FFTuner basic instructions. (for Chrome or Firefox, not IE)

A) Overview

Why FFTuner (piano)? There are many excellent piano-tuning applications available (both for free and for charge). This one is designed to also demonstrate a little music theory (and physics for the 'engineer' in all of us). It displays and helps interpret the full audio spectra produced when striking a piano key, including its harmonics. ***Since it does not lock in on the primary note struck (like many other tuners do), you can tune the interval between two keys by examining the spectral regions where their harmonics should overlap in frequency.*** You can also clearly see (and measure) the inharmonicity driving 'stretch' tuning (where the spacing of real-string harmonics is not constant, but increases with higher harmonics). The tuning sequence described here effectively incorporates your piano's specific inharmonicity directly, key by key, (and makes it visually clear why choices have to be made). You can also use a (very) simple calculated stretch if desired. Other tuner apps have pre-defined stretch curves available (or record keys and then algorithmically calculate their own). Some are much more sophisticated indeed!

B) Graphical User Interface elements defined

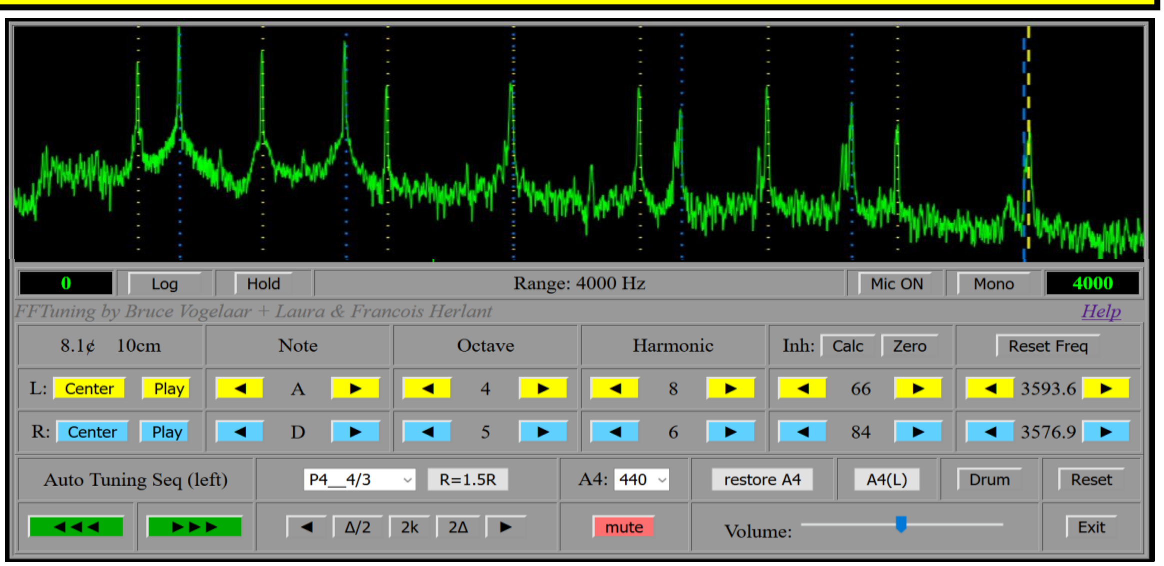

Complete interface.

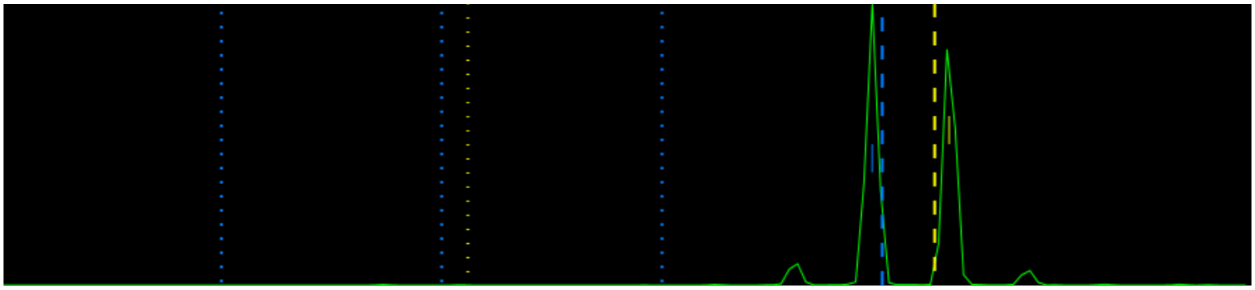

Green:	Fast Fourier Transform of microphone input (linear display in this case)
Yellow:	Left fundamental and harmonics (dotted lines) up to output frequency (dashed line).
Blue:	Right fundamental and harmonics (dotted lines) up to output frequency (dashed line).
Short Vertical Bar:	Yellow/blue bars are centroids of channels near the (left/right) output frequencies. They must appear ***within*** the peak of interest when being used for fine-tuning.

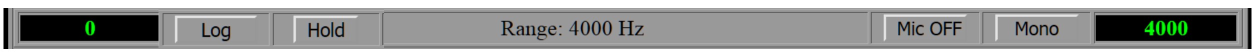

(low frequency)	Low frequency limit of displayed graph (Hz)
Log / Linear	Toggle log or linear amplitude. Auto-ranges to lowest and highest entries within the displayed frequency range.
Hold	If selected, holds the current amplitude range.
Range	**Normal:** range in Hz of displayed graph **Tuning:** During tuning displays four preset zoom levels for each note: a) basic window b) harmonic overlap +/- 1300 cents (just over two octaves) c) harmonic overlap +/- 120 cents (just over two half-steps) d) harmonic overlap +/- 50 cents (just over a half-step)
Mic ON / Mic OFF	Turns microphone on or off. (Switching to OFF will freeze the microphone trace on the graph, which can be useful for later analysis.)
Mono / Stereo	**Mono:** left and right channels are added, and then sent to speakers (good for hearing interference beats). **Stereo:** send left channel to left speaker, right to right speaker. (Note that some computer audio drivers can mix them afterwards as part of an audio 'enhancement' feature, which you can usually disable if you want.)
(high frequency)	high frequency limit of displayed graph (Hz)

Credit, and link to this Help document. If you like the philosophy, and have comments, corrections or suggestions, please email Bruce Vogelaar at vogelaar@vt.edu .

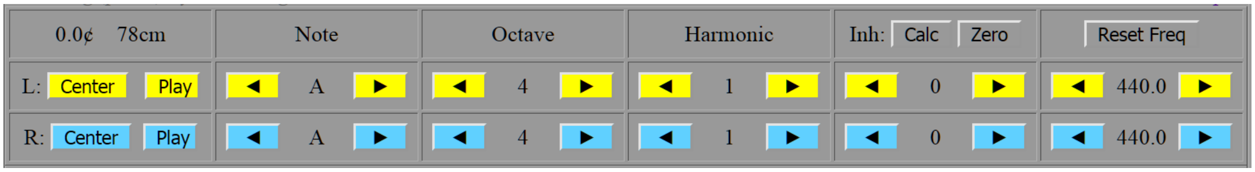

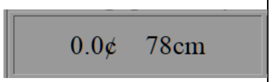	Displays cents between the left and right frequencies = 1200 $\log_2(f_L/f_R)$ Wavelength in air of left frequency sound. (Assumes v_s = 343 m/s.)
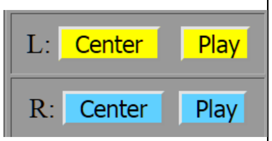	**Center:** centers the display graph on the left (or right) frequency. (Changes frequency range if necessary to force centering.) **Play:** Determines if the output sound includes the left (or right) frequency. (Also must be selected for keyboard up/down arrows to rapidly change frequency.)
Note	Changes the left (or right) note up or down one half-step.
Octave	Changes the left (or right) note up or down one octave.
Harmonic (n)	Gives the left (or right) output frequency as $f_n = nf_0\sqrt{1+\beta n^2}$. Displays fundamental and harmonics up until $(n-1)$ in the graph as dotted lines.
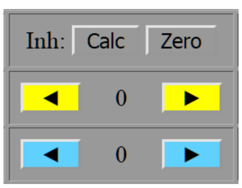 Inharmonicity (β)	Adjusts the inharmonicity parameter $\beta(\times 10^{-5})$ in steps of 3. Used to calculate left (or right) output frequency as $f_n = nf_0\sqrt{1+\beta n^2}$. **Calculate:** calculates the inharmonicity (see section on measuring inharmonicity). **Zero:** resets left (and right) inharmonicity to zero.
	**Reset Frequency:** resets both left and right frequencies to calculated ones. **Displays Left/Right Frequencies:** $f_n = nf_0\sqrt{1+\beta n^2}$ where f_0 is based on the A4 selection and assuming equal temperament. **Arrows:** One can change left or right frequencies manually using arrows. The rate of change per click is proportional to the frequency range of the graph. (Zoom in to change by 0.1 Hz per click, which is the minimum step.) You can also hold down the **up and down keyboard arrows** to rapidly 'repeat' click. [Only channels with 'Play' selected will change.]

