ExpoChirpToolbox: a Pure Data implementation of ESS impulse response measurement

Katja Vetter
katjavetter@gmail.com

Serafino di Rosario
bassik76@gmail.com

Rotterdam/London, July 2011

abstract

Impulse response (IR) measurements for assessing acoustic properties of spaces have been done since long time, with various techniques and test signals. Of these, the Exponential Sine Sweep method, successfully advocated by Angelo Farina [2] in the last decade, has outstanding properties. This paper describes an open source Pure Data implementation of the ESS method. A dedicated Pd class [expochirp∼] is introduced, as the basis of a modular tool chain performing subsequent tasks in measurement and analysis procedures. The test signal is refined in several ways, respective to existing implementations. Mathematical aspects of ESS method and the ExpoChirp implementation are discussed in detail.

keywords

IR measurement, ESS method, room acoustics

1 impulse response measurement with Pure Data

In recent years, auralization and spatialization techniques have gained fair attention of open source programmers, and Pure Data programmers in particular. On the measurement and analysis side of the topic however, one is still dependent on proprietary software, single-platform solutions, or a combination of partial solutions. The ExpoChirp project for Pure Data aims to be a cross platform open source implementation of up to date IR measurement and analysis techniques. ExpoChirp modules show graphical data in all stages of the measurement procedures. This enables users to assess the working, specifications and quality of the implemented routines, which is useful for education purposes.

The ExpoChirp project was initiated in 2010 at the Pure Data forum. So far, we have realized basic modules for test signal generation, test response recording, IR editing, spectral analysis, and convolution filter testing. For all these modules, the Pd class [partconv∼] for partitioned convolution, written by Ben Saylor in 2003, is of crucial importance. The ExpoChirp project has received the interest of potential co-developers, and future efforts will focus on subtopics like RT-analysis, STI-analysis, multi-channel room acoustics measurement, Ambisonic format conversion, non linear system analysis and convolution. Descriptions and tutorials of modules will be published and updated online [1]. This paper discusses principles of the ESS method and the way it is implemented in the toolbox.

1.1 Exponential Sine Sweep method

Our choice for ESS is related to the purpose of the measurement tool: analysis of audio systems, especially (but not exclusively) room acoustics. Simply stated, we want to replace the human ear with a more objective measurement tool, and quantify the results. For human perception, octaves are blocks of interest with equal span, and this is reflected in the exponential sine sweep or chirp, which is used as the test signal. Starting from a frequency well above DC, the frequency doubles repeatedly over identical time intervals. Time and energy is spent on low frequencies more than on high frequencies. With a linear chirp instead, or any other (almost-)white signal, very little information would be gathered on frequencies which are prevalent in perception. For example, over a range of 10 octaves, a white signal spends $1/2^{10}$ of it's time and energy on the lowest of these octaves, that is less than one thousandth.

The ordering of frequencies from low to high in a continuous rise has additional advantages. In the first place, it is easy to deconvolve because it's spectral inverse can be computed with a relatively simple formula. Secondly, the low-to-high arrangement of frequencies has the effect that possible harmonic distortions produced by acoustic / analog test equipment appear before time zero in the deconvolved result, and they do not interfere with the causal test result. A properly designed exponential chirp will even arrange harmonic distortion responses in such a way that each higher order response can be quantified separately, and the 'distortion recipe' identified.

1.2 regular ESS implementation

In 2000, Angelo Farina proposed the IR measurement method [2] which later received the proper name ESS, and which has gained wide popularity. The test signal in this method can be defined by the following formula:

$$x[n] = sin\left[\frac{\omega_1 * N}{ln(\omega_2/\omega_1)} * \left(e^{\frac{n}{N} * ln(\omega_2/\omega 1)} - 1\right)\right] \tag{1.1}$$

where:
ω_1 is the normalized start frequency in radians
ω_2 is the normalized stop frequency in radians
n is the sample index
N is the number of samples

The signal has a constant amplitude over time, and because the frequency rise is exponential, the magnitude spectrum shows a decay, of 6 dB per octave to be precise. This magnitude decay must be compensated when the impulse response is calculated with the deconvolution procedure. Because compensation is done in digital domain, it does not alter the signal-to-noise ratio which favours the low frequency range in the test signal.

