

ODEON APPLICATION NOTE

Investigate and auralise measured 3D Ambisonics impulse responses in ODEON

AR – September 26

1. Introduction

The ODEON Room Acoustics Software contains multiple tools to analyse spatial impulse responses, measured in the ambisonics format. These features are available in the Auditorium and Combined editions, as well as in the free trial version, although the results are limited to the 1 kHz octave band for the latter.

This application note reviews the available spatial analysis tools for ambisonics impulse responses, following an example available online.

2. ReverbData

The online example used in this note comes from [Reverbdata](#) [1]. Reverbdata is an extensive database containing measured 3D Ambisonics impulse responses of significant portuguese spaces (churches, auditoria, theatres, etc.). In this note, we use an impulse response from the [São Carlos National Theatre](#) (Main Opera Hall) measured in 3rd order ambisonics.



Figure 1: The main hall of the São Carlos National Theatre [1].

The ambisonics impulse responses were measured with a Zylia ZM-1-3E third order ambisonics microphone, and they are available in the ACN/AmbiX format, in 1st, 2nd and 3rd order.

Figure 2 shows the source and receiver positions of the measurements in the theatre. In the note we focus mostly on receiver P1, which is situated at the front of the audience area.




Figure 2: Source and receiver positions in the São Carlos National Theatre [1].

3. Quick reminder on ambisonics signals

Ambisonics is a full-sphere surround format that can be used to capture a three-dimensional sound field. The directional information is encoded through a series of channels, corresponding to spherical harmonics, of increasing complexity. The decomposition in spherical harmonics is infinite, but it is commonly truncated to a certain order, known as the ambisonics order. A given order N requires $(N + 1)^2$ channels.

The lower order channels have a physical interpretation:

- Order 0 (channel 0) corresponds to pressure, as recorded by an omnidirectional microphone.
- Order 1 (channels 1, 2 and 3) correspond to 3 perpendicular figure-of-eight microphones. They can be representative of the gradient of pressure, hence particle velocity.

Especially for first order ambisonics, there are two notable formats, namely **B-format** and **AmbiX**. These two formats differ by the scaling and the ordering of the channels.

- **B-format:** FuMa ordering (W, X, Y, Z), MaxN normalisation with -3 dB on channel 0.

- **AmbiX:** ACN ordering (W, Y, Z, X), SN3D normalisation.

The table below summarizes the differences between B-format and AmbiX for first order ambisonics.

Channel	B-format	AmbiX
0	W	$W \cdot \sqrt{2}$
1	X	Y
2	Y	Z
3	Z	X

Note: for higher order ambisonics, the AmbiX format is usually preferred, as it follows the more logical spherical harmonics ordering. You can read more about ambisonics formats here:

https://en.wikipedia.org/wiki/Ambisonic_data_exchange_formats

4. Opening an ambisonics impulse response in ODEON

The measured impulse response should be saved in .wav format. In this example, we use Ambisonics_B_IR_S1_Audience_Pos1_(ACN-SN3D-3).wav, which is a third order Ambisonics file (16 channels).

Click on  and select the corresponding file, or simply drag-and-drop the wav file from Windows Explorer onto ODEON. The main window “Raw Impulse Response” opens. Note that this window only supports first order ambisonics, that is why it will only display the first four channels.

Important! ODEON supports Ambisonics impulse responses in B-format as well as AmbiX. The default format in ODEON is the B-format, so if your impulse response is in AmbiX, you need to change the format in the interface. This is the case of the current example. If you use the wrong format, the spatial analysis will not make any sense.

The title of the graph indicates the selected ambisonics format:





Figure 3: Raw impulse response display. The title indicates that the impulse response is loaded in B-format.

Change the Ambisonics format in the context-sensitive dropdown menu. This will reload the measured impulse response with the proper format, AmbiX in our example.



ODEON “remembers” your last setting, so once you have loaded an Ambisonics impulse response and changed its type, the next loaded impulse responses will all be assigned the same format.

No matter the chosen input format, ODEON always displays the raw impulse response in B-format (W, X, Y, Z) and most of the analysis we will cover is only based on first order ambisonics, except the auralisation tool.

5. Alignment of impulse response

ODEON automatically rotates the response so that the direct sound comes from the front (0 azimuth, 0 elevation). This automatic rotation is essential to calculate specific room acoustic parameters, such as LF_{80} , which assumes that the microphone probe is oriented towards the source.

For this reason, the observed channels in the raw impulse response may not correspond to the channels in the input file (as they would appear in an audio software). This option can be activated or deactivated in the context drop-down menu (*"Front align ambisonics impulse response"*). The title above the raw impulse response graph indicates the current alignment of the impulse response.




