WHITE PAPER

Dirac Live

A technical overview

Dirac 2024

DIRAC

Contents

Chapter 1
Introduction to Dirac Live

Dirac Live is a high-performance digital room correction system. Powered by patented algorithms, it supports sound systems ranging from a single loudspeaker up to complex multichannel systems with many loudspeakers and subwoofers.

What is “room correction”? The accuracy of this term is sometimes questioned, but we use it because it is well established for a family of technologies, products, and strategies to which Dirac Live belongs.

A listening room provides an acoustical environment for your loudspeakers that alters the sound that arrives at your ears. Some of these changes are detrimental to the perceived sound. A digital room correction system works to undo some of the negative changes by digitally processing the audio signal before it reaches the loudspeakers and the room. This is analogous to wearing eyeglasses: the lenses pre-distort the light passing through them before it enters the wearer’s eyes, thereby counteracting the distortion that occurs in the eyes.

There is, however, not just the room to consider, but also the loudspeakers themselves. Very few loudspeakers are perfectly accurate, and different room correction systems differ in their ability to address imperfections in the loudspeakers. With Dirac Live, the term “room correction” can be thought of as meaning “loudspeaker and room correction.” This is a complex acoustical problem. For satisfactory results, it must be considered in time, frequency, and three-dimensional space. Dirac Live consists of three functions that address distinct parts of this problem:

Dirac Live is the core function of the Dirac Live product family. It improves the performance of each loudspeaker individually in the time and frequency domains.

Dirac Live Bass Control (BC) is an advanced bass management solution. It uses phase co-optimization technology to seamlessly integrate subwoofers with loudspeakers and – if multiple subwoofers are used – to minimize seat-to-seat response variation.

Dirac Live Active Room Treatment (ART) uses all loudspeakers and subwoofers in the system to reduce low-frequency room resonances and to further reduce seat-to-seat response variation. (Active Room Treatment is not covered in this version of the white paper.)

Dirac Live is complementary to other methods of improving the sound quality of an audio system, such as loudspeaker placement and passive room treatment. Regardless of the acoustic environment in which it is used, Dirac Live will deliver measurable and audible improvements. If the environment is improved by changes to loudspeaker positioning or the addition of passive room treatment, the Dirac Live application can easily be re-run to generate new correction filters. And because Dirac Live uses acoustic measurements over the whole listening area, its improvements are not limited to a single listener.

This white paper covers the Dirac Live product family's key functions Dirac Live and Dirac Live Bass Control. First, we need to provide some background on loudspeakers and room acoustics.

1.1. How a loudspeaker radiates sound

Imagine for a moment that we have suspended a loudspeaker in mid-air, so that there are no walls, floor, ceiling, or any other nearby objects that can reflect sound.

The loudspeaker radiates sound in all directions, as indicated in Figure 1. Since there are no surfaces that reflect sound, acoustic measurements of the loud-speaker truly would reflect how the loudspeaker produces sound without any influence from a room. They include the effects of the drivers, the crossover, and the shape of the cabinet. These are called the anechoic ("no echoes") responses of the loudspeaker.

A measurement taken directly in front of the loud-speaker is called its on-axis response. The micro-phone must, of course, be a sufficient distance from the loudspeaker that the acoustic signal from all drivers integrates correctly. As indicated in the figure, the response can be thought of as being in time (the impulse response) or in frequency (the frequency response). These are equivalent representations of the loudspeaker's acoustic behavior, each of which reveals various aspects of the loudspeaker's behavior; they will be discussed more in the next chapter.

Measurements at other locations are called off-axis responses. Even though the drivers are on the front of the cabinet, sound wraps around it in different ways depending on frequency.

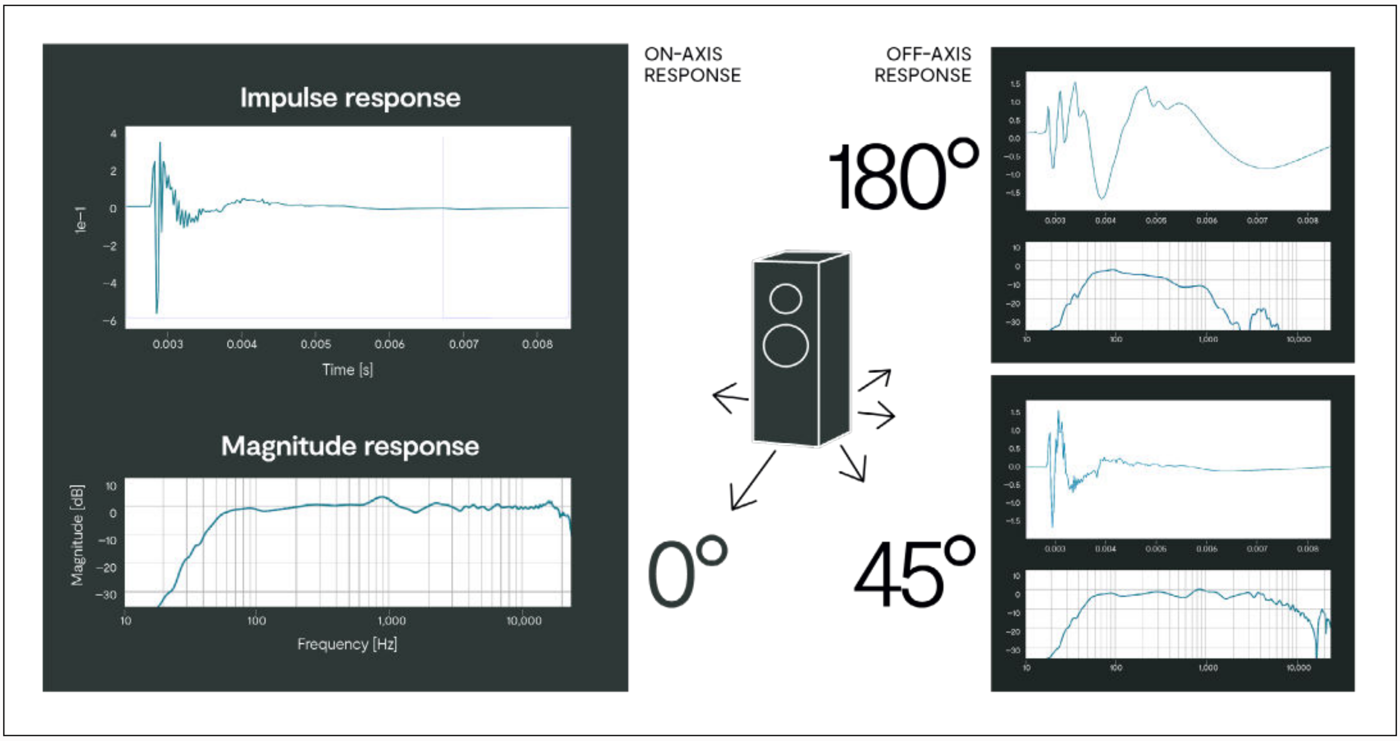

Figure 1. A loudspeaker radiates sound in all directions. Its on-axis response is measured directly in front of the loudspeaker (to the left). Measure-ments at other locations are called off-axis responses (to the right). The off-axis responses are different from the on-axis response in various ways.

Typically, high frequencies are reduced as the measurement microphone is moved further away from the on-axis position, and there may be various bumps and ripples in the response due to the drivers or cabinet effects. Furthermore, interference between adjacent drivers causes notches in the response at certain off-axis angles due to their differing distances from the microphone. The off-axis responses become important when the loudspeaker is placed in a room.

1.2 Reflections and reverberation

Imagine now that we put the loudspeaker into a room and play music through it. Some of the sound from the loudspeaker travels directly to the listener's ears. This direct sound, illustrated in green in Figure 2, will arrive at the listener's ear first. Provided it is a high-quality loudspeaker, the direct sound will be a reasonably accurate representation of the original musical signal.

Sound is also radiated into the room in all directions. This non-direct sound does not just disappear – it is still in the room, and some of its energy will reach the listener eventually, changing what is heard as the "sound" of the system.

Above 200 to 300 Hz (in a typical home listening room), sound is thought of as traveling in rays that reflect off objects. Early reflections from the side wall or ceiling will arrive after the direct sound. Figure 2 illustrates one such reflection in grey. Other reflections may take some time to arrive, as shown in black. These reflections will reflect again off other surfaces, each time being a little weaker, until eventually they die out.

Early reflections arrive too soon after the direct sound for our brains to hear them as separate sounds. They are not like echoes in a canyon. Instead, they tend to contribute to the sense of being in a three-dimensional space and increase the size and solidity of the apparent source of sound. Very early reflections (such as from the edge of the cabinet) and reflections from the same direction as the loudspeaker, however, may be perceived as tonal coloration.

Late reflections give rise to a sense of envelopment. A concert hall is an example of an environment in which a prominent level of late reflections is desirable: the performance "fills the hall." A room for listening to recorded music, however, should have this reverberant field die away more quickly. There is not universal agreement on the desired level of reflection in a home listening room. Some listeners aim for an environment that mimics a studio control room, with substantial amounts of absorbing material to damp reflections as much as possible, so that minute details in the recording can be heard and analyzed.

For recreational listening, however, a managed level of reflection can provide a more enjoyable musical experience. While most audiophiles would find a completely bare room to be overly reflective, normally furnished domestic rooms (carpet, drapes, chairs, shelving) typically have acoustic performance metrics that place them within an acceptable range for listening to recorded music and soundtracks. The choice of whether to invest in additional acoustic treatment is a personal one according to circumstance and preference.

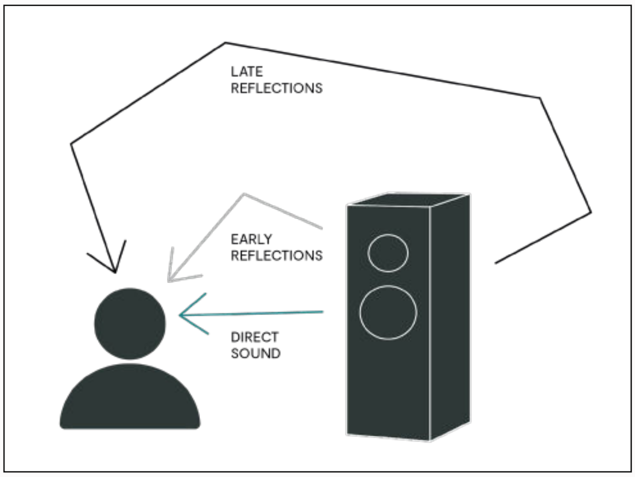

Figure 2. When a loudspeaker is placed in a room, the sound arriving at the listener is not just the direct sound from the loudspeaker, but also reflections from the walls, the floor, the ceiling and other objects.

