

Beolab 28

Technical information

Bang & Olufsen A/S

This manual is for information purposes only and is not legally binding.
April 29, 2021

Contents

1	Introduction	4
1.0.1	Beam Width Control	4
1.0.2	ARC: Active Room Compensation	5
2	Control Parameters	6
2.1	Beam Width Control	6
2.1.1	Narrow	6
2.1.2	Wide	6
2.1.3	Comment	7
2.2	Stereo Pairing	7
2.3	Speaker Role	7
2.4	Active Room Compensation	7
2.5	Volume	7
2.6	Tone Controls	7
2.6.1	Treble	7
2.6.2	Bass	8
3	System	9
3.1	Max Volume	9
3.2	Startup Volume	9
4	Setup Recommendations	9
4.1	Bang & Olufsen television as source	9
4.2	ARC measurements	9
5	Features and Setup	9
5.1	Resonance-based Sound Design	9
5.2	Adaptive Bass Linearisation (ABL) and Thermal Protection	10
6	Appendix 1: Recommendations for Critical Listening	11
6.1	Loudspeaker Configuration	11
6.2	Room Acoustics	12
6.3	Source Devices	12
6.3.1	Optical cables	12
7	Appendix 3: The Influence of Room Acoustics on Loudspeakers	13
7.1	Early Reflections	13
7.2	Room Modes	15
7.3	Reverberation	16
7.4	Solutions	16
7.5	Conclusions	16
8	Appendix 4: Loudspeaker Directivity and Distance Perception in Stereo Imaging	18
8.1	Distance Perception in Real Life	18
8.2	Distance Perception in a Stereo Recording	18
8.3	Combining the Two	18

9

Appendix 6: ABL - Adaptive Bass Linearisation

19

9.1

A General Introduction to ABL

19

Introduction

Unlike almost all other so-called “connected” loudspeakers, the Beolab 28’s acoustical characteristics can be altered by the user to behave differently for different listening situations or listener preferences.

Imagine that you have a pair of Beolab 28 loudspeakers, perfectly positioned in your listening room, with a single chair in the correct location, as is shown in Figure 1.1. You sit in this chair to listen to a recording: experiencing the accurate and precise placement of the instruments within the space in front of you (better known as “stereo imaging”).




Figure 1.1: The sound stage presented by Beolab 28s for active listening when sitting in the sweet spot with properly placed loudspeakers.

Now, imagine that you have the same loudspeakers in the same positions in the same room, but you’ve moved to the sofa as shown in Figure 1.2 (or perhaps you’re still in the same sofa as before as shown in Figure 1.3) and you prefer to hear music in the background while you read a book. In this case, the imaging of the recording is not important: you just want a cloud of sound that does not distract you while you read. Using your controller you simply switch the behaviour of the Beolab 28s to deliver this experience instead.




Figure 1.2: The sound stage presented by Beolab 28s for passive listening when not sitting in the sweet spot.




Figure 1.3: The sound stage presented by Beolab 28s for passive listening when sitting in the sweet spot.

These two scenarios illustrate the primary use cases that the Beolab 28 can deliver. We’ll call the first one **active listening**, since your primary activity is to listen to the recording. We’ll call the second one **passive listening** since, in this case, listening to music is secondary to another activity (in our example, reading) in a stationary listening position or area. In order to be able to do this, there are different adjustable parameters in the Beolab 28. You can choose to alter each of these parameters individually according to your preferences and listening situations, and then save the settings to a Listening Mode for future use.

1.0.1 Beam Width Control

Beolab 28 gives you the power to make these changes using a number of “handles”; controllers that let you change the acoustical behaviour of the loudspeaker. Among these features,

there is one that stands out: Beam Width Control.

When a recording engineer makes a recording in a well-designed studio, he or she is sitting not only in a carefully-designed acoustical space, but a very special area within that space. In many recording studios, there is an area behind the mixing console where there are no reflections from the sidewalls arriving immediately after the direct sound from the loudspeakers. This is implemented either by putting acoustically absorptive materials on the walls to “soak up” the sound so it cannot reflect (as shown in Figure 1.4), or by angling the walls so that the reflections are directed away from the listening position (as shown in Figure 1.5).




Figure 1.4: One way to reduce the problem of side wall reflections is to absorb them at the walls so that there is no reflection. This is a solution often used in recording studios, however, it also results in an unnatural-sounding “dead” room.




