PAMGUARD: An |Overview and Gentle Introduction

Preface:

PAMGUARD has excellent online, context-sensitive help which can be called upon from any part of the program. This is updated as the program develops and is the best source for up to date information. The intention of this document is not to replicate that information but rather to provide an overview of the program and a non technical introduction. Hopefully, the sort of document a user can print out and read browse through to become familiar with PAMGAURD, how it is structured and what it can do, and to help you plan how you can best use it.

Chapter

1

Chapter 1 Overview Structure and Philosophy

What PAMGUARD is designed to do, things you need to know about its fundamental structure and some of the Philosophy behind it.

PAMGUARD differs from many of the other excellent acoustic analysis programs available at the moment by being designed specifically to perform REAL TIME monitoring, detection and localisation of sounds from a moving platform at sea. Many other programs as better suited for careful offline analysis and measurement of sound. You might use these for example to select sections of acoustic files "back in the lab" and make careful measurements and descriptions of interesting sounds on them. If this is what you need to do then PAMGUARD is probably not the program of choice. The particular application that spurred PAMGUARD's development was the need to detect and localize marine mammals in near real time during passive acoustic mitigation exercises at sea. In these situations a single operator needs to quickly detect, classify and localize (DCL) sounds to allow decisions to be made about modifying activities that might pose a risk to sensitive animals (using an airgun in a seismic survey or triggering an explosion for example).

The requirements of this preeminent application has a lead to an emphasis on analysis and localisation of real time data, a capacity to configure and handle data from large and complicated hydrophone arrays (to allow localization of vocalizing animals), routines for dealing with high levels of background noise, a capacity to interface with instruments to receive navigational information and an ability to plot all of this, and acoustic detections, on a map. It is envisaged that usually PAMGUARD will be run at sea analyzing acoustic data provided by a sound card or some other analogue to digital interface. However, the input can also be a digital sound file (e.g. a .wav or .AIIF file) and the program may often be used in this way to review data offline, especially if the acoustic and ancillary navigational data were collected using PAMGUARD in the first place. In addition to a role in assisting PAM for mitigation the program has also ppoven particularly useful for surveing and population monitoring applications. Although PAMMGUARD was written with marine applications in mind it can of course be used to analyse sound recorded in air with microphones.

PAMGUARD has been designed to be highly flexible and almost infinitely configurable so that it can be customized for particular situations and applications.

Specific PAMGUARD configurations can be saved in PAMGUARD Setup Files (psf) and reloaded as required. .

If you start PAMGUARD with a blank configuration file you get a fairly minimal display screen (Figure 1-1). In this form PAMGUARD doesn't do anything very useful but all the functionality is here to build a functioning detection program customized to your particular requirements.

Building up a configuration from a blank psf is a good way to get to know the program. In practice though it makes sense for users to adapt pre-existing .psf files. A few examples ship with PAMGUARD, others can be found at the PAMGUARD website www.pamguard.org.

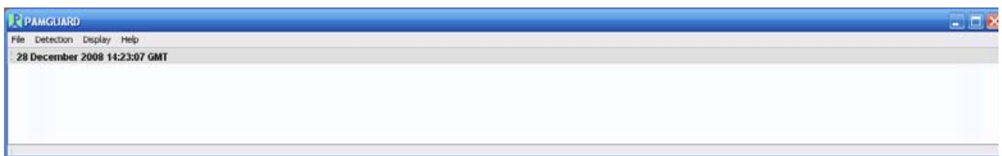

Comment:

Figure 1-1 The somewhat off-putting window which results when PAMGUARD is run with a blank setup (psf) file.

To get PAMGUARD to do something useful we need to add and appropriately configure Pamguard Modules, the key building blocks of the program.

If this process seems daunting to start with don't worry. With the appropriate configuration files, which can be obtained from many different sources, an operator can use PAMGUARD without ever needing to "get under the bonnet" and make adjustments. Once you get to know the program and become more familiar with it though its likely that you will want to make adjustments to get the most out of it.

The highly modular nature of PAMGUARD is a fundamental feature of its design. There are three advantages in this:

- it helps to provide flexibility for the user,
- it makes it easier for many different programmers to contribute useful functionality in the form of plug-in modules that can readily interface with other PAMGUARD units
- it allows sufficient use of computer resources.

Some of the existing modules of different categories are outlined in Table 1

Table 1.1 A Table of PAMGUARD Modules

Maps and Mapping	Sound Processing
NMEA Data Collection	Sound Acquisition
GPS and AIS Processing	Sound Output
Airgun Display	FFT (Spectrogram) Engine
Map	IIRF Filters
Utilities	Decimator
Simulator	Sound Recorder
Database	Signal Amplifier
Displays	Patch Panel
User Input	Spectrogram smoothing
Display Panel	Detectors
Location	Click Detector
3D locator	Whistle Detector
	Ishmael Detectors (Energy sum, spectrogram correlation, matched filtering, localiser)

Modules interact with the "outside world" or with other parts of PAMGUARD by for example, accepting data from a module or passing data on to another module where it may be further processed or displayed and modules can be "connected" in many different configurations.

A level of complication is an inevitable corollary of so much flexibility however, and this can soon seem overwhelming. Fortunately, PAMGUARD incorporates a tool called the **"Data Model"** which depicts the current configuration in an easily understandable diagrammatic form. The Datamodel view also allows you to investigate use of resources by each model, to adjust the configuration of modules and to add new ones. Athough modules can be initiated and setup using the PAMGUARD menu many find the Datamodel the most convenient place to do this.

To load Data Model choose

File > Show Data Model

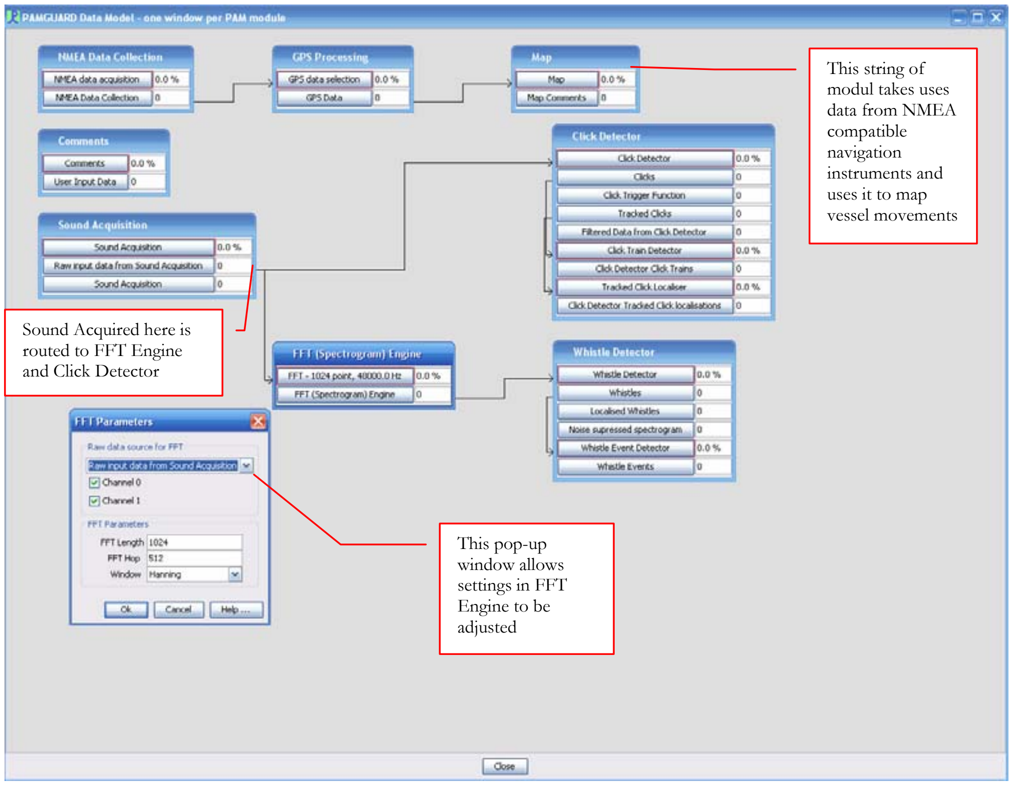

Figure 1-2 A typical Datamodel view for a configuration with NMEA input and an acoustic input through sound acquisition being lead to both a click detector and a whistle detector/

One "feature" of the Data Model view is that it isn't very good to arranging the modules neatly on the pane in an easily viewed configuration. Click and hold to drag them to better locations on the pane for a less cluttered depiction.

Its quite possible to use PAMGUARD with appropriate set up files and never have to adjust them. However, a more expert user can capitalize on its scope for flexibility to tune it for each particular application.

Chapter 2 Windows, Tool Bars and Navigation

An introduction to the Graphical User Interface (Menus windows, panes, user displays, toolbars and Navigation

PAMUARD has an easy to use and fairly generic "windows type" graphical user interface.

