

# Measuring Violin Sound Radiation Using an Impact Hammer

JOSEPH CURTIN

3493 West Delhi, Ann Arbor, MI 48103  
violins@josephcurtinstudios.com

## Abstract

*An impact-hammer-based system for measuring violin and viola sound spectra in a workshop setting is described, along with an overview of some of the practical and theoretical problems involved. In this system, the violin is held vertically while a miniature, pendulum-mounted impact hammer taps the bass side of the bridge, and a sensor in the hammer head records the force of the blow. The resulting sounds are detected by a microphone. Signals from microphone and hammer are fed through an audio interface to a computer with the requisite sound analysis software. The goal is to assess the sound output at one or more microphone positions, per-unit-force at the bridge—a measure of the instrument's radiativity.*

Violinmakers have long relied on rulers, calipers, and gauges to measure the dimensions of their instruments, but until recently the only widely available tool related to violin sound was the tuning fork. Although the mathematical basis for analyzing sound has been understood since the early 1800s, it was the development of digital computers that made the extraordinarily burdensome calculations feasible and eventually trivial. The current availability of sound recording and measurement equipment makes it possible for a violinmaker to assemble a measurement system for about the price of a good new violin bow.

The system presented here is loosely based on one developed by Martin Schleske and described in his paper “Empirical tools in contemporary violin making” [1]. It was further developed for convenience and portability and as a possible standardized approach to making radiation measurements in a workshop setting. Using it in a meaningful way requires a working familiarity with the relevant hardware and software and with the basics of musical acoustics and signal analysis. This article attempts to bring together enough information to give violinmakers a sense of whether a measurement system would be useful in their own work, and if so, how to put together the necessary equipment and then begin using it.

## MEASURING VIOLIN SOUND

Measuring a violin's sound output in a plausible way involves:

- Mounting the instrument;
- Exciting the instrument across its frequency range;
- Recording the resulting sounds;
- Comparing these sounds with the excitation forces in order to calculate an impulse response or a frequency response function;
- Accounting for the violin's increasingly directional radiation patterns for frequencies >850 Hz;
- Dealing with random noise;
- Accounting for the acoustics of the measurement room;
- Calibrating the equipment.

Except for calibration, each of the above is addressed in this article. Also, some of the background concepts, including frequency response and Fourier analysis, are discussed.

## A MEASUREMENT RIG

Figure 1 shows a rig designed for measuring both violins and violas. The distance between the microphone and the central axis of the instrument (here defined as a vertical line rising through the endpin and parallel with the plane of the top plate) can be varied from 5 to 55 cm. The



Figure 1. Sound Radiation Measurement Rig, designed by Joseph Curtin and furniture designer Garry Venable. For more information, see <[www.josephcurtinstudios.com](http://www.josephcurtinstudios.com)>.

instrument can be rotated 360° with respect to the microphone, and the microphone and instrument can together be rotated 360° with respect to the room. A 3-axis positioning stage enables precise positioning of the hammer in relation to the bridge. A photographer's cable release allows triggering of the hammer (Fig. 2) with one hand while operating a computer with the other. The rig is mounted on a speaker stand or camera tripod, using a standard 1/4" or 3/8" stud. To facilitate measurements in concert halls and other places where good instruments are likely to be found, the entire apparatus can be dismantled (Fig. 3) and fitted into a standard carry-on bag. Appendix A lists the associated hardware.

Much the same measurements can be made using a hand-held hammer, an ordinary microphone stand, and a simple frame from which to suspend the instrument. The advantages of a specialized rig are mainly speed and convenience. In my experience, if an instrument cannot be mounted and measured within minutes, it will not get measured at all. Such a rig also offers the

possibility of a standardized approach, allowing makers and researchers using similar equipment to more meaningfully compare results.

## MOUNTING THE INSTRUMENT

Measuring a violin's radiation requires a mounting system that is stable, holds the instrument securely, and yet modifies the instrument's dynamic behavior as little as possible. When a player holds an instrument, the contact inevitably has its effects—principally by adding *damping*. Rather than trying to reproduce this situation in a measurement system, violins are typically mounted with as little outside contact as possible—as if floating freely in space. One way of implementing this is to suspend the instrument from fine elastic bands. Longer, looser bands provide better isolation, but the instrument then tends to bob and weave, making it difficult to fix its position relative to the test equipment. A reasonable compromise is shown in Fig. 4. The instrument is supported from underneath by elastic bands (the kind commonly sold for tying ponytails) on either side of the endpin, and then about two-thirds of the way up the neck.

## DAMPING THE STRINGS

The strings can be damped or undamped during measurement—each provides somewhat different information. Undamped strings show up as a series of sharp peaks and adjoining valleys in the instrument's response curve, and these can obscure the peaks associated with the violin body. Damping the strings will to a large extent




Figure 2. Impact hammer mounting.



Figure 3. Components of disassembled test rig.

remove the string features, but will also add a measure of damping to the whole instrument.

The strings can be damped with a business card or with bits of foam (Fig. 5). I have found that the exact position of the card or foam along the string length can affect the measurements to a surprising degree—particularly the B1 peaks—so it is often worthwhile to try a couple of positions. The string length between bridge and tailpiece can also be damped.



## DRIVING THE BRIDGE

Systems for measuring violin sound differ dramatically in how the instrument is set into vibration, and everything from violinists to electromagnetic bridge drivers have been employed. Each system has its advantages and disadvantages. Violinists, for example, are ideal in terms of the naturalness of their approach, but they fall down badly in terms of consistency; furthermore, they tend to alter their playing to suit the particular instrument being tested. This would not be a problem if there were a sensor capable of measuring the force of the vibrating strings against the bridge while the instrument was being played—and without mass-loading the bridge or otherwise disturbing its normal function. As it is, some sort of bridge-driver is typically used. Ideally, it should:

- Be capable of getting a substantial amount of sound from the instrument, ensuring a good signal-to-noise ratio;
- Contain a sensor that accurately monitors the changing force exerted at the bridge;
- Not couple significant amounts of mass to the



Figure 4. The instrument is supported at either side of the endpin and neck by elastic bands.



