

A low-cost, high-performance, digital signal processor-based lock-in amplifier capable of measuring multiple frequency sweeps simultaneously

Maximiliano Osvaldo Sonnaillon^{a)} and Fabián Jose Bonetto^{b)}
Laboratorio de Cavitación y Biotecnología, San Carlos de Bariloche (8400)-Argentina

(Received 28 March 2003; accepted 2 December 2004; published online 20 January 2005)

A high-performance digital lock-in amplifier implemented in a low-cost digital signal processor (DSP) board is described. This lock in is capable of measuring simultaneously multiple frequencies that change in time as frequency sweeps (chirps). The used 32-bit DSP has enough computing power to generate $N=3$ simultaneous reference signals and accurately measure the $N=3$ responses, operating as three lock ins connected in parallel to a linear system. The lock in stores the measured values in memory until they are downloaded to the a personal computer (PC). The lock in works in stand-alone mode and can be programmed and configured through the PC serial port. Downsampling and multiple filter stages were used in order to obtain a sharp roll off and a long time constant in the filters. This makes measurements possible in presence of high-noise levels. Before each measurement, the lock in performs an autocalibration that measures the frequency response of analog output and input circuitry in order to compensate for the departure from ideal operation. Improvements from previous lock-in implementations allow measuring the frequency response of a system in a short time. Furthermore, the proposed implementation can measure how the frequency response changes with time, a characteristic that is very important in our biotechnological application. The number of simultaneous components that the lock in can generate and measure can be extended, without reprogramming, by only using other DSPs of the same family that are code compatible and work at higher clock frequencies. © 2005 American Institute of Physics. [DOI: 10.1063/1.1854196]

I. INTRODUCTION

A lock-in amplifier is a commonly used instrument for measuring low-level periodic signals in the presence of noise. It works as a very narrow-band filter that measures, ideally, only one Fourier component of the system response. Also ideally, the measurement is not affected by any other frequency. The lock in works synchronously with a reference signal and measures the phase and magnitude changes when the reference signal is passed through an unknown linear system.

At first, lock-in amplifiers were made using analog electronics, which had very good performance for most of the applications. However, fast growing digital electronic technology made possible the implementation of digital lock ins, which have better performance than the analog lock ins in certain aspects, and extends their field of application. Commercial digital lock ins have been available in the market,^{1,2} since a few years ago, and most of them use devices known as digital signal processors (DSPs). DSPs are microprocessors designed to efficiently compute a large number of mathematical operations per second. This and other characteristics make them a very convenient option for processing signals digitally.³

Implementations of digital lock in amplifiers have been previously reported. They include designs with discrete electronic devices,⁴ standard computers,⁵⁻⁷ a general-purpose microprocessor,⁸ and a DSP board.⁹ However, none of them consider systems where X_R and X_C (real and imaginary parts of the electrical impedance, respectively) can change with frequency and time. In order to characterize these types of systems, a lock in capable of measuring multiple frequency sweeps simultaneously is presented in this work.

Our need for designing and implementing a lock-in amplifier arose from a biological application.¹⁰⁻¹³ We need to measure the electrical impedance of a group of growing cells. We apply an excitation that consists of a very small ac current [$\approx 1 \mu\text{A}$ root-mean square (rms)] flowing through the cells. We measure the in-phase and in-quadrature voltage drop across the cells as a function of frequency and time. Thus, with different ac frequencies, we can obtain the frequency response of the cells and compute the system parameters using a mathematical model.¹⁰

More specifically, we need a lock in capable of measuring the in-phase and in-quadrature signals as the frequency of the excitation changes. We call this instrument a multifrequency lock-in amplifier.

As the cells grow their electrical impedance changes with time, and this change has to be measured as well. Using a standard lock-in amplifier, we measured the response to different frequencies by programming the function generator that provides the reference. Using a standard lock in, the reference frequency can be changed too slowly for our pur-

^{a)}Also at: Instituto Balseiro (UNCu), San Carlos de Bariloche (8400)-Argentina, and CONICET; electronic mail: msonnaillon@ing.unrc.edu.ar

^{b)}Also at: Instituto Balseiro (UNCu), San Carlos de Bariloche (8400)-Argentina, CONICET and CAB/CNEA; electronic mail: bonetto@ib.edu.ar

poses. A measurement of the whole frequency range (20 Hz to 20 kHz) takes an unacceptably long time.

