for EQ, a different panel can be selected to be in EQ mode.

controls appear. This means that each operator can have their own assign controls directly in front of them wherever they are working on the surface.

Disadvantages

You have to select the fader before making the adjustment from the assignable panel. This very quickly becomes second nature to digital console users and such concerns disappear after a little hands-on use - some consoles

LAYERS

Layers provide a method of organizing paths on the surface. Each layer stores a different arrangement of paths which can be called up onto the surface at any time.

Figure 1 shows one way to think about layers. It shows 6 different layers, each having a different arrangement of paths on the faders. Switching to a different layer would place this new arrangement of paths and faders onto the surface.

Figure 2 illustrates a simple example of switching layers. Imagine a small console with eight faders. When layer 1 is selected the arrangement of paths A to H are called up and assigned to the faders. The faders now control paths A to H and the levels of the paths can be set as desired.

If layer 2 is selected a different arrangement of paths (I to P) is called

FIGURE 1 - LAYERS





FIGURE 2 - SWITCHING LAYERS



up and assigned to the faders. The fader positions will update to represent the levels of the newly assigned paths and adjustments to these new paths can be made.

As the paths assigned to the faders have changed, paths A to H will now not be affected by and fader movements, but audio will still pass through these paths and through the console.

Selecting layer 1 again causes paths A to H to be assigned to the faders. The faders will update to represent the current levels of paths A to H and allow alterations to these levels. Paths I to P will still pass audio through the surface but will not be accessible for control on the current layer.

Paths assigned to layers that are not visible on the surface will still pass audio through the console if routed and left open.

RUNNING LEVELS AND HEADROOM

All systems have a maximum signal level that they can take in or output. Exceeding the maximum can result in distortion on analog signals and unwanted ticks or clicks, or even silent samples on digital signals.

The difference between the normal operating level and the maximum is referred to as headroom.

Calrec, as specialists in the live broadcast field, have always held high headroom as a fundamentally important design principle to maintain audio quality during unexpected peaks that cannot be controlled.

In the analog domain, signals are often measured in dBu. The standard reference or "line-up" level is usually 0 or +4dBu. Program material is mixed around about the line-up level and the maximum signal level is determined by the design of the electronics. Calrec line level inputs and outputs have a maximum of +28dBu. Calrec microphone level input stages offer up to 36dB of headroom.

In the digital domain, signals are normally measured in dBFS. FS stands for "full-scale" and 0dBFS is the maximum level that can be passed.

Reference level

The normal operating, or "line-up" level chosen should be well below the maximum of any equipment it is to pass through, to ensure the signal is not clipped or distorted in any way if there are unexpected peaks in signal. The lower a signal's level however, the closer it is to the noise floor and the audio quality is reduced, so a high enough reference level is chosen to ensure high quality clean sounding audio with enough headroom to cope with the unexpected.

To ensure uniformity in output, standard reference or operating levels are usually defined by region.

On a digital system, the operating level is chosen as a figure below the maximum of 0dBFS that gives sufficient headroom yet is well above any noise floor. Almost all digital consoles will have analog inputs and outputs, and therefore also need to know the analog reference level being used so that analog inputs and outputs are converted and passed on at the appropriate levels.

Two common reference levels are -18dBFS=0dBu and -20dBFS=+4dBu. These are used in the UK and the US respectively and are also common in other countries, however several other standards are used around the world. This often complicates international programming when signals are sent between countries

Mic input headroom

As a microphone is most likely to be the source of unexpected high signal levels, Calrec designs provide 36dB of headroom through the input stage.

Calrec offer the option to set a lower figure for the mic input headroom. Doing so improves the signal to noise ratio, however in reality, high quality audio designs have a very low noise floor and reducing the mic headroom from maximum has negligible, certainly in-audible impact on noise performance.

The running level and headroom settings should be standards that do not change, however they can be adjusted at any time by settings in software without any hardware changes.

VCA-STYLE GROUPS

VCA-style groups provide an alternate way to control multiple fader levels from a single fader.

Non-VCA groups

'Normal' audio groups are a signal processing path - Audio from multiple inputs is bussed and routed through the audio group where the signals can be adjusted as a whole before routing the audio on to outputs. Being a signal path, there is a finite number of group busses available.

VCA groups


VCA stands for "Voltage Controlled Attenuator", A VCA group is not an audio path, it is a way of adjusting multiple faders at once. Members of a VCA group are referred to as Slaves whilst the fader controlling the whole group is referred to as the Master. Paths assigned to a VCA group are not bussed together, they remain independent and need to be individually routed to outputs. Adjusting a VCA group master fader will adjust the level of the slave faders within that group. A VCA master cannot affect the routing or signal processing of its members, only the fader level. Being a linking of control data, rather than a signal path means the number of VCA groups available is not limited.

Figure 1 shows the effect that moving a master fader has on its' slaves. It also demonstrates that the balance of slave faders within a VCA group can be altered and this new balance will be reflected and updated when controlled by the master.

Note that in some designs, the slave faders do not physically move when the master is adjusted, so that their relative mix is still visual and can be adjusted when the master is closed or is sending slaves past their endstop positions. On such systems there is LED indication to

show that the true level of a slave fader is higher or lower than the physical position. Although there are benefits to the slaves not moving, some users prefer to see them move when the master is adjusted as this gives a clear indication of which faders are in which groups and reassures that the correct audio is being adjusted.

FIGURE 1 - VCA EXAMPLE



Multiple destinations

Because the slaves of a VCA group are not bussed together, the audio of multiple paths can be controlled by one fader even if they are routed to different outputs. This is useful, for example when producing a "full" mix and an "effects only" mix simultaneously on different output busses.

Masters on used faders

Calrec digital consoles allow VCA masters assigned to blank, unused faders or to used faders. Assigning a master to a used fader automatically adds the path on that fader as a slave to the group. The benefit of this is that one of the faders within the group can be the master, rather than having to use up an extra fader.

It is however important to understand the nature of the fader being used as the master. On creating a VCA group, the master fader snaps to the neutral 0dB position. If the master is on a fader with a path assigned, that path is a slave to the group, but adjusting the fader adjusts the master, not the slave. If the slave fader was off when the group was created, it remains off even though the master is at 0. Normal slave faders can be adjusted, however to access and adjust the slave under the master, its' fader assign button should be pressed and held. Display should change to indicate you are now controlling the slave and the path fader can be adjusted.

VCA groups can span A & B faders as well as faders on different layers. It is important to ensure all slaves, including any hidden under the master are set to the correct position.

Intermediates


Intermediates are slaves of masters which are also masters to other slaves. For example, a master (Master A) could be set to be a slave of another master (Master B) meaning it now has a master and slaves and is therefore an intermediate. Alterations to Master B will affect its slaves, including Master A, which will in turn affect Master A's slaves. Figure 2 illustrates this master-intermediate-slave hierarchy.

Cut, AFL and PFL

Slaves in a VCA group will not only follow the master's fader level. They will also follow the status of any Cut, AFL or PFL switches assigned on the master.

For exact product implementation and more details, refer to your product's operator manual.

FIGURE 2 - VCA MASTER-INTERMEDIATE-SLAVE HIERARCHY



There are many different types of audio meters used with varying scales, and reaction times. Audio signals can appear quite different depending on the type of meter being used.

Ballistics

The reaction time of a meter affects the level it displays. Meters do not react instantly with the audio feeding them. Consider a traditional meter with a physical needle that moves across a scale. The needle rises and falls to reflect audio level, it cannot instantly change its position. For example if a meter goes from displaying no audio, to displaying 0dB, it has to move from the off position to the 0dB point. Although a bargraph meter could instantly light up at the 0dB point they don't. Bargraph meters rise and fall in the same way as traditional moving coil meters. The rise and fall is also known as attack and decay. Together they can be referred to as the meters ballistics.

The two most common meter types used in broadcast are known as PPM or VU. In moving coil form, a PPM meter would look very different to a VU meter as they have different scales and background colors. In bargraph form, it may not be so obvious.

PPM meters have a very fast attack time and a very slow decay time. This means the meter can display fast, transient peaks in audio.

VU meters have a much slower attack, but a faster decay. Whilst VU meters do not show fast peaks, they do represent the average audio level

When comparing the two types of meter, if feeding a constant level such as tone, both meters will display the same level. With transient audio, or "program material", a PPM meter will generally display several

dB's higher than a VU meter due to the difference in their ballistics.

Peak-hold and true-peaks

Bargraph meters often have a peak-hold function. Peaks points in level remain displayed for a period of time whilst the rest of the bargraph decays. Traditionally the peak spot conforms to meter ballistics and displays the highest point the meter has reached. Some meters display peaks with different ballistics to the rest of the bargraph, for example a VU bargraph showing peaks in PPM.

