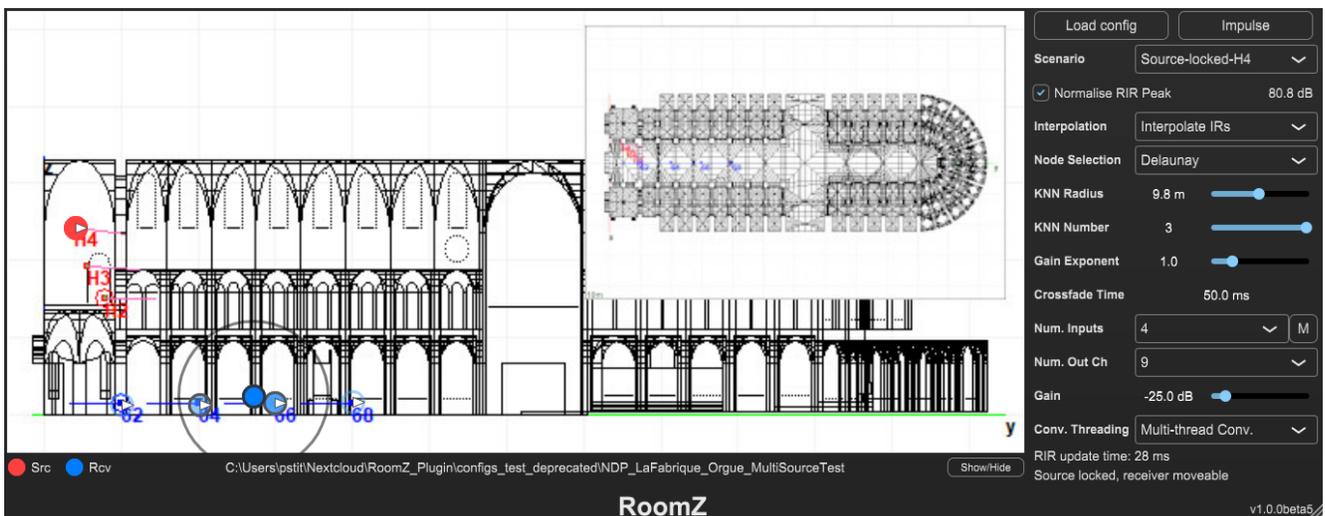


RoomZ Plugin User Manual



25 JANUARY 2023

Authored by: Peter Stitt (SSA Plugins)

