SENSIMETRICS
Auditory Interactivities

User Guide

SENSIMETRICS
Auditory Interactivities

Contents

License Information

Foreword

Auditory Interactivities was initially conceived as a follow-on teaching tool in hearing science to Sensimetrics Corporation's earlier product in speech science, Speech Perception and Production I (SPPI). At the start, we knew only that we wanted to aim the level of the material primarily at undergraduate instruction in audiology, psychology, and bioengineering, and were relatively uncommitted to any particular style or design of presentation. We considered following the tutorial style of SPPI, but both the nature of the material in hearing and our own evolving views led us in a different direction.

The fact that the auditory perceptual phenomena that are described in textbooks and lectures can be experienced directly suggested strongly to us that a new computer-based teaching tool should make those phenomena accessible to students. Such a tool would give students the ability to listen and to explore. The idea of a very experiential presentation fit with our belief that computer-based instructional tools, "courseware," should focus on the presentation of material in ways that cannot be done with traditional modes of teaching. We came to feel that our goal was to contribute to students' understanding of auditory phenomena through animations and interactivity, two qualities that make computer-based instruction more than a textbook or an audio recording. The shorthand statement of that goal is that we wanted to develop a 'personal laboratory', a computer program that allows the student not only to listen to well-controlled acoustic stimuli but also to manipulate them and even to conduct experiments. From this grew the design for the psychoacoustic units, which form the bulk of *Auditory Interactivities*, in which the user is given control of key parameters of the stimulus-generation process.

Although our initial intention was to provide fairly complete coverage of hearing science, it became evident about halfway through the development that that goal could be achieved within budget only if quality and depth were sacrificed. We opted, therefore, to limit the scope to interactivities in signals, acoustics, and psychoacoustics, with nothing in anatomy or physiology. This allowed us to develop fully the block-diagram structure that is used in the psychoacoustic units. It is our hope that holes in the current product will be filled in later versions with the exciting interactivities in anatomy and physiology that we had to leave in the planning stages.

Within the chosen topics there still remained many decisions about which concepts to present and which experiments to implement. In the end we made such decisions based on a combination of importance of the topic and its amenability for interactive presentation. For example, the decibel is an important topic, and one with which students often struggle. But we could think of no engaging interactivity that went much beyond that which has been done in textbooks. For psychoacoustic topics our choices were based on our judgment of which experiments were widely considered to be classics in the field. Others were added because we felt they conveyed special insight about an auditory phenomenon that cannot be given in print. Again, omissions can be remedied in later versions.

We also considered various degrees of generality in the block diagram structure. At one extreme, we could have developed a few large, general-purpose, collections of blocks from which students could construct experiments. For example, all types of masking experiments could be generated with a generalized masking equipment configuration. We felt that that approach would be both too difficult and too unfocused for beginning students. At the other extreme, we could have developed interactivities that were wired for a specific experiment and would be usable for little else. The cost of doing this would be loss of instructors' and students' freedom to create and instrument their own experiments. We decided on a compromise in which experimental interactivities are designed to replicate well-documented phenomena, but with freedom to manipulate the available signal-processing components.

The material in *Auditory Interactivities* can be presented in a variety of settings—in lectures, in labs, or at home. Instructors should note that the *Background* and *Instructions* sections that accompany the interactivities tend to be brief; as a result, support from teacher and textbook may be needed. In the *Background* sections we have tried to convey the important aspects of phenomena without reaching the level of detail available in texts and original research articles. *Instructions* describe how the interactivities can be used in the most relevant ways; they do not describe alternative experiments of interest that can be performed with the same configuration of on-screen components. Instructors are encouraged to deviate from the prescribed experiments.

Many people within and outside Sensimetrics Corporation contributed to the design and implementation of *Auditory Interactivities*. Principal among the outside consultants was Nat Durlach, who helped greatly in the early stages when the important design decisions were made. Others who gave valuable advice were: Bob Berkovitz, Lorraine Delhorne, Richard Freyman, Gerald Kidd, Bill Rabinowitz, Charlotte Reed, John Rosowski, Barbara Shinn-Cunningham, Tino Trahioitis, Neal Viemeister, and Bill Yost. Some of them, along with Julie Greenberg and Brian Moore, graciously provided comments on a beta version of the software; Adrian Houtsma deserves special thanks for his thorough review. Thank you all for your help.

The implementation work at Sensimetrics was led by Robert Beaudoin, a mathematician whose ingenious software design has made possible all of the real-time audio that can be heard in *Auditory Interactivities*. He contributed as well to improving the content throughout the project. His main associate in software design and implementation, and the chief programmer, was Jason Carr. Others who made significant contributions to the code are: Gregor Heinrich, Hilton Miyahira, and Joyce Rosenthal. Jens Jorgensen, our graphic designer, gave the *Auditory Interactivities* its clean, clear appearance.

Finally, we gratefully acknowledge the National Institute on Deafness and Other Communication Disorders for supporting the development of *Auditory Interactivities* through Small Business Innovative Research grants, and Dr. Lynn Luethke of NIDCD for her help and support.

Patrick M. Zurek
Sensimetrics Corporation
Somerville, Massachusetts
November, 2003

Introduction

Auditory Interactivities (AI) consists of a collection of structured interactivities and explanations designed to allow users to experience and study hearing-related phenomena using Windows-compatible personal computers. AI was designed as a supplementary tool for teaching hearing science, and it is assumed that users have access to appropriate texts and other reference material. This courseware employs a high degree of interactivity in its presentation of topics in signals and acoustics, as these relate to hearing, and in psychoacoustics. Because of the degree of interactivity required, what might otherwise be called lessons, exercises, or demonstrations are here called *interactivities*. This relatively new word, which denotes an interaction between the user and a computer, describes precisely the intended use of this courseware.

Software Installation And System Initialization

Installation

Insert the AI CD-ROM into your computer's CD-ROM drive. Using *Windows Explorer*, find the file setup.exe in the top-level folder on the CD-ROM. Double-click on the program icon to start the installation process and follow the on-screen instructions.

