Technical note on collecting spectrographic data

There are many ways of generating a frequency response chart for a guitar. One common way is to drive the guitar mechanically (simulating the drive provided by a plucked string) using a variable frequency electromechanical vibrator, sweeping the frequency over the range of interest and then measuring the vibrations of the top either with an accelerometer attached to the top or as sound using a microphone. A second method, which is gaining in popularity, applies an impulse to the guitar (usually in the form of a blow from a small hammer) and then measures the vibrations of the top, again either with an accelerometer or as sound using a microphone and feeding the signal into a software spectrum analyser. The impulse input is a neat method of applying a drive signal containing all frequencies at one time and has been a method used in industry for many years. This is the method I prefer because it is both simple and quick, allowing a frequency response curve for a guitar to be generated in about 30 seconds.

Until quite recently, collecting spectrographic data demanded the type of equipment only found in serious research laboratories, requiring access to anechoic chambers, accelerometers, impedance heads, etc.. These days it can all be done to quite a degree of sophistication using a home computer with a sound card, a reasonable microphone, some freeware and a bit of cunning.

The major tool in the armoury is a software spectrum analyser. Some of these packages are available as freeware and come with software signal generators, timers, counters, frequency meters, oscilloscopes and a host of other features. At the time of writing the available packages include Audacity, SpectraScope, SpectraPLUS, SpectrumLab, Visual Analyser, Wavesurfer and WavePad amongst many others.

The setup pages vary markedly by application, but all but the most basic will allow you to set two important parameters; the number of samples to collect in the FFT[1] buffer and the frequency at which to collect the data points. These two parameters determine the bandwidth of data that you can analyse, the spectral line resolution and the time taken to fill the buffer.

A guitar builder is primarily interested in the frequency of vibration of the modes that can be influenced by changes to the woodwork and most of these modes occur below 1000Hz. The low frequency modes, in particular the coupled Helmholtz mode, can have a frequency as low as 80Hz, whilst the low E string tuned to E has a frequency of 82.4Hz. A semitone interval at these frequencies is only 5Hz and we are aiming to pitch resonances between semitones. Consequently we need a frequency resolution (spectral line resolution) of better than 1Hz. The spectral line resolution is determined by the sampling frequency divided by the FFT sample size. For example, at the usual compact disc sampling rate of 44.1kHz, an FFT sample size of 65536 will give a spectral line resolution of 0.67Hz. The bandwidth of signals that can be analysed is half the sampling rate (two data points are required to determine the frequency of a signal), so the bandwidth is 22kHz. As we are seldom interested in a bandwidth that wide in guitar construction we can choose a sampling rate of, say, 11.025kHz and a sample size of 16384 which will give the same spectral line resolution of 0.67Hz and a bandwidth of 5512Hz. The time taken to fill the FFT buffer with data is the sample size (buffer size) divided by the sampling rate, which for the latter case is 16384 divided by 11025Hz which equals 1.49 seconds. This would mean that we can tap a guitar once every 1.49 seconds and fill a buffer with the data recorded.

The maximum sample size supported by many spectrum analysers is 65536. If we required very high spectral line resolution we could choose a sampling rate of 8kHz, which would give a spectral line resolution of 0.122Hz, a bandwidth of 4kHz and a time to fill a buffer of 8.2 seconds.

---

[1] FFT = Fast Fourier Transform, a mathematical algorithm which breaks down a sequence of values (a buffer) into "buckets" of values of different frequency ranges. Each "bucket" has a "width" corresponding to the frequency range within the bucket, which is called the spectral line resolution, which determines the resolution to which the frequency data can be read. For example, a spectral line resolution of 6 Hz would mean that you could not resolve the peak of a resonance nominally at 100Hz to better than ±3Hz. The peak could be anywhere between 96Hz and 102Hz.

So which setup parameters should be chosen? I normally use a sample size of 16384 and a sampling frequency of 11.025kHz, which gives a bandwidth of ~5.5kHz, a spectral line resolution of 0.67Hz and a time to fill a buffer of 1.49 seconds.

A similar setup would be to choose a sampling frequency of 8kHz and the same buffer size which would give a spectral line resolution of 0.49Hz, a bandwidth of 4kHz and a time to fill a buffer of 2.0 seconds.

If you want to measure Q values at low frequencies, it is important to consider the setup parameters carefully if the results are to mean anything. The diagonal mode on a guitar half panel may have a frequency as low as 30Hz and let's suppose it has a real Q of 90 and we want to try to measure this.

From Equ. 1.4-25

$$Q = \frac{f_0}{Bandwidth}$$

where the bandwidth is measured at the -3dB points.

For a Q of 90 and centre frequency of 30Hz, the bandwidth at the -3dB points is 0.33Hz. To have any confidence in the Q values we would require at least 5 data points around the resonance, which means the spectral line resolution needs to be 0.33/5 = 0.066Hz. Irrespective of how you select the setup parameters (sample size and sample rate) a typical software spectrum analyser has a limiting spectral line resolution of ~0.12Hz, which means you cannot accurately measure this Q value. If we reverse this calculation we can find a limiting value for what we can measure with reasonable accuracy. With a spectral line resolution of 0.12Hz, we can measure a bandwidth of 5 x 0.12Hz = 0.6Hz. For a frequency of 50Hz (typical of the long grain marimba bar mode of a guitar half panel) we can reasonably measure a Q value of less than ~80. It is necessary to have a very clear understanding of these issues if your Q measurements are to mean anything at all.

When measuring guitar frequency responses we are primarily interested in measuring the frequency of the resonances below 1000Hz. I set the program up to have a frequency resolution of 0.67 Hz and collect the results of 10 taps on the guitar which can be automatically averaged and displayed. Most applications allow the display to be saved in various ways. I save the ones I want as text files which can subsequently be read into a spreadsheet program if further analysis or plotting in a different way is required.

I tap the guitar with a small rubber mallet made from a 19mm (¾") "bouncy ball" and a satay skewer and record the results using a microphone plugged into the sound card of a personal computer. The microphone needs to have a reasonably flat frequency response over the frequency range of interest, namely 50Hz to 5kHz. The response of many inexpensive microphones rolls off quite severely below ~80Hz. You need to be confident that you are primarily measuring the response of your guitar rather than the response of your microphone. Heavier, softer mallets tend to suppress higher frequencies (often a handy thing) whereas lighter, harder mallets elicit more high frequencies.

There are, of course, limitations to this process but for our purposes these haven't presented a serious problem.

The major limitations are:

1) No phase data is collected as only sound pressure measurements via a standard microphone are taken. This has presented no real obstacles to us.
2) The amplitude of plots from one instrument to another cannot be compared because there is no standard intensity for the exciting signal (the hammer tap). This presents no real problems either because we measure the sensitivity of a guitar (its responsiveness) by measuring its monopole mobility (see Section 1.7.2). Further, what we hear as loud instruments tend to be those instruments with high peaks in specific places on the frequency response that the human ear is sensitive to and these are readily apparent even with an uncontrolled excitation.

