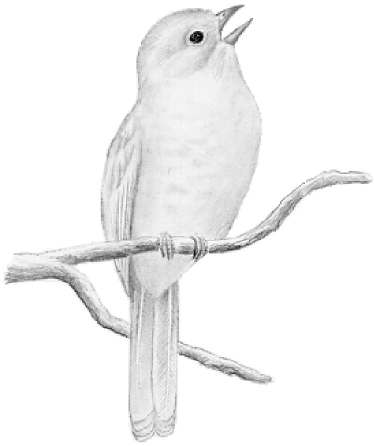

Canary

The Cornell Bioacoustics Workstation

Version 1.2

User's Manual

Bioacoustics Research Program
Cornell Laboratory of Ornithology
Ithaca, New York

Copyright notice

Bioacoustics Research Program
Cornell Laboratory of Ornithology
159 Sapsucker Woods Rd.
Ithaca, NY 14850
USA

Information: http://www.birds.cornell.edu/brp/
User forum: http://canary.RavenSoundSoftware.com

Trademarks

Credits

Canary and the *Canary User's Manual* were developed under NSF grants BIR 8915149 and BIR 9211761 (Principal Investigator: Christopher W. Clark), with additional support from the Cornell Laboratory of Ornithology, and the Higher Education Development Program of Apple Computer, Inc.

Canary was written by Steve Mitchell, Sean Cunningham, Jamie McClellan, and Jacob Montgomery.

The *Canary User's Manual* was written by Russell A. Charif with contributions by Steve Mitchell and Christopher W. Clark. Cover artwork by Cynthia J. Page.

The Canary development project is under the general direction of Christopher W. Clark.

Citation

The correct bibliographic citation for this manual is as follows:

Charif, RA, S Mitchell, and CW Clark. 1995. *Canary 1.2 User's Manual.* Cornell Laboratory of Ornithology, Ithaca, NY.

Contents

Preface Welcome to Canary 1.2

What Canary is

Canary is a software tool for the digital acquisition, manipulation, analysis, and measurement of sound on Macintosh computers. Canary was developed by the Bioacoustics Research Program of the Cornell Laboratory of Ornithology with support from the National Science Foundation to provide a low-cost, user-friendly research and teaching environment tailored to the needs of biologists working with acoustic signals. Many of Canary’s features previously were available only in much more expensive software packages. Canary’s user interface and documentation are designed to make these features accessible to users who may not have extensive backgrounds in acoustics or signal processing, without sacrificing the power and rigor demanded by technically sophisticated users.

The main components of the Canary working environment are *signal windows*, the *command panel*, the *measurement panel*, and the *data log* (Figure 1). A signal window may contain one, two, or three panes, containing waveform, spectrogram, or spectrum views of all or part of a signal. The number of signal windows that may be open at one time is limited only by the available memory. The command panel provides graphic, mouse-operated controls for adjusting the horizontal and vertical scales of the three panes, for showing and hiding each pane, and for controlling sound playback. The measurement panel can be configured to display any combination of signal measurements based on either the current mouse position (e.g., time, frequency, amplitude), a highlighted region (e.g., peak amplitude in a selected section) or signal parameters (e.g., sampling rate, spectrum filter bandwidth). More than eighty individual measurements and parameters are available for display in the measurement panel. Measurements can be transferred from the measurement panel to the data log with a single mouse-click. The data log can then be saved in a variety of formats for easy export to statistical, spreadsheet, word-processing, or other programs.

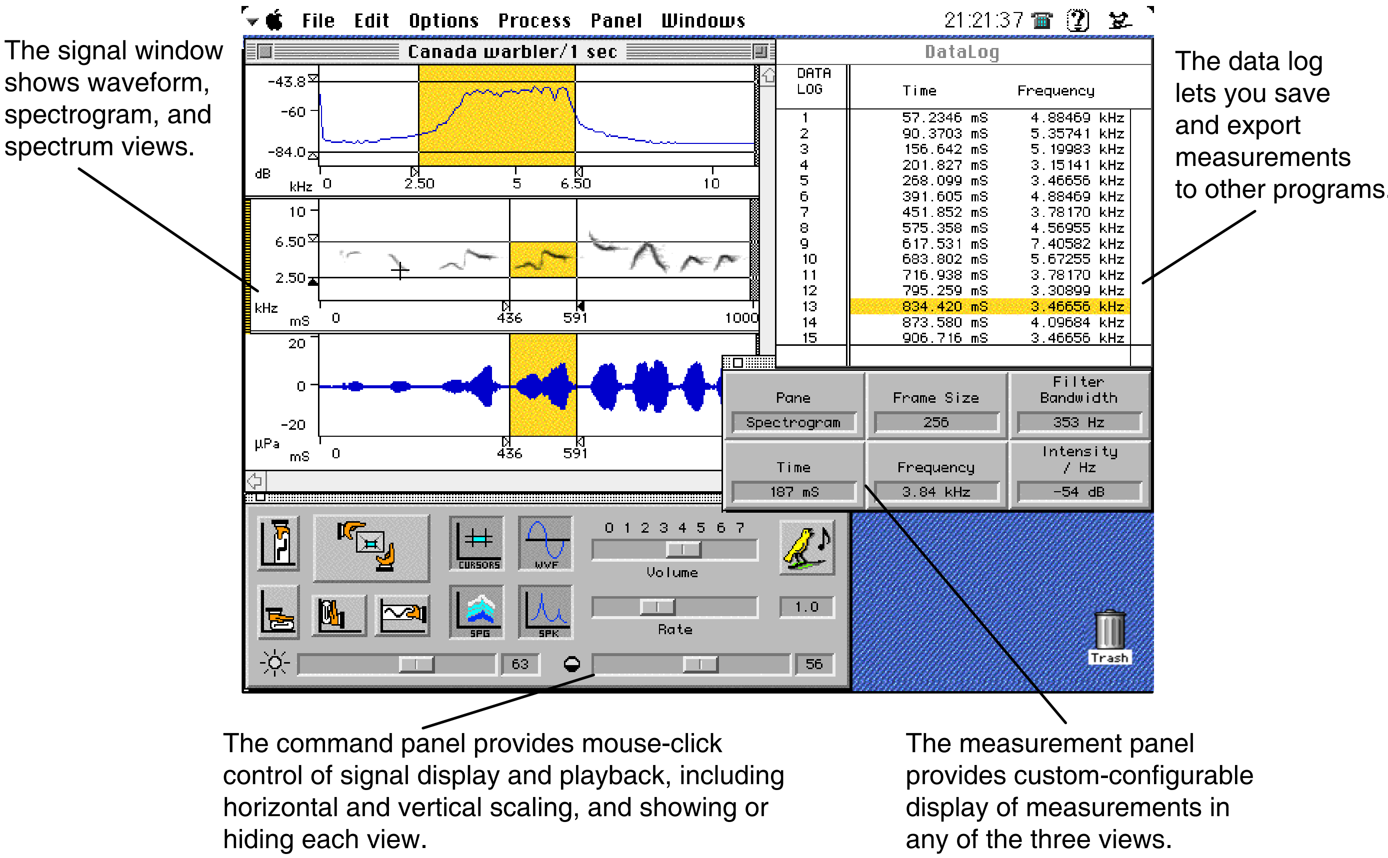

Figure 1. The Canary working environment, showing a signal window, the command panel, the measurement panel, and the data log.

What Canary can do

Canary 1.2 provides the following capabilities:

Data acquisition and playback

Canary can acquire (digitize) audio data using the Macintosh's built-in sound input port (on most newer Macintosh models) or a third-party device such as MacRecorder. Signals can be played back at a variety of speeds.
New in version 1.2:

- Canary can acquire either one- or two-channel signals.
- Signals can be acquired either to memory (RAM) or directly to a disk file.

Signal editing

Canary enables you to cut, copy, paste, delete, or amplify any section of a signal with a few mouse clicks. Sections of a signal that have been cut or copied can be pasted into any location in the same signal or into a different signal. ***New in version 1.2:***

- Filter selected part of a signal to eliminate energy from a specific frequency band.

Spectrum analysis

Canary provides explicit control of all parameters of STFT (short-time Fourier transform) spectrum analysis to calculate and display spectrograms and spectra. ***New in version 1.2:***

- Spectrogram contrast and darkness are fully adjustable from command panel.
- Save named sets of spectrogram or spectrum parameters in preference files for quick retrieval.

Correlations

Canary lets you calculate correlations between pairs of spectrograms or waveforms. Correlations are useful for detecting the presence of one signal within another (possibly noisy) signal. Spectrogram correlations can be used as a rudimentary measure of the similarity between spectrograms. Waveform correlations are useful for determining time delays between occurrences of the same signal in different files.

Batch processing

Canary can automatically perform the same operation on an arbitrarily large number of files and save the results. Simply place all of the files to be processed in one folder, specify the parameters for the process once, and all files will be processed without any further interaction. Batch processes available are spectrogram and spectrum calculation, correlation, and file type conversion.

Printing and graphics export

You can print the contents of the active window directly from Canary to any Macintosh printer. If the printer does not support grayscale printing (e.g., ImageWriter II, Personal LaserWriter SC), spectrogram images are automatically dithered (shades of gray are simulated by varying dot densities). To support the export of PICT graphics to other Macintosh applications, Canary is distributed with the Flash-It shareware screen-capture utility.

Sound level calibration

To facilitate measurements based on absolute signal strength, Canary lets you calibrate signals and copy calibration parameters between signals recorded under the same conditions. ***New in version 1.2:***

- Calibrate signals using either an acoustical or electrical paradigm.
- Specify decibel reference values for sound pressure, intensity, voltage, and power.
- Save a default calibration for automatic application to new signals.

Multiple file formats Canary can read and write sound files in SoundEdit and AIFF (Audio Interchange File Format) formats, as well as Canary's own format, which lets you save a signal with a spectrogram and spectrum in one file. You can also save or retrieve sounds, spectrograms, spectra, or correlation arrays in binary, Text, or MATLAB formats.

What you need to use Canary

Hardware Canary 1.2 will run on any PowerMacintosh or any 68020 or better Macintosh that is equipped with a math coprocessor and adequate memory (RAM; see below). A separate version with identical features, Canary 1.2 LC, is available to run on 680x0 that lack a math co-processor. On Macintosh models that are not equipped with a built-in sound input port, a third-party digitizing device such as MacRecorder is required for sound acquisition.

Canary 1.2 requires a minimum of 2.5 megabytes of RAM to run.[1] This amount of memory is sufficient to make a medium-resolution spectrogram (resolution = 5.8 mS x 43.5 Hz) of a signal slightly more than 2 seconds long, digitized at 22.3 kHz. The actual amount of memory needed for any particular application is roughly linearly dependent on the number and length of signals that are to be open at one time, the sampling rate at which the signals are digitized, and the time and frequency resolution of spectrograms that are calculated. Although there are hundreds of combinations of spectrogram parameters for a given signal that will require different amounts of memory, the following formulas can serve as a guide to the approximate amount of memory needed for signals of different lengths digitized at 22.3 kHz (remember that 1000 Kbytes = 1 Mbyte):

Waveform only:

≈ 1700 Kbytes + (100 Kbytes/sec * signal duration)

Waveform + medium-resolution spectrogram
(e.g., resolution = 5.8 mS x 43.5 Hz):

≈ 1700 Kbytes + (300 Kbytes/sec * signal duration)

Waveform + high-resolution spectrogram
(e.g., resolution = 1.4 mS x 10.8 Hz):

≈ 1700 Kbytes + (2900 Kbytes/sec * signal duration)

These figures should be regarded only as rough indications of memory requirements.

---

[1]This figure refers to memory available for Canary itself, not the total amount installed on your machine. Keep in mind that system software can take up several megabytes, depending on what utilities and system extensions are installed.

Software Canary 1.2 requires Mac System 7.5 or higher. On Mac OS X, Canary runs in classic mode, but not all Canary features will work on Mac OS X. We recommend Raven for sound analysis on Mac OS X.

How to use this manual

What you need to know This manual assumes only minimal knowledge or expertise in acoustics or acoustic analysis. It does assume that you are familiar with the basic features of the Macintosh user interface. For example, you should know how to use the mouse, manipulate windows, open files, and work with dialog boxes, and you should be familiar with some basic Macintosh terminology. If you have never used a Mac before, consult the documentation that was supplied by Apple with your system, or someone who is familiar with the Mac.

Although you do not need an extensive background in acoustics or signal processing to use Canary, an understanding of some fundamental concepts in these areas will enable you to make better use of many of Canary's features. If you are not familiar with the basic principles of digital sound recording and spectrum analysis, you should read Appendices A and B once you've become familiar with basic Canary operations. Appendix A explains how sound is recorded and represented digitally, and includes discussions of sampling rate, Nyquist frequency, aliasing, amplitude resolution, dynamic range, and memory requirements. Appendix B is a nonmathematical introduction to the fundamentals of digital Fourier spectrum analysis (the process that Canary uses for the calculation of spectrograms and spectra), including some of the limitations and trade-offs that are intrinsic to this process. If you plan on using Canary for measuring absolute signal amplitudes (intensities or sound pressures), you should be familiar with the concepts of sound amplitude measurement discussed in Appendix C. Appendix Z is a technical discussion of how Canary performs spectrum analysis and correlations, intended for engineers or others with strong mathematical or signal-processing backgrounds.

How this manual is organized Chapter 1 is a hands-on introduction and quick tour of Canary's major features. We recommend that new Canary users go through Chapter 1 at the computer before reading later chapters, trying out the various features of the program as they're introduced. Chapter 2 discusses signal acquisition (digitizing) in more detail. Most users will want to read Chapter 3, which discusses how to make spectrograms and spectra with Canary. Chapters 4 through 11 go into more detail about specific topics, and can be read or referred to in any order. Chapter 12 is a reference chapter that provides a synoptic description of every command, dialog item (e.g., buttons and checkboxes), graphic control, and icon used by Canary. These later chapters assume that you've read Chapter 1 and are familiar with the features and terms introduced there.

If you need to learn about a process or how to do something (e.g., recording to disk, making a spectrogram...) look in the chapter that deals with that general topic; if you want to know what some individual *thing* does (e.g., what does the “Continuous recording” checkbox do?), the quickest way to find the answer is to look in Chapter 12.

> Although this manual can serve as an introduction and reference for using Canary and can provide a basic background on the conceptual foundations of the program, *the best way to learn about Canary is to experiment with it*, using the program to examine a variety of sounds, especially the type of sounds that you work with. Have fun!

Technical assistance

If you have a question or problem with Canary, consult Appendix D, which covers some of the most common questions about Canary. If you still need help, you can request technical assistance from other users of Canary by visiting the online forum for Canary users at http://Canary.RavenSoundSoftware.com/. The Bioacoustics Research Program does not provide technical assistance by phone or email for the free version of Canary. The software is provided without charge “as is” in the hopes that some researchers may still find it to be a useful tool for sound analysis.

Chapter 1 Getting Started

About this chapter

This chapter introduces the major features of Canary by leading you through a set of examples using one of the sound files supplied with the program. Many details of how these features work are deferred until later chapters. We recommend that you read this chapter at the computer and try the examples as they're discussed.

In this chapter you'll learn how to:

- install Canary on your hard disk
- open an existing sound file
- play back sounds
- make spectrograms and spectra
- work with the waveform, spectrogram, and spectrum displays
- make and save measurements of a signal
- edit a signal
- record new signals
- save new or edited signals.

Installing Canary

NOTE

Canary 1.2 requires System 7.5 or higher to run.

Which version? If you are working on a 680x0 Macintosh with a math co-processor (also known as a floating-point unit or FPU) or on a PowerMac, you should use Canary 1.2. If your Macintosh does not have an FPU, you should use Canary 1.2 LC. If you're not sure whether or not your Macintosh has an FPU, try installing and launching Canary 1.2 as described below. If there is no FPU installed, you will see an error message. You should then remove Canary and install Canary 1.2 LC.

Downloads Canary 1.2 is available on the Bioacoustics Research Program web site at http://www.birds.cornell.edu/brp/ in the section on sound analysis software. The available files are the program executables, examples, and documentation. The executables consist of a single file called "CanaryExectuable.sea.hqx", which when downloaded will become CanaryExectuable.sea". This file is also known as the Canary Installer. To install Canary, download the Canary Installer to your hard disk, and double click on the Installer. The Installer (a "self-extracting archive" created with StuffIt Deluxe) presents a dialog box asking you to specify a folder into which the archive contents will be extracted. When you click on the Save button, the Installer creates an application called "Canary 1.2.4" inside the folder that you specify in the dialog box.

Canary examples are provided on the web site as a self-extracting archive called "CanaryExamples.sea.hqx". When downloaded and run, this archive presents a dialog box asking you to specify a folder into which the archive contents should be extracted. This will create a folder called "Examples 1.2.4" within the folder that you specify.

Canary documentation is provided on the web site as a PDF file that can be viewed and searched online or printed for easy reference. An additional PDF file and readme text file are provided to include extra information that was not available when this manual was written.

If you have an older version of Canary installed on your hard disk, be sure to throw it out before using Canary 1.2 or Canary 1.2 LC.

The ReadMe file (along with any other text files that may be on the Canary web site) contains information about Canary 1.2 that is not in this manual. (If the ReadMe file does not open when you double-click on it, you should be able to open it from within most Macintosh word processors.)

Opening a sound file

Launch Canary by double-clicking on its icon. Canary starts by automatically presenting you with an Open File dialog box. Double-click on the folder named "Examples 1.2"; the dialog box should look like that shown in Figure 1.1. (You can also open a document that was created by Canary by double-clicking on its name or icon in a Finder window, in which case Canary skips the Open File dialog box.)

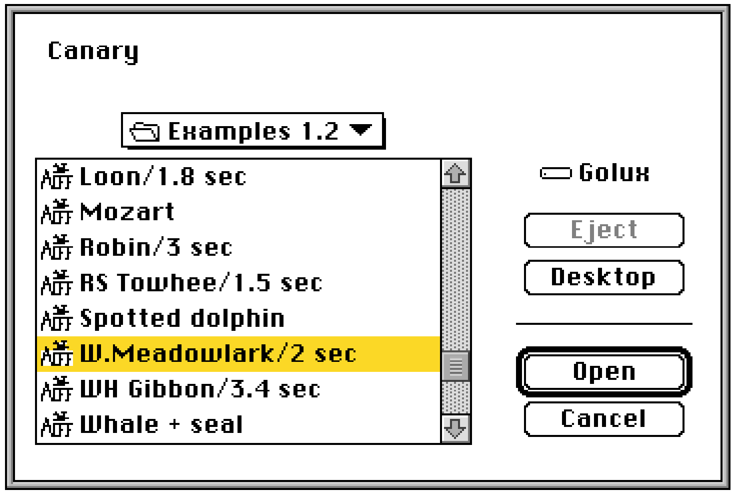

Figure 1.1. Canary's Open File dialog box.

Select the file "W.Meadowlark/2 sec" (which contains a song of a western meadowlark) by double-clicking on its name (or you can click once, then click on the Open button).

The signal window

Two windows will appear on your screen: the signal window and the command panel. The *signal window* shows the waveform (an oscillogram, or graph of amplitude versus time) of the sound stored in the file named W.Meadowlark (see Figure 1.2). The waveform is one of the three *views* of a signal that you can work with in Canary. The units used on the axes are indicated in the lower lefthand corner of the window: in this case, pascals (µPa) for the amplitude (sound pressure) axis and milliseconds (mS) for the time axis.

Canary can display amplitudes of signals using either acoustic units (such as pascals) or electric units (such as volts). In either case, **the absolute amplitude values that Canary displays are meaningful only if the signal has been calibrated**. Chapter 4 explains how to choose electric or acoustic measurements and how to calibrate signal amplitudes.

At the scale of magnification shown, you can't see individual cycles of oscillation of the waveform; what you see is the envelope of the entire signal. Later in this chapter, you'll learn how to zoom in to see fine details of a waveform.

The area above and to the right of the axes is called the *plot area*.

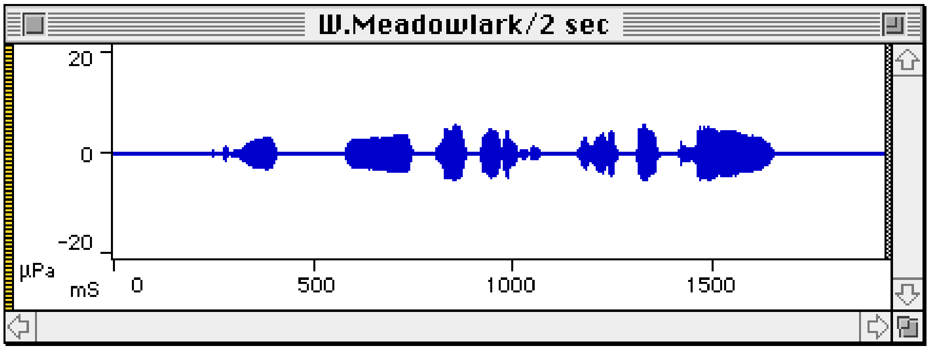

Figure 1.2. Waveform of song of western meadowlark.

The command panel

The second window on the screen is Canary's *command panel*. This window contains controls for manipulating the signal display and for playing back the signal. Figure 1.3 identifies the various buttons on the command panel. The spectrogram brightness and contrast controls appear only once you've made a spectrogram (see below). In the rest of this chapter, we'll explore how these controls work.

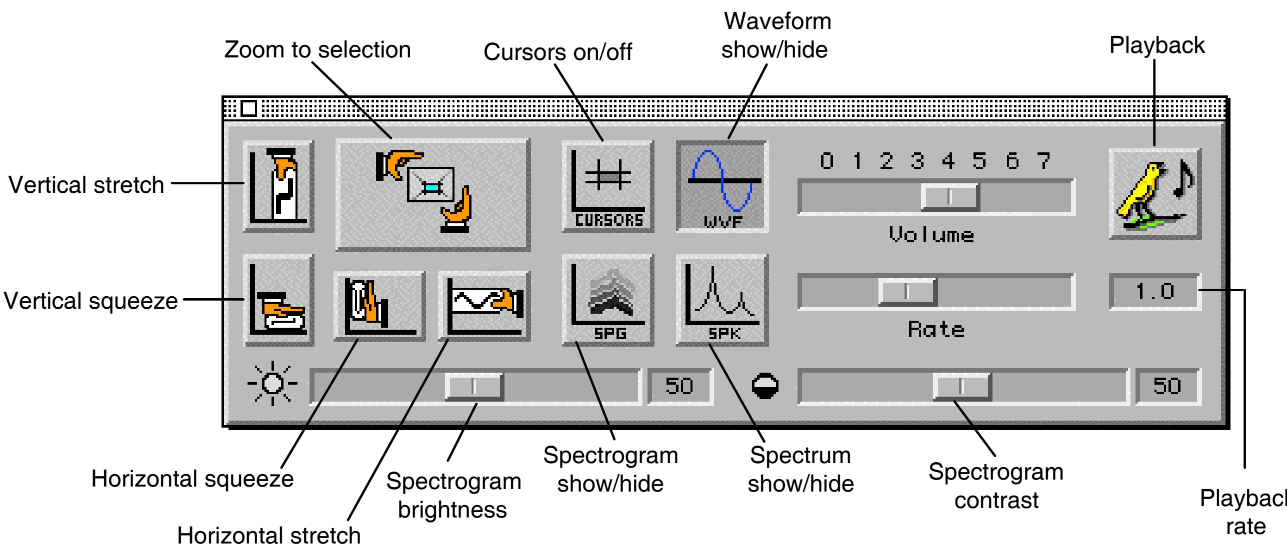

Figure 1.3. The command panel.

Playing back a sound

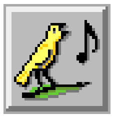

Any time that you have a signal window open, you can play the signal back by clicking on the playback button. Try this now—you should hear the song of a western meadowlark from your Mac's speaker. As the sound is played back, a vertical cursor moves across the waveform in synchrony with the playback. You can adjust the playback volume using the Volume slider control on the command panel (use the mouse to drag the slider from one position to another). You can also slow down or speed up the playback using the Rate control, which has twenty positions ranging from 1/10 normal speed up to the fastest rate that your Macintosh can play.[1] The actual playback rates depend on the sampling rate with which a signal was digitized.

The currently selected playback rate is shown in the display next to the slider. You can also set a playback rate by double-clicking on the rate display and typing any decimal number followed by the *return* or *enter* key. You can reset the rate to 1.0 by holding the *option* key while clicking on the rate slider. Try playing the sound back at different speeds.

[1]As of June 1995, the fastest sample rate at which any Mac can play back sounds is 64 kHz.

Selecting part of a signal

You can use the mouse to select a part of the sound to play. Move the mouse pointer over the waveform; the pointer appears as a crosshair whenever it is on the plot area of the window. Position the pointer at one end of the section you want to play, and while holding the mouse button, drag the pointer across the section you want to play. The section you select will be highlighted. (On color or grayscale monitors, you can use the Macintosh's Color control panel (on the apple menu) to set the color or shade of gray used for highlighting. On monochrome monitors, highlighted sections are displayed in reverse video.) Now if you click on the play button, only the highlighted section is played. Another way to select part of a waveform is to click once at one end of the desired section, and then hold down the *shift* key while clicking at the other end of the selection. You can select the entire signal by double-clicking anywhere on the waveform. (You can also choose Select All from the Edit menu or press Command-A.)

Creating a spectrogram

Click on the spectrogram (SPG) button to create a spectrogram (sonagram) view of the entire signal. A dialog box like the one in Figure 1.4 appears.

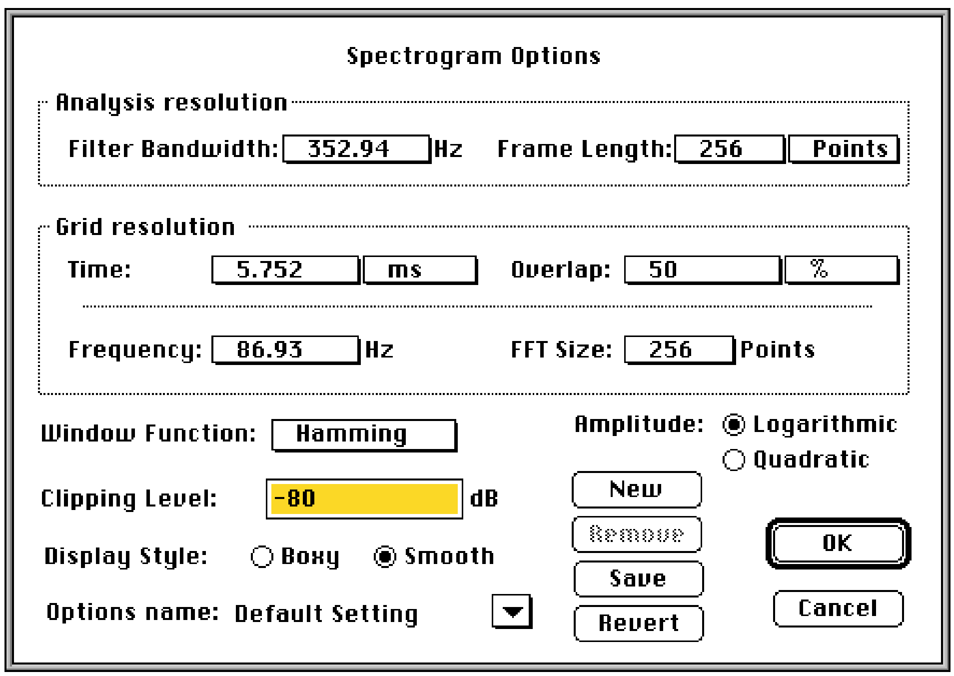

Figure 1.4. The Spectrogram Options dialog box.

This dialog box lets you specify various parameters that affect the final appearance of the spectrogram. For now, don't worry about what these parameters mean; they are explained in detail in Chapter 3. Click OK to make the spectrogram.

A status window will appear indicating the progress of the calculation of the spectrogram. Many factors affect how long it takes to compute the spectrogram, including the processor speed of the Macintosh, the length of the signal, and the parameters that you specify in the Spectrogram Options dialog box. You can switch to another application program and Canary will continue to calculate the spectrogram in the background. (If you want to cancel the calculation of a spectrogram, click on the Pause button in the status window; the Cancel button will then become available.)

The spectrogram pane

When the calculations are completed, the signal window is redrawn, with two separate *panes*. The lower pane shows the waveform; the upper pane shows the spectrogram. The spectrogram displays frequency on the vertical axis and time on the horizontal; the time axes of the waveform and spectrogram panes are aligned with each other. The darkness of the spectrogram at any given time-frequency point is proportional to the logarithm of the sound intensity at that point. The spectrogram is thus a type of three-dimensional plot of intensity level against frequency and time, with the intensity levels represented as shades of gray varying between white (minimum amplitude) and black (maximum amplitude). Later in this chapter you'll see how to adjust the minimum and maximum amplitudes displayed in a spectrogram. The highest frequency that can appear in a spectrogram depends on how the signal was originally digitized (recorded) by the computer (see Appendix A).

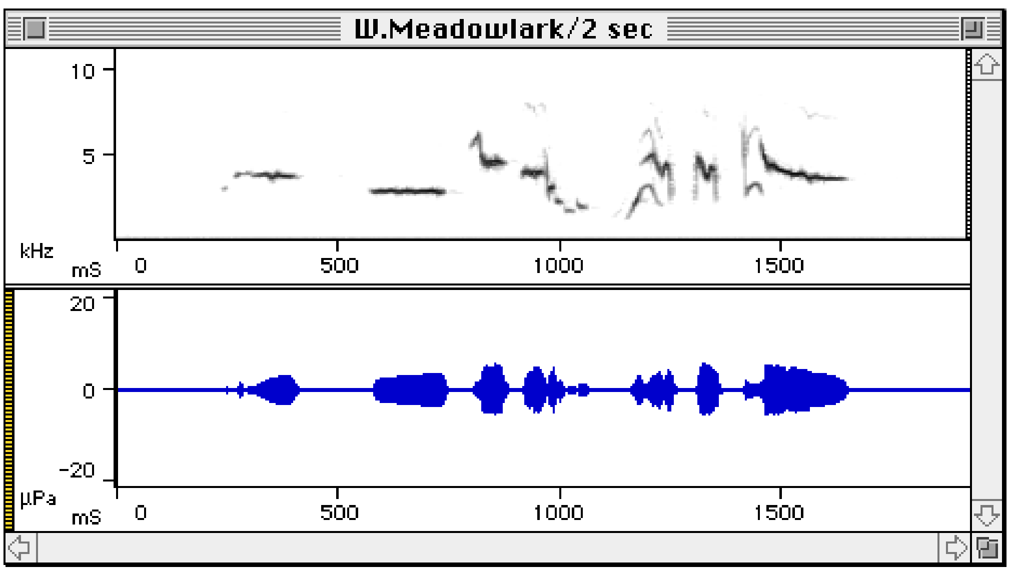

Figure 1.5. The signal window showing spectrogram and waveform panes.

The SPG button on the command panel is now depressed and colored (if you have a color monitor), indicating that the spectrogram display is turned on. If you click the button again, the spectrogram is turned off, and the window is redrawn containing just the waveform pane. Once a spectrogram has been computed, you can turn the spectrogram display off and on again, without the delay of recalculating it. (Canary saves the spectrogram even when it's not shown on the screen.)

Spectrogram contrast and brightness

Whenever a spectrogram is displayed, the contrast and brightness controls on the command panel are available. The contrast control (the righthand slider) adjusts the number of different grayscale values that are shown in the spectrogram between black and white. The brightness control adjusts the overall darkness of the spectrogram: sliding the control to the right lightens the display. If you move the contrast all the way to the right, the spectrogram becomes black and white; all intensity levels above some threshold will be black, and all others will be white. In this case, adjusting the brightness control shifts the white/black threshold intensity. See Chapter 3 for a more detailed explanation of how these controls work.

You can also control contrast and brightness by double-clicking the numeric fields to the right of each control and typing a number between 0 and 100, then pressing *return* or *enter*.

Adjusting the signal display

Moving and resizing windows

You can position the signal window and the command panel wherever you want on the screen in the usual Macintosh way, by dragging the horizontal drag bar at the top of the window (the drag bar is the title bar in the signal window, the stippled bar in the command panel). You can adjust the size of the signal window in the usual way as well: by dragging the size box in the lower right-hand corner, or by clicking in the zoom box in the right-hand end of the title bar. For now, adjust the window so that it takes up about three-quarters of the screen, with the command panel positioned so that you can see the whole signal window. (The command panel always floats in front of the signal window. If it’s in the way, you can get rid of the command panel by clicking in its close box; you can always retrieve it by selecting Command Panel from the Windows menu or typing Command-T.)

Resizing panes

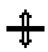

If you position the mouse on the horizontal bar that separates the panes, the mouse pointer changes to the icon shown at left. You can then change the relative vertical sizes of different panes in the signal window by dragging the horizontal bar that separates the two panes.

The active pane

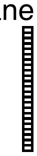

At any time, one of the panes in the signal window is the *active pane*, as indicated by a horizontally hatched vertical bar along its left edge. Some commands apply only to whichever pane is the active pane. You can activate a pane by clicking anywhere in it. To activate a pane without selecting part of the plot shown in it, click below or to the left of the axes.

Showing and hiding panes

Clicking on the command panel’s waveform (WVF), spectrogram (SPG), or spectrum (SPK) buttons will hide the corresponding pane, if that pane is presently displayed. If a pane is hidden, clicking on its button shows the pane. If no spectrogram or spectrum has been calculated yet, clicking the corresponding button brings up the dialog box that lets you specify the spectral analysis options.

Stretching and squeezing the display

The command panel’s stretch and squeeze buttons let you adjust the horizontal and vertical scales of the plot shown in the active pane. To see more detail in the frequency dimension of the spectrogram, first activate the spectrogram pane, then click on the vertical stretch button.

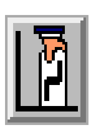

When you click on the vertical stretch button, the spectrogram pane is redrawn with the vertical dimension stretched by a factor of two (in the case of the western meadowlark spectrogram, the vertical axis shown now ranges from 0 to 5.5 kHz). The waveform pane remains unchanged. Notice that the vertical scroll bar has now become active; you can use it to scroll the spectrogram display up and down so that you can see whichever portion of the total frequency range you choose. Scroll bars in the signal window always refer to the active pane; each scroll bar is active only when the active pane does not display all of the available data.

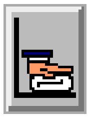

Clicking on the vertical squeeze button compresses the vertical dimension of the active pane by a factor of two, reversing the effect of the previous vertical stretch. If you activate the waveform pane, and click on the vertical stretch or squeeze buttons, the waveform display will be redrawn as appropriate, without any change to the spectrogram pane.

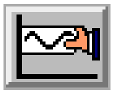

Click on the horizontal stretch button, and notice that the displays in both the waveform and spectrogram panes are stretched, irrespective of which is the active pane. A general rule in Canary is that changing the scale of either dimension of a pane also affects any other pane that shares that dimension. Since the waveform and the spectrogram share the time dimension, any change to the time scale of either pane affects both panes.

You can stretch or squeeze the display as many times as you want. Figure 1.6 shows the signal window after it has been stretched seven times (since each stretch is by a factor of two, the scale has thus been stretched by a factor of 2^7 = 128). At this magnification, individual cycles of the waveform are visible. Compare the time axes in Figures 1.6 and 1.5; you can see that the segment shown in Figure 1.6 is in the second syllable of the song. (If you try zooming in on a single syllable by successively stretching the time axis in this way, you may need to use the horizontal scroll bar to get to the part of the song shown in Figure 1.6.) If you keep zooming in on shorter and shorter segments of the signal, eventually you will see discrete points plotted along the waveform; these are the individual digital samples of which the digitized signal is comprised (see Appendix A for more on digital sampling).

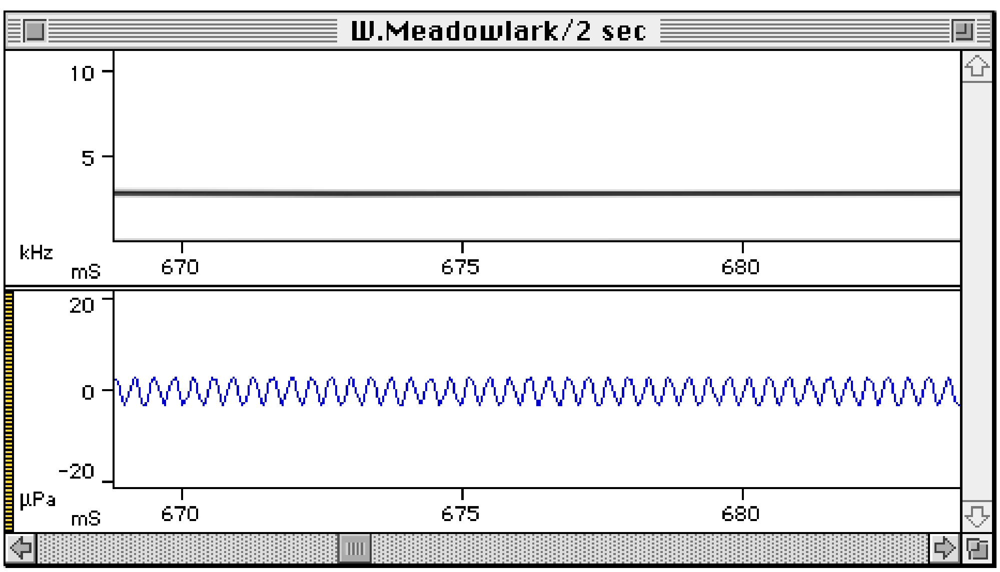

Figure 1.6. Part of a single syllable of the meadowlark song shown Figure 1.5, after seven successive horizontal stretches.

Zooming in to a selection

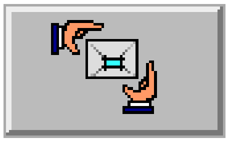

To stretch or squeeze the display so that a particular portion of the signal fills the window without repeatedly using the stretch button, you can use the Zoom to Selection button. Select part of the waveform with the mouse, and click on the Zoom to Selection button: the selected portion will be stretched (or compressed) to fill the pane. If you have already stretched the display (so that the whole signal no longer fits in the pane), you can restore the display scale to its original setting (so that the signal just fills the pane) by double-clicking on the waveform (equivalent to choosing Select All from the Edit menu or typing Command-A), and then clicking the Zoom to Selection button.

In the spectrogram pane, you can select a specific frequency range as well as a specific time interval, and zoom to the selected region.

Boxy vs. smooth spectrograms

If you stretch the horizontal or vertical scale of a "boxy" spectrogram (the default display style) much, you'll find that the spectrogram display, which may have seemed smooth as it was originally drawn, becomes increasingly jagged or boxy. Canary displays one box of a specific shade of gray for each point in time-frequency space where an amplitude has been calculated from the signal. The height and width of the boxes are determined by the parameters set in the Spectrogram dialog box, as discussed in Chapter 3. If you want the display to appear smooth at any magnification, you can select the Smooth display style instead of Boxy in the Spectrogram Options dialog box.

If you've already created a spectrogram, you can still switch back and forth between boxy and smoothed displays without recalculating the spectrogram. Hold down the option key while clicking on the SPG button, and the Spectrogram Options dialog box will reappear. Click on the Smooth display style button (see Figure 1.4), and click OK. The spectrogram will be redrawn, with grayscale values individually interpolated for every screen pixel in the display. Smoothed spectrograms take longer to redraw after any change of scale than boxy ones. See Chapter 3 for further discussion of the use of boxy and smoothed spectrograms.

Views and panes

As mentioned earlier, Canary can display three different *views* of a signal; each view is shown in a separate *pane* of a signal window. (So far, we have discussed only the waveform and spectrogram views; the spectrum view is discussed in the next section.) It is helpful to understand the distinction between views and panes. A view is a representation of a signal or part of a signal. A pane is part of a window, and is simply the graphic region of a window. A pane may show only part of a view, for example, when the time axis of a waveform is stretched. Double-clicking on a pane selects the entire view.

Creating a spectrum

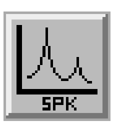

Click once in the waveform pane midway through the second syllable of the song (the constant-frequency whistle at about 600 - 750 msec). Click on the spectrum (SPK) button to create a spectrum view of the signal at this point in time. You can think of a spectrum as being like a thin "slice" through the spectrogram at one moment in time. A dialog box like the one in Figure 1.7 appears.

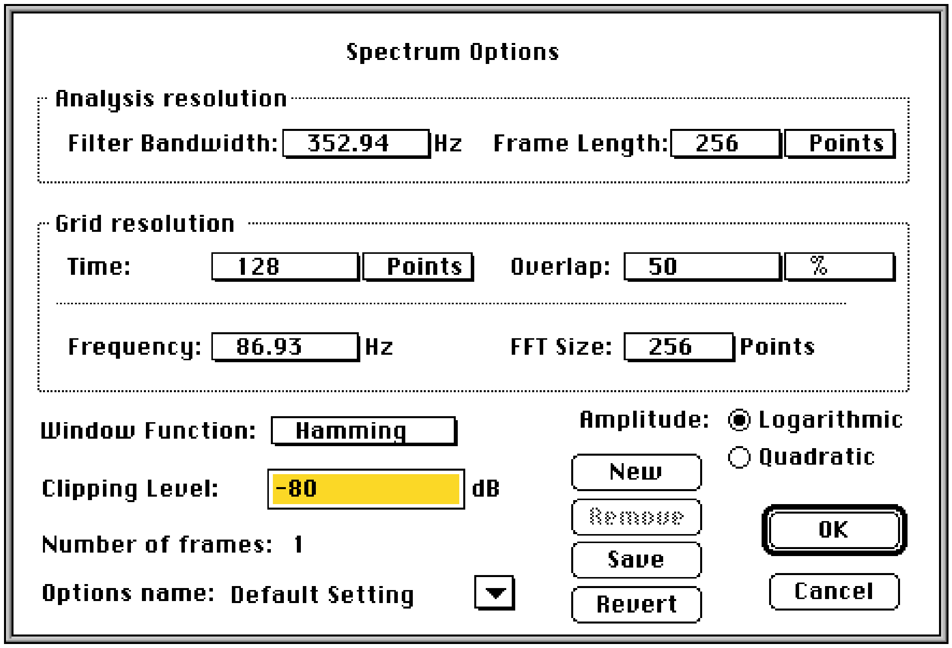

Figure 1.7. The Spectrum Options dialog box.

This dialog box is similar to the Spectrogram Options dialog box that you saw earlier; it's explained in detail in Chapter 3. Click OK to make the spectrum. As before, a status window indicates the progress of the calculations. A spectrum is computed faster than a spectrogram.

The spectrum pane

The signal window is again redrawn, this time with three panes (Figure 1.8). The top pane shows the spectrum of the signal at the time where the blinking insertion point is located in the waveform. The spectrum displays relative amplitude versus frequency.

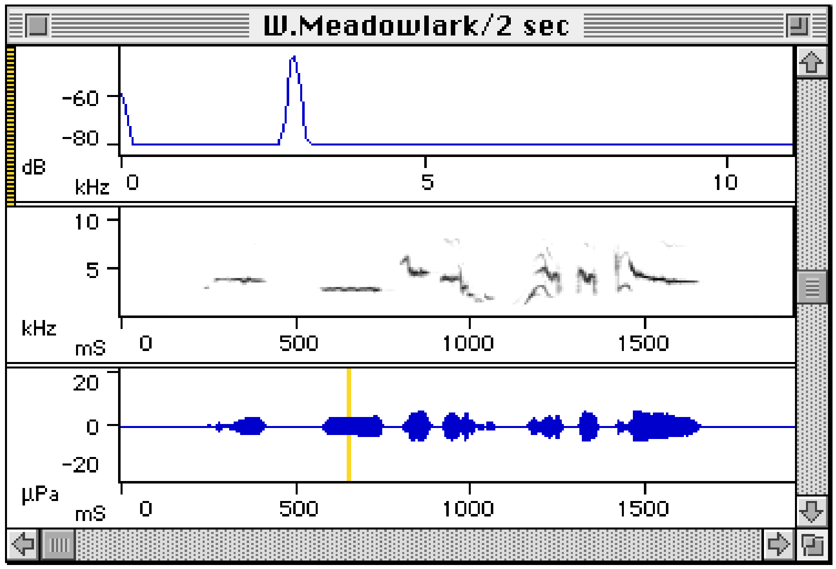

Figure 1.8. The signal window showing spectrum, spectrogram, and waveform panes. The spectrum was made from the highlighted section of the waveform.

The SPK button in the command panel is now depressed and colored, indicating that the spectrum is turned on.

If you click on the spectrum, a thin vertical strip will be highlighted in the waveform, identifying the portion of the signal that was used to make the spectrum (the *spectrum source interval*).

Coupled selections

In general, when you select part of a display in any of the three panes, corresponding portions of the other two panes are highlighted as well. The correspondence between highlighted regions in different panes is based on which dimension (time or frequency) is shared between the panes.

Select part of the spectrum by dragging the mouse pointer across it. In the spectrogram pane, a thin rectangle is highlighted which corresponds to the frequency range selected in the spectrum and the spectrum source interval. In the waveform pane, the spectrum source interval is highlighted.

If you select part of the spectrogram, the corresponding frequency range of the spectrum is highlighted (regardless of whether the selected part of the spectrogram includes the spectrum source interval), as is the appropriate time interval in the waveform.

If you select any part of the waveform, the corresponding segment of the spectrogram is highlighted, and the entire spectrum is highlighted (regardless of whether or not the selected interval includes the spectrum source interval).

Using selection cursors

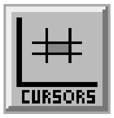

Although you can use the mouse to select portions of a signal, a more precise way of making selections is to use Canary's *selection cursors* (or simply *cursors*). Click once on the waveform display near the middle of the signal; a blinking cursor will appear there. Now click on the CURSORS button to turn the selection cursors on. Clicking on the CURSORS button again turns cursors off.

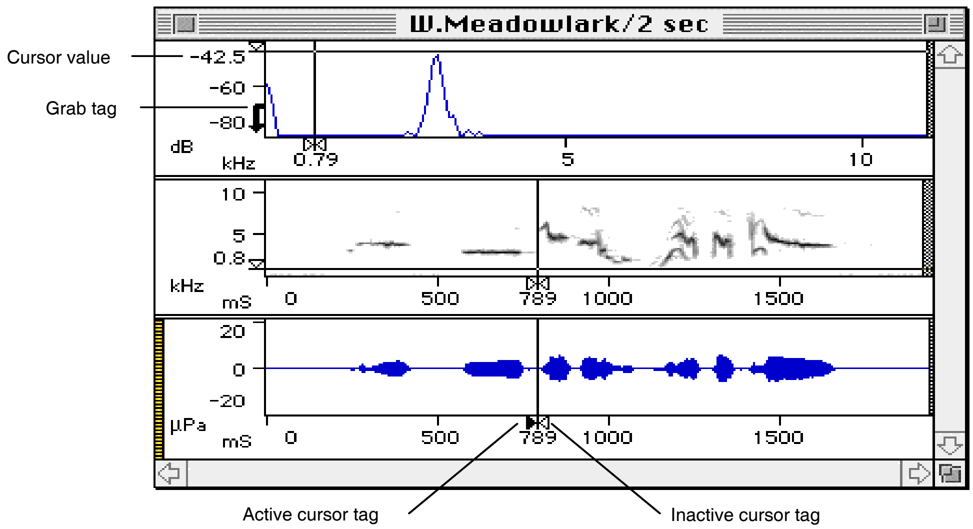

Figure 1.9. Signal window showing selection cursors.

Cursor tags and grab tags

Figure 1.9 shows the window as it appears right after you turn on cursors. At the point where you last clicked in the waveform pane, two cursors (non-blinking vertical lines) appear superimposed on each other. The position of each cursor in time is indicated by a small horizontally pointing arrowhead called a *cursor tag*, just below the time axis. Cursors and cursor tags also appear in the spectrogram and spectrum panes. In those two panes, cursors appear along both the horizontal and vertical axes. A cursor can be placed beyond the end of the visible part of an axis, and so not be shown on the screen. When this happens, a *grab tag* appears at one end of the axis, pointing in the direction of the off-screen cursor.

The active cursor

At any time, one cursor in each direction (horizontal and vertical) in the active pane is the *active cursor*; its tag is filled (black), whereas all others are open. Each active cursor can be moved using the arrow keys on the keyboard. You select which cursor is active in each direction by clicking on one of the tags. Try using the horizontal arrow keys to move the active cursor in the waveform pane. You can also move a cursor by dragging its tag with the mouse.

The *cursor value* is displayed along the axis immediately below or next to the cursor's tag. As you move a cursor, the cursor value is updated continuously to show the cursor's exact position in whatever units are used to label the axis (e.g., seconds or milliseconds on the time axis).

The area between the cursors in the active pane is the current selection and appears highlighted. While cursors are turned on, you can select part of a display either by moving the cursors directly (using the arrow keys or the mouse) or you can hold down the *option* key while dragging the mouse across the region to be selected. Holding down the *option* key causes the cursors to "snap" to the mouse position and to follow the mouse pointer as you drag.

Coupling of cursors between panes

In general, as you move a cursor in one pane, cursors in the corresponding dimensions of other panes move as well. For example, as you move a cursor in the time dimension in the waveform pane, the corresponding time cursor in the spectrogram pane moves in synchrony. No change occurs in the spectrum selection as you move cursors in the waveform pane, since the time dimension is not represented in the spectrum. Similarly, if you move one of the frequency cursors in the spectrogram pane (vertical dimension), the corresponding frequency cursor in the spectrum pane (horizontal dimension) tracks the movement.

Snap cursors to selection

You can make cursors "snap" to outline an existing selection (made with the mouse) by holding down the *option* key while clicking the CURSORS button.

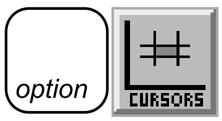

Making and storing measurements

The measurement panel

Activate the waveform pane, and then choose Measurement Panel from the Windows menu. A window like that shown in Figure 1.11 will appear containing the default measurement panel for the waveform pane. The upper row of the measurement panel always shows various *parameters* of the active pane and the signal. The lower row shows the values of *measurements* of the signal for the point where the mouse pointer is located, or for the region selected in the display. As you move the mouse across the signal, the values displayed in the measurement panel change. (If you move the pointer off of the plot area of the active pane, the measurement cells display "--".) Canary lets you make many other measurements of a signal in addition to the ones illustrated in this section; Chapter 6 explains how to configure the measurement panel to display whatever combination of measurements you want.

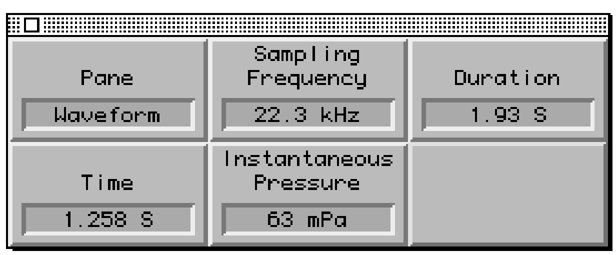

Figure 1.11. The default waveform measurement panel.

Activate one of the other panes, and the measurement panel will be automatically redrawn to display measurements appropriate to the active pane. For example, if you activate the spectrogram pane, the measurement panel looks like the one in Figure 1.12.

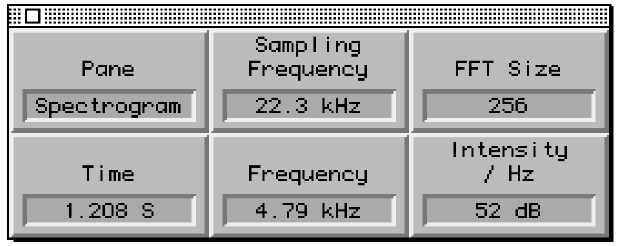

Figure 1.12. The default spectrogram measurement panel.

The data logger

To make a permanent record of values displayed in the measurement panel, use Canary's data log facility. Values in the data log can be exported to other programs (e.g., spreadsheet, statistics, or word processing applications). From the Windows menu, choose Data Log; a new window titled DataLog will appear, displaying several numbered columns. A DataLog menu is also added to the menu bar when the data log is displayed.

Activate the waveform pane, and move the mouse pointer over the waveform so that the measurement panel displays time and amplitude values. You can record the values shown in the lower row of the measurement panel to the data log by holding down the Command key and clicking the mouse button. As soon as you log an entry, the columns in the data logger are automatically labeled; data values are entered with the appropriate units (e.g., mS, mPa). After you log several entries, the data log should look something like Figure 1.13.

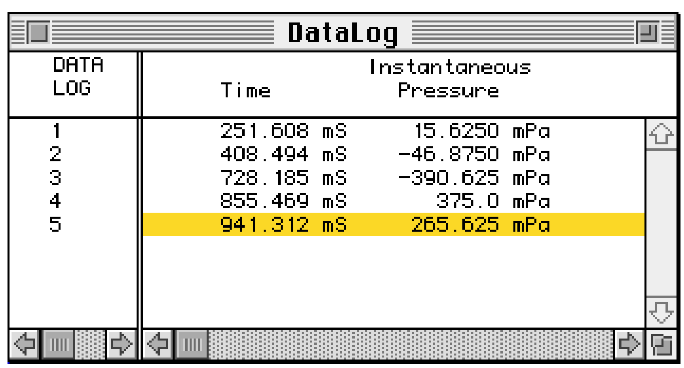

Figure 1.13. The DataLog window after logging some measurements from the waveform pane.

Saving logged data to a text file

You can save the measurements recorded in the data log to a tab-delimited text file that can be opened by most Macintosh word processors, spreadsheet, and statistics programs. Select Save Text Report... from the DataLog menu. The dialog box shown in Figure 1.14 will appear. If you leave the Report Title box checked, the text in the highlighted field is included in the report file. Canary assigns the default report title "LoggedData.Report" but you can change it to anything you want. The remaining items in the dialog box allow you to control the format of entries in the report; these items are discussed in detail in Chapter 6.

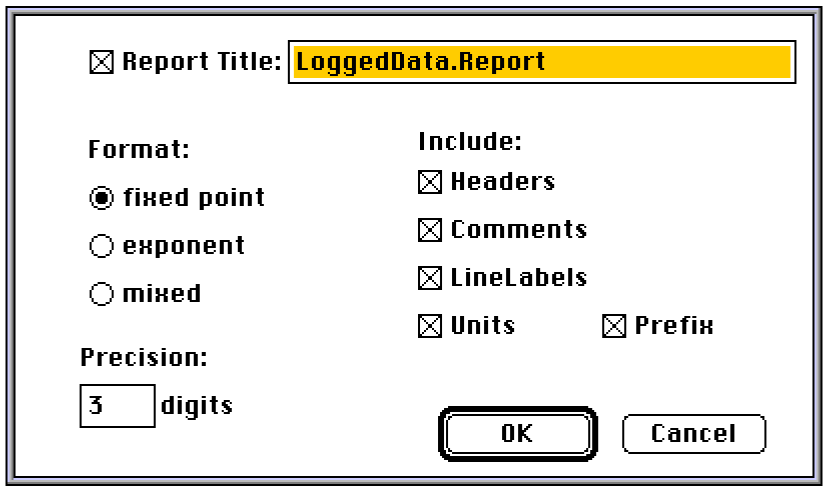

Figure 1.14. The Save Text Report dialog box.

When you click OK, you'll be presented with a standard Save As dialog box so that you can specify the name for the report file, and where you want it to be saved. (The default file name that appears in the Save As dialog box is the same as the Report Title that you gave in the preceding dialog box, but you can change it without affecting the title that is written in the file.)

Once the report file is saved, you can open it from within any program that can read plain text files.

You can also save logged data in other formats, which are discussed in Chapter 10.

Editing a signal

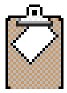

Canary lets you cut, copy, paste, and delete (clear) parts of a waveform using commands on the Edit menu or their standard Macintosh keyboard equivalents. When you cut or copy a selected section, a copy of the selection is put in the clipboard; in a cut, the selected section is then deleted from the signal, and the window is redrawn to reflect the change. The Paste command inserts the contents of the clipboard where the blinking cursor or the current selection is located; if a section of the signal is selected when you paste, the clipboard contents replace the selected section. You can use the Show Clipboard command on the Edit menu to see and hear the contents of the clipboard. Remember that the clipboard can hold only one item at a time; whenever you Cut or Copy, the selection that goes into the clipboard replaces what was there before. You can delete a selected section using either the Clear command on the Edit menu or the *delete* key. You can reverse the last editing operation using the Undo command (Command-Z).

All editing operations take place only in the time dimension of the waveform, even though the corresponding section of the spectrogram will be highlighted when you select sections of the waveform for editing. After you paste data into a waveform, the segment of the spectrogram that corresponds in time to the inserted segment will be blank. In order for a spectrogram to reflect insertions into the waveform, you must recalculate the spectrogram. To recalculate a spectrogram, hold down the *option* key while clicking on the SPG button on the command panel, or choose Make Spectrogram... from the Panel menu (or type Command-G).

Figure 1.15 shows the meadowlark song after a series of editing operations (and after recomputing the spectrogram). As an exercise, you might try to rearrange parts of the song to make it look like the version in the figure.

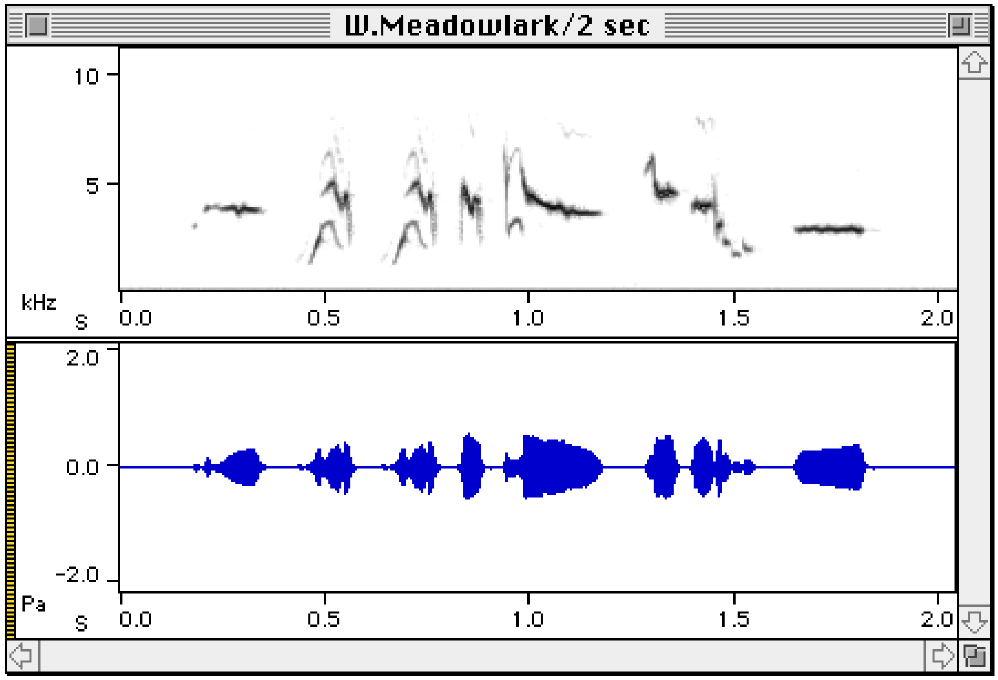

Figure 1.15. Edited western meadowlark song, with syllables rearranged by cutting, copying, and pasting.

You can also cut, copy, and paste between signals in Canary and other application programs that can manipulate sound data.

Filtering a signal

You can filter all or part of a signal to eliminate energy in any specified frequency band. Canary provides two ways of specifying the frequency range to be filtered, and two different algorithms for performing the filtering commands on the Edit menu, Filter Around... and Filter Out... . Filter Around removes all energy in the highlighted time interval that is outside the highlighted frequency band on the spectrogram. Filter Out removes all energy in the highlighted time interval that is inside the highlighted frequency band.

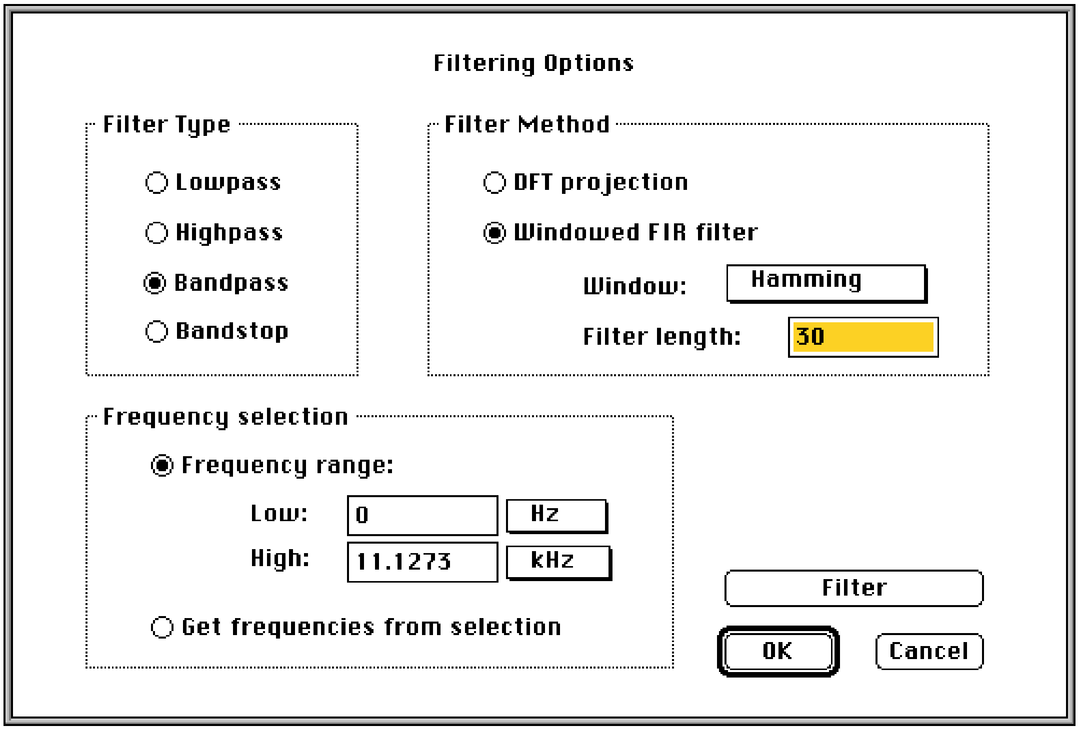

Figure 1.16 illustrates filtering using the example file “Whale + seal”, which contains two bowhead whale calls (the second call is from a more distant animal than the first) and part of a song of a bearded seal. The bowhead calls are in the band between approximately 75 and 200 Hz; the seal call is entirely above 200 Hz (Figure 1.16a). Part (b) of the figure shows the signal after filtering around the 75-200 Hz band over the entire length of the signal: the seal call (and low-amplitude noise not visible in the spectrogram) are eliminated.