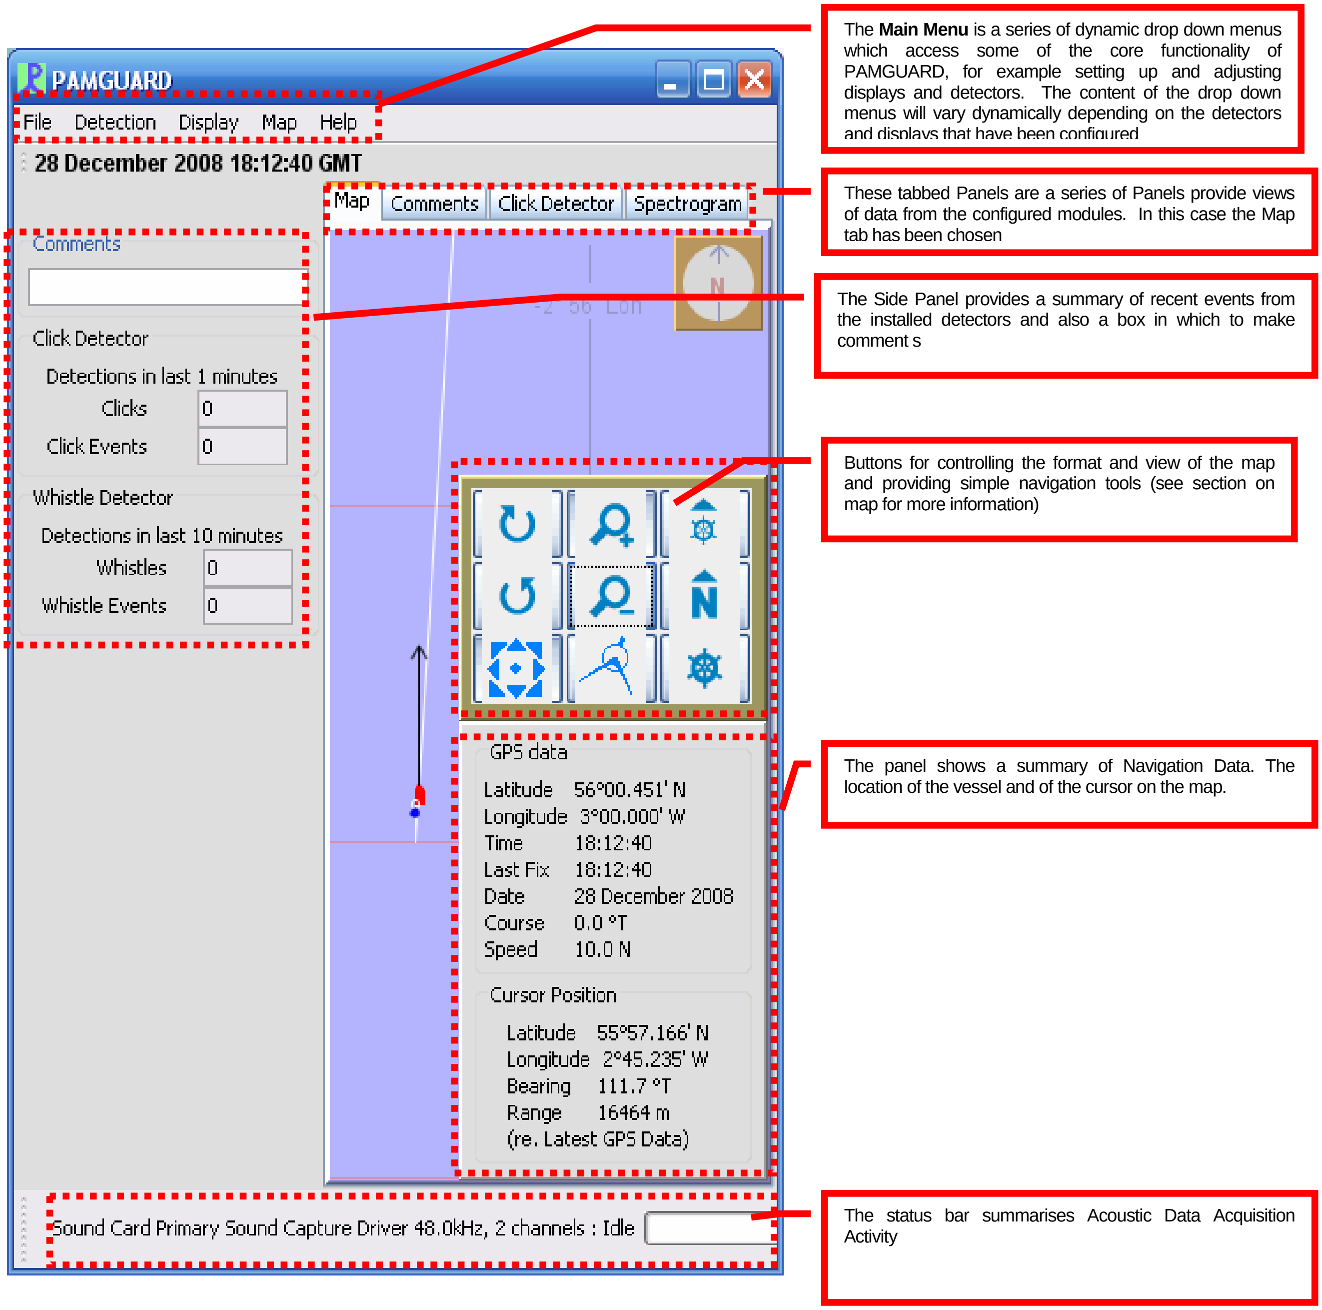

Figure 2-1 Basic PAMGUARD Display and the some of the available frames and controls

Figure 2-1 shows some of the features usually present on a PAMGUARD display. The Main Menu is a series of buttons at the top left of the window which provide access to most of the functionality within PAMGUARD. Most features can be controlled using these, though there are usually also alternative routes for accessing these options too. The Side Panel on the left of the window provides a text box for making quick comments and presents information summarising the recent history of detections from the available detectors. A status bar at the bottom of the window summarises the status of the acoustic acquisition devices that are part of the current PAMGUARD configuration. A series of tabbed panels to the right provide the main display. Each of these tabbed displays is customised to be appropriate to the appropriate detector. There will usually be a single tabbed display for each detector or display module installed as well as for utilities such as a Comments page.

In this case the Map tabbed display is uppermost. In addition to showing vessel tracks and outputs from detectors the Map display includes a Palette of tools for adjusting the format of the map and simple navigation tools such as a distance measuring tool. A panel showing information on the vessel's location (usually from a GPS navigator) and the current cursor position on the map in world coordinates is also provided.

Main Menu Buttons

The main menu is s row of drop down menu buttons which give access to most of t he functionality of PAMUARD

The upper panel in the File Menu Figure 2-2 deals with saving and loading configuration files which determine which modules are loaded and how they interact with each other. Choosing Add Modules opens sub menus which show the available modules arranged in functional groups. In this case the Sound Processing modules are shown expanded. (Modules can also be added from the Data Model View.)

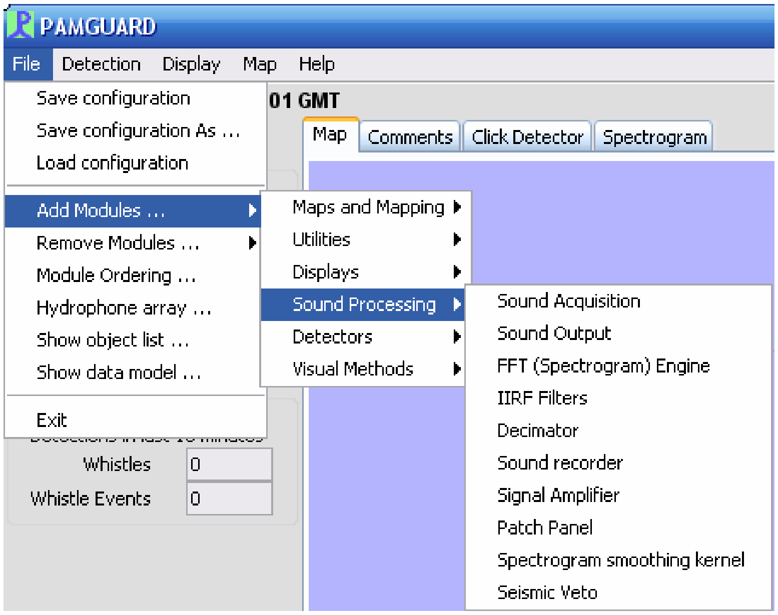

Figure 2-2 The File Menu. Sound processing modules that are avaialbe to be added are shown expanded.

Other buttons allows the user to Remove Modules view and change their ordering, Setup a hydrophone array (see Chapter3) `Show Object List` or `Show Data Model` view (see Chapter 1).

Items under the second Main Menu button `Detection` relate to running detectors and other modules. The first two choices, `Start` and `Stop` are crucially important they turn the active parts of the program on and off. Starting the program initiates acoustic data collection, processing and detection. Stopping the program may be necessary to add new modules or to change some of the setup parameters for existing modules.

`Multithreading` allows the program to take advantage of multi-core processor chips and should normally be activated.

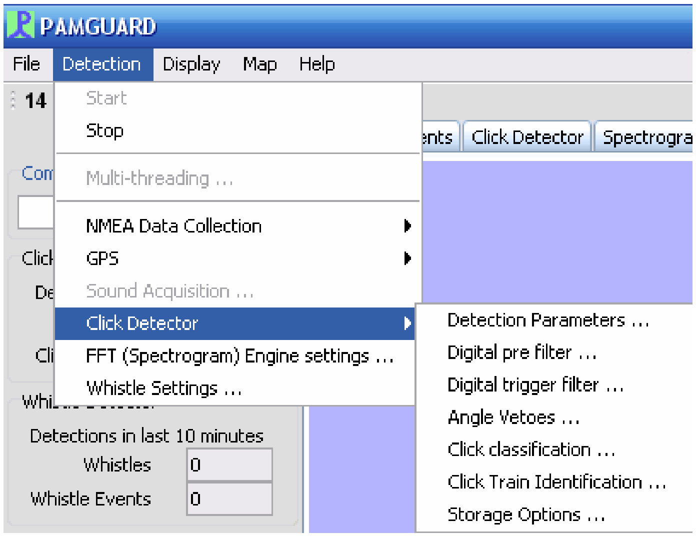

Figure 2-3 The detection menu. In this case the Click Detector has been chosen and the groups of parameters available for adjustment are shown.

The range of menu items that will be shown in the bottom frame for the Detection Menu is dynamic and will depend on the modules that are active in the current instance of PAMGUARD. Choosing any of them gives access to appropriate setup options for that particular module. In Figure 2-2 the click detector module has been chosen. Clicking any of the options shown will bring up an appropriate dialogue allowing those feature and options in the Click Detector module to be adjusted.

The `Display Menu` provides options that control what should be displayed and how Figure 2-4. `Colour Scheme` allows the choice between two alternative scheme that have a general effect on the PAMGUARD display: a "colourful" daytime display or a night time display designed to conserve night vision. Choosing `Symbols` bring up a list of symbols for objects that are currently available to plot on the map and a dialogue that allows them to be customisedFigure 2-5.

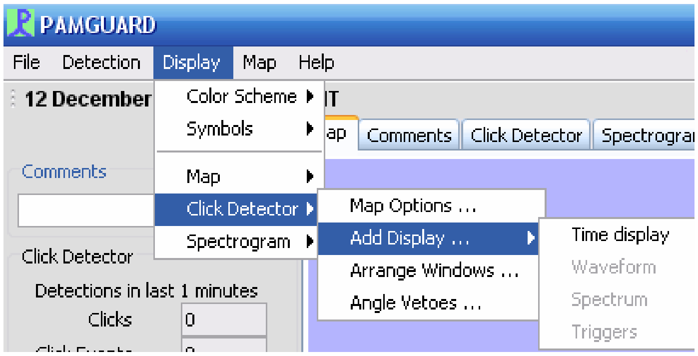

Figure 2-4 The Display Menu. Display options available for the Click Detector are shown expanded.