System Settings

Windows Sounds. As described in the next section, the use of AI requires that your audio system be set up correctly. This may require the Windows volume controls to be set at a higher level than usual. As a result, the standard Windows operating system sounds (Windows start, program error, etc.) may be very loud. This can be especially startling if you are using headphones when operating system sounds are presented. Unless you are genuinely dependent on these Windows sounds for your use of the computer, we suggest that you disable them.

Follow these steps to disable Windows system sounds:

Windows XP:
1. Go to Start > Control Panel > Sounds and Audio Devices
2. Select the Sounds tab.
3. Under Sound scheme select "No Sounds" from the drop-down list.
4. Click the OK button at the bottom of the Sounds and Audio Devices window.

Windows 95/98:
1. Go to Start > Settings > Control Panel > Sounds.
2. Under Schemes select "No Sounds" from the drop-down list.
3. Click OK at the bottom of the Sounds Properties window.

Font Size. This display property must be properly set for AI to display information on the monitor correctly.

This may already be set properly, but if not, follow these steps to change it:

Windows XP:
1. Go to Start > Control Panel.
2. Double click on Display.
3. On the General tab under the Display section DPI setting, select Normal size (96 DPI) from the drop-down list.
4. Click the OK button at the bottom of the Display window.

Windows 95/98:
1. Go to Start > Settings > Control Panel.
2. Double click on Display.
3. On the Settings tab, click on the Advanced button.
4. On the General tab under the Display section, select "Small Fonts" in the Font Size drop-down list.
5. Click the OK button at the bottom of the Advanced window.
6. Click the OK button at the bottom of the Display Properties window.

Changing this selection from Large Fonts to Small Fonts may cause the windows and text in other programs to be slightly smaller than usual.

Screen Area. AI will work with any Screen Resolution (Screen Area on Windows 9X) setting with values 800 x 600 or larger. The Screen Resolution setting determines how large the AI windows appear on your screen. A Screen Resolution setting of 800 x 600 pixels will make the AI interface fill your entire screen. If your vision requires it, such a large low-resolution display is a good choice. With the 800 x 600 setting the lower part of the AI window (including data buttons) may be covered by the Windows taskbar.

If this happens, please take the following steps to make the windows taskbar "Auto hide":

Windows XP:
1. Right-click on the Windows taskbar and select Properties from the popup menu.
2. Check "Auto-hide the taskbar" under the Taskbar Appearance section.
3. Click OK on the Properties window.

Windows 98:
1. Right-click on the Windows taskbar and select Properties from the popup menu.
2. On the Taskbar Options tab, select "Auto hide."
3. Click OK on the Properties window.

Using "Auto hide" will cause the taskbar to appear only when you bring the mouse cursor down to the very bottom of the screen.

A Screen Resolution setting of 1024 x 768 or higher results in a desktop area that is larger than an 800 x 600 display needed for AI. With the higher screen area setting, it is easier to switch to other programs by clicking on their icons because AI's window does not occupy the entire desktop area on your monitor.

Note: Unlike CRT monitors, LCD monitors are sharpest at their native display resolution. While it is possible to use a lower resolution, for example 800 x 600, to display the AI *screen larger on these monitors, it will often look grainy or distorted. Check your monitor (or laptop) user guide for the native resolution of your LCD monitor.*

Audio Setup And Calibration

Selecting Earphones and Loudspeakers

Some of the interactivities require listening with earphones, others require loudspeakers, and others allow signal presentation in either way. For individual listening, earphones are highly recommended whenever possible because:
- they cause minimum disturbance of others in your vicinity;
- they can help prevent interfering environmental noise from reaching your ears;
- it is easier to reproduce acoustic stimuli using earphones than with loudspeakers.

IMPORTANT WARNING!!!

Always observe these precautions when using earphones:

1. Start the sound playing before putting on the earphones. This lets you judge whether the sound is too loud as you bring the earphones to your ears.

2. As always when listening with earphones, be careful to avoid prolonged exposure to high-intensity sound. If you experience ringing in your ears, turn the volume down.

Many different kinds of earphones are available. The major types are: *insert earphones*, which are small enough to fit into the concha or even partly into the ear canal; *supra-aural earphones*, which rest on the pinna; and *circumaural earphones*, which enclose the pinna in a shell-like housing. There can be large differences in the sound levels produced at the ears by different types of earphones even when the volume settings are the same. If ambient noise is a problem where you will be listening with earphones, you may want to consider headsets that have earphones inside sound-attenuating muffs. These devices are sold as a sub-category of hearing protectors called “communication headsets”; an example is the HT-series “Listen-Only” headset from Peltor® Communications.

There are also many types of loudspeakers. They range in size from the small loudspeakers that are often sold with computers, and which often have built-in amplifiers, to large loudspeakers of the kind used in home audio systems, which require an external amplifier.

In general, the interactivities in AI do not require expensive, high-quality earphones and loudspeakers. There are only a few interactivities that are very exacting in terms of distortion and gain control. Warning is given when that is the case.

Maximizing Audio Quality

To present sound to a listener, the representation of the signal stored in the computer must first be converted to an electronic signal, and then to an acoustic signal audible to the listener. Several components and controls affect the quality of these transformations. It is important that their settings are optimized at setup time and maintained afterward.

The main aspects of audio quality that are affected by control settings include dynamic range, linearity and left-right channel balance or matching. Loudspeaker placement can also have important effects which are made more complicated when two loudspeakers are used.

Dynamic Range. You should try to achieve as much range in signal level as possible, while also holding the internal system noise to a minimum, by adjusting those gain controls to which you have access. These gain controls are part of the Windows volume control panel, and should be carefully adjusted when you listen with earphones connected directly to the computer. If you are using an external amplifier, that, too, must be adjusted.

Sound board controls. Some sound boards allow the maximum output to be set by a control using manual switches or a thumb wheel. Check to see if your sound board has such a control. If it does, set it to an easily-replicated setting, such as the maximum value.