Figure 4: Raw impulse response display. The title indicates that the impulse response is loaded in its original orientation.

6. Sound intensity and hedgehog plots

Go to the Decay Curves tab. The plot shows the decay of energy (directly related to the squared pressure) along with the decay in sound intensity (which is estimated from the first-order ambisonics data).

The instantaneous sound intensity is a spatial vector (3D) that is the product of the pressure with the particle velocity. In the context of an Ambisonics signal, the pressure is estimated from the 0th order and the particle velocity from the 1st order. The intensity decay curve shows the evolution of the intensity norm with time. If the intensity curve is close to the energy decay curve, the sound field is very directional, which can be the case in outdoor settings or in a room with strong echoes. The intensity curve will highlight echoes as strong spikes (a flutter echo will have a "sawtooth" aspect).

By clicking on the right-side button, as shown in the screenshot below, you may toggle between displaying the curves only or the curves alongside a 3D hedgehog.




Figure 5 shows the decay curve and the corresponding 3D hedgehog. The original orientation is used (remember to disable Front align to get the same view).




Figure 5: Energy and intensity decay curve (left) and 3D hedgehog (right).

The hedgehog is another visualization of the intensity vector across time. Each spike corresponds to a local maximum of the intensity norm, as detected in the *I, Measured* graph on the left. The direction of the spike corresponds to the direction of the intensity vector at the time of the local maximum, its length is related to the intensity norm, and its color to the time of arrival.

The hedgehog is a useful tool to identify strong early reflections and echoes, where strong spikes can be identified.

In the example, the direct sound is represented by an orange peak from the left of the receiver going from above (40° azimuth, 50° elevation). This corresponds roughly to the direction of the sound source on the stage. The orange colour of the spike indicates that this intensity peak occurs slightly after the direct sound, which could be due to early diffraction effects from the edge of the stage.

The three next orange peaks coming from behind are probably due to early reflections from the seats.




Figure 6: 3D hedgehog at Position 1, showing only intensity peaks before 80 ms.

After these early peaks, the hedgehog shows very few reflections coming from below the receiver. This can be explained by the absorption and obstruction by the audience seats, which lead to very little floor reflection.




Figure 7: 3D hedgehog at Position 1, showing all intensity peaks.

Note: the hedgehog can also be calculated for a specific octave band. Use the up and down arrow keys to browse through frequency bands. If you are using the demo-version, only the 1 kHz band is available.

Various display settings are available to adjust the hedgehog display:

- **Time resolution (R or Shift+R shortcut):** a coarser resolution will cluster more sound events together, leading to fewer peaks. A time resolution of 10 ms is appropriate for studying echo problems and a resolution of 1 ms or finer can be used to analyse individual reflections.

- **Scale colours (Alt+S shortcut):** multiple scales are available; optimized for music, speech or studio, as well as an automatic scale.
- **Peak detection (Ctrl+P shortcut):** if deactivated, all intensity components at each time frame will be displayed in the hedgehog (see Figure 8).




Figure 8: 3D hedgehog at Position 1, showing all intensity values (not only local maxima).

The following options can be used to browse through the hedgehog:

- Clicking on a hedgehog spike will highlight the corresponding time in the Intensity decay curve and vice-versa.
- You can also use the P and Shift+P shortcuts to move to the next intensity peak in time.

Note: if you have an ODEON project of the measured room, the hedgehog can be loaded within the room's 3D view for additional context. This requires that the room is loaded into ODEON when loading the impulse response and that the measured impulse response corresponds to an ODEON job. This is out of the scope of the present example.




Figure 9: Example of a 3D hedgehog in an arbitrary ODEON room.

7. Decay roses




Figure 10: Decay roses display. Original orientation.

Move on to the Decay Roses tab. The decay roses represent the distribution of the cumulative sound energy on an azimuth plane and an elevation plane over time. They illustrate the spatial distribution of sound in the response. The decay roses are available per octave band as well as in broadband – except in the demo version which only shows the 1 kHz result. As time increases, the values at every direction increase as well.




Figure 11: Azimuth decay roses for Position 1 at 5 ms (Blue), 30 ms (Gray) and Infinity (Kaki).

Figure 11 focuses on the azimuth roses at 5 ms, 30 ms and at Infinity. At 5 ms, most of the energy originates from around 320°, which corresponds to the direction of source S1. At this time, mostly direct sound has reached the receiver. At 30 ms, many additional components have been detected in various directions. The entire response represented by the Infinity graph shows a relatively uniform distribution of components over all directions (about -20 dB in the unit of the graph). Nevertheless, a prominent peak is still clearly visible at around 320° azimuth, which shows the dominance of the direct sound at this position.

