

Introduction

This manual will describe the background to and technical features of the data-sonification application “STEREO-Spectro”. Designed as a multipurpose tool, this application may be used by teachers wanting to create an enthusiasm for understanding data; scientists looking for alternative representations of data; musicians interested in new sounds and algorithmic models; and others simply intrigued by space, sound, or both.

For general information on the STEREO/Impact “Sounds of Space” project, please visit our webpage at:

<http://cse.ssl.berkeley.edu/impact/vos/welcome.html>

Specific information and downloads concerning this application may be found at:

<http://cse.ssl.berkeley.edu/impact/vos/advanced.html>

Background Information

The Sounds of Space Project and STEREO-Spectro

The "Sounds of Space" project is being developed jointly by UC Berkeley's Space Science Laboratory (SSL) and Center for New Music and Audio Technology (CNMAT) to explore the connections between science and sound, to compare visual and aural representations of space data, and to promote a better understanding of the sun through stimulating interactive software.

Sounds are generated through real-time processing of data previously collected from the sun. All of the data is multichannel – meaning that several measurements are present for any given moment in time. We take this multiplicity of data and use it to control various parameters of the generated sound. A variation in one channel of the data might be used to control the pitch (how high or low), the amplitude (how loud or soft), or the rhythm for a portion of the overall sonic representation.

We have developed several computational tools to control, map, and sonify these data. The majority of our work has utilized the Max/MSP programming environment which allows for real-time user interaction with the sonification process. Through this interaction the user can become quickly engaged with the data and gain a depth of understanding not as easily attained through passive observation.

Our most recent work has relied on Spectrogram Image Analysis to be able to analyze data stored in diverse formats. By reading the changing color values represented in a spectrogram we are able to determine the number of channels and dynamically reconfigure the program to sonify accordingly. The STEREO-Spectro program utilizes this technology.

Sonification

Data is the outcome of a set of measurements or observations. After scientists place instruments on a spacecraft that flies out to space, these instruments can take measurements of things in space, such as charged particles. These measurements are recorded and sent to Earth. This information is *data* that can help scientists understand more about the space environment, such as the solar wind or regions of Earth's outer most magnetic field (magnetosphere).

Sonification is the act of representing data as sound. An example of sonification is using a wind chime to convert wind speed to music. The higher the wind speed is, the more notes are played in one second. We often use the mathematical verb "to map" to describe converting data to sound: "We mapped the wind speed to the number of notes played in one second by the wind chime."

Different types of sonification can be categorized in three ways*:

Iconic sonification: This type of sonification is when someone maps data to sounds that are associated with certain phenomena. For example, if we gathered weather data, such as cloud cover, temperature, and humidity, to calculate the probability of rain tomorrow, then using the sound of rain to indicate when there is a high probability of rain would be an iconic sonification.

Direct Conversion Sonification: This type of sonification is when someone maps data to sound to listen for patterns that are represented in the data. For example, space scientists map data of waves made up of magnetic and electric fields called electromagnetic waves to sound waves. This direct conversion sonification can be as simple as taking the frequencies of the waves and making sound waves with the same frequencies, which is most useful as long as the frequencies are at pitches that our ears can hear. Earth's whistler wave is such an electromagnetic wave that scientists have been sonifying for over 30 years.

Musical Sonification: This type of sonification is when someone maps data to sound in a musical way. For example, we have created a computer software program that will convert data of very fast particles that have come from the Sun and are captured by an instrument on one of 2 satellites in space, called Helios 1 and Helios 2, to bell-like sounds. Several musicians have used musical sonification of space data to create quartet or orchestra music pieces.

To listen to examples of some of these types of sonifications, please visit our [links](#) page.

Why do we care about Sonification?

There are many reasons why data sonification is useful and interesting. Three such reasons are: 1) to create sounds that people can listen to regarding safety data so that when a situation becomes unsafe, the data changes and people hear this change and know to take action, 2) to help people understand data since people learn many different ways, and 3) to help the blind to do their own scientific research or to learn about science data.

