Bioacoustics
The International Journal of Animal Sound and its Recording, 2004, Vol. 14, pp. 171–178
0952-4622/04 $10

EQUIPMENT REVIEW

Sound Ruler Acoustical Analysis: A Free, Open Code, Multi-platform Sound Analysis and Graphing Package

Those of us interested in acoustic communication have witnessed over the last decade an explosion of new software packages designed for acquiring, analysing, editing, graphing, and playing back animal sounds from personal computers. Thankfully, the number of shareware programs for sound analysis continues to grow, and many of these programs are now able to provide sophisticated and powerful analysis tools at little or no cost. One of the newest such packages to hit the scene is Sound Ruler Acoustical Analysis (Gridi-Papp 2003). Sound Ruler is designed to make the detailed analysis of simple, repetitive sounds quick and easy. As such, Sound Ruler may be of particular interest to those studying acoustic communication in frogs or insects, in which calls or songs are often comprised of rather simple and repeated acoustic elements that are grouped into longer sequences of pulses or chirps (Gerhardt & Huber 2002). Sound Ruler is a free, open-code analysis package that is licensed under the GNU General Public License, and it is designed to run under Windows, MacOS X, or Linux operating systems. Thus, in addition to Sound Ruler being the right price, users of Sound Ruler have a choice of platform and access to the entire code, which is written as a series of modules in the M language. Users are even encouraged by the developers to adapt the Sound Ruler script to fit their specific analysis needs. These features alone make Sound Rule a promising piece of software for sound analysis. In addition, Sound Ruler is fairly well-documented, including some examples and step-by-step instructions for analysing a few frog calls. So, what can Sound Ruler do, and how well does it do it?

Overview of Sound Ruler's Features

Analysis

Sound Ruler has a number of attractive analysis features. First, the program can handle sound files (in the WAV format) of unrestricted size, thus eliminating any need to edit field recordings into chunks smaller than that of the entire cut. Second, adjustable low-pass, high-pass, and band-pass filters can be implemented online (i.e., during the analysis) with a few simple clicks of the mouse, and a special "tuning curve" feature allows the user to build customised filter functions. Hence, there is no need to save separate filtered sound files. Third, Sound Ruler allows for manual, automatic, and interactive recognition of acoustic events at two different time levels. For example, in signals that are comprised of pulsed calls organised into longer call groups, the two levels of analysis can be set by the user to analyse either call groups and calls, or calls and pulses. Fourth, Sound Ruler is able to automatically generate an impressive number of acoustic measurements (more than 50!) for each selected acoustic event, and the program allows the user to perform 2D (frequency x time) cross-correlations at the same time that measurements are made for each acoustic event. The data from these measurements can be saved to a text pad, spreadsheet, or relational database, and imported into practically any statistical package.

Graphing

Sound Ruler's graphing features are a mixed bag. On the one hand, Sound Ruler makes generating attractive figures containing oscillograms, sonograms, or power spectra as easy as a few mouse clicks (Figure 1). Figures can even be updated to show new sounds with a single mouse-click. In addition, Sound Ruler can export any figure as one of seven different graphics formats, including BMP, JPEG, TIFF, and Windows Metafiles, which can then be copied or imported directly into your program of choice. On the other hand, Sound Ruler's graphics capabilities are somewhat limited. For example, arranging or resizing the plots in the figure window is possible, but cannot be done with a simple drag-and-drop, and instead requires the use of cumbersome pull-down menus and scroll bars to resize and reposition the plots. In addition, adjusting the axis limits in the figures cannot be done after the figure is generated, and instead must be done in the main analysis window before generating the figure, and doing so is not as simple as clicking on the axis and entering new values. For those of us who no longer function properly outside the comforts of a click-drag-and-drop world, customising the plots in a figure requires a little effort. Future versions of Sound Ruler that make customising figures even more user friendly would be welcomed.

How Does Sound Ruler Work?

