

Beosound / Beovision Theatre

Technical Sound Guide

Bang & Olufsen A/S
August 18, 2023

Contents

1 Introduction	4
2 Menu Map	5
3 Listening Modes	6
3.1 Factory Defaults	6
3.2 Customisation Parameters	6
3.2.1 Beosonic	6
3.2.2 Frequency Tilt	6
3.2.3 Sound Enhancement	7
3.2.4 Speech Enhancement	7
3.2.5 Bass Management	7
3.2.6 Compression	7
3.2.7 Spatial Control	8
4 Listening Positions	10
4.1 RoomSense	10
4.1.1 Latency	10
4.2 Speaker Role	10
4.3 Speaker Distance	10
4.4 Speaker Level	11
4.5 Speaker Preset	11
4.6 Bass Management	11
4.6.1 Crossover Frequency	12
4.6.2 Re-direction Levels	12
4.7 Phase Compensation	13
4.8 Room Compensation	13
5 Equaliser	14
5.1 Treble	14
5.2 Bass	14
6 Sound Settings	15
6.1 Volume	15
6.2 Default Volume	15
6.3 Max Volume	15
6.4 Loudness	15
7 Sound Info	16
8 Connected Speakers	16
8.1 Internal outputs	16
8.2 Virtual outputs	16
8.2.1 Theory vs. Reality	17
8.3 Connected Speakers	17

9	Frequently-Asked Questions	19
9.0.1	Setup and configuration	19
9.0.2	Daily behaviour	20
10	Appendix 1: A Very Short History Lesson	22
10.0.1	Input Channels and Speaker Roles	23
10.0.2	x.1 ?	24
10.0.3	x.y.4 ?	24
10.0.4	Channel names	25
11	Appendix 2: Dolby Atmos	27
11.0.1	Loudspeaker configurations	27
11.0.2	Channels and Objects	28
11.0.3	Virtual Loudspeakers	29
11.0.4	Up-firing and Side-firing Loudspeakers	29
11.1	Upmixing and Downmixing	29
11.2	Multichannel Configuration Customisation	29
12	Appendix 3: What is a “Virtual” Loudspeaker?	32
12.1	Part 1: Head-Related Transfer Functions	32
12.2	Part 2: Cancelling Crosstalk	34
12.3	Part 3: Virtualisation	36
13	Appendix 4: The Influence of Listening Room Acoustics on Loudspeakers	38
13.1	Early Reflections	38
13.2	Room Modes	40
13.3	Reverberation	42
13.4	Solutions	42
13.5	Conclusions	42

Introduction

Your Bang & Olufsen Beosound Theatre is equipped with an extremely powerful sound processing engine that can be customised to suit almost any configuration or listener preference, including the potential to add up to 16 external loudspeakers to complement its own 11 internal outputs. Internally, the digital signal processing (the computer that decodes, renders, and manages the audio signals coming into and out of your Theatre) uses a proprietary 7.2.4-channel format that is compatible with an enormous variety of loudspeaker configurations that make use of the internal and optional external loudspeakers. In spite of its complexity, it is extremely simple to use compared to other high-end soundbars, since almost all of the parameters are automatically configured when it is 'told' which Bang & Olufsen loudspeakers are connected to it. However, those parameters are all available in the menus to give experienced users the option of customising the settings for different setups.

The controls within the menus for the sound processing can be divided into four general areas:

Global Controls are day-to-day adjustments such as the volume control.

Listening Modes are pre-programmed adjustments to suit different types of program material, and are factory-set for materials such as Movie, Music, or Speech. These can also be customised to suit personal preferences.

Listening Positions are adjustments to suit different listening positions or situations. For example, these can be configured to change between two-channel stereo and surround listening, between one listening position and another, or between individual or group listening situations.

Speaker Connections are used to initially set up your choice of loudspeakers. This is only used as a first-time setup, or when adding new loudspeakers to your configuration.

In addition, The Beosound Theatre give you the option of setting the default Listening Mode and Listening Position for each source (e.g. Live TV, HDMI 1, etc.).

Menu Map

In order to facilitate navigation through the various parameter controls, the menu structure of Beosound Theatre is shown below in Figure 2.1.¹

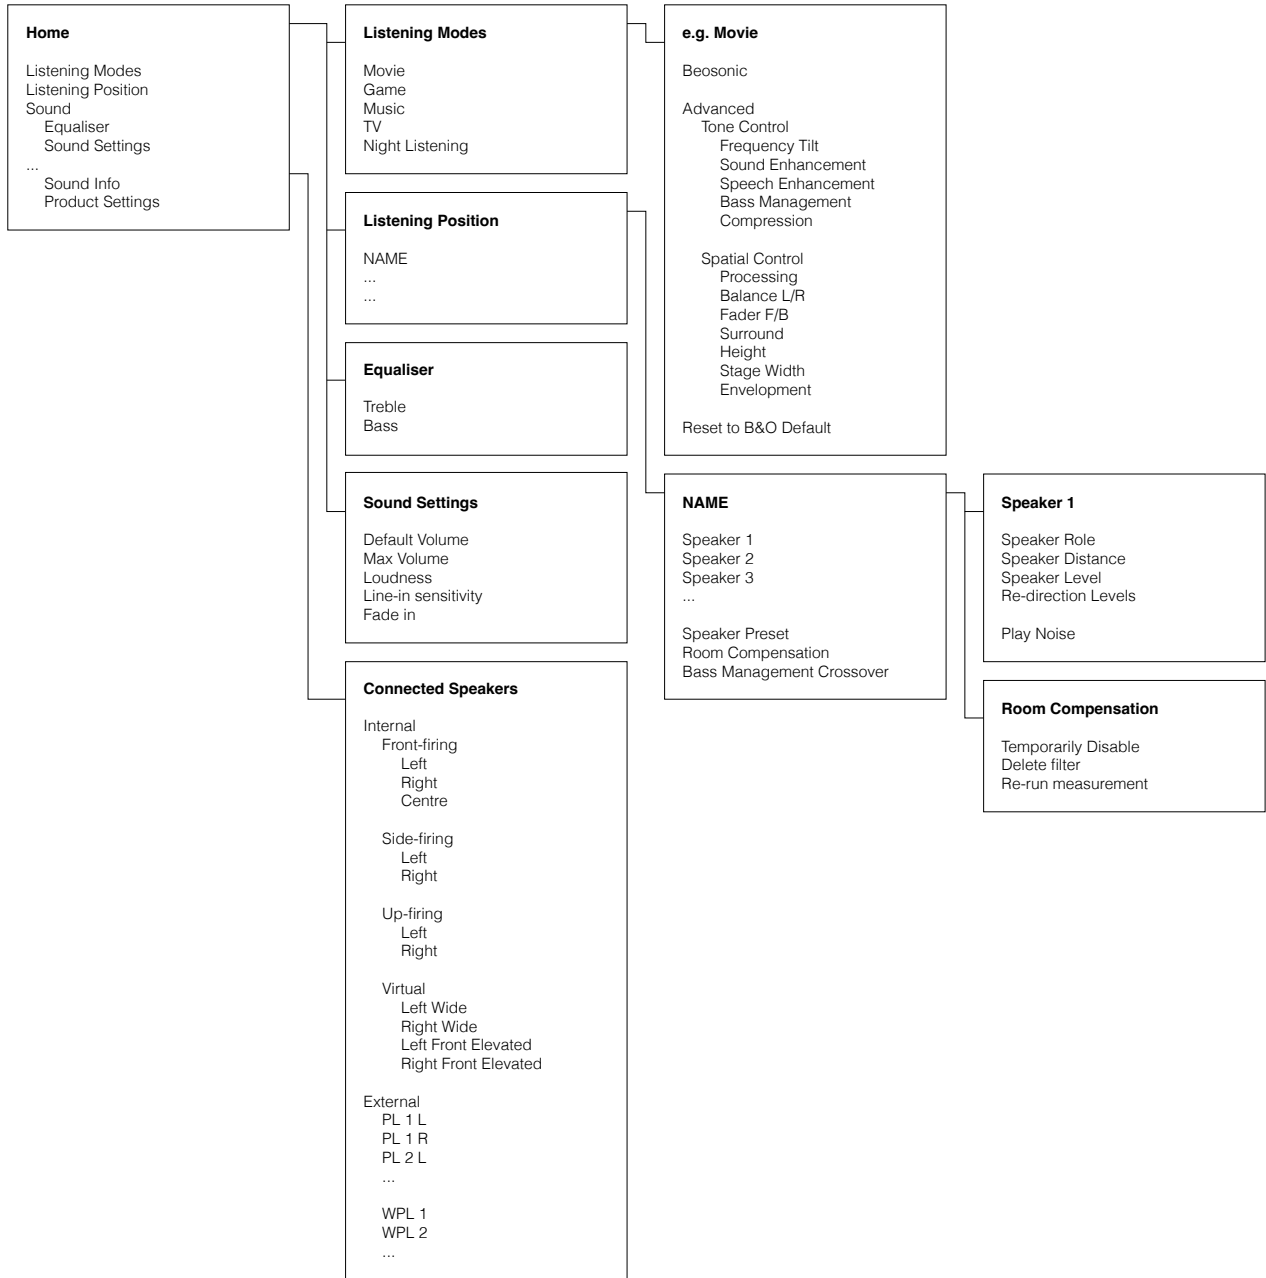


Figure 2.1: Menu map for the Beosound Theatre.

¹Note that the organisation of the controls in the Bang & Olufsen app are subject to change. Consequently, this menu may not directly match the current app version.

Listening Modes

As was briefly described in the Introduction, the Listening Modes on the Beosound Theatre allow you to have different audio settings for different types of signals. For example, you may wish to have a larger sound stage and enhanced bass response while watching movies, but a more purist signal path when listening to music. Listening Modes allow you to have five different presets for these changes.

3.1 Factory Defaults

All Listening Modes have factory-default settings that have been optimised for materials such as Movies, Games, TV, and Music, however, is possible to override these values to customise the settings of all Listening Modes to suit your own preferences.

Movie

The Movie Sound Mode is designed for use when watching movies, either from local media (such as DVD or Blu-ray), streaming sources, or television broadcasts.

Timbral settings are flat and bass management is on. The True Image processing is on and all of its controls are set to the middle position, which will result in a moderate amount of upmixing of two-channel materials to the Centre Front Speaker Role. Dynamic range compression is off to ensure that you experience the full dynamic range of the movie's audio signal.

Music

The Music Sound Mode is designed for use for music sources, either with or without accompanying video.

Note that this mode is not designed as a 'purist' setting. However, it is intended to have a minimal effect on the audio signals, while still up- or down-mixing to all loudspeakers in your current Listening Position. Changing the Spatial Processing to 'Direct' will modify the Sound Mode to ensure that the up- and down-mixing is disabled, if preferred.

Almost all parameters for the Music mode are identical to the Movie mode. The only differences are a slightly reduced value for the Frequency Tilt parameter and an increased value for the Width parameter, which will result in less content in 2.0-channel recordings being upmixed to the Center Front output, thus maintaining an impression of spaciousness in music recordings that contain it.

TV

The TV Sound Mode is designed primarily for use when watching television shows either via broadcast or streaming.

Timbral settings are flat and bass management is on. The True Image processing is on and all of its controls are set to the middle position. Dynamic range compression is set to medium to reduce the level changes encountered during advertising breaks for broadcast materials.

Game

The Game Sound Mode is designed for use for audio with game consoles.¹

The Frequency Tilt and Sound Enhancement settings give a slight bass and treble enhancement and bass management is on. The True Image processing is on and its controls are set to elevate the image and give an increased impression of envelopment and surround. Dynamic range compression is off to ensure that you experience the full dynamic range of the game.

Night Listening

The Night Listening Sound Mode is designed for situations where it is desirable to hear all components of the audio signal without large jumps in dynamics or bass. Consequently, in this mode, the Beosound Theatre's dynamic range compression is set to maximum and the Speech Enhancement is increased slightly.

3.2 Customisation Parameters

3.2.1 Beosonic

Beosonic is a simplified user interface that allows you to explore a range of timbral controls with a single, two-dimensional controller. In the case of Beosound Theatre, it is directly 'connected' to the Frequency Tilt and Sound Enhancement controls, explained in detail below. In fact, you may notice that by moving the position of the dot in the Beosonic screen, the Tilt and Enhance sliders will automatically move accordingly (and vice versa).

3.2.2 Frequency Tilt

Frequency Tilt can be considered to be a combination of Bass and Treble settings in a single parameter. When Frequency Tilt is set to a low value, the low frequency content of your audio signal is increased and the level of the high frequency content is reduced. If the Frequency Tilt is set to a high value, then the opposite will be true.

The Frequency Tilt function will have no effect on the audio signal at its middle setting.

Note not only that Frequency Tilt can have different settings

¹Note that this Sound Mode has no effect on the Input-Output latency of the device.

for different Listening Modes, but also that its position is interdependent with the location of the dot in the Beosonic screen.




Figure 3.1: Frequency response measurements of all settings of the Frequency Tilt control. Note that this response is applied to all input channels.

3.2.3 Sound Enhancement

The Sound Enhancement setting is similar to the Frequency Tilt setting in that it affects the low and high frequency bands with a single slider. Increasing the Sound Enhancement value will increase the level of the bass and treble bands while reducing the midrange. Decreasing the Sound Enhancement value will have the opposite effect.

The Sound Enhancement setting will have no effect on the audio signal at its middle setting.

Note not only that Sound Enhancement can have different settings for different Listening Modes, but also that its position is interdependent with the location of the dot in the Beosonic screen.

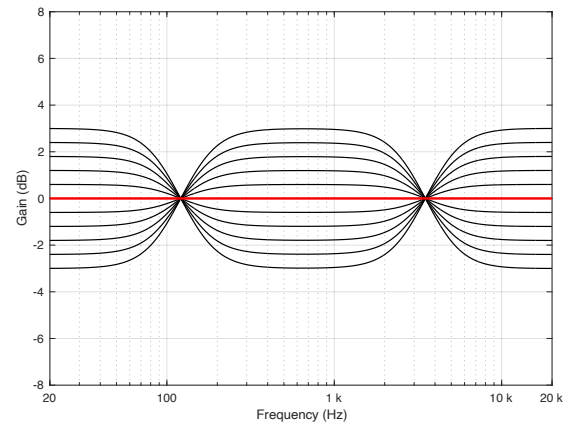


Figure 3.2: Frequency response measurements of all settings of the Sound Enhancement control. Note that this response is applied to all input channels.

