Using Liberty Instruments' PRAXIS for Room Sound Convolution

Rev 9/12/2004

Overview

Room Sound Convolution is an operation that allows you to measure, save, and later recreate a representation of the sound of a room and its audio system. There are three main steps involved in this process. These are

(I). Get the binaural Impulse Response of the room at a listener position.

(II.) Process an available sound recording using this Impulse Response to produce a new recording

(III.) Listening to the new recording that represents what the original would sound like if it were played over the system in that room and heard at the same listener position.

The process assumes that only "linear" effects act on the sound. In other words, distortions such or clipping, loudspeaker compression are not reproduced in the final recording(s). These are assumed to be at reasonably low levels.

I. Getting the Impulse Response

Description: The record of the room's sound is saved as a binaural "Impulse Response" (or IR). You can think of an IR as the result of applying a very brief transient spike or pulse to the room's audio system. The speaker will change the pulse to sound, modifying it based on the loudspeaker's frequency response, and the pulse will then propagate in air to the measurement microphone. Of course, the reproduced spike will travel everywhere else in the room, too, reflecting many times from various surfaces and losing energy all the while until it decays into the noise. These multiple reflections will eventually reach the microphone, contributing to the overall impulse response, as shown in Figure 1.

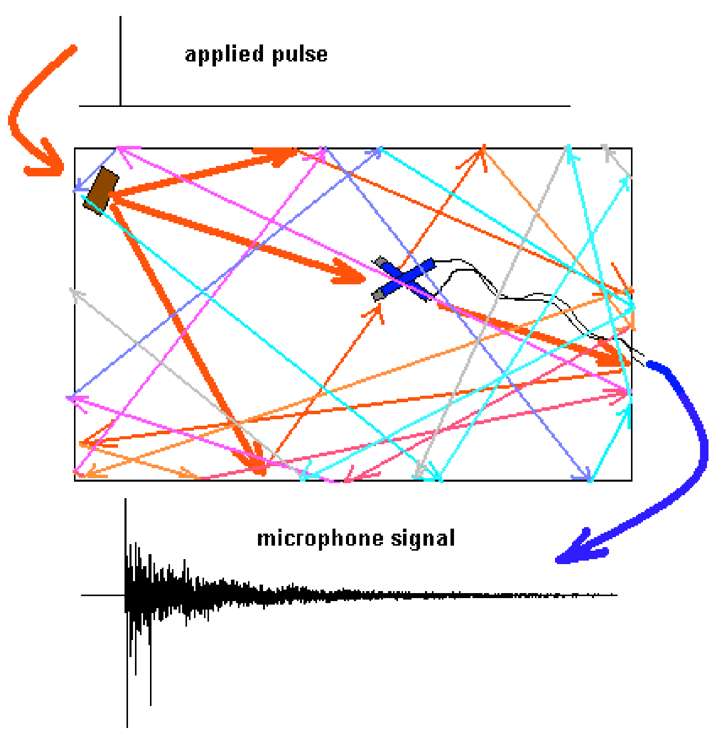

Figure 1: Diagram showing a small portion of the reflections of a reproduced pulse from a loudspeaker in a room arriving at a pair of microphones, and the resulting signal from one of the microphones.

An obvious way to measure the IR would be to use this process literally: apply a narrow spike and record the result using a pair of microphones. But that doesn't work very well. A pulse is very brief and the reverberation in the room can continue for a long time. To keep all the needed reflections well above the ambient noise, the pulse would have to start out at an extremely high level -- higher than a typical sound system could produce or than some microphones could even handle.

So, instead of a pulse, other more spread-out test signals are typically used along with computer software to compute the IR from the original applied signal in a process called "deconvolution". Some typical test signals that are used for this purpose are maximum length sequences (MLS) and sweeps or "chirps". The signals that we will use here are based on logarithmic sweeps, as this has advantages in minimizing low frequency noise and in rejecting unwanted distortion effects.

If you have the Liberty Instruments PRAXIS AudPod hardware, you can easily measure the binaural IR of the room, letting the measurement system handle both the "stimulus" (the signal that drives the sound system) and the "acquisition" (the signal picked up by the microphones, in this case). This is fast, but it does require you to run cables from the AudPod to both the sound system's panel and from your measurement microphones, which may be inconvenient and prone to hum pickup.

