Waves Audio Ltd.

Jordan Valley College

Implementation of Impulse Response Measurement Techniques

"An Intuitive Guide for Capturing Your Own IRs"

Jonathan Sheaffer[1] and Eyal Elyashiv[2]

[1]Jordan Valley College, Tseamch, Israel
jonathan@yarden.ac.il
[2] Waves Audio Ltd., Tel-Aviv, Israel
eyal@waves.com

April 2005

ABSTRACT

The purpose of this paper is to guide you through the different steps of capturing a sweep response within an acoustic space and transforming it into a ready-to-use impulse response file. This file can then be imported into Waves' convolution reverb, IR-1, and be used in the same manner as the factory-included WIR files. For the purpose of this guide, we focus on two measurement techniques; *mono-stereo* and *stereo-stereo (full stereo)*. A discussion on surround IR capturing, B-Format and Ambisonics is not included in this paper, and can be found on www.acoustics.net .

Overview of the measurement process

Creating a ready-to-use IR file is regularly subdivided into seven steps, as described in Figure 1.1.
We discuss two different methods for capturing IRs, which can be easily implemented using standard pro-audio equipment:

(i) *Method A – mono to stereo.*
Using this method, we capture the stereophonic propagation of a monophonic source within an acoustic space.

(ii) *Method B – stereo to stereo* (full stereo).
Using this method, we try to emulate the stereophonic propagation of multiple monophonic sources (a pseudo-stereo image) within an acoustic space.

[Farina & Regev, 2003]

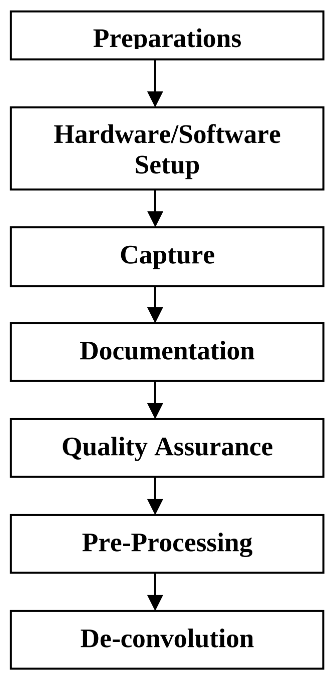

Figure 1.1: The steps for creating an IR file.

1. Preparing for the measurement project

First and above all, one will need to have access to an interesting sounding acoustic space. Wrongfully, acoustical enclosures are sometimes referred to in terms of "Good Sounding" or "Bad Sounding". It should however, be kept in mind that when it comes down to acoustic creativity, there is no such thing. How well an acoustic space sounds, usually relies on one major

factor: “Is the acoustics of the construction suitable for what is intended to be recorded in it?”. For example, recording a heavenly choir in a well designed church will probably produce good sounding results, but attempting to record a heavy-metal concert in that very same location might produce some terrifying noises.
When recording Impulse-Responses, it should be borne in mind that one does not know what these IRs are going to be used for in the future. Respectively, the objective is to generate an IR that was transparently captured at a certain acoustic space, at specific source/listener positions and under certain conditions.

In order to implement the above, the following equipment is required:

1. A multi-track recorder (capable of recording 2 simultaneous tracks). Ideally, this recorder should be able to digitize in 24-bits resolution, and 96 kHz frequency rate.
2. Two condenser matched pairs, with omni-directional or cardoid pickup patterns.
3. If an audio interface that has an embedded analog mixer is not used, there is a need to be able to set the input gain and provide phantom-power. Noteworthy is that the audio signal-flow should be narrowed to a minimum amount of components.
4. A pair of headphones, for monitoring purposes.
5. A loudspeaker or powered monitor (with a wide frequency range).
6. A sweep test signal that will be played through the monitor and into the acoustic space.
7. An inverse-sweep that will be used for deconvolution purposes.

It should be kept In mind that when capturing the IR of an acoustic space, *the IR of the playback and recording system(s) is captured as well*. There is no apparent method for removing these artefacts. Hence, measuring systems should be *as transparent as possible*.
[Farina, 2000]

2. Hardware / Software Considerations

2.1 METHOD-A; Mono-Stereo Setup

As previously mentioned, by using this method we attempt to capture the stereophonic propagation of a single audio signal. The idea behind this is to simulate the way human beings perceive sound. This is originally monophonic, but develops stereophonic characteristics over time due to reflections from the acoustic space being measured.
[Farina & Regev, IBID]

Place a sound source at the source position and place a stereo microphone configuration at the listener position, as described in figures 2.1 and 2.2;