3.2.4 Speech Enhancement

The Speech Enhancement setting allows you to increase the intelligibility of dialogue, making speech and voices easier to understand. This processing has an effect on the spatial balance of signals in a multichannel signal as well as the overall spectral balance of the output.

Note that the Speech Enhancement setting will have no effect on the audio signal when it is at its lowest setting.

3.2.5 Bass Management

The Bass Management processing can be turned on and off as part of the settings for the Listening Modes. Although the factory defaults for all Listening Modes is to have the bass management turned on, it may be more preferable to some to disable this, particularly for the Music mode for example.

Note that, if no loudspeaker outputs have a 'Subwoofer' Speaker Role, and if Bass Management is *disabled*, then the LFE input channel will not be played.

3.2.6 Compression

The Compression processing can be used to reduce the dynamic range of audio signals. This will reduce the difference in level between the quietest and loudest portions of the music; in other words, it makes quiet sounds louder and loud sounds quieter. Consequently, it is designed primarily for a 'night listening' situation where it is desirable to reduce peaks in the signal to avoid waking family members, while still allowing you to hear the quieter moments in the music or

movie. This setting can also be used for a 'party' setting where it is desirable to play music at a more constantly loud level.

There are many instances where it is desirable to reduce the dynamic range of the audio signal. For example, television advertisements are typically much louder than the programme they interrupt, and should be tamed. Films on DVD or Blu-ray often have large differences between the quietest and loudest moments, making it difficult to watch movies late at night without disturbing the rest of the family. At a party, the music should be kept at a constant level.

Consequently, the Beosound Theatre has the ability to reduce (or compress) the dynamic range of audio signals by making quiet passages louder and loud passages quieter. The amount of compression applied to the audio signal for the current Sound Mode is determined using the Compression setting.

3.2.7 Spatial Control

Processing

There are many cases where the number of input channels in the audio signal does not match the number of loudspeakers in your configuration. For example, you may have two loudspeakers, but the input signal is from a multichannel source such as a 7.1-channel stream or a 7.1.4-channel Blu-ray. In this case, the audio must be 'downmixed' to your two loudspeakers if you are to hear all components of the audio signal. Conversely, you may have a full surround sound system with 7 main loudspeakers and a subwoofer (a 7.1-channel system) and you would like to re-distribute the two channels from a CD to all of your loudspeakers. In this example, the signal must be 'upmixed' to all loudspeakers.

Bang & Olufsen's True Image processor accomplishes both of these tasks dynamically, downmixing or upmixing any incoming signal so that all components and aspects of the original recording are played using all of your loudspeakers.

Of course, using the True Image processor means that signals in the original recording are re-distributed. For example, in an upmixing situation, portions in the original Left Front signal from the source will be sent to a number of loudspeakers in your system instead of just one left front loudspeaker. If you wish to have a direct connection between input and output channels, then the Processing should be set to 'Direct', thus disabling the True Image processing.

Note that, in Direct mode, there may be instances where some input or output channels will not be audible. For example, if you have two loudspeakers but a multichannel input, only two of the input channels will be audible. These channels are dependent on the speaker roles selected for the two loudspeakers. (For example, if your loudspeakers' roles are Left Front and Right Front, then only the Left Front and Right Front channels from the multichannel source will be heard.)

Similarly, in Direct mode, if you have a multichannel configuration but a two-channel stereo input, then only the loudspeakers assigned to have the Left Front and Right Front speaker roles will produce the sound; all other loudspeakers will be silent.

If True Image is selected and if the number of input channels and their channel assignments matches the speaker roles, and if all Spatial Control sliders are set to the middle position, then the True Image processing is bypassed. For example, if you have a 5.1 loudspeaker system with 5 main loudspeakers (Left Front, Right Front, Centre Front, Left Surround, and Right Surround) and a subwoofer, and the Spatial Control sliders are in the middle positions, then a 5.1 audio signal (from a DVD, for example) will pass through unaffected.

However, if the input is changed to a 2.0 source (i.e. a CD or an Internet radio stream) then the True Image processor will upmix the signal to the 5.1 outputs.

In the case where you wish to have the benefits of downmixing without the spatial expansion provided by upmixing, you can choose to use the Downmix setting in this menu. For example, if you have a 5.1-channel loudspeaker configuration and you wish to downmix 6.1- and 7.1-channel sources (thus ensuring that you are able to hear all input channels) but that two-channel stereo sources are played through only two loudspeakers, then this option should be selected. Note that, in Downmix mode, there are two exceptions where upmixing may be applied to the signal. The first of these is when you have a 2.0-channel loudspeaker configuration and a 1-channel monophonic input. In this case, the centre front signal will be distributed to the Left Front and Right Front loudspeakers. The second case is when you have a 6.1 input and a 7.1 loudspeaker configuration. In this case, the Centre Back signal will be distributed to the Left Back and Right Back loudspeakers.

The Beosound Theatre includes four advanced controls (Surround, Height, Stage Width and Envelopment, described below) that can be used to customise the spatial attributes of the output when the True Image processor is enabled.

Balance

The Balance setting can be used to re-direct input signals to different output channels in your loudspeaker configuration. For example, setting the Balance all the way to the left on the display will result in signals being directed only to the loudspeakers in your configuration that have a Speaker Role on the left (i.e. Left Front, Left Surround, Left Front Height, etc.)

Fader

The Fader setting can be used to re-direct input signals to different output channels in your loudspeaker configuration. For example, setting the Fader all the way to the left on the display will result in signals being directed only to the speakers in your configuration that have a Speaker Role in the rear (i.e. Left Surround, Right Back etc.)

Note that, if you do not have a surround configuration of loudspeakers (i.e. if you have only a Front Left and Front Right loudspeaker in your current Listening Position) then the Fader B/F setting will not operate correctly and should be set to the middle position (the factory default setting).

Surround

The Surround setting allows you to determine the relative levels of the sound stage (in the front) and the surround information from the True Image processor.

Changes in the Surround setting only have an effect on the signal when the Processing is set to True Image.

Height

This setting determines the level of the signals sent to all loudspeakers in your configuration with a 'height' Speaker Role. It will have no effect on other loudspeakers in your system.

If the setting is set to minimum, then no signal will be sent to the 'height' loudspeakers.

Changes in the Height setting only have an effect on the signal when the Processing is set to True Image.

Stage Width

The Stage Width setting can be used to determine the width of the front images in the sound stage. At a minimum setting, the images will collapse to the centre of the frontal image. At a maximum setting, images will be pushed to the sides of the front sound stage. This allows you to control the perceived width of the band or music ensemble without affecting the information in the surround and back loudspeakers.

If you have three front loudspeakers (Left Front, Right Front and Centre Front), the setting of the Stage Width can be customised according to your typical listening position. If you normally sit in the 'sweet spot', at roughly the same distance from all three loudspeakers, then you should increase the Stage Width setting somewhat, since it is unnecessary to use the centre front loudspeaker to help to pull phantom images towards the centre of the sound stage. The further to either side of the sweet spot that you are seated, the more reducing the Stage Width value will improve the centre image location.

Changes in the Stage Width setting only have an effect on the signal when the Processing is set to True Image.

Envelopment

The Envelopment setting allows you to set the desired amount of perceived width or spaciousness from your surround and back loudspeakers. At its minimum setting, the surround information will appear to collapse to a centre back phantom location. At its maximum setting, the surround information will appear to be very wide.

Changes in this setting have no effect on the front loudspeaker

channels and only have an effect on the signal when the Processing is set to True Image.

Listening Positions

A Listening Position¹ is a configuration of loudspeakers that determine which speakers are playing, to what audio channels they are assigned, their calibration levels and delays, and the room compensation filters used for a given listening position.

4.1 RoomSense

The fastest method of creating a new Listening Position is to use the automatic RoomSense procedure. This uses a microphone placed at the listening position and a series of measurement signals from each loudspeaker to determine suggested values for each loudspeaker output.




Figure 4.1: Before starting the Room Sense procedure, the external microphone should be connected to the Beosound Theatre' microphone input, shown here.

In most cases, the values that are automatically chosen by the RoomSense procedure will be suitable for most listening situations, however, there are limitations to this system, and so it may be desirable to further optimise the settings to suit your particular preferences or listening situations. Section 11.2 details some suggestions for optimisation for different situations.

Note that one limitation of the Room Sense algorithm is that it is unable to detect the vertical position of external loudspeakers that are mounted in or near the ceiling. Consequently, these loudspeakers will likely be assigned incorrect roles that will have to be manually corrected. In addition, in the case where external loudspeakers exist in the "Front Height" positions, since these will be automatically assigned to have "Front" roles instead, then the up-firing outputs of the Beosound Theatre will be assigned to "Front Height" roles. Consequently, the manual correction will require attention to correcting the external loudspeakers' roles and setting the up-firing outputs to "None".

4.1.1 Latency

When running the Room Sense procedure on a Beosound Theatre that is connected to Beolab 50 or Beolab 90 loudspeakers, you will be asked to choose whether to set the system to High or Low latency mode.

If the Listening Position is to be used to sources without video content, then it is recommended that the High Latency option be used. However, if the source includes video content, then the Low Latency option should be used to ensure that lip sync is maintained.²

Information about the implications of the Low and High Latency options in Beolab 50 or Beolab 90 can be found in their respective Technical Sound Guides, available on the Bang & Olufsen website.

4.2 Speaker Role

This menu allows you to enable the loudspeakers that are used in the current Listening Position. In addition, you can set the desired channel allocation for each loudspeaker (or Power Link output channel).

When configuring a 5.1-channel surround system, note that the rear loudspeakers should be set as Left Surround and Right Surround (not Left Back and Right Back).

There are no restrictions on how many copies of a given speaker role that may be distributed in a Listening Position. For example, if you have 10 loudspeakers connected to the Beosound Theatre, it is allowed (although perhaps not advisable) to have 10 'Left Surround' outputs and nothing else.

The positions corresponding to the various Speaker Roles are shown in Figures 11.4 and 11.5.

4.3 Speaker Distance

This setting is used to ensure that the times of arrival of the loudspeakers' signals at the listening position are matched, despite them being placed at different distances from the listening position. The value displayed on the menu should be the distance from the listening position to each loudspeaker. The result of this alignment is that all loudspeakers' signals are individually delayed to match the time of arrival of the sound from the most distant loudspeaker.

Since the listening position can be different for different Listening Positions, these distances may not necessarily be the same from Listening Position to Listening Position. In addition,

¹A 'Listening Position' is equivalent to a 'Speaker Group' on previous Bang & Olufsen televisions and surround processors.

²This may not be source-dependent, but determined by your use case. For example, an Apple TV can be used as an audio-only source when running a music streaming app. In this case, a High Latency Listening Position should be used to ensure the highest audio performance from your system.

the small differences in latency between various Bang & Olufsen loudspeakers, connected either wirelessly or with Power Link cables, are automatically compensated for internally in the system.

4.4 Speaker Level

The Speaker Level setting is used to align the perceived or measured loudness of the loudspeakers at the listening position. Although Bang & Olufsen loudspeakers are all factory-calibrated to give the same output level with the same input signal, different loudspeaker (or listener) placements and different room conditions have an effect on the speaker level at the listening position. As a result, you will most likely require some adjustments to optimise your system. In order to achieve the optimal settings for the Speaker Levels, it is highly recommended that you use a sound pressure level meter. This can either be an app on a smart phone or (preferably) a dedicated unit. Note that, if you are considering purchasing a sound pressure level meter for this purpose, a low-priced device (less than \$100) will produce acceptable results.

The calibration procedure is as follows:

1. Ensure that the correct speaker types have already been entered for each loudspeaker in your system.
2. Set the sound pressure level meter to a 'C' weighting and 'Slow' setting. It should be placed near the listening position with the microphone pointed towards the ceiling.
3. Select a loudspeaker from the menu and use the "Play Noise" option (see Figure 2.1). You should hear a noise signal coming from the loudspeaker you selected.
4. Set the volume (not the Speaker Level) so that the reading on the sound pressure level meter is 65 dB SPL (C).
Note that, in some cases, it is suggested to set the up-firing and side-firing outputs to a lower output level than the front-firing outputs, ranging from 0 to approximately -3 dB (i.e. an output level as low as 62 dB SPL (C) assuming that the reference level is 65 dB SPL (C)). However, this is dependent not only on the acoustical characteristics of the surrounding but also your personal preferences.
5. Select a different loudspeaker from the menu and set its Speaker Level so that it also produces a reading of 65 dB SPL at the listening position.
6. Continue this process for all loudspeakers. Note that subwoofers should also give the same reading when the SPL meter is set to a 'C' weighting.

If you have more than one loudspeaker assigned to a single output channel (for example, if you have two Left Surround loudspeakers) you should calibrate the two loudspeakers using

the same method as all other loudspeakers. The Beosound Theatre automatically compensates the speaker levels for the fact that more than one speaker is used for the same channel. This compensation is not displayed visually.

4.5 Speaker Preset

Some Bang & Olufsen loudspeakers (such as the Beolab 90, 50, and 28) have user-programmable 'Presets' that can be used to customise specific characteristics such as Beam Width or Beam Direction. It is possible to associate a given preset in the loudspeaker with a Listening Position in the Beosound Theatre using the 'Speaker Preset' number.

Any Speaker Preset can be associated with any Listening Position; in other words, it is not necessary that the two numbers match each other.

Note that the Beosound Theatre can only transmit one Speaker Preset value to all Power Link and Wireless Power Link outputs. Consequently, in cases where there are multiple pairs of loudspeakers (e.g. four Beolab 28s) connected to the Beosound Theatre, the parameters within the loudspeakers' presets should be carefully selected to match each other.

If the Speaker Preset value in the Beosound Theatre is set to '0', then no preset value will be transmitted to the loudspeakers and so they will not change presets with a change in Listening Position.

It should also be noted that there are some parameters in the loudspeaker (e.g. Latency Mode) that are automatically overridden by the Beosound Theatre in order to ensure proper integration with other loudspeakers in the configuration.

