

Bell Waveform Analysis Program
Documentation version 1.7
24 October 2004
Copyright © Bill Hibbert, Great Bookham, Surrey

1. Introduction

The purpose of the bell waveform analysis program (Wavanal.exe) is the analysis of bell sounds using a personal computer. It provides facilities for graphical display of recorded bell sounds, identification of partial frequencies, and synthesis of bell sounds from a list of partials.

It allows complete determination of the harmonic character of a bell using the facilities available on any multi-media home PC. It also allows exploration of the way that changing the tuning of a bell changes its resulting sound.

The minimum requirements for use of the program are as follows:

- Windows 95, 98, 2000, XP or NT 4.0
- Sufficient CPU and RAM to run standard applications
- About 95 kbyte of disc for the basic program
- 250kbyte - 1mbyte of disc for each sound file
- Sound card (used to capture and play sounds only)
- Mouse (keyboard-only operation of the program is not possible)
- 640 x 480 VGA screen (800 x 600 SVGA recommended)

Library files mfc42.dll and msvcrt.dll need to be present in the \windows\system directory. If not already loaded, these files can be downloaded from the web (search for mfc42.dll). Between them they occupy about 1 mbyte of disc space. No installation of the Wavanal program is required. Simply copy Wavanal.exe to any convenient directory and run it from there.

The only additional equipment required is a tape recorder or video camera (to capture the sound of the bell or bells) and a cable capable of connecting recorder or camera to the sound card. Mono sound only is used. Using a laptop with sound card and a microphone, it is possible to capture the required sound files direct. Wavanal has the facility to record sounds direct from a PC or laptop microphone.

The remainder of this document suggests how to take recordings of bells for analysis, explains how the program works, documents how to use it and how the various screens work, and finally provides some technical detail on the transform techniques used and some restrictions on the results.

My website www.hibberts.co.uk gives many examples of the use of Wavanal.

2. Capture and preparation of bell sounds

Wavanal can be used either with pre-recorded bell sounds stored on a PC hard-drive, or to record, analyse and save sounds direct from a PC microphone. Sections 2.1 and 2.2 relate to use with pre-recorded sounds. Section 2.3 covers recording from a microphone.

2.1 Taking recordings

Wavanal works with quite poor quality recordings, but the better the recording, the more useful and accurate the result. The key issues are:

- The recorder used
- lack of distortion and external noise
- how the bell is sounded
- length of recording.

Recordings of bells rung singly are needed. It is possible to analyse a recording of bells rung in changes but the task is very time-consuming.

Recording equipment

Reasonable results can be achieved with a cassette tape recorder. However, disadvantages are that speed stability of such recorders is not good, and that automatic level control, if fitted, can lead to very noisy recording. Prior to sounding the bell, and as its sound dies away, the automatic level control kicks in and produces an unacceptable level of hiss.

I use an 8mm video camera to take recordings with very acceptable results - the frequency stability is very good. Other people have also had success using a laptop computer with a microphone to digitise direct to disk.

In all cases, placing the recorder too near the bell will overload it and cause distortion - which can introduce 'false' partials and other strange effects. 3 to 4 metres away from the bell is not too far. A buzzy or sizzly sound in the resulting recording is a sure sign of overload.

Taking the recording

I usually work with bells already hung in the tower. The harmonic content of a bell sound depends critically on where the bell is struck, and to most faithfully capture the intention of the bellhanger, the technique used is to strike the bell once with its own clapper (with the bell down, of course!). To give most accuracy, a recording of at least 3 seconds duration is needed. It is desirable to keep wind, traffic etc. noise to a minimum during the recording.

I also have considerable success with analysis of recordings of bells rung 'up', taken in the ringing room. Attack and decay times and relative intensity of partials will differ from those of the same bell chimed, but results for partial frequency are comparable.

If very accurate determination of frequencies is required, a reference tone (e.g. a handbell or tuning fork of known frequency) should be recorded before and after the bell sound.

2.2 Digitising recordings

The recorded sounds should be digitised using a PC sound card. A cable is required to connect tape recorder to sound card - available from any hi-fi or electronics shop. The aim in digitising the recording is to capture the sound of each bell (if more than one has been recorded) in a separate sound file for processing. The sound file for each bell should include a short period of silence prior to the clapper stroke, followed by 3 to 10 seconds of the bell sound as it slowly dies away.

Capturing the sounds

I use Syntrillion's Cool Edit (shareware) with a 16-bit Soundblaster sound card to digitise recordings. Windows Sound Recorder also works but does not allow editing of the recording. It may be necessary to adjust the recording level in the sound software to achieve distortion free recordings. For best results, digitise as 16-bit mono samples at the maximum sampling rate - 44.1 or 48 kHz. If you have recorded reference tones (handbell or tuning fork) then these should also be digitised.

Files may be saved as raw PCM or in any compression format supported by your PC which can be saved as a .wav or .mp3 file - see the next section.

Once you have digitised the sounds, play the sound files. If they play back OK with no distortion, you are ready to begin analysis.

Sound file formats

Wavanal can read any waveform files with a .wav or .mp3 extension supported by the software on your PC that match the following parameters:

- any digitisation rate
- mono or stereo sound
- 8 or 16 bit samples.

All are converted to 16 bit mono, which is the format used within Wavanal. This conversion is done by Wavanal, not using standard Windows drivers, because:

- it avoids a bug in the windows driver that affects some stereo to mono conversions
- it allows Wavanal to read uncompressed files on a PC with no audio drivers installed.

