AnalookW - an introduction to the software

Chris Corben

Background

AnalookW is the latest in a string of applications which make up the software component of the Anabat system. AnalookW is a Windows based program which is gradually bringing together the diverse components into a single application. Although there are still elements missing from AnalookW, it is already a much friendlier, more capable application than has been available before.

AnalookW is about viewing, analysing and managing data from Anabat Bat Detectors. It can also be used to record new data, and to view bat calls in real time. All Anabat data is stored in files, usually on a hard drive, and accessible to a PC, whether laptop or desktop. Each file consists of up to 15 seconds worth of bat call recordings, and has a maximum length of just 32 Kbytes.

AnalookW uses a multi document interface, which allows the user to display several Anabat files in separate windows at one time. Files can be manipulated in many powerful ways, making it easier to manage the large datasets which often result from long-term passive monitoring.

Setting up AnalookW

AnalookW is distributed as a single ".exe" file. All that is needed to set it up on a PC is to copy that file into a folder somewhere on the hard drive of your computer. Best use is to use two folders, one for executable files of Anabat applications, such as AnalookW.exe, and another for configuration files etc which will be used by various functions available under AnalookW.

One possibility would be to make two folders under the root directory of the C drive, one called Anabat and one called Config. Then put AnalookW.exe in the Anabat folder.

First Use

To start AnalookW, just double click on the AnalookW.exe file. Easier though, to create a shortcut on the desktop, so all you have to do is double click the shortcut. To make a shortcut, drag the AnalookW.exe file from My Computer or Windows Explorer onto the desktop with the right mouse button, then when you release the button, click on Create Shortcuts Here.

When you first start AnalookW, you may find the window is very small and various other options are not how you would like them. You can resize the AnalookW window by

dragging the borders, and you can position where you want it by dragging the title bar. AnalookW will remember these settings.

Closing AnalookW

To close AnalookW, you can click on the close button at the top, right corner, or click on File, Exit.

Working with Anabat Files

There are several ways to load an Anabat file for viewing in AnalookW. You can double click on the filename in an Explorer window (such as My Computer), and AnalookW will open if not already open, and it will open the clicked file. [NOTE: this will not work until after you have opened and closed AnalookW once.]

The other best way to open an Anabat file is to click on File, Open Anabat File, navigate to the filename and open the file in the usual way.

The loaded file will appear in its own window. You can open many windows with different (or even the same) files, and tile them, arrange them etc as you wish. Even if using just one window, it is often better to click on Window, Tile Vert or Window, Tile Horiz so the window fills the application client area. This lets you see the title bar of the window, which disappears if you maximize it. Another option is to click on View, Size in the menu, to choose a preset window size in which to display the file contents.

Changing the file in the window

There are two main ways to change which file appears in the window with the current focus. The easiest is to us the [and] keys, which take you to the previous and next files respectively. You can also use the Q and W keys in the same way - this to make it easier to use your left hand on the keyboard and right hand on the mouse. Yet another way to do this is to click on the left or right single arrow tools on the toolbar. Note that you can jump to the previous or next folder by using the double arrows on the toolbar, but there are risks in that, so I don't recommend it for most purposes. The risk is that if you have marked files, the markings will be lost.

The second way to load a new file into the same window is to use the File, Open here item on the menu, or even easier, to press shift-L. This gives you a much more flexible interface where you can navigate through the folder structure on your hard drive to any file.

Marking Files

You can mark files for deletion, copying or moving, or for header manipulation. There are several ways to accomplish this.

From the File Menu, you can click on Marking which will take you to a sub-menu which lets you perform several actions. Note that these actions are all associated with shortcut keys which illustrate how to carry out the same actions from the keyboard. For example, to mark all files in the current folder, you can press Alt-u, Alt-r. This has the effect of first unmarking all files, then reversing the marks on all files, giving the net effect of marking all files. You can achieve the same thing by clicking on File, Marking, Mark All, or by using the shortcut keys, Alt-f,r,a.