Digital audio can contain much faster transients than analog that do not register on PPM or VU meters. Calrec digital bargraphs show "true peaks". These can be considerably higher than the program content displayed by the meter. Its important to be aware of true peaks, reaching 0dBFS even for only one sample can cause unwanted artifacts in the audio. Different equipment reacts differently to full scale audio. Whilst the console output may sound OK, devices further down the chain may be having a more serious affect.

Its always a good idea to be able to meter and listen to audio from further downstream than the console output to ensure your signal is being delivered as intended.

Phase

As well as showing level, meters can be used to show phase relationship. Phase between two channels is often measured on a scale of -1 to +1. With no audio present a marker is displayed at the central 0 position. The marker moves into the positive region and displays green when signals are in phase, and will move to the negative and display red when signals are out of phase with each other. The phase relationship between two

signals can also be shown on a vector scope which displays in phase content vertically and out of phase content horizontally.

REDUNDANCY AND HOT-PLUGGING

Redundancy is a term used to describe the use of additional hardware or software to automatically take over in the event of a failure.

Calrec systems provide extensive redundancy to minimize single points of failure throughout the system.

Processing hardware

Redundancy of processing is usually achieved by having “hot-spare” or backup hardware fitted. For example, in a Calrec processing rack, there are two Master Controllers, two DSP cards and two routers. On bootup, the system boots on the “primary” cards. The secondary cards boot into a standby mode. If a primary is not fitted, is removed, or fails, the secondary of that type will activate and take over. This is known as a “hot-swap”. A well designed redundant system covers all critical parts of the design and will recover from even the most serious failures quickly with minimal disruption. When planning redundancy, whole communication paths should be considered, rather than individual hardware elements to ensure that a breakdown in any part can be recovered by rerouting data and bringing hardware online. Recovery should be sought by changing as little hardware as possible to reduce the disruption, ie rather than switching to a whole secondary path and changing lots of active hardware, data should be re-routed around the specific problem area. This also retains the redundancy for the rest of the system allowing for further failures to be recovered. Experiencing multiple hardware failures without having time to address them is extremely unlikely, however well designed redundancy gives added confidence to users.

Power supplies

The use of dual power supply units protects against the failure of PSU's and can protect against the loss of AC mains power.

It is common practise to have two separate AC mains power distribution paths within a facility, and for all technical equipment with two PSU's fitted to have one connection from each so systems are maintained if either mains source is lost. Unlike other redundancy, there generally isn't a “hot-spare” PSU. Both PSU's are normally active and sharing the load. If one fails, the loading on the remaining PSU doubles.

Hot-plugging

Hardware that can be fitted or removed whilst the power is on is termed as hot-pluggable. In its most basic form, this means the hardware will not be damaged, for example by plugging a card into a powered card frame.

A hot-pluggable item should not affect the rest of the system when fitted or removed. For example, fitting/removing an I/O card in a modular I/O frame should not interrupt or affect the audio or control of any of the other I/O cards fitted in that rack.

Ideally hot-pluggable items should boot up when connected, automatically sync to the rest of the system and become active without disruption or the need for a reset. For example, a Calrec control surface panel can be replaced without interruption, even if it is a fader panel with audio on open faders going to air. The replacement panel will boot up, sync, and become active in an identical state to when it was removed with no affect on audio or the rest of the system's control.

Each element of the system boots independently and automatically syncs with the rest, therefore control elements (including the Master Controller in the rack) can be replaced without interruption to the audio and vice versa.

Secondary cards can be removed or fitted without any disruption and replacements should come online as hot-spares.

Removal of an active processing element, audio or control will result in a hot-swap to the secondary. Removing active (not hot-spare) audio processing, for example DSP or Router cards, can cause serious splats on audio outputs and should be avoided whilst outputs are active.

AUDIO PRIMER EQUALIZATION


CONCEPTS OF EQUALIZATION

Equalization is a process used to affect the frequency spectrum of an audio signal. Certain frequency components can be boosted or reduced to shape the sound for aesthetic or correctional applications.

Figure 1 shows an example EQ response. In this example plot, the blue line shows the 'shape' that will be applied to the frequency response of the input signal. At the output, the signal will have a cut in the mid range at about 300Hz and a boost applied to all frequencies over approximately 2.5kHz. All frequencies below around 50Hz are rolled off. It is important to be able to understand these plots as it aids greatly in both setting EQ and reading what effect the curve will have on the signal.

The Y axis shows the boost or cut level in Decibels (dB). This is a linear scale. The X axis shows the frequency, measured in Hertz (Hz). Low frequencies are shown at the left and rise over a logarithmic scale to the highest frequency at the right. A logarithmic scale is often more appropriate than a linear scale in this case due to the way that octaves work.

FIGURE 1 - EQ PLOT EXAMPLE



FILTER TYPES

High pass

Put very simply a high pass filter attenuates low frequencies while allowing high frequencies to pass unaffected. A frequency domain plot of a typical high pass filter is shown in Figure 1.

Below the cutoff frequency the signal is rapidly attenuated. The cutoff frequency is measured as the point at which the signal is attenuated by -3dB. In Figure 1 this can be seen to be roughly 100Hz. A small range of frequencies above this point are attenuated.

The rate that the signal is reduced below the cutoff point is determined by the slope of the filter. The slope is a gradient measured in dB per octave. The higher the number the steeper the slope and the greater the reduction.

This type of filter is frequently used to remove low frequency rumble from a source.

Low pass

This filter is the reverse of the high pass. The cutoff frequency remains the point at which the signal is attenuated by -3dB, however it is the signal above this point that is rapidly decreased (Figure 2).

Low pass filters are commonly used to keep unwanted high frequencies out of a woofer feed.

Notch

The notch filter has a very narrow bandwidth and high attenuation (Figure 3). It can be used to pinpoint and remove problematic frequencies while having very little affect on the frequencies around it. The notch frequency is the frequency at the center of the curve. In Figure 3 this is approximately 1kHz.

FIGURE 1 - HIGH PASS FILTER




FIGURE 2 - LOW PASS FILTER





FIGURE 3 - NOTCH FILTER



Common uses include removing the high frequency whistle produced by CRT television sets, or removing problematic resonant frequencies.

SHELF AND BELL CURVES

Low shelf

A low shelf response allows the signal below the corner frequency to be boosted or attenuated. The corner frequency is defined as the point at which the signal falls to 3dB below the boost or gain value. In Figure 1 this point would be around 160Hz. This -3dB point is not strictly correct for all applications (there is no -3dB point for a boost of 1dB for example), however it does serve as a good guide.

The slope of the shelf is measured in dB per octave. The higher the dB value per octave, the steeper the slope and the less the frequency range between the unaffected signal and the signal affected by the applied gain/attenuation.

High shelf

A high shelf is the opposite of a low shelf, boosting or attenuating all frequencies above the corner frequency.

Bell

A bell curve boosts or attenuates its center frequency, and a range of frequencies around it. The range of surrounding frequencies it affects is determined by its Q value. Generally a low Q spans a larger range of frequencies and results in a more 'musical', or pleasant sound. A narrow Q affects fewer surrounding frequencies but can tend to sound less natural, especially at higher gain settings.

FIGURE 1 - LOW SHELF




FIGURE 2 - HIGH SHELF





FIGURE 3 - BELL



The term **Q** relates directly to the bandwidth of a bandpass filter. At a simple level it describes the ratio of the center frequency to the bandwidth.

The Q of a bandpass filter is the center frequency divided by the difference of the upper and lower -3dB frequencies (Figure 1).

FIGURE 1 - EQUATION FOR Q


$$Q = \frac{f_c}{f_{\text{High}} - f_{\text{Low}}}$$

From this equation it can be seen that a filter with a low bandwidth would have a high Q value. A wider bandwidth would have a lower Q value.

When setting the Q value of a filter, for example on an equalizer, the bandwidth would be set in octaves rather than a discrete frequency range. This means that if the same filter is moved across different octaves, its bandwidth, in terms of frequency range in Hertz adjusts accordingly. In lower octaves the bandwidth in Hertz would be small compared to a high octave which spans a greater number of frequencies. So, the 'musical' affect of the filter remains the same across the spectrum.

A Q of approximately 1.4 has a bandwidth of one octave.

FIGURE 2 - BANDWIDTH



AUDIO PRIMER DYNAMICS

COMPRESSOR

A compressor allows the dynamic range of a signal to be reduced. This can be used for technical control to prevent signal overloads or maximise signal intensity, or as a creative tool to shape the dynamic response of a sound.