Wavanal can also read a wide range of compressed formats provided your PC has the right drivers installed. Check ‘Settings | Control Panel | Multimedia | Devices | Audio Compression Codecs’ for a list of supported formats. The benefit of compression is that it considerably reduces the file size without significant loss of quality. The main formats you will encounter are:

Format	Typical compression	Wavanal results	Driver
Raw PCM	none	Very good	None required
A-law and μ -law	2 to 1	Good	Installed as standard
MS ADPCM	4 to 1	Good	Installed as standard
GSM	8 to 1	Good	Installed as standard
MP3	12.5 to 1	Good enough	May need special driver
Speech codecs	-	poor	Installed as standard

You may need to install an MP3 codec (search the web for ‘fraunhofer mp3 codec’) if one is not installed. Wavanal supports both MP3-encoded .wav files and files with the .mp3 extension.

Waveform files written by Wavanal are 16 bit, mono, raw PCM format.

2.3 Use via directly connected microphone

The new release of Wavanal supports recording of sounds direct from the PC microphone, via the 'Record Waveform' facility. The ‘Analyse Direct’ function which was a prototype for the Tuner program has been removed. The documentation on this screen below gives the details of operation. In summary, when this screen is invoked, Wavanal configures the PC microphone to record in 16 bit mono at 44100 samples per second.

The input level can be monitored via Wavanal, but adjustment of the level should be done via the Windows Mixer facility (activated by clicking the microphone icon on the taskbar). Using a few test rings of the bell, the mixer should be adjusted so that the initial bell sound approaches full-scale, but does not activate the ‘overload’ indicator. It may be necessary to reposition the microphone to get good results.

Once a bell sound has been recorded and analysed as described in the screen documentation below, the facilities of Wavanal to re-analyse, and save sound files, transforms and partials can be used in the normal manner.

3. Analysing the sound of a bell

Wavanal provides facilities for detailed analysis of a bell, investigation into the changes over time of the amplitude of each partial etc. If all that is required is determination of a bell's main partial frequencies, proceed as follows:

- digitise the sound of the bell into the PC as described in the previous section
- run Wavanal, and click the 'Analyse Bell' button towards the top right of the screen
- use the 'Open' screen which is displayed to select the bell sound previously recorded
- Wavanal will then process the bell automatically, and display the resulting partials, with the major partials named.

If Wavanal does not name the partials, it is likely that the hum has been missed – this partial is often quiet. Use the 'View Tr, / Get Part.' screen to display the transform, select the hum, use 'Add Partial' to add it to the list and display the partials again with 'View / Edit partials'.

To check that the nominal has been correctly identified, use Wavanal to create a tone of half the nominal frequency. If this tone sounds the same pitch as the bell, all is OK.

Wavanal often displays the note name corresponding to a frequency (e.g. F(2) +22 which corresponds to a frequency of 707.5 Hz. These note names are based on either $A(1) = 440\text{Hz}$ or $A(1) = 435\text{Hz}$. The figure in brackets indicates the octave, and the following figure, lying in the range -50 to +50, indicates the number of cents away from the exact note.

The use of all the screens is explained in some detail below.

4. The Wavanal program

This section of the document explains how the Wavanal program works, and some of the theory behind its operation. For operating instructions and screen descriptions see section 4 below.

4.1 Data storage and basic facilities

Wavanal can store and process four types of information about a bell sound:

- the bell waveform, equivalent to a .wav file
- its transform, the sound amplitude at a range of frequencies
- the partial frequencies and intervals between them
- information about the way each partial's amplitude varies over time.

The sound is stored as a series of values at the chosen sampling rate. Each value gives the amplitude of the sound waveform at that instant. The transform is a list of frequencies and values, and is created by carrying out a series of arithmetic calculations (a Fourier Transform) on the waveform data. The

sound of the bell is assumed to be made up of a set of partial frequencies. Four parameters for each partial describe how its amplitude changes over time.

The Wavanal program provides facilities to independently read each data format from a disc file, write it to file, to view and edit the data in a graphic or tabular format as appropriate, and to carry out transformations between the various formats. Sound files are read and written in .wav format. Sounds loaded into Wavanal can be played if a sound card is fitted to the PC. Transforms and partials are written to disc in comma separated value (CSV) format to allow further processing, charting etc. using a spreadsheet package such as Excel or Lotus. Transform files are saved using file extension .csv. Lists of partials plus associated information are saved using file extension .par. Detailed information on how each partial amplitude varies over time can also be saved in a CSV file with extension .dec for processing outside Wavanal.

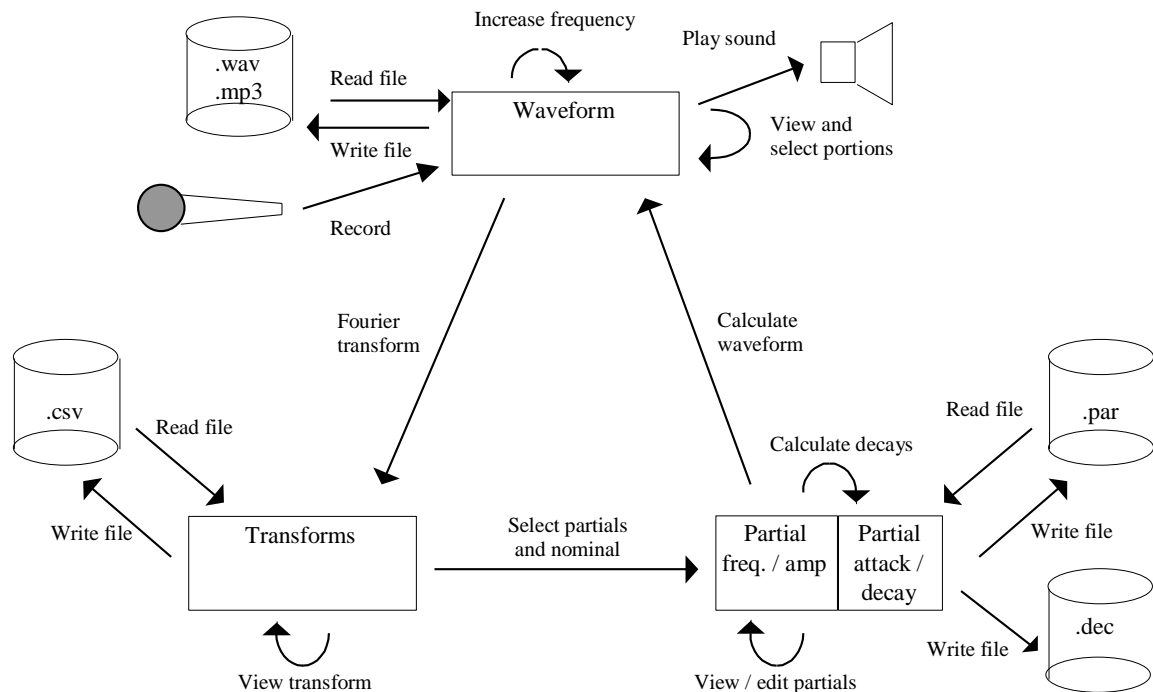
4.2 Compensation for recorder speed variation