If you don't have the PRAXIS AudPod, or if you want to avoid running long cables, you can instead play an audio CD to provide the stimulus sound, as will be described below. You can then record the result using a DAT recorder, or recording software on a laptop computer with a good quality soundcard. The resulting file can then be processed using the PRAXIS software to produce the needed binaural IR file. This processing can be done using PRAXIS in its (free) demo mode.

The process described below may seem long and complex, involving pages of detailed directions. But after you've done it once, you will find that it is not at all difficult – it is much easier to do the process than to describe it.

What you need:

- A CD recording of a series of logarithmic sweeps (known as "Chirps" in PRAXIS jargon). You can make your own CD from files that can be downloaded from Liberty Instruments' web site at www.libinst.com. Or you can use a WAV file that can be played through the PA system from any compatible playback device. You also need single "Ideal" versions of these signals. These are included in the same file download.
- A pair of microphones spaced about 8 inches apart or a dummy head with internal microphones. Some authorities on this subject specify using a Blumlein pair of "figure-of-8" microphones, directed at right angles.
- A stereo recording device that can produce WAV or PX2 files. It should be a linear, non-compressed recorder – do not use an MP3 recorder or devices having automatic level control. The recorder could be a DAT or Flash recorder device with software to convert the output to WAV. Or you can use a laptop computer and a good quality USB soundcard with recording software. Beware that built-in soundcards on most laptops are of rather low quality and will not usually work very well for this application. A PRAXIS system, with a USB soundcard and an Audpod, using "Stimulus = none" and "Acquisition = Time Domain 2ch", can also be used to make the recordings directly.
- The PRAXIS software, version 2.11 or later. This large application can be downloaded at no charge from www.libinst.com and can be used in "Free Mode" (that is, without the AudPod device) to postprocess the files without restriction.

How to do it:

- If you will be using a CD to play the stimulus, you can make your own stimulus CD using a burner and a burning program. The source file is named "Stim44kHz.wav". If you are using NERO to burn the disc, just make an Audio CD project, and drag the file to the project a number of times (to make convenient multiple identical audio tracks), and then burn.
- If you will be playing a WAV file directly for the stimulus, you can use the "Stim44kHz.wav" file if your device works well at 44kHz, or the "Stim48kHz.wav.wav" file for 48kHz operation. The files have 2-second gaps between sweeps.
- Place the microphone pair (or dummy head) at the designated listening position, connected to the recording device you will be using. Keep noisy devices away from the microphones and minimize noise in the area.
- Play the stimulus sounds from the PA system at a moderately loud level, but not so loud that you distort or damage anything (including your hearing!) – **NOTE THAT THE WAV FILES ARE RECORDED AT FULL SCALE, SO THEY CAN BE VERY LOUD! START AT A LOW LEVEL AND TURN IT UP GRADUALLY!**
- Record three of the sweep bursts so that you are sure to get at least one entire burst and the several seconds after it in your recording. The recording should be at either a 48kHz or a 44.1kHz sample rate.
- If you recorded with PRAXIS and its AudPod, the file is already where you need it.

But, otherwise:

  - Convert your recording to the WAV format if necessary, and then transfer that WAV file to the computer that is running the PRAXIS software. Then start PRAXIS. *Note: The first time you start PRAXIS, you can close the "What's New" and "Script Launcher" forms that appear, and click "Ok" on the IO/Mixer matching form.*
  - In PRAXIS, find the form labeled "(Primary Plot)". Use the "File" menu of that form, and select "Open". Where the forms shows "Files of Type", select "WAV (*.wav)" and then browse to find the file you previously recorded and click OPEN to bring it into PRAXIS' Primary Plot.

- On PRAXIS' Primary Plot, there is a menu labeled "Format". Click that, and a row of tabs will appear. These tabs are used to adjust the appearance of the Plot.
- Click on the tab that is labeled "Time", and then double-click on the underlined word "<u>Stop</u>" to display all the data to the end.
- Now, locate a small square near the bottom of the Primary Plot form that contains a picture of a gray lighting bolt. Clicking the mouse on this will cause the display to automatically scale ("AutoScale"), formatting the traces vertically within the plot.
- You should be able to see two tracks (one above, one below) each showing the three bursts that you recorded. It should somewhat resemble the figure below (the bursts will probably be shaped differently):

Figure 2: A recording of three bursts. Arrows point to the blip transients that are used to find the starting points. The third burst has been selected by the Windowing control, causing the unselected parts of the top trace to change to white.

