

Chapter 10 Detection

About this chapter

When users work with large datasets, they often seek to first identify all sounds of a particular type, then they make measurements on the selected sounds and perhaps associate annotations with the sounds to classify them as to type or source. The process of manually browsing through hours, days, or months of recordings can be daunting.

An alternative to manually browsing through large datasets is to use automatic *detection*. Detection is the process of finding specific sounds of interest within recordings. Often, the signals of interest are short in duration compared to the overall recording. Detection can be a useful tool when compared to the alternative process of manually browsing days of recordings in order to find small samples of interesting sounds.

Automatic detection software may be able to reduce the amount of time that it takes to analyze a recording. Detection is accomplished through the use of a specific detector - the mechanism which runs a detection algorithm to detect specific signals of interest. Users must be aware of trade-offs that they will need to make before deciding if detection is the correct approach for their analysis.

In this chapter, you will learn about:

- the four modes of Raven detectors
- trade-offs involved with using detectors
- interactive detection
- full mode detection
- batch detection
- real-time detection
- saving detected selections
- instructions for detectors provided with Raven

The four run-modes of Raven detectors

In order to shorten the amount of time that users take to analyze their recorded datasets, Raven has a detector infrastructure built around Raven's extensible software architecture. Raven provides an Application Programming Interface (API) as well as a reference implementation of a specific detector meant to get users started with detection. Developers

interested in creating their own detectors should contact the Raven team at Raven@cornell.edu for more information.

Detectors in Raven can be run in one of the following four modes:

1. **Interactive detection:** Within a short recording or a set of pages of a long recording, users can test several sets of parameters in order to see how a detector will perform. Available in the View menu and in the context menus of applicable sound views, interactive detection searches through the paged area of a sound file to find areas that match the detector parameters. Results from multiple interactive detectors can be displayed simultaneously, and are color-coded for easy discrimination, allowing the user to run and compare multiple sets of detector parameters.
2. **Full detection:** Once specific parameters are chosen that work for short sections of a recording, Raven can run a full detection process on the entire recording even if only a single page is shown in the window. Located in the Tools menu and through the “Run Full” button in the interactive detector panel (see [Figure 10.1](#)), full detection is useful for bulk processing of large datasets.
3. **Batch detection:** If users want to find similar signals in many different sound files, Raven’s batch detector can be used to run a full detection process on many sound files at once. The same detection parameters will be used on each sound file, and Raven will automatically save a selection table of detection results for each sound file.
4. **Real-time detection:** When an interactive detector is opened in a recorder window, Raven can detect events as they occur by using a real-time detector. Real-time detection identifies signals of interest as they are recorded and marks them as selections. These events can then be reviewed in the recorded file sequence or saved as clips for later review. You can run real-time detection at the same time that you save to file-sequences, or you can post-process the file sequence by running a detector on the full sequence.

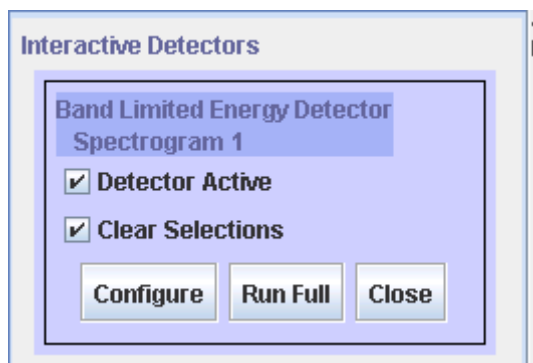


Figure 10.1. The Run Full button in the detector tab on the side panel

Raven comes with two sample detectors; a band limited energy detector, and an amplitude detector. These are described in more detail in [“Full detection mode”](#) on page 261. Both of these detectors require a specific sound view. The energy detector uses the spectrogram view, whereas the amplitude detector uses the waveform view. For interactive detection, the appropriate view must be active in order for the detector to be available through the view menu. When running a detector in full detection mode from the tools menu, Raven will prompt you to open the necessary view if it is not already open.

Interactive detection

Running a detector on an existing sound file First, open the sound file, making sure that the appropriate view is active in the window. Next, open the detector by selecting **View > Interactive Detectors... > [Detector Name]** from the toolbar, or by opening the view's context menu and selecting **Interactive Detectors... > [Detector Name]**.



If the recording is long, choose to open it one page at a time with a page size that is large enough to capture one or more of the signals of interest, but not so large that you run into memory issues. For further information on memory management, go to [“The Memory Manager”](#) in Chapter 11 (page 309)

Once a detector is selected, the configuration dialog will be displayed, allowing the user to set certain detector parameters. The parameters in the upper half of the configuration dialog are the same for all detectors, whereas the parameters in the lower half are specific to the individual detector. These general parameters are described below. For information on parameters specific to the 2 detectors included with Raven, see [“The Band Limited Energy Detector”](#) on page 274 or [“The Amplitude Detector”](#) on page 283. After setting the parameters, and clicking the OK button, the detections (if there are any) should appear in the sound window. [Figure 10.2](#) shows the Amplitude Detector configuration dialog.

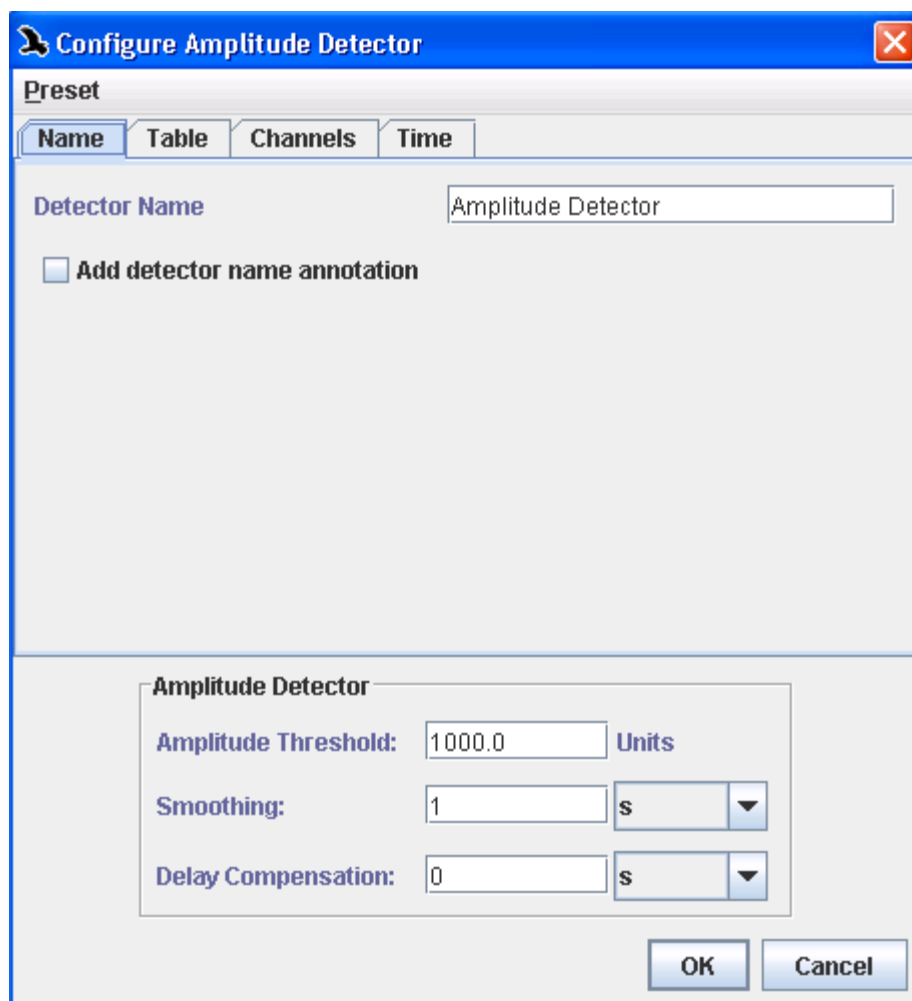


Figure 10.2. The configuration dialog for the Amplitude Detector. After setting the detection parameters, choosing OK will run the detector.

General Detector Parameters These parameters appear in the top half of the Configure Detector dialog and are the same regardless of which detector you are running.

Presets

A preset includes a set of predefined parameter settings. Under the **Presets** menu in the configuration dialog, you can apply an existing preset or save your current settings for future use. Raven comes with a default preset and a BlackCappedVireo preset for both the amplitude and energy detectors. You can also create your own presets for various applications. For more information on managing presets, see [“The Preset Manager” in Chapter 11](#) (page 297).

Naming

Under the **Name** tab of the configuration dialog, you can choose a name that will identify the detector in the side panel and become the name of the

created selection table. You can also opt to add a new annotation column to the created selection table(s) with this name which will remain with the selection if it is saved or moved to a different table. This can be useful if you are running multiple versions of a detector and would like to better differentiate the selections.

Tables

Under the **Table** tab, you can set the detector to automatically save the resulting selections periodically. You can choose to automatically save the selection table at the end of a detector run, or choose to automatically save the table after every *N* detections. You can also select whether or not you would like Raven to create a new selection table after each save. This will result in several smaller selection tables, which may be easier to work with than one large one. Finally, you can configure the naming and location of the resulting files.

In the box labeled **File Names**, you can use a template to customize the names your detector's selection table files will be given. For more information on how to use file naming templates, see "[Sound File\(s\)](#)" in [Chapter 4](#) (page 89). All of the tags listed in that section are available, and you can also use the following tags:

Table 10.1. Additional template tags available in the General Detector Parameters.