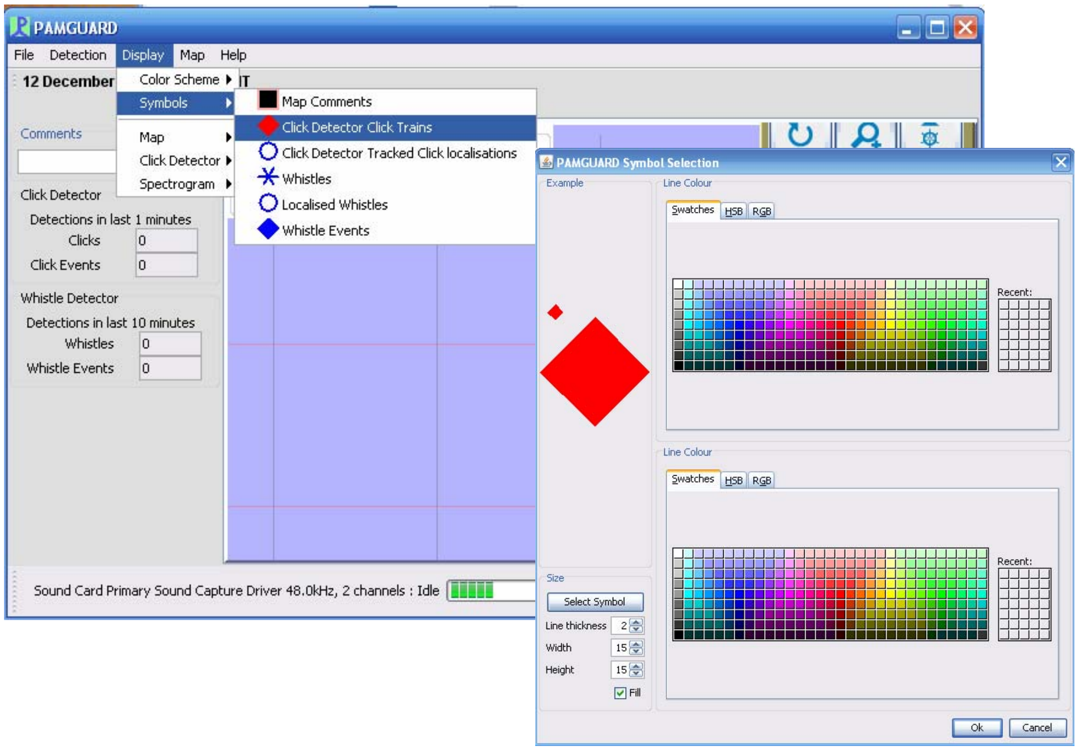

Figure 2-5 The Symbols menu option allows map symbols to be chosen and customized

The lower panel of the Display Menu provide options of altering the way installed modules are displayed. In this particular PAMGUARD installation `Map`, `Click Detector` and `Spectrogram` modules are available. Figure 2-4 shows the `Click Detector` option chosen, which in turns provides options to display detections on the map, initiate a new detector display, arrange the different windows generated by the click detector etc.

Figure 2-6 Map menu options

The Map button in the Main Menu (Figure 2-6) is actually a dynamic button. It is shown in this case because the Map is the chosen Tab dispaly, and is active. This provides easy access to options that could also be accessed through the Map options under the Display menu. If another tab, such as Click Detector, was the chosen one then Buttons providing access to Click Detector functions would appear in this location, e.g. Figure 2-7

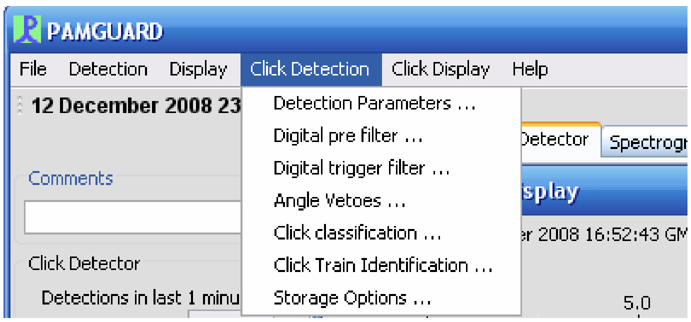

Figure 2-7 Click Detector buttons appear on the Main Menu when a click detector has been incorporated in PAMGUARD and the click Detector Tab is chosen.

The final button in the Main Menu is always the Help Button. This provides information on the current version of PAMGUARD under About Pamguard. Access to the help file that ship with PAMGUARD if **Help** is chosen and a link the help pages on the PAMGUARD website. The help files that are installed on your computer with PAMGUARD provide detailed help on all the PAMGUARD modules and include a search facility to make it easy to find help on the appropriate components.

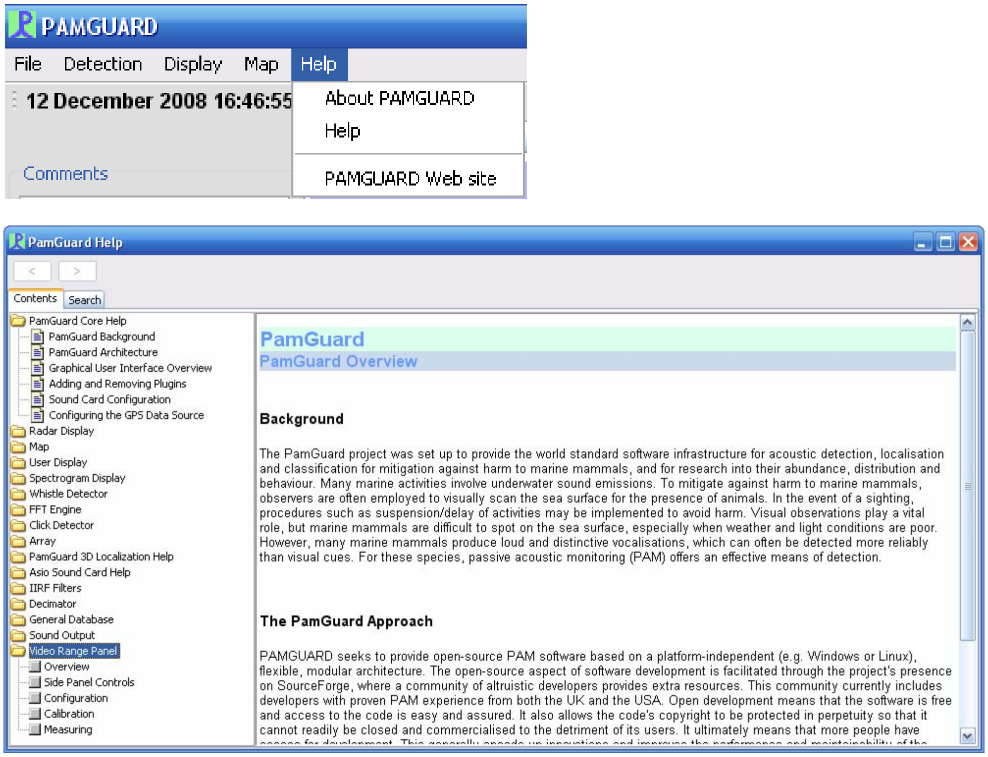

Figure 2-8 Help is always available from the Main Menu Help button.

The other generic route for accessing functionality of different modules and displays is to right click on them. For example, right clicking the map display bring s up a range of option for controlling what should be plotted and how Figure 2-10

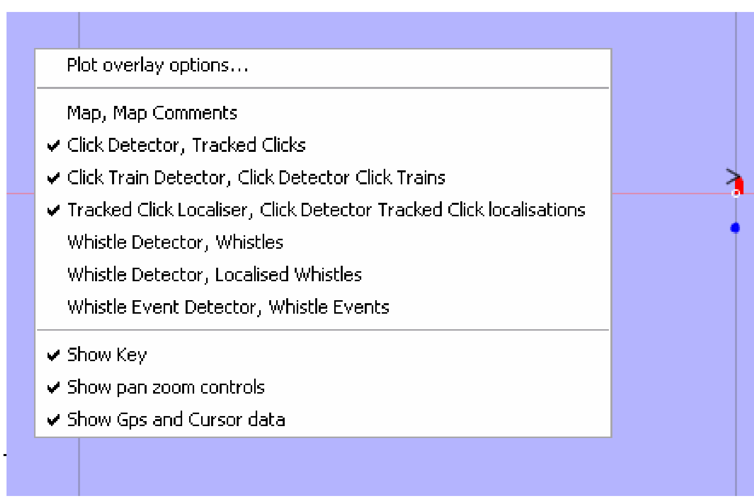

Figure 2-9 Right clicking on the map display allows a range of options controlling what is displayed on maps and how.

Similarly, right clicking in the Click Detector window brings up options giving access to its setup and providing control over its displays .

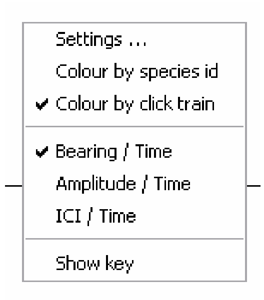

Figure 2-10 Options for setup and display of Click Detector that are offered after right clicking on a Click Detector window

Chapter 3 Inputting Sound and Manipulating Acoustic Data

This chapter introduced the way that sound is input into PAMGUARD, how acoustic data "flows " through the program, some of the ways that acoustic data channels can be manipulated and options for recording and output.

Fundamentally, PAMGUARD is an acoustic analysis program

Acoustic data can be made available to PAMGUARD either through some sort of analogue to digital converter or from pre-digitized sound files.

Three types of analogue to digital converters are currently supported

- WAV Sound Cards
- ASIO Sound Cards
- National Instruments DAQ cards

Audio files can be in .wav, .aiif or .au formats

Once an audio data stream is in PAMGUARD it can be manipulated, modified or reconfigured in many different ways, routed to various detectors or displays and (if an ASIO card is being used) can be routed to an output module so that it an be converted back to an analogue signal and heard.

Specifying a Hydrophone Array

Lets start by considering a situation in which real time data are being captured through an analogue to digital converter. Logically we should start with the hydrophone elements and the array in which they are configured. PAMGUARD has a sophisticated and flexible module which allows a three dimensional array with up to 32 elements to be specified. This is an inbuilt module and it has its own button under the File Menu.

Choose

>File > Hydrophone Array

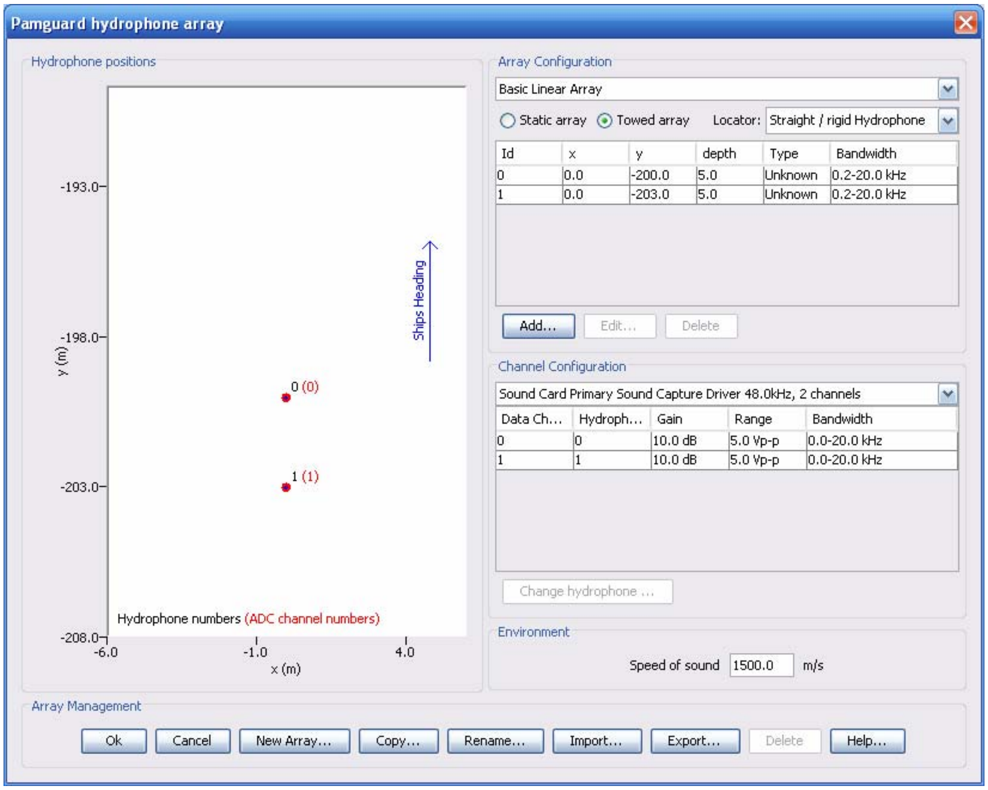

Figure 3-1 Array configuration dialogue

By default PAMGUARD shows a simple linear array with hydrophones 200 and 203 m behind the reference point Figure 3-1.