Threshold

A compressor has a threshold control which sets the level at which the compressor affect the signal. In the opposite way to a gate, any signal below this threshold is unaffected. When a signal exceeds the threshold it is reduced in level (Figure 1). The amount of reduction is set by the ratio.

Ratio

When a signal exceeds the threshold its input to output ratio is altered. Below the threshold the ratio is 1:1. In other words, for every 1dB of level at the input of the compressor, there will be 1dB at the output. When the threshold is exceeded the ratio changes, for example to 2:1. This means that for every 2dB of input over the threshold there will be 1dB of output above the threshold. A ratio of 20:1 means that for 20dB of input over the threshold there will be 1dB of output above the threshold. See Figure 2.

Attack and release

The attack parameter of a compressor sets the amount of time it takes for full gain reduction to be achieved once the threshold has been exceeded. A very short attack time will reduce the signal almost the instant it exceeds the threshold. A longer attack will allow some of the initial transient to pass unaffected before the gain reduction is applied.

The release time defines the amount of time the gain reduction is still applied after the signal has dropped back below the

FIGURE 1 - GAIN REDUCTION




FIGURE 2 - RATIO





FIGURE 3 - ATTACK AND RELEASE TIMES



threshold. Figure 3 demonstrates attack and release times.

Knee

The knee parameter determines the range around the threshold in which gain reduction will be applied. Figure 4 is a representation of knee values.

A hard knee setting means that gain reduction is not applied to a signal until it hits the threshold. As soon as it does, the full amount of reduction is applied. This is a very precise method of applying gain reduction but can sound harsh in some applications.


A soft knee setting will start to apply gain reduction before a signal hits the

threshold, but will not apply full reduction until the signal reaches a certain level above the threshold. This results in a smoother application of gain reduction.

Make-up gain

When a signal has had gain reduction applied, portions of its output will obviously be lower level than its input. In situations such as peak prevention this may be the desired effect. In other situations the output level after reduction may be too quiet and so a make-up gain is required. This simply applies gain to the signal before output.

FIGURE 4 - KNEE



Compression effects

By applying different parameter combinations, a very wide range of effects can be achieved with compression.

Two distinct differences are shown in Figures 5 and 6. Figure 5 shows an input signal affected with a high threshold and a high ratio. In this case, only a few peaks in the signal exceed the threshold and so are the only parts of the signal to be attenuated. The relative levels between peaks and dips has changed in the output signal and is a typical example of peak control.

Figure 6 shows an input signal affected by a low threshold and a low ratio. Here, just about all of the signal is affected and reduced by the same ratio. In this example the peaks and dips in the output signal remain relative to each other compared to the input signal, but the overall level of the signal is lower. Make-up gain could be applied to bring the highest peaks in the output up to match the highest input peaks. This would give the effect of raising the level of the quieter sections rather than turning down the louder sections.

FIGURE 5 - HIGH THRESHOLD, HIGH RATIO





FIGURE 6 - LOW THRESHOLD, LOW RATIO, WITH MAKE-UP GAIN



LIMITER

A limiter prevents a signal from exceeding a certain threshold.


Much like a compressor, a limiter only affects portions of an input signal that exceed a set threshold. A limiter differs from a compressor in the fact that any input signal that exceeds the threshold is reduced to the level of the threshold rather than being reduced by a certain ratio.

Figure 1 shows an input signal before it is affected by a limiter. It then shows the signal at the output of a limiter if it was set with zero attack and zero release. In this case, any signal that exceeds the threshold is immediately reduced to the threshold level. As there is zero release time the limiter also stops immediately. This results in a clipped signal, almost square wave in shape.

The second example shows how increasing the release time can make the limiting less aggressive. As the signal ramps back up to its original level rather than instantly jumping back, the gain does not have to be reduced as much to bring the level down to the threshold level. There is however still an abrupt change in level at the onset of limiting.

By increasing the attack time, this onset clipping can be removed at the expense of allowing a short period of the signal to exceed the threshold.

FIGURE 1 - LIMITER EXAMPLES



A gate filters out quiet portions of signals and allows only louder sections to pass through. They are useful for filtering out unwanted background noise and cleaning up signals.

Threshold

Only signals which are above a certain threshold will be sent to the output of the gate. Signals below this level will be attenuated by a large amount, effectively making the output zero (typically around -80dB). Figure 1 is a representation of this concept. Only the upper portions of the two large peaks in the signal will produce audio at the output, the rest will effectively be silence.

Hysteresis

A hysteresis control effectively sets two thresholds, one sets the opening level of the gate, the other sets the closing level. The hysteresis value is the negative offset in dB from the opening threshold to the closing threshold.

Without hysteresis, if a signal oscillates rapidly around the threshold a chattering effect can occur due to the gate opening and closing rapidly. See Figure 2 for a graphical representation of a gate with the hysteresis control applied. The left side shows the signal before being affected. The right side shows the signal with the gate and hysteresis applied.

Attack

Gates have an attack parameter which sets the amount of time the gate takes to open. When an input signal exceeds the threshold the gate will open with a ramp from full attenuation to no attenuation, the duration of which is specified by the attack time. A short attack time will help to preserve the natural transient attack in the input signal but can sometimes result in a clicking due to the rapid

FIGURE 1 - THRESHOLD




FIGURE 2 - THRESHOLD WITH HYSTERESIS





FIGURE 3 - ATTACK EXAMPLES



transition. A long attack time produces a smoother fade transition but will lose some transient information. Figure 3 shows a representation of some of attack parameters. The black triangles below each graph show the duration of the attack, and the level of the audio signal (from maximum attenuation to original strength).

Delay (or Hold)

If a gate is open and the signal falls below the threshold again (including the hysteresis value if it is switched in), the delay time determines how long the gate stays open before it starts to close.

Release

The release time is the opposite of the attack time. After the hold period, it takes the release time for the attenuation to ramp up to full again and the output signal to return to silence. A longer release leaves a tail out at the end of the signal for a smoother, more natural transition. Figure 4 is a representation of various release times. The black triangles showing the duration of the release and the power of the signal (from original power to full attenuation).

Depth

The depth control determines the amount of attenuation that is applied to signals below the threshold. Although these signals are attenuated, they still maintain their 1:1 input to output gain ratio. Figure 5 shows some examples.

Envelope

Figure 6 describes the gate process through an envelope diagram.

FIGURE 4 - RELEASE EXAMPLES




FIGURE 5 - DEPTH EXAMPLES





FIGURE 6 - GATE ENVELOPE



EXPANDER

An expander performs the opposite function to compressor. Instead of reducing the dynamic range, it increases it.

It is set in a similar manner to a compressor and has similar controls. The main difference is that it affects signal below the threshold rather than above it.

Signal above the threshold has an input to output ratio of 1:1. Signal below the threshold is given an input to output ratio set by the ratio control. Figure 1 illustrates this concept.


Ratio

Consider the ratio 3:1. For every 1dB below the threshold, the signal will be reduced by 3dB. For a very high ratio of 50:1, for every 1dB the signal is below the threshold it will be attenuated by 50dB. This high setting effectively turns the expander into a gate.

Range

The range control, like the gate, determines the amount of gain reduction applied at the selected input to output ratio before the ratio returns to 1:1.

FIGURE 1 - EXPANDER RATIOS



AUDIO PRIMER

SIGNAL FLOW

SIGNAL FLOW

The signal flow describes the route a signal takes through the components of a system, from input to output.

Figure 1 shows a very basic, static route through an audio console. A signal appears at a console channel input, is processed by the channel input controls (gain, phase etc), then proceeds through various processing components before it is sent to the output of the channel.

Inserts

An insert in a signal chain allows the signal to be sent out of the standard signal chain and sent to another external device for further processing. After the required processing has occurred, the signal is returned to the same point in the original chain and continues with the remaining processing.

Figure 2 illustrates this process. If the insert is switched in, the signal is output via the insert send, follows the red dotted path and has external processing applied. It is then returned to the original signal processing chain via the insert return.

If the insert is switched out, the signal would pass from the insert send straight back into the insert return and continue in the chain, even if equipment is still connected to the physical I/O sockets assigned to the insert.

Auxiliary Sends

If a signal is sent from a channel's aux send the signal is effectively split. One leg of the split signal will continue through the original channel's flow from input to output. The other leg of the split signal will leave the original flow at the aux send and will most likely be routed out of the console. This will then have external processing applied and possibly returned to the console via the input of a different channel, where it will start a new flow.