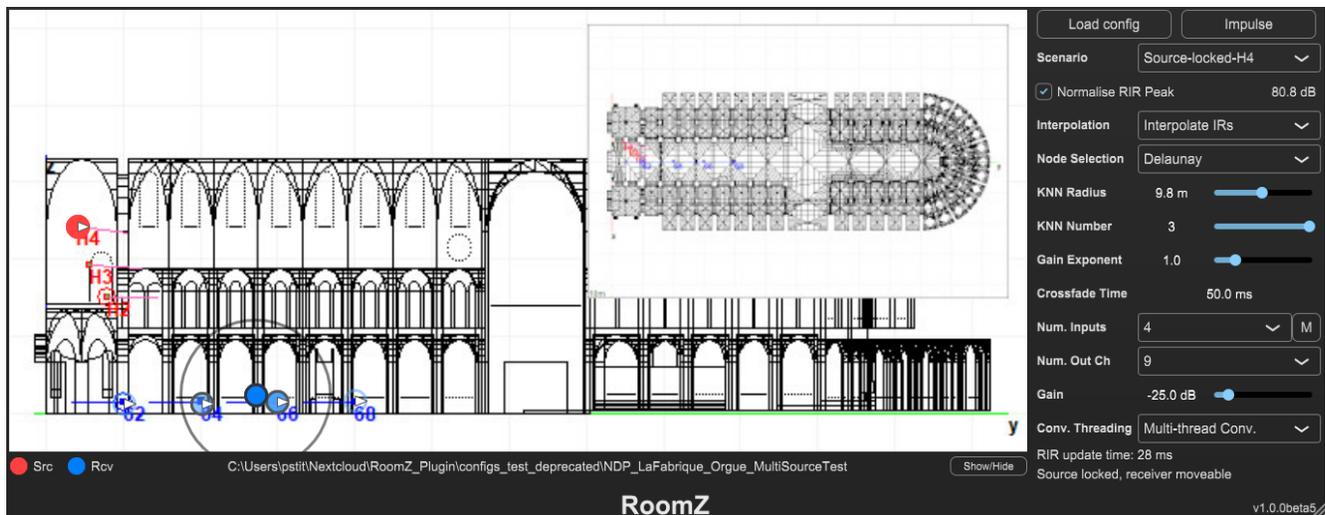
RoomZ Plugin

RoomZ is an auralization plugin that allows for a sound source or a listener to be moved around a grid of room impulse responses (RIRs) in real-time. Each configuration (or “room”) can have a number of “scenarios” that define the grid and the RIRs to be used for the auralization.

RoomZ allows for multi-channel input for the auralization of multiple sources at a time. If a scenario is defined as having fixed source positions, then RIRs between each source and grid point are required. If a scenario has a fixed receiver position, then the individual sources can be moved in real-time around the receiver.

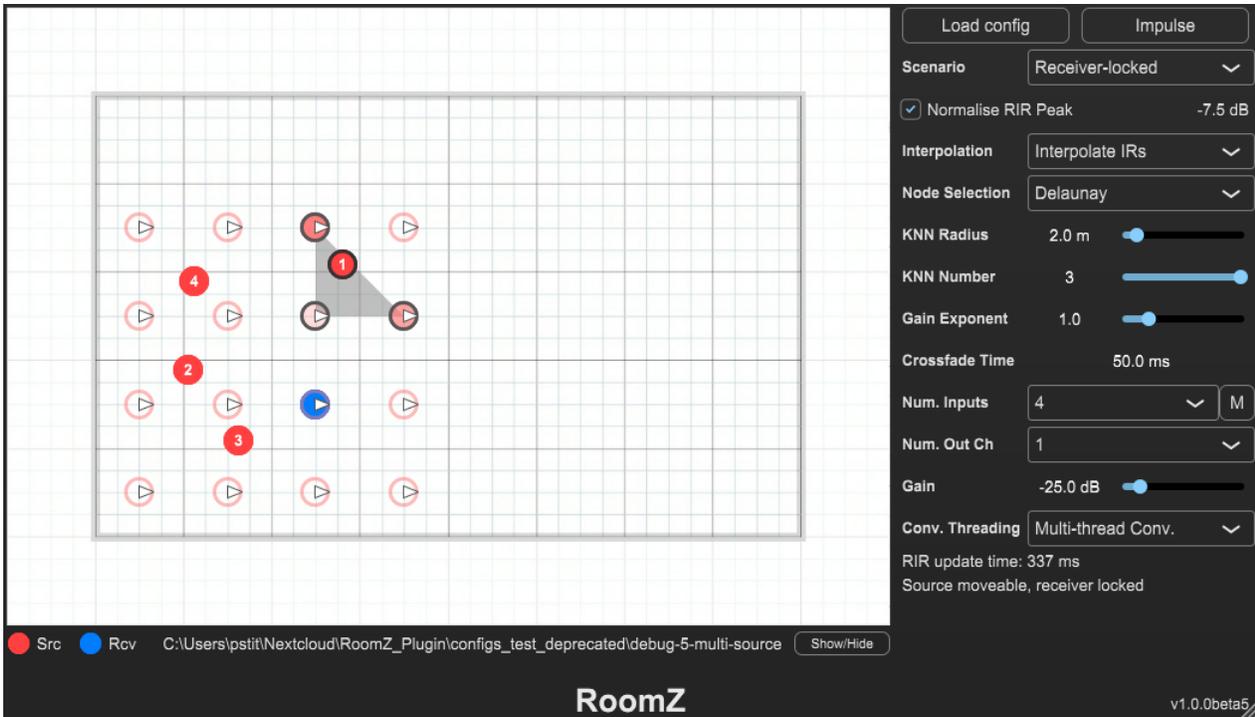
RIRs can be loaded for any spatial audio format up to 64-channels. The plugin itself is agnostic of the format of the impulse responses. The plugin supports a maximum of 8 inputs that can be auralized at any one time.