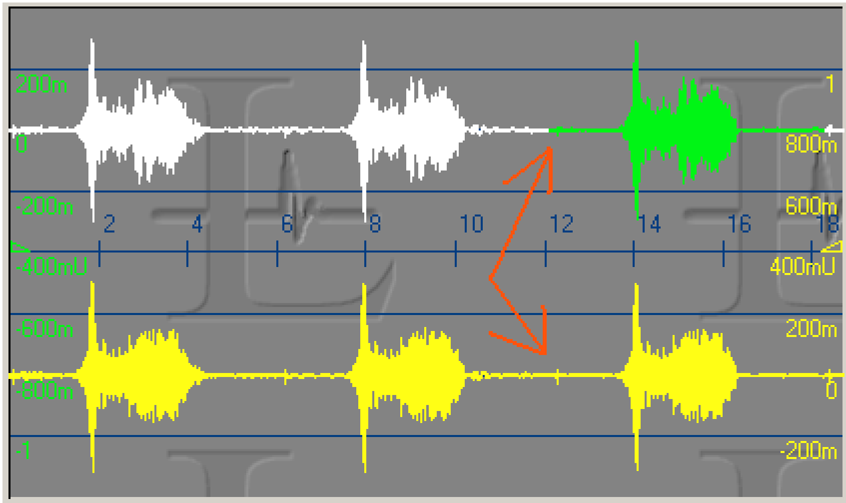

- The beginnings of the stimulus bursts include narrow pulses that are intended to help locate the beginning of the burst. These are needed because the stimulus sweep starts at very low frequencies, below where any loudspeaker is able to play, but the entire sweep response is needed for deconvolution. The added pulses will show as small transient "blips", about 6 seconds apart, both before and after your recorded bursts as shown in Figure 2.
- You need to select the section of this time record that contains ONLY one blip, the following burst, and the time just before the next blip, as shown in figure 2. Choose the burst (of the three) that appears to have the least noise in the flat areas.
- There is a Format tab labeled "Windowing" in the Primary Plot. Click on this tab to bring up the controls for this tab. You will use the Windowing controls to select what part of the tracks that you want to keep, as follows:
    - Double-click the underlined words "<u>Stop At</u>". This will set the right windowing edge to the latest (rightmost) part of the time record.
    - Then click the mouse within the white number-editing area that is to the right of the words "<u>Start At</u>". This will enable you to set the starting edge of the selected section. You can type a number here if you want, but it is easier to do it graphically. After clicking in this area, move the mouse cursor so that the vertical line in it is over one of the traces and SLIGHTLY TO THE LEFT of the first blip transient that you want to include, and click the left mouse button. This selects the starting time for BOTH traces -- but only the top trace will appear to change (to white in the region that is to be "windowed out" or removed).
    - **IMPORTANT**: you need to include some "dead time" (about a tenth of a second or more of no-signal trace) just ahead the first blip that follows your starting edge.

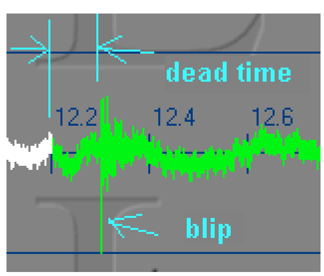

Position the cursor and click the mouse until you have the edge placed correctly. Then press your [Esc] (escape) key to release the mouse cursor back to its normal function.

- Similarly, click the mouse within the white area that is to the right of the words "Stop At", and this time position the cursor over the section JUST TO THE LEFT of the following blip that you do NOT want to include (see figure 2). The region in the top trace following this should then change to white. Press the [Esc] key to release the mouse cursor.