Figure 1.5: An alternative method to reduce the problem of side wall reflections is to re-direct them away from the listening position, producing a “reflection-free zone”. This method is often used in recording studios that are initially designed with the help of an experienced acoustical consultant.

This is different from your living room which has not been designed primarily as an acoustical space. It probably has sidewalls that reflect the energy from your loudspeakers and send that sound to you at the listening position – a situation that is more like that which is shown in Figure 1.6.




Figure 1.6: The direct sound (in red) from the loudspeakers is influenced by the reflections off the side walls (in grey).

In order to get the same acoustical behaviour in your living room that the recording engineer had, we have to reduce the amount of energy that is reflected off the side walls. If we do not want to change the room, one way to do this is to change the behaviour of the loudspeaker by focusing the beam of sound so that it stays directed at the listening position, but it sends less sound to the sides, towards the walls.

This is one of the options that Beolab 28 gives you: to make the beam of sound directed out the front of the loudspeaker narrower to reduce the level of sidewall reflections, so that you

get a more accurate representation of the sound the recording engineer heard when the recording was made.




Figure 1.7: Beolab 28 solves the problem of side wall reflections by reducing the amount of acoustic energy that is radiated towards the side walls: so there is less energy to reflect.

However, if you’re sharing your music with friends or family, depending on where people are sitting, the beam may be *too* narrow to ensure that everyone has similar experiences. In this case, it may be desirable to make the Beolab 28s’ sound beams wider.

This option to change the pattern of the radiation of sound from the Beolab 28 is called *Beam Width Control*.

1.0.2 ARC: Active Room Compensation

In 2002, Bang & Olufsen introduced the Beolab 5 which included *Adaptive Bass Calibration* or *ABC*. This was a system that used a microphone to measure the effects of the listening room’s acoustical behaviour on the sound of the loudspeaker, and then created a filter that compensated for those effects in the low frequency band. As a very simple example, if your room tended to increase the overall apparent bass level, then the Beolab 5’s would reduce their bass level by the same amount.

For the nearly-two decades that have passed since then, Bang & Olufsen has been developing and refining the room compensation systems for our loudspeakers. Beolab 28 includes our latest generation of this algorithm, now called *Active Room Compensation* or *ARC*.

Control Parameters

2.1 Beam Width Control

When two Beolab 28s are stereo paired, their horizontal width of the beam of sound that is radiated from them can be adjusted by selecting from two options:

- **Narrow**
- **Wide**




Figure 2.1: In order to access the Beam Width control, click the Listening Mode Edit button (the circle with the dot).

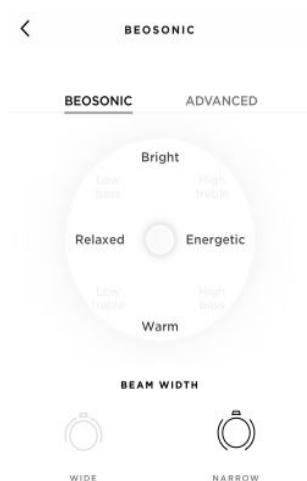


Figure 2.2: The Beosonic and Beam Width control page.

2.1.1 Narrow

Sit in the “sweet spot”: a location in your listening room that is exactly the same distance from each of your

Beolab 28s, and where the two loudspeakers are facing. Using your Beolab 28 interface, set the Beam Width Control to “Narrow”.

While facing a point located at the centre between the two loudspeakers, play “Tom’s Diner”, recorded by Suzanne Vega for her 1987 album “Solitude Standing”. Vega’s voice should appear to float at a position between the two loudspeakers. If her voice does not appear to be located exactly mid-way between the two loudspeakers, it is likely that you are sitting slightly closer to one loudspeaker than the other; in other words, to one side of the sweet spot. Try moving slightly side-to-side and pay attention to the lateral movement of Vega’s voice in space.

This ability for a pair of loudspeakers to deliver the illusion of a sound coming from a location in space between them is called *phantom imaging*, *stereo imaging*, or simply *imaging*.

Change the track to Jennifer Warnes’s recording of “Bird on a Wire” from the Album “Famous Blue Raincoat: The Songs of Leonard Cohen”. In this recording, there are many more instruments and voices, however, it should be very easy to locate the position of each of those sources as coming from somewhere between the two loudspeakers.

Beolab 28 is able to deliver such a precise stereo imaging for active listening because it is able to reduce the amount of mid-frequency band energy in the reflections off the side walls of your listening room. This gives the same result at the listening position as if you used acoustically absorptive materials on your walls, or changed the geometry of your listening room to avoid having early lateral reflections in the listening position.