Figure 5. Damping the strings with a business card and bits of foam.

bridge. Mass-loading mutes the sound and therefore skews the results. So sensitive is the bridge that as little as 0.1 g attached to the top of the bridge will create an audible and measurable difference;

- Excite the instrument more-or-less evenly across its frequency range. A fair bit of unevenness can be accommodated by the analysis process, but only if the minimum excitation level stays well above the level of ambient noise.

Electromagnetic or mechanical 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 known as a *force hammer* or *instrumentation hammer*), which differs from a carpenter's hammer mainly by virtue of a piezo force sensor embedded in its head. Impact hammers weighing anywhere from a few grams to several kilograms have been used for testing everything from printed circuit boards to airplanes. A hammer that is well suited to violin research is the smallest model made by PCB Piezotronics (Model 086C80, pictured in Fig. 2).

For our purposes, impact hammers have a number of advantages: They are off-the-shelf tools with long working lives and can be used for many other measurements, including bridge admittance and modal analysis. Mass loading is negligible. An impact lasts only a fraction of a second, and so measurement cycles tend to be faster than with other methods. A disadvantage is that relatively little energy is delivered with each blow, so care must be taken to achieve a good signal-to-noise ratio.

## MICROPHONES

A microphone is a transducer that converts rapid fluctuations in air pressure into an electrical signal. Most sound measurement systems rely on one or more of them—an exception being Weinreich's reciprocal method [2]. For our purposes, a relatively flat response across the violin's frequency range is important. While this is not difficult to find in a microphone today, some models are designed to add “presence” or “warmth” to the human voice, or to otherwise color the sound, and these should be avoided. There are many inexpensive microphones that are quite flat enough for good violin measurements.

Most good quality microphones—or at least those labeled “condenser”—require phantom



power. This is usually provided by the pre-amp or audio interface (see Appendix A). Calibrated microphones come with individually measured response curves, which were useful when making measurements in absolute rather than relative terms.

Microphones are available with a variety of directional properties. Omnidirectional models pick up sound equally well in all directions, while “shotgun” microphones are highly directional. Omnis are usually preferred for measurement purposes since they have the flattest frequency response. But directional microphones—in particular, cardioid models—are useful in relatively noisy conditions. Cardioids have two trade-offs: their off-axis frequency response is not flat, and their bass response tends to fall off at a distance. I have found that as long as the microphone is pointed directly at the violin, the off-axis response is not a significant problem, at least with the cardioid microphone I use (see Appendix A). With this microphone, the bass is flat at ~15 cm, but by 1 m it is down 5.5 dB at 100 Hz. Since I use a microphone distance between 20 and 37 cm, and since violins and violas radiate very little below ~200 Hz, this does not present a problem. Figure 6 compares measurements done with cardioid and omni microphones (both made by Earthworks) at 37 cm. Note the somewhat smoother line with the cardioid, which picks up fewer room reflections.

## DIGITAL RECORDING

Almost all sound recording and processing is now done digitally. The signal from a microphone or other source is converted into a series of samples, each of which represents the amplitude of the signal at a particular instant in time. The fidelity of the sampled signal to the original depends largely on two factors: *sampling rate* and *sampling depth* (also known as *sampling precision*).

*Sampling rate* is the number of samples taken per second. The standard rate for commercial audio CDs, for example, is 44.1 kHz. The sampling rate places a limit on the highest frequency that can be captured. This high-frequency limit, known as the Nyquist frequency, is exactly half the sampling rate (e.g., 22.05 kHz for CDs). In order to prevent the introduction of artifacts

into the digitized signal, all frequencies above the limit must be filtered out before the signal is digitized. Because analog filters have a more-or-less gradual cut-off rate, the actual high-frequency limit is somewhat lower than the Nyquist frequency. Most sound cards automatically apply the appropriate filtering when the sampling rate is specified. Good sound cards typically support a number of sampling rates, including 44.1 kHz, 48 kHz, 96 kHz, and even higher.

*Sampling depth* refers to the number of digits assigned to each sample. This in turn determines the number of discrete amplitude levels that can be represented. Because computers work with binary numbers, sampling depth is specified by the number of *bits* (binary digits) used—and this increases by powers of two. Four bits allow 16 discrete amplitude levels; eight bits allow 256, and 16 bits—the standard for music CDs—allow 65,536. Sampling depth therefore determines the dynamic range, and this is especially important in capturing a gradual diminuendo. Imagine a low note struck and held on a piano. Using four-bit sampling, the note would be rendered as a bumpy, 16-step journey into silence. Much professional audio recording is now done at 24 bits, allowing 16,777,216 levels. Expressed as dynamic range, this is nearly 17 million to one. Although adequate for almost any conceivable music application, consider that the ear has a dynamic range (from the threshold of hearing to the threshold of pain) of something like 10 trillion to one!

Practically speaking, the impact hammer method described in this article is not likely to yield useful data above 10,000 Hz, and so a 44.1 kHz sampling rate is more than adequate. For the typical signal-to-noise ratios encountered, using anything beyond 16-bit recording is arguably only digitizing noise. That said, 48 kHz sampling at 24 bits is now the norm for pro-audio equipment. Although the increased sampling rate and depth create larger files, that is hardly a problem for violin measurements, given the enormous storage capacity now found on even the smallest computers.

## SOUND ANALYSIS

In everyday life, the term *spectrum* is commonly associated with light. Our eyes perceive differ-

ent 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 19th-century scientist Hermann 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.