Sound Ruler is organised into 5 hierarchically nested levels of time: Analysis-File-Section-Call-Pulse. A new Analysis is started each time the program is started. Any number of sound Files can be included in one Analysis. The user is required to "chunk" each File into Sections of a duration determined by the user. Navigating through the sound File requires selecting different Sections of user-defined duration, which is a little more cumbersome than freely scrolling to any place in the File. Sound Ruler automatically generates an oscillogram of all sounds in a Section (Figure 2A).

After the user is satisfied with the contents of the Section Oscillogram, the program can then be prompted to recognise the Calls and Pulses contained in the Section. Sound Ruler generates automatically an oscillogram, a spectrogram, and a power spectrum of the first recognised Call contained within the Section (Fig. 2B-D). The user has full control over the size of FFT windows, their degree of overlap, and the range of frequencies displayed in the spectrogram and power spectrum. The Calls and Pulses that are recognised by the program are outlined by a blue line in the Section and Call Oscillograms

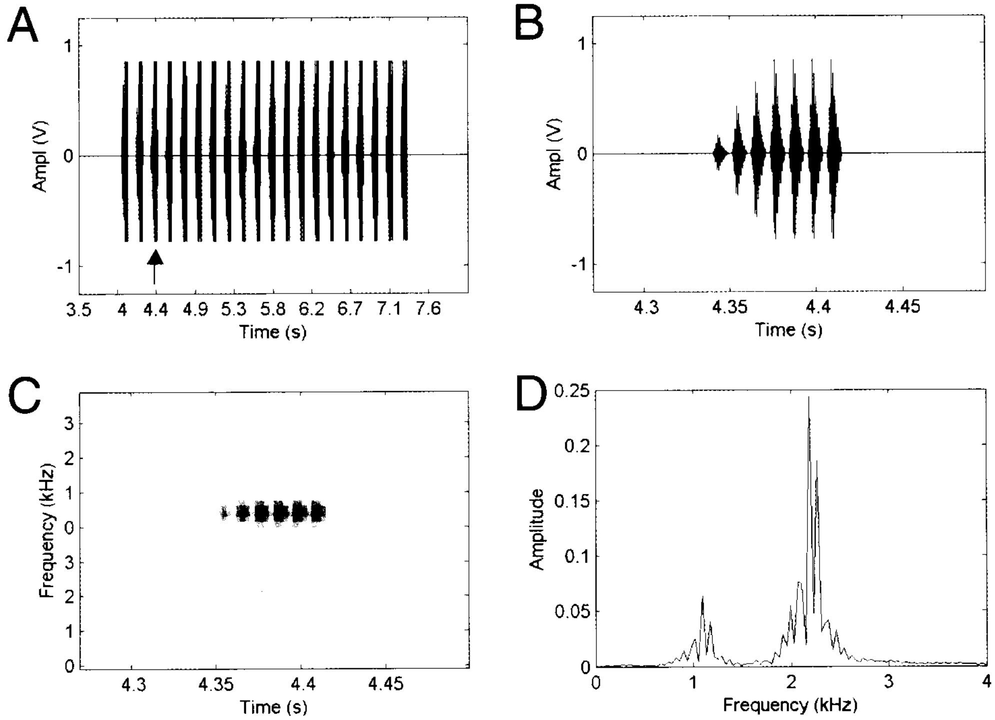

Figure 1. Graphical output of Sound Ruler. (A) An oscillogram of the entire Section showing one synthetic call group from the European treefrog. (B) An oscillogram of the Call section showing a single pulsed call from the call group depicted by the arrow in (A). (C) Spectrogram (FFT 256 points, 78% window overlap) of the pulsed call shown in (B). (D) Power spectrum (FFT 512 points) generated over the peak of the call shown in (B).

(indicated by horizontal arrows in Figures 2A-B), allowing the user to determine the validity of each recognised sound. Calls and Pulses can be recognised based on user-defined expected values (and ranges) of amplitude, duration, interval, or dominant frequency (the latter for Pulses only).