A potential side effect is the sensitivity of the system to an incorrect listening position. You may notice that, in narrow mode, it is more critical that

you are seated at exactly the correct listening position in order to achieve both precise and accurate stereo imaging. Small deviations in listening position may result in noticeable detriments in the spatial representation of your recordings.

2.1.2 Wide

As mentioned above, when the Beolab 28s are set to a narrow beam width, they are somewhat unforgiving of a mis-placement of the listening position. This is particularly noticeable when you are listening to music with friends and family.

Consequently, in more social or passive listening situations, it is likely preferable that the Beolab 28s have a wider beam width. Although this will likely result in more energy in the sidewall reflections, it also ensures that there is a more equal distribution of the direct sound across a wider listening area in the room.

The side effects of the Wide beam width are dependent on the strength of the sidewall reflections, however, in many situations, a number of different effects may be audible.

The first is that the apparent distance to the various sources in the stereo mix will collapse slightly, resulting in the perception that the sources in the recording are roughly the same distance from the listening position as the loudspeakers themselves. This means that (relative to the narrow mode) very close sources will move further away and very far sources will move closer to the listening position.

Secondly, the apparent width of the sources will become slightly larger with less precise left-right locations. You will not have precise locations as in narrow mode: the imaging becomes slightly more “cloudy” or “fuzzy”. This is due to the extra energy reflected off the side walls.

Thirdly, the overall timbre or tone colour of the sound may change as a result of increase influence of the sidewall reflections at the listening position.

2.1.3 Comment

The above examples connecting Beam Widths to listener position are merely that – examples. It should also be said that changing the Beam Width of the Beolab 28 has non-intuitive consequences on the perceived sound of the loudspeakers. Consequently, it may be that you prefer the overall sound of a particular Beam Width, even if you are not sitting “in the beam”.

2.2 Stereo Pairing

Two Beolab 28s can be connected as a pair of loudspeakers for 2-channel reproduction using a network link. This connection not only contains the audio information, but also “metadata” about the loudspeakers’ configurations to ensure that their parameters (such as Beam Width) are matched.

In its factory-default configuration, the Stereo Pairing audio signal is encoded using the LC3plus high resolution mode software licensed by Fraunhofer IIS¹. This is a CODEC running at 48 kHz, 24-bits and a bitrate of 1 Mbps. This is to ensure that a high quality signal can be delivered with a minimum of service interruptions caused by network traffic. In installations with adequate bandwidth in the network connections, the signal can be changed to uncompressed LPCM running at 48 kHz, 24 bit. In both settings, the input-to-output latency of the audio signal is approximately 300 ms to provide adequate buffer time to recover and correct audio dropouts that may occur due to network traffic.

It should be noted that Stereo Pairing does not work for Power Link or Wireless Power Link inputs.

2.3 Speaker Role

When configured as a Stereo Pair of loudspeakers using a network connection, two Beolab 28s function as a synchronised pair of loudspeakers: one “Primary” loudspeaker in which the connection panel for the input signals is enabled and one “Secondary” loudspeaker, in which it is disabled.

Since both the left and right audio channels are input to your Primary loudspeaker, there is no physical way of knowing which loudspeaker is on the left and which is on the right. As a result, the interface allows you to swap the Speaker Role, to ensure that the correct audio channel is reproduced by the correct loudspeaker.

If your only source is a Beovision television, then the Speaker Role should be entered in the television’s Speaker Group menu instead of in the Beolab 28’s Listening Mode.

Note that a Primary / Secondary pair of Beolab 28s cannot share the same Speaker Role. If you wish to send the same audio signal out of both loudspeakers, this will have to be arranged using the source device (e.g. a Beovision television).

2.4 Active Room Compensation

For a general introduction to the effects of room acoustics on the sound of a loudspeaker, please read [Appendix 3: The Influence of Room Acoustics on Loudspeakers](#)

It should be noted that the acoustical behaviour of a room can change considerably when windows or doors are opened and closed. Consequently, for optimal tuning, it is recommended that the ARC measurements are made in the most typical room configuration.

2.5 Volume

The volume of the Beolab 28 is controllable from 0 to 100 in steps of 1 dB. Note that Volume Step 0 is a full mute.