Note that the reference point for the array is normally the location of the GPS antenna and its location on the vessel is defined in the `GPS Options` panel in the `GPS Processing` module. The Y axis is defined by the vessel's heading, hence distances to hydrophones behind the vessel are negative.

More complicated arrays can be built-up by adding new hydrophones

Press the `Add..` button.

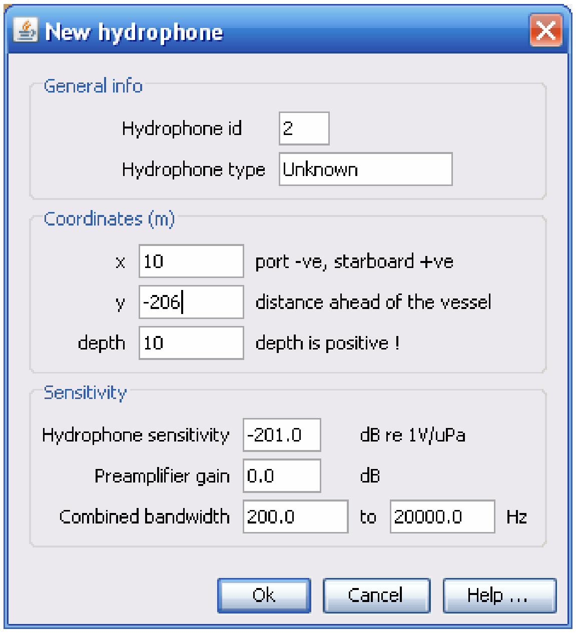

Figure 3-2 New Hydrophone Dialogue

Which causes the `New Hydrophone` dialogue to pop up Figure 3-2.

The location of the hydrophone can be specified, in the coordinates panel. (Note that the y dimension is relative to the ships heading so that distances astern, where most arrays will be, are negative. Depths on the other hand are positive.)

The hydrophone elements' sensitivity and the gain of any associated preamplifier can also be input here as well as their combined bandwidth.

As many hydrophones as necessary can be added to the array or their details can be edited. Their locations are shown in plan view in the hydrophone positions pane with the hydrophone number and ADC channel number (in brackets) displayed next to each one.

Typically the array will be mobile and towed behind a vessel. The `Locator` selection to the right above the hydrophone list allows you to choose between two models of how a towed hydrophone array follows the towing vessel's path. A "Straight/rigid hydrophone" is assumed to extend in a straight line behind the vessel at 180 degrees to the ship's heading. Thus, it will "swing around" as the heading changes. The "Threading Hydrophone" option specifies that the hydrophone follows the ship's path exactly as it is towed, which is generally a more realistic assumption.

Once the hydrophone array has been specified we need to tell PAMGUARD how it is "connected" to the inputs of the available analogue to digital devices. As you might expect, PAMGUARD allows a great deal of flexibility in how this is configured. More than one digitizing device may be present and any hydrophone can connect to any number of the available ADC channel inputs, (which at this stage of course exist as physical connections).

Sound Acquisition

PAMGUARD interfaces to the available acoustic acquisition devices through one or more Sound Acquisition modules and each different device needs a separate Sound Acquisition module..

Sound Acquisition modules can be created by choosing

Files> Add Module> Sound Processing> Sound Acquisition

Or by right clicking on the background in the Data Model view

Once a module has been created, right click on the Sound Acquisition module and choose Setup.

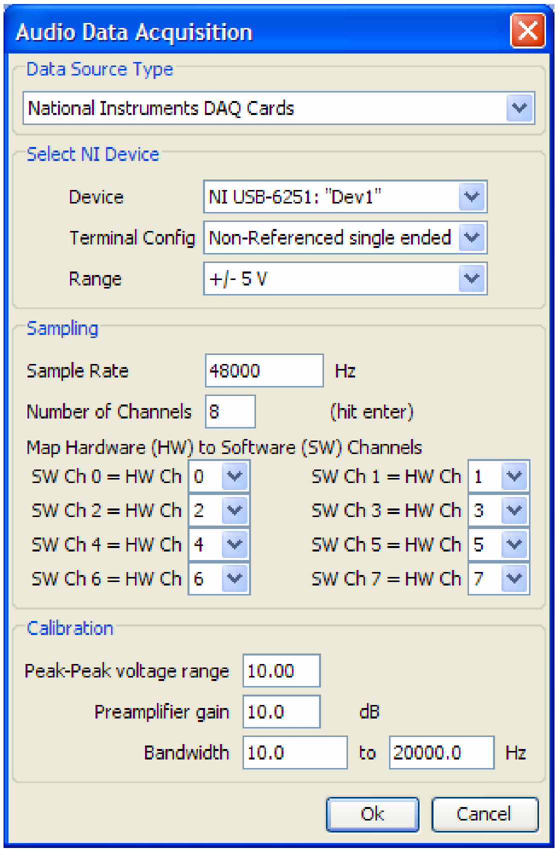

Figure 3-3 Sound Acquisition panel showing a setup for an 8

The first stage in setting up a Sound Acquisition Module is to choose the Data Source Type. The subsequent configurations options will depend on the characteristics of Acquisition Device chosen.

Sound cards are stereo devices and one will normally choose the standard two channel configuration when using one of them.

With a device which can support more than one channel, such as an ASIO card or a National Instruments DAQ, enter the number of channels in the appropriate box (**Number of Channels**) and press return. The form will be redrawn showing drop down selection boxes for each channel. Use these to choose the hardware which will map to each software channel.

The **Hardware Channels** are the sockets (or other connections) providing input into the Acquisition Device, **Software Channels** are the data channels through which acoustic data will be available to other modules within PAMGUARD.

By default, PAMGUARD reuses the same channel numbers (e.g. 0,1,2,3) for each device. Multiple channels from multiple devices can exist in a PAMGUARD configuration but channels originating from different devices can not be "brought together", for example to be recorded as channels in a single data stream or within a detection and localisation module.

Once within PAMGUARD acoustic data in Software Channels can be manipulated in several ways before being recorded, output or processed by a display or detector module..

Signal Amplifier

A Signal Amplifier Module allows a gain (positive or negative) to be applied to each channel independently, and also, if required, for the phase of one or more channels to be inverted. (For example, you might want to switch the phase of a channel if you find that the polarity of the hydrophone to which it is attached is different from that of the others in an array).

In Figure 3-4 Settings windows for Signal Amplifier each evenly numbered channel is given an extra 3dB of gain and its phase is inverted.

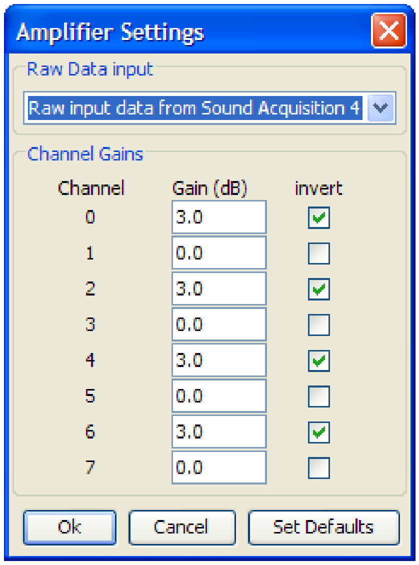

Figure 3-4 Settings windows for Signal Amplifier Module

Patch Panel

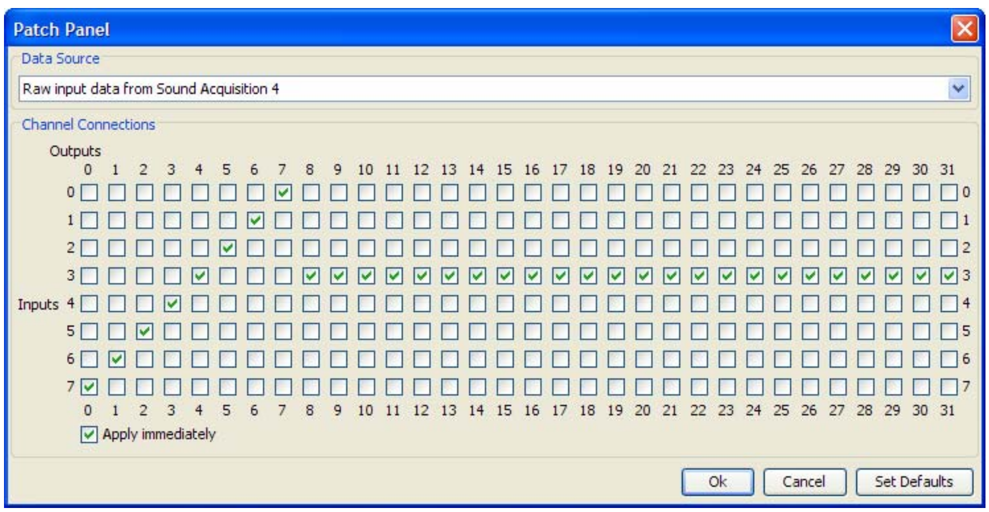

Figure 3-5 Patch Panel Module

Another very useful inline module is a `Patch Panel` . This allows the allocation of channels to be changed and/or to be shared between several channels. Figure 3-5 shows a configuration of a patch panel modules with 8 input channels whose numbering is reversed. The remaining 24 channels are then assigned data from input channel #3 (unlikely to be a very useful configuration in practice but perhaps serves to illustrate a point).

Filter Module

Another module that can be applied "inline" to modify a data stream is a IIFR Filter module. (IIFR stands for Infinite Impulse Filter Response) This allows a range of different filters to be applied to the data stream. Low cut, high cut or band pass filters can be specified using one of two widely utilized filter algorithms, the Butterworth or Chebyshev The response curve of attenuation against frequency for the specified filter (also known as a bode plot) is plotted along with a pole zero diagram. The higher the filter order chosen, the steeper the roll of, but also the greater the amount of processing required. Usually 2nd or 4th order filters are appropriate for general use in PAMGUARD.