8. Room acoustic parameters

Some specific parameters are calculated for spatial impulse responses:

- **Lateral energy fraction (LF_{80} , LFC_{80}):** these parameters indicate the source's perceived width and are based on the early part of the response [2].
- **Late lateral G (L_l):** it represents listener's envelopment and the room's spaciousness. It is based on the late part of the response [2].
- **Steady State Diffusivity:** it compares sound energy with the intensity norm. If the two quantities are equal, the diffusivity is 0 and the sound field is entirely directional [3].

IMPORTANT: these parameters assume that there is only one omni directional source and that the receiver is facing the sound source. Make sure that the response is front-aligned (F shortcut).




Figure 12: LF_{80} at 5 positions in the parterre.

Figure 12 shows the LF_{80} parameter as a function of centre frequency for the 5 positions in the parterre. P1 shows significantly higher values than the 4 other receivers. This can be explained by stronger early reflections from the sides. LF_{80} tends to increase with frequency.




Figure 13: Diffusivity at 5 positions in the parterre.

Figure 13 shows the Diffusivity parameter at the 5 receiver positions on the parterre. Above 250 Hz, the diffusivity is relatively constant across frequency bands. The graph shows a correlation between the distance of the receiver to the source and the resulting diffusivity. P1, which is the closest point shows lower values (around 5 dB). The diffusivity is relatively the same at positions P5, P2 and P4 (6-7 dB). At P3 which is the furthest away, the diffusivity is around 9 dB. Indeed, at short distances, the direct sound and early reflections make the sound field more directional and thus less diffuse. Conversely, at larger distances, the sound field becomes more and more reverberant.

9. Auralization

Last but not least, the **PlayIt!** tool allows you to convolve the full spatial impulse response with an anechoic signal of your choice and auralise it in 3D via headphones. The listener's orientation (azimuth and elevation) can be changed in real time.




Figure 14: Play it! display for binaural auralisation.

The ODEON Play it! supports up to third-order ambisonics, which means that the 2nd and 3rd order channels will also be used in the rendering. A higher order leads to a more precise localization and better out-of-head experience.

HRTF: choose a set of HRTFs for the ambisonics-to-binaural conversion. Read Appendix E of [ODEON's user manual](#) for more information on how to import HRTFs into ODEON.

Note: ODEON provides a couple of sets of HRTFs. They are available with the software's installation.


Headphones: if you have an inverse filter for the headphones in use, you can select it in the corresponding dropdown button. Otherwise, we recommend "Subject_021Res5deg_diffuse.wav".

Anechoic signal: ODEON comes with a collection of anechoic recordings, covering speech, music and miscellaneous sound effects. Change the Signal source/directory and the Signal file to try different source signals.

If you want to use your own recordings, you should place the files in the ODEON WaveSignals folder. Use Options>Explore Wave files folder to open it.

Output device: the dropdown menu lists all audio output devices recognized by your Windows operating system. The Windows icon next to it is a shortcut to the Windows audio settings. We want to listen to a binaural signal, so you should connect and select a pair of headphones.

Listener orientation: the lower part of the "Play it!" display allows you to change both the listener's head azimuth and elevation. Orientation is changed in real time, even while playing a signal. The angles depend on whether your impulse response is in its original orientation or front-aligned. The impulse response orientation can be changed in the dropdown menu "Play it!>Front align..." or using the F-shortcut. The selected orientation is displayed in the interface.



The buttons on the right allow you to either align the head towards the direct sound or revert to the original orientation.

Note on sampling rates: sampling rates may differ between your impulse response, your anechoic signal and your selected HRTFs. In that case, the sampling rate is imposed by the HRTFs, and the other signals will be resampled.

Try loading different ambisonics orders (1, 2 or 3). You will hear that a higher order leads to better quality and localization.

References

- [1] D. Alarcão and A. Couto Pinto, "REVERBDATA - Acoustic Heritage of Significant Portuguese Spaces (Database of Room Impulse Responses - RIR)," [Online]. Available: <https://www.reverbdata.org>.
- [2] ISO 3382-1, "Acoustics - Measurement of room acoustic parameters - Part 1: Performance spaces".
- [3] C. L. Christensen and J. H. Rindel, "Diffusion in concert halls analyzed as a function of time," in *Auditorium Acoustics IOA*, Dublin, 2011.