FIGURE 1 - BASIC SIGNAL FLOW




FIGURE 2 - INSERT SIGNAL FLOW





FIGURE 3 - AUX SEND/RETURN FLOW



PROCESSING ORDER

Certain processing components can be positioned at different places in the signal flow.

On a Calrec console there are multiple positions in the signal flow that certain processing or routing can be placed, for example pre-EQ, post-EQ, post-fade, pre-fade etc. These positions will be detailed in the relevant operator manual. A simple example using pre and post-EQ is used here.

Take the dynamics component for example. In a Calrec console this component contains (amongst other processors) a compressor. In the controls of this compressor is an option for it to be placed pre or post-EQ. A setting of pre-EQ means that the signal will be affected


by the compressor before it is affected by the EQ. Post-EQ is the opposite with the compressor affecting the signal after it has been processed by the EQ.

Figure 1 demonstrates the possible effects of applying a limiter with a pre-set threshold both pre and post EQ. In the top line an input signal has its low frequency content boosted, resulting in a waveform with higher amplitude. If a limiter is applied to this signal, at the threshold shown in red, a large portion of the waveform will be affected. The resultant waveform can be seen to have been affected quite severely.

In comparison, the second line is limited at the same threshold before it has EQ applied. In this case only a few peaks in the waveform are affected. When

the signal has EQ applied the result is significantly different

FIGURE 1 - EFFECT OF LIMITING PRE AND POST EQ



AUDIO PRIMER

SURROUND, PANS & DOWNMIXING

WORKING IN SURROUND

With the increasing demand for multichannel surround content it is important that operators are able to easily visualize and manipulate surround signals.

Calrec consoles provide a range of features that make surround signals easily accessible and compact in terms of surface real estate.

Where a stereo signal is made up of Left and Right (L and R) components, a surround signal is made up of many more components. For a 5.1 surround signal these would be Left (L), Right (R), Center (C), Low Frequency Effects (LFE), Left Surround (Ls) and Right Surround (Rs).

Single fader control of surround channels

When surround signals are patched to a channel in the console, the overall level of that signal can be controlled by a single fader. Moving the fader up will increase the level of all components that make up the surround signal. Altering EQ or dynamics will apply that same processing to all components. The same applies to any routing made.


Spill faders

If control over any of the individual components is required, the surround signal can spill out onto the Spill faders. For a 5.1 signal these faders allow control over the L/R components (as a stereo pair), the C component, the LFE component and the Ls/Rs components (again as a stereo pair).

When using the Spill faders individual level changes can be made and processing applied to each component rather than the overall signal.

The Spill faders on a surface will display and allow manipulation of the signal

FIGURE 1 - SPILL FADERS



attached to whichever fader is currently assigned on the surface. If a different fader is assigned, the Spill panel will update to reflect the state of the signal attached to this newly assigned fader.

SURROUND PANNING

Surround panning requires two dimensions of positioning and therefore becomes more complex than simple one dimensional stereo panning.

Panning signals in surround can be broken up into three main positional aspects:

- Front pan
- Rear pan
- Front to rear pan

Front pan

Similar to stereo panning, front pan varies the position of a signal between the left and right outputs (Figure 1). If LCR panning is selected then the center channel is included in the front pan.

Rear pan

Similar to stereo panning, rear pan varies the position of a signal between the two rear (or surround) outputs (Figure 1).


Front to rear pan

Front to rear pan varies the position of the signal between the front pan position and the rear pan position. If both front and rear pan are central, the signal will be panned from front center to rear center. If the positions of the front and rear pan are different, ie front is panned left and rear is panned right, then varying the front to rear pan will move the signal from front left to rear right. Examples of this are shown in Figure 2.

Divergence

The Divergence control varies the amount of center channel information that is sent to the L and R channels. This in effect widens the center image reducing its point source nature. A divergence value of zero results in no center information sent to the L and R channels. As the value of the divergence control increases, so does

FIGURE 1 - FRONT AND REAR PAN



the center's contribution to L and R. A full divergence value means that no signal is present in the center, and is split fully between L and R. See Figure 3 for an illustration of this process.

FIGURE 2 - FRONT TO REAR PAN





FIGURE 3 - DIVERGENCE



PRECISE POSITIONING WITH LCR

LCR panning allows the positioning of stereo audio to become more accurate with the addition of a discrete center channel.

On a standard stereo output, the signal can be panned between the left and right speakers. Panning centrally sends the signal to both speakers at the same level. Panning a signal to the left increases the level sent to the left speaker and reduces the signal to the right speaker. When the signal is panned hard left, there will be no output from the right speaker.

Figure 1 shows the difference in left and right levels as the signal is panned across the stereo field. If the graph were to continue downwards, at the far left position the right signal level would reach a very low dB value and become effectively silent.

LCR systems

An LCR system extends standard stereo by adding a third center speaker (Figure 2). LCR provides two main advantages, one being enhanced stereo positioning, the other being separation of audio signals.

FIGURE 1 - LR SIGNAL LEVELS





FIGURE 2 - LCR SPEAKER SETUP



Separation of sources


Mono and stereo channels on a Calrec console have the option to be panned using standard stereo or LCR stereo. Standard stereo signals will only be sent to the left and right speakers. LCR signals have the option to also be sent to the center speaker. This means the certain important elements, such as dialogue, can be sent to the center speaker and separated from ambient or background sounds which can be placed in the left and right speakers.

Enhanced stereo positioning

In standard LR stereo, a signal panned centrally is created by sending the same signal to both speakers, creating a 'phantom' center image. Depending on the speaker setup and the position of the listener, it can be difficult to pinpoint the exact position of this image and it may appear wider than a point source. In LCR stereo, as the center signal originates from an actual speaker, the center image becomes much tighter and localization is much easier.

Panning between left and right also becomes more accurate as the signal

FIGURE 3 - LCR SIGNAL LEVELS



is panned between the center speaker and the left or right speaker. Looking at Figure 3 it can be seen that as the signal is panned right of center, the left speaker receives no signal at all.

USING A SURROUND MICROPHONE

Surround microphones provide a simple yet extremely comprehensive method of capturing and manipulating surround information.

The signal from a surround microphone should be fed into an appropriate decoder. This decoder will process the signals and produce discrete outputs for each surround component. In the example shown in Figure 1, a 5.1 surround signal is produced.


These discrete outputs can be sent to the Calrec console and assigned to the inputs of a surround channel.

The surround channel will allow control and provide processing for the surround signal as a whole, for example EQ and dynamics can be applied to all of the surround components simultaneously. This surround channel also allows control over properties including as front pan, rear pan, front-back pan and divergence.

To gain more specific control over an individual component of the surround signal, for example the center component or the LsRs component, the surround spill panel can be used. Parameters of the individual components such as level, EQ, dynamics and APFL can be adjusted here.

This arrangement promotes a very flexible method of working with surround sources. It allows the operator to switch from simple global level control of the whole surround signal, to precise component level manipulation very easily.

FIGURE 1 - SURROUND MICROPHONE INTERFACE EXAMPLE



DOWNMIXING

Downmixing is the method of converting a signal of a given width, into another signal of a narrower width.

For example, a broadcaster may output all their programme audio in 5.1 surround format. If a consumer has access to a 5.1 audio system they can enjoy the full benefit of the surround audio. If another consumer has only a stereo audio system and simply mapped the L and R components of the surround signal to their left and right speakers, they may miss key audio information from the center and surround channels. For example in a film, dialog may appear mainly in the center channel, or in a sporting event the crowd and ambient noise may appear mainly in the surround channels. Figure 1 illustrates this problem.

To solve this problem downmixing is used to convert a wide signal into a narrower signal. Surround signals can be downmixed to stereo and stereo can be downmixed to mono. There are many more complex downmixes that can be made but the theory remains the same.


Stereo to mono

In a very simple stereo to mono downmix, the left and right signals can simply be summed to produce a mono version. There are important mix considerations to ensure that the resultant signal is mono compatible, most importantly phase coherence. For more details on this see the section of this document on Phase.

Surround to stereo

When downmixing a surround signal to a stereo signal, the center and surround signals can be added to the left and right channels to ensure that all the required information is present.

FIGURE 1 - WITHOUT DOWNMIXING



Depending on the type of programme, the contribution of surround and center channels to each of the left and right channels may vary.


For example in a sporting event there may be pitch, commentary and crowd audio in the front channels and more ambient crowd noise in the rear channels. Simply summing the surround channels to the left and right channels may result in an overly loud crowd noise which obscures the commentary or pitch audio.

Downmix levels

To solve these potential problems, individual downmix levels can be set that vary the contribution of each leg of the surround signal to the stereo downmix.