Variations in the record or playback speed of the recorder used, or (possibly) errors in the digitisation rate of the PC can result in incorrect determination of frequencies. If such variations exist the waveform file will be stretched or shrunk, and each partial frequency will be increased or decreased by the same factor. Wavanal provides a facility to correct for this.

To compensate for variations, determine the apparent frequency of the reference tone (see section 2.1) by digitising its sound and transforming it with Wavanal. Calculate the ratio of the apparent frequency thus obtained to the real frequency of the reference tone. Load the bell sound to be analysed, and invoke the 'Increase Frequency' function described below. Wavanal will stretch or shrink the waveform by this factor to ensure the partial frequencies are correct. The modified waveform can be saved under a new filename if required.

4.3 Transformations between data formats

The diagram shows the transformations possible between the three formats.



4.3.1 Waveform to transform

This transformation is carried out by performing a Fourier Transform on the waveform data. The program allows selection of the portion of the waveform to transform, the frequency band to examine, and the frequency interval between each transform. Wavanal is configured with a set of defaults for these values to get you started. Once the transform calculations have been done, Wavanal provides facilities to view the transform graphically and select partial frequencies.

For technical reasons described below, the accuracy with which the frequency of a partial can be determined depends on the length of the sound sample being transformed. Long samples are needed for accurate frequency determination.

The program uses Fourier Transforms without regard to phase, i.e. it generates the amplitude of the sound at each frequency in a selected range. A Hamming window is used to compensate for end effects. The length of the transform is determined by the frequency interval chosen. If a short section of waveform is selected the remaining part of the transform is padded with zeros. If too long a section of waveform is selected the extra data at the end is ignored. The default parameters used in Wavanal have been chosen to minimise the practical impact of these issues.

4.3.2 Transform to partials

The view transform screen has facilities to automatically extract partials from the transform, and for the user to select and deselect individual frequencies as partials. Wavanal has storage for up to 25 partials. If a 26th partial is selected, the partial with the lowest amplitude is replaced. The five 'Simpson' partials are automatically identified from among the partials selected.

The partials information extracted by the view transforms screen consists of the partial frequency and its transform amplitude. Wavanal can also carry out a series of transforms on short sections of the waveform to determine how the amplitude of each partial varies over time.

A screen is provided allowing partial frequencies, amplitudes, attack times and decay rates to be viewed and edited. This screen also allows input of a base frequency from which intervals in cents can be calculated for all partials, and identification of the 'Simpson' partials.

4.3.3 Partial to waveform

The last transformation supported by Wavanal is that from partials to waveform. Using the assumption as to format of partials mentioned in the previous section, a set of calculations is performed to recreate the waveform from the set of partials held in Wavanal. This waveform overwrites any waveform loaded from disc. The program allows input of sampling rate and length of sample before calculation, and creates 16 bit mono samples. Once the waveform has been created, it can be saved to disc or used in further processing in Wavanal as if it were a .wav file.

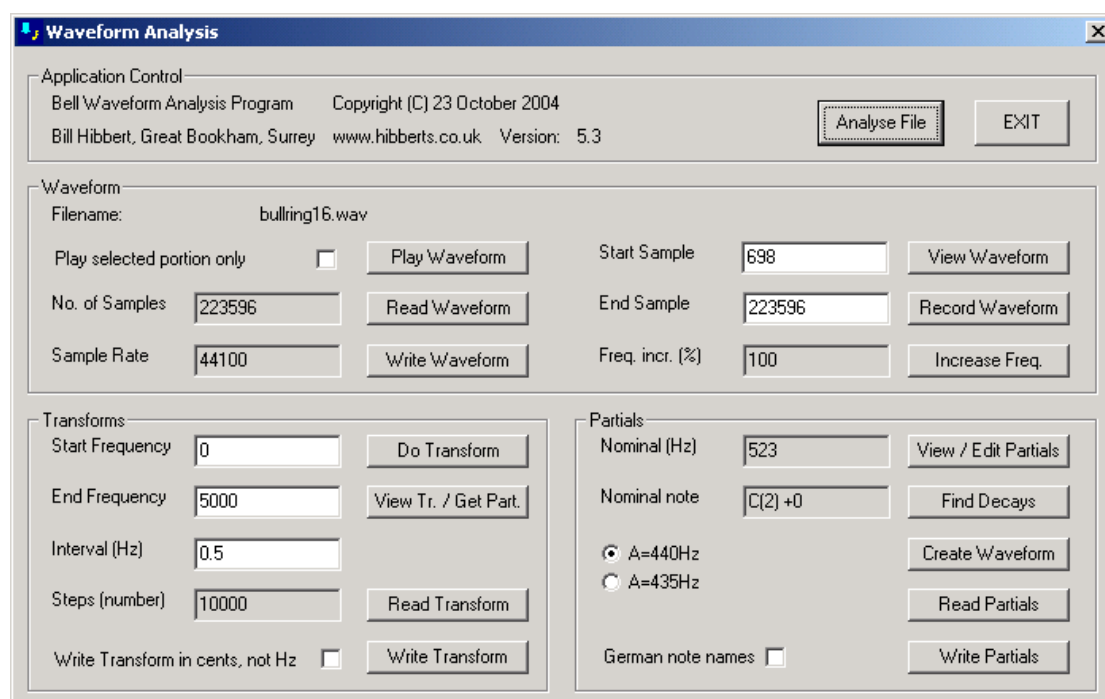
This facility allows the ends results of analysis to be compared aurally with the original. It also allows details of the partials to be changed and the results evaluated.