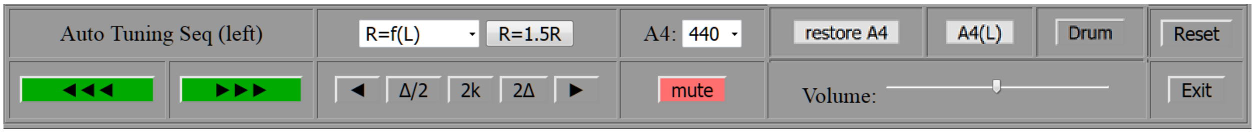

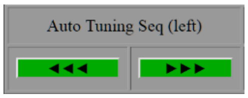	Tune the current note in the left channel (at four zoom levels), and then step up or down to the next note. The interplay between left and right channel is hard-coded to the octave being tuned. (See the 'Basic Tuning Sequence' section.) **Right Arrow:** Clicking the right arrow **(or right keyboard arrow)** will cycle the graph range as follows: basic window *<and up one half-step>*; ±1300¢; ±120¢; ±50¢; repeat **Left Arrow:** Clicking the left arrow **(or left keyboard arrow**) will cycle the graph range as follows: ±50¢ *<and down one half-step>*; ±120¢; ±1300¢; basic window; repeat
	**R=f(L):** This feature demonstrates just and equal temperament tuning. It sets the right channel a selected number of steps above the left channel (from a minor second, to full octave). It also changes the display to show the overlapping harmonics. The audible beat between the left and right channels then shows how far off 'equal' tuning is from 'just' tuning. **R=1.5R:** This feature demonstrates the 'circle of fifths' by raising the right channel by a factor of 3/2 (or a fifth). It divides the resulting frequency by two as needed to keep you in the same octave. Starting from the same note in left and right channels, clicking the button 12 times returns the right one close to the same note. ie $1.5^{12} \approx 2^7$ (off by 23.5¢ the 'Wolf fifth'). It is more interesting if you turn on the speakers.
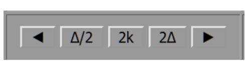	Controls for the graph display region. Shifts you left or right a quarter screen at a time; halves or doubles the range displayed; or sends you directly to 0 – 2kHz. Use in combination with 'center' key for left and right channels.
	- Selects the reference frequency for A4. (440 Hz is default) - You can manually align the left frequency with a known peak in your spectrum. Clicking the 'left freq derived A4', will calculate a temporary A4 such that your settings (note, octave, harmonic, inharmonicity) for the left channel will give you that peak frequency. A4 reports a red 'tmp' value. - Allows you to restore the reference A4 overwriting the temporary value.
	Toggles display of overtones for an ideal circular drum assuming the frequencies in the left and right channels are the fundamentals. (only visible if their respective 'play' button is on)
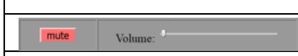	Turns the output speakers on or off and sets their volume.
	**Reset:** returns to default values, except for A4, inharmonicity and microphone status. **Exit:** tries to exit nicely, but on some browsers, pages cannot close windows they did not create.

C) Measuring inharmonicity of a string

Strings tensioned between two posts can vibrate like shown below. An integer number of half wavelengths, $\frac{\lambda}{2}$, must fit between the ends, giving the length, $L = n\frac{\lambda}{2}$. The frequency is then $f_n = \frac{v}{\lambda} = n\frac{v}{2L} = nf_0$ where v is the speed of the wave on the string (proportional to the square root of the tension) and $f_0 = \frac{v}{2L}$. N=1 is called the first harmonic (or fundamental) and $n = 2$ is the second harmonic, etc. For normal tuning, A4 has $f_0 = 440$ Hz.

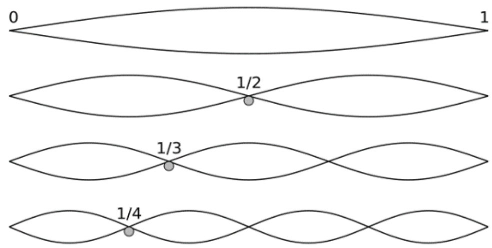

One would thus expect harmonics to be equally spaced in frequency. However, end effects (such as winding stopping - see picture, or stiff short strings with ends effectively clamped), become a larger fraction of the wavelength for the higher harmonics, making those harmonics appear at higher than expected frequencies – numerical solutions to Euler-Lagrange equations. (A tuning fork has its first overtone 6.25 times the fundamental frequency!) This feature has been traditionally parametrized by $f_n = nf_0\sqrt{1+\beta n^2}$ where β is called the inharmonicity. The non-linear behavior makes it mathematically impossible to match more than one harmonic, of say A3 and A4.

To measure the inharmonicity of a given note (even if it's not tuned yet) you just need to know the frequency of two different harmonics of this note.

$$f_1 = n_1 f_0\sqrt{1+\beta n_1^2}$$

$$f_2 = n_2 f_0\sqrt{1+\beta n_2^2}$$

which can be inverted to give:

$$\beta = \frac{1-a}{an_1^2 - n_2^2}; \text{where } a = \left(\frac{n_1 f_2}{n_2 f_1}\right)^2$$

and:

$$f_0 = \frac{f_1}{n_1\sqrt{1+\beta n_1^2}}$$

The process to find $f_1(n_1)$ and $f_2(n_2)$ for a common f_0 is to hit 'Reset', 'Zero' the inharmonicities, and then set the Left and Right channels to the note of interest. Observe the harmonics when you strike the key (you can quickly turn off the mic to hold the waveform if desired), noting that you may not see the fundamental if the frequency is low and your microphone doesn't pick it up (often a problem below 100 Hz). Some harmonics may even be absent, due to where the string was struck and the string's coupling to the soundboard. Adjust the harmonic number on the right channel so the blue dashed line matches the lowest harmonic that is clearly visible. Repeat for the left channel but going out to the 8th harmonic if possible. The observed harmonics may not align with the predicted ones due to tuning and inharmonicity. Next, zoom in on the left channel (hit 'Center' on the left channel, and then the '$\Delta/2$' key several times). Manually adjust the left frequency until it aligns precisely with its peak. Repeat for the right channel. Finally, hit the 'Calc' button in the inharmonicity group to calculate the inharmonicity. (The derived f_0 will also be used to provide a temporary A4 setting.) The measured harmonic lines should now line up with their predictions. (A reported inharmonicity of 9 is really $\beta = 9 \times 10^{-5}$.) When done examining, click 'restore A4'. One can repeat for any note desired (but only if you can see at least two harmonics). Typically one can measure values for A0, A1,…A6, probing the different string types in the piano. These values can be used to help guide your 'stretch' tuning.