In its default settings, Beolab 28 has been calibrated to match the level of other Bang & Olufsen loudspeakers for its Power Link and Wireless Power Link inputs. The Volume control of the Beolab 28 is disabled² for Power Link and Wireless Power Link sources. This restriction is made to prevent incorrect calibration of levels in surround sound configurations.

2.6 Tone Controls

The Tone Controls on the Beolab 28 consist of traditional Treble and Bass controls. These are global adjustments that are applied to all Listening Modes and to both loudspeakers simultaneously.

2.6.1 Treble

The Treble adjustment allows you to change the relative amount of high-frequency sound globally using a high shelving filter with a fixed turnover frequency of 8 kHz and a Q of 0.707. The gain at the turnover frequency is one half the maximum gain applied by the filter in decibels. For example, when the gain of the controller is -4 dB, the gain at 8 kHz is -2 dB.

The Treble control is applied to a global filter and therefore is applied to all Listening Modes. It is also independent of the settings of other equalisation controllers in the system such as the Beosonic control. The range of the controller is from -6.0 dB to +6.0 dB in steps of 0.5 dB.

¹<http://iis.fraunhofer.de/audio>

²In fact, it is fixed at volume step 90 to match legacy Bang & Olufsen loudspeakers.




Figure 2.3: Magnitude Responses, Treble controller. Note that this filter is applied to both loudspeakers simultaneously.

2.6.2 Bass

The Bass adjustment allows you to change the relative amount of low-frequency sound globally using a low shelving filter with a fixed turnover frequency of 120 Hz and a Q of 0.707. The gain at the turnover frequency is one half the maximum change in gain applied by the filter in decibels. For example, when the gain of the controller is +6 dB, the gain at 120 Hz is +3 dB.

The Bass control is a global filter and therefore is applied to all Listening Modes. It is also independent of the settings of other equalisation

controllers in the system such as the Beosonic control. The range of the controller is from -6.0 dB to +6.0 dB in steps of 0.5 dB.




Figure 2.4: Magnitude Responses: Bass controller. Note that this filter is applied to both loudspeakers simultaneously.

System

3.1 Max Volume

The *Max Volume* control allows you to determine the limit of the volume control.

Range 0 - 100
Resolution 1 dB

The *Max Volume* parameter is not available for the Power Link and Wireless Power Link inputs.

3.2 Startup Volume

The *Startup Volume* control allows you to determine the volume level when

the Beolab 28 wakes as a result of a detected signal or a change in source.

Range 0 - 60
Resolution 1 dB

The *Startup Volume* parameter is not available for the Power Link and Wireless Power Link inputs.

Setup Recommendations

4.1 Bang & Olufsen television as source

A current Bang & Olufsen television can automatically switch Beolab 28 Listening Modes as part of the Speaker Group function. However, it should be noted that a given Speaker Group in the television sends only one Speaker Listening Mode value on its Power Link

outputs to all loudspeakers connected to the television. This means that the Listening Mode identification numbers in all Beolab 28, Beolab 50, and Beolab 90 pairs must match for a given configuration corresponding to a Speaker Group in the television.

4.2 ARC measurements

It is important to ensure that there is as little extraneous noise as possible

during the ARC measurement procedure. This includes turning off air conditioning systems or performing the measurements during low-traffic hours, where possible. Short, impact noises such as a single footstep are not problematic. However, continuous sounds such as an idling truck or a humming air conditioning motor will affect the measurement.

Features and Setup

5.1 Resonance-based Sound Design

A very large part of the sound tuning of the Beolab 28, like many other Bang & Olufsen loudspeakers, is based on acoustical measurements performed at many locations around, above and below the loudspeaker.

One of the important aspects of these measurements is to find the behaviour of the loudspeaker in time. For example, if a sound is sent to the loudspeaker, and then stopped suddenly, does the loudspeaker also stop, or does it “ring” at some frequencies (in exactly the same way that a bell rings when struck)? Ringing in the time response of a loudspeaker is an indication that it has a resonance: a frequency at which it “wants” to

move. This resonance has a detrimental effect on the overall sound of the loudspeaker, since it “smears” sounds in time.

For example, if you have a loudspeaker that has a natural resonance at 110 Hz (two octaves below a “Concert A”, to musicians) then it will naturally ring at that note when it is “hit” with an impulsive signal such as a kick drum. If the song that the band is playing is not in A (major or minor), but in B-Flat instead, then there will be a dissonance between the notes played most often in the song, and the note that is “singing along” with the kick drum. This can contribute to the loudspeaker sounding “muddy” or “boxy” (to use only two words...).

