String Theory

Strings seem to be pretty simple, and we usually take them more or less for granted. As is so often the case, that appearance is a bit deceptive, and one clue to that fact is the amount of disagreement you will get about them. In the April 2007 issue of Acoustic Guitar, for example, in an article about steel strings, Dick Boak, of Martin Guitars, is quoted as saying that phosphor bronze strings have a "brighter, more bell-like tone" than 80-20 bronze, while, a paragraph or so down, Dave Cowles of GHS says the 80/20 bronze has "...a brighter tone than phosphor bronze". This sort of disagreement is actually pretty easy to resolve in some ways; it comes down to a matter of taste, and the likelihood is that the two gentlemen simply mean something a bit different when they say 'bright'. There are more substantial disagreements, though. The one that got my attention was that between the physicists and the luthiers. A lot of that probably comes down to differences in purpose, and the differences in approach that come out of that.

Broadly speaking, physicists are most interested in developing equations that will enable them to predict what a system will do. They will often leave out small details if it will help to keep the math simple, and fill them in later if they need to. Luthiers, on the other hand, are generally more interested in a qualitative understanding, and the ability to visualize what's happening. We often find the little details important, and take the big things for granted. As an example, for most physicists the difference between nylon and steel strings is minor, while for us it's very important.

One of the results of these differences in outlook is in the approach we take. Physicists usually look at strings in the 'frequency domain', analyzing the vibrations in terms of the different frequencies the string can produce. If the string is assumed to be an 'ideal' one, with no stiffness, constant tension, no losses, and fixed ends, the math is fairly simple and iterative: the equations that describe the way a string vibrates as a whole at its fundamental frequency also hold for the same string at fractional lengths. That is, if you divide the ideal string into some number, N, of equal length pieces, each piece vibrates like the whole string, but at N times the frequency. To find out what the string is doing in a more complex case, you just solve the same equations for as many fractional lengths as you feel you need, and add up the results. If you need to account for changes in tension, or stiffness or anything else, the corrections can be added in later, although the math can get fairly tricky. Normally, these corrections are 'small' though. A drawback with this approach is that it's not always obvious what the entire string is doing from one instant to the next once it's been plucked.

Another approach that is less common in physics, although just as valid, is that of the 'time domain'. One looks at the shape of the string just at the instant it is released in plucking, say, and watches how that shape changes. If the 'ideal' string is used the forces on the ends are easy to calculate, simply from knowing the tension and angles, and some things like tension changes, that are harder to include in the frequency domain treatment, are rather easy to deal with. Other corrections, such as those for string stiffness, are also reasonably easy to include. For our purposes, this is usually a much better way to do it. However, there is the drawback that it's harder to figure out just what frequencies there are in the signal. The frequency and time domains are just two ways of looking at the same thing, and they are exactly equivalent if you are careful to include all the corrections. Which you start out with will depend on what you think is important. Unfortunately, it is fairly easy to lose sight of the fact that either approach, in its simplest form, leaves out some things that might make a difference to us as luthiers.

This is what gives rise, I think, to one of the disagreements that come up between physicists and luthiers. In the simple frequency domain model, tension is held constant, in part because accounting for the change in tension of the displaced string adds a lot of complication to the math for only a little extra precision in the results. Using this simple model to calculate the force driving the bridge on the guitar would give only a transverse force: if the string is moving up and down relative to the soundboard it is pushing the top in and pulling it out, forcing a motion like that of a loudspeaker. We luthiers know that raising the strings higher off the top of the guitar can change the tone, and one way this might work is if there is a tension change in the vibrating string that torques the bridge top toward the neck. The physicist who is sticking to the simplest model will say the tension change is either not there, or too small to worry about, and can assert that any perceived difference in the tone is probably 'just subjective'. The luthiers end up wondering how much these physics guys really know, and some have decided that the tension change is more important than the transverse force. From time to time people have even solved the equations and said that was the case. There is always the possibility that they made an error in their math, but that can be hard to check if you don't have the skill.

You'd think this would be an easy argument to settle: just look up the experiment in which somebody measured the tension and transverse forces. The trouble was, it seemed as though nobody had. I can think of a couple of reasons for this. One is that the vibrating string has been a staple in physics texts for many years; they feel they understand it already, since the equations have been written, and work pretty well. Besides, the experiment is not all that easy without some good electronic equipment, such as a computer with a sound card, which was not common until fairly recently. The folks who have the equipment tend to want to look at more 'interesting' problems. So I decided to set up the apparatus to measure the tension and transverse forces that a plucked string would develop at the bridge, and, I hoped, settle the argument.

I built up a heavy beam of persimmon wood, in the shape of an inverted "T", that could be clamped to the bench for even more mass and greater stiffness. The end stops were brass levers, the 'nut' moving vertically to measure the transverse force, and the 'bridge' moving horizontally for tension change. Small pieces of PZT-5 piezo ceramic material were sandwiched between the beam and the bridge and nut levers, and provided with the needed electrical contacts. These pickups would then output a signal proportional to the force the string was exerting as it vibrated, which could be fed into my computer, recorded, graphed, printed out, and analyzed. To 'pluck' the string with a known force at a known point a length of fine (about #44) magnet wire from an old relay was looped under the string and pulled upward until it broke. The quality control on this stuff is so good that it always breaks at the same load, within about 3%. What would you expect the output of the pickups to be? A look at the way the string might move over time allows us to make some predictions.

We can start from one thing we know, and one thing that at least makes some sense. We know that a stretched string makes a straight line between the two end points, so that if we pull it aside somewhere in the middle, the two sections on either side of the plucking point will be straight lines, joined at a 'kink'. It is reasonable to think that if you stretch a rope or a long spring, and strike it, so as to make a kink at one point, the kink will run out in both directions from that place at a constant speed until it reaches the ends, where it will reflect and start back. This can be shown to be true within limits by the people who have done the math.

As we pull the string upward, we are, in effect, pulling upward on the saddle, and if we know what the tension is in the string and the upward angle, we can calculate that force. We also know that pulling the string aside stretches it a bit, and, again, if we know a few things about the string and the distance it's pulled aside, we can calculate the tension change as well. Here's a picture of the string pulled aside with a wire just at the instant of the 'pluck'. (#1-a) The 'cycle clock' just tells you where you are in relationship to one complete cycle of vibration.