2 test signal refinement

Even while the exponential chirp has so many qualities which make it an excellent test signal, it is not absolutely perfect. The spectrum is characterized by frequency ripple at the extremes, and overshoot may exceed +5dB in certain cases. In the following sections the cause and nature of the ripple aspect will be examined, and strategies will be proposed to guarantee a minimum-ripple chirp, with overshoot well below +1dB.

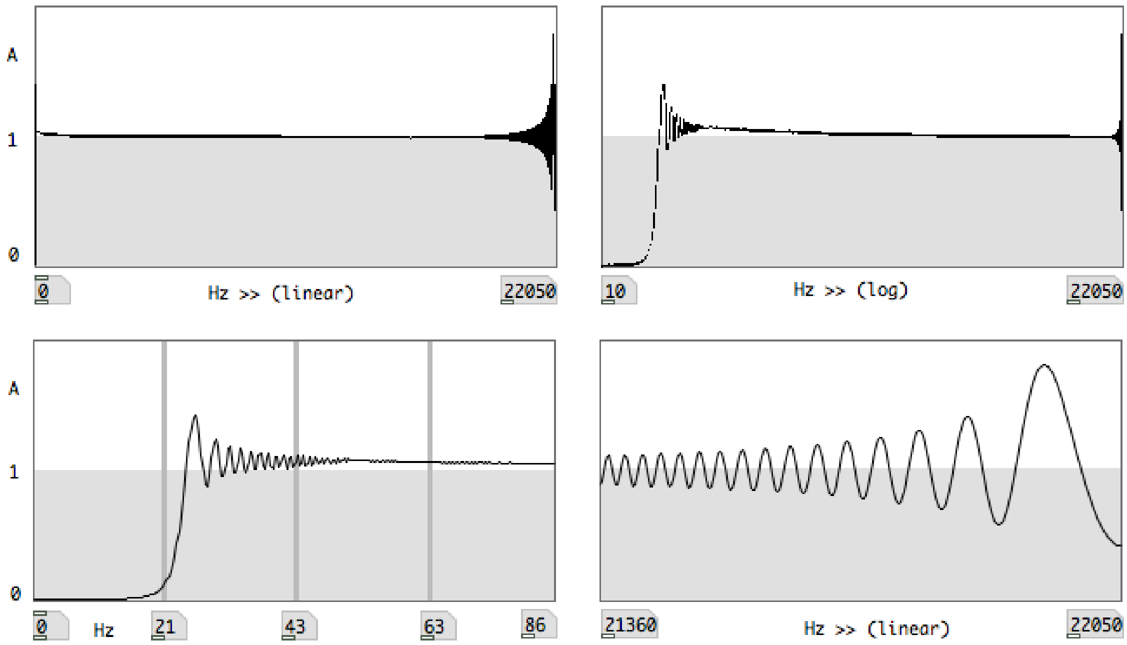

Figure 1: *magnitude spectrum of deconvolved exponential chirp*

The chirp start and stop boundaries represent undesired frequencies, causing partial phase cancellations and summations over the signal which has in itself a constant amplitude and a constant frequency rise. A chirp-start or -stop in cosine phase shows heavier boundary effects than a start or stop in sine phase. The regular formula only takes care of the start phase, while the stop phase is a function of start frequency, stop frequency and chirp length. Figure 1 shows spectrum plots of a deconvolved exponential chirp which happened to end in an unfavourable phase. The deconvolved chirp was generated by Angelo Farina's Aurora plugins, and loaded in a Pd patch for inspection. No windowing was done on the chirp.