4.6 Bass Management

In a perfect sound system, all loudspeakers are identical, and they are all able to play a full frequency range at any listening level. However, most often, this is not an option. Luckily, it is possible to play some tricks to avoid having to install a large-scale sound system to listen to music or watch movies.

Humans have an amazing ability to localise sound sources. With your eyes closed, you are able to point towards the direction sounds are coming from with an incredible accuracy. However, this ability gets increasingly worse as we go lower in frequency, particularly in closed rooms.

In a sound system, we can use this inability to our advantage. Since you are unable to localise the point of origin of very low frequencies indoors, it should not matter where the loudspeaker that's producing them is positioned in your listening room. Consequently, many simple systems remove the bass from the 'main' loudspeakers and send them to a single large loudspeaker whose role it is to reproduce the bass

for the *entire* system. This loudspeaker is called a ‘subwoofer’ since it is used to produce frequency bands below those played by the woofers in the main loudspeakers. The process of removing the bass from the main channels and re-routing them to the subwoofer is called **bass management**.

Note that, although many bass management systems assume the presence of at least one subwoofer, that output should not be confused with an LFE or a ‘.1’ channel. However, in most cases, the LFE channel from your media (e.g. a Blu-ray disc or streaming device) will be combined with the low-frequency output of the bass management system and the total result routed to the subwoofer.

The Beosound Theatre has a new bass management algorithm that is considerably more sophisticated than the simple example described above. It is also a complete departure from earlier Beovision televisions. When the Bass Management is ON, the low frequency content in your audio signal from all input channels will be produced by different loudspeakers in your current configuration at different levels depending on their capabilities. Note that this also means that if you are using a Sound Mode with a multichannel output, the input signal has fewer channels, the Spatial Processing setting is “Direct”, and Bass Management is enabled, then all loudspeakers will contribute to reproducing the low-frequency content.

It is possible to turn the Bass Management on or off for a given Listening Mode. When Bass Management is OFF, then the low frequency content of each input channel will be produced by the loudspeaker with its corresponding Speaker Role.³ This means that if no loudspeaker has a Subwoofer Speaker Role, then the LFE input channel will not be reproduced.

4.6.1 Crossover Frequency

The crossover used in the Bass Management system uses a 4th-order Linkwitz-Riley design. All output signals use the same crossover frequency to ensure that coherent signals on multiple output channels have matched phase responses when passed through the Bass Management processing.




Figure 4.2: Gain response of the crossover used in the Bass Management system, showing an example with the factory-default crossover frequency of 120 Hz.

4.6.2 Re-direction Levels

Selecting this item brings you to the menu where you adjust the level of the two Bass Management low-frequency channels being added back to the individual output channels.

It is possible in the Theatre to use any loudspeaker to reproduce the LFE channel and low-frequency content from the bass management system. This is done using the Bass Redirection. Typically, in a bass-managed 5.1 or 7.1.4 system, the LFE and the bass management output will be directed to a single subwoofer. However, if you desire, it is possible to send that information to any loudspeakers in the current Listening Position by increasing the Redirection Levels.

Note that initial settings in this menu are automatically chosen based on information from the Speaker Connections menu and the Speaker Roles for the current Listening Position. These values are calculated to ensure that the total bass output of the system is correct, with a consideration of the maximum capabilities of the individual loudspeaker models in the Listening Position. This means that, in order to maintain the same overall bass output of the system, if you increase the level of one output in this menu, you should decrease the value of another output by an appropriate (but not necessarily equal) amount. Unfortunately, calculating these linked changes values is rather complex. However, it may be more useful to use the Re-direction Levels as a controller for the overall output level of the bass management system. For example, you may wish to increase all values by 1 dB to increase the bass output level to suit your tastes or the acoustical response in your listening room.

³Note that this is also dependent on the selection of the Spatial Processing parameter.

4.7 Phase Compensation

As mentioned above, Beosound Theatre's bass management system does not merely send the low-frequency content to the most capable loudspeaker as is the case in other products. Instead, it distributes the signals to all loudspeakers in the system simultaneously. One advantage of this strategy is that it ensures that you achieve the maximum possible performance out of the entire system. However, by using multiple loudspeakers placed in many different locations in the listening room, it also helps to control the room's natural resonances called 'room modes'.

One common problem that occurs when creating configurations that include different loudspeaker types is that they do not all have matched magnitude and phase responses. In the high frequency bands, this is not a problem, since the distances between loudspeakers is typically much larger than the wavelengths of the audio signals as they travel through the air. This physical separation has the practical result of randomising the phase relationships between the loudspeakers, and therefore they cannot directly interfere with each other at the listening position.

However, as the frequency decreases, the wavelengths approach and then exceed the loudspeaker spacing. In the case where the loudspeakers produce different signals, this is not an issue. However, when different loudspeakers produce the *same* low-frequency signal simultaneously, their outputs will interfere with each other, either constructively (making things louder) or destructively (making them quieter).

Since the Beosound Theatre's bass management system is built on the distribution of low-frequency signals to all loudspeakers, it becomes necessary to ensure that the phase responses of the various loudspeakers are matched to prevent them from cancelling each other in the listening room. This is done by ensuring that the Phase Compensation ON, thus applying phase correction filters for the specific loudspeakers in your current Listening Position.

It is advisable to enable the Phase Compensation processing, even when Bass Management is disabled. This is because, in almost all multichannel mixes, the lower-frequency content that is distributed to multiple channels is correlated (in other words, the same bass signal is sent to all channels).

Note that it is not currently possible to disable the Phase Compensation filter.

4.8 Room Compensation

When you're sitting in a normal domestic room, listening to the sound coming from anything that is reasonably far away from you, a very large portion of the total sound that you hear is that which is reflected back to you from the various surfaces in the room, and not that which is travelling directly to you from the sound source. This is also true of loudspeakers, even when you are sitting in the 'sweet spot' on your sofa.⁴

This, in turn, means that a large part of why a loudspeaker sounds like it does is the result of the influence of your listening space.

As is described in more detail in Section 13, the effect a room's acoustical behaviour has on a loudspeaker's sound can, at a simple level, be considered under three general headings:

- Early Reflections
- Room Modes
- Reverberation

The test signals that are recorded using the external microphone when creating a Listening Position are also used to measure some aspects of the room's effect on the characteristics of the Beosound Theatre and all external loudspeakers at the microphone's location. These measurements are then used to create filters that reduce the effect of the room modes and the boundary conditions⁵ of the loudspeakers.

It is important to remember that the room compensation filters are created using the measurement performed at the microphone's location. Therefore, the microphone must be placed at the listening position (or the centre of the listening area), since the room's behaviour can differ significantly with changes in listening position.

There may be cases where you prefer to *not* use the room compensation filter.⁶ In these cases, it is possible to temporarily disable the Room Compensation filter, and if preferred, to delete it from the Listening Position without altering other customised parameters.

⁴Luckily, our brains have the ability to focus on the sound source and ignore the contribution of the room using a trick known as the 'cocktail party effect'. However, you can partially disable this by plugging one ear, which will cause the 'sound' of the room to be much more apparent.

⁵This is a customised version of the older 'Free / Wall / Corner' equalisation on older Bang & Olufsen loudspeakers.

⁶For example, Listening Positions that include the Beolab 50 or 90 with previously-customised ARC filters, or Listening Positions that are intended for a wider listening area.

Equaliser

5.1 Treble

The Treble adjustment allows you to change the relative amount of high-frequency sound using a high shelving filter with a turnover frequency of 8 kHz.

This is a global control, meaning that the setting of the Treble control is applied to all Listening Modes and Listening Positions, and is independent of the settings of the Frequency Tilt and Sound Enhancement controls. The range of the controller is ± 8 dB in steps of 0.8 dB.




Figure 5.1: Frequency response measurements of all settings of the Treble control. Note that this response is applied to all input channels.

5.2 Bass

The Bass adjustment allows you to change the relative amount of low-frequency sound using a low shelving filter with a turnover frequency of 120 Hz.




Figure 5.2: Frequency response measurements of all settings of the Bass control. Note that this response is applied to all input channels.

This is a global control, meaning that the setting of the Bass control is applied to all Listening Modes and Listening Positions, and is independent of the settings of the Frequency Tilt and Sound Enhancement controls. The range of the controller is ± 8 dB in steps of 0.8 dB.

Sound Settings

6.1 Volume

It's important to remember that, like all high-end audio devices, the volume setting is **not** an indication of the output level, since this is also dependent on the input signal level. This is why, for example, the same volume setting on a movie and a typical Internet radio station will result in very different output levels, since movies are mastered at a much lower level than pop music.

It is also important to note that, due to non-linear processes such as loudspeaker and thermal protection, the changes in output level of the Beosound Theatre and any connected external loudspeakers is not necessarily *directly* related to changes in the volume control. For example, if you are playing a music track that has been mastered at a high and constant level, and you increase the volume of the Beosound Theatre, at some point, you will reach the maximum possible output. Increasing the volume control beyond this will not result in an increase in the output level, instead, it will merely make the protection algorithms work harder. The volume step at which this transition occurs depends on the input signal itself, so it will be different for different listening materials.

6.2 Default Volume

The Default Volume is the volume setting that is applied when you start up the Beosound Theatre.

6.3 Max Volume

The Max Volume is the highest volume setting that is allowed in normal usage. This can be used to limit the maximum output of the Beosound Theatre.

6.4 Loudness

Sadly, human hearing is imperfect. In particular, one of the issues from which we all suffer is that our perception of the timbre or 'tone colour' of a sound is not constant with listening level. We are less sensitive to low frequencies when they are played at low listening levels. In other words, if you are listening to music at a high level and you turn down the volume, you will notice that, the lower the volume, the less bass you can hear. This is also true of high frequencies, albeit to a lesser extent.

The Loudness setting in your Beosound Theatre counteracts this effect. As you reduce the volume, the levels of the low- and high-frequency bands are automatically increased to

compensate for your reduced perception in the outer frequency bands. The Loudness compensation can be disabled, according to your preference.

Sound Info

Select this menu item to display information about the incoming and outgoing audio signals, including the encoding

type and the number of audio channels in the incoming stream. Also displayed are the current Sound Mode and Listening Position.




Figure 7.1: The 'naked' Beosound Theatre

Connected Speakers

Beosound Theatre has a total of 11 possible outputs, seven of which are 'real' or 'internal' outputs and four of which are 'virtual' loudspeakers. As with all current Beovision televisions, any input channel can be directed to any output by setting the Speaker Roles in the menus. In addition, any output can be disabled.¹

8.1 Internal outputs

On first glance of the line drawing in Figure 7.1 it is easy to jump to the conclusion that the seven real outputs are easy to find, however this would be incorrect. The Beosound Theatre has 12 loudspeaker drivers that are all used in some combination of level and phase at different frequencies to *all* contribute to the total result of *each* of the seven output channels.

So, for example, if you are playing a sound from the Left front-firing output, you will find that you do not only get sound from the left tweeter, midrange, and woofer drivers as you might in a normal soundbar. There will also be some contribution from other drivers at different frequencies to help control the spatial behaviour of the output signal. This Beam Width control is similar to the system that was first introduced by Bang & Olufsen in the Beolab 90. However, unlike the Beolab 90, the Width of the various beams cannot be changed in the Beosound Theatre.

The seven internal loudspeaker outputs are

- Front-firing: Left, Centre, and Right
- Side-firing: Left and Right

- Up-firing: Left and Right

Looking online, you may find graphic explanations of side-firing and up-firing drivers in other loudspeakers. Often, these are shown as directing sound towards a reflecting wall or ceiling, with the implication that the listener therefore hears the sound in the location of the reflection instead. Although this is a convenient explanation, it does not necessarily match real-life experience due to the specific configuration of your system and the acoustical properties of the listening room.

The truth is both better and worse than this reductionist view. The bad news is that the illusion of a sound coming from a reflective wall instead of the loudspeaker can occur, but only in specific, optimised circumstances. The good news is that a reflecting surface is not strictly necessary; therefore (for example) side-firing drivers can enhance the perceived width of the loudspeaker, even without reflecting walls nearby.

However, it can be generally said that the overall benefit of side- and up-firing loudspeaker drivers is an enhanced impression of the overall width and height of the sound stage, even for listeners that are not seated in the so-called 'sweet spot'² when there is appropriate content mixed for those output channels.

8.2 Virtual outputs

Devices such as the 'stereoscope' for representing photographs (and films) in three-dimensions have been around since the 1850s. These work by presenting two different photographs with slightly different perspectives to the two eyes. If the differences in the photographs are the same as the differences your eyes would have seen had you 'been there', then your brain interprets into a 3D image.

¹Note that it may be necessary to enter the "Advanced Settings" menus to do this. This is done by clicking the three dots at the top right of the "Edit Position" screen.

²In the case of many audio playback systems, the 'sweet spot' is directly in front of the loudspeaker pair or at the centre of the surround configuration. In the case of a Bang & Olufsen system, the 'sweet spot' is defined by the user with the help of the Speaker Distance and Speaker Level adjustments.

A similar trick can be done with sound sources. If two different sounds that exactly match the signals that you would have heard had you 'been there' are presented at your two ears (using a **binaural** recording), then your brain will interpret the signals and give you the auditory impression of a sound source in some position in space. The easiest way to do this is to ensure that the signals arriving at your ears are completely independent using headphones.

The problem with attempting this with loudspeaker reproduction is that there is 'crosstalk' or 'bleeding of the signals to the opposite ears'. For example, the sound from a correctly-positioned Left Front loudspeaker can be heard by your left ear and your right ear (slightly later, and with a different response). This interference destroys the spatial illusion that is encoded in the two audio channels of a binaural recording.

However, it might be possible to overcome this issue with some careful processing and assumptions. For example, if the exact locations of the left and right loudspeakers and your left and right ears are known by the system, then it's (hypothetically) possible to produce a signal from the right loudspeaker that cancels the sound of the left loudspeaker in the right ear, and therefore you only hear the left channel in the left ear.³

Using this 'crosstalk cancellation' processing, it becomes (hypothetically) possible to make a pair of loudspeakers behave more like a pair of headphones, with only the left channel in the left ear and the right in the right. Therefore, if this system is combined with the binaural recording / reproduction system, then it becomes (hypothetically) possible to give a listener the impression of a sound source placed at any location in space, regardless of the *actual* location of the loudspeakers.

