An Introduction to Sound Radiation Measurements using an Impact Hammer

By Joseph Curtin

Basics of Sound Analysis

In everyday life, the term *spectrum* is commonly associated with light. Our eyes perceive different frequencies of light in terms of color – white light being an equal mixture of all frequencies within the visible range. A prism is a spectrum analyzer in that it separates a beam of light into its color components. The analysis is reversible, since the beam can be reassembled by passing the spectrum through a second prism.

The nineteenth-century scientist Helmholtz devised a simple spectrum analyzer for sound, using a series of resonators tuned to a closely-spaced series of frequencies. By listening to a sound through each of the resonators and noting which ones responded, he could list the "frequency components" of the sound. We now know that the human ear works in a similar way: tiny hairs in the inner ear respond selectively to narrow bands of frequencies. The brain then “sharpens” this rather crude analysis to achieve the remarkable pitch discrimination of which we are at best capable.

Breaking down a sound into a number of frequency components might seem a natural thing to do with musical sounds, which typically have a pitch, and therefore at least one frequency associated with them. But what does it mean to speak of the frequency components of a brief, non-musical sound of no discernable pitch? As it turns out, *any* sound can be completely characterized by a set of frequency components.

The mathematical foundation for this kind of analysis was developed by the French engineer Jean Baptiste Fourier in the early 1800s. Each frequency component (or *Fourier component)* is a sinusoid of specified frequency, magnitude, and phase. When the original waveform is periodic, the components form a harmonic series – i.e. each is a simple multiple of the lowest. Most acoustical analysis software is based on a computer-optimized version of the original transform: the Fast Fourier Transform (FFT), which works for sampled waveforms.

The results of an FFT come in two parts: 1) the magnitude, which gives the strength of each frequency component, and is typically plotted against frequency as a spectrum, often on a decibel scale, and 2) the phase of each component. The two together form the *complex spectrum*. This is usually presented as two graphs, with phase placed directly beneath magnitude spectrum, as in Fig. 1. The complex spectrum is a complete description of the original waveform. It can be converted back into the original waveform without loss of information.

Frequency Response

Acoustical measurements are grounded in the concept of *frequency response* – or how the response of a system varies with frequency. Where audio amplifiers strive for a flat frequency response, a violin’s response curve is jagged, and this goes to the heart of its character as a musical instrument. Frequency response is typically measured by driving a system with a known input, measuring the output, and then dividing the complex spectrum of the output by that of the input. This yields a *frequency response function* (FRF), which shows the output per-unit-input across the frequency range of interest.

Plotting Frequency Response Functions

In plotting FRFs, amplitude is usually assigned to the vertical (y) axis, and frequency to the horizontal (x) axis. Both linear and logarithmic scales can be used. With a linear amplitude scale, doubling the strength of the signal doubles the height of the curve. With a logarithmic scale, doubling the intensity of the signal raises the curve by 3 decibels. This produces a more compact graph – and one that roughly mimics our subjective sense of loudness. Frequency can also be plotted linearly or logarithmically. With a linear scale, each kilohertz is given equal space. With a logarithmic scale, each octave is given equal space – just as it is on the piano keyboard.