The sports event issue presented earlier can be solved by giving the surround channels low downmix levels, reducing their contribution to the stereo signal and effectively reducing the amount of crowd noise present.

FIGURE 2 - WITH DOWNMIXING



AUDIO PRIMER

MIX-MINUS

MIX MINUS

Mix Minus is a system that allows a comprehensive mix to be sent to multiple listeners each receiving the complete mix, minus their own input.

This is often achieved by using up multiple output busses for each send and ensuring that the routing to each contains everything required in the program mix, without routing the audio that has come from the destination it is to be sent to. This can instead be achieved using a mix-minus system, freeing up output busses and simplifying routing.

Figure 1 shows an abstraction example of a mix minus system. Assume that the sources surrounded by a blue box represent people, either presenters or field reporters. The other sources may represent VT or server audio feeds. These sources are fed to the input of eight channels on a console. Each channel has

the option for its signal to be sent to the console wide mix minus bus.

This mix minus bus can be routed back out to any sources that may require foldback, for example the presenters or field reporters. Each source would be fed the entire mix-minus bus signal, with their own contribution to that mix removed, creating a unique feed.

In Figure 1 the mix sent back to source 1 would consist of sources 2-8. Source 2 would receive a mix of sources 1 and 3-8 and so on...


Why remove a sources own input from its' foldback mix?

There are many reasons to do this. One relates to field reporters, or presenters in studios communicating via long distance systems such as satellite links.

The reporter would need to hear a mix of the show's audio in order to hear cues and communicate with the presenters. The inherent delay in these systems means that it may be a number of seconds before the audio reaches the reporter. It can be very difficult to speak while hearing your own voice returned to you with even a slight delay. By using a mix minus feed to the reporter, their own contribution to the audio is removed before it is fed back to them, eliminating this problem.

Another reason is to eliminate feedback. If a presenter's foldback was monitored on loudspeakers without subtracting their own contribution, some of their original signal may be picked up by the microphone again, thus creating a feedback loop. By using a mix minus system to remove the presenter's own contribution, this feedback loop is broken.

FIGURE 1 - MIX MINUS SCHEMATIC EXAMPLE



AUDIO PRIMER

GPIO – CONTROL LOGIC

GPIO (RELAYS AND OPTOS)

GPIO, “General Purpose Inputs & Outputs”, often referred to as simply GPI - “General Purpose Interface” are simple on/off “logic” connections, rather than audio connections.

Relays and optos open or close under the control of a separate electronic circuit. They allow simple control voltages to be sent between Calrec consoles and other items of equipment.

Opto-isolated Input

An opto is made up of a light emitting device and a light sensitive device combined in one component with no electrical connection between the two. When the component receives an input signal it lights up the emitter. When the receiver detects this light it reacts accordingly. All general purpose inputs (GPI) on Calrec consoles are opto isolated.

Relay

Unlike an opto, a relay does create an electrical connection from input to output. Contacts are made and broken mechanically. General purpose outputs (GPO) on Calrec consoles use relays.

Applications


Calrec consoles can control external signals when certain buttons are pressed or faders are opened. Figure 1 shows some possible applications of GPI/O signals.

Alternatively, other items of equipment can send signals back and activate certain functions on a Calrec console.

Latch and Pulse

When a function is activated, or alternatively de-activated, the type of control signal sent from a GPO can usually be set to one of four options.

FIGURE 1 - APPLICATIONS



Similarly to GPO, GPI receive either latch or pulse signals, depending on the function being controlled.

If the output is set to 'Latch', the relay opens and passes a constant signal at a constant level until the function is de-activated and the relay closes.

Setting the output to 'Pulse' opens the relay for a duration of 100ms then closes it. The relay can be set to pulse when the control message is activated, de-activated or both.

AUDIO PRIMER

APPENDIX A: BASIC AUDIO THEORY

AUDIO WAVES

The sounds we hear are periodic vibrations propagating through a medium (commonly air) which are translated by our ears into useful information.

A source creates a sound wave by oscillating to create movement in the matter of the surrounding medium. A speaker for example oscillates back and forth displacing the air around it.

Figure 1 demonstrates the effect of a speaker oscillating in a sinusoidal movement. As the speaker moves in the positive direction it forces air in that direction. Applying pressure to the molecules in the air moves them closer together and is known as compression. After travelling a certain distance in the positive direction, the speaker slows down and starts to move in the negative direction. This in turn causes the air to move in that same direction and causes a drop in pressure, or rarefaction, and moves the molecules apart. The top line in Figure 1 illustrates this effect.

The two plots at the bottom of Figure 1 compare the velocity of the speaker and the pressure placed on the air in front of it. As the speaker is in the middle of its travel from negative to positive and moving at its highest velocity, it exerts the most pressure on the air in front causing a compression. As it slows down to a standstill it exerts no pressure. During its travel from positive to negative it exerts a negative amount of pressure on the air in front causing a rarefaction.

Speed of sound

The speed of sound varies depending on the medium it is travelling through. The denser the material the quicker sound propagates through it. In air, the material of most interest to audio engineers the

FIGURE 1 - AUDIO WAVE PROPAGATION




FIGURE 2 - WAVELENGTH FORMULA

$$\lambda = \frac{v}{f} \quad f = \frac{v}{\lambda}$$

speed of sound is approximately 340 meters per second.

Wavelength

It may be useful in certain situations to know the physical wavelength of a given frequency, for example when performing acoustical calculations.

The formula for finding the wavelength given the frequency and speed of a sound is shown in Figure 2. The inverse formula for finding the frequency from the wavelength and speed is also given. Speed (v) is required in meters per second, frequency (f) is in Hertz and wavelength (λ) is in meters.

SIMPLE SINE WAVES

Any sound can be broken down and represented by a number of simple sine waves of varying frequency and amplitude. In the opposite way a complex sound can be reconstructed from a similar collection of simple sine waves.

The mathematics and theory involved in these processes is beyond the scope of this document, however it is important to have an understanding of simple sine waves which form the basis of all signals and many of the examples given in this document.

To understand the shape of a sine wave when represented on a plot against time (or angle), think of a circle as shown in Figure 1. Starting at zero degrees trace around the circumference of the circle as shown by the black indicator. By plotting the vertical displacement of the black indicator against its rotation around the circle in degrees, a sine wave shape is created.

The fifth plot in Figure 1 shows that after a full rotation has been completed and the indicator has returned to the same position as in the first plot, rotation around the circle continues but the angle travelled by the indicator exceeds 360 degrees. The second time around the circle however the sine wave pattern remains the same.

Cycles

As the same pattern occurs again and again, a sine wave can be said to be periodic. The degree values around the latter three circles in Figure 1 show that a rotation of 360 degrees is in effect the same as a 0 degree rotation. Similarly a rotation of 450 degrees is in effect the same as a 90 degree rotation. A single complete rotation around the circle is called a cycle. A sine wave is normally

FIGURE 1 - SINE WAVE SHAPE




FIGURE 2 - FREQUENCY FORMULA

$$f = \frac{1}{t} \quad t = \frac{1}{f}$$

described as number of cycles, each from 0 - 360 degrees.

Period

The time it takes for one cycle to be completed is called the period. The period can be measured from any point on a sine wave to the next occurring identical point.

Frequency

Given that a sine wave will repeat the same cycle over and over again, it is possible to deduce the frequency that these cycles will occur. If a cycle is completed every second, we can say that the frequency is 1 cycle per second. If a cycle is completed every millisecond we can say that the frequency is 1000 cycles per second.

The common unit of measure of frequency is Hertz (Hz). 1 Hz equates to 1 cycle per second. 1000Hz equates to 1000 cycles per second. The formula for converting a cycle period to frequency is shown in Figure 2, where t is in seconds. The reverse formula for converting back from frequency to cycle period is also shown.

PHASE AND POLARITY

Phase and polarity are very common and different terms which are all too often used interchangeably. There are very important differences between the two.

Phase

In terms of audio waves phase describes the offset, measured in degrees, that a wave is delayed from its starting position. It is not noticeable when a single wave is played

Figure 1 shows two sine waves, one of which (red) has been delayed by a varying number of degrees.


In the first example the red wave has been put 90 degrees out of phase with the blue wave. If the signals are summed together they produce a wave approximately 45 degrees out of phase from the original but with a slightly larger amplitude.

In the second example the waves are 180 degrees out of phase with each other. When the red wave is at maximum amplitude, the blue wave is at minimum amplitude. In theory summing these waves together would result in a wave with zero amplitude.