Determining the right recognition settings for detecting Calls and Pulses can take a little time, so some patience is required. However, Sound Ruler has a rather handy feature that allows the user to sample one Call and then let the program determine the most appropriate settings for detecting a Call. Once the program has recognised the appropriate signal units for the analysis, the user can simply scroll through the Section (Figure 2E), accepting or rejecting the recognised elements and their measurement values, which are output to a Results window on the screen (Figure 2F). The user can add comments for each measured signal in a Comments window at the bottom of the screen (Figure 2). Sound-Ruler's on-screen appearance (Figure 2) is somewhat daunting at first, but with a little practice it becomes quite easy to navigate between Sections, between Calls within a Section, and between Pulses within a Call.

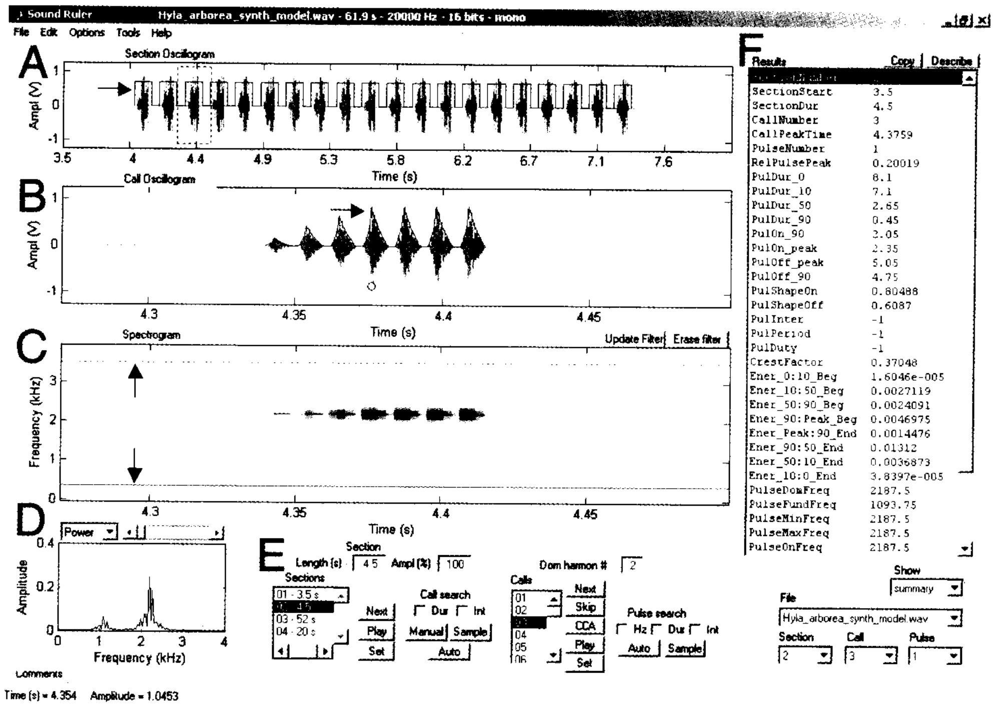

Figure 2. Sound Ruler's on-screen appearance. The large block letters have been added to indicate (A) the Section oscillogram, (B) the Call oscillogram, (C) the Spectrogram, (D) the power spectrum, (E) the control panel for navigating between Sections and between Calls within a Section, and for determining and adjusting the recognition settings, and (F) The Results window. The solid lines in (A) and (B), indicated by horizontal arrows, outline the recognized calls. The dashed-line rectangle in (A) shows the Call that is currently depicted in the Call oscillogram in (B). In (C), the vertical arrows show the limits of a band-pass filter function

How Well Does Sound Ruler Work?