Decimator

If the sample rate of an acoustic DataStream needs to be reduced this can be done by passing it through a Decimator module. One simply enters the desired (reduced) sampling rate that is required after the data has passed through the decimator (Figure 3-6a). The other setting that is essential to attend to when using a decimator is a filter Figure 3-6b. When sampling an analogue signal using a data acquisition device it is essential to apply a filter of around half the sampling rate (the Nyquist Frequency) to the analogue signal to avoid aliasing . In the same way it is important to apply a similar low pass digital filter to the digitial data stream when re-sampling at a lower rate. The Filter dialogue allows you to specify an appropriate filter. A low pass 4th or 6th order Chebyshev filer with a cut off frequency of 0.8 times the Nyquist frequency (0.4 times the sampling rate), is recommended in the PAMGUARD help.

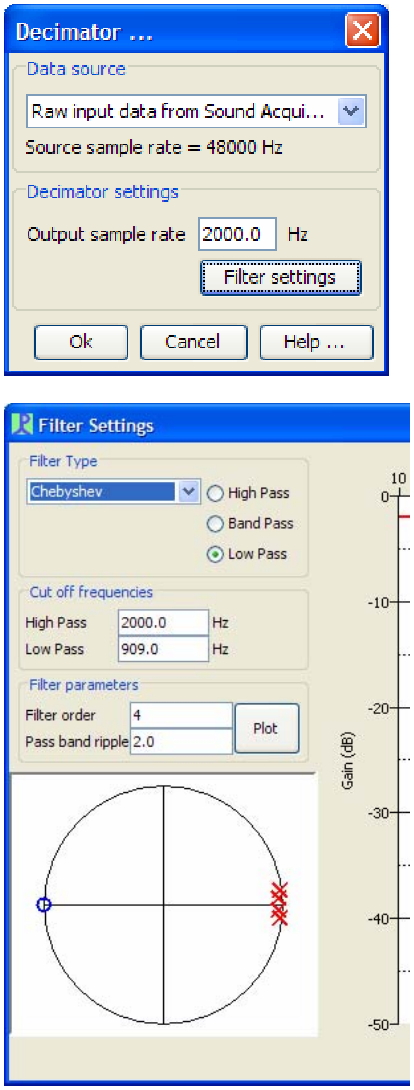

Figure 3-6 Setup window for a Decimator Module and the filter settings necessary to avoid aliasing

On final point to make about acoustic data channels in PAMGUARD. PAMGUARD modules often calculate arrival bearings for sounds based on time of arrival of signals between a pair of hydrophones (which are usually close to each other in space). When a sound acquisition device provides more than two acoustic data channels then any pairs of channels that should be considered to be grouped in this way, known as `User Defined Groups`, can be set up in the `Detections Parameters` panel for the detector, which will be found under `Settings`.

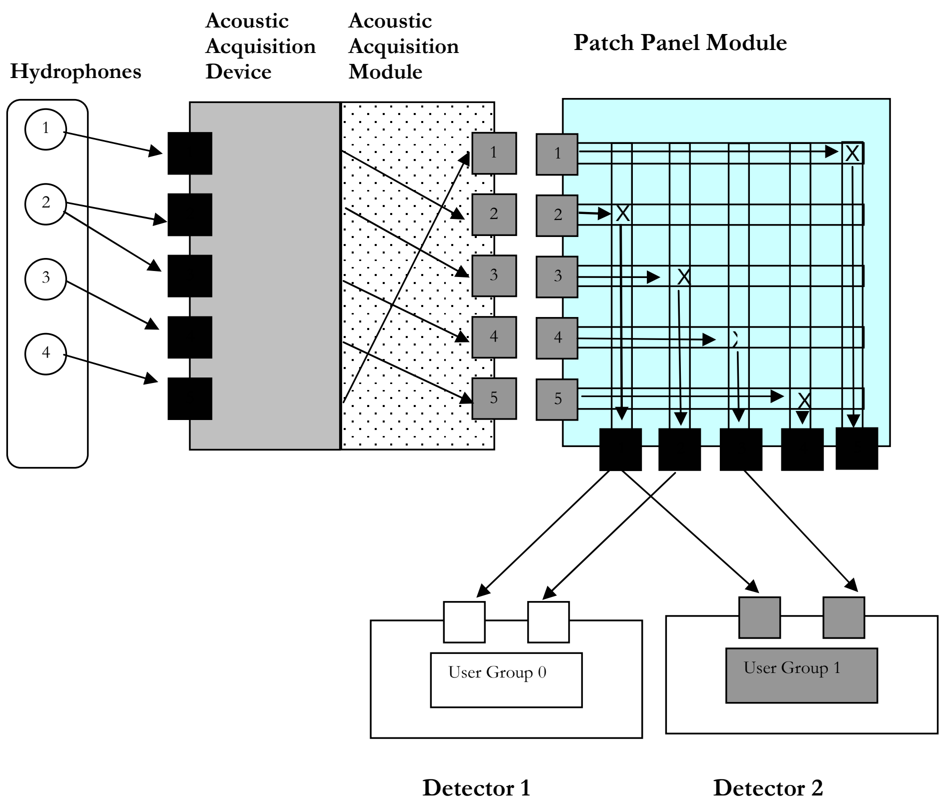

Error! Reference source not found. shows a setup in which4 hydrophones are attached to 5 inputs on a sound acquisition device with hydrophone 2 providing input to both hardware channels 2 and 3. The Sound Acquisition module rearranges the mapping of hardware to software channels while a subsequent Patch Panel module changes this back so that the numbering of the software and hardware channels are once again the same. Finally, User group 1 of Detector 1 combines

The ultimate "destination" of an acoustic data stream may be a display or a detector and these are discussed in later chapters. However, two other destination modules for an acoustic data stream in PAMGUARD are a `Recorder` module and a `Sound Output` module.

Recorder Module

The recorder module allows digital audio data to be stored as files on the computer's hard drive. As you might expect, there are a number of options and settings available making this a powerful and versatile utility.

During the initial setup of a recorder the source of acoustic data must be specified. This might be the raw output from a sound acquisition device, or data that has been modified by some later module, after it has been decimated or filtered for example. The recorder can be set to start recording automatically every time the PAMGUARD program is started.

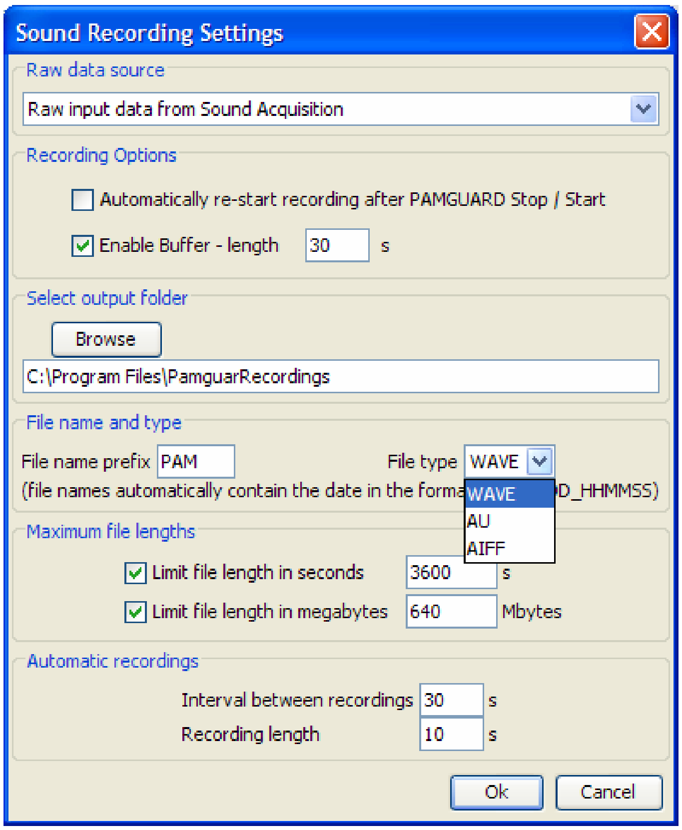

Figure 3-7 Setup panel for the Sound Recorder

The Buffer is a very useful feature. When this is activated a rolling buffer, with a length specified by the user, is maintained and any recording then starts with the data that are in the buffer at the time it is initiated. Thus, if a 30second buffer is specified then the 30seconds of sound that was received before the record button is pressed are saved, as well as the audio record from then on. This can be very useful, for example in a situation where you wish to capture a rare event., you can wait until you detects or hear it before starting a recording and still capture it fully.

The next option on the setup panel is the location in which you want the sound files to be saved. You can also specify a descriptive 3 letter prefix for file names (the rest of the file name will be made up of the time and the date). The format in which acoustic data should be written can also be chosen (wav, au or AIFF formats are available).

The next two options relate to the length of each file and the size of the directories in which these files should be ordered. For many applications it is sensible to keep files to a manageable length and it can also be convenient to arrange files into directories that have a size just a little less than the capacity of your likely storage medium, e.g. either a CD or a DVD.

The final two options on this panel allow the recorder to be setup to automatically collect sample recordings of a certain length at a specified time interval. (In this case, the time between recordings is the time from the start of one recording to the start of the next.

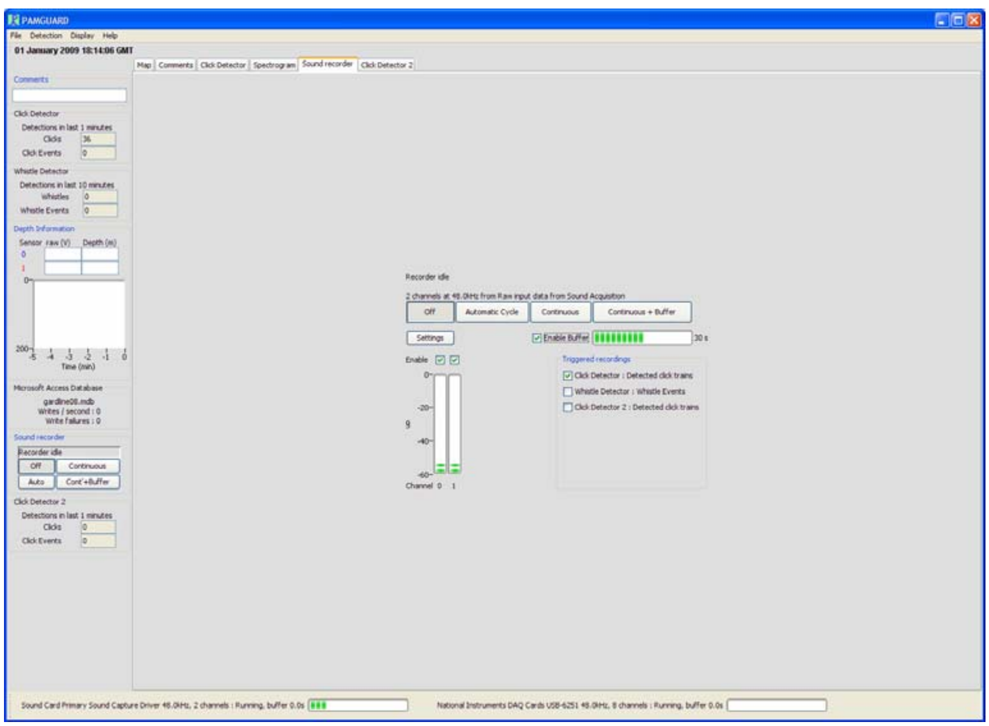

Figure 3-8 A recorder Panel a 30 second buffer has been specified and this about half full at the time that the screen shot was captured.