Analyzing a musical sound in terms of its frequency components might seem a fairly natural thing to do, since a musical sound usually has pitch, and therefore at least one frequency associated with it. But what does it mean to speak of the frequency components of a brief, non-musical sound of no discernable pitch—for example, the resonant thud made by tapping a violin bridge with a tiny hammer? Before answering this, it is useful to review some terminology.

A *waveform* is a graphic or numerical representation of a wave. By extension, anything that varies over time in a quantifiable way—from barometric pressure to stock market values—can be plotted as a waveform.

A *periodic* waveform is one that repeats itself at regular intervals. The waveform can be of any finite length.

A *period* is the time required for a single iteration of a periodic waveform. Period is the reciprocal of frequency—thus the period of concert A is 1/440 of a second.

As it turns out, any sound—or more precisely, any *waveform*—can be completely characterized by a set of frequency components. The mathematical foundation for doing this was developed in the early 1800s by the French engineer Jean Baptiste Fourier. In his honor, the frequency components are often referred to as Fourier components and the whole process as Fourier analysis. Fourier analysis allows any

waveform to be converted into a series of frequency components. Furthermore:

- Each frequency component is a sinusoid, and can be completely described by its frequency, amplitude, and phase;
- If the waveform is periodic, the frequency components form a harmonic series, meaning the frequency of each is a simple multiple of the lowest—e.g., 100 Hz, 200 Hz, 300 Hz, and so on;
- The frequency of the lowest, or fundamental, is equal to the reciprocal of the period of the waveform. Thus a waveform that repeats itself every 10 seconds has a fundamental of 0.1 Hz;
- The shape of the waveform determines the amplitude and phase of the components, while the length of the waveform determines the frequencies of the components;
- The components extend in frequency from the lowest (fundamental) on up to infinity. In practice, only components within a useful frequency range are calculated;
- The analysis is reversible, meaning that the original waveform can be reconstructed without loss of information.

## THE FAST FOURIER TRANSFORM (FFT)

Most sound analysis software relies on the Fast Fourier Transform, or FFT—a computer-optimized transform that works only for sampled waveforms, with the further restriction that the number of samples analyzed in a single calculation must be a power of two. The number of samples is known as the FFT size, and the seemingly arbitrary sizes offered by sound analysis software (512, 1024, 2048, etc.) are all in fact powers of two. If the waveform being analyzed is shorter than the FFT size, a string of zeros must be added to make up the difference.

Waveforms to be analyzed using the FFT algorithm must be periodic. This is not a particularly stringent limit, since *any* waveform of finite length can be considered a single iteration of a periodic waveform. For example, a performance of an opera could be recorded and endlessly repeated. However, if a waveform does not end and begin at zero, there will be a discontinuity in the form of a vertical line created when beginning

and end are joined to form a periodic waveform. This vertical line forms a transient, which produces its own set of frequency components in the analysis. To avoid this situation, the two ends of a waveform can be brought to zero by means of a “smoothing window.” Smoothing windows such as the Bartlett, Blackman, and Hamming are useful when analyzing segments of an extended waveform—but at the cost of introducing some artifacts into the results. Fortunately, impulse measurements begin and end at something very close to zero (depending on noise levels), and so do not require smoothing windows. (Note that in a software menu, choosing a “uniform” or “rectangular” window is equivalent to choosing no window.)

If we think of the frequency components as pixels that together form an image of the violin’s sound output, the more pixels there are, the more clearly will the details of the curve be defined. A small FFT size means widely spaced frequency components. For example, with a 48 kHz sampling rate, an FFT size of 2048 has frequency components every 23.438 Hz, and thus a “frequency resolution” of 23.438 Hz. This is clearly insufficient to accurately delineate a peak at 280 Hz (for example). On the other hand, an FFT size of 32,768 has a frequency resolution of 1.465 Hz, and so could define the peak much more precisely.

The results of an FFT come in two parts: magnitude and phase. The magnitude, which gives the strength of each frequency component, is typically plotted against frequency as a spectrum, using either a linear or decibel scale for magnitude. Spectrum and phase together form the *complex spectrum*, which is a complete description of the original waveform. An inverse FFT allows the complex spectrum to be transformed back into the original waveform without loss of information (other than very small errors due to the rounding of decimal places). The spectrum and phase are usually presented as two graphs, with phase placed directly beneath magnitude spectrum, as in Fig. 6.

### CHOOSING FFT SIZE

The usual strategy in analyzing an impact hammer measurement is to take the signals associated with a single impact and analyze them with a single FFT. For this to work, it is crucial that the total number of samples in the measurement signal (beginning with the silence just before the hammer blow, and ending after the violin’s response has died away) is smaller than the FFT size. The FFT size and the sampling rate together define a “time window.” Calculated by dividing the FFT size by the sampling rate, the time window is simply the length of time spanned




Figure 6. Cardioid (red curve) and omni (blue curve) microphones at 37 cm from the central axis of a typical violin, directly in front of the bridge. Note the somewhat smoother curve with the cardioid, which detects fewer room reflections.

by a single FFT at the chosen sampling rate. At 48 kHz, for example, an FFT size of 16,384 has a time window of ~0.341 second. This is quite long enough for a violin with damped strings. With strings undamped, a larger FFT size must be chosen.

The FFT size is typically selected from a software menu. There is no problem with choosing larger sizes, except that calculation times are longer. An FFT size of 32,768 gives a time window of ~0.7 seconds and a frequency resolution of ~1.5 Hz. An FFT size of 1,048,576 gives a frequency resolution of ~0.05 Hz and a time window of nearly 22 seconds.

## FREQUENCY RESPONSE

Frequency response is a general term for how a system's response varies with frequency. Audio amplifiers strive for the flattest possible frequency response. The violin, by contrast, has a very jagged one—and this goes to the heart of its character as a musical instrument. Frequency response is measured by dividing the output of a system by its input. This gives a record of the output per-unit-input across the frequency range of interest.

