TUNER

A PC-based Spectrum Analyser for the tuning of bells

User Documentation

Bill Hibbert
Great Bookham, Surrey, UK
July 2002
web: www.hibberts.co.uk
email: bill@hibberts.co.uk

Contents list

Bell Tuning Program
Documentation version 1.0
6 August 2002
Copyright © Bill Hibbert, Great Bookham, Surrey

1. Introduction

This program (Tuner.exe) is a PC-based spectrum analyser specially developed to assist in the tuning of church bells. It provides facilities for recording of bell sounds as they are being tuned, spectrum analysis, selection and identification of partials, and display of spectrum and tuning information. Recordings and information on partials can be saved to disc. All the information required during the tuning of a bell is displayed on a single screen for convenience and speed of operation.

The minimum requirements for use of the program are as follows:
Windows 95, 98, 2000 or NT 4.0
Sufficient CPU and RAM to run standard applications
About 70 kbyte of disc for the basic program
A sound card and microphone, necessary for capture of the bell sound
Mouse (keyboard-only operation of the program is not possible)
800 x 600 or larger VGA screen
100kbyte - 1mbyte of disc for each sound file saved to disc.
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 Tuner program is required. Simply copy Tuner.exe to any convenient directory and run it from there.

The remainder of this document explains the program's operation and provides technical documentation on various aspects.

My website www.hibberts.co.uk gives much additional information on the recording and analysis of bell sounds. You can also email me on bill@hibberts.co.uk for further information and advice, to report program bugs, or make suggestions for improvements.

I am very grateful to the following people who have contributed significantly to the development of Tuner:
- Ben Kipling of Hayward Mills, who made the initial suggestion for the program and helped with its early development
- Andrew Higson of Taylors, who provided valuable feedback and suggestions after using an early version of the program
- Stephen Ivin, who provided advice on the latest version.

This documentation relates to version 1.0 of the program.

2. Principles of operation

Tuner uses the PC sound card with an attached microphone to capture sounds for analysis. No special facilities in the sound card other than basic input are used, and very good results have been reported with low-cost microphones, or those built into a laptop.

Whenever Tuner is running, data is input from the microphone via the sound card. A level meter in the program allows the recording level to be monitored. (No program other than Tuner must be accessing the sound card.) The standard Windows Mixer is used to set the recording level as appropriate.

The user can set various trigger conditions and parameters for the spectrum analysis. Tuner starts with a set of defaults which should be adequate for most purposes. When triggered, Tuner records a period of sound, analyses and displays the spectrum, attempts to identify the five main partials, and displays information about them.

The user can scale and scroll the spectrum display to allow details (such as doublets) to be examined. Additional frequencies can be selected as partials if they have been missed. If Tuner has difficulty in automatically identifying the partials, a frequency range for the nominal can be entered by the user to force identification of the partials.

Tuner supports a calibration factor for the sound card where needed for the most accurate measurements.

A second screen allows details of all partials to be displayed. Both the recording, and the identified partials, can be saved to disc for archiving or later study.

2.1 Recording equipment

As mentioned above, Tuner works with any PC sound-card and compatible microphone, and low cost equipment seems to give very good results. The microphone used must give reasonable response down to the lowest frequency of interest (which can be below 100 Hz for the hums of large bells) but linearity of response is not a prime requirement.

The sound of a bell, even struck gently, can overload most microphones if they are too close. Overload can lead to poor analysis and even spurious partials. This can sometimes be avoided by changing the sound-card configuration to reduce the input amplitude. However, with cheap microphones it is likely that the microphone will need to be a little distance from the bell – say 5 to 10 metres away.

The Windows Mixer can be used to set an appropriate recording level. It may also be necessary to make sure the right input is selected. For details see Appendix 1.

If you have a PC with multiple sound cards, note that Tuner selects the first for recordings. Operation of the program in this configuration has not been tested.