When you filter a signal, the spectrogram does not change until you recalculate the spectrogram (by *option*-clicking the SPG button, typing Command-G, or selecting Make Spectrogram... from the Panel menu).

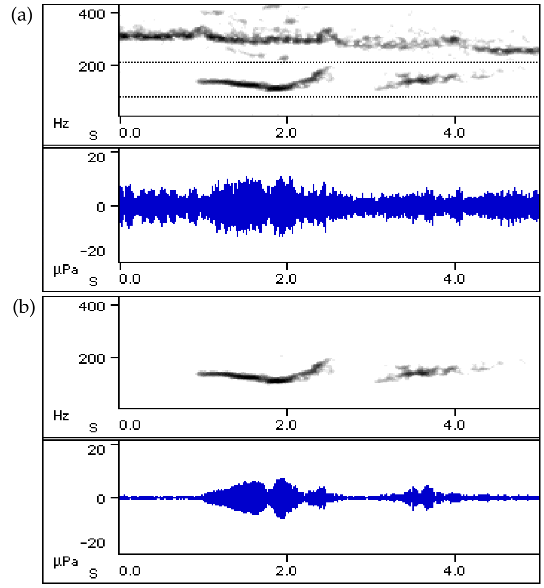

Figure 1.16. Example of filtering around a signal. **(a)** Recording of bowhead whales and bearded seal from the Alaskan Arctic. There are two whale calls, between approximately 75 and 200 Hz (dotted lines). The call that descends from 315 to 250 Hz over the entire signal is part of a song of a bearded seal. **(b)** Same signal after performing a Filter Around operation on the selection encompassed by the dotted lines in (a).

Saving a signal

data in Canary can save several file formats. A file saved in Canary's own format can contain any combination of the three views (waveform, spectrogram, or spectrum) of a signal. A file saved in any of the other formats can contain only one of the three views. Canary files are sometimes larger than files in other formats.

When saving signals that you have recorded or edited, it's a good idea to save the waveform (which is included automatically if you save the entire window), even if what you're most interested in is the spectrum or the spectrogram. You can always generate spectrograms and spectra from a waveform, but you can't reconstruct the waveform from the spectral displays, nor can you play back or edit them.

To save all of the views of a signal in a Canary file, choose Save Window... or Save Window As... from the File menu. If you choose Save As... or if the window has not been given a name yet, Canary's version of the Macintosh standard file dialog box appears, so that you can specify a name and a location for the file to be saved (Figure 1.17). If you choose Save Window (or type Command-S), and the window has been previously saved, it is simply re-saved without any dialog box.

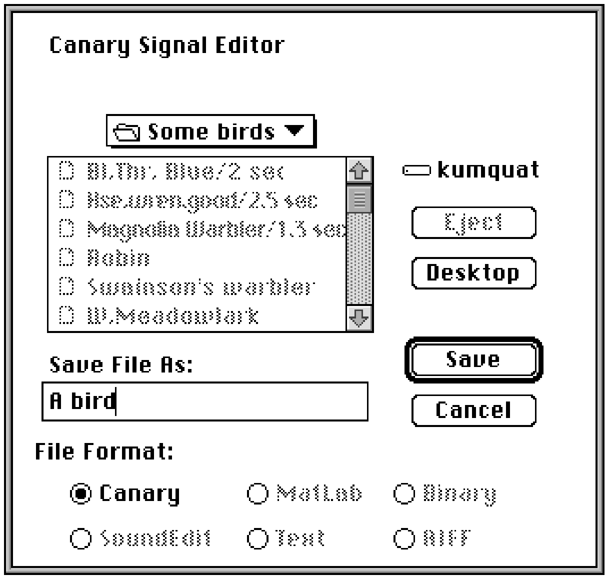

Figure 1.17. Canary's Save As... dialog box. The buttons at the bottom let you specify the file format.

To save a single view of a signal, activate the corresponding pane, and choose Save *Pane*... or Save *Pane* As... from the File menu (the menu will actually show Save Sound, Save Spectrogram or Save Spectrum, depending on which pane is active).

Depending on which pane is active, different combinations of radio buttons will be available at the bottom of the dialog box to let you specify the format of the file to be saved. Canary format is the default file type for all three panes. For sound (waveform) files, SoundEdit or AIFF formats result in the smallest files and are the fastest to open in Canary. File formats are discussed further in Chapter 10.

Working with more than one signal

Canary lets you work with more than one signal at a time. The actual number of signals you can have open at once is limited by the amount of memory available to Canary, the length of the signals, and how many spectrograms and spectra you've built. Each pane in a signal window requires memory for the storage of the data displayed in it, so a signal window that has never had a spectrogram or a spectrum computed for it takes less memory than the same signal would with either of these views.[1]

---

[1]The issue in determining memory requirements is whether these views have been computed, not whether they are currently displayed. Once they've been computed, Canary saves the resulting data even if you turn off the spectrogram or spectrum panes, so that the pane can be displayed again later without re-computation.

When more than one signal window is open, the command panel always refers to the active window (the one window in which the title bar is striped). As usual in the Mac world, you activate a window by clicking anywhere in it with the mouse. You can also activate a window by choosing its name from the Windows menu (this is handy if you can't find a window because it's hidden behind others).

You can paste portions of signals copied or cut from one window into another, provided that the two signals were recorded with the same sampling rate. (Sampling rate is discussed in Appendix A.)

See the section on "Calibrations and signal editing" in Chapter 4 for a discussion of pasting data between signals that are calibrated differently.

Recording a signal

This section provides a brief introduction to recording sounds with Canary. Chapter 2 discusses recording options in more detail.

If your Macintosh is equipped with a sound recording device, you can use Canary to record new signals. Recent Macintosh models (including Performa, Centris, Quadra, Powerbook, and PowerMac models) come with a built-in sound recording port that can take input either from a microphone (supplied with the Macintosh) or from a line input connection, which can be connected to any standard audio output device. If your Macintosh does not have a built-in sound input port, you can add recording capability with a MacRecorder (from MacroMedia, Inc.).[1]

Close any signal windows that are presently open on the screen, and select Record... from the Process menu.[2] A dialog box similar to the one shown in Figure 1.18 will appear. Depending on what model of Macintosh you are working on, the dialog box you see may differ slightly from the one shown in the figure. The different forms of controls that appear on different Mac models are discussed in Chapter 2.

---

[1]Although Canary may also work with other third-party recording devices for the Macintosh, MacRecorder is the only device with which the program has been thoroughly tested.

[2]You don't have to close open windows before recording; however, doing so makes more memory available, which allows you to record a longer signal.

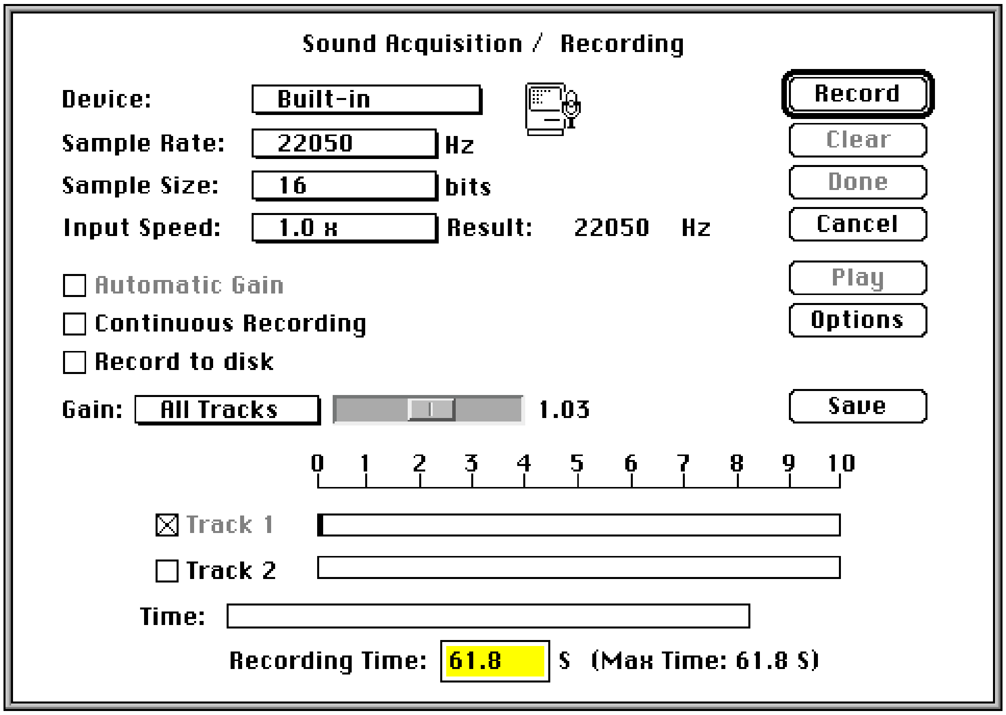

Figure 1.18. The Sound Acquisition/Recording dialog box. This dialog box may look slightly different on different models of Macintosh.

Choosing a recording device

The box labeled Device displays the name of the currently selected recording device. Press the mouse button on this box to display a pop-up menu listing all of the recording devices available on your machine (there may only be one such device). Sample Rate, Sample Size, and Input Speed are discussed in Chapter 2; for now, you can leave them set as they are.

The Automatic Gain, Continuous Recording, and Record to disk options are discussed in Chapter 2.

Recording time You can specify the length of the recording to be made by entering any value into the Recording Time field, up to the Max Time shown. The time available for recording is determined by the total amount of memory (RAM) available to Canary and by how much of that memory is presently free (i.e., not occupied by other signals, spectrograms, or other data). If you fill all of the available recording time, you probably will not have enough free memory left to make a spectrogram. Limiting the recording length to about half of the maximum time usually allows room for a spectrogram using Canary's default spectrogram parameters.

Recording level The bars labeled Track 1 and Track 2 indicate the strength of the incoming signal, like a recording level meter or LED display on a tape recorder. The Track 2 bar is active only if Track 2 is checked (not all Macintosh models are capable of stereo recording; see Chapter 2). The level indicator is active as soon as the dialog box appears on the screen. Use this indicator to adjust the level of the signal to get the best recording. If your Macintosh is equipped with adjustable gain control, you can use the slider in the dialog box to adjust the recording level. You can also set the level by adjusting the output level of the device (e.g., tape recorder, MacRecorder) you're recording from. If you use the built-in microphone, adjust the loudness of the sound source or the distance between the microphone and the source. You'll get the best recording by adjusting the level so that the strongest parts of the signal almost fill the level indicator bar. If the level indicator fills all the way, your recording will be distorted.

Recording Once you've adjusted the level, click on the Record button to start recording. As you record, the Time indicator bar fills to show how much of the specified recording time has been used, and a digital display appears at the right-hand end of the indicator bar to show you exactly how much time has been recorded. While recording, you can click the Pause button (which replaces the Record button) at any time to suspend recording and display the recorded signal in a new window. While recording is paused, you can Play the sound that has been recorded so far, Resume recording, or Clear the recording. Clicking the Cancel button stops recording and discards whatever sound has been recorded so far. Clicking the Done button closes the dialog box and brings up the waveform of the new signal in an untitled window.

Once the signal appears in the new window, you can work with it as you would with any previously saved signal.

Quitting and preferences

You can quit Canary at any time by choosing Quit from the File menu or typing Command-Q. When you quit Canary, the program asks if you want to save changes to the current *preference file*. You can use preference files to customize various how Canary operates. Preference files are discussed in Chapter 8. For now, click the No button when Canary asks if you want to save preferences. (You can specify that you want Canary always to save or never to save preferences, without asking, when quitting the program; see Chapter 8.)

Other capabilities of Canary

This chapter has introduced most of Canary's basic operations. If you've read the entire chapter, you should be ready to start using Canary for your own work. Other parts of this manual provide more detail on many of the features already discussed, and also discuss more advanced features.

In later chapters you'll find information on the following:

- Recording sounds directly to disk, and other recording options (Chapter 2).
- Fine-tuning spectrograms and spectra (Chapter 3).
- Calibrating signal amplitude measurements (Chapter 4).
- Working with multi-track sounds (Chapter 5).
- Customizing measurement panels (Chapter 6).
- Calculating correlations between pairs of signals or spectrograms (Chapter 7).
- Customizing display and other options and saving preferences (Chapter 8).
- Printing and exporting graphics from Canary (Chapter 9).
- Using various file formats (Chapter 10).
- Batch processing operations on multiple files (e.g., making spectrograms of all the files in a folder automatically) (Chapter 11).

Chapter 2 Signal Acquisition

About this chapter

This chapter describes how to acquire (digitize) signals with Canary. In order to make good recordings with Canary, you should be familiar with the basic concepts of digital sound representation discussed in Appendix A. Failure to understand concepts such as sample rate and sample size can lead to digitized recordings that are distorted in various ways.

Choosing and setting up a recording device

Canary can acquire sounds from the Macintosh's built-in sound input port or from a third-party sound input device. **Incorrect connection of an audio device to the Macintosh may yield distorted signals and may result in damage to the audio device or the computer.** Appendix F explains how to connect an audio device (e.g., tape recorder, CD player, or stereo system) for sound input to your Macintosh.

In order to use a particular sound input device, there must be an appropriate *sound input device driver* installed in the System Folder. A sound input device driver acts as the interface between the sound input hardware and the Sound Manager routines that are part of the Macintosh operating system. The driver for a device is supplied by the device's manufacturer. Device drivers should be placed in the Extensions folder inside the System Folder, and are loaded into memory automatically when the Macintosh starts up.

To acquire a new sound, select Record... from the Process menu or press Command-R. The Sound Acquisition / Recording dialog box (Figure 2.1) should appear. On some Macintosh models, an error occurs if there is no recording device physically connected to the Macintosh when you select Record... .

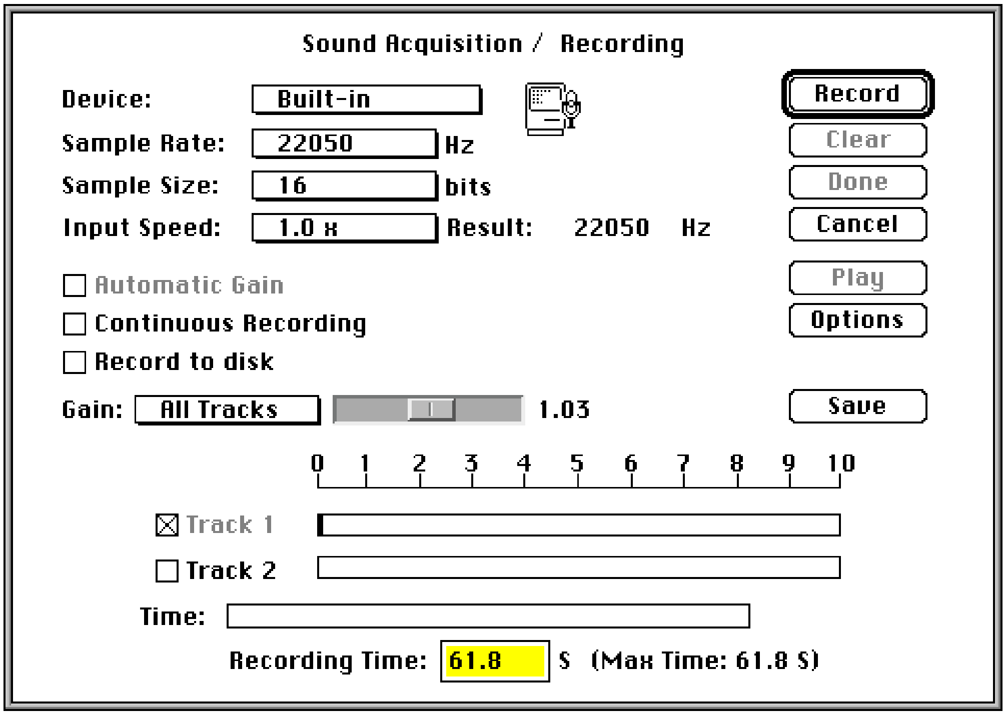

Figure 2.1. The Sound Acquisition / Recording dialog box. Some features of the dialog box are different on different models of Macintosh (see text).

Recording device

The box labeled Device displays the name of the currently selected recording device. Press the mouse button on this box to display a pop-up menu listing all the recording devices available on your machine. To choose a different device, select its name from the pop-up menu.

If a device that is connected to your machine does not appear in the menu, the device's driver was not loaded when the Macintosh started up, in most cases because the driver is not installed in the Extensions folder. If a device's name appears but is dimmed (gray), the driver is present but the device is not physically available (e.g., a MacRecorder is not plugged in).

Options Clicking the Options button brings up a dialog box that allows you to set various options for the selected recording device, if any options are available. Whether or not there is an Options dialog box depends on the recording device; there is no Options dialog box for the built-in recording port on some Macintosh models.

Sample rate The Sample Rate pop-up menu displays all of the sample rates available for the selected recording device (there may be only one sample rate available). Be sure to choose a sample rate that is at least twice as high as the highest frequency in the recording you want to acquire, in order to avoid signal distortion due to aliasing. See Appendix A for a discussion of the tradeoffs involved in choice of sample rate.

Sample size The Sample Size pop-up menu displays all of the sample sizes available for the selected recording device (there may be only one sample size available). A larger sample size results in a better quality digitized signal (higher signal-to-noise ratio), but may also increase the storage requirements for signals saved to disk files. See Appendix A for a discussion of the tradeoffs involved in choice of sample size.

Input speed It is sometimes convenient or necessary to acquire sounds from a tape that is being played back at some speed other than the original recording speed. This is especially true when working with signals that contain energy at frequencies greater than half the sampling rate (the Nyquist rate; see Appendix A). For example, if the highest sampling rate available is 22 kHz, the highest frequency signals that can be acquired accurately are those up to 11 kHz. If you need to acquire a signal at 15 kHz, you can shift the signal down in frequency before acquiring by playing the tape at half speed during acquisition. This would bring the signal down to 7.5 kHz, which is well within the 11 kHz effective bandwidth of a 22 kHz sample rate. If you are working with signals that are too low in frequency to be audible, it may be convenient to play the tape back at a high speed during acquisition, so that you can hear the signals as they are acquired.

The Input speed pop-up menu lets you specify the speed at which a tape will be played during acquisition, relative to the speed at which it was recorded. For example, if you play back a tape at half speed to shift high-frequency signals to below the Nyquist frequency, as described above, you should select 0.5 x from the pop-up menu. If you play a tape at four times normal speed, select 4 x, and so on. If the tape is being played at normal speed during acquisition, the speed selector should be left at 1 x. Canary needs to know the input speed in order to calculate frequencies correctly. The number labeled Result: is the effective sampling rate of the original signal at normal speed.

The fastest effective sampling rate at which the Macintosh can play sounds is 65.5 kHz.

Stereo vs. mono recording

On Macintosh models that support stereo sound input, you can choose to record either in mono or stereo by clicking the checkbox labeled Track 2. If Tracks 1 and 2 are both on, the recording will be stereo; if only Track 1 is on, it will be mono. Chapter 5 discusses working with stereo signals.

Setting the recording level

The types of gain (recording level) controls that are available in the dialog box vary among different Macintosh models. Some of the controls shown in Figure 2.1 may not be available on your machine.

Automatic gain control

Some models of Macintosh have an optional automatic gain control (AGC). If AGC is available on your Macintosh, you can turn it on or off by checking or unchecking the Automatic Gain checkbox. If AGC is not available, Automatic Gain will be dimmed (gray), and you cannot turn the checkbox on. Automatic gain control prevents extreme fluctuations in signal level by attenuating high-amplitude portions of a signal. Since automatic gain control can result in a recording that is not an accurate representation of the original signal, it is not recommended for scientific work.

Adjustable gain

Some models of Macintosh allow you to adjust the recording gain when acquiring a signal. If your Macintosh has adjustable gain control, the dialog box will display the controls just below the Record to disk checkbox; if no gain controls are shown in the dialog box, the Macintosh does not have adjustable recording gain. On some Macintoshes with adjustable gain, you can use the pop-up menu to set the gain for tracks 1 and 2 independently; on other models, the gain for both tracks is always the same.

The Macintosh's built-in gain control allows you to adjust the gain in a number of steps between a minimum setting of 0.5 and a maximum of 1.5. The number of settings available and their values depends on which model of Macintosh you are using.[1]

[1]The gain values displayed (which are provided by the Macintosh operating system and are not controlled by Canary) are not necessarily linear. For example, the amplitude of a digitized signal acquired with a gain setting of 1.0 is not necessarily twice the amplitude obtained by acquiring the same source signal with a gain setting of 0.5.

Setting the recording level

The two horizontal bars labeled Track 1 and Track 2 are recording level meters. The Track 2 bar is active only if Track 2 is checked. The meter indicates the strength of the incoming signal, like a recording level meter or LED display on a tape recorder. The numerical scale is in arbitrary units that are proportional to the logarithm of the signal amplitude. The scale is useful for making repeated recordings at the same level. The meter, which is active as soon as the dialog box appears on the screen, incorporates a "peak bar" that shows the highest level the meter has registered in the preceding 1.5 seconds (this facilitates adjusting the level for short transient peaks).

You'll get the best recording by adjusting the level so that the strongest part of the signal almost fills the meter bar. If a signal's amplitude is high enough to completely fill the meter bar, the recorded signal will be distorted. If the gain controls in the dialog box don't allow you to obtain the desired recording level, or if your Macintosh is not equipped with adjustable gain, you will need to adjust the output level of the device (e.g., tape recorder, MacRecorder) that you're acquiring from. (If you're acquiring from the Macintosh's microphone, adjust the loudness of the sound source or the distance between the microphone and the source.)

If you plan to make calibrated amplitude measurements from the signals you record, be sure to acquire the calibration signal and all data signals using the same gain settings. Chapter 4 contains further information about calibrated amplitude measurements.

Recording to memory

Canary allows you to record either to active memory (RAM) or to a disk file. This section describes recording to memory; the next section discusses recording to disk.

The recording buffer; recording time

Canary sets aside a certain proportion of the free memory available to it as a *recording buffer*, into which recorded sound is stored. You can specify the size of the recording buffer by entering the amount of recording time you wish to use (up to the maximum time shown) in the field labeled Recording Time at the bottom of the dialog. The maximum time available is influenced by the size of Canary's memory allocation, the number and logical size of windows that are open[1], the sample rate selected, and the number of tracks being recorded.[2]

To make more memory available to record a longer signal, you can close any open windows. If you still need more memory, quit Canary and increase its memory allocation (see Appendix D for information on how to set the memory allocation). Keep in mind that higher sampling rates result in greater memory requirements (see Appendix A for a discussion of choosing an appropriate sampling rate).

If you fill the maximum recording time, you probably won't have enough memory available to make a spectrogram. Using half of the available recording time allows sufficient memory for a spectrogram using Canary's default spectrogram parameters. Changing the spectrogram parameters can either increase or decrease the amount of memory needed for the spectrogram (see Chapter 3).

---

[1]The logical size of a window is the amount of memory required to store the window's waveform, spectrogram, and spectrum panes, and is independent of the physical size of the window on the monitor.

[2]The maximum time available for recording is not affected by the sample size because Canary's internal representation of a signal always uses 32 bits per sample, irrespective of the size of the samples output by the digitizer.

Continuous recording If the Continuous recording checkbox is on, recording "wraps around" the buffer: when the specified recording time is reached, recording continues, overwriting the earliest data already saved into the buffer. This feature is useful for acquiring signals without needing to carefully position a recording at the sound of interest before you begin acquiring. You can thus "browse" a tape with continuous recording switched on, and then stop acquiring after the sound of interest has been acquired (but before it gets overwritten!). When you finish acquiring in Continuous recording mode, Canary rearranges the sound in memory so that it is in the right order. Continuous recording is available when you use the Macintosh's built-in sound input port; it may or may not be available with third-party sound input devices.

Making a recording to memory Click on the Record button to start recording. The Time bar near the bottom of the dialog box fills from left to right to indicate the proportion of the specified recording time that has been filled; the digital display at the right end of the bar shows the number of seconds of sound that have been recorded. If Continuous recording is selected, the Time bar fills, and then a vertical line moves from left to right through the Time bar as new data overwrite old data. If Continuous recording is off, recording stops automatically when the recording buffer is full.

While sound is being recorded to memory, the Record button is replaced by a Pause button. When recording is paused, you can Resume recording, Play the sound that is already recorded, or Clear the recording buffer by clicking the appropriate buttons in the dialog box.

Clicking the Done button at any time closes the Signal Acquisition dialog, and opens an untitled Canary signal window containing the signal that was recorded. Clicking the Cancel button closes the Signal Acquisition dialog and discards any sound that may have been recorded.

Signals recorded to memory are recorded with Canary's current default amplitude calibration. See Chapter 4 for a discussion of amplitude calibrations.

Recording to disk

Destination file If you click on the Record to disk checkbox, a standard file dialog appears, asking for the name and location of a file to be created for the sound you will record. When recording to disk, the only file format available is AIFF. After you specify the file, the name and location of the file are displayed in the dialog.

Recording time After you select Record to disk and specify the name of the file to be saved, you can specify the Recording Time at the bottom of the dialog. You can record a file of any length, limited only by the space available on your disk; however, large files may be too large to open, or too large to allow you to make a spectrogram with the amount of RAM that your system has. You may have to experiment to determine the length of the largest file that you can use.

Making a recording to disk Click on the Record button to start recording. The Time bar near the bottom of the dialog box fills from left to right to indicate the proportion of the specified maximum time that has been filled. Recording stops automatically when the specified recording length is reached.

While sound is being recorded to disk, the Record button is replaced by a Pause button. When recording is paused, you can Resume recording, or Close the file (which saves the sound recorded to it). You cannot Play sound recorded to disk from within the Signal Acquisition dialog.

Clicking the Done button at any time while recording to disk saves and closes the file, and closes the Signal Acquisition dialog. Clicking the Cancel button before a file has been closed deletes the file from the disk and closes the Signal Acquisition dialog.

Once a signal has been recorded to disk, you can open it from Canary's File menu or by double-clicking on the file's icon in the Finder.

Signals recorded to disk are not assigned any calibration information when they are saved. The first time such a file is opened in Canary, it is assigned the default calibration that is in effect when the file is opened.

Saving preferences

Clicking the Save button immediately saves the settings for Device, Sample Rate, Sample Size, Input Speed, Automatic Gain, Continuous Recording, Gain, and Recording Time in the current preferences file. Preferences are discussed in more detail in Chapter 8.

Chapter 3 Spectrum Analysis

About this chapter

This chapter is a reference for working with spectra and spectrograms in Canary. The chapter assumes that you have read Chapter 1, and that you are acquainted with basic concepts involved in short-time Fourier transform (STFT) analysis of time-varying signals. Terms such as frame (equivalent to “aperture” or “window”), filter bandwidth, frame overlap, and window function are explained here only briefly. If you are not already familiar with these concepts, we recommend that you read Appendix B, which provides the conceptual background needed to make full use of Canary’s spectrum analysis capabilities. For an explanation of how to make multiple spectrograms or spectra at a time (batch processing), see Chapter 11.

Because of the close relationship between spectra and spectrograms, we cover both types of analysis in this chapter, first discussing those parameters common to both, then discussing particular issues specifically relevant to spectra and spectrograms separately. Unless otherwise indicated, the term *spectrum analysis* in this manual refers to the calculation of either a spectrum or a spectrogram.

Canary’s spectra and spectrograms can display signal amplitudes in either acoustical or electrical units, depending on which calibration paradigm is selected for a signal. (When you first use Canary, the acoustical paradigm is the default choice.) This chapter discusses spectrum analysis in terms of acoustic signals. See Chapter 4 for a discussion of the two calibration paradigms and the correspondence between acoustical and electrical amplitude units.

Spectrum analysis parameters

There are two convenient ways to think about the production of spectrograms and spectra and the relationship between the two analyses. On the one hand, a spectrogram can be thought of as a series of spectra of successive short time segments, or *frames*. Alternatively, a spectrogram can be considered as a plot of the time-varying output amplitude of a bank of bandpass filters. A single spectrum corresponds to a cross-section or slice of the spectrogram. Therefore, most of the parameters that must be specified are the same for spectrograms and spectra (Figure 3.1). The following sections briefly explain each parameter. For a more detailed explanation of the relationship between spectrograms and spectra, see Appendix B.

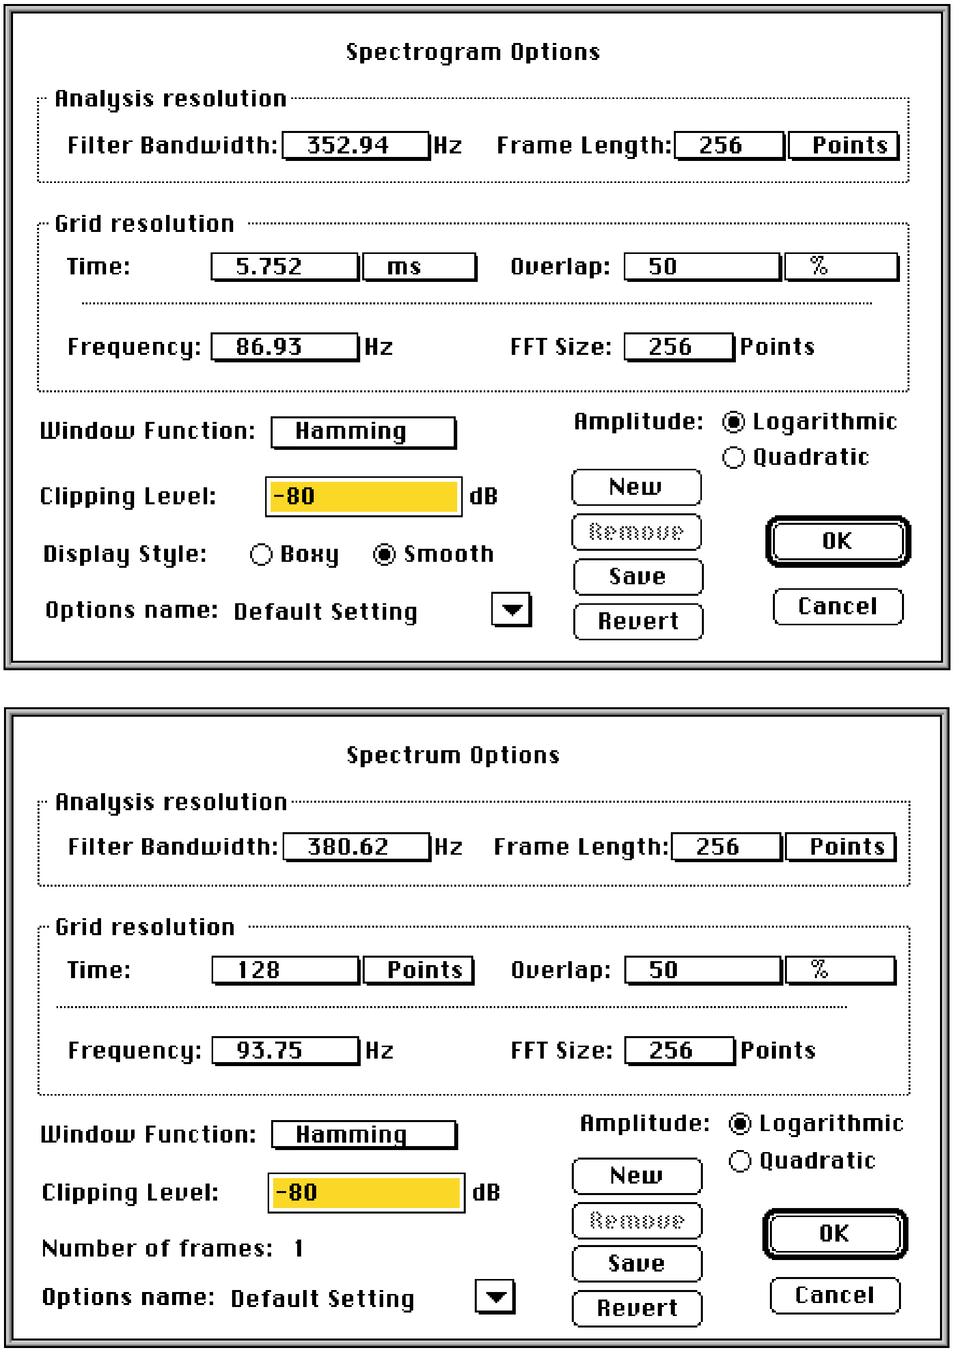

Figure 3.1. The Spectrogram Options and Spectrum Options dialog boxes.

Analysis resolution: Filter Bandwidth and Frame Length

The Filter Bandwidth and Frame Length parameters provide two alternative (and equivalent) ways to control the tradeoff between analysis resolution in the time and frequency dimensions. Filter bandwidth is the bandwidth (in Hz) of the individual analysis filters in the filterbank simulated by the short-time Fourier transform (STFT) with the selected frame length (see Appendix B). The filter bandwidth is also influenced by the choice of window function (see discussion of the Window Function parameter below). Frame length is the duration (measured in mS or points[1]) of each successive interval over which the spectral composition of a signal is estimated. The Filter Bandwidth and Frame Length fields in the dialog box are coupled to each other: changing one of them automatically changes the other. Improved analysis resolution in either dimension (time or frequency) can only be obtained at the expense of resolution in the other dimension.

For computational efficiency, Canary requires that the number of individual points (digital samples) in each frame be a power of 2 (i.e., 16, 32, 64, ...). Pressing the mouse button on the frame length field pops up a menu from which you can choose the desired frame length. You can choose whether to have the frame length displayed in points or milliseconds using the pop-up menu that initially reads Points. The Frame Length parameter cannot be set to a value larger than the FFT Size.

In a spectrum, where time is not represented, you may want to use a narrow filter bandwidth (large frame length) to achieve greater frequency resolution. Keep in mind, however, that the spectrum that is calculated is the *average* spectrum over the duration of a frame. Long frame lengths may thus “blur” spectral peaks in a sound when frequency is changing rapidly.

In a spectrogram, where you are interested in frequency variations with time, the “best” choice of frame length depends in part on the nature of the signal, and on what features you are most interested in observing or measuring. If you are most concerned with precise frequency measurements, you will probably want to choose a small filter bandwidth (hence long frame length and poorer time resolution). If you want better time resolution, select a shorter frame length; the filter bandwidth will then be larger (poorer frequency resolution).

Figures 3.2 and 3.3 show pairs of spectrograms and spectra that differ only in analysis resolution. See Appendix B for further discussion and more examples of the effect of varying filter bandwidth and frame length.

[1]A point is a single sample taken by the A/D converter.

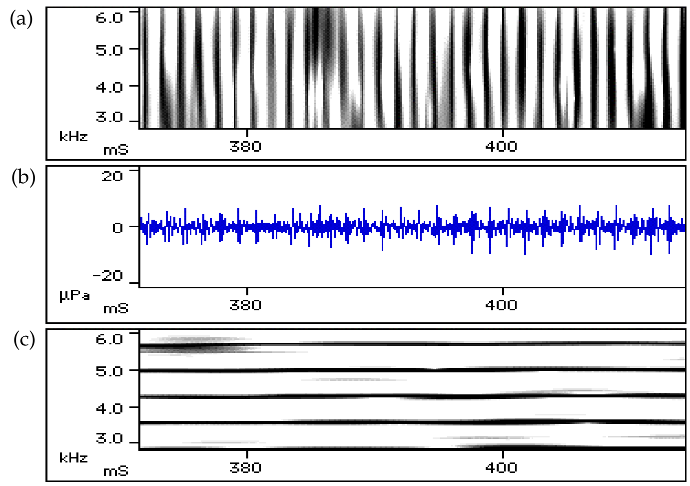

Figure 3.2. Effect of varying analysis resolution on spectrograms. The signal is part of a rapid series of clicks produced by a spotted dolphin, digitized at 48 kHz. The period between clicks is about 1.4 mS, corresponding to a frequency of about 720 Hz. The two spectrograms differ only in analysis resolution. In both spectrograms, grid resolution = .33 mS x 46.9 Hz, window = Hamming. **(a)** Filter bandwidth = 6090 Hz (frame length = 32 points = .67 mS), overlap = 50%. In this representation, each click appears as a broad-band vertical stripe on the spectrogram because the frame length is short enough to resolve individual clicks. **(b)** Waveform. When played at normal speed, the signal sounds to a human like a buzz. **(c)** Filter bandwidth = 190 Hz (frame length = 1024 points = 21.3 mS), overlap = 98.4%. In this representation, individual clicks cannot be resolved because each frame encompasses about 15 clicks; instead the click repetition frequency shows up as a series of horizontal bands spaced 720 Hz apart (the click repetition frequency).

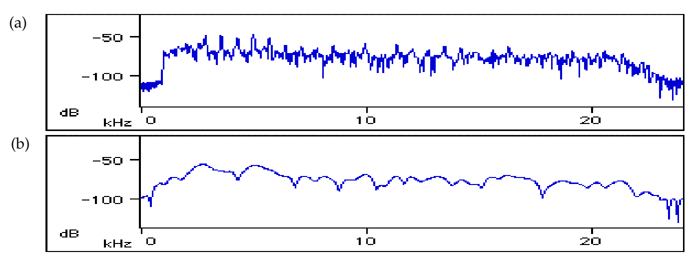

Figure 3.3. Effect of varying analysis resolution on spectra. The two single-frame spectra are taken at about 370 mS into the signal shown in Figure 3.2; they differ only in analysis resolution. In both spectra, frequency grid resolution = 23.4 Hz, window = Hamming. **(a)** Filter bandwidth = 95.6 Hz (frame length = 2048 points = 42.7 mS). **(b)** Filter bandwidth = 1.52 kHz (frame length = 128 points = 2.7 mS).

Grid resolution

You can specify the time and frequency resolution of the grid of points at which spectral densities are calculated using the Time and Frequency parameters in the box labeled Grid resolution.[1] In a boxy spectrogram, these two parameters control the width and height of the grid boxes, respectively. In a spectrum, the Frequency parameter controls the number and spacing of the calculated points along the frequency axis (i.e., the frequency filters). For multi-frame spectra, the Time parameter controls the spacing between points in time at which spectral densities are calculated (and then averaged to obtain the final displayed spectrum). For single-frame spectra, the Time parameter is irrelevant. (See the discussion of “Number of frames to average” in the section on “Spectra” later in this chapter.)

Frequency grid resolution (FFT size)

The frequency resolution of the grid depends on the sampling rate (which is fixed for a given digitized signal) and the FFT size. The relationship is

frequency grid resolution = (sampling frequency) / FFT size

where frequency resolution and sampling frequency are measured in Hz and FFT size is measured in points. The Frequency and FFT Size parameters in the dialog box are coupled to each other, so you can specify the desired frequency resolution either by choosing a value directly from the Frequency pop-up menu, or by choosing the FFT Size. FFT size can be displayed either in points (the default) or mS, depending on your choice from the units pop-up menu. The number of points in the FFT is constrained to be a power of two. The FFT Size parameter cannot be smaller than the Frame Length.

Figures 3.4 and 3.5 show pairs of spectrograms and spectra that differ only in frequency grid resolution.

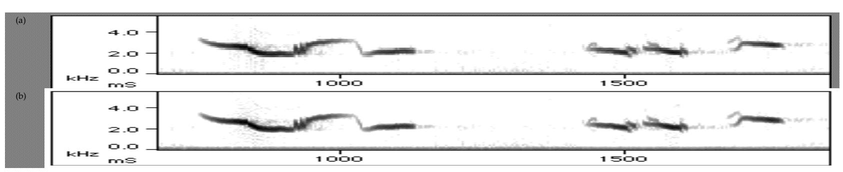

Figure 3.4. Effect of frequency grid resolution (FFT size) on spectrograms. The signal is part of a song of an American robin, digitized at 22.3 kHz. The two spectrograms differ only in frequency grid resolution. In both spectrograms, filter bandwidth = 706 Hz (frame length = 128 points = 5.8 mS), time grid resolution = 2.9 mS (frame overlap = 50%), window = Hamming. **(a)** Frequency grid resolution = 172 Hz (FFT size = 128 points). **(b)** Frequency grid resolution = 43.5 Hz (FFT size = 512 points).

[1]The grid resolution, which determines the “graininess” of a spectral display, should not be confused with the analysis resolution, which determines the relative precision of time and frequency determination at each point. See Appendix B for further explanation.

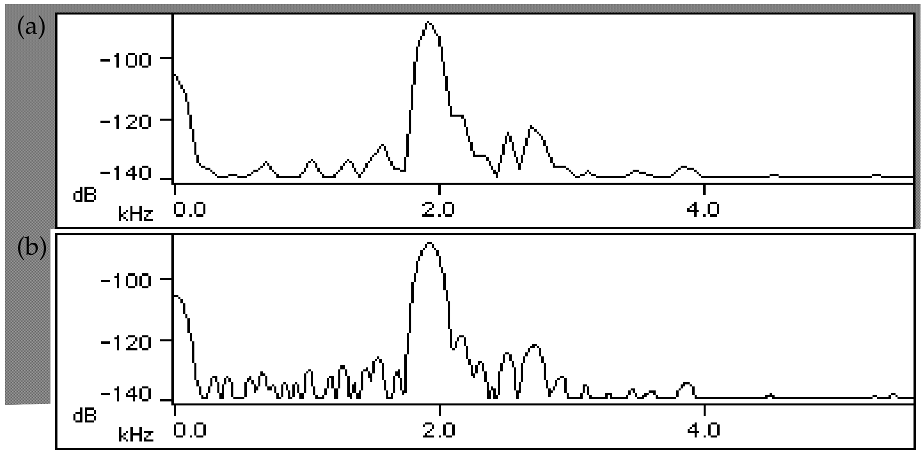

Figure 3.5. Effect of varying frequency grid resolution (FFT size) on spectra. The two single-frame spectra were taken at approximately 880 mS along the time axis shown in Figure 3.4 and differ only in frequency grid resolution (FFT size). In both spectra, filter bandwidth = 353 Hz (frame length = 256 points = 11.5 mS), time grid resolution = 5.7 mS (frame overlap = 50%), window = Hamming. **(a)** Frequency grid resolution = 86.9 Hz (FFT size = 256 points). **(b)** Frequency grid resolution = 10.9 Hz (FFT size = 2048 points).

Time grid resolution (frame overlap)

The time interval (in mS) between successive gridpoints is the time between the beginnings of successive frames. Frames can be overlapping, contiguous (0% overlap) or separated by time intervals that are omitted from the analysis (negative overlap).

Frame overlap is usually expressed as percent of frame length. For example, an overlap of 50% means that each frame begins halfway through the preceding fame. An overlap of -100% means that one frame of data is skipped between successive frames that are analyzed; -300% skips three frames, and so on. The relationship between time grid resolution and frame overlap is given by

$$\textit{time grid resolution} = \textit{frame length} * (100\% - \textit{overlap}\%)$$

where frame length is measured in mS. The Time and Overlap choices in the dialog box are coupled so that you can specify time resolution either directly, by choosing a value from the Time pop-up menu, or indirectly, by specifying an Overlap value. Using the units pop-up menu, you can specify the measurement units for time resolution as either mS (the default) or points. Measurement units for overlap are percent (of frame length), mS or points.

For computational efficiency, Canary requires that non-negative overlaps be equal to $100(1 - (1/2^n))$, where $n \geq 0$ (e.g., 0%, 50%, 75%, 87.5%, 93.8%,...). Negative overlaps must be equal to $-100(2^n-1)$ (e.g., -100%, -300%, -700%,...), where $n \geq 0$.

Remember that a spectrogram or spectrum made with negative frame overlaps ignores some of the available data, and can give an extremely misleading picture of a signal. Negative frame overlaps should generally be avoided unless you have some specific reason for wanting to omit some parts of a signal from analysis.

Figures 3.6 and 3.7 show pairs of spectra and spectrograms that differ only in time grid resolution.

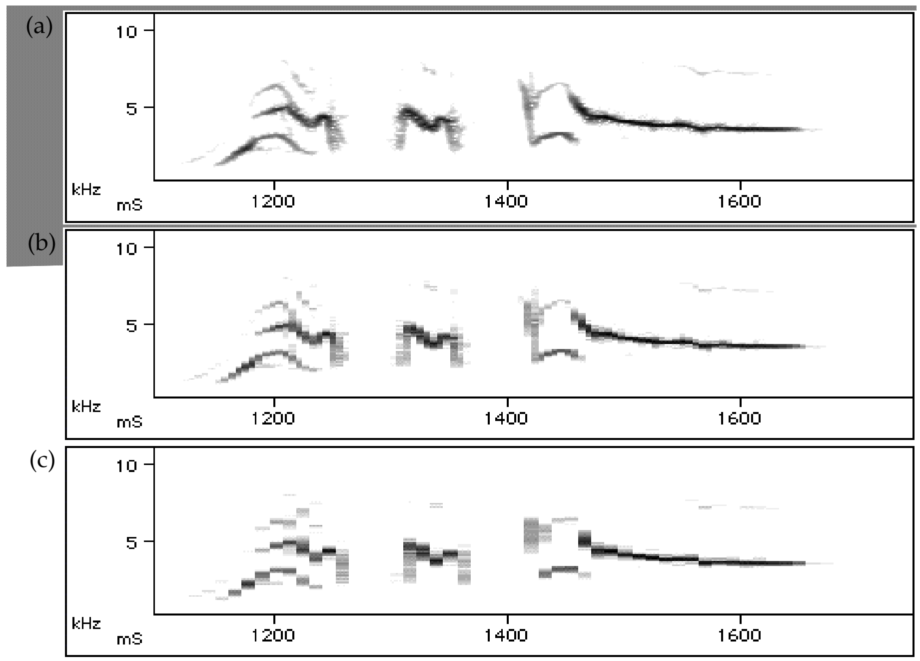

Figure 3.6. Effect of varying time grid resolution in spectrograms. The signal is part of a song of a western meadowlark, digitized at 22.3 kHz. The two spectrograms differ only in time grid resolution (frame overlap). In all three spectrograms, filter bandwidth = 353 Hz (frame length = 256 points = 11.5 mS), frequency resolution = 22.7 Hz (FFT size = 1024 points), window = Hamming. **(a)** Time grid resolution = .719 mS (frame overlap = 93.8%). **(b)** Time grid resolution = 5.73 mS (frame overlap = 50%). **(c)** Time grid resolution = 11.5 mS (frame overlap = 0%).

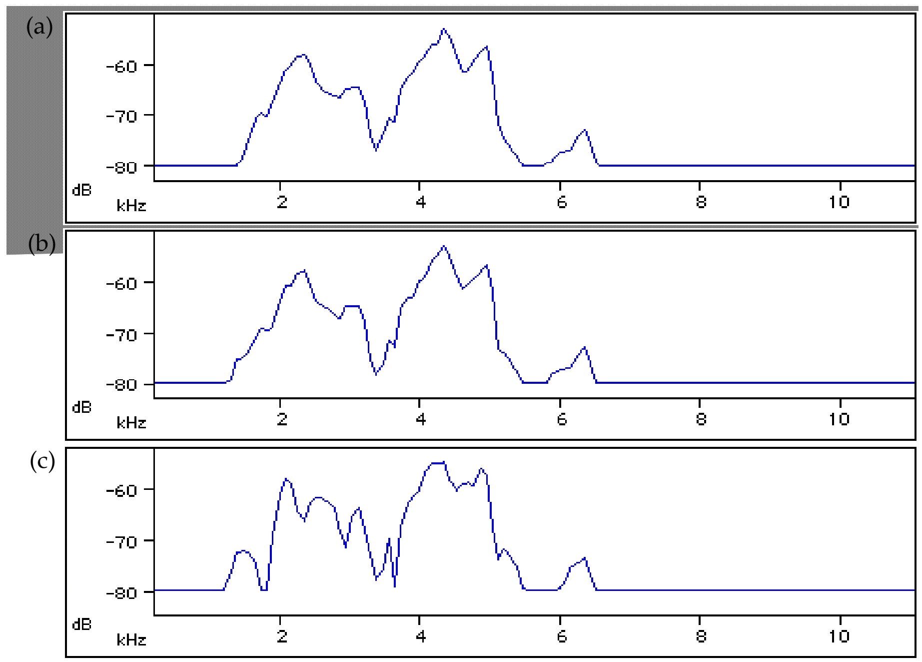

Figure 3.7. Effect of varying frame overlap on multi-frame spectra. All three spectra are of the first syllable shown in Figure 3.6, from 1159 mS to 1268 mS into the signal. All were made with filter bandwidth = 705.9 Hz (frame length = 128 points = 5.8 mS), frequency resolution = 173.9 Hz (FFT size = 128 points), and window = Hamming. **(a)** Overlap = 93.8%; spectrum includes 136 frames. **(b)** Overlap = 50%; spectrum includes 17 frames. **(c)** Overlap = 0%; spectrum includes 9 frames.

Finer grid resolution, in either the time or frequency dimension, comes at the cost of calculation speed. Finer frequency grid resolution increases the memory requirement for both spectra and spectrograms. Finer time resolution increases memory requirements for spectrograms only.

Window Function

Each frame of data is multiplied by a window function before its spectrum is calculated. Window functions are used to reduce the magnitude of spurious frequencies that appear as sidelobes flanking each analysis frequency in a spectrum. These sidelobes appear as a result of analyzing a finite (truncated) section of a signal. A window function can reduce these sidelobes by "tapering" the portion of the waveform that appears in each frame.

Canary provides five different window functions (rectangular, Hamming, Hanning, Bartlett, and Blackman). Each window function is characterized by the magnitude of the sidelobes relative to the center lobe. The difference in decibels between the center lobe magnitude and the magnitude of the largest sidelobe is called the *sidelobe rejection* (Figure B.15). In a spectrogram, differences among windows in sidelobe rejection result in different amounts of gray "fringing" in the frequency dimension around black or very dark areas.

For a given frame length, different window functions will result in different filter bandwidths (Figure B.16). In terms of a spectrogram, this means that the vertical thickness of a line drawn to represent a pure tone will depend on which window function is used.

Figures 3.8 and 3.9 illustrate the effect of different window functions on spectrograms and spectra of the same signal.

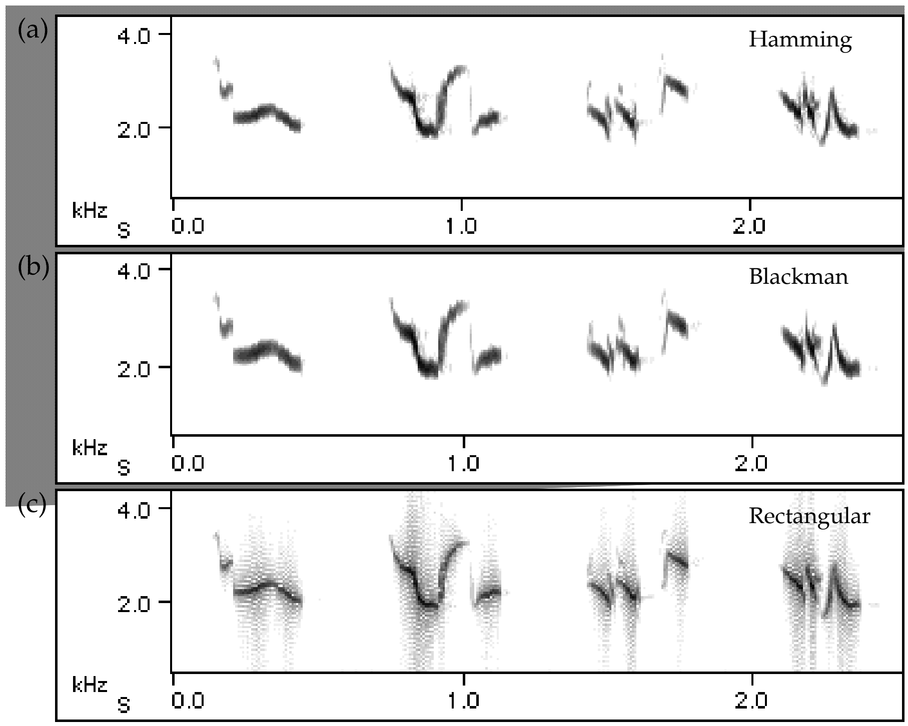

Figure 3.8. Effect of choice of window function on spectrograms. The signal is part of a song of an American robin, digitized at 22.3 kHz. All three spectrograms have the same clipping level, frame length = 256 points (= 11.5 mS), time resolution = 5.8 mS (frame overlap = 50%), and frequency resolution = 10.9 Hz (FFT size = 2048 points). Filter bandwidths: **(a)** 353 Hz, **(b)** 538 Hz, **(c)** 174 Hz.

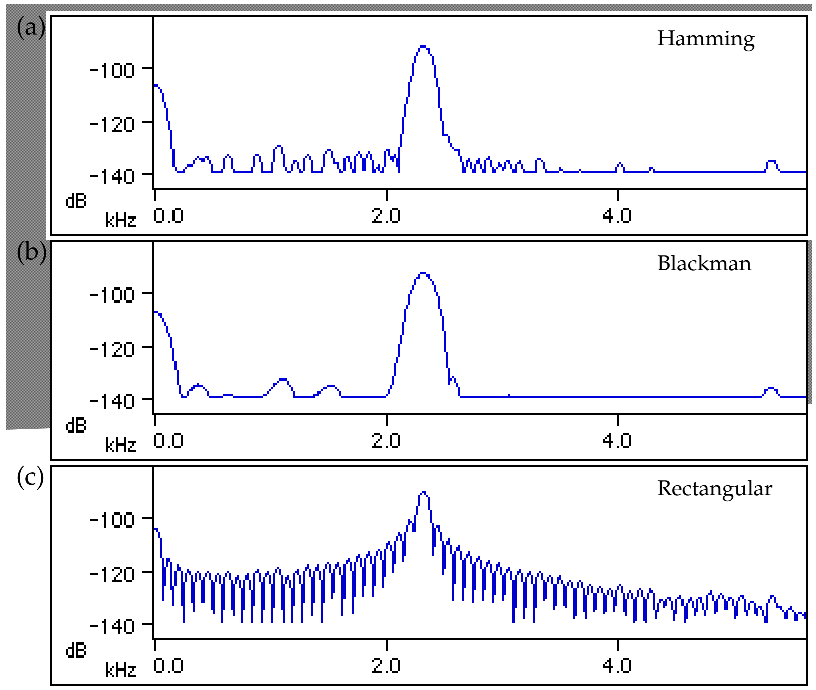

Figure 3.9. Effect of window function on spectra. The single-frame spectra were made at a point approximately 0.30 sec along the time axis of the signal shown in Figure 3.8. All three spectra have the same clipping level, filter bandwidth = 353 Hz (frame length = 256 points = 11.5 mS), and frequency resolution = 10.9 Hz (FFT size = 2048). Filter bandwidths: **(a)** 353 Hz, **(b)** 538 Hz, **(c)** 174 Hz.

The appearance of sidelobes in spectra of finite signals, and the use of window functions to reduce their magnitude, is discussed further in Appendix B.

Clipping Level

The Clipping Level parameter allows you to specify a "noise floor" below which any amplitude value is ignored. In a spectrum, any frequency component less than the clipping level is displayed as being equal to the clipping level. In a spectrogram, any value below the clipping level is displayed as white. The Clipping Level parameter is in units of intensity dB relative to the dB reference value specified in the Set Calibration dialog. When you first use Canary, the factory default clipping level is -80 dB.

There are several reasons for specifying a noise floor. First, because of the finite precision of the digitizing process, a digitized sound always contains some error and has a limited dynamic range. For signals digitized with 8-bit samples, the dynamic range is limited to 48 dB; for 16-bit signals, the dynamic range is 96 dB.[1] Thus, any amplitude value in an 8-bit spectrum or spectrogram that is more than 48 dB below the highest peak in the signal must

[1]The dynamic range of a digitized sound is 6 dB/bit.

be noise introduced by the digitizing process, and should be disregarded.[1] The noise floor can also be useful for removing noise that was present before the digitizing process (for example, from a recording with low-level wind or other broad-band noise).

A more pragmatic reason for noise clipping is that very small frequency components show up on a log scale as large negative dB levels. (Recall that the logarithm of zero is minus infinity.) The noise floor allows Canary to ignore very small frequency components. If you choose a quadratic rather than logarithmic display, the clipping level is automatically set to zero (which is displayed as -∞ dB).

If the clipping level is set too low, excessive noise will be displayed in the spectrum or spectrogram along with the signal. If it is set too high, portions of the signal will not be visible. You may need to experiment with different clipping levels in order to find a value that produces a satisfactory display.

Since Canary calculates intensity values based on a signal's amplitude calibration information, clipping level depends on calibration. The clipping level required to achieve a given level of noise rejection for a signal will change if you change the signal's calibration.

The noise floor can also be used to eliminate spectral sidelobes (which show up as gray fringes around strong signal components in spectrograms).

You can raise the noise floor in a spectrogram display after it has been drawn by raising the lower amplitude (horizontal) cursor in the spectrum pane (see the discussion of "Contrast and brightness" in the "Spectrograms" section later in this chapter). Note that this method does not change the noise floor in the spectrogram data stored in memory, however. If you are only using spectrograms for visual examination and display, the distinction between the noise floor of the display and the noise floor of the spectrogram data is unimportant. However, if you plan to do any quantitative analysis (e.g., correlations) using the spectrogram data, remember that the only way to change the noise floor is to recalculate the spectrogram, specifying a different Clipping Level in the spectrogram dialog box.

If you want to lower the noise floor in a spectrogram (either the display or the data), you must recalculate it with a lower Clipping Level.

Figures 3.10 and 3.11 show spectrograms and spectra that differ only in clipping level.

---

[1]If the highest spectral peak in a signal is smaller than the digitizer's maximum output level, the dynamic range between the peak and noise introduced by digitizing will be less than 6 dB/bit. However, it can never be more than 6 dB/bit.

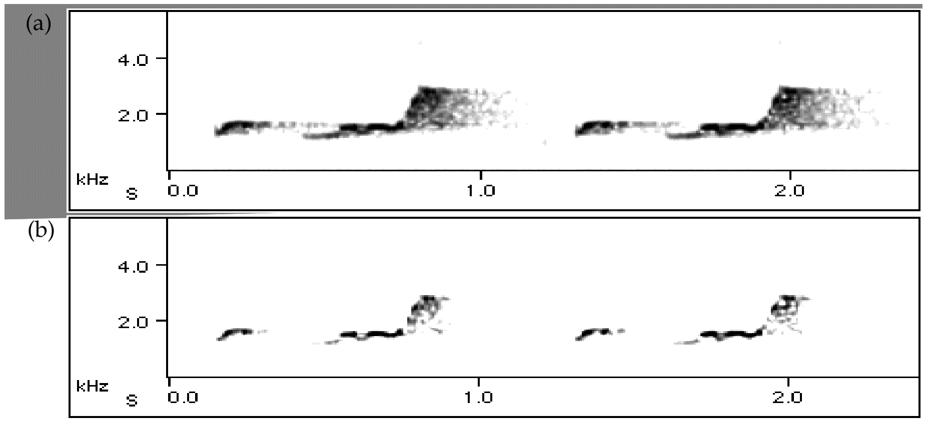

Figure 3.10. Effect of clipping level on spectrograms. The signal is part of a song of a whip-poor-will. **(a)** Clipping level = -74 dB (using Canary's factory default amplitude calibration). The temporal smearing of each syllable is a result of reverberation from forest vegetation. **(b)** Clipping level = -55 dB. Notice the reduction in the darkness of the reverberation smears.

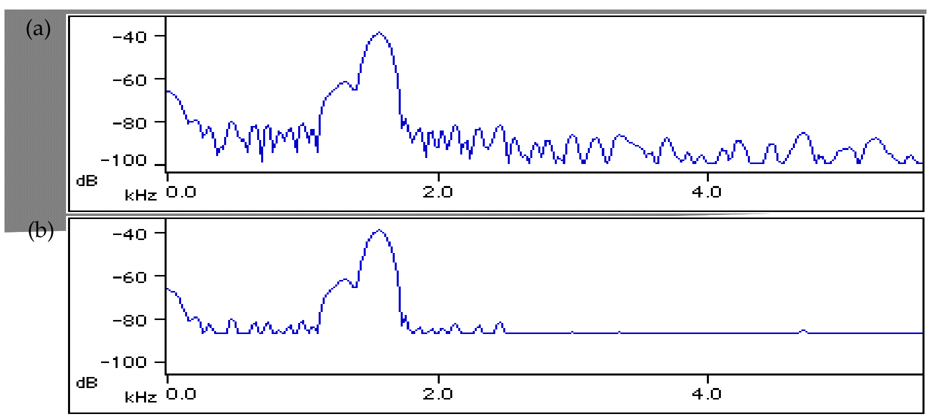

Figure 3.11. Effect of clipping level on spectra. Both spectra are of a single 256-point frame approximately 0.7 sec into the whip-poor-will song shown in Figure 3.10. The signal was digitized with 8-bit resolution, and thus has a dynamic range of 48 dB (see text). **(a)** Clipping level = -100 dB, using Canary's factory default calibration. The highest peak has an amplitude of -39 dB, so all amplitude values below -87 dB (=-39-48 dB) represent noise in the digitizing process. **(b)** Clipping level = -87 dB.

Logarithmic vs. quadratic amplitude axis

The Amplitude radio buttons allow you to choose either a logarithmic or quadratic[1] amplitude axis. In a logarithmic spectrum, intensity levels are plotted in dB relative to the intensity dB reference value specified in the Set Calibration dialog; a quadratic spectrum plots intensities in $W/m^2/Hz$. For a spectrogram, the choice of amplitude axis determines whether the grayscale values are proportional to the intensity at each point or the log (base 10) of the intensity.

Logarithmic displays are usually more useful for spectra and spectrograms of biological sounds because of their large dynamic range.

When you select Quadratic amplitude axis, the Clipping Level parameter is automatically set to $-\infty$ dB.

Figures 3.12 and 3.13 show spectrograms and spectra plotted with logarithmic and quadratic amplitude axes.

---

[1]If you choose a quadratic amplitude axis, the values displayed are intensity spectrum values (in $W/m^2/Hz$) rather than intensity spectrum levels (decibels/Hz). If the signal is calibrated using the electrical paradigm, quadratic amplitudes are power spectrum values (J/Hz). These spectra are considered “quadratic” because the values shown are proportional to the square of the signal amplitude (sound pressure or voltage).