1.3 Modes

At low frequencies where wavelengths are longer, sound in the room creates standing waves. These are resonances that occur at specific frequencies according to the dimensions of the room. Each frequency at which a standing wave can occur is called a mode. In a typical home listening room, modes are significant up to about 200 to 300 Hz.
Figure 3 illustrates a cross-section of a room. The horizontal axis represents the distance between two walls and the vertical axis indicates the strength of each mode – the sound pressure level (SPL).

A sound source – such as a loudspeaker or subwoofer – excites a mode according to its locs-tion along it. In other words, the mode resonates.

For example, a sound source at the wall excites all modes the maximum amount. At the midway point, the second-order mode (black) is excited the maximum amount, but the fundamental and third-order modes are not excited at all. The phase of the excitation is also determined by the location of the sound source, as suggested by the "+"and "–" symbols. In addition, the SPL of the mode as measured by a microphone (or heard by a listener) varies with its location along the mode.

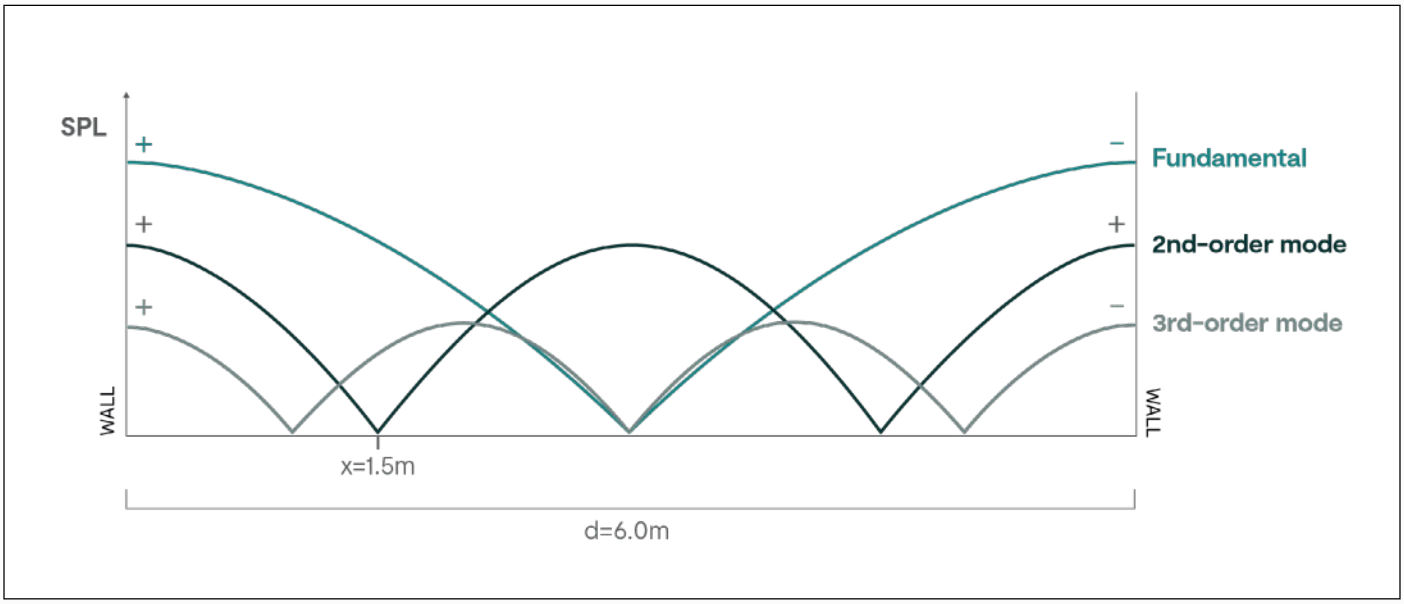

Figure 3. An illustration of the first three room modes along one dimension of a room. Each mode resonates at a specific frequency according to its wavelength. There are many more modes higher in frequency, as well as modes in different dimensions.

The lowest mode frequency is determined by the room's longest dimension and has a wavelength equal to double that dimension. For a room 6 meters long, for example, this is 28.5 Hz. More modes occur at multiples of this frequency (57 Hz, 85.5 Hz, 114 Hz, ...). There are also modes for the room width and height, and modes involving two or three of the room dimensions. And when there are multiple sound sources, each excites the modes differently. Suffice it to say, it is quite complex.
The above discussion assumed the simplest case: a rectangular room. Many rooms are not rectangular and calculating the modes becomes even more complex. Regardless of shape, typical listening rooms have dozens of modes up to 200 to 300 Hz. Strong modes produce peaks and dips in the low frequency room response and can take a long time to decay, resulting in boomy or muddy sound.

1.4 The in-room frequency response

To understand more about how the room affects the sound from a loudspeaker, we can measure the loudspeaker in the room. Figure 4 shows a typical example of a frequency response measurement taken with the microphone at the listening position. This is quite different from the anechoic response of the loudspeaker in free space shown in Figure 1.

What has happened? At low frequencies, we see the effect of the room modes as peaks and dips in the response. Some bass notes may seem to be missing and some may "boom." The response in this frequency range can be dramatically different when measured at distinct locations in the room.

At higher frequencies, we see a wild response in light green due to interference between reflections in the room. It is not as bad as it looks – most of the reflections responsible for these high-frequency details arrive late relative to the direct sound. Furthermore, they arrive from all directions, which makes it easier for our hearing mechanism to separate them from the direct sound; they are perceived as envelopment or reverberation. Our ears do not really hear this wild roller-coaster but something more like the smoothed response shown in dark green.

Recall from Figure 1 that the loudspeaker radiates different frequencies differently in different directions. This, together with the reflection and absorption characteristics of the room, leads to the imbalance in the smoothed response of Figure 4. It may sound a little bright, dull or "hollow." You may not even realize how much the sound of your loudspeakers has been changed by the room until you have corrected it.

Note that it is normal for the in-room response to reduce at high frequencies, as most loudspeakers radiate less sound into the room at high frequencies. In addition, high frequencies are absorbed more easily. A neutral-sounding in-room response typically has a boost in the low bass and a gradual fall-off to the highest frequencies.

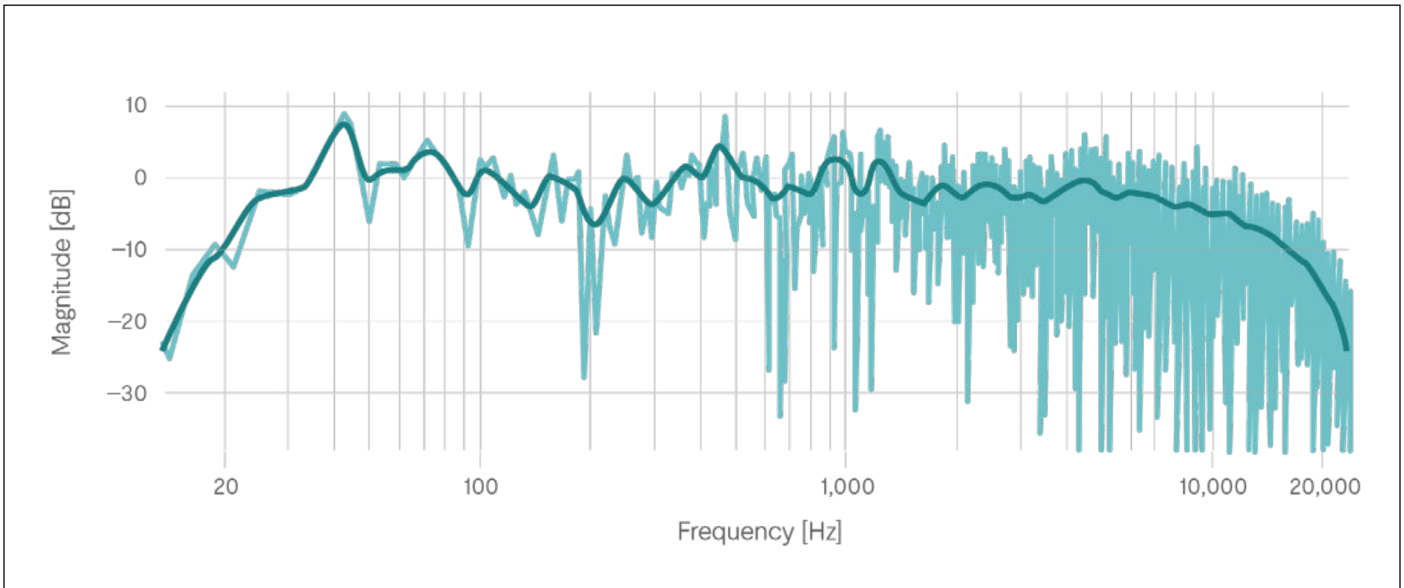

Figure 4. A typical frequency response measured at the listening position. The dark green plot is smoothed while the light teal plot is unsmoothed.

1.5 Measuring the room

A single in-room measurement is not sufficient to perform high-quality correction. Specifically, a correction based on a single measurement is not robust – it has no safeguard against changing the response in ways that may make the sound worse elsewhere.

Dirac Live therefore requires that several measurements be taken within a user-selected listening area. Figure 5 illustrates an example of a small home theater. The grey area indicates the chosen listening area, and the dotted circles represent the measurement locations that define this area. The first measurement taken defines its center. While the listening area is usually rectangular, it does not have to be. For superior results, the vertical height of the measurement microphone must also vary between locations.

To take these measurements, the user places the microphone at different locations, guided by the Dirac Live software application running on a computer or mobile device. The recommended number of measurements will vary with the size of the listening area.

When the measurements are complete, the algorithms running in the Dirac Live application go to work. The first step is to compute the impulse response for each combination of loudspeaker and microphone location. The algorithm then constructs a mathematical model that represents the acoustic behavior of the loudspeaker and room. The exact way this model is used to generate correction filters depends on which functions are being used: Dirac Live, Dirac Live Bass Control, and Dirac Live Active Room Treatment.