Farina has discussed this problem of exciting the sound system with a step function when the chirp ends in an unfortunate phase [3]. A fade-out window can be applied to force a smooth decay of the high spectrum end, but such a fade-out is in the deconvolved chirp represented by filter coefficients with substantial value around the central pulse, which should be a unit impulse ideally. To avoid both these effects, step function and time domain artifact, Farina proposes to manually cut the chirp at the latest zero-crossing before it's termination. With an alternative chirp definition however, the problem can be solved in a more systematical way, as will be shown below.

2.1 designing a phase-controlled exponential chirp in Pure Data

Chirp generator design in Pd started out with a small inconvenience unrelated to the above mentioned ripple issue. Numerical data flow between Pd 'objects' currently allows only one data type: single precision (32 bit) floating point format. The desired length of the chirp, up to a million samples, requires a much better resolution of increments in the exponential phase curve, which is an intermediate calculation step. Meanwhile, to properly represent the resulting chirp signal, 32 bit floating point is fortunately sufficient. Therefore a new Pd class [expochirp∼] was written in C, calculating with double precision internally, while rendering the chirp and it's inverse as 32 bit audio signals.

Initially, the chirp was calculated following the regular formula as written in equation 1.1. Test patches were made to assess the chirp and it's deconvolution in graphical detail, and conspicuous ripples similar to those shown in figure 1 became a source of irritation. Although ripples are only present at the low and high frequency extremes, which are not always of concern, we would rather like to provide a reliable test signal for all possible measurement situations. Part of the solution was found in an alternative chirp definition, which makes the chirp end in sine phase at Nyquist by definition, thus minimizing high frequency overshoot.

For the purpose of this discussion, the traditional formula as presented in equation 1.1 is repeated here, and components of it are distinguished and labeled.

$$x[n] = sin\left[\frac{\omega_1 * N}{ln(\omega_2/\omega_1)} * \left(e^{\frac{n}{N} * ln(\omega_2/\omega 1)} - 1\right)\right] \tag{2.1}$$

$$\boxed{\frac{\omega_1 * N}{ln(\omega_2/\omega_1)}} \qquad \text{scaling factor} \tag{2.2}$$

$$\boxed{e^{\frac{n}{N} * ln(\omega_2/\omega 1)}} \qquad \text{exponential curve} \tag{2.3}$$

$$\boxed{-1} \qquad \text{constant} \tag{2.4}$$

The exponential curve has remarkable values at indices n=0 and n=N:

$$e^{\frac{0}{N} * ln(\omega_2/\omega_1)} = e^0 = 1 \qquad \text{and} \qquad e^{\frac{N}{N} * ln(\omega_2/\omega_1)} = e^{ln(\omega_2/\omega_1)} = \omega_2/\omega_1$$

We observed that it would be advantageous if (ω_2/ω_1) were an integer, and the scaling factor an integer multiple of $(2 * \pi)$. The chirp will then start and stop in the same phase; phase zero, if we drop the constant -1. The most obvious way to make (ω_2/ω_1) an integer, is by creating chirps with an integer number of octaves P. Start frequency ω_1 becomes $(\omega_2/2^P)$, and (ω_2/ω_1) reduces to 2^P.

Since the optimal approximation of a unit impulse (after deconvolution) is obtained with the normalized stop frequency at π radian (the Nyquist frequency), we use π as the default for ω_2. The chirp definition in [expochirp~] is given as an equation set (2.5).

$$x[n] = sin\left[\frac{(\pi/2^P) * L}{ln(2^P)} * e^{\frac{n}{N} * ln(2^P)}\right] \qquad \text{and} \qquad \frac{(\pi/2^P) * L}{ln(2^P)} = M * \pi * 2 \tag{2.5}$$

where:
P = an integer number of octaves
L = ideal chirp length (floating point value)
N = chirp length (= L rounded to integer)
M = a positive non-zero integer

This is implemented in ExpoChirpToolbox by letting the user select the desired number of octaves and a maximum chirplength first, from where L and N are computed. With the chirp ending in sine phase, overshoot and ringing at the high spectrum side are systematically kept at a minimum, well under +1dB (see figure 4 at the end of the article).