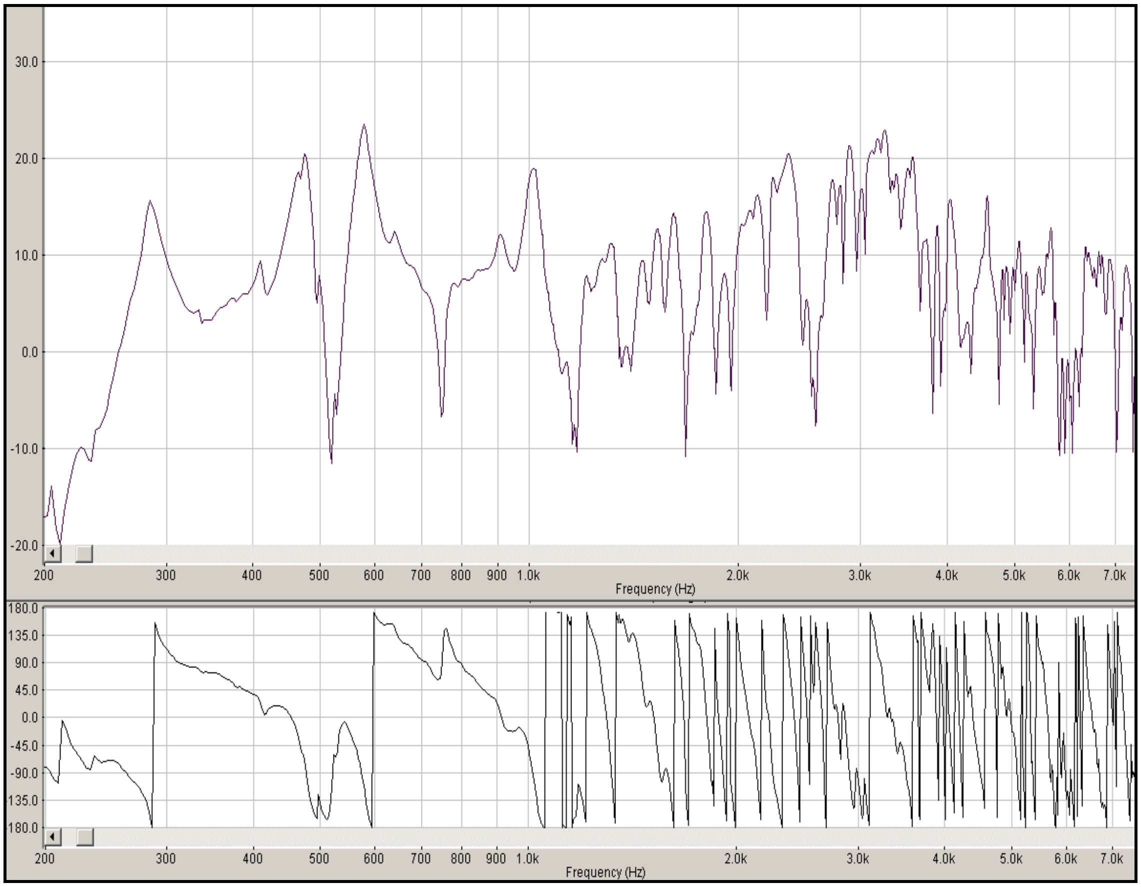

Figure 1) Spectrum (above) and phase (below) for the sound output of a typical violin, driven by an impulse at the bridge. Frequency is plotted here logarithmically, with the amplitude in decibels (dB). Phase is plotted in degrees, from plus to minus 180°. Note that the vertical discontinuities in the phase plot, while suggesting sudden jumps in phase, are in fact artifacts of the scaling: a jump from +180° to -179°, for example, represents a shift of just one degree.

The jaggedness of a violin's response curves makes it difficult to assess how much energy is concentrated in a given frequency band. To facilitate this, the spectrum can be broken down into a number of bands, and the average energy in each band calculated – a process known as *band averaging.*

Acousticians have long employed third-octave bands to study everything from speech defects to room acoustics. While the human ear is capable of the subtlest discrimination in the pitch of a note, third-octave bands roughly capture the resolution with which the ear perceives tone color. Third-octave bands can be plotted as bar-graphs, where the height of each bar represents the average amplitude of the signal in that band. Though much detail is lost, they give a good sense of the general distribution of energy.

There are many other ways to create band averages, including octave and 1/6th octave bands. George Bissinger, in his VIOCADEUS project, uses 250 Hz band-averages across the violin's entire range, except for the lowest, which contains A0, and was averaged over 10 Hz to either side of the peak, which is typically around 280 Hz).

The method that best reflects our current understanding of the hearing process is the *Bark Scale*, which divides the audible range into bands demarcated by the following frequencies, in Hertz: 0 – 100 – 200 – 300 – 400 – 510 –630 – 770 – 920 – 1080 – 1270 – 1480 – 1,720 – 2,000 – 2,320 – 2,700 – 3,150 – 3,700 –4,400 – 5,300 – 6,400 – 7,700 – 9,500 – 12,000 – 15,500. These bandwidths are based on psycho-acoustic studies of human hearing, and are intended as a refinement of the 1/3rd octave approach. Note that it would be a mistake to assume that two instruments with identical Bark scale profiles necessarily sound the same. Consider, for example, that one instrument might have a much less peaky response curve than another, and yet have the same amount of energy in each frequency band.