Tag	Definition
<f>	Filename of the sound file on which the detector will run
<t>	Integer sequence number of the table. Only relevant if you choose to create a new table after every automatic save.

Channels and Timing

For a multi channel sound (or multi channel recorder) you can choose to run the detector on all channels or a specific subset. You can also choose to run the detector on the entire sound, or on a shorter portion of the entire sound.

Displaying detector information

Once you've run the detector, controls for the detector are located on the **Detection** tab in the side panel. Also, if you make the selection table view of the current sound visible, the name of the active detector will appear on a new tab. The tab will be color-coordinated with the detector's selection boxes in the sound window and will display detections (selections created by the detector) in the selection table. See [Figure 10.3](#) for an example.

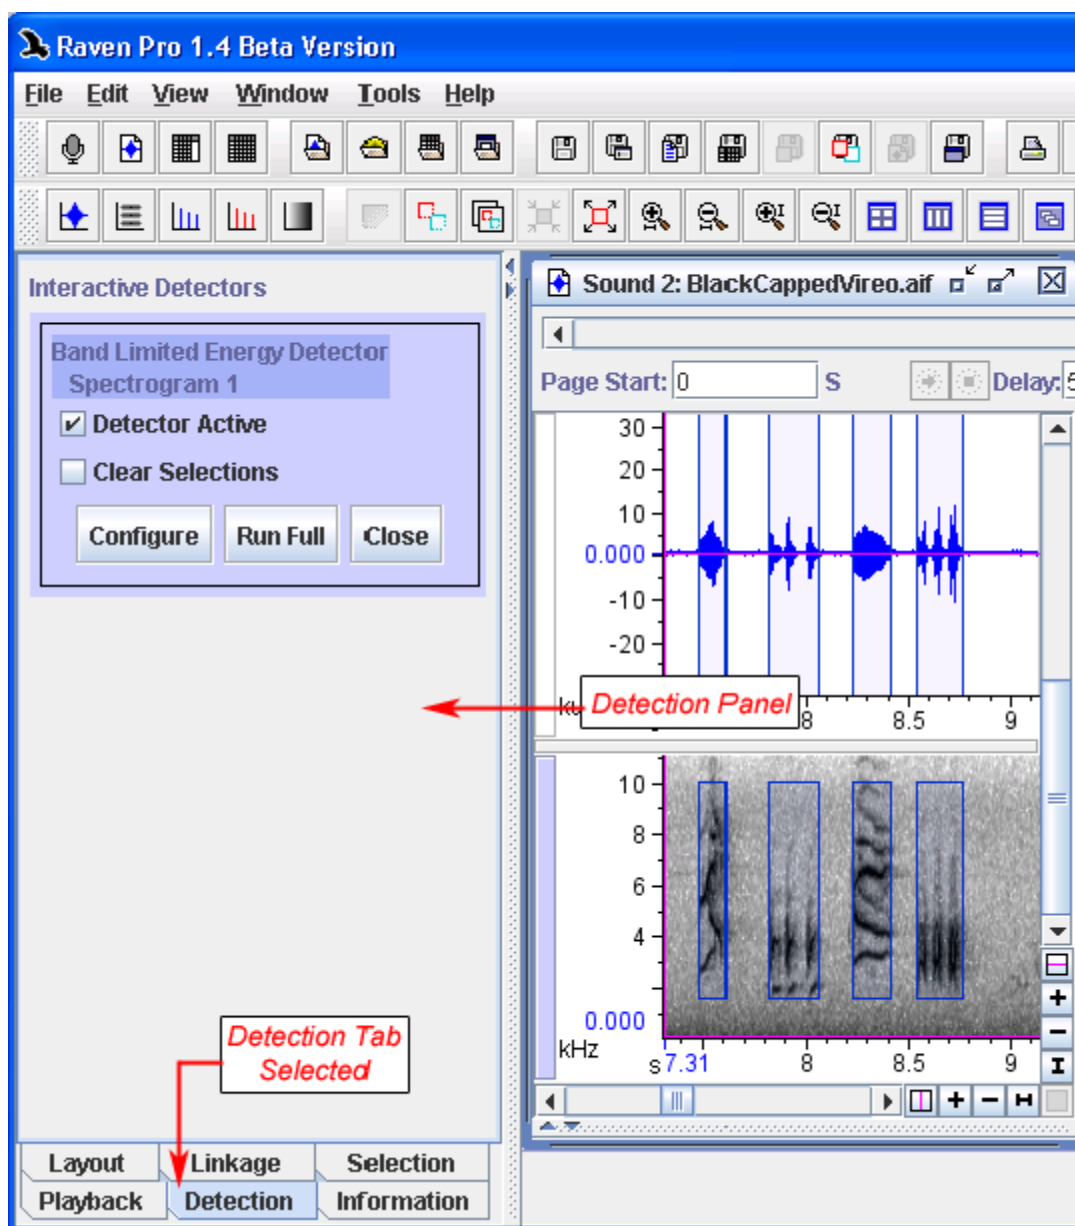


Figure 10.3. The detection tab, activated in the side panel, shows information about the current detector. Also, detection information is shown on its own tab in the sound window's selection table.

You can add and remove measurements from the detection selection table in the same way as you would with the default selection table. You can also rename the tab by opening the tab's context menu and selecting

Rename Table. See “Selection Tables” in Chapter 6 (page 150) for more information.



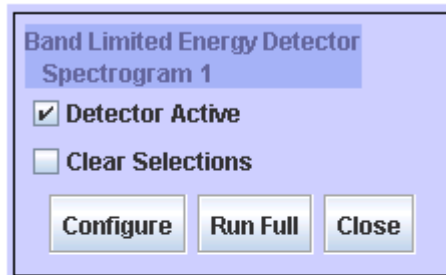
After running a detector, the detector selections tab will become the current active tab. This means that if you make a manual selection in the sound window, the data will be added to the detector tab and not the default selections tab. To add selections to the default tab, you must make it the active tab by clicking on it before making selections

Running an interactive detector in a paged sound

To test a detector on a long sound, first open the sound in a paged sound window, then configure an interactive detector as described previously. This will run the detector on the current page. To see the results on different pages in the sound recording, use the page scroll bar at the top of the screen to advance one page in the recording. The detector will then create selections for that page based on the input detection parameters.

Once you are satisfied with the results of a detector, you can choose to run it on an entire recording. To do this, you can press the Run Full button in the corresponding box on the detector tab in the side panel

a) Paged Sound



b) Non Paged Sound




Figure 10.4. The detector box as it appears for (a) a paged sound, and (b) a non paged sound. For a paged sound, the detector active and clear selection options and the Run Full button appear.

Detection options

If you display the contents of the detector tab in the side panel, you will see some boxes containing basic information about the detector(s)

currently open. These boxes also give you the ability to access or change some other detector information.

Configure: This button allows you to change the detector parameters via the standard configure dialog box. Once the parameters are changed and the OK button is pressed, the detector will be run again on the currently displayed page or the entire sound if it is all in memory. Note that reconfiguring the detector will cause all current selections to be deleted.

Run Full: Pressing this button causes Raven to run the detector, with the current set of parameters, on the entire sound recording. The results are stored in the selection table associated with this detector. While the detection is running, the **Configure** button will be grayed out. When the detection is completed, the **Configure** button will become accessible again.

Detector Active: Checking this item box means that the detector will be run on each page as it is opened. Unchecking the box means that the detector will not run on subsequent pages as you display them.

Clear Selections: Interactive detectors are one way to test a set of parameters. If testing parameters is your only goal, you may not want to have all selections from every page saved. Checking the Clear Selections box forces Raven to clear all old selections from the detector's selection table each time the detector runs. Note that reconfiguring the detector will clear all selections, even if the Clear Selections box is unchecked.



The Run Full button, Detector Active box, and Clear Selections box are only available when a detector is configured in a paged sound. Otherwise, you will not be able to select these options

Running multiple interactive detectors

You can configure multiple interactive detectors and run them simultaneously in Raven. Each will be given its own tab in the selection table with a corresponding color for the selections in the sound window (Figure 10.5).