In the third example, the red wave has been delayed by 360 degrees, in other words it has been delayed by a whole cycle. The two waves are perfectly in phase, just as they would be had there been no delay applied at all. Summing these waves results in a perfectly phase coherent wave of double amplitude.

The fourth example shows the red wave with a delay of 540 degrees. As discussed in the sine wave section of this document, 540 degrees is equivalent to 180 degrees, so the same result occurs as in the second example.

FIGURE 1 - TWO SINE WAVES PUT OUT OF PHASE BY VARYING DEGREES



Because the red wave is being delayed by a degree value, which relates directly to a position in the cycle, the frequency of the waves is irrelevant. As long as they are both the same, these waves could be of any frequency and the results would remain the same.

Delaying by time


If there are two sine waves of the same frequency and one is delayed by a time value (for example 1ms) instead of by an amount of degrees relative to a cycle, the frequency of the waves determines the amount that they will be in or out of phase.

Figure 2 shows some examples of this behavior. A 1kHz sine wave has a period of 1ms and so if it is delayed by 1ms it remains perfectly in phase with the original wave. If a 500Hz sine wave with a period of 2ms is delayed by 1ms it ends up 180 degrees out of phase with the original.

Multi-frequency and non periodic waveforms

The concept of phase does not transfer simply to a signal made up of more than one frequency. It is easy to see how two simple sinusoids can be out of phase but

FIGURE 2 - PHASE VARIES WITH FREQUENCY FOR CONSTANT TIME DELAY



consider there are two identical signals, each made up of a 1kHz sine wave and a 650Hz sine wave. If the 650Hz component of the first wave is delayed, the 650Hz components of both waves are out of phase but the 1kHz components are still in phase.

Non-periodic waveforms are also problematic. If the signal was broken down into its simple sine wave components then each of these could be examined for its periodicity, however this still does not determine the overall period of the signal.

So how is it that a button on an audio console allows a signal of any number of frequencies or any non-periodic length to be put 180 degrees out of phase?

The answer is, it doesn't. It inverts the polarity of a signal.


Polarity

Inverting the polarity of a signal means to flip its amplitude so that the positive components become negative and the negative components become positive.

If a sine wave has its polarity inverted its plot can look as if it has been put 180 degrees out of phase which is possibly where confusion can occur, however a plot of a non-periodic waveform will clearly reveal this inversion.

Technically the phase reverse, or 180 degree phase button on a console should be labelled 'polarity inversion' but convention over time has led to this ambiguous use of the term phase.

FIGURE 3 - POLARITY INVERSION



DECIBELS

Decibels are the logarithmic unit of measure used to describe the ratio of one sound pressure level to another.

In terms of measuring sound level in relation to human hearing, a reference level is defined at approximately the threshold of perception. The decibel value describing the sound level is a ratio of this reference level and the sound level being measured.

The equation to calculate the decibel ratio in terms of sound pressure level (dB SPL) is shown in Figure 1. P2 is the reference value and P1 is the pressure level being measured.

FIGURE 1 - DB PRESSURE EQUATION

$$dB = 20 \log \left(\frac{P_1}{P_2} \right)$$


dB scales on the Calrec surface

When dB levels are used on a Calrec surface, they do not refer to the ratio of sound pressure level. Each scale has a slightly different application. Take for example the dB scale next to a fader as shown in Figure 1. The fader allows the operator to apply gain or attenuation to a signal.

When the fader is set at the 0dB position this means there is no gain or attenuation applied and so there is effectively 0dB difference between the input and output signals.

If the fader is set at the 10dB position below 0dB, this means that the output signal would be 10dB less than the input signal.

FIGURE 2 - SURFACE DB SCALES



In this case the dB ratio will not refer to sound pressure level, but to a generic dB ratio.

In the meter example in Figure 2, the level of the signal is measured in dBFS or dB full scale. For this scale the reference value is set at the highest signal level possible in a fixed point digital system. A signal's level is measured at a dB value below this reference, so the R signal in the diagram could be said to be at a level of ~-6dBFS, or 6 decibels below full scale.

EQUAL LOUDNESS CONTOURS

The equal loudness contours shown in Figure 1 represent the level at which pure steady tones across the frequency spectrum must be presented for the listener to perceive a constant loudness.

The unit of loudness is the phon. At 1000Hz, 1 phon is equal to a level of 1dB SPL. For a steady tone at any other frequency to have a loudness level of 1 phon, it must be perceived to be at the same loudness as the tone at 1000Hz.

Looking at the lowest blue line in Figure 1 (labelled as 10 phons) it can be seen to have a level of 10dB SPL at 1000Hz. This line gives a level of ≈ 20 dB SPL at 200Hz. This means that for a steady tone at 200Hz to be perceived at the same loudness as a steady tone at 1000Hz with a level of 10dB SPL it must have a level of 20dB SPL, 10dB higher.


Examining the 10 phons line again it is clear that the levels of very low frequencies need to be raised significantly to achieve the same loudness as the 1000Hz tone.

Working up to the 110 phons line it can be seen that the loudness contours start to flatten off at lower frequencies, meaning that less boost must be applied for them to have the same perceived loudness as the 1000Hz reference. The same can be seen at the higher end of the frequency range but with less pronounced effects.

So as the perceived loudness level increases it could be said that our ears become more sensitive to high and low frequencies (in relation to the 1000Hz reference).

This is an important notion to consider, especially when working in a critical listening environment. If, for example, you

FIGURE 1 - EQUAL LOUDNESS CONTOURS



had created an audio mix at high listening levels and decided that the overall spectral balance was as desired. You may find that when played back on a system at a lower listening level, the mix may be lacking in low and high end detail compared with other successful mixes.

The reverse could also apply. If a mix is created at very low listening levels, the high and low frequency components may have been boosted to compensate for the lack of hearing sensitivity at these low levels. When played back at an 'average' listening level the low and high components may appear excessively loud.

Some consumer products such as CD players or amplifiers include a switchable or variable 'loudness' control. This control compensates for the lack of sensitivity at low levels by emphasizing the high and low frequency components of the audio signal.

OCTAVES

An octave is the interval between a certain frequency and another that is half, or double it's wavelength.

For example, Figure 1 shows the wavelength and frequency of the note A3 (the A below middle C) and the next four A notes above it. It also contains the octave width, which is the number of Hz until the next A note is reached.

It is clear from the diagram that as the frequency increases, the frequency range of an octave also increases. Even though the range of frequencies in higher octaves is greater than that in lower frequencies, the human ear responds to frequency ratios rather than actual frequencies so still recognizes it as an octave.

Linear vs. Logarithmic scales

Look at a frequency range from 100Hz to 1000Hz plotted on line with a linear scale (upper line in Figure 2). As the frequency increases, the width of the octave on the plot would also increase. A logarithmic scale means that octaves on the plot have the same physical width across the frequency range (lower line in Figure 2).

Working with audio in octaves and viewing information on a logarithmic scale therefore, in most cases, produces a more 'musical' and useful representation than a linear scale.

FIGURE 1 - OCTAVES





FIGURE 2 - LINEAR VS LOGARITHMIC SCALES



FREQUENCY RESPONSE

The frequency response of a piece of equipment shows it's ability to capture or reproduce sound across the frequency spectrum.

An item of equipment can be measured by the application of a known test signal to it's input. The output of the equipment is then captured and the compared to the known signal.


A technically perfect or 'transparent' system would output a signal that is identical to the input system. In reality this is extremely unlikely for technical as well as aesthetic reasons. Technical reasons could include the efficiency of signal conversion, the physical performance of system components (particularly mechanical parts) or the quality of components used. Aesthetic reasons could include deliberate manipulation of the signal to achieve a certain character of sound.

The frequency response of the equipment being measured is detailed on a logarithmic plot in the frequency domain as shown in Figure 1. In this example, the frequency response of a loudspeaker has been measured using a specially designed measurement microphone with an extremely flat frequency response of its own.

The known measurement signal used here was a pure tone at constant amplitude, swept up through the frequency range from 20Hz to 20KHz. The software analyzing the measured signal compares the output level and the phase shift with the input signal and produces a graph of the results (Figure 1).

A reference can be used to see how the measured response deviates from the input signal and to produce a meaningful statement about the results. The red line

FIGURE 1 - FREQUENCY RESPONSE OF A LOUDSPEAKER



in Figure 1 represents the output from a perfectly transparent, and fictional, system.

In the frequency range from 70Hz-13KHz (highlighted in Figure 1), the measured signal deviates from the reference signal by approximately ± 4 dB. It has a gentle low frequency roll off from 90Hz down to 20Hz and slight peaks at 14KHz and 17KHz.