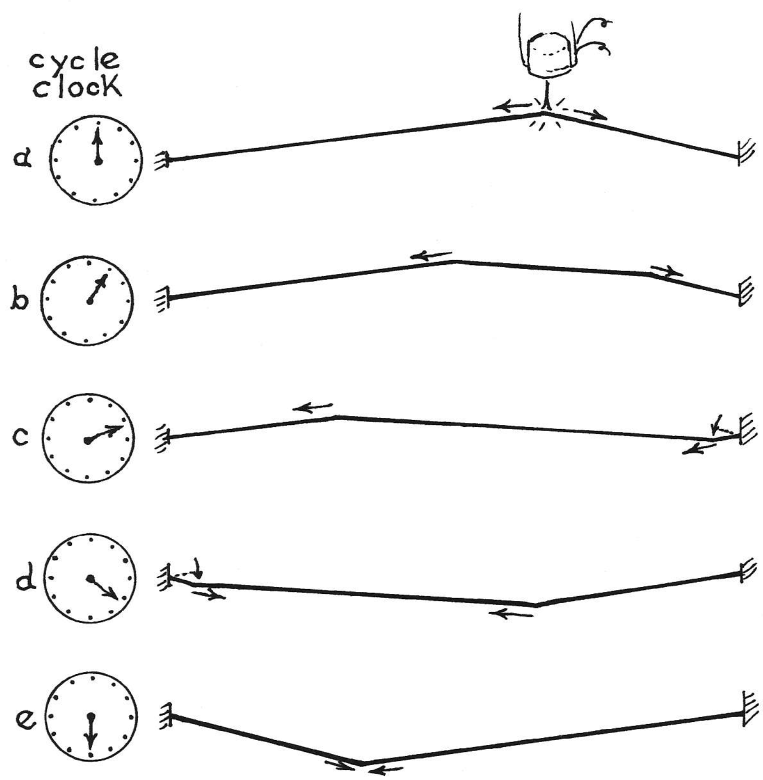

A short time later it looks like this. (#1-b) The 'kink' from the pluck point has moved a certain distance toward the ends, so the string is now three straight segments joined by two kinks.
It's a little shorter than it was before we released it, so the tension has dropped some. Note, however, that the up angle between the saddle and the string has not changed, so the upward force is almost exactly the same as it was to begin with.

After a little bit the first kink reaches the end of the string and reflects off the saddle. The angle flips from 'up' to 'down', and the angle is a bit smaller, so the force on the saddle has gone suddenly from an upward to a somewhat smaller downward one. The string now is even shorter than it was, so the tension has dropped a bit more, but, now that both kinks are running the same direction, the length won't change until the other kink reaches the end. (#1-c)

When that happens, and the two kinks are running toward each other again, the tension will rise a bit. (#1-d)

As the two kinks cross through each other we can see that the string is in the same shape it was just before the pluck happened, but upside down and backwards, so to speak. (#1-e) This defines one half cycle of the complete vibration. From here the recording runs along in the same way for the next half cycle, until it is in the same shape it was to begin with, and the whole thing starts over.

Thus, for every cycle of vibration of this upward pluck, the transverse force will be 'up' and at a fairly high level until the kink hits the saddle, and then it switches to a lower level 'down' force until the next kink hits it, and it switches to 'up' again. It is a square wave, with a duty cycle that depends on where the string was plucked. This description, of course, assumed a wire pulled upward for the pluck ; when playing you'd normally start by pushing the string down, and the transverse force wave form would be flipped, but of the same shape.

The tension of a vibrating string will always be higher than the static tension of the string. It starts out high, drops until the first kink hits the saddle, stays level until the second kink hits the nut, and then rises again until the two kinks cross, when it starts to fall. There is a triangular 'bump' in the force twice for every cycle of the transverse vibration, and these bumps are separated by a level force somewhat greater than the resting tension of the string. So, in an ideal case, this is what we'd expect to see. (#2)

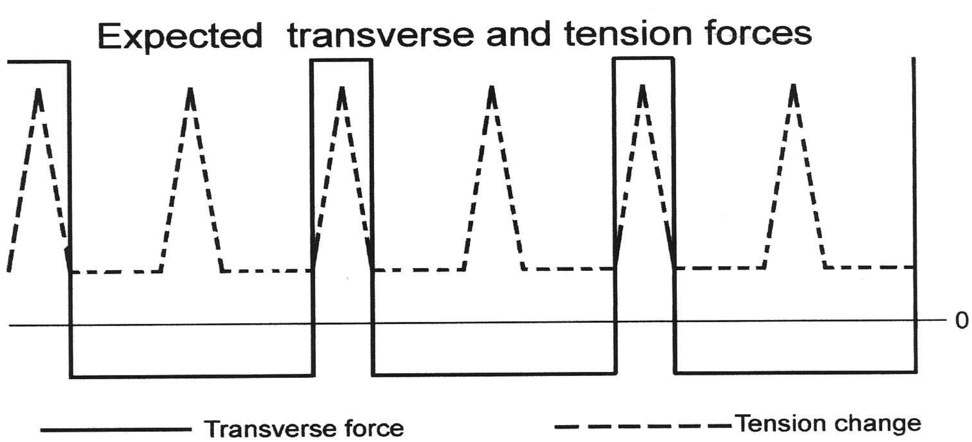

Here's what I actually saw. (#3)