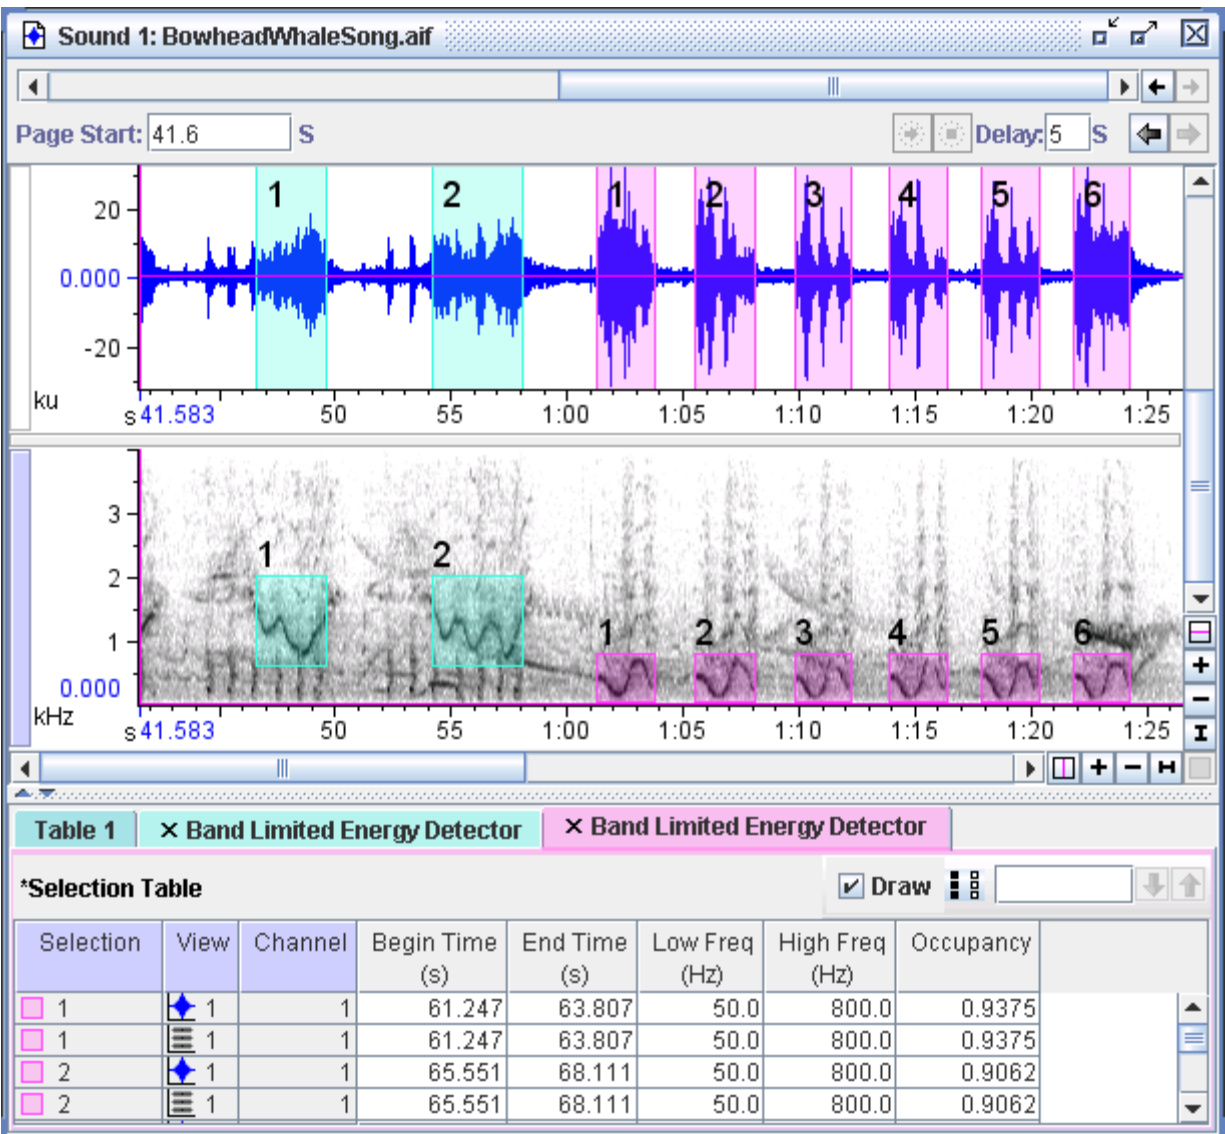


Figure 10.5. Two detectors that have been run on the same spectrogram view. The first detector (green selection table tab and green selections) is in a higher frequency range than the second detector (pink tab and selections).

Full detection mode

Once you are satisfied with the results of a detector, you can choose to run it on an entire recording. To do this, you can press the Run Full button in the corresponding box on the detector tab in the side panel. Or, you can use the Tools > Detector menu to start running a detector on an entire file or file sequence.

An alternative is to bypass interactive mode and run the detector on the full sound initially. Use the **Tools > Detector** menu to choose the detector you want. The dialog in [Figure 10.6](#) will appear, allowing you to choose an existing file or open a new file (Raven will open any sound that is not already open). This mode of detection can be used to post-process recordings; by first recording to a file sequence, you can then run a detector on the entire sequence at once.




Figure 10.6. The detection input dialog for the Amplitude detector. This allows you to choose the full sound on which you'd like to run the detector. For other detectors the dialog may look different: it may ask specifically for a waveform or spectrogram signal.

Batch detection

Once you come up with a set of detector parameters that are effective at finding a certain type of signal, you may wish to run the detector on many sound files at once to look for that kind of signal. Raven's batch detector allows you to automatically run the same detector on an arbitrarily large number of files and save the results. Once you set up the batch detection process, it runs in the background so that you can continue to do work in Raven while it progresses.

To run the batch detector, choose **Tools > Batch Detector...** from the menu bar. This will display the batch detection dialog (Figure 10.7).




Figure 10.7. The batch detection dialog.

Choosing files To perform a batch detection run, begin by choosing the files on which you'd like to run the detector. Click the **Add** button to select the files you'd like to add. You can select a single file by clicking on it, or multiple files by holding down **Control** or **Shift** while you click. You can add more files by clicking **Add** again. The files you choose do not have to be from the same folder, and they do not need to be in the same format, as long as they are all compatible with the detector parameters that you specify.

After you've added files, their names will show up in the **Files** box. If you'd like to remove one of the files you added, click on its name in the **Files** box and then click **Remove**. You can also remove all the files currently in the **Files** box by clicking **Clear All**.

When choosing files, you can add either individual sound files or list files. For more information about list files, see ["Using a list file" in Chapter 7](#)

(page 195). If you add a list file, the detector will run back-to-back on all the files in the list, and save a single selection table with the results from all of those files.

Choosing a detector Once you've chosen your files, choose which type of detector you'd like to run using the dropdown menu labeled **Detector**. You can choose between the Amplitude Detector and the Band Limited Energy Detector. For more information about those detectors, see ["The Amplitude Detector" in Chapter 10 \(page 283\)](#) and ["The Band Limited Energy Detector" in Chapter 10 \(page 274\)](#).

Configuring the detector After you've chosen your files and your detector, you can configure the detector by clicking the **Configure Detector** button. This will display the detector configuration dialog. The dialog displayed will differ based on whether you've selected the Amplitude Detector or the Band Limited Energy Detector. If you'd like to use a detector configuration you've set up before, or save the detector configuration you're currently using for future use, you can use the Preset menu (see ["Presets" in Chapter 10 \(page 256\)](#)).

For more information about how to configure the detector you've chosen, please see ["Configuring the Amplitude Detector" in Chapter 10 \(page 283\)](#) or ["Configuring the Band Limited Energy Detector" in Chapter 10 \(page 274\)](#).

Saving the detection results To choose where to save the results of your batch detection run, click on the **Table** tab in the top panel of the detector configuration dialog ([Figure 10.8](#)). Note that the checkbox marked **Automatically save table after detector run** is already checked. During the batch detection run, the detector will create a selection table of results for each file you select, and save each of those tables in the place that you specify. Choose the folder in which those results will be saved by clicking the **Browse** button next to the **Save Directory** box. Then enter a template for the tables' names in the **File Names** box. We recommend that you use the symbol `<f>` somewhere in your template; in the file names that are created, that `<f>` will be replaced with each sound file's file name. For more information about file name templates, see ["Tables" in Chapter 10 \(page 257\)](#).



Warning: if you do not include the `<f>` symbol in your template, then multiple selection tables may be given the same name. In that case they will write over each other, and some of the results from your batch detection run will be lost.

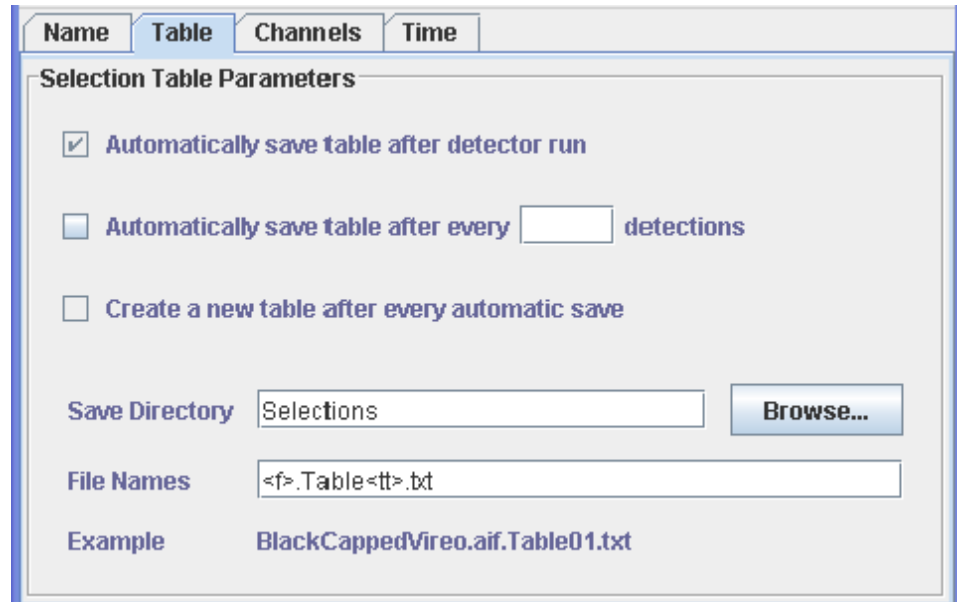


Figure 10.8. The Table tab of the detector configuration dialog.

- Configuring the view** If you choose the Band Limited Energy Detector, then the Configure View button will become available. Clicking it will display the spectrogram configuration dialog. Configuring the view is optional; if you choose not to configure the view, the batch detection run will use the default spectrogram configuration. If you do choose to configure the view, see [“Configuring spectrographic views” in Chapter 5](#) (page 112) for more information.
- Running the batch detector** Once you’ve chosen a list of files and a detector, configured the detector, and optionally configured the view, you’re ready to begin the batch detection run. Click the OK button at the bottom right of the batch detection dialog. If the detector configuration you’ve chosen is not compatible with the first file in your list of files, or if you’ve configured the view in a way that makes your detector configuration invalid, a warning will pop up asking you to change your configuration. Once these issues are resolved, the batch detection run will begin.

