

Stochastic multi-channel lock-in detection

This content has been downloaded from IOPscience. Please scroll down to see the full text.

2014 New J. Phys. 16 013005

(<http://iopscience.iop.org/1367-2630/16/1/013005>)

View [the table of contents for this issue](#), or go to the [journal homepage](#) for more

Download details:

IP Address: 155.198.12.147

This content was downloaded on 11/01/2017 at 15:56

Please note that [terms and conditions apply](#).

You may also be interested in:

[A multichannel phase-sensitive detection method using orthogonal square waveforms](#)

G E Harrison, M A Player and P G H Sandars

[An algorithm to remove fringe jumps and its application to microwave reflectometry](#)

A Ejiri, K Shinohara and K Kawahata

[A low-cost, FPGA-based servo controller with lock-in amplifier](#)

G Yang, J F Barry, E S Shuman et al.

[ON THE USE OF SHOT NOISE FOR PHOTON COUNTING](#)

Jonas Zmuidzinis

[Discovery of Peculiar Periodic Spectral Modulations in a Small Fraction of Solar-type Stars](#)

Ermanno F. Borra and Eric Trottier

[Noise in measurements obtained by sampling](#)

D R White and J F Clare

[ON THE FOURIER AND WAVELET ANALYSIS OF CORONAL TIME SERIES](#)

F. Auchère, C. Froment, K. Bocchialini et al.

[Standard deviation of the mean and other time series properties of voltages](#)

Thomas J Witt and N E Fletcher

[Comparison of three digital fringe signal processing methods](#)

S Svitlov, P Masyk, Ch Rothleitner et al.

Stochastic multi-channel lock-in detection

J J Hudson¹, M R Tarbutt, B E Sauer and E A Hinds

Centre for Cold Matter, Physics Department, Imperial College London, London SW7 2AZ, UK
E-mail: jony.hudson@imperial.ac.uk

Received 16 July 2013, revised 8 November 2013

Accepted for publication 5 December 2013

Published 2 January 2014

New Journal of Physics **16** (2014) 013005

doi:[10.1088/1367-2630/16/1/013005](https://doi.org/10.1088/1367-2630/16/1/013005)

Abstract

High-precision measurements benefit from lock-in detection of small signals. Here we discuss the extension of lock-in detection to many channels, using mutually orthogonal modulation waveforms, and show how the choice of waveforms affects the information content of the signal. We also consider how well the detection scheme rejects noise, both random and correlated. We address the particular difficulty of rejecting a background drift that makes a reproducible offset in the output signal and we show how a systematic error can be avoided by changing the waveforms between runs and averaging over many runs. These advances made possible a recent measurement of the electron's electric dipole moment (Hudson *et al* 2011 *Nature* **473** 493).

1. Introduction

Lock-in detection is a well-known method for picking out a small modulated signal of interest in the presence of a large background. After multiplying the signal-plus-background by the modulation waveform, the dc component is proportional to the signal while the background now alternates and can be filtered out. An early account of this idea dates back to 1934 [2], when it was used to detect the balanced state of an ac bridge. Subsequent improvements to the electronics by Michels, Curtis and Redding [3, 4] resulted in a practical, narrow-band detector of

¹ Author to whom any correspondence should be addressed.



Content from this work may be used under the terms of the [Creative Commons Attribution 3.0 licence](https://creativecommons.org/licenses/by/3.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

high sensitivity, now widely known as a lock-in amplifier². This became a ubiquitous laboratory tool after it was commercialized by Robert Dicke and Princeton Applied Research [5].

The lock-in technique has been extended to measure several modulated signals simultaneously. In the context of high-precision measurement, Harrison, Player and Sandars (HPS) [6] describe signal modulation with complex, multi-frequency waveforms chosen to be mutually orthogonal. More recently in the context of optical tomography [7–9] consider phase-sensitive detection at several frequencies. Careful choice of the frequencies and filter time constants allows the signals to be extracted independently.