2.2 instantaneous frequency curve

Because of the rounding of L to N in the phase-controlled exponential chirp, there will be a small deviance in the resulting frequency range, as compared to the desired range. The instantaneous frequency curve must be inspected to see how important the error is. The frequencies are represented by the first derivative of the phase curve. The derivative of an exponential with base e, and a constant in the exponent, is simply that exponential multiplied by the constant:

$$(e^{ax})' = ae^{ax}$$

The constant in the exponent is $ln(2^P)/N$ in our case. The frequency curve is:

$$x'[n] = \frac{(\pi/2^P) * L}{ln(2^P)} * \frac{ln(2^P)}{N} * e^{\frac{n}{N} * ln(2^P)} = (\pi/2^P) * \frac{L}{N} * e^{\frac{n}{N} * ln(2^P)} \tag{2.6}$$

This shows that every frequency is multiplied by L/N, where it would normally be $N/N = 1$. The frequency deviance is not a shift but a frequency stretch. Chirp length N is very large in real test conditions, typically 100,000 to 1,000,000 samples. The frequency deviance then, is in the order of a factor 0.00001. The advantage of a phase-controlled chirp outweighs this minor frequency stretch for practical purposes, but for cases with extremely short chirps the importance of the impurity should be considered.

With the phase-controlled chirp comprising an integer number of octaves P, the chirp has P equally-sized segments which all start and end in sine phase. As a consequence, frequencies with octave intervals or other harmonic intervals will appear with identical phase (more on this in section 2.4). In figure 2, the first two octaves of such a chirp are shown. Each subsequent octave segment has a doubled number of cycles.

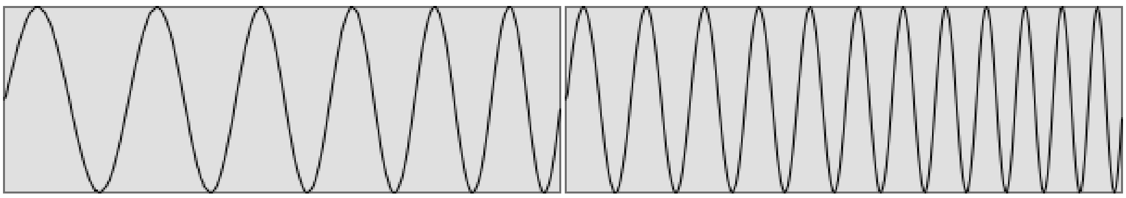

Figure 2: *first two octaves of a phase-controlled exponential chirp*

2.3 inverse of the exponential chirp

If each subsequent equally-sized time interval in the chirp has a doubled number of cycles, then it must be the case that there is an amplitude decay from the low to high frequencies, not in time domain but in the spectrum. For any frequency, the second harmonic has half the amplitude of the fundamental, a 6 dB per octave decay. To deconvolve this chirp, not only must the frequencies be conjugated (by time reversal of the signal, pivoting over index N), but the amplitude effect must be inverted as well. The general definition of a 6 dB per octave decay is 2^{-p} where p is the octave index. This is a curve starting

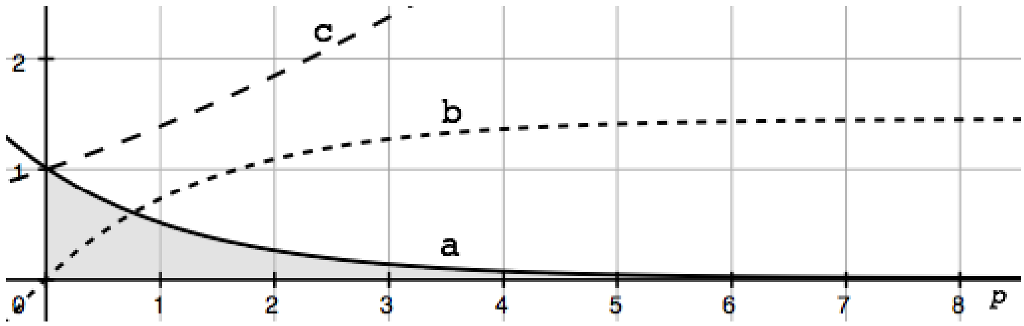

Figure 3: *curve (a) 2^{-p} decay, (b) integral of the decay curve (eq. 2.7), (c) amplitude compensation factor (eq. 2.8)*