3) You need to have some idea where the plate's antinodal regions are for the modes of vibration you are interested in and then you need to tap on these antinodal regions in order to get a response from them.

It might appear that it is possible to significantly influence the relative height of various modal peaks, and that is true. For example, you will get very little excitation of a cross dipole if you only tap on the centreline of the instrument. Conversely, if you tap only on the antinodes of an anticipated mode, a quick identification of the mode is possible, because the height of the peak representing that mode will be exaggerated. However, if you tap systematically in a reasonably constant pattern you will get very reproducible results. One useful technique is to tap only on the bridge saddle immediately over each string. This ensures that you are exciting the top in a very similar way to how a string would excite the top and will give you results that are more consistent. If it is important to you to more accurately simulate the string's action, this can be done by tapping the saddle at each string's position in directions parallel to the string and perpendicular to the string, making a total of 12 taps if all strings are included.

To examine repeatability, we can take a series of 5 spectrographs, each an average of 10 taps, read these into a spreadsheet (and plot the average of the 5, if desired). The degree of repeatability can be seen from the plot below of the spread across the 5 individual spectrographs. If plots of the same guitar are taken on different occasions and compared, I find that they are readily identifiable as the same guitar. Whilst this might not satisfy an acoustics laboratory, it is plenty good enough for our needs as we are working primarily in the frequency domain and so the relative amplitude of the various plots is of lesser importance to us. We are mostly interested in checking the frequencies of the resonant peaks. It is far less time consuming obtaining data this way than setting up a frequency sweeping exciter with an accelerometer attached in an anechoic chamber with the guitar in a special holding device and suffering the vagaries of auto-correlation in the output signal due to the smooth nature of the frequency sweep. A plot such as the one below (5 sets of 10 taps read into a spreadsheet and plotted) takes 7-10 minutes to complete. A single set of taps will produce a plot in less than 30 seconds. The wider spread around the 100-120 Hz mark is due to changing the angle of the guitar relative to the microphone so that a little more or less sound hole output has been read. The guitar is held by a person as if it were being played (i.e. no fancy holding jigs).

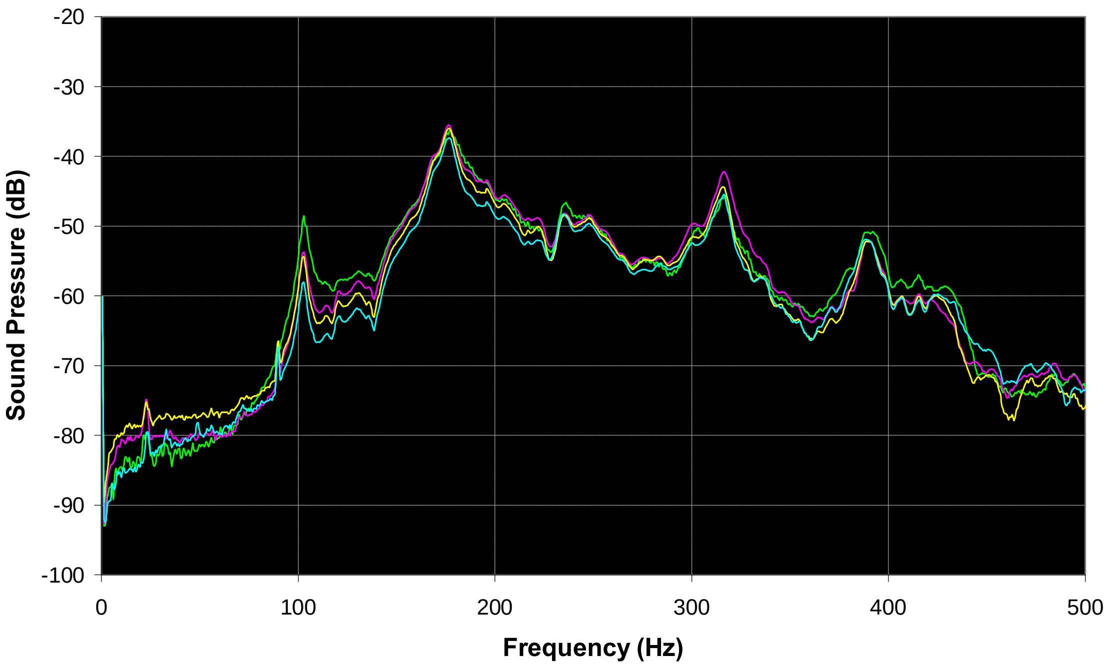

Repeatability of 5 tap test plots (OO style guitar)

Set up instructions for Visual Analyser

Most of the dedicated spectrum analysis packages that I currently know of are only available for Microsoft Windows. Audacity, which is a sound recording and wave editing package, has an embedded spectrum analyser, but it has limited functionality and an inability to zoom in on particular features of a spectrum. However, it does work in a Macintosh environment. WavePad is another Mac compatible package retailing for US $30-$50. Of the other packages, my preference is for an early version of Visual Analyser, v 9.0.6, which has a very functional spectrum analyser along with a programmable signal generator, frequency meter and an oscilloscope. Later versions have a number of annoying bugs and are less useable for guitar acoustics. Visual Analyser version 9.0.6 works on all Windows platforms from Windows XP through to Windows 10 with just one or two minor differences, the main one seeming to be that cross-hair cursors for the Spectrum window do not appear in some releases of Windows (but the normal mouse cursor always does).

All versions of Visual Analyser can be downloaded at no cost from the Sillanumsoft website: http://www.sillanumsoft.org/download.htm. Version 9.0.6 is at the bottom of the page. Install it on your computer. The following instructions regarding the set-up of VA refer specifically to Windows 7, although the procedure is very similar in all other versions of Windows.

Find the launch icon for Visual Analyser (on the Desk top or via the Start button). For the initial set-up, right click on the VA start icon and select "Run as administrator". You may see a window asking you to select input and output devices. Accept the defaults for now. They can be changed later.

You will then be presented with this screen:

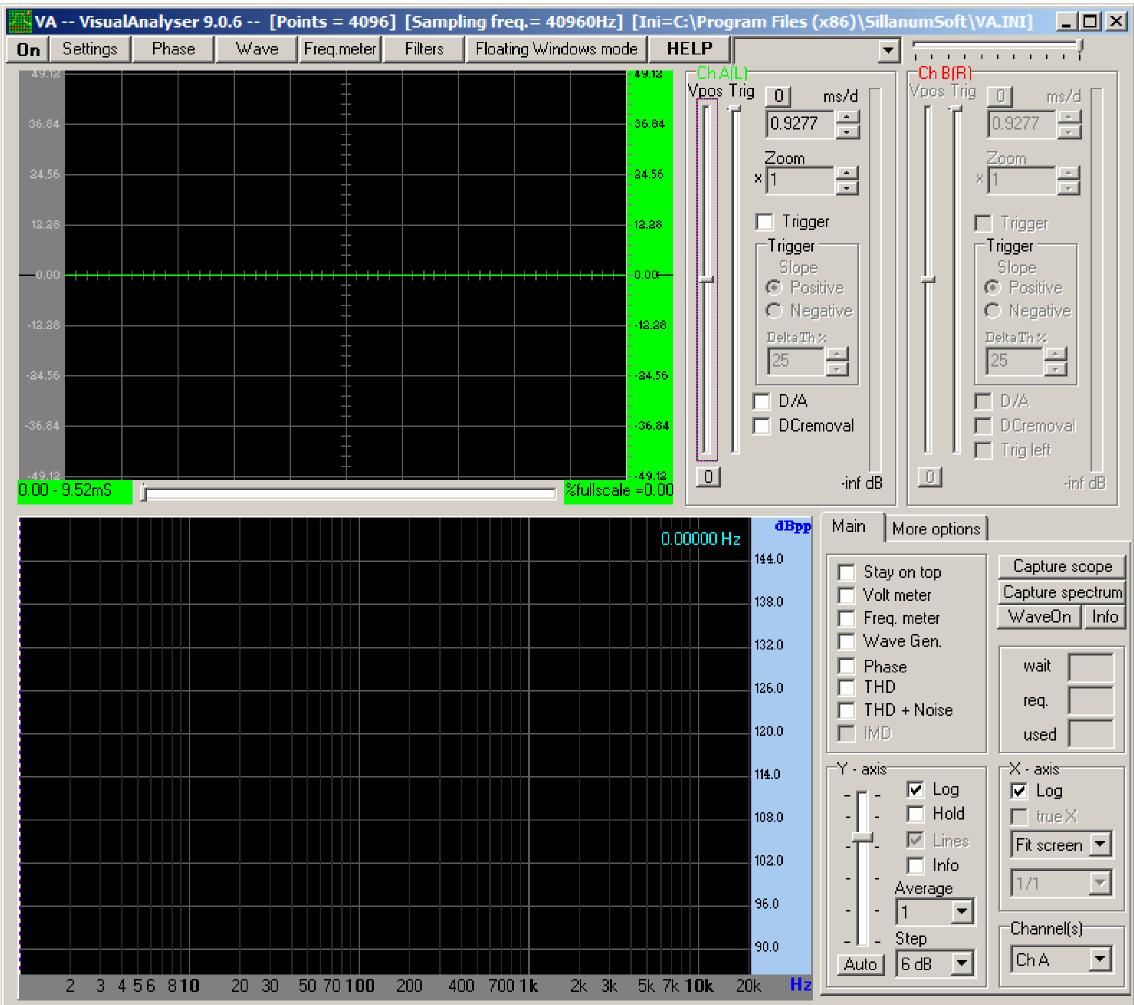

Select the 'Settings" tab (upper left):

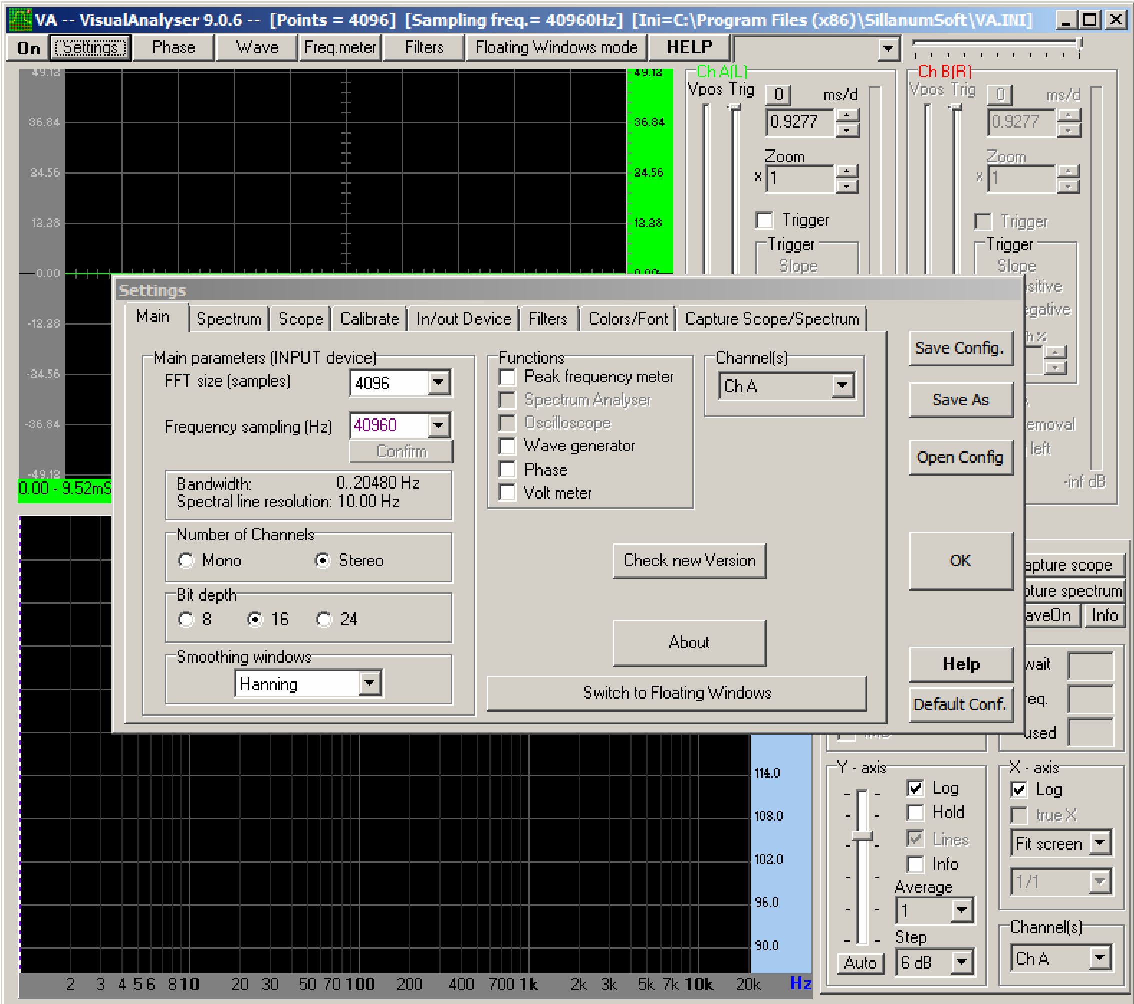

Under "FFT size (samples)", select 16384 from the drop down menu. Likewise, under "Frequency sampling" select 11025. Press "Confirm" and the bandwidth will change to 0..5512Hz and the spectral line resolution will change to 0.67Hz. Under "number of channels (centre left) select "Mono". Under "Smoothing window" (lower left) select "None" from the drop down menu. The "Main" tab should now look like this:

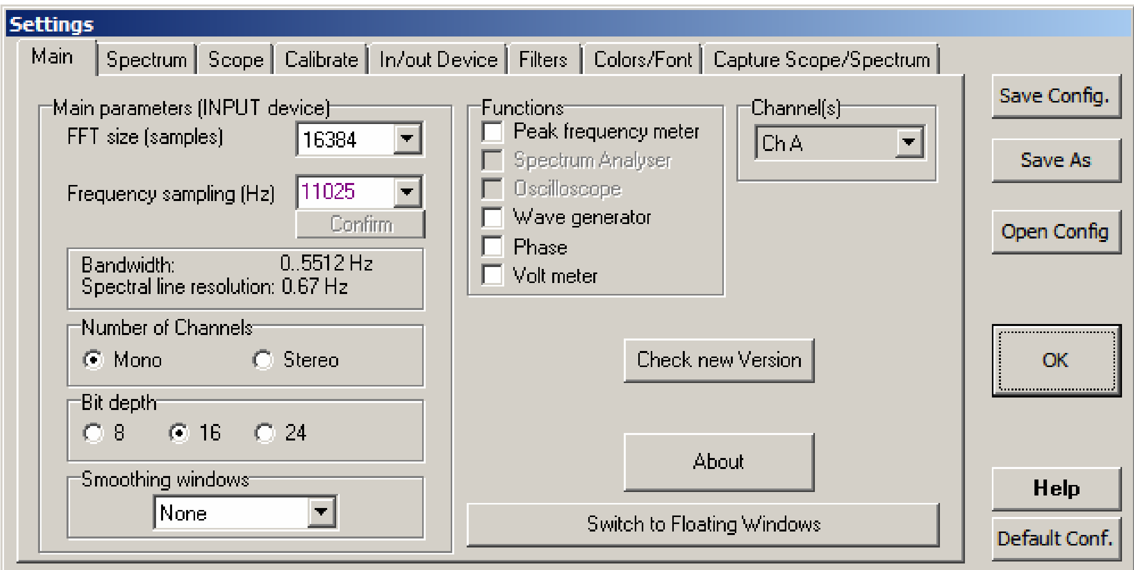

Now select the "Spectrum" tab. On this tab, uncheck "Log X axis", which will leave this tab looking like this:

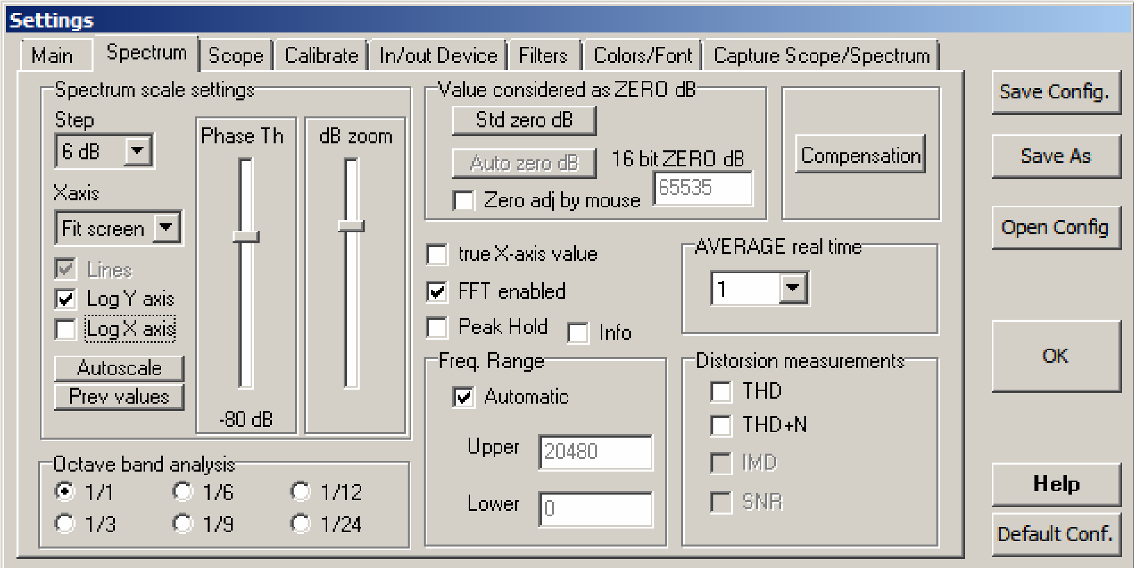

Now select the "In/out Device" tab. If no input and output devices are shown, press "Detect". Select the input and output devices you prefer. The output device only matters if you want to generate output from the "Wave" signal generator facility (for producing Chladni patterns for example) but we are mainly interested in the input device. My preference is to use an external microphone, which should be automatically detected if you plug it in. Select this option from the drop down menu, leaving the In/out tab looking like this:

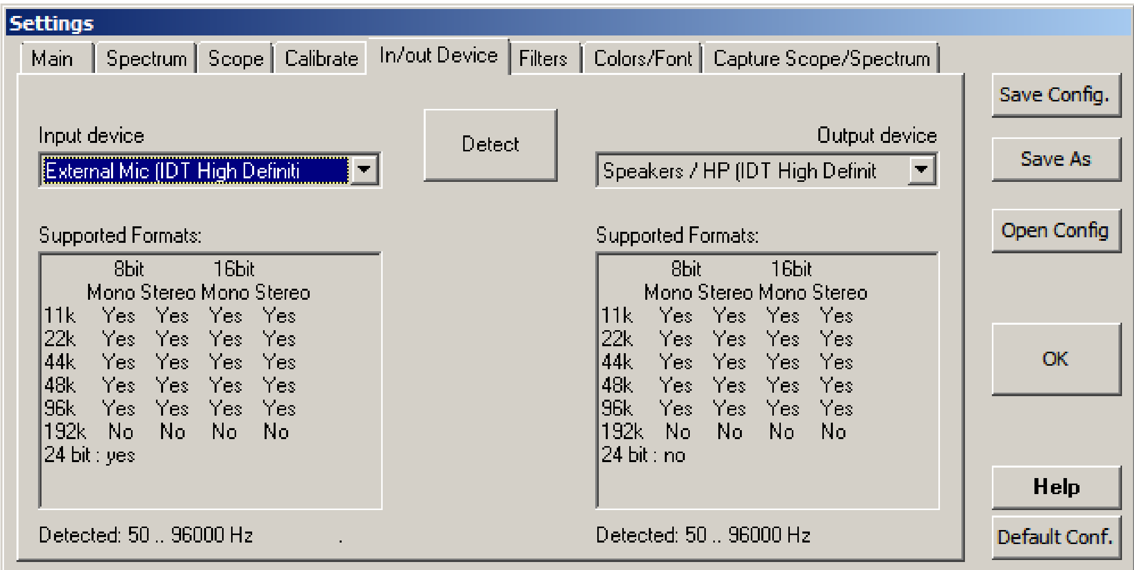

Now select the "Capture Scope/Spectrum" tab. Ignore the scope settings. In the "Capture spectrum" panel, select 10 from the drop down menu. This means that 10 buffers of information will be averaged to produce the spectrum plot, one buffer every 1.49 seconds, each buffer containing, ideally, only one tap.