*J. Keller's definitions, 2003, private communication

Technical Notes

STEREO-Spectro was developed in the Max/MSP environment. Max/MSP is a graphically-based object oriented programming language that works well for integrating multiple media and performing real-time analysis and synthesis of data, sound, and visuals.

The main data structure of STEREO-Spectro is based on the analysis of color values in a spectrogram. Using Max/MSP's video/image analysis tools (known as "Jitter"), the program maps the color values for each bin or channel of data into a floating-point number to use in the subsequent processing algorithms.

How to Use STEREO-Spectro

Getting Started

There are two versions of the program that will run on either a Macintosh running OSX 10.3 or a PC running Windows XP or Windows 2000. The program runs best with a minimum memory of 512 MB RAM and processor speed of 1 GHz but it will run with 318 MB RAM and 600 MHz.

To begin, download the STEREO-Spectro application as well as the accompanying “Image Data Files” from:

<http://cse.ssl.berkeley.edu/impact/vos/advanced.html>

These have been combined and compressed using StuffIt in the case of Mac, and using Zip in the case of Windows. Unzip or Unstuff the program and data files (this may happen automatically depending on your system’s setup).

To start the program, double click on the program icon (“STEREO_Spectro_ver_2.4-build” for Mac; “Stereo_pc.exe” for Windows).

The Work Environment

This program is constructed as a “real-time” interactive environment in which you have a good deal of control over a number of independent processes. The program is built as a stand-alone “runtime” application in the Max/MSP programming language. If you are familiar with Max/MSP, the tools and icons should be familiar to you. If you are not familiar with Max/MSP do not worry – most of the features are either self-explanatory or at least easily explained.