with value 1 at index 0, and applying the curve to the inverse chirp causes an amplitude loss respective to the original chirp, which must be compensated. An exact amplitude compensation factor can be derived from the integral of the amplitude decay curve over the interval $[0, P]$. This integral is:

$$\int_0^P dp(2^{-p}) = -\frac{2^{-P}}{ln(2)} + \frac{1}{ln(2)} \tag{2.7}$$

From equation 2.7, the multiplicative inverse of the normalized integral over the interval $[0, P]$ is derived. This yields the constant C which compensates amplitude loss over the decay curve 2^{-p} (eq. 2.8). Figure 3 shows the functions in a plot.

$$C = P * \left(\frac{ln(2)}{1 - 2^{-P}}\right) \tag{2.8}$$

The time-reversed chirp, amplitude decay and amplitude compensation factor together form the inverse chirp definition as formulated in equation 2.9.

$$x^{-1}[n] = x[N - n] * (2^{\frac{P}{N}})^{-n} * \left(\frac{P * ln(2)}{1 - 2^{-P}}\right) \tag{2.9}$$

This calculates an inverse chirp with average amplitude equal to the forward chirp, being $1/\sqrt{2}$. When convolving chirp with inverse chirp, the result is an approximate pulse with height N/2. A 2/N normalisation factor is applied after convolution, in order to get an IR with the unit impulse as a reference.

2.4 nonlinear system aspects and phase-controlled chirp

It must be mentioned that the notion of a phase-controlled exponential chirp is not in itself new. Prof. Angelo Farina pointed us to the fact that the use of a phase-synchronized exponential swept-sine signal is introduced by Antonin Novak for the purpose of nonlinear system identification [4]. If higher order impulse responses (products of harmonic distortions) are to be accurately traced, harmonics in the chirp must be phase-aligned, otherwise phase cancellations and summations may corrupt the analysis result. Comparison of Novak's chirp definition and the method proposed in this paper reveals that in both cases the calculation of an ideal chirp length (given the other parameters) effectuates the phase alignment of harmonics. Since Novak's definition does not require an integer number of octaves and may end in any phase, it does not address the ripple issue.

In [expochirp∼], the overall phase alignment is a consequence of our method to minimize overshoot at the high spectrum end. At the moment of this writing, we have not yet worked on routines for nonlinear system identification in the ExpoChirp toolbox. However, with future extensions of the toolbox in mind, it is a fortunate coincidence that [expochirp∼] generates a chirp suitable for this purpose.

2.5 further ripple reduction: fade-in window

Though the phase-controlled chirp definition minimizes frequency ripple at the high spectrum end of the chirp, this intervention did not do anything for the low frequency ripple. Even when the signal starts and stops in sine phase, the overshoot at the start frequency is around +3dB while at the stop frequency it is below +1dB, in the cases we tested. A fade-in window can help to produce a smoother spectrum. Remarkably, a fade-in window does not raise coefficient values in the deconvolved chirp, in contrast with the case of a fade-out window as it was mentioned earlier.