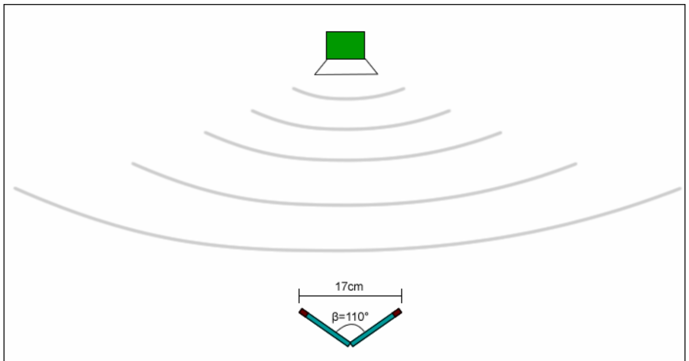

Figure 2.1: *Mono–Stereo* IR capturing using a single audio source, and ORTF mic configuration.

Figure 2.1 describes an ORTF mic configuration that captures sound waves propagated from the single audio source, and reflected off the walls, ceiling and floor.
We won’t go into discussion on stereo microphone techniques, but in case you are not familiar with the ORTF configuration, here are the basics:

- Use two condenser microphones, with cardioid pickup patterns, and a frequency response as flat as possible.
- The microphones are crossed with 110 degree angle between them.
- The distance between the two capsules should be exactly 17 cm (or 6.69 inches).
- The distance between the microphones and the source signal will directly reflect on the balance between the direct sound, early reflections and diffused reflections.
- Scientifically, the ORTF configuration is the closest simulation of how our ears perceive sound, using only two microphones.

[Martin, 2003].

Alternatively, a different approach may be endorsed. Figure 2.1 describes the same mono-stereo method, but this time using conventional spaced-omni AB configuration:

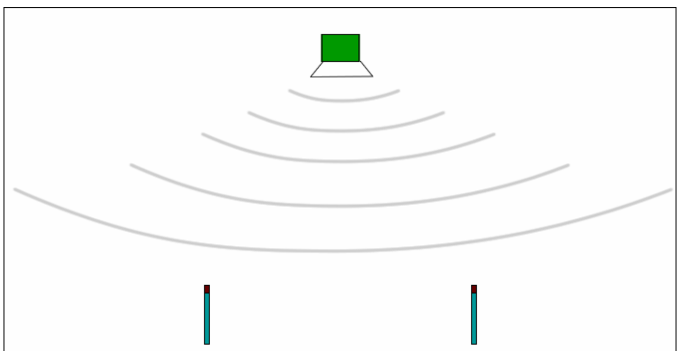
Figure 2.2: *Mono–Stereo* IR capturing using a single audio source, and AB mic configuration.

For the AB technique, two omni-directional microphones should be set up parallel to each other, approximately 1.5 meters apart. The distance between the microphones is determined by the size of the sound source. This technique usually produces a good, wide stereo image. The image is created by the level, time and phase differences due to microphone placement.

Once the desired microphone configuration has been decided upon, the recording system should be set up to simultaneously playback the sweep tone through the loudspeaker while recording the inputs of both microphones into two different channels. We suggest naming the recorded files as follows:

- 'Center_L.WAV' – for the left microphone
- 'Center_R.WAV' – for the right microphone

Before the capturing process commences, gain structure should be calibrated so that input signals are recorded at headroom level in the digital domain (about -4dBFS). After the capturing process, the recorded IR should look somewhat like figure 2.3:

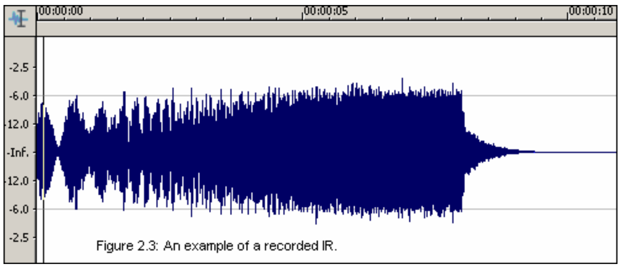

Figure 2.3: An example of a recorded IR.

Only once the playback/recording system has been correctly calibrated, the IR should be captured.

2.2 METHOD-B; Stereo-Stereo Setup

While mono-stereo capturing is an appropriate solution for artificially reverberating a monophonic source (such as a single musical instrument in a mix, or anechoic speech for ADR), many applications require a more detailed simulation of the propagation of sound waves projected from a stereophonic panorama.