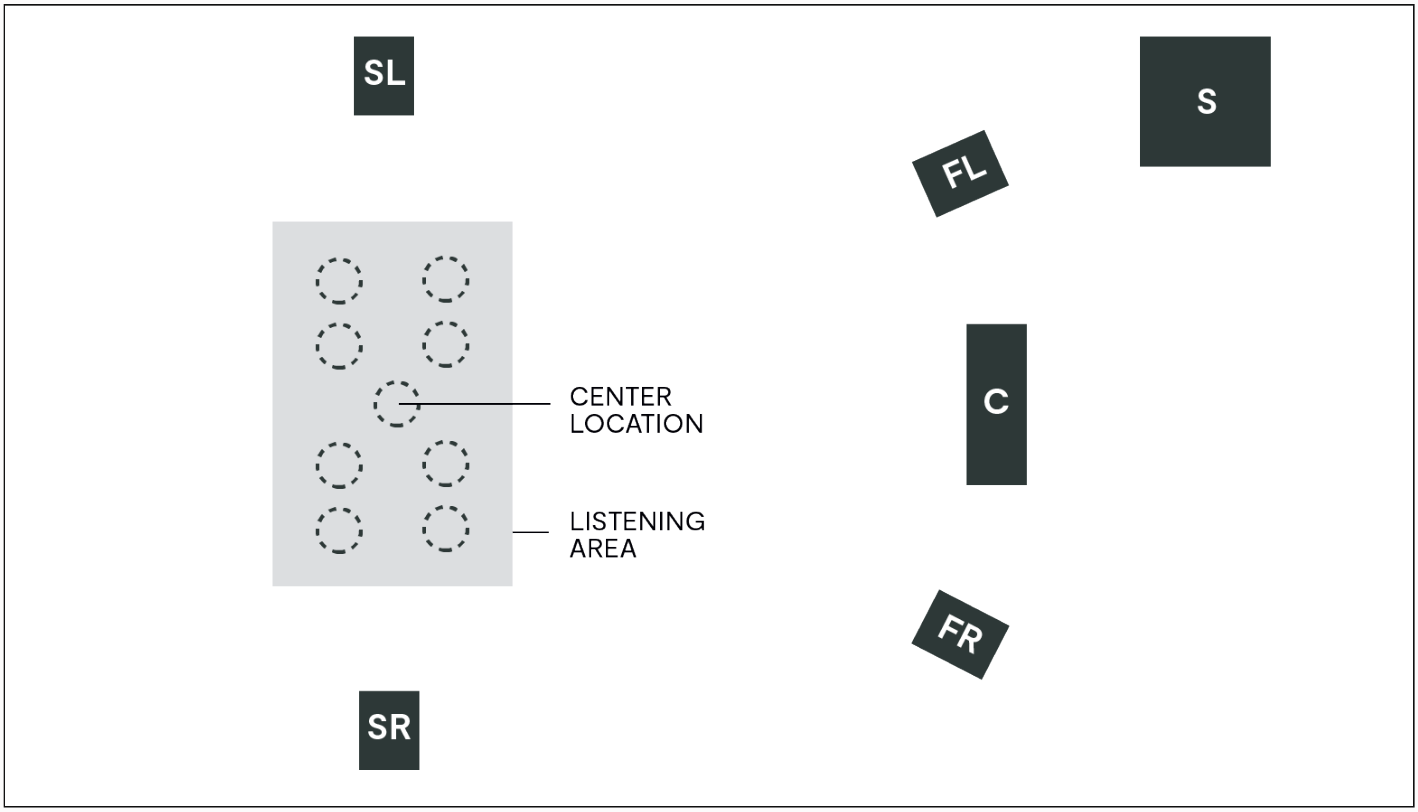

Figure 5. An illustration of measurement locations and the listening area in a small home theater.

Chapter 2
Dirac Live

Dirac Live is the core function of the Dirac Live product family. Dirac Live analyzes the acoustic measurements of the room and corrects each loudspeaker and subwoofer as a single unit.

Figure 6 is a simplified illustration of the measurement data used by Dirac Live to correct the front left loudspeaker, highlighting three measurement locations. At each location, Dirac Live has recorded the acoustic behavior of the loudspeaker and the room. At the most basic level, Dirac Live adjusts the time delay between all loudspeakers and the subwoofer so that the direct sound from each arrives at the center of the listening area at the same time. In addition, the trim on all loudspeakers and the subwoofer is adjusted for the same acoustic level at the listener. However, Dirac Live does much more than this.
In the time domain, it uses the acoustic measurement data to correct the impulse response of the loudspeaker and room, removing phase distortion and providing improved clarity, transient response, and spatial imaging. In the frequency domain, the response is adjusted to smooth out undesirable peaks and provide consistent tonality around the whole stage.

The correction is carefully designed to provide the best subjective experience possible: it applies evenly across the whole listening area and not just at the measurement locations; phase distortion is corrected only if it is found across the whole area, to avoid creating further errors; filter gain is limited to prevent overdriving amplifiers or loudspeakers; and by design, the algorithm does not create digital processing artifacts such as “pre-ringing” that are audible as coloration.

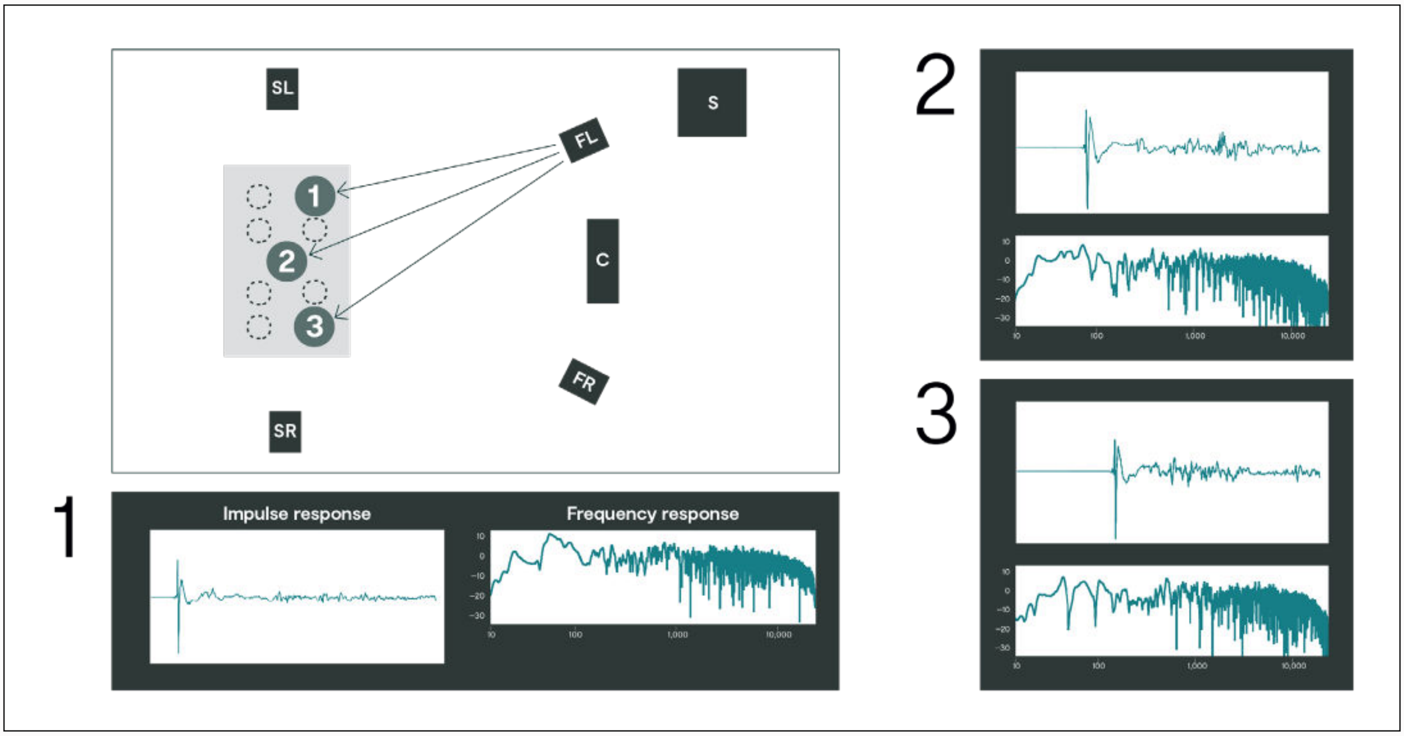

Figure 6. Illustrating some of the in-room measurements that Dirac Live uses to correct the front left speaker.

2.1 Correcting the impulse response

The impulse response is a measure of how a loudspeaker and room respond in time to an input signal. Since this is at the core of how Dirac Live works, this section will first provide some background information and theory, then explain how impulse response correction works in Dirac Live. Suppose we put a measurement microphone in the room and play music. The loudspeaker creates changes in air pressure that are picked up by the microphone's diaphragm, which are in turn converted into an electrical signal for analysis. These are time domain signals – the acoustic signal represented by the plot proceeds in time from left to right. The waveform in Figure 7b, while completely authentic (it is from a measurement in one of our listening rooms), does not tell us anything especially useful – we do not need to know what the music signal looks like while we are playing music. Instead, we want to know the relationship between what we put in and what we get out, i.e., the response of the loudspeaker and room.

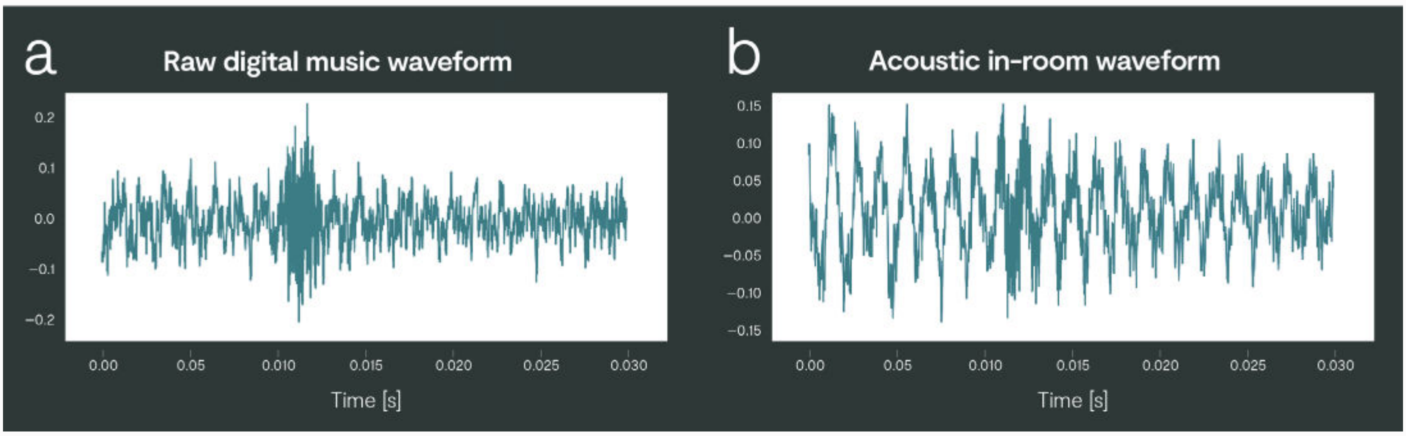

Figure 7. a) A short segment of a music signal. b) The same segment as picked up by a measurement microphone in the room.