5. Description of major screens

This section of the document is divided into two sections. The first section covers all areas of Wavanal needed for tonal analysis of bells, including automatic identification and naming of the main partials and the new ‘Analyse Direct’ facility. The second section on advanced functions describes those parts of Wavanal used for investigation into the changes over time of partial amplitude, and the simulation of bell sounds. The information in the second section can be ignored unless you are interested in this area of experimentation.

5.1 Tonal analysis using Wavanal

5.1.1 Main window



The main window is always displayed when Wavanal is running. It provides for entry and display of key parameters and control of the various functions.

There are four areas of the main screen, corresponding to the three types of data used by the program, plus a section about the application itself.

Pushbuttons

Some pushbuttons are disabled (i.e. greyed out) if the data required for that function has not been loaded or created.

Each of the three types of data can be read from or written to disc using the ‘Read’ and ‘Write’ buttons in each section. Hitting these buttons displays the standard dialog box allowing the user to browse the disc, select files, and enter filenames. Various error messages are displayed if a file cannot be read or

written. The 'Write transform in cents, not Hz' checkbox enables writing the transform file to disc in a special format used to produce plots of transform intensity with normalised frequency and amplitude.

The 'Analyse File' pushbutton provides automatic determination of the partials and naming of the major partials for a bell. It uses several of the functions described below as processing takes place. The success of automatic processing depends on the quality of the recording. If partials are missed or more detailed investigation is needed the remaining screens below should be used.

The 'View Waveform', 'Record Waveform', 'View Transform / Get Partial', 'View / Edit Partial', 'Find Decays' and 'Create Waveform' buttons invoke the screens described below. The 'Play Waveform' button plays the waveform currently held in Wavanal through the sound card if one is fitted. If the 'Play selected portion only' option is chosen, only the portion of the waveform defined by 'Start sample' and 'End sample' will be played. The 'Do Transform' pushbutton starts the process of transform calculation using the parameters from the Waveform and Transforms sections of the main window. The 'Increase Freq.' button brings up a screen that prompts for a correction factor by which the waveform should be stretched or shrunk to compensate for recorder speed variations. The 'EXIT' button prompts the user for confirmation before closing the application.

Data Fields

Information is displayed about each type of data held in Wavanal. Data can be directly entered by the user into those fields with a white background. Fields with a grey background are derived or calculated from other data.

The 'Filename' field shows the name of the sound file currently loaded, if any, and indicates if the waveform has been modified. 'No. of Samples' shows the length of the waveform currently loaded. 'Sample Rate' shows the digitisation rate of the waveform in samples per second. 'Start Sample' and 'End Sample' are used to select a portion of the waveform for processing. Wavanal tries to set 'start sample' to the moment of clapper strike when the sound file is read. The user can either enter values directly to these fields, or use the 'View Waveform' screen to do the selection. 'Freq. incr.' shows the percentage by which the waveform has been stretched or shrunk if the 'Increase Freq.' function has been used.

The 'Start Frequency', 'End Frequency' and 'Interval (Hz)' fields are used to set the parameters for the transform calculation. All values are in Hertz (i.e. cycles per second). The 'Steps (number)' field is calculated from the other three if valid data has been entered.

The 'Nominal (Hz)' field displays the nominal frequency of the bell. This nominal frequency is identified automatically when partials are selected from the 'View Transforms / Get Partial' screen, and can be changed manually on the 'View / Edit Partial' screen. The nominal note name is also displayed.

Note names are calculated from a base of either A=440Hz, or A=435Hz, as selected using the radio buttons. If the check-box for German note names is selected, then the note a semitone above A is called B, and that a semitone below C is called H. If German note names are not selected, these two notes are called Bb and B respectively. The default values, each time the program starts, are A=440 with German note names disabled.

5.1.2 Record Waveform



Whenever this screen is displayed, Wavanal is monitoring input from the PC sound card, with the input parameters set as 16 bit mono at 44100 samples per second. If a suitable input device cannot be found, an error message is displayed when the screen is invoked.

The volume of the input sound is shown on the 'volume' indicator. If an instantaneous overload occurs, the 'overload' indicator is set for two seconds. Recording can either take place immediately, or when the input sound reaches a trigger threshold, set by the slider, dependant on the setting of the 'wait for trigger' clickbox.