The implementation of our digital lock-in amplifier allowed us to add some improvements that are important for our experiment and may be useful for other experiments. These improvements include the capability of measuring multiple simultaneous frequencies (the system is considered linear) that can change in time as sweeps. Then, we are able to obtain the frequency response of the system as a continuous function with a measurement that takes a short time.

The article is organized as follows. First, a brief description of the basic theory of lock-in amplifiers is given. Then, a description of how the proposed lock in works and how it was implemented is presented. Finally, some representative experimental results, that validate the advantages of using this lock in instead of using a standard one, are given.

II. HOW A LOCK IN AMPLIFIER WORKS

The theory of lock-in amplifiers is well known and is explained in other works.^{8,9} Here, we give a brief introduction to the basic principles.

A basic lock-in amplifier consists of a reference signal, an amplifier, a phase sensitive detector (PSD), and a low-pass filter. The reference signal is passed through an unknown system that produces a modification of the signal phase and magnitude. The system also adds noise to the signal. This signal is amplified to a level adequate for the PSD, which multiplies it by the internal reference. The result of this product is the sum of a component with the double frequency, a component of zero frequency (dc signal), and the noise that was added in the system. Filtering the ac frequencies and only keeping the dc component, the level of the signal that has exactly the same frequency and phase as the reference signal can be measured.

If the input signal is multiplied by a sinusoid that is in quadrature (90° shifted) with the reference, the signal level that is in quadrature can be measured. With both measurements (in-phase and in-quadrature), the magnitude and phase of the measured signal with respect to the reference can be computed.

Since the low-pass filter is not ideal, the filter passes a narrow bandwidth of noise that negatively affects the measurements. In order to reduce noise, it is convenient to have very low cut-off frequencies in the low-pass filters. However, these low cut-off frequencies represent large time constants that make the response of the lock in too slow. A convenient trade-off between noise level and measurement speed must be found for each experiment.

III. HOW THIS LOCK IN WORKS

This lock-in implementation works as N standard lock ins connected in parallel, where N depends on the DSP computing power, and is equal to 3 for this particular implementation. The lock in generates N independent references with programmable frequencies. Each reference has a frequency that changes logarithmically with time while the measured values are stored in memory. In spite of working as N lock ins, this lock in uses a single analog channel for the output

references and a single analog input channel for measuring the resulting signals. Thus, a system can be characterized using the N components simultaneously. This reduces the total measurement time.

The capability of making a sweep of the reference frequency makes possible the measurement of the frequency response as a continuous function, instead of for only a few discrete Fourier components. However, there is a limit on the frequency rate of change that depends on the system frequency response and the time constant of the low-pass filter at the lock-in output.

If the unknown system has a frequency response with a sharp roll off, then the output of the lock in will have to change rapidly (the sharper the roll off, the more rapid the change). In the case of too large time constants, the low-pass filter will attenuate this change, producing measurement errors. So, the user will have to reduce either the time constant of the filters (with the consequence that the lock in will not behave as a narrow bandwidth filter) or the sweep rate (increasing the total measurement time). Usually, users choose a trade off between both.

In a standard lock in, the sweep rate limitation depends mainly on the time constant of its phased locked loop (PLL) at the reference input, which is usually slow. This is one of the advantages of the presented lock in compared to a commercial one.

Another feature of the proposed lock in is the capability to autocalibrate itself before making each measurement. This means that it automatically measures the frequency response of the analog circuitry in the working spectrum of frequency. Thus, the phase and magnitude distortion can be taken into account and compensated for the characterization of the unknown system. This distortion is produced by the amplifiers, the antialiasing and the smoothing filters of the audio codec. They do not have a flat response in the pass-band frequencies for the range of interest.

The calibration is performed making a measurement with the output reference of the lock in connected directly to the input. Thus, any magnitude or phase distortion is attributed to the analog circuits of the codec, and can be eliminated by making a division of the measured frequency response of the unknown system and the known frequency response of the lock-in circuits. After making this calibration, the lock in switches a relay to connect the input to the unknown system, and then starts the measurements. A block diagram of this procedure is shown in Fig. 1. Mathematically:

$$Y_C(s) = X(s) \times H_1(s) \Rightarrow H_1(s) = \frac{Y_C(s)}{X(s)} \quad (1)$$