You can also manipulate file markings from the Open Here dialog. If you go shift-L, you can select one or many files and then right click to open a context menu which gives you more flexibility for file marking, since you can see the whole list of files and select just the ones you want.

Note that in both the File Menu and the context menu in the File Open Here dialog, you have another option, which is Mark if Labelled. A labelled file is any file with an entry in the Species field of the header. More of that later, but there is a lot of power in this option.

Once you have marked a file or files, you can delete them, copy them to another folder or move them to another folder through the Edit, Copy if Marked, Move if Marked or Delete if Marked menu items.

You can also make changes to the headers of all marked files by clicking on Edit, Global Header Change. This is a very useful technique for data management which will be dealt with in more detail later.

Another useful feature is to move between marked files with the - and + keys, which take you directly to the previous and next marked files respectively.

Viewing Files

There are many options for how to view the data in Anabat files using AnalookW. There are two essential points to remember about bat calls and how they look on the screen. The first point is that the appearance of any given bat call depends a great deal on the magnification at which the call is viewed. The second is that most bat call sequences have a low duty cycle, meaning that the gaps between calls are much longer than the calls themselves. This means that if you increase the magnification enough to see adequate detail in a single bat call, you will probably only have one or two calls visible at once. For many purposes, it is important to see more calls than this, because then you can see how the calls vary, and that variation is the single most important feature of bat calls and how to tell them apart.

OK, there's a third essential point. To best understand bat calls, you need to understand the temporal relationships between calls as well as the details of the calls themselves. So you need to feel free to move between different types of views and to comprehend how

they relate to each other. If you grasp this point, it will take you a long way towards understanding how to get the most out of bat acoustics.

So there are two broad classes of views - Truetime in which all the temporal details are preserved, and Compressed views in which the gaps between calls are not displayed.

The Toolbars

There are two toolbars which can be used to affect how the screen looks, one above the other. The upper toolbar can be turned off via the view menu. It contains standard Windows tools plus a few extras. The lower toolbar is always activated when a file is loaded, and provides feedback about the current view state, as well as letting you change it.

Escape

One of the most useful keys is the Escape key. This returns the screen to a known state in which the entire file contents are visible. If you get lost and aren't sure what is going on, pressing Escape will be a good starting point. You will end up in a Truetime view starting at the beginning of the file, with the entire contents spread out across the screen. From there, you can move to a particular part of the file and magnify the view to see what interests you.

Truetime/Compressed

Toggling between Truetime and Compressed views is just a matter of pressing the Spacebar. You can achieve the same thing with a tool on the toolbar - it's the tool with three curving lines (representing bat calls) and you will see it responds to pressing of the Spacebar.

Horizontal Resolution

Horizontal resolution is most easily altered by pressing one of the Function Keys 1 to 10. These keys control buttons on the toolbar, which can also be clicked to change the resolution. The toolbar buttons have the additional value of telling you which magnification is currently in use.

Another toolbar button is the All button. If this is clicked, it will make sure the entire file is visible on the screen. It has the same effect as the Escape key.

The left four buttons on the Magnification toolbar indicate a series of dots in vertical columns. These allow you to use four different modes which have a somewhat different effect than either of the regular Truetime or Compressed modes. These are called Dots Per Pixel (DPP) modes, and they cause a fixed number of dots to be drawn on each pixel column, irrespective of the time between those dots. Although these modes can be used in Truetime, they are inherently Compressed type modes in that the spaces between calls are

not drawn. They differ from the normal Compressed mode in that the time axis, even within a call is not fixed, but varies according to the frequency of the signal. This makes sense, because higher frequencies produce more dots per second than lower frequencies. The effect of DPP modes is that the higher frequencies are more magnified than the lower frequencies. This has some advantage, because bat calls at lower frequencies tend to have longer durations than those at higher frequencies. So DPP modes tend to even this out a bit. You cannot measure time intervals or the durations of calls or parts of calls in DPP modes, because the magnification is not fixed and the time scale is not linear.

When new bat calls are being recorded, DPP is the only mode available for display, as the recording process works better if it isn't necessary to calculate time intervals. Therefore, use of DPP when viewing a saved file lets you see more accurately how the calls would have looked when you first recorded them.