Typically manufacturers will quote a frequency response similar in structure to:

20Hz-20KHz ± 3 dB

A system with a 'flatter' frequency response would have a lower deviation from the reference.

It should be noted however, that a good frequency response does not necessarily translate to mean a good overall performance or sonic quality. It merely describes the equipment's accuracy in reproducing the frequency range of an input signal.

POLAR PLOTS

Polar plots aid in determining the directional frequency response of audio signal emitters or detectors.

Figure 1 shows an example of the directional response of a microphone in a single plane at two different frequencies.

Zero degrees on the diagram is the point directly in front of the microphone, referred to as 'on-axis'. The diaphragm of the microphone would be at the center point of the plot.


The outer most concentric circle represents the highest level that is picked up by the microphone. Each circle inside this represents a drop in the detected level.

The curves drawn on this plot illustrate the level of signal picked up by the microphone as the sound source is moved around it at a consistent distance. The solid curve shows the response of the microphone when presented with a continuous 1000Hz tone. It is clear that the rear of the microphone is much less sensitive to a signal of this frequency than the front. Alternatively it could be said that the on-axis response at 1000Hz is much greater than the response at 180° off-axis.

The dashed curve on the plot shows the response of the microphone when presented with a continuous tone at 125Hz. While the level detected at 180° off-axis is still lower than it is on-axis, the overall response is much more consistent than at 1000Hz.

In reality these responses are not limited to just one plane as shown here. The response of a piece of equipment should be considered in three dimensions.

FIGURE 1 - MICROPHONE RESPONSE




A reverse version of the plot shown here can be used, for example when working with loudspeakers. In this situation the plot would show the level of signal the device would emit in all directions at a certain frequency.

Certain response curves have specific names as shown in Figure 2. These names can help to quickly describe the response of a piece of equipment without going into detailed description or needing to see a diagram of the polar plot.

It can be said that a cardioid pattern has maximum response on axis and maximum rejection at 180° off axis.

FIGURE 2 - RESPONSE TYPES



M/S (Mid/Side or Sum/Difference) stereo provides a versatile stereo image with extended width while remaining 100% mono compatible.

M/S stereo is based around a matrix combination of two microphones (Figure 1). To achieve best results the capsules should be placed as near to each other as possible.

The first is a forward facing microphone, usually with a cardioid polar pattern although other patterns can be used for varying results. This microphone will pick up the discrete center M signal (mid or mono).

The second is a bi-directional microphone, usually a figure of 8 pattern, oriented on the lateral axis (90° off axis) which will pick up the S signal (side or stereo).

Resultant polar patterns

By summing the outputs of the two microphones patterns, a polar pattern as shown in Figure 2 is created. This is effectively the polar response of a hyper-cardioid microphone pointed around 60° off axis to the left.

Obtaining the difference between the two microphones results in a polar pattern similar to a hyper-cardioid microphone at around 60° off axis pointing to the right (Figure 3).

At this point it seems that all that has been achieved is a coincident pair, something which could have been more easily achieved using a pair of microphones set up in this manner in the first place. However the M/S pair gives the valuable addition of a discrete mono signal.

Decoding the matrix

As this mono and stereo information is contained within just two microphone

FIGURE 1 - M/S MATRIX POLAR PLOT




FIGURE 2 - SUM PLOT





FIGURE 3 - DIFFERENCE PLOT



signals it must be decoded in the correct manner. On a Calrec console this is as simple as assigning the M and S signals to a stereo path and then pressing the M/S button. The width control allows the width of the stereo signal to be varied.

This decoding works by sending the M and S signals to both left and right outputs. The S signal sent to the right output however, has its polarity reversed.

Assuming the M and S signals are at the same level, when the two left signals are summed, they result in what would have been picked up by the polar plot shown in Figure 2, the sum. When the M and inverted S signal in the right are summed, the resultant signal is what would have been picked up by the polar pattern in Figure 3, the difference. This produces a standard width stereo signal (Figure 5).

By varying the balance of the M and S signals the stereo width can be altered. When M is full and S is at minimum, both left and right receive the same signal, so the output is effectively mono. See Figure 4.

When S is at full and M is at minimum, the signals at the left and right outputs are perfectly out of phase resulting in a very wide, but unnatural sounding stereo spread. See figure 6.

Mono compatibility

Take for example a standard stereo signal recorded with a spaced pair of microphones. The signals captured at each microphone will never be identical due to the distance between them and the different directions they are pointed. This signal would lose information when played in mono as when the signals are summed, only identical signals remain at full power. Signals that are not identical may either

FIGURE 4 - FULL M, NO S




FIGURE 5 - EQUAL M AND S





FIGURE 6 - FULL S, NO M



be reduced or augmented depending on their phase alignment.

In M-S stereo the S signals sent to the left and right outputs are perfectly out of phase and will also cancel each other out if the signal is collapsed to mono, however the M component remains intact. This discrete center signal is the reason that M-S stereo signals are truly mono compatible.

Microphone considerations


A number of manufacturers produce M-S microphones containing both the M and S capsules in one body, providing each output from a separate feed. These microphones greatly simplify capturing signals using M-S stereo and often produce superior results due to the small size and close proximity of the capsules.

Delay displaces a signal by a given amount of time and can be introduced naturally by certain pieces of equipment, or correctively or creatively by the operator.

In a complex broadcasting system where audio and video are gathered from many different sources, there is a possibility that certain signals will become out of sync with each other. Every digital system introduces a certain amount of delay to a signal due to the analog to digital conversion, the internal processing, the transfer of information around the system and the resultant digital to analog conversion. The input to output delay inside a Calrec console is so small as to be insignificant, however other systems may produce issues.

In Figure 1, the top line represents a real life event which has a visual and audio component, for example a door closing. This event is recorded for a live show by cameras and microphones. If the visual signal is processed in a HD system, it is very likely that there will be a relatively large time delay from when the video enters the system, to when it leaves (see the second line). The audio system processes and outputs its signal with much less delay. When the signals are combined there can be a noticeable synchronization issue. This issue can be corrected by delaying the audio signal by a certain amount so it is back in line with the processed video signal (see line three).

FIGURE 1 - AUDIO TO VISUAL DELAY



BASICS OF DIGITAL SAMPLING

Unlike analog signals which are continuous, digital signals are discrete. This means that at an exact point in time the value of a digital signal can be measured precisely.

When measuring the value of an analog signal, for example the amplitude of an electrical signal after a period of time, a reading is taken accurate to two decimal places. The result may be 2.64V at a time of 10 seconds. To get a more accurate reading the amplitude measurement could be taken to three decimal places, possibly reading 2.639V. The accuracy of the time of the measurement could also be increased by being accurate to milliseconds rather than seconds. This pursuit of accuracy could conceivably continue forever.

FIGURE 1 - A/D-D/A PROCESS





FIGURE 2 - SAMPLE RATE



In a digital system the incoming continuous signal is periodically sampled, each sample having a discrete value. This means that for any given sample at any time interval, a definite value can be retrieved at any time and the value and accuracy of that sample will always be the same no matter when it is retrieved.

Sample rate

The sample rate determines the number of samples which are processed by the digital system each second. Normally measured in KHz, a sample rate of 48KHz states that 48000 samples are processed every second. Figure 2 illustrates the results of analog to digital (A-D) conversion by sampling an input signal at different sample rates.

The first plot shows a steady tone which is used as an input signal. The exact amplitude and frequency are not important at this point.

The vertical lines ending in dots in the second plot represent the value of the samples taken at regular intervals defined by the sample rate. In this plot the sampling rate is eight times the frequency of the input tone. The red line connecting the points shows the waveform which would be reconstructed by a basic digital to analog (D-A) converter if it received these points. Here the resulting waveform is identical to the input waveform.

The third plot demonstrates the limits of the Nyquist-Shannon sampling theorem which states that a tone of a given frequency can be reconstructed by the D-A converter if the sample rate is more than twice the maximum frequency of the signal being sampled. The resulting waveform is again identical to the input waveform.


The fourth plot confirms this theory by attempting to sample the input signal with a sample rate less than twice the input signal frequency. In this case the frequency that is reconstructed is of a lower frequency than the input signal. This is known as aliasing.

It should be clear now that if a sample rate of 48KHz is used, the theoretical maximum frequency that can be reproduced by a digital system is 24KHz, which is half the sampling frequency. This frequency is known as the Nyquist frequency.

Bit depth

As important as the resolution of the frequency of the sampled waveform is the resolution of the amplitude. This resolution

FIGURE 3 - QUANTIZATION



is determined by the bit depth of the digital signal.

