

Frequency ratios of spectral components of musical sounds

Judith C. Brown^{a)}

*Physics Department, Wellesley College, Wellesley, Massachusetts 01281
and Media Lab, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139*

(Received 11 August 1994; revised 3 August 1995; accepted 10 August 1995)

A recently developed high resolution frequency tracker [J. C. Brown and M. S. Puckette, "A high resolution fundamental frequency determination based on phase changes of the fourier transform," *J. Acoust. Soc. Am.* **94**, 662–667 (1993)] has made it possible to measure the ratios of the frequencies of the upper harmonics of a sound with respect to its fundamental frequency with high accuracy. Calculations were carried out on digitized sounds produced by a clarinet, alto flute, voice, piano, violin, viola, and cello. The sounds produced by the stringed instruments included examples played pizzicato and bowed both with and without vibrato. Measured ratios were exactly equal to integers for all instruments except the piano and strings played pizzicato. Anomalous behavior was observed for the fundamental frequency for vibrato sounds played by stringed instruments with the frequency deviation exceeding the extrema of the other harmonics divided by their harmonic number by about 1% on average. Piano inharmonicity was proportional to harmonic number squared in agreement with Fletcher [J. Acoust. Soc. Am. **36**, 203–209 (1964)]. The major limitation on this calculation was found to be instrumental fluctuations. © 1996 Acoustical Society of America.

PACS numbers: 43.75.St

INTRODUCTION

A knowledge of the exact ratios of the frequencies of the partials of sounds produced by musical instruments is important for an understanding of the underlying physics of the production mechanisms of the instruments producing these sounds. Also important is the application to the production of synthetic sounds which may be used in musical compositions for computers (Fletcher *et al.*, 1962, 1965, 1967).

Recent fundamental frequency tracking methods, which are more accurate and efficient, and advances in computer speed and memory have made it possible to make rapid measurements on the frequencies of the Fourier components of sounds produced by a variety of musical instruments. Preliminary results in this study were reported by Brown (1994).

I. BACKGROUND

Early work by Fletcher *et al.* (1965) reported that the frequencies of the partials in steady tones produced by members of the string family "were found to be harmonic—that is, integral multiples of the fundamental frequency." These measurements were made on a Sonograph so they were not highly precise frequency measurements. In a later paper on violin vibrato tones (Fletcher and Sanders, 1967), it was reported that the frequencies found for all the harmonics "were similar in shape and extent." In his textbook on musical acoustics, Benade (1976) states that there is a wide class of instruments whose frequency ratios are related by precisely whole numbers, and these are the instruments producing sustained sounds. No experimental results are cited.

More recently Ando and Yamaguchi (1993) have measured the statistical fluctuations of the note C₅ produced by a

number of instruments both with and without vibrato. They find that, for a given note, the standard deviations of all the harmonics of a sound are nearly equal and "conjecture" that the reason for this is that the frequencies of the harmonics vary synchronously with the fundamental. This paper is interesting because it measures the inherent fluctuations of sounds produced by musical instruments.

Schumacher (1992) states that sounds produced by stringed instruments are aperiodic with the origin lying in the fundamental mechanisms of sound production, such as bow hair inhomogeneity for the bowed instruments. It is thus of great interest to follow the fluctuations in frequency for the fundamental and at the same time determine experimentally whether the higher harmonics exhibit identical fluctuations at integer ratios.

Thus while harmonic ratios have been discussed widely over the past decades, there has been no systematic effort to measure the frequencies of a variety of instruments purported to produce sounds with harmonics in integer or near integer ratios and report them along with an assessment of the accuracy of the measurements.

II. ACCURACY OF FREQUENCY DETERMINATION METHODS

Since the goal of this study is to determine how closely the measured ratios of harmonics are to integers, it is important to have the most accurate possible value of the relative frequencies. We will examine three frequency domain methods of frequency tracking.

A. Quadratic fit

We assume that the fast Fourier transform (FFT) of a sound has been taken and that the bin position of the fundamental frequency of a sound is known. The quadratic fit

^{a)}E-mail: brown@media.mit.edu

(Smith and Serra, 1987; Serra, 1989; Brown and Puckette, 1993) has been widely used to obtain a greater accuracy for the frequency than a bin width. With this method, one fits a parabola through the point corresponding to a local maximum of the FFT magnitude plot and the points on either side of this maximum. Then the position corresponding to the maximum of this parabola is called the “best” approximation to the true frequency.

The error in the quadratic fit method depends on the exact position of the component within the bin and the window used, and this error is present even for a single sinusoidal component (Brown and Puckette, 1993) except for the case of a Gaussian window (Serra, 1989). The maximum error for a Hanning window is 0.057 bin widths, and this occurs for a component which has a frequency corresponding to one-third of a bin width away from the center frequency; it is independent of which bin the component falls in. So the fractional error ($=0.057/\text{bin number}$) would be greatest for the first bin, where it is 5.7%, and it varies inversely with bin number. The errors for this method can be reduced at computational expense by resampling so that the maximum falls near a bin center frequency and by zero padding to give a smaller bin width. For an example, see Ando and Yamaguchi (1993).

B. Phase vocoder

The phase vocoder (Flanagan and Golden, 1966) was originally formulated as a means of data compression in communications. It has been extremely useful, as well, as a means of obtaining more accurate values for frequencies of Fourier components than those which can be obtained from the Fourier magnitude spectrum. Taking values of the real and imaginary parts of the Fourier component in a particular bin, one calculates the phase for that bin. After a time advance of H samples, another FFT is computed, and the phase is again calculated for the same bin. The frequency in radians per sample is the phase change divided by the time advance: $\omega = (\phi_2 - \phi_1)/H$.

There is no theoretical error for this method for a single Fourier component. In practice, errors arise from interference by other Fourier components and from extraneous noise (Brown and Puckette, 1993). A more serious problem for the study of time-varying frequency components is that the time resolution for this method is equal to the window size of the FFT plus the hop size. If the hop size is small, as is desired for good temporal resolution, then many FFT's must be calculated, and this is expensive computationally.

