3D Audio

Welcome to 3D Audio !

This program enables you to create true three-dimensional sound experiences. For an introduction both on the topic of spatial perception and this program, select the Quick Start.

The program is working with Windows95 and WindowsNT 4.0.

For best visualisation results you will need DirectX3 be installed on your system. You can obtain a free copy of the runtime binaries from MS via WWW:

http://www.microsoft.com/DIRECTX/RESOURCES/downloads/Dx3busa.exe

Note that this file doesn't work with Windows NT 4.0! When using NT4.0, you need Service Pack 3 installed on your system; it includes the DirectX3 drivers. You may download from MS:

http://www.microsoft.com/isapi/support/bldqpage.idc?ProductPage=q_servpk

If you use the program with Windows NT 4.0, please read the NT 4.0 technical note!

This program is shareware and may be used for seven days for evaluation; if this period is exceeded it will stop working. If you like to unlock your copy you must register. If you want to distribute the program have a look at the distribution status.

You may download a demonstration project (pingpong.zip) for 3D-Audio from our WWW-Site:

http://www.audiophile.com/climax

Contact the authors

You may contact us at :

E-Mail	:	climax@audiophile.com
WWW	:	http://www.audiophile.com/climax
Mail	:	Climax Software Solutions PO Box 210416 10504 Berlin Germany

Name : _____
Company : _____
Address : _____
_____
_____
_____
E-mail : _____

Your preferred username
for registering 3D Audio : _____

Payment method

[ ] Cash
[ ] Eurocheque in German marks.
[ ] Direct deposit
[ ] Credit card : [ ] VISA [ ] Diners [ ] AMEX
Credit card no. : _____
Expiration date : _____
Name on card : _____
Billing address : _____

I understand the license terms completely. I understand that improper payment will not be returned. Users registering the beta version 0.99 will receive a free copy of the golden release.

Date : _____

Signature : _____

License terms

The received password may be used by one single user only. The license is confined to private use only; for commercial or institutional use, contact the authors.

Disclaimer of warranty

THIS SOFTWARE AND THE ACCOMPANYING FILES ARE SOLD "AS IS" AND WITHOUT WARRANTIES AS TO PERFORMANCE OF MERCHANTABILITY OR ANY OTHER WARRANTIES WHETHER EXPRESSED OR IMPLIED. NO WARRANTY OF FITNESS FOR A PARTICULAR PURPOSE IS OFFERED. IN NO EVENT WILL THE AUTHORS BE LIABLE FOR SPECIAL, DIRECT, INDIRECT OR CONSEQUENTIAL DAMAGES, LOSSES, COSTS, CHARGES, CLAIMS, DEMANDS OR CLAIM FOR LOST PROFITS, FEES OR EXPENSES OF ANY NATURE OR KIND.

The user must assume the entire risk of using the program.

Choose the fully qualified path for definition files. Provided for compatibility; need not be set.

Press to select the definition file path by means of a file dialog.

Choose the name of the PCM-file to write the computed data to.

Select the output file by means of a file dialog.

Select the monaural 16-Bit big-endian PCM-File to use as sound data for the current source.

Select the PCM-File to use for this source by means of a file dialog.

Specify the .HTF-file containing the head related transfer function as impulse responses to use.

Select the .HTF-file to use by means of a file dialog.

Use the internal set of head responses instead of specifying an external file.

Specify the samplerate of both the impulse data and the source sound files.

Give the length of the internal synthesis buffer. Must be larger than the maximum delay occurring in the scene.

The length of one atomic computation block in samples. Positions and Doppler-shifts are updated every block.

Specify the time that this source is to start relative to the beginning of the simulation.

Determine the power of the sound assumed for this source. Allows you to adjust the relative volume of the sources.

The path-loss exponent for this source. For free space this would be -2.0, but smaller values may yield more realistic results.

The order of reflections to compute for this source. Zero means only the direct beam is computed, one means reflections over one wall are computed etc.