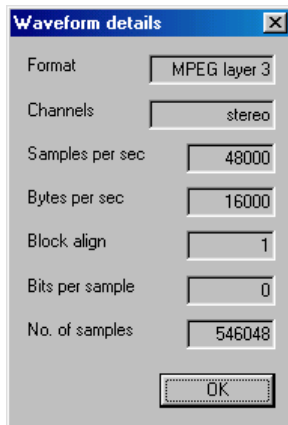
No recording takes place until the 'Start Recording' button is clicked. Recording then starts immediately or after a trigger, as selected. Recordings can be cancelled at any time by clicking the 'Cancel Record' button. The 'status' display can show the following values:

- Stopped – if Wavanal is waiting for the user to click 'Start Recording'
- Waiting for trigger – if Wavanal is awaiting a trigger
- Recording – if Wavanal is recording the sound.

The 'Samples' display shows the number of samples collected to date while Wavanal is recording.

The 'Exit' button returns to the main Wavanal screen.

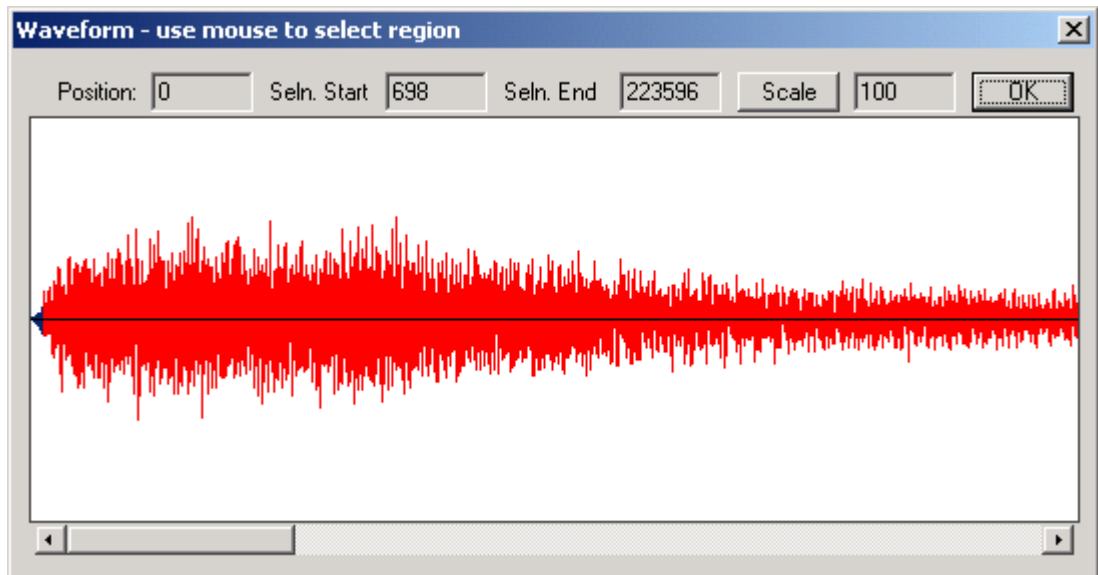
5.1.3 Waveform details



This screen is displayed when a waveform file is opened from the main screen. The screen displays technical information from the file header. Click on the OK button to return to the main window. If the file header is not that of a valid .wav or .mp3 file or your PC cannot convert the file format, error messages are displayed instead of this screen.

Wavanal reads any sound file formats for which codecs are installed. To read MP3 compressed files, you may need to install an MP3 codec - see the section above on sound file formats for details. Both mono and stereo files can be read, at any digitisation rate. Raw PCM files must be 8 or 16 bit. The number of samples is estimated from the file header, the exact number of samples loaded may be different for compressed files.

5.1.4 View waveform



This screen is invoked by the 'View Waveform' pushbutton. It allows the waveform to be viewed and a portion selected for processing.

The 'Position' field shows the start point (i.e. left-hand side) of the display. 'Selection Start' and 'Selection End' show the start and end of the selected

portion. The units of all three fields are samples. The selected portion of the waveform is shown in red, the unselected portions in blue. Clicking the ‘Scale’ button repeatedly cycles through the possible values of the scale field to the right and re-scales the graph accordingly.

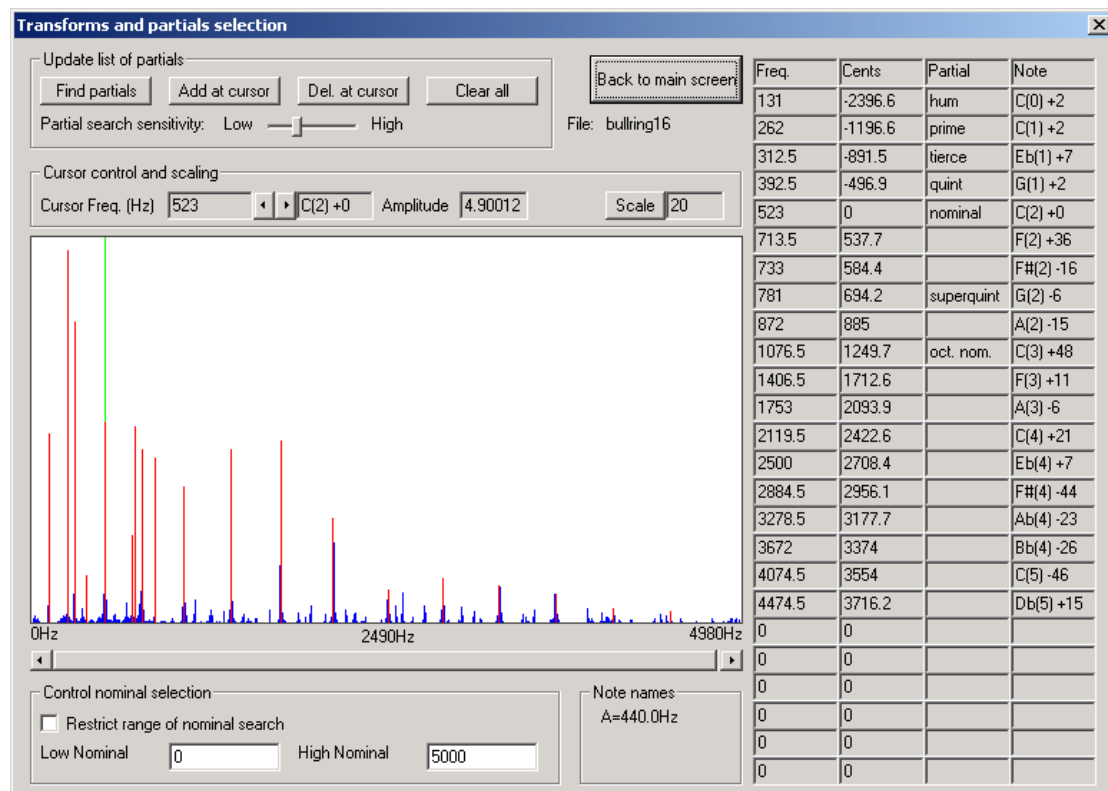
The slider bar beneath the graph allows the display to be scrolled.