User Interface



The user interface consists of two regions. To the left is the map of the room that has been loaded. Sound sources are indicated by a red circle with a white triangle showing the orientation of the source. Receivers are blue circles, also with a white triangle indicating the orientation. The path to the loaded configuration is displayed below the map image.

To the right of the interface are the controls used to control the auralization. More detail is given on each of the individual controls later. This area can be hidden by clicking the Show/Hide button to the bottom-right of the map.



The above image shows the plugin in a fixed-receiver configuration with 4 sound sources active. The numbers on each of the source (red circles) is displayed. The grey triangle indicates the grid triangle of the last moved source (source-1, in this case). The strength of the contribution of the grid points used for the interpolation of the source is indicated by their transparency.

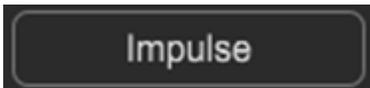
Loading a Configuration

The first thing you will need to do once you have loaded RoomZ into your DAW is to load a configuration file. The configuration file contains information about the layout of the sources or receivers in a sound scene. A configuration can be loaded in one of two ways:

- Click the “Load config” button towards the top right of the plugin user interface. This will open a search window. Select either the folder containing your configuration or the config.xml file itself. Confirm your choice and the configuration will load.
- Drag the folder containing your configuration or the config.xml from the operating system file viewer on to the user interface of the plugin.

User Controls

RoomZ has a large number of user controls that can affect the auralization of the sound scene and the efficiency of the plugin. The default settings aim to give good results for most use cases, but you might find that you have better results with other settings. This will particularly depend on the arrangement of the sources and receivers in your configuration.



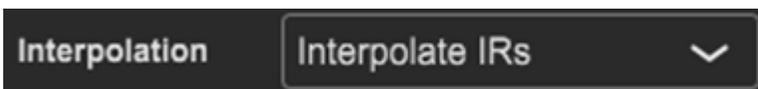
Impulse: this will cause a single impulse to be injected into the first channel of the inputs. This is like clapping in the room being auralized.



Scenario: Your configuration might contain multiple scenarios. These are different arrangements of sources and receivers. The dropdown Scenario menu allows you to select between them. By default, the first scenario listed in config.xml will be loaded after a new configuration is loaded.



Normalise RIR Peak: When this is selected the RIRs will have the peak value across all RIR files and channels normalize to -10dB i.e. all files are scaled by the same amount to ensure that the loudest sample peak in all RIRs is no louder than -10dB. If this is not selected, then no gain is applied to the RIRs. Changing this parameter will cause the scenario to be reloaded. When activated, the gain applied to the RIRs is displayed in decibels.

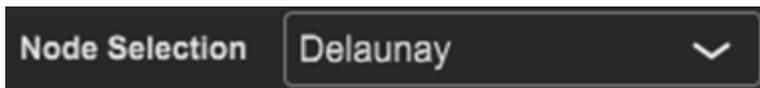


Interpolation: This controls how the interpolation between different points in the grid is handled when auralizing them.

- **Interpolated Convolution Buffers:** This uses a convolution engine for each of the currently active grid points. The outputs of each convolution engine are then

summed together with the appropriate weights of each grid required for the interpolation. This method is more CPU intensive but has lower latency update of the spatial information in the signal.

- **Interpolated IRs:** This method takes the active grid points and creates a RIR based on their weights. This composite RIR is then loaded to a single convolution engine which is used for auralization. This method uses less CPU but has higher spatial latency since a new RIR must be recalculated and loaded for every movement of the source/receiver.

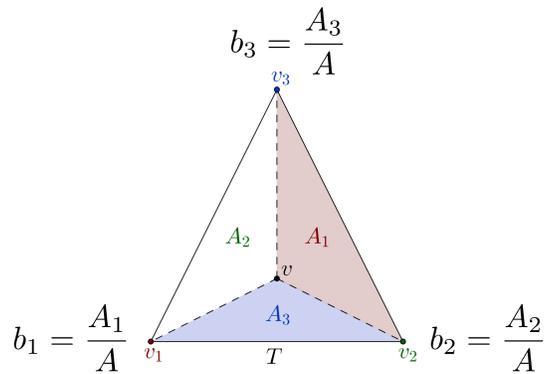


Node Selection: This parameter controls how the nodes in the grid are selected to use for the interpolation.