This value allows you to change the radiation characteristic of the source, which is given as function of the angle between the direction of emission and the direction of the orientation of the source:

Exp[- EmissionFactor * Angle]

A value of zero yields a constant radiation pattern; values above concentrate the radiation to the direction of orientation.

Define the length that this source is to run in seconds. If the value exceeds the length of the input file, looping is performed.

Check this to cause the length the source runs to be exactly that of the input file.

Define the size of the box-like room that the scenery is to take place in. Values are given in meters.

The size in X-direction of the box-like room that the scenery is to take place in given in meters. This value corresponds to the width of the room.

The size in Y-direction of the box-like room that the scenery is to take place in given in meters. This value corresponds to the depth of the room.

The size in Z-direction of the box-like room that the scenery is to take place in given in meters. This value corresponds to the height of the room.

Define the reflection properties of the virtual room. The given values specify the reflection coefficients:
a value of e.g. 0.7 means that reflected waves are attenuated by 30%.

Check this to apply a modified computation, which attempts to incorporate the impact of the distance of the source on the used head related transfer function. Quite experimental parameter, which doesn't yield a big improvement.

Motion dialog

This dialog allows you to define atomic steps of continous motion. By creating of sequence of such motions, virtually any complex motion can be implemented.

This node contains a list of elementary motion instructions to define the path the source or head will follow during simulation.

The angular velocity for turning motion given in rad/sec.

The time that the desired motion is to last.

When checked the desired motion is performed in one single step. Use this to initially position the sources.

When checked the orientation of the source will be directed towards the virtual head during the motion.

Define the end-point of the motion giving its X, Y, Z - coordinates.

Define an absolute translation, i.e. a movement to a given point. The specified point will be reached after the given motion time.

Define the speed of the motion as velocities in X, Y, Z - coordinates.

Define the the acceleration of the motion in X, Y, Z - coordinates.

Define a relative translation, i.e. a motion with given speed and/or acceleration.

Define the angle that the orientation of the source is to be turned to. The angles are given as "phi" (angle between direction and the XY-plane) and "theta" (angle between the X-axis and the direction measured counterclockwise).

The angular speed of the rotation of the direction. The angles are given as "phi" (angle between direction and the XY-plane) and "theta" (angle between the X-axis and the direction measured counterclockwise).

The angular acceleration of the rotation of the direction. The angles are given as "theta" (angle between direction and the XY-plane) and "phi" (angle between the X-axis and the direction measured counterclockwise).

The ciruclar motion of turning will take place on a plane in the three-dimensional space. The point defined by this (X,Y,Z) vector specifies the center of that circle when projected to the plane of motion.

The ciruclar motion of turning will take place on a plane in the three-dimensional space. This (X,Y,Z) vector defines that plane in that it will be located perpendicular to the plane.

Turning means moving the source with given rate on a circular path perpendicular to a given axis. The latter is defined by a direction and a position vector.

Absolute rotation means turning of the direction of the source to a given absolute angle.

How to register ?

The price for registering 3D Audio is **DEM 50** (approximately US-$ 30) for a single-user license. You will receive a password which you can use to unlock your copy of the program. This license is limited to private use only! For commercial or institutional use, contact the authors.

To register, submit the completed registration form to us. Payment may be done by credit card. Users residing in Europe may alternatively send an eurocheque for DEM 50 or can directly transfer the money to a bank account; contact us for the bank account number.

As a last option you can send cash, but this is entirely at your own risk!

Enter your name as given when registering your copy.

Enter the password you received for your username.

The button for the greedy people...

Press this button to register your copy of 3D Audio after entering your name and the received password.

Press this button to view the terms for registering.

Select the way the Direct3D output is displayed. Possible are either in a window or fullscreen.

Check this to let the Direct3D output be displayed in a window.

Check this to let the Direct3D output be display fullscreen. The combo box to the right may be used to select the desired resolution and color depth.