for the calibration process, and

$$\begin{aligned} Y_M(s) &= X(s) \times H_1(s) \times H_2(s) \\ \Rightarrow H_2(s) &= \frac{Y_M(s)}{X(s) \times H_1(s)} \\ &= \frac{Y_M(s) \times X(s)}{X(s) \times Y_C(s)} = \frac{Y_M(s)}{Y_C(s)} \end{aligned} \quad (2)$$

for each measurement, where $X(s)$ is the input signal, $H_1(s)$

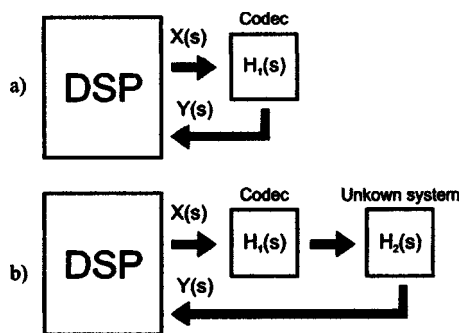


FIG. 1. Block diagram for the calibration and measurement procedure. The determination of the code transfer function is performed using configuration (a). In (b), the configuration to measure the system transfer function is shown.

is the transfer function of the codec, $H_2(s)$ is the transfer function of the unknown system, and $Y(s)$ is the output signal that is measured by the DSP.

All of the parameters of the lock in, such as the rate of change of the frequency sweep, its frequency limits, the number of measurements to store in memory, and the filters' time constant, are programmed through a personal computer (PC) to fit the specific requirements of the experiment.

IV. IMPLEMENTATION

For the implementation of the multifrequency lock-in amplifier, we used a low-cost DSP board from Analog Devices (SHARC DSP EZ-KIT Lite). It has a 32-bit floating point DSP running at 40 MHz (ADSP-21061) with 1 Mbit of on-chip memory [static random access memory (RAM)], and boots through a nonvolatile memory (erasable-programmable read only memory), where the lock-in program resides. It also has a stereo audio codec with the following on-chip features:

- Two 16-bit ADCs and two 16-bit DACs, with a maximum sample rate of 48 kilosamples per second (kps),

- Programmable gain amplifiers and antialiasing filters for the inputs,
- Smoothing filters and programmable attenuators for the outputs, and
- A dynamic range of 70 dB and a total harmonic distortion plus noise (THD+N) of 0.04%.

These characteristics make this codec suitable for our range of frequencies and required resolution. External analog-to-digital converters (ADCs) and digital-to-analog converters (DACs) with their associated electronics (e.g., antialiasing and smoothing filters) can be connected to the DSP to extend the resolution and the frequency range of the lock in.

Once the codec converts to digital, each sample (every $1/f_s$ s) generates an interrupt that runs a subroutine with the necessary computations for the N parallel lock ins. For each lock in, the DSP performs the following operations:

- Generates the reference signals in phase and in quadrature (sine and cosine),
- Multiplies the input signal with these two references,
- Filters the ac components of the calculated products using a low-pass filter with a low value of cut-off frequency, and
- Stores the values of the resulting dc signal which constitutes the output measurement.

A schematic block diagram of the lock in is shown in Fig. 2.

The reference signals are generated internally by the DSP with 32-bit of accuracy, using direct digital synthesis (DDS).¹⁴ This technique is used in almost all commercial digital signal generators and digital modulators, and has the advantages of requiring a short computing time and producing a sinusoidal wave form that is very stable and precise in frequency and amplitude.

For the DDS, we use a 32-bit counter that increments—at fixed time intervals—an amount that depends on the desired wave frequency. This counter is used as the index of a table that has the values of a complete sine cycle.