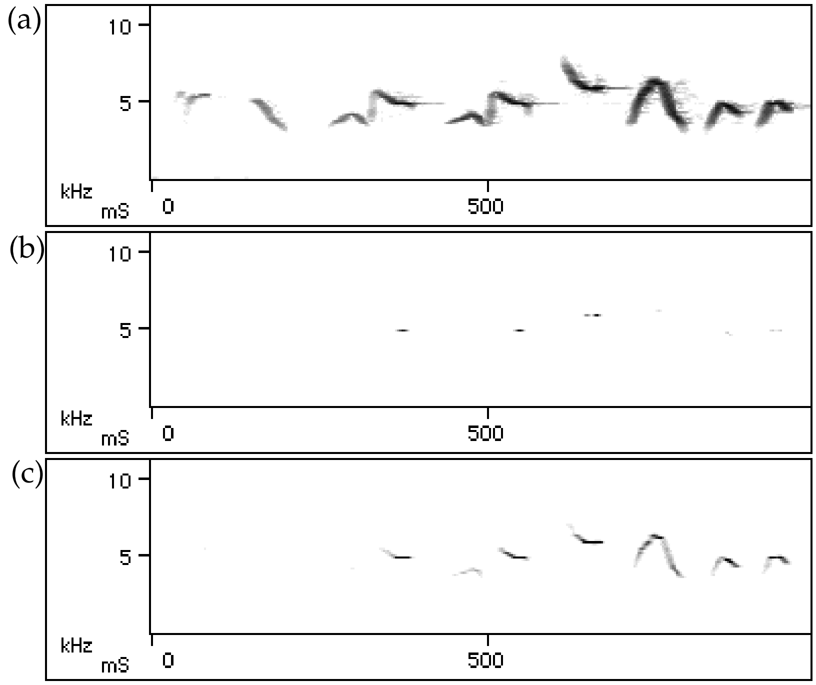

Figure 3.12. Logarithmic and quadratic spectrograms of a song of a Canada warbler, digitized at 22.3 kHz. Filter bandwidth = 176.5 Hz (frame length = 512 points = 23.0 mS), time resolution = 2.9 mS (87.5% overlap), frequency resolution = 43.5 Hz (FFT size = 512 points), window = Hamming. **(a)** Logarithmic spectrogram. **(b)** and **(c)** quadratic spectrograms. The amplitude ceiling of (b) is the same as (a); the ceiling of (c) is lower.

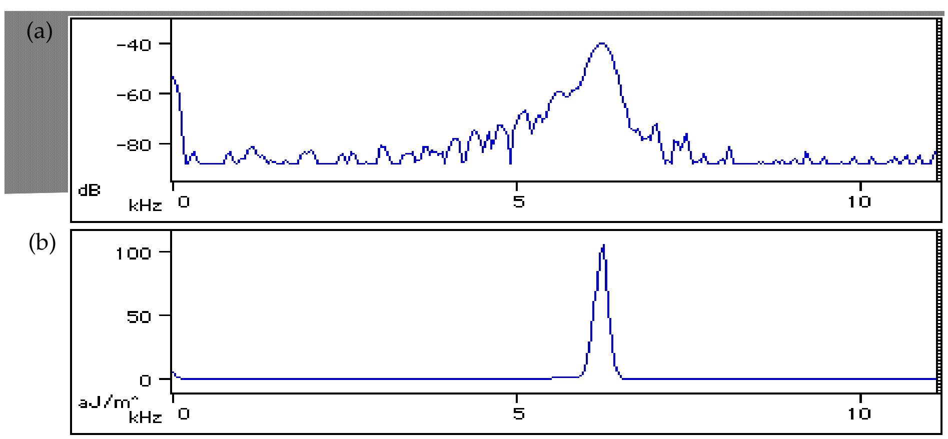

Figure 3.13. Logarithmic and quadratic single-frame spectra taken at the frequency apex of the antepenultimate syllable of the Canada warbler song shown in Figure 3.12 (755 mS into the signal). Filter bandwidth = 352.9 Hz (frame length = 256 points = 11.5 mS), frequency resolution = 43.5 Hz (FFT size = 512), window = Hamming. **(a)** Logarithmic spectrum, clipping level = -88 dB. **(b)** Quadratic spectrum.

Named options sets

Canary allows you to name and save sets of spectrogram or spectrum options. Once you've saved a set of options, you can restore them all with a single mouse click.

The name of the currently selected set of options is shown next to Options name: in the lower left corner of the dialog box. Pressing the vertical arrow button next to the options name pops up a menu showing the names of all of the options sets saved in the current preference file (see Chapter 8 for a complete discussion of preference files). When you first start to use Canary, there is only one named option set, called Default.

Any time that you select a given set of options from the popup menu, all of the settings that were in effect the last time that the set was selected are instantly restored.

Clicking the New button creates a new set of options that is initially named Untitled, and allows you to edit the name of the set. You can then set the various options any way that you want.

Clicking the Save button saves the definitions of all of the options sets into the current preference file. If you do not click Save, and if you do not save preferences before quitting Canary, definitions of named options sets will be forgotten.

Whenever you change the spectrogram or spectrum parameters while a given options set is selected, you are changing the definition of that set. You can restore all named options sets to the last definitions that were saved by clicking the Revert button.

You can remove a named options set by selecting its name from the popup menu and then clicking the Remove button.

Spectra

Significance of the spectrum values

The values shown in a spectrum represent either the spectrum intensity level (in dB) or the spectrum intensity (in watts/m^2 Hz) at each frequency, depending on whether a Logarithmic or Quadratic amplitude axis was selected. Spectrum intensity level and spectrum intensity are defined as the intensity level or intensity in a band one hertz wide centered at the specified frequency. The numerical values for intensity level or intensity at each frequency (which are displayed by the spectrum measurement panel) depend on the current calibration settings. **The intensity or intensity level values on a spectrum are meaningless** unless the signal has been properly calibrated according to the procedures outlined in Chapter 4.[1]

Number of frames to average

If part of the waveform is selected (highlighted) when you request a spectrum, the number of frames selected is shown in the spectrum dialog box. The number of frames that a given selection encompasses depends on both the frame length and the frame overlap. Consequently, the frame count shown in the dialog box may change if you change either of these parameters. If more than one frame is selected, Canary calculates the spectrum for each frame independently, then averages the spectra of all the frames. The averaged spectrum is then displayed. If the selected part of the signal is at least one frame long, but is not an integral number of frames in length, Canary uses just those frames that are completely contained in the selection (i.e., the selection is truncated to an integer number of frames). If the selection encompasses less than one whole frame, the selection is zero-padded to make one full frame, which is then used for the spectrum.

If no part of the waveform is selected, Canary makes a spectrum of a single frame, which begins at the blinking insertion point.

Once a spectrum is displayed, you can double-click on it to highlight the interval in the waveform from which the spectrum was made (the spectrum source interval). This function is particularly useful with single-frame spectra if you want to see exactly how much of the signal was actually used for the spectrum. You can determine how many frames were used to make an existing spectrum using the spectrum measurement panel (see Chapter 6).

[1]If you are interested only in differences in decibel levels between parts of one signal, you do not have to calibrate the signal.

Spectrograms

Significance of the grayscale values

The grayscale values shown in a spectrogram represent either the spectrum intensity level (in dB) or the spectrum intensity (in watts/m^2 /Hz) at each point, depending on whether a Logarithmic or Quadratic amplitude axis was selected. The numerical values for intensity level or intensity associated with each point (which are displayed by the spectrogram measurement panel) depend on the current calibration settings. **The intensity or intensity level values in a spectrogram are meaningless** unless the signal has been properly calibrated according to the procedures outlined in Chapter 4.[1]

Contrast and brightness

The basic operation of the spectrogram contrast and brightness controls is discussed in Chapter 1. This section provides a more detailed explanation of how these controls work. The information in this section is not needed to use the controls effectively; it is provided for those who are interested in a quantitative explanation of how these controls affect the spectrogram image.

To see exactly how the spectrogram and brightness controls work, first make a logarithmic spectrogram, then click anywhere on the waveform and make a logarithmic spectrum. Finally, click on the CURSORS button to turn selection cursors on. The window should look something like Figure 3.14.

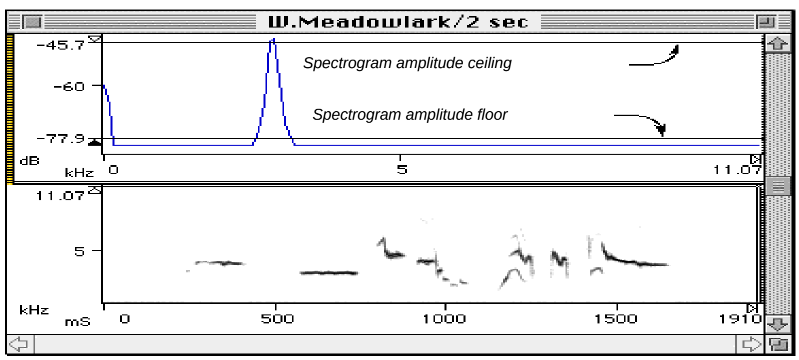

Figure 3.14. The spectrum's amplitude cursors delimit the amplitude "ceiling" and "floor" of the spectrogram. All amplitudes above the ceiling are displayed as black in the spectrogram; all amplitudes below the floor are displayed as white. The amplitude ceiling and floor can be manipulated together using the contrast and brightness controls on the command panel, or independently by moving the amplitude cursors in the spectrum pane. In this illustration the waveform pane has been hidden.

The amplitude (vertical axis) cursors in the spectrum pane delimit the range of amplitudes represented by the scale of gray shades (from white to black) displayed in the spectrogram. The position of the lower cursor in the spectrum determines the amplitude "floor" of the spectrogram, below which all

[1]If you are interested only in differences in decibel levels between parts of one signal, you do not have to calibrate the signal.

amplitudes appear as white. The upper cursor determines the spectrogram's "ceiling" amplitude, above which all values are represented by black. When a spectrogram is first drawn, the amplitude floor is set to the spectrogram's clipping level, and the amplitude ceiling is set to the highest amplitude value on the spectrogram.[1]

Increasing the contrast moves the floor and ceiling closer to each other, reducing the number of shades of gray in the spectrogram. If the contrast is set to its maximum (100%) value, the floor and ceiling are equal, and the spectrogram becomes black and white (no grays). Adjusting the brightness shifts the floor and ceiling up (decreasing brightness) or down, while preserving the distance between them. The initial contrast and brightness settings that are determined when the spectrogram is calculated are arbitrarily assigned values of 50%.

You can adjust the floor and ceiling of the spectrogram independently by moving the amplitude cursor in the spectrum up or down.

The number of shades of gray in a spectrogram may be limited by the "pixel depth" of your monitor. You can specify the number of colors or shades of gray that your monitor can display using the Macintosh "Monitors" control panel (accessible through the apple menu). The choices that are available through the control panel depend ultimately on the display hardware installed in your machine.

[1]The spectrogram's amplitude ceiling is equal to the highest point in the spectrum only if the spectrum happens to have been made from the time interval that includes the highest value in the spectrogram.

Boxy vs. smooth

Canary can display spectrograms in either of two styles, “boxy” or “smooth.” In a boxy spectrogram, each actual data point on the spectrogram grid is represented by a rectangular gray box. The width and height of the boxes depend on the grid resolution in the time and frequency dimensions respectively. Grid resolution is determined in the time dimension by frame length and frame overlap and in the frequency dimension by FFT size, as discussed in the section above on “Grid resolution.” The left edge of each box represents the beginning of a frame; the top edge is located at the analysis frequency whose amplitude is given by the darkness of the box. The size and visibility of the boxes on the screen depends on the size of the entire spectrogram on the screen, which in turn depends on the window size and the signal length (which affects the time dimension only), and on the degree of zoom invoked (by using the zoom and stretch controls). Figure 3.15 shows a boxy spectrogram of the entire time and frequency range of a 1.5-second signal and a portion of the spectrogram after a zoom; individual boxes are indistinguishable in the full-scale spectrogram, but are clearly visible in the magnified portion.

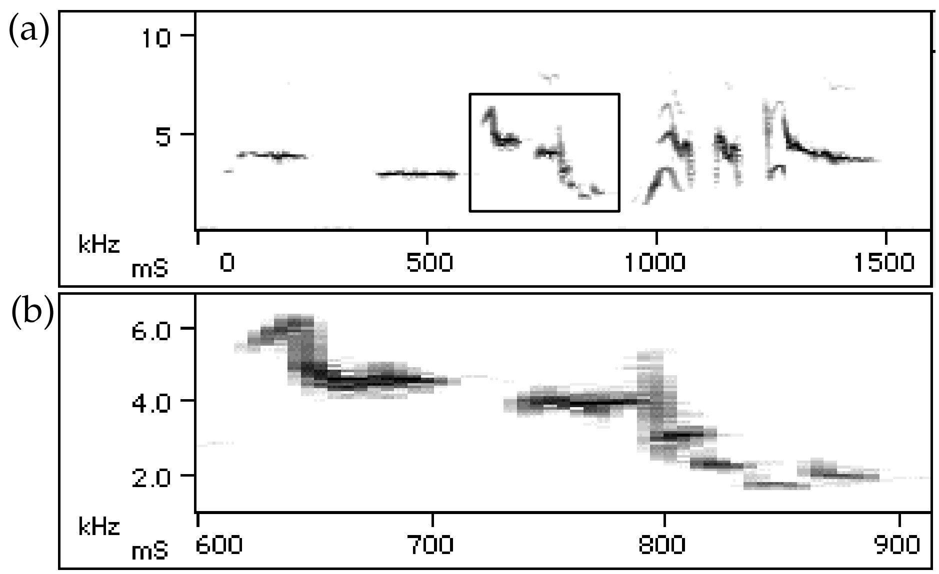

Figure 3.15. A boxy spectrogram at two different magnifications. **(a)** Song of a western meadowlark, digitized at 22.3 kHz. Filter bandwidth = 176 (frame length = 512 points = 23.0 mS), time resolution = 5.8 mS (overlap = 75%), frequency resolution = 43.5 Hz (FFT size = 512 points), window = Hamming. **(b)** Magnified view of section outlined in (a).

In a smooth spectrogram, the darkness of each individual screen pixel is determined by bilinear interpolation between the amplitude (or log amplitude) values calculated at the grid points. Each time the spectrogram is resized, the grayscale values for individual screen pixels are recalculated. Thus no matter how much you stretch a smooth spectrogram, you will not see sharp-edged boxes as you would with a boxy spectrogram. Display smoothing is not a substitute for finer grid resolution, however. Both may make a spectrogram more esthetically pleasing, but only finer grid resolution will reveal structural details of the signal that are invisible in a boxy low-resolution spectrogram (Figure 3.16). Smoothed spectrograms take longer to redraw than boxy spectrograms.

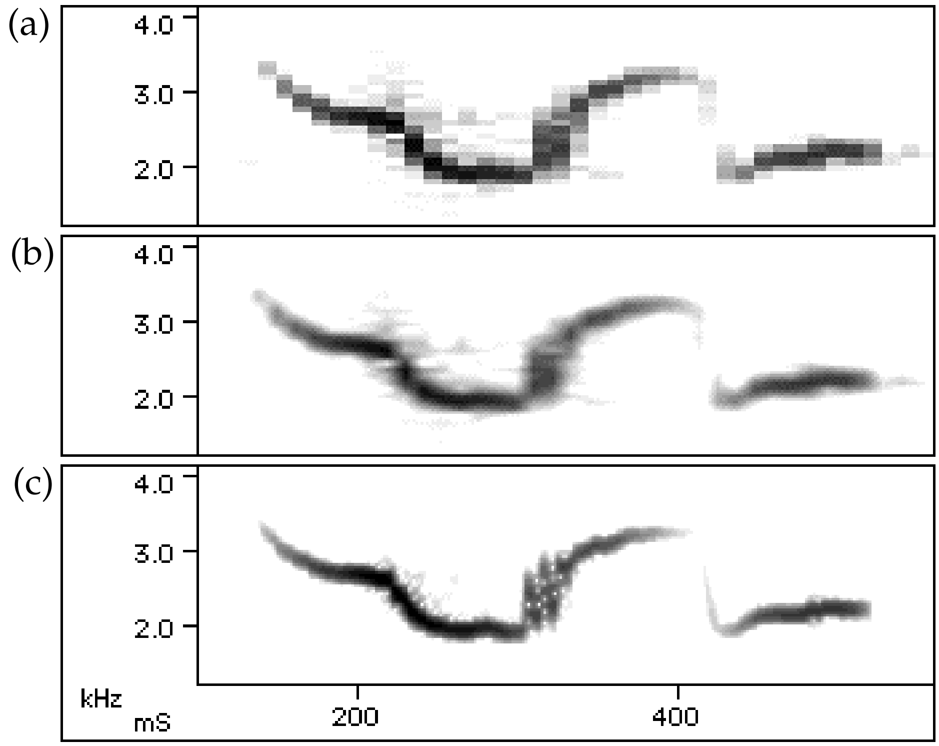

Figure 3.16. Difference between smooth-style display of spectrogram calculated on low-resolution grid and boxy display calculated on a higher-resolution grid. The signal is one syllable from an American robin song, digitized at 22.3 kHz. In all three spectrograms, filter bandwidth = 352.9 Hz (frame length = 256 points = 11.5 mS), window = Hamming. **(a)** Boxy spectrogram on low-resolution grid: 11.5 mS x 86.9 Hz (frame overlap = 0%, FFT size = 256 points). **(b)** Same spectrogram as (a), but with smooth display. **(c)** Boxy spectrogram with higher-resolution grid: 1.4 mS x 21.7 Hz (frame overlap = 87.5%, FFT size = 1024 points). Notice the rapid frequency modulation about halfway through the syllable that is detectable in (c) but not in (a) or (b).

Once a spectrogram has been calculated, you can switch back and forth between boxy and smooth displays without recalculating the entire spectrogram. Typing Command-G or option-clicking on the command panel's SPG button will bring up the spectrogram dialog box. If you change the display style without changing any other parameters, the spectrogram is immediately redrawn in the selected style.

Time and memory requirements

Higher quality spectrograms require more time and memory to calculate. The time and memory required are roughly proportional to the length of the sound, the sampling rate, the ratio *FFT Size / Frame Length*, and *1 / (1-Overlap)*. (For example, a spectrogram created with 93.8% overlap requires sixteen times as much memory and time as a spectrogram created with 0% overlap.)

Scaling and labeling of axes

Automatic scaling and labeling

When a file is first opened, or when a new sound is first recorded, Canary scales the axes of whatever panes are initially drawn to fit into a window 200 pixels tall and 504 pixels wide. The window can be resized at any time either by dragging its size box, or by clicking the zoom box in its title bar. The unit of measurement for each axis (e.g., µS, mS, Hz, kHz, etc.) is always selected to yield a maximum value between 0 and 1999.

Stretching and squeezing axes

The stretch and squeeze buttons on the command panel rescale the horizontal or vertical axis of the active pane by a factor of two. You can also use the zoom button to zoom directly to a selection, as discussed in Chapter 1.

When axes of any of the three panes are compressed, Canary displays a gray stippled background in time or frequency regions where there are no data. Thus, regions beyond the end of the signal in the time dimension are stippled in the waveform and spectrogram, and regions above the Nyquist frequency are stippled in the spectrogram and spectrum.[1]

[1]The Nyquist frequency is the highest frequency that can be represented in a digitized signal with a given sampling rate, and is equal to one half the sampling rate. See Appendix XX for further discussion.

Manual scaling

You can explicitly specify a scale factor for the horizontal and vertical axes of any of the three panes by activating the desired pane, and then selecting Manual Scaling... from the Options menu (Figure 3.17). Manual scaling affects only the active window unless you click the Apply to all windows checkbox. The Apply to all windows option is useful when comparing windows that were originally drawn at different scales. The scale that you specify applies until you resize the window, or change the scale by using stretch, squeeze, or zoom controls. See Chapter 8 for further discussion of manual scaling.

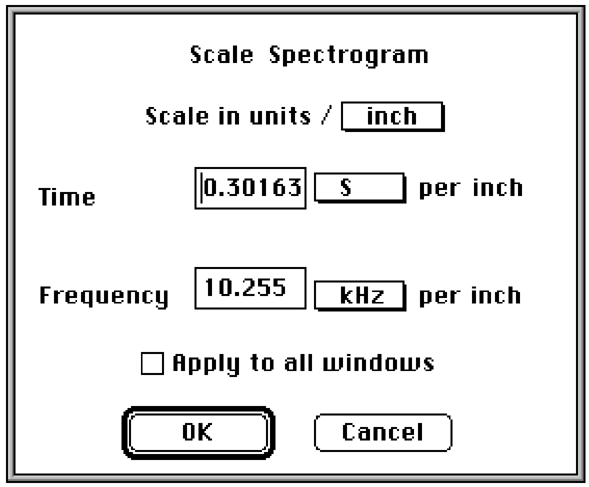

Figure 3.17. The manual spectrogram scaling dialog box.

Selections and cursors

Coupled selections

In general, when you select part of a display in any of the three panes, either by dragging the mouse pointer, shift-clicking, or using the cursors, corresponding portions of the other two panes are highlighted as well. The correspondence of highlighted areas in different panes depends on which dimensions are shared between the pane in which the selection was made and the other panes, as explained below.

Selections in the spectrum

Selections in the spectrum can only be made in the frequency dimension: you cannot select an amplitude range (but see below on the use of spectrum amplitude cursors). When you click anywhere above and to the right of the axes of a spectrum, the spectrum source interval is highlighted in the waveform. When you select a frequency range the source interval is highlighted in the waveform, and the corresponding time-frequency area is highlighted in the spectrogram.

Spectrum amplitude cursors

The spectrum amplitude cursors that appear on the vertical axis of the spectrum delimit the range of amplitude values that correspond to the range of grayscale values in the spectrogram. All amplitudes below the lower cursor appear in the spectrogram as white; amplitudes above the upper cursor appear as black. See the discussion of spectrogram contrast brightness controls earlier in this chapter for further discussion of spectrum amplitude cursors.

Selections in the waveform

Selections in the waveform can be made only in the time dimension: you cannot select an amplitude range. When you click anywhere above and to the right of the axes of a waveform, the entire spectrum is highlighted. When you select a time interval, the corresponding time interval in the spectrogram is highlighted, as is the entire spectrum.

Selections in the spectrogram

Selections in the spectrogram can be made in both the time and frequency dimensions. When you click anywhere above and to the right of the axes of a spectrogram, the blinking insertion cursor appears at the corresponding time in the waveform. When you select a time-frequency range in the spectrogram, the corresponding time interval is highlighted in the waveform, and the corresponding frequency range is highlighted in the spectrum. Note, however, that the displayed spectrum does not necessarily correspond to the selected part of the spectrogram unless the time interval selected is the spectrum source interval. If a spectrogram is displayed, and you want a spectrum of some part of it, select the desired time-frequency range, and make a new spectrum. The spectrum will be made from the selected time interval as discussed under “Number of frames to average” in the section on “Spectra” above; when the spectrum appears, the frequency range selected in the spectrogram will be highlighted.

Background processing

When a spectral analysis starts, Canary hides the command and measurement panels and displays a status box. You can continue to do other work with Canary while a spectrogram or spectrum is being calculated. Clicking in any signal window will bring that window forward, send the status window into the background, and activate the command and measurement panels and the track palette (if they were showing before you started the spectral analysis). All of Canary's controls and menus can be used normally while a spectral analysis is being calculated in the background. You can even have more than one spectral analysis running at once (provided sufficient memory is available). A spectral analysis will take longer to complete, however, if you are doing other work at the same time.

You can switch to another application program while Canary is calculating a spectral analysis. You may notice some degradation in the performance of other applications while Canary is doing a spectral analysis in the background. Such degradations will be more noticeable on slower machines.

You can adjust the tradeoff between computation speed and machine responsiveness during spectral analyses by selecting Speed... from the Options menu, and then selecting one of the three options under Pause Button Responsiveness. Depending on whether you select Pause Button Faster, Computations Faster, or No Pause Button, Canary allows more or less time for tasks other than calculating the spectral analysis. This choice affects the machine's responsiveness to all user interactions (not just clicking the Pause button), and the speeds of any other applications that are running at the same time as Canary.

Changing analysis parameters

You can recalculate a spectrum or spectrogram with different parameters either by choosing Make Spectrum... or Make Spectrogram... from the Panel menu, using the command key equivalents (Command-K or Command-G), or by holding down the *option* key while clicking on the SPK or SPG button on the command panel.

Spectrogram and spectrum measurements

Chapter 6 describes all of the measurements that you can make on spectrograms and spectra.

Chapter 4 Signal Amplitude Calibration

About this chapter

This chapter discusses how to use Canary for making calibrated absolute amplitude measurements. If the only measurements you need to make are time and frequency measurements, or measurements of the relative amplitudes of parts of a signal, you do not need to set the signal calibration as discussed below.

> **BEWARE!**
>
> *Amplitude measurements provided by Canary are* ***meaningless*** *unless the signal has been properly calibrated!*

In most cases, making calibrated amplitude measurements requires the use of calibrated microphones and tape recorders, and/or the use of a calibrated sound source to generate a calibration signal of known amplitude (see “Recording a calibration signal” below).

This chapter assumes a working knowledge of the relationships among sound pressure, intensity, and characteristic impedance in elastic media, and among voltage, power, and impedance in electric circuits, and with decibel measurements. If you are not familiar with these concepts, or need to review them, see Appendix C.

Acoustical and electrical paradigms

Canary allows you to make amplitude measurements from waveforms, spectrograms, and spectra using either an acoustical or an electrical paradigm.

In the acoustical paradigm, the waveform represents a time-varying pressure at some point in space in an elastic medium such as air or water. In a plane or spherical free-progressive sound wave, the squared effective (RMS) sound pressure is proportional to the average sound intensity over a given time interval at a given point in space. The proportionality constant relating squared sound pressure[1] and sound intensity is the inverse of the characteristic impedance of the medium. Characteristic impedance is the product of the medium’s density and the speed of sound in the medium. Sound pressure, sound intensity, and characteristic impedance are thus related by

---

[1]The term “sound pressure” here and elsewhere in this manual means RMS sound pressure, unless otherwise noted.

$$I = \frac{p^2}{\rho c} \qquad (4.1)$$

where I = intensity (in watts/m^2), p = RMS pressure (in pascals; one pascal equals one newton/m^2), ρc = the characteristic impedance of the medium (in mks rayls), ρ = the density of the medium (in kg/m^3), and c = the speed of sound in the medium (in m/sec).

In the electrical paradigm, the waveform represents a time-varying voltage in an electric circuit. In a given time interval, this voltage delivers some average power to the circuit. The power delivered by a given voltage depends on the impedance of the circuit, according to the relationship

$$P = \frac{v^2}{R} \qquad (4.2)$$

where P = power (in watts), v = RMS voltage, and R = impedance (in ohms).

A signal is considered to be an *acoustic signal* or an *electric signal* depending on which paradigm is selected in the Set Calibration dialog box, discussed later in this chapter.

Canary can express any amplitude-related quantity either in acoustical or in electrical units. Table 4.1 shows the correspondence between quantities expressed using the two paradigms.

Table 4.1. Correspondence between acoustical and electrical quantities.

SOUND		ELECTRICITY	
Quantity	**Units**	**Quantity**	**Units**
pressure	pascals	voltage	volts
intensity	watts/m^2	power	watts
characteristic impedance	rayls	line impedance	ohms
energy flux density	joules / m^2	energy	joules

The remainder of this chapter discusses the calibration procedure for sounds. The procedure for electrical signals is the same, except that electrical units are substituted for acoustical units according to the correspondences in Table 4.1.

Signal calibration in Canary: an overview

Canary's internal representation of a digitized signal consists of a sequence of *sample values* output by an analog-to-digital converter sampling an input voltage at some regular sampling frequency (see Appendix A). The sample values are integers ranging from $-(2^{n-1})$ to $2^{n-1}-1$, where n is the number of bits per sample. Thus, 8-bit samples range from -128 to 127, and 16-bit sample values range from -32768 to 32767. Each sample value is proportional to the instantaneous input voltage to the A/D converter at the time the sample was taken. This input voltage, in turn, may be proportional to an instantaneous sound pressure as measured by a microphone or other pressure transducer.

Information used in calibration

In order to display amplitude measurements in meaningful acoustical or electrical units, Canary makes use of four types of calibration information:

- A *paradigm* selection that specifies whether measurements are to be expressed in acoustical or electrical units.
- A *calibration factor* that is used to convert the dimensionless sample values in the digitized waveform into pressure values. Canary calculates the calibration factor directly from information that you supply about the pressure of a signal, or indirectly from information you supply about intensity and characteristic impedance, as discussed below under “Setting calibration parameters”.
- An *impedance* value that defines the relationship between the squared RMS pressure of a signal and the signal’s average intensity.
- *Reference values* for use in decibel measurements of pressure and intensity.

Making calibrated measurements

There are four steps to making calibrated amplitude measurements with Canary:

1. Record a calibration signal with the same system used for recording data.
2. Digitize the calibration signal and create a calibration document.
3. Set the calibration parameters in Canary for the calibration document.
4. Copy the calibration values from the calibration document window to a window containing the signal to be measured.

Each of these steps is discussed further below.

Recording a calibration signal

Canary’s calibration protocol relies on using a *calibration signal*. When working with acoustic signals, the calibration signal can be any signal for which you know the RMS or peak pressure and/or intensity of the signal and the characteristic impedance of the medium.[1] **The calibration signal should be recorded using the same equipment and settings as the signals to be measured.** Typically, the calibration signal is a pure tone.

One approach to recording an acoustic calibration signal is to record a calibrated sound source that produces a known RMS sound intensity or pressure at a specified distance. The recorded signal is then known to

[1]If both pressure and intensity are known, you do not need to know the impedance, because it can be calculated from the other two quantities.

represent the specified intensity or pressure, even if no calibration information is available for the recorder or microphone directly.

Alternatively, one can record a signal from an uncalibrated source, using a calibrated microphone and recorder. Some calibrated recorders (for example, the Nagra IV-SJ) provide a direct readout of the sound pressure of a recorded signal via the built-in level meter. A microphone-recorder combination might also be calibrated so that a given output voltage from the recorder is known to represent a certain sound pressure at the microphone. In this case, you would need to use an oscilloscope to determine the output voltage of the recorder while playing back the calibration signal, and then calculate the corresponding sound pressure.

Using either of these approaches, you would know the amplitude of the recorded calibration signal; this information is supplied to Canary as discussed below.

The characteristic impedance of the medium can be calculated from measurements of temperature and barometric pressure in air, or from measurements of temperature, depth, and salinity in water. See Appendix B for further information on calculating characteristic impedance.

Setting the calibration parameters for the calibration signal

For acoustic signals, you use Canary's Signal Calibration dialog box to specify either (a) the effective (RMS) or peak pressure of the calibration signal and the characteristic impedance of the medium, (b) the average (RMS) intensity of the calibration signal and the characteristic impedance of the medium, or (c) the RMS or peak pressure and average intensity of the signal. Canary then calculates the third quantity using Equation 4.1.

The Set Calibration dialog also allows you to specify the reference values for dB measurements of pressure and intensity, as discussed under "Decibel reference values" in the next section of this chapter. (See Appendix C for a discussion of reference levels for dB measurements.)

Copying calibrations between signals

Once the calibration values have been set for a calibration signal, they can be copied from the calibration signal and pasted to other signals that were recorded under the same conditions using the same equipment and recording settings.

Setting calibration parameters

In order to set the signal calibration parameters, you must create a *calibration document*. A calibration document is simply a Canary signal window that contains a calibration signal of known intensity and effective pressure. The calibration document should contain *only the calibration signal*, without any leading or trailing period of silence or noise.

It is important that you use the same gain control settings in the Sound Acquisition/Recording dialog when acquiring the calibration signal and the signals from which you plan to make measurements. In particular, Automatic Gain Control should always be off when acquiring any signals that are to be used for calibrated amplitude measurements. (See Chapter 2 for further information on setting the gain for signal acquisition.)

Choosing the measurement paradigm

With the calibration document as the frontmost window in Canary, select Set Calibration from the Calibration submenu (on the Options menu). The Signal Calibration dialog (Figure 4.1) will appear. Using the pop-up menu at the top of the dialog, choose either Sound or Electricity. The setting of this menu determines how values are labeled in the dialog, and also selects whether Canary displays measurements in acoustical or electrical units.

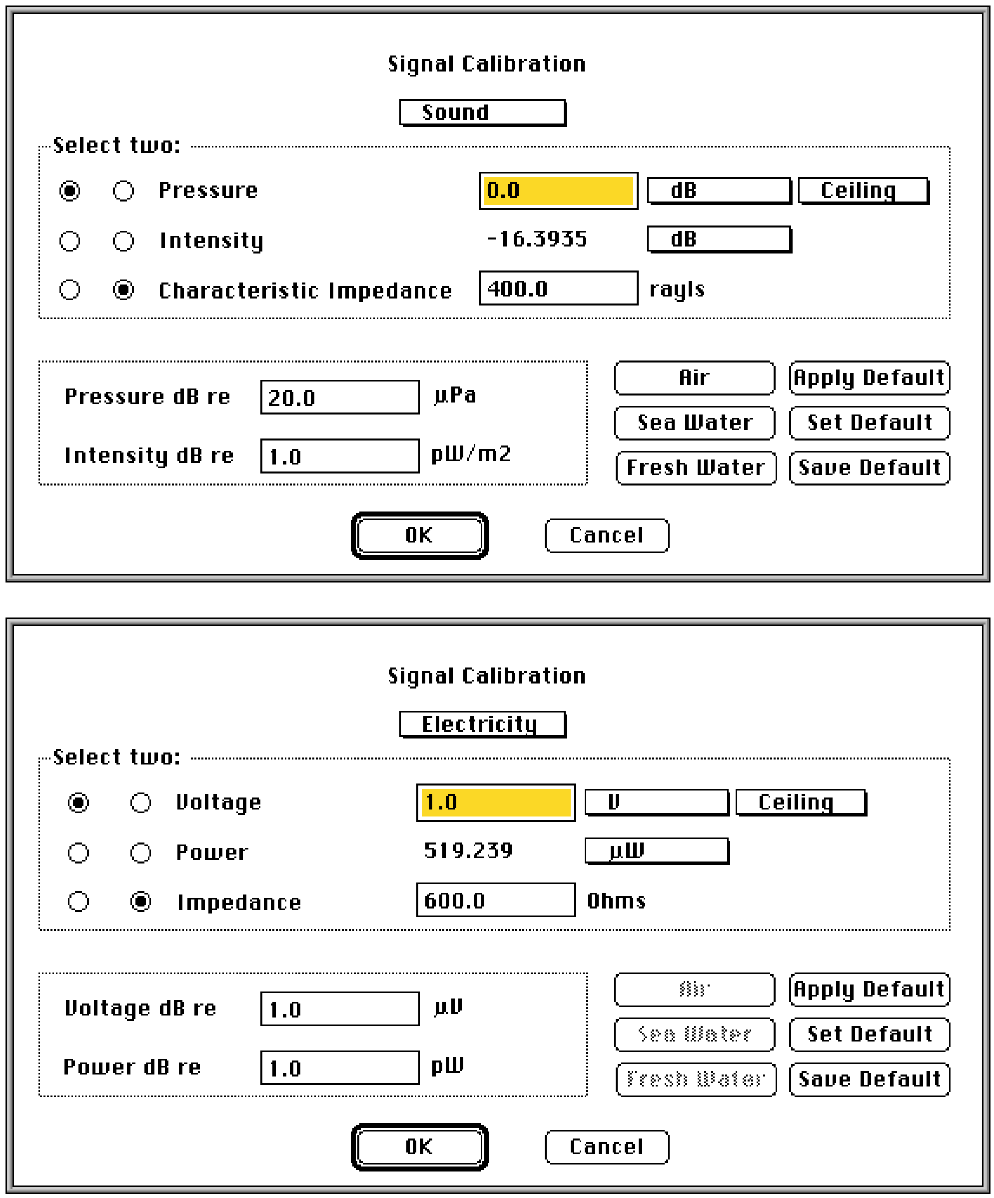

Figure 4.1. The Signal Calibration dialog box. The upper panel shows parameters for calibrating sound signals, the lower panel shows parameters for calibrating electrical signals.

Selecting calibration parameters

Select two calibration parameters by clicking two radio buttons next to the terms Pressure, Intensity and Characteristic Impedance. The values shown for the two parameters that you select can be edited. The value of the third parameter is calculated by Canary from the other two, according to Equation 4.1 or 4.2.

Pressure

The pop-up menu next to the value shown for Pressure indicates whether pressure is expressed in absolute pressure units (pascals, abbreviated Pa) or in decibels (dB). If you choose dB, the value in the Pressure field is taken as dB relative to the pressure reference value specified in the Pressure dB re field in the lower left part of the dialog box (see below). You should select the units before entering a value into the Pressure field.

A second pop-up menu lets you specify whether the value in the Pressure field will be the RMS, Peak, or Ceiling value for pressure. Figure 4.2 illustrates the relationship between Ceiling, Peak, and RMS values. If you choose RMS, the value is taken as the RMS pressure value of the entire signal. If you choose Peak, the value is the pressure value of the individual sample with the largest absolute value. If you choose Ceiling, the value is the pressure represented by a sample in which all bits are equal to 1.[1] You should select RMS, Peak or Ceiling before entering a value into the Pressure field.

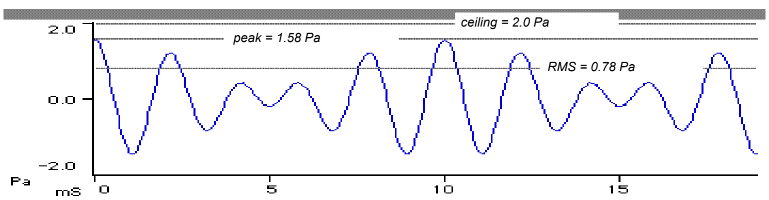

Figure 4.2. Relationship between ceiling, peak, and RMS values of a signal. Canary lets you assign an absolute amplitude value to any of these three values. This signal has been calibrated to set the ceiling equal to 2.0 Pa (=100 dB re 20 μPa).

If you selected Pressure as one of the two parameters to be set (by clicking on either of its radio buttons), enter the RMS, peak, or ceiling pressure of the entire signal. If you are using a pure-tone calibration signal of known amplitude, as discussed above, you will probably want to specify the RMS or peak pressure.

If you did not select Pressure as one of the two values to be set, the pressure shown will be the pressure that Canary has calculated from the current values for Intensity and Impedance.

[1]The ceiling value thus corresponds to 128 for 8-bit samples, or 32768 for 16-bit samples.

Intensity

The pop-up menu next to the Intensity field lets you choose whether to express Intensity in absolute units (watts/m^2) or in dB. If you choose dB, the value in the Intensity field is taken as dB relative to the intensity reference value specified in the Intensity dB re field in the lower left part of the dialog box (see below). You should select the units before entering a value into the Intensity field.

If you selected Intensity as one of the two parameters to be set (by clicking on either of its radio buttons), enter the average intensity of the entire signal.

If you did not select Intensity as one of the two values to be set, the intensity shown will be the intensity that Canary has calculated from the current values for Pressure and Impedance.

Characteristic impedance

The characteristic impedance of an elastic medium (in mks rayls) is equal to ρc, where ρ = the density of the medium (in kg/m^3) and c = the speed of sound in the medium (in m/sec). Appendix C includes information on estimating ρ and c in air and water.

If you selected Characteristic Impedance as one of the two parameters to be set (by clicking on either of its radio buttons), enter the characteristic impedance of the medium in which the signal was recorded. Alternatively, you can use the buttons marked Air, Sea water, and Fresh water to have Canary fill in approximate impedance values for these three media (see below).

If you did not select Characteristic Impedance as one of the two values to be set, the impedance shown will be the impedance that Canary has calculated from the current values for Pressure and Intensity.

Decibel reference values

The fields Pressure dB re and Intensity dB re let you specify reference values (in pascals and watts/m^2 , respectively) for decibel measurements of pressure and intensity. These reference values are used in calculating and interpreting dB values elsewhere in the signal calibration dialog, and also apply to dB values shown by Canary in the spectrum pane of a signal window and in the measurement panel.

For sound transmitted through air, the usual pressure standard is 20 µPa; the usual intensity standard is 1×10^{-12} W/m^2, or 1 pW/m^2. In sea water, the usual standards are 1 µPa for pressure, and 6.7×10^{-19} W/m^2 for intensity.

Air, sea water, fresh water buttons

When the Sound paradigm is selected, three buttons labeled Air, Sea water, and Fresh water are available in the dialog box. Clicking on any one of these buttons automatically selects a standard characteristic impedance value, and standard dB reference values for pressure and intensity (Table 4.2).

Table 4.2. Values assigned for characteristic impedance, reference pressure, and reference intensity by the Air, Sea Water, and Fresh Water buttons in the Signal Calibration dialog box.

	Air	Sea Water	Fresh Water
Characteristic impedance:	400 rayls	1.5458 Mrayls	1.46389 Mrayls
Reference pressure:	20 µPa	1 µPa	1 µPa
Reference intensity:	1 pW/m^2	.6469 aW/m^2	.6831 aW/m^2

The characteristic impedance value set by the Air button is calculated (using Equation 4.1) from the reference pressure of 20 µPa and the reference intensity of 1 pW/m^2 (=10^{-12} W/m^2). The resulting value (400 rayls) is actually slightly lower than the characteristic impedances that would typically be encountered over common ranges of temperature and barometric pressure at sea level.[1] Appendix B includes formulas for calculating characteristic impedance from temperature and pressure.

The characteristic impedance values that are set by the Sea water and Fresh water buttons are for temperature = 15°C, depth = 1 m and salinity = 35 and 0 parts per thousand for sea water and fresh water respectively. The values are derived from estimates of density calculated according to Millero et al. (1980), and sound speed calculated according to Mackenzie (1981) (see references at end of Appendix B). The reference intensity for dB values is calculated (using Equation 4.1) from the characteristic impedance value and a reference pressure of 1 µPa.

Default calibrations

Canary maintains a set of default calibration parameters that are used to calibrate new signals. A new signal is any signal that is digitized within Canary or that is imported by opening a file created by another program (including older versions of Canary that do not use all of the calibration information described in this chapter).

[1]At temperature = 20°C and pressure = 1000 mbar, the characteristic impedance of air is 408 mks rayls. Intensities calculated from a given pressure value using ρc = 400 rayls are about 0.1 dB higher than intensities calculated using ρc = 408 rayls.

Factory default calibration

When you first use Canary, the default calibration is set as shown in Figure 4.1. These parameter values are called the "factory" default calibration. The factory default is used whenever Canary does not find a preference file called "Canary.Prefs" when the program launches. The factory default calibration is set to a ceiling of 0 dB so that new signals will have negative intensity and pressure levels until they have been properly calibrated. See Chapter 8 for more on preference files and factory defaults.

Set default

To set the default calibration parameters equal to the ones that are currently displayed, click the Set default button. New signals will be calibrated using these parameters until you quit the program, re-set the default calibration, or load a new preference file.

If you set a default calibration using a ceiling value, that ceiling value will be assigned to all new signals that are recorded within Canary and to imported signals that have the same number of bits per sample as the orginal calibration signal. However, new signals that are *imported* (e.g., from some other program) with a different sample size will be assigned a ceiling value that differs from the calibration signal by 6 dB per bit difference in sample size. For example, if you set the default calibration for a 16-bit signal to be a ceiling value of 0 dB, new imported 8-bit signals will be assigned a ceiling of -48 dB (= (16 - 8 bits) × 6 dB/bit).

Save default

To save the current calibration parameters into the current preference file, click on the Save default button. The current calibration parameters will be saved even after you quit Canary and will be applied to all new signals whenever the current preference file is in use. If you have not explicitly loaded a new preference file since Canary was launched, the calibration parameters will be saved into Canary's default preference file. See Chapter 8 for information on preference files.

Apply default

To apply the default calibration parameters to a signal, click the Apply default button. The default calibration parameters are the values that are stored in the current preference file, unless another set of values has been set using the Set default button.

Copying and pasting calibrations

Once you have specified a set of calibration parameters for a calibration signal, you can apply the same calibration to other signals that were recorded with the same sample size.[1] With the calibration document as the active window, select Copy Calibration from the Calibration submenu on the Options menu. Click on the window that you wish to calibrate (or select its name from the Windows menu), and then select Paste Calibration from the Calibration submenu on the Options menu. The signal in the second window will be

[1]Although Canary will allow you to paste calibrations between signals that were recorded with different sample sizes, the results are generally not what was intended.

recalibrated instantly, causing the window and the measurement panel (if it is showing) to be redrawn to reflect the changed calibration.

If you have multiple signal windows open, you can recalibrate them all at once by selecting Paste Calibration To All from the Calibration submenu on the Options menu.

Calibrations and signal editing

When you copy or cut part of a signal, the signal's calibration information is copied to the clipboard with the data. When you then paste data from the clipboard into another signal, this information is preserved. For example if you copy a signal having a peak intensity of 100 dB, and paste it into another signal, the pasted data will still have a peak of 100 dB, even if the two signals were calibrated differently. (The pasted data may look different in the waveform window if the drawing scale of the destination window is different from that of the source window.)

Although Canary will allow you to paste data between signals with different calibration paradigms (acoustic and electric), the amplitude of the resulting signal would be impossible to interpret.

Setting the dB reference using a spectral peak

When making measurements of the relative spectrum levels within a signal, it is often useful to assign a particular dB value (usually 0) to a peak in a spectrum, and then measure spectrum levels relative to this peak. You can reset the dB reference value for a spectral peak by selecting part of the spectrum, and then choosing Set dB reference... from the Calibration submenu on the Options menu. The dialog that appears displays the dB level of the highest peak in the selected band of the intensity spectrum, and allows you to edit the value displayed. If you change the value and click OK, the reference value for dB intensity levels is re-set so that the selected peak will have the dB intensity level that you specify.

When you reset the intensity dB reference value in this way, the reference value for pressure dB is automatically reset to preserve the dB difference between the average intensity and the RMS pressure of the entire signal.

Making calibrated amplitude measurements

Once you have properly calibrated a signal (either directly using the Signal Calibration dialog or indirectly by copying and pasting a calibration from another signal), you can use Canary's measurement panel to make reliable amplitude measurements from waveform, spectrogram, and spectrum views of the signal. For further details on measurements, see Chapters 6 and 3.

Chapter 5 Multi-track Documents

About this chapter

This chapter provides a reference for working with multi-track documents. For information about recording two-track (stereo) signals, see Chapter 2.

About multi-track documents

Canary can work with documents that contain multiple *tracks*.. Each track, which can contain a different sound, is shown in a separate pane of the signal window. Figure 5.1 shows a two-track window containing a single call of a bearded seal recorded in the Arctic Ocean by two hydrophones located about a kilometer apart from each other. The seal was closer to the hydrophone recorded in Track #1 (the upper track), so the sound appears in Track #1 a fraction of a second earlier than it appears in Track #2.

If you are using Canary to acquire multi-track data, the recording is limited to two tracks (i.e., stereo), which have a common sampling rate. Canary can import files created with other programs containing up to 32 tracks.[1] All tracks in a multi-track document must contain the same number of samples, and must be sampled at the same rate.

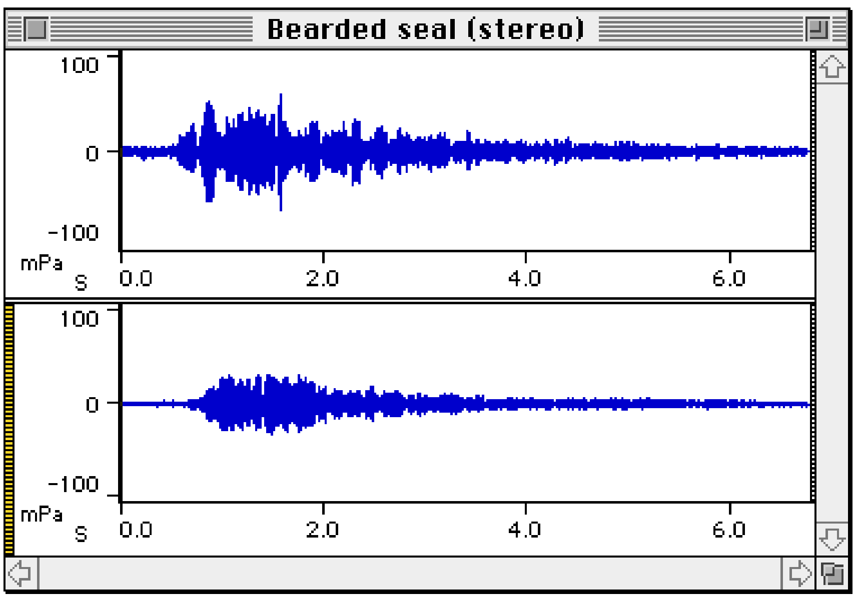

Figure 5.1. A 2-track signal window showing waveforms only.

[1]Canary can import AIFF or MATLAB files containing up to 32 tracks; SoundEdit files can have up to two tracks.

The number of tracks in a document is determined when data are first recorded or pasted into the document, or when the document is created in another program. You cannot add or remove tracks once the number of tracks in a document has been determined. You can, however, copy or save selected tracks (see the sections on "Editing multi-track signals" and "Saving selected tracks" later in this chapter).

If you request a spectrogram or spectrum in a multi-track document, the analysis is always done for all tracks.

The track palette

When you activate a window containing a multi-track document, Canary displays a small control window called the *track palette* (Figure 5.2). You can hide a document's track palette by clicking the palette's close box. You can display the track palette if it is hidden by choosing Track Palette from the Window menu.

Figure 5.2. The track palette for a stereo (2-track) signal.

Showing and hiding tracks

When you first open or record a multi-track document, all tracks are displayed. Tracks are numbered consecutively, with Track #1 displayed at the top of the window. (See the section on "Display options for multi-track documents" later in this chapter for an explanation of how to display track numbers in the signal window.) You can toggle a track between being shown and being hidden by clicking on its rectangular Track *N* button in the track palette. One track must always be showing in any document.

Track selections

The small square buttons on the righthand side of the track palette are *track selection* buttons. When a track is selected, its selection button is depressed and contains a black circle. The track selection buttons are used to specify which tracks will be affected by editing commands and the Save sound tracks... command, as discussed in the sections on "Editing multi-track signals" and "Saving selected tracks" later in this chapter. Track selections also control which tracks are played by the Play Selection command or button.

Multiple selections

You can automatically select all tracks by clicking the All button in the track palette.

If you hold the *option* key down while clicking on a track's show/hide button, the track's button is marked with a small circle (Figure 5.3). The window is not redrawn to change which tracks are displayed until you click the Apply button. Such “multiple selections” are useful for preventing unnecessary redrawing of the window with documents containing more than two tracks.

Figure 5.3. The track palette showing three tracks designated by holding the *option* key while clicking on the track buttons. Clicking the Apply button in the palette causes all three tracks to be hidden at once.

Display options for multi-track documents

Grouping views by track

Canary can arrange the views (waveform, spectrogram, and spectrum) in a multi-track window in two different ways. If you select Group track panes in the Display Options dialog box (Options menu), Canary arranges all of the views for a given track in adjacent panes (e.g., shown from top to bottom in the window: *SPK-1, SPG-1, WVF-1, SPK-2, SPG-2, WVF-2*). If Group track panes is not selected (the default), all of the views of each type are arranged together (e.g., *SPK-1, SPK-2, SPG-1, SPG-2, WVF-1, WVF-2*).

Track labels

Canary's track/pane labels are useful for identifying the track numbers of tracks displayed on screen (Figure 5.4), especially when only some of the tracks in a document are being displayed. You can turn on track/pane labels from the Display Options dialog box (Options menu). See Chapter 8 for further discussion of track/pane labels.

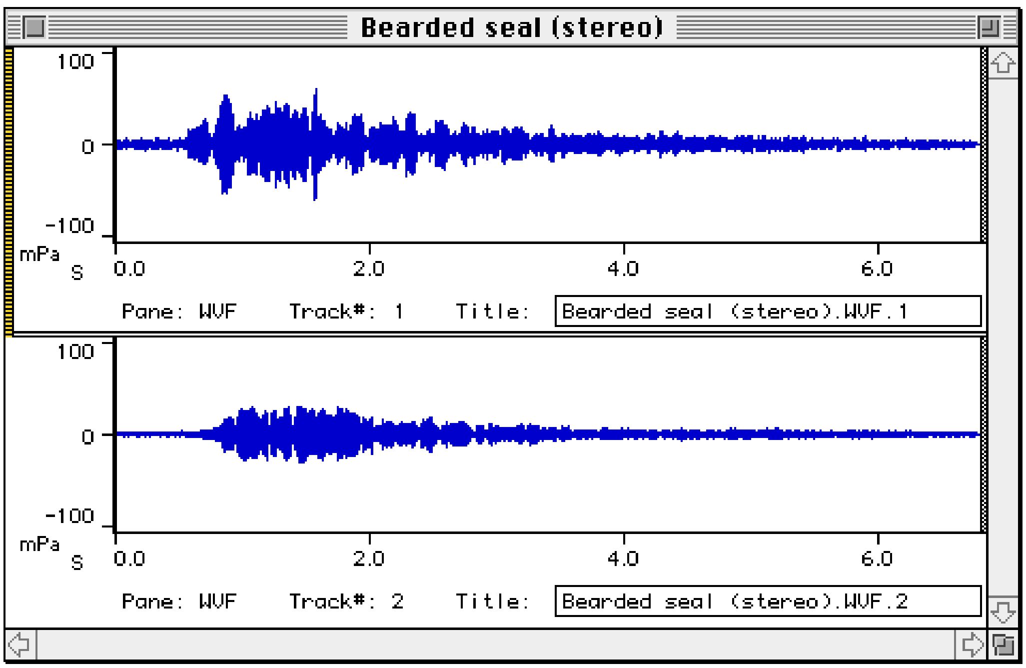

Figure 5.4. Two-track waveform showing track/pane labels.

Playing back multi-track sounds

You use the track palette's selection buttons to select one or more tracks of a multi-track document for playback. No track selection is equivalent to selecting all tracks. If the Macintosh is equipped for stereo playback, and Stereo Playback is selected in the Multi-track Options dialog, stereo tracks will be played back via the Macintosh's left and right sound output channels. If the Macintosh lacks stereo playback capability, or if Mono Playback is selected in the Multi-track Options dialog, the selected tracks are mixed and played together.

Editing multi-track signals

Whenever you choose one of Canary's signal-editing commands (Cut, Copy, Paste, Clear, Amplify, or Filter) while working with a multi-track document, the command is performed only on the tracks that are selected on the track palette. **If no tracks are selected at all, Canary performs the operation on all tracks.**

Constraints on multi-track editing

All of the tracks in a multi-track document must always remain the same length. If you perform an editing operation that would change the length of selected tracks (i.e., Cut, Paste, or Clear), Canary adjusts the length of the selected tracks or the non-selected tracks by "zero-padding" (inserting silence) as described for each command below.

Clear

When you select an interval in a waveform and choose Clear from the Edit menu (or press the *delete* key), the highlighted data are deleted from the selected tracks. Zeros are appended to the end of the selected tracks to preserve their original length.

Cut

When you select an interval in a waveform and choose Cut from the Edit menu (or press Command-X), the highlighted data are copied to the clipboard and deleted from the selected tracks. Zeros are appended to the end of the selected tracks to preserve their original length.

Copy

When you select an interval in a waveform and choose Copy from the Edit menu (or press Command-C), the selected data are copied to the clipboard. You can view and play the sound in the clipboard by choosing Show Clipboard from the Edit menu.

Paste

When pasting data into a multi-track document, the number of tracks selected must equal the number of tracks in the clipboard. Track numbers are ignored when pasting into an existing multi-track document, but the order of tracks is preserved. For example, if you copy tracks 1 and 2 from one document (the "source" document) and paste them into tracks 2 and 4 in a second document (the "destination" document), the data from the source Track 1 will be pasted into Track 2 in the destination document, and the Track 2 data from the source document will be pasted into Track 4 in the destination document.

When pasting data into a new document (one that contains no data), all of the tracks in the clipboard are pasted, thereby determining the number of tracks in the new document. The tracks are renumbered consecutively starting at Track 1.

When pasting data into selected tracks of a multi-track document, Canary zero-pads the end of each non-selected track in order to make all tracks the same length as the selected tracks after the paste.

Saving selected tracks

You can save selected tracks of a multi-track document by choosing Save Sound Tracks... from the File menu. Only tracks whose selection buttons are on will be saved into the file you specify. The tracks in the new saved document will be in their original order, but will be renumbered consecutively starting at Track 1.

Files saved in SoundEdit format cannot have more than two tracks.

Multi-track correlations

Canary's correlator can be used to generate cross-correlation functions of waveforms or spectrograms of individual tracks within multi-track files. You can calculate cross-correlations of tracks within a single file by simply selecting the same file as both the first and second inputs to the correlator. See Chapter 7 for discussion of the correlator.

Multi-track correlations (i.e., correlations in which at least one document is a multi-track document) are displayed in a multi-pane correlation window showing all possible cross-correlations (Figure 5.5).

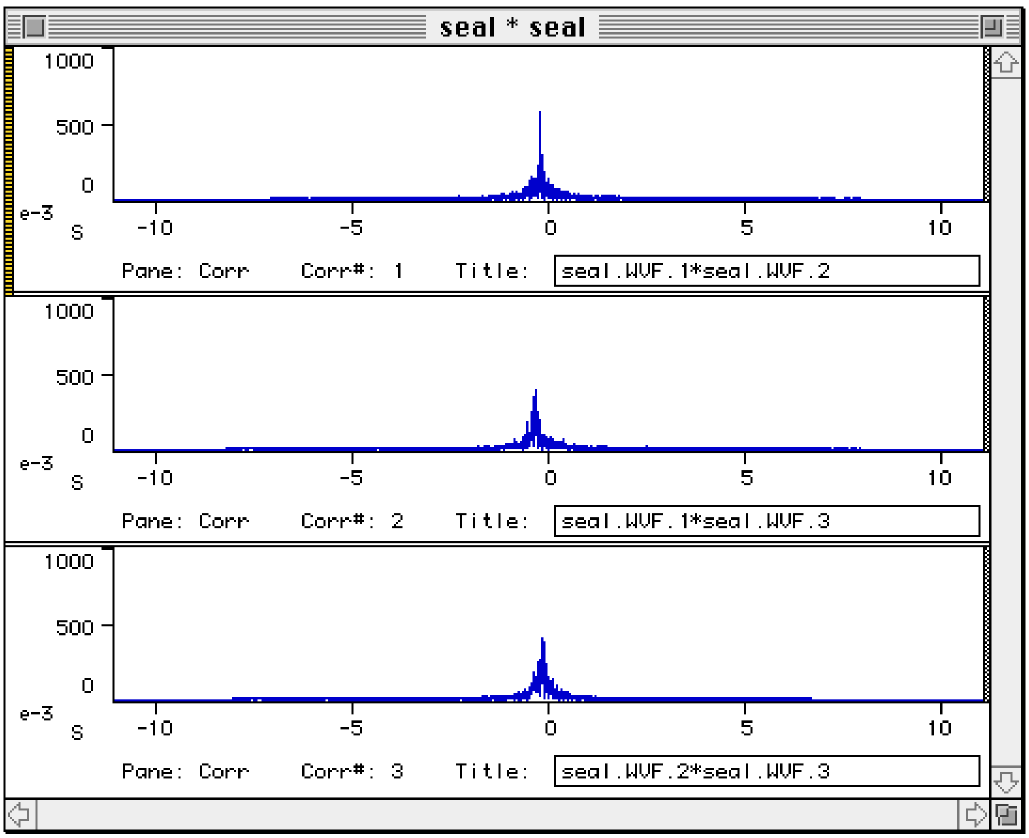

Figure 5.5. A correlation window showing the three correlations obtained by correlating a three-track file against itself. The number after "WVF." in the title assigned to each pane indicates which tracks are used in the correlation.

For correlations of a file containing *N* tracks with itself, the number of correlations shown is $^1/_2[N(N\text{-}1)]$. (Canary automatically omits the correlation of each track with itself.) For correlations between two files containing N_1 and N_2 tracks, the number of correlations is N_1N_2.

To determine which correlation function in a multi-track correlation window corresponds to any given pair of tracks, turn on the Show Track/Pane Labels switch in the Display Options dialog box (Figure 5.5).

Chapter 6 Measurements

About this chapter

Canary's measurement panels can be configured to display any combination of measurements and parameters for each of the three signal views and for correlation plots. Over 80 individual measurements and parameters are available for display in the measurement panels. Measurements can be transferred from a measurement panel to the data log with a single mouse-click. The data log can then be saved in a variety of formats for easy export to statistical, spreadsheet, word-processing, or other programs.

This chapter first explains the basic operation of the measurement panels and shows how to configure panels to display the desired measurements. Two reference sections then briefly describe each of the parameters and measurements available. The final section discusses how to transfer measurements to the data logger and export of data to other application programs.

The measurement panel

The measurement panel can be displayed by selecting Measurement Panel from the Windows menu, or by typing Command-M. Selecting Measurement Panel or typing Command-M (or clicking in the panel's close box) while the measurement panel is showing hides the panel.

Canary provides separate measurement panels for waveform, spectrogram, and spectrum signal views, and for correlation functions. At any one time, only the measurement panel for the active pane (which may be one pane of a signal window or a correlation window) is displayed. When you switch panes by clicking on a different pane, the panel is rebuilt for the new active pane. The data that are displayed for each pane depend on how the panel has been configured, as discussed in the section on "Configuring the measurement panel" later in this chapter. In this manual, references to "the measurement panel" mean whichever measurement panel is currently displayed.

The measurement panel contains one or more *cells* that display either of two types of data, called *parameters* and *measurements* (Figure 6.1). Parameters, which are displayed in the upper row of the panel, are characteristics of an entire signal, view, or correlation, and are unaffected by the position of the mouse pointer or the location of a selection. For example, parameters for the waveform pane include signal duration and sampling rate. Measurements, which are displayed in the lower row of the panel, can be either *point measurements* or *range measurements*. Point measurements refer to the location of the mouse. As you move the mouse within the active pane, point measurement cells are continually updated to reflect the mouse position. If the mouse moves out of a pane's plot area, point measurement cells display "--".

Range measurements refer to the selected region in the active pane, and are updated only when you change the selection (not whenever you move the mouse). For example, in the waveform measurement panel (Figure 6.1) Time and Amplitude are point measurements; ΔTime is a range measurement.