The central octaves are usually tuned first (setting the 'temper'), and their higher harmonics are sharp compared to an equal tempered scale (due to inharmonicity), and thus the higher octaves need to be tuned a bit sharp to match. Likewise, the low octaves have higher harmonics which would sound sharp in the central octaves, and so the lower octaves need to be tuned a bit flat. This is the origin of the 'stretched' tuning. It's an imperfect game, where what 'sounds' good is hard to avoid as a criterion. The longer the strings, the less inharmonicity there typically is, making concert grands sound better, and spinets poorer. (Electric pianos could in principle have zero inharmonicity, but it turns out people like a little bit anyway.) In pictures, it looks like this:

Measuring the inharmonicity of A3 on my piano, finding $\beta = 0.00027$

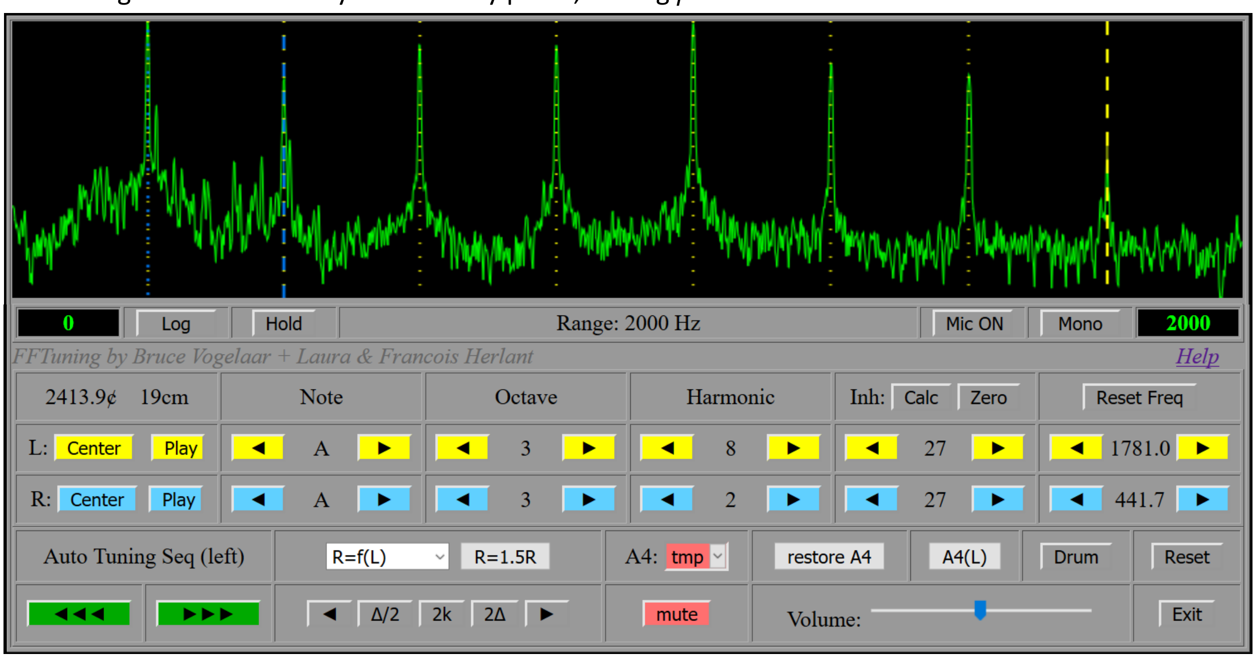

Tuning the A notes, from A0 to A7. Here are all the harmonics, with perfectly even spacing.

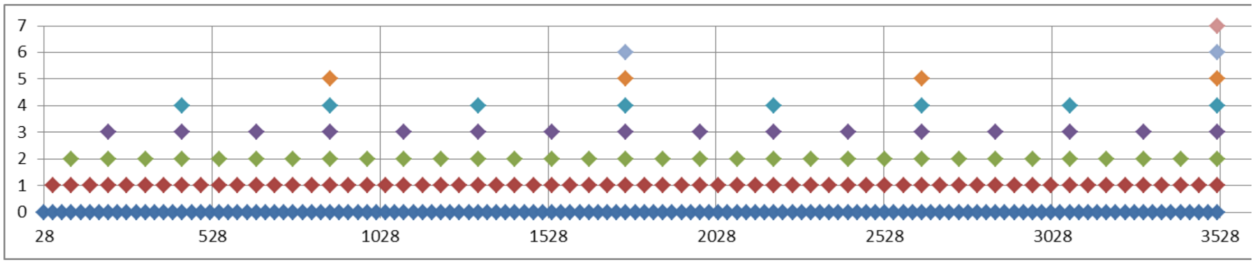

A0 to A7 on vertical-axis, frequency on horizontal-axis.

Here is taking the logarithm (base 2) of the above:

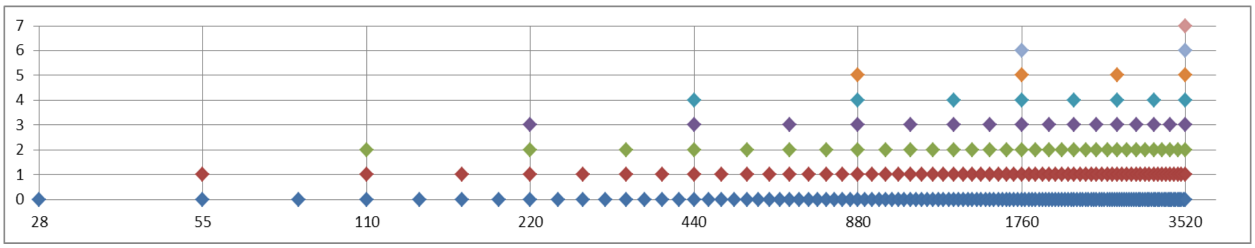

Here is keeping only the nearest octaves (to make it easier to see):

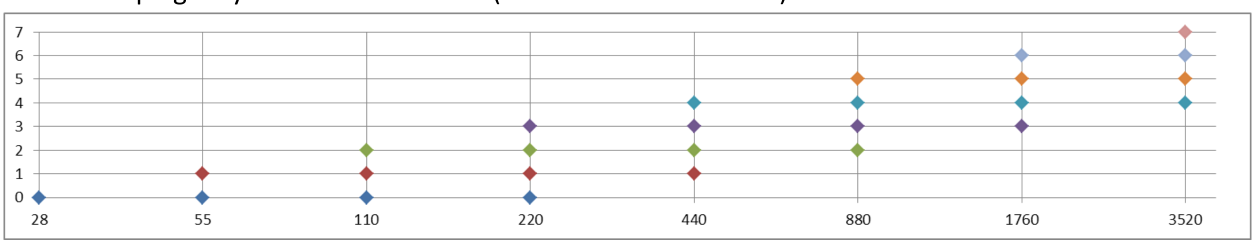

and now adding inharmonicity of $\beta = 0.001$:

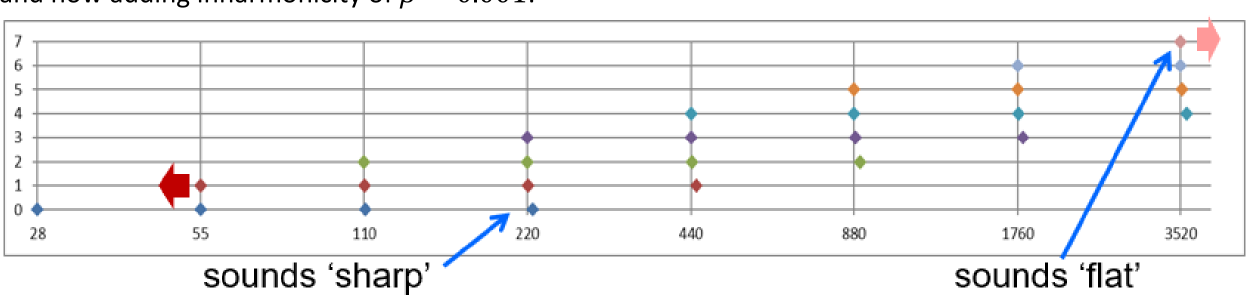

To align at 440 Hz, A0 and A2 fundamentals must be pushed to a lower frequency, while at 3520 Hz, A5, A6 and A7 fundamentals must be pushed to higher frequencies. Resulting in a ‘stretched’ tuning:

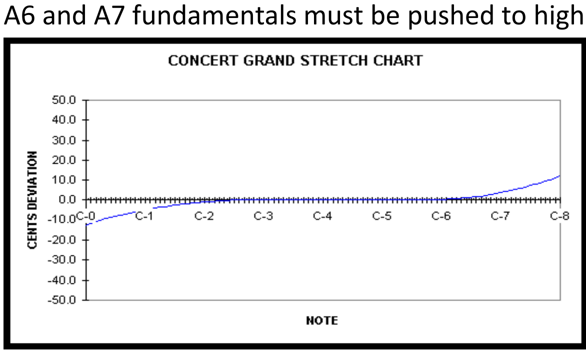

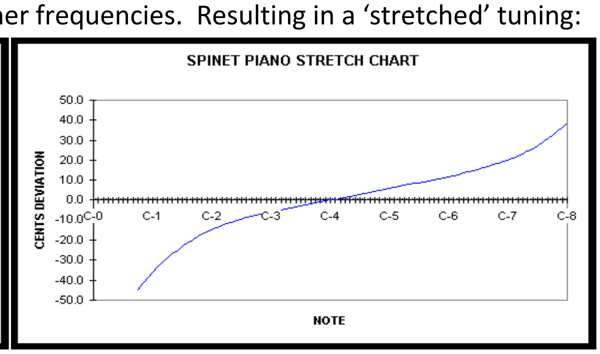

D) Demonstration of just and equal temperament scales

Notes that sound harmonious when played together have overlapping harmonics. Octaves, an overlap of the second harmonic of the lower note with the fundamental of upper one, are universally pleasing. To the Western ear, the 5th is next most important, leading to the "circle of fifth", where increasing the frequency by a factor of 3/2 twelve times, and shifting down by 7 octaves (ie: dividing by 2^7) returns you to within 23.5¢ of where you started (called the Wolf fifth). This hints at maybe a 12-note octave…