Frequency Scale

The frequency scale (vertical axis) is equally important for determining how calls look on screen. Unlike the horizontal scale, the vertical scale is also affected by the size of the screen you are using, which means that call appearance is influenced by window size. This results from the fact that the vertical axis covers the same range of frequencies independent of the vertical height of the window. You can see this effect by grabbing the upper or lower edge of the window and moving it up and down. The frequency range remains the same.

You can change the vertical magnification from the menu by clicking on View, Freq Scale. You can get to the same point by right clicking on the screen area to bring up a context menu. These menus give you a wide choice of preset magnifications in either logarithmic or linear scales, and they also let you set the scale to whatever you like with the Custom option.

While the choice of Logarithmic or Linear is up to you, there are real advantages to using a logarithmic scale. Many people find the idea of logarithms scary, but you don't have to understand the maths. The fact is, our senses are all logarithmic, so this is the natural way to look at bat calls! This can be illustrated by a simple point: If we are looking at coins, one centimetre is a large distance, easily perceived. If we are looking at the heights of people, one centimetre is barely perceivable and only in a direct comparison. If we are looking at tall buildings, a centimetre is completely irrelevant and totally invisible. We perceive things in proportion to the context, not as absolutes. In the same way, one kHz makes a big difference if we are looking at bat calls of 7 or 8 kHz, but is almost meaningless if the calls are at 55 and 56 kHz. A logarithmic scale reflects this reality by expanding the low frequencies and shrinking the high frequencies. As a natural consequence, we can see a very wide range of values on a single scale at a resolution which has intuitive meaning.

The Magnify button

Sometimes you might want to zoom into a particular signal to look at it in more detail. To do this, click on the magnify button on the top toolbar (it looks like a small magnifying glass). You can then drag a rectangle around the area you want to magnify, and when you release the button, the screen will zoom in on that area. To revert to the previous scale, just press the revert button on the lower toolbar. It has the letters F and T with arrows pointing left, and it will be depressed if zoom is in effect.

There are two modes for the magnify function. On the lower toolbar, there is a magnifying glass with an F inside it. If this button is depressed, then zooming in magnifies the view in both the horizontal and vertical axes. If not depressed, zooming will only affect the horizontal axis.

Linking windows

For some purposes, it may be useful to look at two or more files and compare how they look at different magnifications. This can be achieved by linking the files so when you change viewing conditions, such as the horizontal resolution, the change affects all windows at once. This function is toggled by a button on the lower toolbar which has two small, overlapping rectangles. If this button is depressed, then multiple windows will be linked.

Moving within a file

There are several ways to move within a file. For the most part, this consists of bringing a chosen item (e.g. bat call) to the left hand edge of the screen, so you can examine it more closely.

The simplest way to move within a file is to just press the left and right arrow keys. This moves you one bat call to the left or right respectively. But note that the software doesn't necessarily make a good decision about what is a bat call, and it may take you several key presses to get to where you want to be.

You can also jump bigger distances (always to the right) by using the keys 1 to 0 and a to f. These jump to specific parts of the screen. For example, if you have a screen with 15 seconds along the horizontal axis, pressing 3 will shift you to the right by 3 seconds . Pressing f would bring the right hand edge of the screen to the left hand edge.

You can move directly to the beginning of the file by pressing the Home key, or to the end of the file by pressing End. Pressing Escape also brings you to the start of the file (as does Home) but it also changes the resolution and display mode so all the file is shown.

You can use the mouse to move a particular point to the left edge of the screen by clicking on the horizontal scale just below the horizontal axis. This is especially useful if you are looking at the calls at a high magnification (horizontal resolution) because then it is easier to see where to click.

Manipulating Header Data

Anabat files contain header information which can contain various kinds of data helpful to the management of those files. This data can be displayed by checking the Header option under the View menu.

The names of the header fields are carry-overs from the earliest versions of Anabat, and some may not make much sense today. You can use them however you like, except that the Species field has some very important characteristics which have big implications for data management, so it is worth getting to know how the Species field works and what it can do for you.