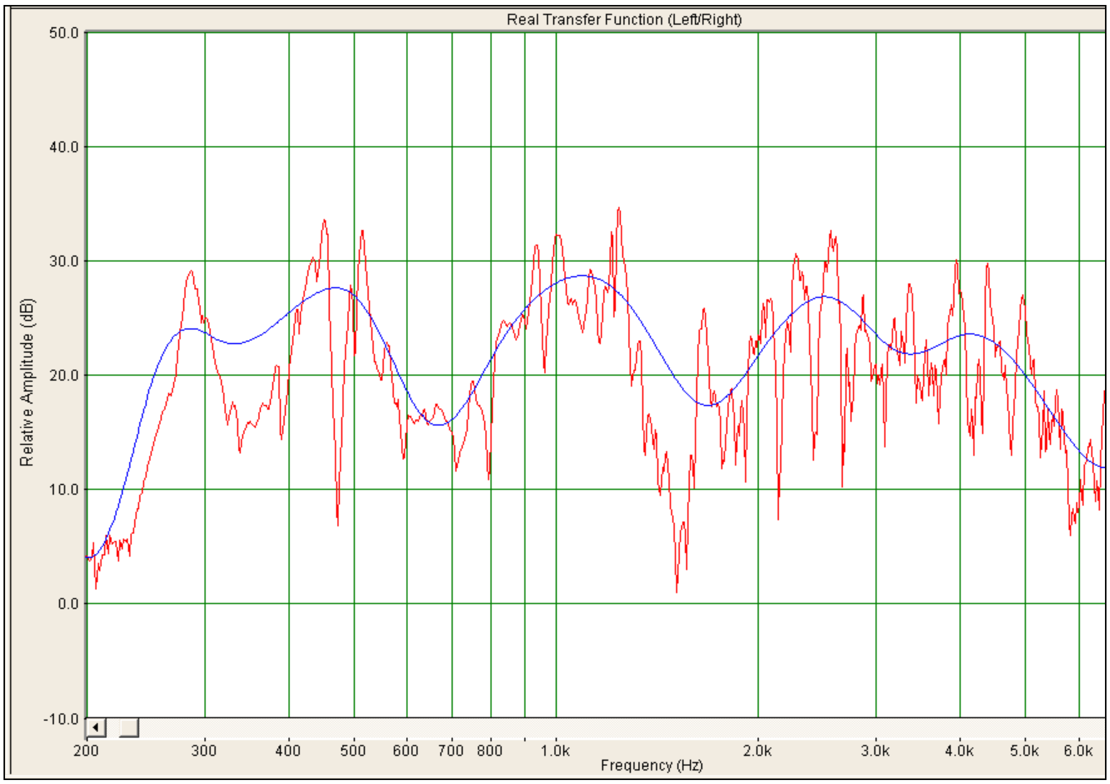

Figure 2) FRF for the "Plowden," plotted logarithmically (red) and in 1/3rd octave bands

Fig. 2 shows an FRF for the "Plowden" Guarneri del Gesu (measured in a relatively small, noisy room). The red line shows amplitude against frequency, plotted logarithmically. The blue line shows the same measurement in smoothed 1/3rd octave bands.

Measuring Sound Radiation

To measure a violin's sound output effectively, we must do at least four things:

1) Excite the instrument across its frequency range
2) Pick up the resulting sound at one or more microphone positions

3) Compare the sound output with the excitation force
4) Deal with the acoustics of the room in which the measurements were made.

Looking first at excitation, under normal playing conditions, the body of the instrument is driven by a fluctuating force exerted on the bridge by the bowed string. For measurement purposes, the bridge is typically driven by a transducer. Everything from bowing machines to coil & magnet drivers have been used. Ideally, a driver should have the following characteristics:

1) It should be capable of getting a substantial amount of sound from the instrument. This ensures a good signal-to-noise ratio.
2) It must not couple significant amounts of mass to the bridge. So sensitive is the bridge to mass-loading that as little as .2 grams will produce an audible and measureable change in the instrument's sound.
3) It should contain a sensor that monitors the force being exerted. This provides the input signal (by which the output signal will be divided).
4) It should excite the instrument more or less evenly across its frequency range. A fair bit of unevenness can be dealt with in the analysis process, but only if the minimum excitation level stays well above the noise floor.