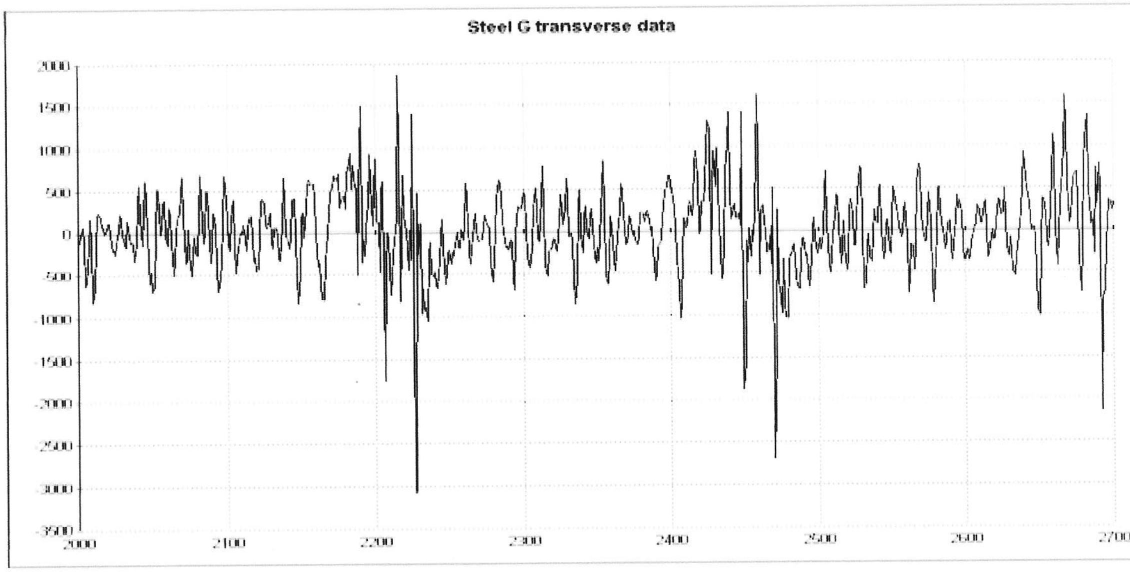

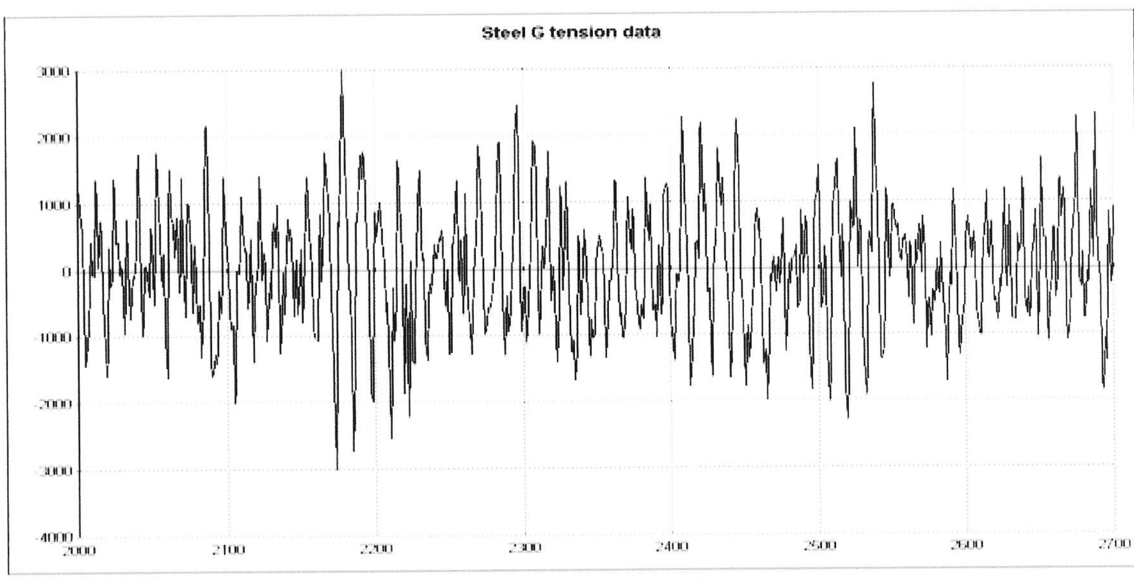

Even without the high frequency stuff, it's not much like what I expected. Still, a frequency analysis showed that it had all the right frequencies (with some added highs) in the right proportions, so I went ahead and tried different strings at different tensions and looked at things in the frequency domain. I actually wrote up a short report, dropped it off for a friend, and was driving home when I realized what the problem was.

Electrically the piezo elements I was using for sensors are tiny capacitors; well under 1 Pico farad. The relatively low input impedance of the soundcard on my computer would drain off their charge so fast that, even though I was sampling as 48,000 times per second, the signal I was seeing was not the force, but rather a measure of how much the force had changed since the last reading. Mathematically it was the differential of the force, and what I had to do to find the force itself was to integrate, or sum over, all the values.

Fortunately, the old DOS Fourier Transform program that I often use, FFT4WAV3, has a utility that will export any numerical file it can open as comma delimited text, which can be opened by any spreadsheet. From there it was easy to sum over the data and reconstruct the actual force waveforms, and here they are.

(#4)

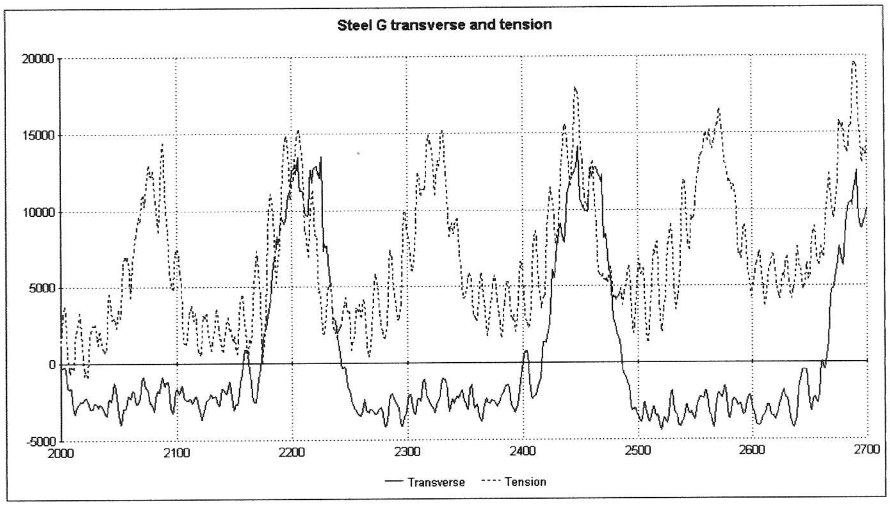

I'll note here that my soundcard only allowed for monaural input, so the tension and transverse waveforms were from separate plucks, and had to be aligned and zeroed separately to get the graphs. The numbers at the bottom are simply sample numbers, and reflect the time after the nominal beginning of the pluck signal.

There is still a lot of high frequency stuff in there, particularly in the tension signal, but, as it is quite regular, it can be smoothed out by averaging over a number of cells in the spreadsheet equal to one period on a running basis. Here is the smoothed data.

(#5)

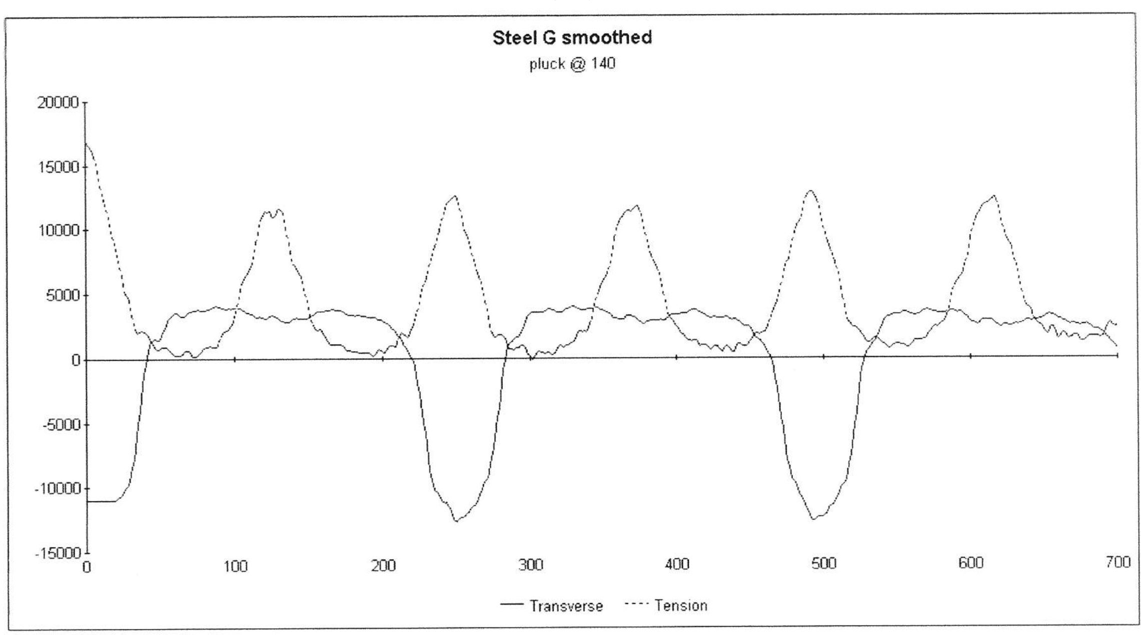

This looks pretty much like what we'd expected to see, just rounded over at the corners by the loss of high frequency from the mathematical smoothing process, and with the transverse force flipped over to make it easier to see.