The recording format used by Tuner, and the format of sound files saved to disc, is 44,100 samples per second, 16 bit, mono, raw (uncompressed) PCM.
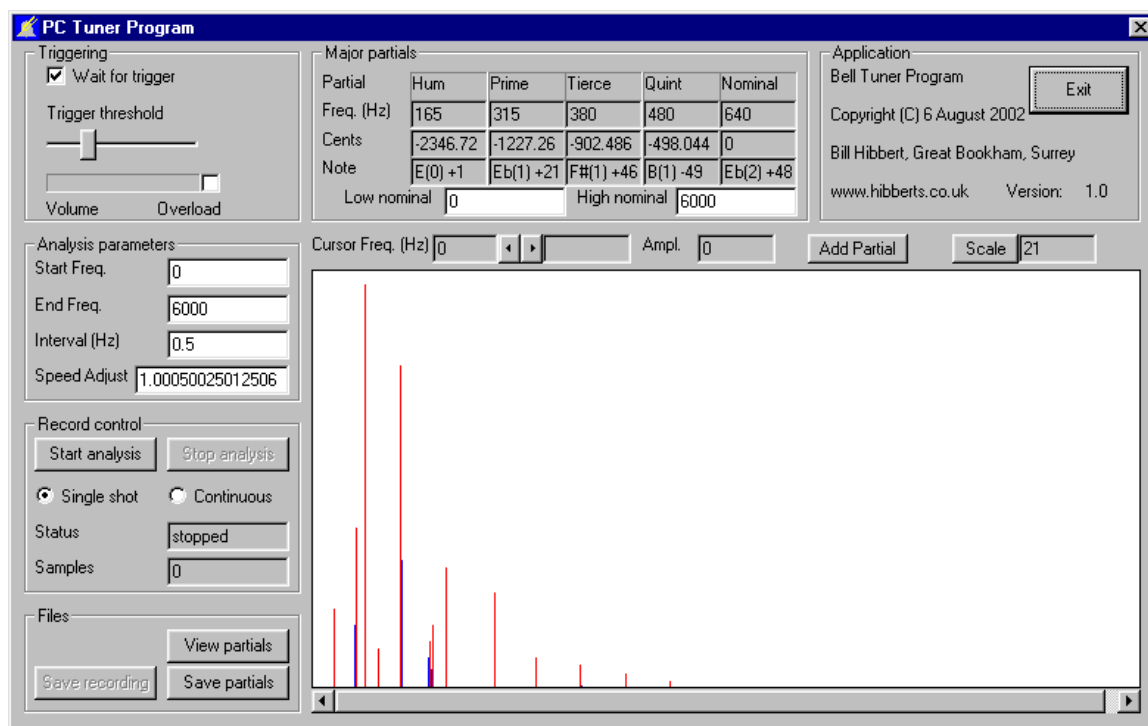
I have had some trouble on PCs, especially laptops, with mains hum, which shows as a set of strong frequencies in the analysis at mains frequency (50 or 60Hz) and its multiples. There are two ways to avoid this effect:

- On laptops or other battery powered machines, disconnect the mains when taking recordings (this is rather impractical as other than a one-off expedient)
- Connect a reasonably heavy earth lead between the PC chassis and ground on the electrical outlet supplying the machine.

You should consider consulting a qualified electrician before making changes to the mains wiring of your PC.

3. Operation of the Tuner program

When Tuner is run, the following screen is displayed.

The screen is divided into a number of areas, each of which will be described in turn. Note that a 'dummy' spectrum and its partials are displayed to allow familiarisation with the program without the need for bell recordings.

If instead of the main screen you see one of the following error messages:

there is some problem with Tuner accessing the soundcard on your PC. Make sure that a soundcard is fitted, it is recognised by Windows (e.g. you can play sounds through it), it supports the sound-file format used by Tuner, and that no other program is accessing it at the same time as Tuner.

The remaining parts of this section describe the various areas of the main screen, as follows:
- Application
- Triggering
- Analysis parameters
- Record control
- The spectrum display, cursor operation, scaling etc.

- Major Partials
- Files.

Various pushbuttons and fields on the screen are greyed-out and inoperable if their use does not make sense at a particular point. Data can be entered in fields with a white background. Fields with a grey background are display only.

3.1 The Application area

This area of the screen displays information on the program and its version. The 'Exit' button closes the program, as does the Window Close button at the top right. If either of these is clicked, the user is prompted for confirmation before program close in case data needs to be saved to disc.

3.2 The Triggering area

Tuner monitors the soundcard continuously and displays the instantaneous volume of the input sound on the 'Volume' display. This allows operation and input volume of the microphone to be checked. The 'Overload' marker is set for a brief period every time the input sound exceeds a preset level. If this marker appears when analysing a bell sound, either turn the input level down with Windows Mixer (see Appendix 1) or move the microphone further away from the bell.

If the 'Wait for trigger' checkbox is selected, Tuner waits for a trigger threshold in the input sound before starting to record. The threshold is set with the 'Trigger threshold' slide control. The scale of this slider is the same as the volume display immediately below it.

To use triggering, choose 'Wait for trigger' and set the trigger threshold above the volume of ambient noise but below the volume of the initial sound when the bell is struck. Then, once recording has been enabled as described below, it will start when the bell is struck.