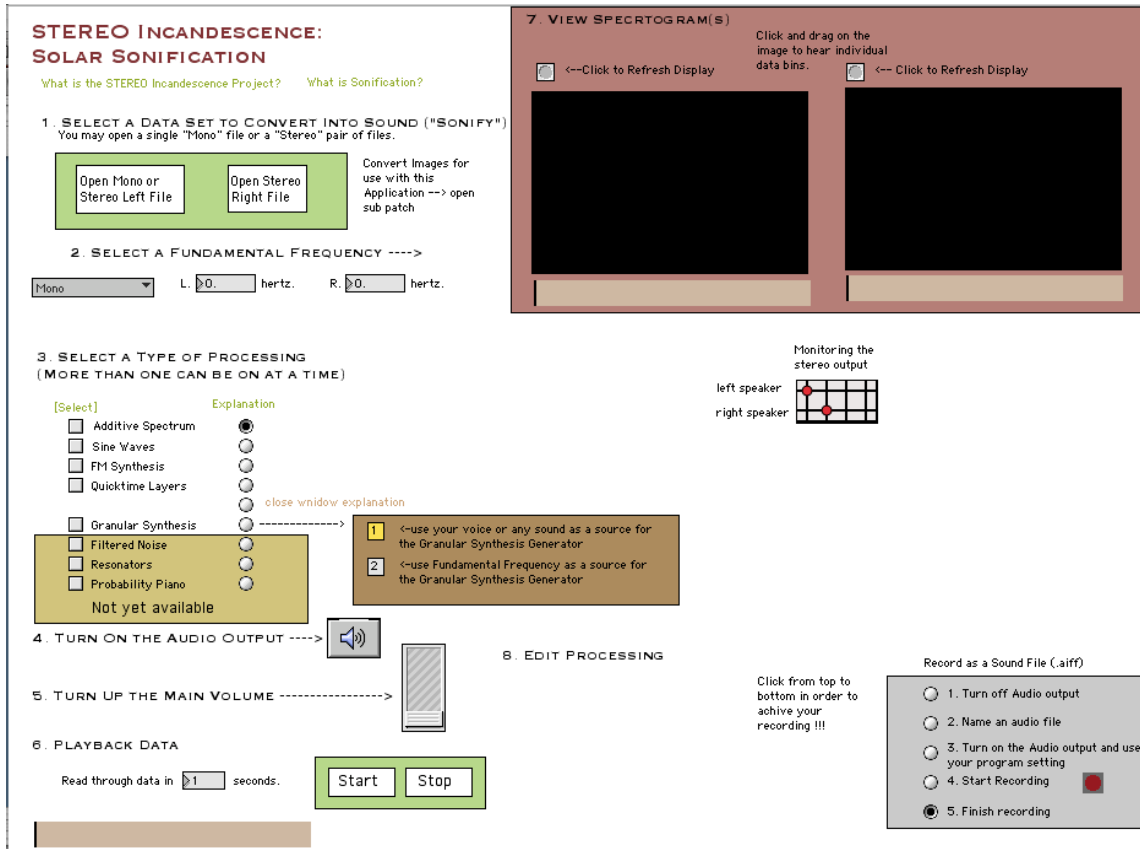


Figure 1.1 – A full screen shot of the STEREO-Spectro environment.

The program is organized so that you follow instructions step-by-step down the left side of the window. Each major task is given a number between 1 and 7. You may always jump backward to edit your previous choices. In general, though, it is best to proceed sequentially through the steps.

Loading a Data File

You may download Image Data Files from our website. These files have been formatted to work with STEREO-Spectro. When selecting data, you may choose to load and process a single file (the “Mono” mode) or a pair of related files (the “Stereo” mode).

1. SELECT A DATA SET TO CONVERT INTO SOUND ("SONIFY")
You may open a single "Mono" file or a "Stereo" pair of files.



Figure 1.2 How to load data files.

To load a mono file, click on the box labeled “Open Mono or Stereo Left File” in the section of the program marked “1. Select a data set to convert into sound”. Then locate the desired file in the open dialog box that appears. Click “open” to select that file and to proceed. A spectrogram image of the data should appear in the first box in step #7.

To load a stereo pair of files, first click on the box labeled “Open Mono or Stereo Left File” and choose the first file. This file will be conceptually the “left” file. Repeat for the “right” file by clicking on the box labeled “Open Stereo Right File”.

You may, at any time, select one or two new files to sonify by simply re-clicking on the appropriate open box.

Creating your own Data Files

STEREO-Spectro can read data from almost any source – given that it is formatted properly. We have structured the program to read color image files in the standard image formats *.gif* and *.jpg*.

The formatting of the image should correspond to the structure of the data. For example, if you have data that has 8 channels of information and 1000 time points for that data you will need to create an image with the pixel dimensions of 1000 x 8. While these images may look funny in a photo application (long and skinny), STEREO-Spectro will convert it into a fairly traditional looking spectrogram.

The color values at each x and y coordinate will represent the value at that location. In general, red represents a high value and blue represents a low value.

Setting a Fundamental Frequency

Before you can sonify the data you must set a fundamental frequency upon which to base the processing. You can think of this as the “key” that your sonification will be in. A higher value in hertz (cycles per second) will result in a higher pitch.




Figure ?? Setting the Fundamental Frequency.

When you load a data set the program will automatically set the appropriate left or right fundamental frequency to 200.0 hertz. You may change this value by moving the cursor to the number box for left or right and clicking and dragging up and down. You will see the values rise and fall with your movement. For finer control, click and drag starting from the right half of the number box. This will allow you to control the decimal portion of the number.

The fundamental frequencies may be different for the left and right – provided that you are in “Stereo” mode. In fact, having a difference in the fundamental frequency may likely help you to aurally distinguish between the two.

Choosing the Processing Type

The next step toward sonification is to decide how you want the computer to convert your data into sound. At the present time there are five functional “types of processing”. The bottom three options are not currently available on STEREO-Spectro but may be found in the application “Incandescence 1.1”. You may find a downloadable copy of “Incandescence 1.1” on the same webpage as for this application.