When the batch detection run begins, Raven will display the Progress Manager ([Figure 10.9](#)). The first panel in the Progress Manager, labeled “Starting X detectors...”, keeps track of the batch detection run as a whole. When the progress bar in that panel is complete, that means that individual detection runs have been started on each of the files you’ve chosen. If some of the files have to be skipped, either because Raven is unable to open them or because they are not compatible with the detector configuration you’ve chosen, their names will appear as a list of skipped files in the first panel.

As each individual detection run begins, a new panel for that run will appear in the progress manager. That panel will track how many detec-

tions have been made, how much data has been processed, and the detection speed so far. When the progress bar in one of those panels is complete, that means the detector has finished running on the file listed at the top of the panel.




Figure 10.9. The Progress Manager at the end of a batch detection run. Note the list of skipped files in the first panel.

When all the progress bars have completed, the batch detection run is finished. The results of each detection are stored in the folder that you selected in the **Table** tab of the detector configuration dialog.

At any point during the batch detection run, you can continue using Raven for other tasks. If you'd like to make the Progress Manager visible again, select **Window > Progress Manager** from the menu bar.

Real-time detection

Running a detector within a recorder You can also run a detector within a recorder. To do this, first open a recorder window, and then configure an interactive detector that you would like to run. Next, start the recorder and watch as the detector finds signals according to the configured parameters, and creates corresponding selections in the recording.

There are three ways to process detections from a recording:

1. **Post-processing:** This process involves running a recorder to a file sequence, and then running a full detector on the file sequence after the recorder has finished running. For this, you must record to a file sequence completely, open the file sequence in a sound window, and then run a detector on the file sequence.
2. **Real-time processing while saving the entire recording:** In this case, you will save the entire recording but will run the detector while the recorder is running. You can then save the resulting selections as a selection table that can be opened in the saved sound, allowing you to review the detections within the larger context of the sound at a later time. To set up this type of recording session, run the recorder to record to a file sequence, making sure to record to file (red button).



To ensure that the times in the selection table correspond to the times in the file sequence, it is important to start recording directly to a file sequence. Starting to record a file sequence after recording to memory for a short time will result in a discrepancy between these times as the time shown in the selection table corresponds to the time since the start of recording, not the time since the start of the file sequence.

3. **Real-time processing without saving the entire recording:** Here, you set the recorder to record to memory and you also configure the recorder to save any selections as sound clips. With this method, you are not saving the entire recording; you are only saving those selections that the detector identifies. Unlike saving the detections as a selection table within a saved file sequence, this method saves only the detections, without the surrounding signal. Therefore, this method will not allow any post-processing on non-selected data, so you should only use this technique when you are confident that the detector is working as anticipated to avoid losing data.

Example of running a detector within a recorder This example shows, step by step, how to run a recorder with a configured detector that saves clips to files, and then how to open and browse the results.

First, configure a recorder to record from your selected audio input device. Configure it to record to memory and to save clips to files. The exported clip files will be saved as individual sound files and as a list file (see [Figure 10.10](#)).

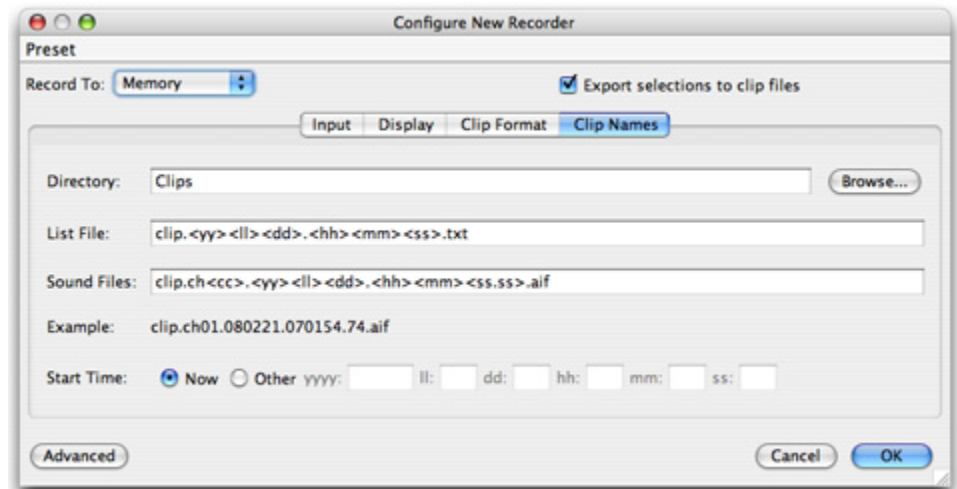


Figure 10.10. Configuring a recorder to record to memory and save clip files to disk. Note that the names of the files are saved to a list file.

Next, configure a detector to find the signals that are of interest to you. In this example we configure a detector to find short whistles. These settings are shown in [Figure 10.11](#). (For more information on these parameters, see the next section, “[The Band Limited Energy Detector](#)” on page 274.). After you configure the detector, a new selection table will be automatically created in the recorder window and a box for the detector will appear in the Detection tab of the side panel as shown in [Figure 10.12](#).

The figure displays two screenshots of the 'Configure Band Limited Energy Detector' dialog box, showing the 'Target' and 'Noise' tabs.

Target Tab (Target Signal Parameters):

- Minimum Frequency: 3000 Hz
- Maximum Frequency: 4500 Hz
- Minimum Duration: 0.09868 s
- Maximum Duration: 0.99846 s
- Minimum Separation: 0.09868 s

Noise Tab (Signal-to-Noise Ratio Parameters):

- Minimum Occupancy (%): 30.0
- Return Occupancy %
- SNR Threshold (dB): 10.0

Noise Tab (Noise Power Estimation Parameters):

- Block Size: 1.99692 s
- Hop Size: 0.49923 s
- Percentile: 20.0

Figure 10.11. The Target and Noise Tabs in the Configure Band Limited Energy Detector dialog with parameters to find short whistles

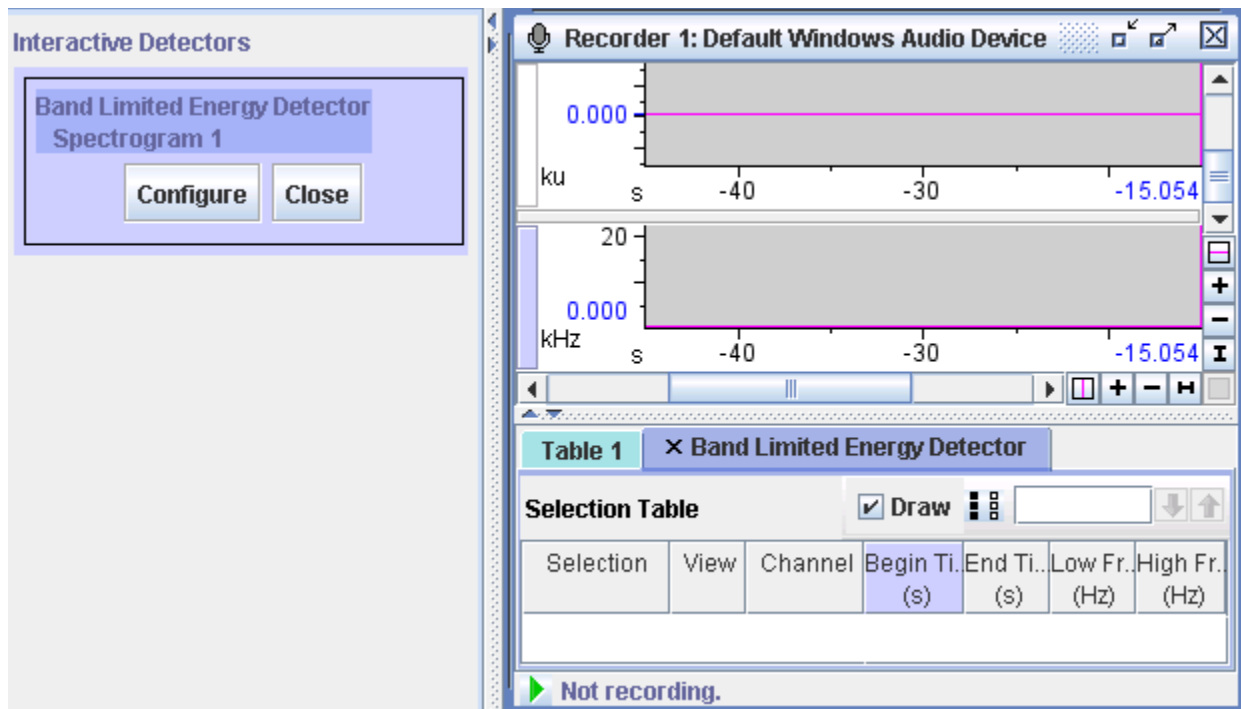


Figure 10.12. A configured recorder with a selection table created for the detector.