To do this, imagine that the loudspeaker is given a theoretical "perfect" input signal. The acoustical signal measured by the microphone will then show us how the loudspeaker and room change that perfect input into something else. In the time domain, this theoretical signal is a very narrow (and very tall) pulse called the Dirac delta function, or just the unit impulse, illustrated in Figure 8a.[1] Below in Figure 8b is a typical example of a measured impulse response. There is a short delay because it takes time for sound to travel from the loudspeaker to the microphone. The main peak of the impulse is "smeared" in time by phase shift in the crossover, present in most loudspeakers. There may be other ripples due to resonances in the drivers and reflections from the cabinet edge. Following the main peak is a series of reflections from the room.

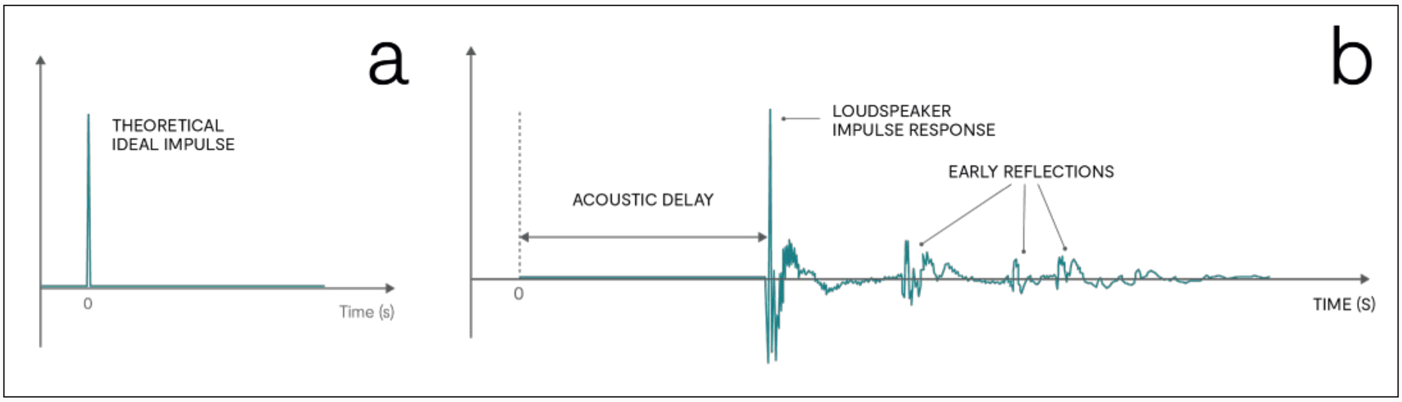

Figure 8. a) The theoretical input signal, a unit impulse. b) A typical measured impulse response of a loudspeaker and room at one location.

We note here that it is not practical to generate a unit impulse and play it through an audio system. Instead, the Dirac Live application generates sweep signals that contain all frequencies up to half the sample rate, and mathematically converts the recorded signal to the impulse response.

Figure 8b is different from Figure 8a. What can – and should – Dirac Live correct? First, we note that the measured impulse response varies between locations in the listening area. It is impractical to achieve perfect correction across all locations; thus, a negligible amount of variance must be accepted at specific locations. The room correction problem is therefore formulated mathematically as a system of equations that aims to minimize the total error between the corrected impulse responses and a set of target impulse responses. The target impulse response at each location is naturally the unit impulse, delayed by the acoustic propagation time from the loudspeaker to that location. In addition, a short processing delay is required for the filters. This is illustrated in Figure 9a.

There are different ways of solving these equations, but only some are acceptable for a high-fidelity audio system. The solution is therefore constrained in several ways. In the time domain, the key constraint is that there must be minimal pre-ringing. This refers to energy in the corrected impulse response that appears before the main impulse peak; it is a time domain artifact that can add an unnatural, artificial character to the sound if not strictly controlled.

A key realization at the core of Dirac Live is that pre-ringing is avoided if we only correct phase distortion that occurs in all measurement locations. A careful analysis is required to allow this principle to be realized in a real acoustic environment. This is explained more in the sidebar Mixed phase room correction.

[1] Technically speaking, Dirac Live is a discrete-time system, so the relevant signal is the unit sample function. For the purposes of this white paper, we felt that a continuous time description would be more approachable.

Figure 9b shows a measured impulse response at one location after correction. Some reflections are still present – they cannot usefully be removed because they do not occur consistently across the whole listening area. In some cases, very early reflections such as off the loudspeaker stand or the edge of the loudspeaker cabinet can be corrected.

It's undeniable that Figure 9b looks more like the Dirac delta function than Figure 8b. Improved transient response and increased clarity and delicacy of sound are attributable to this improvement. In any evaluation of the effect of impulse response correction, it must also be remembered that we are listening to more than one loudspeaker. In the case of stereo reproduction, differences in phase between the two loudspeakers – whether caused by the room or by the loudspeakers themselves – disrupt the illusion of stereo perception, with the result that an instrument or voice becomes more diffuse and harder to locate in the stereo soundstage. Imaging in multichannel reproduction is likewise adversely affected by differences in phase between loudspeakers.

Furthermore, it is not realistic to expect identical impulse responses from two or more loudspeakers even of the same model, due to manufacturing tolerances in the drivers and in the components that make up the crossover. And complete symmetry in home listening rooms is impractical. By applying impulse response correction on all loudspeaker channels – without introducing any negative processing artifacts – Dirac Live improves the imaging and spatial fidelity of any audio system.

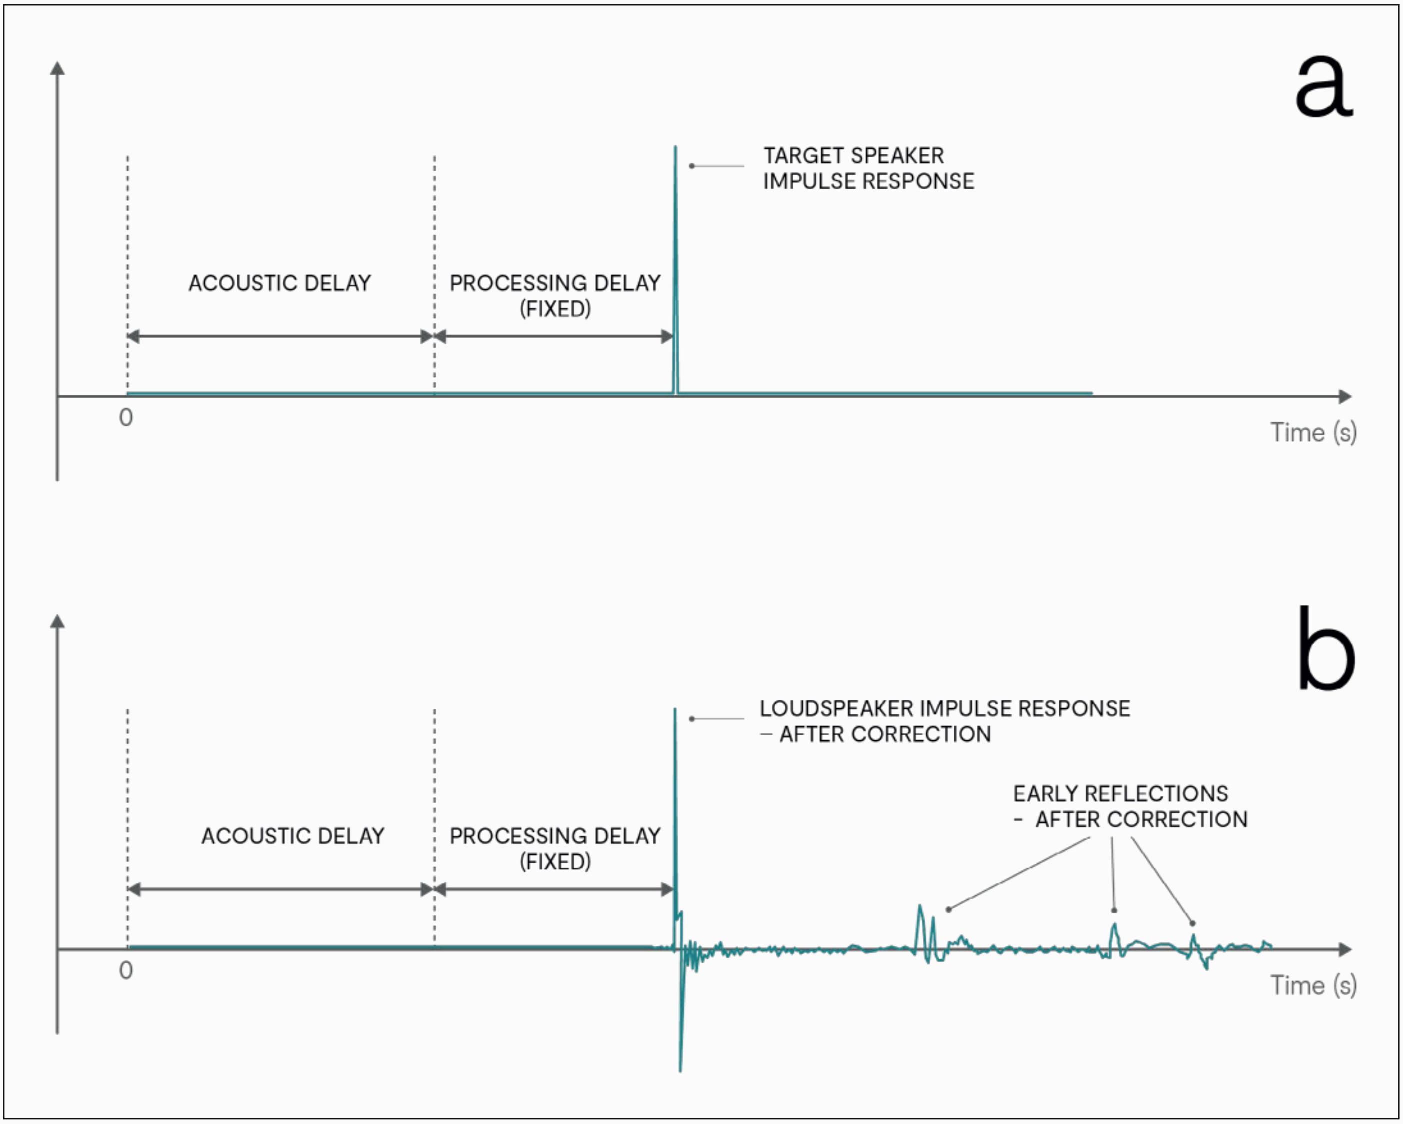

Figure 9. a) The target impulse response at one location. b) A measured impulse response at one location, after correction.