The Capture Scope/Spectrum window should look like this:

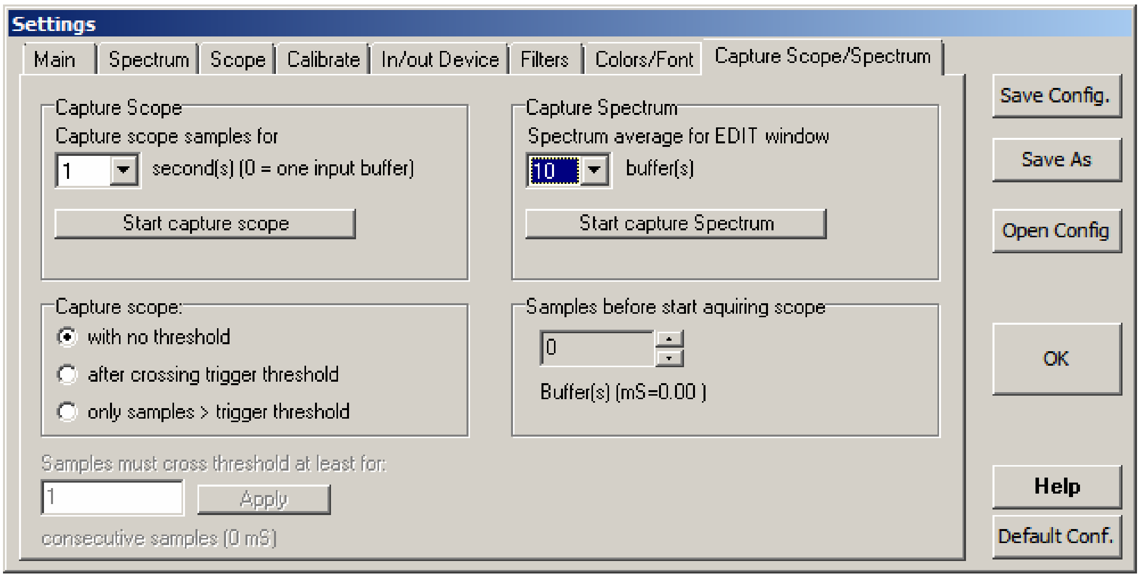

Press OK and the selected settings will load into VA when you press the On button in the main window. Check that this happens by pressing the On button and then checking the settings.

The settings now have to be saved, which is only allowed if you have Administrator rights, which is why the program on this first occasion was started in "Run as administrator" mode. In the Settings window, Press the "Save Config" button which will save a file called VA.ini in the Sillanumsoft Program Files (x86) folder. Other configurations can be saved (in Administrator mode) by using the "Save As" button, if you so wish. Next time VA is launched in normal user mode, you will be asked if you wish to use the saved settings. Press "Yes" and the configuration you saved will be loaded automatically.

With VA launched and the configuration loaded, press the "On" button. You should see a screen with a green spectrum plot in the lower panel looking like this:

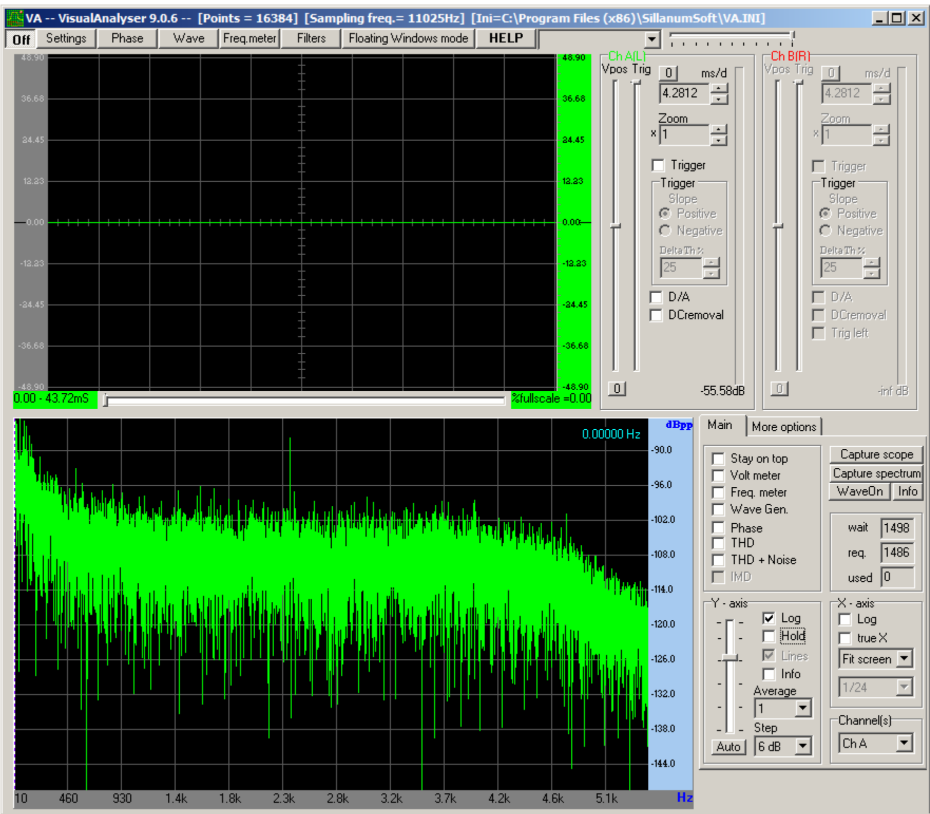

If you don't see the green display, press the "Auto" button on the "Y-axis" panel (lower right)

The Y axis slide bar alters the "gain" of the display (not the gain of the input signal), whilst clicking on and moving the blue bar on the right of the display moves the position of the Y -axis. For my equipment, the positions shown below work well.