Select the desired resolution and color depth for fullscreen display.

Define the render mode used by Direct3D.

Renderer output will be a wireframe model.

Renderer output is unlit flat mode.

Renderer output will be in flat mode.

Renderer output will be Gouraud mode; this looks best but is the most demanding in terms of computational complexity.

The mode of color emulation used by Direct3D.

Emulation mode will be "ramp". This is faster than RGB and looks reasonably good.

Emulation mode will be RGB. Very heavy load for the computer...

This window gives you an idea of how the selected Direct3D options will look like. If this display is active (click with the mouse on it), you can influence the movement of the displayed room by means of the num-locked keypad and <Pos1>, <End> and <Insert>. See the actual help file for a detailed list.

Here you may change the textures used for walls, floor and ceiling of the virtual room.

Select the wall that you want to change the texture of.

Brings up a file dialog where you may specify a .BMP-file to use as new texture for the selected wall. Note that the dimension of the picture must be a power of 2!

Check this to revert back to the default texture of the selected wall.

Select the object that you want to modify and define a new model for it.

Use the combo box to select the object that you want to modify by its name.

Brings up a file dialog which lets you specify a new Direct3D model (.X-file) to use for the selected object.

Check this to use the default model for the selected object.

Lets you change the color of any of the used Direct3D objects.

Displays the currently selected color of the current object.

Brings up a dialog to select a color.

Check this to revert back to the default color for the currently selected object.

Use this notebook to select that property of the object which you want to edit. You may modify the spatial position, orientation and scaling of the used Direct3D model to make it fit for later display purposes.

Specify a (X,Y,Z) vector to use as an offset for the currently selected object.

Specify a (X,Y,Z) vector with scaling factors that will be applied to the model.

Specify a (X,Y,Z) vector of angles that the Direct3D object will be turned prior to display.

Check this to revert back to default settings for the transform options.

Is there anything to say about About?

Press this button to access a window for entering your registration code.

Press this when you're bored enough by the window.

Tree window

This window displays the parameters and settings of the simulation in a tree-like structure. Dive into the structure and edit the nodes parameters by using the according context-menu (right mouse-button).

This node contains basic data concerning the simulation, such as the output file name, the samplerate of the data etc.

This node contains the definition of the virtual room that the scenery is to take place in.

This node contains information on the virtual head, i.e. its position and motion in time.

Contains basic information on the virtual head; actually just a flag which defines if the head is moved during the simulation.

Signals if the virtual head is animated and thus must be updated during the simulation. This flag cannot be changed by the user, but automatically changes with the definition of the "motion" node.

This node contains the data defining a sound source.

Contains basic information on the source, such as the PCM-file associated with it, the power etc.

Contains the definition of the motion of the source or the head during the simulation.

Signals if the source is animated and thus must be updated during the simulation. This flag cannot be changed by the user, but automatically changes with the definition of the "motion" node.

A dummy entry which signals the end of the list of sources.

Each node represents one motion instruction where the order of the statements gives the order of the complete, compound motion that will finally be performed. There are various types of motion supported, which are described in the actual help file.

The parameters of the parent motion instruction. The type of possible parameters depends on the type of motion selected. See the help file.

A dummy entry which signals the end of the list of motion instructions.

Basic properties dialog

This dialog lets you control the basic properties of the various objects, such as the room, the virtual head and the sound sources. The actual elements of the dialog depend on the type of object worked on.

Options menu

This menu contains several entries which let you control how the rendered scenery will be displayed on the screen. These submenu entries are available:

Use Direct3D	Check this to let the output be displayed using Direct3D. Naturally, Direct3D must be installed on your system.
Use wireframe	Check this to let the output be a simple wireframe model.
Direct3D options	Brings up a dialog to define a number of parameters when using the Direct3D display, such as desired emulation and rendering modes or the used models for displaying sources and the head.
Try to use Direct3D	Instruct the program to check for Direct3D. Invoke this operation if you installed Direct3D after 3DAudio.
Snapshot	Store the current position of the MDI-child windows.