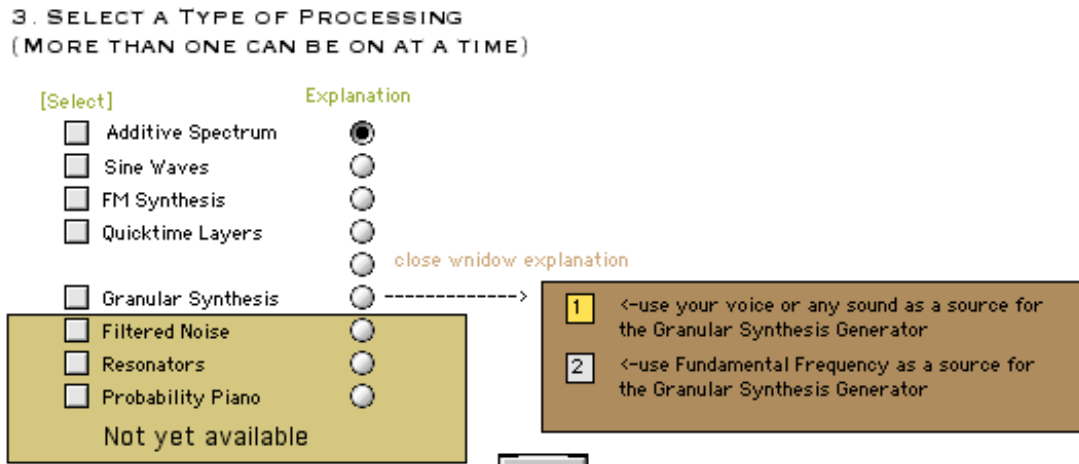


Figure ?? Selecting the type of processing.

Each of the five functional processing types have a square “check-box” directly to the left of the name of the process. This check-box is used to turn that processing type on or off. To turn it on, click once – a check mark should appear indicating that that process is now on. To turn it off, click again.

On the right side of each named process is a circular “radio-button”. These buttons, when pushed, will open another window with text describing the method that a particular process will use to convert the data into sound. In the case of the “Sine Waves” process, there will also be additional parameters that you can control in this window.

It is possible to select more than one type of processing at a time. If you do select more than one, the program may crash if your computer does not have enough memory or CPU speed. If this does happen, simply restart the application and be more frugal in your choices!

Additive Spectrum: This process creates a harmonic spectrum and varies the volume of individual parts of that spectrum based on the color values. The redness at an individual point controls the volume of a note that is harmonically related to the fundamental frequency. The “blueness” and “greenness” control the volume of frequencies slightly off from the red frequencies. This means that when there is more blue or green present the sound will have a somewhat fuzzy quality.

Sine Waves: A sine wave is the simplest type of sound that a computer can generate. In this process, the color value for each of the 8 bins of data is mapped to the frequency (pitch) of 8 sine waves. Each sine wave rises and falls with the changing data. Additionally, the 8 sine waves have different starting pitches in relation to the fundamental frequency. The higher in the spectrogram the bin is, the higher the starting pitch. This allows you to discern the different bins.

When you click on the explanation radio button for “Sine Waves” you have the option to control the range that the sine waves will sweep. A number lower than 1.0 will invert the sweep (making highs into lows and vice versa). The number 1.0 will cause there to be no sweep (always staying the same). Numbers higher than 1.0 will make for larger and larger sweeps. The number 2.0 will make the data sweep between the starting frequency and a frequency twice that size.

FM Synthesis: FM Synthesis is a method of generating sounds by causing more than one sound wave interact with another. Each of the 8 y-axis bins is a different starting frequency with the top bin at the fundamental frequency. Frequency modulation (FM) synthesis works using the same wave principles as radio with a carrier wave and a frequency modulation of the carrier wave. How red the color determines the amplitude (loudness) of the carrier wave, which is the fundamental frequency (main pitch). How green the color determines the frequencies added to the carrier frequency (i.e. the harmonicity), in a five octave range. How blue the color determines the complexity of the sound by changing the modulation index (how much the carrier is affected by the modulating frequencies).