The sample taken at the intervals defined by the sample rate is actually a word of a certain bit depth that represents the amplitude of the input signal at that point in time. The greater the bit depth, the more values are available to represent signal level and therefore the greater the resolution of the amplitude, giving a wider dynamic range and increasing signal to noise ratio.

Quantization

The discrete amplitude value of each sample is a result of the quantization process inherent to digital systems. Because each word has a fixed number of bits, and those bits can each be set to only a 1 or a 0, the accuracy of the value is fixed. In a 2 bit word the value can only be a 0, 1, 2 or 3. It cannot be 0.3 or 2.556.

So given the problem of measuring an exact value from a continuous input signal (is the value measured as 3 really 2.995 or 3.114...), what happens if the value of the amplitude of the input signal at a certain sampling interval falls somewhere between two of the values in the digital word?

The answer is quantization. The value may be rounded up or down (depending on the quantization model used) to the nearest value available to the digital word. Looking at Figure 3 it can be seen that as the word length decreases, so does the number of amplitude values available and therefore the accuracy of the sampled signal is reduced, and therefore the dynamic range and signal to noise ratio are also reduced. The black dots on the curve show where the original continuous signal has been sampled. The red line shows the resulting discrete signal after being quantised. Using a 4 bit wordlength, as shown in the diagram above left, the digital signal can only be represented in 16 different amplitudes / levels. Increasing the bit-rate to 16, as used in CD quality audio provides over 65,000 different amplitude values that can be represented.

Dynamic range

The dynamic range of a system typically describes the ratio, in dB, between the lowest and highest signals that the system can operate on. Generally speaking, as the bit depth of a system increases by 1 bit, the dynamic range increases by 6dB. Given that a CD is a 16 bit medium, it's dynamic range is approximately 96dB. A 24 bit system would have a dynamic range of around 144dB.

DIGITAL CLOCKING

In a digital audio system where discrete signals are sent at regular time intervals, it is important that each item of equipment is synchronized to send and receive information at exactly the same time.

The top half of Figure 1 shows a situation where there are multiple items of equipment each sending and receiving information at intervals set by their own internal clocks. While the internal clocks of each item may be set to the same sample rate, for example 48kHz, it is extremely unlikely that they will be starting their clock pulses at exactly the same time. It is also unlikely that all four clocks will have the same accuracy. Even the tiniest of differences in the clock pulse spacing can lead to a drift in phase over a long period of time and result in out of sync signals.

Master clock


If a master clock source is used, all other items of equipment can be synchronized to this source and as such will all send and receive information at exactly the same time. The lower half of Figure 1 shows this in practice.

Synchronization signal

The structure and implementation of the synchronization signal itself depends on the type of clock signal used. Common sync signals used include Wordclock, AES3 or Tri-Level. We will focus on Wordclock here.

Wordclock is a bi-level signal (meaning it oscillates between two levels) that represents a square wave. Each full cycle is known as a pulse and it is the point in this pulse where the level drops from the highest point to the lowest point that indicates the point of synchronization. Each item of equipment that is clocked by this master signal listens for this point in the pulse and it is at this exact moment

FIGURE 1 - INTERNAL AND MASTER CLOCKS



that they send and/or receive a defined amount of digital data.

Clock accuracy

No clock signal generator can generate a signal at exactly the frequency required. If the desired frequency is 48000Hz, a clock generator may output a signal anywhere in the range of 47999.9Hz to 48000.1Hz. This is not such a problem in practice as long as all equipment is synchronized to the same clock source, and thus all working slightly above or below the desired frequency.


Jitter

In basic terms, jitter can be used to describe timing variations in sample clocks used in digital audio systems. In the case of a clock with considerable jitter, samples would be taken at non-uniform time intervals. Take Figure 2 for example, which shows a continuous input signal sampled at uneven time intervals due to considerable jitter in the clock feeding the AD converter. The discrete amplitude values obtained by each sample are stored in the digital medium. If a continuous signal is then reconstructed

using a DA converter with a much lower jitter clock, the amplitude values sampled previously will be reconstructed with a different spacing. In this case the resulting waveform will be different to the input waveform.

A situation such as this may arise when a continuous signal is converted to digital using a clock with high jitter and then stored in a digital medium, for example on a portable solid state recorder. If this recorded audio is transferred and played back on a totally separate system in a studio with a low jitter clock feeding the D/A converter, the resulting continuous signal will be different to the original.

FIGURE 2 - POSSIBLE EFFECT OF JITTER IN A CLOCK SIGNAL



BINARY NUMBERS

It is often useful to be able to understand numbers written in a binary form. Setting certain parameters in a Calrec system requires this understanding.

Simple explanation

Binary numbers are strings of 0 and 1 values, for example '10110011'. The length of a binary number is given in 'bits'. The previous example is an eight bit binary number, deduced by the number of digits.

The right most bit can represent the decimal value of 1. The next bit on the left can represent double that value, so it can represent a decimal number 2. The next bit on the left can represent a decimal value of 4 and so on.

If a bit is assigned the binary value 1, it represents its decimal value. If it is assigned the binary value 0, it represents a decimal 0. To get the decimal value of the whole binary number, add together the decimal representation of each bit assigned a binary value of 1. See Figure 1 for examples of binary representations of decimal numbers from 0 to 10.


FIGURE 1 - COUNTING IN BINARY

Binary Value	Decimal Representation	Decimal Value
	8 4 2 1	
0 0 0 0	0 + 0 + 0 + 0	0
0 0 0 1	0 + 0 + 0 + 1	1
0 0 1 0	0 + 0 + 2 + 0	2
0 0 1 1	0 + 0 + 2 + 1	3
0 1 0 0	0 + 4 + 0 + 0	4
0 1 0 1	0 + 4 + 0 + 1	5
0 1 1 0	0 + 4 + 2 + 0	6
0 1 1 1	0 + 4 + 2 + 1	7
1 0 0 0	8 + 0 + 0 + 0	8
1 0 0 1	8 + 0 + 0 + 1	9
1 0 1 0	8 + 0 + 2 + 0	10

Mathematical Explanation

Each bit in the binary number (or 'word') represents a decimal value of 2, raised to a certain power. 2² for example is the same as 2x2, which equals 4. 2⁶ is the

FIGURE 2 - BINARY NUMBER EXAMPLES



same as writing 2x2x2x2x2x2 which equals 64.

The top line in Figure 2 details the decimal values that each bit in a binary word represents. This diagram shows values up to a 16 bit word. Each word, no matter what length, starts with its lowest decimal value at the right most side and extends to the left. The red marks indicate the length of each word, and which decimal values each bit represents.

The first bit has an index of 0 and represents the value of 2⁰, which evaluates to 1. The second bit has an index of 1 and represents a value of 2¹, which evaluates to 2. The third bit has

an index of 2 and represents a value of 2², which evaluates to 4. The formula for finding the decimal value at a certain index is 2ⁱ.

If a word is N digits long, it is an N bit word. The maximum decimal value that word can represent is 2^N-1. For example if the word has 4 digits, it is a 4 bit word. The maximum decimal value it can represent is 2⁴-1, or 2x2x2x2 which equals 16, minus 1, which gives 15. The range of decimal values this word can represent is 0-15 inclusive. Figure 3 shows a table of word lengths and the maximum decimal values they can represent.

FIGURE 3 - MAX DECIMAL VALUES

Word Length	Maximum Decimal Value
1 bit	1
2 bit	3
3 bit	7
4 bit	15
8 bit	255
16 bit	65535
24 bit	16777215
32 bit	4294967295

If a bit at a given index (i) in the word is assigned a binary value of 1, it represents the decimal value of 2^i (more specifically $2^i \times 1$). If it is assigned the binary value of 0, it represents a decimal 0 (more specifically $2^i \times 0$). The total decimal value of the word can be found by summing all values of 2^i where the binary value at i is 1.

Figure 2 shows various examples of different word lengths and their associated values. The first is a 2 bit word written as 10. The first index (0) has the binary value of 0. This means that $2^0 \times 0 = 0$. The second index (1) has a binary value of 1. In this case $2^1 \times 1 = 2$. Summing these two results together gives the final value of 2.

The second example (1011) can be converted to a decimal number using the following formula:

$$(2^0 \times 1) + (2^1 \times 1) + (2^2 \times 0) + (2^3 \times 1) = 11$$

Calrec Audio Ltd

Nutclough Mill
Hebden Bridge
West Yorkshire
England UK
HX7 8EZ

Tel +44 (0)1422 842159
Fax +44 (0)1422 845244
Email Enquiries@calrec.com

calrec.com

(926-134 Iss.5)