Project menu

The project menu contains the items to invoke a simulation and control its run:

Compute	Start computing the current project. If another computation is in progress, it will be canceled.
Simulate	Start simulating the current project. This means that the scenery will be rendered and displayed, but no audio data are computed. This can be used to verify the motion sequences.
Priority	Enables you to adjust the priority of the thread computing the audio data.
Pause	Pause the current computation.
Resume	Resume the paused computation.
Stop	Stop the current computation.
Open outputfile	Open the PCM-file that contains the output data; this can be performed while the simulation is running.

Main menu

The contents of the main menu depend on the type of MDI-child that is currently active. These submenus are defined:

File The file menu contains items that allow you to load and store simulation project files as well as sound files. The structure of the menu and the meaning of the items should be well known.

Project (only available for project and render display window)

Audio (only available for audio-waveform display window)

Options

Window Now this should really be known to you!

Help Contains access to this help file as well as to the registration dialog, where you may enter the password to unlock your copy of 3D Audio.

Audio menu

The audio menu is available when a waveform window is the current MDI-child. It contains these entries:

- Change
- Play
- Pause
- Resume
- Stop
- Loop

Audio change menu

The entries of this menu allow you to change the parameters of the current waveform, such as samplerate, bits per sample and mono/stereo. Note that changing doesn't mean the data are changed, but only the assumed format! These parameters may be changed:

Samplerate Available are a number of usual rates.

Bits per sample Can be either 8 or 16.

Stereo Check this for stereo waveforms.

Visualisation window

This window displays the scenery you created. The rendering will either be done by Direct3D or it will be a simple wireframe model. In either case you can switch the viewpoint interactively:

Wireframe view

Press the left mouse button within the window and drag the mouse to change the view.

Direct3D view

This view doesn't support mouse dragging, but uses keys instead:

<PgUp>	Rotate the viewpoint around the center of the room by a small amount.
<PgDn>	Rotate the viewpoint around the center of the room by a small amount in the opposite direction.
<Pos1>	Start rotating the viewpoint around the origin of the coordinate-system.
<End>	Start rotating the viewpoint around the origin of the coordinate-system in the opposite direction.
<Delete>	"Freeze" viewpoint, i.e. stop moving, rotating etc.
<Insert>	Reset the viewpoint to the default.
<Csr-Keys>	Move the viewpoint in space.
<0>	Attach the viewpoint to the head.
<1..9>	Attach the viewpoint to source 1..9 if they exist.

When navigating within the preview window of the "Direct3D Options"-dialog, the cursor-keys are not available. Use the num-locked keypad instead.

Audio window

This window displays the waveform of a loaded audio file. When this type of window is the active MDI-child, the audio menu is available, which allows to play the sound etc.

Progress and verbose window

This window displays the current spatial position of the objects during computation. Additionally, progress bars indicate the current degree of completion for each of the sources.

During initialisation of the computation, the upper part of the window serves as log-display to dump error messages if any problems occurr.

Start playing the current waveform.

Pause playing the current waveform.

Resume playing the current waveform.

Stop playing the current waveform.

Select loop mode, when playing a waveform.

Quick start

3D Audio enables you to create sound files that convey the notion of true spatial placement of the sound's source. A prerequisite to make this possible is, however, that you listen to the created sounds over headphones, since using loudspeakers would imply the real acoustic scenery of the room you're in being mixed with the virtual one.

There are various influences which determine the way we localize sounds. The most important of these factors are empirically modeled by the program to allow the simulation of a sound arriving at the head from a given direction. Since in the real world we always are surrounded by objects reflecting the sound waves and thus providing more information to the ear, the simulation has to allow for the surroundings by means of introducing a virtual room, where the virtual head and an arbitrary number of sources, each associated with a monaural PCM-file, may be positioned. In this way the reflections from the walls can be incorporated, where you can specify the maximum order of reflections to allow for. And since static scenes are always a little boring, you may freely animate both the head and the sources! The animated scene can also be viewed as a rendered picture sequence if you have Direct3D installed on your system; if not, you still may view a simple wireframe animation. This visualisation is very useful to verify the motion instructions and can therefore also be perfomed without actually computing the audio output data.