C. Single frame approximation

We can take advantage of the accuracy of the phase vocoder without sacrificing computational efficiency or temporal accuracy by making an approximation which we call the single frame approximation. Here we calculate the phases for a single FFT frame. Then a standard approximation is made for these Fourier components after a time advance of one sample. From this, the corresponding phases can be calculated to give the frequency of any component of the FFT. This method is equivalent to the phase vocoder with a hop of

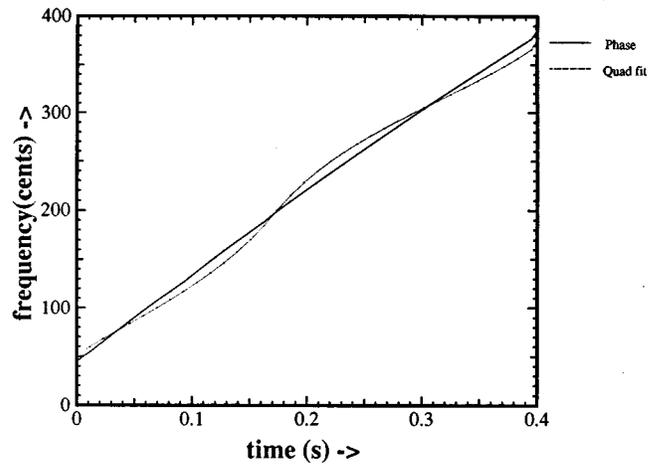


FIG. 1. Comparison of frequency tracking by quadratic fit (small dots) and single frame approximation (solid line) for a ramp frequency generated in software. Measured frequency (in cents with respect to A 440) is plotted against time in seconds.

one sample but without the computational expense of the second frame (Brown and Puckette, 1993). An added advantage is that maximum accuracy is achieved for M samples (with window N and hop size H subject to the constraint $N + H = M$) if $H = 1$ (Puckette and Brown, 1996).

In Fig. 1 are shown the results of the single frame approximation (SFA) and those of the quadratic fit applied to a synthetic sound with a linearly increasing frequency generated with C sound (Vercoe, 1993). There are five harmonics present with amplitudes decreasing as $1/k$, where k is the number of the harmonic. The frequency range is chosen to cover slightly more than one FFT bin change so the quadratic fit returns a frequency value going below the true value, coming back to the true value halfway to the next bin center, and then going above the true frequency. The SFA returns the frequency values with an error less than 0.5 cents (where 100 cents $\approx 6\%$) for the arithmetic average of the frequencies over the duration of the window.

The deviations by the quadratic fit are small, and for most applications would be negligible. However, for our test of harmonicity, the quadratic fit would not be appropriate without correcting for the position within the bin. The results which are given below were done using the single frame approximation. The initial results were checked for consistency with the quadratic fit.

III. CALCULATIONS

A. Description

Calculations were carried out on a Decstation 5000/200 with all programs written in C. The sounds were recorded in private research studios or taken from scientific recordings such as the McGill series or the Japan Audio Society CD with an acoustically neutral background. They were recorded with a sampling rate of 44.1 kHz, with the exception of the violin sample which was recorded at 32 kHz. Most were resampled at 11.025 kHz for computational speed, and the fundamental frequency was determined using the high reso-

lution fundamental frequency tracker described by Brown and Puckette (1993). See also Brown (1991, 1992), and Brown and Puckette (1992). The frequency was converted to a fractional FFT bin, and the original FFT was then tested for maxima at integers times the position of the fundamental. If a maximum was found, the frequency of that component was determined using the single frame approximation. If no maximum was found, the two adjacent bins on either side were checked for maxima, and, if found, the frequency of that bin was recorded. If no maximum was found, then a large negative frequency value was returned which would go off scale in the graphs. All frequency values were stored for graphing by Drawplot, which is public domain X window software.

Amplitude plots were done by taking the FFT and adding up the magnitude squared for a given Fourier component and the bins on either side up to half the bin number of the fundamental. For example if the fundamental fell in bin 9, then the intensity of the 3rd harmonic would consist of the sum of the squares of bins 23 through 31. These plots are expressed in terms of intensity level given by $IL = 10 \log(A^2)$, and since the numbers are 16-bit integers the maximum amplitude is 2^{15} which gives a range of 90 dB on the curves.

Frequency measurements were made with a Hanning window of 25–100 ms and a time advance or hopsize of about 6 ms. This window was chosen because the single frame approximation relies on the Hanning window. There are roughly 175 frequency measurements per second for each harmonic. Frequencies are plotted in cents with the frequency of each harmonic divided by its harmonic number. Thus if all curves coincide, this means exact integer ratios to within 0.1%, which is the visual resolution of the curves. Results are presented graphically rather than in a table of averages with standard deviations. This is because important information on frequency fluctuations is preserved in the graphs, and this would be lost by taking numerical averages.

B. Synthetic sounds

To determine the limits on our method we generated test sounds using Csound (Vercoe, 1993). Each sound consists of five harmonics which are frequency and amplitude modulated with a successive phase shift for the amplitude modulation. The amplitude of the fundamental is 300 out of a possible maximum of 32 000 (16-bit digitization) so this is a very weak signal. Successive amplitudes are proportional to the inverse harmonic number with all amplitudes modulated by 33% of their maximum value. In Fig. 2 are found the results of the frequency tracker on the synthetic sound with no added noise. In Fig. 3 are found frequency tracking results on the previous sound but with random noise with an amplitude of 300 added, i.e., an amplitude equal to that of the fundamental and a factor of 5 greater than that of the weakest harmonic. The noise amplitude was doubled to 600 (not shown) showing the successive breakdown of the frequency tracking results due to poor signal to noise. Amplitude results are found in Figs. 4 and 5 corresponding to Figs. 2 and 3.

Test file with 5 harmonics

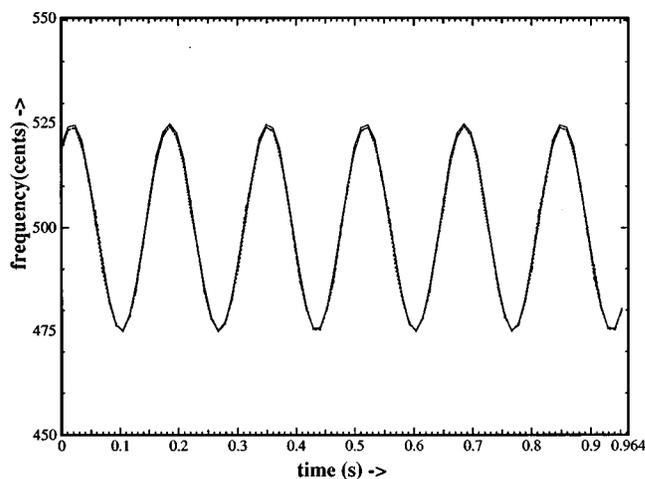


FIG. 2. Measured frequency (in cents with respect to A 440) of fundamental and harmonics 2–5 divided by harmonic number plotted against time for a sound generated in software. Amplitudes are proportional to inverse harmonic number and are modulated by 33% with a successive phase offset. The distance between two vertical markers is 10 cents.