Clicking and dragging with the mouse within the graph selects a portion of the waveform. If the mouse is placed in the border of the window on either side of the graph while dragging, the display scrolls. Releasing the mouse button completes the selection.

The ‘OK’ button returns to the main window.

Note that if your waveform shows an initial period of steady amplitude followed by the expected exponential fall-off then it is probable that you recorder has overloaded when the bell was first struck, or the sound was digitised with the volume set too high. The periodic fluctuations in amplitude visible in the waveform above are due to beating between partial doublets. This effect is often seen in older bells.

5.1.5 View transform / get partials



This screen allows the results of a transform to be viewed, and partials selected. The graph shows transform amplitude (on the vertical axis) against frequency (on the horizontal axis). A cursor, coloured green, is initially at the left-hand side of the screen and can be moved either by clicking on the graph, or using the arrows to the right of the ‘Cursor Freq. (Hz)’ field. The ‘Cursor

Freq. (Hz) field and 'Amp'de' fields show the frequency and amplitude of the largest transform at the current cursor position. The note name corresponding to the frequency is also displayed.

The 'Scale' button and the slider bar at the bottom of the display work in an identical fashion to those on the 'View Waveform' screen. The display is centred on the cursor if the scale permits.

The 'Find Partial', 'Add at cursor' and 'Del. at cursor' pushbuttons are used to select partials. Frequencies currently selected as partials are highlighted in red. The 'Find Partial' button attempts to find up to 25 partials from the transform provided enough data is present. The selection uses some rule-of-thumb judgements and the results should be checked visually before leaving the screen. The 'Add at cursor' button adds the biggest transform at the cursor position to the list of partials. The 'Del. at cursor' button deletes all partials at the current cursor position from the list of partials. 'Clear all' deletes all the partials.

The 'Find Partial', 'Add at cursor' and 'Del. at cursor' pushbuttons also cause the five 'Simpson' partials including the nominal to be identified if possible. If this identification fails (usually because one of the lower partials is too quiet to be selected, or has a frequency very different from the norm) the nominal frequency is set to zero. The nominal can be forced to be in a range of frequencies by entering values in 'Low Nominal' and 'High Nominal' and selecting 'Restrict range of nominal search'. The slider below the partial selection buttons can be used to change the sensitivity of the partial search; slide it to the left if too many partials are selected, or to the right if too few, and click 'Find Partial' again.

Note names are based on A=440Hz or A=435Hz as selected on the main screen and shown in the 'Note names' area. If German note names are selected (Bb = B, B = H) then a comment to this effect also appears in this part of the screen.

'Back to main screen' returns to the main window.

5.1.6 View / Edit partials

Edit Partial

Nominal Frequency (Hz) Note

Freq. (Hz)	Splash Amp.	Tail Amp	Attack	Decay	Cents	Partial	Note
131	0	4.59923856	0	0	-2396.6	hum	C(0)+2
262	0	9.02469744	0	0	-1196.6	prime	C(1)+2
312.5	0	7.33033351	0	0	-891.5	tierce	Eb(1)+7
392.5	0	1.19576101	0	0	-496.9	quint	G(1)+2
523	0	4.90012833	0	0	0	nominal	C(2)+0
713.5	0	2.15396168	0	0	537.7		F(2)+36
733	0	4.77637160	0	0	584.4		F#(2)-16
781	0	4.22306745	0	0	694.2	superquint	G(2)-6
872	0	4.01356342	0	0	885		A(2)-15
1076.5	0	3.34678831	0	0	1249.7	oct. nom.	C(3)+48
1406.5	0	4.22419993	0	0	1712.6		F(3)+11
1753	0	4.45627443	0	0	2093.9		A(3)-6
2119.5	0	2.57773904	0	0	2422.6		C(4)+21
2500	0	0.83080577	0	0	2708.4		Eb(4)+7
2884.5	0	1.11444567	0	0	2956.1		F#(4)-44
3278.5	0	0.95977385	0	0	3177.7		Ab(4)-23
3672	0	0.73226158	0	0	3374		Bb(4)-26
4074.5	0	0.39441775	0	0	3554		C(5)-46
4474.5	0	0.32779274	0	0	3716.2		Db(5)+15
0	0	0	0	0	0		
0	0	0	0	0	0		
0	0	0	0	0	0		
0	0	0	0	0	0		
0	0	0	0	0	0		
0	0	0	0	0	0		
0	0	0	0	0	0		
Noise	0	0	0	0			

OK Recalc A=440.0Hz Clear

This screen allows display and editing of partials and associated information, and selection of the nominal frequency. The information displayed is of two types. The columns Freq (Hz), Tail Amp, Cents and Partial contain data which is derived from the transform via the find, add and delete partials buttons on the View Transform / Edit Partial screen. The columns Splash Amp, Attack and Decay allow additional information to be manually entered about the time profile of the partial, allowing accurate simulation of bell sounds. The final column displays the note name corresponding to each frequency. The note names are based on either A=440Hz or A=435Hz as selected on the main screen and shown on the right hand side. If German note names are selected, this is also show.

