

Interactive Reverberation Modelling System

Dr Jez Wells – Department of Music, University of York

Introduction

The Interactive Reverberation Modelling System (IRMS) is an artificial reverberator that combines the ‘photo realism’ of sampled acoustic impulse responses (AIRs) with the extensive parametric control offered by algorithmic simulators. To date no reverberation system has combined AIRs with such extensive and plausible control of parameters.

It was originally developed for an audio-visual installation work presented at the National Centre for Early Music, commissioned by Vespertine. The audio component of this work, solely devised and produced by the author of this software, used AIR measurements taken at different positions within the church, processed by the IRMS (amongst other techniques developed by the author) and distributed over 17 loudspeakers.

The system models AIRs as sets of atoms which are decaying sinusoids, each with their own frequency, amplitude, phase and decay rate. For very high modelling accuracy (often perceptually indistinguishable from the original sound) a novel combination of spectral modelling and matching pursuits, devised by the author, is used to derive the parameters of each atom. The system offers control over reverberation time, room size and density, parameters that are typically confined to algorithmic reverbs or are offered in a much cruder form. Reverberation times can be increased to 20 times their original length, without acoustically implausible artefacts.

The demonstrator downloaded with this document implements the modelling and interaction of stereo AIRs. It demonstrates how AIRs can be modelled and then interacted with, however it only permits audition (not saving) of AIR transformation. A commercial plugin version is currently in development and the technology is now being deployed to other areas of sound design in a Knowledge Transfer Partnership. The system has been used on a recording engineered and edited by the author, released on Naxos¹. For the entire recording the system is used subtly to enhance the quality of reverberation, but is used more overtly in *Cippyn*, co-produced by the author, to completely transform the space in which the music situated.

IRMS demonstrator

The reverb modelling system was devised, implemented and used by the author within Matlab. A standalone demonstration system with a graphic user interface has been created to allow the capabilities of the system to be investigated and auditioned. The modelling part of system (which analyses the AIRs) is computationally intensive, however the ‘modelled pursuits’ analysis method devised by the author is significantly faster than other high-resolution methods (e.g. ESPRIT) for the estimation of damped sinusoids whilst offering similar quality. The ‘interactor’ part of system (which allows previously created models to be loaded, modified and auditioned) does not operate in real-time because of the underlying architecture used to produce the demonstrator, however a ‘comp. time’ box does show how quickly the AIR is recalculated from the model and, for many AIRs and settings, this can be seen to be lower than the length of the AIR (and so has the potential for ‘within real-time’ operation, e.g. within a DAW plugin).

¹ *A Hundred Years of British Piano Miniatures* – Duncan Honeybourne (piano). Details of recording and sleeve notes etc. available at https://www.naxos.com/catalogue/item.asp?item_code=GP789 Available for streaming via Spotify (<https://open.spotify.com/album/0PwC5fdWWy7N6uZZyrEkiy>).

The demonstrator comes in two parts: the ‘reverb modeller’ and the ‘reverb interactor’. The installation packages for each of these are large (just less than 1 GB) because they include the Matlab runtime environment that enables them to run outside of the Matlab application. Two brief guides to their use follow. Any difficulties with installation and/or operation should be reported to the author (jez.wells@york.ac.uk).

Reverb modeller

The interface for the application is shown below. The top left plot is the time-domain waveform of the input AIR, the top right plot is the waveform of the modelled audio. The bottom right shows how the modelling algorithm has progressed: the horizontal axis is the number of modelling iterations (equivalent to the number of components in the final model) and the vertical axis shows the reduction, in dB, in the residual energy (i.e. the energy left in the original signal after the individual components have been extracted). The modelling algorithm (provided the analysis method is selected to be ‘iterative’) continues to iterate all the while the residual energy is being reduced or until the number of model components is equal to $N/4$, where N is the number of samples in the AIR. The reasoning behind this stopping limit is a rule of thumb that the amount of data required by the model should not exceed that required by the original AIR, and there are four parameters to each component of the model. Sometimes there will be an increase in residual energy: at this point the analysis discards the most recent component (the one that caused the increase in residual energy) and halts.




Figure 1: reverbModeller user interface

An example use of the system is as follows:

1. Launch the modeller.
2. Press ‘IR Open’.
3. Navigate to the location of the ‘French 18th Century Salon.wav’ file that is included with the demonstrator (this file is from the Voxengo collection, which is made freely available at <https://www.voxengo.com/impulses/>) and open this file
4. After a few seconds this AIR’s waveform will be seen in the top-left window and can be auditioned by pressing the ‘Play’ button on the left-hand side of the screen.
5. Press the ‘Trim’ button. This removes any silence from the start of the AIR and estimates its reverberation time (RT60). This must be done before the next step.

6. Keeping the default parameters that are set in the bottom-left of the screen, press the 'Analyse' button. This begins the analysis process which can take a few minutes (a pre-modelled version of this AIR, for immediate use the 'interactor' part of the system, is also provided to save time).
7. Once the modelling is complete the waveform of the resynthesized AIR will appear in the top-right display. The bottom-right display is for debugging/analysis purposes only but shows how the iterative analysis process has reduced in the energy in the difference between input and the modelled versions of the signal (the lower the energy in the difference between them, the more similar they are).
8. The modelled version AIR can be auditioned by pressing the play button on the right-hand side of the screen (for comparison, the original AIR can still be heard by pressing the left-hand play button).
9. Before exiting the modeller, make sure the model has been saved by pressing the 'Model Save' button.