Quicktime Layers: Each of the 8 y-axis bins is a different instrument determined by the fundamental frequency and the sound files on your machine. All instruments have the same fundamental frequency. If there is red in the color, the instrument will play a note a major third above the fundamental frequency. If there is green in the color, the instrument will play a note a perfect fifth above the fundamental frequency. If there is blue in the color, the instrument will play a note a major 7th above the fundamental frequency. Each of these notes is reiterated each time data is read -- making the resulting sound somewhat panicked. Listen for larger shifts in the patterns of the data.

Granular Synthesis: Granular synthesis is a process in which a sound is broken up into small “grains” and then re-assembled. This allows the computer to sustain shorter sounds, and to change the pitch of a sound without changing the overall length. In this case each of the 8 bins of data controls a parameter of a granular synthesis engine. These parameters include grain size, time between grains, the pitch of the grains, amount of variation, etc.

You have the option of performing granular synthesis on a sine wave based on the fundamental frequency or any sound coming in through your computer’s built-in or external microphone. To choose between the two, select either the box marked “1” or “2” next to the granular synthesis explanation radio button. When selecting your voice, the computer won’t output any sound unless sound is coming into the computer. It will process your voice according to the parameters at a given moment in the data.

Turning on the Audio Output and Setting the Volume

Once you have selected the type of processing you are ready to turn on the audio.

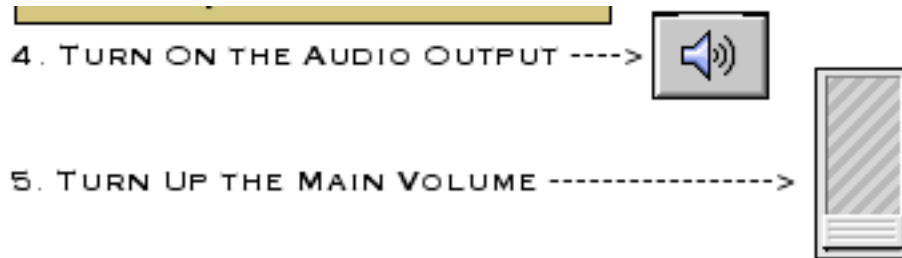


Figure ?? The Audio Controls.

Turn on the audio by clicking the grey-rectangular button with a speaker icon on it. You will know that it is on when the grey is darkened. To turn the audio off, click this button again.

Once the audio is turned on, you may adjust the programs volume by raising or lowering the sliding button to the right of step #5. Slightly above the mid-point is a good starting point.

In order to hear sound you must also have your computer's internal sound set properly. Make sure that your computer's sound isn't muted and that the volume is at a good level.

Playing Back the Data

You may now play back the data and test how the sonification sounds. To do this, first set the duration in which you want the computer to read through the data. You can change the duration (measured in seconds) by clicking and dragging the number box in step #6 up or down. When you load a new set of data this number is automatically set to 30 seconds – meaning that it will play through the data in 1/2 of a minute. The lower this number is, the faster it will read through the data. For more “detail oriented listening” input higher values – in the range of 60-120 seconds. In that range you have a better chance of correlating what you hear with what you see.

6. PLAYBACK DATA



Figure ?? The Playback Controls.

Once you have set the duration, you may click on the “Start” box to start reading through the data from the beginning to the end. You may click “Stop” at any time – but each time you click “Start” it will begin again from the beginning.

As the program is playing back the data you will notice a tan slider box in the lower left-hand corner scrolling through keeping track of where you are. You can also control playback from this box – instructions are found in the next section of this manual.

Controlling Stereo Playback

You may control the way the program routes the sonification to the left and right speakers. There are two ways to do this. The first method is the easiest and should be used unless there is a special configuration that you require that is not accounted for.

To the left, next to the area where you set the fundamental frequency, you will see a small rectangle with the word “mono” or “stereo” in it. This gets set automatically based on the number of data files you have opened.

If you click on this box a pop-up menu appears and you can choose and of the following options:

Mono: The sound of the file is split and goes evenly out the left and right speakers.

Stereo: The left portions of both files go out the left speaker and the right portions of both files go out the right speaker.

Left Graph: The left spectrogram goes out both speakers (same as “Mono”).