This is why the measurement-based portion of the filtering of all current Bang & Olufsen active loudspeakers is primarily designed to counteract the

natural resonances in the system. So, for example, if the woofer in the Beolab 28 has a natural resonance at 110 Hz, then that resonance is mirrored with an equal, but opposite phase behaviour in the Digital Signal Processing engine. The total result of the filter in the DSP and the behaviour of the woofer is that there is no unwanted ringing in the entire system. This, in turn, means that the loudspeaker’s response is controlled in the time domain and therefore also in the frequency domain as well.

This is only possible with an extensive set of measurements of each loudspeaker driver’s mechanical and acoustical behaviour and a custom-created set of filters for it.

5.2 Adaptive Bass Linearisation (ABL) and Thermal Protection

Almost all loudspeakers in the Bang & Olufsen portfolio (including Beolab 28) feature Adaptive Bass Linearisation or “ABL”. This is an algorithm that was first patented by B&O in 1991 and is custom-tuned for each of our products. Its purpose is to ensure that, when the physical limits of a component of the loudspeaker are reached (for example, a woofer is approaching its maximum excursion or a power amplifier is close to clipping) the loudspeaker prevents that limit from being reached and the transition to that limit is “softened” (depending on the component in question).

In addition, Beolab 28’s processing continually monitors the individual temperatures of many internal components including:

- Individual loudspeaker driver magnets
- Power Amplifier modules
- DSP circuit boards
- Power Supply circuit boards

Using this information, combined with the power that the amplifiers deliver to the loudspeaker drivers, the temperatures of many more components within Beolab 28 are calculated using customised thermal models of the loudspeaker’s components and construction.

If the temperature of a component

inside the loudspeaker approaches its “thermal limit” (the temperature at which it stops working due to overheating) the signal processing of the Beolab 28 adjusts the signals to protect the component. The exact type of adjustment depends on the particular component that is approaching its limits. As a simple example, if a tweeter voice coil is calculated to be approaching its limit, then its gain is reduced to attempt to protect it from destruction.

It is important to state that this does *not* mean that the Beolab 28 is indestructible, but it does make it more difficult to destroy.

More information can be found in [Appendix 6: ABL - Adaptive Bass Linearisation](#).

Appendix 1: Recommendations for Critical Listening

6.1 Loudspeaker Configuration

The Beolab 28 in a Bang & Olufsen system provides you with a range of parameters that can be used to adjust the timbral and spatial presentation of your recordings for various listening rooms, loudspeaker placements and listening positions. However, it is always best to start with an optimal configuration in your listening room.

First consider the relationship between the loudspeakers and the listening position itself. The two loudspeakers and the listening position should be the three corners of an equilateral triangle. This means that the distance from each loudspeaker to the listening position should be the same as the distance between each loudspeaker. This also means that the loudspeakers will be 30° away from the front, centre location, directly in front of the listening position.

Secondly, the two loudspeakers should be “toed-in” by 30° . This means that they should be slightly rotated so that they are both facing the listening position. This is shown in Figure 6.1.




Figure 6.1: A “perfect” loudspeaker configuration with Beolab 28s. Both loudspeakers are aimed at the listening position. The distance from the listening position to each loudspeaker is the same as the distance between the two loudspeakers.

The next consideration is symmetry within your listening room. It is

commonly recognised that the best stereo imaging will be achieved if the listening configuration (the triangle formed with the listener and the two loudspeakers) is placed in left-right centre of the room. Therefore, the side walls will both be the same distance from the listening position, and the loudspeakers will have the same distance to its adjacent walls. This is to say that the distance from the left loudspeaker to the left wall is the same as the distance between the right loudspeaker and the right wall. The distance to the front wall (behind the loudspeakers) should be the same for both loudspeakers, but certainly does not have to be the same as the distance to the side walls.




Figure 6.2: An optimal placement for the loudspeakers with respect to adjacent walls. Note that the distance between each loudspeaker and its closest side wall are identical, and that the distances from the loudspeakers to the front wall are also matched.




Figure 6.3: A less-optimal placement for the loudspeakers with respect to adjacent walls. Note that the distances from the loudspeakers to the front wall are matched, however, the distance between each loudspeaker and its closest side wall are not identical.




Figure 6.4: Another example of a less-optimal placement for the loudspeakers with respect to adjacent walls. The distances from the loudspeakers to the front wall are matched, however, the right loudspeaker lacks a side wall nearby.