Windows volume controls. The full array of Windows volume controls is made visible by clicking on the small loudspeaker icon in the Windows taskbar at the bottom of the desktop. To adjust these controls optimally, follow these steps while the environment is quiet:

1. Right-click the loudspeaker icon and select *Open Volume Controls* to access the Balance and Volume controls for the audio devices on your system. The array also includes a master volume control, usually labeled “Volume Control” (sometimes labeled “Master”). Note also that clicking the loudspeaker icon once with the left mouse button brings up a single slider labeled “Volume Control.” That slider is also the master volume control.

2. For optimum use of AI it is best to mute all devices other than “Wave” and ”Volume Control” by clicking their “Mute” checkboxes because these unused channels may add noise. Of course, if you wish to use any of those devices later, it will be necessary to uncheck the appropriate “Mute” checkbox.

3. Make sure that the Balance slider control on “Wave” and “Volume Control” is at the center position of each control.

4. Adjust the vertical slider of the “Wave” volume control to its highest (maximum) position.

5. Go to Interactivity 3-1-1, which generates a single sine wave signal. Adjust the sine wave frequency to 100 Hz and the level to 40 dB_N (the dB_N scale is explained below). Switch the audio on (by clicking on the loudspeaker icon in the AI Control Bar), then put on the earphones. You may hear the low-frequency tone. If not, slowly increase its level until you hear the tone.

6. After you have adjusted the level so that you hear the 100 Hz tone, reduce the level using the “signal L” slider until you can no longer hear the tone. Listen very carefully at this point. There may be other faint sounds, higher in frequency and with a scratchy quality, that you can hear pulsing on and off with the same temporal pattern as the tone, even though the 100 Hz tone itself is no longer audible. This sound is called *quantization noise*. In addition to this noise, you may hear a constant faint hissing (random noise) or hum in

the background, produced by the electronic components in your audio system. Adjust the "Volume Control" so that all of these noises, both pulsing and constant, are just below the level at which they become audible.

7. Write down and save the resulting settings, using the ticks on the volume scale for reference. This will allow you to reset the controls if they are changed later.

8. Click on the 'X' in the upper right corner of the Volume Control window to exit.

These volume control settings will remain as they have been set until they are changed.

When these controls have been set correctly, there should be a dynamic range of more than 80 dB between the level of a mid-frequency (2000 Hz) sine wave that is just audible and the maximum level that can be produced without waveform clipping, which occurs at 87 dB_N (explained below).

Amplifier controls. If you are using loudspeakers you will need an external amplifier. In addition to a volume control, your amplifier may have tone controls (treble and bass), which change the frequency response of the amplifier, and a balance control, which affects the relative levels of the left and right signals. Adjust the amplifier's volume control so that a low-level sine wave signal, at 0-10 dB_N, is barely audible. Note or mark the volume control setting so that it can be reset easily if necessary. Usually, balance and tone controls have middle positions for exact balance between left and right channels and no treble or bass emphasis. Use those positions. Record the amplifier settings when you are satisfied with them.

Linearity. Without electronic measurement equipment one must rely on listening tests to detect distortion. Gross signal distortion of certain types can be detected easily by listening.

Detecting distortion. Go to Interactivity 3-1-1 and set the sine wave frequency to about 1000 Hz. Start at a low level of about 20 dB_N and gradually increase the level. There should be no change in the quality of the tone as the level is increased. If you hear a change in the character of the sound as its level increases (without setting off the clipping indicator), especially a high-pitched component that adds a rough or raspy quality, there is probably harmonic distortion present.

Finding the source of the distortion requires replacing components, one at a time, until the source of the distortion is found. For example, if you try different earphones and the distortion appears at a noticeably higher volume level setting, or disappears completely, you can conclude that the original distortion was at least partly due to the earphones. If there is no change in the perceived distortion with different earphones, then you can conclude that the problem is in the sound board.

Automatic gain control. Some inexpensive audio amplifiers have automatic gain control or loudness compensation circuits designed to change their amplification or frequency response characteristics with changes in signal level. Although this processing does not usually generate audible distortion, it does make the output signal characteristics level dependent, which can interfere with some interactivities. On some amplifiers this automatic gain or loudness control can be switched off. The presence of such controls is difficult to detect by listening alone, and would have to be determined by consulting the amplifier's Operating Manual, or by electronic measurement. The inexpensive and widely used Radio Shack *Realistic SA-155* mini-amplifier, for example, has built-in loudness compensation that cannot be eliminated without changing the amplifier's internal wiring.

Most of the interactivities in AI that use loudspeakers will not be affected by such gain controls. Those that might be affected by such controls contain a warning.

Left-right channel balance. The left and right channels of your audio system should be matched as closely as possible. Errors in matching can come in the form of amplitude or delay differences between channels. Further, these differences can potentially be frequency-dependent, although some of the most common problems lead to frequency-independent errors. Gross errors in left/right matching will disrupt any of the interactivities that require two-channel output.

The gain difference between left and right channels from the sound board is controlled by the Windows balance settings, as described above. Most amplifiers have a balance control with a middle (equal) position that is clearly marked.

To check that both of your earphones are matched in amplitude, go to Interactivity 5-1-1. Select a sine wave source at about 500 Hz, with on = 500 msec, and R/F = 100 msec. The intracranial auditory image of this tone should appear to be in the center of your head, and not lateralized noticeably toward left or right.

A common problem with some types of earphones is a polarity difference between left and right earphones. If the two phones are not wired consistently, then the same electrical signal that causes a pressure increase at one ear will cause a pressure decrease at the other ear. The left and right acoustic signals are then said to be "out of phase." This condition might be detected with the test just described using a 500-Hz tone. The image of an out-of-phase 500-Hz tone is broader than that of an in-phase tone. However, this is a fairly subtle difference.