Right Graph: The right spectrogram goes out both speakers.

LG RG per Channel: The left portion of the left spectrogram goes out the left speaker and the right portion of the right spectrogram goes out the right speaker (i.e. what you see on the left you hear on the left and vice versa).

No MSP Process: Means that the speakers are shut off. You will only hear “Quicktime Layers” as they do not use the same audio engine.

Sampling the Data

Besides the playback method described above, there are a number of ways you can hear what is happening in STEREO-Spectro.

Scrolling and Scrubbing: You may click on any part of the tan box beneath where you set the duration for playback to hear the data at that point. In addition you may “scroll” forward or backward by clicking and dragging on this box to freely control the speed of playback – going slower over some spots, faster over others.

In addition you can “scrub” back and forth over a single region in order to focus in on the details of that section.

Controlling the playback from this box will cause both the left and right channels to playback (if you are in Stereo mode).

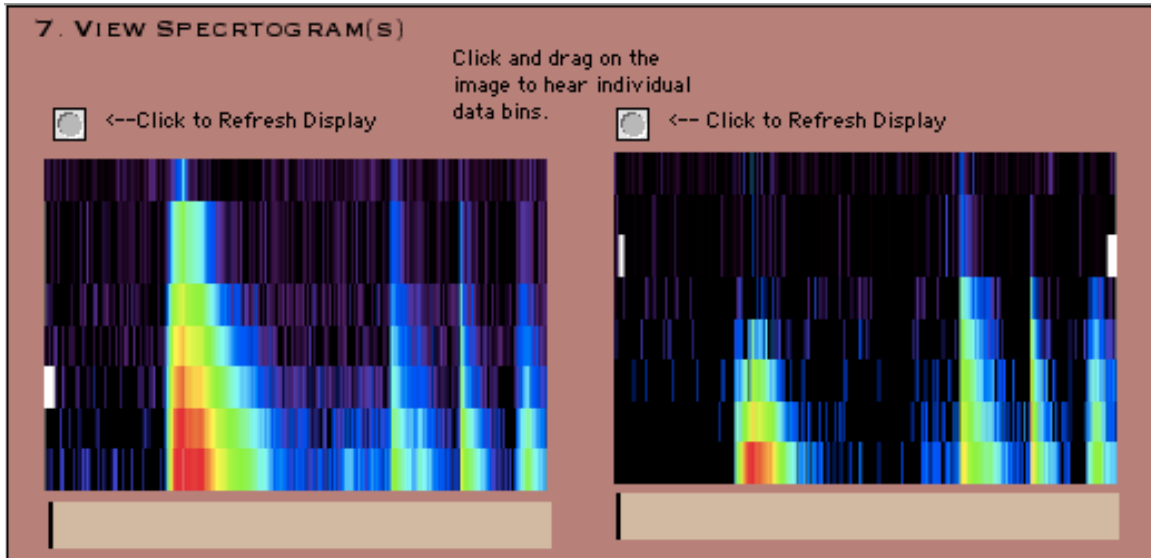


Figure ?? The Spectrograms and their controls.

You may control each of the channels of data individually by clicking and dragging on the tan box directly below each of the spectrograms. This way you can hear the left without the right and vice versa.

Clicking and Dragging on the Spectrogram: You may also click directly on the spectrogram to “train your ear” to what it is that you are hearing. You will hear the sound associated with wherever you click. Notice that the sounds may be different moving up and down as well as right to left. You may do this while automatic playback is happening or not.

This feature will not work for the “Granular Synthesis” mode of sonification.

Recording the Sound of the Data

You may record your sonification as a standard *.aiff* audio file. This type of file may be then played back through, Quicktime, iTunes, or most audio applications.



Figure ?? Record your sonification.

To begin, setup up the sonification to your satisfaction. Then follow the step-by-step instructions found in the grey box in the lower right-hand corner of the screen.

When you click on step #4 the recording begins and the red light will become illuminated. You may then click on “Start” and “Stop” to control the playback, or scrub through the data yourself. You may even click and drag on the spectrogram itself – all sound that you make with the program will be recorded.