While the input and output of an amplifier are clearly defined, the situation is more complicated for a violin. The bridge is the obvious input for the “signal” from the bowed string. And yet each string notch is at a slightly different location. Furthermore, the string applies forces in three directions: mainly laterally (i.e., side to side in the plane of the bridge), but also vertically, and to some degree longitudinally, i.e., pulling the bridge toward and away from the fingerboard. Since each of the four string notches has these three degrees of freedom, a complete monitoring of the string signal would require 12 separate channels. In practice, a single, lateral excitation force is typically applied, as this is the best single-channel approximation of the direction of the string forces. That said, including vertical and perhaps out-of-plane excitation in a measurement cycle would undoubtedly lead to a more complete characterization of the instrument's radiativity.

A violin's output is even more complicated than its input. Though the instrument radiates more-or-less omnidirectionally (i.e., with equal

strength in all directions) at low frequencies, above ~850 Hz radiation becomes increasingly directional, so that at high frequencies the sound field consists of numerous “quills” or “beams” of sound, whose directions can change radically with small changes in frequency. (For more on this, see Weinreich [3].) Clearly, a large number of microphone positions would be needed for a detailed record of the sound field.

All of the above makes it virtually impossible to chart a violin's frequency response in a comprehensive way. Instead, researchers rely on more specific and limited measurements known as *frequency response functions*, or FRFs. An FRF is a response curve whose input and output conditions are carefully described. For our purposes, an FRF typically represents the sound pressure at a particular microphone position in response to a lateral force at the bass corner of the bridge. An FRF can be thought of as a snapshot—as one particular view of a three-dimensional landscape. By taking measurements at a number of different microphone positions, a more-or-less detailed picture of the total sound field can be developed.

## TRANSFER FUNCTION

Assuming a satisfactory hammer and microphone signal have been acquired, the two signals must be processed in order to yield a single FRF. This is done by means of a *transfer function*. When the software divides the complex spectrum of the microphone signal by the complex spectrum of the hammer signal, the result is a *complex transfer function*, which contains both the magnitude and phase of the frequency components. (Note that for presentation purposes, the phase plot is often omitted and the spectrum shown on its own.) When the phase information is left out of the transfer function calculation, the result is a *real transfer function*, which can differ markedly from its complex counterpart due to cancellations resulting from phase differences.

Many free or inexpensive sound analysis programs will perform FFTs on a single channel. A transfer function requires software capable of handling at least two channels simultaneously. A menu typically gives a choice of which channel is to be divided by which. Thus if the microphone goes into the left channel and the hammer into



the right, a left-over-right transfer function must be selected. The beauty of a transfer function is that, within certain practical limits, hammer blows of different strengths and spectral content will all produce the same results.

## PLOTTING FREQUENCY RESPONSE FUNCTIONS

In plotting FRFs, amplitude is usually assigned to the vertical (y) axis, and frequency to the horizontal (x) axis. Either linear or logarithmic scales can be used. With a linear amplitude scale, doubling the amplitude of the signal doubles the height of the curve. With a logarithmic scale, doubling the amplitude raises the curve by 6 dB. 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 7 presents the sound spectrum (amplitude and phase) of a typical violin.

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. 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—or as a smooth curve, as in Fig. 8. Though much detail is lost, a good sense of the general distribution of energy is obtained.

The band-averaging method that best reflects




Figure 7. The spectrum (amplitude and phase) of the sound output of a typical violin, measured in an anechoic chamber. The units of the amplitude (upper) and phase (lower) graphs are dB and degrees. An impact hammer tapped the bass corner of the bridge. The microphone was positioned 37 cm from the violin's central axis. The vertical discontinuities in the phase plot, while suggesting sudden jumps in phase, are in fact artifacts of the scaling. For example, a jump from to  $-179^\circ$  to  $+180^\circ$  represents a shift of just  $1^\circ$ .




Figure 8. The smooth line represents 1/3-octave band averaging of the jagged spectral curve of a typical violin. The band average gives a better sense of the amount of energy in each frequency band.

our current understanding of the hearing process is the *Bark Scale*, which divides the audible range into the following bands (in Hertz): 0 – 100 – 200 – 300 – 400 – 510 – 630 – 770 – 920 – 1,080 – 1,270 – 1,480 – 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 psychoacoustic studies of human hearing and are intended as a refinement of the 1/3-octave approach.

There are many other ways to create band averages, including one-octave and 1/6-octave bands. George Bissinger, in his VIOCADEUS project [4], uses 250-Hz band averages across the violin's entire range. It would be a mistake to assume that two instruments with identical band averages would 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.

## IMPULSE MEASUREMENTS

If the bridge is tapped sideways, the resulting sound will contain contributions from all the instrument's vibrational modes—or at least those which radiate sound, are within the hammer's frequency range, and can be excited by side-to-side motions of the bridge. Similarly, a vertical

blow to the top of the bridge excites those modes which can be excited by a vertical force.

Figure 9 shows the hammer and microphone signals for a single impact to the bass corner of the bridge of a typical violin (with strings damped). The hammer signal rises from zero to full value and then falls to zero again in about half a millisecond. The microphone signal rises quickly and then falls exponentially, taking ~0.125 sec. to fall below the noise floor.

In Fig. 10 a delay of about a millisecond between the beginning of the hammer signal and the beginning of the microphone signal represents the time taken for the sound waves to travel the ~34 cm between the violin and microphone.

Figure 11 shows a typical measurement cycle, consisting of four impacts at each of 12 microphone positions. At position one, the microphone is directly in front of the bridge and 37 cm from the central axis of the instrument. For each subsequent position, the instrument is rotated by 30° with respect to the microphone. A *complex average* of the impacts at each position yields a single measurement in which random noise has to some extent been averaged out.