The charts so far have been from a plain steel string, 25.6" vibrating length, tuned to G=196 Hz. This is a very low pitch for a plain steel string. Remember that the transverse force will be determined by the tension on the string and the angle it makes with the saddle, and will be fairly low for such a slack string. On the other hand, the tension change will largely depend on the material and diameter of the string, and how much it's displaced, and won't vary with the initial tension. The tension in a given string might increase one pound when it's displaced 1/8" from rest at a certain point, and that change will be much the same if the string has five pounds or twenty pounds of tension on it to begin with. Since the tension signal is a measure of that change the ratio of the transverse and tension forces will be expected to vary for a given string depending on the pitch.

Material makes a difference, too. A material like steel, which has a high Young's modulus, will not stretch much as it is displaced, so the tension will increase a lot. Nylon, with a lower Young's modulus, stretches more, and undergoes less tension change, so the ratio of the transverse to tension signals will be higher for nylon than steel. This chart shows those ratios at different tensions relative to the theoretical breaking point of the material.

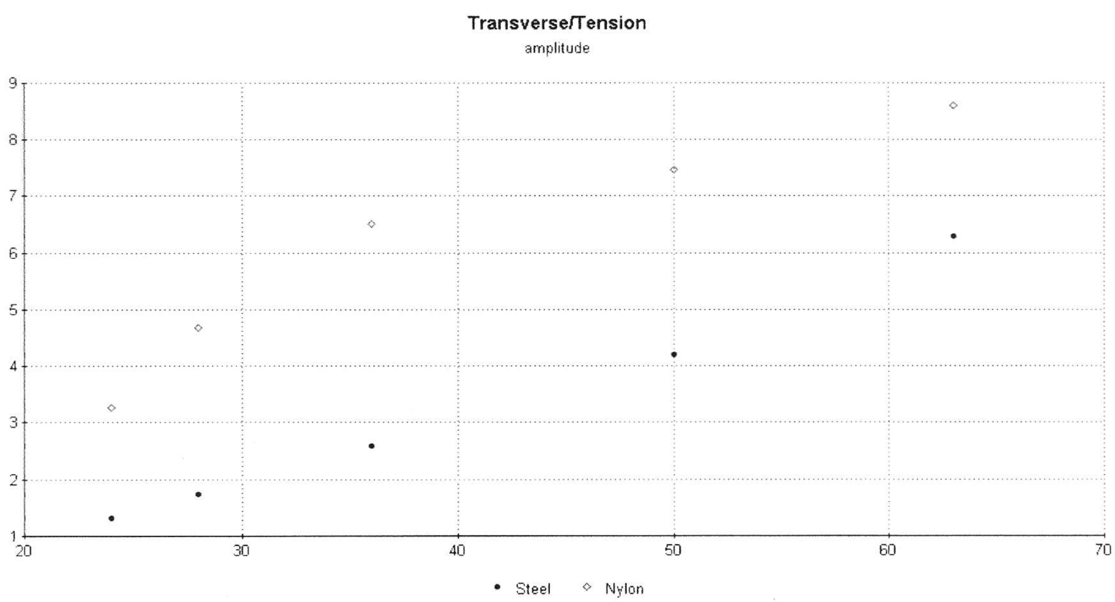

Note the dogleg in the nylon plot. There are several different types of bonds between and within the molecules that make up nylon, and these act like springs that stretch at different rates. Think about picking up a weight by lifting a handle that is attached to a bungee cord, and a strong steel spring. As the load goes up the looser bonds get saturated, and the Young's' modulus of the material rises in steps, altering the Tr/Te relationship. Just to show the other extreme from the slack steel string, here is a chart of the smoothed transverse and tension signals from a plain nylon string, tuned to about high E.

(#6)

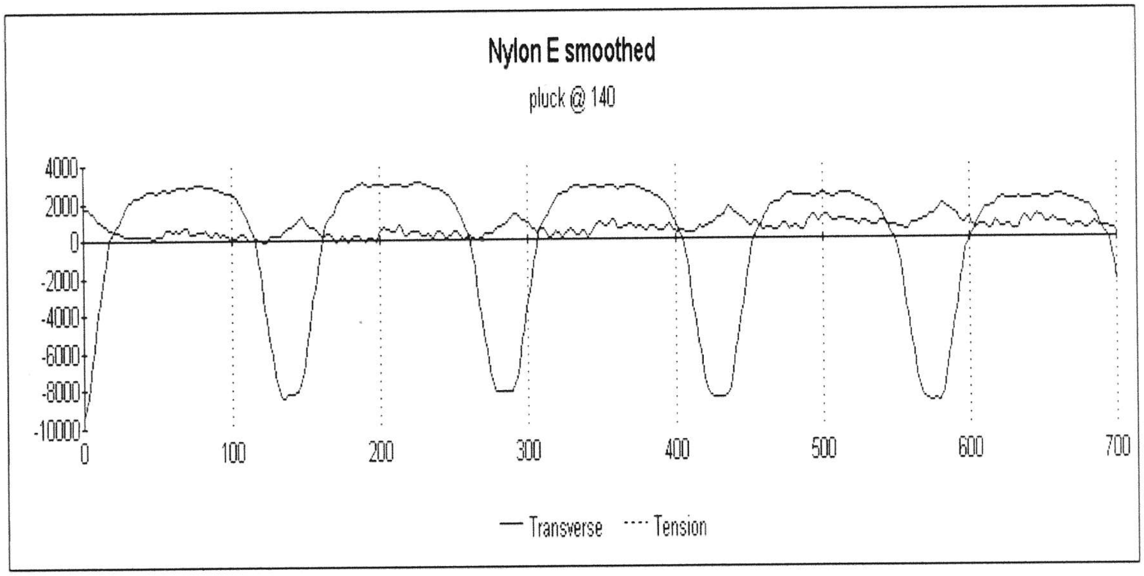

In the unsmoothed chart the high frequency component of the tension change signal pretty well swamps the tension change itself.

I should note here that those Tr/Te ratios were obtained from wire break plucks, which are of constant force. For the slacker strings the displacement would be greater than for tighter ones, and that, in turn, would tend to give a relatively greater tension change. The ratios for constant displacement would be a little more alike at different tensions. Another variable is the plucking point. For a given force pulling the string aside you will get more displacement if you pluck toward the center of the string, so, again, the closer to the center the pluck, the lower the ratio of transverse to tension change amplitude.

At this point I'd found the information I set up the experiment to find, but there seem to be some other things I had not expected to see. While I was at it, it seemed logical to look at these, and the string rig gave me the chance to look at a couple of other things I already knew about.