In order to simulate this effect, we need to record sweep tones coming from two different source locations (the physical boundaries of the stereophonic panorama) using two different microphones at a time. In sum, this produces four recorded impulses:

1. 'LL.WAV' – Left to Left. This simulates the signal projected from the left boundary and perceived by the left microphone (or the listeners's left ear).
2. 'LR.WAV' – Left to Right. This simulates the signal projected from the left boundary and perceived by the right microphone.
3. 'RL.WAV' – Right to Left. This simulates the signal projected from the right boundary and perceived by the left microphone.
4. 'RR.WAV' – Left to Right. This simulates the signal projected from the right boundary and perceived by the right microphone.

Eventually, an accurate stereophonic image of the reverberation in the measured acoustic space is portrayed (this is why this method is referred to as 'full stereo' – we accurately capture the complete stereophonic panorama.). Should this approach be implemented, the capturing process should be repeated twice, once for each IR recording.
The first time at left boundary and the second time at right boundary loudspeaker positions.
[Farina & Regev, IBID]

Figures 2.4 and 2.5 show the two steps of capturing an IR using the stereo-stereo method:

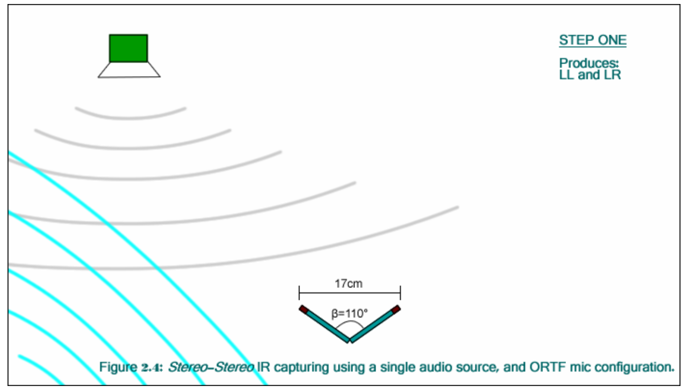

Figure 2.4: *Stereo–Stereo* IR capturing using a single audio source, and ORTF mic configuration.

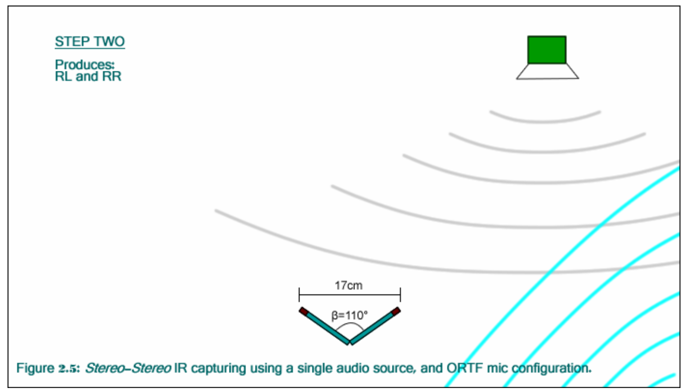

Figure 2.5: *Stereo–Stereo* IR capturing using a single audio source, and ORTF mic configuration.

Much like Method-A, the same procedure should be followed. It is important to calibrate gain structure only once - gain settings should not be changed between the two steps in this process.
Since the IR1 plug-in may input monophonic signals even when working in *full stereo* mode, when capturing IRs using this technique, a monophonic source coming from the centre of the phantom-stereo image should also be recorded (as described in method-A).

We suggest naming the recorded files as follows:

- 'Center_L.WAV' – left microphone (mono)
- 'Center_R.WAV' –right microphone (mono)
- 'Left_L.WAV' – left microphone (left source)
- 'Left_R.WAV' – right microphone (left source)
- 'Right_L.WAV' – left microphone (right source)
- 'Right_R.WAV' – right microphone (right source)

3. Capturing the Impulse Response

Once the system has been tested, the IR may safely be recorded. Before making the actual recording, it is important to make sure that any of the setup has not been changed (microphone stands tend to tilt over time, EQ settings might not be bypassed, etc…)

While already present at the measured location, we recommend capturing impulses at a variety of listener positions. For example, inside a concert hall, these positions may vary from first row seats, through the center of audience and up to the balcony. The more IRs captured, the better are the chances of recreating a truly unique and interesting sound in the future.
For each set of IRs, a separate folder should be created. This helps distinguishing between the different recordings.

4. Documentation Guidelines

It is very important to collect all relevant documentation while present at the measured location. These documents usually include:

- Digital photos of the measurement process microphone configurations, loudspeaker positions, etc…)
- Digital photos of the construction itself (indoors and outdoors).
- Blueprints of the architectural construction (if available).
- Physical measurements of distances and angles, between microphones and loudspeakers.
- Microphone and loudspeaker manufacturers and models that were used for the project.
- General information about the location, or a description why is the location special (historically, acoustically, or even politically).
- Any other relevant information to your discretion.