Turn on the recorder and begin recording a sound that includes some high whistles. Notice that as the detector identifies these whistles, they are marked in the spectrogram view, appear in the selection table, and are also exported to clip files in real time

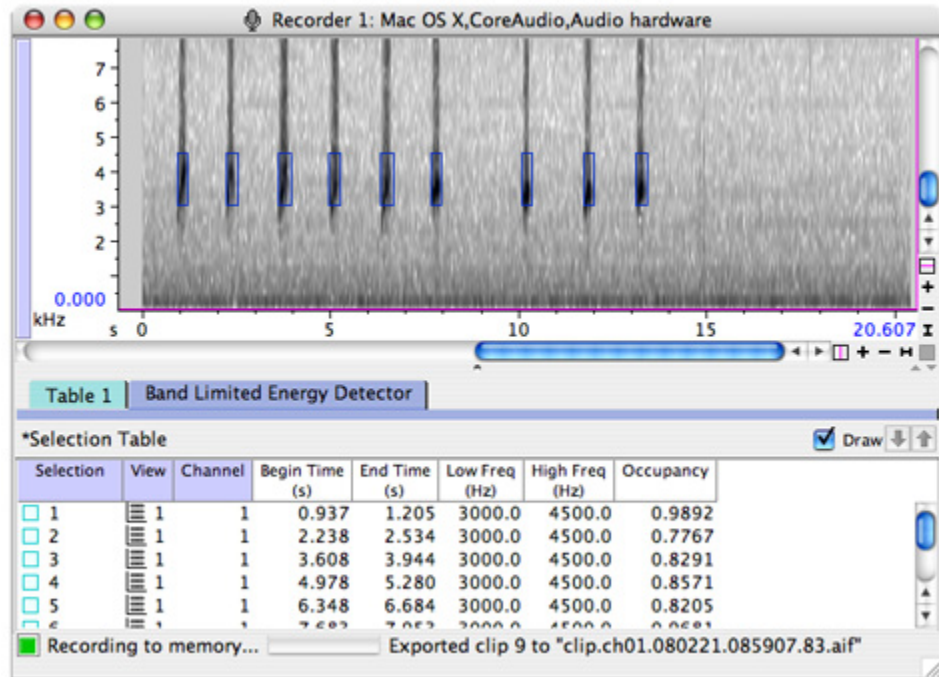


Figure 10.13. A running detector saving detected clips to files.

To view the clip files and confirm that they are saved correctly, open the saved list file by going to **File > Open Sound Files...** and selecting the list file. If the default naming conventions were used when configuring the detector, this file is saved to the Clips folder within the Raven folder and will be named with the date and time of recording (see [Figure 10.14](#)). When prompted, choose “Open as file sequence in one window” to view all the detected selections in sequence. The clips should be opened in a new sound window similar to the one in [Figure 10.16](#).

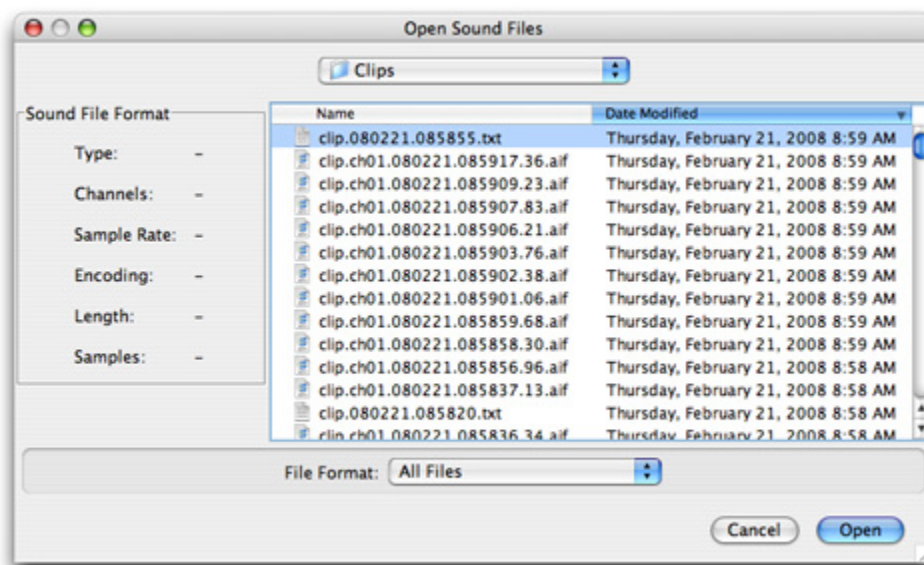


Figure 10.14. Selecting a clip list file to open.

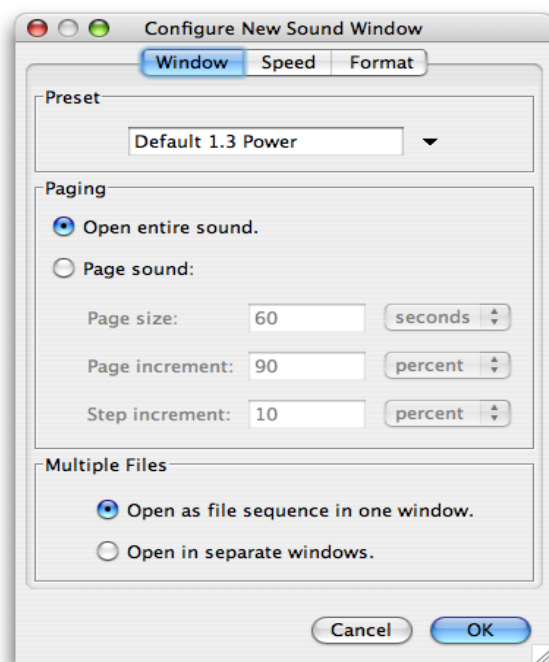


Figure 10.15. Choosing to open the list as a file sequence as opposed to opening each sound in a separate window.

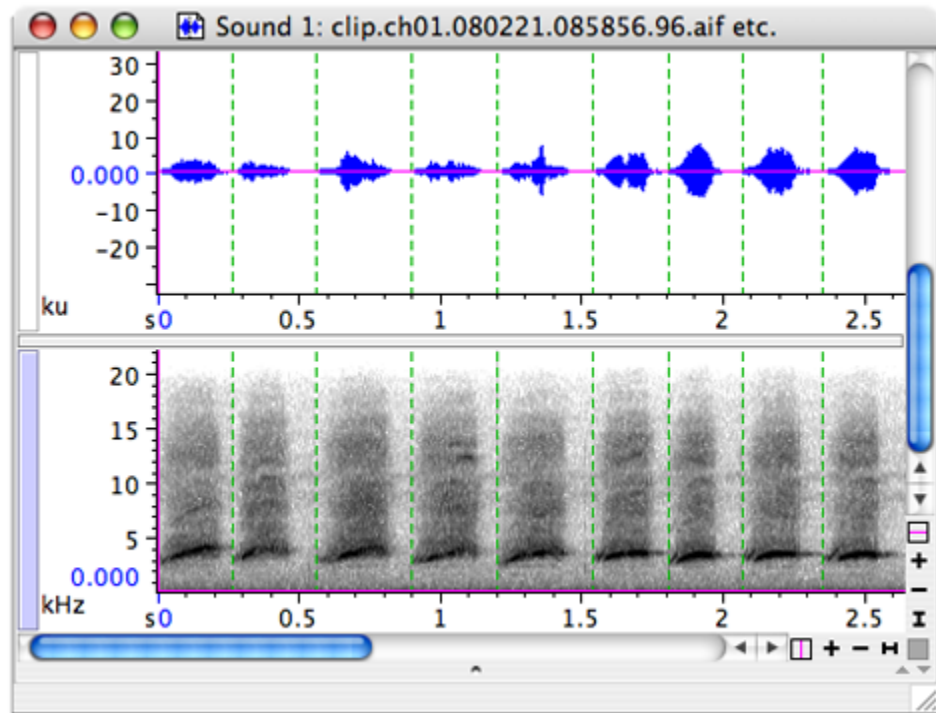


Figure 10.16. The clip files, opened in a file sequence.

Saving detected selections

Once you are satisfied with the selections created by your detector, you can save your selections as a text file by selecting **File > Save Selection Table name] As...** or **File > Save Selection Table "[Table name]"**. In the Save dialog, you must provide a name and path for the text file to be saved. The default folder is the Selections folder within the Raven folder.

To save sound clips of the selections, you must select **File > Save All Selections As...** or you can use the shortcut **<Ctrl-Shift-A>** (Windows, Linux), or **<Command-Shift-A>** (Mac OS). Detections can also be saved while the detector runs by configuring the recorder to export selections as clip files and then running the detector. As the detections get marked as selections, the recorder exports these selections as clips.

In addition to saving the detections, you may wish to save the detector settings as well, in case you wish to run the detector again on the same or different data. Some Raven detectors provide a preset menu to allow you to save parameters. If a preset menu is not provided, then you should record these settings on your own.

The Band Limited Energy Detector

Band Limited Energy Detector (Energy Detector)

The Energy Detector supplied in Raven is a time-frequency energy detector. It estimates the background noise of a signal and uses this to find sections of signal that exceed a user-specified signal-to-noise ratio threshold in a specific frequency band, during a specific time. First, the detector makes a rolling estimate of the background noise, which is then used to identify potential detections. Next, the detector compares individual frames of a signal to the running noise estimate, selecting those that exceed the specified threshold as candidates for detection. Finally, Raven analyzes the distribution of these selected frames to identify regions that satisfy the detector's duration and occupancy parameters. The energy detector runs on a spectrogram view only. In order to run it, the associated view must be open or Raven will prompt you to open the required view before running the detector.