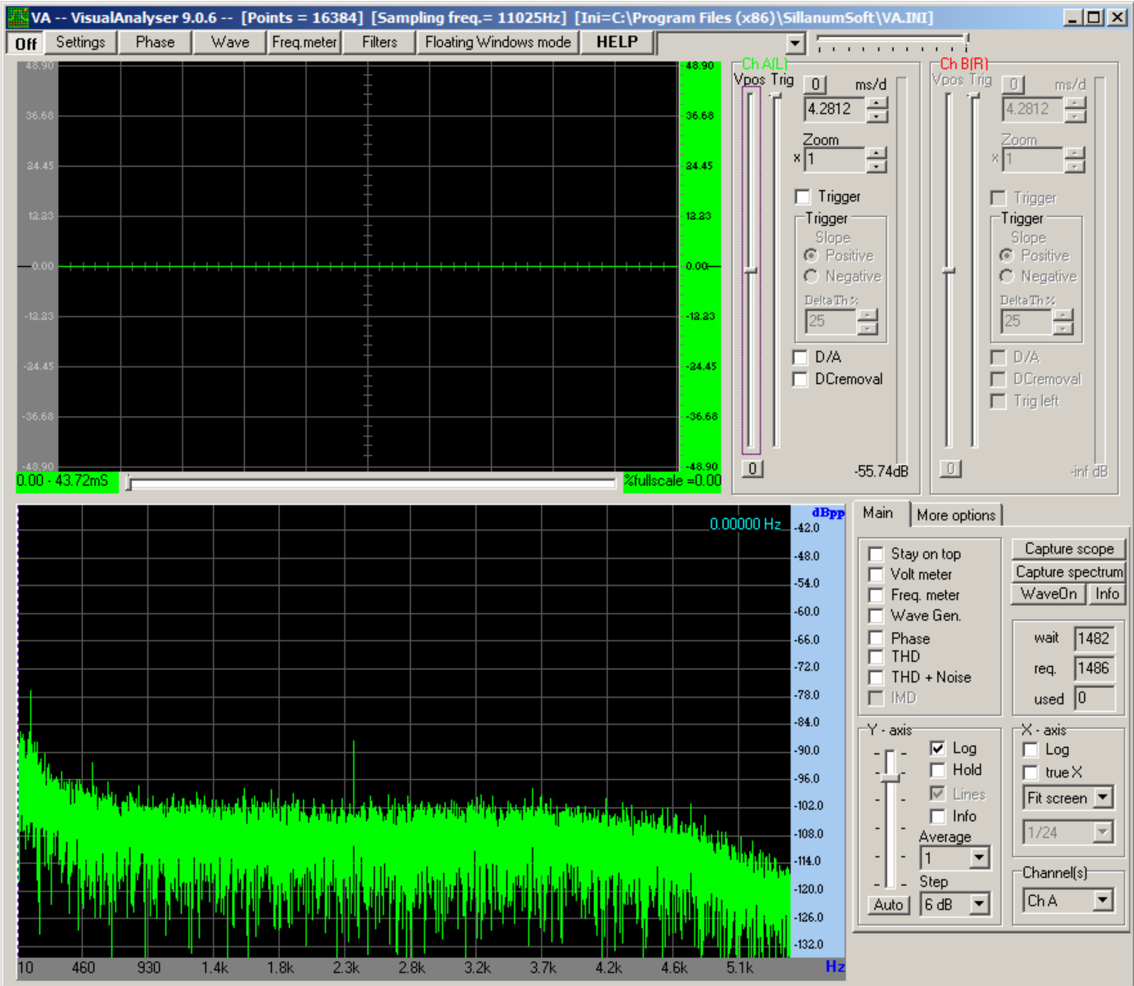

You will notice a drop down box and a horizontal slider, upper right of the main window. In Windows XP, these select the input device and the input signal level. In later versions of Windows these are disconnected. The input device is selected from the "Settings" menu, "In/out device" tab and the signal level is set via the computer's sound card controller. Access to the sound card controller tends to vary by make of computer but often it can be accessed via the speaker icon in the tool tray (typically bottom right of the computer's screen). Left clicking this icon gives access to recording devices (e.g. microphones) whose level can be set by selecting the device in use, clinking on properties and setting the level. Otherwise, levels can be set via the audio card controller accessed via "Control Panel". It is very important not to set the signal level too high, which will cause wildly spurious result in the spectrum analyser. Err on the low side rather than the high side.

Troubleshooting data collection

There are three categories of problem: no signal, distorted signal and noisy signal.

No signal: The most common cause of this problem is that the software controlling the sound card is not correctly set up. The recording device (microphone) has to be selected, enabled and the signal level set. The ways of doing this vary significantly across operating systems, computer brands and time (software versions) such that I am unable to give definitive instructions on how to do this. The "help" facility on the computer usually has sufficient information to get you started. Generally, better results are produced by external microphones rather than a laptop computer's built-in microphone, so be sure to select "external microphone" if presented with that option. Also make sure that the microphone is switched on, has live batteries (if required) and a good electrical connection to the appropriate input. If you can't see any signal input on your spectrum analyser, attempt a normal audio recording, which may be simpler to troubleshoot. Select "Programs" or "All Programs", select "Accessories", (possibly followed by "Entertainment") and finally select "Sound Recorder" which will open a new window. Click on the "Record" button and start talking into the microphone to record a short sound clip. Click on the "Stop" button when you are finished. Save the file and double click on it to play it back (or click the "Play" button if that option is available) to listen to the recording and the sound quality. Once you know you have a signal, you can then look elsewhere for why the spectrum software is not displaying it.

Distorted signal: A typical manifestation of this problem is that spurious peaks appear in the spectrum plot in places where you would not expect to find them, or that the spectrum produced looks nothing like a regular guitar spectrum. Frequently, this is due to saturating (overloading) the input and the spectrum analysis is subsequently performed on a clipped signal. The clipped signal will resemble a square wave, with lots of high odd harmonics which appear in the spectrum plot. The solution, of course, is to reduce the signal levels and better set them to some mean level which won't result in clipping. Monitoring an audio signal through earphones and adjusting the level to avoid distortion is one way of checking your levels.

Noisy Signal: The three major categories of noisy signals are electromagnetic interference, background acoustical noise and "grass" caused by tapping technique.

Electrical noise can be caused by poor grounding, poor shielding in cables and connectors, and electromagnetic interference due to mains hum (50 or 60Hz), fluorescent or LED lighting, dimmer switches (thyristors) and switching transients due to inductive loads (motors, household appliances) switching on and off. Good screening and the elimination of potential root causes is the path to proceed along. It is worth investing the time to recognise the difference between analogue ground problems familiar to many people (which usually sound like a powerful hum at mains frequency (50 or 60Hz, depending on where you live) and digital ground problems which sound like a scratchy hiss. They are resolved in much the same way, usually by reducing the number of paths to ground in the connected systems.

Acoustical noise problems are usually due to obvious background sounds, but also be wary of sounds transmitting through structure, for example the sound of your computer's fan transmitting to your microphone via the table they both sit upon. Elimination of background noise and a block of foam plastic beneath the microphone can reduce most problems to insignificance.

"Grass" on a frequency response chart is low amplitude, spikey waves on the plotted line. The problem is caused by capturing the sound of two taps in the same input buffer due to tapping too frequently, or by capturing a double bounce of the hammer on the guitar rather than a single tap. The buffer update frequency (which should be at ~1.5s intervals) can often be seen as either a flicker of the computer screen and/or an increment to a timer bar. Avoiding the bounce problem is a matter of tapping technique. After the initial tap, the hammer must be allowed to recoil. The force required to initiate the tap should be applied such that the hammerhead accelerates, but be removed before impact, because if it isn't, the head will be forced down again, giving the double bounce. The way to achieve this is to accelerate the hammer head, then keep a very relaxed grip on the hammer handle and let the hammerhead look after itself. It will naturally bounce just once, so you must allow it to do just that.