A sensitive check for consistent polarity between earphones can be done using acoustic cancellation. Hold the two earphones so that their diaphragms are facing and are as close to each other as possible without touching. Bring the pair of phones in this position up to one ear so that you can hear sound coming from them. Play the 500-Hz tone in Interactivity 5-1-1, first with ITD = 0 and then with ITD = 1000 μsec. If the earphones are wired correctly, the level of the 500-Hz tone you hear should be much louder with ITD = 0 than with ITD = 1000 μsec. If the earphones are wired incorrectly, the tone should be louder with ITD = 1000 μsec. This simple check relies on the fact that in-phase tones reinforce one another, while out-of-phase tones cancel. Out-of-phase earphones should be repaired or replaced, as they would disrupt many of the interactivities on binaural and spatial hearing.

In addition to matching left and right earphones, you should also make sure that the earphone that you use as the source of the "left" signal is consistently used on the left ear, and that it is connected to the computer's "left" audio output. To check this, go to Interactivity 6-1-1, which allows you to switch off the left or right earphone channel. As is customary in audiology and hearing research, the left and right earphone channels have blue and red labels, respectively, on the interactivity screens.

Loudspeaker placement. The placement of loudspeakers is important. They should be symmetrically positioned to the listener's left and right. For all interactivities that require the use of loudspeakers, they should be oriented at an angle between 30˚ and 60˚ relative to the listener's mid-sagittal plane, and at equal distances from the listener. For most interactivities, a distance between 0.5 and 1 meter will be best.

The acoustic environment in which listening is done with loudspeakers is also important. If possible, try to avoid large reflecting surfaces that would produce strong asymmetric reflections.

You should also perform the same channel polarity and consistency checks with your loudspeakers that are described above for earphones. To check the polarity of a pair of loudspeakers, place them so that they are facing each other, a few centimeters apart. From Interactivity 5-1-1 play a 500-Hz tone and toggle the interchannel delay between 0 and 1000 μsec. If the sound is louder with 0 than with 1000 μsec, then the connections are correct as they are. If the sound is louder with the 1000 μsec delay, then reverse the connections to one of the loudspeakers. Then check again with 0 and 1000 μsec delays to make sure it is now correct.

To check channel consistency, use Interactivity 6-1-1 (with your loudspeakers for this purpose rather than earphones as pictured). Turn off one channel or the other to check that your loudspeakers' left and right configuration is consistent with the program. If it is not, swap the loudspeakers' positions.

Biological Calibration and Repeatability Checks

Because the electroacoustic conditions established by users of AI will differ greatly, and because it cannot be assumed that the user has access to electroacoustic measurement equipment, it is not possible to build into the program an absolute calibration of the sound levels delivered to the ears. Without precise measuring equipment, it is not even possible to be certain that electroacoustic conditions are repeatable from one session to another. The best that can be done in the absence of measuring equipment is to employ certain "biological" calibration and reliability checks.

A biological calibration. Interactivity 3-1-1, on measurement of absolute threshold, describes how to use the thresholds of listeners known or assumed to have normal hearing to relate dB_N measurements (described below) to dB SPL. It is possible to estimate dB SPL levels to within an accuracy of 5-10 dB with the conversion factors that result.

Biological repeatability checks. Interactivity 3-1-1 can also be used to check repeatability. Thresholds should be repeatable within 5-10 dB from one measurement to another. It is suggested that a threshold curve be established early with recorded system settings and transducers. Perform this check occasionally and whenever you are suspicious that a change may have taken place.

Interactivity 3-4-5, on monaural diplacusis, allows another opportunity to assess repeatability. Monaural diplacusis is a phenomenon in which a pure tone is heard as being rough in specific idiosyncratic frequency-intensity regions. At a given frequency where diplacusis is present, there is an abrupt disappearance of the roughness as level is increased. Measurement of this transition level is normally highly repeatable, within a few dB, and can be used as a consistency check. However, not everyone experiences monaural diplacusis; if you do not, you will have to find a listener who does and is willing to serve as your calibrator.

The Decibel Scales

The decibel scales used in AI are based on calculations made from the digital signals stored in the computer or generated by the program.

In Units 1 and 2 of AI, arbitrary decibel scales are used based on signal amplitudes normalized to unity.

Units 3-7 use two decibel scales. The first is dB_N, which is the RMS amplitude of the digital representation of a signal, expressed in decibels. The digital signals created (or read in) by AI take on integer values ranging from –32,767 to 32,767. The choice of unity as the reference value is a natural one for this decibel scale. If a digital signal, s[k], is K samples in length, then its dB_N level is

$$dB_N = 10\log_{10}\left(\frac{1}{K}\sum_{k=1}^{K} s^2[k]\right)$$

The second decibel scale is used primarily for pulse train signals, for which an RMS measurement is not appropriate. This scale is called $dB_{N\text{-pk}}$ and is defined for a signal s[k] to be

$$dB_{N-pk} = 20\log_{10}\left(\max_k \left|s[k]\right|\right)$$

A full-scale square wave would have an RMS level of 90.3 dB_N, as well as a peak level of 90.3 $dB_{N\text{-pk}}$. A full-scale sine wave would have an RMS level of 87.3 dB_N, but a peak level of 90.3 $dB_{N\text{-pk}}$.

Using The Interactivities

Structure and Navigation

The *Auditory Interactivities* are listed in the online Table of Contents, which is the first screen encountered when the program is started. The contents are organized into Units that are associated with major topics in hearing. Units are subdivided into Lessons, each of which consists of one or more interactivities.

The interactivities are independent of one another; consequently, there are no restrictions on the order in which they can be carried out. When used for teaching, of course, consideration should be given to the sequence of presentation.

Each interactivity consists of a *Main Screen*, a *Background* section and an *Instructions* section. In addition, an interactivity in which data can be collected (described below) will have an associated window for the display of results.

The *Main Screen* of an interactivity displays graphically the components needed to set up and control the interactivity. It also contains navigation controls and an audio on-off switch. An interactivity in which data can be collected also has buttons for storing and displaying data. The toolbar icon for navigating to the *Main Screen* is .

Background gives a brief summary of the topic of the interactivity and its relevance to the study of hearing. It is accessed by clicking the *Background* button 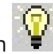 in the toolbar.

Instructions explain how the interactivity can or should be used. The *Instructions* window is accessed by clicking the *Instructions* button  in the toolbar.