Entering data into the non-Species fields, many files at a time

In typical use, you would choose to use the Tape, Date and Loc fields for data which is the same for large numbers of files. I use the Tape field to give an indication of how the files were recorded. You would also typically keep the Notes field empty except for comments which apply to a single file, or a few closely related files - for example, a series of files recorded from the same bat as it was released after capture. The Species field should be kept for special entries which classify the file according to criteria of special interest to the user. Usually, the Species field would contain information about the identity of the bat or bats in the file.

There are also some fields which are used for GPS positions. These behave differently from the other fields, and will be dealt with separately.

The simplest way to enter data into any of the fields of many files is to mark the files using any of the methods mentioned above, and then use the Global Header Change option in the Edit menu. This brings up a dialog labelled "Alter Fields for Global Header Change". The key to understanding the use of this dialog is that any field left blank in this dialog will remain unaffected by the Global Header Change. This is VERY important! If you put anything in any field of the Alter Fields for Global Header Change dialog, then that field will be changed to the new value for all the currently marked files. THERE IS NO WAY TO GET THE OLD FIELD CONTENTS BACK! You can do a lot of harm this way if you aren't careful. So practice on files which are safely backed up elsewhere.

You can also clear the contents of any field for all the marked files by entering just a single character, the tilde (~) into the field in the Alter Fields for Global Header Change dialog. The tilde must be the first and only character for this to work.

Entering data into the non-Species fields, for a specific file

If the header data is displayed on the screen, you can simply edit it directly by clicking in the field of interest and entering whatever text you like. If you then press Escape your changes will be lost and the header will revert to its former state. But if you press Enter,

or move to another file with [,], q or w, the changes will be permanently saved into the file.

Another way to start editing directly in the header data is to press the unshifted t key. This takes you straight into the header at the Species field. You can then move to other fields with the TAB or up or down arrow keys.

Shortcuts

It is often helpful to be able to enter the same data into the headers of random files. For this purpose, you can store the header data into one of two buffers, assigned to the function keys F11 and F12. Assign the data to a buffer while the header is editable (the fields are drawn on white rather than grey), by pressing ctrl-F11 or ctrl-F12 (holding down the control key while pressing one of the two function keys). To enter the data assigned to a buffer into another file, press t to enter edit mode, then either F11 or F12 and the entire buffer contents will be transferred to the header and all the old header data will be lost. You can undo this move straight away if you press Escape to exit the edit mode without saving the new header.

You can also use the F11 and F12 buffers to enter data into the Alter Fields for Global Header Change dialog. This can be useful, as it allows you to easily copy the header data in one file into a number of other files. There is a subtle distinction between using the buffers in this way and using them as previously described. If a field is blank in the buffer, it will cause that field to be cleared if you use it to edit the header of a single file directly, but if you use it to setup the contents of the Alter Fields for Global Header Change dialog, it won't have that effect, because an empty field in the Alter Fields for Global Header Change dialog will leave that field unchanged when the Global Header Change is performed.

Entering data into the GPS position fields - using AnaPocket

If you recorded bat calls using AnaPocket and with a GPS active, then the GPS data will have been automatically entered into the file header every time you saved a file. But there are other ways to get data into the GPS fields. The simplest of these is to enter the data directly in the Anahead tool.

To activate the Anahead tool, click on Tools, Anahead. This will open up the Anahead window, which contains three panels. If you can't see all the panels, click on All in the View menu - you will probably have to adjust the boundaries of the three panels to make them more useful.

The left panel is the tree view and the middle panel shows Anabat files in the folder selected in the tree view. The third panel contains information from a selected positions file. More about that later. Any actions you take in Anahead will affect only the files you select in the middle panel. All the files will be selected when you open Anahead. If you

click on Enter Manual in the Positions menu, you can enter the latitude, longitude and elevation directly into the dialog. Note that because the whole GPS interface isn't fully implemented yet, you should ALWAYS use the WGS84 datum. Latitude and Longitude are entered in decimal degrees, and elevation in metres. When you click on OK, all the selected files will get the same GPS position. Note: THERE IS NO WAY TO UNDO THIS PROCESS!