To check the resolution of our method, in Fig. 6 we have results on a sound with five partials where the frequency of each successive higher partial is offset from an integer ratio by 2 cents relative to the harmonic number minus 1. These differences are clearly perceptible.

C. Winds

1. Clarinet

The clarinet is an interesting instrument to begin with as it shows some fairly general characteristics of the results (see Figs. 7 and 8). The frequency is varying, and all curves coincide to within 2 cents (0.12%) for the first 200 ms. At this point the amplitude of the fifth harmonic drops below 40 dB,

Test file with 5 harmonics

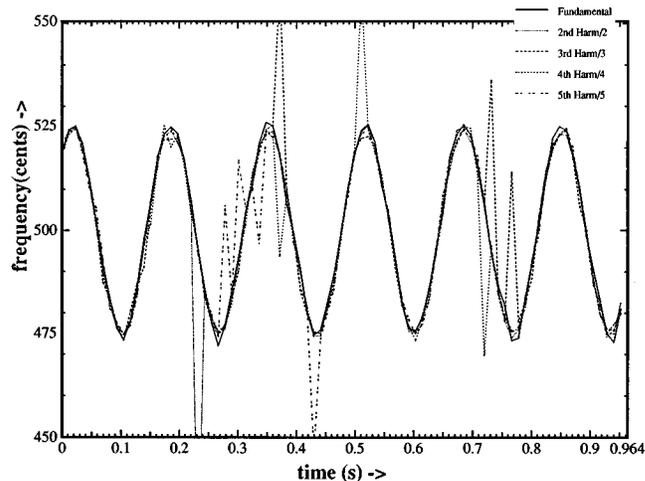


FIG. 3. Measured frequency (in cents with respect to A 440) of fundamental and harmonics 2–5 divided by harmonic number plotted against time for the sound of Fig. 2, but with random noise of the same amplitude as that of the fundamental added.

Test file with 5 harmonics

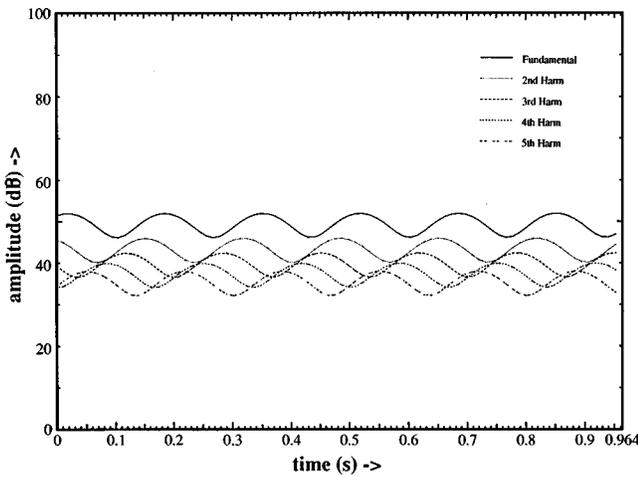


FIG. 4. Amplitude in dB of individual harmonics plotted against time for the sound of Fig. 2.

and errors are made for its frequencies as can be seen by the departure from the composite curve. The amplitudes continue to decrease and the frequency results on the 3rd and 4th harmonics become erratic as well.

After 450 ms the amplitudes increase (although all amplitudes except for that of the fundamental remain below 40 dB) and the frequency values all return to a single curve indicating integer ratios for the harmonics. We can thus infer that the clarinet does have harmonic frequency components, but a certain sound level is needed to measure them accurately. The 40-dB level can be considered as a rough indicator of this level since errors first begin to occur here. Forty dB means an amplitude of 100 out of a maximum 32 000, so this is a low-level signal, and noise and spill over of other components probably give rise to the errors in frequency determination.

Test file with 5 harmonics differing by 2 cents

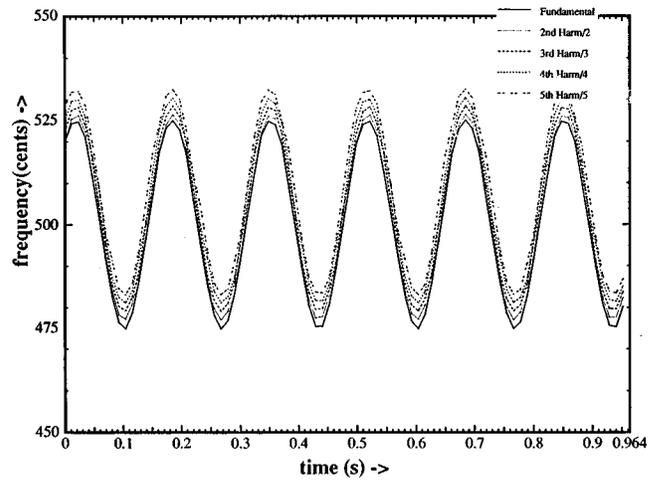


FIG. 6. Measured frequency (in cents with respect to A 440) of fundamental and harmonics 2–5 divided by harmonic number plotted against time for a sound generated in software. The frequency of each successive higher harmonic is offset from an integer ratio by 2 cents times the harmonic number minus 1. The distance between two vertical markers is 10 cents.

2. Alto flute

The alto flute is another member of the wind family (see Figs. 9 and 10). This sound was played without vibrato, and it shows fairly clearly the small fluctuations that are characteristic of any musical note generated by a human performer. Again we find exact integer ratios as shown by the coincidence of curves in Fig. 9 for the first three harmonics. There are a number of errors in harmonics 4 and 5, especially in the 4th harmonic. These results indicate perfect harmonicity; this is a consequence of the periodic production mechanism for the flute, often called the air reed.

Amplitudes are higher for this sound as seen in Fig. 10. They are on the order of 45 dB for the higher harmonics. For other instruments perfect frequency measurements were pos-

Test file with 5 harmonics

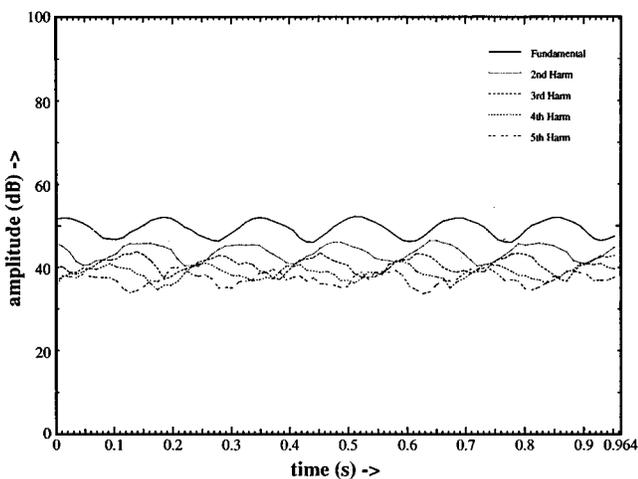


FIG. 5. Amplitude in dB of individual harmonics plotted against time for the sound of Fig. 3.