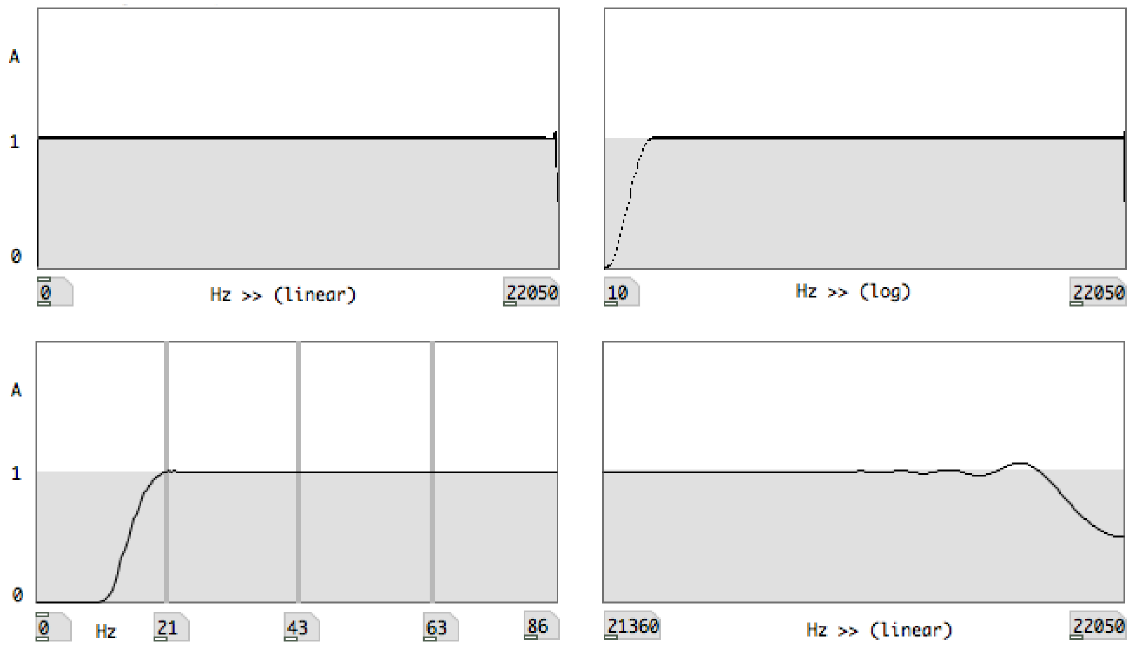

Figure 4: *magnitude spectrum plots of deconvolved fade-in-windowed and phase-controlled exponential chirp*

Experiments with a sine-shaped fade-in learnt that the window must be fairly long to be sufficiently effective, in the order of a second. At the low frequency side, this is not problematic because the chirp can be easily extended downwards. We decided to implement a optional fade-in window stretching over exactly the first octave in the chirp. For test signals with ten or more seconds length, this has a very beneficial effect. Figure 4 shows spectrum plots of a phase-controlled chirp with the one octave long fade-in. The chirp was calculated to comprise 11 octaves, of which the first one is modulated by the window. The transition band is less steep and the flat spectrum region starts exactly 10 octaves below Nyquist. With this systematic window definition, the exact start point of the flat spectrum range is always known.

3 conclusion

The test signal refinements presented in this paper illustrate that it is very well possible to design a meticulous IR measurement procedure within the Pure Data environment. At the moment of this writing, ExpoChirpToolbox merely offers basic functionality: test signal generation, IR capturing and editing, spectral analysis. In order to match the operational speed and extensive functionality of existing tools like the Aurora plugins or Easera software, a lot of work still has to be done. Since the first results are rather promising, the authors express the hope that the ExpoChirp project will keep evolving, and be of use for professionals and students by virtue of it's openness and quality.

4 acknowledgements

The ExpoChirp project is powered by the work of Miller Puckette and the Pure Data community, providing the wonderful Pure Data dsp environment and the means of communication required to set up collaboration projects. We are also indebted to Angelo Farina for his extensive publications on the ESS method, and for his comments on the ExpoChirp project.

references

[1] www.katjaas.nl/expochirp/expochirp.html

[2] Angelo Farina: "Simultaneous Measurement of Impulse Response and Distortion with a Swept-Sine Technique", *pre-print of the 108th AES Convention*, 19-22 February 2000.

[3] Angelo Farina: "Impulse Response Measurements", *paper for the 23rd Nordic Sound Symposium*, 27-30 September 2007.

[4] Antonín Novák, Laurent Simon, František Kadlec and Pierrick Lotton: "Nonlinear System Identication Using Exponential Swept-Sine Signal", *IEEE Transactions on Instrumentation and Measurement*, 2009-10-30.