8.2.1 Theory vs. Reality

It's been said that the difference between theory and practice is that, in theory, there is no difference between theory and practice, whereas in practice, there is. This is certainly true both of binaural recordings (or processing) and crosstalk cancellation.

In the case of binaural processing, in order to produce a convincing simulation of a sound source in a position around the listener, the simulation of the acoustical characteristics of a particular listener's head, torso, and (most importantly) pinnae (a.k.a. 'ears') must be both accurate and precise.⁴

Similarly, a crosstalk cancellation system must also have accurate and precise 'knowledge' of the listener's physical

characteristics in order to cancel the signals correctly; but this information also crucially includes the exact locations of the loudspeakers and the listener (we'll conveniently pretend that the room you're sitting in does not exist).

In the end, this means that a system with adequate processing power can use two loudspeakers to simulate a 'virtual' loudspeaker in another location. However, the details of that spatial effect will be slightly different from person to person (because we're all shaped differently). Also, more importantly, the effect will only be experienced by a listener who is positioned correctly in front of the loudspeakers. Slight movements (especially from side-to-side, which destroys the symmetrical time-of-arrival matching of the two incoming signals) will cause the illusion to collapse.

Beosound Theatre gives you the option to choose Virtual Loudspeakers that appear to be located in four different positions: Left and Right Wide, and Left and Right Elevated. These signals are *actually* produced using the Left and Right front-firing outputs of the device using this combination of binaural processing and crosstalk cancellation in the Dolby Atmos processing system. If you are a single listener in the correct position (with the Speaker Distances and Speaker Levels adjusted correctly) then the Virtual outputs come very close to producing the illusion of correctly-located Surround and Front Height loudspeakers.

However, in cases where there is more than one listener, or where a single listener may be incorrectly located, it may be preferable to use the 'side-firing' and 'up-firing' outputs instead.

For a more detailed explanation of virtualisation, see Appendix 3 on page 32, below.

8.3 Connected Speakers

As mentioned at the start of this section, Beosound Theatre on its own has 11 internal outputs.

In addition to these, there are 8 wired Power Link outputs and 8 Wireless Power Link outputs for connection to external loudspeakers, resulting in a total of 27 possible output paths. And, as is the case with all Beovision televisions since Beoplay V1, any input channel (or output channel from the True Image processor) can be directed to any output, giving you an enormous range of flexibility in configuring your system to your use cases and preferences.

The 'Connected Speakers' is used to indicate the type or model of loudspeaker connected to each output of the Theatre. All current Bang & Olufsen loudspeakers are listed in this

³Of course, the cancelling signal of the right loudspeaker also bleeds to the left ear, so the left loudspeaker has to be used to cancel the cancellation signal of the right loudspeaker in the left ear, and so on...

⁴For the same reason that someone else should not try to wear my glasses.

menu, including some discontinued models such as the Beolab 1 and Beolab Penta loudspeakers.

In addition to Bang & Olufsen loudspeakers, an additional option is provided. 'Other' is used to indicate a loudspeaker that is not included in the list of Bang & Olufsen loudspeakers, for example, a loudspeaker from a different manufacturer or a Bang & Olufsen loudspeakers that is not in the list.

Note that the Theatre automatically enters settings in other menus and makes decisions regarding the appropriate signal processing⁵ based on the information entered in this menu.

Note that Beolab loudspeakers connected using Wireless Power Link will have the speaker type automatically set by the system. This cannot be overridden in the user menus.

⁵Bass Management, selection of Phase Compensation filters, and latency compensation, for example.

menu, including some discontinued models such as the Beolab 1 and Beolab Penta loudspeakers.

In addition to Bang & Olufsen loudspeakers, an additional option is provided. 'Other' is used to indicate a loudspeaker that is not included in the list of Bang & Olufsen loudspeakers, for example, a loudspeaker from a different manufacturer or a Bang & Olufsen loudspeakers that is not in the list.

Note that the Theatre automatically enters settings in other menus and makes decisions regarding the appropriate signal processing⁵ based on the information entered in this menu.

Note that Beolab loudspeakers connected using Wireless Power Link will have the speaker type automatically set by the system. This cannot be overridden in the user menus.

⁵Bass Management, selection of Phase Compensation filters, and latency compensation, for example.

Frequently-Asked Questions

9.0.1 Setup and configuration

My external loudspeakers have a Free/Wall/Corner switch. Does it matter which setting I should use?

It is recommended to put all switches in “Free” mode before running the room compensation measurement. This is because the “Wall” and “Corner” settings typically reduce the output level of the loudspeaker in the low frequencies. If the Beosound Theatre’s room compensation measurement detects too little bass as a result, then it will apply gain to its outputs. It is smarter (on a system level) to ensure that the loudspeaker is naturally producing as much low frequency level as possible (in “Free” mode) and consequently *reducing* the output levels from the Beosound Theatre where appropriate.

Which settings should I use on my external subwoofer?

When connecting a **Beolab 2** to a Beosound Theatre, it is important to ensure that the Position is set to ‘Free’ as shown in Figure 9.1.




Figure 9.1: The recommended settings of the Beolab 2 when used with the Beosound Theatre are shown in red. Additional Beolab loudspeakers should be connected directly to the Beosound Theatre instead of to the Beolab 2’s outputs.

When connecting a **Beolab 11** to a Beosound Theatre, it is important to ensure that the subwoofer’s input is set to “Power Link”, the Bass Management is set to “External”, and that the Position is set to ‘1’ as shown in Figure 9.2.




Figure 9.2: The recommended settings of the Beolab 11 when used with the Beosound Theatre are shown in red.

When connecting a **Beolab 14** to a Beosound Theatre, it is important to set the controls as shown in Figure 9.3.




Figure 9.3: The recommended settings of the Beolab 14 when used with the Beosound Theatre are shown in red. This diagram assumes that the subwoofer is connected using a Wired Power Link; this may not be the case in your setup.

When connecting a **Beolab 19** to a Beosound Theatre, it is important to ensure that the Low Pass Filter (LP Filter) on the subwoofer is ‘OFF’, and that the Position is set to ‘Free’ as shown in Figure 9.4.




Figure 9.4: The recommended settings of the Beolab 19 when used with the Beosound Theatre are shown in red. This diagram assumes that the subwoofer is connected using a Wired Power Link and that you are using the Left channel of that connection; this may not be the case in your setup.

What is the procedure when connecting Beolab 50s or Beolab 90s?

There are some important issue to remember when connecting Beolab 50s or Beolab 90s to a Beosound Theatre.

The first is to ensure that the configuration of the loudspeakers should be completed before running the automatic setup measurement on the Theatre. Consequently, you should set the loudspeakers’ Beam Width, Beam Direction, run the ARC measurements (and create the ARC filter for the Preset, and make any desired modifications of the Parametric Equalisation *before* beginning the microphone-based setup on the Beosound Theatre.

Secondly, you should ensure that the Beolab 50s or Beolab 90s are in the appropriate Preset before running the automatic setup measurement on the Theatre.

If you make any changes to the parameters on the Beolab 50 afterwards, it is recommended that you re-run the Beosound Theatre’s Room Compensation procedure again.

Can I connect a Bluetooth transmitter for a wireless connection to headphones or hearing aids?

It is possible to connect a third-party Bluetooth transmitter to the Beosound Theatre using one of the Power Link outputs and

an appropriate adapter cable. For content with dialogue, it is probably best to set both output channels' Speaker Types to . This will ensure that the dialogue is routed to the transmitter. For music content, it is probably better to set the Speaker Roles to Left Front and Right Front instead. If the latency (delay) of the transmitter-to-receiver path is known, then this can be time-aligned with the loudspeakers in the system by adjusting the Speaker Distance for the output channels. The correct distance can be calculated using the equation below.

$$Distance_m = ms * 0.343 \quad (9.1)$$

where ms is the latency of the transmission system in milliseconds and $Distance_m$ is the appropriate compensation setting in the Speaker Distance menu in metres.

Why is the automatic Speaker Role assignment not using all of my loudspeakers?

There are many cases where it is better to use one loudspeaker than another for a given Speaker Role. An obvious case is where a pair of 'main' front loudspeakers are placed to either side of the Beosound Theatre. In this case, the Left Front and Right Front roles will be assigned to the external loudspeakers instead of the internal Left and Right Front-firing outputs.

However, it is possible to edit any Listening Position to include or remove loudspeakers in the configuration to suit your preferences.

9.0.2 Daily behaviour

Why are some sources so much louder/quieter than others?

The intention of Beosound Theatre, like all Bang & Olufsen loudspeakers, is to reproduce the recorded signals as faithfully as possible. This statement is broad and therefore includes multiple dimensions. For example, it means that you should be able to hear all frequency bands in the music or film, with the same spectral balance as was heard in the studio when the content was mixed. In the case of correctly-configured playback systems, it also means that you should hear the spatial effects encoded in the recording. This includes not only accurate and precise locations of individual voices and instruments, but also the spaciousness and envelopment of reverberant spaces when it's appropriate. Finally, the dynamic range of the music or movie should also be reproduced: the loud explosions of the movie and the sforzando moments in the orchestral piece should be much louder than quiet footsteps, whispers, and the pianissimo moments.

This last sentence can cause a problem when you switch from source to source; whether that means two different devices such as an HDMI source and the internal Internet radio, or two different apps such as a movie streaming service to a music streaming service on the same device. Then, the goal to

remain faithful to the dynamic range of the audio signals reveals the differences in the levels at which the signals themselves were created.

For example, in general, the dialogue (speech) in a movie is mixed at a level low enough to leave room to get louder when an explosion occurs. Similarly, a concerto for violin and orchestra will have a range from the quietest moment of the solo violin to the loudest moment of the entire orchestra. Modern pop, rock, and country music (as well as other genres) typically has very little dynamic range (the difference between the quietest and the loudest moments) and therefore is typically mixed to be as loud as the explosions in the film or the loudest moments of the orchestra piece.

This means that, for example, if you are watching a movie, and, during a scene that consists of only dialogue, you switch sources to an Internet radio station playing pop music, you will experience a sudden jump upwards in the audio level, despite the fact that you did not change the volume setting. The radio station will appear to be much louder than the video stream.

Similarly, if you are listening to your favourite classical piece of music for solo lute and you switch to what happens to be a final battle scene in a superhero movie, you will also experience a sudden jump upwards in the audio level. However, now, the radio station is much quieter than the video stream.

In the past, the simple solution to this problem was to have an offset gain (or level control) for each source. So, for example, you could set the level of the CD player lower than the level of the DVD player (because your CD collection consisted of only heavy metal music, but you only watched art films). However, since we now use different apps on a single HDMI device, this offset is not the solution, since the Beosound Theatre can't 'know' which app you are using on an external device.

There are a number of different solutions to this problem, and each is applicable to different users.

- *Fix the problem manually.* It's important to remember that the volume control is not a predictor of the output level; it's an indication of what the Beosound Theatre is doing to the input signal. The volume setting in any sound system is analogous to the amount that you're pushing down on the accelerator pedal in your car. The more you push, the faster (louder) you'll go; but your actual speed (audio output level) is dependent on other external factors, such as whether you're going up or down a hill (the level of the input signal).
- *Increase the 'Compression' setting in the Listening Modes.* This will reduce the overall dynamic range of all input signals as described on page 7, and thus reduce the jump in levels when changing sources or apps on a single source. However, the 'price' of this change is that the dynamic range of all signals will be reduced all the time. Therefore, the explosions will not be much louder

than the dialogue in a superhero movie, for example.

Why is there sound coming out of a loudspeaker when there's no signal in its input channel when the Sound Mode is set to use "Direct" in the Spatial Processing?

When Bass Management is ON, there will be sound reproduced by all loudspeakers that are part of the Listening Position, even when the Spatial Processing is set to "Direct".

Why can't I hear the LFE input channel?

If Bass Management is turned off and you do not have an output assigned with a 'Subwoofer' Speaker Role, then the LFE channel will not be routed to any output.

However, if Bass Management is turned on and/or if you have a Subwoofer Speaker Role assigned to an output, then the LFE channel will be reproduced.

Appendix 1: A Very Short History Lesson

Once upon a time, in the incorrectly-named ‘Good Old Days’, audio systems were relatively simple things. There was a single channel of audio picked up by a gramophone needle or transmitted to a radio, and that single channel was reproduced by a single loudspeaker. Then one day, in the 1930s, a man named Alan Blumlein was annoyed by the fact that, when he was watching a film at the cinema and a character moved to one side of the screen, the voice still sounded like it was coming from the centre (because that’s where the loudspeaker was placed). So, he invented a method of simultaneously reproducing more than one channel of audio to give the illusion of spatial changes.¹ Eventually, that system was called ‘stereophonic’ audio.

The word ‘stereo’ first appeared in English usage in the late 1700s, borrowed directly from the French word ‘stéréotype’: a combination of the Greek word στερεό (roughly pronounced ‘stereo’) meaning ‘solid’ and the Latin ‘typus’ meaning ‘form’. ‘Stereotype’ originally meant a *letter* (the ‘type’) printed (e.g. on paper) using a *solid plate* (the ‘stereo’). In the mid-1800s, the word was used in ‘stereoscope’ for devices that gave a viewer a three-dimensional visual representation using two photographs. So, by the time Blumlein patented spatial improvements in audio recording and reproduction in the 1930s, the word had already been in used to mean something akin to ‘a three-dimensional representation of something’ for over 80 years.

Over the past 90 years, people have come to mis-understand that ‘stereo’ audio implies only two channels, but this is incorrect, since Blumlein’s patent was also for three loudspeakers, placed on the left, centre, and right of the

cinema screen. In fact, a ‘stereo’ system simply means that it uses more than one audio channel to give the impression of sound sources with different locations in space. So it can easily be said that all of the other names that have been used for multichannel audio formats² since 1933 are merely new names for ‘stereo’ (a technique that we call ‘rebranding’ today).

At any rate, most people were introduced to ‘stereophonic’ audio either through a two-channel LP, cassette tape, CD, or a stereo FM radio receiver. Over the years, systems with more and more audio channels have been developed; some with more commercial success than others. The table below contains a short list of examples.