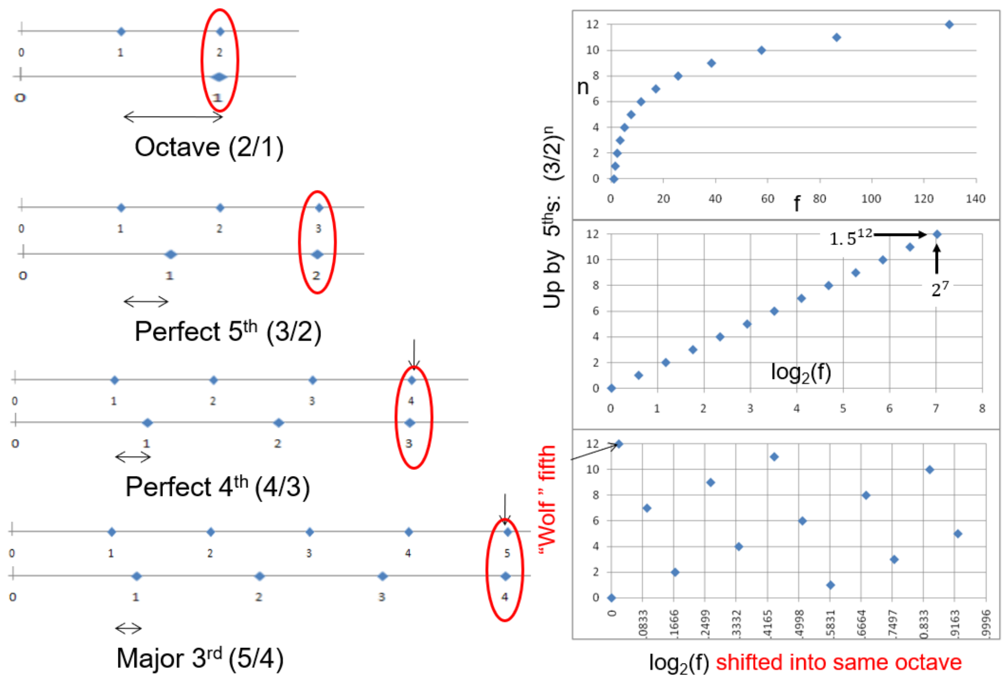

To demonstrate this, 'Reset' FFTuner, turn on the audio for both the left and right channels, and then click the 'R=1.5R' button twelve times. You will hear the tones indicated in the lower right graph above versus the starting tone, and wind up hearing the beating Wolf fifth at the end.

Here might normally follow a history of different scales from pentatonic to heptatonic, but they are ***completely*** out of our field of knowledge. We do note that all 'just' scales suffer from having a preferred starting note, rendering a very different sound when transposing a step up or a step down. To avoid this, the equal temperament was developed, where the frequency ratio between all half-steps were equal. Twelve notes per octave means that this ratio is $\sqrt[12]{2}$ so that after twelve half-steps, you are an octave (factor of two) higher.

One might ask if some other number of notes could have been chosen. Consider the graph of $\log_2$ of ideal (or just) ratios:

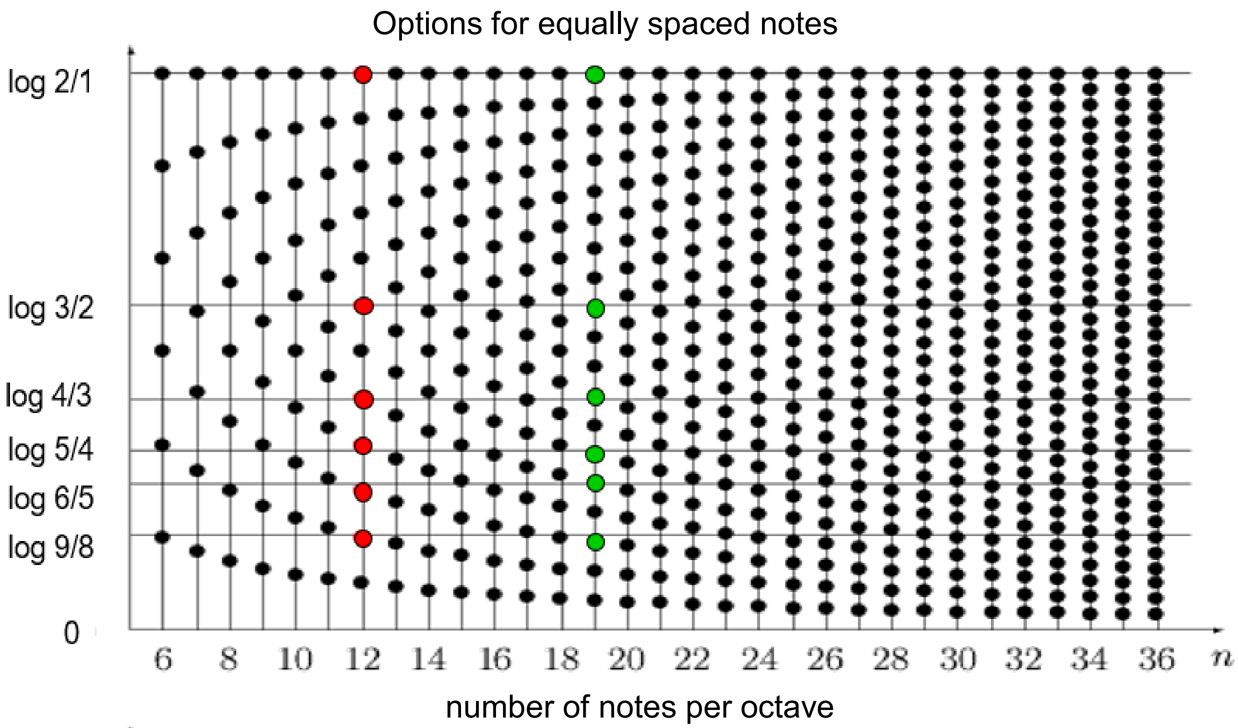

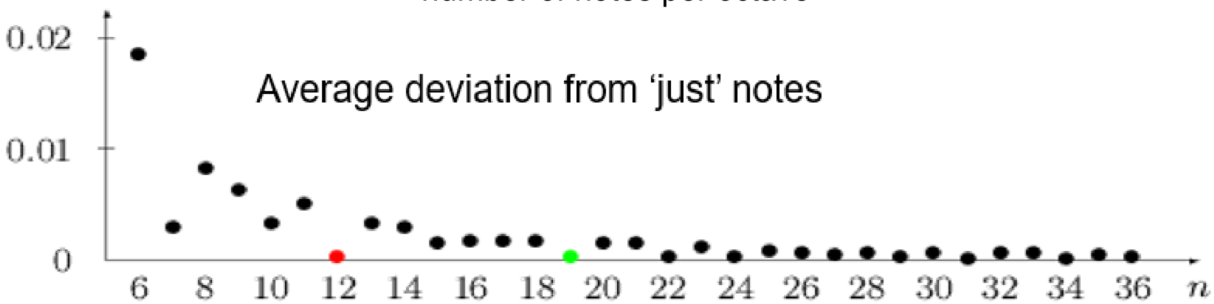

The western world has chosen 12 equal tempered steps; but there could have been 19 just as well…

True Equal Temperament Frequencies (Hz) for standard 88 key piano.