Since this detector makes a rolling estimate of the background noise in the band of interest, it is more useful when there is more uniform or "white" noise, because the estimate can provide an adaptive threshold from which real signal can be estimated. In more detail, the noise estimation process begins by summing up the energy in the band for each frame in the previous "block size" seconds. Then, it takes this list of energies and sorts them before picking one according to the "percentile" parameter. For example, a "30" in the "noise estimate percentile" parameter will take the energy value 3/10th of the ways up from the bottom of the list. Likewise, a value of "50" would result in the median value being chosen. This number is then taken to be the noise estimate for that block, at which time the "block" is advanced according to the hop size and the process is repeated. So, if the noise block size is 5 seconds and the hop size is 1 second, the next noise estimate will re-use the last 4 seconds, plus one more. This gives a smooth "rolling average".

A few consequences of this procedure are that a broad band setting will result in a very small range of energy values for a block, resulting in a very small difference between the noise estimate and the "highest" energies. Also, to make a selection, the in-band energy needs to exceed the noise estimate (by the SNR threshold amount) and then drop below it again according to the duration/separation parameters. So, a noise estimate that is too low can result in fewer selections, as can one that's too high. This makes the "percentile" number very important.

Configuring the Band Limited Energy Detector

To configure and run the energy detector, go to **View > Interactive Detectors > Band Limited Energy Detector**. This will open the Configure Band Limited Energy Detector dialog box, which allows you to adjust the detector parameters to isolate specific signals. In the upper portion of the dialog box, you can set the general detector settings as described in ["General Detector Parameters"](#) on page 256. The lower portion of the dialog contains settings specific to the energy detector, which are described below.

Target signal parameters

Under the **Target** tab, you can specify a range of frequencies on which to run the detector and set length requirements for the detected signals.

Minimum and Maximum Frequencies (Hz, kHz): These determine the range of frequencies in which the detector searches. The frequency band should be wide enough to include any potential detections, but narrow enough to exclude large bands of unnecessary noise. Since potential detections are compared with the total noise level in the band, a wider band setting can cause the detector to be less effective.

Minimum and Maximum Duration (sec, msec, samples): These parameters specify the length of signal that can be considered a single detection. These values should be set to include the full length of detections you wish to identify. However, choosing a higher value for the minimum duration reduces the likelihood that the detector will identify a short segment of noise as a target signal. Analogously, choosing a lower maximum duration prevents sustained tones in the frequency range from being classified as detections.

Minimum Separation (sec, msec, samples): The minimum separation corresponds to the amount of time between adjacent detections. If your desired selections include short groups of sounds, this parameter can be used to determine whether the individual elements are grouped together as a single detection, or classified as multiple smaller ones.

Noise Estimation Parameters

The settings in the **Noise** tab control how Raven estimates the background noise level. They also allow you to specify how loud a detection must be with respect to the background noise.

Minimum Occupancy (%): The minimum occupancy defines the percentage of samples within a selection that must exceed the background noise by the SNR threshold in order for the signal to be considered a detection. This prevents a brief section of louder noise from being considered a detection when only a small portion of it exceeds the SNR threshold.

Signal to Noise Ratio (SNR) Threshold (dB): The SNR threshold establishes how far above the background noise a sample must be to be considered a signal. Whereas the minimum occupancy determines how much signal must be present in the detector's range, the signal to noise ratio defines the energy level that constitutes a signal above the background noise.

Noise Block Size (sec, msec, samples): The block size corresponds to the width of the block used to calculate the background noise level. Choosing a block size substantially longer than the duration of a potential detection prevents Raven from calculating noise from a selection that contains mostly signal. However, if the background noise in your sound varies

considerably, a shorter block size will better account for this variation. Note that if you attempt to run the detector on a segment of time smaller than the noise block size, it will return no results.

Noise Hop Size (sec, msec, samples): The hop size determines how far the noise block steps each time it recalculates. As a result, a shorter hop size will result in more smoothing of the background noise levels. Choosing a hop size longer than a typical detection, but shorter than the block size generally provides the best results.

Noise Percentile (%): The noise percentile setting determines the percentage of ranked noise that is counted as background noise. A higher noise percentile setting results in a higher noise estimate, whereas a lower noise percentile results in a lower noise estimate. The appropriate value for a given application will depend on the specific characteristics of your sound.

Exclusion band parameters

Raven can exclude selections that might otherwise be considered detections through the use of an exclusion band. You can choose to ignore selections that exceed the background noise by a given level within a specified frequency range. This may be useful if your sound contains signals that closely resemble the desired detections.

Minimum and Maximum Frequencies (Hz, kHz): These determine the range of frequencies that are included in the exclusion band. You can choose to exclude signals within a subset of the overall detection range, or those with energy in a specific range outside the range of the detector. It is possible to introduce new detections by using an exclusion band within the range of the detector. This is because some selections that did not occupy enough of the bandwidth before may now exceed the minimum occupancy within the reduced frequency range.

SNR Threshold (dB): The SNR threshold in the exclusion band affects how sensitive the detector is to signals in the exclusion band. A low threshold value excludes potential detections which contain even a weak signal in this range, whereas a higher threshold value only excludes selections with considerable energy in the exclusion band.

Bandwidth filter parameters

The bandwidth filter in the energy detector allows you to target selections with particular frequency distributions. Using these settings, you can choose how wide of a frequency band individual selections should occupy. This feature can be useful when searching for detections that occupy a narrow bandwidth, but may be located across a larger frequency range.

Minimum and Maximum Bandwidth (Hz, kHz): These values correspond to the allowed bandwidths for individual detections. Setting

these parameters determines the minimum and maximum range that can be occupied by the amount of energy specified in the energy percentile setting. For example, specifying an energy percentile of 50% and minimum and maximum bandwidths of 100 Hz and 200 Hz respectively would require that 50% of the energy in each detection be concentrated in a range wider than 100 Hz, but narrower than 200 Hz. However, this span may occur anywhere in the detector's range

Energy Percentile (%): The energy percentile corresponds to the percentage of the total energy in a signal that must be concentrated within the specified range of bandwidths.

Band Limited Energy Detector Example To demonstrate the use of the Band Limited Energy Detector, we will walk through an example, demonstrating how to configure an effective set of detector parameters for the file `BlackCappedVireo.aif`. Raven comes with these parameters saved as the `BlackCappedVireo.energy` preset. [Figure 10.17](#) shows the Band Limited Energy Detector configuration dialog with this preset loaded.

The figure displays two screenshots of the Band Limited Energy Detector configuration dialog. The top screenshot shows the 'Target Signal Parameters' tab with the following settings:

Parameter	Value	Unit
Minimum Frequency	1500	Hz
Maximum Frequency	10000	Hz
Minimum Duration	0.05224	s
Maximum Duration	0.59791	s
Minimum Separation	0.05224	s

The bottom screenshot shows the 'Signal-to-Noise Ratio Parameters' and 'Noise Power Estimation Parameters' tabs. The SNR parameters are:

Parameter	Value
Minimum Occupancy (%)	40.0
Return Occupancy %	<input checked="" type="checkbox"/>
SNR Threshold (dB)	10.0

The noise power estimation parameters are:

Parameter	Value	Unit
Block Size	1.99692	s
Hop Size	0.49923	s
Percentile	20.0	

Figure 10.17. The Band Limited Energy Detector configuration dialog with the BlackCappedVireo.energy preset loaded.

Choosing the minimum and maximum frequencies is straightforward: choose an example of a typical event you wish to detect. Make a selection which encompasses a significant fraction of the energy of this event, and read the Low Freq and High Freq measurements from the selection table. In the case of the Black-capped Vireo sound, the goal is to detect several different types of events in one detector run. To achieve this, draw

selections around each of the event types you wish to detect, and set the minimum frequency equal to the lowest Low Freq measurement obtained, and the maximum frequency equal to the highest High Freq measurement obtained. See [Figure 10.18](#).




Figure 10.18. Finding the minimum and maximum frequencies. Note that the least Low Freq measurement is around 1500 Hz and the greatest High Freq measurement is around 10000 Hz. These correspond to the preset's minimum and maximum frequency values.

To choose values for the minimum and maximum duration parameters, add the Delta Time measurement to your selection table. Find the shortest duration event you wish to detect and draw a selection around it. Read the Delta Time measurement from the selection table, and use this value for the minimum duration (for an example, see [Figure 10.19](#)). To obtain the

maximum duration parameter, repeat the procedure with the longest duration event you wish to detect.