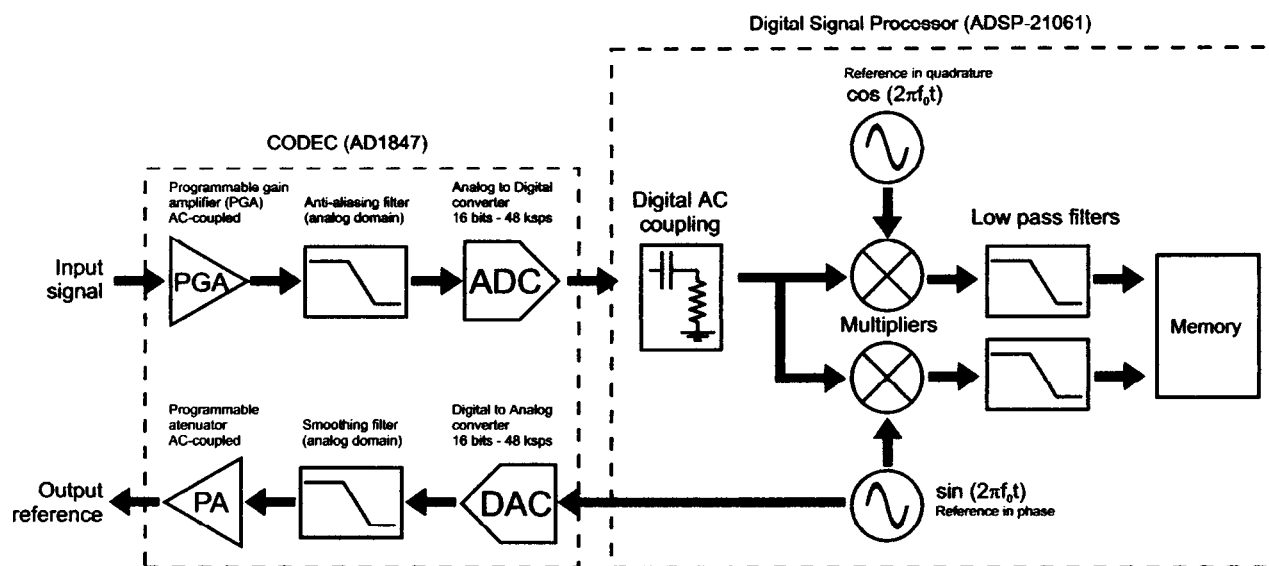


FIG. 2. Schematic block diagram for the DSP-based lock-in amplifier. The dotted-line shows the operations performed by the DSP.

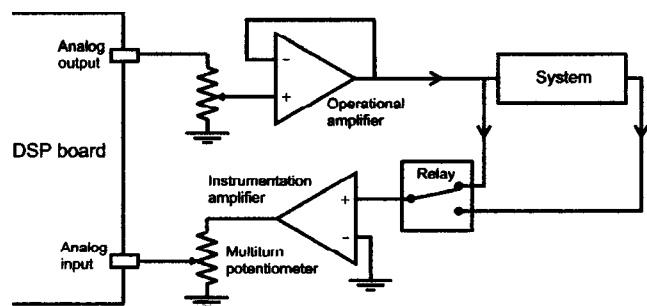


FIG. 3. Schematic diagram of the external circuit used to adequate the signal level to the experiment. Both amplifiers are low noise. The relay is controlled by a DSP digital output signal.

Because we cannot use a table of 2^{32} values (this table would be much larger than the available memory), we only use a 64-element table and linearly interpolate the intermediate values. This produces a wave form that is almost free of harmonics and distortion (with 64 elements, we have numerically computed a THD+N of less than -150 dB).

We chose a wave table instead of a digital oscillator because we want to vary the reference frequency with time. A digital oscillator is simpler, but not as flexible as the wave-table implementation. Using a wave table, in order to change the sine wave frequency, we just need to change the increment of the counter (a few instructions). Using a digital oscillator, we need to recompute the coefficients (many more instructions) for each desired frequency.

The values of the N references are calculated, added, and then converted to a single analog signal using the audio codec that is included in the DSP board. In the analog domain, a standard low-noise operational amplifier (TL071) is used as a current buffer to provide low output impedance. A multiturn potentiometer is used as an attenuator to adjust the signal level injected to the experiment.

Once the reference signal (composed of $N=3$ added sine waves) is passed through the system, it is amplified using a single-chip low-noise instrumentation amplifier (AD621). This amplifier has a programmable gain of 10 or 100. We use potentiometer to attenuate and adjust the signal level to the input range of the codec. The ADC converts this measured signal to a digital value that is used in the processing routine. A schematic diagram is shown in Fig. 3.