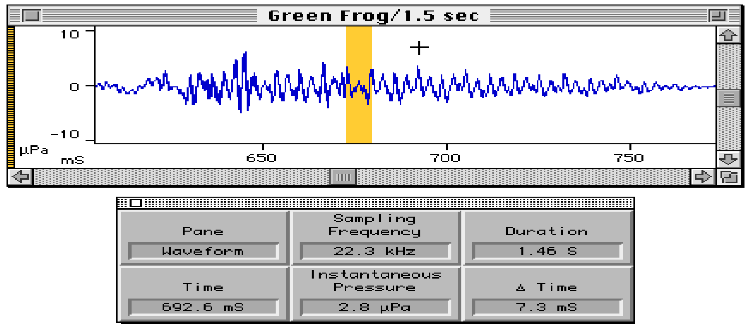

Figure 6.1. A signal window and the default waveform measurement panel.

Configuring the measurement panel

Selecting Measurement Panel... from the Options menu brings up the Measurement Panel Configuration dialog box (Figure 6.2), which lets you specify which measurements and parameters are displayed by each of the four measurement panels.

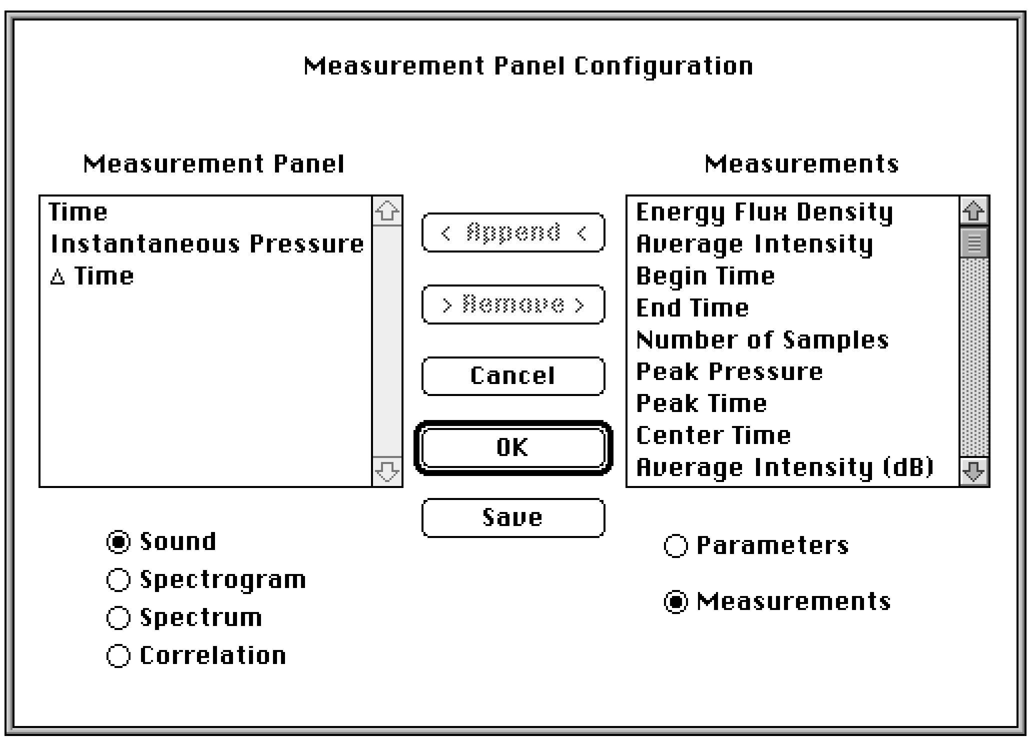

Figure 6.2. The Measurement Panel Configuration dialog box.

The radio buttons on the lower left are used to specify which of the four panels is to be configured. The two list boxes display either parameters or measurements, depending on which radio button is selected. The left list displays the measurements or parameters that are currently included in the selected measurement panel. The right list displays the measurements or parameters that are available for inclusion in the panel. To add a measurement or parameter to the panel, select its name by clicking on the right list, and then click the Append button. To delete a measurement or parameter from the panel, select its name by clicking on the left list, then click the Remove button.

To insert a measurement or parameter before a measurement that is already on the measurement panel, select the item in the left list before which you want to insert the new item. Select the item(s) on the right list that you want to insert, and click on the Insert button (which replaces the Append button when an item is highlighted in the left list).

You can select more than one item in either list by holding down the *shift* or Command key while clicking the mouse. Shift-clicking selects all of the items on the list between the first and second items clicked; Command-clicking selects only those items on which you click (and not the intervening items).

Measurements are displayed from left to right in the panel (and in the data log) in the order that their names appear on the measurement list in the configuration dialog box.

The Save button saves the current configuration for all measurement panels to the current preference file.

Parameters

This section briefly describes all of the parameters available for all four measurement panels. Parameters are listed alphabetically. The name of each parameter is followed by a parenthesized list of the panels for which the parameter is available.

Boxy or Smooth? ***(spectrogram)***

Labelled “Display Style” in measurement panel; displays “Boxy” or “Smooth”.

Calibration Factor ***(all)***

Calibration factor of the active signal. The calibration factor is the number by which the raw sample values of the digitized signal are multiplied to obtain a pressure or voltage measurement. Units: pascals (for acoustic signals) or volts (for electric signals).

Clipping Level ***(spectrogram, spectrum)***

Clipping level of the active spectrogram or spectrum pane. Displays “-∞” if the amplitude scale is quadratic. Units: dB.

Complex Envelope *(correlation)*

Displays “Yes” or “No” to indicate whether the complex envelope of the correlation function (rather than the correlation function itself) is plotted.

Duration *(waveform, spectrogram)*

Duration of the entire signal. Units: seconds.

FFT Size *(spectrogram, spectrum)*

Number of points in the FFT used to construct the active spectrogram or spectrum. The number of frequency bins in the spectrogram or spectrum equals half the FFT size.

Filter Bandwidth *(spectrogram, spectrum)*

Filter bandwidth corresponding to the frame length and window function of the active spectrogram or spectrum. Units: Hz.

Filtered *(correlation)*

Displays “Yes” or “No” to indicate whether the signals were bandpass filtered before being correlated.

Frame Size *(spectrogram, spectrum)*

Number of points in a single STFT frame of the active spectrogram or spectrum.

Frequency Resolution *(spectrogram, spectrum)*

The frequency grid resolution of the active spectrogram or spectrum. The vertical spacing between gridpoints in a spectrogram, or the horizontal spacing between points in a spectrum. Units: Hz.

From Frequency *(correlation)*

The lower frequency limit of the passband in a frequency-filtered correlation. (See also To Frequency) Units: Hz.

Hop Size *(spectrogram, spectrum)*

The time grid resolution of the active spectrogram or spectrum, expressed in points. The number of samples between gridpoints in a spectrogram, or between the beginnings of successive frames whose spectra are averaged to obtain the displayed spectrum. (See also Time resolution.) Units: points.

Length in Samples *(correlation)*

Number of points in the correlation function. For a waveform correlation, this equals the sum of the number of digitized samples in the signals that were correlated, minus one. For a spectrogram correlation, Length in Samples equals the sum of the number of frames in the spectrograms that were correlated, minus one.

Length in Samples *(waveform)*

Number of digitized samples in the entire signal. The number of samples in a signal equals the signal's duration times the sampling rate.

Log Scale? *(spectrogram, spectrum)*

Labelled "Amplitude" in measurement panel. Type of amplitude scale of the active spectrogram or spectrum pane; displays "Quadratic" or "Logarithmic".

Normalized? *(correlation)*

Displays "Yes" or "No" to indicate whether the plotted correlation values are normalized.

Number of Frames *(spectrogram, spectrum)*

The number of STFT frames from which the active spectrogram or spectrum was calculated. For a spectrogram, equivalent to the number of (possibly overlapping) "time slices," which is determined by the signal length, frame length, and frame overlap. For a spectrum, Number of Frames is determined by the duration of the selection from which the spectrum was made, the frame length, and the frame overlap.

Overlap *(spectrogram, spectrum)*

The percentage overlap between successive STFT frames in a spectrogram or spectrum.

Pane *(all)*

Displays "Waveform", "Spectrogram", "Spectrum", or "Correlation", depending on which measurement panel is currently displayed.

Sample Size *(waveform)*

The number of bits per sample in the signal when it was originally digitized or imported into Canary.

Sampling Frequency *(all)*

Number of digital samples per second with which the signal was digitized. Units: Hz.

Start Time *(spectrum)*

The time (measured from the beginning of the signal) at which the spectrum source interval begins. Units: seconds.

Stop Time *(spectrum)*

The time (measured from the beginning of the signal) at which the spectrum source interval ends. Units: seconds.

Time Resolution *(spectrogram, spectrum)*

The time grid resolution of the active spectrogram or spectrum. The horizontal spacing between gridpoints in a spectrogram, or the time interval between the beginnings of successive frames whose spectra are averaged to obtain the displayed spectrum. (See also Hop Size.) Units: seconds.

To Frequency *(correlation)*

The upper frequency limit of the passband in a frequency-filtered correlation. (See also From Frequency) Units: Hz.

Window Function *(spectrogram, spectrum)*

The name of the window function applied to each STFT frame of data in calculating the spectrum or spectrogram.

Measurements

This section briefly describes all of the measurements available for all four measurement panels. Measurements are listed alphabetically. Each measurement is labelled as to whether it is a point or range measurement. The name of each measurement is followed by a parenthesized list of the panels for which the measurement is available.

Note that dB measurements less than 0.1 dB are displayed as 0 dB.

BEWARE!

Amplitude measurements provided by Canary are ***meaningless*** *unless the signal has been properly calibrated!*

Amplitude Ceiling *(spectrogram, spectrum)*

The amplitude ceiling of a spectrogram is controlled by the upper amplitude cursor in the signal's spectrum pane. All amplitudes (i.e., intensities for acoustic signals, powers for electric signals) above the amplitude ceiling are displayed as black. (Some amplitudes below the ceiling may also be black.) The amplitude ceiling can be adjusted directly by moving the upper amplitude cursor in the spectrum, or indirectly using the spectrogram brightness and contrast sliders on the command panel. Units: dB for logarithmic spectrograms; watts/m^2/Hz for quadratic spectrograms of acoustic signals; watts/Hz for quadratic spectrograms of electric signals. See also Amplitude Floor, Dynamic Range.

Amplitude Floor ***(spectrogram, spectrum)***

The amplitude floor of a spectrogram is controlled by the lower amplitude cursor in the signal's spectrum pane. All amplitudes (i.e., intensities for acoustic signals, powers for electric signals) below the amplitude floor are displayed as white. (Some amplitudes above the floor may also be white, since the floor may be below the clipping level.) The amplitude floor can be adjusted directly by moving the lower amplitude cursor in the spectrum, or indirectly using the spectrogram brightness and contrast sliders on the command panel. Units: dB for logarithmic spectrograms; watts/m^2/Hz for quadratic spectrograms of acoustic signals; watts/Hz for quadratic spectrograms of electric signals. See also Amplitude Ceiling, Dynamic Range.

Average Intensity ***(acoustic waveform, spectrogram)***

(Range) The energy flux density in the selected portion of the waveform or spectrogram (in joules/m^2) divided by the duration of the selection (in seconds). For spectra, the energy flux density in the selected portion of the spectrum divided by the duration of the spectrum source interval. For electric signals, Average Intensity is replaced by Average Power. (See also Average Intensity (dB).) Units: watts/m^2.

Average Intensity (dB) ***(acoustic waveform, spectrogram, spectrum)***

(Range) The average intensity in the selected portion of the waveform, spectrogram, or spectrum, expressed in decibels. For electric signals, Average Intensity (dB) is replaced by Average Power (dB). (See also Average Intensity.)

Average Power ***(electric waveform, spectrogram, spectrum)***

(Range) The energy in the selected portion of the waveform or spectrogram (in joules) divided by the duration of the selection (in seconds). For spectra, the energy in the selected portion of the spectrum divided by the duration of the spectrum source interval. For acoustic signals, Average Power is replaced by Average Intensity. (See also Average Power (dB).) Units: watts.

Average Power (dB) ***(electric waveform, spectrogram, spectrum)***

(Range) The average power in the selected portion of the waveform, spectrogram, or spectrum, expressed in decibels. For acoustic signals, Average Power (dB) is replaced by Average Intensity (dB). (See also Average Power.)

Begin Time ***(waveform, spectrogram, correlation)***

(Range) The time (from the beginning of the signal) at which the selection begins. For correlations, time is measured from the beginning of the first correlation file. Units: seconds.

Center Time ***(waveform, spectrogram)***

In the waveform, the amplitude-weighted central time of the selected interval. Center Time for a waveform selection is computed as

$$\frac{\sum_{i=1}^{n} t_i \cdot x_i^2}{\sum_{i=1}^{n} x_i^2} \tag{6.1}$$

where n is the number of samples in the selection, t_i is the time (in sec from the beginning of the signal) of the ith sample in the selection, and x_i is the value of the ith sample in the selection. In the spectrogram, the amplitude weightings are limited to the selected frequency band. Units: sec.

Correlation Peak *(correlation)*

(Range) The maximum value of a correlation function in the selected range.

Correlation Value *(correlation)*

(Point) The value of the correlation function at the point in time corresponding to the horizontal position of the mouse pointer.

Dynamic Range *(spectrogram, spectrum)*

The difference between the amplitude floor and ceiling of a spectrogram. Units: dB for logarithmic spectrograms; watts/m^2/Hz for quadratic spectrograms of acoustic signals; watts/Hz for quadratic spectrograms of electric signals. See also Amplitude Ceiling, Amplitude Floor.

End Time *(waveform, spectrogram, correlation)*

(Range) The time (from the beginning of the signal) at which the selection ends. For correlations, time is measured from the beginning of the first correlation file. Units: seconds.

Energy *(electric waveform, spectrogram, spectrum)*

(Range) The total energy in the selection. For a stationary signal (one with a spectrum that does not change over time), the energy in a given time interval equals the average power (in watts) times the length of the interval (in seconds). For a waveform, the energy is calculated as

$$\left(\sum_{t=t_1}^{t_2} x_t^2\right)\frac{1}{R_L f_s} \tag{6.2}$$

where t_1 and t_2 are the sequence numbers of the first and last samples in the selection, x_t is the voltage of the tth sample, R_L is the line impedance (in ohms), and f_s is the sampling rate (in Hz). For a logarithmic spectrogram, energy is calculated as

$$\left(\sum_{t=t_1}^{t_2}\sum_{f=f_1}^{f_2}\left(W_0 \cdot 10^{\left[X_{t,f}/10\right]}\right)\right)\frac{\Delta f}{R_L} \tag{6.3}$$

where f_1 and f_2 are the lower and upper frequency limits of the selection, t_1 and t_2 are the beginning and ending frame numbers of the selection, W_0 is the power dB reference value, $X_{t,f}$ is the energy in frame t at frequency f (in decibels), and Δf is the frequency bin size (which is equal to the sampling rate divided by the FFT size). For a quadratic spectrogram, the energy is calculated as

$$\left(\sum_{t=t_1}^{t_2}\sum_{f=f_1}^{f_2}\left(X_{t,f}\right)\right)\frac{\Delta f}{R_L} \tag{6.4}$$

where all symbols are defined as above, except that amplitude values are measured in joules. For a logarithmic spectrum, energy is calculated as

$$\left(\sum_{f=f_1}^{f_2}\left(W_0 \cdot 10^{\left[X_f/10\right]}\right)\right)\frac{\Delta f}{R_L} \tag{6.5}$$

where f_1 and f_2 are the lower and upper frequency limits of the selection, and X_f is the energy in the source interval at frequency f, in decibels. For a quadratic spectrum, energy is calculated as

$$\left(\sum_{f=f_1}^{f_2}\left(X_f\right)\right)\frac{\Delta f}{R_L} \tag{6.6}$$

where symbols are defined as above, except that amplitude values are in joules. Note that the total energy of a signal displayed for a spectrogram is exactly equal to the total energy of the waveform only if the spectrogram was calculated using a rectangular window and zero frame overlap. For acoustic signals, Energy is replaced by Energy Flux Density. Units: joules.

Energy per Hertz *(electric spectrogram, spectrum)*

(Point) For a spectrogram, the energy per hertz at the time and frequency indicated by the position of the mouse pointer. For a spectrum, the energy per hertz at the frequency indicated by the horizontal position of the mouse pointer. If the spectrogram or spectrum is logarithmic, the value shown (in dB) is the “energy spectrum level”; for a quadratic spectrogram or spectrum, units are joules/Hz. In a quadratic spectrogram or spectrum, Energy/Hz is equal to Power/Hz times the frame length (in seconds). For acoustic signals, Energy/Hz is replaced by Energy Flux Density/Hz.

Energy Flux Density *(acoustic waveform, spectrogram, spectrum)*

(Range) The total energy flux density in the selection. For a stationary signal (one with a spectrum that does not change over time), the energy flux density in a given time interval equals the average intensity (in watts/m^2) times the length of the interval (in seconds). For a waveform, the energy is calculated as

$$\left(\sum_{t=t_1}^{t_2} p_t^{\,2}\right)\frac{1}{(\rho c) f_s} \tag{6.7}$$

where t_1 and t_2 are the sequence numbers of the first and last samples in the selection, p_t is the pressure of the *t*th sample, ρc is the characteristic impedance of the medium (in mks rayls)[1], and f_s is the sampling rate (in Hz). For a logarithmic spectrogram, energy is calculated as

$$\left(\sum_{t=t_1}^{t_2}\sum_{f=f_1}^{f_2}\left(W_0 \cdot 10^{\left[X_{t,f}/10\right]}\right)\right)\frac{\Delta f}{(\rho c)} \tag{6.8}$$

where f_1 and f_2 are the lower and upper frequency limits of the selection, t1 and t2 are the beginning and ending frame numbers of the selection, W0 is the power dB reference value, Xt,f is the energy flux density in frame t at frequency f (in decibels), and Δf is the frequency bin size (which is equal to the sampling rate divided by the FFT size). For a quadratic spectrogram, the energy is calculated as

$$\left(\sum_{t=t_1}^{t_2}\sum_{f=f_1}^{f_2}\left(X_{t,f}\right)\right)\frac{\Delta f}{(\rho c)} \tag{6.9}$$

where all symbols are defined as above, except that energy flux density values are measured in joules/m². For a logarithmic spectrum, energy is calculated as

$$\left(\sum_{f=f_1}^{f_2}\left(W_0 \cdot 10^{\left[X_f/10\right]}\right)\right)\frac{\Delta f}{(\rho c)} \tag{6.10}$$

where f_1 and f_2 are the lower and upper frequency limits of the selection, and X_f is the energy flux density in the source interval at frequency *f*, in decibels. For a quadratic spectrum, energy is calculated as

$$\left(\sum_{f=f_1}^{f_2}\left(X_f\right)\right)\frac{\Delta f}{(\rho c)} \tag{6.11}$$

where symbols are defined as above, except that amplitude values are in joules/m². Note that the total energy flux density of a signal displayed for a spectrogram is exactly equal to the total energy of the waveform only if the spectrogram was calculated using a rectangular window and zero frame overlap. For electric signals, Energy Flux Density is replaced by Energy. Units: joules/m².

Energy Flux Density per Hertz ***(acoustic spectrogram, spectrum)***

---

[1]The characteristic impedance of an elastic medium (such as air or water) is equal to ρc, where ρ = the density of the medium (in kg/m³) and *c* = the speed of sound in the medium (in m/sec).

(Point) For a spectrogram, the energy flux density per hertz at the time and frequency indicated by the position of the mouse pointer. For a spectrum, the energy flux density per hertz at the frequency indicated by the horizontal position of the mouse pointer. If the spectrogram or spectrum is logarithmic, the value shown (in dB) is the "energy flux density spectrum level"; for a quadratic spectrogram or spectrum, units are joules/m^2/Hz. In a quadratic spectrogram or spectrum, Energy Flux Density/Hz is equal to Intensity/Hz times the frame length (in seconds). For electric signals, Energy Flux Density/Hz is replaced by Energy/Hz.

Frequency ***(spectrogram, spectrum)***

(Point) For a spectrogram, the frequency corresponding to the vertical position of the mouse pointer. For a spectrum, the frequency corresponding to the horizontal position of the mouse pointer. Units: Hz.

High Frequency ***(spectrogram, spectrum)***

(Range) The upper frequency limit of the selected region. Units: Hz.

Instantaneous Pressure ***(acoustic waveform)***

(Point) The pressure at the point in time indicated by the horizontal position of the mouse pointer. Units: pascals. (See also Instantaneous Voltage.)

Instantaneous Voltage ***(electric waveform)***

(Point) The voltage at the point in time indicated by the horizontal position of the mouse pointer. (See also Instantaneous Pressure.)

Intensity per Hertz ***(acoustic spectrogram, spectrum)***

(Point) For a spectrogram, the intensity per hertz at the time and frequency indicated by the position of the mouse pointer. For a spectrum, the intensity per hertz at the frequency indicated by the horizontal position of the mouse pointer. If the spectrogram or spectrum is logarithmic, the value shown (in dB) is the "intensity spectrum level"; for a quadratic spectrogram or spectrum, units are watts/m^2/Hz. In a quadratic spectrogram or spectrum, Intensity/Hz is equal to Energy Flux Density/Hz divided by the frame length (in seconds). For electric signals, Intensity/Hz is replaced by Power/Hz.

Low Frequency ***(spectrogram, spectrum)***

(Range) The lower frequency limit of the selected region. Units: Hz.

Number of Frames ***(spectrogram)***

(Range) The number of whole overlapping STFT frames included in the selection.

Number of Samples ***(waveform)***

(Range) The number of digitized samples in the selection.

Peak Energy per Hertz ***(electric spectrogram)***

(Range) The maximum energy/Hz in the selected region. Units: dB for logarithmic spectrograms, joules/Hz for quadratic spectrograms. For acoustic signals, Energy per Hertz is replaced by Energy Flux Density per Hertz .

Peak Energy Flux Density per Hertz ***(acoustic spectrogram)***

(Range) The maximum energy flux density/Hz in the selected region. Units: dB for logarithmic spectrograms, joules/m^2/Hz for quadratic spectrograms. For electric signals, Energy Flux Density per Hertz is replaced by Energy per Hertz.

Peak Frequency ***(spectrogram, spectrum)***

(Range) The frequency at which the highest amplitude in a selection occurs. Units: Hz

Peak Intensity per Hertz ***(acoustic spectrogram, spectrum)***

(Range) The maximum intensity/Hz in the selected region. Units: dB for logarithmic spectrograms, watts/m^2/Hz for quadratic spectrograms. For electric signals, Peak Intensity/Hz is replaced by Peak Power/Hz.

Peak Location ***(correlation)***

(Range) The time of the maximum value of a correlation function in the selected range. Time is measured from the beginning of the first correlation file. Units: seconds.

Peak Power per Hertz ***(electric spectrogram, spectrum)***

(Range) The maximum power/Hz in the selected region. Units: dB for logarithmic spectrograms, watts/Hz for quadratic spectrograms. For electric acoustic, Peak Power/Hz is replaced by Peak Intensity/Hz.

Peak Pressure ***(acoustic waveform)***

(Range) The pressure in the selected region that has the maximum absolute value. Units: pascals. For electric signals, Peak Pressure is replaced by Peak Voltage.

Peak Time ***(waveform, spectrogram)***

(Range) The time (from the beginning of the signal) at which the highest amplitude in the selection occurs. Units: seconds.

Peak Voltage ***(electric waveform)***

(Range) The voltage in the selected region that has the maximum absolute value. Units: volts. For acoustic signals, Peak Voltage is replaced by Peak Pressure.

Power per Hertz ***(electric spectrogram, spectrum)***

(Point) For a spectrogram, the power per hertz at the time and frequency indicated by the position of the mouse pointer. For a spectrum, the power per hertz at the frequency indicated by the horizontal position of the mouse pointer. If the spectrogram or spectrum is logarithmic, the value shown (in dB) is the "power spectrum level"; for a quadratic spectrogram or spectrum, units are watts/Hz. In a quadratic spectrogram or spectrum, Power/Hz is equal to Energy/Hz divided by the frame length (in seconds). For acoustic signals, Power/Hz is replaced by Intensity/Hz.

RMS Pressure ***(acoustic waveform)***

(Range) The root-mean-square pressure (sometimes called "effective pressure") of the selected part of the signal. RMS pressure is equal to

$$\sqrt{\frac{1}{n}\sum_{i=1}^{n} p_i^{\,2}} \qquad (6.12)$$

where n is the number of samples in the selection, and p_i is the pressure (in pascals) of the ith sample in the selection. Units: dB relative to the dB pressure reference specified in the signal's calibration. For electric signals, RMS Pressure is replaced by RMS Voltage.

RMS Voltage ***(electric waveform)***

(Range) The root-mean-square voltage (sometimes called "effective voltage") of the selected part of the signal. RMS voltage is equal to

$$\sqrt{\frac{1}{n}\sum_{i=1}^{n} v_i^{\,2}} \qquad (6.13)$$

where n is the number of samples in the selection, and v_i is the voltage of the ith sample in the selection. Units: dB relative to the dB voltage reference specified in the signal's calibration. For acoustic signals, RMS Voltage is replaced by RMS Pressure.

Time ***(waveform, spectrogram, correlation)***

(Point) The time point in the signal or correlation function at the horizontal position of the mouse. For correlations, time is measured from the beginning of the first correlation file. Units: seconds.

ΔIntensity / Hz ***(acoustic spectrum)***

(Range) The intensity/Hz at the left edge of the selection minus the intensity/Hz at the right edge of the selection. Units: dB for logarithmic spectra; joules/m^2/Hz for quadratic spectra. For electric signals, ΔIntensity/Hz is replaced by ΔPower/Hz.

ΔFrequency *(spectrogram, spectrum)*

(Range) The difference between the upper and lower frequency limits of the selected region. Units: Hz.

ΔPower / Hz *(electric spectrum)*

(Range) The power/Hz at the left edge of the selection minus the power/Hz at the right edge of the selection. Units: dB for logarithmic spectra; joules/Hz for quadratic spectra. For acoustic signals, ΔPower/Hz is replaced by ΔIntensity/Hz.

ΔTime *(waveform, spectrogram, correlation)*

(Range) The duration of the selection. Units: seconds.

The data logger

The data logger enables you to record measurements displayed in the measurement panel and save them for export to other programs. Selecting Data Log from the Windows menu opens a window named DataLog, which displays the logged measurements, and adds a DataLog menu to the menu bar.

Logging measurements

Clicking the mouse in any pane of a signal window while holding down the command key copies all of the measurement values from the measurement panel and records them as a single entry in the DataLog window. Parameters are not copied to the data log when an entry is made; you can record parameters using the Signal Parameters... command, which is discussed below in the section on "Comments." Measurements can be entered whether the measurement panel is showing or hidden.

The DataLog window

The DataLog window is divided into two vertical panes. The left pane contains entry sequence numbers that are recorded automatically each time you make an entry by command-clicking. The right pane contains data, comments, and column headers.

Each entry appears as a series of measurements in successive columns on a single row in the data log. Measurements are recorded from left to right in the order that they appear in the measurement panel.

Column headers appear across the top of the pane to identify the measurements in the currently highlighted entry. An entry is automatically highlighted when it is first entered, and remains highlighted until you record a new entry or click on another entry already in the log.

New entries are always appended to the end of the log, irrespective of which entry was highlighted when the new entry is made.

You can close the DataLog window by clicking in its close box (at the left end of the title bar). If you have recorded any data since the last time you saved the log, Canary asks if you want to save the log before closing. If you close a DataLog window without saving, its contents will be lost, since choosing Data Log from the Windows menu always brings up a new (empty) data log window. Saving the data log and opening saved logs are discussed later in this chapter.

Mixed measurements

The entries in a data log may contain different sets of measurements ("mixed measurements"). For example, a data log can include entries from different panes. Figure 6.3 shows a data log containing measurements made in both the waveform and spectrogram panes (see the discussion below on how to add comments to a data log, like those shown in the figure). Mixed measurements can also occur even if all of the entries are from one pane, if the pane's measurement panel is reconfigured between entries.

Remember that if a data log contains mixed measurements, the column headers that are displayed may not apply to all of the entries (since the displayed headers refer only to the highlighted entry).

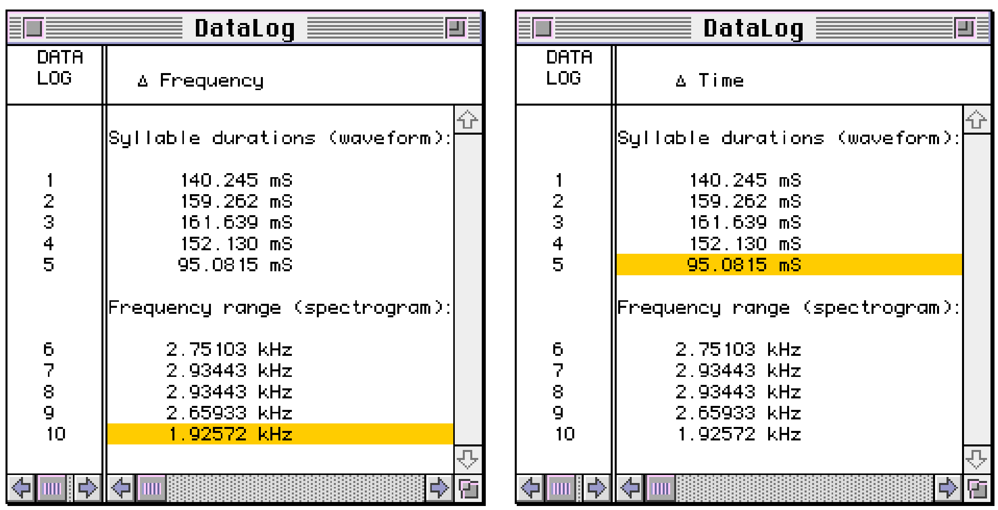

Figure 6.3. A data log containing mixed measurements. Both windows contain the same data, but a different entry is highlighted in each. The column header shows which measurements are recorded in the highlighted entry.

Comments

Selecting Add Comment... from the DataLog menu brings up a dialog box that lets you enter a text comment to be appended to the DataLog window (Figure 6.4). When you click OK, whatever text you entered in the dialog box is appended to the log, following the last measurement entry already recorded. Subsequent entries are added after the comment. In Figure 6.3, for example, comments are used to identify the two sets of data from the waveform and spectrogram panes.

Figure 6.4. The Add Comment dialog box.

Selecting Signal Parameters... from the DataLog menu brings up the Add Comment dialog box and automatically enters the parameter values displayed on the measurement panel into the text box (Figure 6.5). Clicking OK appends the parameters to the log as a comment.

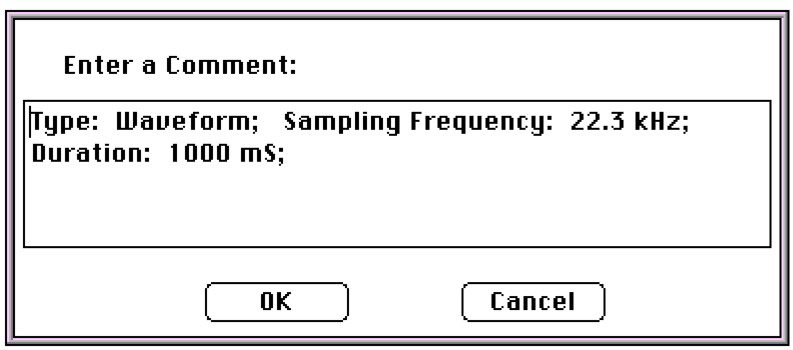

Figure 6.5. The Signal Parameters... command from the DataLog menu automatically fills in the text entry box of the Add Comment dialog box with the parameters displayed by the measurement panel.

Deleting entries from the datalog

To delete an entry from the data log, select the entry in the DataLog window, and choose Clear from the Edit menu. Deletions from the data log cannot be undone.

Saving the data log

Canary can save a data log in Text or MATLAB file formats, or as a text report file. A data log saved in Text or MATLAB format can be re-opened with Canary later, but is generally not useful for exporting data to other programs. A data log saved as a text report can be read by most word processing, spreadsheet, and statistics programs, but cannot be reopened by Canary.

To save a data log in Text or MATLAB format, select Save Log... from the DataLog or File menu, then choose the desired format in the standard Save File dialog box.

To save a data log as a text report to be imported into another program, select Save Text Report... from the DataLog menu (the Save Text Report... command also appears on the File menu when the DataLog window is the active window). The dialog box shown in Figure 6.6 then allows you to specify the layout of the file to be saved.

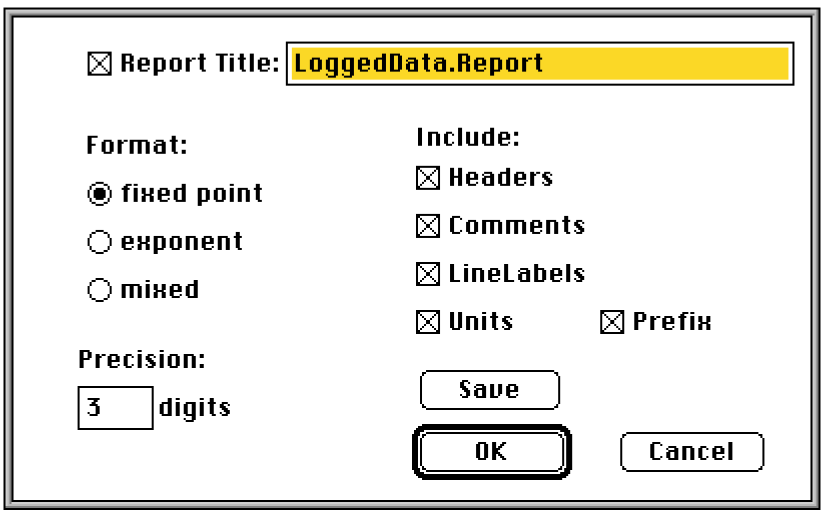

Figure 6.6. The Save Text Report dialog box allows you to specify the format of a text report file.

If Report Title is checked, the specified text is included on one line at the top of the file that is written. You can enter whatever text you want for the report title. Note that the report title is not the name of the file to be saved (although it may be the same); the file name is specified in a standard Save File dialog box that appears after you click OK in the Save Text Report dialog box.

The three radio buttons under Format: and the Precision field allow you to specify the format used for the numeric entries in the data log.

The five checkboxes under Include: allow you to specify how rows, columns, and individual numeric entries in the array should be labelled:

Headers: If this box is checked, the file will include text labels at the top of each column identifying the measurements that appear in the column. If the log contains mixed measurements, a new header line is placed before each row that contains different measurements from the row above it (see example below).

Comments: Comments placed in the log using the Add Comment... or Signal Parameters... commands will be included in the text report only if this box is checked.

LineLabels: If this box is checked, the first column of each data line in the file will contain the sequence number for the entry, as shown in the first column of the DataLog window.

Units: If this box is checked, each individual measurement will be labelled with appropriate units (e.g., S, V, dB, etc.). If the Prefix box is also checked, the unit label may be preceded by a prefix (e.g., µS, mV; see below). If the Units box is unchecked, the column headers (if present) will include the unscaled units (e.g., S or V, but never mS or mV), and individual measurements will appear as unlabelled numbers. If units are used with no prefixes, small values may be written as zeros (see following paragraph).

Prefix: Appends the appropriate standard prefix before the unit label on each entry in the array. This checkbox is disabled if the Units checkbox is not checked. (See Appendix E for standard metric prefixes and their abbreviations.) The prefix is chosen so that the integer part of a fixed-point measurement is between 0 and 999. Thus, for example, a value of .000999 sec would be written as 999 µS, whereas .001000 sec would be written as 1 mS. Note that if you omit prefixes and specify few digits of precision, small values may appear as zeros. For example, a value of 3.67 mS would appear as 0.003670 in 6-digit fixed point output, but would be written as 0.00 seconds if 2-digit fixed-point output were requested.

In a text report file, each row of the data log is written as a single line of text, terminated by a carriage return character (called a “hard return” or “paragraph mark” in some word processors). Entries in successive columns within a line are separated by tab characters. Most spreadsheet and data analysis application programs can import such a “tab-delimited text” file. If you import a text report file into a word processor, you may initially find that the data (and the columns headers, if present) don’t line up in columns as expected. You can align the data properly by adjusting the tab stop positions and line length (margin) settings. See your word processor’s documentation for further details.

Figure 6.7 shows a portion of a text report file generated from the data log shown in Figure 6.3.

```
Syllable durations (waveform):

       Δ Time (S)
1      0.140
2      0.159
3      0.162
4      0.152
5      0.095

Frequency range (spectrogram):

       Δ Frequency (Hz)
6      2751.030
7      2934.432
8      2934.432
9      2659.329
10     1925.721
```

Figure 6.7. Text from a text report file saved from the data log shown in Figure 6.3. The report file was generated with 3-digit fixed-point numeric format, and included headers, comments, line labels, and no units.

Opening a saved data log

You can open a data log saved in Text or MATLAB format by selecting Open Log... from the DataLog menu. (If the DataLog menu is not on the menu bar, first choose Data Log from the Windows menu). The window in which a data log opens is always titled "DataLog", irrespective of the name of the log file that was opened. Once a data log has been opened, you can add comments or measurements to it, and resave it.

Remember that data logs saved as text reports cannot be opened by Canary.

Chapter 7 Correlation

About this chapter

This chapter explains the operation of Canary's correlator, which performs quantitative comparisons between spectrograms or waveforms. The first section explains how correlation functions are calculated and shows examples of both types of correlation. The second section explains the mechanics of selecting inputs to the correlator and specifying options common to both spectrogram and waveform correlations. The final sections discuss issues and options specific to each type of correlation, including a discussion of the limitations of correlations as a measure of similarity between spectrograms. For an explanation of how to perform multiple correlations at a time (batch correlations), see Chapter 11.

It is difficult to generalize about how some options and parameters will affect any particular correlation function. For example, a change to the clipping level of the spectrograms being correlated may either increase or decrease the correlation values observed or leave them unchanged, depending on the particular signals being correlated. Since the behavior of the correlator depends very much on the data to which it is applied, the best way to learn about the correlator is to read this chapter and then experiment with correlations using the type of data with which you normally work.

Background

The correlator takes either two spectrograms or two waveforms as input and incrementally slides them past each other in time, calculating a correlation coefficient between the inputs at successive discrete time offsets. The output of the correlator is a plot of correlation value versus the time offset between the signals. For example, Figure 7.1 shows spectrograms and a correlation plot comparing two renditions of a particular syllable type by a northern mockingbird. Figure 7.2 shows waveforms representing the same signal recorded at two different microphones and a correlation plot comparing them.

The time axis of a correlation plot displays time relative to whichever signal is specified as the first input to the correlator. Thus, a peak at a positive time offset indicates that the first signal is ahead of the second one.

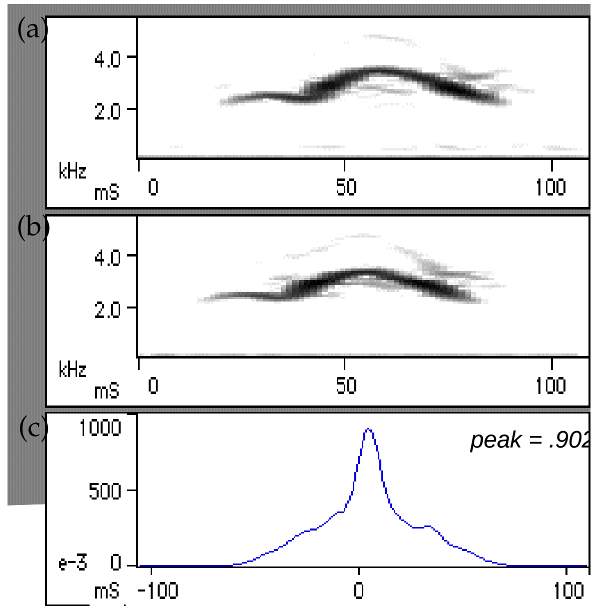

Figure 7.1. Spectrogram correlation between two renditions of a syllable type in the song of a northern mockingbird. Renditions **(a)** and **(b)** were produced during a single song bout but separated by several minutes in time. **(c)** Correlation between spectrograms (a) and (b).

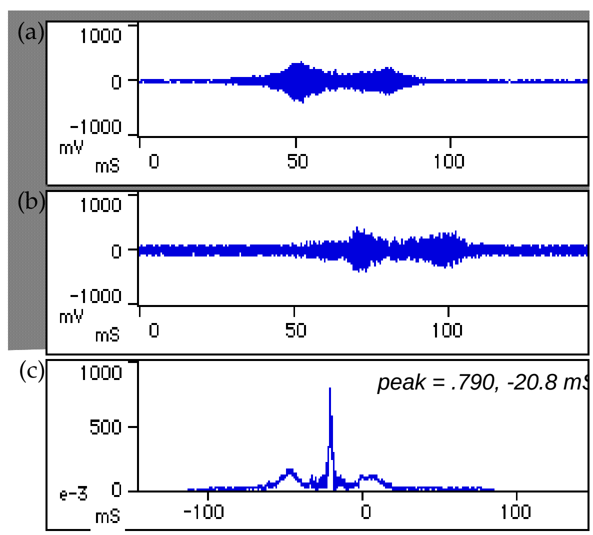

Figure 7.2. Waveform correlation between a single syllable from a mockingbird song simultaneously recorded at two different microphones. **(a)** Waveform from microphone 1. **(b)** Waveform from microphone 2, showing more noise than the signal from microphone 1. **(c)** Complex envelope of correlation between (a) and (b). The time delay of -20.8 mS indicates that the bird was 7.2 m closer to microphone 1 than to microphone 2 (assuming the speed of sound is 344 m/S).

Often the only point of interest in a correlation plot is the maximum value of the correlation function. Depending on the application, one might be interested just in the peak value of the correlation, or just in the time at which it occurs, as discussed separately for spectrograms and waveforms below. For example, if a correlation is being used as a measure of similarity between two spectrograms, the important value is the peak correlation; in most cases, the time offset at which the peak occurs is of little interest. If one is cross-correlating waveforms recorded synchronously from one sound source by two different microphones to determine delays between the arrival times of the sound at the two microphones (e.g., for use in computing the position of the sound source relative to the microphones), the exact value of the peak correlation is of less interest than the time offset at which it occurs.

Spectrogram correlations

For each time offset Δt, Canary calculates either a normalized or non-normalized correlation value $C_{\Delta t}$ between two spectrograms. If the Normalize option is selected, each correlation value is calculated using the formula

$$C_{\Delta t} = \frac{\sum_{t=1}^{n} \sum_{f=1}^{FFT} \left(X_{t,f} \cdot Y_{t+\Delta t, f} \right)}{\sqrt{\left(\sum_{t=1}^{n} \sum_{f=1}^{FFT} \left(X_{t,f} \right)^2 \right) \left(\sum_{t=1}^{n} \sum_{f=1}^{FFT} \left(Y_{t+\Delta t, f} \right)^2 \right)}} \tag{7.1}$$

where n equals $(N_1 + N_2) - 1$, and N_1 and N_2 are the numbers of frames in the two spectrograms. *FFT* equals the number of frequency bins, which must be the same for the two spectrograms being correlated. $X_{t,f}$ and $Y_{t+\Delta t,f}$ are the amplitude values in the two spectrograms at frequency f and times t and $t+\Delta t$, respectively. The normalized correlation value for spectrograms can vary between 0 and 1. A correlation of 0 means that the non-zero values in the two spectrograms do not coincide at all; a correlation of 1 indicates that the two signals are identical (given the time offset Δt).

Successive correlation values are calculated by incrementing the value of Δt in steps equal to the time grid resolution (which must be the same for both spectrograms), in effect sliding the two spectrograms past each other in time. If normalization is turned off, only the numerator of Expression 7.1 is used (see the discussion of “Normalization” in the section “Using the correlator” later in this chapter).

The peak correlation value can provide a quantitative measure of one type of similarity between spectrograms. However, **correlations have serious limitations as a way of comparing spectrograms; the type of similarity that they measure may or may not be appropriate or useful in a given context**. The interpretation and limitations of spectrogram correlations are discussed in the section “Spectrogram correlations” later in this chapter.

Waveform correlations

For each time offset, Canary calculates either a normalized or non-normalized correlation value between two waveforms. If the Normalize option is selected, each correlation value is calculated using the formula

$$C_{\Delta t} = \frac{\sum_{t=1}^{n}\left(x_t \cdot y_{t+\Delta t}\right)}{\sqrt{\left(\sum_{t=1}^{n} x_t^{2}\right)\left(\sum_{t=1}^{n} y_{t+\Delta t}^{2}\right)}} \qquad (7.2)$$

where n equals $(N_1 + N_2) - 1$, and N_1 and N_2 are the numbers of digitized samples in the two waveforms. x_t and $y_{t+\Delta t}$ are the values of sample numbers t and $t+\Delta t$ of the two waveforms, respectively. If the two signals differ in length, the shorter signal is zero-padded to the length of the longer signal. The correlation value for waveforms can vary between -1 and 1. (If you choose to plot the complex envelope of a waveform correlation, the values that are plotted vary between 0 and 1, as discussed later in this chapter.) A correlation of 0 means that the signals are orthogonal[1]; a correlation of 1 indicates that the two signals are identical; a correlation of -1 indicates that the signals are identical in magnitude, but opposite in phase.

Successive correlation values are calculated by incrementing the value of Δt in steps equal to the inverse of the sampling frequency, in effect sliding the two waveforms past each other in time. If normalization is turned off, only the numerator of Expression 7.2 is used, as discussed under “Correlator options” in the section “Using the correlator” in this chapter.

Waveform correlations are useful for determining the time offset at which two signals most closely match each other. For example, to determine the position of a sound source in space, an array of microphones in a known geometry can be used to record onto separate but synchronized recording tracks. The time offsets of the correlation peaks between signals on these synchronized recordings then indicate the delays between the arrival times of the sound at the different microphones. These time delays can then be used to calculate the location of the sound source relative to the positions of the microphones.

Using the correlator

This section discusses general aspects of using the correlator that are common to both waveform and spectrogram correlations. Later sections discuss issues particular to each type of correlation.

[1]Whether or not two signals are orthogonal depends on their frequency content and on their relative phase. For example, sinusoidal signals of very different frequencies are orthogonal, as are signals of the same frequency that are 90° out of phase with each other.

Selecting which data to correlate

Choosing Correlator from the Process menu brings up the correlator's first file selection dialog box (Figure 7.3). If the first signal that you want to correlate is already open in Canary (i.e., if a window exists for the signal), click on the Existing... button; otherwise, select a file from the scrolling list as you would in a standard open file dialog box and click on the Open button (or double-click on the file name).

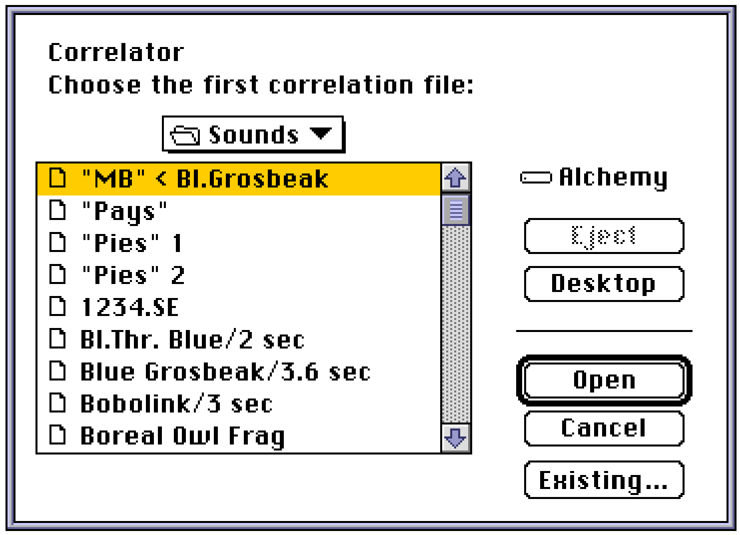

Figure 7.3. The Correlator Input dialog box.

If you click on the Existing... button, another dialog box appears, which lists the names of the waveforms or spectrograms that are presently open in Canary (Figure 7.4). A pair of radio buttons at the bottom of the dialog lets you choose whether to list waveforms or spectrograms. You select a waveform or spectrogram from the list either by clicking on its name and then clicking OK, or by double-clicking on the name.

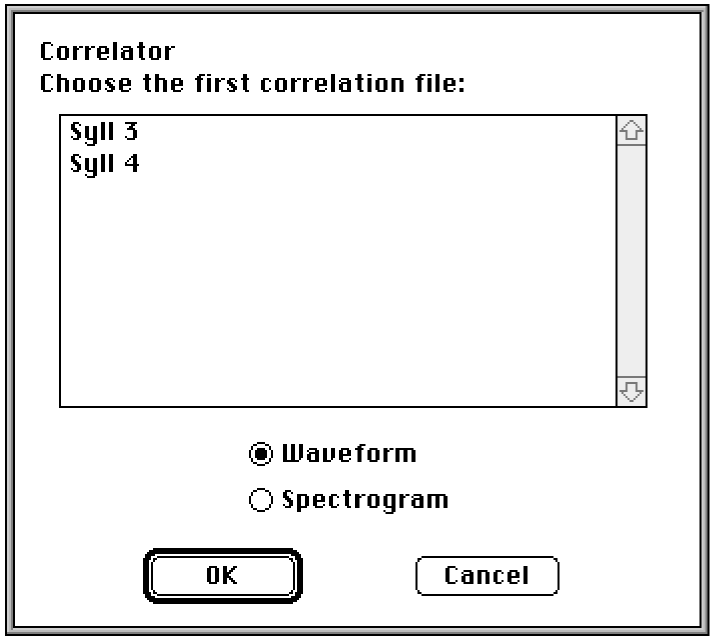

Figure 7.4. The correlator’s Existing Window file selection dialog box.

If you select a saved Canary file that contains both a waveform and a spectrogram, a dialog box appears that lets you specify which view you want to use as input to the correlator (Figure 7.5). If the selected file contains only a waveform or only a spectrogram, this dialog box does not appear.

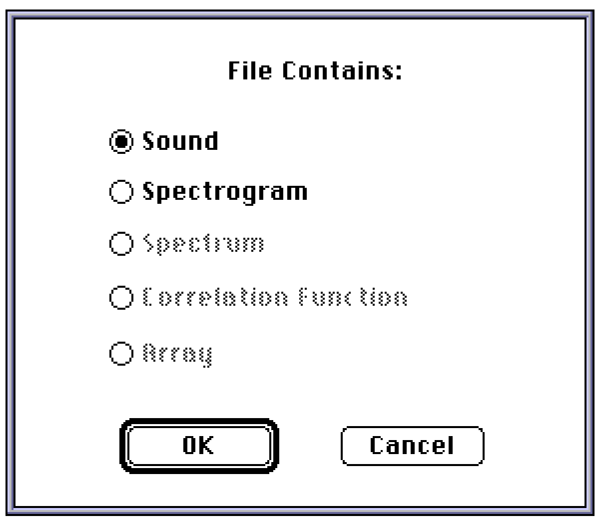

Figure 7.5. The correlator's view selection dialog box. This dialog box appears only when you select a file that contains both a waveform (sound) and a spectrogram. The choice of Sound or Spectrogram tells the correlator which view to use as input.

Once you have selected the first file (or a signal in an open window) to correlate, the file selection dialog box appears again, prompting this time for the second file to be correlated. Again, you can choose either a file saved on disk or a signal that is presently open.

Waveforms to be correlated must have the same sampling rate. Spectrograms to be correlated must match each other in all of the parameters specified in the Spectrogram Options dialog box, or an error will occur.

Correlator options

After you have selected both files to be correlated, the Correlator Options dialog box appears (Figure 7.6). If you selected waveforms or files that contain waveforms, the options dialog box lets you select whether to correlate the waveforms themselves or spectrograms of those waveforms (which the correlator will make from the selected waveforms). If you selected spectrograms or files that contain spectrograms only, the Waveform choice is dimmed.

Clicking the OK button initiates calculation of the correlation function, and displays a status window titled Correlator, which shows the progress of the calculation. Clicking the Pause button in the status window makes the Cancel button available. Clicking the Cancel button aborts the calculation.

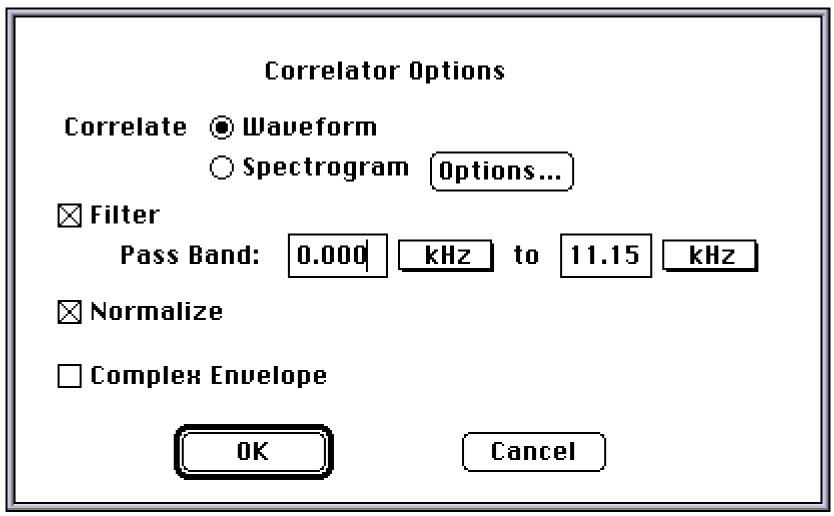

Figure 7.6. The Correlator Options dialog box.

Spectrogram options

The spectrogram Options... button brings up the Spectrogram Options dialog box that lets you specify the parameters for calculating the spectrogram (analysis and grid resolution, window function, amplitude scale, and noise floor), as discussed in Chapter 3. If you selected waveforms in the file or view selection dialog boxes, the options that you specify in this dialog box are used to make new spectrograms that are then used as input to the correlator. If you selected spectrograms in the file or view selection dialogs, the spectrogram Options... button displays the parameters of the first spectrogram that was selected, but does not allow you to change any parameters. In order to correlate two spectrograms, all of their parameters must be identical, or an error occurs.

Filtering

If you click on the Filter checkbox, fields appear in the dialog box that let you specify the lower and upper frequency limits of a bandpass filter that is applied to the data before the correlation is performed. You can use the pop-up units menus to specify frequencies either in kHz or Hz.

In most applications, filtering is advisable both for waveform and spectrogram correlations. By selecting a frequency band that just includes the actual signal(s) in which you're interested, you reduce the effect of any other noise in the file on the correlation values. You can determine the appropriate filter band to use by adjusting the frequency cursors in a spectrogram to include only the lower and upper frequencies of your signal.

The effect of filtering on the correlation function (for both waveforms and spectrograms) depends very much on the particular signals being correlated. If neither file contains much energy outside the frequency band occupied by the signals, the filtered and unfiltered correlations are usually extremely close to each other; either one may be slightly higher than the other, depending on the particular files being correlated. If either input file does contain much energy outside of the frequency band of interest, the difference between the filtered and unfiltered correlations can be much larger. Again, filtering may either raise or lower correlation values, depending on the particular signals being correlated.

Normalization

If Normalize is checked, the sum of the products of the data values from the two signals is divided by the square root of the product of the sums of values from the two signals, as indicated in Expressions 7.1 and 7.2. The units in the numerator and denominator cancel and the correlation value is scaled to a dimensionless value. For spectrograms, which contain only non-negative amplitude values, the normalized correlation value is always between 0 and 1. For waveforms, which can contain positive, negative, and zero values, the normalized correlation varies between -1 and 1 (but see the discussion of the Complex Envelope option in the section "Waveform correlations" later in this chapter).

If Normalize is off, the correlation is calculated as the sum of the products of the data values from the two signals (i.e., just the numerator of Expressions 7.1 and 7.2). A non-normalized correlation is given in arbitrary units.

In a normalized correlation, the overall amplitude level of each signal does not affect the correlation value: if two signals are identical except that the amplitude of one is half that of the other, their peak correlation will still be 1 (Figure 7.7c). In a non-normalized correlation, the overall amplitude level does affect the value of the correlation (Figure 7.7d, e). If your application requires that the similarity measure incorporate information about the absolute amplitude levels of the signals being compared, you should turn Normalize off; otherwise it should probably be left on.

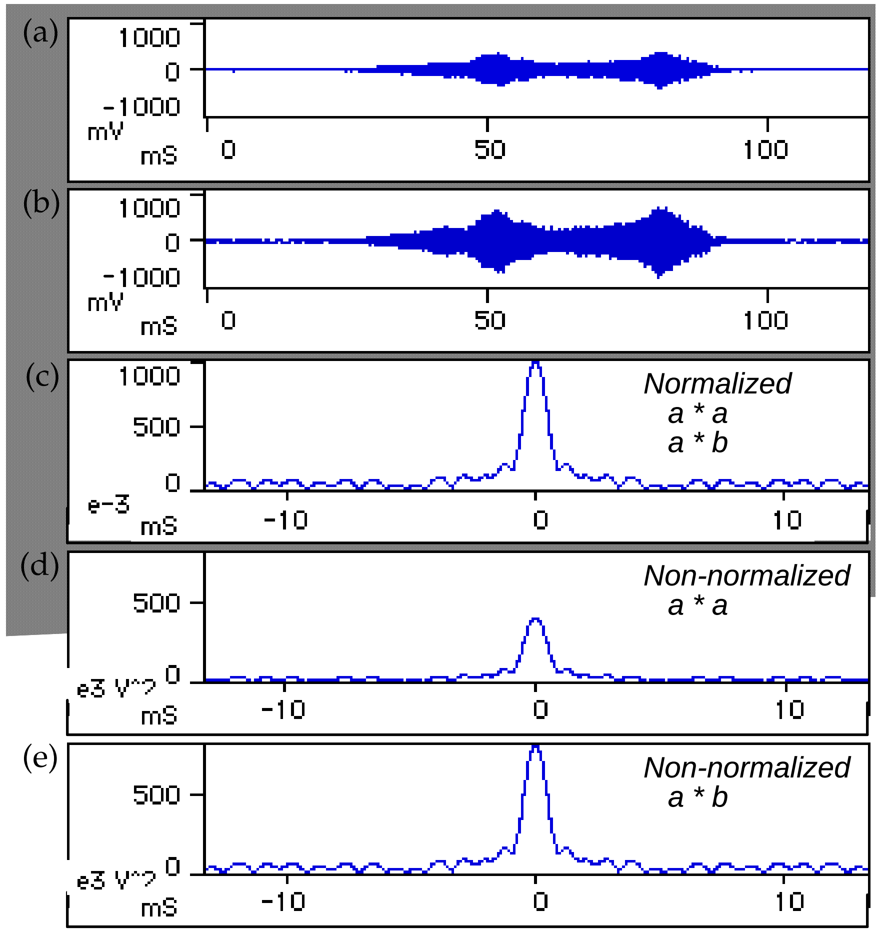

Figure 7.7. Comparison of normalized and non-normalized correlations. The complex envelope of the waveform correlations is shown. **(a)** Waveform of a syllable from a mockingbird song. **(b)** Same waveform amplified by a factor of 2. **(c)** Normalized correlation; the correlation between (a) and (b) is identical to the correlation between (a) and itself. **(d)** Non-normalized correlation between (a) and itself. **(e)** Non-normalized correlation between (a) and (b).

Complex envelope

The complex envelope option applies only to waveform correlations, and is discussed in the section “Waveform correlations” later in this chapter.

Spectrogram correlations

Interpreting spectrogram correlations

The peak value of a spectrogram correlation function can provide an objective, well-defined, repeatable, and comparable measure of the similarity of two spectrograms. (The correlation measurement panel can be configured to display the peak value in a selected part of a correlation function; see Chapter 6.) However, **care should be exercised in the interpretation of spectrogram correlations because the type of "similarity" that they measure is simple and narrowly defined, and may or may not be appropriate to the research question being asked.**

Spectrogram correlation is not a tool for generalized "pattern recognition". The ranking of similarities among a group of spectrograms based on peak correlations may conflict with a human's subjective ranking based either on listening to the sound or on examining the spectrograms. For example, spectrogram correlations are insensitive to similarities in "shape" of two spectrograms if they differ in frequency, but such similarities are striking to a human observer. Figure 7.8 shows correlations between two pairs of syllables from within the song of a painted bunting. To a human, syllables 2 and 6 look and sound more "similar" than syllables 2 and 5 because they have a more similar pattern of frequency modulation. But because the frequency ranges of syllables 2 and 6 barely overlap, their peak correlation is much lower than that between syllables 2 and 5, which overlap more extensively in frequency (even though they differ more in pattern of frequency modulation).

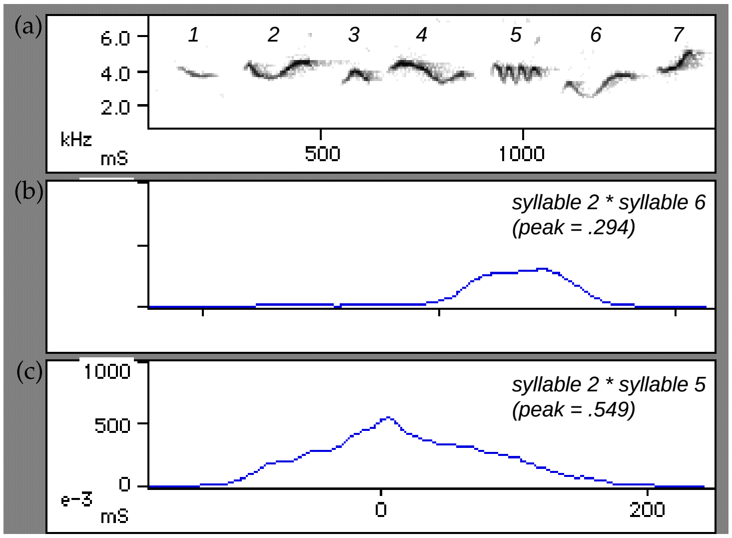

Figure 7.8. Two sounds that are similar in “shape” (frequency contour) may have a low correlation if they are offset from each other in frequency**. (a)** Spectrogram of part of a song of a painted bunting, with syllables numbered for reference. Syllables 2 and 6 have similar patterns of frequency modulation, but overlap only slightly in frequency range. Syllables 2 and 5 are less similar in their pattern of frequency modulation, but overlap more extensively in frequency range. **(b)** Spectrogram correlation of syllables 2 and 6. **(c)** Spectrogram correlation of syllables 2 and 5.

The usefulness of spectrogram correlations as a measure of similarity thus depends very much on the specific context in which they are being used. The best way to develop a feel for how to interpret spectrogram correlations is to experiment with correlating a variety of spectrograms.