Entering data into the GPS position fields - using a CF card

If you recorded GPS data along with bat calls using a CF Card in an SD1 or Storage ZCAIM, then entering GPS data into the headers is rather different. Unfortunately, at present, the CFCread program does not automatically put the GPS data into the Anabat files it generates, so you have to do this yourself. However, you don't have to do it manually. Again, the Anahead tool is what you need to use.

Go to Tools, Anahead and you will be faced with a 3 pane window. In the middle pane, select the files you want to work with. Click on Positions, Choose File and select the GPS.TXT file which covers the time period of the files you selected in the middle pane. The GPS.TXT file will have been produced automatically by CFCread when the CF card was downloaded. Select the appropriate track in the right pane, if there were more than one, and then click on Positions, Track. This will interpolate the position for each file, based on the file time, from the track file, and insert that position into the file header data. Please note - the GPS system needs further work - at the moment you should ONLY use the WGS84 datum.

It's important to make sure the selected GPS track file contains fixes both before and after the times of the first and last files selected. If it doesn't, you will get an error. In that case, unmark the files which are outside the time window of the GPS track and repeat the process.

Generating a track file from GPS data recorded using Anapocket

If you collected GPS data using AnaPocket, you will end up with a file called GPS.abg. This file contains all the GPS data collected by AnaPocket, not just the data corresponding to when Anabat files were recorded. You can convert the GPS.abg file into a text file containing the track data by using the Tools, Converter, PDA to GPSU function. The result will be a text file which can be read directly using the program GPS Utility. This inexpensive software can be obtained from:

www.gpsu.co.uk

It offers many useful features for manipulating GPS data, changing datums, simple mapping etc.

Extracting GPS data from the GPS position fields

You can generate a text file containing a waypoint for each Anabat file. Again, use Anahead and select the files from which you wish to extract GPS data. Then click on Positions, Make Waypoint File. The result will be a file which contains a separate waypoint for each Anabat file - in other words, for each bat recorded.

Combining Track and Waypoint Data

Using an external GPS package, such as GPS Utility, mentioned above, you can combine a waypoint file containing waypoints corresponding to bat positions, with the track data collected by AnaPocket or on a CF card. The result is a file containing both the waypoints and the track - in other words, all the GPS data you need to plot both the locations of the bats and where you went looking for them. This data can be plotted directly by GPS Utility. It is even possible, using GPS Utility, to overlay this data on an image of a map or aerial photo.

A spectacular way to use this data is to employ GPS Utility to save the combined waypoint and track file to the GPX file format. You can then dump this file straight into Google Earth, just by dragging it onto the Google Earth image. Google Earth will then rotate its globe and zoom in to plot your data directly onto its images!

The Species Field

The Species field gets its own section, because it has vastly more significance than the other fields. Using it wisely can be of enormous benefits for data management, so it is well worth getting to understand it well.

The basic concept is that you can label files in the Species field to indicate the species which are represented in that file. Any file containing ANY content in the Species field is said to be labelled. You could use this label for other criteria if you like. I tend to combine criteria, by using a 4 letter code for the species, in capitals, but prefixing it with a single character which says something about how the identity was established. So if a species with the code SAFL was identified acoustically, I would enter aSAFL, whereas if I saw the bat, so was able to identify it by sight, I would use vSAFL. Other codes can be used - I use rSAFL for a bat recorded upon release from captivity, xSAFL for one recorded exiting from a roost, and so on. I also use the code to indicate the general nature of a call of interest if I don't know the identity - so q45 would mean an unknown species giving calls at about 45 kHz. A very useful code is Q - meaning just that this is something interesting which I would like to look at again.

To make the most of these codes, the essential requirement is that you only use codes which can form valid folder names. This means the codes should not include characters like ? or * or : or a comma or a period, though @ # $ and % should be OK. A space or comma will always be used to indicate the end of a label, so don't ever use those. Keep the labels reasonably short, and never more than 8 characters. You might want to put several labels in one file, for example, because there are several species in the one file.