## IMPULSE RESPONSE

For measurement purposes, the term *impulse*




Figure 9. Signals from hammer (below) and microphone (above) for a single impact. The impact is about half a millisecond in length, while the response of the instrument (with its strings damped) sinks below the noise floor within ~0.125 second.




Figure 10. Close-up of hammer and microphone signals from Fig. 9. The ~1-ms delay between the beginnings of the hammer and microphone signals represents the time taken for the sound waves to travel the ~34-cm distance between violin and microphone.




Figure 11. A typical measurement cycle, consisting of four impacts at each of 12 microphone positions. Taking the complex average of the impacts for each position tends to cancel out random noise, yielding a single frequency response function (FRF) for each position.

refers to a large force delivered over a very short time period. The term *impulse response* refers to a system's response to a perfect impulse. Physicists define a *perfect impulse* (known as a *Dirac* or *delta* function) as one that is infinitely short in duration and whose amplitude is infinite, so as to have an area of one unit (see Appendix B). In terms of digital signals, *infinitely short* means less than the length of a single sample, and *infinite amplitude* means the maximum that can be captured by the chosen sampling depth. For our purposes, then, an impulse response can be thought of as the sound that would be made if the violin were tapped with an impact shorter than one sample, and with sufficient force for the hammer signal to reach the maximum value on the amplitude scale. Because this is not feasible in practice, an impulse response is calculated by comparing the hammer and microphone signals, as described below.

It happens that an infinitely short impulse contains an equal amount of energy at all frequencies, meaning that an instrument's *impulse response* is the precise equivalent to its frequency response. But where a frequency response function describes the instrument's behavior in the *frequency domain* (i.e., as amplitude plotted against frequency), an impulse response gives exactly the same information in the *time domain*

(amplitude against time). An impulse response can therefore be calculated by taking the complex spectrum provided by the transfer function between microphone and hammer, and then translating this back into the time domain using an inverse FFT.

Because an impact hammer delivers a relatively short impulse, the unprocessed microphone signal tends to resemble the calculated impulse response. Note, however, that an impulse response can be derived from any measurement system, whether it relies on impulse excitation or not. All that is necessary is that there is an accurate record of both input force and output response, and that the excitation force contains a reasonable amount of energy across the frequency range of interest.

## HAMMER SPECTRUM

If an impact hammer could deliver a blow shorter in length than a single sample, then it would contain equal amounts of energy at all frequencies within the range defined by the sampling rate. In practice, a blow of this brevity would require both hammer and bridge to be very hard indeed. The reason 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 the high-frequency content in the excitation force, and therefore in the sound produced. To demonstrate this, tap a tabletop with your fingertip, and then with a house key. The brighter click made by the key indicates the greater high-frequency content.

All of this is in a sense moot, for as long as the sensor in the hammer gives an accurate history of the changing force between hammer and bridge, any unevenness in the frequency content of the impact can be accounted for in the analysis process. If, however, the hammer signal contains very little energy above, let's say, 2 kHz, then the instrument will be only weakly excited at higher frequencies. The resulting sound levels will be so low they tend to get lost in background sounds. For this reason, a relatively hard hammer tip is required.

Practically speaking, by varying the hardness of the hammer tip, the energy in the impact can be either distributed over a wide frequency band (hard tip) or concentrated at relatively low frequencies (soft tip). A problem with very hard tips is the possibility of damaging the bridge—indeed, the pointed metal tip of the smallest PCB impact hammer leaves a visible dent. A second consideration is the increased possibility of double bounces, discussed below. A tip made from high-molecular-density polyethylene—the white, translucent polymer much used in kitchen

cutting boards—works well. It is softer than the hammer's metal tip, harder than the PCB hammer's soft tip, and gives good data up to ~9 kHz. Figure 12 shows spectra for soft, medium, and hard tips when used to strike a massive metal block. When tapping a violin bridge, the relative softness of the maple comes into play, and so using tips harder than the polymer does not much improve the frequency range. To get around this, one could perhaps attach a tiny piece of hard material to the bridge at the point of impact, although I have not yet tried this.

## 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, causing a second collision. Because this can happen within milliseconds, double bounces are often difficult to hear. Figure 13 shows the hammer signal for a small double bounce. The initial peak falls to zero (indicating loss of contact between hammer and bridge), but then rises again to form a secondary peak.

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




Figure 12. Hammer spectra using soft, medium, and hard hammer tips against a massive metal block. Upper curve: PCB metal tip; middle curve: custom polymer tip; bottom curve: PCB soft (red) tip. (Note: the cut-off at 20 kHz is due to the sampling rate, not the hammer tip.)



Figure 13. Hammer signal from a double bounce in time domain (below) and frequency domain (above), showing the characteristic oscillation in the hammer spectrum.

Fig. 8. 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. In all other cases, as Gabriel Weinreich demonstrates in Appendix B, a double or even triple bounce will yield exactly the same results as a single bounce.