One of the already known things is 'inharmonicity'; the fact that the overtones of a string are not generally at exact multiples of the fundamental frequency. If the only force that was straightening the string out when it was displaced were the tension, the overtones would be at the proper harmonic intervals, but since real strings have some stiffness, the upper partials tend to become more and more sharp the higher you go. Here is a chart showing how much the upper partials of a steel A string (Thomastic 'Plectrum' series) are sharp from the expected pitches when the sting is mounted on the rig and plucked vertically. (#7) One would expect that thicker and slacker strings would show more inharmonicity, and that is likely one of the problems with the nylon G string. Indeed, when comparing plucks on a classical guitar, the 9th partial of the G string was sharp from the

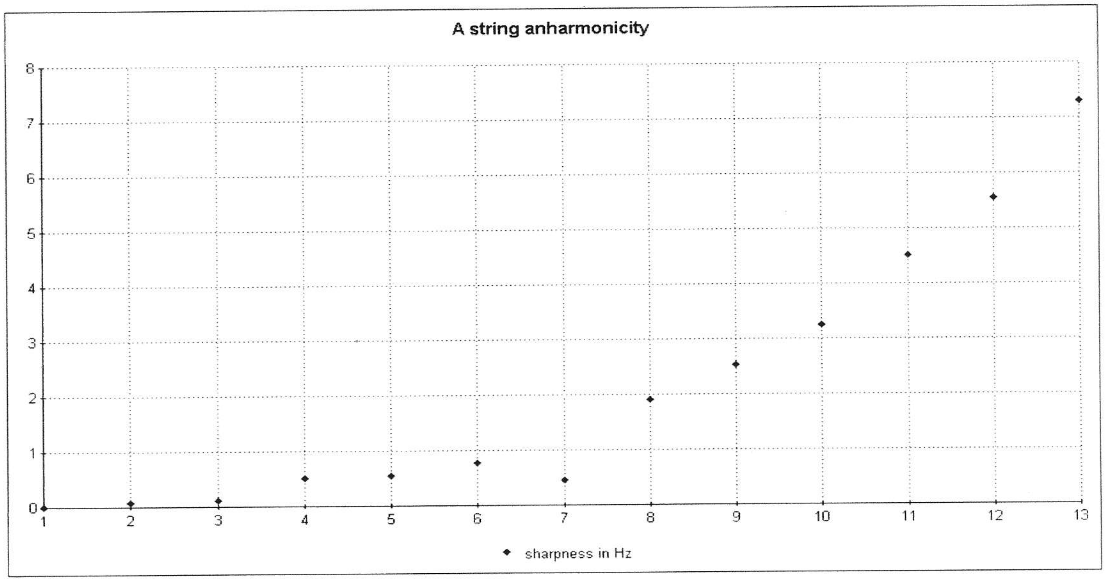

true harmonic pitch by more than twice as much (6.9 Hz) as that of the E string (3.1 Hz). Since our impression of pitch is probably gained from a sort of average of all of the partials we hear, this could account for the somewhat 'fuzzier' pitch sense we tend to get from the G than from the E string; everything is not pointing to the same pitch.

Another known complication of real strings, as compared to ideal ones, is damping. An ideal string loses no energy as it vibrates, but internal friction, and loss due to air drag slow down real strings. High frequencies tend to be more effected by both of these. Nylon has higher internal damping than steel, as a material, and nylon strings, being fatter in general, have more air drag. Thus, while the initial levels of the transverse force in both nylon and steel E strings are the same after a wire pluck, the nylon string falls off much faster. Averaged over the first second or so, the nylon string only has about half the energy of the steel string, but a higher proportion of the nylon string's energy summed over that time is in the fundamental. After vibrating for a second the nylon string has very little energy above 6000 Hz or so, while the steel string doesn't start to drop off until it hits 8000 Hz or higher, and it still has some power at 10000-12000 Hz. This has obvious implications for the design of guitars. One problem in making a good classical guitar is to preserve the small amount of high frequency energy in the strings, while in steel string instruments the issue is to get enough bass to balance out all of the high-end energy and brightness.

That pretty well covers the main points about the transverse and tension change signals in the string itself. Later we'll have a look at some of the implications of these things in guitar design and construction. For now, what about that high frequency stuff in the tension change signal? Here is a chart of the smoothed tension change signal from the plain steel string, tuned to G=196, with the high frequency part that was subtracted superimposed on it.

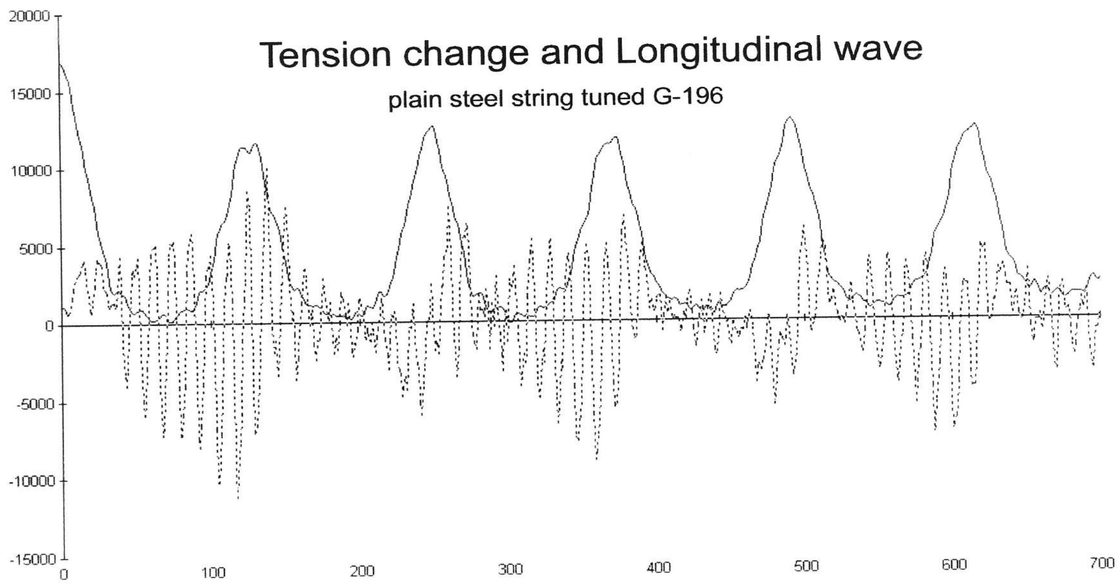

The first thing to note is that, although the high frequency comes and goes in step with the tension change signal, it is itself pretty regular: it's not just noise, but a signal based on a particular frequency. Given the sample rate, and the number of cells in the spreadsheet that were averaged over to smooth it out, it looks as though the frequency of that wave is around 4000 Hz. This turns out to be just right for a longitudinal compression wave along the length of the string, analogous to a standing wave in air in a long pipe. How could such a wave get going? One clue is the fact that it's present right from the start; something in the pluck feeds it.

All of these plucks were started by pulling the string upward at a point about 4/5 of the way along the length (140mm up on a 650mm string length). If you were to pinch the string hard between your thumb and finger, and pull it straight up at that point, the short section near the bridge would be a bit tighter than the longer section at the nut end, simply because it has made a greater angle of displacement. Since the real string is not pinched it will move a little bit lengthwise to equalize the tension between the two ends, thus shifting the center of gravity of the string toward the bridge. A half cycle later, when the shape of the string is the inverted reflection of its initial shape, the center of gravity will have shifted just as far toward the nut. Thus the whole string is vibrating along it's length at the fundamental frequency. This feeds energy into the longitudinal compression wave. If the longitudinal wave is at some integral multiple of the fundamental, the energy input will either be exactly in phase or exactly out of phase with it, and the longitudinal wave will either die out quickly, or build up to a very high amplitude. Since the match is seldom exact, the wave form usually comes and goes, as we see here. The interaction is not all one way, either: the same smoothing operation that was used on the tension signal also smoothed out the transverse waveform, so there is coupling between the longitudinal and transverse waves. The main influences on the pitch of the longitudinal wave are the length of the string, and the Young's modulus and density of the material. The tension and diameter of the string itself should not have a direct effect.