Clarinet F5

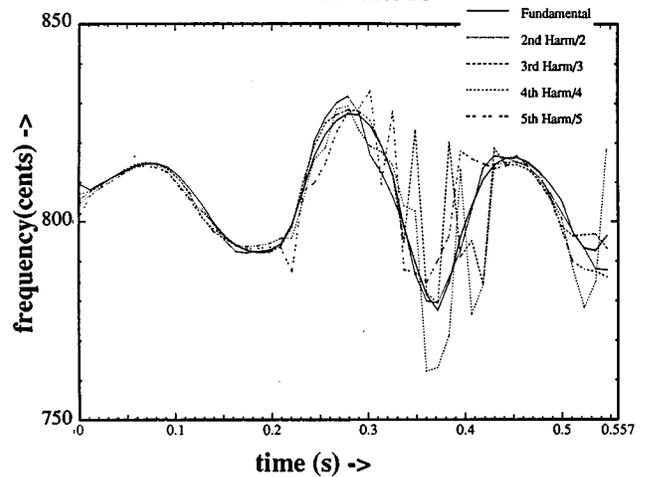


FIG. 7. Measured frequency (in cents with respect to A 440) of fundamental and harmonics 2–5 divided by harmonic number plotted against time for a clarinet playing the note F_5 .

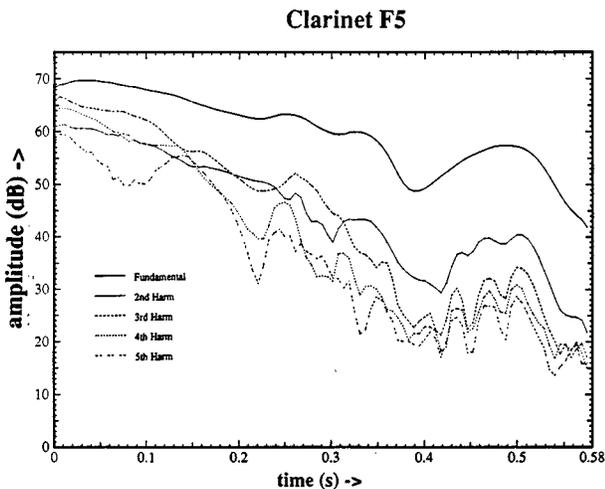


FIG. 8. Amplitude in dB of individual harmonics plotted against time for a clarinet playing the note F_5 .

sible at this amplitude level. The difference is probably due to the noise generated by the breathy production mechanism in the flute.

D. Voice

The voice, along with the bowed strings, is an instrument where harmonic components are expected since sound production by periodic glottal pulses is definitely phase locked (see Figs. 11 and 12). The amplitudes of the harmonics dropped off very rapidly so clean frequency measurements were only possible up to the 3rd harmonic, but results are near perfect and support the expectation of exact integer ratios of harmonics.

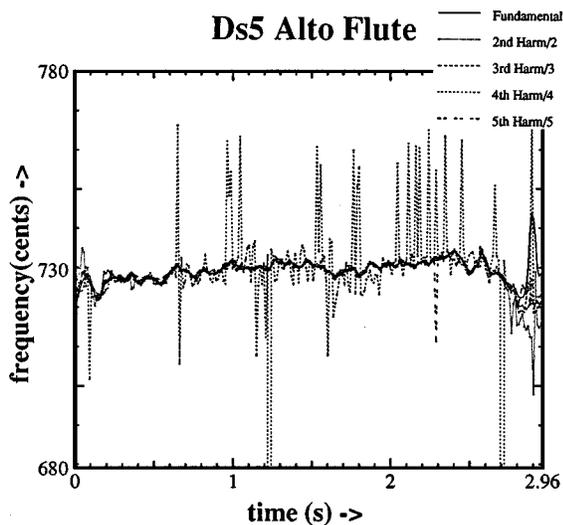


FIG. 9. Measured frequency (in cents with respect to A 440) of fundamental and harmonics 2 through 5 divided by harmonic number plotted against time for an alto flute playing the note $D\#_5$.

Ds5 Alto Flute

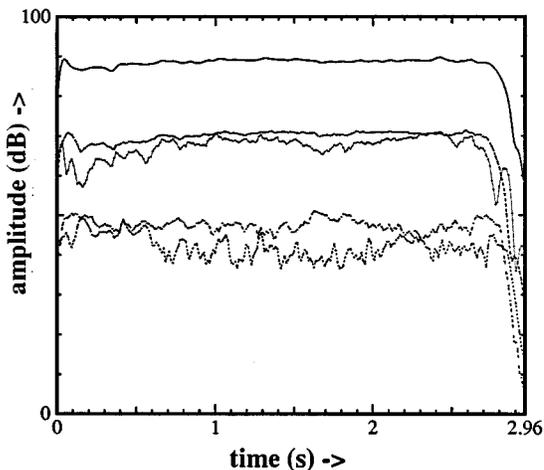


FIG. 10. Amplitude in dB plotted against time for an alto flute playing the note $D\#_5$. At time 2 s the harmonics are identified in order of decreasing amplitude as fundamental, 3rd harmonic, 2nd harmonic, with a drop to 5th harmonic and 4th harmonic.

E. Piano

There have been a number of studies of inharmonicity for piano sounds going back to Fletcher (1964) and references therein. Results (Figs. 13 and 14) were in agreement with previous studies, with partial ratios deviating from integers, and this deviation increasing with increasing partial number. In Fig. 15 the deviation from exact harmonicity for this note and for the note $F_5\#$ is plotted against harmonic number squared. Excellent agreement was found with the equation $f_k = k f_1 (1 + B/2k^2) / (1 + B/2)$ from Fletcher (1964) for the notes $D_4\#$ and $F_5\#$ with values of $B=0.00046$ and $B=0.0011$, respectively. This equation was adapted to give the ratio of f_k to $k \cdot f_1$ in cents plotted in Fig. 15.