In order to obtain a good ac coupling and eliminate the offset of the analog circuits and the ADC, the input signal is ac coupled with a digital resistance/capacitance (RC) network. This network is a first-order high-pass filter with a cut-off frequency of 0.1 Hz, which almost completely eliminates the dc component of the signal. For the internal in-phase reference, we use the already computed sine value, and for the in-quadrature reference, we use the same table with a shift in the index value that corresponds to 90° ($64/4=16$ elements). The multiplication is the simplest part of the algorithm: We use the processor's multiplication instruction that is computed in only one clock cycle. For the low pass filtering, we implemented a recursive filter (also called an infinite impulse response, or IIR) whose coefficients are calculated and can be modified by the PC. A recursive filter is computed with the following difference equation:

TABLE I. Minimum cut-off frequency ratio in recursive filters as a function of the maximum number of poles using single precision (see Ref. 15).

| Cut-off frequency (f_c/f_s) | 0.01 | 0.05 | 0.25 | 0.45 | 0.49 |
|---------------------------------|------|------|------|------|------|
| Maximum poles | 2.0 | 6.0 | 20.0 | 6.0 | 2.0 |

$$y[n] = a_0 \times x[n] + a_1 \times x[n-1] + a_2 \times x[n-2] + \dots + b_1 \times y[n-1] + b_2 \times y[n-2] + \dots \quad (3)$$

This equation consists of products and additions. These operations represent a low computational load for the DSP because it can compute them in parallel. For example, a two-pole IIR filter is performed in 10 cycles or $0.25 \mu\text{s}$ @ 40 MHz.

The problem with using a recursive filter with a very low (or very high) cut-off frequency is that the “ a ” coefficients become too small compared to the “ b ” coefficients. When computing the filter, the finite word length produces a high truncation error that affects the filter performance and can make it unstable. Furthermore, the use of high-order filters also produces this difference in the magnitude of the coefficients. Some typical limit values of the cut-off frequency divided by the sampling frequency when using single precision (32 bits) are shown in Table I.¹⁵ If we use a low-pass filter of eighth order, we will be limited to a cut-off frequency ratio of 0.05 , which is 2400 Hz at a sample rate of $48\,000$ Hz.

In order to increase the number of poles without affecting the filter performance, multiple two-pole filters can be arranged in cascade configuration¹⁵ with the cut-off frequency limit of a second-order filter. However, this is not a complete solution of the problem because there still exists a high minimum cut-off frequency limit (480 Hz @ 48 ksp/s). One possible solution to this problem is to increment the word length in order to reduce the quantization effects, as described in other works.⁹

We implemented a different algorithm that has an arbitrary small low cut-off frequency and very sharp roll-off in the frequency response, without degrading the filter performance. This algorithm consists of three steps. First, we filter the signal with a cascade configuration of two-pole filters. After that, we reduce the sampling rate of the signal (this is called decimation or downsampling).¹⁶ And finally, we filter the downsampled signal with a filter similar to the one used in the first step. This filter has the restrictions shown in Table I but the sampling frequency is much smaller now. Therefore, the filter cut-off frequency is much smaller in Hz.

To reduce the sampling rate by an integer factor (L), we take every L th sample as the new downsampled signal. We must take into account that when reducing the sampling rate, the maximum frequency present in the signal must be below $f_s/2$ (Nyquist theorem) to avoid the aliasing effect. So the high-frequency signals must be filtered before the downsampling is carried out.

In the proposed lock-in, we implemented an eighth-order Chebyshev low-pass filter in cascade configuration with a relative cut-off frequency of $5/1000$ (240 Hz @ 48 ksp/s). Then, we reduced the sampling rate by a factor of 100 (can




FIG. 4. Simplified block diagram of the algorithm used to reduce the cut-off frequency of the low-pass filters. Both the low-pass filter and the downsampling stage represent a low computational load to the DSP.

be configured by the user) to obtain a cut-off frequency of 2.4 Hz in the second filter, with a very sharp roll-off (theoretically -160 dB/decade). A simplified block diagram of this strategy is shown in Fig. 4.

In the DSP software, we used the IEEE extended precision data format (floating point, 40 bit) that allows computing and storing the intermediate results using 40-bit data words. Thus, even though the filter coefficients are stored in 32-bit words, we could reach f_c/f_s values of 0.005 with two-pole filters in cascade.

The number of low-pass filters and downsampling stages can be increased in order to obtain time constants of several minutes or even more. The minimum cut-off frequency is reduced quadratically with the number of stages (e.g., if we reduce twice the sampling rate by 100, we reduce the minimum frequency by 10 000). However, as in a conventional lock in, a long time constant in the output filters produces a slow time response and, consequently, it can not respond to fast variations of the signal. So, the user must find a trade off between the noise level and the measurement speed.

After filtering the signal, the processor stores, in the memory, the X (in-phase) and Y (in-quadrature) values that correspond to each frequency. The available memory allows storing more than 10 000 measurements. Using the synchronous serial port available in the board, we can connect low-cost nonvolatile memories used as permanent storage devices to save the measurements without a PC.