To test this out an easy way was found to excite the longitudinal wave. By applying a little rosin to the string, and rubbing along its length with a rosined rag, it is easy to get the longitudinal wave going, and to record the pitch for analysis. Here is a chart of the longitudinal pitches of sets of strings, two each of steel and nylon, mounted on guitars.

(#9)

D'Addario 'Pro Arte' hard tension		LaBella 2001 X-hard
E	629	596
A	818	780
D	1070	1020
G	1816	1700
B	2027	1896
E	2466	2418
No-name steel on Dread		Elixer Med.
E	1450	1400
A	1600	1484
D	1950	1808
G	2350	2232
B	3860	3855
E	unknown	3786

The plain steel E and B strings (one E string broke) have longitudinal pitches around 4000 Hz, as predicted. The wound strings have lower pitches; evidently the windings add a load to this vibration. The longitudinal pitch change in the plain nylon strings has to do with the stretching of the molecular bonds as the string is stretched that I mentioned before. As the tension rises the effective Young's modulus of the nylon also goes up, raising the longitudinal pitch.

Assuming you stay in standard tuning, the transverse pitches of the strings will always be the same, no matter what scale length you use. However, the longitudinal pitches will change with the scale length, so the relationship between the longitudinal and transverse pitches will be altered. I believe this was the gist of the thinking behind Ralph Novack's 'fanned fretting'.

It was pointed out to me recently by a friend that classical guitar D strings often buzz when there is no obvious reason for them to do so. It turns out that the seventh partial of the D string is in the range of the longitudinal pitch, at around 1050 Hz or so. Chart #10 shows a 'split peak' at the seventh partial frequency of a not particularly buzzy D string on a guitar, possibly caused by the longitudinal/transverse couple. The low, broad peak around 1055-1060 is the air resonance between the top and back plates; it is spread out because they are somewhat flexible, and the body is tapered. Still, the close proximity of the air resonance to the string partial does enhance the output of that pitch.

This problem seems to be somewhat random, with the exact pitch of the longitudinal mode probably depending on the ratio of the wrap density to the core size. Small variations in

Chart #10

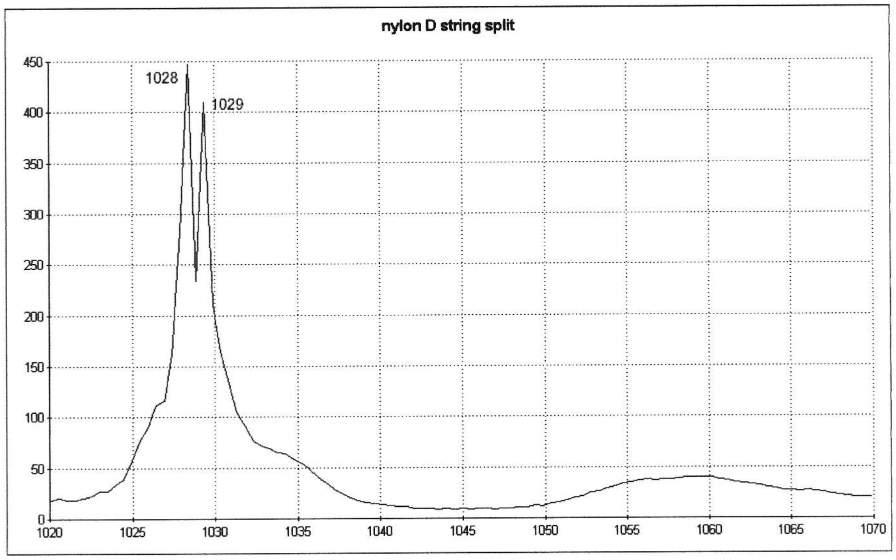

the wrap could put the longitudinal pitch in or out of the right range. It is possible that simply twisting a buzzy string could alter the longitudinal pitch, since it would wind up, or unwind, the wrap a bit, as well as giving a twist to the fibers of the core, and thus changing its effective Young's modulus. If your classical guitars suffer from this problem frequently, a small increase in the scale length might help. It has been noted that sometimes the buzz goes away as the string ages, and it's probable that the load of dirt that normally builds up between the windings would drop the longitudinal pitch.

There is one other way in which a string might vibrate that has been posited to have an effect on the tone of the guitar. Fred Dickens pointed out years ago that a plucked string normally rolls off the end of the finger or pick, and so is given a torsional vibration. By putting tiny flags on the strings he observed that the timbre of the note tended to change when the torsion mode died out. It is known that this mode has an effect on violin tone, being the cause of the infamous string squeak that gives beginners so much trouble.

It proved difficult to activate this mode with a bowing machine on my rig: the torsion vibration tends to phase lock at some multiple of the transverse frequency that is close to it's natural pitch. However, by gluing a piece of glitter to the string, and using a bright white LED as a strobe, driven by my signal generator, I could project a dot of light on a screen of frosted Mylar arched over the string. Adjusting the strobe frequency as the string was plucked made it possible to isolate the torsion frequency on some strings. These vibrations die out very quickly. There is some evidence in spectrograms of an effect in the sound just after the initial pluck, but it is hard to isolate. It thus seems as though this is a part of the attack transient, possibly affecting tone color, but of much less importance to guitar players than violinists.

So far we've been looking mostly at the string on a more or less rigid mounting. We've seen how the transverse vibration can give rise to tension changes at multiples of twice the fundamental frequency, and also to a longitudinal vibration at a higher pitch. All of this may be interesting, but it's also a bit theoretical unless we can use this knowledge to make better sounding guitars. Thus a few experiments were done to see what effect differences in the setup or construction of the guitar would have in relation to these string modes.

Unlike the string on the test rig, a string on a guitar 'sees' a moving end, at least at the bridge. If the bridge didn't move, the guitar would make no sound. On the other hand, the movement of the bridge will also affect the way the string vibrates, most likely in ways that will cause problems. So there are two questions we can ask right away: how well can the string actually move the bridge in different directions, and how much does that movement affect the string?

Chart 11 shows how much the bridge on one classical guitar moves at different frequencies, given a constant driving force. 'Vertical' means that the force was directed perpendicular to the top, and 'longitudinal' is parallel to the top along the length of the strings. The admittance 'across' the top was lower in general than either of the other levels, and renders the graph confusing. In all cases, the motion was measured at the bridge location, and the 'across' and 'longitudinal' motions were measured at the string height. The vertical axis has been shown in dB, a logarithmic scale, which exaggerates the differences at low levels: otherwise the lines are too close together to resolve.