5. Quality Assurance

Before disconnecting the system and leaving a location, it is recommended to make a short assessment of the quality of the measurements. It is always a good thing to listen to the recordings on location and make sure the recordings are:

- not too quiet
- not too loud or distorted
- not cancelled when summing to mono
- digitized using the maximum sample rate and bit depth available

We also recommend importing the IRs into the convolution reverb module and making sure that everything is in working order. Taking a few minutes to check the results is much easier than reassembling the whole setup after figuring out that something went wrong.

6. Pre-processing and Deconvolution

Once the IRs have been captured, one needs to prepare for the de-convolution process. Waves' IR1 version 2 includes a built-in de-convolution algorithm and an embedded IR editor. It is always recommended to pre-process the IRs using an external audio editor. But for the purpose of this guide, we edit and de-convolve using IR1's internal features.
During this process, we take raw .wav files and turn them into finalized .WIR ("Waves Impulse Response") files that are natively supported by the IR1 convolution reverb.

Accordingly, a few procedures should be followed:

1. .WAV files should be imported into the IR1 environment, according to the LL-LR-RL-RR configurations.
2. The starting point of the direct source should be located.
3. The ending point of the reverb tail and the beginning of the noise floor should also be located.
4. IR ending should be trimmed and faded out.

Run the IR1 v2 plug-in through an audio host application. On the "LOAD" menu you will notice two options:

1. "Import impulse response from file" – this option allows importing a custom-made IR from a file. The IR1 assumes that this file has already been de-convolved and is ready for use.
2. "Import sweep response from file" – this option imports recorded sweeps into the IR1 environment, and de-convolves the sweep files with the inverse-sweep that is located inside the…\IR1Impulses V2\Inverse Sweeps folder.

Since the sweeps haven't been de-convolved yet, we choose the second option from the LOAD menu, as described in figure 6.1:

Figure 6.1: IR1v2 "Load" menu.

Sweep responses should be loaded in this specified order:

1. L-L
2. L-R
3. R-L
4. R-R

If Method-B has not been used, only two sweep files are available. Each file should be loaded twice, for example:

- for L-L, 'Center_L.WAV'
- for L-R, 'Center_R.WAV'
- for R-L, 'Center_L.WAV' again
- for R-R, 'Center_R.WAV' again

After loading these files, IR1 will automatically begin the deconvolution process. This takes some time, depending on processing power of the deconvolving computer.
Once IR1 has finished deconvolving the files the following message is displayed:

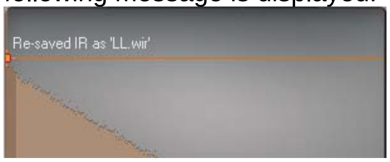

This means that IR1 has successfully de-convolved the sweeps, and has created a .WIR file (in our example, LL.WIR). We suggest renaming this file (making sure the .WIR extension is kept), and copying it into the appropriate impulses folder.
If "direct-off" mode is chosen, there are no further operations. The IR1 deconvolution algorithm will automatically omit the direct source and the floor noise after the reverb tail. Using this mode, the 'direct' source is the dry signal inputted into the IR1 module.

However, if "direct-on" mode is chosen, a few more procedures should be followed. Using this mode, the 'direct' source is the actual convolved version of the dry source inputted into the IR1 module.

Since the original excitation signal is a log-sweep, some non-linear responses are measured due to harmonics from the loudspeakers and microphones. These responses are deconvolved into signals earlier in time than the linear measured impulse response, and are referred to as *Pre-Responses*.
[Ben-Hador & Neoran, 2004]

In order to assure accurate convolution, these pre-responses should be removed from the IR using the "convolution start" parameter.

Bibliographic References

[1] Ben-Hador, R. & Neoran, I. (2004) – "Capturing Manipulation and Reproduction of Sampled Acoustic Respones"; 117th AES Convention, San-Francisco.

[2] Farina, A. & Regev, A. (2003) – "Recording Concert Hall Acoustics for Posterity"; AES 24th International Conference on Multi-Channel Audio; p. 9-12.

[3] Martin, G. (2003) – "Introduction to Sound Recording", stereo microphone techniques; p.490-491, www.tonmeister.ca

[4] Farina, A. (2000) – "Simultaneous Measurement of Impulse Response and Distortion with a swept-sine Technique"; AES 108th Convention, Paris.