Logarithmic vs. quadratic amplitude axis

Spectrogram correlations can be performed using spectrograms with either logarithmic or quadratic amplitude axes (henceforth referred to as logarithmic or quadratic spectrograms), as long as both spectrograms are of the same type. Although logarithmic spectrograms are often more visually informative because of their wide dynamic range, quadratic spectrograms may sometimes be more appropriate for purposes of correlation. If quadratic spectrograms are used, the values being correlated at each point in time are the spectrum intensities (in watts/m^2/Hz) at each frequency; if logarithmic spectrograms are used, the values being correlated are the spectrum intensity levels (in dB).[1] For spectrograms whose peak normalized correlations are close to 1, the two types of spectrograms yield very similar results. (For identical signals, of course, either type of spectrogram yields a correlation of 1.) For any given pair of signals, the correlation values from quadratic spectrograms are often lower than those from logarithmic spectrograms (Figure 7.9).

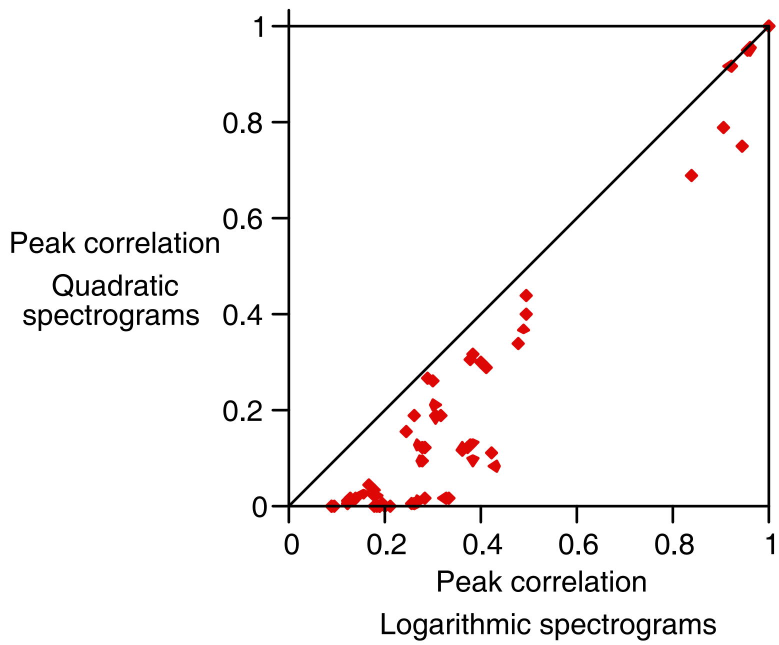

Figure 7.9. Peak normalized correlation coefficients for spectrograms of 78 pairs of signals. One point is plotted for each pair of signals, showing the peak values when quadratic versus logarithmic spectrograms were correlated. The diagonal is the line of equal peak correlations.

In a given set of signal pairs being correlated, the rank orderings of peak correlation coefficients obtained using quadratic and logarithmic spectrograms are not necessarily exactly the same, although they are generally

[1]If the electric calibration paradigm is selected, the values that are correlated are spectrum powers (in watts/Hz) or spectrum power levels (in dB).

similar. For example, in one set of 12 signals that were correlated in all 78 unique pairs (including the 12 autocorrelations), the 22 correlations whose values exceeded 0.5 when using log spectrograms were ranked in the same order when using quadratic spectrograms from the same signals. Below this correlation level, the exact rankings differed somewhat. Thus whether a given pair of spectrograms has a higher correlation than another pair can depend in part on whether log or quadratic spectrograms are used. You should therefore consider carefully whether logarithmic or quadratic spectrograms are more appropriate for correlations given your application, and experiment with both types if there is any question about which to use.

Clipping level

Spectrogram correlations are calculated on the basis of all of the information in a spectrogram, including noise. Sources of noise in a spectrogram include acoustic noise in the original recording before it was digitized (e.g., from wind or other environmental sources), and noise introduced as an artifact of the limited precision and dynamic range of the digitizing process, as discussed in Appendix A and Chapter 3.

You can reduce the amount of noise in a logarithmic spectrogram by setting an appropriate clipping level when the spectrogram is made. To determine an appropriate clipping level, first make a spectrogram with a low clipping level, and then use the amplitude cursors in the spectrum pane to adjust the noise floor of the spectrogram display upward, as explained in Chapters 1 and 3.[1] Note the noise floor setting at which your signal shows up clearly with minimal noise visible on the spectrogram. Finally, recalculate the spectrogram using this noise floor as the Clipping Level parameter.

For normalized correlations, any change to the clipping level of the input spectrograms (in either direction) can either raise or lower the peak correlation value, depending on the signals being correlated. For non-normalized correlations, raising the clipping level always reduces the values of the correlation function.

Grid resolution

The size of the time increment between successive points in a spectrogram correlation function equals the time grid resolution of the input spectrograms.

The calculation time required for a spectrogram correlation is directly proportional to the number of points in the spectrograms, and thus depends on the time and frequency grid resolution.

[1]If there is initially no noise visible as gray regions in the spectrogram that are obviously not part of the signal, try a lower clipping level.

Waveform correlations

Interpreting waveform correlations

Waveform correlations are most useful for determining timing relationships between inputs that are (or might be) versions of the same original signal. Examples include the same sound recorded at two different locations, or a sonar pulse and its echo. In such contexts, the time of the peak correlation is generally of more interest than the correlation value itself. (The correlation measurement panel can be configured to display the time of the peak value in a selected part of a correlation function; see Chapter 6.)

In most applications, waveform correlations are less useful than spectrogram correlations for assessing the degree of "similarity" between signals in a way that is intuitively satisfying, in part because (unlike spectrogram correlations) they are sensitive to phase differences that our auditory system does not detect.

Complex envelope

If Complex Envelope is checked in the correlator options dialog box, Canary displays the complex envelope of the correlation function, rather than the correlation function itself. The complex envelope varies between 0 and 1. The relationship between the complex envelope and the correlation function itself is illustrated in Figure 7.10. Figure 7.11 compares a correlation plot between two signals to the correlation's complex envelope.

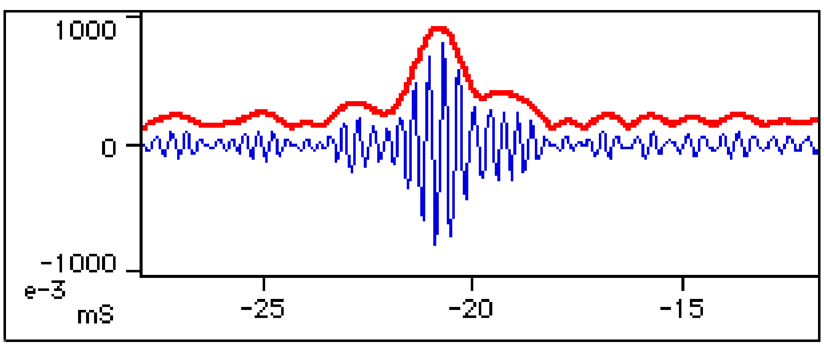

Figure 7.10. The complex envelope of part of a waveform correlation function. For clarity of illustration, the complex envelope (heavy line) is shifted slightly upwards in the figure from the raw correlation plot (thin line).

Correlation functions always contain high-frequency oscillations, which are related to the frequencies present in the signals being correlated. If the signals are approximately sinusoidal (i.e., at any moment, most of the energy in each signal is concentrated at a single frequency, as in the frequency-modulated whistles common in bird song), their correlation function will itself be close to an amplitude-modulated sinusoid (Figure 7.10). In this case, the complex envelope is roughly equivalent to the amplitude envelope of the absolute value of the correlation function. If the signals being correlated are spectrally complex (with energy distributed over many frequencies, as in human speech), their correlation function contains high frequency oscillations that are generally not sinusoidal.

Taking the complex envelope removes much of the high-frequency oscillation in a correlation function, which can make it easier to visually identify the peak of a waveform correlation. Since the qualitative relationship between the appearance of the complex envelope and the raw correlation plot depends somewhat on the signals being correlated, you should experiment with the type of signals that you work with in order to get a feel for the relationship between the two types of plot.

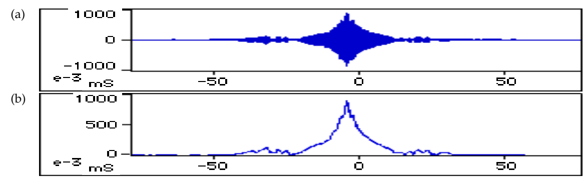

Figure 7.11. (a) Waveform correlation function between two successive calls in a sequence made by a northern flicker; filter = 0.7 to 6.8 kHz; peak value = 0.890, at -4.3 mS. **(b)** Complex envelope of the correlation function shown in (a); peak value = 0.910, at -4.3 mS.

Chapter 8 Preferences and Options

About this chapter

The first section of this chapter discusses how to customize Canary's default settings using preference files.

The remainder of this chapter discusses the following commands on Canary's Options menu:

Manual scaling...	Control horizontal and vertical scaling of signal and correlation windows.
Display...	Control spectrogram highlighting and other display parameters.
Speed...	Control calculation speed and program responsiveness.
General...	Control program startup and exit behavior.

Signal calibration, multi-track options, and measurement panel configuration, which are also accessed via the Options menu, are discussed in Chapters 4, 5, and 6, respectively.

Preferences

Canary allows you to save various options and settings in *preference files*. You can have as many different preference files as you want, and you can load a new preference file any time while you're running Canary. Preference files enable you to recall configurations of options quickly and easily. For example, if you work on two projects that use different combinations of measurements on the measurement panel, you can save the measurement panel configurations in two separate preference files. Then you can instantly reconfigure the measurement panel just by loading the appropriate preference file.

Information saved in preference files

Many of Canary's dialog boxes contain a Save or Save default button that allows you to record the settings of all the dialog's controls in the current preference file. Once the settings are recorded, the dialog box will display those settings whenever the dialog is first opened after that preference file is loaded.

The Spectrogram Options and Spectrum Options dialogs allow you to save multiple named sets of options in a single preference file, and then recall them by name from a popup menu (see Chapter 3).

Canary saves the settings of all of these dialog boxes, the settings of the command panel's volume and rate control sliders, and the positions of the command, measurement, and track palette panels when you save preferences from the File menu (see below) or when quitting Canary.

Default preferences

When Canary is launched, it opens the default preference file, named "Canary.Prefs", which is located in the Preferences folder inside the System Folder. If this file does not exist (e.g., the first time that you run Canary 1.2), Canary creates it and saves a set of "factory default" preferences in it.

You can restore Canary's default preferences to their factory settings by choosing Revert to Defaults from the File menu. Alternatively, you can throw away the Canary.Prefs file while Canary is not running.

Current preference file

You can see what preference file is currently in use by selecting About Canary... from the apple menu.

Saving changes to the current preference file

At any time you can save all preference information to the current preference file by selecting Save Preferences from the File menu.

When you quit the program, Canary asks if you want to save changes to the current preference file, unless you disable this query in the General Options dialog (see the section "General options" later in this chapter).

You can save the current settings to the default preference file (Canary.Prefs) at any time (even if you are using some other preference file) by choosing Default from the submenu that appears when you choose Save Preferences As... in the File menu.

Creating a new preference file

To create a new preference file, choose Save As... from the submenu that appears when you choose Save Preferences As... in the File menu. (Preference files can be saved anywhere on your disk; the only preference file that must be in the Preferences folder is Canary.Prefs.) Canary saves all of the current settings into the new file, and makes the new file the current preference file.

Loading a preference file

To load a preference file, choose Open... from the Load Preferences... submenu on the File menu. All control settings and panel positions are reset immediately to the default values specified by the preference file you select.

Selecting Default from the Load Preferences... submenu loads the preferences saved in the default preference file (Canary.Pref).

You can automatically load a preference file other than Canary.Prefs when launching Canary by double-clicking on the preference file's icon in a Finder window.

Reverting to default preferences Select Revert to Defaults from the File menu to revert to the factory default preferences.

Manual scaling

Canary provides two ways of changing the horizontal and vertical scale of a signal view or a correlation function. As illustrated in Chapter 1, you can use the stretch, compress, and zoom buttons on the command panel. Alternatively, you can specify a particular scale for a view or correlation function using the Manual Scaling... command on the Options menu. For example, selecting Manual Scaling... while the active pane contains a waveform brings up the Scale Waveform dialog box shown in Figure 8.2.

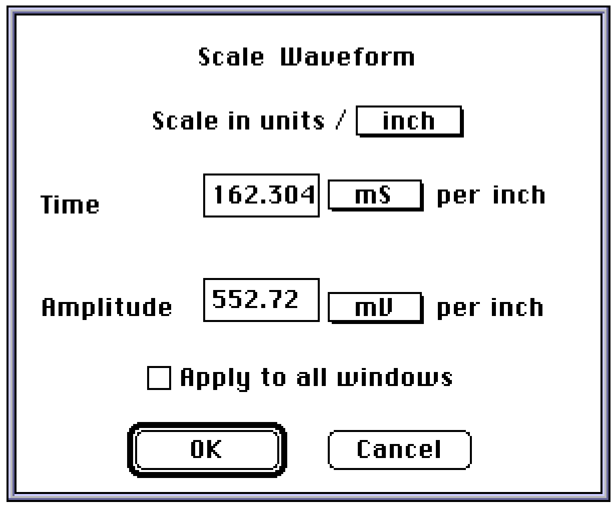

Figure 8.2. The Scale Waveform dialog box.

Similar dialog boxes appear for scaling spectrogram and spectrum panes, and for correlation functions. In all of the manual scaling dialogs, a pop-up units menu lets you choose inches, centimeters, or pixels as the display unit. In each dialog, you specify the scale for the horizontal and vertical axes by typing in the number of dimensional units (i.e., seconds, pascals, hertz, decibels, joules/m^2/Hz, or correlation units) per display unit. Pop-up menus let you choose a multiple of the basic dimensional unit (e.g., mS, S, µS for time axes). The values that are displayed when the dialog box first appears are the scale at which the plot is currently drawn.

Checking the Apply to all windows box applies the specified scale to all windows that contain a pane of the same type as the pane whose scale is displayed in the dialog box. This feature is particularly useful for comparing panes that are initially drawn at different scales (e.g., because the signals are of different durations).

Changing the scale of one pane in a signal window always changes the scale of at least one other pane because the scales of panes are linked by their common dimensions. For example, changing the time scale of a spectrogram automatically changes the time scale of the waveform in the same window, so that the plots in the two panes remain aligned with each other.

Resizing a window (either by clicking on its zoom box or by dragging its size box) rescales the contents of the window. Thus, if you manually set the scale for a window and then change the window's size, you will have to reset the scale manually if you want the resized window to have the same scale (in units per centimeter) as you initially set.

Display options

Selecting Display... from the Options menu brings up the Display Options dialog box (Figure 8.3). The display options that you select remain in effect until you change them, load a new preference file, or quit Canary.

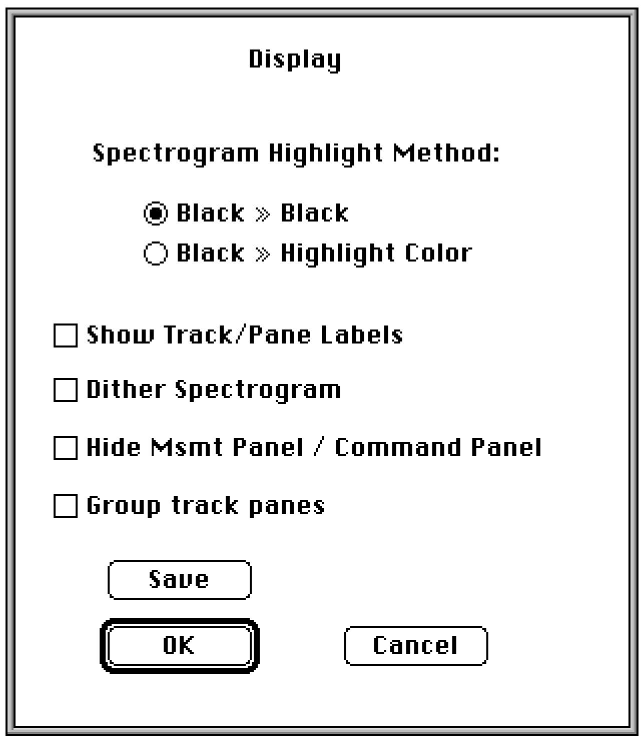

Figure 8.3. The Display Options dialog box.

Spectrogram Highlight Method

Canary provides two methods of displaying highlighted regions of spectrograms. If Black » Black (the default) is selected, white regions of the spectrogram appear in the highlight color (hue, saturation, and brightness) specified by the Macintosh Color control panel[1]; darker areas are drawn with the same hue and saturation, but less brightness, so that black areas appear black. If Black » Highlight Color is selected, black regions of the spectrogram appear in the highlight color, and lighter regions appear with the same hue and brightness as the highlight color, but lesser saturation.

Show Track/Pane Labels

When Show Track/Pane labels is selected, Canary displays a label at the bottom of each pane of a signal window (Figure 8.4). The label indicates the pane type (WVF, SPG, or SPK), the Track number, and a pane title. By default, Canary assigns a pane title consisting of the file name followed by "`.XXX.n`", where `XXX` is the pane type, and `n` is the track number. You can edit the title by highlighting text with the mouse, typing, and pressing ***return*** or ***enter***. Titles are limited to 31 characters in length.

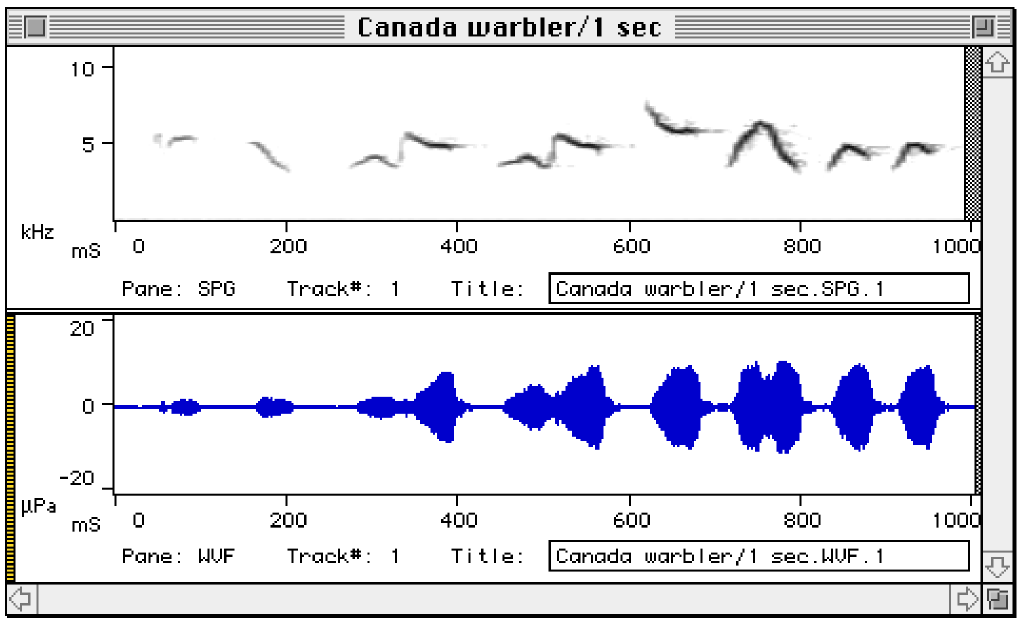

Figure 8.4. A signal window with track/pane labels displayed.

[1]The Color control panel is in the folder named "Control Panels" inside the System Folder. The Control Panels folder is usually accessible under the apple menu.

Dither Spectrogram When **Dither Spectrogram** is selected, Canary uses varying densities of white and non-white pixels to simulate grayscale values in spectrograms. Spectrogram dithering is useful mainly on monitors displaying fewer than sixteen shades of gray.[1] Spectrogram dithering involves some loss of precision, because the shade of each pixel is in part determined by the values of neighboring pixels. Dithered spectrograms take longer to redraw than non-dithered spectrograms.

Hide Measurement and Command Panels When the box labeled **Hide Msmt Panel / Command Panel** is checked, the measurement and command panels are hidden when Canary is switched into the background (i.e., when you activate another program, including the Finder). If the box is unchecked, the panels remain visible even when Canary is in the background.

Group Track Panes Canary can arrange the views (waveform, spectrogram, and spectrum) in a multi-track window in two different ways. If you select **Group track panes**, Canary arranges all of the views for a given track in adjacent panes (e.g., shown from top to bottom in the window: *SPK-1, SPG-1, WVF-1, SPK-2, SPG-2, WVF-2, ...*). If **Group track panes** is not selected (the default), all of the views of each type are arranged together (e.g., *SPK-1, SPK-2, SPG-1, SPG-2, WVF-1, WVF-2, ...*).

Speed options

Selecting **Speed...** from the **Options** menu brings up the Speed Options dialog box (Figure 8.5), which lets you adjust the speed with which Canary performs background calculations and the speed of waveform drawing. The speed options that you select remain in effect until you change them, load a new preference file, or quit Canary.

[1]Within limits determined by your display hardware, the number of colors or shades of gray that your monitor displays is determined using the Monitors control panel (in the folder named "Control Panels" inside the System Folder). The Control Panels folder is usually accessible under the apple menu.

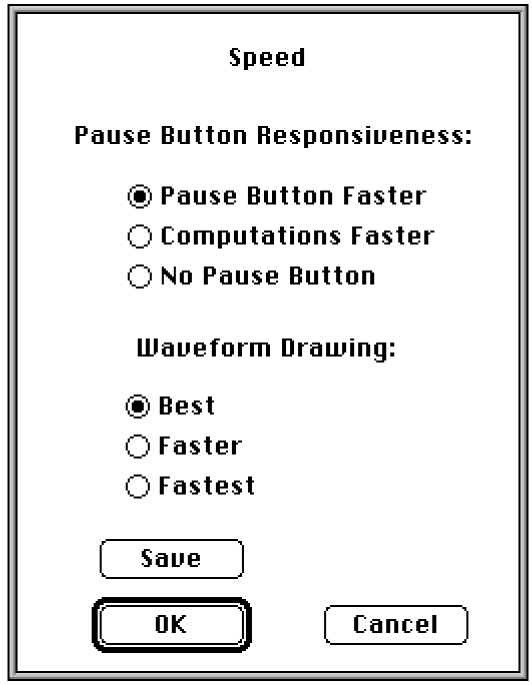

Figure 8.5. The Speed Options dialog box.

Pause Button Responsiveness

When Canary is performing calculations that can be done in the background (computing spectrograms, spectra, correlations, and running batch processes), there is a tradeoff between computation speed and the responsiveness of the machine in interacting with the user. Depending on whether you select Pause Button Faster, Computations Faster, or No Pause Button, Canary allows more or less time for tasks other than background computation. This choice affects the responsiveness of the machine to all user interaction (not just clicking the Pause button), and the speed of any other applications that are running at the same time as Canary. **If you select No Pause Button, you will not be able to do anything with the Macintosh while Canary is performing computations.**

Waveform Drawing

Canary draws waveforms by drawing straight lines between points on the screen that represent successive samples in a digitized sound. If Best waveform drawing is selected, every point in the digitized signal is used to draw the waveform. If Faster or Fastest is selected, Canary skips some points, which speeds up the drawing process at the expense of precision. Although very sharp, brief peaks in the waveform may be missed in Fastest mode, in practice the waveforms drawn in Best and Fastest modes are often indistinguishable from each other. You should experiment with the drawing modes to see how they affect the appearance of the signals that you work with. The speed difference between Best and Fastest modes is most apparent with long signals on slow machines.

General options

Selecting General... from the Options menu brings up the General Options dialog box (Figure 8.6), which lets you specify whether Canary should display an Open File dialog when the program first starts up, and what action Canary should take regarding changes to the current preference file when quitting. The options that you select remain in effect until you change them, load a new preference file, or quit Canary.

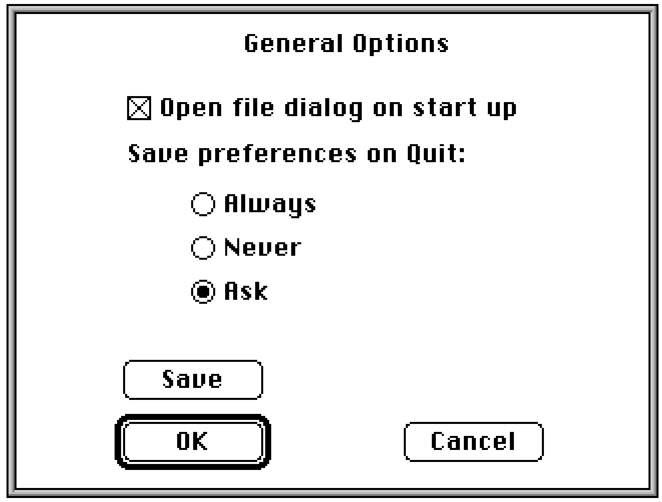

Figure 8.6. The General Options dialog.

Chapter 9 Printing and Graphics Export

About this chapter

This chapter explains how to print the contents of a signal or correlation window from within Canary. It also introduces the Flash-It shareware utility program, which enables you to copy any portion of a Canary window (or anything else on the Macintosh screen) for export to word-processing or graphics programs. Further details on the use of Flash-It are provided in the documentation included with the program.

Printing the contents of a Canary window

To print the contents of the active signal or correlation window, select Print... from the File menu. Selecting Print... first brings up the Print dialog box for the printer that is currently selected by the Chooser (under the apple menu; see your Macintosh system documentation for more information about the Chooser). Figure 9.1 shows the Print dialog box for an Apple LaserWriter printer. The Print dialog box allows you to specify the number of copies to be printed, and lets you set various options for your printer (e.g., number of copies). The Print dialog box may look slightly different from the one in Figure 9.1, depending on the printer driver software installed on your system.

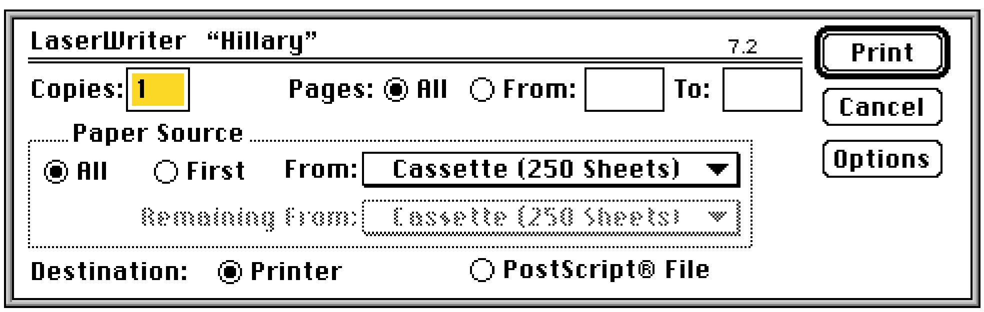

Figure 9.1. The Print dialog box for an Apple LaserWriter printer. The dialog box for your printer may differ.

When you click the OK or Print button, Canary sends an image of the contents of the current window to the selected printer. The window borders (i.e., the title bar and scrollbars) are not printed.

If the window contains a grayscale spectrogram, and if grayscale printing was selected in the Print dialog box (or the Printer Options dialog activated by the Options button in the Print dialog), Canary prints the window with the best grayscale resolution possible with the selected printer. If black-and-white printing was selected in the Print dialog box, or if the printer is incapable of grayscale printing (e.g., Apple ImageWriter II, Apple Personal LaserWriter

SC), the spectrogram image is automatically dithered. Dithering uses varying densities of black dots to represent shades of gray. The quality of the printed image depends on your printer.

Some printing options (e.g., paper size, vertical or horizontal orientation of the image on the page) are controlled using the Page Setup... command on the File menu. Choosing Page Setup... brings up a dialog box like the one shown in Figure 9.2. The exact appearance of the Page Setup dialog box depends on which printer driver you are using. See your printer documentation for an explanation of options in the Page Setup dialog box. Page setup options that you set while using Canary apply to all documents until you quit Canary or set new options.

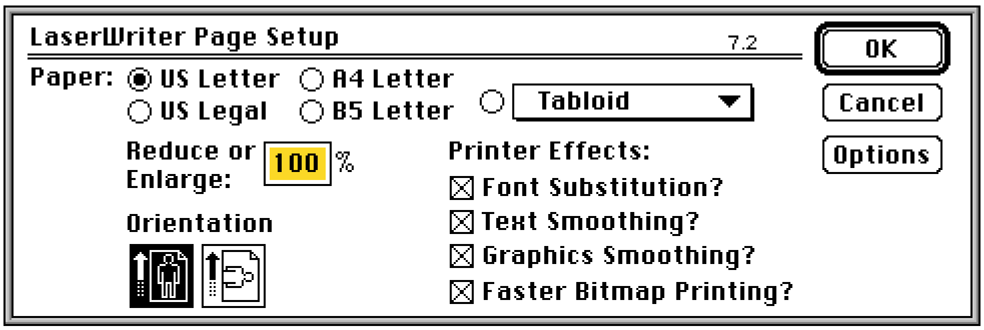

Figure 9.2. The Page Setup dialog box for an Apple LaserWriter printer. The dialog box for your printer may differ.

To print the contents of the DataLog window, either save a text report from the log and print it from a word processing or spreadsheet program, or use Flash-It to print a graphic image of the DataLog window (see below). Using Flash-It will print only the part of the data log that is shown on the screen. If you want to print all of the data in the log, and the log contains more data than will fit in the window, print a text report of the log. Text reports from data logs are discussed in Chapter 4.

Exporting graphic images with Flash-It

Flash-It is a utility program that is distributed with Canary. You can use Flash-It to copy any portion of the image displayed on your Macintosh screen, then paste the image into a document in another application program, such as a word processor or graphics program.[1] Flash-It can also send a copied screen image directly to a printer. Many of the illustrations in this manual were made using Flash-It.

[1]Flash-It saves screen images in PICT format.

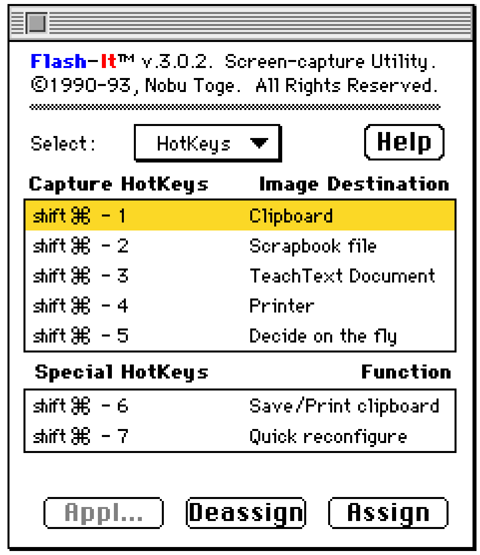

Figure 9.3. The Flash-It control panel.

You control how Flash-It operates using the Flash-It control panel (Figure 9.3). The control panel includes a Help button that brings up on-screen documentation. Flash-It also comes with a documentation file that you can read with most Macintosh word processors.[1]

Flash-It can copy either a region of the screen that you select with the mouse, or the topmost window, or the entire screen. See the Flash-It documentation for further details.

If you use Flash-It's "copy topmost window" feature, you should bear in mind that Canary's command and measurement panels are both considered windows by the Macintosh, and that they always "float" in front of all other windows within Canary. Thus if you want to use Flash-It to copy a signal or correlation window, you should first hide the command and measurement panels.

Flash-It is shareware. It is not free, and its price is not included with the price of Canary. It is distributed with Canary with the consent of the author with the following understanding: you are entitled to try the program at no charge, but if you find it useful, and continue using it, you should send the very modest shareware fee to the author as described in the program's

---

[1]If the Flash-It™ documentation file does not open when you double-click on it, try starting a word processor and opening the documentation file using the Open... command on the word processor's File menu.

documentation file. People who put time and effort into writing high-quality programs like Flash-It, and who choose to distribute their products as shareware (thereby sparing consumers from prices inflated by marketing and distribution costs) deserve your support. If you use Flash-It, send in the shareware fee!

Chapter 10 File Formats

About this chapter

Canary enables you to save and retrieve data in a variety of formats, which makes it easy to transfer information between Canary and other application programs, including word processors, spreadsheets and statistics programs, and other programs that work with sound data.

This chapter begins by making a distinction between the types of data stored in files and the format in which the data are stored. Subsequent sections discuss how to open files saved by Canary or other programs, and how to save files in the various formats. The last section includes a discussion of factors that might influence the choice of format for saving files.

Background

Data types and file formats

Canary can open and save files containing six different types of data. The *data type* determines whether Canary interprets the information in a file as representing a sound (waveform), a spectrogram, a spectrum, a correlation function, a correlation array, or a data log.

Each file is characterized by its *file format*. The file format determines how data are organized or structured within the file, irrespective of what the data represent (the data type). Canary can read and write six different file formats, but not all of them can be used for every data type. Table 10.1 shows which formats are available for each type of data. The Canary file format allows sound, spectrogram, and spectrum data to be saved together in a single file, as discussed later in this chapter. All of the other formats require that each file contain only one type of data.

Table 10.1. File formats available for each type of data used by Canary.

		DATA TYPE					
Icon	**FILE FORMAT**	Sound	Spectrogram	Spectrum	Correlation function	Correlation array	Data log
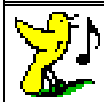	Canary	•	•	•	•		
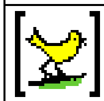	MATLAB	•	•	•	•	•	•
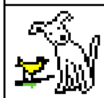	SoundEdit	•					
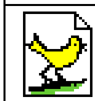	Text	•	•	•	•	•	•
	AIFF	•					
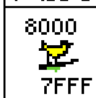	Binary	•					

Two files that contain the same data saved in different formats are indistinguishable from each other when opened in Canary. However, the format in which a file is saved determines what other application programs can open the file. Files containing the same data in different formats may differ in size as well. The effects of file formats on file size and compatibility with other programs are discussed later in this chapter.

For an explanation of how to convert files from one format to another, see the discussion of batch file conversion in Chapter 11.

Each file format saved by Canary is represented by a unique icon in Finder windows displayed in icon view (Table 10.1). Except for those formats that can only contain sound (waveform) data (SoundEdit, AIFF, and Binary formats), you cannot tell from a file's icon what type of data it contains.

Opening files

Choosing Open... from the File menu brings up a standard Macintosh Open File dialog box. The scrolling list box in the dialog box displays the names of files that are saved in any of the six file formats that Canary can read. Selecting a file that was saved by Canary will, in most cases, cause the file to open immediately. In some cases (mainly with files that were saved by other programs), Canary may not be able to determine what type of data the file contains. In this case, a dialog box appears that allows you to specify the data type (Figure 10.1). Depending on which data type you indicate, Canary may then ask you to specify various parameters needed to interpret the data in the file (Figures 10.2, 10.3, and 10.4).

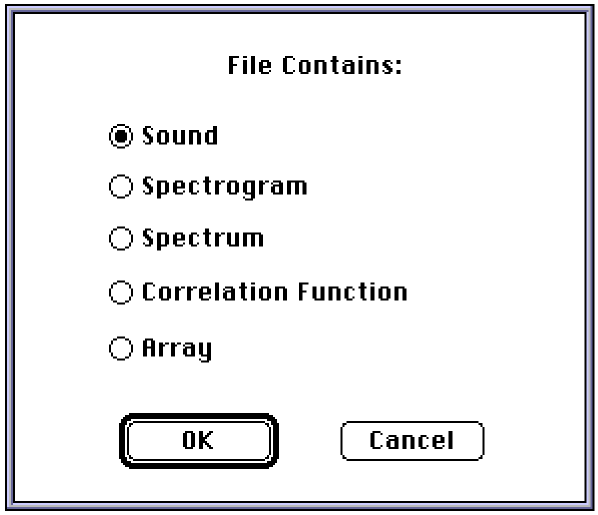

Figure 10.1. The File Contents dialog box, which appears when Canary is unable to identify the type of data in a file that is being opened.

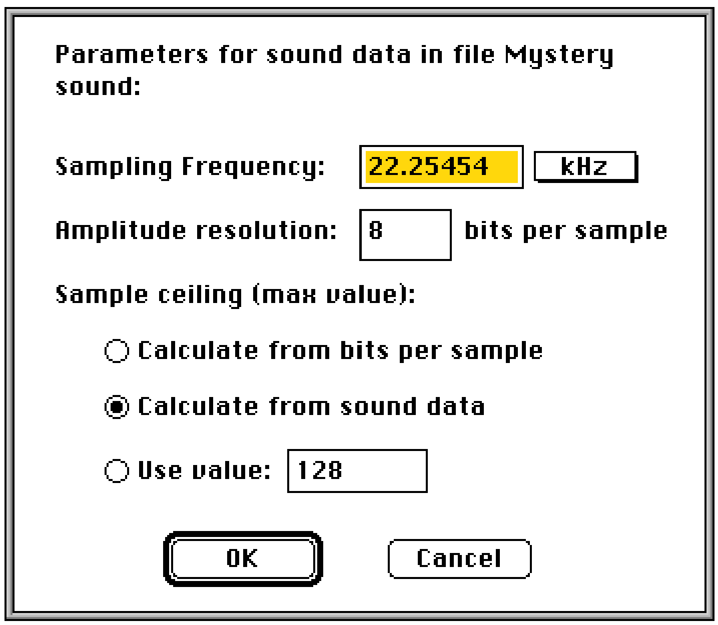

Figure 10.2. The Sound Parameters dialog box, which appears when Canary attempts to open a sound file that does not include data describing how the sound was digitized.

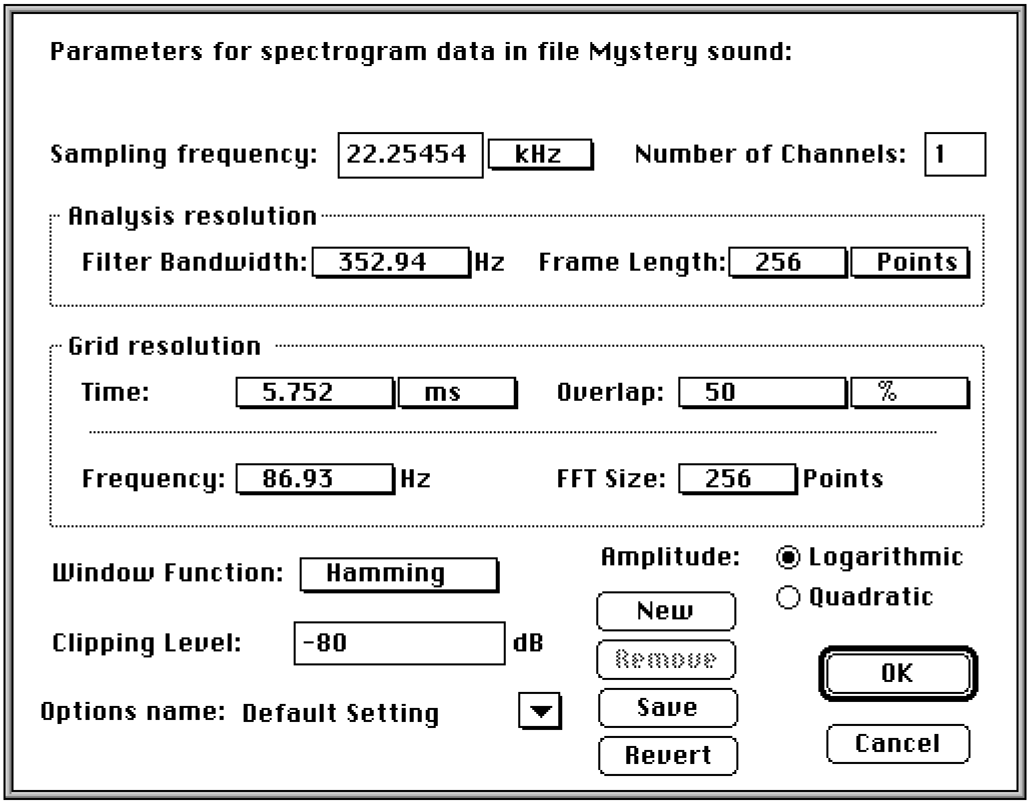

Figure 10.3. The Spectrogram Parameters dialog box, which appears when Canary attempts to open a spectrogram file that does not include data describing the parameters with which the spectrogram was calculated. A similar dialog box appears when opening a spectrum file that lacks parameter data.

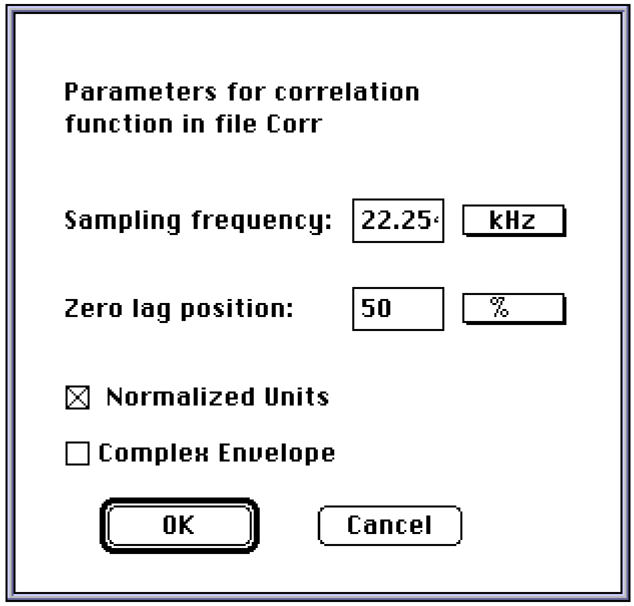

Figure 10.4. The Correlation Parameters dialog box, which appears when Canary attempts to open a file containing a correlation function that does not include data describing the correlation parameters.

Saving files

When the frontmost window is a signal window, the File menu contains four commands for saving the contents of the window: Save Window, Save Window As..., Save *view* As..., and Save Sound Tracks... where *view* is replaced by the name of the active pane (Waveform, Spectrogram, or Spectrum). Save Sound Tracks... is available only if the window contains a multi-track signal, and one or more tracks are selected on the track palette (see Chapter 5 for a discussion of saving selected tracks).

Save Window (which is the default choice activated by Command-S) saves whatever views of the signal are currently available into a single Canary-format file. The spectrogram and spectrum views will be included in the file if they have been calculated, even if you have hidden them. Save Window As... is equivalent to Save Window, except that it forces the standard file dialog box to appear, so that you can change the name of the file that is saved. When you save a window, the only format available is Canary's own format.

Save *view* As... allows you to save just the view in the presently active pane (i.e., the pane that has the hatched vertical bar at its left-hand edge). You specify which format should be used to save the file by clicking on one of the format choices at the bottom of the standard Save File dialog box. The type of the pane that you are saving determines which file format choices are available (as indicated in Table 10.1); the names of unavailable formats are dimmed in the Save File dialog box.

When saving single views, remember that a saved waveform can always be used to recalculate a spectrogram or spectrum, but that a spectrogram or spectrum cannot be used to reconstruct either of the other two views.

When the frontmost window is a correlation array or a data log, the File menu contains three commands for saving the array or log: Save Log, Save Log As..., and Save Text Report.... The Save Log... and Save Text Report... commands also appear on the DataLog menu. Save Log (which is the default choice activated by Command-S) saves the contents of the data log or correlation array into a file in MATLAB or Text format. Save Log As... is equivalent to Save Log, except that it forces the standard file dialog box to appear, so that you can change the name of the file that is saved. Files saved with the Save Log command can be reopened by Canary later, but are often inconvenient for exporting data to other programs because they are not organized into columns and lack column and line labels. Save Text Report... saves the data log or correlation array in a text format that can include column and line labels. Text report files can be imported into most word processing, spreadsheet, and statistics programs, but cannot be reopened by Canary. See the discussions of text report files in Chapters 6 and 11 for further details.

Which file format to use?

The choice of format for saving files depends on considerations of compatibility with other programs, file size, and convenience.

Compatibility

If you plan to use a file created by Canary with another program, the format(s) that the other program can read may determine which format you use to save your files. For example, most other sound-processing programs can read sound files in AIFF format, but may not read any of the other formats in which Canary can save sounds. If you want to use MATLAB to analyze or manipulate sound or spectrogram files created with Canary, you should save the files in MATLAB format.[1]

SoundEdit format should be used only if the sampling rate with which a sound was recorded is equal to one of SoundEdit's standard sampling rates (22 kHz, 11 kHz, 7 kHz, or 5 kHz).[2] Signals digitized at non-standard rates should be saved in one of the other formats.

Size

[1]MATLAB (The Mathworks, Natick, Massachusetts) is a powerful matrix manipulation program that can be used to perform virtually any type of digital signal processing function (e.g., filtering, various kinds of spectral analysis). Using MATLAB's signal processing tools requires considerably more mathematical sophistication than Canary demands.

[2]Signals digitized at non-standard rates that are saved in SoundEdit format will not have the correct sampling rate when reopened by Canary (or any other program).

The different file formats available can yield files of substantially different sizes for the same data (Tables 10.2, 10.3). The exact size ratios between files of different formats containing a given set of data may vary somewhat depending on the particular data being saved, but the figures in Tables 10.2 and 10.3 are representative of the relative sizes of files in the different formats. For signals digitized at higher than 8-bit amplitude resolution, the size differences among sound files in different formats are smaller.

FILE FORMAT	SIZE (KBYTES)
AIFF	24
SoundEdit	24
Binary	45
Canary	89
MATLAB	176
Text	218

Table 10.2. Sizes of files containing the same 1.0-second sound (digitized at 22 kHz, 8-bit samples), saved in different formats. File formats are listed in order of increasing size.

FILE FORMAT	SIZE (KBYTES)
Canary	174
MATLAB	347
Text	476

Table 10.3. Sizes of files containing the same spectrogram of a 1.0-second sound (digitized at 22 kHz), saved in different formats. File formats are listed in order of increasing size.

Convenience

The Canary file format is the only format that allows you to save any combination of Canary's three views of a signal (waveform, spectrogram, and spectrum) in one file. If you plan to use a signal's waveform, spectrogram, and spectrum together repeatedly, saving all of the views (or any two) in a Canary-format file eliminates having to reconstruct the spectrogram and spectrum views each time you open the signal.

When you open a window file that was saved in Canary format (i.e., a file saved with the Save Window command), the window that appears initially displays all of the views that were saved in the file.

Chapter 11 Batch Processing

About this chapter

Canary's batch processes provide a mechanism for automatically performing the same operation on an arbitrarily large number of files and saving the results. For example, the batch spectrogram process lets you make spectrograms of many files at a time, using a set of spectrogram parameters that you specify only once. The batch processes available are spectrogram and spectrum calculation, correlation, and file type conversion. All batch processes can run in the background, so you can continue to do other work with Canary or another application program while a batch process is running.

This chapter begins by discussing how to specify input files and output destinations for Canary's batch processes. The remainder of the chapter discusses each of the batch processes individually. The batch spectrogram and spectrum processes are included together in a single section on spectrum analysis.

Input and output

For each of Canary's batch processes, you must specify which files are to be used as input and where the output files should be saved. The input files must all be stored in the same folder, and all of the output files will be saved in a single folder (which must be different from the input folder). This section discusses those aspects of input and output specification that are common to all four types of batch processes. Later sections discuss each type of process individually.

When you choose Batch from the Process menu, a submenu appears showing the four types of batch processes (Figure 11.1).

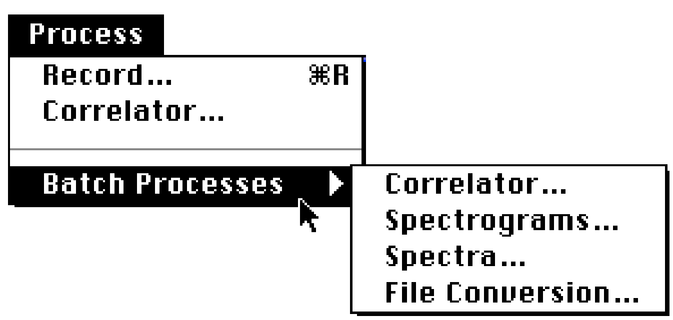

Figure 11.1. The Batch Processes submenu.

When you select any process from the submenu, a dialog box appears asking you to select a folder to be used as input to the process (Figure 11.2). If you select batch correlation, two input dialog boxes appear in succession, as discussed under “Correlation” later in this chapter.

Although the prompt message in the dialog box differs depending on which type of batch process is being performed, all of the input selection dialog boxes work the same way. You specify the folder containing the files to be used as input to the batch process by opening that folder, then clicking on the Select button. The folder that is selected is thus the folder whose name appears *above* the scrolling list in the dialog box (e.g., “Sounds” in Figure 11.2). You can also select a folder by double-clicking on the name of any file in the folder. If the item that is highlighted in the scrolling list is the name of a folder rather than the name of a file, the Select button is replaced by an Open button; double-clicking on a folder name opens that folder. Once you have opened a folder (so its contents are shown in the list), the file that is highlighted in the list has no effect on the selection, since you’re selecting a folder, not a file.

You specify which files in the folder are to be used as input by entering a file name pattern in the designated field of the dialog box. Initially, when the dialog box first appears, the file name pattern field contains the single character “*”, which is a “wildcard” character that can stand for any string of characters (including spaces or the null string). Thus, if you leave the file name pattern as “*” when you click on the Select button, all files in the folder are processed. If you were to change the file name pattern to “Cricket *”, Canary would process all files in the folder whose name matched the pattern (e.g., “Cricket 1”, “Cricket 2”, “Cricket frog”), and ignore files with names that did not match (e.g., “frog”, “Field cricket”). You can specify a pattern containing more than one “*”. For example, the pattern “*warbler*” matches any file name containing the string “warbler”. Upper- and lower-case letters are regarded as the same in the file name pattern.

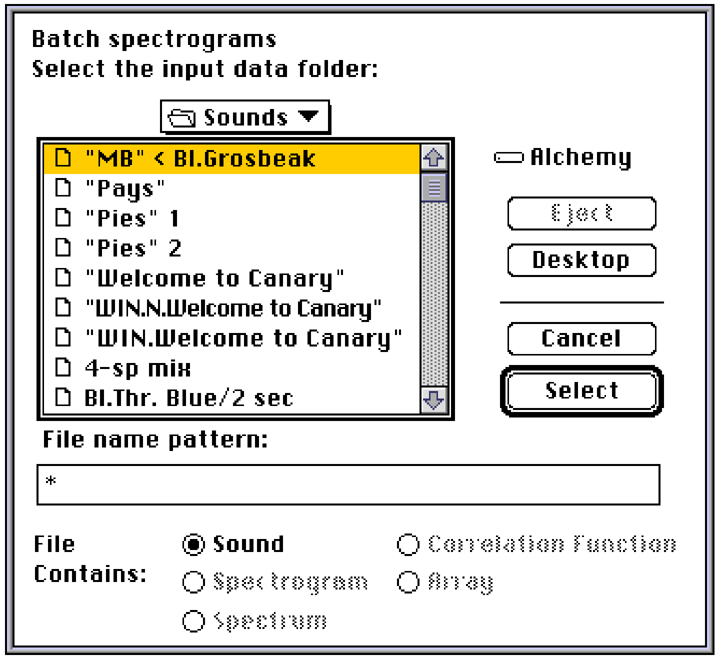

Figure 11.2. The input selection dialog box for batch spectrograms.

After you select the input folder (or the two input folders if you're doing a batch correlation), the output folder selection dialog box appears (Figure 11.3). You select the folder in which you want the output of the process to be saved by opening the folder and clicking on the Select button, or by double-clicking on the name of any file in the folder. The Select button is replaced by an Open button if the highlighted item in the scrolling list is the name of a folder rather than a file. The format in which the data are saved is determined by which radio button is clicked at the bottom of the dialog box. The file formats available depend on which process is being run, as discussed later in this chapter.

For batch spectrograms and spectra, the output folder selection dialog contains a checkbox labeled Attach source data to result. If this option is checked, the resulting file contains the waveform data for the sound, as well as the spectrogram or spectrum. This option is available only for files being saved in Canary format.

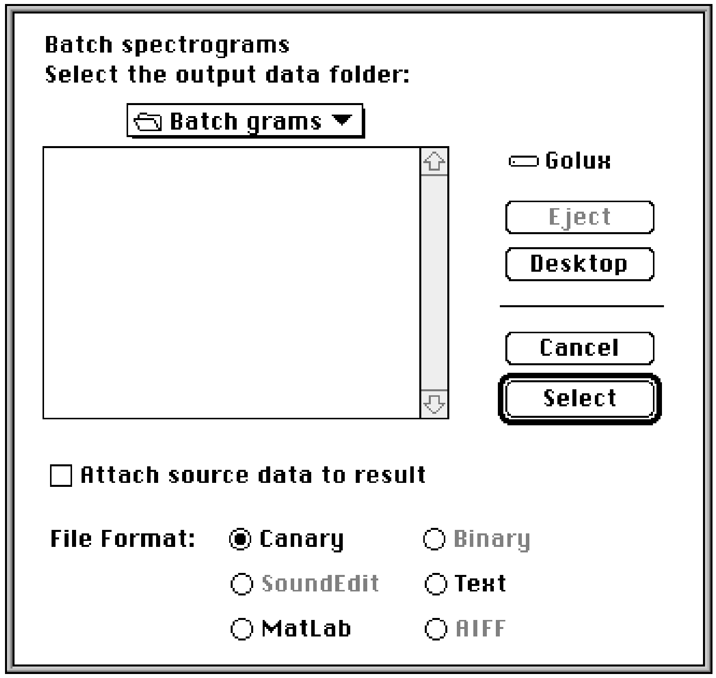

Figure 11.3. The output folder selection dialog box for batch spectrograms.

Correlation

The batch correlator computes correlations for each possible pair of files in the two input sets. The output is a rectangular array containing the peak correlation value from each pair of files. A correlation array can be viewed in its own window within Canary (Figure 11.4). It can also be saved in a variety of formats for export to other application programs, as discussed below.

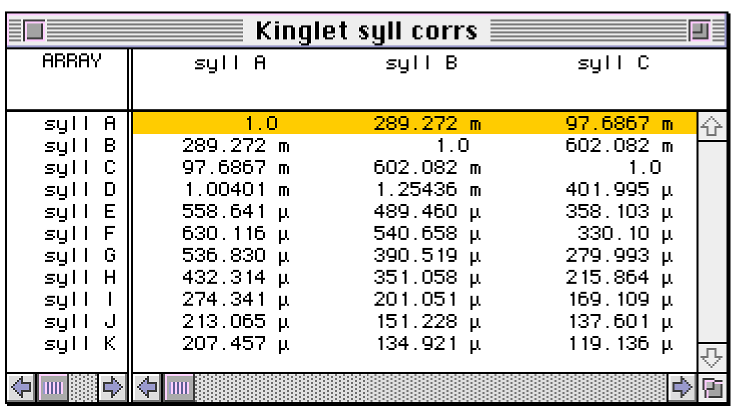

ARRAY	syll A	syll B	syll C
syll A	1.0	289.272 m	97.6867 m
syll B	289.272 m	1.0	602.082 m
syll C	97.6867 m	602.082 m	1.0
syll D	1.00401 m	1.25436 m	401.995 µ
syll E	558.641 µ	489.460 µ	358.103 µ
syll F	630.116 µ	540.658 µ	330.10 µ
syll G	536.830 µ	390.519 µ	279.993 µ
syll H	432.314 µ	351.058 µ	215.864 µ
syll I	274.341 µ	201.051 µ	169.109 µ
syll J	213.065 µ	151.228 µ	137.601 µ
syll K	207.457 µ	134.921 µ	119.136 µ

Figure 11.4. A correlation array as shown in Canary.

Input specification

When you select Correlator from the Batch submenu, Canary prompts you for two sets of input files (Figure 11.5), which may be the same. You specify whether you want to correlate waveforms or spectrograms using the radio buttons at the bottom of the dialog box. (Although the prompt next to the buttons is “File Contains:”, a file could contain both a sound and a spectrogram, if it’s a Canary-format file.)

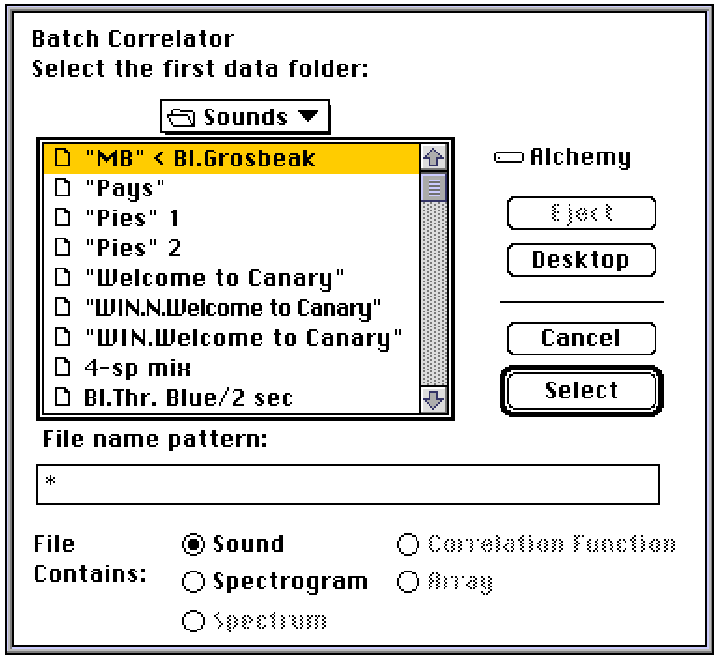

Figure 11.5. The first input selection dialog box for the batch correlator. The second input selection dialog box is identical except for the prompt message above the scrolling list of files.

One important difference between the batch correlator and the interactive correlator (discussed in Chapter 7) is that the input files for batch spectrogram correlations must already contain spectrograms; the interactive correlator, in contrast, can automatically make spectrograms from files containing waveforms only. Thus if you have two sets of sounds on which you want to run batch spectrogram correlations, you should first use the batch spectrogram process to make spectrograms of the sounds. You can then use the resulting spectrogram files as input to the batch correlator.

Output specification

After you have specified both sets of input files, Canary asks where and in what format to save the correlation array to be produced (Figure 11.6). Correlation arrays can be saved either in MATLAB or Text format; MATLAB files are usually smaller than text files containing the same data. A correlation array saved as text appears (when opened in a word processor or a spreadsheet program) as a single list of numbers, rather than a rectangular table like that shown in Figure 11.4. If you want a correlation array saved as a rectangular text table that can be read into a word processing, spreadsheet, or other application program, use the Save Text Report... command discussed below.

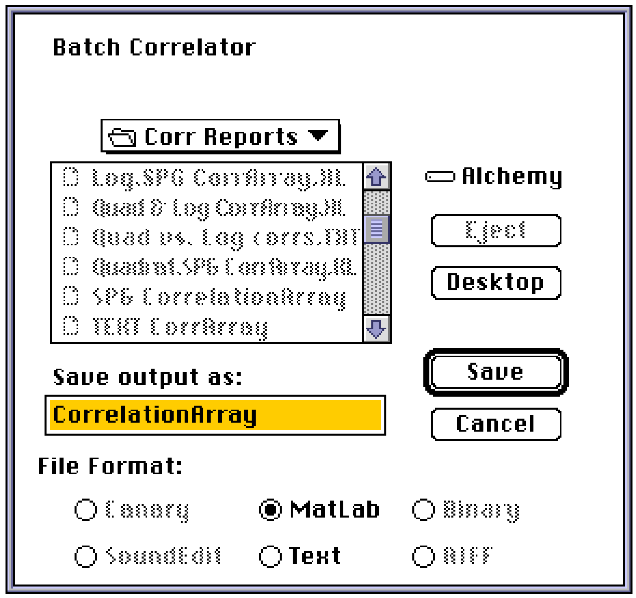

Figure 11.6. The batch correlator output dialog box.

Correlator options

After you specify the output destination, the batch version of the Correlator Options dialog box appears (Figure 11.7). Either the waveform or spectrogram button is checked automatically, depending on which type of input you selected in the previous two dialog boxes. For spectrogram correlations, the Options... button displays the parameters (i.e., resolution, window function, etc.) used to compute the spectrogram in the first input file. All of the spectrograms to be used as input must have the same parameters, or an error will occur.

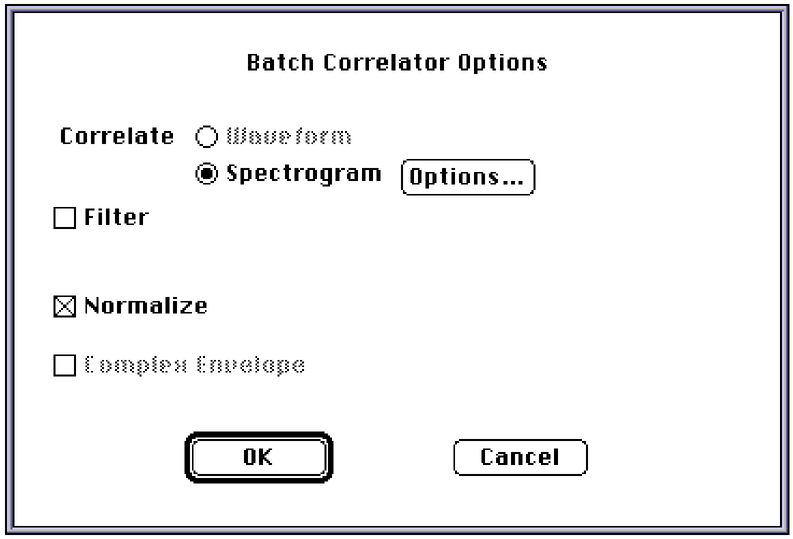

Figure 11.7. The batch spectrogram Correlator Options dialog box.

The Filter, Normalize, and Complex Envelope checkboxes are discussed in Chapter 7.

When you click on the OK button, Canary starts calculating the correlations and displays a status window titled “Batch Correlator”, which shows the progress of the calculations. The progress bar in the status window shows the progress of the entire batch process; the text in the window shows which two files are currently being correlated. Clicking the Pause button in the status window makes the Cancel button available. Clicking the Cancel button aborts the batch correlation.

If you are correlating the contents of a folder with itself, Canary calculates only the unique correlations (i.e., it does not calculate both X*Y and Y*X, since these are the same); however, the resulting array does contain the redundant values.

The order in which files are processed for a batch correlation (and hence the order in which they are listed in the correlation table) depends on the sizes of the files and memory availability and is generally not predictable.