Once a recorder has been set up it appears as a selectable tab on the main screen Figure 3-8 shows a typical recorder display.

On the recorder display, the first rows of buttons allow the recorder to be turned Off, or to be started in one of three ways.

- Automatic Cycle To start making sample recording on the time schedule set up by the user
- Continuous To record continuously, until turned off by the user,
- Continuous and Buffer To start recording continuously including the sound already stored in the buffer – provided of course a buffer has been specified during setup and the buffer has been enabled (by ticking the box below).

Each acoustic file made by the Recorder is given a unique name consisting of the three letter prefix specified during setup and the date and start time. If the buffer is enabled then the start time is adjusted to reflect the time minus the length of the buffer. Thus, the real time that any sound was received can be calculated as the start time in the file name plus the length of time into the file that the acoustic event occurs.

For convenience, a button is provided here to allow the settings dialogue to be accessed directly.

At the bottom of the panel to the left is a bar level meter for each channel. To the right are tick boxes to allow recordings to be triggered by any existing detectors. This can be very useful if for example one needs to capture samples of the acoustic record that triggered certain detectors to be able to verify their performance.

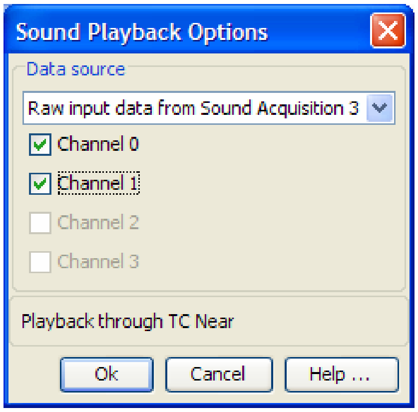

Figure 3-9 Setup options for a Sound Playback Module. The Module is configured to play channels 0 and 1 through the output of a TC Near ASIO sound card.

Sound Playback Module

Another "destination" for an acoustic datastream might be a sound output module – a digital to analogue conversion stage to make it audible as an analogue signal once more. Figure 3-9 shows the setup panel for a `Sound Playback` module. If the original digital acoustic data source is a sound file then the playback module will play back through the default sound output device for the computer (this can be reset using the Windows Control Panel.) If an ASIO sound card is being used as the sound Acquisition Device then the output can be any of the output channels of that device. The data source for playback can be any stage in the acoustic data stream. For example, this could be directly from a sound acquisition device or it could be after a filtering, amplification or decimation stage. PAMGUARD can be configured with Seismic Veto Modules which effectively gate out very loud signals (these will be explained later). If one of these has been configured appropriately then its effect can also apply to the output from a Sound Playback module making acoustic monitoring in environments subject to very intense intermittent sounds both easier and safer for the operator.

Chapter 4 Visual Display of Acoustic Data: Waveforms and Spectrograms

In biological acoustics important information s is often contained in the way that the frequency content of the sound varies with time. Any waveform, however complicated, including acoustic waveforms or sound, can be described as a combination of a series of values of the amplitude and phase of different discrete frequencies (In biological acoustic applications phase is usually ignored, and we will do so here.) Our own auditory systems conduct a, largely mechanical, frequency analysis of received sounds and we are thus very used to perceiving the acoustic world as one in which different emphasized frequencies and the way that these vary with time, has significance. The distribution of energy between different frequency ranges within a waveform can be calculated using a process known as Fourier Analysis which essentially performs a transform between the time and frequency domain. On a computer this is usually achieved using a computationally efficient mathematical formula known as the Fast Fourier Transform (FFT). The FFT algorithm can be applied to a sample of N regularly spaced amplitude values (for example a sequence of amplitude values describing the waveform of a digitized sound) to obtain a description of the distribution of frequencies within that sample or !time slice" (Note: N must be an exact power of 2 e.g., 4, 8,….512, 1024 etc. In PAMGUARD the number of samples in this sample is termed the "FFT size".). The output of an FFT transform of such a sample could be represented as a plot of level against frequency, often called a **Spectrum plot** (e.g. the bottom panel in Figure 4-1)

Biological signals, such as the vocalizations of marine mammals, are typically highly dynamic; their frequency structure. Levels at different frequencies vary with time and the manner in which this temporal variation occurs is important and significant and something that is important to be able to display and measure.. The usual way to represent this is to take a series of relatively short time slices or FFT samples,, calculate the distribution of frequencies in each and stack these time slices next to each other in series, to form a grid. Usually frequency values are distributed vertically on the y axis and time scrolls horizontally on the x axis., The amplitude value in each frequency, time grid cell is typically coded using either a grey or a colour scale. (For example higher values could have a darker shade of grey or a "hotter" colour.) Displays of this type are called spectrograms and are widely used by biologists.

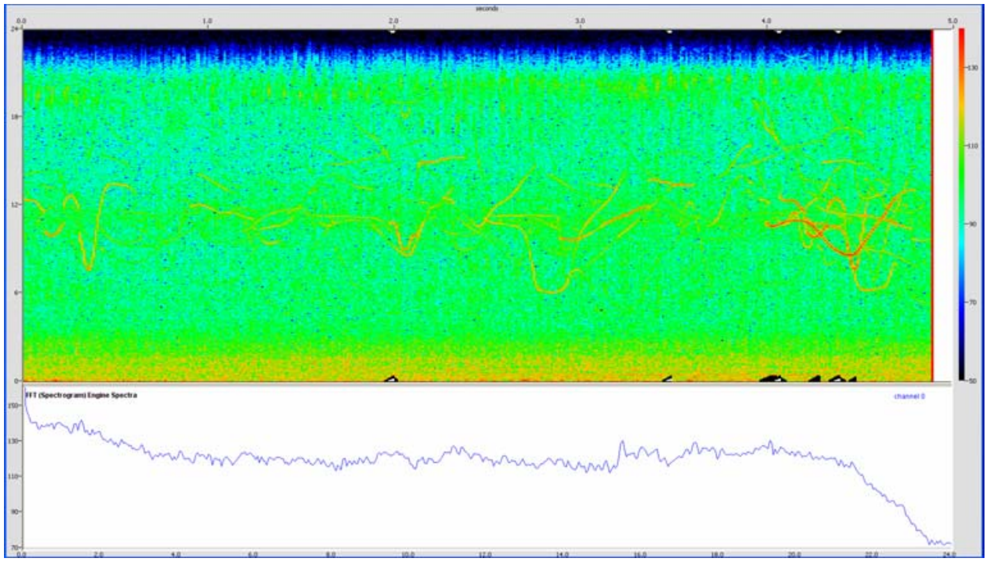

Figure 4-1 Specotorgorgram of dolphin whistle (top pane) and a spectrum plot of a single one of the time slices that make it up in the bottom pane.

Here we don't go into the details of how FFTs work but will accept the process as a "black box". One characteristic is that the number of frequency bins is determined by the length of the time slice analyzed by the FFT (the FFT size). In fact for any sample of N values with a time interval of T between each one the FFT will provide 1+N/2 frequency bins distributed over the frequency spectrum extending up to half of the sampling rate of the waveform samples

(1 / (T/2).

Thus, to obtain a better frequency resolution one must analyse a larger (and longer) sample of sound (use a greater FFT size). However, its clear that the temporal resolution of the FFT can never be better than the FFT length so there is an unavoidable trade off, as these time slices are made longer (to improve frequency resolution) the temporal resolution within the time frequency grid decreases. This trade off between frequency and temporal resolution, sometimes referred to as an example of the uncertainty principle. One consequence of this is that FFT size is an important consideration when specifying how the FFT calculation to be performed and spectrograms of the same sound can look very different depending on how the FFT calculation was specified.

PAMGUARD employs two "tricks" help to provide spectrograms with better resolution and a less "blocky" appearance. The first is to use FFT sequences that overlap with each other. This increases the processing required but can provide improved time resolutions. In PAMGUARD the amount of overlap is specified by the "hop length", the number of samples between the start of each FFT sequence. If hop length was the same as FFT length then FFT samples would be continuous with no overlap. A typical hop length is half the FFT length, doubling the number of FFTs that must be calculated and providing a 50% overlap.

The second procedure that can improve spectrogram appearance is called "Spectrorgram Kernel Smoothing" this is a separate module that sits between the FFT engine and a Spectrogram display

Choose

Files> Add Module> Sound Processing> Spectrogram smoothing kernel

Then choose Spectrogram smoothing kernel as input in the settings for the spectrogram display.

The Effects of some of these are shown factors are illustrated in the following figures which show a 1 second spectrogram of 1 second of dolphin whistles.

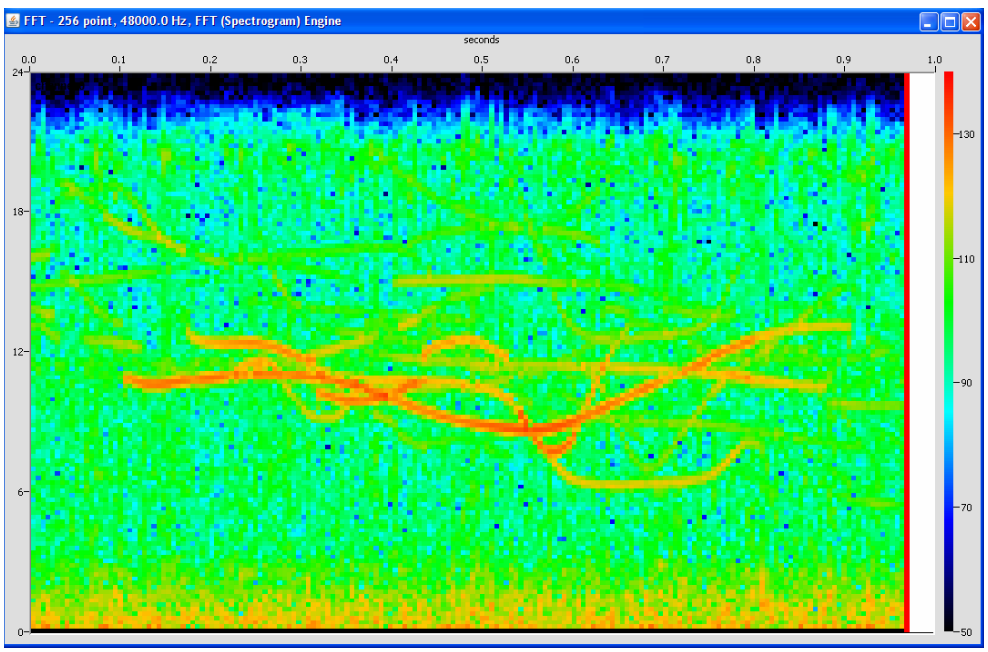

Figure 4-2 Sperctrogram Dolphin whistles Hamming Window 48000kHz sample rate. FFT size 512, hop length 512. Good temporal resolution but frquecny resolution is quite poor. (The whistle contour is quiete "thick")