Multiple labels can be included in one file by separating them with commas. A space ends labelling, so anything to the right of a space will not have this functionality.

Dispersing labelled files

One of the best features of the labels, and the reason why valid folder names should always be used, is that you can disperse labelled files into multiple folders each containing all the files which have a particular label. So if you have files labelled SAFL and LACI, and some with both (i.e. SAFL,LACI) then you could dispersed these into folders labelled SAFL and LACI. That way, if you want to see all the SAFL you recorded last night, you can just look in the SAFL folder and there they are! This process will take up extra room on your hard disk, because some or many files can appear in more than one folder. But hard disks are cheap and the benefits are well worth it!

To disperse labelled files, click on Disperse in the File menu. This opens a dialog which lets you select where you want the new folders to go. When you have chosen it, you will be taken to another dialog which lets you decide whether or not to delete the dispersed files from the folder they came from. Normally, you would choose NO, so copies of the files remain in the original location. If you choose YES and regret it, it may take you a while to copy all the files back from their new locations to the original folder! So it's safer to hit NO for most purposes.

Other things to do with labelled files

You can choose to mark all of the labelled files in a folder. To do that, just click on the Mark if Labelled option in the File menu. This would allow you to copy or move the labelled files, but you could also follow this action by using Alt-r to reverse the marks, then delete the un-labelled files.

You can also move quickly between labelled files by pressing the shift { and shift } buttons. These take you directly to the previous and next labelled files respectively.

Using the Species List

There are many ways to get data into the Species field of a file, but the easiest and most powerful is to use a Species list. This can be seen by checking the Species List option in the View menu. It consists of a table of buttons each of which can contain a label. All you have to do to enter a label is click on one of these buttons. That's all - the label is both entered and permanently saved into the file with just one mouse click.

If you want to enter labels for more than one species, just click on more than one button, and each additional label will be added with a comma as a separator . You can have any number of different species lists saved in different species list files, which have the extension "ars". It is a good idea to save these into a folder such as the Config folder discussed above under "Setting up AnalookW".

Getting labels into a species list is a bit tedious, but the simplest way is to right click on any of the buttons and enter the label. If you make a mistake, click on Revert and it will return to the original value. Once you've set up the list the way you want it, click on Save and give it a name and location. If you make any changes to your species list, you must save them with the Save button. You can load a different species list with the Load button.

If you are making gross changes to a Species list file, you can edit the file itself directly. The simplest way is to click on the Edit button in the Species list. This opens a very basic text editor which shows the contents of the file. Although very basic, it does allow you to use cut and paste to move things around in big chunks, which can save a lot of time. Just remember to make sure you end up with the right format. You could also use Notepad to edit the file, which gives you a few extra abilities, such as find and replace.

Remember to make sure you end up with a file which has the structure it needs. There should be comment lines like:
//Button 00
marking the start of each line of buttons, and the file should end with:
//End
However, these aren't strictly necessary, and are only there to help you organise the contents. The important thing is to make sure there aren't more than ten lines describing buttons between the //Button xx lines, or 40 lines with button names altogether. If you don't adhere to these rules, you will end up with label buttons where you didn't expect them, and some labels may not appear in your Species list. To make things easier, always keep ten lines between comment lines, and make sure every button has a line corresponding to it, even if it is just a blank line to indicate a button without a label.

You can mess up a Species List file so you lose some label data, so it is very wise to make a backup copy of a Species List before editing it this way. To do that, simply click on the Save button in the Species List, and give it a different name (e.g. add "_bak" to the filename, but make sure the extension stays as ".ars").

Other ways to manipulate labels

There are a few other keys on the Species List which can be helpful. The first of these is Replace. This works across the entire current folder to replace all examples of a specific label with a different label. Suppose you labelled a hundred files with LACI, but then realised you were half asleep and should have labelled them SAFL. All you need to do is replace LACI with SAFL by clicking the Replace button, and entering the relevant text.

Undo lets you undo the most recent label change you made to the current file.

Clear gets rid of everything in the Species field of the current file.