Partials are displayed in frequency order. A partial is deleted by setting its frequency to zero, and additional partials can be added at the end. The list is sorted next time the screen is called up. The cents values are calculated from the nominal frequency when the 'Recalc' button is pressed provided the nominal is non-zero.

The nominal frequency can be directly entered in the field at the top of the screen. Any value can be entered in this field, allowing the ‘Recalc’ button to be used as a general purpose cents calculator.

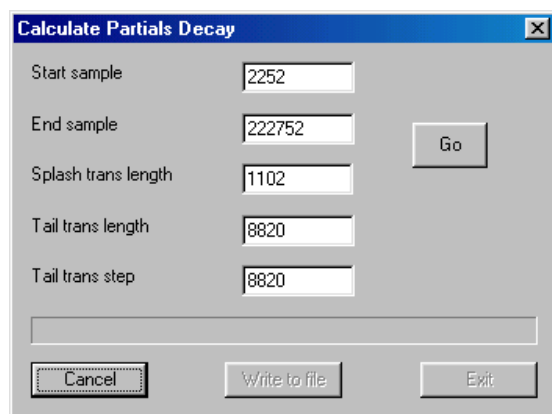
The ‘Clear’ button deletes all partials from the list. The ‘OK’ button returns to the main window.

The Splash Amp, Tail Amp, Attack and Decay parameters are used when creating a waveform. For an explanation of the meaning of these parameters see the section below on advanced functions.

5.2 Advanced functions

As described above, the functions described in this section will be of use if you are interested in the detail of how individual partials vary over time, or in the simulation of bell sounds. These advanced functions are not automatic in operation and require some knowledge to use.

5.2.1 Find decays



This function carries out a sequence of transforms over periods of time and produced an array of values showing the time variation in the partial amplitude. This array of values cannot be accessed from within Wavanal. It is intended to be written out to file as a CSV file for processing and examination in a spreadsheet package. Using the ‘find decays’ function does not change the values (e.g. for partial amplitude) displayed by Wavanal.

The dialog above is displayed when the ‘Find Decays’ button is clicked. It is used to set up parameters for the calculation of partial decay rates.

The ‘Start sample’ value is picked up from the main window. ‘Start sample’ should be set, via the ‘View Waveform’ screen, to the exact moment the clapper hits – for most waveforms, Wavanal will have done this automatically when the soundfile was loaded. The ‘End sample’ is set to cover five seconds of the waveform.

The function carries out a sequence of short transforms at the start of the sound (to capture detail around the clapper strike) followed by a set of longer

transforms for the rest of the time period. The start and end points of the transforms, based on the parameters in the screenshot above, are:

start to start + splash
start to start + 2 * splash

etc. until the end point is just less than length, and then:

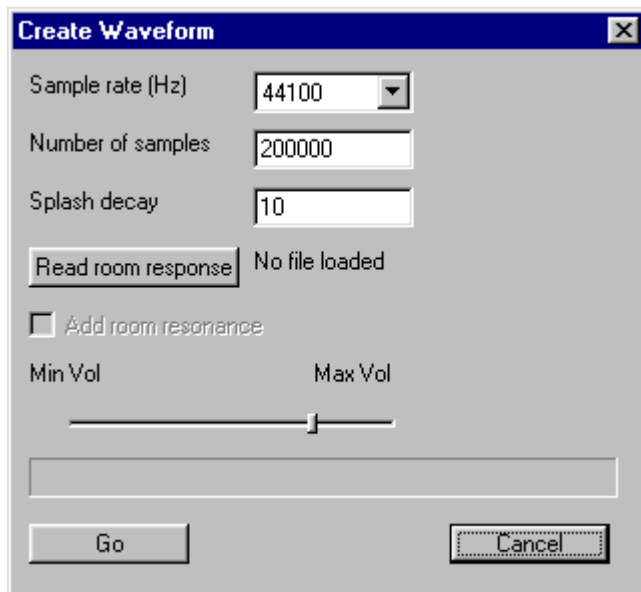
start to length
start + step to length + step
start + 2 * step to length + 2 * step

etc. until the end point is less than or equal to end. If splash is set to zero only the second, longer, set of transforms is carried out.

The 'Go' button starts the calculation. The 'Cancel' button can be clicked at any point to terminate the calculation. When the calculation is complete, the 'Cancel' button is disabled and the 'Write to file' and 'Exit' buttons are enabled. 'Write to file' writes the calculation results to disc as a .csv file with file extension .dec for more detailed examination.

'Exit' returns to the main window.

5.2.2 Create waveform



This dialog is displayed when the 'Create Waveform' button is clicked. The sample rate selector box is used to pick the digitisation rate at which the waveform will be produced. The 'Number of samples' field sets the length of the generated waveform. Only 16-bit waveforms can be produced.

The 'Go' button starts the calculation. Any existing waveform held in Wavanal is deleted when the calculation starts. The 'Cancel' button can be clicked at any point to abandon the calculation.

If you want to understand how the room response and resonance functions are used, contact the author.