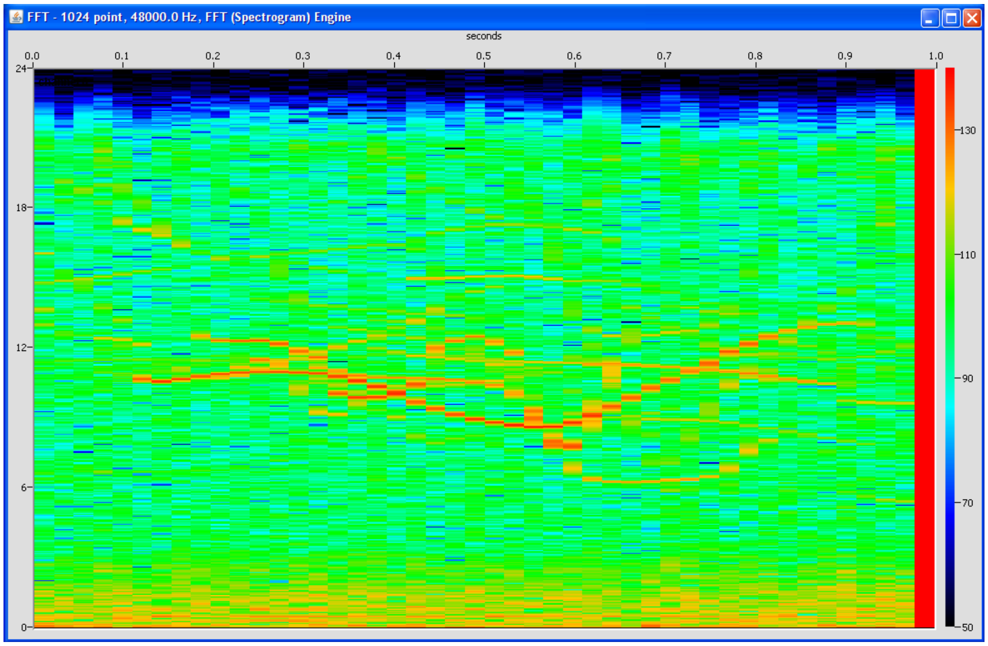

Figure 4-3 Spectrogram Dolphin whistles Hamming Window 48000kHz sample rate. FFT size 1024, hop length 1024. The freqeuncey resolution is has improved but temporal resolution has been sacrificed.

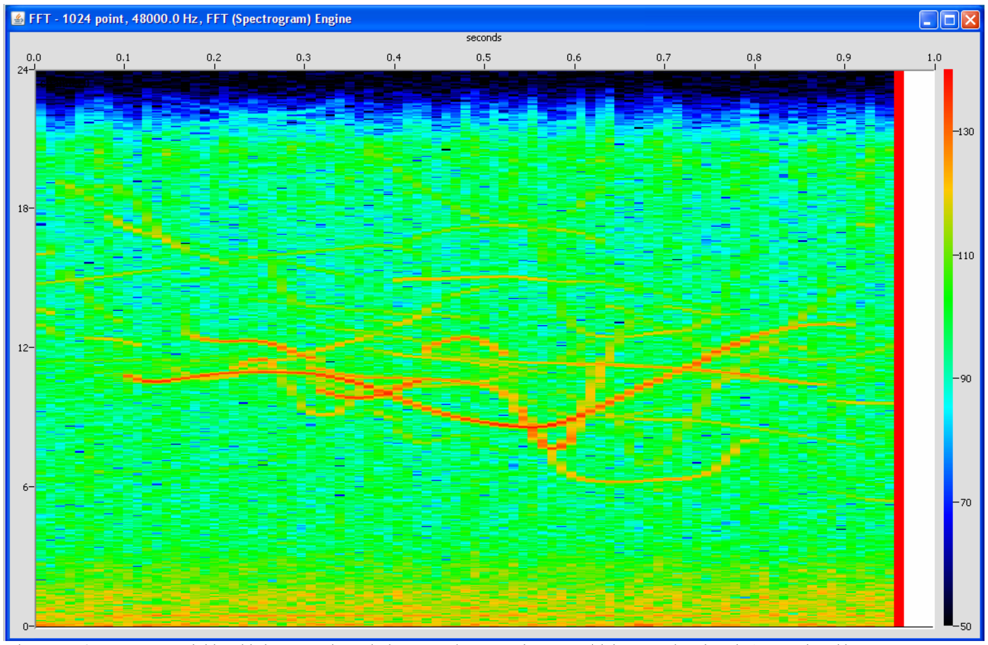

Figure 4-4 Spectrogram Dolphin whistles Hamming Window 48000kHz sample rate. FFT size 1024, hop length 512. Reduced hop length has resulted in improved temporal resolution.

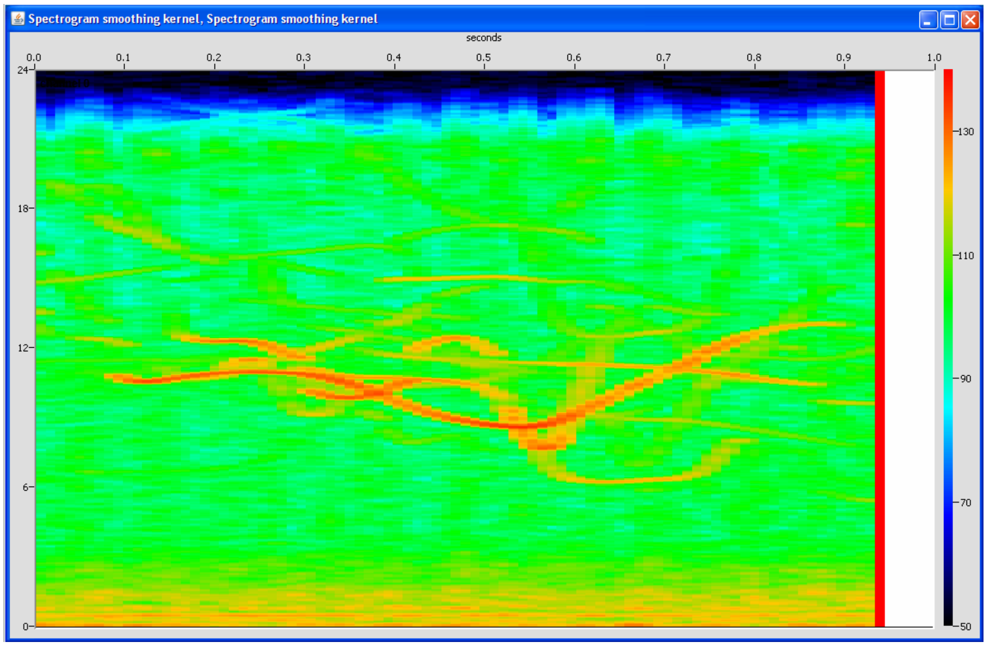

Figure 4-5 Spectrogram Dolphin whistles Hamming Window 48000kHz sample rate. FFT size 1024, hop length 512.. Application of kernel smoothing results in a smoother, less blocky i image

Adding a Spectrogram in PAMGUARD

A spectrogram display depicts a sequence of FFT calculations and PAMGUARD expects the Data Model to include an appropriate FFT engine before it can draw a spectrogram..

An FFT engine can be initiated and setup by selecting

Add Modules> Sound Processing> FFT (Spectrogram) Engine

From either the File menu or right-clicking in the Data Model view.

Once it has been initiated the FFT Engine it should be setup to take an appropriate source of data – usually raw data from a Sound Acquisition or some later processing module and appropriate values for FFT length, Hop Length and Window should be chosen.

Typically in acoustic analysis programs the FFT parameters are set from within the part of the program that draws the spectrogram. Having a separate FFT Engine in PAMGUARD may seem rather cumbersome. There is a very good reason for this however. Calculating FFTs is processor-intensive and many other modules may need to make use of the frequency domain data provide by the FFT engine (whistle detectors and the Ishmael detectors for example). Maintaining the FFT engine as a discrete module means that it can provide data to other modules in addition to one (or more) spectrograms, and this obviates the need for repeating the computer intensive task of calculating an FFT more often than is absolutely necessary.

Having initiated an FFT engine we are ready to add a spectrogram to display it.

As with any display for a module in PAMUARD we must first provide a location in which it can reside. This takes the form of a User Defined Display Panel (UDP). A UDP is a container for windows continuing other displays and once formed it exists as a panel with its own tab allowing it to be selected for display by the user.

To add a new UDP select

Add Modules > Displays > User Display Panel

From either the File menu or right clicking in the Data Model view.

A New Module Name dialogue will appear and it is useful to give it a sensible descriptive name.

The new (blank) UDP will appear on the main display with its own tab to allow selection.

To add a display to the UDP select the new UDP then choose

User Display .> New Spectrogram

(Note: UDPs are not shown in the Data Model view and spectrograms can not be added from there)

The Spectrogram Parameters Dialogue will appear which contains a number of tabbed forms covering different areas of the setup process.

The first thing to specify is the Source of Data, this may be any one of the available FFT engines or sources of modified FFT data (such as noise suppressed, noise gated or kernel smoothed).

Spectrogram displays will be drawn in a series of panels on the UDP. You can choose as many of these as required or desired.

Enter the number of panels required, hit enter and the form will redraw to provide a selection box to determine which data channel will map to which panel,

Choose the Scales Tab to bring up a setup form allowing adjustment of the scales applied to the frequency range displayed, the amplitude range (and whether it should be coded as a grey scale or as a colour scheme) and the time range for the display.

The frequency range specifies the upper and lower range for a display. You might want to set this to show just the frequency range in which a particular signal of interest was likely to occur. A user might typically have more than one spectrogram display active providing a good view over several different frequency ranges.

The amplitude range sets the upper and lower amplitudes over which the colour or shading scale will apply. Thus, if a grey scale is being used every cell with a level equal to or less than the minimum value will be white and every cell with a value equal to or greater than the maximum value will be shaded black.

There are two options for the time scale. Either to specify the number of pixels to allocate to each FFT sample. A value of one will provide the longest possible time display while depicting every slice on the available screen. On very high resolution screens a value greater than one might make for an easier to view image..

The alternative option is to specify the length of time that should be displayed in the window and allow PAMGUARD to calculate the appropriate scaling between FFT slices and pixels.