Some of the information in that list may be surprising, such as the existence of a 7-channel audio system in the 1950s, for example. Another interesting thing to note is the number of multichannel formats that were not ‘merely’ an accompaniment to a film or video format.

One problem that is highlighted in that list is the confusion that arises with the names of the formats. One good example is ‘Dolby Digital’, which was introduced as a name not only for a surround sound format with 5.1 audio channels, but also the audio encoding method that was required to deliver those channels on optical film. So, by saying ‘Dolby Digital’ in the mid-1990s, it was possible that you meant one (or both) of two different things. Similarly, although SACD and DVD-Audio were formats that were capable of supporting up to 6 channels of audio, there was no requirement and therefore no guarantee that the content be multichannel or that the LFE channel actually contain low-frequency content. This grouping of features under one name still causes confusion when discussing the specifics of a given system, as we’ll discuss below in the section on Dolby Atmos.

¹Patent GB394325A: ‘Improvements in and relating to sound-transmission, sound-recording and sound-reproducing systems’

²like ‘quadraphonic’, ‘surround sound’, ‘multichannel audio’, and ‘spatial audio’, just to name a few obvious examples

Format	Introduced	Channels	Note
Edison phonograph cylinders	1896	1	
Berliner gramophone record	1897	1	
Wire recorder	1898	1	
Optical (on film)	ca. 1920	1	
Magnetic tape	1928	2	
Fantasound	1940	3 / 54	3 channels through 54 loudspeakers
Cinerama	1950s	7.1	Actually 7-channels
Stereophonic LP	1957	2	
Compact Cassette	1963	2	
Q-8 magnetic tape cartridge	1970	4	
CD-4 Quad LP	1971	4	
SQ (Stereo Quadraphonic) LP	1971	4	
IMAX	1971	5.1	
Dolby Stereo	1975	4	Also known as 'Dolby Surround'
Compact Disc	1983	2	
DAT	1987	2	
Mini-disc	1992	2	
Digital Compact Cassette	1992	2	
Dolby Digital	1992	5.1	
DTS Coherent Acoustics	1993	5.1	
Sony SDDS	1993	7.1	
SACD	1999	5.1	Actually 6 full-band channels
Dolby EX	1999	6.1	
Tom Holman / TMH Labs	1999	10.2	
DVD-Audio	2000	5.1	Actually 6 full-band channels
DTS ES	2000	6.1	
Dolby Surround 7.1	2010	7.1	
NHK: Ultra-high definition TV	2005	22.2	
Auro 3D	2005	9.1 to 26.1	
Dolby Atmos	2012	up to 24.1.10	

Looking at the column listing the number of audio channels in the different formats, you may have three questions:

1. "Why does it say only 4 channels for Dolby Stereo? I saw Star Wars in the cinema, and I remember a LOT more loudspeakers on the walls."
2. "What does the '.1' mean? How can you have one-tenth of an audio channel?"
3. "Some of the channel listings have one or two numbers and some have three, which I've never seen before. What do the different numbers represent?"

10.0.1 Input Channels and Speaker Roles

A 'perfect' two-channel stereo system is built around two matched loudspeakers, one on the left and the other on the right, each playing its own dedicated audio channel. However, when better sound systems were developed for movie theatres, the engineers (starting with Blumlein) knew that it

was necessary to have more than two loudspeakers because not everyone is sitting in the middle of the theatre. Consequently, a centre loudspeaker was necessary to give off-centre listeners the impression of speech originating in the middle of the screen. In addition, loudspeakers on the side and rear walls helped to give the impression of envelopment for effects such as rain or crowd noises.

It is recommended but certainly not required, that a given Speaker Role should only be directed to one loudspeaker. In a commercial cinema, for example, a single surround channel is most often produced by many loudspeakers arranged on the side and rear walls. This can also be done in larger home installations where appropriate.

Similarly, in cases where the Beosound Theatre is accompanied by two larger front loudspeakers, it may be preferable to use the three front-firing outputs to all produce the Centre Front channel (instead of using the centre output only).

10.0.2 x.1 ?

The engineers also realised that it was not necessary that all the loudspeakers be big, and therefore requiring powerful amplifiers. This is because larger loudspeakers are only required for high-level content *at low frequencies*, and we humans are terrible at locating low-frequency sources when we are indoors. This meant that effects such as explosions and thunder that were loud, but limited to the low-frequency bands could be handled by a single large unit instead; one that handled all the content below the lowest frequencies capable of being produced by the other loudspeakers' woofers. So, the systems were designed to rely on a powerful **sub-woofer** that driven by a special, dedicated **Low Frequency Effects** (or **LFE**) audio channel whose signals were limited up to about 120 Hz. However, as is discussed in the section on Bass Management, it should not be assumed that the LFE input channel is *only* sent to a subwoofer; nor that the *only* signal produced by the subwoofer is the LFE channel. This is one of the reasons it's important to keep in mind that the LFE *input* channel and the subwoofer *output* channel are separate concepts.

Since the LFE channel only contains low frequency content, it has only a small fraction of the bandwidth of the main channels. ('Bandwidth' is the total frequency width of the signal. In the case of the LFE channel it is up to about 120 Hz.³ In the case of a main channel, it is up to about 20,000 Hz; however these values are not only fuzzy but dependent on the specifications of distribution format, for example.) Although that fraction is approximately 120/20000, we generously round it up to 1/10 and therefore say that, relative to a main audio channel like the Left Front, the LFE signal is only 0.1 of a channel. Consequently, you'll see audio formats with something like '5.1 channels' meaning '5 main channels and an LFE channel'.⁴

LFE ≠ Subwoofer

Many persons jump to the conclusion that an audio input with an LFE channel (for example, a 5.1 or a 7.1.4 signal) means that there is a 'subwoofer' channel; or that a loudspeaker configuration with 5 main loudspeakers and a subwoofer is a 5.1 configuration. It's easy to make this error because those are good descriptions of the way many systems have worked in the past.

However, systems that use bass management break this direct connection between the LFE input and the subwoofer output. For example, if you have two large loudspeakers such as Beolab 50s or Beolab 90s for your Lf/Rf pair, it may not be necessary to add a subwoofer to play signals with an LFE channel. In fact, in these extreme cases, adding a subwoofer could result in *downgrading* the system. Similarly, it's possible to play 2.0-channel signal through a system with two smaller loudspeakers and a single subwoofer.

Therefore, it's important to remember that the 'x.1' classification and the discussion of an 'LFE' channel are descriptions of the *input* signal. The *output* may or may not have one or more subwoofers; and these two things are essentially made independent of each other using a bass management system.

10.0.3 x.y.4 ?

If you look at the table on Page 23, you'll see that some formats have very large numbers of channels, however, these numbers can be easily mis-interpreted. For example, in both the '10.2' and '22.2' systems, some of the audio channels are intended to be played through loudspeakers *above* the listeners, but there's no way to know this simply by looking at the number of channels. This is why we currently use a new-and-improved method of listing audio channels with three numbers instead of two.

- The first number tells you how many 'main' audio channels there are. In an optimal configuration, these should be reproduced using loudspeakers at the listeners' ear heights.
- The second number tells you how many LFE channels there are.
- The third number tells you how many audio channels are intended to be reproduced by loudspeakers located above the listeners.

For example, you'll notice looking at the table below, that a 7.1.4 channel system contains seven main channels around the listeners, one LFE channel, and four height channels.

³In fact, different formats have different bandwidths for the LFE channel, but 120 Hz is a good guess.

⁴This is similar to the way rental apartments are listed in Montréal, where it's common to see a description such as a 3½; meaning that it has a living room, kitchen, bedroom, and a bathroom (which is obviously half of a room).

B&O Speaker Role	Standard Format			
	2.0	5.1	7.1	7.1.4
Left front	x	x	x	x
Right front	x	x	x	x
Centre front		x	x	x
LFE		x	x	x
Left surround		x	x	x
Right surround		x	x	x
Left back			x	x
Right back			x	x
Left front height				x
Right front height				x
Left surround height				x
Right surround height				x

It's worth mentioning here that there are many other standard formats; the table above only lists some of the ones in most common usage. For example, some movies were released with audio in a 6.1-channel format, which had a single Centre Back channel instead of the Left Back and Right Back. Also, multichannel music released on SACD and DVD-Audio is typically mastered in a 5.0-channel format, and intended to be played through five large, matched loudspeakers instead of using a mis-matched system (e.g. with smaller surround loudspeakers) and no subwoofer to play an LFE channel.⁵

10.0.4 Channel names

As was already mentioned, originally, there was just one audio channel in an audio recording or transmission. This evolved

into two: a *left* and a *right* channel. However, today we have a wide variety of audio formats with large numbers of audio channels, defined by different companies and organisations. Unfortunately for the consumer, the result is different names for the same audio channels when switching between different formats.

The Speaker Roles listed in the Beosound Theatre use the same names that have appeared in Bang & Olufsen televisions since the introduction of the Beoplay V1 and the Beovision 11. This was done to maintain consistency within B&O's own portfolio, however, it may be necessary to 'translate' these names to other naming systems occasionally. As a result, the table below is included to (hopefully) alleviate some confusion.

B&O Speaker Role		Alternative Name	
	Abbreviation		Abbreviation
Left Front	Lf	Left	L
Right Front	Rf	Right	R
Centre Front	Cf	Centre	C
<i>Subwoofer</i>	<i>Sub</i>		
Left Surround	Ls	Left Side Surround	Lss
Right Surround	Rs	Right Side Surround	Rss
Left Back	Lb	Left Rear Surround	Lrs
Right Back	Rb	Right Rear Surround	Rrs
Left Front Height	Lfh	Left Top Front	Ltf
Right Front Height	Rfh	Right Top Front	Rtf
Left Surround Height	Lsh	Left Top Rear	Ltr
Right Surround Height	Rsh	Right Top Rear	Rtr

⁵Typically using the standard loudspeaker configuration described in ITU-R BS.775, for example.

It is worth noting here that the logic behind the Bang & Olufsen naming system is either to avoid having duplicate letters for different role assignments, or to reserve options for future formats. For example, 'Back' is used instead of 'Rear' to prevent confusion with Right, and Height is used instead of Top because another, higher layer of loudspeakers may be used in future formats.⁶

⁶The 'Ceiling' loudspeaker channel used in multichannel recordings from Telarc is an example of this.

Appendix 2: Dolby Atmos

In 2012, Dolby introduced its Dolby Atmos surround sound technology in movie theatres with the release of the Pixar movie, 'Brave', and support for the system was first demonstrated on equipment for home theatres in 2014. However, in spite of the fact that it has been 10 years since its introduction, it still helps to offer an introductory explanation to what, exactly, Dolby Atmos is.¹

From the perspective of audio / video systems for the home, Dolby Atmos can most easily be thought of as a collection of different things

1. a set of recommendations for loudspeaker configuration that can include loudspeakers located above the listening position
2. a method of supplying audio signals to those loudspeakers that not only use audio **channels** that are intended to be played by a single loudspeaker (e.g. Left Front or Right Surround), but also audio **objects** whose intended spatial positions are set by the mixing engineer, but whose actual spatial position is 'rendered' based on the actual loudspeaker configuration in the customer's listening room.
3. a method of simulating the spatial positions of 'virtual' loudspeakers
4. the option to use loudspeakers that are intentionally directed away from the listening position, potentially increasing the spatial effects in the mix. These are typically called 'up-firing' and 'side-firing' loudspeakers.

In addition to this, Dolby has other technologies that have been enhanced to be used in conjunction with Dolby Atmos-encoded signals. Arguably, the most significant of these is an upmixing / downmixing algorithm that can adapt the input signal's configuration to the output channels.

11.0.1 Loudspeaker configurations

Dolby's Atmos recommendations allow for a large number of different options when choosing the locations of the loudspeakers in your listening room. These range from a simple 2.0.0, traditional two-channel stereo loudspeaker configuration up to a 24.1.10 large-scale loudspeaker array for movie theatres. Figures 11.1 to 11.5 show a small sampling of the most common options.




Figure 11.1: Standard loudspeaker configuration for two-channel stereo. (Lf / Rf)




Figure 11.2: Standard loudspeaker configuration for 5.x multi-channel audio. The actual angles of the surround loudspeakers at 110° shows the reference placement used at Bang & Olufsen for testing and tuning. Note that the placement of the subwoofer is better determined by your listening room's acoustics, but it is advisable to begin with a location near the centre front loudspeaker.

¹For more in-depth explanations, <https://www.dolby.com/technologies/dolby-atmos> is a good place to start and <https://www.dolby.com/about/support/guide/speaker-setup-guides> has a wide range of options for loudspeaker configuration recommendations.




Figure 11.3: Recommended loudspeaker configuration for most 7.x channel audio signals. The actual angles of the loudspeakers shows the reference placement used at Bang & Olufsen for testing and tuning.




Figure 11.5: Loudspeaker positions associated with the speaker roles available in the Theatre, showing a full 7.x.4 configuration.




Figure 11.4: Loudspeaker positions associated with the speaker roles available in the Theatre, showing a full 7.x.4 configuration.

11.0.2 Channels and Objects

Typically, when you listen to audio, regardless of whether it's monophonic or a stereo² signal, you are reproducing some number of audio **channels** that were mixed in a studio. For example, a recording engineer placed a large number of microphones around a symphony orchestra or a jazz ensemble, and then decided on the mix (or relative balance) of those signals that should be sent to a loudspeakers in the left front and right front positions. They did this by listening to the mix through loudspeakers in a standard configuration with the intention that you place your loudspeakers similarly and sit in the correct location.

Consequently, each loudspeaker's input can be thought of as receiving a 'pre-packaged' audio channel of information.

However, in the early 2000s, a new system for delivering audio to listeners was introduced with the advent of powerful gaming consoles. In these systems, it was impossible for the recording engineer to know *where* a sound should be coming from at any given moment in a game with moving players. So, instead of pre-mixing sound effects (like footsteps, for example) in a fixed position, a monophonic recording of the effect (the footsteps) was stored in the game's software, and then the spatial position could be set at the moment of playback. So, if the footsteps should appear on the player's left, then the game console would play them on the left. If the player then turned, the footsteps could be moved to appear in the centre or on the right. In this way different sound **objects** could be 'rendered' instead of merely being reproduced. Of course, the output of these systems was still either loudspeakers or headphones; so

²remember that 'stereo' merely implies 'more than one channel'

the rendered sound objects were mixed with the audio channels (e.g. background music) before being sent to the outputs.