Bridge-drivers are typically fed with computer-generated signals. These can take the form of a sinewave swept across the frequency range in a kind of continuous glissando. Alternatively, a mixture of all frequencies, such as white noise, can be used. A third approach, known as impulse excitation, delivers all frequencies at once in the form of a sudden blow. This is most conveniently done with an impact hammer (also know as a *force hammer*), which for our purposes has the following advantages:

1) An impact takes only a fraction of a second, and so measurement cycles are faster than with most other methods.
2) Mass-loading is negligible.
3) A calibrated force-sensor is built into the hammer head.
4) An impact hammer can be for a wide variety of vibration measurements.
5) It is a readily-available, off-the-shelf tool.

On the other hand, an impact hammer has the following disadvantages:

1) It is fairly expensive.
2) Not much energy is delivered to the bridge, so care must be taken to achieve a good signal-to-noise ratio
3) Double-bounces can be a problem (more on this below).

It happens that an impact of infinitely short duration contains equal amounts of energy at all frequencies. In order to produce infinitely short impact, however, both hammer and bridge would need to be infinitely hard. The reason for this becomes clear if we imagine the converse – a hammer tip in the form of a soft spring. Upon impact, the spring would compress and rebound only gradually, thus prolonging the contact between hammer and bridge. The result would be a reduction of high-frequency content in the excitation force, and therefore in the sound produced. To demonstrate this, tap a tabletop – first with a fingertip, and then with a house key. The brighter click made by the key indicates the greater high-frequency content.

If all of the above suggests that a very hard hammer tip is ideal for measuring violins (whose frequency range goes to 10 kHz and beyond), one problem with hard tips is the possibility of damaging the bridge (the metal tip of the smallest PCB impact hammer leaves a visible dent). A second problem is the increased possibility of double-bounces.

A double-bounce occurs when the bridge deflects in response to an impact, but then rebounds toward the retreating hammer and catches up with it, producing a second collision. Because this happens within milliseconds, double-bounces are difficult to hear. They can be seen by looking at the hammer signal in the *time domain* (i.e. as amplitude against time): The initial peak falls to zero, indicating loss of contact between hammer and test object, but then rises again to form a small secondary peak (see Fig. 3).

Double-bounces are treated with great suspicion by many researchers, and so it is useful to understand the circumstances under which they can be a problem. If the hammer signal is plotted in the frequency domain, a double-bounce shows up as an oscillation in the response curve, as in Fig.3 below. If the second bounce were as strong as the first, the oscillation would swing down to zero. The smaller the second bounce compared with the first, the smaller the oscillation. Providing that even the lowest dip is well above the noise floor in the frequency range of interest, there is no problem. If, however, a dip drops to the level of the noise floor, then the results at that frequency will be unreliable. There is only one other case where a double bounce is a problem, and that is when the second bounce and/or the instrument's response fall outside of the time window of the chosen FFT. Otherwise, a double or even triple bounce will yield exactly the same results as a single bounce.