2.2 Correcting the in-room magnitude response

In the previous section, we looked at the response of the speaker-room system in the time domain. We can also look at it in the frequency domain. A frequency domain signal has two parts: magnitude, or the strength of the signal at each frequency, and phase, which represents the relative timing at each frequency. There is a precise mathematical relationship between signals in the time and frequency domains, so they can be converted back and forth with no loss.

Figure 10 illustrates several frequency domain plots. The unit impulse is shown in Figure 10a: it has infinite bandwidth and flat magnitude and phase. This is the frequency domain version of our "perfect" time domain input signal. Figure 10b shows the on-axis anechoic frequency response of a typical high-quality loudspeaker. The loudspeaker has a finite bandwidth and there are some bumps and ripples. The on-axis magnitude response alone is often referred to as the loudspeaker's "frequency response." The phase response exhibits an overall deviation from flat, with a notable change of slope around the midrange/tweeter crossover frequency.

Figure 10c shows a typical magnitude response when the loudspeaker is measured in a room. This is clearly quite differentdifferent from Figure 10b, and the two should never be confused. (We have not shown the phase response here as it is "messy.")

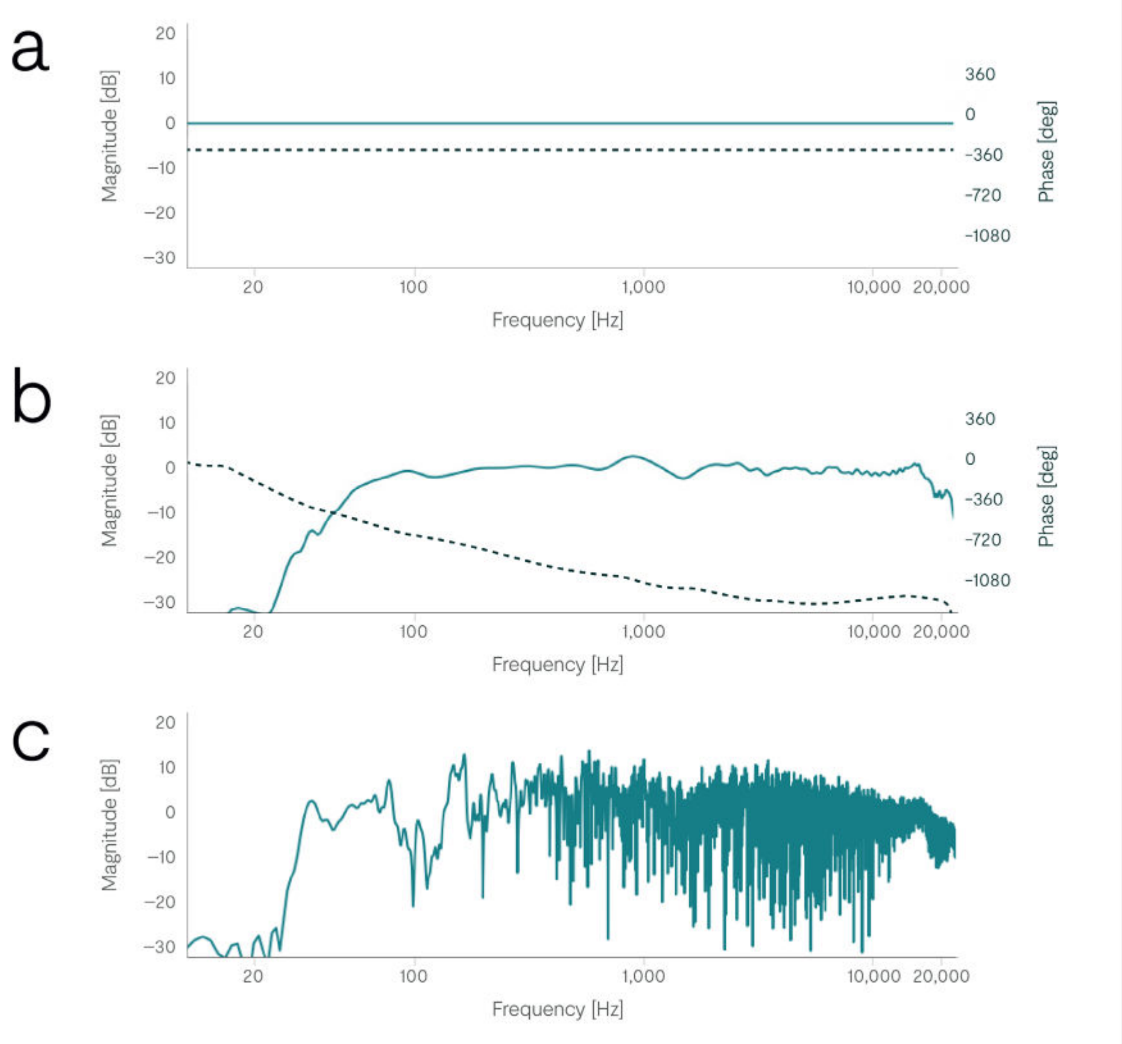

Figure 10. a) The frequency domain plots of a unit impulse. b) A typical on-axis anechoic response of a high-quality loudspeaker. c) A typical in-room magnitude response of the same loudspeaker.

Because the unit impulse has a flat frequency domain magnitude, we can consider this to be the magnitude response target for the error minimization algorithm. In the case of the magnitude response, however, the filter that is created acts to correct the average of all locations to flat. (In contrast to the phase response, for which only common errors are corrected.)

For various practical and acoustic reasons, however, a flat response is not the desired result. For example, magnitude response correction must be limited to the passband of the loudspeaker to avoid excessive gain, which – at low frequencies in particular – would reduce system headroom and introduce distortion. Mathematically, limiting the frequency range is handled by introducing a frequency-dependent weighting factor into the error minimization equation. As mentioned on page 7, a neutral-sounding in-room response typically has a boost in the low bass and a gradual fall-off to the highest frequencies. This is accomplished with a minimum phase filter overlaid on the main correction filter. Figure 11a shows a typical example. Because there is no universally-correct in-room magnitude response – it depends on the loudspeaker, room, and listener preference – this filter is made available to the user for editing as the “target curve” in the Dirac Live application.

The notion of a user-adjustable target curve can be extended to the frequency limits discussed above, which can also be adjusted in the Dirac Live application. In cases where the loudspeakers have well-controlled off-axis responses and the room has a moderate amount of absorption, for example, a user may prefer to limit magnitude response correction to the modal region – up to 200 or 300 Hz, say – as illustrated in Figure 11b.