Finally, there are four Save buttons and their associated Buf buttons at the right hand edge of the Species List. By using a Save button, you can store the entire set of labels in

the current Anabat file to a buffer, then recall that set of labels into another file with the corresponding Buf button. This lets you copy the a label or set of labels from one file to maybe many other files. There is an important difference between the two upper buffers (with Bufx+) and the lower two buffers (Bufx-). The plus sign after the buffer name indicates that these buttons ADD the buffer contents to the current label in a file. The minus signs indicate that these buttons REPLACE the current label with the buffer contents. If this sounds confusing, just play with it (using copies of backed up files, so you can't do any harm!) until you understand it.

Recording

It is quite simple to record directly into AnalookW if you have a serial port and a serial ZCAIM (such as a CF Storage ZCAIM or an SD1 detector). Just connect the machine to the ZCAIM in the same way you would if you were using CFCread to communicate with the ZCAIM (e.g. to set its time).

Then all you have to do is click on the Record button (a square with a red dot in it) on the upper toolbar and AnalookW will find the ZCAIM and start recording in the current window, or it will open a new window for recording if there is no file open. You can save a file by clicking on the Save button (next to the Record button, with a green disk being pointed at by an arrow), and you can stop recording by clicking on the Stop button, which contains a black rectangle.

Relevant recording parameters can be changed under the Record tab in Tools, Options.

If you start recording with Monitor under the Record menu, the software itself will save files whenever certain criteria (controlled by the Monitor Mode Parameters in the Recor Options) are met. This lets you save all the bat files without having to even be there, but it carries the risk that some sequences of bat calls will be broken up across two or more files.

Buddy window

When recording, it's often useful to be able to look at the last file which was saved without stopping recording. This can be achieved using a Buddy Window. Just click on Make Buddy under the Record menu, and a new window will pop into existence. Whenever you save a file while recording, that file will be automatically shown in the buddy window.

Making measurements

There are many ways to make measurements of call characteristics, and more are under development. Here are some of the most important features.

The Cursor

Make sure the Status Bar is visible by checking the Status Bar option in the View menu. Now you can see, in the lower right hand corner of the screen a small window (Cursor Window) which contains information about the position of the cursor, amongst other things.

As you move the cursor around the screen, the Cursor Window will always show a time and a frequency. The time is the absolute time at the position of the cursor, from the beginning of the file. The frequency is just the position of the cursor relative to the vertical scale. you should always use a Truetime screen when making measurements.

There is also another parameter, shown as "st= nnnn", which gives a value representing the position of the first point displayed at the left edge of the screen in relation to the file contents. Its primary use is to let you tell someone else where a particular event lies within a file.

Measuring Time and Frequency Displacements and Slopes
If you want to measure the difference in time or frequency between two points, you can select a rectangular area of the screen (provided the Selection Tool is activated - it is a dotted rectangle on the upper toolbar). Then the Cursor Window will show a different display which gives the width of the rectangle in seconds, the height of the rectangle in kHz and the slope of the straight line from first point where the rectangle was started to the second point where the left button was released, in Octaves Per Second. Note that for any selected rectangle, there will be two possible slopes, one from the top left to the bottom right, and the other, which has the same value but opposite sign, from the bottom left to the top right. However, the slope always depicted in the Cursor Window will be that from the first point clicked to the second point. There will never really be any ambiguity, because only one of the rectangle's diagonals will have a slope with the same sign as that in the Cursor Window.

Slope Display

You can activate the Slope Display by clicking on the Slope option in the View menu. You can then toggle the Slope Display on and off by pressing the v key.

The Slope Display shows the instantaneous slope of each displayed call on a scale of: -1000 to 1000 Octaves Per Second (OPS). The Slope is calculated by effectively sliding along the call looking at the slope between the start and end of a line which includes a certain number of dots. The number of dots depends on the Transitions parameter, which can be changed by right clicking on the Slope Display and clicking in the resulting context menu.

Transitions are the number of zero-crossings in the original bat call, not in the frequency-divided signal which is actually analysed. S0, the number of dots over which the slope is calculated is equal to the Transitions parameter divided by the Division Ratio. The reason

for doing it this way is that the results are then independent of the Division Ratio being used.