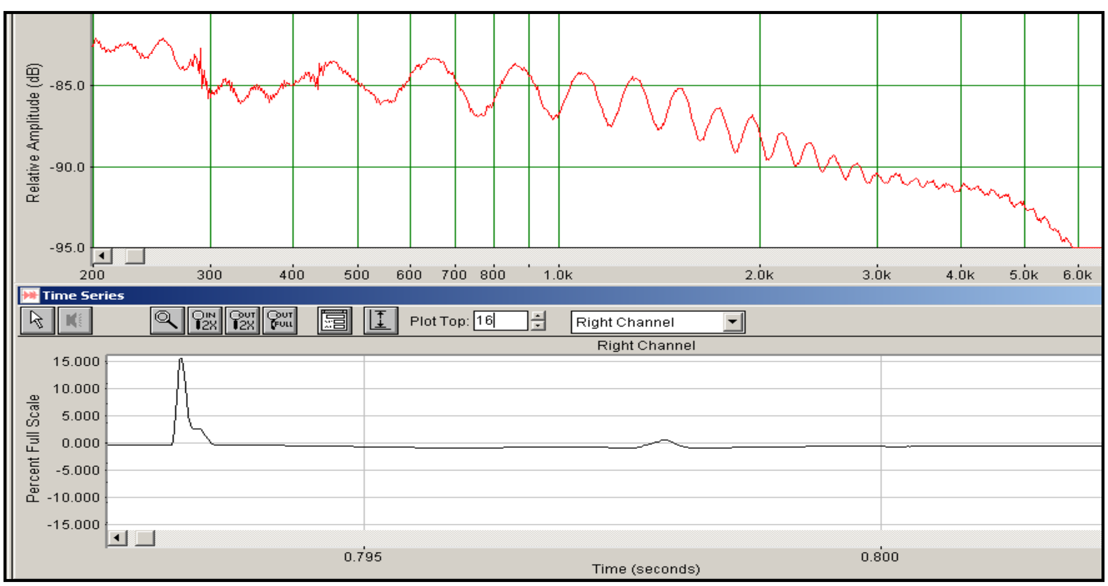

Figure 3) The lower (time domain) window shows the hammer signal from a double-bounce. The upper (frequency domain) window shows the resulting oscillation in the hammer spectrum.

Double-bounces can be avoided by using a softer hammer tip, which prolongs the contact time, so that the secondary impact becomes, in effect, part of the initial one. As can be seen in Fig. 4, the softer the tip, the less energy is delivered at high frequencies. This progressively lowers the signal-to-noise ratio, effectively limiting the highest frequencies that can be measured.

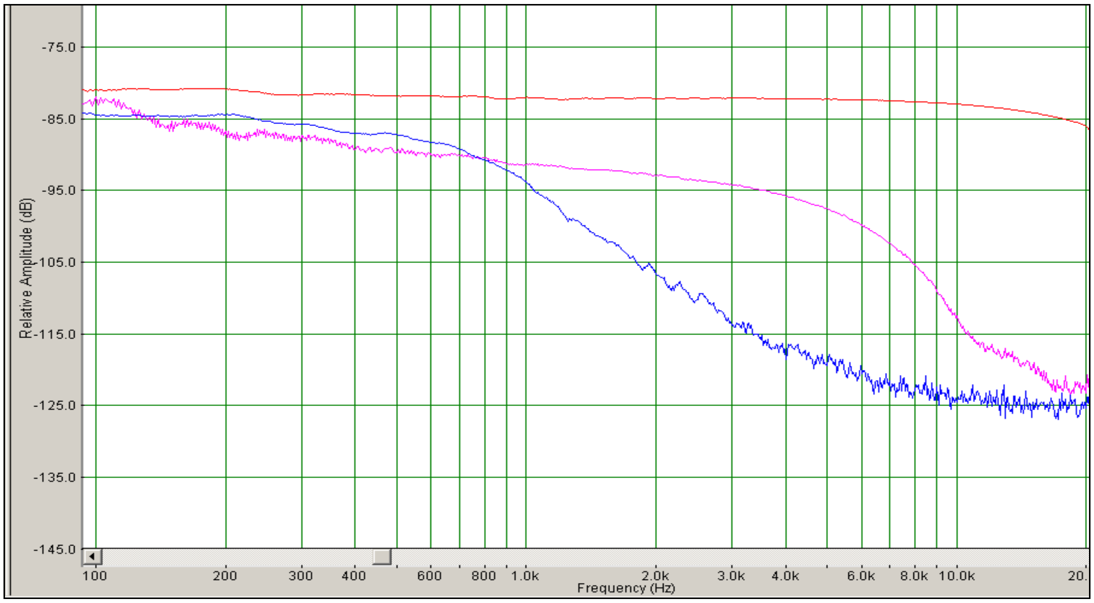

Figure 4) Spectrum of hammer blow for three different tips. Red – PCB metal tip; Magenta – custom polymer tip; Blue – PCB soft (red) tip.

If we tap the side of a violin's bridge, the sound produced will contain contributions from all of the instruments' resonances – or at least those which:

1) radiate sound
2) are within the hammer's frequency range
3) can be excited by a side-to-side motion of the bridge.

The radiated sound is picked up by a microphone, whose signal is fed to one channel of a stereo soundcard. The hammer signal is fed to the other channel, and the software then calculates the transfer function. The beauty of this process is that, within certain practical limits, hammer blows of different strengths will all give the same results: an FRF representing the sound-pressure at the microphone position, per-unit-force at the bridge.

Pictured below is a rig designed to measure violin sound radiation in workshop or laboratory settings. The hammer taps the bass corner of the bridge, while a sensor in the hammer-tip records the force-history of the blow. The microphone, which can be moved 360º around the instrument, is held in the plane of the bridge at a chosen distance from the central axis of the instrument – here defined as a line rising vertically through the end-pin. To increase the signal-to-noise ratio, the complex average of several impacts is taken for each microphone position.