V. EXPERIMENTAL RESULTS

We chose a RC network to carry out the measurements because it is a simple nontrivial system and has a frequency

TABLE II. Frequency ranges for the $N=3$ frequency sweeps used in the measurements presented.

| Component number | Minimum frequency (Hz) | Maximum frequency (Hz) |
|------------------|------------------------|------------------------|
| 1 | 20 | 200 |
| 2 | 200 | 2000 |
| 3 | 2000 | 20 000 |

response that varies enough in magnitude and phase. This allows one evaluate the performance of the lock in in the whole frequency range and amplitude level. For the following measurements, the lock in first performs the above described autocalibration and then measures the frequency response of the system. Each reference component starts with the minimum frequency, rises up to the maximum, and then decreases to the minimum again, taking two measurements at each frequency value. In these measurements, the three components ($N=3$) have limits equally spaced logarithmically, as shown in Table II. Each reference component takes 256 measurements, giving 768 (256×3) values in the frequency spectrum from 20 Hz to 20 kHz.

The values of the electronic components are $R = 1980 \Omega$ for the resistor and $C=0.177 \mu\text{F}$ for the capacitor, giving an RC constant of $350 \mu\text{s}$ and a cut-off frequency of 455 Hz. However, these values depend on the temperature and are different for each run we performed.

In Fig. 5, we show the Bode plot (magnitude and phase) of the frequency response of the RC network. We fitted the experimental results with the theoretical curve of an ideal RC circuit. This curve fits very well and it is drawn below the experimental points. The calculated error of the RC constant is less than 0.1%. The good performance of the lock in in the whole frequency spectrum is also evident. The difference with the ideal curve at high frequencies can be attributed to the resistive and inductive components of the capacitor.

The lock-in parameters used for making these measurements were:

- (a) Filters time constant: 105 ms, and
- (b) Sweep time: 20 s.




FIG. 5. Measured frequency response of the RC network and the ideal curve that fits this data. The calculated error in the RC constant is 0.1%.

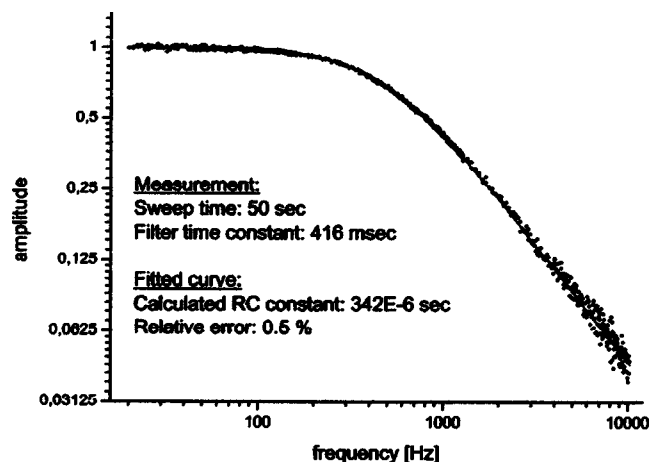


FIG. 6. Measured frequency response of the RC network with white noise added artificially to the system. The error in the calculated RC constant is 0.5%.

In order to minimize errors in measurements, the user has to find a trade off between the filters' time constant (reduces noise) and the frequency sweep speed (makes the measurement faster). With this time constant and the measured system, we obtained accurate measurements with sweep times of 20 s and more.

We made a similar measurement using a programmable function generator (Hewlett Packard 33120A) and a standard lock-in amplifier (Stanford Research Systems SR530). The minimum sweep time was approximately 150 s because of the slow response of the PLL circuit in the lock-in reference input.

The second plot (Fig. 6) is the frequency response of the same system with white noise added artificially. The noise was generated by a digital function generator (HP33120A). We also show the ideal expected curve. The signal and the noise were added using a low-noise operational amplifier, as is described in Ref. 17. In this measurement, we chose the following parameters:

- (a) Filters time constant: 416 ms, and
- (b) Sweep time: 50 s.

The filter time constant is longer because of the higher-noise level.

The signal level was 14 mV rms and the white-noise level was 282 mV rms, representing a signal-to-noise ratio (S/N) of -26 dB. We note that the relative error (the plot is in logarithmic scale) increases when the amplitude of the system response is reduced at high frequencies (e.g., $S/N = -54$ dB @ 10 kHz). With this measurement, the calculated error in the RC constant is 0.5%. As in a standard lock in, the influence of noise can be reduced by increasing the time constant, but with the consequent need for reducing the sweep speed as described above.