## NOISE

Both electronic noise and noise from the environment need to be dealt with. Electrical noise is most problematic at early stages of amplification, prior to analog/digital conversion. Microphones use shielded cables and balanced outputs to screen out most electrical noise. For relatively short cable lengths, this works very well. The interior of a computer tends to be electrically noisy, and so outboard sound cards (typically connected by USB or Firewire interfaces) are generally preferred to those installed inside the computer. Environmental noise can be minimized by turning off heating, air conditioning,

fluorescent lights, computer fans, telephones, etc., and otherwise using a quiet room for measurements. In general, the effects of random noise on a measurement can be reduced by:

- Driving the bridge as hard as possible without overloading the hammer sensor;
- Adjusting the gain on both hammer and microphone so that their signals are, for the loudest anticipated blows, just shy of an overload;
- Averaging the results from many impacts. If the complex average (which takes phase into account) is taken for all the impacts at each microphone position, random noise tends to be averaged out. The cancellation effect is proportional to the square root of the number of measurements used. Thus, averaging 16 separate hammer blows will tend to reduce noise by a factor of four;
- Removing the "dead time" between impacts. Because the impulse response of a violin is relatively short, the time between individual impacts (with its associated noise) tends to far exceed the time periods of interest. To improve the signal-to-noise ratio, this dead time is usually avoided by *triggering*, which typically works as follows: A trigger in the software menu is set so that when the hammer signal reaches, let's say, 10% of its maximum possible value, the processor begins the measurement. First, it backtracks a chosen number of samples, e.g., 40, to ensure that the entire

hammer signal is included. A single FFT of sufficient size to capture both the impact and response is then calculated, after which the trigger resets itself. Triggering can be done during acquisition of the original signal, or later using a recording of the entire measurement cycle. Triggering during acquisition ensures much smaller files. Post-measurement triggering means that a record of the time between impacts is retained, and this is sometimes useful for assessing background noise and room acoustics.

## COHERENCE

Both random noise and tiny drifts in the measurement conditions ensure that even nominally identical measurements yield slightly different results. A way of testing the consistency of a measurement series is to have the software compute a coherence function. The lower graph in Fig. 14 is a coherence function for a series of hammer blows at a single microphone position. If the measurements were all identical, they would have a coherence value of 1.0 across the whole frequency range. As it is, much of the range is close to one, but with “icicles” here and there of lower coherence. These icicles tend to coincide with valleys in the FRF, where the rela-

tively low amplitude levels allow random noise to significantly influence the measurements, and thus reduce the coherence between them. For the same reason, the drop off in excitation level with increasing frequency (due to the relative softness of the hammer tip) leads to a corresponding drop in coherence. Generally speaking, coherence is an indication of the quality of the measurement.

## ROOM ACOUSTICS AND MICROPHONE DISTANCE

Room acoustics are probably the biggest obstacle to clean measurements. One solution is to build some approximation of an anechoic chamber. Something the size of a pantry or walk-in closet works for violins and violas. An alternative approach is to hang absorbent materials in a kind of tent around the measurement area. Irregular objects such as furniture and wall hangings help in diffusing wall-reflections, and specially designed diffusers and absorbers can be made or bought. Generally speaking, the larger and less reverberant the room, the better it is for measurement purposes. For an excellent nontechnical account of room acoustics, including sound adsorption and diffusion methods and materials, see Ref. [5].

Figure 15 compares a single violin measured




Figure 14. Frequency response function (above) and coherence function (below) for a series of impacts at a single microphone position.




Figure 15. A single violin measured in an anechoic chamber (red curve) and in a relatively small, reverberant room (blue curve). The microphone distance for both was 37 cm from the instrument's central axis. The graphs represent the average magnitude of all 12 microphone positions.

in an anechoic chamber (red line) and a relatively small, relatively reverberant room (blue line).

Schleske [1] reduces the effects of room acoustics by rotating both microphone and instrument with respect to the room. This tends to average out some of the room modes. He finds that the complex average of 20 readings taken at  $18^\circ$  increments gives reasonably repeatable measurements from one room to another.

The closer the microphone is to the instrument, the greater the proportion of the sound being measured will be coming directly from the instrument, rather than from room reflections. Close microphone positions therefore lessen the influence of room acoustics and environmental noise. On the other hand, to be sure that we are measuring sound that actually reaches the so-called far field, the microphone distance should ideally be at least one wavelength of the lowest frequency being measured. As violins and violas do not radiate much below about 200 Hz, this means  $\sim 1.7$  meters. While feasible in an anechoic chamber, this kind of microphone distance in a normal environment would mean prohibitive levels of reflected sound and room noise. Clearly, a compromise must be struck.

Schleske finds that 50 cm in a workshop setting is adequate to avoid most near-field cancellation. George Bissinger uses a distance of 1.2

meters (from the surface of the top) when measuring violins in an anechoic chamber, but moves in to 30 cm when working outside the chamber. I have used 37 cm from the central axis of the instrument as the default microphone position. More recently I have moved this up to 20 cm from the center.

Figure 16 compares measurements taken in my workshop (a relatively live room of about 24 ft x 30 ft) using microphone distances of 55, 37, 20, and 15 cm. Finding the optimum microphone distance should be governed by the particular aspects of the sound being studied. For example, when trying to measure what the violinist hears while playing the instrument, it makes sense to locate the microphone about where the violinist's left ear would be.

## DIRECTIONALITY

Because of the violin's directional characteristics, a number of readings must be taken around the instrument in order to get a reliable assessment of total radiation. Bissinger [4] uses 266, arranged spherically in an anechoic chamber. Langhoff [6] found that eight microphone positions provided a good estimate of total radiation. (By way of confirmation, the instrument was rotated with respect to those positions,






Figure 16. Measurements made at a single microphone angle (directly in front of the bridge), but at varying distances from the central axis of the violin. Red curve: 15 cm, magenta: 20 cm, blue: 37 cm; and light blue: 55 cm.

and the average signal did not change by more than 1 dB.) Schleske finds that the average of six readings, taken at 30° intervals around the instrument, in plane with the bridge, gives a reasonable estimate. (For a full measurement cycle, he does each of these six readings 20 times—rotating both microphone and instrument with respect to the room, as described above—for a total of 120 individual measurements.)

I have found that a single reading taken directly in front of the bridge gives a useful indication of an instrument's behavior up to ~800 Hz. To get a better sense of what is happening at higher frequencies, I have followed Schleske in using 12 positions, spaced evenly around the instrument in the plane of the bridge. How adequate is this as an estimate of total sound output? It is probably very good below ~850 Hz, where omnidirectional radiation predominates. More testing is needed to determine the adequacy at higher frequencies.