The Sound Radiation Rig – a tool for measuring the sound radiation of violins and violas. For further information, see www.josephcurtinstudios.com.

The i*nstrument is supported by elastic bands on either side of the endpin, and then about 2/3rds of the way up the neck.*

Many possible radiation measurements can be made for a given violin. The bridge can be excited at different driving points and in different directions – and each will excite a different balance of body modes. Because of increasingly directional radiation above about 1 kHz, shifting the microphone position can significantly change the measurements. Furthermore, the response of the instrument itself may change with temperature, humidity. All this makes it virtually impossible to measure violin sound radiation in a comprehensive way. Instead, researchers rely on a number of simplifications. The bridge is generally excited with a lateral force, while a limited number of microphone positions are used to sample the sound field. Martin Schleske uses 12 positions, spaced evenly around the instrument in the plane of the bridge; Bissinger uses 266 arranged spherically around the instrument.

The measurements cited below are all for a single, good-quality, modern violin whose strings have been damped. Measurements were done in an anechoic chamber to avoid problems with reflected sound. Twelve microphone positions were used – one directly in front of the bridge, the others spaced at 30° intervals around the instrument. The microphone is 37 cm from the center of the instrument.

Figure 5 shows an overlay of 5 microphone positions – one directly in front, the others at 30° and 60° to either side. Though there is good correspondence at low frequencies, the spread above about 1 kHz indicates radiation patterns that change significantly with frequency.

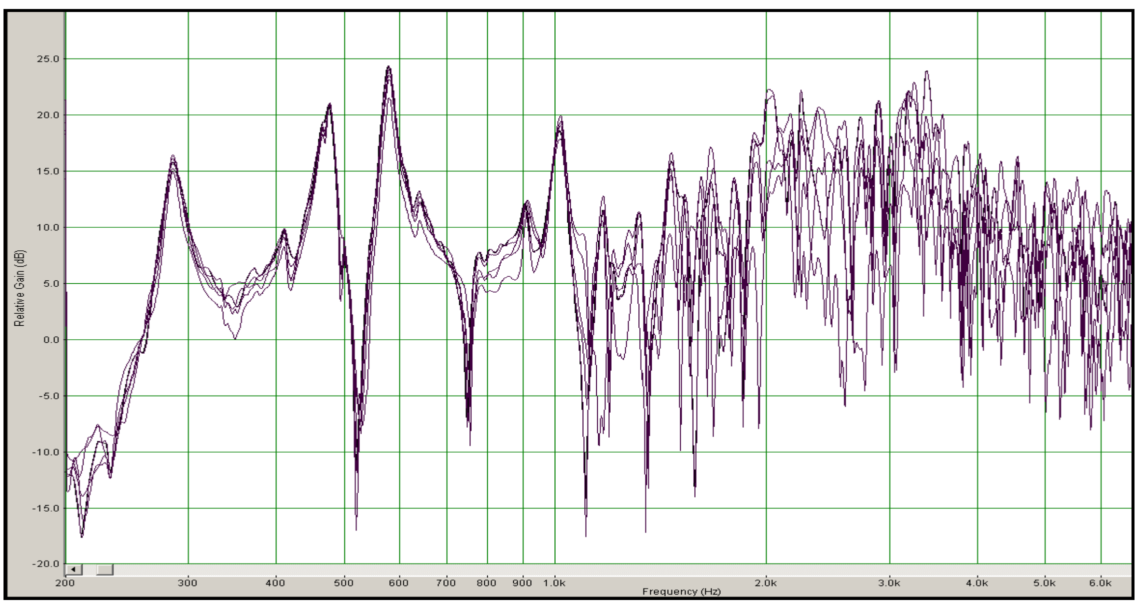

Figure 5) Overlay of 5 front microphone positions

There are many ways of deriving single response curves from multiple readings. The first graph in Fig. 6 shows the average magnitude for all 12 microphone positions, while the second graph shows a single, front-central position. Which curve better represents the sound of the instrument?