If 'Wait for trigger' is disabled, there is no trigger threshold. Recording starts as soon as it is enabled. If continuous recording is selected, Tuner begins the next recording as soon as it has displayed the previous spectrum. Enabling of recording and selection of continuous recording are described below in the Record control section.

3.3 The Analysis parameters area

The parameters in this area are set to defaults which should give reasonable results when Tuner is run.

The 'Start Freq.' and 'End Freq.' parameters set the start and end frequencies between which Tuner will carry out a spectrum analysis and look for partials. They also provide the start and end frequencies for the spectrum display until this display is scaled and / or scrolled by the user. If it proves necessary to

examine a frequency outside this range, it will be necessary to re-record the bell after changing these parameters. The largest value of End Freq. permitted is 22,000Hz.

The 'Interval (Hz)' parameter sets the spacing between adjacent analysis frequencies, and also by implication sets the length of the recording taken. If 'Interval' is set to 0.5 Hz, for example, partials will be determined to the nearest 0.5 Hz, and adjacent frequencies on the spectrum display will be 0.5 Hz apart. The length of the recording is the reciprocal of Interval, and in this case will be 2 seconds.

To avoid excessive use of PC resources and long recording and analysis times, the smallest value of 'Interval' supported is 0.05Hz (requiring a 20s recording), and the maximum number of transform steps, i.e. the quantity (End Freq – Start Freq) / Interval, is 80,000. To look at a spectrum with a small interval such as 0.05Hz it will be necessary to reduce the frequency range before recording the bell.

Interval can be set to any value you wish above 0.05Hz, but for speedy operation the largest value giving the required measurement accuracy, and sensible frequency steps, is recommended.

The parameters in this area of Tuner must be set while recording is not enabled, and changes are disabled while recording and analysis is in progress. Erroneous values in Start Freq., End Freq. or Interval (or in Low nominal and High nominal) cause an error message to be displayed if a recording is attempted.

The Speed Adjust field and frequency calibration of your PC are described in Appendix 2. They can be ignored unless you require very accurate frequency measurements.

3.4 The Record control area

In this area, the user starts and stops the recording and analysis process (though Tuner always monitors the input from the microphone) and can see the progress of the recording.

Tuner has two recording modes, single shot and continuous. In single shot mode, one recording is taken and the results analysed, either as soon as the 'Start Analysis' button is clicked, if 'Wait for trigger' is not selected, or after a trigger is received.

In continuous mode, repeated recordings are taken and analysed (waiting for a trigger if selected) until the 'Stop analysis' button is clicked. Tuner displays a progress box when it has taken a recording and is doing the spectrum analysis. Clicking 'Cancel' on this progress box has the same effect as clicking 'Stop analysis'. In either case, the current recording may be lost.

The 'Status' display shows whether recording is stopped, waiting for a trigger or in progress. The 'Samples' display shows the number of samples collected as recording takes place.

While recording and analysis are in progress, many of the fields and buttons on the screen are disabled.

3.5 The spectrum display, cursor operation, scaling etc.

On start-up, Tuner displays a dummy spectrum with its associated partials. When a recording and its analysis is complete, the actual spectrum is displayed, based on the analysis parameters selected. The display is of amplitude (in the vertical direction) against frequency (in the horizontal direction), both with linear scales.

The spectrum is drawn in blue. Frequencies which Tuner has selected as partials are shown in red. A cursor (which may be placed anywhere in the graph area by clicking with the mouse) appears as a long green line.

The cursor frequency appears in the box marked 'Cursor Freq. (Hz)'. The cursor may be moved in the graph area using the left and right arrow buttons to the right of the frequency display. The cursor moves repeatedly if the button is held down.

The next field shows the note corresponding to the cursor frequency. The note is displayed as note name (octave) +- error in cents. The note names assume concert pitch (i.e. 440Hz is A in octave 1). The errors in cents are given assuming equal temperament - cents are hundredths of a semitone, for the uninitiated. Note names are not given for frequencies below about 33 Hz.

The 'Ampl.' Field shows the amplitude of the spectrum at the cursor frequency in arbitrary units. The main use of this field is in spotting very low amplitude partials as the cursor moves across them.

If Tuner has missed a partial (see the next section for the implications of this) it can be selected by placing the cursor on it and clicking the 'Add Partial' button. Tuner will then add the partial, if relevant, to the 'Major partials' display. This is especially useful for finding quints, and sometimes hums, missed by Tuner.