Correlation array windows

When all of the correlations have been calculated, Canary displays a window containing the correlation array (Figure 11.4). The title of the window is the same as the name specified for the batch correlation output file. The correlation array shows the peak correlation value for each pair of signals.

For example, the array in Figure 11.4 contains the correlations among spectrograms of thirteen syllables from the song of a golden-crowned kinglet (Figure 11.8). When the correlation array window is not large enough to display all of the rows or columns of the array, the vertical and horizontal scrollbars can be used to scroll the array within the window.

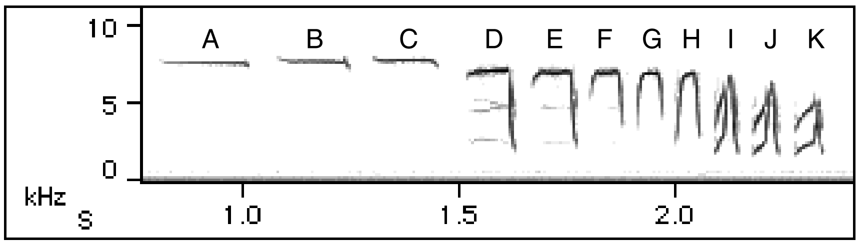

Figure 11.8. Partial song of a golden-crowned kinglet, showing 13 discrete syllables. Figure 11.4 shows part of an array of correlations among spectrograms of these syllables.

The correlation window is divided into two panes by a double vertical line. Each pane has an independent horizontal scrollbar. You can change the relative sizes of the two panes by positioning the mouse pointer on the line between the panes, and then dragging the line to a new position. (When you move the mouse pointer onto the pane divider, the pointer changes to the icon shown to the left.)

Saving a text report of a correlation array

To save a correlation array in a text file to be imported into a word processor, spreadsheet, or most other types of application programs, select Save Text Report... from the File menu (this command is available only when the frontmost window is a correlation array window). The dialog box shown in Figure 11.9 then allows you to specify the format of the file to be saved.

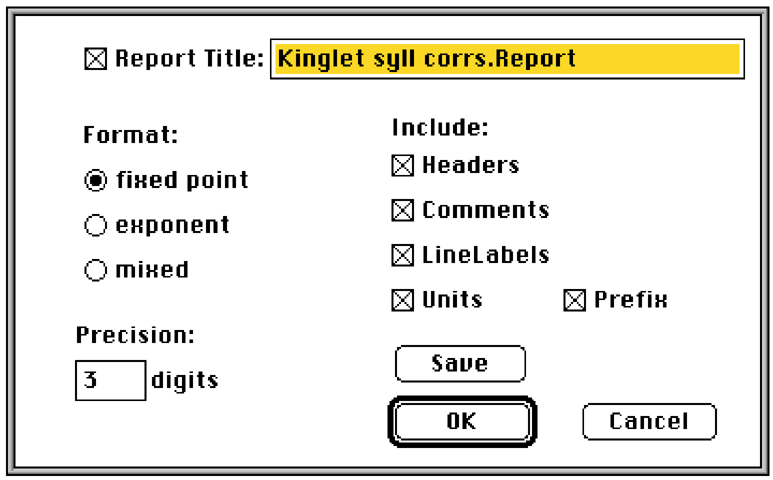

Figure 11.9. The Save Text Report... dialog box.

If Report Title is checked, the specified text is included on one line at the top of the file that is written. You can enter whatever text you want for the report title. Note that the report title is not the name of the file to be saved (although it may be the same); the file name is specified in a standard Save File dialog box that appears after you click OK in the Save Text Report dialog box.

The three radio buttons under Format: and the Precision field allow you to specify the format used for the numeric entries in the correlation array.

The five checkboxes under Include: allow you to specify how rows, columns, and individual numeric entries in the array should be labelled:

Headers: The first row of the file will contain text labels that identify which file from the first input set corresponds to each column.

Comments: This checkbox has no effect on the file. (Correlation text report files cannot include comments.)

LineLabels: The first column of the file will contain text labels that identify which file from the second input set corresponds to each row.

Units: This checkbox has no effect on the file.

Prefix: Appends the appropriate standard prefix before the unit label on each entry in the array. This checkbox is disabled if the Units checkbox is not checked. (See Appendix E for standard metric prefixes and their abbreviations.) The prefix is chosen so that the integer part of a fixed-point measurement is between 0 and 999. Thus, for example, a value of .000999 sec would be written as 999 µS, whereas .001000 sec would be written as 1 mS. Note that if you omit prefixes and specify few digits of precision, small values may appear as zeros. For example, a value of 3.67 mS would appear as 0.003670 in 6-digit fixed point output, but would be written as 0.00 seconds if 2-digit fixed-point output were requested.

In a text report file, each row of the correlation array is written as a single line of text, terminated by a carriage return character (called a "hard return" or "paragraph mark" in some word processors). Entries in successive columns within a line are separated by tab characters. Most spreadsheet and data analysis application programs can import such a "tab-delimited text" file. If you import a text report file into a word processor, you may initially find that the data (and the columns headers, if present) don't line up in columns as expected. You can align the data properly by adjusting the tab stop positions and line length (margin) settings. See your word processor's documentation for further details.

The table below shows a portion of a text report file generated from the correlation array shown in Figure 11.4. The text of the table was copied from the report file, pasted into the word processing file for this chapter, and the tab stops positioned for proper alignment of the columns.

	syll A	syll B	syll C
syll A	1	0.289	0.0977
syll B	0.289	1	0.602
syll C	0.0977	0.602	1
syll D	0.001	0.00125	0.000402
syll E	0.000559	0.000489	0.000358
syll F	0.00063	0.000541	0.00033
syll G	0.000537	0.000391	0.00028
syll H	0.000432	0.000351	0.000216
syll I	0.000274	0.000201	0.000169
syll J	0.000213	0.000151	0.000138
syll K	0.000207	0.000135	0.000119

Spectrum analysis

The batch spectrogram and batch spectrum processes allow you to specify a set of spectrum analysis parameters that are used to make spectrograms or spectra of all of the files specified in the input file selection dialog box. The output spectrograms or spectra are then saved in the location and format that you specified in the output specification dialog box. Selection of input files and output file location and format are discussed in the section "Input and output" earlier in this chapter. Spectra and spectrograms can be saved in Canary, MATLAB, or text formats.

When making batch spectrograms, Canary names the output files "SPG of *inFile*", where *inFile* is the name of the input file. Spectra are named "SPK of *inFile*". If a file named "SPG of *inFile*" or "SPK of *inFile*" already exists in the output folder, it will be overwritten.

After you have specified the input and output locations and file formats, the batch version of the Spectrogram or Spectrum Options dialog box appears. Figure 11.10 shows the Batch Spectrogram Options dialog box; the Batch Spectrum Options dialog box is similar to it.

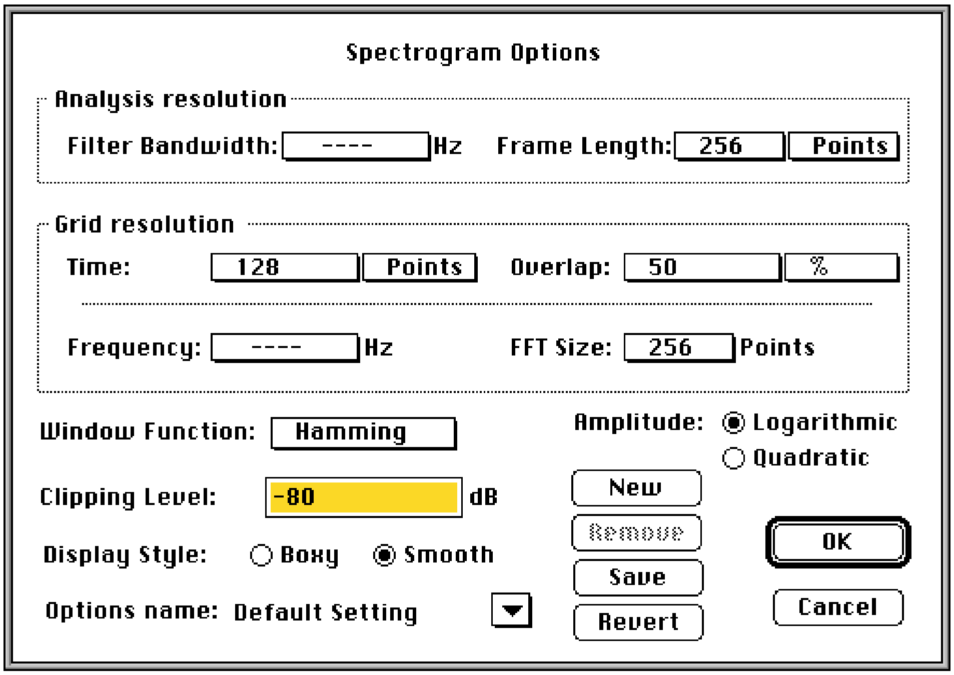

Figure 11.10. The Batch Spectrogram Options dialog box. Canary cannot display values for Filter Bandwidth and Frequency grid resolution because the signal's sample rate is used to derive these parameters from Frame Length and FFT Size respectively. Since the sample rate may differ among the input waveform files, the Filter Bandwidth and Frequency grid resolution of the resulting spectrograms may also differ.

The batch versions of these two dialog boxes are similar to the dialog boxes discussed in Chapter 3, except that the filter bandwidth and frequency grid resolution fields display "----", and cannot be changed. Recall from Chapter 3 that parameters that appear next to each other in these dialog boxes (i.e., Filter Bandwidth and Frame Length; Time grid resolution and frame Overlap; and Frequency grid resolution and FFT Size) are linked to each other; changing one parameter from each pair automatically changes its partner. When you select spectral analysis parameters for a single signal (rather than a batch), Canary uses the signal's sampling frequency (i.e., the rate at which it was digitized; see Appendix A) to calculate the filter bandwidth corresponding to a given frame length and the frequency grid resolution corresponding to a given FFT size. But when multiple signals are to be processed, Canary cannot display filter bandwidth and frequency grid resolution values that would apply to all of the signals because the input signals may have different sampling rates. For the same reason, you cannot change the units of the Frame Length field from Points to mS in the batch version of the dialog boxes.

Batch spectra are averaged over the entire input signal, using the frame overlap given in the Batch Spectrum Options dialog box (see Chapter 3 for an explanation of spectrum averaging).

Once you click OK in the options dialog box, spectrogram calculation begins. A status window named "Batch spectrogram" appears that shows the progress of the calculations. Buttons in the window let you pause or cancel the processing.

File conversion

The batch file converter reads files saved in any one of the formats that Canary can read (see Chapter 10), and saves a new copy of each file in a different folder in a format that you specify. You tell the file converter what type of data are contained in the input files using the radio buttons in the file converter's input selection dialog box (Figure 11.11). All of the input files must contain data of the same type (e.g., sound data (waveforms), spectrograms, correlation functions); the file formats of the input files may differ, however. For example, if the input files contain sound data, some might be saved in Canary format, some in SoundEdit format, and others in MATLAB format. If an input file is a Canary window file containing more than one type of data, the file converter uses only the data of the type specified by the selected radio button.

After you specify the input files, a standard batch process output dialog box appears to let you specify where and in what format to save the converted files, as discussed in the section "Input and output" earlier in this chapter. Each output file generated by the file converter has the same name as the input file from which it was created. The output folder that you specify must be different from the input folder, so that the original input files do not get overwritten. However, if the output folder already contains a file with the same name as a file being written by the file converter, it will be overwritten.

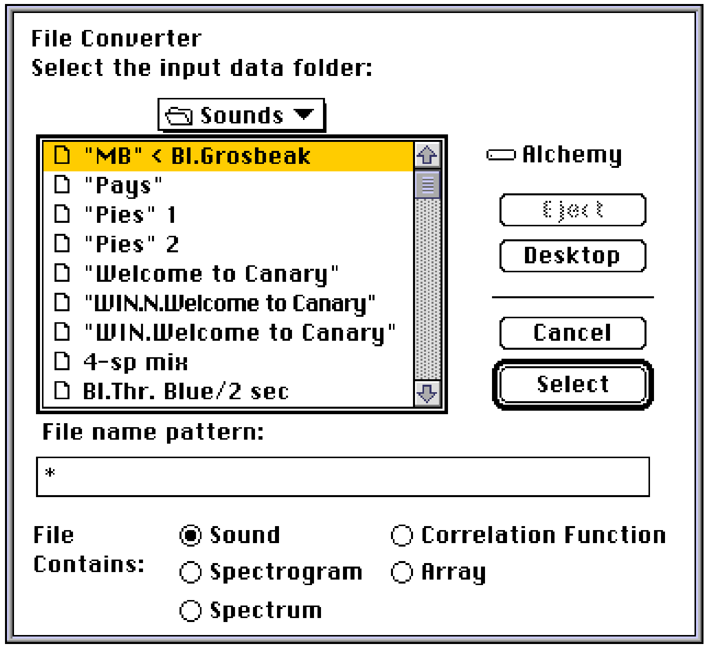

Figure 11.11. The file converter's input selection dialog box.

Once you click the Select button in the output selection dialog box, file conversion begins. A status box titled "File Conversion" displays the progress of the conversion process; buttons in the box allow you to pause or cancel the conversion.

Chapter 12 Canary Reference

About this chapter

This chapter provides a synoptic description of every command, dialog box item (e.g., buttons and checkboxes), graphic control and icon used by Canary. The chapter is most useful once you are familiar with Canary and the concepts underlying basic operations of the program (e.g., spectrum analysis or correlation). This background is provided by other chapters and appendices in this manual.

The chapter is divided into four sections, dealing with menu choices, items in dialog boxes, graphic controls (including mouse icons), and file icons. Each section begins with a brief explanation of how the section is organized.

Menu choices

This section contains an entry for every command on each of Canary's menus. The heading for each entry lists the name of the menu and the name of the command as it appears on the menu, separated by a slash (e.g., File / Open..., Edit / Copy). If there is a Command key equivalent for the command, it is also given in the heading for the entry. Entries are listed alphabetically by menu name and command.

Edit / Amplify...

Amplifies the selected portion of a waveform by the factor that you specify in the dialog box that appears when you select this command. Amplifying part of a signal by zero silences it. The window is immediately redrawn to reflect the change to the waveform. The spectrogram pane is not automatically updated after an Amplify operation; to update the spectrogram, you must recalculate it by clicking on the command panel's SPG button, choosing Make Spectrogram from the Panel menu, or typing Command G.

Edit / Clear

Deletes the selected portion of the waveform in the current window. The window contents are redrawn to reflect the deletion in the waveform and spectrogram panes. When deleting data from selected tracks of a multi-track signal, silence is appended to the selected tracks so that all tracks in the signal remain the same length.

Edit / Copy   Command C

Copies the selected portion of a waveform to the clipboard, overwriting any previous contents of the clipboard. A signal in the clipboard can be pasted into the waveform pane of any Canary window and is available to any other application that can read sound data. When data are copied to the clipboard, the calibration factor from the source signal is copied as well, so that the data's calibration is preserved when the data are pasted into another document (see Paste).

Edit / Cut   Command X

Copies the selected portion of a waveform to the clipboard and deletes it from the signal window. Any previous data in the clipboard are overwritten. A signal in the clipboard can be pasted into the waveform pane of any Canary window and is available to any other application that can read sound data. The window contents are redrawn to reflect the deletion in the waveform and spectrogram panes. When deleting data from selected tracks of a multi-track signal, silence is appended to the selected tracks so that all tracks in the signal remain the same length.

Edit / Filter Around...
Removes all energy from a signal that is included in the selected time interval *outside* of the selected frequency band. Requires a selection in the spectrogram pane. See also Edit / Filter Out... .

Edit / Filter Out...
Removes all energy from a signal that is included in the selected time interval *inside* of the selected frequency band. Requires a selection in the spectrogram pane. See also Edit / Filter Around... .

Edit / Paste Command V
Pastes sound data from the clipboard into the waveform pane of the current window. If part of the signal in the current window is selected, the data that are pasted in replace the selected section. If there is no selection, the insertion is done at the blinking insertion point or (if cursors are on) at the point where the two cursors are superimposed on each other. The window is redrawn to reflect the insertion in the waveform pane. If the signal into which data are being pasted is a multi-track signal, the number of tracks selected by on the track palette must equal the number of tracks in the clipboard. If a Paste operation increases the length of some track(s), silence is appended to the remaining tracks to keep all tracks the same length. If a Paste operation shortens the length of some track(s), silence is appended to those tracks to keep all tracks the same length. After a Paste operation, the portion of the re-drawn spectrogram that is synchronized with the newly inserted data is blank; to update the spectrogram, you must recalculate it by *option*-clicking on the command panel's SPG button, choosing Make Spectrogram... from the Panel menu, or typing Command G. Sound data is placed in the clipboard by a Cut or Copy command in Canary or in some other sound-processing program (e.g., SoundEdit).

Edit / Select All Command A
Selects all of the data in selected tracks of the active pane of the current window. Equivalent to double-clicking on the plot area (above and to the right of the axes) of the current window.

Edit / Show Clipboard
Displays a window that shows any sound data in the clipboard as a waveform. The clipboard window includes a Play Sound button that plays whatever sound is in the clipboard. The Clipboard window remains visible until you close it by clicking in its close box or selecting the Close command from the File menu.

Edit / Undo Command Z
Undoes the effect of the most recent Cut, Paste, Delete (Clear), or Amplify operation.

File / Close Command W
Closes the current window. If changes have been made to the window contents since the last Save Window command, Canary will ask if you want to save changes.

File / Load Preferences... / Default Preference
Loads the preferences stored in the file "Canary.Prefs" in the Preferences folder, which is inside the System Folder.

File / Load Preferences... / Open...
Loads preferences stored in the file that you specify in response to the standard Open File dialog that appears when you choose this command.

File / New   Command N
Creates a new signal window containing only a waveform pane. A signal copied from another window can be pasted into the new window. To record a new signal with Canary, use the Record command on the Process menu, not the New command.

File / Open...   Command O
Brings up a standard Macintosh Open File dialog box that lets you choose a file to open. The dialog box lists all files that Canary can read. If the file is of a type that can potentially contain more than one kind of data (e.g., a text or MATLAB file, which could contain a waveform, spectrogram, spectrum, or correlation array) and the file has never been saved by Canary, you will be prompted to specify which type of data it contains. Depending on which type of data you indicate, you will be prompted to specify various parameters, such as sampling rate for waveforms, and FFT size for spectrograms.

File / Page Setup...
Brings up the Page Setup dialog box for the printer that is currently selected by the Chooser (apple menu).

File / Print...
Brings up the Print dialog box for the printer that is currently selected by the Chooser (apple menu). When you click OK in the Print dialog box, the contents of the current window are printed according to the options specified in the Page Setup and Print dialog boxes.

File / Quit   Command Q
Closes all windows and quits Canary. Before closing any window whose contents have changed since the last Save, Canary will ask if you want to save changes.

File / Revert To Defaults
Resets all preferences to their “factory default” settings (irrespective of what preferences are stored in the default preference file “Canary.Prefs”).

File / Revert Sound
Restores the sound (waveform, spectrogram, spectrum) in the active window to its last saved form, in effect undoing all editing changes (Cut, Paste, Delete, or Amplify) since the last Save.

File / Save Preferences
Saves all current preference settings to the current preference file. Preference settings include the settings of controls in all dialog box that contain a Save button, and the position onscreen of the command and measurement panels and the track palette. The name of the current preference file can be displayed by choosing About Canary from the apple menu. See also File / Save Preferences As... .

File / Save Preferences As... / Default Preference
Saves all current preference settings to the default preference file, “Canary.Prefs” in the Preferences folder inside the System Folder.

File / Save Preferences As... / Save As...
Saves all current preference settings to the file that you specify in response to the standard Save File dialog that appears when you choose this command.

File / Save Sound As...
This command appears on the menu only when the waveform pane is the active pane. Saves the waveform only into a file of whichever format you specify using the radio buttons in the Save File dialog box that appears. Available formats: Canary, SoundEdit, AIFF, MATLAB, Binary, Text.

File / Save Sound Tracks...
Saves the waveforms of only those tracks of a multi-track signal that are selected on the track palette into the file that you specify in a standard Save File dialog. This command is available only when some track(s) of a multi-track document are selected on the track palette.

File / Save Spectrogram As...
This command appears on the menu only when the spectrogram pane is the active pane. Saves the spectrogram only into a file of whichever format you specify using the radio buttons in the Save File dialog box that appears. Available formats: Canary, MATLAB, Text.

File / Save Spectrum As...
This command appears on the menu only when the spectrum pane is the active pane. Saves the spectrum only into a file of whichever format you specify using the radio buttons in the Save File dialog box that appears. Available formats: Canary, MATLAB, Text.

File / Save Text Report...
Brings up the Save Text Report dialog box to enable you to specify the layout of a text report file to be saved from a data log or from a correlation array. This command is available only when the frontmost window is a data log or correlation array.

File / Save Window Command S
Saves all panes of the current window into one Canary-format file. If the signal has previously been saved in this format, it is re-saved under the same name, overwriting the previously saved version. Otherwise, a standard Save File dialog box appears, allowing you to specify a name and location for the file. After saving a window under a new name, the window is re-named to the specified file name. Note that spectrogram and spectrum panes are saved if they have been calculated, even if they are hidden.

File / Save Window As...
Same as Save Window, but always brings up the Save File dialog box, allowing you to specify a name and location for the file.

Options / Calibration / Copy Calibration
Copies the calibration data of the signal in the current window. Calibration data consist of a paradigm selection, a calibration factor, an impedance value, and dB reference values. Once a calibration has been copied, it can be pasted to a signal in another window.

Options / Calibration / Paste Calibration
Calibrates the signal in the current window according to a set of calibration data (paradigm selection, calibration factor, impedance, and dB reference values) copied from another signal.

Options / Calibration / Paste Calibration To All
Calibrates all presently open signals according to a set of calibration data (paradigm selection, calibration factor, impedance, and dB reference values) copied from another signal.

Options / Calibration / Set Calibration...
Brings up the Signal Calibration dialog box, which lets you select an acoustic or electric calibration paradigm. For acoustic signals, you specify a calibration pressure and characteristic impedance of the medium, or average intensity and characteristic impedance, or a calibration pressure and average intensity for the signal. Calibration pressure can be specified either as RMS, peak, or ceiling pressure. For electric signals, you specify a calibration voltage and line impedance, or average power and line impedance, or a calibration voltage and average power for the signal. Calibration voltage can be specified either as RMS, peak, or ceiling voltage.

Options / Calibration / Set dB Reference...
When there is a selection in the spectrum pane, Set dB Reference... brings up a dialog box that displays the dB level of the highest peak in the selection, and allows you to edit the value. Changing the value that is displayed resets the dB reference value for intensity (or power). The dB reference value for pressure or voltage is also reset to preserve the number of dB difference between dB measurements of RMS pressure and average intensity of any given portion of the signal. If there is no selection in the spectrum, this command is not available.

Options / Display...
Brings up the Display Options dialog box, which lets you set the spectrogram highlight method, specify whether to display track/pane labels, whether to dither spectrograms, whether to hide the measurement and command panels when Canary is switched into the background, and whether to group waveform, spectrogram and spectrum panes for different tracks.

Options / General...
Brings up the General Options dialog box, which allows you to specify whether Canary should automatically present an Open File dialog when the program first starts up, and whether Canary should automatically save preferences when the program quits.

Options / Manual Scaling...
Brings up the Manual Scaling dialog box for the active pane of the current window. The dialog box lets you specify the number of measurement units (e.g., mS, pascals, volts, Hz, dB) per drawing unit (e.g., inch, centimeter, pixel) on the screen. A checkbox in the dialog box enables you to scale all windows presently open to the specified scale.

Options / Measurement Panel...
Brings up the Measurement Panel Configuration dialog box, which lets you specify which measurements are to be displayed for each of the four measurement panels. You select which panel to configure (waveform, spectrogram, spectrum, or correlation) by clicking on the appropriate radio button. The dialog box displays the items that can appear on the selected panel in two scrolling list boxes. The left list shows the items that are currently on the panel; the right list shows the other items that are available. The items shown on the lists are either *parameters* (i.e., characteristics of the data in the selected pane, such as sample rate or filter bandwidth) or *measurements* (i.e., actual data values such as time, amplitude, or frequency), as specified by the radio buttons under the right list box. You can add an item to the measurement panel by selecting its name in the right list, then clicking the Append button between the two lists; you remove an item by selecting its name in the left list and clicking the Remove button. You can insert an item before an item that is already on the

measurement panel by selecting the item in the left list before which you want to insert; then select the item(s) on the right list that you want to insert, and click on the Insert button. You can select more than one item at a time by holding the □ key while clicking on successive items.

Options / Multi-Track...
Allows you to choose stereo or mono playback for multi-track signals. If you select Stereo Playback, the two tracks of a stereo (two-track) signal) are played back through the Macintosh's two stereo output channels, if the Macintosh is equipped for stereo playback. If you choose Mono Playback, or if the Macintosh lacks stereo playback capability, the two channels are mixed for playback.

Options / Speed...
Brings up the Speed Options dialog box, which lets you adjust tradeoffs between computation speed and machine (pause button) responsiveness, and between drawing speed and drawing precision.

Panel / Cursors On/Off
Turns selection cursors on or off in the current window. Equivalent to clicking the command panel's CURSORS button.

Panel / Make Spectrogram...   Command G
Brings up the Spectrogram Options dialog box. Equivalent to clicking on the command panel's SPG button if no spectrogram has been calculated yet or *option*-clicking if a spectrogram already exists. Spectrogram calculation begins when you click OK in the Spectrogram Options dialog box.

Panel / Make Spectrum...   Command K
Brings up the Spectrum Options dialog box. Equivalent to clicking on the command panel's SPK button if no spectrum has been calculated yet or *option*-clicking if a spectrum already exists. Spectrum calculation begins when you click OK in the Spectrum Options dialog box.

Panel / Horiz. Zoom In   Command]
Stretches the horizontal scale of the active pane by a factor or two. Equivalent to clicking on the command panel's Horizontal Stretch button.

Panel / Horiz. Zoom Out   Command [
Compresses the horizontal scale of the active pane by a factor or two. Equivalent to clicking on the command panel's Horizontal Squeeze button.

Panel / Play Selection   Command P
Plays the selected part of the signal at the rate and volume determined by the slider controls on the command panel. Equivalent to clicking the command panel's Play button. For multi-track signals, only those tracks that are selected on the track palette are played; if there are no selections on the track palette, all tracks are played.

Panel / Snap Cursors
Turns selection cursors on in the current window and positions them to outline the current selection. Equivalent to *option*-clicking the command panel's CURSORS button.

Panel / Vert. Zoom In Command =
Stretches the vertical scale of the active pane by a factor or two. Equivalent to clicking on the command panel's Vertical Stretch button.

Panel / Vert. Zoom Out Command -
Compresses the vertical scale of the active pane by a factor or two. Equivalent to clicking on the command panel's Vertical Squeeze button.

**Panel / Zoom to Selection Command **
Rescales the active pane so that the selected region exactly fills the pane. Equivalent to clicking the command panel's Zoom button.

Process / Batch Processes / Correlator
Starts the series of dialogs through which you specify information required to run a batch correlation. Two successive dialog boxes let you specify the two sets of files to be used as input to the correlator. Each set of input files must reside in a single folder. Asterisks ("*") are used as wildcard characters to specify a filename pattern that is used to select the files used for input; the filename pattern "*" matches any file. A third dialog box lets you specify where and in what format to save the correlation array that will be produced. Finally, a fourth dialog box lets you specify the correlation options (filtering, normalization, and complex envelope). Calculation of the correlations begins when you click OK in the Correlator Options dialog box.

Process / Batch Processes / File Conversion...
Starts the series of dialogs through which you specify information required to run a batch file conversion process. The first dialog box lets you specify a set of files to be converted to a particular format. All of the input files must reside in a single folder and all must contain the same type of data (i.e., either sounds (waveforms), spectrograms, spectra, correlation functions, or correlation arrays), although the input files may be saved in different formats (e.g., SoundEdit, AIFF, Canary). Asterisks ("*") are used as wildcard characters to specify a filename pattern that is used to select the files used for input; the filename pattern "*" matches any file. A second dialog box lets you specify a folder and a format in which the converted files will be saved. The output folder must be different from the input folder. Each converted file has the same name as the input file from which it was made. Any existing file with that name in the output folder is overwritten without warning. File conversion begins when you click the Select button in the output selection dialog box.

Process / Batch Processes / Spectra...
Starts the series of dialogs through which you specify information required to run a batch spectrum process. The first dialog box lets you specify a set of files containing waveforms from which spectra will be made. All of the input files must reside in a single folder. Asterisks ("*") are used as wildcard characters to specify a filename pattern that is used to select the files used for input; the filename pattern "*" matches any file. A second dialog box lets you specify a folder and a format in which the spectra will be saved. Each spectrum is saved as a separate file named "SPK of *inFile*" where *inFile* is the name of the waveform file from which the spectrum was made. A third dialog box lets you specify the parameters to be used in calculating the spectra (frame length, frame overlap, FFT size, window function, clipping level, and amplitude scale). Calculation of spectra begins when you click OK in the Spectrum Options dialog box.

Process / Batch Processes / Spectrograms...
Starts the series of dialogs through which you specify information required to run a batch spectrogram process. The first dialog box lets you specify a set of files containing waveforms from

which spectrograms will be made. All of the input files must reside in a single folder. Asterisks ("*") are used as wildcard characters to specify a filename pattern that is used to select the files used for input; the filename pattern "*" matches any file. A second dialog box lets you specify a folder and a format in which the spectrograms will be saved. Each spectrogram is saved as a separate file named "SPG of *inFile*" where *inFile* is the name of the waveform file from which the spectrogram was made. A third dialog box lets you specify the parameters to be used in calculating the spectrograms (frame length, frame overlap, FFT size, window function, clipping level, display style, and amplitude scale). Calculation of spectrograms begins when you click OK in the Spectrogram Options dialog box.

Process / Correlator
Starts the series of dialogs through which you specify information required to perform a single correlation between two waveforms or spectrograms. Two successive dialog boxes each allow you to choose either a saved waveform or spectrogram, or one that is presently open (i.e., one for which a window exists). A third dialog box lets you choose whether to correlate waveforms or spectrograms (if waveforms were selected as input to the correlator) and specify options for filtering, normalization and (for waveform correlations) whether to plot the complex envelope of the correlation function.

Process / Record   Command R
Brings up the Record dialog box. The Record dialog box lets you choose a recording device, a sample rate, a sample size, and an input speed, and lets you specify other options that may be available for your recording device. Depending on the model of the Macintosh, you may be able to select one or two tracks for recording, and adjust the recording gain. You may record a signal either to RAM or to a disk file. When recording to RAM, you may record for a fixed time interval or record continuously while the dialog is open. You may pause and resume recording; while recording to RAM is paused, you may play back the sound that has been recorded so far.

Windows / Command Panel   Command T
Displays or hides the command panel.

Windows / DataLog
Displays the DataLog window.

Windows / Measurement Panel   Command M
Displays or hides the measurement panel.

Windows / Redraw
Completely redraws the current window. Useful if the window graphics become corrupted due to a bug in Canary.

Windows / Track Palette
Displays or hides the track palette. This menu item is available only if the current window is a multi-track window.

Windows / *(window name)*
The windows menu includes the name of each window that is presently open in Canary. Clicking on the name of a window activates that window, making it the frontmost window on the screen.

Dialog fields, checkboxes, and buttons

This section includes an entry for each field, checkbox, and button in all of Canary's dialog boxes. The heading for each entry lists the name of the dialog box and the name of the item, separated by a colon (e.g., "Correlator Options dialog box: Filter checkbox").[1] In some cases, a group of items (e.g., radio buttons that are mutually exclusive with each other) is discussed under a single entry. In a few cases, a single entry is provided for an entire dialog box (e.g., the Print dialog box) and items within the dialog box are not discussed individually. This section does not discuss those dialog boxes that are standard to Macintosh application programs (e.g., the Open and Save dialog boxes). Entries are listed alphabetically by dialog box and item name.

Amplify dialog box: Amplification Factor **field**
Specifies a factor by which the selected part of a signal is amplified. Amplifying by zero silences the selected section of the signal.

Batch Correlation Input dialog box: Open **button**
See Batch Process Input dialog box.

Batch Correlation Input dialog box: File Contains **radio buttons**
The radio buttons specify whether to correlate sounds (waveforms) or spectrograms. If Spectrogram is selected, the input files must already contain spectrograms (they may contain sounds as well); the batch correlator, unlike the interactive correlator, will not make spectrograms as needed if an input file contains only a sound.

Batch Correlation Input dialog box: File name pattern **field**
See Batch Process Input dialog box.

Batch Correlation Input dialog box: Select **button**
See Batch Process Input dialog box.

Batch Correlation Output dialog box
See Batch Process Output dialog box.

Batch File Conversion Input dialog box
See Batch Process Input dialog box.

Batch File Conversion Output dialog box
See Batch Process Output dialog box.

Batch Process Input dialog box: File name pattern **field**
The File name pattern field specifies which files within the selected folder are to be used as input to the batch process. The field uses one or more asterisks ("*") as "wildcard" characters that can stand for any string of characters (including spaces or the null string). Thus the file name pattern "*"

---

[1] In general, the dialog name under which items are listed is either the title displayed at the top of the dialog (e.g., "Spectrogram Options"), or the name of the menu command that brings up the dialog (e.g., "Save Text Report"). A few dialogs that do not display titles have names that describe the dialog's function, but that do not correspond precisely to a single command (e.g., Batch Spectrogram Input dialog).

matches any file name; "*frog" matches any file name whose last four characters are "frog"; "*whale*" matches any file name containing the string "whale".

Batch Process Input dialog box: Open button
Opens the folder whose name is highlighted in the scrolling list box. To select a folder whose name is highlighted to be used as input to a batch process, click the Open button (or double-click on the folder's name), then click on the Select button or double-click on the name of a file in the folder (see discussion of Select button). If the item that is highlighted in the list box is the name of a file (rather than the name of a folder), the Open button is replaced by a Select button.

Batch Process Input dialog box: Select button
Selects the folder whose name appears above the scrolling list box (and whose contents are displayed in the list box) as a folder to be used as input to a batch process, and closes the dialog box. All files in the folder whose names match the pattern specified in the File name pattern field are used as input. Double-clicking on the name of a file in the list box is equivalent to clicking on the Select button. If the item that is highlighted in the list box is the name of a folder (rather than the name of a file), the Select button is replaced by an Open button.

Batch Process Output dialog box: Attach source data to result checkbox
This checkbox appears only for batch spectrogram and spectrum processes. If this option is checked, the resulting file contains the waveform data for the sound, as well as the spectrogram or spectrum. This option is available only for files being saved in Canary format.

Batch Process Output dialog box: Open button
Opens the folder whose name is highlighted in the scrolling list box. To save the output of the batch process in a folder whose name is highlighted, click the Open button (or double-click on the folder's name), then click on the Select button or double-click on the name of a file in the folder (see discussion of Select button). If the item that is highlighted in the list box is the name of a file (rather than the name of a folder), the Open button is replaced by a Select button.

Batch Process Output dialog box: Select button
Selects the folder in which the output of the batch process will be saved and closes the dialog box. Double-clicking on the name of a file in the list box is equivalent to clicking on the Select button. If the item that is highlighted in the list box is the name of a folder (rather than the name of a file), the Select button is replaced by an Open button.

Batch Spectrogram Input dialog box
See Batch Process Input dialog box.

Batch Spectrogram Output dialog box
See Batch Process Output dialog box.

Batch Spectrum Input dialog box
See Batch Process Input dialog box.

Batch Spectrum Output dialog box
See Batch Process Output dialog box.

Correlator Input dialog box: Existing... **button**
Clicking the Existing... button specifies that an open window exists for the signal to be used as input to the correlator and brings up a dialog box listing all open windows. You then select the waveform or spectrogram to be correlated from this list.

Correlator Options dialog box: Complex Envelope **checkbox**
The Complex Envelope checkbox is available only if waveform (as opposed to spectrogram) correlation is selected. If the box is checked, Canary plots the complex envelope of the waveform correlation function. If the box is unchecked, the correlation function itself is plotted. The complex envelope removes high frequency oscillations in the correlation function, which in some cases facilitates visual identification of the correlation peak.

Correlator Options dialog box: Filter **checkbox**
If the Filter box is checked, fields appear that let you specify the upper and lower frequency limits of a bandpass filter that is applied to both signals before they are correlated.

Correlator Options dialog box: Normalize **checkbox**
If Normalize is checked, the sum of the products of the data values from the two signals is divided by the square root of the product of the sums of values from the two signals, as indicated in Equations 7.1 and 7.2. The units in the numerator and denominator cancel and the correlation value is scaled to a dimensionless value. For spectrograms, which contain only non-negative amplitude values, the normalized correlation value is always between 0 and 1. For waveforms, which can contain positive, negative, and zero values, the normalized correlation varies between -1 and 1. If Normalize is off, the correlation is calculated as the sum of the products of the data values from the two signals (i.e., just the numerator of Equation 7.1 or 7.2 is used). Strictly speaking, a non-normalized correlation is in units that are the square of the units of the view being correlated (i.e., pascals or volts for waveforms, $J/m^2/Hz$ or J/Hz for quadratic spectrograms, or dB for logarithmic spectrograms).

Correlator Options dialog box: Pass Band **fields**
The two Pass Band fields, which are available only if the Filter checkbox is checked, let you specify the upper and lower frequency limits of a bandpass filter that is applied to both signals before they are correlated. Pop-up units menus next to each field specify whether the values in the fields are in Hz, kHz, etc.

Correlator Options dialog box: Spectrogram Options **button**
If the inputs to the correlator are waveforms or files that contain waveforms and the Spectrogram radio button is selected, the Options... button brings up the standard Spectrogram Options dialog box, allowing you to specify analysis resolution, grid resolution, and other spectrogram parameters.

Correlator Options dialog box: Waveform / Spectrogram **radio buttons**
If the inputs to the correlator are waveforms or files that contain waveforms, the radio buttons specify whether to correlate the waveforms themselves or to make and correlate spectrograms of the waveforms. If the inputs are spectrograms, the Waveform radio button is disabled.

Display Options dialog box: Dither Spectrograms **checkbox**
If Dither Spectrograms is checked, Canary simulates shades of gray in spectrograms by varying densities of black and white pixels. Canary dithers spectrograms automatically on monitors that

cannot display more than four shades of gray. Use this setting to turn dithering on manually on other monitors.

Display Options dialog box: Group track panes **checkbox**
If Group track panes is checked, Canary displays the waveform, spectrogram, and spectrum panes together for each track of a multi-track document (e.g., shown from top to bottom in the window: *SPK-1, SPG-1, WVF-1, SPK-2, SPG-2, WVF-2*). Otherwise, Canary groups panes by type, so that waveforms of all tracks are adjacent, as are spectrograms and spectra of all tracks (e.g., *SPK-1, SPK-2, SPG-1, SPG-2, WVF-1, WVF-2*).

Display Options dialog box: Hide Msmt Panel / Command Panel **checkbox**
If Hide Msmt Panel / Command Panel is checked, the panels are hidden when Canary is switched into the background.

Display Options dialog box: Show Track/Pane Labels **checkbox**
If Show Track/Pane Labels is checked, Canary displays at the bottom of each pane a label that identifies the pane type, the track number, and an editable title.

Display Options dialog box: Spectrogram Highlight Method **radio buttons**
Two methods are available for highlighting portions of a spectrogram (i.e., sections that have been selected directly or sections that are highlighted because a corresponding time or frequency range has been selected in the waveform or spectrum panes). If Black » Black (the default) is selected, white regions of the spectrogram appear in the highlight color (hue, saturation, and brightness) specified by the Macintosh Color control panel; darker areas are drawn with the same hue and saturation, but less brightness, so that black areas appear black. If Black » Highlight Color is selected, black regions of the spectrogram appear in the highlight color and lighter regions appear with the same hue and brightness as the highlight color, but lesser saturation.

Manual Scaling (Scale Waveform, Spectrogram, Spectrum) dialog box: Amplitude **field**
Specifies the vertical scale factor for a waveform pane in terms of number of pascals (mPa, µPa, etc.) or volts (mV, µV, etc.) per inch, centimeter or pixel.

Manual Scaling (Scale Waveform, Spectrogram, Spectrum) dialog box: Apply to All Windows **checkbox**
Checking the Apply to all windows box applies the specified scale to all windows that contain a pane of the same type as the pane whose scale is displayed in the dialog box.

Manual Scaling (Scale Waveform, Spectrogram, Spectrum) dialog box: Energy Amplitude **field**
Specifies the vertical scale factor for a spectrum pane in terms of number of decibels or watts/m^2/Hz (for acoustic signals) or watts/Hz (for electric signals) per inch, centimeter or pixel. Decibels are used if the spectrum is logarithmic, watts if it is quadratic.

Manual Scaling (Scale Waveform, Spectrogram, Spectrum) dialog box: Frequency **field**
Specifies the horizontal scale factor for a spectrum pane or the vertical scale for a spectrogram pane in terms of number of hertz (kHz, etc.) per inch, centimeter of pixel.

Manual Scaling (Scale Waveform, Spectrogram, Spectrum) dialog box: Time **field**
Specifies the horizontal scale factor for a waveform or spectrogram pane in terms of number of seconds (mS, µS, etc.) per inch, centimeter of pixel.

Measurement Panel Configuration dialog box: Append button
Adds the selected measurements or parameters (in the scrolling list on the right side of the dialog box) to the measurement panel. If an item in the left list is highlighted, the Append button is replaced by an Insert button.

Measurement Panel Configuration dialog box: Insert button
Inserts the selected measurements or parameters (in the scrolling list on the right side of the dialog box) in the measurement panel before the item that is highlighted in the left list. If no item is highlighted in the left list, the Insert button is replaced by an Append button.

Measurement Panel Configuration dialog box: Parameters / Measurements radio buttons
The selected button specifies whether the dialog's list boxes display parameters or measurements for the selected measurement panel. Parameters are characteristics of an entire signal, view, or correlation, and are unaffected by the position of the mouse pointer or the location of a selection. Measurements refer to the location of the mouse or to the selected region in the active pane.

Measurement Panel Configuration dialog box: Remove button
Removes the selected measurements or parameters from the measurement panel (scrolling list on the left side of the dialog box).

Measurement Panel Configuration dialog box:
Sound / Spectrogram / Spectrum / Correlation radio buttons
The selected button determines which of Canary's four measurement panels is to be configured.

Page Setup dialog box
The Page Setup dialog box lets you specify the size of paper to be used, whether to print horizontally or vertically on a page, and various other options, which depend on which printer you are using. See your printer documentation for an explanation of options in the Page Setup dialog box. Page setup options that you set while using Canary apply to all documents until you quit Canary or set new options.

Print dialog box
The Print dialog box displays the type of the printer that is currently selected by the Chooser, (e.g., "LaserWriter"), allows you to specify the number of copies to be printed, and lets you set various options for your printer (e.g., number of copies). See your printer documentation for an explanation of options in the Print dialog box.

Record dialog box: Automatic Gain checkbox
If Automatic Gain is checked, the Macintosh's automatic gain control (which is available only on some Macintosh models) is engaged to prevent strong signals from overloading the recording. **Automatic gain control is strongly not recommended for any scientific work.**

Record dialog box: Cancel button
Clicking the Cancel button closes the dialog box and discards any data recorded to memory. If a signal was being recorded to disk and was not closed, clicking the Cancel button deletes the file from the disk.

Record dialog box: Clear button
Deletes any data that have already been recorded into the RAM recording buffer. The Clear button is available only while recording to memory is paused; the Clear button is replaced by the Close button when recording to disk.

Record dialog box: Close button
When recording to a disk file, the Close button closes the current disk file, saving whatever data have been recorded into the file. The button is available only when recording to disk is paused. The Close button is replaced by the Clear button when recording to memory.

Record dialog box: Device pop-up menu
The box labeled Device displays the name of the currently selected recording device. Pressing the mouse button on this box displays a pop-up menu from which you can select any one of the available recording devices (there may be only one such device). The menu displays all devices for which device driver software is installed on the Macintosh.

Record dialog box: Done button
Clicking the Done button closes the dialog box. If sound data have been recorded to memory, a new signal window appears displaying the waveform of the recorded signal. If sound data have been recorded to a disk file, the data are saved and the file is closed.

Record dialog box: Gain pop-up menu and slider control
The Gain pop-up menu and slider control appear only on those Macintosh models that have adjustable gain control. On Macintosh models that allow independent gain control of two recording tracks, the menu lets you choose whether to adjust the gain for All Tracks, Track 1, or Track 2. The slider lets you adjust the gain between 0.5 and 1.5 in a number of steps that depends on the Macintosh model. Due to limitations in the Macintosh sound input software, the actual signal gain may not be a linear function of the gain values displayed; in other words, a signal recorded with a gain of 1 does not necessarily have twice the amplitude of the same signal recorded with a gain of 0.5.

Record dialog box: Input Speed pop-up menu
The Input Speed menu lets you specify the speed at which sound is being played into the Macintosh relative to the "real" speed of the sound. Thus, if sounds are being played at half normal speed (for example to lower their frequencies to below the Nyquist frequency for the selected sampling rate), select .5 x from the menu. Choices are .125 x, .25 x, .5 x, and 1 x, and Other... . Choosing Other... brings up a dialog in which you can enter any other value.

Record dialog box: Options button
Clicking the Options button brings up a dialog box that allows you to set various options for the selected recording device, if such a dialog box is available. Whether or not there is an Options dialog box depends on the recording device; there is no Options dialog box for the Macintosh built-in recording port.

Record dialog box: Pause button
Clicking the Pause button while recording temporarily suspends recording without closing the dialog box. Data that have already been recorded to memory or disk are preserved. The Pause button replaces the Record button while recording.

Record dialog box: Play **button**
Clicking the Play button plays back the sound already in the recording buffer while recording to memory. The Play button is available only when recording to memory is paused.

Record dialog box: Record **button**
Clicking the Record button initiates recording. During recording, the Record button is replaced by a Pause button. Recording continues until you click the Pause, Done, or Cancel button, or until the specified recording time has been filled.

Record dialog box: Record to Disk **checkbox**
When the Record to disk box is checked, the recorded signal is saved to the file that you specify in the standard Open File dialog that appears when you initially check the box. If the checkbox is left unchecked, the signal is recorded to memory (RAM).

Record dialog box: Recording Time **field**
The Recording Time field is an editable text field into which you can type the maximum number of seconds that a recording should last. Recording automatically stops when the specified recording time is reached.

Record dialog box: Resume **button**
Clicking the Resume button while recording is paused re-starts recording and appends incoming sound data to the data already recorded into memory or a disk file. The Resume replaces the Pause button while recording is paused.

Record dialog box: Sample Rate **pop-up menu**
The box labeled Sample Rate displays the currently selected sample rate. Pressing the mouse button on this box displays a pop-up menu from which you can select any one of the sample rates available for the selected device.

Record dialog box: Sample Size **pop-up menu**
The box labeled Sample Size displays the currently selected sample size (number of bits per sample). Pressing the mouse button on this box displays a pop-up menu from which you can select any one of the sample sizes available for the selected device.

Record dialog box: Stop Play **button**
Stops playback of sound from the recording buffer. The Stop Play button replaces the Play button while sound is being played back.

Record dialog box: Track 1, Track 2 **checkboxes**
When the Track 1 and Track 2 checkboxes are both checked on Macintosh models that are capable of stereo recording, the recording will be made in stereo; if only the Track 1 box is checked, the recording will be mono.

Save Text Report dialog box: Comments **checkbox**
If the file being saved is a data log, comments placed in the log using the Add Comment... or Signal Parameters... commands will be included in the text report only if this box is checked. If the file being saved is a correlation array, this checkbox has no effect on the file. (Correlation text report files cannot include comments.)

Save Text Report dialog: Format radio buttons
The Format radio buttons specify whether numeric entries in a text report file should be written in fixed-point, exponential, or mixed formats. In exponential format, numbers are written as n.ddde±xx, where n and x are any non-zero decimal digits and d is any decimal digit. The number of d's is determined by the dialog's Precision field. The signed two-digit number xx to the right of e indicates a power of ten by which the value should be multiplied. If the Prefix checkbox is checked, the prefix is chosen so that the exponent is between 0 and 2 (see the discussion of the Prefix checkbox). In fixed-point format, the number of digits to the left of the decimal point depends on whether the unit Prefix checkbox is checked. The number of digits to the right of the decimal point is determined by the dialog's Precision field. If unit prefixes are included, the prefix is chosen so that the number of non-zero digits to the left of the decimal point is between one and three. If prefixes are not included, a fixed-point measurement has as many digits as necessary to represent the measurement in unscaled units (e.g., seconds, pascals, Hertz); if the measurement is less than 1, a single zero is written to the left of the decimal point. If Mixed format is selected, Canary uses exponential format if the exponent is less than -4 or greater than the value of the Precision field; otherwise, fixed-point format is used. For example, if Precision is 3 digits and unit prefixes are included, the value 123.4567 mS would be represented in exponential notation as 1.235e+02 mS and in fixed-point notation as 123.457 mS.

Save Text Report dialog box: Headers checkbox
If the file being saved is a data log and this box is checked, the file will include text labels at the top of each column identifying the measurements that appear in the column. If the file is a data log containing mixed measurements, a new header line is placed before each row that contains different measurements from the row above it. If the file being saved is a correlation array and the Headers box is checked, the first row of the file will contain text labels that identify which file from the first input set corresponds to each column.

Save Text Report dialog box: LineLabels checkbox
If the file being saved is a data log, checking the LineLabels checkbox causes the first column of each data line in the file to contain the sequence number for the entry, as shown in the first column of the DataLog window. If the file being saved is a correlation array, checking the LineLabels checkbox causes the first column of each data line in the file to contain a text label identifying which file from the second input set corresponds to each row.

Save Text Report dialog box: Precision field
The Precision field specifies the number of digits to be written to the right of the decimal point for all numeric formats.

Save Text Report dialog box: Prefix checkbox
Allows one of nine standard prefixes to be appended before the unit label on each measurement in a data log. This checkbox is disabled if the Units checkbox is not checked. (See Appendix E for standard metric prefixes and their abbreviations.) The prefix is chosen so that the integer part of a fixed-point measurement is between 1 and 999 (or, equivalently, so that the exponent of an exponential-format measurement is between 0 and 2). Thus, for example, a value of .000999 sec would be written as 999 µS, whereas .001000 sec would be written as 1 mS. Note that if you omit prefixes and specify few digits of precision, small values may appear as zeros. For example, a value of 3.67 mS would appear as 0.00367 S in 5-digit fixed point output, but would be written as 0.00 S if 2-digit fixed-point output was requested.

Save Text Report dialog box: Report Title field and checkbox
If Report Title is checked, the specified text is included on one line at the top of the text report file that is written. You can enter whatever text you want for the report title. Note that the report title is not the name of the file to be saved (although it may be the same); the file name is specified in a standard Save File dialog box that appears after you click OK in the Save Text Report dialog box.

Save Text Report dialog box: Units checkbox
If this box is checked, each individual measurement will be labeled with appropriate units (e.g., S, V, dB, etc.). If the Prefix box is also checked, the unit label may be preceded by a prefix (e.g., µS, mV). If the Units box is unchecked, the column headers (if present) will include the unscaled units (e.g., S or V, but never mS or mV) and individual measurements will appear as unlabeled numbers. If units are used with no prefixes, small values may be written as zeros (see discussion of Prefix checkbox).

Scale Spectrogram dialog box
See Manual Scaling dialog box.

Scale Spectrum dialog box
See Manual Scaling dialog box.

Scale Waveform dialog box
See Manual Scaling dialog box.

Set Calibration dialog box: Air button
When the Sound calibration paradigm is selected, the Air button sets characteristic impedance to 400 rayls, pressure dB reference value to 20 µPa, and the intensity dB reference value to 1 pW/m^2.

Set Calibration dialog box: Apply Default button
Sets the values of all parameters in the dialog to the default values specified in the current preference file.

Set Calibration dialog box: Characteristic Impedance field
The characteristic impedance of the medium in which a sound was recorded, measured in mks rayls. Characteristic impedance is equal to the product of ρ, the medium's density (in kg/m^3) and c, the speed of sound in the medium (in m/sec). The characteristic impedance (ρc) of the medium determines the relationship between RMS pressure (p, in pascals) of a signal and the signal's average intensity (I, in watts/m^2) according to the formula

$$I = \frac{p^2}{\rho c} \tag{12.1}$$

When the Electricity rather than Sound paradigm is selected, Characteristic Impedance is replaced by (electrical) impedance.

Set Calibration dialog box: Fresh Water button
When the Sound calibration paradigm is selected, the Fresh Water button sets characteristic impedance to a typical value for fresh water, 1.4639 Mrayls, pressure dB reference value to 1 µPa (the standard value for underwater measurements), and the intensity dB reference value to .6831

aW/m² (the value that results in equal dB values for pressure level and intensity level of a sound, given the characteristic impedance that is assigned).

Set Calibration dialog box: Impedance **field**
The electrical line impedance, measured in ohms. The impedance (*R*) determines the relationship between RMS voltage (*v*, in pascals) of a signal and the signal's average power (*P,* in watts) according to the formula

$$P = \frac{v^2}{R} \tag{12.2}$$

When the Sound rather than Electricity paradigm is selected, Impedance is replaced by Characteristic Impedance.

Set Calibration dialog box: Intensity **field**
The average intensity of the entire signal. The value is displayed either in watts/m², or in dB relative to the Intensity dB reference value specified elsewhere in the dialog box (depending on which units are selected on the Intensity units pop-up menu). If the radio buttons for Pressure and Characteristic Impedance are selected, Intensity is calculated according to Equation 12.1. Otherwise, the value entered for Intensity is used in conjunction with a value entered for Pressure or Characteristic Impedance to calculate the third quantity. When the Electricity rather than Sound paradigm is selected, Intensity is replaced by Power.

Set Calibration dialog box: Intensity dB re **field**
The reference value, in watts/m², for all dB measurements of intensity (see Appendix C for a discussion of dB reference values) for an acoustic signal. This value is used in calculating dB intensity values for the current signal within the Set Calibration dialog box, and on the measurement panel. When the Electricity rather than Sound paradigm is selected, Intensity dB re is replaced by Power dB re.

Set Calibration dialog box: Power **field**
The average power of the entire signal. The value is displayed either in watts, or in dB relative to the Power dB reference value specified elsewhere in the dialog box (depending on which units are selected on the Power units pop-up menu). If the radio buttons for Voltage and Impedance are selected, Power is calculated according to Equation 12.2. Otherwise, the value entered for Power is used in conjunction with a value entered for Voltage or Impedance to calculate the third quantity. When the Sound rather than Electricity paradigm is selected, Power is replaced by Intensity.

Set Calibration dialog box: Power dB re **field**
The reference value, in watts, for all dB measurements of power (see Appendix C for a discussion of dB reference values) for an electric signal. This value is used in calculating dB power values for the current signal within the Set Calibration dialog box, and on the measurement panel. When the Sound rather than Electricity paradigm is selected, Power dB re is replaced by Intensity dB re.

Set Calibration dialog box: Pressure **field**
The RMS, peak, or ceiling pressure of the entire signal, depending on which selection is made on the pop-up menu (see Set Calibration dialog box: RMS / Peak / Ceiling pop-up menu). The value is displayed either in pascals, or in dB relative to the Pressure dB reference value specified elsewhere in the dialog box (depending on which units are selected on the Pressure units pop-up menu). If the

radio buttons for Intensity and Characteristic Impedance are selected, Pressure is calculated according to Equation 12.1. Otherwise, the value entered for Pressure is used in conjunction with a value entered for Intensity or Characteristic Impedance to calculate the third quantity. When the Electricity rather than Sound paradigm is selected, Pressure is replaced by Voltage.

Set Calibration dialog box: Pressure dB re **field**
The reference value, in pascals, for all dB measurements of pressure (see Appendix C for a discussion of dB reference values) for an acoustic signal. This value is used in calculating dB pressure values for the current signal within the Set Calibration dialog box, and on the measurement panel. When the Electricity rather than Sound paradigm is selected, Intensity dB re is replaced by Power dB re.

Set Calibration dialog box: RMS / Peak / Ceiling **pop-up menu**
Selects whether the value displayed for Pressure or Voltage is the signal's ceiling value, peak value, or RMS value. The ceiling value is the largest value that can be represented with the signal's sample size, i.e., the value represented by a sample in which all bits equal 1. The peak value is the sample with the largest absolute value in the signal. The RMS (root-mean-square) value is the square root of the mean of the squared sample values. RMS, peak, and ceiling pressures are related by the formulas

$$p_{RMS} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}\left[\left(\frac{p_{ceiling}}{2^{s-1}}\right)x_i\right]^2} \tag{12.3}$$

$$p_{peak} = \left(\frac{x_{max}}{2^{s-1}}\right)\cdot p_{ceiling} \tag{12.4}$$

where N is the number of samples in the signal, x_i is the ith sample in the signal, x_{max} is the value of the largest sample in the signal, and s is the number of bits per sample (2^{s-1} is the largest value that can be represented by a sample of size s bits).

Set Calibration dialog box: Sea Water **button**
When the Sound calibration paradigm is selected, the Sea Water button sets characteristic impedance to a typical value for sea water, 1.5458 Mrayls, pressure dB reference value to 1 µPa (the standard value for underwater measurements), and the intensity dB reference value to .6469 aW/m^2 (the value that results in equal dB values for pressure level and intensity level of a sound, given the characteristic impedance that is assigned).

Set Calibration dialog box: Save Default **button**
Saves all of the current settings of the Set Calibration dialog box as the default calibration in the current preference file. These settings can then be recalled whenever the current preference file is loaded by clicking the Apply Default button. (Compare Set Default button.)

Set Calibration dialog box: Set Default **button**
Specifies that the current settings of the Set Calibration dialog box be used as the default settings (until you define different defaults, load a preference file containing different defaults, or quit Canary). These settings can then be recalled by clicking the Apply Default button. Set Default differs from Save Default in that the defaults are temporary and are not saved after you quit Canary (i.e.,

they are not saved to a preference file, unless you later explicitly save preferences from the File menu). (Compare Save Default button.)

Set Calibration dialog box: Sound / Electricity **pop-up menu**
Chooses the calibration paradigm to be used with the current signal. The choice of calibration paradigm controls the units chosen for displaying amplitude quantities.

Set Calibration dialog box: Voltage **field**
The RMS, peak, or ceiling voltage of the entire signal, depending on which selection is made on the pop-up menu (see Set Calibration dialog box: RMS / Peak / Ceiling pop-up menu). The value is displayed either in volts, or in dB relative to the Voltage dB reference value specified elsewhere in the dialog box (depending on which units are selected on the Voltage units pop-up menu). If the radio buttons for Power and Impedance are selected, Voltage is calculated according to Equation 12.2. Otherwise, the value entered for Voltage is used in conjunction with a value entered for Power or Impedance to calculate the third quantity. When the Sound rather than Electricity paradigm is selected, Voltage is replaced by Pressure.

Set Calibration dialog box: Voltage dB re **field**
The reference value, in volts, for all dB measurements of voltage (see Appendix C for a discussion of dB reference values) for an acoustic signal. This value is used in calculating dB voltage values for the current signal within the Set Calibration dialog box, and on the measurement panel. When the Sound rather than Electricity paradigm is selected, Voltage dB re is replaced by Pressure dB re.

ignal Calibration dialog box
See Set Calibration dialog box.

Sound Acquisition / Recording dialog box
See Record dialog box.

Spectrogram / Spectrum Options dialog box: Amplitude **radio buttons (**Logarithmic / Quadratic**)**
The Amplitude radio buttons select either a logarithmic or quadratic energy amplitude axis. In a logarithmic spectrum, intensity or power amplitudes are plotted in decibels relative to the reference value specified in the signal's calibration information; a quadratic spectrum plots intensity or power amplitudes in $J/m^2/Hz$ or J/Hz, depending on whether the acoustic or electric calibration paradigm is selected. For a spectrogram, the choice of amplitude axis determines whether the grayscale values are proportional to the intensity (power) amplitude at each point or the log of the intensity (power) amplitude.

Spectrogram / Spectrum Options dialog box: Clipping Level **field**
The Clipping Level parameter specifies a "noise floor", below which any amplitude value is ignored. In a spectrum, any frequency component less than the clipping level is displayed as being equal to the clipping level. In a spectrogram, any value below the clipping level is displayed as white. The Clipping Level parameter is in units of intensity dB (for acoustic signals) or power dB (for electric signals) relative to the reference intensity or power specified in the signal's calibration information.

Spectrogram Options dialog box: Display Style **radio buttons (**Boxy / Smooth**)**
In a boxy spectrogram, each actual data point on the spectrogram grid is represented by a rectangular gray box. The width and height of the boxes depend on the grid resolution in the time and frequency dimensions respectively. In a smooth spectrogram, the darkness of each individual screen pixel is determined by bilinear interpolation between the amplitude (or log amplitude) values calculated at the grid points. Each time the spectrogram is resized, the grayscale values for individual screen pixels are recalculated. Smoothed spectrograms take longer to redraw than boxy spectrograms.

Spectrogram / Spectrum Options dialog box: FFT Size **pop-up menu**
FFT (Fast Fourier Transform) size determines the number and size of frequency bins in the spectrogram or spectrum. See the discussion of the Frequency pop-up menu for further details.

Spectrogram / Spectrum Options dialog box: Filter Bandwidth **pop-up menu**
Specifies the bandwidth (in Hz) of the individual analysis filters in the filterbank simulated by the short-time Fourier transform (STFT) with the selected frame length (see Appendix B). The filter bandwidth is also influenced by the choice of window function. The Filter Bandwidth and Frame Length parameters in the dialog box are coupled to each other: changing one of them automatically changes the other.

Spectrogram / Spectrum Options dialog box: Frame Length **pop-up menu**
Frame length is the duration (measured in mS or points) of each successive interval over which the spectral composition of a signal is estimated. The Filter Bandwidth and Frame Length parameters in the dialog box are coupled to each other: changing one of them automatically changes the other. The frame length (in points) cannot be larger than the FFT size.

Spectrogram / Spectrum Options dialog box: Frequency **pop-up menu**
Controls the frequency resolution of the spectrogram or spectrum. For a boxy spectrogram, specifies the height (in Hz) of each box. For a spectrum, specifies the spacing between points along the frequency axis. The frequency resolution of a spectrogram or spectrum depends on the sampling rate (which is fixed for a given digitized signal) and the FFT size. The relationship is

frequency resolution = (sampling frequency) / FFT size

where frequency resolution and sampling frequency are measured in Hz, and FFT size is measured in points. The Frequency and FFT Size parameters in the dialog box are coupled to each other; changing one changes the other.