The *Background* and *Instructions* are provided on the CD-ROM in a file named Back&Instr.pdf. Opening that file will allow you to view *Background* or *Instructions* in a separate window for access while you also have the *Main Screen* displayed.

The *Results* window shows the collected data in graphic and tabular form. It is accessed by clicking on the *show plot* button at the bottom of the *main screen*. This button and its companions, *record settings*, *clear results*, and *save data to a file*, are available only in those interactivities for which AI supports data collection.

The user can go directly to a specific interactivity by clicking on its title in the Table of Contents, the initial screen seen at start-up. The Table of Contents can also be accessed from any of the Main Screens by clicking on the Table of Contents button  in the toolbar.

Sequential navigation through the interactivities within a unit is also possible by using the right and left (forward and back) arrow buttons  in the toolbar. The Table of Contents must be used to go from an interactivity in one unit to one in a different unit.

In most interactivities, the acoustic stimulus is available for presentation whenever the *main screen* is active. The default program condition, however, has the sound turned off whenever an interactivity is entered. Sound is switched on by clicking the loudspeaker icon  .

Block Diagrams

Many of the interactivities allow exploration of psychoacoustic phenomena through manipulation of schematic equipment diagrams. These diagrams depict the signal generation and processing components or "blocks" required for the interactivity. Detailed descriptions of the functions of these components are given in a separate section of this Guide entitled "Block Diagram Components".

Slider Controls

Parameters in the interactivities are usually set with slider controls, or "sliders." The parameter controlled by a slider can be changed in three ways. First, a single mouse click on the arrow at either end of the slider results in a small step change in the parameter. Holding the mouse button down while the cursor is on one of these arrows causes repeated small-step changes. Second, a single mouse click within the sliding area results in a large step change in the parameter, and holding down the mouse button with the cursor in that area causes repeated large step changes. Third, the slider's "thumb"—the square element that slides—can be dragged with the mouse to a new position and released.

Changes to slider values take effect nearly instantaneously, with the current value of the parameter normally displayed at the right of the slider. When the thumb is clicked and dragged, the displayed parameter value is not yet the actual parameter value, but only corresponds to the current position of the thumb. The parameter is changed to the displayed value when the thumb is released.

For ease of manual control, one of the sliders is always controlled by the left and right arrow keys on the computer keyboard. By pressing the Tab key, the slider controlled by the arrow keys can be selected. Pressing the Tab key repeatedly will cycle the selection through all of the sliders of an interactivity. The thumb blinks on the slider that is currently active.

In addition to reducing the need for manipulation of the mouse, use of keyboard keys allows adjustments to be made without requiring visual feedback from the monitor screen, which may be desirable in some experiments.

In a few interactivities, some of the sliders are not controllable by the user, either because their values are derived from other sliders, or because the parameter is under direct program control. In these cases, the slider is grayed out and an arrow appears through it, indicating that it is being controlled by the program.

Data Collection

Some of the interactivities allow data to be collected and stored. However, even for these interactivities, data collection is not required. The user is always free to listen to the stimuli presented and adjust parameters at will, disregarding the data collection option.

Each interactivity that allows data collection has a pair of parameters that are specified as the "designated" *independent* (x) and *dependent* (y) parameters. The sliders for these two parameters are highlighted. The x and y parameters have been chosen as those most likely to be used when replicating a classic experiment on the interactivity topic. The remaining parameters are "non-designated".

To reduce further the need for mouse manipulation, two keyboard keys have been assigned to control the slider for the designated y variable, which is likely to be manipulated most frequently by the user. The comma/less-than key (,<) decreases the y variable and the period/greater-than key (. >) increases the y variable. A small step is made by typing one of the keys directly, while a large step is made by holding down the Shift key and typing.

The interactivities that allow data collection all use some variant of the method of adjustment. Using this method, the listener adjusts a stimulus parameter to meet some perceptual criterion, for example, "just audible," or "just noticeably different." Suggestions for use of this method can be found in "Guidelines for Data Collection".

Data are recorded every time the *record settings* button is clicked. This and the other data buttons are described in Table I. The recorded data consist of the values of all parameters (i.e., every slider setting) at the time the button is clicked.

Using designated parameters. To use designated parameters, the experimenter should first set up the non-designated parameters. The designated x parameter can then be set to a desired value by the experimenter, and the listener can start adjusting the designated y parameter. When satisfied with the adjustment, the listener clicks the *record settings* button, adding all the current parameter values to a data buffer. The experimenter then can change the x parameter and the listener can make another adjustment. The procedure can be repeated as many times as required.

Table I. Data Buttons And Their Actions

Button	Action
record settings	Saves all parameter values and adds the designated x and y values to the plot. Pressing this button also switches off the sound.
show plot	Switches from the *Main Screen* to the *Results* window, showing a plot of the latest set of designated-parameter pairs recorded with consistent non-designated parameters. The *Results* window also contains a comment box and the data table.
store data in file	Opens a 'Save As' dialog box for writing the data and related information to a text file. The data buffer is not cleared. Sound is switched off.
clear results	Clears the data buffer.
show interactivity	Switches from the *Results* window to the *Main Screen*.

After a series of measurements has been made, the data in the data buffer can be saved by clicking the button labeled *store data in file*. The Windows "Save As" dialog box will appear to allow naming the file in which the data will be stored. Before beginning another run with a different set of non-designated parameters, the experimenter should clear the data buffer by clicking on *clear results*.

Viewing the data. At any time during data collection the data in the data buffer can be viewed in the *Results* window by clicking the *show plot* button at the bottom of the *Main Screen*. The complete set of data in the buffer is available in tabular form at the right side of the *Results* window.

The data are plotted at the left in the *Results* window, with the designated y variable plotted on the vertical axis and the designated x variable plotted on the horizontal axis. The form of this plot for any interactivity (axis ranges, tick marks, log and linear axes, etc.) is fixed. Further, all data points are represented by the same symbol (that is, if there is an exact replication of an x,y pair, only one data point will be visible). The plot in the *Results* window can be printed by copying the screen to the clipboard, done by pressing the *Print Screen* key on your keyboard, and then pasting the clipboard into any word processing program. The quality of the resulting printout is often not very good with this method, however. Better results are obtained by importing the saved data into a spreadsheet or graphing program and creating the plot with that program.