To start, load one of the examples and examine the simulations parameters by diving into the tree-like hierarchy. Let it run with "Project / Compute", load the result by selecting "Project / Open outputfile" and listen to it. Then start changing a few parameters to get a feeling for them. The usage of the tree-structure should be intuitive; you can change the parameters by double-clicking on them, or you may invoke the context-menu with the right mouse button. Some operations, like inserting new sources, are only possible via the context-menu.

The virtual room

The virtual room has a regular box-like shape. Its extent and the reflecting properties of each wall can be chosen arbtirarily. The figure illustrates the orientation of the room within the coordinate-system as well as the definition of the angles.

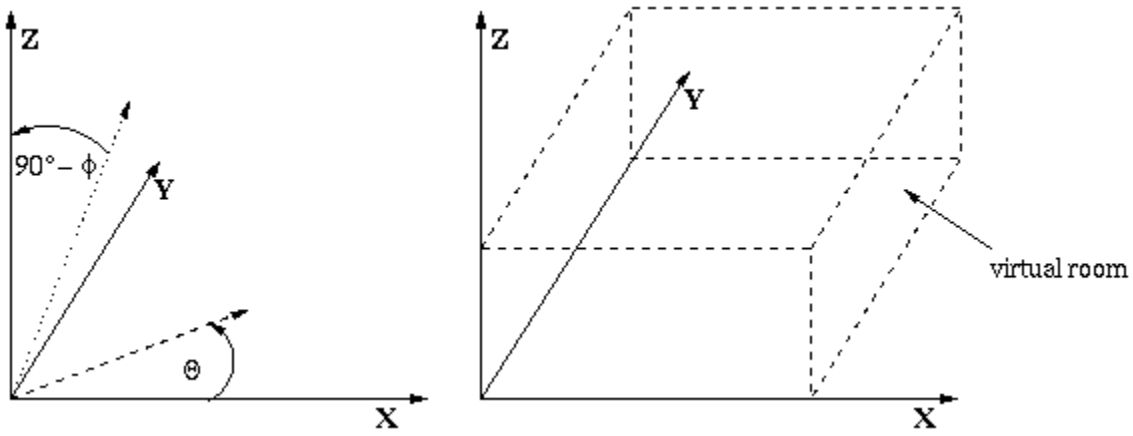

The reflection coefficients of the walls can be given in the range from 0.0 to 1.0, where a coefficient of 1 means total reflection a 0 total absorbtion. Thus you can omit a wall by setting its coefficient to zero. To simulate a surrounding with just a floor, set all coefficients except that of the floor to zero.

Since the programs time to finish the calculation should be finite, you can specify the reflection order to incorporate: an order of N means computing the line-of-sight wave and all waves reflected via at maximum N walls. Note that each reflection order introduces additional reflected waves from various directions, so that each of these waves has to be processed using a different impulse repsonse. Hence the simulation time dramatically increases with N! Using reflections of first order significantly increases the naturalness of the result; reflections of second order perform a little better, at the cost of about four times as much computational time. The highest order to use within finite time seems to be seven. Results really sound great, but the time for computing is about 5000 times the length of the source's sound on a modern PC...