Figure 7 shows a measurement of the same system with a time variation of the RC value. This variation is carried out by changing the temperature (from 15°C to 70°C) of the capacitor in a total time of 400 s. Each RC value of the plot is calculated in the same way as described in the previous




FIG. 7. Measured time variation of the RC constant. Each RC value is calculated in the same way as in the previous figures. The estimated error bar is smaller than the data points. The error in time is 5 s (length of the sweep time).

plots (we obtain the frequency response and then fit it with the RC ideal curve). In this figure, we show the flexibility of our lock in to measure how the frequency response of a system changes with time. The parameters used for these measurements are:

- (a) Filters time constant: 41 ms, and
- (b) Sweep time: 5 s.

This short time constant produces bigger errors in the measurements but allows faster frequency sweeps. With these measurements we verified that the proposed lock in works correctly and the measurements are accurate and precise.

To conclude, we enumerate the characteristics of our current lock-in implementation and the possible improvements:

- (a) Range of frequencies: 20 Hz to 20 kHz. It can be extended using external ADCs and DACs.
- (b) Number of simultaneous frequencies: 3. It can be extended to 12 or more using faster code-compatible DSPs (e.g., ADSP-21161).
- (c) Maximum low-pass filters time constant: Our experiment does not need very long time constants. The maximum time constant is 416 ms. To achieve time constants of several seconds or minutes, more low-pass filters and downsampling stages can be added to the output filter.
- (d) Maximum number of measurements it can store in RAM memory: More than 10 000, and a nonvolatile memory can be connected to the serial port of the DSP.

We believe that the DSP-based lock-in amplifier described in this article is an effective solution for applications where a characterization of a linear system as a function of time and frequency is required.

ACKNOWLEDGMENTS

The authors thank the Electronic Division of Centro Atómico Bariloche and Mario Barbaglia. They also thank

CONICET for partial support of one of the authors (M.O.S.), and FUESMEN for financial support of the project.

- ¹Stanford Research Systems, SR850 DSP Lock-in Amplifier User's Manual, also in http://www.srsys.com/html/body_sr850.html
- ²Signal Recovery, Dual Phase DSP Lock-in Amplifier Model 7265 User's Manual, also in <http://www.signalrecovery.com/7265.htm>
- ³D. Skolnick and N. Levine, Why Use DSP?, Analog Dialogue (Volume 31, Number 1, 1997), also in <http://www.analog.com/library/analogDialogue/archives/31-1/DSP.html>
- ⁴P. A. Probst and B. Collet, Rev. Sci. Instrum. **56**, 466 (1985).
- ⁵J. Berger and D. S. Tannhauser, Rev. Sci. Instrum. **54**, 1781 (1983).
- ⁶P. A. Probst and A. Jaquier, Rev. Sci. Instrum. **65**, 747 (1994).
- ⁷F. De Rosa, Rev. Sci. Instrum. **57**, 1693 (1986).
- ⁸F. Barone, E. Calloni, L. DiFiore, A. Grado, L. Milano, and G. Russo, Rev. Sci. Instrum. **66**, 3697 (1995).
- ⁹L. A. Barragán, J. I. Artigas, R. Alonso, and F. Villuendas, Rev. Sci. Instrum. **72**, 247 (2001).
- ¹⁰M. A. Garcia, M. Tadey, and F. J. Bonetto, LAAR (to be published).
- ¹¹M. Garcia, M.S. thesis in Nuclear Engineering, Instituto Balseiro, June, 2000.
- ¹²M. Tadey and F. J. Bonetto, Revista CNEA **4**, 16 (2001).
- ¹³E. Ibarra, M.S. thesis, Instituto Balseiro, December, 2002.
- ¹⁴Analog Devices Inc., A technical tutorial on Direct Digital Synthesis (1999), URL: <http://www.analog.com>
- ¹⁵S. W. Smith, *The Scientist and Engineer's Guide to Digital Signal Processing*, 2nd ed. (California Technical Publishing, San Diego, CA, 1999), Chap. 20, p. 339. Also in <http://www.dspguide.com>
- ¹⁶S. J. Orfanidis, *Introduction to Signal Processing* (Prentice-Hall, New York, 1996), Chap. 12, p. 644.
- ¹⁷Analog Devices Inc., AD630 datasheet.