Data file organization. The body of a data file is organized as a matrix, or table, with the interactivity's parameters for the columns and sequential records as rows. A cell in the table will contain the value of a specific parameter for a specific instance of a *record settings* button click.

Preceding the body of the table is a preamble, which includes the version number of the program, the title of the interactivity, the name of each parameter and its associated column number, the date, the start and end times of the run (more precisely, the times of the first and last *record settings*), and any text comment entered before the file is saved.

The order of assignment of parameters to columns follows that given in the *Instructions* section for the interactivity. Consistent with that ordering, the designated x parameter will always be the first column and the designated y parameter the second column.

For example, the text file containing results from Interactivity 3-1-1 would look like this:

Sensimetrics *Auditory Interactivities* v.1.0.0
Unit 3: Monaural Perception Of Stationary Signals
Lesson 1: The Range of Hearing
Interactivity 1: Threshold of Hearing
column 1: x = signal F (Hz)
column 2: y = signal L (dB_N)
column 3: on (msec)
column 4: ISI (msec)
column 5: R/F (msec)
date: June 16, 2003
starting time: 12:46pm
ending time: 12:49pm
comment:

3242.4	65	500	750	25
3242.4	65	500	750	25

Using non-designated x or y parameters. The user is free to adjust any parameter at any time, and to store those settings as recorded data. The sequence of parameter adjustments need not follow the scenario described above for designated parameters. This allows considerable flexibility in the design of experiments.

The only consequence of varying non-designated parameters is in the graphic display in the *Results* window. Because this plot is specifically designed for the designated parameters, plotting the values of the designated x and y values when other parameters varied would be inappropriate.

The way in which the program prevents mixing of plotted data (designated x and y pairs) across a change in a non-designated parameter is by plotting only the designated x and y pairs recorded since the most recent change in a non-designated parameter. In other words, all of the data plotted at any time must have been obtained with the same set of non-designated parameters. Consider as an example the following body of a data table showing the settings for the first six recordings.

50	75	1000	500	300	25	70	1000
70	76	1000	500	300	25	70	1000
100	78	1000	500	300	25	70	500
140	83	1000	500	300	25	70	500
200	85	1000	500	300	25	70	500
400	84	1000	500	300	25	70	500

After the first point was recorded, the *Results* window plot would display a data point at (50,75). After the second recording it would add a point at (70,76). For the third recording a non-designated parameter was changed (relative to the previous set), so the previous data points are not displayed and only the new one, at (100,78) is shown. Note that the previous points are retained in the data buffer and will be written to the output file along with the subsequent points if *a store data in file* command is executed.

The designated x and y pair for each of the next three trials would be added to the plot after they were recorded. The display after the sixth record would show the last four designated x and y pairs. Because of this program behavior, it is important that printouts of the plot in the *Results* window be carefully annotated to indicate the relation of the data points on the printout to those in the data file. If the data are not stored in a file, the values of the non-designated parameters should be recorded by hand because there will be no other record.

Data that result from variations in non-designated parameters can be imported into other software packages for plotting. Remember, however, that there is no re-adjustment of the columns in the body of the data file. If the values of the manipulated *independent* and *dependent* parameters are to be extracted from the data file, be sure that the correct columns are selected.

Memory for parameters and data. *Auditory Interactivities* will remember parameter settings and collected data for the last several interactivities visited. The exact number of data groups remembered depends on availability of Windows graphics resources, which depends on which version of the Windows operating system is being used, and how many other graphic-intensive programs are running concurrently with AI—among other factors. Given this uncertainty, it is best to save data that you may want to keep before leaving an interactivity.

Other than the data that have been saved to files, there is no memory for parameter settings or data after the user exits *Auditory Interactivities*.

Block Diagram Components

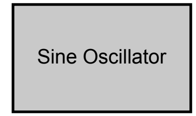

Generates a continuous sine wave with parameters frequency (Hz) and level (dB_N).

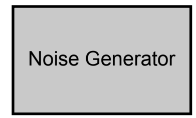

Generates white gaussian noise with parameter RMS level (dB_N).

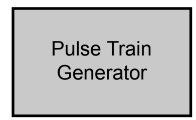

Generates a periodic train of impulses with parameters fundamental frequency (Hz) and peak level ($dB_{N\text{-}pk}$).

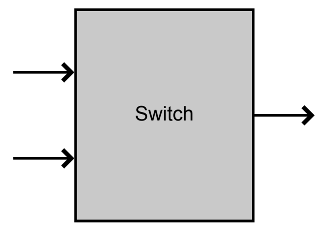

Provides switching control of multiple input signals. Switches are used extensively in the interactivities to turn signals on and off with precise timing. However, it is not obvious from inspection of the block diagram on the *main screen* how a switch is controlled by the displayed parameters. One must refer to the timing diagram in the *Instructions* for an explanation of the relation between the switch sliders and signal switching parameters. The most-frequently used switch parameters are the on-time of a signal, the rise and fall times of a signal, and the time interval between signal presentations. The rise and fall portions of the envelope follow the rising and falling portions, respectively, of a Hanning function.

Usually, switches operate in a cyclic mode in which the switching operations are repeated after a specified cycle time. A switch that operates in a cyclic mode can be set to have its output continuously on by setting ‘on’ to its maximum value, and setting both ISI and R/F to zero. A few interactivities overide such cyclic repetition and present only one cycle when commanded by the user. The *Instructions* should be consulted for a description of switch behavior for an interactivity.

Switches appear to be different in different interactivities because, in order to simplify, only the most relevant subset of the underlying switching parameters for each interactivity is made available to the user.

A gain block applies a specified gain to the input signal. The parameter is the gain in dB.