The third tab is titled Plug Ins. This allows a series of additional displays to be "plugged in" underneath the Spectrogram Display. The options available here will depend on the current PAMGUARD configuration but they typically include Raw input data from the Acquisition Device (depicted as a waveform), Raw input data after any installed sound processing modules (e.g. filter modle), output from click detectors, the spectrum plot from an FFT engine for the latest FFT slice. Plugin displays can be selected from any detector or data stream, they are not restricted to the data-stream that is contributing the FFT display in the Spectrogram.

As we have seen, a Spectrogram Display can have multiple panels but these will always come from the same datastream and be subject to the same scaling facotrs, so that are mainly useful for displaying specrograms of different data channels. Multiple Spectrogram Displays can be added to each UDP and each one of these can have the same or different data sources and have quite different settings applied to them. Thus they could be used to make spectrograms configured to display different frequency bands for sample.

Chapter 5 Acoustic Detectors

What this chapter is about.

Introduction

PAMGUARD makes extensive use of a range of acoustic detectors to detect and classify probable marine mammal vocalisations. Once made detections are usually shown in appropriate displays, plotted on maps and stored to databases. Although the operator can review these in typical field conditions signals may be arriving so frequently that a high level of reliable automated detection is essential for the operation of a real time system.

Marine mammal acoustic signals are both diverse and variable. The cover many octaves, from the low infrasonic moans of baleen whales (10s of Hz) to the high ultrasonic clicks of odontocetes, which may have the majority of their energy over 100kHz. Some signals such as small odontocete whistles are highly variable at source. In addition of course the environment affects acoustic signals as they propagates through it so that characteristics at a distant received may be quite different from those at source. It is impossible therefore to make effective detectors that are tightly matched to particular sterotyped signals and the strategy in PAMGUARD has been to make more generic detectors for common and characteristic types of marine mammal signals, such as whistles and clicks of odontocetes and the low frequency tonal calls of baleen whales.

An additional difficult feature of the environment in which PAMGUARD must operate is that it is likely to be quite noisy and

Any detection system can make two types of error.

- Type I errors are false positives when a signal is wrongly detected when it is not really present. For example a particular segment of noise may have triggered the detector and been erroneously classified as signal.

- Type II errors are false negatives, they occur when the signal really was present but was not detected.

A trade off between these two types of error is inevitable and depends on the signal to noise ratio.

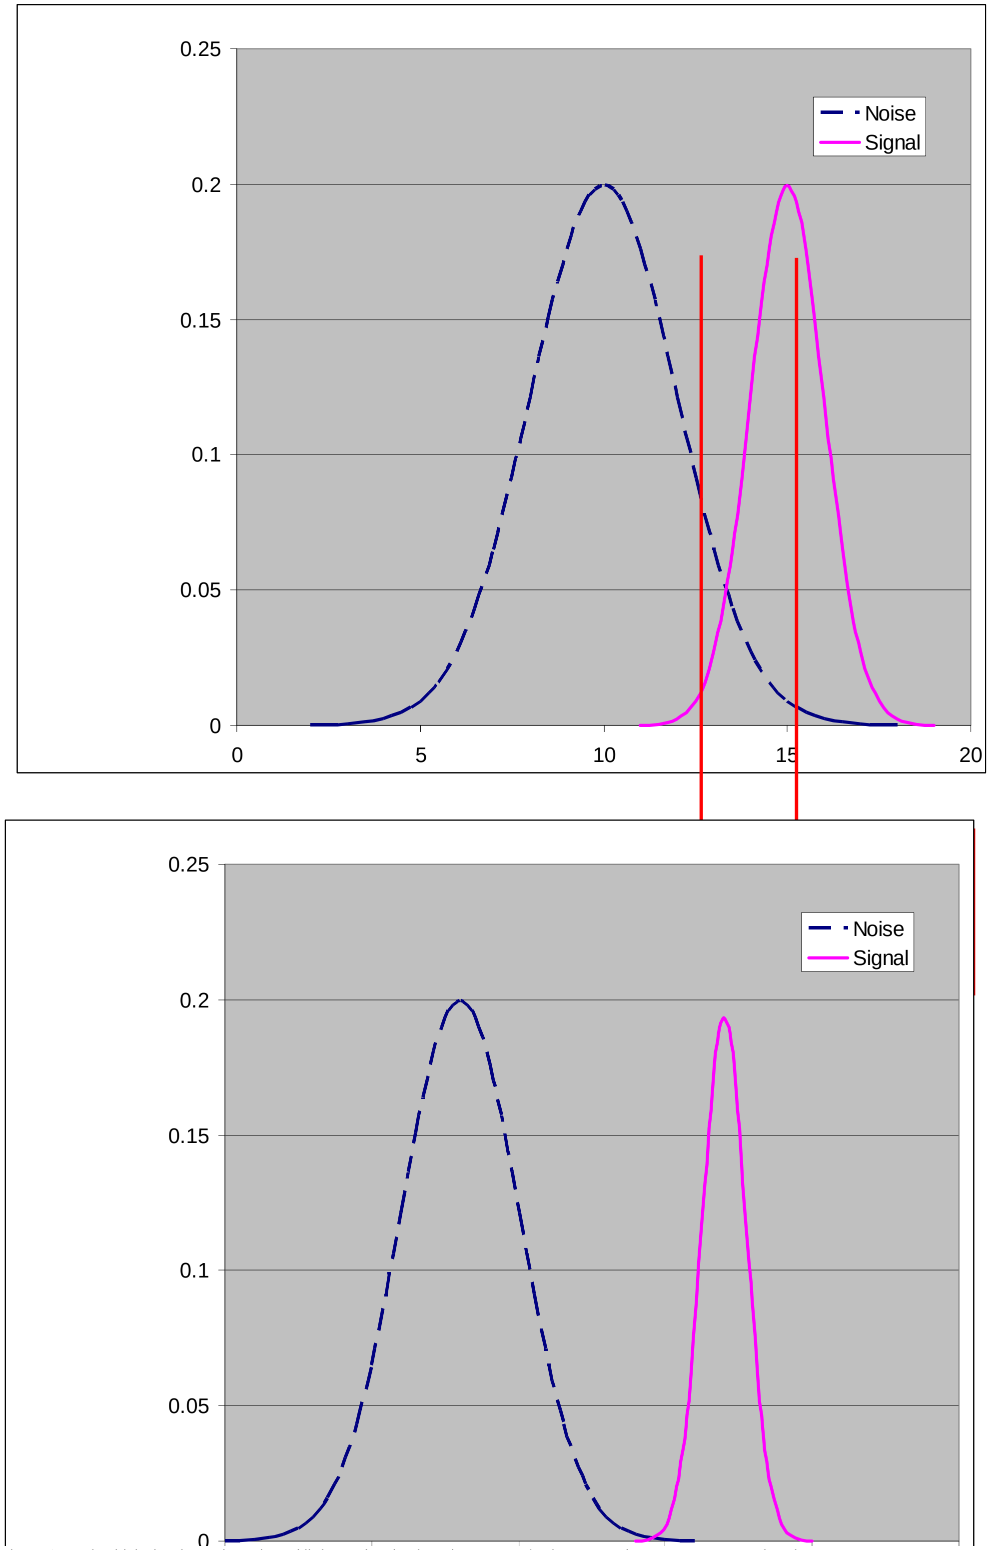

Figure 5-2 Wth a high signal to noise ratio and little overlap the detection process both type I and type II errors are reduced

The way to improve this situation is to improve the signal to noise ratio. In Improvement can often be achieved by transforming acoustic data so that the amplitude of signals and noise diverge. The simplest way in which this might be done is by using filter to attenuate the signal at frequencies which are prominent in the noise but not in the signal. More sophisticated procedures might use templates that "search for" particular wave forms or particular temporal patterns of frequency emphasis (we will see examples of these later).

A common way of depicting these trade offs between type I and type II errors is called a ROC curve and example of a set of ROC curves for a detector working in a range of signal to noise ratios is shown in Figure 5-2.

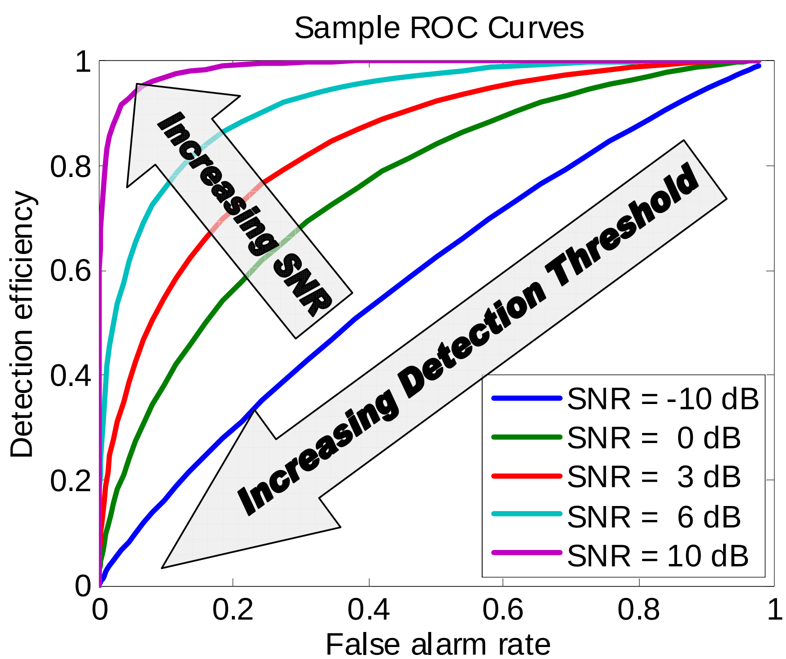

Figure 5-3 ROC curves, plots of false alarm rate against detection efficiency for different signal to noise (SNR) ratios. At low SNR a high detection rate can only be achieved at the cost of high false alarm rates. At a a high SNR false alarm rates remain low, even when low threshold are set to allow high signal detection rate. (Figure from Gillespie)

This clearly shows the trade off between detector efficiency and false alarm rate, and also the improvements that result from operating in improved signal to noise ratios.

As we've seen, this trade off between false positives and false negatives is unavoidable in any detection problem and leaves the operator having to make decision about the threshold values to choose in any particular situation. This is largely driven by the costs and consequences of making type I and type II errors. Often false positive errors can be detected by a skilled operator, but this takes time and effort that is limited in real time applications.

Encounters with cetacean groups usually results in the reception of a large number of signals, so missing a few of these may not be critical when the goal is to reliably detect a vocalising group rather than every individual signal. In these situations PAMGUARD allows other characteristics of signals to be used for detecting false positives. Detections from a group will tend to occur on consistent bearings for example and for some species, vocalizations may occur at regular recognizable patterns. We'll see examples of this in later Chapters.

In real world situations background noise is dynamic and will fluctuate with time. This means that a threshold set to give a low false positive rate in quiet conditions may result in a deluge of detections if the noise levels increase. To mitigate against this PAMGUARD continuously measures background noise and can adjust trigger levels for different detectors in relation to a statistical measure of current background noise levels.

Chapter 6 Index