	0	1	2	3	4	5	6	7	8
C		32.70	65.41	130.81	261.63	523.25	1046.50	2093.00	4186.01
C#		34.65	69.30	138.59	277.18	554.37	1108.73	2217.46	
D		36.71	73.42	146.83	293.66	587.33	1174.66	2349.32	
D#		38.89	77.78	155.56	311.13	622.25	1244.51	2489.02	
E		41.20	82.41	164.81	329.63	659.26	1318.51	2637.02	
F		43.65	87.31	174.61	349.23	698.46	1396.91	2793.83	
F#		46.25	92.50	185.00	369.99	739.99	1479.98	2959.96	
G		49.00	98.00	196.00	392.00	783.99	1567.98	3135.96	
G#		51.91	103.83	207.65	415.30	830.61	1661.22	3322.44	
A	27.50	55.00	110.00	220.00	440.00	880.00	1760.00	3520.00	
A#	29.14	58.27	116.54	233.08	466.16	932.33	1864.66	3729.31	
B	30.87	61.74	123.47	246.94	493.88	987.77	1975.53	3951.07	

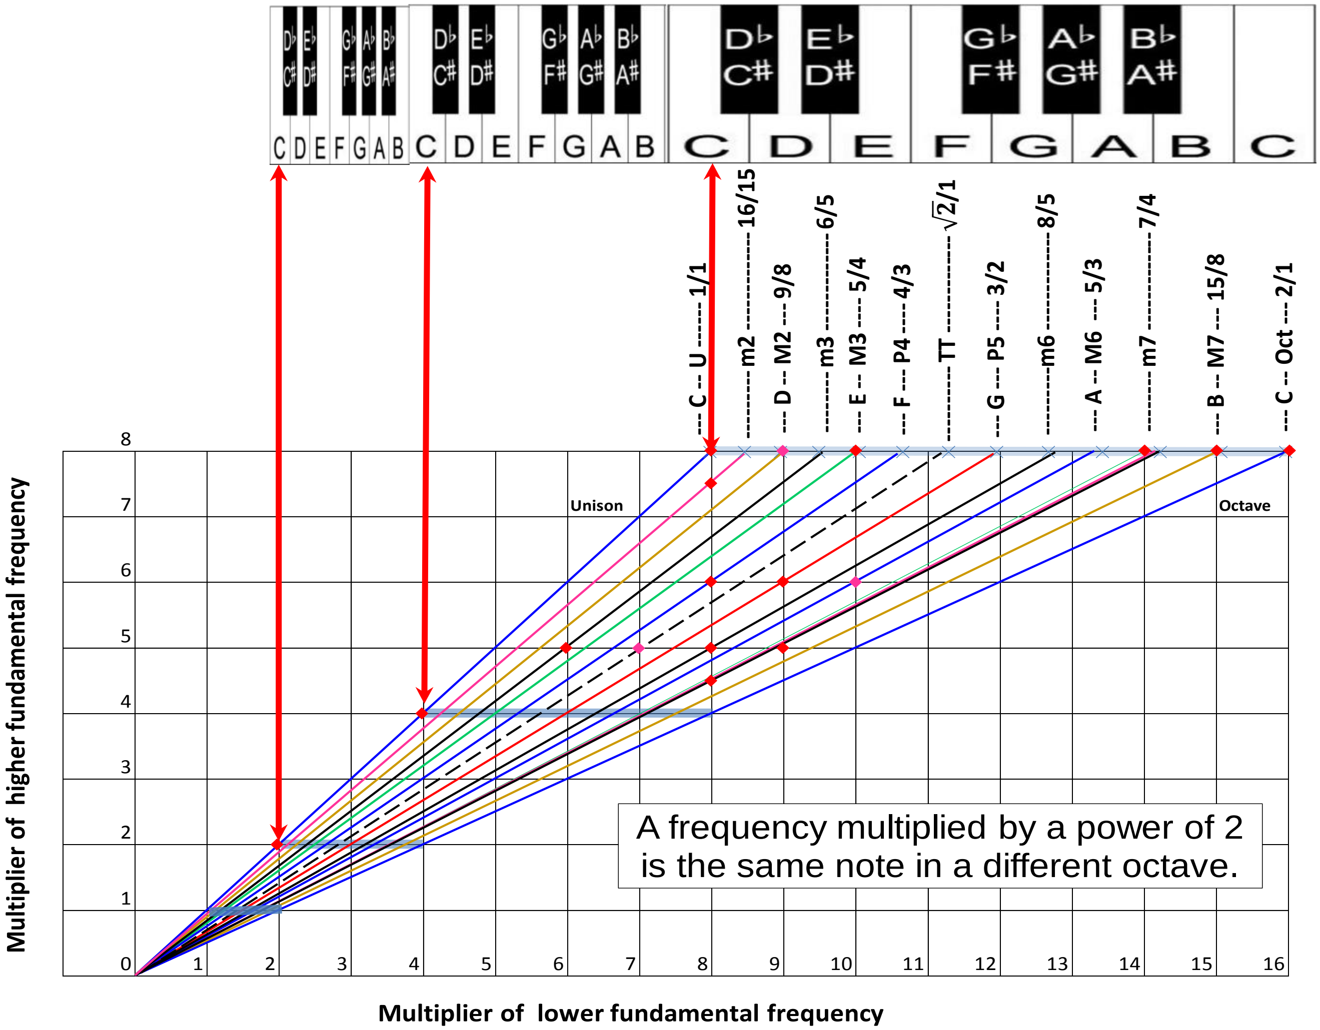

The 'R=f(L)' feature demonstrates just versus equal temperament tuning. It sets the right channel a selected number of steps above the left channel (from a minor second, to full octave). It also changes the display to show the overlapping harmonics. The audible beat between the left and right channels then shows how far off 'equal' tuning is from 'just' tuning. Here is an example of a perfect fourth (4/3):

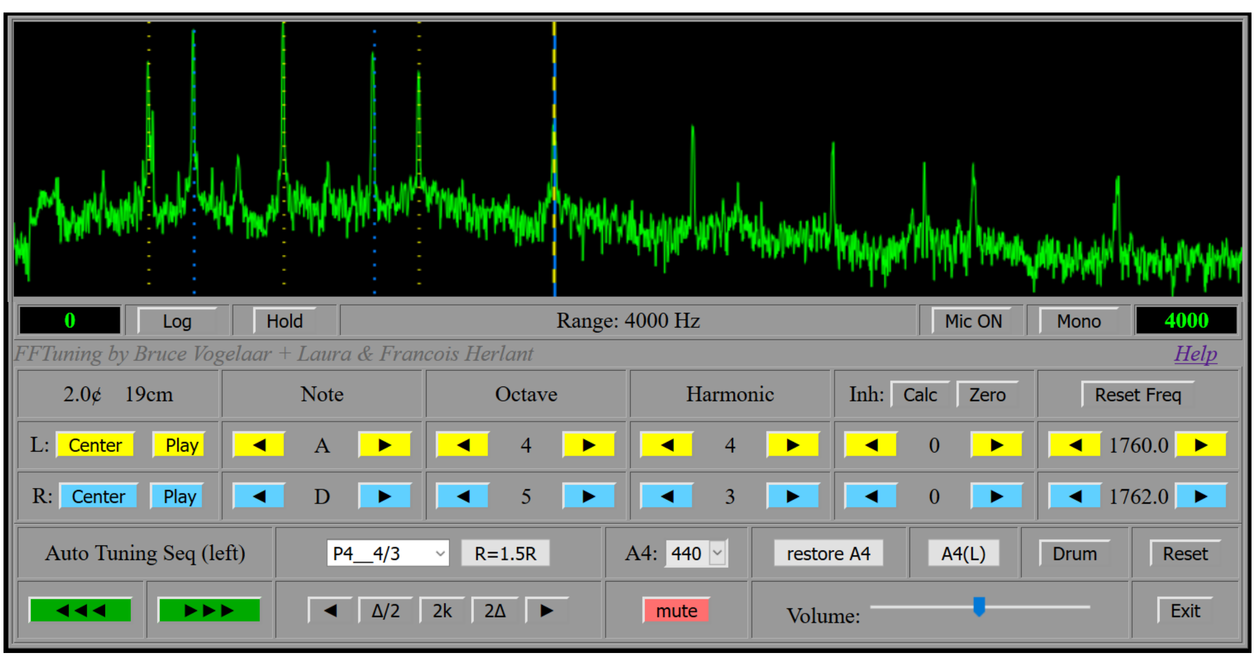

E) Basic tuning sequence

Ultimately, this needs to be as fast and painless as possible. You are encouraged to measure the inharmonicity of several reference keys first and enter them based on the region you're tuning (base, tenor, treble), but they only provide guidelines, and actual tuning requires looking at the overlap of harmonics. Here is the general approach: tune octaves 3, 4, and 5 to their exact frequencies (setting the 'temper'). This eliminates 'stretch' for this region as a reasonable first estimate. For octaves 0, 1, and 2, tune their 8^{th}, 4^{th}, and 2^{nd} harmonics respectively to match the same note in octave 3. For octaves 6 and 7, tune their fundamentals to match the 2^{nd} and 4^{th} harmonics respectively from octave 5.