Figure 6.5: An example of a worst-case placement of loudspeakers with respect to the listening room. No two distances between a loudspeaker and an adjacent wall match each other.

It should be noted that the primary casualty of poor loudspeaker placement in a listening room will be the spatial representation of your recordings. The precision and accuracy of the stereo imaging, as well as the sensation of envelopment from the recording will be adversely affected by early reflection patterns that are not matched for the two loudspeakers. This problem is minimised by using Beolab 28's narrow beam width, however, even this mode can benefit from correct loudspeaker placement in the room.

6.2 Room Acoustics

The Beolab 28 has two features that can overcome some detrimental effects of your room's acoustical behaviour (**Beam Width Control** and **Active Room Compensation**). However, the sound of any loudspeaker can be optimised by improving the room's acoustics.

One of the main acoustical problems in rooms is that of room modes or resonances. These occur because the room behaves very much like an organ pipe, naturally "singing" at specific frequencies that are determined by the dimensions of the room. Without correct acoustical treatment, these resonances are almost unavoidable. It is preferred to ensure that the resonances in the room's three dimensions (length, width, and height) do not overlap each other. This means that the better listening rooms have complex relationships between these three dimensions. For example, a "worst case" for a listening room would be a cube, where all three dimensions are identical, thus all resonances have the same frequencies. A next-worst case is one where a dimension is a multiple of another, for example, a room that is 9m x 6m x 3m. In a best case, the ratios of the room's dimensions would have non-simple values (e.g. 1 : 2.16 : 2.96 – so, as an example, 3m x 6.48m x 8.88m).¹

A second issue in many rooms is that of hard, reflective surfaces; particularly in locations where the sound from the loudspeaker is directly reflected to the listening position. There are two ways to alleviate this problem: absorption and diffusion. In order to absorb a sound wave so that it does not reflect off a surface, an absorptive material such as fibreglass insulation or acoustical foam must be placed on the surface, or in the path taken by the reflection. A reflection can be diffused by making the reflective surface irregular. For example, placing a bookcase at the point of reflection will help as a diffusor if the books are

arranged in random heights and depths.

Finally, it is wise to absorb the sound waves that would be reflected off the floor (e.g. with carpet or a rug) and ceiling (using absorptive ceiling tiles). This will also help to reduce the overall reverberation time of the room.

6.3 Source Devices

When connecting an audio source to the Beolab 28, there are some basic, general rules that should be followed in order to get the optimal performance from your system. Note that these are general rules, so there are exceptions.

- If possible, the source should be connected to the Beolab 28 using a digital audio connection.
- If the source device has a volume control it should be disabled and the Beolab 28's volume control should be used instead
- If the source has two analogue outputs: one volume-regulated and the other at a fixed level, the fixed-level output should be used
- If you are connecting a source using a line-level analogue input, check the source device's datasheet to find its maximum output level and set the value appropriately on the Beolab 28. If the maximum output of your device is greater than the Beolab 28's maximum possible input level of 6.5 V RMS, then it is recommendable that the source device's output level be reduced if possible, either within its own settings or using an external attenuator. Figure 6.6 shows the necessary attenuation to reduce various voltage levels to 6.5 V RMS.




Figure 6.6: The minimum attenuation required to apply to a source with a given Maximum Output level to reduce it to 6.5 V rms in order to prevent clipping the Beolab 28 analogue inputs.

6.3.1 Optical cables

It is recommended that high-quality optical cables are used for the Beolab 28, particularly for longer cable runs. This is due to the fact that there is attenuation (dimming or loss of light intensity) of the optical signal on the plastic or glass fibre in the cable. This attenuation is proportional (in dB) to the length of the cable. Therefore, in order to ensure that the optical receiver on the Beolab 28 has an adequate signal at its input, the light attenuation on the cable should be minimised either by using short cables and/or high-quality optical fibre.

¹See "Room dimensions for small listening rooms" by Dr. Trevor Cox for a good introduction to this topic.

Appendix 3: The Influence of 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**

7.1 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 7.1: 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.




Figure 7.2: 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.

For example, take a look at Figure 7.2. You can see that, at the very low end, the reflection boosts the level of the loudspeaker by 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 scaled logarithmically (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 7.3.




Figure 7.3: 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 7.4.




Figure 7.4: 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 rooms 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 7.5, 7.6 and 7.7 which have identical distances as the simulations in Figures 7.2, 7.3, and 7.4 for comparison.