For settings with a time dimension, Raven automatically rounds values entered in seconds or milliseconds to the nearest frame size. For example, a value of 1 second may automatically be rounded to 1.00136 seconds. This change will appear immediately after the affected setting is entered.,

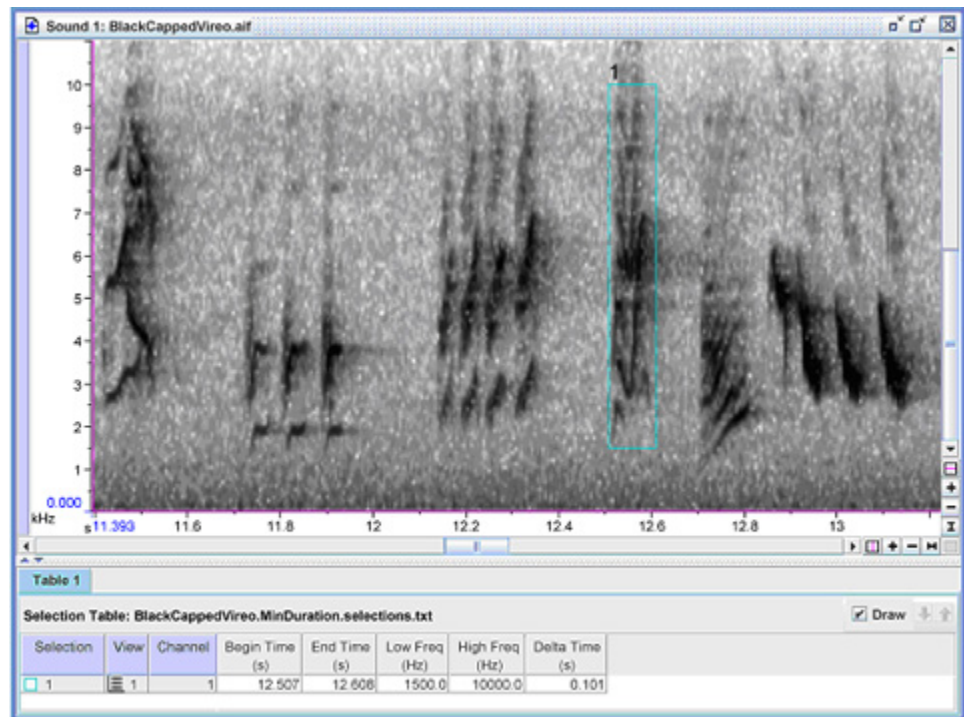


Figure 10.19. Finding the minimum duration. Note the minimum Delta Time measurement of 0.101 seconds.

To choose a value for the minimum separation parameter, find the two closest distinct events and draw a selection between them. Read the Delta Time measurement from the selection table, and use this value for the minimum separation.

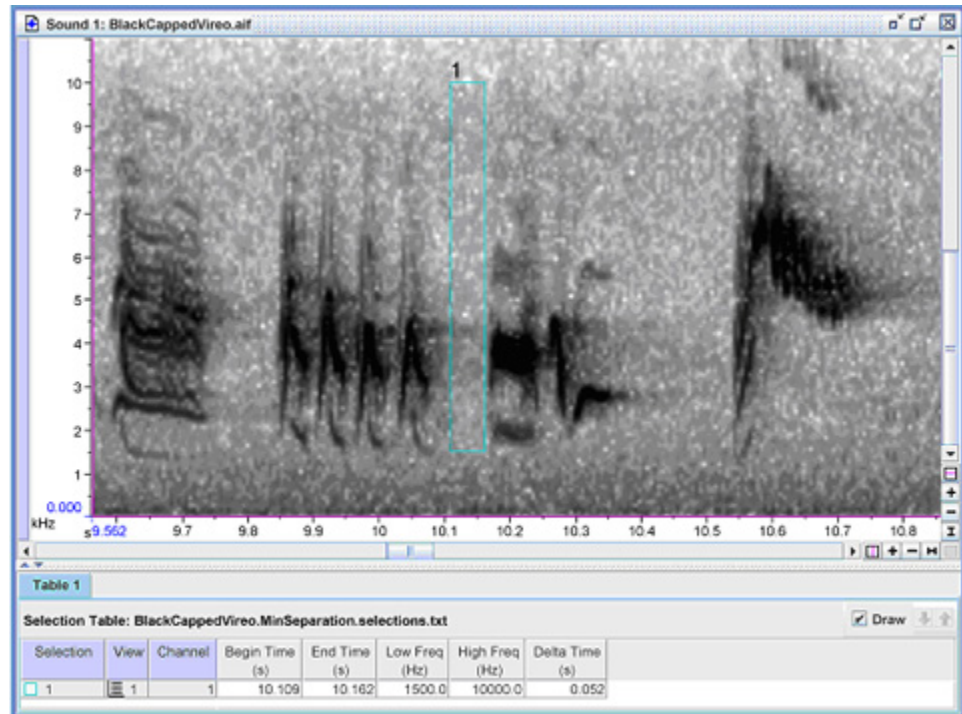


Figure 10.20. Finding the minimum separation.

Skip over the minimum occupancy and SNR threshold parameters for now, and set the Noise Power Estimation parameters. Pick a block size that is longer than the maximum duration by at least a factor of 3, and longer if the events cluster together. The goal is to have a long enough block size that the data being used for noise estimation is guaranteed to always include noise. If you make the noise block too short, it may result in noise estimates based primarily on only signal, rather than signal and noise together. In the BlackCappedVireo.energy preset, the block size is about three times the maximum duration.

The hop size parameter determines how much smoothing is performed in the noise estimation computation. Choose a hop size value somewhere between the maximum duration and the noise block size. The BlackCappedVireo.energy preset has the hop size set to 0.499 seconds, slightly less than the maximum duration of 0.597 seconds.

An appropriate value for the noise percentile parameter depends on the characteristics of the noise and clutter in your data. A higher percentile will result in a higher noise estimate. For the BlackCappedVireo.energy preset, the default 20th percentile was used, and appears to give a reasonable noise estimate. You may wish to play with higher and lower values to see which gives the best result for your data.

Once the Noise Power Estimation parameters have been determined, go back to the SNR threshold parameter. For an approximation of this

parameter, you can measure the average power of a typical noise block. [Figure 10.21](#) shows a typical detection and a typical noise block. The difference in average power is 10.5 dB. The BlackCappedVireo.energy preset rounds this to 10 dB.

Note that the noise block in [Figure 10.21](#) actually contains a large amount of signal. This is why the SNR threshold as computed here is so low. Since Raven uses all the sound in each block to estimate background noise, both the distribution of signal within the sound and the block size will affect the calculated SNR. In this example, the calculated SNR is significantly lower than the true SNR due to the rapid succession of calls.

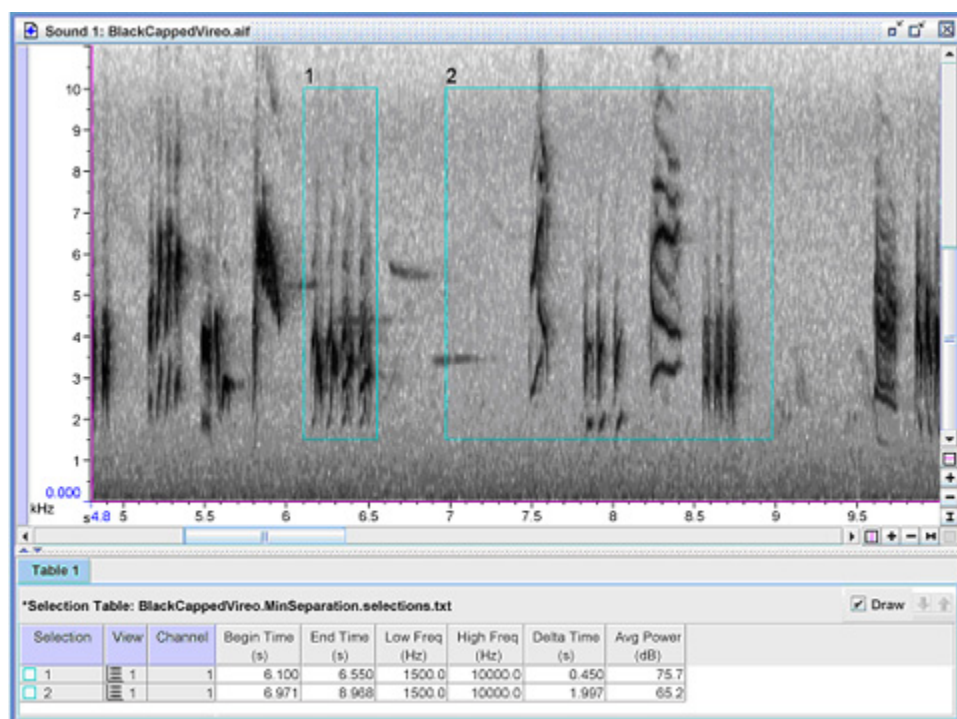


Figure 10.21. Measuring the estimated signal-to-noise ratio. Selection 1 contains only signal. Selection 2 represents a typical block of data used to estimate the noise power.

Accept the default minimum occupancy, and check the Return Occupancy checkbox. Then run the detector. Examine the detections and determine if the performance was satisfactory; if not, adjust the minimum occupancy downward and run the detector again. Sorting the selection table by occupancy can serve as a guide to setting the minimum occupancy parameter. Note that the BlackCappedVireo.energy preset has a minimum occupancy of 40%, yet it yields detections with occupancies no lower than 72.97%. The preset's minimum occupancy could have been set as high as 72% and returned the same results.

When citing use of Raven's Band Limited Energy detector in scientific publications, please refer to it as follows:

Mills, H. G. 2000. Geographically distributed acoustical monitoring of migrating birds. *Journal of the Acoustical Society of America*, 108(5, part 2): 2582.

The Amplitude Detector

The Amplitude Detector detects regions of a signal where the magnitude of the waveform's envelope exceeds a threshold value. To approximate the envelope, the detector computes the absolute value of the waveform and passes it through a lowpass smoothing filter. The envelope signal is then tested against the threshold. Raven does not display the actual envelope data.

Configuring the Amplitude Detector

To configure the detector, first choose the amplitude threshold based on the waveform amplitude. Then set the smoothing time constant equal to the most rapid rise or fall of the waveform that you would like the envelope detector to follow (see [Figure 10.22](#)). After running the detector, you may notice detections that occur slightly after the waveform actually crosses the desired threshold. To minimize this delay, set the delay compensation parameter to the size of the offset.