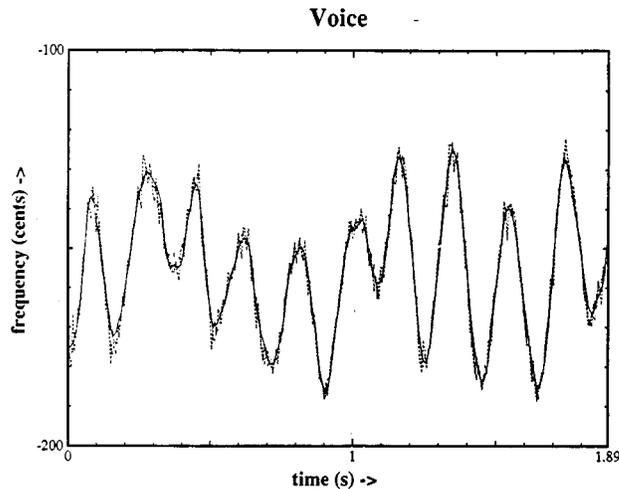


FIG. 11. Measured frequency (in cents with respect to A 440) of fundamental and harmonics 2 and 3 divided by harmonic number plotted against time for a voice singing with vibrato.

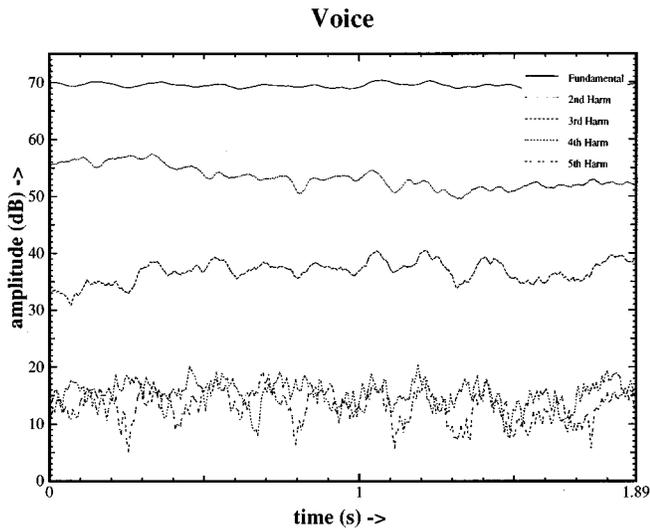


FIG. 12. Amplitude in dB plotted against time for harmonics 1 to 5 for a voice singing the note G_4 with vibrato.

F. Strings

1. Violin

Vibrato. In Fig. 16 are found the results for a violin played with vibrato. The sounds with vibrato played on stringed instruments were the most difficult for the fundamental frequency tracker for several reasons. First, the frequency is constantly changing due to the frequency modulation, and only an average can be measured due to the finite number of samples in the FFT window. Second, it is the motion of the performer's finger on the string which is causing this change in effective length, and there will always be unwanted fluctuations in bow pressure since humans are not mechanically perfect. Third, there may be more bow noise due to the varying conditions.

Measurements were possible for the first nine harmonics, although there were a few errors for the eighth harmonic,

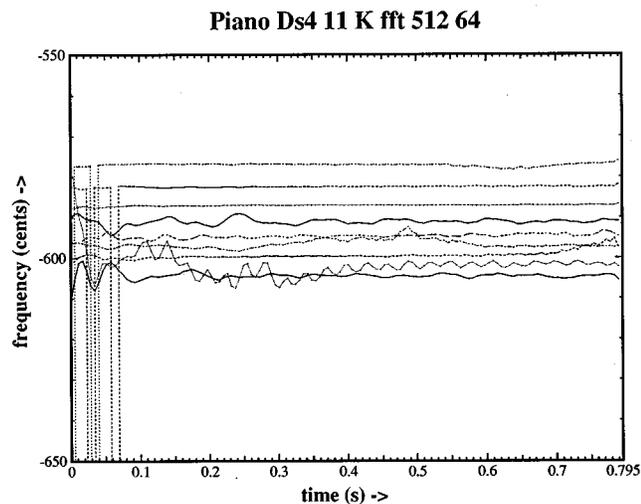


FIG. 13. Measured frequency (in cents with respect to A 440) of fundamental and partials 2 through 9 divided by partial number plotted against time for a piano playing the note $D\#_4$. The partials can be identified at time 0.65 s where the frequency deviations increase with increasing partial number.

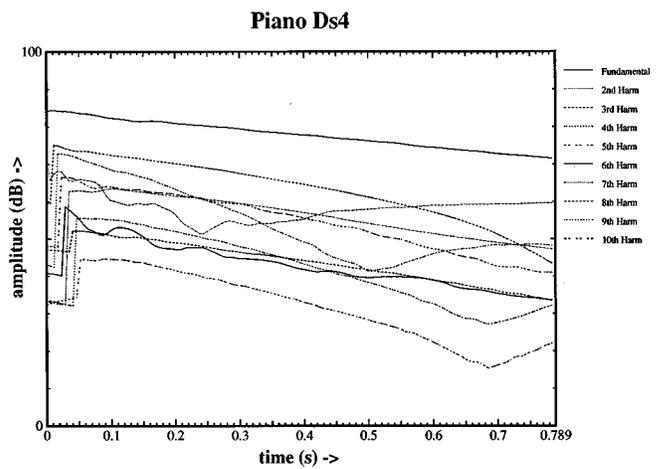


FIG. 14. Amplitude in dB plotted against time for a piano playing the note $D\#_4$.

the component with the lowest amplitude (see Fig. 17). All frequencies of harmonics with the exception of the fundamental are in ratios of the quotients of integers corresponding to their harmonic number to within a few cents. The fundamental is off by differences ranging from 5 to 25 cents, i.e., up to 1.5%, which is certainly greater than experimental error. We will return to this anomaly in a later section.

2. Viola

Pizzicato. In Fig. 18 are found the frequencies of the partials of a viola being played pizzicato for the note G_3 . As was found for other viola notes and a violin note played pizzicato, the frequencies are not in the ratio of exact integers. Differences are small, and the data are poor for the pizzicato notes where the amplitudes die off rapidly. So while a quantitative comparison to a formula such as that of Fletcher is not possible, it is nevertheless clear that inharmonicity is present.