Chart #11

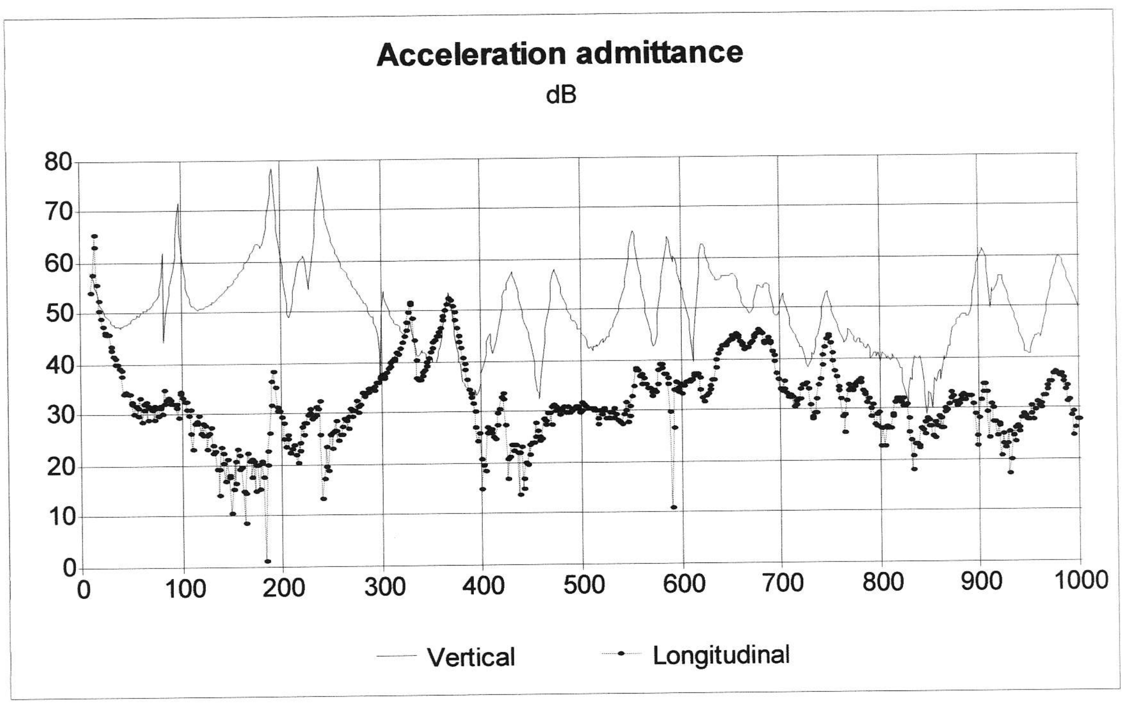

Generally, motion in the 'vertical' direction is much greater, with the longitudinal motion only approaching it in the range between about 350-400 Hz, where there are top and air resonances that can rock the bridge. That makes sense; if the bridge twists upward under the string forces that's a problem, so we design against that. A similar test done on a steel string guitar (a Dreadnought, built from a kit by one of my students several years ago) showed only about 1/3 as much motion in the vertical direction from the same applied force as the classical guitar had. The Dread also lacked the prominent 'long dipole' peak in the longitudinal chart; evidently the added stiffness of the X bracing reduces that mode significantly.

Keep in mind that the transverse force of most guitar strings is greater than either the twice per cycle tension change or the high frequency longitudinal tension/compression force. Clearly the most effective way to drive the bridge is with the transverse force of a vertical string vibration. In most cases plucking the string results in motion at some angle, with both perpendicular and parallel components to it, and the job of figuring out the rate of energy transfer from the string to the top is more difficult, although, of course, the guitar figures it out just fine.

The rate at which energy is going to be transferred from the string to the top will depend on the relative impedance of the two. Mechanical impedance is the ratio of Force over Velocity at a given frequency, so the curves we just looked at are, roughly, the inverse of impedance curves of the top for the different directions. Where the acceleration is high, the impedance is low. The 'characteristic impedance' of a string is proportional to the square root of the product of the mass per unit length and tension. Since steel strings are more massive, and hence also under more tension than nylon strings, they have higher impedance and can transfer more energy to the top, which makes up for the lower response of the Dread top at a constant force.

It is sometimes said that strings for a guitar should be chosen so that the tension on each one when it is tuned to pitch will be equal, to give a uniform 'feel'. If you do that, though, it turns out that the impedances of the different strings will vary, and they will thus tend to drive the top unequally for the same amplitude. In reality, string sets seem to be made up so that the tensions are allowed to vary somewhat, so as to make the impedances come out to be more nearly equal. Some uniformity of feel is given up to gain more uniformity of response.

Older instruments, such as lutes, often used strings that were relatively slack, and would thus have a higher ratio of tension to transverse force. It is interesting to note that these instruments also generally place the bridge far down on the soundboard, rather than in a more centered position. The centered bridge on a modern guitar will not be likely to produce much sound in a rocking motion, simply because the top areas above and below it are nearly equal, and tend to cancel each other out. It is possible that the design of the lute takes better advantage of the relatively greater tension change signal of the strings.

All this time we've been talking about the string as if the top was not moving, and this is not the case. If the top didn't move, there would be no sound produced, but motion of the top and bridge also can effect the way the string vibrates. The strongest effect is likely to be at the pitch of the 'main top' resonant mode, where there is the most motion at the bridge; normally close to the fundamental pitch of the open G string, but any partial of any string can be affected.

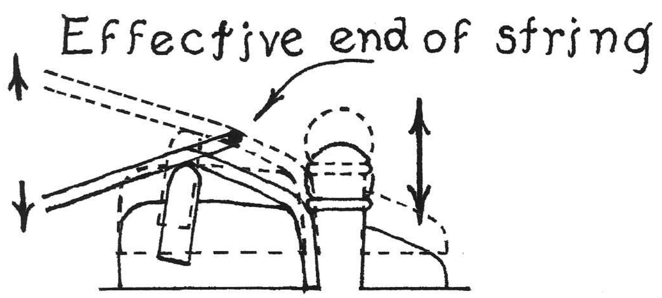

Fig 12 Moving Bridge

Suppose for a moment that the pitch of the 'main top' resonance is at 198 Hz, just above the pitch of the open G string at 196 Hz. In that case the bridge will always be moving in the direction of the string pull, at least for vertical motion. Figure 12 shows what happens; when the string is moving vertically, it will be, in effect, a little longer than it 'should' be, and the pitch of that vibration will consequently be lower than it would be, given the mass and tension of the string. However, when moving parallel to the top the bridge position will be much more 'fixed', and the string will 'see' the stationary point in the correct place, thus making the correct pitch. The same string potentially has two different pitches, depending on how it's moving, and, since a normal pluck will give some motion in both directions, it will make both sounds. It is easy to see that if the top vibration is lower in pitch than that of the string, which reverses the phase of the motion, the 'vertical' string pitch will be higher than it 'should' be. This spectrum plot of an open G string, on a guitar with it's 'main top' resonant mode at 195 Hz, clearly shows a split in the fundamental mode of the string. In this case, other measurements found that the 'horizontal' vibration produces the peak at 194 Hz, and the 'vertical' at 199, as predicted. Both the ear, and an electronic tuner, 'hear' the string as being in tune at 195.9 Hz, close to the pitch of the 'dip', but the note 'beats' somewhat, although the effect is masked by the higher partials. (Figure 13)