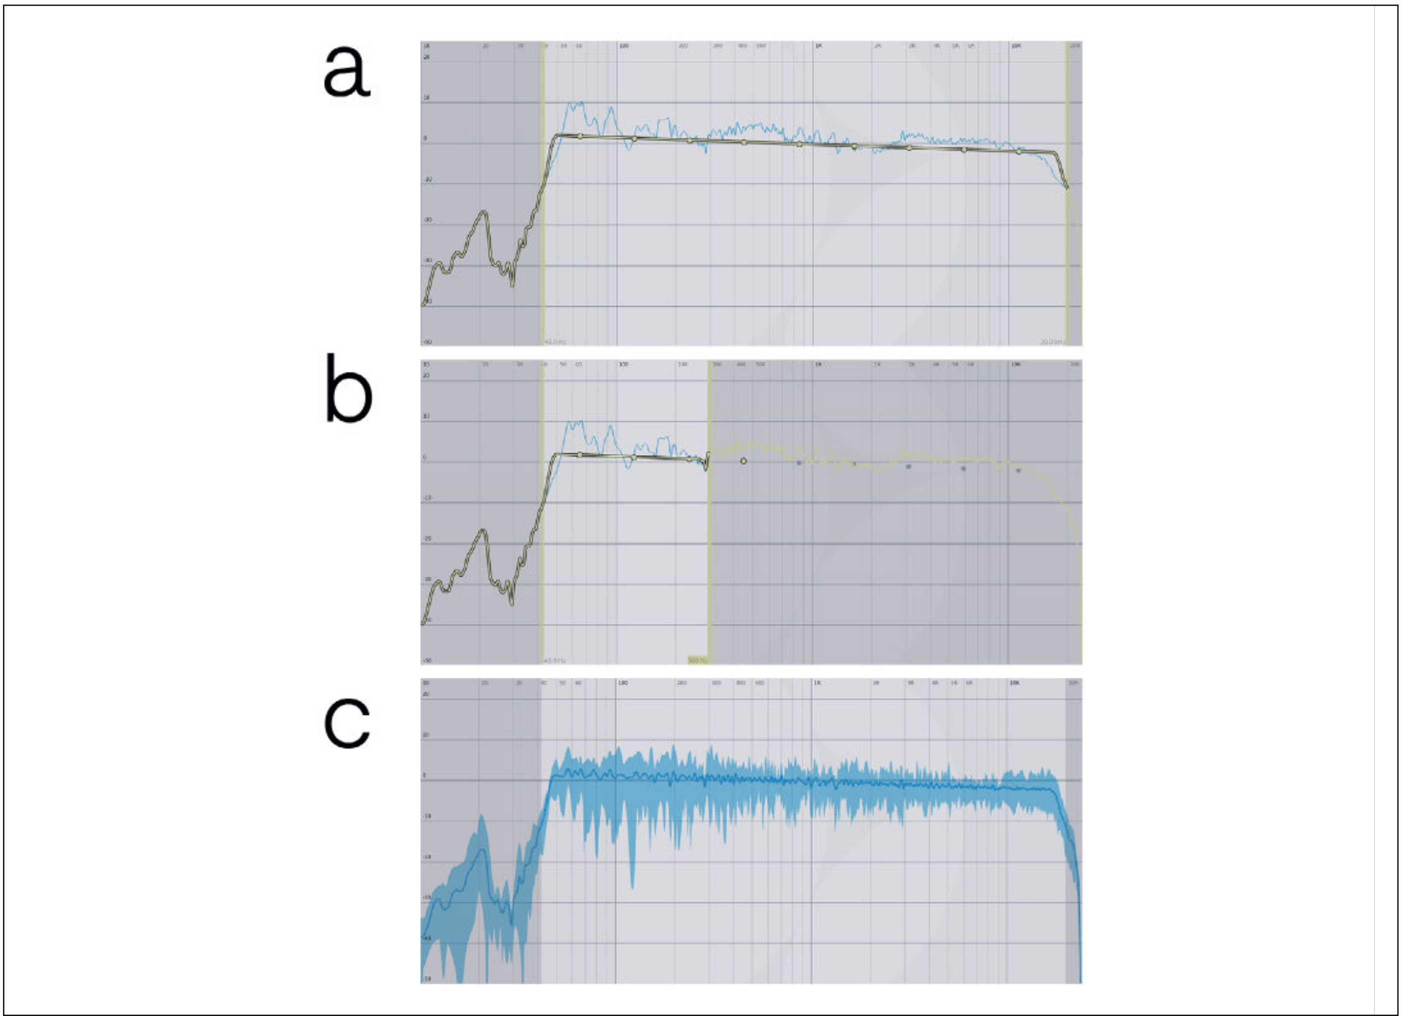

Figure 11. a) A magnitude target curve with default correction range. b) With the correction range adjusted to the modal region. c) An actual corrected in-room magnitude response (full range) showing average response in dark blue and the range of response variation in lighter blue.

Finally, the magnitude response correction filter is gain-limited and smoothed in frequency. The first reduces the likelihood of overdriving amplifiers or loudspeakers. The second increases the robustness of the filter, ensuring that the correction is consistently useful over the listening area rather than optimized just for the measurement locations.

To illustrate with a concrete example, Figure 11c shows the corrected average of nine measurement locations in a typical listening room. The corrected response still varies in different locations; the range of this variation is shown in lighter blue. This implies, of course, that at any one location the measured magnitude response after correction will usually not be identical to the average response.

Subjectively, the improvement from magnitude response correction depends on the frequency range. At low frequencies, the sound of the system is dominated by room modes, so significant changes can be made. The correction flattens out obnoxious peaks and resonances, resulting in a much more even sounding bass with no boom or “one note bass.” Typically, and provided the loudspeakers and/or subwoofer are capable enough, bass response is extended a little further.

At higher frequencies, magnitude response correction takes more of a tonal shaping role, evening out the perceived tonality of the system. Excess brightness or midrange energy may be toned down and hollowed out parts of the spectrum filled in. In a multichannel system, consistency of imaging around the stage is improved by making all loudspeakers sound more similar.

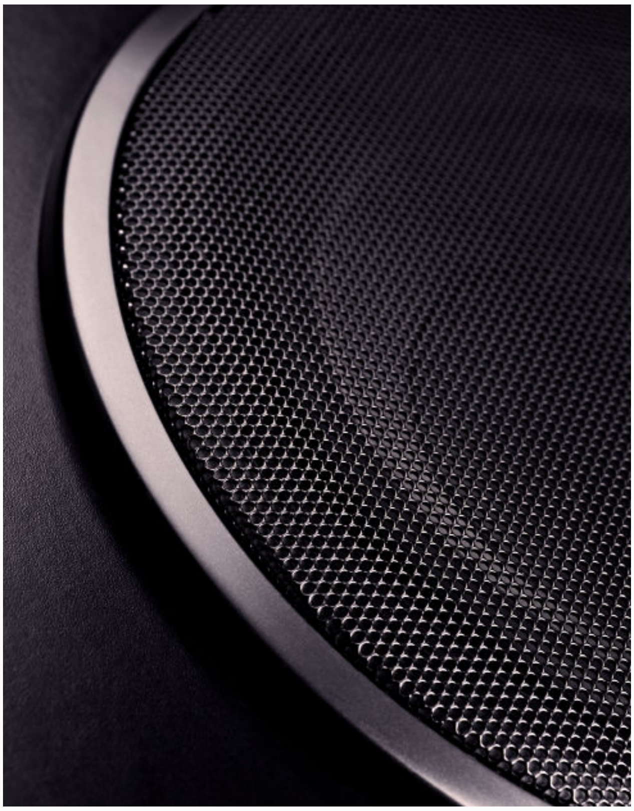

TECHNICAL SIDEBAR

2.3 Mixed phase room correction

Dirac Live is a mixed phase room correction system. What does this mean?

Figure 12a shows a typical in-room magnitude response. There is an infinite number of possible impulse responses that correspond to this magnitude response, determined by the phase response. There is, however, a unique phase response for which the speaker-room system is minimum phase. Figure 12b shows the impulse response in this case. A minimum phase system has, for a given magnitude response, minimum energy delay at all frequencies. It can be inverted – that is, corrected – by a minimum phase filter. Minimum phase filters are common and easy to implement in analog circuitry and digital processing: if you have used a graphic or parametric equalizer in your car or home theater, for example, you have used a minimum phase filter.

Speaker-room systems are not, however, minimum phase. They are mixed phase. To illustrate, Figure 12c is the actual measured impulse response corresponding to Figure 12a. The difference relative to Figure 12b is due to excess phase: phase change in addition to the minimum phase response. Excess phase arises in several ways, but we can start with the loudspeaker, which typically introduces excess phase in the crossover. The reflections and modes in the room also create excess phase. In general, summation of multiple components such as the many reflections and modes can create excess phase even though individual components are minimum phase.

Correcting a mixed-phase system with a minimum-phase filter can lead to objectionable audible artifacts. Dirac Live, however, directly addresses the mixed phase nature of the speaker-room system in the mathematical model that it constructs from the acoustic measurements and that it uses to create the correction filters. Importantly, to avoid introducing pre-ringing into the corrected responses, the algorithm corrects excess phase only if it is common to all measurement locations. Excess phase that occurs in only some locations is therefore not corrected.

By correcting the common excess phase, the impulse response across all locations more closely approximates the ideal unit impulse, evident from a comparison between Figure 9b with Figure 8b. The result – when combined with the minimum phase magnitude response correction described elsewhere in this chapter – is a practical mixed phase correction that is optimally effective over the chosen listening area while conforming to our design principle of not creating objectionable audible artifacts.

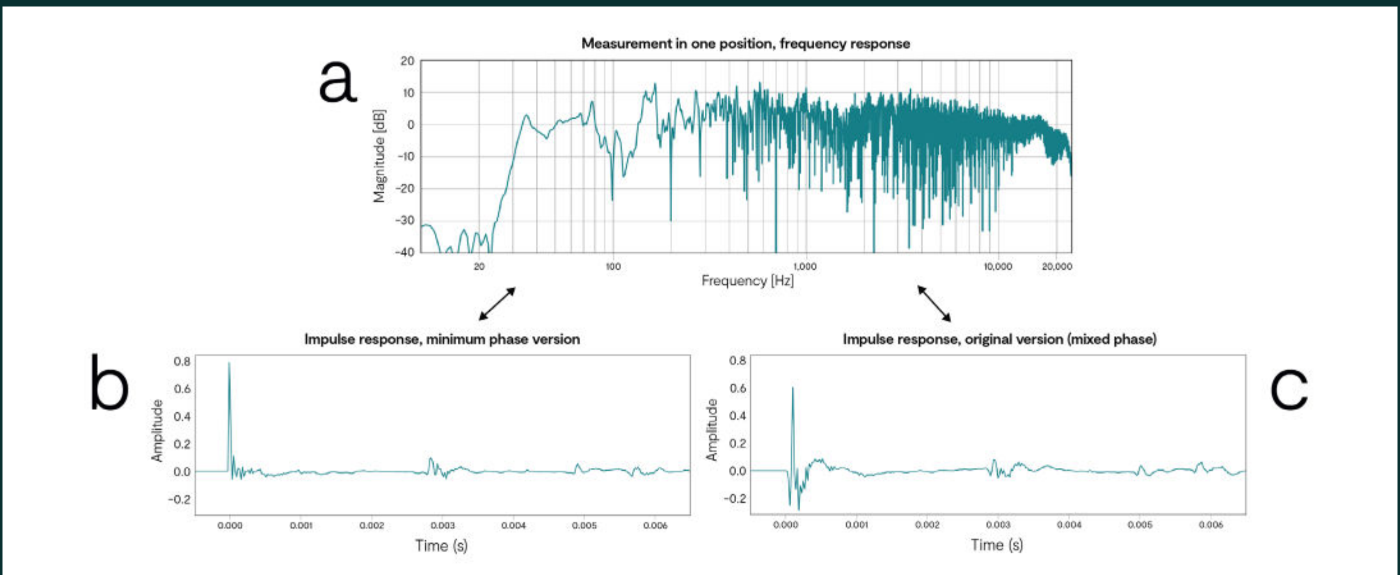

Figure 12. a) A measured in-room magnitude response. b) The impulse response if the system is minimum phase. c) The actual measured impulse response.