This component applies a bandpass filter, with parameters low-frequency cutoff (LFC) and high-frequency cutoff (HFC). In some interactivities the filter parameters are presented in terms of center frequency CF and bandwidth BW. In such cases BW and CF are translated internally so that

LFC = CF - 0.5BW
and
HFC = CF + 0.5BW.

LFC and HFC denote frequencies at which the filter's magnitude response is 6 dB below that at the center frequency.

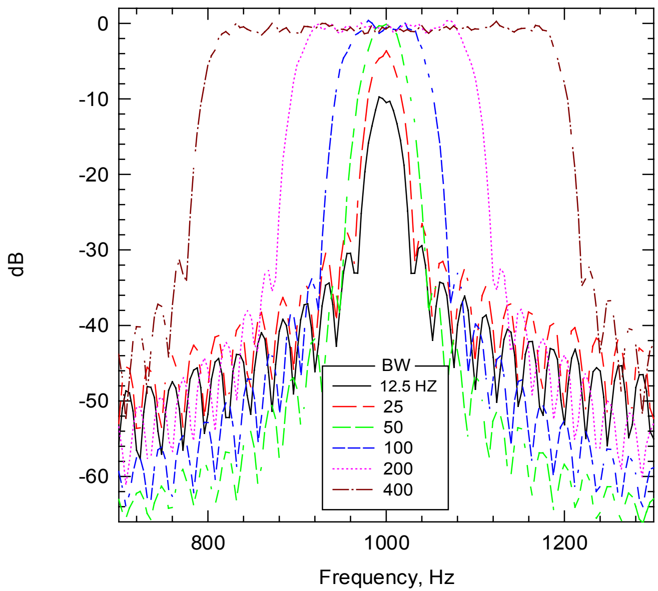

Figure 1. Frequency response of the bandpass filter measured at six bandwidth settings.

The magnitude response of the bandpass filter to white noise was measured using a spectrum analyzer (Figure 1). The center frequency of the filter was fixed at 1000 Hz and the bandwidth was varied from 12.5 Hz to 400 Hz. The shape of the bandpass filter depends only on the distance in Hz from the center frequency, and does not vary appreciably with center frequency.

Note that because of the non-ideal (non-rectangular) shape of the transition band, the peak spectral level at the center of the band is reduced for bandwidths less than about 50 Hz. This error can be summarized in terms of its effect on overall output level. Ideally, the relation between bandwidth (BW) and band level out of the bandpass filter (dBout) with white noise input should be

$$dB_{out} = N_0 + 10\log_{10}(BW)$$

where N_0 is the spectrum level (i.e., the noise level in a 1-Hz band) in decibels. Measured values of dBout are plotted in Figure 2(a) along with the ideal relation. The error between measured and ideal is plotted in Figure 2(b). The error is more than 1 dB for bandwidths less than 50 Hz. For bandwidths larger than 50 Hz the error is negligible.

This non-ideal behavior of the bandpass filter, although applicable only at narrow bandwidths, must be taken into account in experiments such as the Critical Band in Interactivity 3-5-1. If bandwidths below 50 Hz are used, the effective band levels should be estimated by compensating for the errors described in Figure 2.

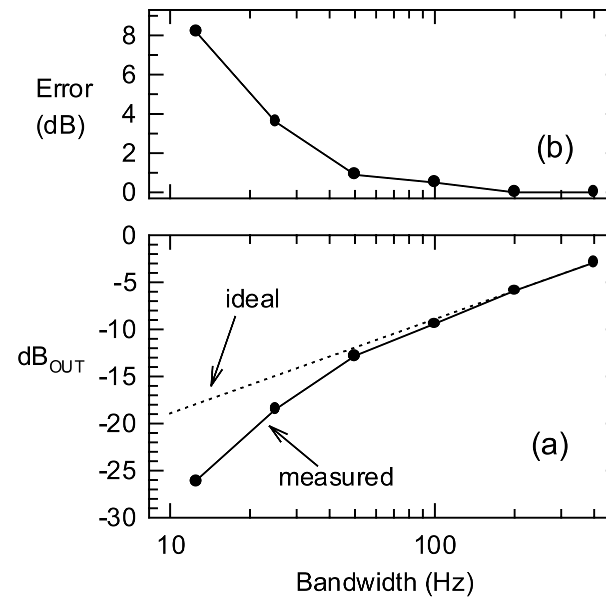

Figure 2 (a) Theoretical (ideal) and measured band levels out of the bandpass filter as a function of bandwidth; (b) difference between ideal and measured band levels.

Band-stop Filter

Applies a notch filter. The parameters are the low-frequency cutoff and the high-frequency cutoff (or center frequency and notch width). Measured frequency responses for a range of notch widths are shown in Figure 3. The responses are essentially independent of center frequency, and depend only on the distance in Hz from the center frequency.

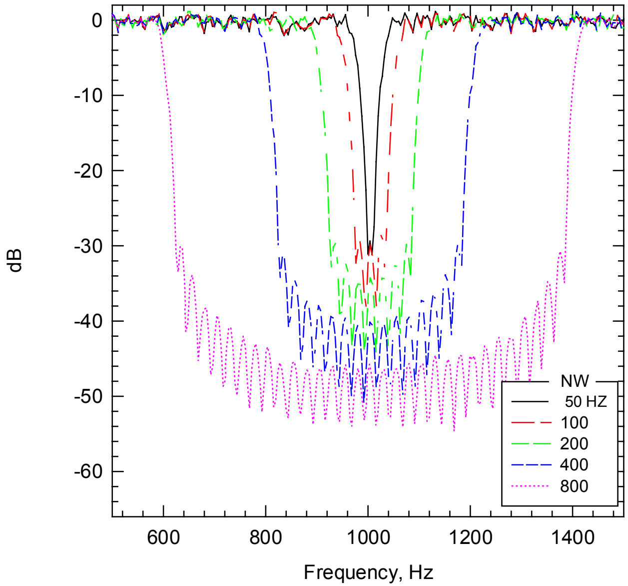

Figure 3. Frequency response of the notch filter at five different notch width settings of the filter.