- Now, you will PostProcess the data to leave only the approximately 6 second long section that you have just selected. To do this, find the form in PRAXIS that is labeled "Liberty – Praxis- (MAIN FORM)". At the right in this form is an area labeled "PostProcess" with a drop-down box just beneath. In the drop-down, select the PostProcess type called "Time Dom. Length". In another nearby drop-down box labeled "Source", select "Primary Plot". In the drop-down labeled "Target", select "<create new plot>". There is a PostProcesss section button with a calculator icon that is labeled either "Configure" or "(hide configure)". Click on it one or two times until it reads "(hide configure)" and you should see a form labeled "(Time Lengths)" appear.
- In the PostProcess configure form "(Time Lengths)", choose the setting "Use windowed section only". Under "Gain", select "Specify dB", and set Ch1 and Ch2 to "0dB". Then click the "Do PostProcess" button. You should see a new plot appear, containing only the time region that you have previously selected. You can save the shortened data that is in this plot if you wish, using the "File, Save" menu of the new plot, saving as file type "px2". Keep this Plot on the screen.
- Now, you need to load a copy of the original swept burst into the Primary Plot. Use its "File, Open" menu to browse and select the previously downloaded file "Reference48k.px2" if you recorded at 48kHz rate or the file "Reference44k.px2" if you recorded at 44.1kHz rate.
- The next PostProcess operation will derive the impulse responses from the data contained in these two plots. Go back to the Main Form, and in the top PostProcess drop-down box, select the PostProcess type "Time Dom. Math". Select the Source to be the Primary Plot, and the Target to be the plot (probably "Plot 2") that contains your shortened recording.
- Click the button with the calculator icon, as needed, to again display the PostProcess configuration form, which will now be labeled "(Time Math)". In that form, choose "Deconv. Target via Src ref" – that means "deconvolve the Target plot traces using the Source plot traces as a reference". Again choose "Specify dB" and "0dB" for both channels, and then click the "Do PostProcess" button. Depending on the speed of your computer, this process could take a rather long time, about 9 seconds on a 2.4GHz machine. When it is finished, the plot containing your recording should change – it may appear to be only two flat lines, until you use the "lightning bolt" AutoScale button of that plot to resize the data vertically. You can use the Format menu of the plot to access a "Time" configuration tab if you want to zoom in and look at your IR set.

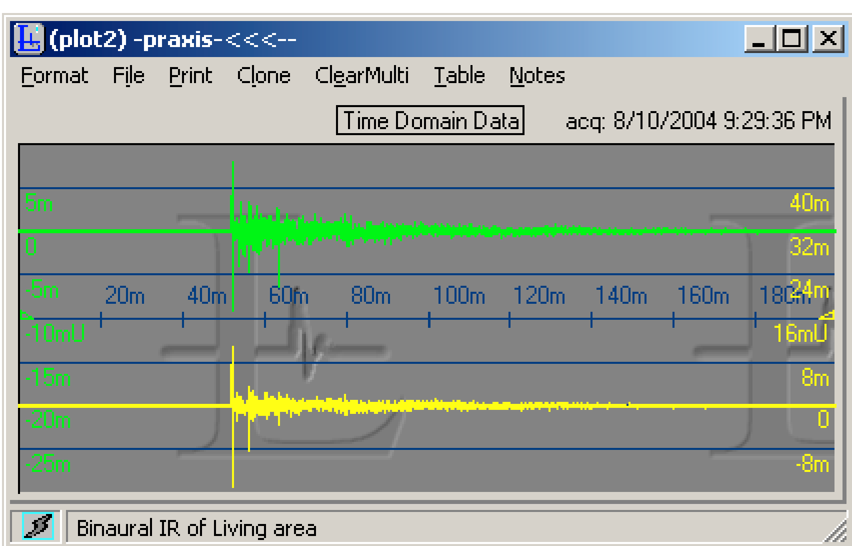

- It is a good idea to save this IR at this time. Use the “File, Save” menu to do so, and make sure to use a name that indicates the sample rate that you used, for future reference! You can also provide a description of the data (how it was made and what it represents) using the “Notes” menu of the plot before you save it. The “Notes” form also has a menu labeled “About this Data” which can be used to find some basic information such as size and sample rate.
- You may want to later use your IR with files made with other sample rates than you just used. If so, you can make a resampled version of the IR for these other rates. To do so, load a file that contains time data having the rate you want to use into the Primary Plot, and set the Primary Plot to be the PostProcess Source. Make the Plot containing your IR the Target, and use the “Resample Targ to Src rate” setting of the “Time Dom. Math” PostProcess to do the needed sample rate conversion. Then save the resampled data file under an appropriate filename.
- The rest is easy. Now you are ready to use your binaural impulse response to “play back” the room sound.