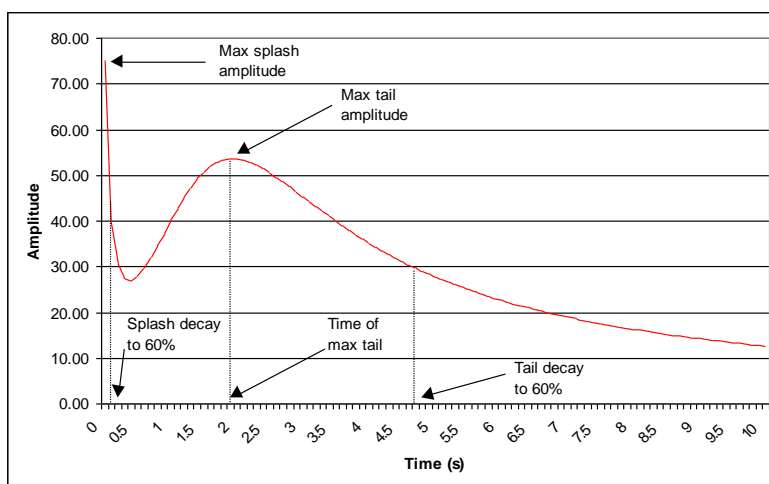
5.2.3 Modelling the amplitude of partials

To explain how waveforms are generated it is necessary explain the use of the remaining parameters for each partial which were skipped over in the description of the ‘View / Edit Partials’ screen.

The amplitude of a partial can start high as the clapper strikes and decay to a low value - the splash. It can then rise again to a higher value, and decay again over a much longer period of time – the tail. Some partials have both splash and tail, and others one or the other.

Partials from an actual bell can also display cyclic variation in amplitude due to doublets. It is not necessary to model this behaviour via additional parameters, as the partial that is the second of the doubletted pair can be added to the mix instead.

This parameterised model of partial amplitude has been arrived at as the minimum set of information which allows reasonable simulation of bell sounds to take place. An example appears below:



The bell sound is divided into splash (the clapper strike) and tail (the elongated sound that comes later). The parameters held by Wavanal for each partial are the splash amplitude, tail amplitude, attack and tail decay. The splash decay is defaulted to a pre-determined value (currently 0.25s).

The various parameters relate to the graph as follows:

Parameter	Units	Physical explanation
Splash Amp	Amplitude	The maximum amplitude of the splash
Tail Amp	Amplitude	The maximum to which the tail rises
Attack	Seconds	The time at which the tail is maximum
Splash decay	1 / Seconds	The inverse of the time taken for the splash to fall to 60% of its maximum (defaults to 0.25s)
Tail decay	1 / Seconds	The inverse of the time taken for the splash to fall to 60% of its maximum

To simulate the sound of a bell, it is necessary to:

- use find decays to create a .dec file for the partial structure of the bell
- plot graphs (using a spreadsheet package) of the partial time variations
- read the various parameters off the graphs as above
- manually increase the splash amplitudes for the rim partials.

To explain about rim partials goes beyond the scope of this documentation. See www.hibberts.co.uk for details on this. With practice, very good simulations can be achieved, suggesting that the internal dynamics of the bell have been successfully modelled.

6. Restrictions on transform algorithm used

Wavanal uses a modified Fast Fourier Transform calculation to derive transform (i.e. frequency and amplitude) information from the waveform. The transforms are carried out without regard to phase. There are three restrictions of which the user needs to be aware, though the default settings used in Wavanal should preclude problems due to these effects in any but extreme cases.


6.1 Nyquist effects

Nyquist's theorem is that transform information is only correct at frequencies up to half the sampling frequency. For example, a sound digitised at 44kHz will give unreliable transform information at frequencies above about 20kHz. A sample digitised at 11kHz will give erroneous transforms at 5,500 Hz and above.

Experience suggests that restricting transforms to rather less than half the sampling frequency can be advisable, but that bells have no significant partials above about 5 kHz. Therefore, digitising at 22 kHz, or 44 kHz for best results, is recommended.

6.2 Frequency resolution and end effects

The results of the transform process depend on the length of the waveform sample used. A waveform file containing a single frequency will give a transform whose shape is shown in the diagram.



The transform will be centred approximately on the actual frequency in the waveform file. The spacing between the ripples is given by:

(ripple spacing in Hz) = (sample frequency in Hz) / (length of transform in samples)

5000 samples of a waveform digitised at 44 kHz will give ripples of 8 or 9 Hz, which sets the precision at which a partial frequency can be estimated. 200,000 samples of a waveform digitised at 22 kHz (i.e. about 10 seconds of sound) gives a ripple spacing of about 0.1 Hz. Therefore, long sound samples are needed to give accurate frequency determination.

6.3 Frequency interval and number of samples selected

The data length used in the transform calculation is determined by the frequency interval selected. For example, for an interval of 0.5Hz, two seconds of data is required. For an interval of 0.1Hz, ten seconds of data will be required.

If the actual period of time in the waveform differs from this, one of two actions occurs:

- if less data is selected than is required (e.g. if a short section of waveform is being investigated) the remaining data values input to the transform are set to zero. If the selected sample length is significantly short, the peaks in the transform will be broadened, and the accuracy of frequency determination will be less than anticipated
- if more data is selected than required, the extra data beyond the end of the required length is ignored.

Care should be taken in interpreting results if the length in seconds of the selected portion of the waveform differs significantly from the reciprocal of the frequency interval.

7. Program restrictions

The current version of the program (version 5.7 dated 15 October 2004) has no significant restrictions. It supports all common .wav and .mp3 compressed and uncompressed formats. It uses a new way of selecting partials from the bell spectrum which works better with noisy recordings.