TBC Display

You can activate the Time Between Calls Display by clicking on the TBC option in the View menu. You can then toggle the TBC Display on and off by pressing the v key.

The TBC Display shows you a histogram of the distribution of different TBCs for the sequence displayed on the screen. There are a number of ways of showing this data, but they fall into two main categories, depending on whether the vertical axis is the time between calls or the rate at which the calls are being produced. So calls with a TBC of 100 ms would have a call rate of 10 calls per second. These display parameters can be chosen from a context menu called up by right clicking in the TBC display.

You can also choose the Min Gap parameter. If two calls are separated by less than the Min Gap, they will be treated as one call for the purposes of the TBC display.

You can also choose the width of each bar in the display with the Pixels Per Bar parameter.

Cycles Display

You can activate the Cycles Display by clicking on the Cycles option in the View menu. You can then toggle the Cycles Display on and off by pressing the v key.

The cycles display (I have sometimes referred to it as an Antislope Display) effectively gives an indication of the frequencies where the signals on the screen spend most time. So the peak in a Cycles Display corresponds to the frequency where the absolute slope of a call is at its minimum.

There is a context menu which lets you choose between a number of parameters which affect the display. The first of these is Smoothing, which combines the values in adjacent frequency bands to smooth out some of the raggedness which can occur without smoothing. Note that a consequence of smoothing is that it is possible for a very distinct peak to be shifted very slightly in frequency.

The next parameter determines whether times or cycles are displayed. Because higher frequencies produce more cycles per second than lower frequencies, the cycles option tends to give more emphasis to higher frequencies while the time option tends to give more emphasis to lower frequencies. For most purposes, they are similar, but the time option is much more susceptible to noise at low frequencies, so the cycles option is often less distracting.

At the right hand edge of the horizontal axis, there is a single frequency parameter which gives the frequency of the peak of the graph. To its left, there is a value which represents

the total number of cycles (or milliseconds) within the peak frequency band. The band is the frequency range represented by one pixel in the vertical axis. Note that therefore this value will be affected by the vertical scale and the window size. It is pretty meaningless anyway, unless just a single call is being displayed.

Further left are three values which correspond to the three vertical bars drawn across the display. These bars are at values of one half, one quarter and one eighth of the peak value, and the values attributed to them are the bandwidths in kHz of the main peak at each of those points. These values are primarily affected by how steep the call is - a steep call will have broader bandwidths while a flat call will have narrower bandwidth.

Getting Anabat data as text

Anabat files contain zero-crossings data in a very terse, efficient form which is why the files are so small. For some purposes, though, you might want to see what the data actually contains in terms of the XY co-ordinates of the points on the screen. To do this, activate the Select tool, which is the dotted rectangle in the upper toolbar. Then drag a rectangle around the part of the display which you want to examine (e.g. a single call - easier in Truetime mode) then right click on the rectangle to bring up the context menu, and click on Copy. This copys the data from the dots inside the rectangle onto the clipboard in text form. You can then just paste it into any text editor (e.g. Notepad) or a variety of other applications. You can even paste the data straight into a spreadsheet such as Excel.

The text form of the data is just two columns, the first being the time in seconds since the beginning of the pasted data, the second being the frequency in Hertz.

Making Images of Anabat File Contents

At this stage, there is no simple way to print directly from AnalookW. You can try the File, Print menu standard printing interface, but will probably find the image too small to be useful. A much more powerful way to print is to copy the screen contents to the clipboard (by pressing the Prnt Scrn key), then edit it after pasting it into an external image editor. Even Windows Paint will work quite well for this.

There are a few things to be aware of when doing this. Firstly, it helps to use a standard window size to give consistency between your images. Select the window size by clicking Size in the View menu. Bear in mind that a larger size gives more resolution, but many of the screen elements are of fixed size, and may seem too small in the resulting image.

Secondly, you may want to change the screen colours before printing - e.g. to make a black and white image. This can be done in the Display tab under Tools, Options.