Make sure to click on step #5 to end the recording session – otherwise the program will keep recording.

Troubleshooting

If you are having trouble getting sound out of STEREO-Spectro check the following:

- Is the volume on your computer turned up? Your computer's volume will override the volume control within STEREO-Spectro.
- Is the audio output turned on (step #4) and volume turned up (step #5) within STEREO-Spectro?
- Look at the "Status" window. You can view the status window by selecting "Status" from the "Window" pull-down menu or by pressing Command-M. Below you will find a screen shot of a Status window. Most of the lines you do not need to worry about – they simply state that various processes were completed without problem. In the Macintosh example below, there is a problem in the first line that will keep you from hearing any audio other than "Quicktime Layers". The line reads "•error: couldn't load preferred audiodriver CoreAudio" – meaning that your system cannot connect it's audio driver to this application. If you restart the application, this error may go away allowing the program to function properly.

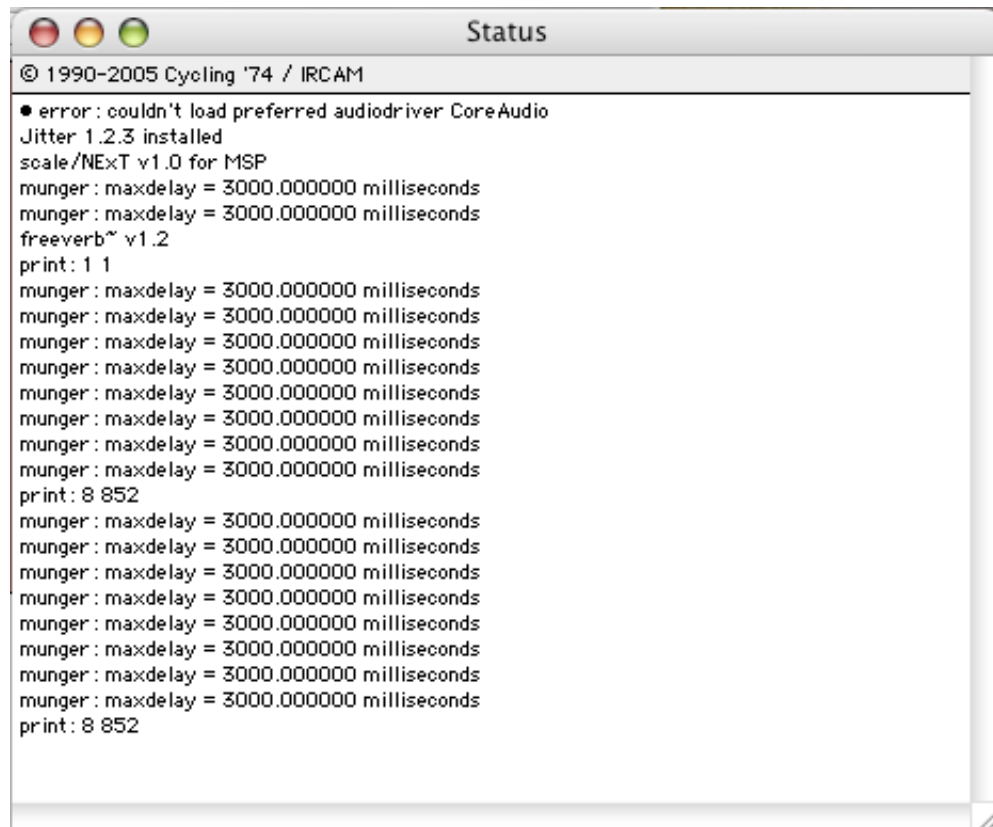


Figure 2.2 The Status window with a CoreAudio error.

If STEREO-Spectro crashed:

- Try using fewer processing types at one time.

If the Spectrogram image has disappeared:

- Click on the appropriate button labeled “Click to Refresh Display”. If that still doesn’t work, it is likely that you don’t have a data file loaded. Open a new file by clicking the appropriate open box in step #1.