The software is set up to pre-populate settings to simplify this process. Start with a note in octave 4, then use the tuning arrows (or keyboard arrows – which provide auto-repeat):

Right Arrow: Clicking the right arrow **(or right keyboard arrow)** will cycle the graph range as follows: basic window <*and up one half-step*>; harmonic overlap ±1300¢; ±120¢; ±50¢; repeat

Left Arrow: Clicking the left arrow **(or left keyboard arrow)** will cycle the graph range as follows: harmonic overlap ±50¢ <*and down one half-step*>; ±120¢; ±1300¢; basic window; repeat

For octaves 3, 4, and 5, the left and right notes will be the same, and the zoom region includes their fundamental frequency. Tune each string so that the peak is centered on the reference frequency. Check unisons as desired (they should be pure if you centered the peaks well).

For octaves 0, 1, and 2, you will tune a specific harmonic when plucking a string of the left note so it matches the fundamental when playing the same note in octave 3 (the right channel). The zoom region is on this overlap. This builds-in the actual inharmonicity. (note: entering the inharmonicity value on the left channel alone lets you visualize the stretch in the lower octaves)

For octaves 6 and 7, you will tune by plucking a string of the left note so it matches the correct harmonic when playing the same note in octave 5 (the right channel). The zoom region is on this overlap. This again builds-in the actual inharmonicity. (note: entering the inharmonicity values on the right channel alone - say from A5 - lets you visualize the stretch in the upper octaves)

Simple, but you will encounter problems. You may decide that notes in octave 7 should match 2^{nd} harmonics from octave 6, instead of 4^{th} harmonics of octave 5. But they don't. In fact, due to inharmonicity you can't have it both ways. You may also find that tuning the notes in octave 1 by matching their 4^{th} harmonics to fundamentals in octave 3, results in flawed intervals in octave 1. That's because the inharmonicity of the strings might be different. You might find that the 8^{th} harmonic of octave 0 is highly suppressed, but the 7^{th} is very clear (in which case, simply use it instead). You might even find that A3 beats with A4, since you tuned them to exact frequencies, but there is actually some inharmonicity… For the upper frequencies, you'll find all sorts of odd interferences of one string with another, and even splitting of frequencies within a single string (probably due to orthogonal modes of oscillation, with slightly different boundary conditions).

Yikes! You simply can't do a perfect job without compromise unless the inharmonicities of all the strings are zero. So, hearing what sounds good is still kind of important, but FFTuner will get you very close, and allow you to see exactly what's going on/wrong so you don't walk too far away from ideal.

You can of course define a stretch and tune to it. If you want that coded in as an option, let me know.

(PS: the unisons are easy to set even one string at a time, given the resolution of the FFT. Unfortunately, the current Javascript API does not let you change the FFT parameters, and thus resolution.)

An "aural" tuner faces similar challenges. When setting the initial temperament of the central octave, they tune for pure harmonic overlaps, and then de-tune them slightly to produce the specified beats (converting a 'just' tuning into an 'equal' temperament tuning). Here is the idea:

Equal temperament beatings (all figures in Hz)												
261.626	277.183	293.665	311.127	329.628	349.228	369.994	391.995	415.305	440.000	466.164	493.883	523.251
0.00000			14.1185	20.7648	1.18243		1.77165	16.4810	23.7444			C
		13.3261	19.5994	1.11607		1.67221	15.5560	22.4117			B	
	12.5781	18.4993	1.05343		1.57836	14.6829	21.1538			B♭		
11.8722	17.4610	.994304		1.48977	13.8588	19.9665			A			
16.4810	.938498		1.40616	13.0810	18.8459			A♭				
.885824		1.32724	12.3468	17.7882			G				Fundamental	
	[illegible]	11.6539	16.7898			F♯					Octave	
1.18243	10.9998	[illegible]			F						Major sixth	
10.3824	14.9580			E							Minor sixth	
14.1185			E♭								Perfect fifth	
		D									Perfect fourth	
	C♯										Major third	
C											Minor third	

From C, set G above it such that an octave and a fifth above the C you hear a 0.89 Hz 'beating'

Interval	Approximate ratio	Beating above the lower pitch	Tempering
Unison	1:1	Unison	Exact
Octave	2:1	Octave	Exact
Major sixth	5:3	Two octaves and major third	Wide
Minor sixth	8:5	Three octaves	Narrow
Perfect fifth	3:2	Octave and fifth	Slightly narrow
Perfect fourth	4:3	Two octaves	Slightly wide
Major third	5:4	Two octaves and major third	Wide
Minor third	6:5	Two octaves and fifth	Narrow

These beat frequencies are for the central octave.

This can be demonstrated simply with FFTuner. Hit 'Reset', set the left channel to C4, and then use the 'R=f(L)' key to select a Perfect Fifth. The graph will load the third harmonic of C4 into the left channel and the second harmonic of G4 into the right channel. The difference is 2.0 cents. Turn on the audio output, and you will hear the 0.9 Hz beating an octave and a fifth above C4. Try hitting these two keys on the piano, and you should be able to hear this beating. (It was trying to learn how to hear these beatings that prompted writing FFTuner in the first place.)

What puzzles me, is that inharmonicity can dominate these precise beatings in any case, and you still have the problems of ‘stretch’ to deal with and lots of compromises.

F) Tuning hardware and techniques

You will need a tuning hammer at a minimum. (A specialized tapered-square wrench for turning tuning pins on a piano. Do not try your crescent or box wrench.) To sound one string at a time (for notes with more than one), there are several options. First (and simplest for the lower octaves) is to simply ‘pluck’ each string by hand using your fingernail (while holding the sustain pedle down). For the higher octaves, use a guitar pick attached to a long rod by some means (I used a pick between two chop-sticks rubber-banded together and then taped.) A second method is to dampen the unwanted strings and simply strike the key as normal, holding it down for the sustain. There are proper tools for doing this shown below (rubber wedge, felt strip, spreading-tweezers), but you can also simply use two heavy-duty cable-ties as shown below muting the outer strings of a set of three.

The felt strip lets you mute lots of strings at once (by wedging it between the strings of different notes) so you can use single strings to tune all the keys first. Then, at the end, bring in the unisons.

We suggest you first open the top of the piano, and mark the tuning pins for most notes using a marker, pencil, tape strip, or some other technique. This will help avoid turning the wrong pin. See example below (and above).

You will likely want to use a laptop, and an external microphone which you can locate near the strings you're tuning. Laptops have arrow keys which makes changing zoom and notes far easier than touching little screen buttons on your phone. (Don't forget that you can make the webpage full-screen, and use <ctrl><plus> or <ctrl><minus> to make the application window larger or smaller.)

Microphones on cell phones and even laptops typically have very poor low-frequency response. This may not be fatal, since the lowest octaves are tuned by matching their harmonics to higher octaves (so you may never need to actually see the fundamental). To get the low frequencies reliably, you may want to obtain a bass (or kick-drum) microphone with phantom (48V) power.

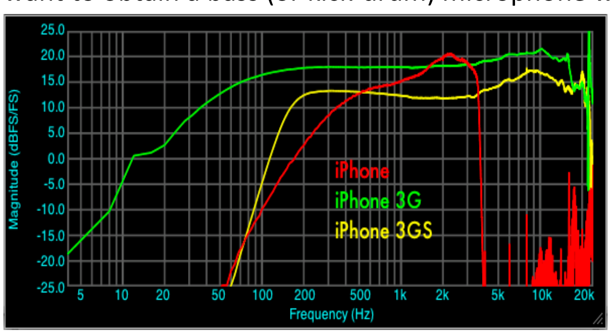

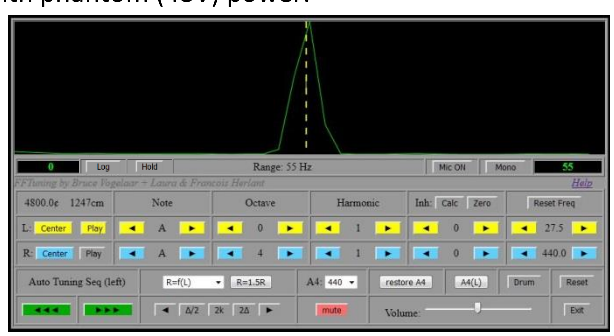

To check, simply launch FFTuner and have it produce A0 @ 27.5 Hz (ideally into a subwoofer), open a duplicate instance of FFTuner into another window and see if you can see the peak like in the figure above. (Check by changing the note and watching the peak move to ensure you're not seeing just a room background. I used a $30 condenser mic with 48V phantom power; the peak was not visible at all with any of my unpowered mics.)