I evaluated the Sample feature, which can be used to let the program determine the most appropriate settings for recognising signal units, using the calls of two species of frogs, the European treefrog, *Hyla arborea* (Hylidae), and the North American bullfrog, *Rana catesbeiana* (Ranidae). The European treefrog produces discrete call groups of 10-20 calls comprised of short (7-10 ms) pulses (Figure 1; Friedl & Klump 2002). The Sample feature was able to determine settings that recognised most calls on the very first try. After an hour of additional tinkering, I was able to adjust the settings so that all Calls and their individual Pulses were recognised. I also had little problem getting the program to recognise entire call groups and the individual calls as the two levels of analysis. Occasionally the pulses of calls from frogs calling in the background were also recognised, but Sound Ruler allows the user to exclude such cases from the analysis. In contrast to the treefrogs, bullfrogs produce relatively long (500-900 ms) call notes organised into calls that are spaced at intervals of several minutes, each containing 3-7 notes spaced about 500-700 ms apart (Bee & Gerhardt 2001). Although I only spent one entire day trying to get the Call Recognition feature to recognise natural and synthetic bullfrog calls, I was unsuccessful in getting the program to recognise either individual call notes or entire calls. This may have been a question of time constraints (i.e., impatience on my part), as the program documentation claims that Sound Ruler can recognise such signals.

I evaluated the performance of Sound Ruler by analysing a sample of synthetic frog calls, in which I determined all of the relevant acoustic properties. I used a custom-designed synthesis program (available from J. J. Schwartz) to create a synthetic call group of the European treefrog. The call group contained 20 calls, each comprised of 7 pulses with identical temporal properties (Figures 1 & 3). The synthesised values for a number of acoustic properties commonly analysed in frog communication studies were manually verified using different software and are reported in Table 1. I used Sound Ruler's automatic call recognition feature to analyse one call group in each of two conditions. First, I analysed the synthetic call group in the absence of any background noise (Figure 3A). Second, I created a condition in which I added background noise to the synthesised call groups by digitally mixing the natural calls of two different *H. arborea* males (recordings provided by Thomas Friedl) at low amplitudes (Figure 3B). This latter condition simulated the situation in which calls from a target male were recorded under more natural conditions that included other nearby calling individuals. The calls depicted in Figure 3B would likely be considered "good recordings" by most frog researchers working in populations with moderate to high densities of calling males. I used Sound Ruler's feature of limiting the signal units that are analysed to exclude recognised pulses that were clearly from one of the background calls. In ambiguous cases, I made the same decision about accepting or excluding a pulse that I would have made if I were manually analysing natural calls.

As is clear in Table 1, Sound Ruler did a nice job in the no-noise condition. For most of the temporal call properties analysed, the measured value deviated from the synthesised value by less than 1 ms, and for spectral properties, the values deviated by less than 1% of the synthesised values. Moreover, Sound Ruler's measurements of the 140 identical pulses did not vary between pulses or calls. Hence, in what approximated perfect recording conditions with an optimal signal-to-noise ratio, Sound Ruler's accuracy and precision were quite good. Sound Ruler's performance in measuring temporal properties deteriorated considerably, however, when synthetic calls were embedded in natural levels of background noise. In general, the program was still quite good at correctly recognising the presence of pulses in the calls, although it often recognised pulses from background calls too, which I was able to exclude from the analysis. The deterioration in performance was due to the program's inability to correctly identify the onsets and offsets of individual pulses. The accuracy and precision of measured values could have been somewhat improved if I had excluded analyses of the first pulse, which was often overlapped to a large degree by background calls. Of course, lower performance in the noise condition is not a problem unique to Sound Ruler.