## SAMPLE MEASUREMENTS

The measurements cited below are all for a single, good-quality, modern violin with its strings damped. Measurements were done in an anechoic chamber, and 12 microphone positions were used, as described above. The microphone distance was 37 cm from the central axis.

Figure 17 shows all 12 positions averaged in

two different ways. The line that starts dropping in level at ~1 kHz is the complex average for all 12 positions. This gives an estimate of the *monopole* radiation—i.e., the omnidirectional component of the total sound radiation. Monopole radiation clearly predominates below ~1 kHz. The other line is the average magnitude for all 12 positions. This provides an estimate of the total sound radiated in the plane of the bridge. Above ~1 kHz, monopole radiation for this violin is ~10 dB below the total (multipole) radiation.

Figure 18 shows an overlay of five 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 ~1 kHz indicates radiation patterns that change significantly with frequency.

Figure 19 gives the average magnitude of five front and five rear microphone positions. At high frequencies the sound is radiated mainly forward from the instrument. The difference in the levels of the A0 peak (280 Hz) are due to the close position of the microphone relative to the acoustic center of the A0 mode, which is slightly in front of the *f*-holes.

There are many ways of deriving single response curves from multiple readings. The upper graph in Fig. 20 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




Figure 17. Estimated total radiation (upper curve) and monopole radiation (lower curve). The first was arrived at by taking the average magnitude for 12 microphone positions, and the second by the complex average, which allows for cancellation due to phase differences.




Figure 18. Overlay of five front microphone positions.

the instrument?

The averaged (top) measurement provides an estimate of the total radiated sound—assuming that radiation in the plane of the bridge is representative. While its overall shape is similar to the single-microphone measurement, note that above 1 kHz, the peak-to-valley heights are far less, having been smoothed over in the averaging process. The pronounced “spikiness” of the single-microphone measurement reflects an important aspect of the instrument’s sound.

The single-microphone measurement represents what might be heard by listening with one ear held 37 cm from a violin suspended 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 that 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.

## ACKNOWLEDGMENTS

Thanks to Martin Schleske, who introduced me to impact hammer measurements; to George Bissinger, for sharing his extensive experience in measuring violin radiation; to Fan Tao, who was




Figure 19. The average magnitude of five front (upper curve) and five rear microphone positions (lower curve). This shows that at high frequencies the sound is radiated mainly forward from the instrument.




Figure 20. The top graph shows the average magnitude for 12 microphone positions. The graph below shows the magnitude for a single, front-center microphone position.

instrumental in the development of the rig; and to Jim Woodhouse, for helping me understand some of the underlying physics. My special thanks goes to Gabriel Weinreich, Emeritus Professor of Physics at the University of Michigan, for his many years of patient mentoring, for his crucial role in refining the measurement system presented here, and for his permission to include the material he contributed in Appendix B. Many thanks also to the anonymous reviewers who made valuable suggestions.

## REFERENCES

- [1] M. Schleske, Empirical tools in contemporary violin making: Part II: Psychoacoustic analysis and use of acoustical tools, *Catgut Acoust. Soc. J.*, Vol. 4, No. 5 (Series II) pp. 43-49 (2002).
- [2] G. Weinreich and E.B. Arnold, Method for measuring acoustic radiation fields, *J. Acoust. Soc. Am.*, Vol. 68 (2), pp. 404-11 (1980).
- [3] G. Weinreich, Directional tone color, *J. Acoust. Soc. Am.*, Vol. 101, pp. 2338-46 (1997).
- [4] G. Bissinger, Contemporary generalized normal mode violin acoustics, *Acustica*, Vol. 90, pp. 590-99 (2004).
- [5] F.A. Everest, *Master Handbook of Acoustics*, 4th ed. (McGraw-Hill, New York, 2000).
- [6] A. Langhoff, Measurement of acoustic violin spectra and their interpretation using a 3D representation, *Acustica*, Vol. 80, pp. 505-15 (1994).
- [7] SpectraPlus™ acoustical analysis software; <www.SpectraPlus.com>.

## APPENDIX A: Hardware

Listed below is the hardware used in the rig described above, along with associated software and electronics.

**Impact hammer:** PCB Force Hammer (Model 086C80) & power supply (e.g., Model 480E09). Total cost: ~\$1,200 at time of printing. A variety of power supplies are available, including a rechargeable one, which is useful as battery changes are inconvenient with their design. They also offer a 2-channel USB-powered unit. Though small and very convenient, it is notice-

ably noisier than battery-powered versions.

**Sound card:** I use the Edirol UA25 (~\$225) USB audio interface. Like most others of its kind, it has two channels, provides the phantom power and pre-amps needed by condenser microphones, is small and portable, and is powered by the computer via the USB connection. Therefore, a separate power supply is not needed. Many good-quality external sound cards are available. Some use a FireWire interface, which is more common on Mac computers than PCs.

**X-Y-Z Axis Metric Stage:** Available from Edmond Optics, Stock No. K6-041, current price ~\$490.

**Microphone:** I have used an Earthworks SR77 calibrated small-diaphragm cardioid condenser microphone, priced at ~\$880. Many less-expensive microphones also will work fine. Note that omnidirectional microphones are usually recommended for measurement purposes (see section above on microphones).

**Pro-audio speaker stand:** Ultimate Support, model TS 80B, 1-1/2" shaft: ~\$80. For added portability, use lightweight camera tripod legs, such as *Velbon 540A* graphite tripod legs, ~\$275. The unit is incredibly light at 2.7 lb, and it folds down to 16.5".