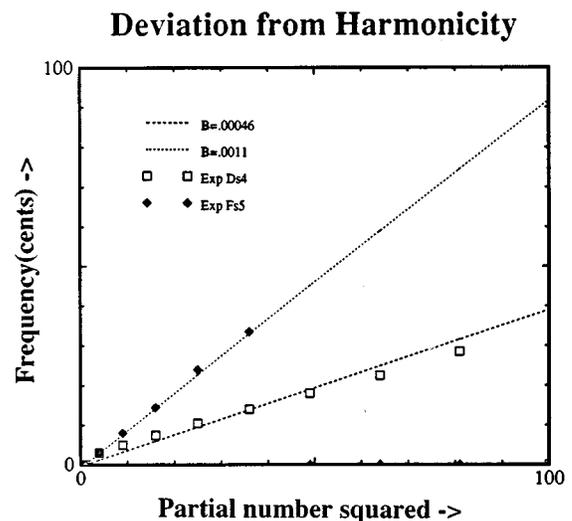


FIG. 15. Deviation from harmonicity for the piano notes $D\#_4$ and $F\#_5$ against partial number along with continuous lines representing the theoretical curves from Fletcher (1964).

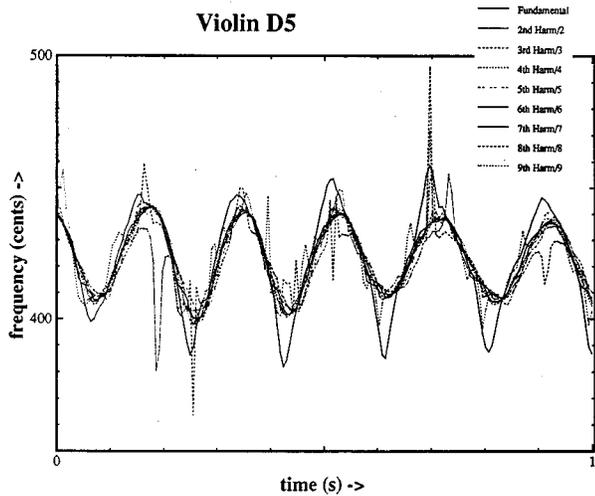


FIG. 16. Measured frequency (in cents with respect to A 440) of fundamental and harmonics 2 to 9 divided by harmonic number plotted against time for a violin executing the note D_5 with vibrato.

Amplitude plots are found in Fig. 19. The overall amplitude (not shown) of the time wave drops by a factor of 8.5 in 0.8 s. The fundamental only drops by about 5 dB during the 0.8-s duration while the second harmonic changes dramatically from 80 to 40 dB at time 0.8 s after which it increases. There are errors in the frequency reported for this harmonic in the region of its minimum.

Vibrato. Graphs of the frequency and amplitude results for the viola executing vibrato did not differ significantly from those for the violin vibrato (Figs. 16 and 17) and are not included. The measured frequencies for a viola playing the note D_5 show overall harmonicity for components up to five after which low amplitudes cause excessive errors for the frequency tracker. Amplitudes are below 40 dB for most

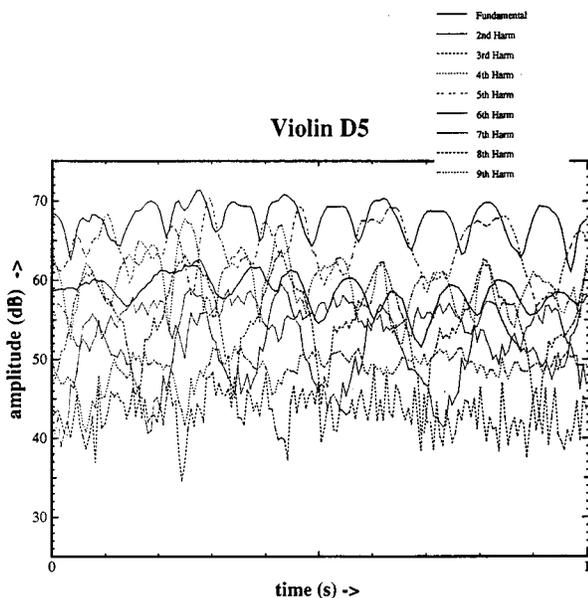


FIG. 17. Amplitude in dB plotted against time for a violin executing the note D_5 with vibrato.

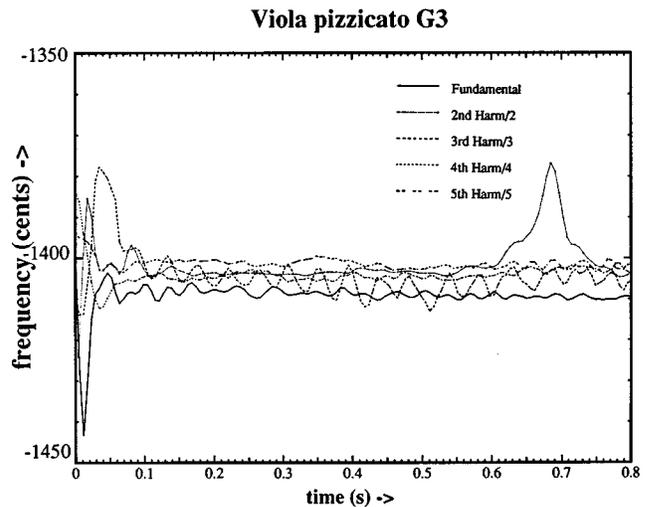


FIG. 18. Measured frequency (in cents with respect to A 440) of fundamental and partials 2 to 5 divided by partial number plotted against time for a viola playing the note G_3 pizzicato.

of the upper harmonics, and frequency measurements are not as clean as desired. There are periodic occurrences of “no frequency returned” for the 4th harmonic, which go off scale on the low-frequency side. These correlate nicely with amplitude minima. As in the case of the violin vibrato, the excursions of the fundamental exceed those of the other components, here by as much as 15 cents.

3. Cello

Vibrato. The graphs of Fig. 20 show two C_3 notes being played with vibrato on a cello. The first note is C_3 at the top of an ascending scale, and the next note is C_3 at the beginning of the descending scale. This can be seen clearly in the amplitude plot of Fig. 21. The region of the bow change between notes is very noisy as are the frequency measurements in this region.