Figure 7.5: 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 7.6: 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 7.7: 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 7.7 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.

7.2 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 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 7.8: The concept of the effect of a room mode and Active Room Compensation. See the associated text for an explanation.

Figure 7.8 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

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

²Do a search for "tritone" or "diabolus in musica".

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 28

7.3 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.

7.4 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 28) 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 28's Active Room Compensation which will correct the effects of boundaries (walls) and room modes on the timbre of the loudspeaker. Using measurements made with the built-in microphone, the ARC algorithm 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 28s. 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 28: 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 28 can deliver a more accurate and precise representation of the spatial presentation of your stereo recordings.

7.5 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.

Appendix 4: Loudspeaker Directivity and Distance Perception in Stereo Imaging

8.1 Distance Perception in Real Life

Go to the middle of a snow-covered frozen lake with a loudspeaker and a friend. Sit there and close your eyes and get your friend to place the loudspeaker some distance from you. Keep your eyes closed, play some sounds out of the loudspeaker and try to estimate how far away it is. You will be wrong (unless you're VERY lucky). Why? It's because, in real life with real sources in real spaces, distance information (the information that tells you how far away a sound source is) comes mainly from the relationship between the direct sound and the early reflections from walls in your listening room. If you don't have any early reflections, then you don't have any distance information. Add the early reflections and you can very easily tell how far away it is.¹

8.2 Distance Perception in a Stereo Recording

Recording engineers have a basic trick for controlling the apparent distance to a sound source in a stereo recording using the so-called "dry-to-wet" ratio; in other words, the relative levels of the direct sound and the reverberation. To be honest, this is a bit of an over-simplification, but it's at the level of knowledge one would typically have if one were just starting out recording a budding rock band in a garage.

Many classical recordings are made with a pair of microphones. An instrument that is on the left side of the pair will produce a sound that is slightly louder and/or slightly earlier in the left microphone than in the right microphone. This means that, when you sit in the sweet spot and listen to the stereo recording, you will hear that source on the left side of the stereo image. This effect is true not only for

the direct sound of the instruments arriving at the microphone pair, but also for the acoustic reflections off the various surfaces in the recording space. So, if the recording engineer has been paying attention, the distance information (the relationship between the direct sound and the reflections) has been captured in the recording. This means that when you listen to the recording, you not only can tell where the instruments are from left to right, but also their relative distances.

8.3 Combining the Two

So, we know that early reflections tell your brain how far away the sound source is. Now think to a loudspeaker in a listening room:

Case 1: If you have a listening room that has no sidewalls, then there are no early reflections, and, regardless of how far away the loudspeakers are, a sound source in the recording without early reflections (e.g. a close-mic'ed vocal) will sound closer to you than the loudspeakers.

Case 2: If you have a listening room with early reflections, and the loudspeakers are less directional (such as Beolab 28s with their Beam Width set to Wide or Omni) then the early reflections from the side walls tell you how far away the loudspeakers are. Therefore, the close-mic'ed vocal track from Case 1 cannot sound any closer than the loudspeakers: your brain is too smart to be told otherwise.

Case 3: If you have a listening room with sidewalls and therefore early reflections, but the loudspeakers are directional such that there is no energy being delivered to the side walls, then the result is the same as in Case 1. This time, the level of the early reflections has been reduced because of loudspeaker directivity instead of

wall absorption, but the effect at the listening position is the same. This is the case with Beolab 28 when its Beam Width is set to Narrow.

The conclusion is that, in order to get an accurate and precise representation of the spatial properties in a stereo recording, you should try to minimise the levels of the early reflections from the sidewalls in your listening room. However, this means that you are optimising the sound for the sweet spot, on-axis to both loudspeakers. When listening with friends, it may be necessary to widen the loudspeakers' Beam Widths.

¹This has been proven in various listening tests. For example, go check out "Psychoacoustic Evaluation of Synthetic Impulse Responses" by Per Rubak & Lars G. Johansen as a starting point.

Appendix 6: ABL - Adaptive Bass Linearisation

9.1 A General Introduction to ABL

Almost all loudspeakers made by Bang & Olufsen include Adaptive Bass Linearisation or ABL. This includes not only our “stand alone” loudspeakers (the Beolab series) but also our smaller “Beosound” loudspeakers and televisions. The only exceptions in the current portfolio are our passive loudspeakers and some headphone/earphone models.