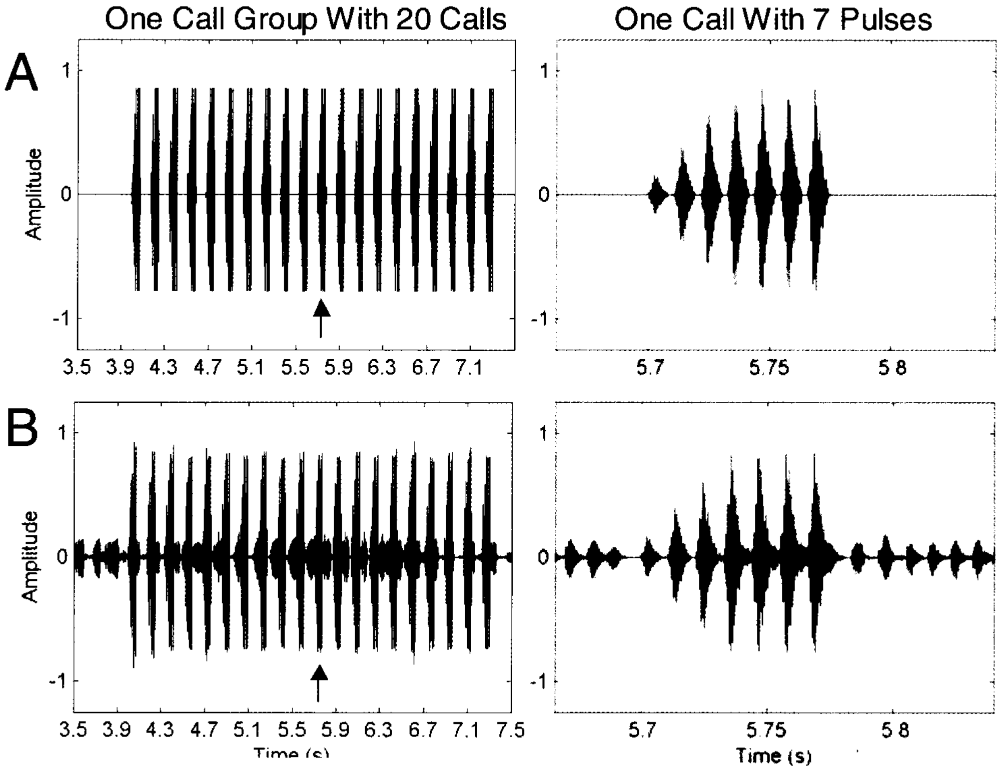

Figure 3. Oscillograms of a synthetic call group (left) and a single synthetic call (right) from a European treefrog *Hyla arborea* in two noise conditions. The arrows in the left oscillograms indicate the position of the call depicted in the right oscillograms. The oscillograms on the left and right depict the same call group and call, respectively. In (A), the synthetic call group was analysed in the complete absence of background noise, giving an optimal signal-to-noise ratio. In (B), the synthetic call group was digitally mixed with the natural calls of two *H. arborea* males to simulate field recordings with more realistic signal-to-noise ratios.

Humans and other automatic recognition programs have similar problems in noise conditions, although humans might be able to make better informed judgements than the program about the onsets and offsets of signals. Thus, an ability to manually adjust the times of the automatically recognised onsets and offsets in Sound Ruler might be a useful implementation in future versions.

Notice in Table 1 that many of the coefficients of variation in the measured properties in background noise were more than 30%, even though there was no variation in the actual signal. In frog communication, accurately measuring the within-individual and among-individual variation in acoustic signals has been fundamental in understanding the mechanisms of sexual selection by female choice (e.g., Gerhardt 1991; Howard & Young 1998) and of

TABLE 1

Synthesised values of call properties for a single *Hyla arborea* call group containing 20 pulsed calls with 7 pulses each in comparison to mean (CV) values measured by Sound Ruler with and without natural background noise.

Call Property	Synthesised Values	Mean (CV) Measured Values	
		Without Noise	With Noise
Number of pulses / call	7	7 (0)	7.5 (0.24)
Pulse duration (ms)	9.0	8.1 (0)	6.9 (0.26)
Pulse duration between 50% amplitude points (ms)	4.0	2.65 (0)	1.1 (0.36)
Pulse rise time (ms)	3.0	2.4 (0)	2.3 (0.48)
Pulse fall time (ms)	6.0	5.1 (0)	4.2 (0.31)
Pulse period (ms)	11.0	11.0 (0)	10.1 (0.39)
Inter-pulse interval (ms)	2.0	3.9 (0)	3.8 (0.79)
Pulse duty cycle	0.82	0.65 (0)	0.76 (0.88)
Dominant frequency (Hz)	2200	2189 (0)	2229 (0.05)
Fundamental frequency (Hz)	1100	1094 (0)	1114 (0.05)
Relative amplitude of fundamental frequency (dB)	-6	-6 (0)	-5.9 (0.10)