Chapter 3
Dirac Live Bass Control

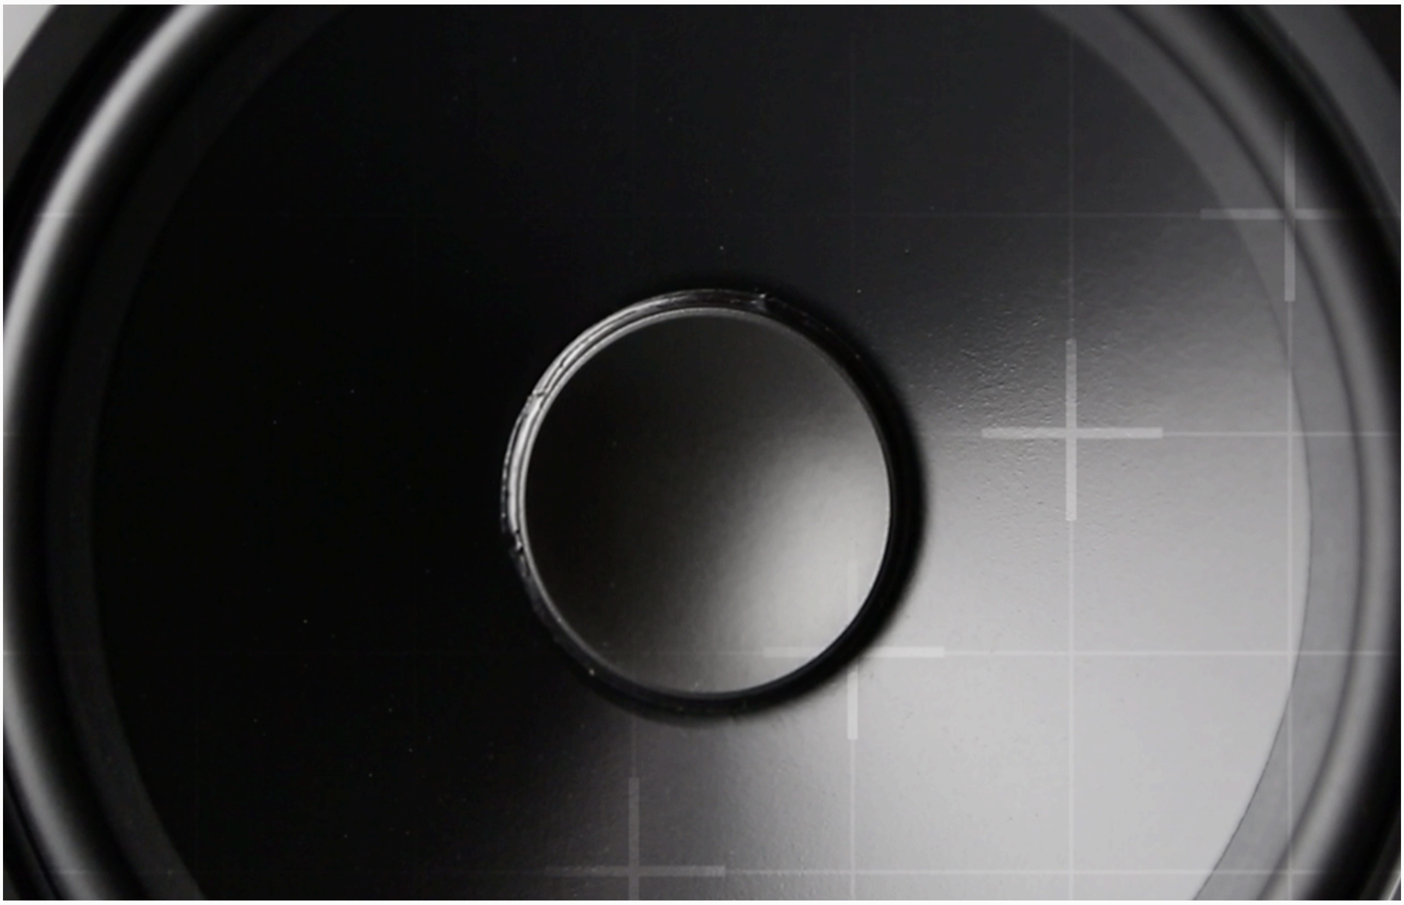

Dirac Live Bass Control (BC) is a highly advanced bass management system. It extends the concept of bass management commonly implemented in home theatre equipment to provide seamless integration with the loudspeakers and phase-matched responses. If multiple subwoofers are used, BC controls them individually to reduce seat-to-seat variation in low bass response.

Figure 13a illustrates a theatre system without bass management. Each loudspeaker channel is sent unaltered to a full range loudspeaker. In addition, the low frequency effects (LFE) channel is sent to a subwoofer. In this system configuration, it is assumed that each loudspeaker can reproduce the full 20 Hz to 20 kHz audible range. In a home environment, this arrangement is usually not practical due to the cost and size of true full range loudspeakers, giving the use of BC a clear advantage. For most home speakers, filtering out the lowest frequencies and sending them instead to a loudspeaker specifically designed for them – the subwoofer – will result in better quality output from that loudspeaker and more, cleaner bass. These low frequencies are summed with the LFE channel to produce a mono subwoofer signal. This concept, called bass management, is illustrated in Figure 13b.

The design of the bass management filters is crucial to obtaining accurate integration between the subwoofer and loudspeakers. While traditional bass management systems rely on simple input from the user – such as loudspeaker distances and the selection of crossover frequency – BC uses a comprehensive set of acoustic measurements to both inform the user and as input data to its advanced filter design algorithm.

Bass management has other advantages in a home environment, especially when the integration is as easy and as seamless as it is with BC. For example, the loudspeakers can be smaller and less expensive; or conversely, they can be higher quality for the same price. Furthermore, the location of the subwoofer can be changed if necessary to improve low-frequency output, which may not be feasible with large full range loudspeakers.

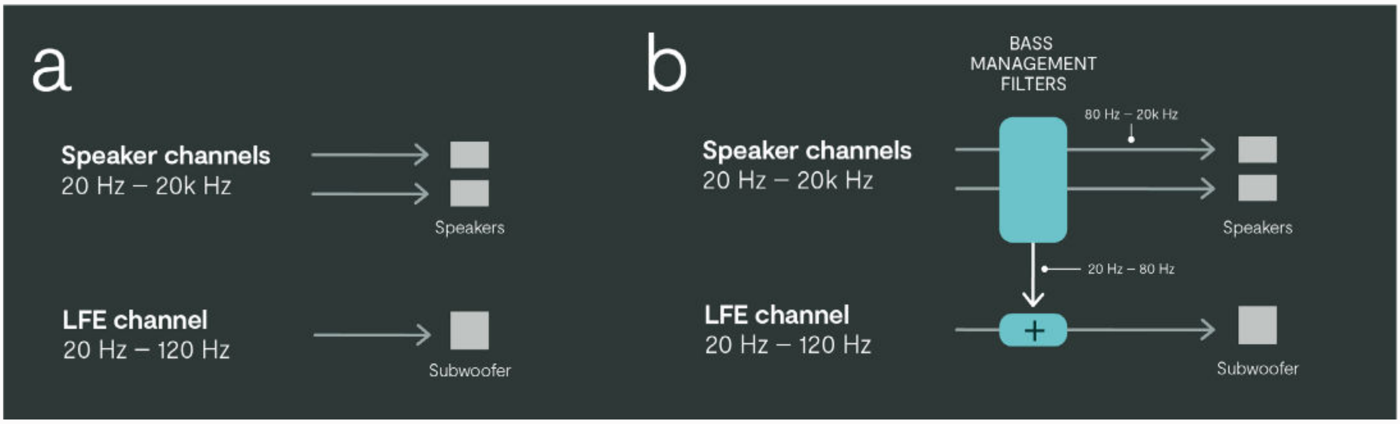

Figure 13. a) An example system configuration with full range loudspeakers. b) An example system configuration with bass-limited loudspeakers and bass management.

3.1 Accurate subwoofer integration

Simple bass management systems typically allow the user to change the time delay of the subwoofer and loudspeakers by setting the distance of each from the listener. Even if these delay settings are accurate, this is not enough to prevent response issues: the loudspeakers and subwoofer are in different locations in the room and there is no guarantee that they will sum to an optimum response. Furthermore, there is no way for the user to truly know how accurate the result is.

Dirac Live Bass Control removes the guesswork. The first step is a simple one: acoustic measurements integrated into the Dirac Live application, so the user can visualize exactly what is happening with bass in the listening room and has the information needed to choose the crossover frequency between the loudspeakers and the subwoofer. The second step is automated: BC, in conjunction with Dirac Live, uses the measurements it has taken to optimize the crossover between the subwoofer and loudspeakers and tame the room modes. This is done partly by setting the time delays on each channel – as for the simple bass management example given above, but fully automated. However, the real key to correct integration between the loudspeakers and subwoofer is matching the phase response at the crossover. This is done with the use of all-pass filters. (For more technical detail, refer to 3.3 All-pass filters.)

In the case of left-right loudspeaker pairs, this is done for the left and right loudspeakers together to ensure that they have matching phase responses at the center of the listening area. This compensates for mismatches that can occur if the room and loudspeaker placement are not fully symmetrical – as is often the case in a home environment. The center loudspeaker is treated as a special case: it is phase-matched with the front left and right loudspeakers.

Figure 14 illustrates an example of successful integration done with BC. The target curve is shown as a dashed black line and the actual response average in a thick green line. (This is one loudspeaker, but measurements of the other loudspeakers show a similar clean result.) The thin green line shows the average response without BC and crossover phase matching.

3.2 Reducing seat-to-seat variation in bass response

BC, when used with multiple subwoofers, opens up the ability to address sound variations between seats.

To illustrate, Figure 15 shows measurements of a system with one subwoofer, at three listening seats. The target curve is shown as a dashed line. There is clearly variation between the seats. For example, seat 2 has a dip of -30 dB at 50 Hz, while seat 3 has a dip of -5 dB at 70 Hz.

This variation between seats is simply a consequence of the modal behavior of the room and the different path lengths from the subwoofer to each seat. Variations depend on the size, shape, and construction of the room, on the location of the subwoofer, and on the location of the seats. These variations can be reduced by using BC. As explained on page 6, the location of a sound source along a mode affects how the mode is excited in both amplitude and phase. In principle, then, multiple subwoofers will excite the modes differently, potentially exciting more modes in a more uniform way.

Simply adding subwoofers, however, will not give optimum results. Best results require that each subwoofer be controlled individually so the combined response has the least variation between seats. This requires a computing algorithm based on acoustic measurements. The method used by BC is to apply all-pass filters to change the phase relationships between the subwoofers. It does not require that subwoofers be placed in any particular location. More diversity, however, is better, so they should be spread out in the room and not (for example) all located in the same area or along the same wall. This leads to increased flexibility for loudspeaker placement in the home.

The search algorithm used is called genetic optimization. This takes a considerable amount of computing power and is not something that can be done by hand. In brief, it starts with default sets of filters with a good spread of parameters and calculates the response variation across the listening area for each of them. The best are then combined to create new sets, then the algorithm repeats. When there is no further improvement, the algorithm stops, with the best-performing set of filters being considered the optimized result. After this result is obtained, BC calculates the optimal integration with the loudspeakers, as described earlier. This time, however, it treats the co-optimized subwoofers as a new mono bass channel with optimized properties.