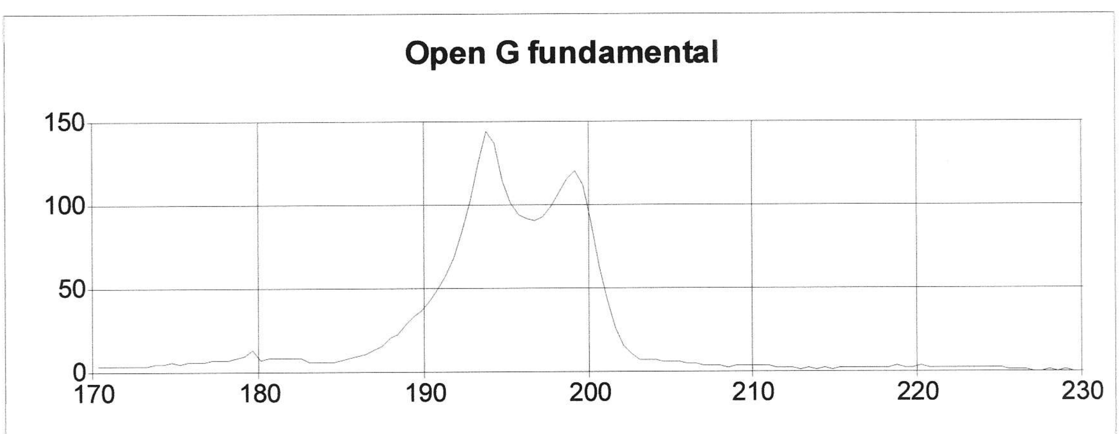

(Figure 14)

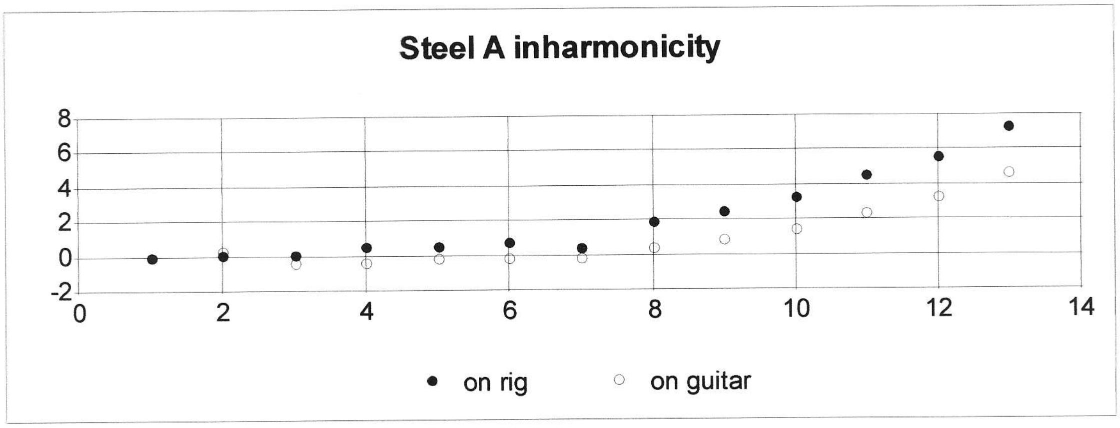

Figure 14 shows the inharmonicity of that Thomastic A string on the more or less rigid rig, and on a guitar. On the rig, if the fundamental is tuned correctly all of the partials are a little sharp of their proper pitches, as we would expect. On the guitar, with the fundamental tuned properly, the second partial is noticeably sharp, and all of the others, while following the same general curve, are flatter than they were on the rig. It is not until you get to the eighth partial that they actually go sharp of the proper pitches. The 'main air' resonant mode on this guitar, which pushes strongly on the top, is lower than the open A string fundamental pitch, and may pull that mode sharp. To get the pitch to come out right on the tuner the string needs to be slacked off a bit, resulting in a fundamental that is correct. The second partial is reacting with the 'main top' mode, which is at 209 Hz, a little below A=220, and is also pulled sharp. Other notes on that string will not be likely to be so strongly affected by top or air modes, and will thus reflect the true tension of the string by sounding a bit flat. In this way, bridge motion can affect the intonation.

I have noticed that solid body guitars often have the node lines of the lower three modes of the neck-body system near the bridge. The mass being concentrated near the lower end tends to move those lines downward, and bunch them together in one spot. If the bridge is on that spot the string force won't be able to move the body of the guitar much: it's like pushing a see-saw at the fulcrum. Solid body basses, on the other hand, tend to have the bridge as far down on the top as possible, on the most active area for all of the lower order modes. Any note that coincides with a body vibration will move the bridge, affecting the intonation as we have seen, and absorbing some of the string's energy. This may be one reason why basses suffer more from intonation problems and dead notes.

One of the issues that prompted this whole series of experiments was the way the sound is altered by changes in saddle height, and I was able to do some experiments with this on both steel and nylon string guitars. Basically, as one might expect, the main difference was that raising the saddle yielded a sound from the guitar that had more of the octave-doubled tension change signal in it. The change seemed to be disproportionate; a rather small increase in saddle height put in a lot more of the frequency-doubled signal. It was not possible in the time I had to do this to control for the break angle of the strings over the saddle, but the observed effect is easy to account for simply by the fact that the string being further off the top would have more leverage to torque the bridge when the tension changed.

I also ran some tests using different saddle materials: bone and polyethylene, to see what the difference would be between hard and soft saddles. Although it was easy to hear a difference, it proved harder to find consistent differences in the recorded spectra of guitars. Perhaps I'm looking for the wrong thing.

One well-known guitar design has the strings angled upward from the top by about five degrees at the bridge. One claim I have heard for this is that it might work like a harp, so, to begin with, I looked at the sound from a small harp. Plucking a string exactly in the center should result in a signal that has no energy in the transverse force in the even numbered partials. This is the inverse of plucking the string and touching it at the twelfth fret, which suppresses the odd numbered partials which must be moving at that point. Plucking in the center feeds no energy into the even partials, which must be stationary where you are forcing them to move. On the other hand, the tension change signal, which is pitch doubled, would contain only even number partials. Tests on the rig conformed this. Since a harp string pulls upward on the soundboard, as well as pushing it sideways with the transverse force, plucking a harp string in the center should yield a sound that contains all of the partials. In fact this is the case. The decay rates of the odd- and even-order partials are different, which suggests that they are coupled differently to the soundboard. The same experiment tried on a guitar with the neck angled so that the strings pulled upward on the soundboard at a five degree angle did not show this effect to any great degree. Apparently the angle is not great enough.

One of the main differences that is cited between archtop and flat top guitars is that the archtop bridge cannot transmit the torque of the tension change signal to the top effectively. However, it has recently been suggested that violins might nonetheless benefit from the tension change, since the back angle of the neck causes tension changes to push downward on the top. Given the outcome of the experiment cited in the last paragraph, it was felt that this was unlikely to be the case for an archtop guitar, but the experiment was tried anyway. Again, it seems as though there is not enough back angle to transmit the tension change to the top with any great effect. Nor is it likely we'll be able to achieve enough of an angle: harp strings usually pull upward on the soundboard at an angle of at least 30 degrees, and often more.

As you can see, this is a work in progress. It is often the case that one starts out to answer a simple question, and finds that it is in fact much more complicated than it seemed at first. Obviously more work needs to be done on saddle material. The question of the effect of break angle needs to be settled, and a close look needs to be taken at bridge mass and stiffness. And how about the issue of PBroze vs 80/20? Much of that, though, is taking us well beyond the string itself, and so I will close.