The spectrum display can be scaled and scrolled. It is initially displayed full scale (i.e. with start and end frequencies at the left- and right-hand sides of the display). The 'Scale' button cycles through a set of scale values, attempting to centre the display on the cursor where possible (useful for looking at an individual partial at different scales). If the graph is at other than full scale, the slider at the bottom of the graph scrolls it.

3.6 The Major Partials area

When Tuner completes analysis of a recording, it attempts to identify the five main partials. When the identification is successful the results are displayed in this area of the screen. Ignoring the 'Low nominal' and 'High nominal' fields for the moment, the information displayed down the column for each partial is as follows:

- The partial name
- The partial frequency in Hz
- The interval in cents from this partial to the nominal
- The note name for the partial frequency.

The note name display follows the same rules as those described above for the cursor on the spectrum graph.

Tuner is quite successful at identifying partials correctly, but can fail or pick the wrong frequencies in the following circumstances:

- The recording is noisy, so there are too many likely candidates
- The hum, prime, tierce or nominal were too quiet to be selected as partials
- The intervals between partials are very different from the common values.

If Tuner has missed a partial, say the hum or quint, it can be selected from the spectrum display. The major partials area should then display all the partials.

The 'Low nominal' and 'High nominal' are used to force Tuner to select a nominal in a particular range if the automatic identification is failing. If these fields hold the same values as 'Start Freq.' and 'End Freq.' in the Analysis parameters area (which are the default values), Tuner uses automatic identification. If 'Low nominal' and 'High nominal' are set to a restricted range, Tuner takes the most prominent partial in this range as the nominal and attempts to find candidates for the other major partials based on this selection.

The user is allowed to enter a range for the nominal frequency to allow continuous detection of the nominal while the bell is tuned. Note that identification of the major partials may still fail if there is not a suitably prominent partial, or more than one, in the selected range.

To return to automatic identification of the nominal, reset 'Low nominal' and 'High nominal' to the same values as 'Start Freq.' and 'End Freq.'. Entering erroneous values in 'Low nominal' or 'High nominal' causes an error message to be displayed if a recording is attempted.

Note that the quints of many bells are very quiet when they are struck on the soundbow. Strike the bell elsewhere to be certain of the identification of this partial.

3.7 The Files area

The three buttons in this area allow the user to:
- View a complete list of up to 25 partials, with frequencies, amplitudes, intervals from the nominal, partial name and note name
- Save this list of partials to disc
- Save the current recording to disc.

When saving files, the user is prompted for a filename. The list of partials is saved as a comma-delimited file with the extension '.par'. This format is used to allow the partials to be loaded into a spreadsheet or other program for further analysis, investigation etc. (The partials data is slightly different from that saved by Wavanal).
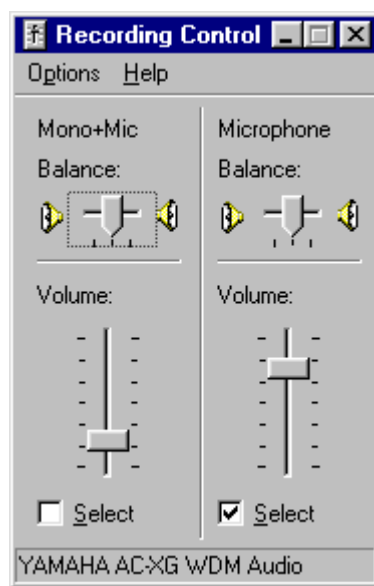
The recording is saved as an uncompressed PCM file in Windows .wav format. The recording parameters used are 44,100 samples per second, 16 bit, mono. The length of the recording and the .wav file in seconds is the inverse of the Interval parameter – i.e. a 0.5Hz interval gives a 2s recording, and to achieve a 10s recording it is necessary to set Interval to 0.1Hz. If the Speed Adjust parameter is not 1.0, the recording will have been resampled in accordance with the parameter value.

Wavanal can of course be used for further analysis of sound files saved by Tuner.

Appendix 1

Windows Mixer, recording levels and input selection

PCs have a recording level control and the facility to select different sound-card inputs for recording. These facilities are in the Windows Mixer which is accessed by double clicking the loudspeaker icon on the Start bar. There are different versions of this program, but typically what appears first is the playback mixer. To access the record mixer follow Options : Properties and click the 'Recording' radio button. Then use either the microphone or line-in slider depending on which input socket you are using. The settings appear to be saved automatically when they are changed. The picture below shows the record control on my laptop.

Selecting Mono+Mic instead of Microphone allows me to record louder sounds without overload. I understand that with some sound cards and microphones it is possible to plug the microphone into the 'line-in' socket to achieve the same effect.