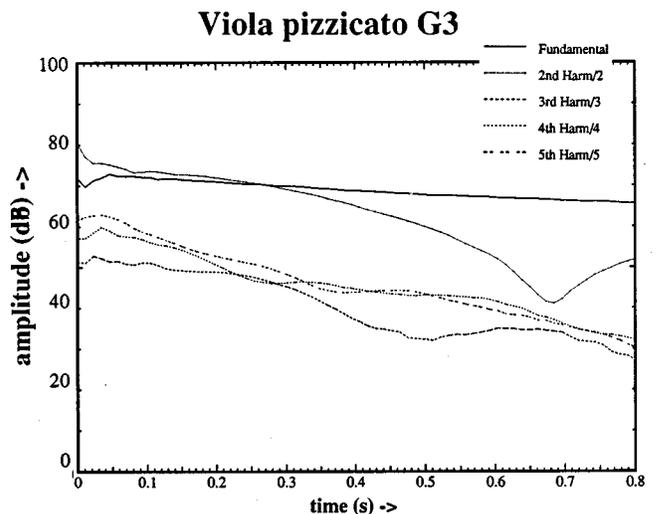


FIG. 19. Amplitude in dB plotted against time for a viola playing the note G_3 pizzicato.

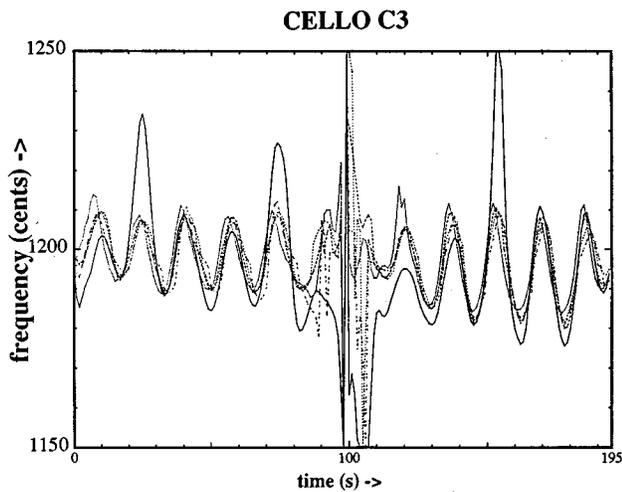


FIG. 20. Measured frequency (in cents with respect to $C_2=65.41$ Hz) of fundamental and harmonics 2 to 5 divided by harmonic number plotted against time for a cello executing the note C_3 with vibrato. The fundamental deviates from the other harmonics divided by harmonic number at the extrema.

All of the harmonics are in exact integer ratios with the exception of the fundamental. The frequencies for this vibrato note show the same behavior seen for vibrato executed by the other stringed instruments in that the excursions of the fundamental exceed those of the other components. Here, in

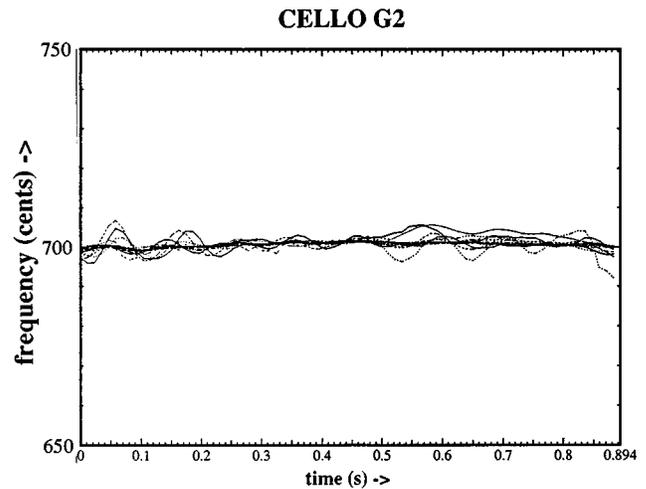


FIG. 22. Measured frequency (in cents with respect to $C_2=65.41$ Hz) of fundamental and harmonics 2 to 15 divided by harmonic number plotted against time for a cello playing the note G_2 .

addition, there is highly anomalous behavior at the amplitude minima of the fundamental (see Fig. 20). The frequency of the fundamental rises to as much as 40 cents above the frequencies of the other harmonics divided by their harmonic numbers.

Open string. The final result found in Fig. 22 is a real tour de force. This sound was recorded at MIT with

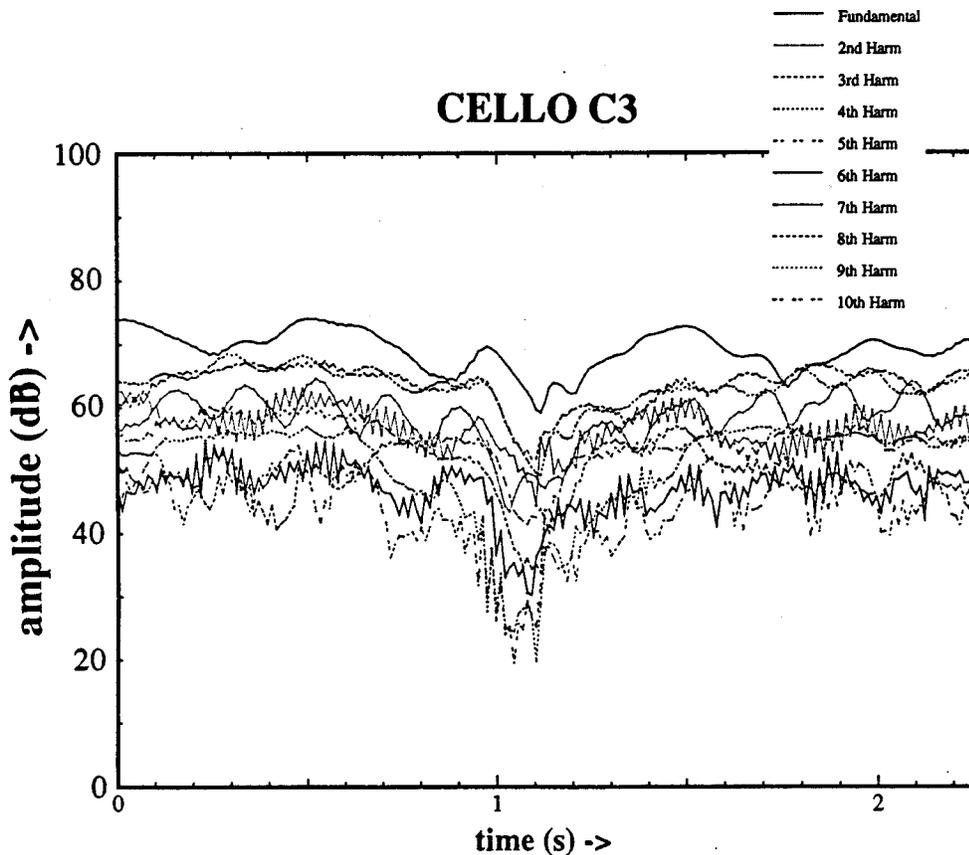


FIG. 21. Amplitude in dB plotted against time for a cello executing the note C_3 with vibrato.