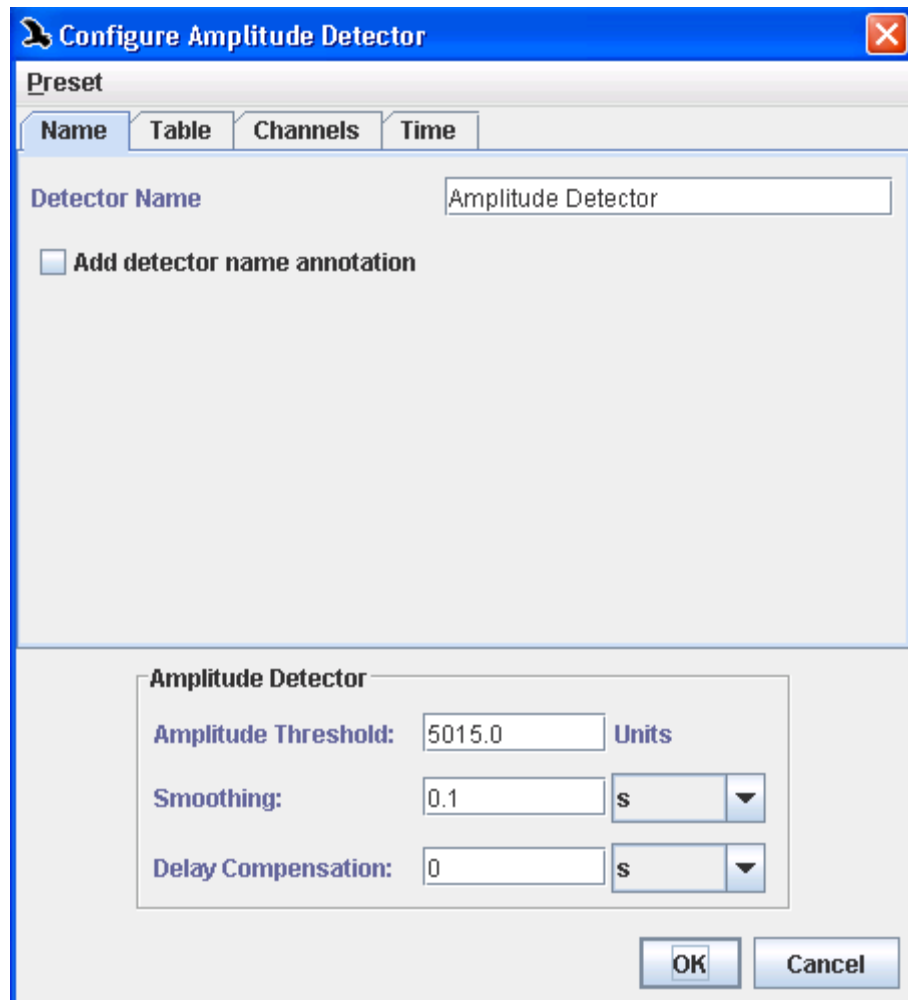


Figure 10.22. The dialog box to configure an amplitude detector.

To demonstrate the use of the amplitude detector, we will walk through an example. First, open the ChestnutSidedWarbler.aif file (included in the Raven Examples folder). Next, we will choose the amplitude threshold by moving the vertical position marker until it intersects a set of events typical of those which you would like to detect, but above the noise and clutter you would like to reject. Then, read off the amplitude threshold value from the y-axis on the left. In this example, the reading is 5.015 kUnits (or 5015 Units). See [Figure 10.23](#).




Figure 10.23. The ChestnutSidedWarbler.aif sound window with the vertical position marker (in the waveform view) positioned at 5015 Units. This reading will become the amplitude threshold for our detector.

Next, make a selection which encloses the most rapid rise or fall of the waveform that you would like the envelope detector to follow. Ideally, the rise should be from zero to the maximum value of the waveform. In this example, we've used the Delta Time measurement to compute the width of the selection for us. The time width, τ , is .09575 s which we will use as the value of the smoothing time constant. See [Figure 10.24](#).




Figure 10.24. After selecting the most rapid rise/fall of the waveform (selection 1, in red) we use the Delta Time measurement (highlighted in purple in the selection table) to obtain a value for the smoothing time constant.

After running the amplitude detector on ChestnutSidedWarbler.aif with the previously determined settings (amplitude threshold of 5015 Units, smoothing .09575) you should get something similar to Figure 10.25. The detections generally correspond to the regions we wish to detect, but are slightly delayed as illustrated in Figure 10.26. The delay compensation parameter minimizes this delay by shifting the detections earlier by a specified amount. To use this feature, make a selection spanning the time between when the waveform crosses the amplitude threshold and the start of a detection and use the delta time measurement as the delay compensation parameter.

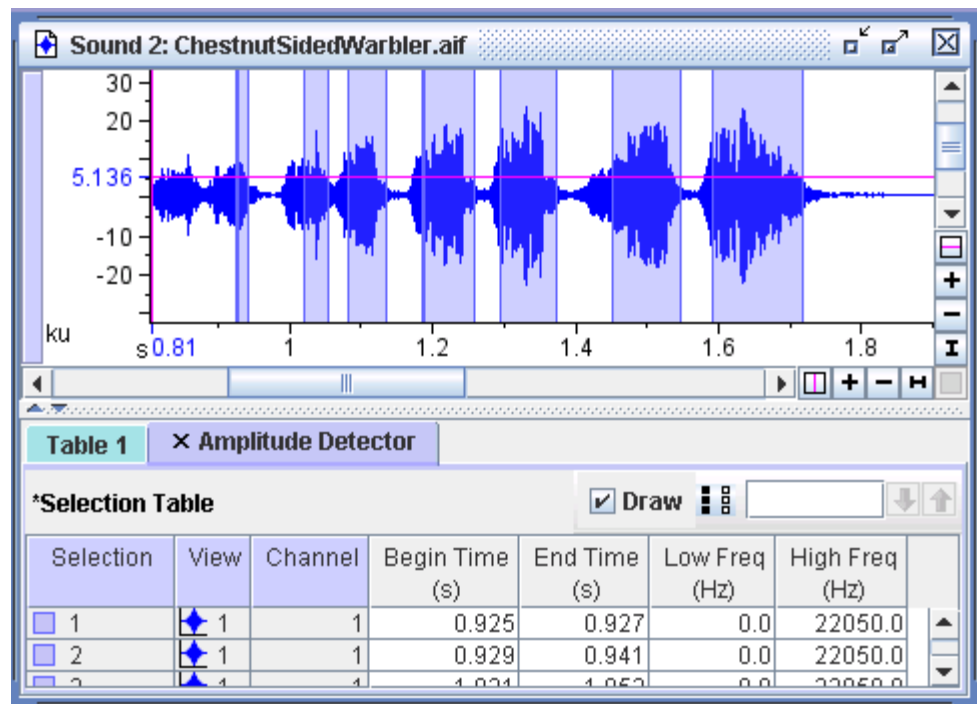


Figure 10.25. The selections made by the amplitude detector.

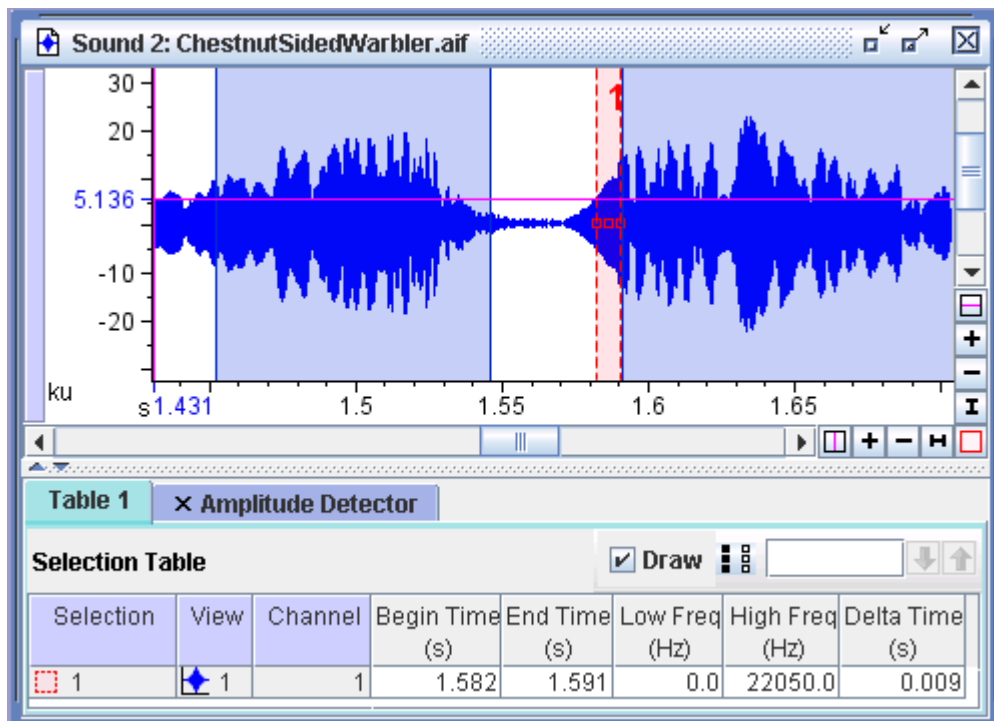


Figure 10.26. A magnified view of the selections made by the amplitude detector. Notice the delay between the time when the waveform crosses the threshold and the start of the detection.

Once the delay compensation parameter has been set to the appropriate value, the resulting selections should better align with the actual waveform, as shown in [Figure 10.27](#).




Figure 10.27. Selections made by the amplitude detector using delay compensation parameter of 0.009 seconds. The detections now correspond more accurately to the shape of the waveform.