The advantage of a channel-based system is that there is (at least theoretically) a one-to-one match between what the recording or mastering engineer heard in the studio, and what you are able to hear at home. The advantage of an object-based system is that it can not only adapt to the listener's spatial changes (e.g. the location and rotation of a player inside a game environment) but also to changes in loudspeaker configurations. Change the loudspeakers, and you merely tell the system to render the output differently.

Dolby's Atmos system merges these two strategies, delivering audio content using both channel-based and object-based streams. By delivering audio channels that match older systems, it becomes possible to have a mix on a newly-released movie that is compatible with older playback systems. However, newer Dolby Atmos-compatible systems can render the object-based content as well, optimising the output for the particular configuration of the loudspeakers.

11.0.3 Virtual Loudspeakers

Dolby's Atmos processing includes the option to simulate loudspeakers in 'virtual' locations using 'real' loudspeakers placed in known locations. Beosound Theatre uses this Dolby Atmos processing to generate the signals used to create the four virtual outputs described in detail above in Section 8.2.

11.0.4 Up-firing and Side-firing Loudspeakers

A Dolby Atmos-compatible soundbar or loudspeaker can also include output beams that are aimed away from instead of towards the listening position; either to the sides or above the loudspeakers (this concept is described in more detail in Section 8.1).

These are commonly known as 'up-firing' and 'side-firing' loudspeakers. Although Beosound Theatre gives you the option of using a similar concept, it is not merely implemented with a single loudspeaker driver, but with a version of the Beam Width and Beam Direction control used in other high-end Bang & Olufsen loudspeakers. This means that, when using the up-firing and side-firing outputs, more than a single loudspeaker driver is being used to produce the sound. This helps to reduce the beam width, reducing the level of the direct sound at the listening position, which, in turn can help to enhance the spatial effects that can be encoded in a Dolby Atmos mix.

11.1 Upmixing and Downmixing

There are cases where an incoming audio signal was intended to be played by a different number of loudspeakers than are available in the playback system. In some cases, the playback uses fewer loudspeakers (e.g. when listening to a two-channel stereo recording on the single loudspeaker on a mobile phone). In other cases, the playback system has more loudspeakers (e.g. when listening to a monophonic news broadcast on a 5.1 surround sound system). When the number of input channels is larger than the number of outputs (typically loudspeakers), the signals have to be **downmixed** so that you are at least able to hear all the content, albeit at the price of spatially-distorted reproduction. (For example, instruments will appear to be located in incorrect locations, and the spaciousness of a room's reverberation may be lost.) When the number of output channels (loudspeakers) is larger than the number of input channels, then the system may be used to **upmix** the signal.

The Dolby processor in a standard playback device has the capability of performing both of these tasks: either upmixing or downmixing when required (according to both the preferences of the listener). One particular feature included in this processing is the option for a mixing engineer to 'tell' the playback system exactly how to behave when doing this. For example, when downmixing a 5.1-channel movie to a two-channel output, it may be desirable to increase the level of the centre channel to increase the level of the dialogue to help make it more intelligible. Dolby's encoding system gives the mixing engineer this option using 'metadata'; a set of instructions defining the playback system's behaviour so that it behaves as intended by the artist (the mixing engineer). Consequently, the Beosound Theatre gives you the option of choosing the **Downmix** mode, where this processing is done exclusively in the Dolby processor.

However, there are also cases where you may also wish to upmix the signal to more loudspeakers than there are input channels from the source material. For these situations, Bang & Olufsen has developed its own up-/down-mixing processor called **True Image**, which is discussed on page 8.

11.2 Multichannel Configuration Customisation

In a perfect loudspeaker configuration, all loudspeakers are the same distance from the listening position. They have all been calibrated to have the same loudness at the listening position. Also, they are all large, full-range loudspeakers (and therefore, you do not need a subwoofer). However, in the real world, not only do most people *not* have a system like this, but typically there is no single 'listening position' since the system is used by more than one person at the same time.

Consequently, it may be beneficial to manually optimise some of the parameters in a Listening Position to accommodate a

wider listening area. If you do follow the procedures described below, it's important to re-run Beosound Theatre's Room Compensation measurement in order to include the changes that were made to the




Figure 11.6: A reasonably good five-channel setup in the real world with a 'correct' calibration scheme.

If we were to calibrate this system 'perfectly' using the centre of the sofa as our reference 'sweet spot' as shown in Figure 11.6, then we would apply a delay to the Centre Front loudspeaker to make the time of arrival of its signals match the Left Front and Right Front loudspeakers (usually done by setting the Speaker Distance). We would also apply a delay to the surround loudspeakers to do the same. It would also probably be helpful to drop the Speaker Levels of the centre and surround loudspeakers to match the Left Front and Right Front signals because they're closer, and therefore louder.

However, let's consider what happens if you sit on the left side of that sofa. Now, the Left Surround loudspeaker is very close to your left ear; and that has some serious implications on your experience. Firstly, since sound pressure doubles with every halving of distance, then sitting on the left side of the sofa means that you'll get a noticeable boost in the signal level from that one loudspeaker. In addition that signal arrives at your listening position much earlier than the other channels. The same problem, albeit on a much smaller scale, happens with the centre loudspeaker. If its time-alignment delay is calibrated using the centre position, then, if you're sitting on the left side of the sofa, then the Left Front loudspeaker's signal will arrive before the Centre Front. The end result of this is that, if you're sitting on the side of the sofa, you'll have too much from one of the surround loudspeakers and the intelligibility of the dialogue will be reduced a little.

Can we manually improve the calibration to improve the experience for everyone on the sofa?




Figure 11.7: A reasonably good multichannel setup in the real world with an alternative calibration scheme.

One solution is to calibrate both the distance and the level for each loudspeaker relative to the closest person in your listening area. So, for example (in this case), the distances and levels of the Left Front and Left Surround loudspeakers are measured at the left position on the sofa, but the Centre Front loudspeaker is calibrated to the centre of the sofa. The result of this is that the centre speaker will be delayed; but less than it would have been if you had calibrated it as in Figure 11.6, because the Left Front loudspeaker is closer to the person on the left side of the sofa than to the person in the centre of the sofa. Also, the Surround loudspeakers will be delayed much more than they would have been using the scheme in Figure 11.6. However, they'll still be symmetrical. Also, this will result in the centre channel being a bit louder and the surround channels being a little lower in level; both of which are technically incorrect for the person in the sweet spot, but at least it's a mistake in the right direction; so you're improving intelligibility of the dialogue

If you do calibrate the system this way, you'll be incorrectly calibrated at the sweet spot, but your friends on the sides of the sofa will be much happier; and you won't notice too much. In the case of Beosound Theatre or any other Bang & Olufsen television, you can make this configuration just one of your available Speaker Groups or Listening Positions; you can always use another one for a 'perfect' calibration for the sweet spot when you're home alone.

If, after aligning your system using this method, you still find that the dialogue is a little hard to understand, and the surround levels are a little high (this is often the case when your sofa and the surround loudspeakers are all situated against the same wall) you should not be afraid to do the following:

- make the centre channel one or two milliseconds early. You can do this by telling your Beosound Theatre that it's about 30 to 60 cm farther away than it really is.
- raise the level of the centre channel 1 or 2 dB
- drop the level of the surrounds as much as necessary; it's not unusual to have to drop them by as much as 6 dB

if you're against the same wall with them. (Note that this can be done using Beosound Theatre's 'Fader' adjustment instead. This will merely control the relative levels of the Front and Surround / Back loudspeakers; so it's a one-slider solution rather than doing it manually for each loudspeaker output.)

If your listening area is larger, the technique is the same; you calibrate any given loudspeaker in the system to the closest listening position, and then tweak to taste.

The main message here is 'just because your system is configured 'correctly' doesn't mean that it can't sound better'. Don't be more afraid to tweak the adjustments on your calibration than you would be to add cream and sugar to your coffee, or salt and pepper to your meal in a restaurant. As Duke Ellington once said: "If it sounds good, it *is* good."

Appendix 3: What is a “Virtual” Loudspeaker?

12.1 Part 1: Head-Related Transfer Functions

Without connecting external loudspeakers, Bang & Olufsen’s Beosound Theatre has a total of 11 independent outputs, each of which can be assigned any Speaker Role (or input channel). Four of these are called “virtual” loudspeakers – but what does this mean? There’s a brief explanation of this concept earlier in this Technical Sound Guide for the Theatre. However, let’s dig into this concept a little more deeply.

To begin, let’s put a “perfect” loudspeaker in a free field. This means that it’s in a space that has no surfaces to reflect the sound – so it’s an acoustic field where the sound wave is free to travel outwards forever without hitting anything (or at least appear as this is the case). We’ll also put a “perfect” microphone in the same space.



Figure 12.1: A loudspeaker and a microphone (the circle) in a free field: an infinite space completely free of reflective surfaces.

We then send an impulse; a very short, very loud “click” to the loudspeaker. (Actually a perfect impulse is infinitely short and infinitely loud, but this is not only inadvisable but impossible, and probably illegal.)

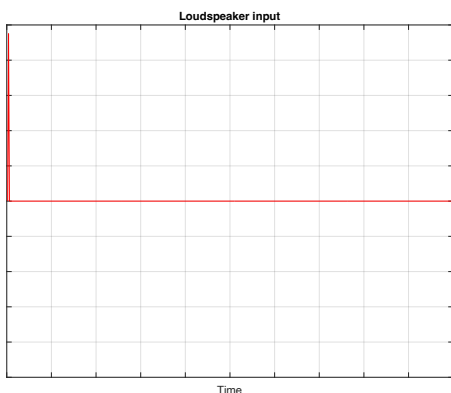


Figure 12.2: The “click” signal that’s sent to the input of the loudspeaker.

That sound radiates outwards through the free field and reaches the microphone which converts the acoustic signal back to an electrical one so we can look at it.

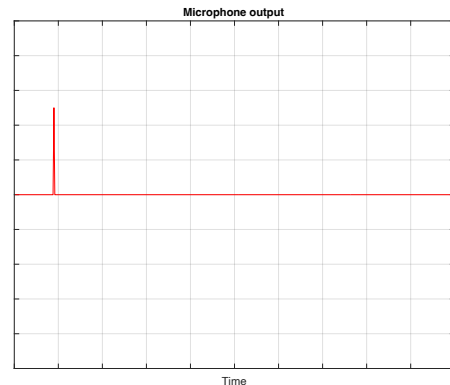


Figure 12.3: The “click” signal that is received at the microphone’s location and sent out as an electrical signal.

There are three things to notice when you compare Figure 3 to Figure 2:

- The signal’s level is lower. This is because the microphone is some distance from the loudspeaker.
- The signal is later. This is because the microphone is some distance from the loudspeaker and sound waves travel pretty slowly.
- The general shape of the signals are identical. This is because I said that the loudspeaker and the microphone were both “perfect” and we’re in a space that is completely free of reflections.

What happens if we take away the microphone and put you in the same place instead?



Figure 12.4: The microphone has been replaced by something more familiar.

If we now send the same click to the loudspeaker and look at the “outputs” of your two eardrums (the signals that are sent to your brain), these will look something like this:

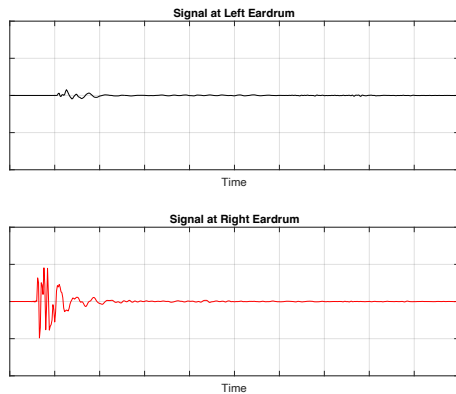


Figure 12.5: The outputs of your two eardrums with the same “click” signal from the loudspeaker.

These two signals are obviously very different from the one that the microphone “hears” which should not be a surprise: ears aren’t microphones. However, there are some specific things of which we should take note:

- The output of the left eardrum is lower than that of the right eardrum. This is largely because of an effect called “head shadowing” which is exactly what it sounds like. The sound is quieter in your left ear because your head is in the way.
- The signal at the right eardrum is earlier than at the left eardrum. This is because the left eardrum is not only farther away, but the sound has to go around your head to get there.
- The signal at the right eardrum is earlier than the output of the microphone output (in Figure 3) because it’s closer to the loudspeaker. (I put the microphone at the location of the centre of the simulated head.) Similarly the left ear output is later because it’s farther away.
- The signal at the right eardrum is full of spikes. This is mostly caused by reflections off the pinna (the flappy thing on the side of your head that you call your “ear”) that arrive at slightly different times, and all add together to make a mess.
- The signal at the left eardrum is “smoother”. This is because the head itself acts as a filter reducing the levels of the high frequency content, which tends to make things less “spiky”.
- Both signals last longer in time. This is the effect of the ear canal (the “hole” in the side of your head that you should NOT stick a pencil in) resonating like a little organ pipe.

The difference between the signals in Figures 2 and 4 is a measurement of the effect that your head (including your shoulders, ears/pinnae) has on the **transfer** of the sound from the loudspeaker to your eardrums. Consequently, we geeks

call it a “**head-related transfer function**” or **HRTF**. I’ve plotted this HRTF as a measurement of an impulse in time – but I could have converted it to a frequency response instead (which would include the changes in magnitude and phase for different frequencies).

Here’s the cool thing: If I put a pair of headphones on you and played those two signals in Figure 5 to your two ears, you might be able to convince yourself that you hear the click coming from the same place as where that loudspeaker is located.