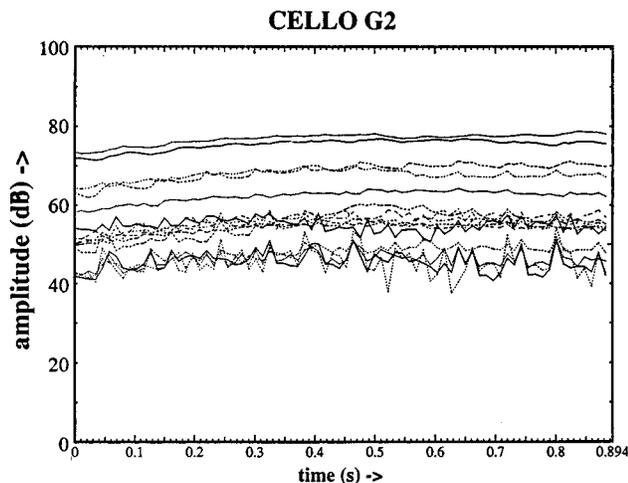


FIG. 23. Amplitude in dB plotted against time for a cello playing the note G_2 . Amplitudes decrease roughly in order of harmonic number.

Yo-Yo Ma playing a G_2 which is an open string and therefore has minimum frequency fluctuations. The amplitudes in Fig. 23 are above 40 dB out to harmonic number 15 and we have perfect harmonicity to within 3 cents or so. There are no errors in the frequency determinations out of over 1800 values. Measurements indicating exact harmonicity within ± 3 cents were possible up to the 25th harmonic.

IV. CONCLUSIONS

Continuously driven instruments such as the bowed strings, winds, and voice have phase-locked frequency components with frequencies in the ratio of integers to within the currently achievable measurement accuracy of about 0.2%. Since frequency fluctuations greater than the measurement accuracy are inherent in any sound produced by a human performer, improvement of the measurements is unnecessary. In fact comparison to measurements on synthetic sounds, where the accuracy is an order of magnitude smaller, indicates that frequency fluctuations are the limiting factor in this study rather than the accuracy of the frequency tracker.

The larger frequency excursions displayed by the fundamental frequency at its extrema for sounds produced by stringed instruments played with vibrato is perplexing. We explored the possibility that this might be due to the deviation from a sinusoidal modulation frequency by synthesizing a sound with a triangular modulation frequency. This displayed no such effect. We are unable to find a mechanism for this behavior, but think it unlikely that it could be an artifact of the measurement software since the synthetic sound displayed no such effect.

Impulsively driven instruments such as the piano and strings played pizzicato have partials which deviate from integer ratios. Here there is a brief excitation followed by an

independent decay for each component. The mechanism for this deviation is the stiffness of the strings, and the theory is worked out in detail in Fletcher (1964).

Finally, it is generally accepted that machine perception is inferior to that of the human perceptual system. In the case of pitch perception the human perceives a sound with a complex spectrum as having a single pitch at the frequency of the fundamental. We have demonstrated that a computer can do this as well, and in addition is capable of extracting frequencies of the higher harmonics with as high precision as that of the fundamental.

ACKNOWLEDGMENTS

I would like to thank Professor Miller Puckette for his helpful comments, particularly on the amplitude measurements, during a stay at IRCAM. I am also grateful to Professor Thomas Rossing for his suggestion that the triangular modulation might have caused the anomalous behavior of the fundamental frequencies of the vibrato results.

- Ando, S., and K. Yamaguchi (1993). "Statistical study of spectral parameters in musical instrument tones," *J. Acoust. Soc. Am.* **94**, 37–45.
- Benade, A. H. (1976). *Fundamentals of Musical Acoustics* (Oxford U.P., New York).
- Brown, J. C. (1991). "Calculation of a Constant Q Spectral Transform," *J. Acoust. Soc. Am.* **89**, 425–434.
- Brown, J. C. (1992). "Musical Fundamental Frequency Tracking using a Pattern Recognition Method," *J. Acoust. Soc. Am.* **92**, 1394–1402.
- Brown, J. C. (1994). "Measurement of harmonic ratios of sounds produced by musical instruments," *J. Acoust. Soc. Am.* **95**, 2889.
- Brown, J. C., and Puckette, M. S. (1992). "An Efficient Algorithm for the Calculation of a Constant Q Transform," *J. Acoust. Soc. Am.* **92**, 2698–2701.
- Brown, J. C., and Puckette, M. S. (1993). "A high resolution fundamental frequency determination based on phase changes of the fourier transform," *J. Acoust. Soc. Am.* **94**, 662–667.
- Flanagan, J. L., and Golden, R. M. (1996). "Phase Vocoder," *Bell Syst. Tech. J.* **45**, 1493–1509.
- Fletcher, H. (1964). "Normal vibration frequencies of a stiff piano string," *J. Acoust. Soc. Am.* **36**, 203–209.
- Fletcher, H. and Sanders, L. C. (1967). "Quality of Violin Vibrato Tones," *J. Acoust. Soc. Am.* **41**, 1534–44.
- Fletcher, H., Blackham, E. D., and Stratton, R. (1962). "Quality of Piano Tones," *J. Acoust. Soc. Am.* **34**, 749–761.
- Fletcher, H., Blackham, E. D., and Geertsen, O. N. (1965). "Quality of Violin, Viola, Cello and Bass-Viol Tones," *J. Acoust. Soc. Am.* **37**, 851–63.
- Puckette, M.S. and Brown, J. C. (1996). "Accuracy of frequency estimates using the phase vocoder," *IEEE Trans. Signal Process.* (submitted).
- Schumacher, R. T. (1992). "Analysis of aperiodicities in nearly periodic waveforms," *J. Acoust. Soc. Am.* **91**, 438–451.
- Serra, X. (1989). "A System for Sound Analysis/Transformation/Synthesis Based on a Deterministic Plus Stochastic Decomposition," Ph.D. thesis, Stanford University, Stanford, CA.
- Smith, J. O., and Serra, X. (1987). "An Analysis/Synthesis Program for Non-Harmonic sounds Based on a Sinusoidal Representation," *Proceedings of the 1987 International Conference on Computer Music*, Urbana, Illinois, pp. 290–297.
- Vercoe, B. (1993). "Csound: A Manual for the Audio Processing System and Supporting Programs with Tutorials," Massachusetts Institute of Technology, Cambridge, Massachusetts.