- **Delaunay:** If the grid in the scenario is somewhat regularly laid out then the plugin will attempt to split it into triangles. If this mode is selected, then the triangle that contains the source/receiver position will have its vertices used for the interpolation. If this mode is selected but the source/receiver position is outside of a valid triangle, then the KNN method will be used instead.
- **KNN:** This method selects the K (from 1 to 3) nearest grid nodes and applies weights for the interpolation based on the proximity to the source/receiver position.

Note on interpolation methods gains

- **Delaunay:** If the node selection is set to “Delaunay”, gain values for the three selected RIRs will be computed using barycentric coordinates. The figure below illustrates such computation, where b_i represents the gain applied to RIR at position v_i . This method ensures smooth transition between neighboring triangles, as the gain applied to a triangle summit / RIR will always decrease to 0 before it is discarded.



- **KNN:** If the node selection is set to “KNN”, gain values for the three selected RIRs will be computed based on the inverse of their distance to the current source/receiver position.

KNN Radius 9.8 m

KNN Radius: This is maximum distance that a grid node can be from the source/receiver position and be used for the interpolation. Note that if your radius is too small for your grid you might end up selecting no grid points and the plugin will output silence.

KNN Number 3

KNN Number: The maximum number of grid nodes that will be used by the KNN algorithm. For example, setting this to 1 will always use only the nearest grid point (assuming the KNN Radius is large enough to find at least one).

Gain Exponent 1.0

Gain Exponent: This applies an exponent to the gains of the nodes used for the interpolation when they are being calculated. Effectively, this governs the weight of nearer grid points vs further ones. A value of less than 1 will mean that all active grid nodes have more similar gain values. The extreme value of 0 means that all selected active nodes will have the same gain. For values greater than 1 the closest grid node will be given more weight, essentially pulling the source/receiver towards that grid node position.



Crossfade Time: This controls how long it takes for the convolution engines to change from one RIR to a newly loaded one. Longer values will lead to a smoother sound but greater spatial latency. Shorter values will increase spatial reactivity of the plugin at the expense of possible zipper noises while moving the source/receiver. The time can be changed by typing directly in the box. Alternatively, click-and-hold the numerical display and drag the mouse vertically to increase/decrease the value.



Num. Inputs: Select the number of input channels that can be used. In a receiver-locked scenario this will govern the number of moveable sources there are in the sound scene. In a source-locked with multiple defined source positions this will govern the number of channels that are used. In this case you will most likely want to set this to the number of sources in the scenario.

- The “M” button beside the Num. Inputs dropdown opens an overlay menu that lets you select which input channels (if any) you want to mute. If any channels have been muted, then the button will indicate this by becoming red.



Num. Out Ch: This sets the number of channels in the RIR are actually auralized. For example, if you load a third-order Ambisonics RIR you can select to render only first order by setting the value of this parameter to 4.



Gain: Applies a gain to all input channels **before** they are sent to the convolution engines.



Conv. Threading: This selects whether work for the convolution engines is done in the main plugin audio thread or if some are run on background processes.

- **Single Thread Conv:** All convolution processes are run on the audio thread. This is the safest option but for long RIRs or those with many channels this could lead to CPU overs, resulting in audible clicks in the audio. If your audio is running without glitches then this should be your preferred option.
- **Multi-Thread Conv:** This option should be selected for long RIRs in particular. It is also useful if rendering a large number of input sources, or multiple sources with a large number of RIR channels. Use this option if you are hearing glitches.



RIR Update Time: Displays how long it took to update the RIR due to the last movement of the source/receiver position. This is only provided for information so you can see if the spatial latency is becoming unacceptable when rendering long RIRs.



Show/Hide: This button can be used to hide the user controls, maximizing the space available for the source/receiver movement map.