Appendix 2

Soundcard calibration and the Speed Adjust field

Casual users or those not requiring accurate frequency measurements can ignore this area of Tuner's operation.

Accurate determination of frequencies by programs such as Tuner depends on the rate at which the soundcard digitises the sound from the microphone. Importantly, accuracy does not depend in any way on the speed of the PC's processor clock. Practical experience has shown that sound cards show slight variations from the theoretical rate used by Tuner of 44,100 samples per second and that the frequency errors can be detected. For instance, the PC on which this document is being written reports a precisely known frequency of 600Hz as 599.7Hz. Some of the very cheapest sound cards have proved to be the most accurate!

Tuner can compensate for this using the Speed Adjust field and a calibration frequency via the following procedure:
1. Set the Interval to 0.05Hz (and the Start and End frequencies to some restricted range including your calibration frequency)
2. Record the calibration frequency and note the exact frequency as measured by Tuner
3. Calculate the sum (actual frequency / measured frequency) to many decimal places (I recommend 10 or 15 places)
4. Enter this figure in the Speed Adjust field
5. Tuner then automatically resamples each recording to give the correct frequency (as may be verified by re-recording the reference frequency).

As a guide, the Speed Adjust value for this PC (which measures 599.7Hz for a 600Hz reference) is 1.000500250125063.

This speed correction is used until Tuner exits. To make the calibration permanent, follow the following procedure:
1. Using Notepad or another text editor, create a file containing the Speed Adjust value (note that the value can be copied and pasted from the Windows Calculator to save retyping it)
2. Save this file under filename 'correct.txt' (without the apostrophes!) in the same directory as Tuner.exe.

Tuner will now load this Speed Adjust value every time it starts. If the file 'correct.txt' does not exist, or does not contain a valid number between 0.9 and 1.1, Tuner uses the default value of 1.0.

When a sound file is saved from Tuner, the resampled version is written to disc, so that if (for example) it is transferred to another PC the sampling rate and frequencies will be correct. If the sound is played back on the PC which recorded it, and the output data clock of the soundcard has the same error as

the input clock, the frequencies played will be fractionally incorrect. It is not thought that there are any practical implications of this.

The only remaining issue is that of obtaining a reference tone with precisely known frequency with a useful duration of at least 20s. Tuning forks and bells of known frequency are possibilities, but the author has had excellent results with the American National Standards reference tones, which are very accurate and can be accessed via a telephone call. Details are given in Appendix 3.

Appendix 3

Frequency calibration using the US NIST standards

The US National Institute of Standards and Technology broadcasts standard frequencies, linked to the equipment which produces the US time standard. These broadcasts can be accessed via telephone and serve as high integrity, easily accessible and cheap frequency standards.

A description of the service is available from the NIST Time and Frequency Division website http://www.boulder.nist.gov/timefreq/ from which the following details are taken. To hear a simulcast of the WWV shortwave broadcast, call (303) 499-7111 (from the UK, this is 001-303-499-7111). This is not a toll-free call, except in the local Boulder/Denver, Colorado area. WWV broadcasts precise 1-second time intervals (the interval between the "ticking sounds" you hear) and standard audio frequencies of 440, 500, and 600 Hz. Part of each minute is taken up with time and weather announcements, but the tones last for tens of seconds, giving ample opportunity to measure their frequency.

The schedule for transmission of the tones repeats every hour, at the following times. Only at around quarter to the hour is it necessary to wait more than a minute or so for a reference tone. NIST disconnect telephone calls after 3 minutes. Here is the schedule for the WWV transmissions:

Tone	Minutes past the hour
600 Hz	1, 3, 5, 7, 11, 13, 17, 19, 21, 23, 25, 27, 31, 33, 35, 37, 39, 41, 53, 55, 57
500 Hz	4, 6, 8, 9, 10, 12, 14, 15, 16, 18, 20, 22, 24, 26, 28, 32, 34, 36, 38, 40, 42, 52, 54, 56, 58
440 Hz	2
None	0, 29, 30, 43, 44, 45, 46, 47, 48, 49, 50, 51, 59

Though the sound over the telephone is noisy the tones give a very clear reference frequency, as you will see if you analyse them. To use the tones, I turn the record level on the PC to maximum and hold the microphone close to the telephone earpiece. I have made successful calibrations with ease using both mobile and ordinary phones.

As a check on the NIST transmission, I checked the timing of the one-second pulses by counting samples over a 55 second recording. The broadcast frequency of 600 Hz was accurate to at least 5 significant figures. The NIST frequency standard gives a very precise frequency calibration for the price of a telephone call.