Although this sounds magical, don’t get too excited right away. Unfortunately, as with most things in life, reality tends to get in the way for a number of reasons:

- Your head and ears aren’t the same shape as anyone else’s. Your brain has lived with your head and your ears for a long time, and it’s learned to correlate your HRTFs with the locations of sound sources. If I suddenly feed you a signal that uses my HRTFs, then this trick may or may not work, depending on how similar we are. This is just like borrowing someone else’s glasses. If you have roughly the same prescription, then you can see. However, if the prescriptions are very different, you’ll get a headache very quickly.
- In reality, you’re always moving. So, even if the sound source is not moving, the specific details of the HRTFs are always changing (because the relative positions and angles to your ears are changing) but my system doesn’t know about this – so I’m simulating a system where the loudspeaker moves around you as you rotate your head. Since this never happens in real life, it tends to break the simulation.
- The stuff I showed above doesn’t include reflections, which is how you determine distance to sources. If I wanted to include reflections, each reflection would have to have its own HRTF processing, depending on its angle relative to your head.

However, hypothetically, this can work, and lots of people have tried. The easiest way to do this is to not bother measuring anything. You just take a “dummy head” -a thing that is the same size as an average human head (maybe with an average torso) and average pinnae* – but with microphones where the eardrums are – and you plunk it down in a seat in a concert hall and record the outputs of the two “ears”. You then listen to this over earphones (we don’t use headphones because we want to remove your pinnae from the equation) and you get a “you are there” experience (assuming that the dummy head’s dimensions and shape are about the same as yours). This is what’s known as a **binaural recording** because it’s a recording that’s done with two ears (instead of two or more “simple” microphones).

If you want to experience this for yourself, plug a pair of headphones into your computer and do a search for the

“Virtual Barber Shop” video. However, if you find that it doesn’t work for you, don’t be upset. It just means that you’re different: just like everyone else.¹ Typically, recordings like this have a strange effect of things sounding very close in the front, and farther away as sources go to the sides. (Personally, I typically don’t hear anything in the front. All of the sources sound like they’re sitting on the back of my neck and shoulders. This might be because I have a fat head (yes, yes... I know...) and small pinnae (yes, yes... I know...) – or it might indicate some inherent paranoia of which I am not conscious.)

12.2 Part 2: Cancelling Crosstalk

In Part 1, I talked at how a binaural recording is made, and I also mentioned that the spatial effects may or may not work well for you for a number of different reasons.

Let’s go back to the free field with a single “perfect” microphone to measure what’s happening, but this time, we’ll send sound out of two identical “perfect” loudspeakers. The distances from the loudspeakers to the microphone are identical. The only difference in this hypothetical world is that the two loudspeakers are in different positions (measuring as a rotational angle) as shown in Figure 1.




Figure 12.6: Two identical, “perfect” loudspeakers in a free field with a single “perfect” microphone.

In this example, because everything is perfect, and the space is a free field, then output of the microphone will be the sum of the outputs of the two loudspeakers. (In the same way that if your dog and your cat are both asking for dinner simultaneously, you’ll hear dog+cat and have to decide which is more annoying and therefore gets fed first...)

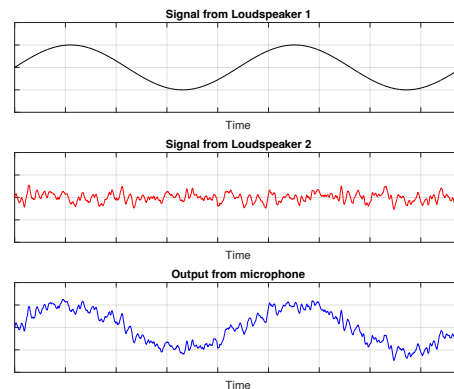


Figure 12.7: The output from the microphone is the sum of the outputs from the two loudspeakers. At any moment in time, the value of the top plot + the value of the middle plot = the value of the bottom plot.

IF the system is perfect as I described above, then we can play some tricks that could be useful. For example, since the output of the microphone is the sum of the outputs of the two loudspeakers, what happens if the output of one loudspeaker is identical to the other loudspeaker, but reversed in polarity?

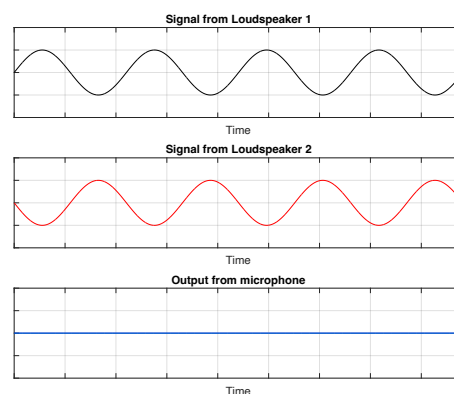


Figure 12.8: If the output of Loudspeaker 1 is exactly the same as the output of Loudspeaker 2 except for polarity, then the sum (the output of the microphone) is always 0.

In this example, we’re manipulating the signals so that, when they add together, you nothing at the output. This is because, at any moment in time, the value of Loudspeaker 2’s output is the value of Loudspeaker 1’s output * -1. So, in other words, we’re just subtracting the signal from itself at the microphone and we get something called “perfect cancellation” because the two signals cancel each other at all times.

Of course, if anything changes, then this perfect cancellation won’t work. For example, if one of the loudspeakers moves a little farther away than the other, then the system is broken, as shown below.

¹Of course, depressingly typically, it goes without saying that the sizes and shapes of commercially-available dummy heads are based on averages of measurements of men only. Neither women nor children are interested in binaural recordings or have any relevance to such things, apparently...

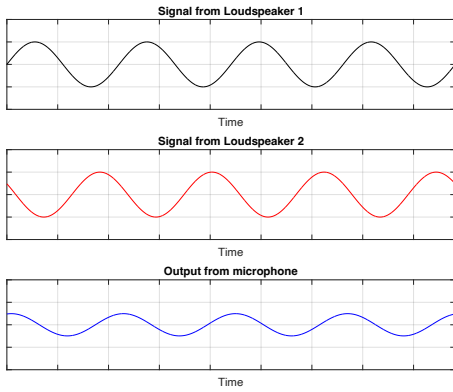


Figure 12.9: A small shift in time in the output of Loudspeaker 2 causes the cancellation to stop working so well.

Again, everything that I've said above only works when everything is perfect, and the loudspeakers and the microphone are in a free field; so there are no reflections coming in and ruining everything.

We can now combine these two concepts:

1. using binaural signals to simulate a sound source in a location (although this would normally be done using playback over earphones to keep it simple) and
2. using signals from loudspeakers to cancel each other at some location in space

to create a system for making virtual loudspeakers.

Let's suspend our adherence to reality and continue with this hypothetical world where everything works as we want... We'll replace the microphone with a person and consider what happens. To start, let's just think about the output of the left loudspeaker.




Figure 12.10: The output of the left loudspeaker reaches both ears with different time/frequency characteristics caused by the HRTF associated with that sound source location.

If we plot the impulse responses at the two ears (the "click"

sound from the loudspeaker after it's been modified by the HRTFs for that loudspeaker location), they'll look like this:




Figure 12.11: The impulse responses of the HRTFs for a sound source at 30° left of centre.

What if we were able to send a signal out of the right loudspeaker so that it cancels the signal from the left loudspeaker at the location of the right eardrum?

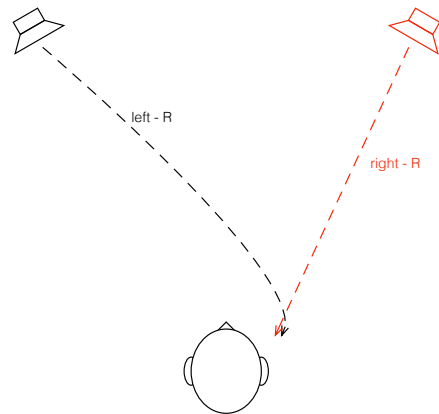


Figure 12.12: What if we could cancel the signal from the left loudspeaker at the right ear using the right loudspeaker?

Unfortunately, this is not quite as easy as it sounds, since the HRTF of the right loudspeaker at the right ear is also in the picture, so we have to be a bit clever about this.

So, in order for this to work we:

- Send a signal out of the left loudspeaker. We know that this will get to the right eardrum after it's been messed up by the HRTF. This is what we want to cancel...
- ...so we take that same signal, and
 - filter it with the inverse of the HRTF of the right loudspeaker (to undo the effects of the HRTF of the right loudspeaker's signal at the right ear)

- filter that with the HRTF of the left loudspeaker at the right ear (to match the filtering that's done by your head and pinna)
- multiply by -1 (so that it will cancel when everything comes together at your right eardrum)
- and send it out the right loudspeaker.

Hypothetically, that signal (from the right loudspeaker) will reach your right eardrum at the same time as the unprocessed signal from the left loudspeaker and the two will cancel each other, just like the simple example shown in Figure 3. This effect is called crosstalk cancellation, because we use the signal from one loudspeaker to cancel the sound from the other loudspeaker that crosses to the wrong side of your head.

This then means that we have started to build a system where the output of the left loudspeaker is heard ONLY in your left ear. Of course, it's not perfect because that cancellation signal that I sent out of the right loudspeaker gets to the left ear a little later, so we have to cancel the cancellation signal using the left loudspeaker, and back and forth forever.

If, at the same time, we're doing the same thing for the other channel, then we've built a system where you have the left loudspeaker's signal in the left ear and the right loudspeaker's signal in the right ear; just like a pair of headphones!

However, if you get any of these elements wrong, the system will start to under-perform. For example, if the HRTFs that I use to predict your HRTFs are incorrect, then it won't work as well. Or, if things aren't time-aligned correctly (because you moved) then the cancellation won't work.

12.3 Part 3: Virtualisation

In Part 1, I talked about how a binaural audio signal can (hypothetically, with HRTFs that match your personal ones) be used to simulate the sound of a source (like a loudspeaker, for example) in space. However, to work, you have to make sure that the left and right ears get completely isolated signals (using earphones, for example).

In Part 2, I showed how, with enough processing power, a large amount of luck (using HRTFs that match your personal ones PLUS the promise that you're in exactly the correct location), and a room that has no walls, floor or ceiling, you can get a pair of loudspeakers to behave like a pair of headphones using crosstalk cancellation.

There's not much left to do to create a virtual loudspeaker. All we need to do is to:

- Take the signal that should be sent to a right surround loudspeaker (for example) and filter it using the HRTFs that correspond to a sound source in the location that this loudspeaker would be in. REMEMBER that this signal

has to get to your two ears since you would have used your two ears to hear an actual loudspeaker in that location.

- Send those two signals through a crosstalk cancellation processing system that causes your two loudspeakers to behave more like a pair of headphones.

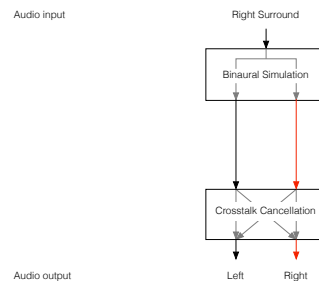


Figure 12.13: A block diagram of the system described above for the Right Surround loudspeaker as an example.

One nice thing about this system is that the crosstalk cancellation is only there to ensure that the actual loudspeakers behave more like headphones. So, if you want to create more virtual channels, you don't need to duplicate the crosstalk cancellation processor. You only need to create the binaurally-processed versions of each input signal and mix those together before sending the total result to the crosstalk cancellation processor, as shown below.




Figure 12.14: You only need one crosstalk cancellation system for any number of virtual channels. This is good because it saves on processing power.

So, there are some important things to realise after having read this:

- All "virtual" loudspeakers' signals are actually produced by the left and right loudspeakers in the system. In the case of the Beosound Theatre, these are the Left and Right Front-firing outputs.
- Any single virtual loudspeaker (for example, the Left Surround) requires BOTH output channels to produce sound.

- If the delays (aka Speaker Distance) and gains (aka Speaker Levels) of the REAL outputs are incorrect at the listening position, then the crosstalk cancellation will not work and the virtual loudspeaker simulation system won't work. How badly it doesn't work depends on how wrong the delays and gains are.
- The virtual loudspeaker effect will be experienced differently by different persons because it's depending on how closely your actual personal HRTFs match those predicted in the processor. So, don't get into fights with your friends on the sofa about where you hear the helicopter...
- The listening room's acoustical behaviour will also have an effect on the crosstalk cancellation. For example, strong early reflections will "infect" the signals at the listening position and may/will cause the cancellation to not work as well. So, the results will vary not only with changes in rooms but also speaker locations.

Finally, it's worth noting that, in the specific case of the Beosound Theatre, by setting the Speaker Distances and Speaker Levels for the Left and Right Front-firing outputs for your listening position, then you have automatically calibrated the virtual outputs. This is because the Speaker Distances and Speaker Levels are compensations for the ACTUAL outputs of the system, which are the ones producing the signal that simulate the virtual loudspeakers. This is the reason why the four virtual loudspeakers do not have individual Speaker Distances and Speaker Levels. If they did, they would have to be identical to the Left and Right Front-firing outputs' values.

Appendix 4: The Influence of Listening Room Acoustics on Loudspeakers

A room comprised of large flat reflective surfaces with little acoustical absorption has a very different acoustical behaviour from a recording or mastering studio where the final decisions about various aspects of a recording are made. Consequently, this must have an effect on a listener's perception of a recording played through a pair (assuming stereo reproduction) of loudspeakers in that room. The initial question to be asked is "what, exactly, are the expected effects of the room's acoustical behaviour in such a case?" The second is "if the room has too much of an effect, how can I improve the situation (e.g. by adding absorption or changing the physical configuration of the system in the room)?" The third, and possibly final question is "how can a loudspeaker compensate (or at least account) for these effects?"

The effect a room's acoustical behaviour has on a loudspeaker's sound can, at a simple level, be considered under three general headings:

- Early Reflections
- Room Modes
- Reverberation

13.1 Early Reflections

Early reflections, from sidewalls and the floor and ceiling, have an influence on both the timbre (tone colour) and the spatial characteristics of a stereo reproduction system. We will only discuss the timbral effects in this section.




Figure 13.1: The sound arriving at a listener from a loudspeaker in a room with only one wall. Note that the sound arrives from two directions - the first is directly from the loudspeaker (in red). The second is a "first reflection" off the wall (in blue).