Complex Tone Generator

This component generates a signal that is a sum of a set of sine waves determined by the following parameters:

N = number of components
L = low harmonic number
H = high harmonic number
DF = frequency shift (Hz)
F0 = fundamental frequency (Hz)
signal L = RMS level of output (dB_N).

The Complex Tone Generator generates N adjacent harmonics of the fundamental frequency F0, between the lowest (L) and highest (H) harmonic selected. If N is larger than the number of harmonics that exist in the selected range, the specified range will take precedence over the number of harmonics selected. For example, if N = 8, L = 3 and H = 7, five harmonics (the third through the seventh) will be selected. However, if N, the number of harmonics, is less than the number existing within the range (which is equal to H – L + 1), then the program will select N adjacent harmonics randomly from the specified range.

In addition, each sine component can be shifted in frequency by DF Hz.

The Complex Tone Generator can operate as a stand-alone generator, but it can also be driven by a “melody generator”. This component, which shows up in Interactivity 3-4-1 as a choice of melodies, supplies relative pitch sequences to the Complex Tone Generator. These relative values are translated via F0 into frequency values. When driven in this way, the location of the adjacent harmonics is chosen randomly for each note in the sequence when N < (H – L + 1).

Mixer

Forms output signal from the sum of input signals.

Delay

Delays a signal by the time selected. Both positive delays (input leads output) and negative delays (output leads input, also known as an advance) are possible.

new

.wav Player

This component plays a .wav file repeatedly, that is, “in a loop.” An audio file is specified by clicking ‘new’, which stops all audio until the file has been selected and loaded. The .wav Player measures the RMS level of the signal, allowing its level to be set to a specific value. To be played by the .wav Player, a file must have a sample rate of 22050 Hz. Files with a duration less than 30 seconds will be looped in their entirety. If a file is longer than 30 seconds, only the first 30 seconds will be looped. For a few interactivities (for example, 3-7-1 and 6-1-2), the maximum duration is set to less than 30 seconds.

On-off switch. The state of the switch is changed by clicking it. This switch operates as an on-off switch without altering the source signal. This is to be distinguished from the behavior of the loudspeaker icon in the toolbar, which turns all audio processing for the interactivity off and on, and resets all internal timing (such as switch timing cycles and file reading).

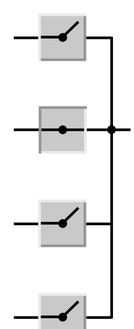

When the outputs of multiple on-off switches are connected to each other they function as “radio buttons” which means that only one can be on at any time. Clicking on a switch that is off will turn it on, and turn off any switch that was on.

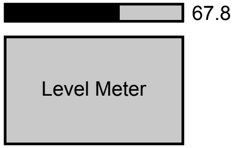

The level meter shows signal level graphically as well as the RMS level (dB_N). The level is calculated using a window typically 100 msec in length.

The clipping indicator shows when the left or the right channel output signal exceeds the numerical range of the system, causing clipping of the signal.

Special purpose filter. These include the speech-spectrum filter (Interactivity 6-1-4) and the HRTF filters (Interactivity 6-1-1, for example). The filters are described in the respective *Background* and *Instructions* for those interactivities.

Guidelines For Data Collection

The method used for data collection throughout AI is the method of adjustment. This is a very simple method from some points of view, but complicated from other viewpoints. The main simplification that it offers is that every interactivity in AI involving data collection can use the same method. In addition, data can be collected using non-designated parameters by the same method as that used for the designated parameters.

In addition to being simpler in some ways, because the method of adjustment asks listeners questions about their internal sensations, it allows measurement of some phenomena that require introspection. An example of such a subjective phenomenon is monaural diplacusis, which can be measured in Interactivity 3-4-5.

While the method of adjustment seems simple and direct, there are drawbacks to its use. First, because it is entirely subjective, there is no way in which a listener's response (i.e., the adjustment) can be considered correct or incorrect. For many of the psychoacoustic phenomena covered in AI, the method of adjustment would not be considered the most rigorous means possible for collecting psychophysical data, for publication, for example. Specifically, for tasks that can be structured so that the listener's response can be scored as correct or incorrect, it is preferable to use forced-choice methods. Forced-choice methods have the advantage of not requiring listeners to make decisions about the magnitudes of sensations or about differences between the sensations produced by two sounds which may be known to be different. Rather, in a forced choice task a listener is asked a specific verifiable question about the stimulus, e.g., "Was the signal presented in the first or the second interval?" The accuracy of the listener's answers, often expressed as the percentage correct in a block of trials, as a function of a stimulus variable, such as signal-to-noise ratio, is measured and interpreted in terms of auditory phenomena. Using the method of adjustment, on the other hand, a listener would be asked to set the signal-to-noise ratio, for example, so that the signal is just detectable.

Another concern with the method of adjustment is that it requires considerable motivation on the part of the listener to listen carefully and cooperate with the experimenter. The listener must be willing to exercise diligence and patience, while being efficient in searching with the adjusted variable and homing in on the final adjustment. Forced-choice tasks require motivation as well, but they require no interaction between listener and stimulus. Sounds are presented and the listener responds. For this reason, you should try to find and maintain conscientious, motivated listeners. Always make sure that the listener is comfortable and has frequent rest breaks.

A potential problem with the method of adjustment is possible anticipation by the listener of a perceived pattern in the final adjustments. For example, a listener who finds that the final adjusted value is six or seven steps down from the starting value may be biased to settle in that region. To prevent this anticipation, randomize the starting values of the adjusted variable over a wide range.

It is also good practice to randomize the sequence of values of the independent variable presented to the listener. Again, the listener may develop an expectation based on past adjustments or on knowledge of the phenomenon being studied.

Finally, we strongly urge you to observe one of the fundamental rules of any experimental science: REPLICATE! By repeating your measurements you will have a measure of the variability of the replicated conditions—the inherent error of measurement—and from that you will be able to judge the significance of any differences in conditions. This requirement applies to all methods of measurement, of course, not only the method of adjustment. Replication will establish the appropriate degree of confidence in the validity of your measurements.