vocally mediated neighbour-stranger discrimination in territorial species (Bee & Gerhardt 2001; Bee 2003). Given Sound Ruler's apparent difficulty in accurately and precisely measuring calls embedded in what might be considered a decent signal-to-noise ratio in field recordings, one would have to be extremely careful in using this program to analyse their recordings for the purpose of estimating within-individual and among-individual variation in acoustic signals. Users should always verify with manual measurements that Sound Ruler's automatic recognition and analysis features are making reliable and satisfactory measurements.

Summary

Sound Ruler is an analysis package that will likely be of interest to researchers who routinely spend loads of time measuring the properties of short, repeated acoustic elements, like the pulses of frog calls and cricket songs. It is questionable whether researchers working with more complex signals, like birdsong or primate vocalisations, will find this program that useful. If the right settings can be determined, and if the signal-to-noise ratios in field recordings is high enough, then Sound Ruler has the potential to accurately and precisely measure acoustic properties in a small fraction of the time that would be required for manual measurements. Hence, some students and researchers might find that Sound Ruler will really speed up their work. The program does, however, have some drawbacks, and there appear to be a few remaining bugs in the code. For example, in the course of getting to know the program and performing the analyses described above, I crashed the program 9 times. While Sound Ruler's ability to recognise acoustic events has great potential, the program is less practical if the user wishes to make manual measurements. The program does have a mouse-logging feature, but I did not find this feature that useful. The program could be improved by giving the user control over adjustable, sliding cursors that can be positioned at any location in the sound file with the option of reading in an on-screen Results window the cursor positions, and the difference between the two cursor positions, in both the spectral and temporal domains. In addition, it would be nice if the user could over-ride the program's automatic determination of the onsets and offsets of signals, while still allowing the program to compute the more than 50 measurements it is capable of over the user-defined duration of the signal. In conclusion, the utility of Sound Ruler will depend on the type of signals to be analysed, and the signal-to-noise ratios of the recordings. If the signals are rather simple and repetitive, and the field recordings of superb quality, then Sound Ruler may facilitate signal analysis by drastically reducing the time spent taking measurements. Given that the program is currently free and has an open code, the program is definitely worth checking out.

REFERENCES

Bee, M.A. (2003). A test of the "dear enemy effect" in the strawberry dart-poison frog (*Dendrobates pumilio*). *Behavioral Ecology and Sociobiology*, **54**, 601–610.

Bee, M.A. & Gerhardt, H.C. (2001). Neighbour-stranger discrimination by territorial male bullfrogs (*Rana catesbeiana*): I. Acoustic basis. *Animal Behaviour*, **62**, 1129-1140.

Friedl, T.W.P. & Klump, G.M. (2002). The vocal behaviour of male European treefrogs (*Hyla arborea*): Implications for inter- and intrasexual selection. *Behaviour*, **139**, 113-136

Gerhardt, H.C. (1991). Female choice in treefrogs: static and dynamic acoustic criteria. *Animal Behaviour*, **42**, 615-636.

Gerhardt, H.C. & Huber, F. (2002). *Acoustic communication in insects and frogs*. Chicago: Chicago University Press.

Gridi-Papp, M. (2003). SoundRuler, acoustic analysis and graphing. http://soundruler.sourceforge.net/.

Howard, R.D. & Young, J.R. (1998). Individual variation in male vocal traits and female mating preferences in *Bufo americanus*. *Animal Behaviour*, **55**, 1165-1179.

Mark A. Bee
AG Zoophysiology and Behaviour
Institute for Biology
Carl von Ossietzky Universität
D-26111 Oldenburg
Germany