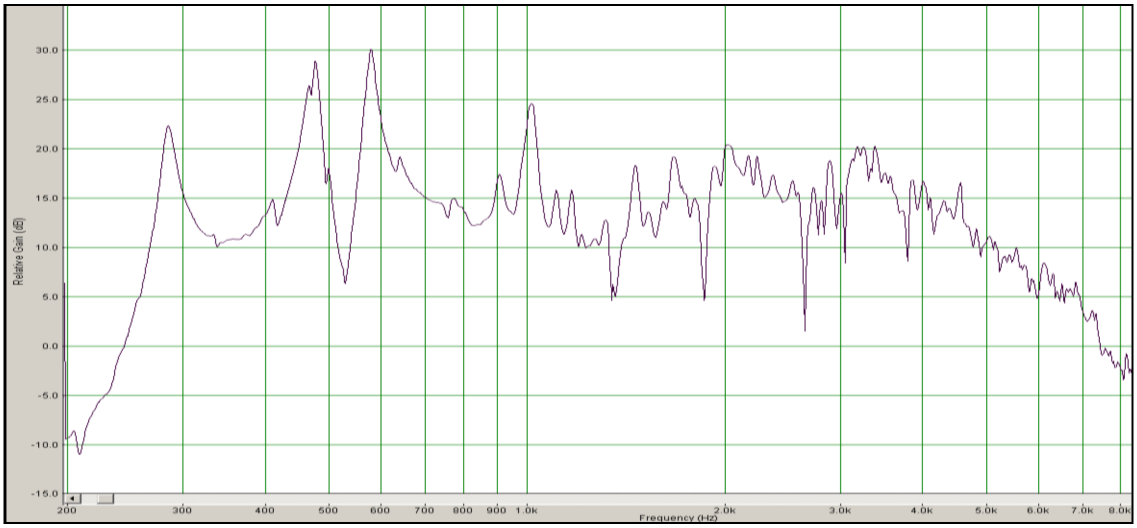

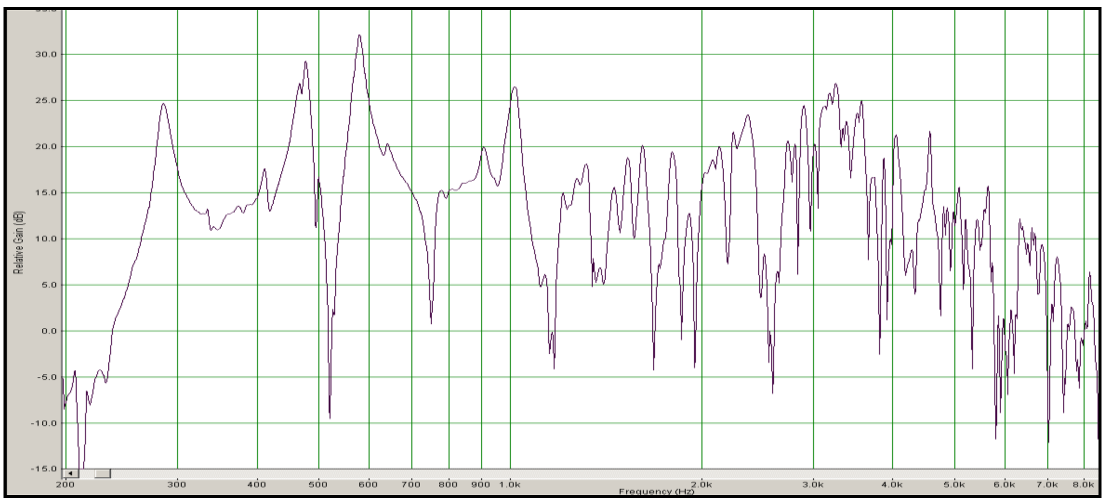

Figure 6) The top graph shows the average magnitude for 12 microphone positions. The graph below shows the magnitude for a single, front-center mic position.

While the two graphs have a very similar overall shape, note that above 1 kHz, the peak-to-valley heights are far greater for the single microphone measurement – and this "spikiness" is important for sound quality. The averaged measurement provides an estimate of the total radiated sound – assuming that radiation in the plane of the bridge is representative. The single-microphone measurement represents what might be heard by listening with one ear held 37cm from the violin in an anechoic chamber. This is further away than the violinist's ear, but much closer than a typical listener's – though neither player nor listener is likely to be found in an anechoic chamber! Without further laboring the point, there is no single radiation measurement which fully captures the sound of a violin. The meaningfulness of any particular measurement depends on how well the measurement conditions are specified. The usefulness depends on how much light is shed on the question at hand.

Note: an expanded version of this article will appear in VSA Papers