Figure 16 shows the same room as Figure 15, but with three subwoofers controlled by BC. Below 100 Hz, the worst deviation from the target response is at seat 2, by 8 dB, and all three seats are mostly within +/- 3 dB up to 100 Hz.

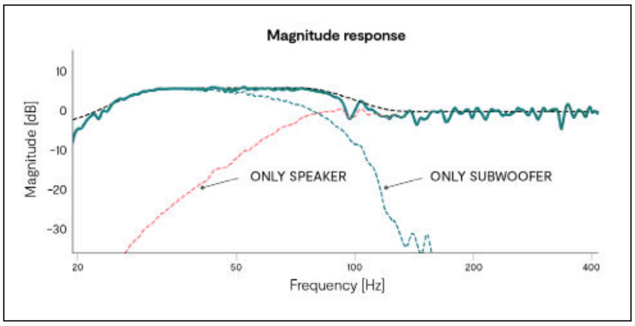

Figure 14. Measured results of integration with Dirac Live BC, showing average responses across the listening area.

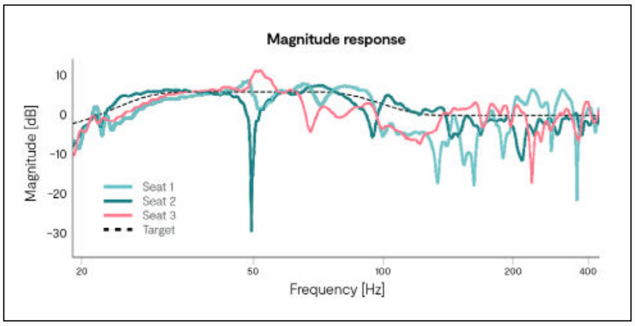

Figure 15. Measurements of bass response at three seats.

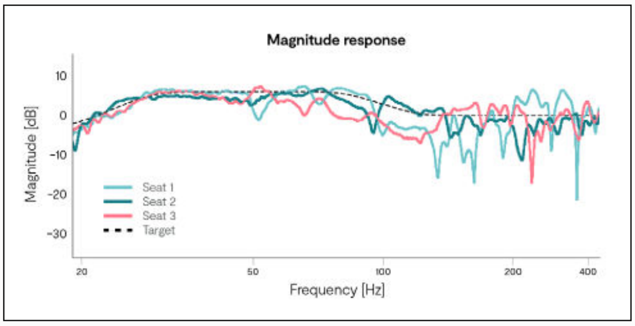

Figure 16. Frequency response measurements at three seats, with three subwoofers optimized and integrated by Dirac Live BC.

TECHNICAL SIDEBAR

3.3 All-pass filters

BC makes extensive use of all pass filters to accomplish its results. An all pass filter changes the phase of a signal without changing its magnitude. Shown in pink (a) in Figure 17 is the flat phase response of an unaltered signal, while green (b) is the phase response with an all pass filter applied. The small waveforms show a single cycle of each to illustrate the phase relationships in the time domain.

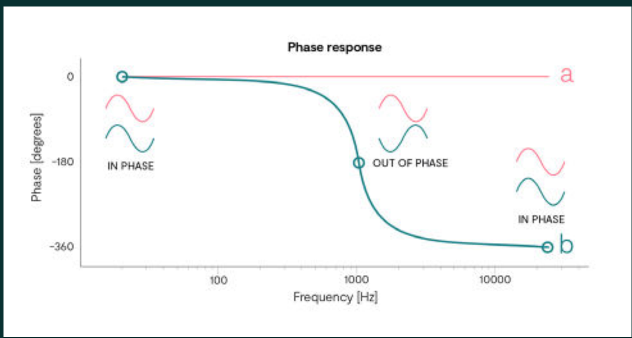

Figure 17. Illustrating the phase response of an all pass filter. In pink is shown a sound source with no filter. Its phase response is flat with respect to frequency. In green is shown a sound source with an all pass filter, which has phase that changes with frequency. The small waveforms show a single cycle of each.

If an all pass filter is applied to a single sound source, its in-room magnitude response will not change. If, however, there is more than one sound source and an all pass filter is applied to one of them, the combined in-room response of the sound sources will change.

To illustrate, Figure 18 shows measurements of two identical subwoofers in a room. The green trace shows the combined magnitude response when both subwoofers are fed with an identical input signal. The pink trace shows the combined magnitude response when one of the subwoofers has its input signal altered with an all pass filter centered at 60 Hz. Even though each subwoofer alone still produces the same SPL at all frequencies, their combined in-room response changes because of their altered phase relationship.

Note that the purpose of this example is simply to illustrate the basic principle by which BC accomplishes its task; it is not intended to be an example of improved integration or reduced variation. BC uses an iterative algorithm to mathematically evaluate different solutions and search for the best overall result.

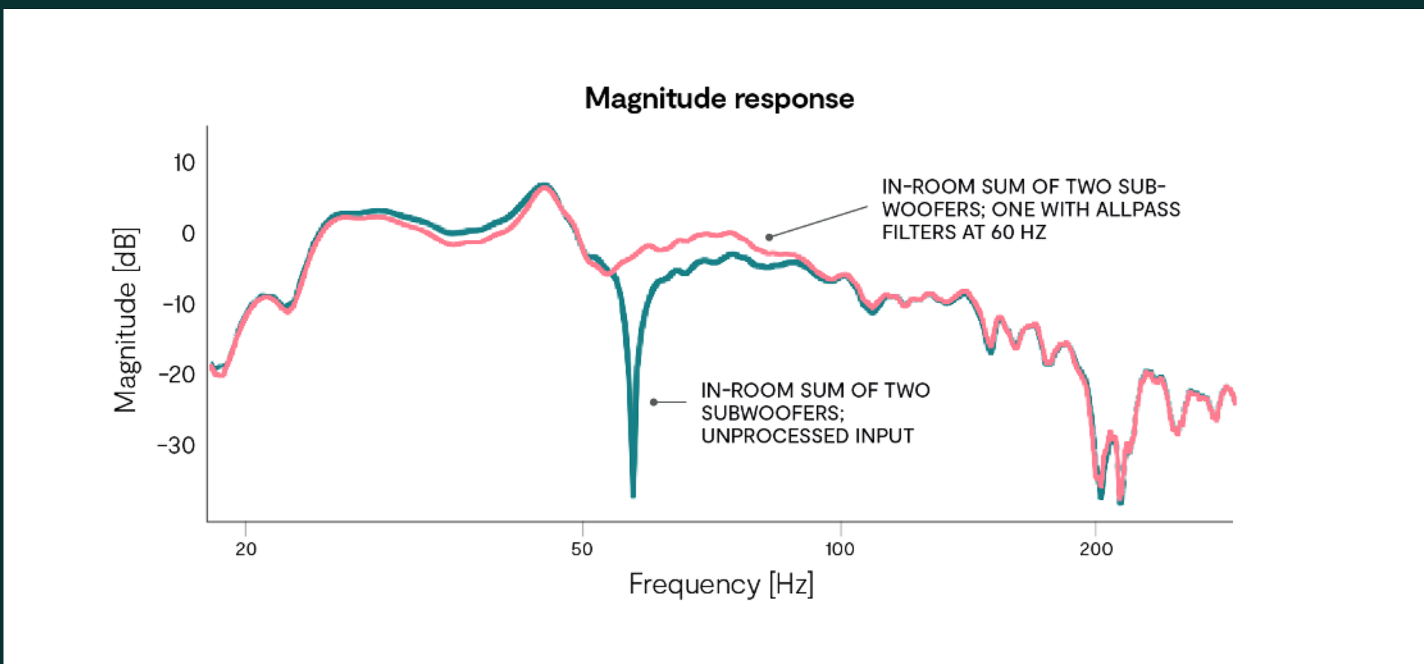

Figure 18. Illustrating how an all pass filter changes response: in green, the combined response of two subwoofers fed with the same input signal; in pink, the combined response of the same two subwoofers when their phase relationship is changed with an all pass filter.

Summary

Dirac Live is an advanced mixed phase room correction system that corrects for undesirable changes to the audio signal caused by imperfect loudspeakers, room modes and reflections. It consists of three functions, of which Room Correction and Bass Control are covered in this white paper.

Dirac Live Room Correction is the core function of the Dirac Live product family. It corrects the time domain behavior of each loudspeaker (the impulse response) and smooths peaks in the frequency magnitude response due to room modes. This technology goes beyond traditional room correction systems by setting a rigorous standard for audio quality based on scientific principles.

Dirac Live Bass Control (BC) is an advanced bass management solution that seamlessly integrates one or more subwoofers with the loudspeakers. If more than one subwoofer is used, BC adjusts the phase of each subwoofer individually to reduce seat-to-seat variation in low bass response. BC's algorithms are designed to optimize the performance of multiple subwoofers, advancing beyond simple bass management systems.

Dirac Live Active Room Treatment (ART) uses all loudspeakers and subwoofers in the system to further improve room correction. By having all loudspeakers cooperate, it reduces problematic low-frequency room resonances. In addition, SPL variation across the listening area is reduced, up to its maximum effective frequency (typically 150 to 300 Hz). ART's adaptability makes it suitable for various room sizes and configurations, offering a scalable solution for different acoustic environments.

These three functions build upon each other, as each is progressively more advanced. Nonetheless, all three are seamlessly integrated into the Dirac Live application running on a computer or mobile device. The availability of each function will, however, vary according to the audio processor.

The design decisions behind Dirac Live consider the subjective effects of every change to the audio signal. The improvements brought by Dirac Live include:

- **Improved clarity and soundstage**
- **Improved bass quality**
- **Reduction of room resonances**

The specific type and level of improvement depends not only on the Dirac Live functions being used but also the particulars of the listening environment: room size and construction, number and quality of loudspeakers and subwoofers, loudspeaker placement and room treatment.

The Dirac Live software application guides the user to an optimized listening experience. Its easy-to-use interface provides the user with an intuitive level of control over its powerful correction algorithms, creating an unmatched combination of acoustic performance and ease of use.

In summary, Dirac Live offers a scientifically grounded, user-friendly solution for improving audio quality in many settings. Its modular approach allows for tailored optimization, making it a versatile choice for those seeking both performance and ease of use.

Thank you for reading

Dirac Research
Phone number: +46 18 4108210
Email: info@dirac.com
www.dirac.com

DIRAC