Localization of sounds

There are a number of cues influencing and supporting our spatial perception of sounds. The actual process of spatial hearing is not yet fully understood, but these factors appear to be the most important cues:

- *Interaural delay time* , the delay between a sound reaching the closer and the farther ear. This delay is zero for sources directly behind, ahead or above and reaches its maximum of about 0.63ms for a sound from the far left or right. This quantity depends primarily on the direction of the sound, but also on distance and frequency.
- *Head shadow*, being the effect imposed on a sound having to pass through or around the head to reach the farther ear. The most striking effect is an attenuation, but there is also a dependency on distance and frequency of the source.
- *Frequency response of the outer ear*, which is significant for frequencies above 4kHz. It is supposed to be important for determining the direction of the sound.
- *Reflections from the shoulders*, which are present for frequencies below 3kHz. Also depending on the direction of the sound.
- *Head motion*, which is subconsciously performed to gather information by slightly moving the head to and fro. This naturally cannot be implemented in this program.
- *Vision*; we so strongly depend on out vision, that any auditory cues on a sources position is discarded when we see it.
- *Early echoes and reverberation*, which are primarily a characteristic of the surrounding.

The first four mentioned cues form the so called head-related transfer function. The impact of these cues can be empirically modeled as done by this program.

Parts of the enumeration above have been taken from:

D.A. Burgess, "Techniques for Low Cost Spatial Audio", Tech Report, Georgia Institute of Technology.

The programs model of spatial perception

The most important influences of spatial perception of acoustic signals are attributed the term head-related transfer functions. These are modeled empirically within this program, represented by a set of impulse responses which have been measured by a microphone placed inside a dummy head. Since these responses strongly depend on the direction of the sound source relative to the head, a set of responses has been measured to cover most of the range of possible directions. Applying these data to a sound file, i.e. convoluting an impulse response corresponding to a certain direction with the sound data, yields a vague impression of spatial positioning of the sound. This vagueness is due to the fact that the response measurements took place inside an anechoic chamber; if you've ever been in such a place you know that acoustic perception is very confusing and odd in there. Therefore a virtual room is introduced to allow for signal reflections from the walls. These reflections supply both additional information for localisation and the substance for a more realistic impression.

Animation of the objects

Each of the, i.e. the virtual head and the sources, may be given a set of instructions to describe a motion to be performed during simulation. These movements concern both the position and the orientation of the objects. The initial state of each object is a position in the origin of the coordinate-system with an orientation along the X-axis. Any of the movements is given a length as a period of time that the motion is to last. If the time is zero, the motion will be performed in one atomic step; this option is useful for initial positioning of the objects.The following movement instructions are supported:

- Absolute translation, which means specifying an end-point that will be reached after the given time by the object by moving along a straight line.
- Relative translation, which means giving a speed and optionally an acceleration. The object will move according to these values for the given time.
- Absolute rotation, which means giving an end-orientation for the object. This orientation will be reached after the given time. Note that rotation implies no change of position!
- Relative rotation, meaning changing the orientation of the object with a given angular speed and acceleration. Note that rotation implies no change of position!
- Turning, which means animating the object on a circular path, located in plane which is defined by an axis perpendicular to it. The center of the circle is given by a vector by projecting it on the mentioned plane.

When applied to a source all possibilities, except instructions on rotation, can be chosen such that the orientation of a source will always point towards the virtual head during the motion. Normally, a translation leaves the orientation unchanged. Changing of the orientation naturally only affects the simulation if the source has a non-uniform radiation pattern. This can be controlled by means of the source's emission factor.

Note that with speed too high, more precisely when approaching the speed of sound, the computation may get somewhat unphysical!

NT 4.0 technical note

- There seems to be a problem with the screensaver and the Direct3D output in Window NT 4.0. If you've installed the service pack 3, make sure to turn off the screensaver. If you do not turn it off, 3D-Audio will lock if the screensaver appears!

Distribution status

You may freely distribute the unregistered version of 3D-Audio, provided that all the files are included and are unmodified and that no files have been added to the package. Please distribute it by copying the original zip file. You may not ask any money for the distribution.

If you want to put this component on a freeware/shareware CD-ROM or other compilation, please contact us before doing so, to be sure that you are not including old or incomplete stuff in the compilation. We would also appreciate it if we could get a free copy of the CD-ROM.