Let's start by assuming that you have a loudspeaker that has a magnitude response that is perfectly flat - at least from 20 Hz to 20 kHz. We will also assume that it has that response regardless of which direction you measure it in - in other words, it's a perfectly omnidirectional loudspeaker. The

question is, "what effect does the wall reflection have on the measured response of the loudspeaker?"

Very generally speaking, the answer is that you will get a higher level at some frequencies (because the direct sound and the reflection add constructively and reinforce each other) and you will get a lower level at other frequencies (because the direct sound and the reflection work against each other and "cancel each other out"). What is potentially interesting is that the frequencies that add and the frequencies that cancel alternate as you go up the frequency range. So the total result looks like a comb (as in a comb that you use to comb your hair, if, unlike me, you have hair to comb). For example, take a look at Figure 13.2.




Figure 13.2: Distance to loudspeaker = 2 m. Distance to wall = 1 m. Wall is perfectly reflective and the loudspeaker is perfectly omnidirectional. The red line is the magnitude response of the direct sound. The blue line is the magnitude response of the reflected sound. The black line is the magnitude response of the combination.

You can see that, at the very low end, the reflection boosts the level of the loudspeaker by a approximately 5 dB (or almost two times the level) at the listening position. However, as you go up in frequency, the total level drops to about 15 dB less before it starts rising again. As you go up in frequency, the level goes up and down. This alternation actually happens at a regular frequency spacing (e.g. a notch at multiples of 200 Hz) but it doesn't look regular because the X-axis of the plot is logarithmic (which better represents how we hear differences in frequency).

What happens if we move the wall further away? Well, two things will happen. The first is that the reflection will be quieter, so the peaks and notches won't be as pronounced. The second is that the spacing of the peaks and notches in frequency will get closer together. In other words, the effect starts at a lower frequency as can be seen in the example in Figure 13.3.




Figure 13.3: Distance to loudspeaker = 2 m. Distance to wall = 3 m. Wall is perfectly reflective and the loudspeaker is perfectly omnidirectional. The red line is the magnitude response of the direct sound. The blue line is the magnitude response of the reflected sound. The black line is the magnitude response of the combination.

Conversely, if we move the wall closer, we do the opposite (the problem gets worse, but starting at a higher frequency), as can be seen in Figure 13.4.




Figure 13.4: Distance to loudspeaker = 2 m. Distance to wall = 0.25 m. Wall is perfectly reflective and the loudspeaker is perfectly omnidirectional. The red line is the magnitude response of the direct sound. The blue line is the magnitude response of the reflected sound. The black line is the magnitude response of the combination.

So, if you have a room with only one wall which is perfectly reflective, and you have a perfectly omnidirectional loudspeaker, then you can see that your best option is to either put the loudspeaker (and yourself) very far or very close to the wall. That way the artefacts caused by the reflection are either too quiet to do any damage, or have an effect that starts at too high a frequency for you to care. Then again, most room have more than one wall, the walls are not perfectly reflective, and the loudspeaker is not perfectly omnidirectional.

So, what happens in the case where the loudspeaker is more directional or you have some absorption (better known as

“fuzzy stuff”) on your walls? Well, either of these cases will have basically the same effect in most cases since loudspeakers are typically more directional at high frequencies - so you get less high end directed towards the wall. Alternatively, fuzzy stuff tends to soak up high frequencies. So, in either of these two cases, you’ll get less high end in the reflection. Let’s simulate this by putting a low pass filter on the reflection, as shown in Figure 13.5, 13.6 and 13.7 which have identical distances as the simulations in Figures 13.2, 13.3, and 13.4 - for comparison.




Figure 13.5: Distance to loudspeaker = 2 m. Distance to wall = 1 m. Wall is absorptive and/or the loudspeaker is directional at high frequencies. The red line is the magnitude response of the direct sound. The blue line is the magnitude response of the reflected sound. The black line is the magnitude response of the combination.




Figure 13.6: Distance to loudspeaker = 2 m. Distance to wall = 3 m. Wall is absorptive and/or the loudspeakers is directional at high frequencies. The red line is the magnitude response of the direct sound. The blue line is the magnitude response of the reflected sound. The black line is the magnitude response of the combination.




Figure 13.7: Distance to loudspeaker = 2 m. Distance to wall = 0.25 m. Wall is absorptive and/or the loudspeaker is directional at high frequencies. The red line is the magnitude response of the direct sound. The blue line is the magnitude response of the reflected sound. The black line is the magnitude response of the combination.

What you can see in all three of the previous plots is that, as the high frequency content of the reflection disappears, there is less and less effect on the total. The bottom plot is basically a proof of the age-old rule of thumb that says that, if you put a loudspeaker next to a wall, you'll get more bass than if it's farther from the wall. Since there is not much high frequency energy radiated from the rear of most loudspeakers, Figure 13.7 is a pretty good general representation of what happens when a loudspeaker is placed close to a wall. Of course, the exact behaviour of the directivity of the loudspeaker will be different - but the general shape of the total curve will be pretty similar to what you see there.

So, the end conclusion of all of this is that, in order to reduce undesirable artefacts caused by a wall reflection, you can do any combination of the following:

- move the loudspeaker very close to the wall
- move the loudspeaker farther front the wall
- sit very close to the wall
- sit farther away from the wall
- put absorption on the wall

However, there is one interesting effect that sits on top of all of this - that is the fact that what you'll see in a measurement with a microphone is not necessarily representative of what you'll hear. This is because a microphone does not have two ears. Also, the direction the reflection comes from will change how you perceive it. A sidewall reflection sounds different from a floor reflection. This is because you have two ears - one on each side of your head. Your brain uses the sidewall reflections (or, more correctly, how they relate to the direct sound) to

determine, in part, how far away a sound source is. Also, since, in the case of sidewall reflections, your two ears get two different delay times on the reflection (usually), you get two different comb-filter patterns, where the peaks in one ear can be used to fill in the notches in the other ear and vice versa. When the reflection comes from the floor or ceiling, your two ears get the same artefacts (since your two ears are the same distance to the floor, probably). Consequently, it's easily noticeable (and it's been proven using science!) that a floor or ceiling reflection has a bigger timbral effect on a loudspeaker than a lateral (or sideways) reflection.

13.2 Room Modes

Room modes are a completely different beast - although they exist because of reflections. If you pluck a guitar string, you make a deflection in the string that moves outwards until it hits the ends of the string. It then bounces back down the string, bounces again, etc. etc. As the wave bounces back and forth, it settles in to a total result where it looks like the string is just bouncing up and down like a skipping rope. The longer the string, the lower the note, because it takes longer for the wave to bounce back and forth on the string. You can also lower the note by lowering the tension of the string, since this will slow down the speed of the wave moving back and forth on it. The last way to lower the note is to make the string heavier (e.g. by making it thicker) - since a heavier string is harder to move, the wave moves slower on it.

The air in a pipe behaves exactly the same way. If you "pluck" the air in the middle of a pipe (say, by clapping our hands, or coughing, or making any noise at all) then the sound wave travels along the pipe until it hits the end. Whether the end of the pipe is capped or not, the wave will bounce back and travel back through the pipe in the opposite direction from whence it came.¹ As the wave bounces back and forth off the two ends of the pipe, it also settles down (just like the guitar string) into something called a "standing wave". This is the pipe's equivalent of the skipping rope behaviour in the string. The result is that the pipe will "resonate" or ring at a note. The longer the pipe, the lower the note because the speed of the sound wave moving in air in the pipe stays the same, but the longer the pipe, the longer it takes for the wave to bounce back and forth. This is basically how all woodwind instruments work.

What's interesting is that, in terms of resonance, a room is basically a big pipe. If you "pluck" the air in the room (say, by making sound with a loudspeaker) the sound wave will move down the room, bounce off the wall, go back through the room, bounce off the opposite wall, etc. etc. (Of course, other things are happening, but we'll ignore those.) This effect is most obvious on a graph by putting some sound in a room and stopping suddenly. Instead of actually stopping, you can see

¹Whether the pipe is closed (capped) or open only determines the characteristic of the reflection - there will be a reflection either way.

the room “ringing” (exactly in the same way that a bell rings when it’s been hit) at a frequency that gradually decays as time goes by. However, it’s important to remember that this ringing is always happening – even while the sound is playing. So, for example, a kick drum “thump” comes out of the speaker which “plucks” the room mode and it rings, while the music continues on.




Figure 13.8: The concept of the effect of a room mode and Active Room Compensation. See the associated text for an explanation.

Figure 13.8 shows the concept of the effect of a room mode and how it’s dealt with by Active Room Compensation. The sound coming out of the loudspeaker is shown on the top plot, in black. The response of the loudspeaker and a single room mode is shown below, in red. You can see there that the room mode keeps “ringing” at one frequency after the sound from the loudspeaker stops.

There are two audible effects of this. The first is that, if your music contains the frequency that the room wants to resonate at, then that note will sound louder. When you hear people talk of “uneven bass” or a “one-note-bass” effect, one of the first suspects to blame is a prominent room mode.

The second is that, since the mode is ringing along with the music, the overall effect will be muddiness. This is particularly true when one bass note causes the room mode to start ringing, and this continues when the next bass note is playing. For example, if your room rings on a C#, and the bass plays a C# followed by a D – then the room will continue to at C#, conflicting with the D and resulting in “mud”. This is also true if the kick drum triggers the room mode, so you have a kick drum “plucking” the room ringing on a C# all through the track. If the tune is in the key of F, then this will not be pretty.²

In order for the loudspeaker to compensate for the effect of the room mode, it has to not only produce the signal it should (shown in black) but it must also produce a signal that

counter-acts the ringing in the room mode. This is shown in the lower plot in blue. As can be seen there (most easily in the ringing after the signal has stopped), the loudspeaker’s compensation signal (the blue curve) is the mirror image of the room’s “misbehaviour” (in red). If you add these two curves together, the result is that they cancel each other out, and the result is the black curve.

If you would like to calculate a prediction of where you’ll have a problem with a room mode, you can use the following equation:

- metric version:
frequency = $172 / (\text{length in m})$
- imperial version:
frequency = $558 / (\text{length in feet})$

This calculation will produce the fundamental frequency of the room mode in Hz for the dimension of the room represented by “length”. Your most audible modal problems will be at the frequencies calculated using either of the equations above, and multiples of them (e.g. 2 times the result, 3 times the result, and so on).

So, for example, if your room is 5 m wide, your worst-case modes (for the room’s width) will be at $172 / 5 = 34.4$ Hz, as well as 68.8 Hz, 103.2 Hz and so on. Remember that these are just predictions – but they’ll come pretty close. You should also remember that this assumes that you have completely immovable walls and no absorption – if this is not true, then the severity of the actual problem will vary accordingly.

Sadly, there is not much you can do about room modes. There are ways to manage them, including, but not exclusive to the following strategies:

- make sure that the three dimensions of your listening room are not related to each other with simple ratios
- put up membrane absorbers or slot absorbers that are tuned to the modal frequencies
- place your loudspeaker in a node – a location in a room where it does not couple to a problematic mode (however, note that one mode’s node is another mode’s antinode)
- sit in a node – a location in a room where you do not couple to a problematic mode (see warning above)
- use room correction DSP software such as ARC in the Beolab 90

²Do a search for “tritone” or “diabolus in musica”.

13.3 Reverberation

Reverberation is what you hear when you clap your hands in a big cathedral. It's the collection of a lot of reflections bouncing from everywhere as you go through time. When you first clap your hands, you get a couple of reflections that come in separated enough in time that they get their own label – “early reflections”. After that, there are so many reflections coming from so many directions, and so densely packed together in time, that we can't separate them, so we just call them “reverberation” or “reverb” (although you'll often hear people call it “echo” which is the wrong word to use for this).

Reverb is what you get when you have a lot of reflective surfaces in your room – but since it's so irregular in time and space, it just makes a wash of sound rather than a weird comb-filter effect like we saw with a single reflection. So, although it makes things “cloudy” – it's more like having a fog on your glasses instead of a scratch, or a soft-focus effect on a kitschy photograph of a field of flowers.

13.4 Solutions

As we've seen, if your listening room is normal, you have at least these three basic acoustic problems to deal with. Each problem has a different solution...

The first solution has already been started for you. The final tuning of every Bang & Olufsen loudspeaker (including the Beolab 90) is voiced in at least four rooms with very different acoustical behaviours ranging from a very “dead” living room with lots of absorptive and diffusive surfaces to a larger and very “live” space with a minimalistic decorating, and large flat surfaces. Once we have a single sound design that is based on the common elements those rooms, we test the loudspeakers in more rooms to ensure that they'll behave well under all conditions.

The second solution is Beolab 90's Active Room Compensation which will correct the effects of boundaries (walls) and room modes on the timbre of the loudspeaker at the listening position(s). Using measurements of the characteristics of the loudspeaker at the listening positions, the ARC algorithm then creates a filter that is used to “undo” these effects. For example, if the loudspeaker is close to a wall (which will generally result in a boosted bass) then the filter will reduce the bass symmetrically. Similarly, ringing caused by room modes will be actively cancelled by both Beolab 90s. That way, the loss in the filter and the gain due to the room will cancel each other.

The third solution is unique to the Beolab 90: Beam Width Control. This allows you to customise the relative levels of the direct sound and the reflected sound at the listening position. The result of this is that, even if you have acoustically reflective side walls, the Beolab 90 can still deliver an accurate

and precise representation of the spatial presentation of your stereo recordings.

13.5 Conclusions

Of course, this section does not cover everything there is to know about room acoustics. And, of course, you can't expect a loudspeaker to sound exactly the same in every room. If that were true, there would be no such thing as a “good” concert hall. A room's acoustical behaviour affects the sound of all sound sources in the room. On the other hand, humans also have an amazing ability to adapt: in other words you “get used to” the characteristics of your listening room.

However, there is no debate that, due to many issues (the first two that come to mind are frequency range and directivity) two different loudspeakers will behave differently from each other in two different rooms. In other words, if you listen to loudspeaker “A” and loudspeaker “B” in a showroom of a shop, you might prefer loudspeaker “A” – but if you took them home, you might prefer loudspeaker “B”. This would not be surprising, since what you hear is not only the loudspeaker but the loudspeaker “filtered” by the listening room. This is exactly why, even with automated room compensation algorithms, some fine tuning may be necessary to achieve a sound that best suits your room and your tastes.