Spectrogram / Spectrum Options dialog box: New **button**
Creates a new named options set, initially named Untitled. You can change the name of the options set by editing its name.

Spectrogram / Spectrum Options dialog box: Options name **field**
The Options name field displays the name of the currently selected set of spectrogram options. You can edit the name of any named options set except for the default settings.

Spectrogram / Spectrum Options dialog box: Options name **pop-up menu**
The pop-up menu next to the Options name field lists the names of all currently defined options sets, and lets you select one of them. Selecting an item from the menu sets all of the dialog parameters to the values specified for that set of options. Named options sets are saved in preference files, so loading a new preference file may change the list of options sets available.

Spectrogram / Spectrum Options dialog box: Save **button**
Clicking the Save button immediately saves all of the named options sets that are presently defined into the current preference file.

Spectrogram / Spectrum Options dialog box: Overlap **pop-up menu**
Specifies the overlap (as percent of frame length or in number of points) between successive frames within which the spectral composition of a signal is estimated. The overlap displayed in the dialog box changes if either the frame length or the time grid resolution are changed. Changing the frame overlap changes the time grid resolution. Overlap, frame length, and time grid resolution are related according to the formula

*time grid resolution = frame length * (100% - overlap%)* .

(See also the discussion of the Time pop-up menu.)

Spectrogram / Spectrum Options dialog box: Remove **button**
Removes the currently displayed set of options from the list of named options sets. Removing a set of options does not affect the current preference file unless you then click the Save button (or later select Save Preferences from the File menu).

Spectrogram / Spectrum Options dialog box: Revert **button**
Clicking the Revert button restores the settings in the dialog box to those specified for the options set whose name is currently displayed in the Options name field.

Spectrogram / Spectrum Options dialog box: Time **pop-up menu**
For a spectrogram, controls the time resolution of the spectrogram grid, i.e., the width (in mS or points) of each box in a boxy spectrogram. For multi-frame spectra, the Time parameter controls the spacing between points in time at which spectral densities are calculated (and averaged to obtain the final displayed spectrum). For single-frame spectra, the Time parameter is irrelevant. The Time and Overlap parameters are coupled to each other: changing one changes the other (see discussion of Overlap parameter).

Spectrogram / Spectrum Options dialog box: Window Function **pop-up menu**
Specifies a window function by which each frame of data is multiplied before the frame's spectrum is calculated. Window functions are used to reduce the magnitude of spurious frequencies that appear as sidelobes flanking each analysis frequency in a spectrum. Canary provides five different window functions, each characterized by the magnitude of the sidelobes relative to the center lobe. In a spectrogram, differences among windows in sidelobe rejection result in different amounts of vertical gray "fringing" around black or very dark areas. For any given frame length, different window functions result in different filter bandwidths. In terms of a spectrogram, this means that the traces drawn using different window functions will be of slightly different vertical thicknesses for the same signal.

Spectrum Options dialog box: Number of frames
Displays the number of frames selected for a spectrum analysis. If no part of the waveform is selected, one frame is analyzed, beginning at the blinking insertion point in the waveform. If part of the waveform is selected, the number of frames that a given selection encompasses depends on both the frame length and the frame overlap.

Speed Options dialog box: Pause Button Responsiveness **radio buttons**
Depending on whether Pause Button Faster, Computations Faster, or No Pause Button is selected, Canary allows more or less time for tasks other than background computation. This choice affects the machine's responsiveness to all user interaction (not just clicking the Pause button) and the speed of any other applications that are running at the same time as Canary.

Speed Options dialog box: Waveform Drawing **radio buttons**
If Best waveform drawing is selected, every point in the digitized signal is used to draw the waveform. If Faster or Fastest is selected, Canary skips some points, which speeds up the drawing process at the expense of precision. Although very sharp, brief peaks in the waveform may be missed in Fastest mode, in practice the waveforms drawn in Best and Fastest modes are often indistinguishable from each other.

Graphic controls and mouse icons

This section contains an entry for each graphic control item and mouse icon used by Canary. Each entry shows the control item or icon and identifies it by name. Command panel buttons are shown first, followed by mouse icons, and graphic elements that appear within signal windows. This section does *not* include entries for graphic controls that are standard throughout the Macintosh environment (e.g., window scrolling controls).

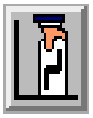

Command Panel: Stretch Vertical Axis (Vertical Zoom In)
Clicking this button stretches the vertical axis of the active pane by a factor of 2. Any other pane that shares the dimension (time or frequency) affected is also rescaled correspondingly.

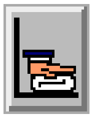

Command Panel: Squeeze Vertical Axis (Vert. Zoom Out)
Clicking this button compresses the vertical axis of the active pane by a factor of 2. Any other pane that shares the dimension (time or frequency) affected is also rescaled correspondingly.

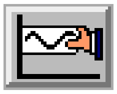

Command Panel: Stretch Horizontal Axis (Horiz. Zoom In)
Clicking this button stretches the horizontal axis of the active pane by a factor of 2. Any other pane that shares the dimension (time or frequency) affected is also rescaled correspondingly.

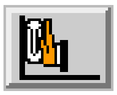

Command Panel: Squeeze Horizontal Axis (Horiz. Zoom Out)
Clicking this button compresses the horizontal axis of the active pane by a factor of 2. Any other pane that shares the dimension (time or frequency) affected is also rescaled correspondingly.

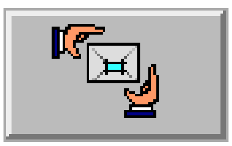

Command Panel: Zoom To Selection
Clicking this button rescales the axes of the active pane so that the selected region fills the pane. Other panes that share dimensions with the active pane are also rescaled accordingly.

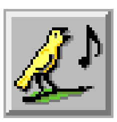

Command Panel: Play Selection
Plays the selected portion of a sound at the rate and volume determined by the slider controls on the command panel. If there is no selection, plays the entire sound.

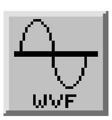

Command Panel: Show/Hide Waveform
Clicking this button alternately displays and hides the waveform pane. The waveform pane cannot be hidden if it is the only pane currently displayed.

Command Panel: Show/Hide Spectrogram
If no spectrogram exists for the current window, clicking this button brings up the Spectrogram Options dialog box. If a spectrogram exists but is not displayed, clicking the button shows the spectrogram. Holding the *option* key down while clicking the button forces the Spectrogram Options dialog box to appear, even if a spectrogram already exists. Clicking the button when a spectrogram is displayed hides the spectrogram. The spectrogram pane cannot be hidden if it is the only pane currently displayed.

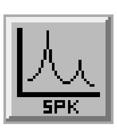

Command Panel: Show/Hide Spectrum
If no spectrum exists for the current window, clicking this button brings up the Spectrum Options dialog box. If a spectrum exists but is not displayed, clicking the button shows the spectrum. Holding the *option* key down while clicking the button forces the Spectrum Options dialog box to appear, even if a

spectrum already exists. Clicking the button when a spectrum is displayed hides the spectrum. The spectrum pane cannot be hidden if it is the only pane currently displayed.

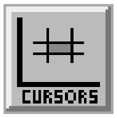

Command Panel: Cursors On/Off
Clicking this button toggles selection cursors on or off in the active window.

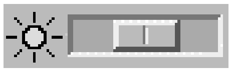

Command Panel: Spectrogram Brightness
Sliding the brightness control to the right makes the entire spectrogram image lighter; sliding to the left darkens the image. Brightness can also be adjusted by typing a number between 0 and 100 into the box to the right of the slider control.

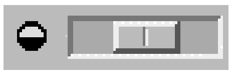

Command Panel: Spectrogram Contrast
Sliding the contrast control to the right increases the contrast in the entire spectrogram (i.e., reduces the range of gray values between white and black). Sliding to the left reduces contrast. Contrast can also be adjusted by typing a number between 0 and 100 into the box to the right of the slider control.

Signal Window: Select / Log data (mouse pointer)
The mouse pointer adopts the crosshair shape whenever it is over the plot area of the active pane. Dragging the mouse or clicking and shift-clicking select part of the view shown in the pane. Double-clicking selects the entire view (including parts that are beyond the edges of the pane). Command-clicking copies measurements to the data log. Point measurement cells on the measurement panel are updated continuously while the mouse pointer appears as a crosshair.

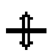

Signal Window: Change pane size (mouse pointer)
The mouse pointer adopts this shape when placed on the double horizontal line between two panes in a signal window; you can then use the mouse to drag the line up or down to change the relative sizes of the two panes.

Correlation Window: Change pane size (mouse pointer)
The mouse pointer adopts this shape when placed on the double vertical line between the two panes in a correlation window; you can then use the mouse to drag the line left or right to change the relative sizes of the two panes.

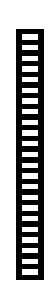

Signal Window: Active Pane Indicator
The hatched bar appears at the left edge of the active pane in the active (frontmost) signal window. The active pane is the one to which the window's scrollbars and the measurement panel apply. You can make a different pane active by clicking on it.

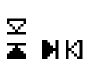

Signal Window: cursor tags
When a signal window's cursors are on, cursor tags appear along the horizontal axis of the waveform pane and along both axes of the spectrogram and spectrum panes. Numbers below or to the left of the cursor tags indicate the exact position of the cursor, in whatever units label the axis. In the active pane, the active cursor in each direction (horizontal and vertical) is filled (black), the inactive cursors are open (white). The active cursor can be moved

using the horizontal or vertical arrow keys. An inactive cursor is activated by clicking on it. Cursors in inactive panes are always open.

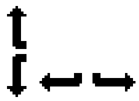

Signal window: cursor grab tags
When a signal window's cursors are on and a cursor is beyond the edge of a pane, a cursor grab tag appear at the end of the axis, pointing in the direction of the cursor. Clicking on a grab tag brings the cursor to the end of the axis and makes it the active cursor.

Track palette
The track palette is displayed only when the active window is a multi-track window. Clicking on a Track N button toggles the track between displayed and hidden. Clicking on a selection button (to the right of the Track N buttons) toggles a track between being selected and un-selected; when the track is selected, its button displays a black circle. With signals containing more than two tracks you can mark multiple tracks to be displayed or hidden by option-clicking the Track N buttons; clicking the Apply button then shows or hides all of the marked tracks at once.

File icons

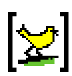

Any file saved by Canary in MATLAB format. The file may contain a sound, a spectrogram, a spectrum, a correlation function, a correlation array, or a data log.

Sound (waveform) file saved by Canary in SoundEdit format.

Canary window file. The file may contain any one, two, or three views (sound, spectrogram, or spectrum) of a signal.

Sound (waveform) file saved by Canary in binary format.

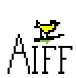

Sound (waveform) file saved by Canary in AIFF (Audio Interchange File Format).

File saved by Canary in Text format. The file may contain a sound, a spectrogram, a spectrum, a correlation function, a correlation array, or a data log.

Appendix A Digital Representation of Sound

About this appendix

This appendix provides a brief explanation of how sound is represented digitally. An understanding of the basic principles introduced here will be helpful in using Canary.

Digital sampling

Before a continuous, time-varying signal such as sound can be manipulated or analyzed with a digital computer, the signal must be *acquired* or *digitized* by an *analog-to-digital (A/D) converter*.[1] The A/D converter repeatedly measures or samples the instantaneous voltage amplitude of an input signal at a particular sampling rate, typically thousands or tens of thousands of times per second (Figure A.1). The digital representation of a signal created by the converter thus consists of a sequence of numeric values representing the amplitude of the original waveform at discrete, evenly spaced points in time.

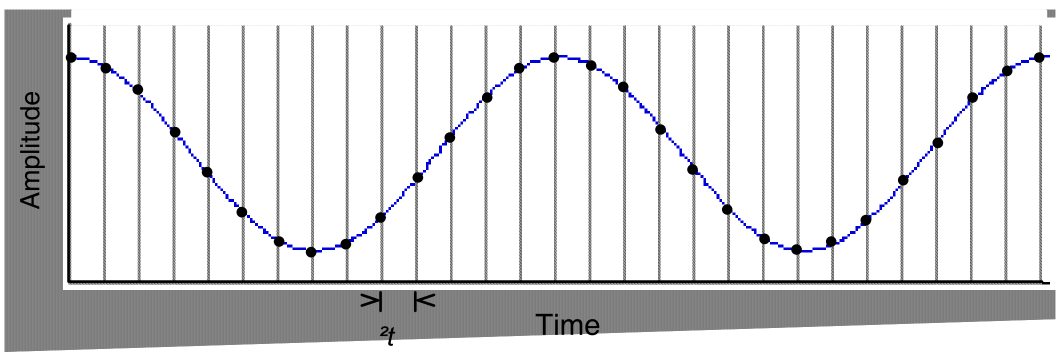

Figure A.1. Sampling to create digital representation of a pure tone signal. Measurements of the instantaneous amplitude of the signal are taken at a sampling rate of $1/\Delta t$. The resulting sequence of amplitude values is the digitized signal.

The precision with which the digitized signal represents the continuous signal depends on two parameters of the digitizing process: the rate at which amplitude measurements are made (the *sampling rate* or *sampling frequency*), and the number of bits used to represent each amplitude measurement (the *sample size*).

---

[1]Recent Macintosh models (including Quadra, Performa, Centris, PowerMac, and Powerbook models) come equipped with a built-in A/D converter, which takes its input from the Mac's microphone jack. A/D converters can also be purchased from third-party manufacturers (e.g., MacRecorder from MacroMind Paracomp) for use with these or other Mac models.

Sampling rate

Canary's Sound Recording dialog enables you to choose the sampling rate at which a signal is to be digitized. The choices available are determined by the A/D converter hardware and the program (called a *device driver*) that controls the converter; most converters have two or more sampling rates available.[1] The highest frequency available with the Macintosh built-in A/D converter depends on which model of Macintosh you are using. Commercial digital audio applications use higher sampling rates (44.1 kHz for audio compact discs, 48 kHz for digital audio tape). Once a signal is digitized, its sampling rate is fixed.[2]

The more frequently a signal is sampled, the more precisely the digitized signal represents temporal changes in the amplitude of the original signal. The sampling rate that is required to make an acceptable representation of a waveform depends on how rapidly the signal amplitude changes (i.e., on the signal's frequency). More specifically, the sampling rate must be more than twice as high as the highest frequency contained in the signal. Otherwise, the digitized signal will have frequencies represented in it that were not actually present in the original at all. This appearance of phantom frequencies as an artifact of inadequate sampling rate is called *aliasing* (Figure A.2). The highest frequency that can be represented in a digitized signal without aliasing is called the *Nyquist frequency,* which is half the frequency at which the signal was digitized. The highest frequency in a spectrogram or spectrum calculated by Canary is always the Nyquist frequency of the digitized signal. If the only energy above the Nyquist frequency in the analog signal is in the form of low-level, broadband noise, the effect of aliasing is to increase the noise in the digitized signal. However, if the spectrum of the analog signal contains any peaks above the Nyquist frequency, the spectrum of the digitized signal will contain spurious peaks below the Nyquist frequency as a result of aliasing. The usual way to guard against aliasing is to pass the analog signal through a low-pass filter (called an anti-aliasing filter) before digitizing it, to remove any energy at frequencies greater than the Nyquist frequency. (If the original signal contains no energy at frequencies above the Nyquist frequency or if it contains only low-level broadband noise, this step is unnecessary.)

---

[1]The device driver for the built-in Macintosh sound input port is included as part of the system software. Device drivers for other A/D converters are supplied with the converter, usually as a file that must be placed in the System folder.

[2]Some digital signal processing programs (e.g., MATLAB's Signal Processing Toolbox) enable you to resample a digitized signal at a lower rate than the original sampling frequency by discarding some samples, or to increase the nominal sampling rate by interpolating samples.

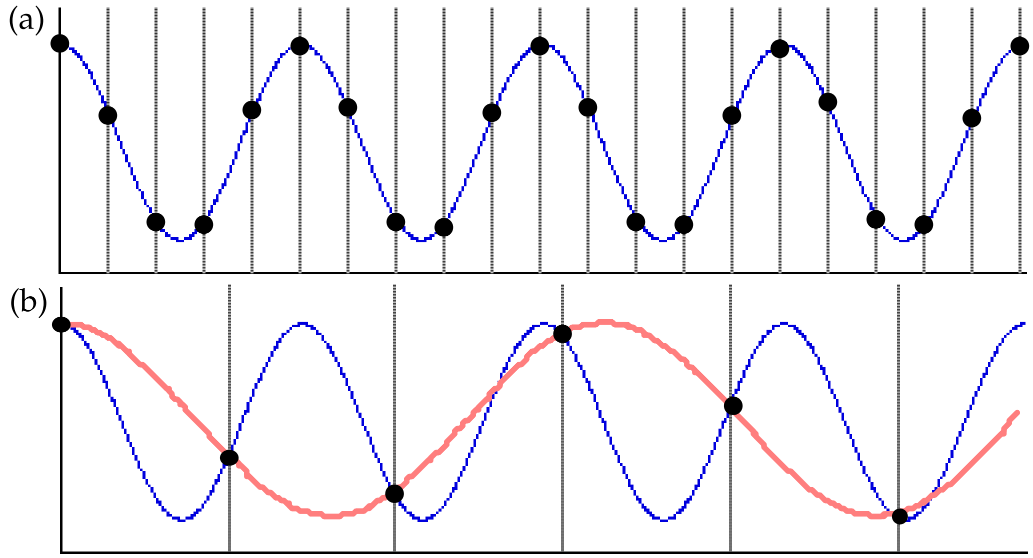

Figure A.2. Aliasing as a result of inadequate sample rate. The same analog waveform is shown in both figures. Vertical lines indicate times at which samples are taken. **(a)** Sampling frequency approximately five times the signal frequency. **(b)** Sampling frequency approximately 1.5 times the signal frequency. The resulting digitized signal (gray waveform) exhibits aliasing: it portrays a waveform of lower frequency than the original analog signal.

In order to interpret a sequence of numbers as representing a time-varying signal, one needs to know the sampling rate. Thus, when a digitized signal is saved in a file format that is designed for saving sound information, information about the sampling rate is usually saved along with the actual data points comprising the signal. If you try to open a file with Canary that contains sound data, but no information about the sampling rate, Canary asks you for the sampling rate.

Sample size (amplitude resolution)

The precision with which a sample represents the actual amplitude of the waveform at the instant the sample is taken depends on the *sample size* or number of bits used in the binary representation of the amplitude value. Some A/D converters can take samples of one size only; others allow you to choose (usually through software) between two or more sample sizes. Some Macintosh models provide only 8-bit sampling capability; others allow you to choose between 8-bit and 16-bit samples. An 8-bit sample can resolve 256 (=2^8) different amplitude values; a 16-bit converter can resolve 65,536 (=2^{16}) values. Sound recorded on audio CDs is stored as 16-bit samples. When a sample is taken, the actual value is rounded to the nearest value that can be represented by the number of bits in a sample.

Since the actual analog value of signal amplitude at the time of a sample is usually not exactly equal to one of the discrete values that can be represented exactly by a sample, there is some error inherent in the process of digitizing (Figure A.3), which results in *quantization noise* in the digitized signal. The more bits used for each sample, the less quantization noise is contained in the digitized signal.

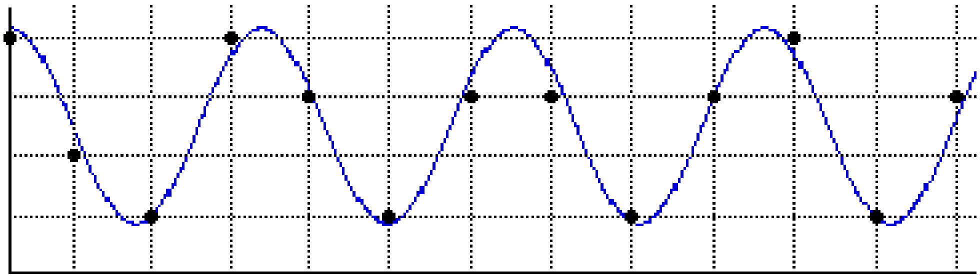

Figure A.3. Digitizing error with a hypothetical 2-bit sample size. 2-bit samples can represent only four different amplitude levels. At each sample time (vertical lines), the actual amplitude levels are rounded to the nearest value that can be represented by a 2-bit sample (horizontal lines). The amplitude values stored for most samples (black dots) are slightly different from the true amplitude level of the signal at the time the sample was taken.

The sample size determines the maximum dynamic range of a digitized sound. Dynamic range is the ratio between the highest amplitude and the lowest non-zero amplitude in a signal, usually expressed in decibels. The dynamic range of a digitized sound is 6 dB/bit.[1]

Storage requirements

The increased frequency bandwidth obtainable with higher sampling rates and the increased dynamic range obtainable with larger samples both come at the expense of the amount of memory required to store a digitized signal. The minimum amount of storage (in bytes) required for a digitized signal is the product of the sample rate (in samples/sec), the sample size (in bytes; one byte equals 8 bits), and the signal duration (seconds). Thus, a 5-second signal sampled at 22.3 kHz with 8-bit precision requires about 110 Kbytes of storage. The actual amount of storage required for a signal may exceed this minimum, depending on the format in which the samples are stored. For reasons of programming efficiency, Canary always uses a 32-bit format for its internal representation of a digitized signal, irrespective of the sample size with which the signal was digitized. The amount of memory (RAM) required by Canary to store a signal while working with it is thus four times the minimal requirement for a signal sampled with 8-bit resolution. When a signal is saved as a disk file, its storage requirements depend on the file format used and may be less than the amount of storage required by Canary's internal representation. For example, if a sound is recorded with 8-bit precision and saved in SoundEdit or AIFF format, it will be saved with 8-bit precision, (even though Canary uses 32 bits while working with the sound).

No matter what file format is use, digitized sound files take up a lot of storage space. If you plan on storing many or long digitized signals and need to save storage space, you might consider using a data compression program such as Compact Pro (shareware, from Cyclos, PO Box 31417,

[1]The dynamic range of a signal in decibels is equal to 20 log(V_{max}/V_{min}), where V_{max} and V_{min} are the maximum and minimum voltage in the signal. For a digitized signal, $V_{max}/V_{min} = 2^n$, where n is the number of bits per sample. Since $\log(2^n) = n * .3$, the dynamic range of a digitized signal is 6 dB per bit.

San Francisco, CA 94131-0417 USA; available from many users' groups and bulletin boards), StuffIt (Aladdin Systems, Inc.), or Disk Doubler (Salient Software, Inc.). These programs can often compress digitized sound files by 50% or more (the amount of compression will vary from file to file). Compressed files must be expanded before they can be used with Canary.

Appendix B A Biologist's Introduction to Spectrum Analysis

About this appendix

This appendix provides some conceptual background for making and interpreting spectrograms and spectra with Canary. It introduces the short-time Fourier transform (STFT), the mathematical technique used by Canary for making spectrograms. We treat the STFT here as a black box, but one that has controls on its outside that affect its operation in important ways. One aim of this appendix is to convey enough qualitative understanding of the behavior of this box to allow intelligent use of its controls. Specific details of the controls are covered in Chapter 3. A second aim of this appendix is to explain some of the limitations and tradeoffs intrinsic to spectrum analysis of time-varying signals. More rigorous mathematical treatments of spectral analysis, at several levels of sophistication, can be found in the references listed at the end of the appendix.

Several approaches are available for explaining the fundamentals of digital spectrum analysis. The approach taken in this appendix is geared specifically to spectrum analysis with Canary; thus some of the terms and concepts used here (e.g., "analysis resolution" and "grid resolution") may not appear in other, more general discussions of spectrum analysis, such as those listed at the end of the appendix.

The discussions in this appendix assume a basic understanding of how sound is recorded and represented digitally. If you are not already acquainted with concepts such as sampling rate and sample size, you should read Appendix A before proceeding.

Time-domain and frequency-domain representations of sound

Any acoustic signal can be graphically or mathematically depicted in either of two forms, called the *time-domain* and *frequency-domain* representations. In the time domain, the amplitude of a signal is represented as a function of time. Figure B.1a shows the time-domain representation of the simplest type of acoustic signal, a pure tone. Such a signal is called a *sinusoid* because its amplitude is a sine function of time, characterized by some frequency, which is measured in cycles per second, or Hertz (Hz). (In terms of real-world physical quantities, the amplitude may represent a measurement such as the pressure exerted by vibrating air or water molecules, or a voltage at some point in an electric circuit.) In the frequency domain, the amplitude of a signal is represented as a function of frequency. The frequency-domain representation of a pure tone is a vertical line (Figure B.1b).

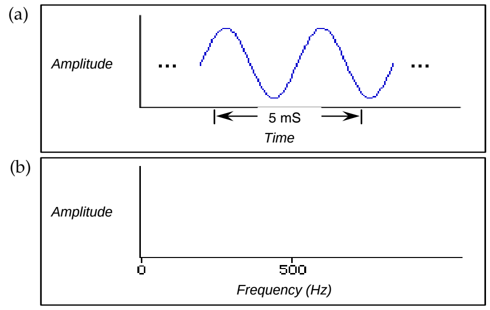

Figure B.1. Time-domain and frequency-domain representations of an infinitely long pure sinusoidal signal with a frequency of 500 Hz. **(a)** Time domain. **(b)** Frequency domain.

Any sound, no matter how complex, can be represented as a sum of pure tones (sinusoidal components). Each tone in the series has a particular amplitude, relative to the others, and a particular phase relationship (i.e., it may be shifted in time relative to the other components). The frequency composition of complex signals is usually not apparent from inspection of the time-domain representation. Spectrum analysis is the process of converting the time-domain representation of a signal (which is the representation directly produced by most measuring and recording devices) to a frequency-domain representation that shows how different frequency components contribute to the sound.

The complete frequency-domain representation of a signal consists of two parts. The *magnitude spectrum* (Figure B.2b) contains information about the magnitude of each frequency component in the entire signal. The *phase spectrum* (not shown) contains information about the phase or timing relationships among the frequency components, but in a form that is not easily interpreted. Since the phase spectrum is rarely of practical use in most bioacoustic work and is not provided by Canary, it is not discussed further here. Henceforth, unless otherwise noted, we use the term "spectrum" to refer to the magnitude spectrum alone.

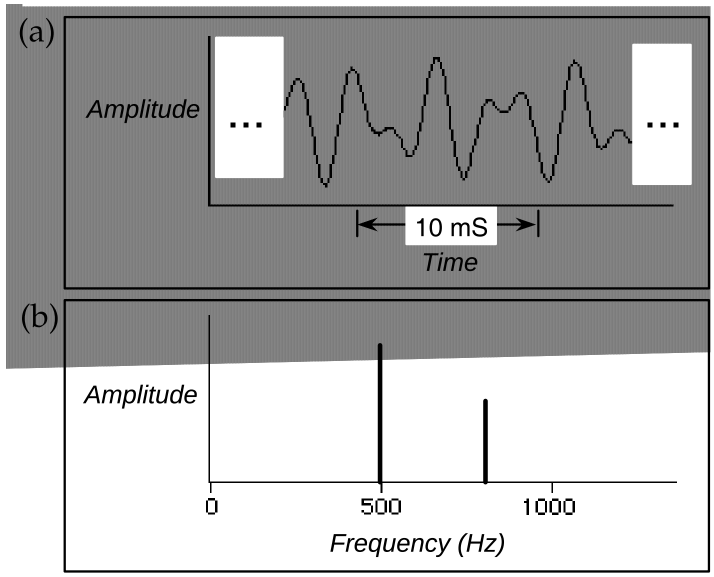

Figure B.2. Time-domain and frequency-domain representations of an infinitely long sound consisting of two tones, with frequencies of 490 Hz and 800 Hz. **(a)** Time domain. **(b)** Frequency domain.

The *Fourier transform* is a mathematical function that converts the time-domain form of a signal (which is the representation directly produced by most measuring and recording devices) to a frequency-domain representation, or spectrum. When the signal and spectrum are represented as a sequence of discrete digital samples, a version of the Fourier transform called the *discrete Fourier transform (DFT)* is used. The input to the DFT is a finite sequence of values— the amplitude values of the signal— sampled (digitized) at regular intervals. The output is a sequence of values specifying the amplitudes of a sequence of discrete frequency components, evenly spaced from zero Hz to half the sampling frequency (Figure B.3). Canary implements the DFT using an algorithm known as the *fast Fourier transform (FFT).*

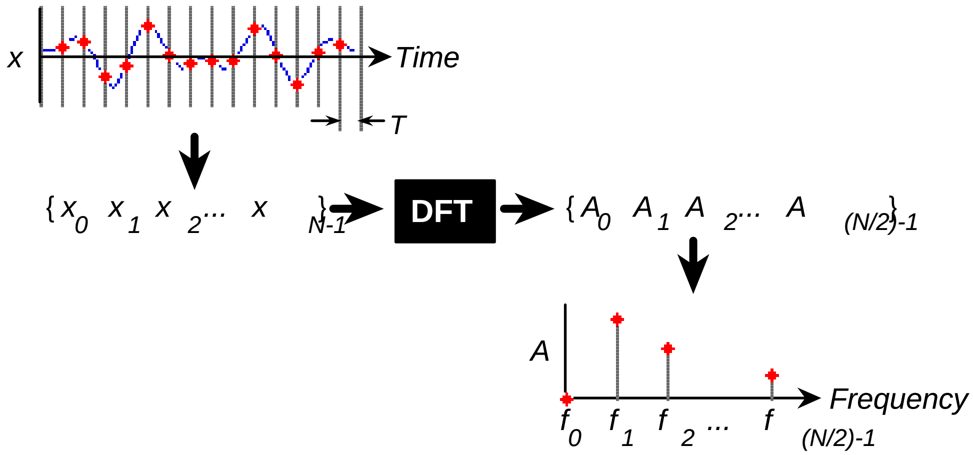

Figure B.3. Schematic representation of the discrete Fourier transform (DFT) as a black box. The input to the DFT is a sequence of digitized amplitude values (x_0, x_1, x_2, ... x_{N-1}) at *N* discrete points in time. The output is a sequence of amplitude values (A_0, A_1, A_2, ... $A_{(N/2)-1}$) at *N/2* discrete frequencies. The highest frequency, $f_{(N/2)-1}$, is equal to half the sampling rate (= $^1/_{(2T)}$, where T is the sampling period, as shown in the figure). The output can be plotted as a magnitude spectrum.

In practice, a spectrum is always made over some finite time interval. This interval may encompass the full length of a signal, or it may consist of some shorter part of a signal.

Spectral analysis of time-varying signals: spectrograms and STFT analysis

An individual spectrum provides no information about temporal changes in frequency composition during the interval over which the spectrum is made. To see how the frequency composition of a signal changes over time, we can examine a sound *spectrogram*. The spectrograms produced by Canary plot frequency on the vertical axis versus time on the horizontal; the amplitude of a given frequency component at a given time is represented by a grayscale value between white and black (Figure B.4).[1] Spectrograms are produced by a procedure known as the *short-time Fourier transform (STFT)*.

[1]There are other ways of representing amplitude, such as by color, or by using contour lines, but grayscale spectrograms are most widely used by biologists.

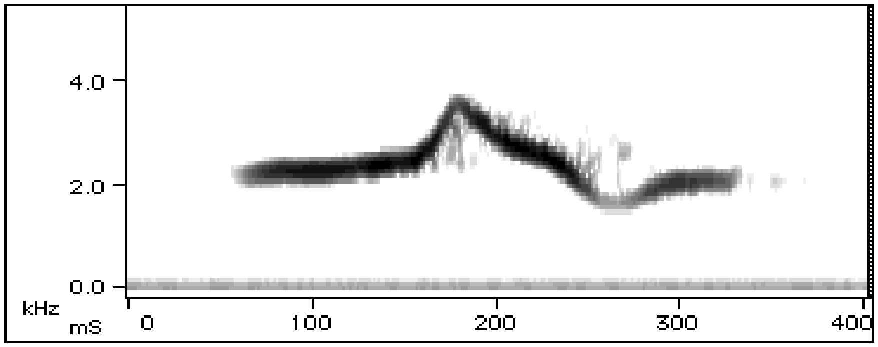

Figure B.4. Sound spectrogram of one syllable from song of a rose-breasted grosbeak, digitized at 22.3 kHz.

There are two convenient ways to describe the operation of the STFT. One approach is to think of the STFT as dividing the entire signal into successive short time intervals or *frames* (which may overlap each other in time). Each frame is used as the input to a DFT, generating a series of spectra (one for each frame) that approximate the "instantaneous" spectrum of the signal at successive moments in time. To display a spectrogram, the spectra of successive frames are plotted side by side with frequency running vertically and amplitude represented by grayscale values (Figure B.5a). We call this the "spectral slice" model of STFT analysis. A given STFT can be characterized by its *frame length*, usually expressed as the number of digitized amplitude samples that are processed to create each individual spectrum.

An alternative description considers the STFT as equivalent in function to a bank of bandpass filters, each centered at a slightly different analysis frequency. The output amplitude of each filter is proportional to the amplitude of the signal in a discrete frequency band or *bin*, centered on the analysis frequency of the filter. To display a spectrogram, the time-varying output amplitudes of filters at successive analysis frequencies are plotted above each other, with amplitude again represented by grayscale values (Figure B.5b). We call this the "filterbank" model of STFT analysis. A given STFT can be characterized by its *bandwidth*, the range of input frequencies around the central analysis frequency that are passed by each filter. All of the filters of a single STFT have the same bandwidth, irrespective of analysis frequency.[1]

These two descriptions of STFT analysis are related in specific ways that are discussed further below. Depending on the context, one or the other of these models may be a more convenient way to think about the STFT. Canary's controls and measurement panels are designed to facilitate considering spectrograms and spectra from either perspective. In the remainder of this appendix, we will refer to both models in discussing how Canary generates spectrograms and spectra.

[1]There are other time-frequency representations (for example, the wavelet transform) that employ different filter bandwidths at different center frequencies.

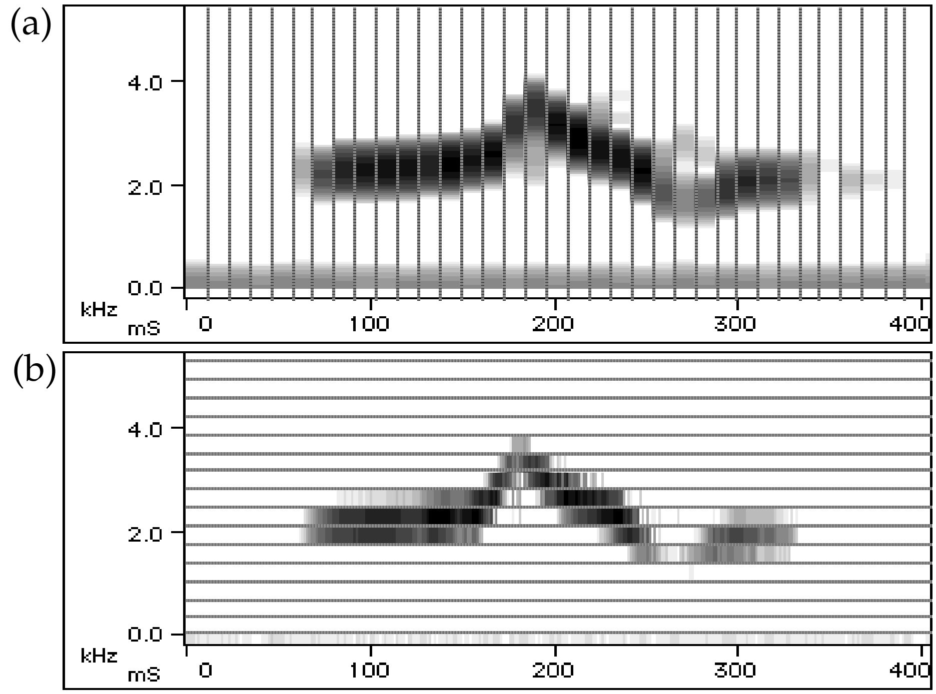

Figure B.5. Two ways of considering a sound spectrogram. Both spectrograms are of the same signal shown in Figure B.4, but with different horizontal and vertical resolution. **(a)** Each vertical bar represents the spectrum of a single short time interval or frame, and approximates the “instantaneous” spectrum at a point midway through the frame. **(b)** Each horizontal bar represents the amplitude of the time-varying output of one bandpass filter.

Frame length, filter bandwidth, and the time-frequency uncertainty principle

The frame length of a STFT determines the *time analysis resolution (Δt)* of the spectrogram. Changes in the signal that occur within one frame-length of each other (e.g., the end of one sound and the beginning of another, or changes in frequency) cannot be resolved as separate events. Thus, shorter frame lengths allow better time analysis resolution.

Similarly, the bandwidth of a STFT determines the *frequency analysis resolution (Δf)* of the spectrogram: frequency components that differ by less than one filter-bandwidth cannot be distinguished from each other in the output of the filterbank. Thus a STFT with a relatively wide filter bandwidth will have poorer frequency analysis resolution than one with a narrower bandwidth.

Ideally we might like to have very fine time *and* frequency analysis resolution in a spectrogram. These two demands are intrinsically incompatible, however: the frame length and filter bandwidth of a STFT are inversely proportional to each other, and cannot be varied independently. Although a short frame length yields a spectrogram with finer time analysis resolution, it also results in wide bandwidth filters and correspondingly poor frequency analysis resolution. Thus a tradeoff exists between how precisely a spectrogram can specify the spectral (frequency) composition of a signal and how precisely it can specify the time at which the signal exhibited that particular spectrum.[1]

The relationship between frame length and filter bandwidth applies to spectra as well as spectrograms. The spectrum of a single frame at a particular point in time can be thought of as a cross-section or vertical slice through the output of a filterbank. The bandwidth of the filters in the bank is determined by the length of the frame.

Figure B.6 illustrates the relationship between frame length and filter bandwidth. The two spectra, of a 1000 Hz pure tone digitized at 22.3 kHz, were made with different frame lengths and thus different bandwidths. Spectrum (a), with a frame length of 1024 points (46.0 mS)[2], shows a fairly sharp peak at 1000 Hz because of its relatively narrow bandwidth filter; spectrum (b), with a frame length of 256 points (11.5 mS), corresponding to a wider bandwidth filter, has much poorer frequency resolution.

---

[1]The spectrogram is one of many different types of time-frequency representations (TFRs) that show how the frequency spectrum of a signal changes over time. The TFR with the highest resolution is the Wigner distribution. Spectrograms and all other (reasonable) TFR's are smoothed (blurred) versions of the Wigner distribution. The smoothing of the spectrogram is controlled by the length and shape of the spectrogram's windowing function. The uncertainty principle gives a lower bound on the amount of blurring that takes place when passing from the Wigner distibution to the spectrogram. Although it might be tempting to use the Wigner distribution without smoothing, there are practical disadvantages to this. See the recent book by Cohen for further discussion.

[2]The frame length of a STFT can be expressed either in "points" (i.e., the number of digital samples in the frame), or in seconds. The time between successive points is equal to the inverse of the sampling rate (1/fs), so the frame length in seconds equals the number of points in the frame divided by the sampling frequency.

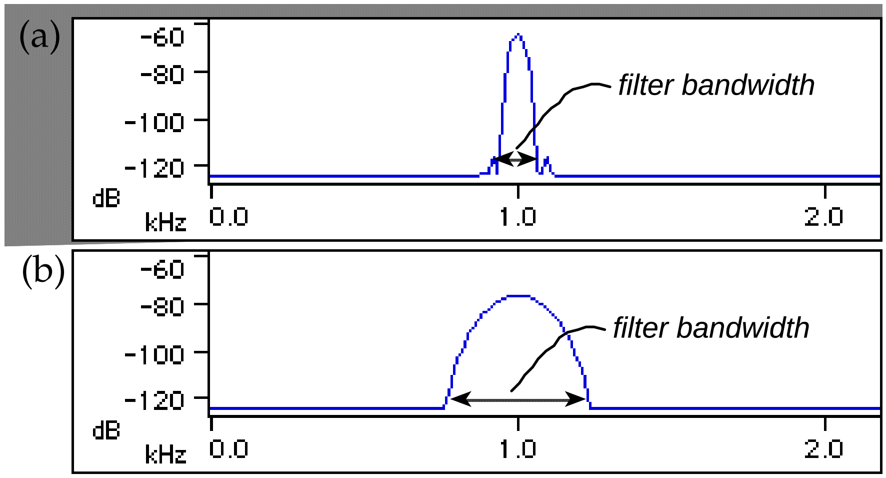

Figure B.6. Relationship between frame length and filter bandwidth.[1] Each spectrum is of a single frame of a 1000 Hz tone, digitized at 22.3 kHz. In both spectra, FFT size = 2048 points, window function = Blackman, clipping level = -130 dB.
(a) Frame length = 1024 points = 46.0 mS; filter bandwidth = 135 Hz.
(b) Frame length = 256 points = 11.5 mS; filter bandwidth = 540 Hz.

Making spectrograms

A spectrogram produced by Canary is a two-dimensional grid of discrete data points on a plane in which the axes are time and frequency. At each gridpoint, an estimate of the amplitude of sound energy is plotted as a grayscale value. In a spectrogram displayed in "boxy" mode, the gridpoints are at the corners of the boxes (Figure B.7). The grayscale value in each box reflects the amplitude at its upper left corner.

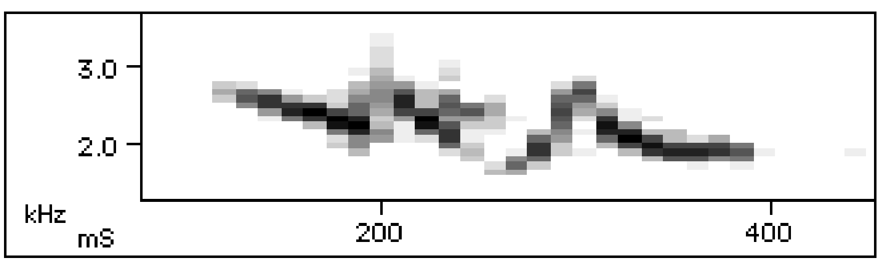

Figure B.7. Low-resolution boxy spectrogram of part of a song of an American robin, digitized at 22.3 kHz. The grayscale value in each box represents an estimate of the energy amplitude at the time-frequency point that is at the upper left corner of the box. Filter bandwidth = 353 Hz, frame length = 256 points (= 11.5 mS). Grid resolution = 11.5 mS x 86.9 Hz.

---

[1]Filter bandwidths are often measured as the width of the band between the frequencies where the amplitude of the filter's output is 3 dB below the peak output frequency. The arrows indicating the filter bandwidths in this figure are placed at a lower amplitude for clarity of illustration.

Canary lets you independently specify the spacing between gridpoints in the horizontal and vertical directions, and thus the width and height of the boxes in a boxy spectrogram (Figure B.8). These spacing values are called, respectively, the *time grid resolution* and *frequency grid resolution* of the spectrogram. Canary's Spectrogram Options dialog box lets you specify time and frequency grid resolution directly, or indirectly by specifying the amount of overlap between successive frames, and the FFT size, respectively. The relationships between time grid resolution and frame overlap, and between frequency grid resolution and FFT size are discussed below. See Chapter 3 for a detailed discussion of how to control these parameters in Canary.

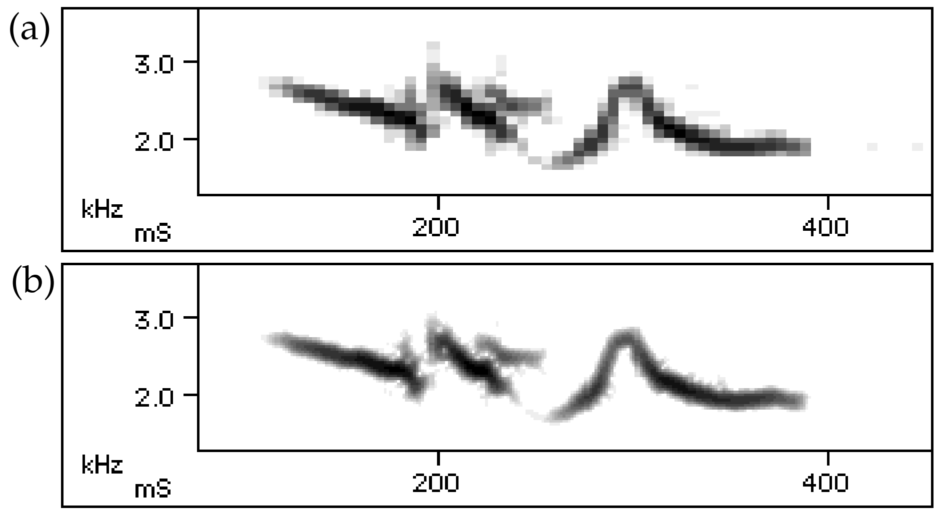

Figure B.8. Boxy spectrograms of the same signal as in Figure B.7, with the same analysis resolution (filter bandwidth and frame length), but different grid resolutions.
(a) Grid resolution = 5.8 mS x 86.9 Hz. **(b)** Grid resolution = 1.4 mS x 10.9 Hz.

Grid resolution should not be confused with *analysis resolution*. Analysis resolution for time and frequency are determined by the frame length and filter bandwidth of a STFT, respectively. Analysis resolution describes the amount of smearing or blurring of temporal and frequency structure at each point on the grid, irrespective of the spacing between these points. The following sections seek to clarify the concepts of analysis resolution and grid resolution by showing examples of spectrograms that illustrate the difference between the two.

Analysis resolution and the time-frequency uncertainty principle

At each point on the spectrogram grid, the tradeoff between time and frequency analysis resolution is determined by the relationship between frame length and filter bandwidth, as discussed above. According to the uncertainty principle, a spectrogram can never have extremely fine analysis resolution in both the frequency and time dimensions.

For example, Figure B.9 shows two spectrograms of the same signal that differ only in frame length and filter bandwidth. The signal consists of two repetitions of a sequence of four tones. Each tone is 20 mS long and has a frequency of 1, 2, 3, or 4 kHz. In spectrogram (a), with a frame length of 64 points (= 2.9 mS; filter bandwidth = 1412 Hz), the beginning and end of each tone can be clearly distinguished and are well-aligned with the corresponding features of the waveform. However, the frequency analysis resolution is poor: each tone appears as a bar that is nearly 800

Hz in thickness. In spectrogram (b), the frame length is 512 points, or 23 mS (filter bandwidth = 176 Hz), or about as long as each tone in the signal. Most of the frames therefore span more than one tone, in some cases including a tone and a silent interval, in other cases including two tones and an interval. The result is poor time resolution: the beginning and end of the bars representing the tones are fuzzy and poorly aligned with the actual features of the waveform (compare, for example, the beginning time of the first pulse in the waveform with the corresponding bar in the spectrogram).

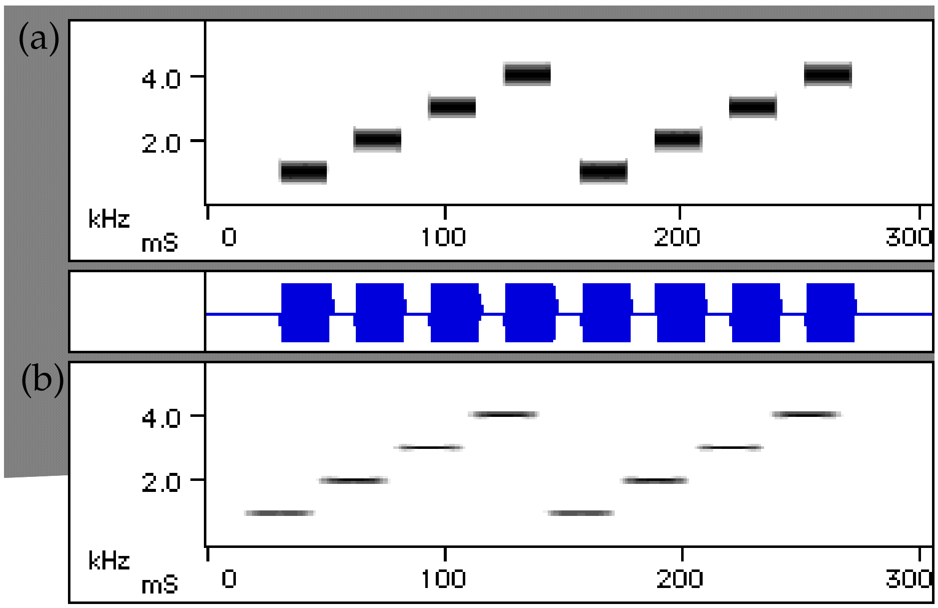

Figure B.9. Effect of frame length and filter bandwidth on time and frequency resolution. The signal consists of two repetitions of a sequence of four tones with frequencies of 1, 2, 3, and 4 kHz. Each tone is 20 mS in duration. The interval between tones is 10 mS. Both spectrograms have the same clipping level, time grid resolution = 1.4 mS, frequency grid resolution = 43.5 Hz (FFT size = 512 points), and window function = Hamming.

(a) Wide-band spectrogram: frame length = 64 points (= 2.9 mS), filter bandwidth = 1412 Hz.

(b) Narrow-band spectrogram: frame length = 512 points (= 23.0 mS), filter bandwidth = 176 Hz.

The waveform between the spectrograms shows the timing of the pulses.

What is the “best” analysis resolution to choose? The answer depends on how rapidly the signal’s frequency spectrum changes, and on what type of information is most important to show in the spectrogram, given your particular application. For many applications, it may be best to start with an intermediate frame length (e.g., 256 or 512 points) and filter bandwidth. If you need to observe very short events or rapid changes in the signal, a shorter frame may be better[1]; if precise

[1]If the features that you’re interested in are distinguishable in the waveform (e.g., the beginning or end of a sound, or some other rapid change in amplitude), you’ll achieve the best precision and accuracy by making time measurements on the waveform rather than the spectrogram.

frequency representation is more important, a longer frame may be better. If you need better time *and* frequency resolution than you can achieve in one spectrogram, you may need to make two spectrograms: a wide-band spectrogram with a small frame for making precise time measurements, and a narrow-band spectrogram with a larger frame for precise frequency measurements.

Time grid resolution and frame overlap

Time grid resolution is the time between the beginnings of successive frames. In a boxy spectrogram, this interval is visible as the width of the individual boxes (Figures B.7 and B.8). Successive frames that are analyzed may be overlapping (positive overlap), contiguous (zero overlap), or discontiguous (negative overlap). Overlap between frames is usually expressed as a percentage of the frame length.

Figure B.10 illustrates the different effects of changes to frame length and time grid resolution. Each pulse in the signal is a frequency-modulated tone that sweeps upward in frequency over a range of 380 Hz centered at 1, 2, 3, or 4 kHz. Spectrograms (a) and (b) both have a frame length of 512 points (= 23.0 mS; filter bandwidth = 176 Hz). (a) was made with 0% overlap (grid resolution = 23.0 mS), whereas (b) was made with an overlap of 93.8% (grid resolution = 1.4 mS). In the low-resolution spectrogram (a), each box is as wide as a frame, which in turn is about the same size as each pulse in the signal. The result is a spectrogram that gives an extremely misleading picture of the signal. Spectrogram (b), with a greater frame overlap, is much “smoother” than the one with less overlap, and it reveals the frequency modulation of each pulse in the signal. It still provides poor time analysis resolution, however, because of its large frame length— notice the fuzzy beginning and end to each bar on the spectrogram and the poor alignment with the corresponding features in the waveform. Comparison of the spectrograms in Figure B.10 demonstrates that improved time grid resolution is not a substitute for finer time analysis resolution, which can be obtained only by using a shorter frame (Figure B.10c).

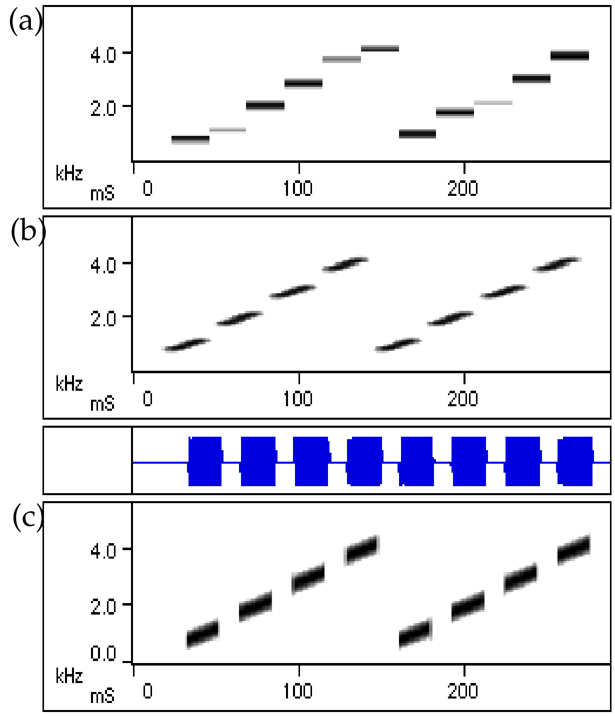

Figure B.10. Different effects on spectrograms of changing frame length (time analysis resolution) and time grid resolution. The signal is two repetitions of a series of four frequency-modulated tones, each 20 mS long, with 10 mS between tones. Each tone sweeps upward in frequency through a range of about 380 Hz centered around 1, 2, 3, or 4 kHz. Spectrograms (a) and (b) have the same frame length, but (b) has finer time grid resolution. (b) and (c) have the same grid resolution, but (c) has a shorter frame length (finer time analysis resolution). In both spectrograms, filter bandwidth = 176 Hz (frame length = 512 points = 23.0 mS), frequency grid resolution = 43.5 Hz (FFT size = 512 points).
(a) Frame length = 512 points = 23.0 mS (filter bandwidth = 176 Hz); Time grid resolution = 23.0 mS (overlap = 0%).
(b) Frame length = 512 points = 23.0 mS (filter bandwidth = 176 Hz); Time grid resolution = 1.4 mS (overlap = 93.8%).
(c) Frame length = 64 points = 2.9 mS (filter bandwidth = 1412 Hz); Time grid resolution = 1.4 mS (overlap = 50%).
The waveform between the spectrograms shows the timing of the pulses.

Frequency grid resolution and FFT size

Frequency grid resolution is the difference (in Hz) between the central analysis frequencies of adjacent filters in the filterbank modeled by a STFT, and thus the size of the frequency bins. In a boxy spectrogram, this spacing is visible as the height of the individual boxes (Figures B.7 and B.8). Frequency grid resolution depends on the sampling rate (which is fixed for a given digitized signal) and a parameter of the FFT algorithm called *FFT size*. The relationship is

frequency grid resolution = (sampling frequency) / FFT size where frequency grid resolution and sampling frequency are measured in Hz, and FFT size is measured in points.[1] Thus a larger FFT size draws the spectrogram on a grid with finer frequency resolution (smaller frequency bins, vertically smaller boxes). The number of frequency bins in a spectrogram or spectrum is half the FFT size.[2]

Figure B.11 illustrates the different effects of changes to filter bandwidth and frequency grid resolution. Spectrograms (a) and (b) both have a filter bandwidth of 1412 Hz (frame length = 64 points = 2.9 mS). However, the frequency grid resolution in (a) is 348 Hz, whereas in (b) it is 43.5 Hz. Spectrogram (b), with finer grid resolution, is "smoother" than (a), but it still provides poor frequency analysis resolution because of its wide bandwidth— the bars representing the pulses in the signal are still quite thick in the vertical dimension. Only by using a narrower bandwidth (Figure B.11c) can we get finer analysis resolution.

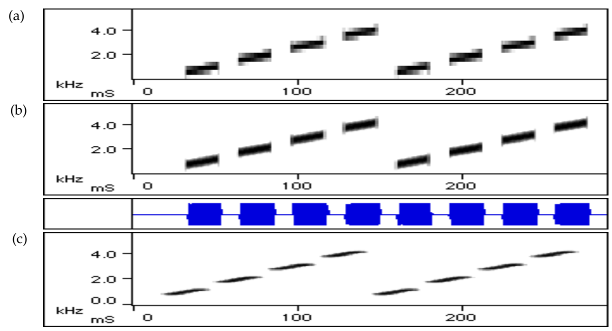

Figure B.11. Different effects on spectrograms of changing filter bandwidth (frequency analysis resolution) and frequency grid resolution (FFT size). The signal is the sequence of frequency-modulated tones shown in Figure B.10. Spectrograms (a) and (b) have the same filter bandwidth, but (b) has finer frequency grid resolution. (b) and (c) have the same grid resolution, but (c) has a narrower bandwidth. In both spectrograms, filter bandwidth = 1412 Hz (frame length = 64 points = 2.9 mS), time grid resolution = 1.4 mS.
(a) Filter bandwidth = 1412 Hz (frame length = 64 points = 2.9 mS); Frequency grid resolution = 348 Hz (FFT size = 64 points).
(b) Filter bandwidth = 1412 Hz (frame length = 64 points = 2.9 mS); Frequency grid resolution = 43.5 Hz (FFT size = 512 points).
(c) Filter bandwidth = 176 Hz (frame length = 512 points = 23.0 mS); Frequency grid resolution = 43.5 Hz (FFT size = 512 points).
The waveform between the spectrograms shows the timing of the pulses.

Spectral smearing and sidelobes

The spectrogram (or a single-frame spectral slice) produced by a STFT is "imperfect" in several respects. First, as discussed above, each filter simulated by the STFT has a finite band of

---

[1]A *point* is a single digital sample.

[2]Ordinarily, the FFT size of a discrete Fourier transform equals the frame size. Canary allows you to specify a larger FFT to obtain finer grid resolution. This is achieved by zero-padding the selected frame length up to a frame whose length is equal to the FFT size.

frequencies to which it responds; the filter is unable to discriminate different frequencies within this band. According to the uncertainty principle, the filter bandwidth can be reduced— thus improving frequency resolution— only by analyzing a longer frame, which reduces temporal resolution.

Second, the passbands of adjacent filters overlap in frequency, so that some frequencies are passed (though partially attenuated) by more than one filter (Figure B.12). Consequently, when a spectrum or spectrogram is constructed by plotting the output of all of the filters, a signal consisting of a pure tone becomes "smeared" in frequency (Figure B.12c).

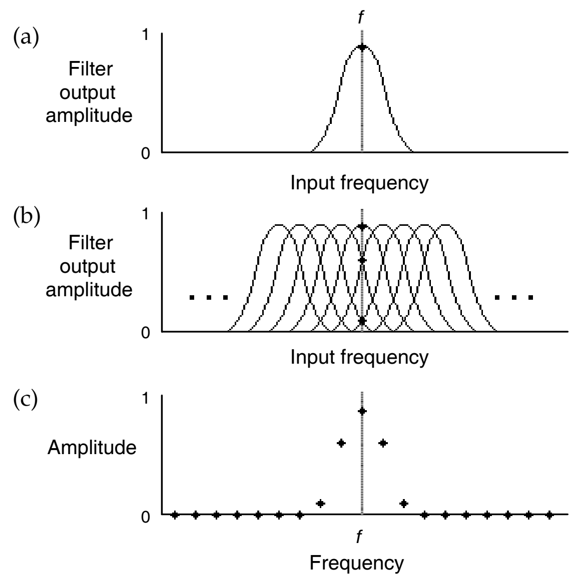

Figure B.12. Spectral smearing resulting from overlapping bandpass filters.
(a) A single hypothetical bandpass filter centered at frequency *f*. For clarity of illustration, sidelobes to the main passband are not shown (see text and Figure B.13).
(b) A set of overlapping filters. Each curve shows the filter function of one filter in a bank simulated by a STFT. Frequency *f* falls within the passbands of the filter centered at *f*, and of two filters on either side.
(c) Spectrum of a pure tone signal of frequency *f* produced by the filterbank shown in (b). The spectrum consists of one amplitude value from each filter. Because the filters overlap, the spectrum is smeared, showing energy at frequencies adjacent to *f*.

Third, each filter does not completely block the passage of all frequencies outside of its nominal passband. For each filter there is an infinite series of diminishing "ripples" in the filter's response to frequencies above and below the passband (Figure B.13a). These ripples arise because of the onset and termination of the portion of the signal that appears in a single frame. Since a spectrum of a pure tone made by passing the tone through a set of bandpass filters resembles the frequency response of a single filter (Figure B.12), a STFT spectrum of any signal (even a pure tone) contains frequency ripples. In a logarithmic spectrum, these ripples show up as "sidelobes" (Figure B.13b).

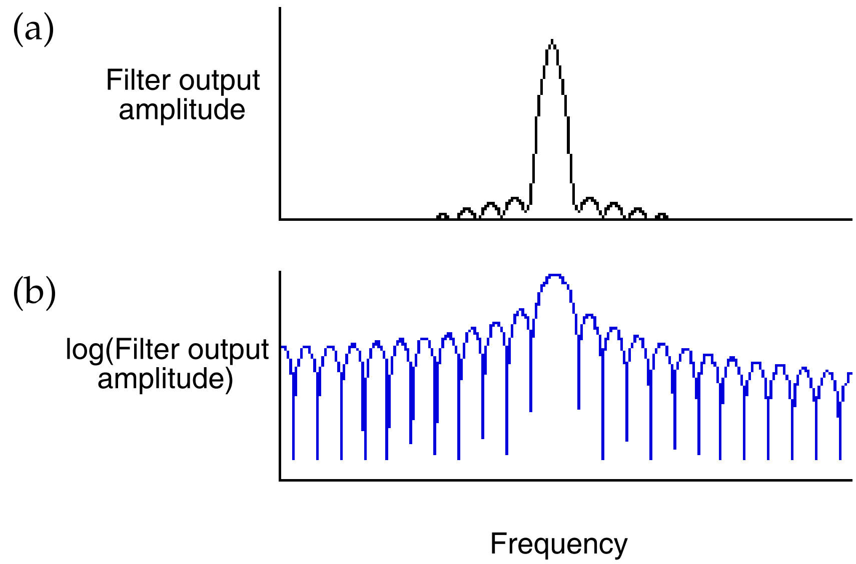

Figure B.13. Frequency response of a hypothetical bandpass filter from a set of filters simulated by a short-time Fourier transform, showing ripples or sidelobes above and below the central lobe, or passband. The magnitude of the sidelobes relative to the central lobe can be reduced by use of a window function (see text). Note that a spectrum produced by passing a pure tone through a set of overlapping filters is shaped like the filter frequency response (see Figure B.12). **(a)** Linear plot. **(b)** Logarithmic plot.

Window functions

The magnitude of the sidelobes (relative to the magnitude of the central lobe) in a spectrogram or spectrum of a pure tone is related to how abruptly the signal's amplitude changes at the beginning and end of a frame. A sinusoidal tone that instantly rises to its full amplitude at the beginning of a frame, and then instantly falls to zero at the end, has higher sidelobes than a tone that rises and falls smoothly in amplitude (Figure B.14).