II. <u>Processing recorded sound files with the IR</u>

Description: In the steps above, you took a measurement of what a room and sound system did to a wave file of a sweep burst, and used PRAXIS to deconvolve it, obtaining an IR for the room at the listener position. Now, we will do the opposite process and convolve some recordings with that IR to impress that room’s (and sound system’s) sound onto them.

What you need:
- Some “dry” WAV file recordings for which you want to hear the room sound applied. These need to be at the same sample rate as the IR file that you generated (if they are not, go back and resample your IR to make a version for this new sample rate, as described above). A “dry” recording is one that has little or no “room sound” of its own. If a recording containing room effects is used, the end result will be that of one room’s sound played through another, which is not a very natural situation! Ideally, these files should be mono, but stereo files can also be used (or the PRAXIS “Time Dom. Channels” PostProcess can be used to make mono versions from either channel). The files should be no longer than necessary, perhaps 5 or 10 seconds.
- The PRAXIS software, which you can again use in Free Mode for these purposes.

How to do it:
- Start PRAXIS, and load your IR into the Primary Plot using its “File, Open” menu.
- Open another Plot within PRAXIS. There are several ways to do this:
    - Use the “View, *New Secondary Plot* “ menu of the Main Form to create a new Plot form.
    - Or use the “Clone” button of an existing Plot form to make a duplicate of the plot and the data it contains.
- In the new Plot Form, load your “dry” WAV file, using the “File, Open” menu, and selecting “WAV” file types, to browse to the folder where you stored your WAV files.
- In the PRAXIS Main Form, find the section to the right labeled “PostProcess”, and in its drop-down box select “Time Dom. Math”. Select the Primary Plot to be the Source and your new secondary plot (probably “Plot 2”) to be the Target. In the PostProcess configuration form, select “Convolve Source, Target”, and for Gain select “Specify dB” and “0dB” for both channels.
- Then click “Do PostProcess” and wait until the process finishes and the Target plot updates (the time required will depend on the length of your files).
- The resulting file can be saved, if desired. To save to the PRAXIS file format (which allows Notes and formatting information to be saved with it), use the “File, Save” menu for the Plot. To save it as a WAV file, without Notes or formatting, use the “File, Save to WAV file” menu.

III. <u>Listening to the Processed time files</u>

You will most likely want to listen to the result of your room sound convolution using headphones. Otherwise, the sound of the room you are in can alter the sound. But fairly good results are also possible using near-field monitor loudspeakers.

There are several ways to do this. If you save the file to a WAV file format, you can use any number of player applications (including Windows' "Media Player") to play the file back through your soundcard and headphones.

A more convenient way that allows you to stay within the same application is to use PRAXIS for playback as well. As you've probably already noticed, PRAXIS can make a number of Plots forms available for you to work with. Whenever a PRAXIS Plot form contains time domain data (such as recordings or Impulse Responses), there is a menu option "File, Listen" that becomes available for that Plot. This menu option opens a small mini-player for the data contained in that time domain file.

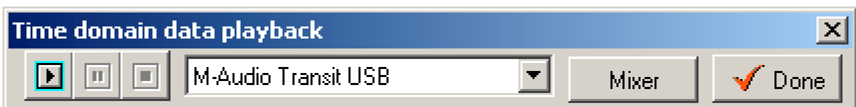

This player lets you:

- Select the soundcard device you want to use to play back. If you use the same soundcard device that PRAXIS is configured to use (or the only device in the system!), you will be asked to verify this.
- Open the Windows Mixer for the selected soundcard. You may need to use this to adjust the volume or to enable the WAV playback for the soundcard.
- Start playing the file
- Pause the playback
- Stop the playback.
- The "Done" button closes the player. You cannot work with any of the other PRAXIS forms until this player is closed

IV. <u>Continuing further</u>

You can also use the Impulse Responses obtained this way to do other kinds of analysis with PRAXIS, including loudspeaker quasi-anechoic frequency response, room frequency response (best if smoothing is applied) and room parameters such as RT60, etc. While these are being left "as an exercise for the reader", PRAXIS does contain an extensive HELP file, and there is a free downloadable manual in pdf format which can help you perform these operations.

<u>Acknowledgements</u>

Thanks to Curtis List for suggesting this as a PRAXIS feature and to Pat Brown of SynAudCon for his work on this technique and his comments and review of this paper.