There are many webpages describing how to use the tuning hammer. Please read several of them first. In general the amount of actual turning is very small (like several degrees). It is probably best to start by reducing the tension ever so slightly while plucking (striking) the string, to make sure you see the frequency decrease in the spectrum, and then tensioning to just above the desired level, and then finally nudging the tension back down to obtain its final value. The direction of this last step helps hold the tune longer. You can BREAK strings! If the frequency is not changing when you think it should, double check that your hammer is on the right pin! Also, please read the Disclaimer section below. Call a professional tuner when in doubt. Pianos come apart by design, so a broken string is not the end of the world, but a really unwanted pain. Taking the ‘action’ out of a piano to fix a broken key mechanism is also more straight-forward that you might imagine, but I’ve only done it on my own pianos. If your piano has real value to you, get things done professionally. If a tuning pin really won’t stay where you want it, there are many webpages describing steps to take.

G) Guitars

Guitars have twelve frets per octave and are equal-tempered. The six strings on a classical guitar are tuned to E2, A2, D3, G3, B3, and E4. In practice, you only need to set one of these notes, typically E2 (6^{th} string to 82.4 Hz). Then, since each fret is separated by a semitone you press just above the fifth fret to play A2 with the sixth string, which can then be used to tune the 5^{th} string, etc (see lower right diagram).

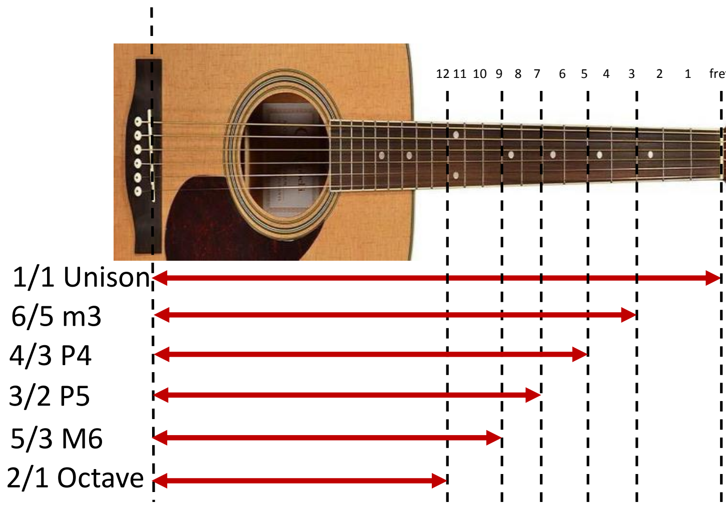

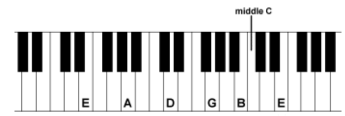

String	Note	Frequency
6	E 2	82.4
5	A 2	110.0
4	D 3	146.8
3	G 3	196.0
2	B 3	246.9
1	E 4	329.6

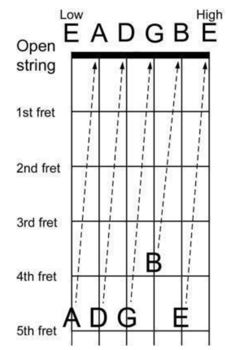

The spacing between frets is interesting. Since we know for the fundamental that $f = \frac{v}{2L}$ (where v is proportional to the square root of the tension), the spacing between any two frets, x, should be:

$$\frac{f_{n+1}}{f_n} = \frac{L_n}{L_n - x} = \sqrt[12]{2}$$

where L_n is the distance from the bridge to fret n. This then simplifies to

$$x = \left(1 - 2^{-1/12}\right)L_n = 0.05613 L_n$$

But, this is only true if the tensions are the same! When you press on the string between frets, you change the tension by an amount which depends upon which frets, and which strings, are involved. So even here reality must be accounted for, and bridge angles come into play (note in the picture above, that the bridge is indeed slightly angled).

You can easily see this effect using FFTuner. Select E4, second harmonic, center on it, and then zoom in. When you pluck the string one, you will see the second harmonic show up. Tune until it is centered on the predicted value. Next, play E5 by pressing the string just above fret 12 (which divides the string in half). Note that the frequency is now higher. That is because the string is tauter. (You could also press just below fret 12, pluck between there and the nut, and the frequency will high by the same amount.)

Another interesting demonstration is to gently touch the string at ½ (fret 12) , 1/3 (fret 7), 1/4 (fret 5) its length while plucking it. This will create a 'node' (zero displacement) at that location, so you will hear frequencies at 2, 3, or 4 times the fundamental only. This is called 'playing natural harmonics.' You can also easily see the effect with FFTuner.

This brings up 'tuning with harmonics'. The method above using the fifth fret (a perfect fourth), involves actually fretting one of the strings (which raises its tension). However, we know that for a perfect fourth, the forth harmonic of the lower note should be at the same frequency as the third harmonic of the higher note. With FFTuner it is easy to see and tune to these harmonics, without even having to dampen the other harmonics. Try tuning strings 5 and 6 this way. Then fifth-fret string 6 (making it now also an A) and you'll see its third harmonic is too high. Pick your poison!

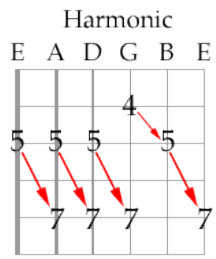

H) Drums

'Tuning' a drum kit is very individualized, and there are many articles on the web describing the process. But they sometimes use words with well-defined technical meanings in rather creative ways. The drum shell defines the basic parameters: larger diameter means lower frequencies possible; shallower depth means 'brighter' sound. There is a top 'batter' head, and often a bottom 'resonance' head (usually thinner), and the column of air between these two planes forms a resonant cavity. It is a strongly coupled system, with pressure waves from the batter head exciting the resonant head (along with the shell). The efficiency of converting strike energy into sound waves increases with total surface area, and for certain vibration modes. The various tensioning of the heads, coupled to a shell, can then produce a cumulative complex sound (composed of multiple discrete frequencies) with fast decaying 'pitch' sounds and longer resonant sounds, more like a 'musical note'. The details of each drum are unique.

The oscillation modes on an ***ideal drum head in a vacuum*** look like this (note that frequencies correspond to zeros of Bessel functions, and are not simple multiples of the fundamental):

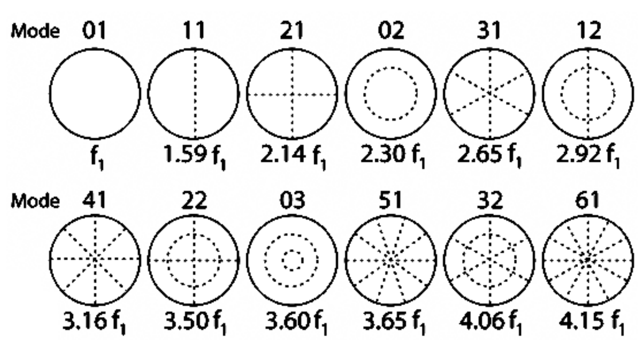

But your drum is not in a vacuum, nor is it supported by an ideal rigid circular frame. The result is that your spectrum spacing will not be the above.