Using Visual Analyser

The reason I continue to use Visual Analyser is that from the very first tap I did I was able to recognise the shape of a guitar frequency response curve from the charts that were published in some of the key guitar research papers dating back to the early 80's. Here is such a frequency response curve compared to a much more recent curve from Visual Analyser:

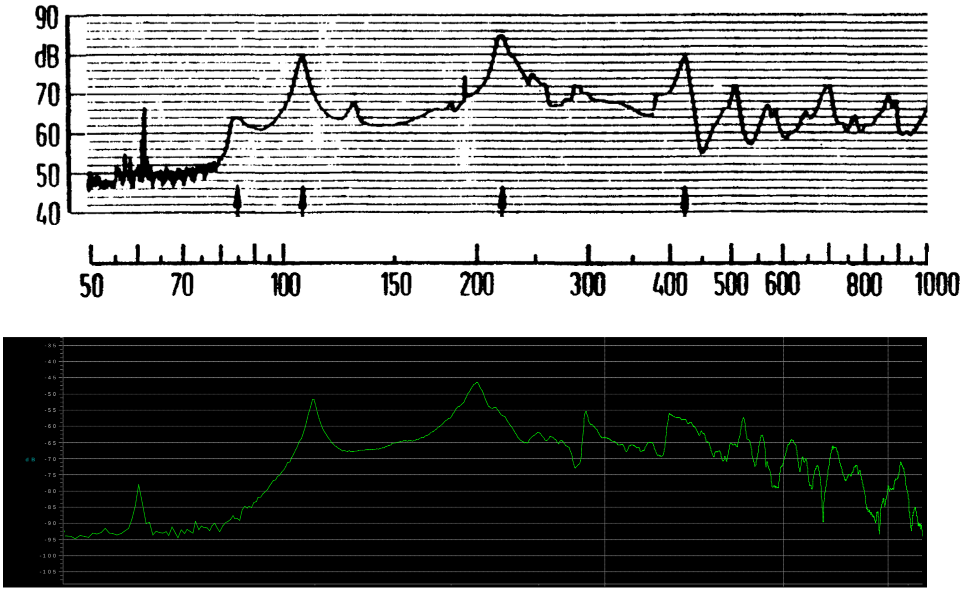

*Upper chart from "*Quality aspects of the guitar tone*", Jürgen Meyer, 1983; lower from Visual Analyser. The peaks at 60Hz are spurious electrical noise. Note the log scales here*

To my mind, that is what a typical frequency response of a guitar should look like, because it is very similar in its characteristics to those that the 80's researchers had spent years to produce and spent thousands of dollars on their early specialist equipment to produce them. Whilst the equipment we use today has changed significantly from that used in the 80's, the characteristic frequency response of a guitar hasn't changed, so should still look much the same. What the frequency response looks like depends not only on the guitar, but the mic you use, how you position it, how you tap the instrument, what you use to tap it with and how you set up the spectrum analyser.

To produce the frequency response curve from VA (above), the guitar was held as if it was to be played, ensuring that the back of the guitar was not damped against the player. The microphone was placed approximately 1 metre away, pointing midway between the bridge and the sound hole, on the guitar centreline. The mic used was a Shure PG57 dynamic mic, but many similar alternative mics could be used. Dynamic mics will plug directly into most computer sound cards without needing an interface, pre-amp or phantom power. To guard against spurious data capture, avoid resonant rooms and nearby guitars, which may sympathetically resonate.

The tapping hammer was made from a 19mm bouncy ball and a satay skewer. The flexibility of the satay skewer ensures that the hammer has a naturally frequency of just a few Hertz, well below any resonance on a guitar and therefore unlikely to produce spurious results.

With the strings damped (by hand, with foam etc.) the guitar is ready to be tapped with the hammer. With Visual Analyser "On", the "Capture Spectrum" button (centre right, main screen) is selected. This starts the collection of 10 buffers of information at 1.49s intervals. The objective is to record 10 clean, single taps, one tap per buffer, i.e. a tap every ~1.5s starting the tapping within 0.5s of hitting the capture button. After 10 taps, the Spectrum window will automatically open, displaying the averaged results from the 10 taps.