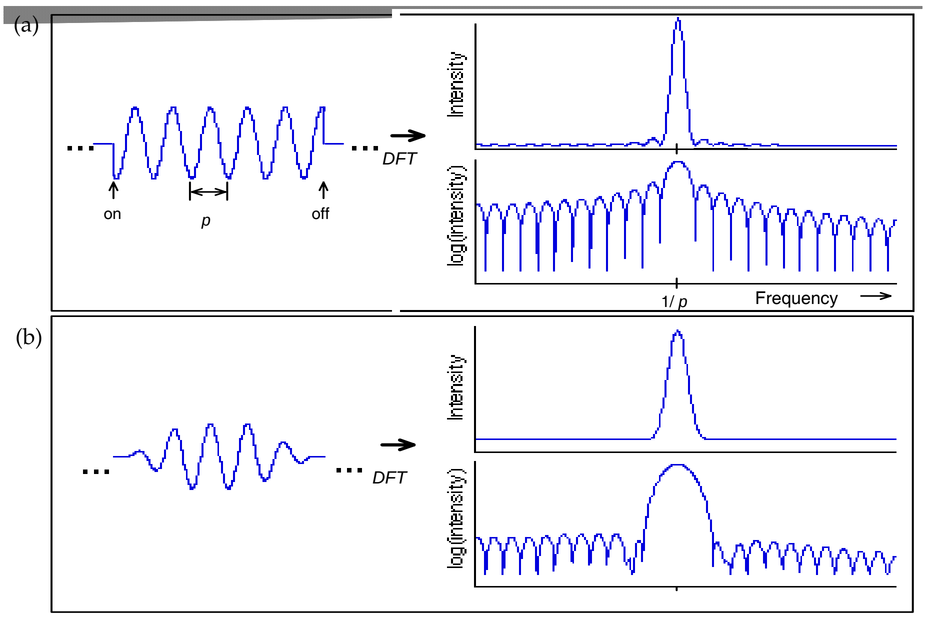

Figure B.14. Relationship between abruptness of onset and termination of signal in one frame and spectral sidelobes. Each panel shows a signal on the left, and its spectrum on the right.

(a) A single frame of an untapered sinusoidal signal has a spectrum that contains a band of energy around the central frequency, flanked by frequency ripples, as if the signal had been passed through a bank of bandpass filters like the one shown in Figure B.13; the ripples appear as sidelobes in the logarithmic (lower) spectrum.

(c) A single frame of a sinusoidal signal multiplied by a “taper” or window function, has smaller sidelobes; the ripples are too small to be visible in the linear (upper) spectrum.

The magnitude of the sidelobes in a spectrum or spectrogram can be reduced by multiplying the frame by a *window function* that tapers the waveform as shown in Figure B.14.[1] Tapering the waveform in the frame is equivalent to changing the shape of the analysis filter (in particular, lowering it sidelobes). Canary supplies five window functions to choose from. Figure B.15 shows spectra of a pure tone made with each of the available window functions. These are also the shapes of the resulting analysis filters. Each window function reduces the height of the highest sidelobe to some particular proportion of the height of the central peak; this reduction in sidelobe magnitude is termed the sidelobe *rejection*, and is expressed in decibels. Given a particular frame length, the choice of window function thus determines the sidelobe rejection, and also the width of the center lobe. The width of the center lobe in the spectrum of a pure tone is the filter bandwidth. For example, the rectangular window function has a narrower filter bandwidth for a given frame length than the Hamming window function, but the Hamming window has lower sidelobes. Figure B.16 shows filter bandwidths corresponding to various frame lengths for each of the five window functions, in order of increasing sidelobe rejection.

---

[1]Window functions are also sometimes called “tapers”.

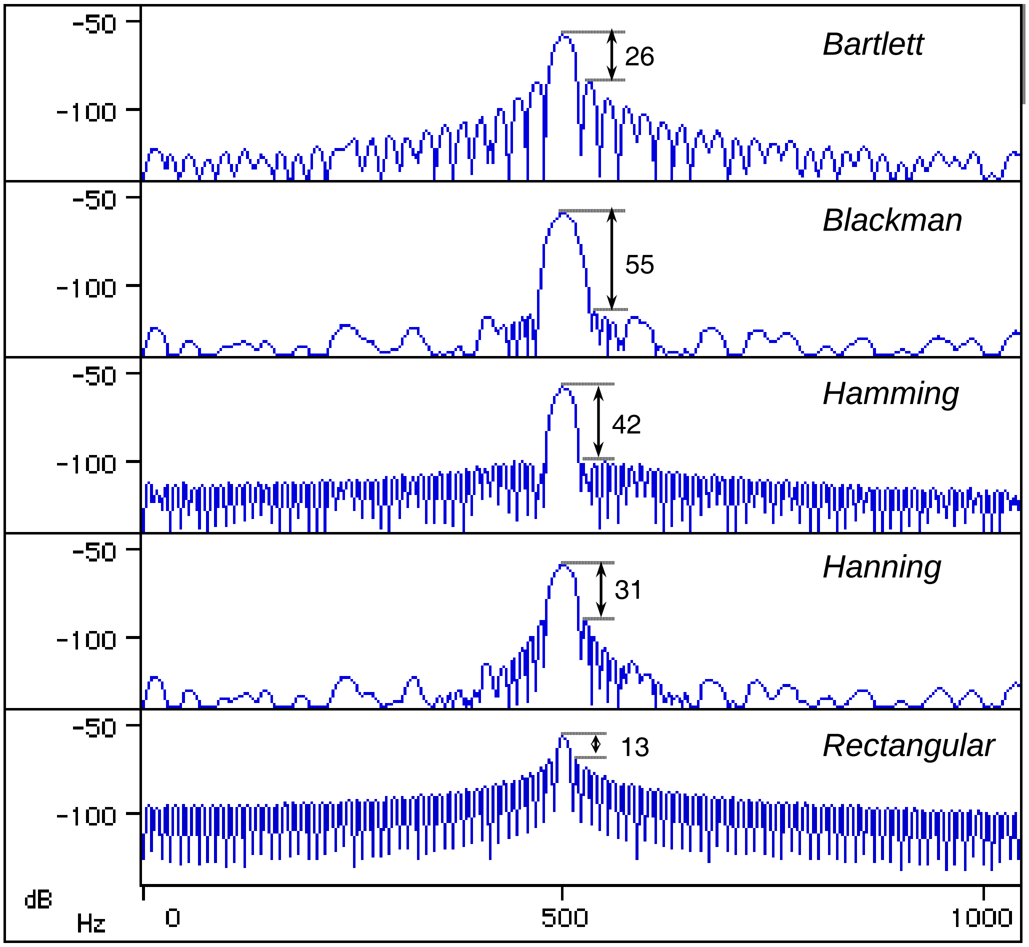

Figure B.15. Single-frame spectra of a 500 Hz tone made with five different window functions. Frame length = 2048 points, FFT size = 8192 points for all spectra. The vertical arrows indicate the sidelobe rejection in dB for each window function.

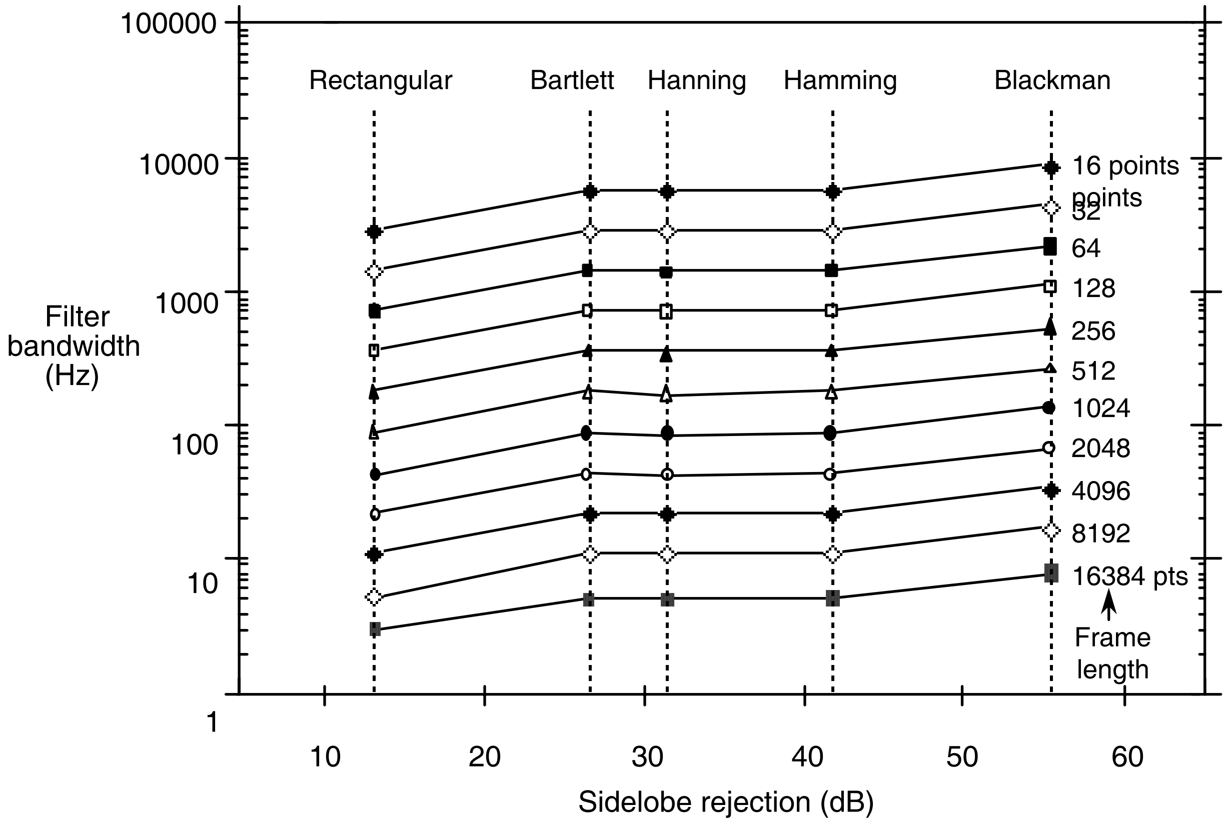

Figure B.16. Filter bandwidths corresponding to different frame lengths for each of Canary's five window functions, in order of the windows' increasing sidelobe rejection, given a sample rate of 22.3 kHz.

For further reading

The books and articles listed below can provide entry at several levels into the vast literature on spectrum analysis and digital signal processing.

Beecher, M. D. 1988. Spectrographic analysis of animal vocalizations: Implications of the "uncertainty principle." *Bioacoustics* 1:(1): 187-207.

Includes a discussion of choosing an "optimum" filter bandwidth for the analysis of frequency-modulated bioacoustic signals.

Cohen, L. 1995. *Time-frequency analysis.* Prentice-Hall, Englewood Cliffs, NJ.

Hlawatsch, F. and G.F. Boudreaux-Bartels. 1992. Linear and quadratic time-requency signal representations. *IEEE Signal Processing Magazine*, 9(2): 21-67.

A technical overview and comparison of the properties of a variety of time-frequency representations (including spectrograms), written for engineers.

Jaffe, D. A. 1987. Spectrum analysis tutorial. Part 1: The Discrete Fourier Transform; Part 2: Properties and applications of the Discrete Fourier Transform. *Computer Music Journal*, 11(3): 9-35.

An excellent introduction to the foundations of digital spectrum analysis. These tutorials assume no mathematics beyond high school algebra, trigonometry, and geometry. More advanced mathematical tools (e.g., vector and complex number manipulations) are developed as needed in these articles.

Marler, P. 1969. Tonal quality of bird sounds. In: *Bird Vocalizations: Their Relation to Current Problems in Biology and Psychology* (ed. R. A. Hinde), pp. 5-18. Cambridge University Press.

Includes an excellent qualitative discussion of how the time and frequency analysis resolution of a spectrum analyzer interact with signal characteristics to affect the "appearance" of a sound either as a spectrogram or as an acoustic sensation.

Oppenheim, A.V. and Schafer, R.W. 1975. *Digital Signal Processing.* Prentice-Hall, Englewood Cliffs, NJ. xiv + 585 p.

A classic reference, written principally for engineers.

Rabiner, L.R. and Gold, B. 1975. *Theory and Application of Digital Signal Processing.* Prentice-Hall, Englewood Cliffs, NJ. xv + 762 p.

Another classic engineering reference.

Yost, W.A. and Nielsen, D.W. 1985. *Fundamentals of Hearing: An Introduction.* 2d ed. Holt, Rinehart and Winston, New York. x + 269 p.

A good general text on human hearing that includes some discussion of the elementary physics of sound and an appendix that introduces basic concepts of Fourier analysis.

Appendix C Sound Amplitude Measurements

About this appendix

This appendix provides a summary of the relationships among sound power, sound intensity, and sound pressure. These quantities are sometimes confused, in part because all three are often expressed as levels using a decibel (dB) scale. First we define each quantity and explain how they are related to each other. We then explain the use of a dB scale to express relative levels of power, intensity, and pressure. An understanding of the basic principles introduced here will be helpful in using Canary.

This appendix is not intended to be a comprehensive review of any aspect of the physics of sound. The goal is to provide a minimal level of understanding needed to use Canary effectively.

The references cited at the end of this appendix provide further background.

Sound power, intensity, and pressure

Sound consists of traveling waves of alternating compression and rarefaction in an elastic medium (such as air or water), generated by some vibrating object (a sound source).

Sound power

A sound source transfers acoustic energy to the surrounding medium at some rate. The average amount of acoustic energy radiated in all directions by a source per unit time is called the *sound power* of the source. Since the usual unit of measurement for energy is the joule, power is usually expressed in joules per second, or watts. One watt equals one joule per second.

Because sound power is a characteristic of a sound source, its value does not depend on where an observer or a measurement instrument is located relative to the source. The power of a sound source may vary over time.

Sound intensity

Consider a sound source radiating sound at a constant power uniformly in all directions. If no sound energy is lost as it radiates away from the source, the total power passing through the surface of any sphere centered on the source is the same, irrespective of the size of the sphere. At greater distances, the same amount of power is distributed over spheres with progressively larger surface areas, resulting in a lower density of power per unit area. This density of power passing through a surface perpendicular to the direction of sound propagation is called *sound intensity*, and is usually expressed in watts per square meter.

The surface area of a sphere of radius r is equal to $4\pi r^2$. Therefore the intensity I (in watts/m^2) of sound at distance r (in m) from a source that is radiating acoustic power equally in all directions is given by

$$I = \frac{W}{4\pi r^2} \qquad \text{(C.1)}$$

where W is the sound power of the source (in watts).

Sound pressure

Sound pressure is the (usually small) alternating incremental change in pressure from ambient pressure that results from a sound. When no sound is present in a medium (i.e., there is no propagating pressure change), we say that sound pressure is zero, even though the medium does exert some static ambient pressure. The dimensions of pressure are force per unit area. The usual unit of sound pressure is the pascal (abbreviated Pa); one pascal equals one newton per square meter.[1] Since the smallest audible sound pressures in air are on the order of 10^{-6} Pa, sound pressures are usually expressed in μPa.

The RMS magnitude of the pressure change that results from sound of a given intensity depends on a property of the medium known as the *characteristic impedance*.[2] Characteristic impedance is equal to the density of the medium ρ (in kg/m^3) times the speed of sound in the medium c (in m/sec). The units of characteristic impedance are mks rayls, named after the famous acoustician Lord Rayleigh. 1 mks rayl equals 1 $kg/(m^2$ sec).

Pressure and intensity are related by

$$I = \frac{p^2}{\rho c} \qquad \text{(C.2)}$$

where p is the RMS or root mean square pressure in Pa, ρc is the characteristic impedance of the medium, and I is intensity in W/m^2. The RMS pressure is equal to the square root of the average of the squared pressure.[3]

Sound pressure is the quantity that is directly measured by most sound measurement or transduction devices, such as sound level meters and microphones.

Sound levels: the decibel scale

Sound power, sound intensity, and sound pressure are all different physical quantities with different dimensions. But all are commonly expressed in decibels, which is sometimes a cause of confusion. Decibels are dimensionless units used to express the logarithm of the ratio between a given value and some specified reference value; some authors require that the values used in the ratio be powers.

Sound levels: definition of decibel measurements

In general, the term “level” in acoustics refers to the logarithm of the ratio of two quantities.

---

[1] Some older acoustic literature uses pressure units of dynes per square centimeter (dyn/cm^2). One pascal equals 10^{-5} dyn/cm^2.

[2] Characteristic impedance is also sometimes called *specific acoustic resistance* (Urick 1983).

[3] If the sound is a constant-amplitude sinusoidal tone, RMS pressure is equal to the peak pressure divided by $\sqrt{2}$.

For a given sound power W, the *sound power level* in dB is given by

$$\text{Sound Power Level} = 10\log_{10}\left(\frac{W}{W_{ref}}\right) \tag{C.3}$$

where W_{ref} is some reference power, which should be clearly stated.

For a given sound intensity I, the *sound intensity level* in dB is given by

$$\text{Sound Intensity Level} = 10\log_{10}\left(\frac{I}{I_{ref}}\right) \tag{C.4}$$

where I_{ref} is some reference intensity. The commonly used reference intensity in air is 10^{-12} W/m^2 (or 1 pW), which is approximately equal to the threshold of audibility of a 1000 Hz tone to a human. For a given sound in air, the intensity level calculated using this reference level is usually within 0.1 dB of the sound pressure level calculated using the standard reference pressure of 20 µPa (see below). The reference intensity for sea water is the intensity that corresponds to the standard reference pressure of 1 µPa (see below), equal to .65 aW/m^2 (=.65 × 10^{-18} W/m^2).

For a given RMS sound pressure p, the *sound pressure level* in dB is given by

$$\text{Sound Pressure Level} = 10\log_{10}\left(\frac{p^2}{{p_{ref}}^2}\right)$$

$$= 20\log_{10}\left(\frac{p}{p_{ref}}\right) \tag{C.5}$$

where p_{ref} is some reference pressure. Squaring the pressure values when calculating sound pressure level ensures that the numerical dB values for intensity and pressure will be the same for a given measurement, provided that the reference values for intensity and pressure are chosen appropriately. The standard reference pressure in air is 20 µPa, which is approximately equal to the threshold of audibility of a 1000 Hz tone to a human. Use of 20 µPa and 10^{-12} W/m^2 as the reference values for pressure and intensity in air yields dB pressure and intensity levels that are usually within 0.1 dB of each other. (The exact difference depends on the value of ρc, which depends on temperature and pressure; see below).

Why the decibel scale is useful

There are two reasons why the decibel scale is often a more convenient way of expressing power, intensity, and pressure than using the corresponding physical units. First, the values that commonly occur for the physical units of sound power, intensity, and pressure all span very large numerical ranges. For example, acoustic power outputs of sound sources range from approximately .000000001 watt for a whispering human voice to 40,000,000 watts for a Saturn rocket taking off. The range of sound intensities between a barely audible 1000 Hz tone and the same frequency at the intensity threshold of pain is .000000000001 watt/m^2 to 1 watt/m^2. Sound pressures range from 20 µPa for sounds at the threshold of human audibility to 100,000,000 µPa for a jet engine at a distance of 25 m. It is often inconvenient to work with such large measurement ranges. The dB scale compresses these very large ranges to more manageable ones. For example, if we take 20 µPa as our reference level for sound pressure, the range of sound pressures between a barely audible sound and the jet engine at 25 m is 0 to 134 dB.

The second reason why the decibel scale is useful is that the ability of the human auditory system to discern a difference in intensity between two tones is roughly logarithmically related to the intensity ratio of the tones (Moore, 1989). The "just noticeable difference" or *jnd* in sound intensity of pure tones (typically around 0.2-0.4 dB) is nearly constant across most of the human hearing range when expressed in decibels. The jnd in intensity or pressure would vary hugely over the range of audible intensities if expressed in watts/m^2 or µPa.

Note that positive dB levels indicate the measured value is greater than the reference value; negative dB levels indicate that the measured value is less than the reference.

Some pitfalls to avoid in using dB measurements

Perhaps the most common mistake made in using dB levels is the failure to state the reference value used. The statement that a sound was measured with an intensity or pressure level of 87 dB, without specifying the reference intensity or pressure, is analogous to the statement that the mass of some object is 46%. Even though the values of 20 µPa and 10^{-12} W are common standards for determining pressure and intensity levels in air, the value used should always be explicitly stated. The preferred way of expressing sound pressure levels is to write (for example) that a given sound has a "sound pressure level of 87 dB (re 20 µPa)".

The standard pressure and intensity reference values for dB levels in air and water have been chosen to result in the same (or very similar) dB levels for pressure and intensity *within a given medium*. This is convenient because it allows one to know for example, that a sound with a pressure level of 100 dB in air (re 20 µPa) also has an intensity level of approximately 100 dB in air (re 1 pW/m^2), without having to convert the pressure dB level to absolute units, calculate the intensity, and then calculate the intensity level in dB.

When comparing dB levels of sounds in different media, remember that (1) the "standard" reference values for calculating pressure and intensity dB levels may be different, and (2) differences in characteristic impedance mean that the relationship between pressure and intensity may be very different in different media. Thus a sound pressure of 100 dB in air (re the standard value of 20 µPa) is not the same pressure as 100 dB in sea water (re the standard value of 1 µPa). The former is a pressure of 2 Pa, while the latter is 0.1 Pa. Furthermore, a sound pressure of 2 Pa in air has an intensity of 10 mW/m^2, which is an intensity level of 100 dB (re 1 pW/m^2); but a sound of 2 Pa pressure in water would have an intensity of .0026 mW/m^2, which is an intensity level of 126 dB (re .65 aW/m^2). When comparing pressures or intensities of sounds in different media, it is simplest to use absolute pressure or intensity values, rather than dB levels.

Characteristic impedance

The characteristic impedance of an elastic medium is the product of the medium's density (ρ) and the speed of sound in the medium (c). The density of air (in kg/m^3) is approximately equal to

$$\rho = 1.29\left(\frac{273}{T+273}\right)\left(\frac{P}{0.76}\right) \qquad \text{(C.6)}$$

where T = temperature in °C, and P = barometric pressure in meters of mercury (Beranek, 1986).

The speed of sound in air (in m/sec) is approximately equal to

$$c = 331.4\sqrt{1 + \frac{T}{273}} \quad \text{(C.7)}$$

where T = temperature in °C (Beranek, 1986). Over the range -30° to +30°C this is approximately equal to

$$c = 331.4 + 0.607\ T \quad \text{(C.8)}$$

For water, see Millero, et al. (1980) for formulas for density, and MacKenzie (1981) for formulas for speed of sound.

For further reading

Beranek, L.L. 1986. *Acoustics*, revised edition. Published for the Acoustical Society of America by the American Institute of Physics. xii + 491 p.

Beranek, L.L. 1988. *Acoustical Measurements*, revised edition. Published for the Acoustical Society of America by the American Institute of Physics. xiv + 841 p.

Mackenzie, K.V. 1981. Nine-term equation for sound speed in the oceans. *J. Acoust. Soc. Am.* 70(3): 807-812.

Millero, F. J., C-T. Chen, A. Bradshaw, and K. Schleicher. 1980. A new high pressure equation of state for sea-water. *Deep-Sea Res.* 27A: 255-264.

Moore, B.C.J. 1989. *An Introduction to the Psychology of Hearing*, 3rd ed. Academic Press, New York. 350 + xvi p.

Urick, R. J. 1983. *Principles of Underwater Sound*, 3d edition. McGraw-Hill, New York. 422 + xiii p.

Appendix D Troubleshooting

About this appendix

This appendix lists solutions of some of the most common questions that are likely to arise with Canary. If you encounter a problem that is not discussed here or in the relevant chapter of the manual, you can request technical assistance by posting to the Canary online user forum at http://canary.RavenSoundSoftware.com/. The Bioacoustics Research Program is not actively supporting Canary since the software was made available for free on the BRP web site.

Known Bugs

- **Double-clicking a Canary file launches the program, but the file doesn't open.**

  This is a known bug that occurs only on PowerMac machines. To open the file, either use the Open command from Canary's File menu, or double-click the file's icon again once Canary is running.

- **Switching off one of the two tracks in the record dialog results in the two stereo input channels being mixed on the remaining track, instead of turning one channel off completely.**

  This is a known bug that afflicts some Macintosh models with stereo recording capability. The solution is simply to provide a zero (silent) signal on the channel that does not carry the desired signal.

Speed and memory

- **Canary never seems to have enough memory.**

  Some background about Macintosh memory allocation

  Whenever you launch Canary (or any other program), the Macintosh operating system allocates a certain amount of memory (RAM) for the exclusive use of that program. The amount of memory that is available for allocation to a program depends first on how much memory is physically installed in the machine and second, on how much of that memory is already allocated to other programs (including the opearting system). To see how much memory is installed in your machine ("Total Memory"), how much is available ("Largest Unused Block"), and how much is allocated to each program that is currently running, select About This Macintosh... from the Finder's apple menu (Figure D.1).

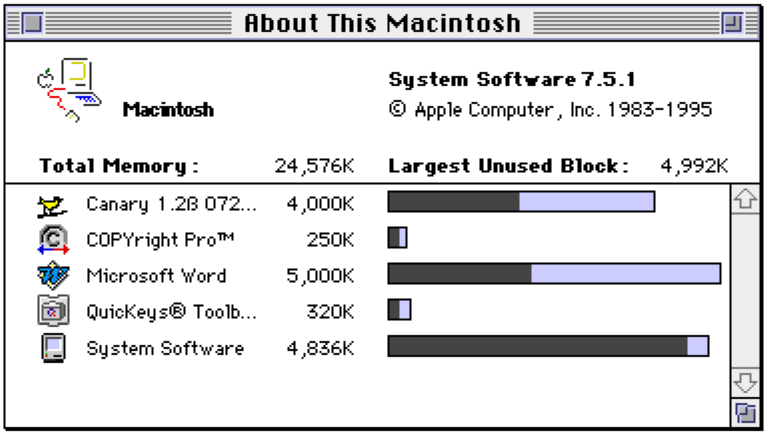

Figure D.1. The Finder's About This Macintosh window, which displays information about memory usage and availability. For each program4that is running, the window shows the amount of memory allocated (e.g., 3000K for Canary).The dark portion of the horizontal bar next to each program name indicates the proportion of the allocated memory that the program is currently using.

The amount of memory that the system attempts to allocate when a program is launched is shown in the program's Info window, in the box labelled Current size or Preferred size under Memory (Figure D.2). To view Canary's Info window, select the Canary program icon in the Finder, then choose Get Info... from the Finder's File menu. (The Memory parameters are in the lower righthand corner of the Info box.) The Info window also displays a Suggested size or Minimum size. If the largest block of memory that is available when a program is launched is less than the Minimum size (Suggested size in System 7.0), the program will not run and the Finder will display a message stating that there is insufficient memory available.[1] If the available memory is greater than the minimum, but less than the Preferred size value (Current size in System 7.0), the system allocates whatever memory is available and launches the program. You can check how much memory was actually allocated to a program in the About This Macintosh window (Figure D.1).

---

[1]The Suggested size of a program is the minimum size suggested by the software manufacturer. You can force the Finder to launch a program with less than its suggested minimum memory requirement by setting the Preferred or Current size value to less than the Suggested size. In general, though, this is not a good idea; the behavior of a program running with less than its recommended minimum memory (if it runs at all) can be unpredictable.

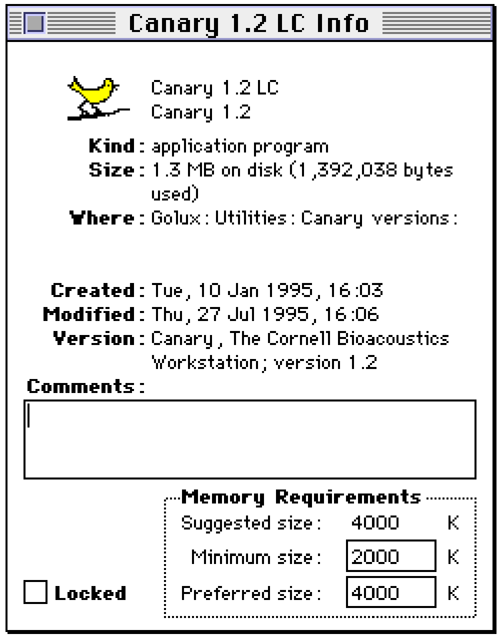

Figure D.2. Canary's Info window. The exact layout of the window depends on which version of the sytem software is running.

You can increase the amount of memory that Canary requests from the operating system by simply editing the Preferred size (or Current size in older versions of the System 7) value in the Info window. Note that you cannot change the Preferred size value while Canary is running (although you can view the Info window at any time).

Solutions to inadequate memory conditions

Below are several potential solutions to inadequate memory conditions. Note that you must quit from Canary before you can try any of these (with the exception of #1), since you cannot change a program's memory allocation while it is running.

1. Close unnecessary windows. Every window in Canary uses up some of the allocated memory. How much memory a window uses depends principally on the signal length, on whether a spectrogram has been calculated for the signal, and on the spectrogram's grid resolution. Closing unnecessary windows frees up memory within Canary's allocation block.

2. Increase the amount of memory that Canary requests. Change the Preferred size parameter in Canary's Info window, as discussed above.

3. Make sure you're using 32-bit addressing. This is relevant only if you have more than 8 megabytes of RAM installed. The Macintosh can operate in either of two "addressing modes", 24-bit and 32-bit. In 24-bit mode, the system can only use up to 8 megabytes (irrespective of how much RAM is installed in the machine). In 32-bit mode, the system uses all of the memory you have installed. System 7 allows you to choose the addressing mode, via the Memory control panel (usually on the apple menu). The only reason to ever turn 32-bit addressing off is that some older Mac programs are incompatible with 32-bit addressing. Note that if you change the addressing mode with the Memory control panel, you must restart the Mac for the change to take effect.

4. Quit from all unnecessary programs. Every program that's running on your machine reduces the amount of memory that can be allocated to Canary. Remember that the amount of memory actually allocated to Canary will be less than what is requested if insufficient memory is available when Canary is launched.

5. Install a memory-management utility such as RAM Doubler from Connectix, Inc. RAM Doubler uses sophisticated memory management techniques to enable the Mac to use memory more efficiently. With RAM Doubler you can allocate more memory to a program than is physically installed in your machine.

6. Buy and install more memory. If you can afford it, do it. You'll never regret having more memory.

7. Use Virtual Memory. This should probably be your last resort and is an option only on a machine equipped with a paged memory management unit (PMMU).[1] Virtual memory increases the memory available to the system by using part of a hard disk as an extension to the system's (RAM) memory. With virtual memory, the amount of memory that can be allocated to Canary is not limited by the amount of RAM installed in your machine. Canary will function under virtual memory, but there are some notable problems. Because virtual memory uses "slow" disk space as a substitute for "fast" chip memory, routines that manipulate memory in Canary are markedly slower. This can create problems for processes that require real time speeds, such as recording or playback. Various factors determine whether virtual memory will have a detrimental effect on real time operations, the most crucial being the physical memory to logical memory ratio determined by the virtual memory settings in the Memory control panel. Satisfactory performance requires additional "disk" memory not to exceed the total amount of physical memory on the machine. In general, you can use larger amounts of "disk" memory with Canary to study and manipulate already digitized sounds, but not to perform recording or playback. Even then, be prepared to wait much longer than usual for many operations. (You can read the manual while waiting!)[2]

- **Drawing or redrawing a spectrogram, or rebuilding the measurement panel inexplicably slows down.**

  This is a consequence of Canary running low on memory. In order to let you keep working in a low memory situation, Canary releases memory that it usually reserves for speeding up graphic operations (like drawing spectrograms). You can restore graphics to normal speed by closing some windows to free up memory. If there are no unneeded windows to close and the slow drawing speed is unbearable, quit Canary and increase the program's memory allocation, as discussed above.

---

[1]If virtual memory is available on your machine, the Memory control panel includes virtual memory controls.

[2]The MacRecorder driver version 1.02 is NOT compatible with virtual memory.

- **Canary seems to be “hanging” for an inordinate amount of time when making selections in a spectrogram, then continuing normally (i.e., the watch hands are spinning but nothing seems to be happening).**

Certain routines that do calculations for some range measurements can take long periods of time, especially when working with a large selected portion of a spectrogram. You can avoid these delays by removing the following measurement cells from the spectrogram measurement panel:

- Average Intensity (or Power)
- Average Intensity (or Power) (dB)
- Energy Flux Density (or Energy)
- Peak Energy Flux Density (or Energy) per Hz
- Peak Frequency
- Peak Intensity (or Power) per Hz
- Peak Time

If any of these cells are included in the measurement panel, the program is forced to recalculate their values every time a new selection is made, which takes a long time.

Spectrograms

- **The spectrogram of a signal appears to be shifted slightly to the left of the waveform: features of a sound that are visible in the waveform appear in the spectrogram at a slightly earlier time.**

This is an inevitable consequence of the frame length of the spectrogram. Consider a spectrogram with a frame length of 512 points made from a signal sampled at 22.3 kHz. The interval between successive points is thus 1/22300 = .045 mS. Therefore, a 512-point frame has a length of 512(.045) = 23 mS. Each frame of the spectrogram thus reflects all events in the waveform that occur within 23 mS after the start of the frame. So, for example if a recording contains a period of silence and then a sudden-onset signal, the signal may show up as much as 23 mS earlier in the spectrogram than it does in the waveform. If you use a longer frame, you’ll see a greater apparent “shift” in the spectrogram; a shorter frame results in less of a shift.

You can see this effect more clearly if you make a boxy (rather than smooth) spg with 0% frame overlap, and stretch the time axis till you can see the individual boxes clearly. This will enable you to see how each frame of the spectrogram covers an interval of the waveform. Again, you will see the effect most clearly if you have part of a signal that is silent (you can make this by selecting a part of the waveform and amplifying by zero), and then a sudden onset of a non-zero signal.

- **Printed spectrograms are dithered.**

Make sure your printer is capable of grayscale printing and that you select Color/Grayscale (rather than black and white) printing in the Print dialog box.

Canary will automatically dither spectrograms if the chosen printer cannot handle grayscale or is set to black and white.

Note: The Dither Spectrogram option in the Display Options dialog box pertains only to Canary's onscreen graphics and has no bearing on printing.

- **Gray areas on printed spectrograms seem grainy.**

  In general, Canary has no control of the final appearance of the printed output. Canary simply sends the image of the signal window contents directly to the printer and the Macintosh operating system and the printer software and hardware handle the actual printing. The main determinant of printed output quality (especially grayscale resolution) is the make and model of your printer; there is considerable variation among printers. If the spectrograms that your printer is producing seem grainy and inadequate, the most likely reason is that your printer is lacking. Try a different printer, you may be surprised.

Opening files

- **Canary asks for parameters every time I open a particular sound, even though this sound has been opened in Canary before.**

  Canary needs to save this sound in order to attach parameters to it. Be certain to save this sound in Canary even if it has not been modified.

- **The file I want to open does not appear in the Open File dialog box.**

  Canary does not recognize the file type of this file. The file type is a four-character string that is attached to each and every item in the Macintosh file system when the file is saved. Canary recognizes the following file types, each corresponding to a file format that Canary understands.

Format	Type
SoundEdit	`'FSSD'`
MATLAB	`'MATW'`
Text	`'TEXT'`
AIFF	`'AIFF'`
Binary	`'Bin '`
Canary	`'CNRY'`

Confusion about a file's type may arise because of the "Kind" designation displayed by the Finder for each file. The Kind displayed for a file is based on what program created the file, *not* on the file's format or type. (Since many programs can save files in multiple formats or types, the Kind does *not* necessarily tell you the file's type. For example, SoundEdit can save files in a format called Instrument, which has a type of `'DEWF'`. `'DEWF'` files are not readable by Canary. However, the Kind that the Finder shows for both `'DEWF'` and `'FSSD'` files created by SoundEdit is "SoundEdit™ document". Thus, the Finder's Kind designation does not necessarily tell you whether or not Canary can read the file.) Although the Macintosh operating system does not provide any way for you to determine the file type of a file directly, some third-party utility programs (e.g., DiskTop) allow you to do so. Also, most application programs use different Finder icons for files that they save with different types (e.g., see Canary's icons in Chapter 7), but you will have to consult the program's documentation or do some experimenting to figure out the correspondence between icons and file types.

If the file you want does not show up in Canary's Open File dialog box, open it in the program that originally created it and re-save it one of the formats listed above.

Miscellaneous

- **When I launch Canary a message appears stating "Canary requires a floating-point coprocessor."**

  If you are using a 680x0 (non-PowerMac) Macintosh without a floating-point co-processor (also known as a math co-processor or floating-point unit (FPU)), you should use Canary 1.2 LC.

- **Sometimes the graphic display in one or more panes of a signal window gets corrupted.**

  Occasionally a window's graphics become garbled or improperly drawn. In most cases, corrupted window graphics can be fixed by selecting the Redraw command from the Windows menu.

- **A recording device (e.g., MacRecorder) that is connected to the Macintosh does not appear on the Device pop-up menu in the Record dialog box.**

  Every recording device should be supplied with a program called a *device driver*, usually in the form of a System Extension which enables the Macintosh to use the device. In order for Canary to use a recording device, its driver must be installed in the System folder. See the documentation that came with your MacRecorder (or other device) for instructions about installing the device driver.

- **Increasing the playback rate doesn't seem to do anything.**

  The Macintosh can play back sounds at any rate up to 64 kHz. This ceiling effectively limits the range of available playback rates. For example, if you have a sound sampled at $\approx$22 kHz, at best the Macintosh can play the sound back at about three times its original rate. Setting the rate slider to a larger factor will not increase the actual playback rate.

- **Background colors seem to change for no reason.**

  The sometimes bizarre behavior of colors within Canary is a byproduct of the way the Macintosh manages the color display when the hardware cannot display the full range of possible colors at once. The number of colors that a color monitor can display at once depends on the display card and the amount of video RAM (VRAM) installed in the Macintosh. Basically, colors change because there are not enough colors available to meet the needs of all the running applications. This can be avoided by using display hardware that uses 24 bits rather than 8 bits to represent the color of each pixel (yielding 16 million rather than 256 possible colors).[1]

[1]24-bit and 8-bit color representation schemes are sometimes called “direct” and “indexed” color, respectively.

Appendix E Metric System Prefixes

Canary's measurements are expressed in mks metric units. Measurements that are more than two orders of magnitude larger or smaller than the basic unit are displayed with one of the standard prefixes listed in the table below.

Prefix	Abbreviation	Multiplier
exa-	E	10^{18}
peta-	P	10^{15}
tera-	T	10^{12}
giga-	G	10^{9}
mega-	M	10^{6}
kilo-	k	10^{3}
milli-	m	10^{-3}
micro-	μ	10^{-6}
nano-	n	10^{-9}
pico-	p	10^{-12}
femto-	f	10^{-15}
atto-	a	10^{-18}

Appendix F Using the Macintosh Built-in Sound Input Port

About the Macintosh built-in sound input port

As of this writing (August 1995), two different types of sound input port are available on Macintosh computers. Both types are marked with a small microphone icon and accept a stereo miniplug (1/8″ phone plug). However, the two types of input port are wired differently.

> ***WARNING!***
>
> *Before you connect any sound input device other than the microphone that came with the computer, it is critical that you determine which type of sound input port your Macintosh has. Connecting audio equipment (e.g., a tape recorder, CD player, or stereo system) to the sound input port incorrectly may damage the audio equipment and/or the computer!*

Determining the type of the sound input port

The simplest way to tell which type of sound input port a Macintosh has is to examine the microphone that came with the machine. The "classic" sound input port is found on Mac models (including the Mac LC, Mac IIsi and non-AV Quadras) that come with a disc-shaped Apple Omni-Directional Microphone that is about 3.5 cm (1.5 inches) in diameter. These models accept only one channel of input sound. The "PlainTalk" sound input port is used on Macintosh models (including AV Macs and PowerMacs) that come with an Apple PlainTalk Microphone, which is roughly triangular in shape and about 5 cm (2 inches) wide. The PlainTalk sound input port supports two input channels (stereo).

Line input to a "classic" Macintosh sound input port

Although it is possible to connect a line output cable directly from an audio device to the classic sound input port using a stereo (three-conductor) miniplug or adapter, doing so may yield distorted signals and may damage the audio device, because the wiring of the classic port does not conform to standard conventions for stereo mini-plugs. Use of a mono (two-conductor) mini-plug in the classic port may damage the computer.

Problem: Distortion of the digitized signal

The classic port is designed for a mic level input voltage. However, standard audio line level signals (such as those produced from the Line Out jack of a tape recorder or CD player) are much higher. A line level signal can saturate the Mac's digitizer, resulting in severely distorted signals.

Problem: Damage to audio devices and the Macintosh

A stereo mini-plug has three contacts: a tip contact, a barrel contact, and a ring contact (between the tip and the barrel). A mono miniplug has only a tip and barrel contact. In most audio applications, the barrel is used as a ground, and the other contact(s) carry the one or two audio signals. In the classic Macintosh sound input port, the barrel is used as a ground and the one audio channel is carried by the tip contact. The ring contact is used to supply a voltage to the Apple Omni-Directional Microphone from the Mac. This voltage can potentially damage an audio device when applied to the device's line output circuits.

If a mono (two-conductor) instead of a stereo (three-conductor) mini-plug is used in the classic port, the voltage intended for the stereo plug's ring connector is short-circuited via the plug's barrel contact, which could potentially damage the Mac.

Solution 1: Use an Apple sound input adapter

Some of the earliest Macs equipped with sound input were shipped with an attenuating line-input adapter (Apple part #590-0618-A) that could be used to connect a standard audio patch cable from the line output jack of an audio device to the Mac's sound input port. The adapter consists of a 19 cm (7.5 inch) cable with a stereo miniplug at one end, and two phono (RCA) jacks at the other end. Apple later discontinued shipping this adapter, and it is now unavailable. If you have access to such an adapter, you can use it to attenuate a line level signal to a level appropriate for the sound input port. The adapter also isolates the audio device from the voltage supplied to the ring contact on the mini-plug.[1]

If you do not have access to the original Apple adapter, see Solutions 2 and 3 below.

Solution 2: Build an adjustable attenuator cable

Using a combination of inexpensive parts available from an electronics store, you can build an adapter cable that provides adjustable attenuation of a line-level signal, and also isolates an audio device from the voltage that the Mac applies to the sound input port.

The Radio Shack Stereo Headphone Volume Control (RS catalog # 42-2459) is a short cable with a stereo mini-plug at one end, a stereo mini jack at the other end, and an in-line potentiometer that controls the degree of attenuation of a signal. Plug the mini-plug into the sound input port (microphone jack) on the Mac. The Radio Shack Stereo Y-adapter (catalog #274-369) has a stereo mini-plug on one side, and two color-coded phono (RCA) jacks on the other side. Plug the Y-adapter into the jack of the headphone volume control. The Y-adapter's red phono jack will now carry the voltage that the Macintosh supplies to power the classic microphone. **Do not connect the line output of any audio device to the red phono jack!** Connect the line output from the audio device to the **white** phono jack. You can use the dial on the headphone volume control to adjust the input level.

Solution 3: Build a fixed attenuator cable

The Radio Shack Phono Jack To Mini Stereo Plug Adapter (catalog #274-378) has a stereo mini-plug on one side, and a single phono jack on the other. The phono jack's center contact connects

---

[1]Although you can find off-the-shelf adapters (at places like Radio Shack) that look equivalent to the Apple adapter (i.e., they have the same connectors at their ends), such an adapter would not attenuate the signal, and it would not isolate the audio device from the voltage that the Mac supplies to the microphone port. **Do not use such an adapter just because it looks equivalent!**

to the mini-plug's tip contact. Plug the adapter's mini-plug into the sound input port (microphone jack) on the Mac. The Radio Shack Audio Attenuator Cable (catalog #42-2461A) has a single phono plug on one end, and a mono mini-plug on the other end. Plug the attenuator cable's phono plug into the phono jack on the Phono-Miniplug adapter. The attenuator cable's mono miniplug can then be connected to the line output of an audio device (you may need another adapter to make this connection). The attenuator cable reduces a line level signal to a level that is acceptable to the Macintosh; the Phono-Miniplug adapter isolates the audio device from the voltage that the Mac applies to the sound input port (because the ring contact on the adapter's mini-plug is not connected to anything).

Line input to a PlainTalk sound input port

If your Mac is equipped with a PlainTalk sound input port, you can connect a stereo line-level audio signal directly to the sound input port, using a stereo miniplug. If the signal you are acquiring is a mono (one-channel) signal, you must still use a stereo miniplug to connect your sound source to the Macintosh.[1] If the line-level signal is higher than desired, you can use the Radio Shack Stereo Headphone Volume Control (catalog # 42-2459) to reduce the level of the input signal to the computer.

[1] If you use a mono miniplug on a Macintosh equipped for stereo sound input, Canary's recording gain controls will not work properly, and the signals you record may be distorted.

Appendix Z: Under the Hood

This appendix is a technical description of the signal processing algorithms used by Canary. It is not meant to be tutorial; instead, it provides a reference for the engineer who needs to know the details of Canary's processing algorithms.

The Spectrogram

Canary's spectrogram is based on the complex short time Fourier transform (STFT) [2,4,6,8] of a real valued signal $x[i]$,

$$X(n\Delta, k) = \sum_{m=0}^{M-1} x[n\Delta + m]\, u[m]\, e^{-\frac{j2\pi km}{N}},$$

where $u[m]$ is the length M *analysis window* (e.g., Hamming, rectangular, etc), M is the *frame size*, N is the *FFT size*, and Δ is the *frame increment* ("hop size"). The STFT is defined for $-N/2+1 \leq k \leq N/2$ and $0 \leq n\Delta \leq \tilde{L} - M$, where $\tilde{L}$ is the number of samples in x.

Canary's "overlap" is

$$100\% \cdot \frac{M - \Delta}{M}.$$

Notice that overlap can be negative if $\Delta > M$. The number of frames produced from a signal of length $\tilde{L}$ is

$$N_f = 1 + \left\lfloor \frac{\tilde{L} - M}{\Delta} \right\rfloor.$$

The spectrogram is computed for $0 \leq n \leq N_f - 1$.

The spectrogram $S(n, k)$ is the squared magnitude of the STFT, scaled by factors that adjust for the proper units. The units of Canary's spectrogram are watts/hertz in the "electric" calibration paradigm, and watts / meter^2 / hertz in the "acoustic" paradigm. The STFT of a real signal is conjugate symmetric with respect to frequency; this means that for any n and k we have $|X(n, -k)| = |X(n, k)|$. In Canary 1.1, the value $|X(n, k)|$, properly scaled, was reported. This had the disadvantage of not reporting the signal energy contained in the negative frequencies ($|X(n, -k)|$). Therefore, Canary 1.2's spectrogram $S(n, k)$ includes both the negative and positive frequency energies. This is a change from Canary 1.1. That is, for $k > 0$, $S(n, k)$ is twice as large in Canary 1.2 than in Canary 1.1.

Energy measurement in the spectrogram

We now present precise formulas for how the spectrogram values and energies are computed. First we assign symbols to the quantities of interest.

$$
\begin{aligned}
f_s &= \text{sampling frequency (Hz)} \\
\Delta t &= \text{sampling period (sec)} \\
&= \frac{1}{f_s} \\
\tilde{L} &= \text{length of waveform (samples)} \\
x[i] &= \text{waveform samples, } 0 \leq i \leq \tilde{L}-1 \\
R &= \text{impedance} \\
p[i] &= \text{instantaneous power} \\
&= \frac{1}{R} x[i]^2 \\
E &= \text{total energy} \\
&= \sum_{i=0}^{\tilde{L}-1} p[i]\Delta t \\
N &= \text{FFT size} \\
M &= \text{frame length} \\
\Delta &= \text{frame increment (samples)} \\
\Omega &= \text{overlap (\%)} \\
&= 100\% \cdot \frac{M-\Delta}{M} \\
dt &= \text{frame increment (sec)} \\
&= \Delta \cdot \Delta t \\
df &= \text{frequency bin size (Hz)} \\
&= f_s/N \\
N_f &= \text{number of frames} \\
&= 1 + \left\lfloor \frac{\tilde{L}-M}{\Delta} \right\rfloor \\
L &= \text{number of waveform samples used} \\
&= M + (N_f - 1) \cdot \Delta \\
T &= \text{duration of waveform used (sec)}
\end{aligned}
$$

$$\begin{aligned} &= L \cdot \Delta t \\ u[i] &= \text{window function, } 0 \le i \le M-1 \\ U^2 &= \sum_{m=0}^{M-1} u_m^2 \end{aligned}$$

In a moment we will have use for these symbols as well:

$$\begin{aligned} \phi[i] &= \text{sample weighting function, } 0 \le i \le \tilde{L}-1 \\ \lambda &= \text{“energy fudge factor”} \\ \eta[k] &= \begin{cases} \frac{1}{2}, & \text{if } k = 0 \\ 1, & \text{if } k = 1, 2, 3, \ldots \end{cases} \end{aligned}$$

We take as our basic discrete Fourier transform the following operation.

$$X(n,k) = \sum_{m=0}^{M-1} x[n\Delta + m]u[m]e^{-\frac{j2\pi km}{N}}$$

for $-\frac{N}{2}+1 \le k \le \frac{N}{2}$ and $0 \le n \le N_f - 1$. Parseval's identity yields the following

$$\frac{1}{N} \sum_{k=-\frac{N}{2}+1}^{\frac{N}{2}} |X(n,k)|^2 = \sum_{m=0}^{M-1} x[n\Delta + m]^2 u[m]^2$$

from which we have

$$\begin{aligned} \frac{1}{N} \sum_{n=0}^{N_f-1} \sum_{k=-\frac{N}{2}+1}^{\frac{N}{2}} |X(n,k)|^2 &= \sum_{n=0}^{N_f-1} \sum_{m=0}^{M-1} x[n\Delta + m]^2 u[m]^2 \\ &= \sum_{i=0}^{L-1} \left[\sum_{l=0}^{\frac{M}{\Delta}-1} u[(i + l\Delta) \bmod M]^2 \right] x[i]^2 \\ &= \sum_{i=0}^{\tilde{L}-1} \phi[i]^2 x[i]^2 \end{aligned}$$

where the sample weights $\phi[i]^2$ are the coefficients of the $x[i]^2$ in the above equation. A typical graph of $\phi[i]^2$ is shown in Figure 1. We define $\phi[i] = 0$

Figure 1: A typical plot of $\phi[i]^2$, using a Bartlett window for $u[i]$.

$$\begin{aligned}\tilde{L} &= 750 \\ M &= 256 \\ \Delta &= 64 \\ N_f &= 8 \\ L &= 704\end{aligned}$$

for $L \leq i \leq \tilde{L}-1$. From the definition of $\phi[i]$ we have the following equation.

$$\sum_{i=0}^{L-1} \phi[i]^2 = N_f \cdot U^2$$

It is sometimes desirable to choose the spectrogram parameters to make the graph of $\phi[i]^2$ as rectangular as possible. Accordingly we define the quantity λ by the equation

$$\lambda \cdot \frac{\sum_{i=0}^{L-1} \phi[i]^2 x[i]^2}{\sum_{i=0}^{L-1} \phi[i]^2} = \frac{1}{L} \sum_{i=0}^{L-1} x[i]^2$$

From this we derive the following approximation for the energy E.

$$\begin{aligned}
E &= \sum_{i=0}^{\tilde{L}-1} p[i]\Delta t \\
&= \sum_{i=0}^{L-1} p[i]\Delta t + \sum_{i=L}^{\tilde{L}-1} p[i]\Delta t \\
&= \frac{1}{R} \cdot \frac{1}{f_s} \sum_{i=0}^{L-1} x[i]^2 + E_L \\
&= \frac{1}{Rf_s} L \cdot \frac{1}{L} \sum_{i=0}^{L-1} x[i]^2 + E_L \\
&= \frac{L}{Rf_s} \cdot \lambda \frac{\sum_{i=0}^{L-1} \phi[i]^2 x[i]^2}{\sum_{i=0}^{L-1} \phi[i]^2} + E_L \\
&= \frac{L\lambda}{Rf_s N_f U^2} \sum_{i=0}^{L-1} \phi[i]^2 x[i]^2 + E_L \\
&= \lambda \frac{(L/f_s)}{RN_f U^2} \frac{1}{N} \sum_{n=0}^{N_f-1} \sum_{k=-\frac{N}{2}+1}^{\frac{N}{2}} |X(n,k)|^2 + E_L \\
&= \lambda \sum_{n=0}^{N_f-1} \sum_{k=0}^{\frac{N}{2}-1} \frac{2\eta[k]T}{R\Delta N_f U^2} |X(n,k)|^2 \left(\frac{f_s}{N}\right) \left(\frac{\Delta}{f_s}\right) \\
&\quad + \lambda \sum_{n=0}^{N_f-1} \frac{T}{R\Delta N_f U^2} |X(n, \frac{N}{2})|^2 \left(\frac{f_s}{N}\right) \left(\frac{\Delta}{f_s}\right) + E_L
\end{aligned}$$

$$
\begin{aligned}
&= \lambda \sum_{n=0}^{N_f-1} \sum_{k=0}^{\frac{N}{2}-1} \hat{S}(n,k) df\, dt + \lambda \sum_{n=0}^{N_f-1} \hat{S}(n, \frac{N}{2}) df\, dt + E_L \\
&= \lambda \tilde{E} + \lambda \sum_{n=0}^{N_f-1} \hat{S}(n, \frac{N}{2}) df\, dt + E_L
\end{aligned}
$$

where

$$\tilde{E} = \sum_{n=0}^{N_f-1} \sum_{k=0}^{\frac{N}{2}-1} \hat{S}(n,k) df\, dt$$

is the energy value computed from the spectrogram and reported in the measurement panel as *Energy* [joules] (*Energy Flux Density* [joules / meter^2]),

$$E_L = \sum_{i=L}^{\tilde{L}-1} p[i] \Delta t$$

is the energy in the waveform between samples L and $\tilde{L}$, and

$$\hat{S}(n,k) = \frac{2\eta[k]T}{R\Delta N_f U^2} |X(n,k)|^2$$

is the measurement titled *Power / Hz* [watts / hertz] (*Intensity / Hz* [watts / meter^2 / hertz]). We also define the measurement *Energy / Hz* [joules / hertz] (*Energy Flux Density / Hz* [joules / meter^2 / hertz]) as

$$S(n,k) = \hat{S}(n,k) \cdot dt$$

For logarithmic (decibel) measurements, Canary reports $\max[L_C, 10 \log_{10}(S(n,k)/W_0)]$ and $\max[L_C, 10 \log_{10}(\hat{S}(n,k)/W_0)]$, where W_0 is the power reference level from the Calibration dialog box and L_C is the *clipping level*, from the Spectrogram dialog box.

The approximation $E \approx \tilde{E}$ has three sources of error. The first is the approximation

$$\lambda \approx 1$$

which is inaccurate because the spectrogram does not "satisfy the marginals" [2,6]. Other time-frequency distributions, such as the Wigner distribution, do satisfy the marginals, but are harder to compute and interpret than the spectrogram. Another source of error is the approximation

$$E_L \approx 0$$

which is inaccurate if the length of the waveform is not an integral number of frames, and if there is significant signal energy at the end of the waveform.

The third source of error is the approximation $\hat{S}(n, \frac{N}{2}) \approx 0$, which states that the Nyquist frequency bin holds negligible energy, and which holds for any properly digitized waveform. Waveforms with $\hat{S}(n, \frac{N}{2})$ appreciably different from zero are undersampled and will exhibit aliasing.

The Spectrum

Canary's spectrum is in fact an "averaged periodogram" [3,5,7,9,10]. Using the notation from the previous subsection, the formula for the spectrum is

$$S(k) = \sum_{n=0}^{N_f - 1} \hat{S}(n, k) dt$$

for the *Energy / Hz* [joules / hz] (*Energy Flux Density / Hz* [joules / meter^2 / hz]) measurement and

$$\hat{S}(k) = \frac{1}{T} S(k)$$

for the *Power / Hz* [watts / hz] (*Intensity / Hz* [watts / meter^2 / hz]) measurement.

Waveform correlations

In this subsection we consider the cross correlation of two input signals x and y, considered as vectors in $\mathbf{R}^N$, where N is the size of the larger vector, and the smaller is zero padded to size N. The formula for the basic cross correlation is

$$R_{xy}[n] = \sum_{m=-N}^{N} x[n+m]\, y[m],$$

where any negative vector indices evaluate to zero. If we let $x_n[m] = x[n+m]$, then

$$R_{xy}[n] = < x_n, y >,$$

where the bracket notation $< \cdot, \cdot >$ indicates the usual Euclidean inner product.

Normalized correlations

Canary's normalized correlation pre-normalizes the vectors x and y so that the maximum correlation magnitude value is 1. The formula for normalized correlations is

$$\hat{R}_{xy}[n] = \frac{< x_n, y >}{||x||\,||y||},$$

where $||x|| = \sqrt{\sum_m x[m]^2}$. If we use the notation $\hat{x} = x/||x||$, we have $\hat{R}_{xy}[n] = R_{\hat{x}\hat{y}}[n]$.

Correlations are best known in the signal processing literature as a theoretical tool for analyzing the spectra of wide-sense stationary stochastic processes [3,5,9,10]. Their use as a pattern recognizer in deterministic signals arises from the following relation for normalized correlations.

$$\hat{R}_{xy}[n] = 1 - \frac{1}{2}||\hat{x}_n - \hat{y}||^2$$

Therefore, a match between x_n and y will result in a peak value of 1 in $\hat{R}_{xy}$ at n. A value of $\hat{R}_{xy}[n] = 0$ means that x_n and y are orthogonal.

Filtered correlations

Canary's correlations are computed by taking the FFTs of the input vectors, multiplying one by the conjugate of the other, and computing the inverse FFT of the result. For filtered correlations, before multiplying the FFTs of the input signals, the frequency bins outside of the specified pass band are zeroed.

For normalized filtered correlations, the norms $||x||$ and $||y||$ are computed from the FFTs after zeroing the stop bands.

Complex envelope

The complex envelope of a real signal is the magnitude of the corresponding "analytic" signal. The analytic signal corresponding to $x[n] = a[n]cos(\omega n)$ is

$$\begin{aligned} z[n] &= a[n]e^{j\omega n} \\ &= a[n]cos(\omega n) + ja[n]sin(\omega n) \\ &= x[n] + j\mathcal{H}(x)[n], \end{aligned}$$

where $\mathcal{H}(x)$ is the *quadrature signal* or *Hilbert transform* [7,8,9] of x. The complex envelope of x is $|z[n]| = a[n]$. The original use of the complex envelope correlation was to find the peak of the correlation function of two signals containing the same frequency chirp component.

Spectrogram correlations [1]

If we let $S_x(n,k)$ and $S_y(n,k)$ denote the spectrograms of two signals x and y, the spectrogram correlation is

$$R_{S_x S_y}[n] = \sum_{m,k} S_x(n+m,k) S_y(m,k).$$

Normalized spectrogram correlations divide the spectrograms by $\sqrt{\sum_{n,k} S(n,k)^2}$ before correlating. Filtered spectrogram correlations zero out the stop band in the spectrograms before correlating.

Notice that the gross shape of most logarithmic spectrograms is roughly the same, with most of the interesting information contained in the upper dB values. This makes logarithmic spectrogram correlations roughly triangularly shaped, unless a high clipping level is used.

Filtering

Canary 1.2 has a filtering function, which allows the user to band-pass or band-reject a selected section of the spectrogram. Canary does not perform a traditional filter design to perform its filtering. Instead, Canary performs a projection operation. The selection of waveform data is FFT'd (this may be a very large FFT), and the corresponding frequency bins are zeroed. The boundary frequency bins are not zeroed, but are tapered according to the non-zero values of the Fourier transform of the Hamming window. This is similar in spirit to windowed bandpass filter design [7,8,11].

Calibration

Canary 1.2 has an improved calibration facility. In Canary 1.1, all decibels were relative to unity. In Canary 1.2, the user can specify independent dB reference values for both voltages and power. There are two calibration paradigms, "acoustic" and "electric", which perform analogously, with a simple renaming of the quantities and units involved. The waveform is measured in volts, while spectrogram and spectrum values are measured in power or energy. The conversion factor is the impedance, according to Ohm's law:

$$P = \frac{V^2}{R}$$

where P is power, R is impedance, and V is rms voltage.

In Canary 1.1 files with floating point formats (text files, Matlab files, and Canary files) held uncalibrated values, and so were hard to export. In Canary 1.2, exported text and Matlab files contain calibrated values.

References

[1] Altes, R.A., "Detection, estimation, and classification with spectrograms," *J. Acoust. Soc. Am.* 67, 1232-1246 (1980).

[2] F. Hlawatsch and G. F. Boudreaux-Bartels, "Linear and Quadratic Time-Frequency Signal Representations," *IEEE Signal Processing Magazine,* Vol 9. No 2, pp. 21-67, April 1992.

[3] Steven M. Kay, *Modern Spectral Estimation: Theory and Application,* Prentice-Hall, Englewood Cliffs, NJ, 1988.

[4] Jae S. Lim and Alan V. Oppenheim, *Advanced Topics in Signal Processing,* Prentice-Hall, Englewood Cliffs, NJ, 1988.

[5] S. Lawrence Marple, Jr., *Digital Spectral Analysis with Applications,* Prentice-Hall, Englewood Cliffs, NJ, 1987.

[6] W. Mecklenbrauker, "A Tutorial on Non-Parametric Bilinear Time-Frequency Signal Representations," in *Time and Frequency Representation of Signals and Systems,* G. Longo and B. Picinbono, ed, pp. 11-68, Springer-Verlag, New York, 1989.

[7] Alan V. Oppenheim and Ronald W. Schafer, *Digital Signal Processing,* Prentice-Hall, Englewood Cliffs, NJ, 1975.

[8] Alan V. Oppenheim and Ronald W. Schafer, *Discrete-Time Signal Processing,* Prentice-Hall, Englewood Cliffs, NJ, 1989.

[9] Athanasios Papoulis, *Signal Analysis,* McGraw-Hill, New York, 1977.

[10] Athanasios Papoulis, *Probability, Random Variables, and Stochastic Processes, 3rd ed,* McGraw-Hill, New York, 1991.

[11] T. W. Parks and C. S. Burrus, *Digital Filter Design,* Wiley, New York, 1987.

INDEX

D

E

F

G

H

I

L

M

N

O

P

Q

R

S

T

U

V

W

Z