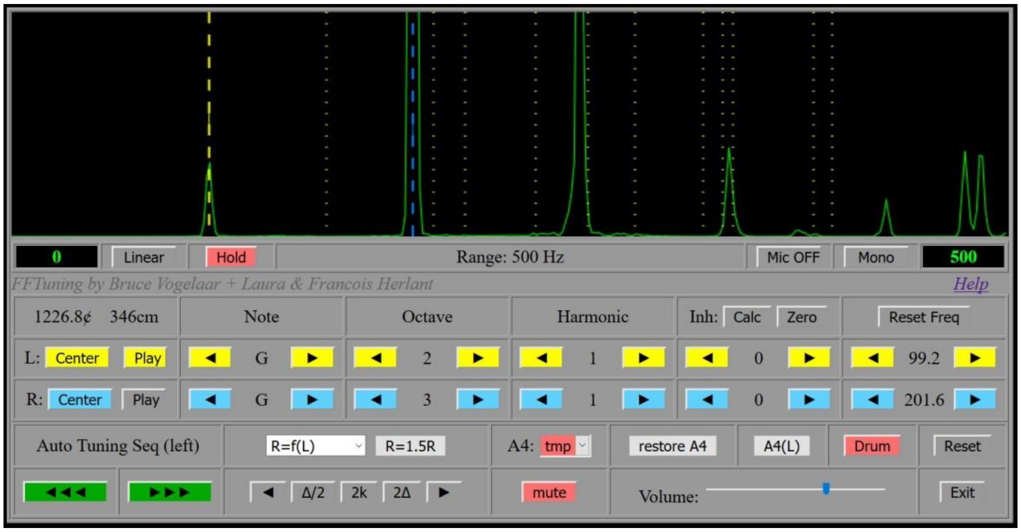

FFT can be used to demonstrate this. Note that toggling 'Drum' in FFTuner displays the expected overtones for a single ideal head in a vacuum (clearly not the case at all). I took my snare and removed the bottom head (and springs) to identify the modes of just the batter head. If the shell was tilted on its side, the dominant frequency (mode 0,1) was at about 99 Hz, but it decayed away rather quickly (you can use the linear scale and 'hold' feature of FFTuner to see this more clearly). If you bring a large piece of cardboard about an inch from the bottom side, the 0,1 frequency drops, revealing the impact of air loading (not being in a vacuum). Laying the shell flat on the floor almost eliminates the 0,1 peak. Leaving the shell flat on the ground, the second peak (mode 1,1) near 202 Hz was a single peak when the head is evenly stretched, but splits into two frequencies if there is non-uniform tension (symmetry breaking). One can identify the 1,1 mode by lightly touching with two fingers (or use a round pencil pressed on its side) along a diameter to define the node line of this mode, which dampens all but this mode of oscillation when striking. If the peak is split, you can excite one or the other frequency alone, depending upon the angle of the diameter you held still. You can next identify the 2,1 peak by using three finger touches: one at the center, and two others 90 degrees apart, then striking between the three fingers (or using two pencils 90 degrees apart). The peaks near 283 Hz become relatively larger.

If you strike the drum at its center, you will mostly excite mode 01 (but it converts to sound well, and dies out relatively quickly). If you strike near the edge, you will excite higher modes. The high modes also die out quickly (converting their energy into heat), but between these extremes the sustain is longest for the 1,1 mode. The math is beautiful, but messy. In fact, the details of coupling can move resonant frequencies significantly, so focusing too much on any but the fundamental can seem transient. ***Having a microphone designed for low-frequency response so you can see the "fundamental" can help reduce confusion.*** Also, exactly ***where*** you place your microphone relative to the drum and ***when*** you look will ***significantly*** impact the peaks appearing in the spectra.

Tuning is an iterative process, and the key is to explore a bit first. (Let us assume your drums have two heads which have already been seated.) Find out where your drums sound best individually. Start with a loosely-tightened batter head, and dampen (or remove) the resonance head (I used a soft pillow). Tune the drum head to itself first by making sure the tension in the head is uniform in all directions using rimshots near each lug. Listen for higher or lower pitches, and turn opposite lugs with the drum key as necessary to get them balanced. You can also use FFTuner as described above. Then slowly tension or release the entire head to change the fundamental note. Tighten diagonal opposites when doing so:

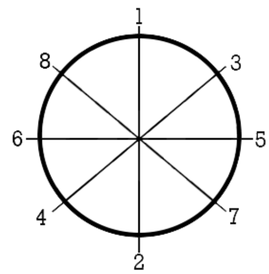

Repeat the process for the resonant head, setting it initially to the same fundamental. Move the fundamental of both heads up and down in pitch until the drum produces its best volume and longest resonance. You may want to note the frequency of a stronger mode for future reference using FFTuner. Pictured is an example for a 16" x 16" tom I was tuning, with the resonant head damped.

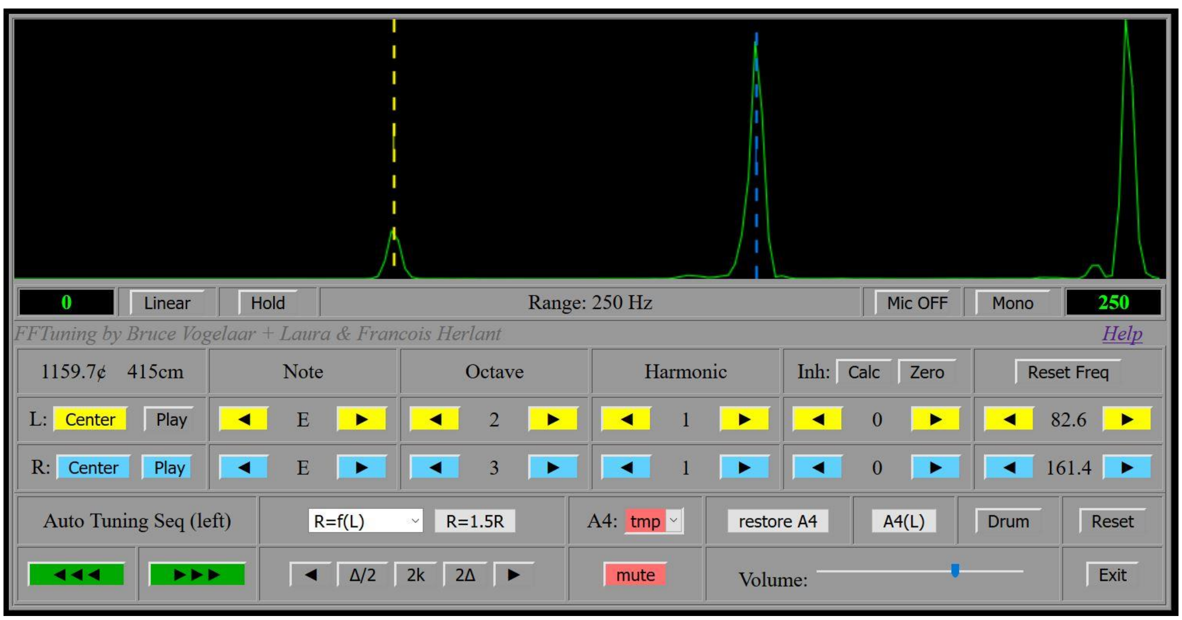

Next, try tuning the resonance head higher or lower in frequency than the batter head (picking one of the larger features to scale). You can try ratios associated with various steps, such as m3, P4, P5, etc. Doing so will change the resonant characteristics of the drum. Higher should be more 'lively' with shorter resonance time, and lower more 'depth', also with shorter resonance time. You might find the overall fundamental also changes frequency a bit. For a floor tom, the height above the floor can change the frequency of the fundamental!

Finally, you might want to select and then tune for the interval between different drums. One might choose P5, P4, or M3, but it is really up to you. There is no rule about this, other than having a head tuned to itself is pretty important (uniform tension). Also, tension will impact stick rebound. There is nothing like trying it out when you have some time to kill – you will quickly discover there is no 'correct' method, since modes are ‘witchy’, and ultimately you’ll need to choose what sounds good for you and your style of playing (versus a piano string which is wonderfully and consistently very harmonic).

The kick drum and snare have special requirements, and you should read elsewhere about tuning them.

I) Disclaimer

We are not professional piano tuners. We suggest you engage a professional tuner if you can. We do not make any warranties about the completeness, reliability, or accuracy of this information. Any action you take upon the information in this program or document is strictly at your own risk, and we will not be liable for any losses and/or damages in connection with the use of this program or document.

PS: in Chrome, on your laptop, FFTuner can take the FFT of its own sound output. In Firefox, you would need to launch a duplicate window, so one instance can hear the other instance. It is not currently understood why this is the case. The peak starts, but quickly disappears, as though being self-filtered. Any thoughts you may have and want to share are welcome…