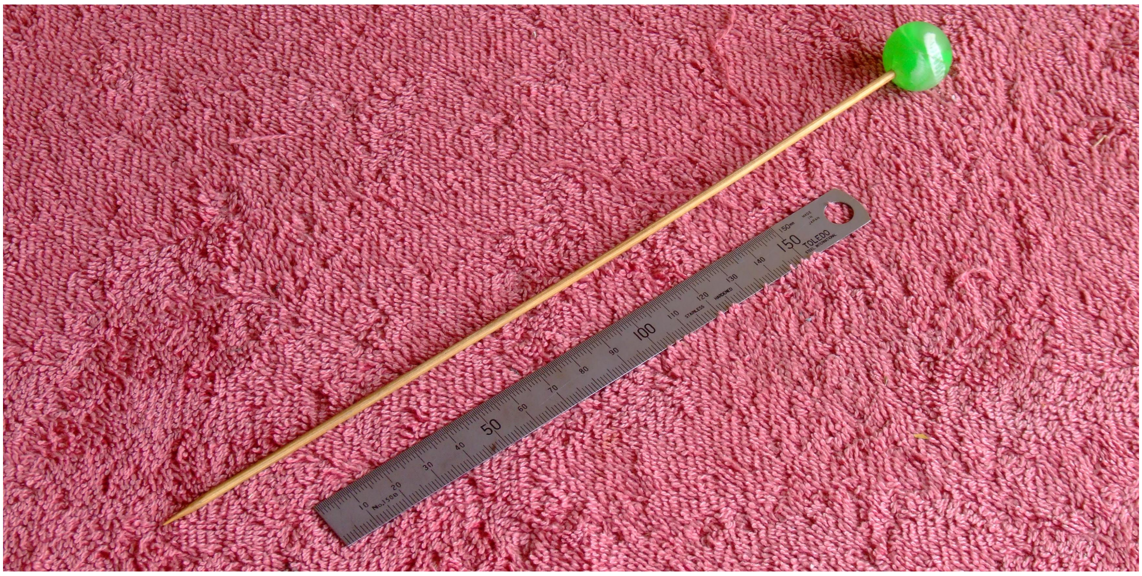

"Bouncy ball" and satay skewer tapping hammer

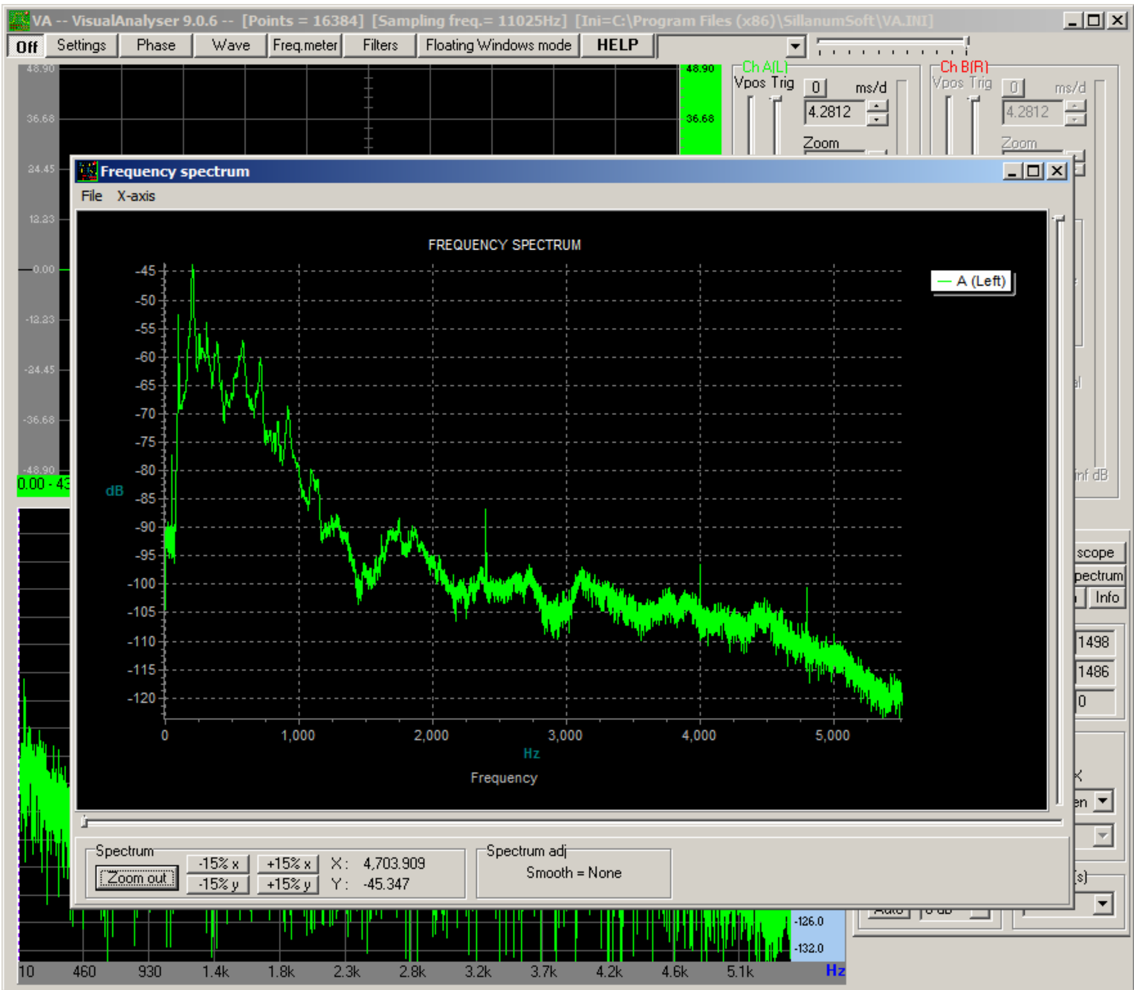

The spectrum capture window appears automatically after 10 taps

Right clicking on the spectrum chart allows the chart to be moved around the window. Normal clicking will allow a frame to be drawn around any area of interest and when the mouse button is released, that area is zoomed to fill the spectrum window. I typically zoom in on the 0 to 1000Hz range and carefully frame it in the spectrum window. The spectrum window can be enlarged to fill your computer screen in the normal Windows way. Holding the cursor over a peak on the chart allows its frequency to be read (X: at the bottom of the screen). To save the results, click on File/Save as text/ in the spectrum window. This will save a text file of the raw data (amplitude in dB vs. frequency) AND a *.wmf file (a screen shot) of whatever is in the spectrum window. That is why it is important to frame the spectrum chart correctly, so that you can build a library of comparable screen shots. The files can be saved in a folder of your choice.

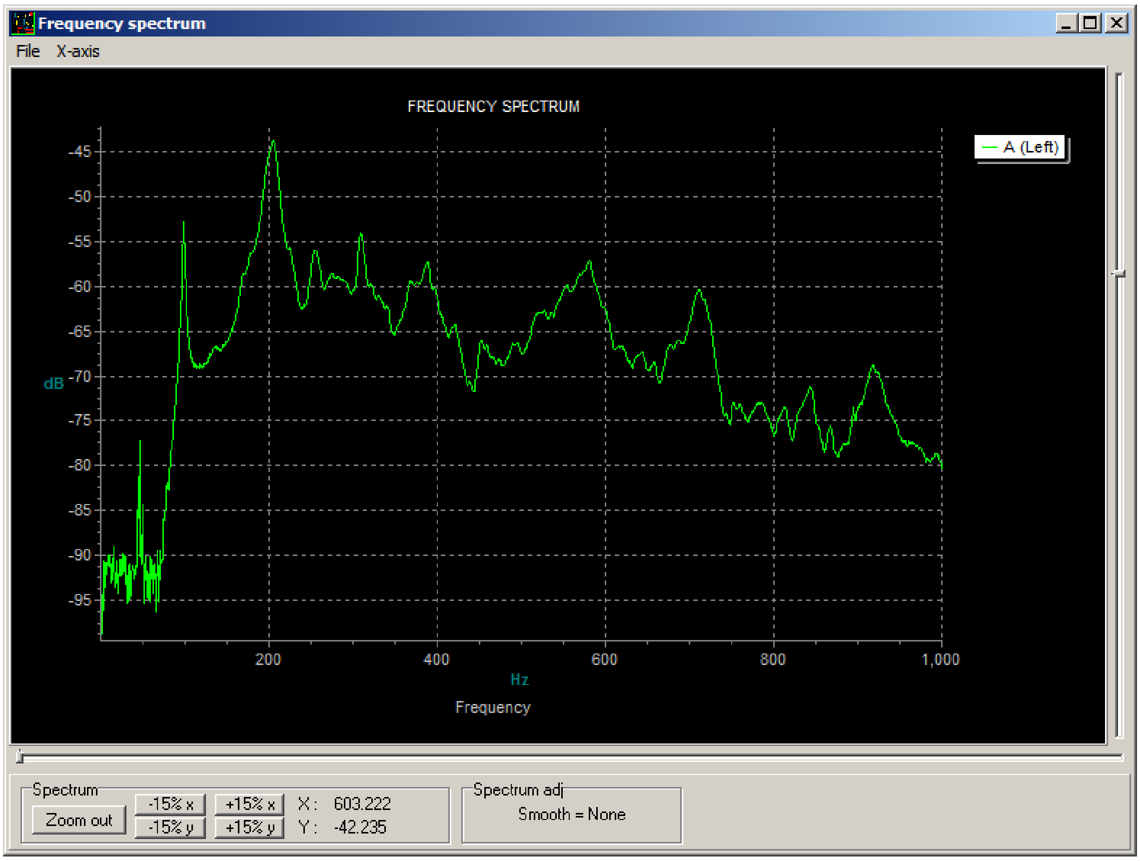

A selectively framed spectrum chart ready to be saved

To view saved data, locate the folder it was saved in and double click on the chosen file. A *.txt file will open in the default application as will a *.wmf file. A *.wmf file will open as a picture, so cannot be further zoomed or manipulated as if it was the initial spectrum window. If opening saved data using Visual Analyser is important to you (it's not to me) you will need to use one of the later versions of Visual Analyser that facilitate that (and also suffer the software bugs the later versions contain).

Data in the *.txt file can be imported into other applications in the typical ways for that application. For Excel, for example, the data can be imported and parsed using the import wizard.