**Software:** SpectraPlus™ acoustical analysis software [7]. The 6-option package is recommended (\$995) or the full 10 options (\$1,295). An alternative, lower-priced, and in many ways preferable software is that written by George Stoppani; see <www.stoppani.co.uk>.

## APPENDIX B: The Double Bounce in Impact Hammer Measurements

As described in this paper, a standard way to obtain the radiativity impulse response of a violin is to use an impact hammer to exert a side-ways impulsive force on the bridge (mimicking the approximate direction of the force exerted by the bowed vibrating string) and record the signal from a microphone located in a standardized position (for example, at a distance of 30 cm above the bridge in a direction perpendicular to the plane of the violin). If the impact of the hammer were truly impulsive (that is, zero length in



time; or, for a sampled signal, limited to a single sample), the unprocessed microphone signal would immediately yield the desired impulse response. In reality, however, the impact lasts more than one sample, often even showing more than one peak. The exact shape of the impact is given to us by the output of the piezoelectric sensor in the tip of the hammer. Thus, the experimental data consist of two sampled signals: the hammer impact, lasting perhaps 1 ms or somewhat more, and the microphone signal which, even if the open strings are damped, still lasts an appreciable fraction of a second.

## SAMPLE RESULTS

The following examples are not, strictly speaking, experimental in the physical sense. Rather, they are obtained by first inventing a numerical form for a possible impulse response; then inventing various possible hammer signals; and finally computing the corresponding expected microphone signals by an actual point-by-point convolution, without using any FFT shortcuts. We then try to recover the original impulse response. If our mathematics is correct, the recovery should be perfect, except for small deviations due to the accumulation of rounding-off errors.

Figure 21 shows the layout for this and all successive pictures: the right half is the assumed impulse response; the left shows the assumed hammer and microphone signals (in this picture, there are none). In this imagined case, this impulse response has the simple shape of a damped sine function, except that in the beginning there is a distinct delay due to the transit time of the acoustic pulse from the violin to the microphone. (As a rule of thumb, sound travels through air with a speed of ~1 millisecond per foot.)

Figure 22 shows, in the left pane, an assumed hammer signal (green) and the resulting microphone signal (red), computed by point-by-point convolution. The (green) hammer signal is particularly simple here, being a single-sample impulse; as a result, the microphone signal is identical to the impulse response, except that it is timed relative to the hammer signal and slightly diminished in amplitude by the same factor by which the green impulse fails to reach the top of

the diagram.

We next ask to *compute* the impulse response back from the red and green curves considered as measured data, without taking advantage of the hammer signal being a perfect impulse, and realizing that the point-by-point convolution is not only laborious, but it leads to no direct way of being reversed. At this point we can, however, make good use of the passage from time domain to frequency domain. The logic is the following: in the frequency domain, the microphone signal  $M(f)$  has a component at frequency  $f$ , which is a simple product of the hammer signal  $H(f)$  component at that frequency and the impulse response component  $IR(f)$  at that frequency:

$$M(f) = H(f) \bullet IR(f) . \quad (1)$$

This equation can be solved for  $IR$  simply by dividing both sides by  $H(f)$ :


$$IR(f) = M(f) \div H(f) . \quad (2)$$

By using three FFTs—one to obtain  $M(f)$ , the second to obtain  $H(f)$ , and the third to transform  $M/H$  back to the time domain—we get the impulse response, which we pretended not to know. What we did in Fig. 2 is to plot it as a light blue curve right on top of the black curve shown in Fig. 1. The fact that no black shows through demonstrates that any mathematical error involved in this computation is less than the thickness of the curve.


Figure 23 shows a situation similar to Fig. 2, but with the hammer impulse coming at a later time, so that its tail does not fit into the region covered by the FFT. The computed impulse response (still light blue) begins correctly, but becomes suddenly zero where the red curve ends. The light blue curve then jumps to the x-axis, making the black “original” visible until it, too, joins the axis.

A more interesting case is when the hammer takes two bounces. In this case, no inversion of the convolution is possible directly, and the FFT method of deconvolution is the only one that works well.

We conclude with a more complex, yet in practice often encountered, case in which there are not only two bounces but also neither bounce is limited to the width of a single sample (Fig. 24).



*Figure 21. Impulse response to a hammer impact, imagined as a damped sine function, delayed by the transit time of the acoustic pulse from the violin to the microphone.*



*Figure 22. Single-sample hammer signal (green), the resulting microphone signal (red), computed by point-by-point convolution, and the computed impulse response (light blue).*




Figure 23. Similar situation as in Fig. 22, but with the hammer impulse coming at a later time.




Figure 24. Similar to Fig. 22, but a more complex case involving two bounces with neither bounce limited to the width of a single sample.

As one might have expected with a little thought, the pattern of the red response curve is quite complex in the beginning, up until the moment when the second bounce terminates. In fact, the only reason it appears simple afterwards is that the actual impulse response is a damped sinusoid, and there is a mathematical theorem that the sum of any number of damped sinusoids all with the same frequency and the same damping is itself a damped sinusoid with that frequency and that damping. Had the underlying impulse response been more complex, the red curve in Fig. 24 also would have been more complex. And yet the problem would have yielded just as smoothly to the FFT method.

The reader's attention is especially called to the fact that in the last two figures the hammer

took more than one bounce. The reason this happens in practice can be rather complicated, yet it is important to note in these examples that, contrary to statements that you may encounter in other sources, *the presence of multiple bounces does not make it any more difficult to obtain the correct impulse response.*

Given experimental data from an impact hammer experiment, provided that the data have sufficient precision, it is always possible to obtain an impulse response by (1) taking an FFT of the hammer data, (2) taking an FFT of the microphone data, (3) doing a complex division of the result of (2) by the result of (1) term by term, or (4) doing an inverse FFT on the result of (3). The presence of multiple hammer impacts is not an obstacle to this method.