Reverb interactor

This part of the system enables models to be loaded, modified and auditioned as well as being applied to a sample sound that the user can load. The user can load an AIR model (.irm extension), with the model audio (resynthesized AIR) time-domain waveform displayed in the top left plot. A sound to reverberate with the AIR can be loaded and displayed in the top right plot. It is best to use very dry or completely anechoic sounds for auditioning (anechoic files can be freely obtained from a variety of sources, such as <https://users.aalto.fi/~ktlokki/Sinfrec/sinfrec.html>). This can be heard without any reverberation applied ('Play sound' button) or with reverberation ('Convolve!'). There is a slider to enable the relative amounts of dry and convolved sound to be adjusted. The relative level of early reflections to late reverberation ratio can also be controlled (although this is only relevant if early reflections have been separately detected in the original modelling process – by checking the 'ER separation' radio button in reverbModeller). The bottom left plot shows where in frequency the room modes are in the model and how much energy they contain. The bottom right plot shows modal energy plotted into descending order and is for analysis and debugging purposes only.

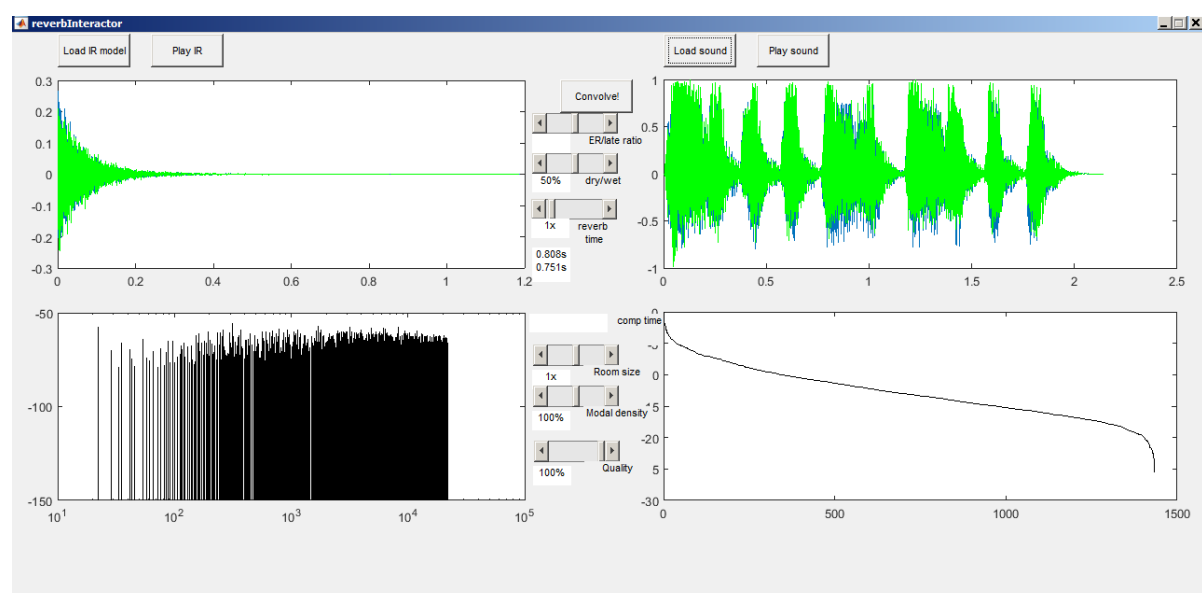


Figure 2: reverb Interactor user interface

There are three sliders that are designed to change the AIR in similar ways to controls found on a typical algorithmic reverberator:

Reverb time

This adjusts the decay time of the reverberation time. The slider sets a multiplier of the original decay rate, between a range of 0.2 and 20x. To take account of air absorption effects, high frequency components decay relatively faster at longer reverb times (greater than x1) and relatively slower at shorter reverb times (less than x1). The RT60 is displayed in the small box just below the reverb time slider. It should be noted that setting a reverb time multiplier of, say, 2 will not necessarily produce an exact doubling of the reverb time due to the difference in behaviour at higher frequencies and it just offers an approximation of how the reverb time will be changed.

Room size

As rooms become larger so the wavelengths that can fit between surfaces become greater, changing the distribution of modes and stretching it downwards. Increasing the room size usually gives rise to a warmer and more evenly distributed sound.

Modal density

As room geometry become more complex or surface materials cause greater diffusion, so the apparent modal density increases. A greater modal density creates a less sparse, 'ringy' sound.

Quality

This enables a trade-off between the accuracy and computational cost of the synthesis model. It works on the basis that once a sinusoidal component has fallen to a certain amplitude it no longer needs to be synthesized. This slider effectively controls the level of this amplitude threshold. Very high settings can lead to longer computation times than are necessary. Very low settings can lead to audibly truncated IRs being synthesized.

Model synthesis

Whenever any of these four sliders (reverb time, room size, modal density, quality) are adjusted the AIR is completely resynthesized from the updated model parameters. Whenever the AIR is resynthesized the time taken per channel for the resynthesis is shown in the 'comp. time' box. This resynthesis is now faster than real-time for all tested short and medium term and for many long AIRs. (Please note, the time displayed in the comp.time box and the actual wait to access the resynthesized AIR may be quite different, any additional delay is due to the demonstrator user interface and is not due to model computation).

An example use of the interactor is as follows:

1. Click the 'Load model IR' button and select the French 18th Century model (.irm file).
2. Audition the modelled AIR by pressing the 'Play IR' button.
3. Load a sound to audition it with by pressing the 'Load sound' button and selecting a suitable .wav file. This sound can be auditioned by pressing the 'Play sound' button.
4. Play the loaded sound through the AIR model by pressing the 'Convolve!' button.

5. Adjust the 'reverb time', 'room size' and modal density parameters to taste, and audition the results using the 'Play IR' or 'Convolve!' button.

March 2019