There is no one technical definition for ABL, since it is in continual evolution. In fact it may change from product to product as we learn more and as different products require different algorithms. Speaking very broadly, however, we could say that it reduces the low frequency content sent to a loudspeaker driver (e.g. the woofer) when the loudspeaker is asked to play loudly, but even this is partially inaccurate.

It is important to note that it is *not* the case that this replaces a “loudness function” which may (or may not) be equalising for Equal Loudness Contours (sometimes called “Fletcher-Munson Curves”). However, since (generally) the bass is pulled back when things get loud, it is easy to assume this to be true.

When we are doing the sound design for a loudspeaker (which is based both on measurements and listening), we ensure that we are operating at a listening level that is well within the “linear” behaviour of the loudspeaker and its components. (Typically, the sound design is done at a standard playback level where a -20 dB FS full-band pink noise produces 70 dB SPL (C) at the listening position.)

This means that

- the drivers (usually the woofers) aren't being asked to move too far (in and out)
- the amplifier is operating within

its limits

- the power supply is operating within its limits, and
- nothing (neither the power supply, nor the amplifiers, nor the voice coils) is getting so hot that the loudspeaker's behaviour is altered.

This is what is meant by “linear”: it's fancy word for “predictable”. In addition, it should be stated that if we were listening to loudspeakers at high levels daily, we would get increasingly bad at our jobs due to hearing loss.

So, we do the tuning at a listening level where we know things are behaving: remember that we always do it at the same calibrated level every time for every loudspeaker so that we don't change sound design balance due to shifts associated with equal loudness contours. (If you voice a loudspeaker when it's playing loudly, you'll wind up with a loudspeaker with less bass than if you voiced it quietly. This is because you're automatically compensating for differences in your own hearing at different listening levels.)

After that tuning is done, then we go back to the measurements to see where things will misbehave. For example, in order to compensate for the relatively small cabinet behind the woofer(s) in smaller loudspeakers, we increase the amount of bass that we send to the amplifiers for the woofers as part of the sound design. If we left that bass boost in the tuning when you turn up the volume, the loudspeaker would go up in smoke – or at least sound very bad. This could be because:

- the woofer is being pushed or pulled beyond its limits, or
- because the amplifier clips or
- the power supply cannot supply more current or
- something else.

After the tuning process is complete, we put the loudspeaker in a small torture chamber roughly the size of a clothes closet, put on some dance music (or modified synthetic test signals) and turn up the volume. While that's playing, we're continually monitoring the signal that we're sending to the loudspeaker, the driver excursion, the demands on the electronics (e.g. the amp's, DAC's, power supply, etc.) and the temperature of various components in the loudspeaker, along with a number of other parameters.

Armed with that information, we are able to “know” how those parameters behave with respect to the characteristics of the music that is being played (e.g. how loud it is, in various frequency bands, for how long, in both the short term and the long term). This means that, when you play music on the loudspeaker, it “knows” the following:

- how hot it is at various locations inside,
- the loudspeaker drivers' excursions,
- amplifier demands,
- power supply demands,
- and so on. (The actual list varies according to product: these are just some typical examples...)

So, when some parameter gets close to a maximum (e.g. the amplifier starts to get too hot, or the woofer is nearing maximum allowable excursion) then *something* will be pulled back.

What is pulled back? It depends on the product and the conditions at the time you're playing the music. It could be a band of frequencies in the bass region, it could be the level of the woofer. In a worst-case-last-ditch situation, the loudspeaker might even be required to shut itself down to protect itself from you (or the guests attending your party). Of course, there is no

guarantee that you cannot destroy the loudspeaker somehow, but we do our best to build in enough protection to cover as many conditions as we can.

How is it pulled back (i.e. how quickly and by how much)? That also depends on the product and some decisions we made during the sound design process, as well as what kind of state-of-emergency your loudspeaker is in (some people are very mean to

loudspeakers...).

Note that all this is done based on the signals that the loudspeaker is being asked to produce. So it doesn't know whether you've turned up the bass or the volume. It just knows you're asking it to play this signal right now and what the implications of that demand are on the current conditions (voice coil temperature, for example) This is similar to the fact that the seat belts in

my car don't know why the car is stopping quickly: maybe it's because I hit the brakes, maybe it's because I hit a concrete wall. The seat belts just lock up when they're asked to move too quickly. Your woofer's voice coil doesn't know the difference between Eminem and Stravinsky with a bass boost. It just knows it's hot and it doesn't want to get hotter.