In this paper we extend the work of HPS. First, in order to establish a theoretical framework, we review the case when a single parameter is modulated. Then, in section 3 we revisit and extend the HPS discussion of multi-channel modulation and detection. In section 4, we introduce and analyse the problem of reproducible backgrounds and in section 5 we present a new, stochastic technique that solves the problem. We close with a summary in section 7.

2. Single channel demodulation

Consider a simple experiment with a scalar output signal, s , and a binary control parameter, c , which we set to the value +1 or –1 over each of N equal time intervals, labelled by the integer j . We sometimes refer to the set of values c_j as a switching pattern or waveform. We write the signal measured in each time interval as

$$s_j = \beta_j + \alpha c_j \quad (1)$$

with $j = 0, \dots, N - 1$. Here β_j is a background term, which may be time-varying, and α is the modulated part of the signal that we are trying to extract, which we will assume to be constant. Note that this form of the signal does not assume a linear response to the control parameter c_j . We obtain an estimate, $\hat{\alpha}$, of the modulation strength, α , by looking at the part of the signal correlated with c_j

$$\hat{\alpha} = \frac{1}{N} \sum_{j=0}^{N-1} s_j c_j. \quad (2)$$

The error in this estimate of α is

$$r_\alpha = \hat{\alpha} - \alpha = \frac{1}{N} \sum_{j=0}^{N-1} \beta_j c_j. \quad (3)$$

Our goal is to devise a scheme that gives the best estimate of α , least affected by the presence of the background β .

If β is constant during an experiment, then any switching waveform that is equally often positive and negative will reject it perfectly. We will call a waveform that has this property *balanced*. It is more realistic though to let the background vary in time, and then the cancellation of β is normally imperfect. To optimize the choice of c_j we must consider in more detail how β_j varies. Let us define the discrete Fourier transform of β ,

$$(\mathcal{F}\beta)_k = \frac{1}{\sqrt{N}} \sum_{l=0}^{N-1} \beta_l e^{-i\frac{2\pi k l}{N}}, \quad (4)$$

² Michels coined the name ‘lock-in amplifier’ in [3], then tried to change it in [4] to ‘synchronous detector’ but the original name is still widely used.

and write equation (3) in terms of these Fourier components

$$r_\alpha = \frac{1}{N} \sum_{j=0}^{N-1} \left(\frac{1}{\sqrt{N}} \sum_{k=0}^{N-1} (\mathcal{F}\beta)_k e^{i\frac{2\pi k}{N}j} \right) c_j, \quad (5)$$

where the large-bracketed quantity is nothing more than β_j . Rewriting this as

$$r_\alpha = \frac{1}{N} \sum_{k=0}^{N-1} (\mathcal{F}\beta)_k \left(\frac{1}{\sqrt{N}} \sum_{j=0}^{N-1} c_j e^{i2\pi\frac{jk}{N}} \right) = \frac{1}{N} \sum_{k=0}^{N-1} (\mathcal{F}\beta)_k (\mathcal{F}^{-1}c)_k, \quad (6)$$

we see that each Fourier component of the background contributes to the error in proportion to the (inverse) Fourier component of the modulation waveform at that frequency. In the laboratory, the Fourier spectrum of the background is often dominated by the low-frequency components, typified by the well-known $1/f$ power spectrum. In general terms, then, it is often desirable to have the Fourier spectrum of the modulation waveform reside at as high a frequency as possible. The simple modulation $c_j = (-1)^j$ achieves this, indeed that is the square-wave version of the standard lock-in detection scheme [2–4]. We shall see, though, that switching waveforms more complex than this fast square wave have some advantages when we come to consider multi-channel detection in section 3.

Equation (6) describes how a particular time-dependent background, β_j , will produce an error in a particular measurement of $\hat{\alpha}$. To properly consider the impact of *noisy* backgrounds we need to consider what happens when we repeat the measurement a number of times. A noisy background is one that differs from one experimental realization to the next, in contrast to a *reproducible* background which is the same for each experimental realization. We extend our treatment to multiple experimental realizations by allowing β_j to be a stochastic variable. To evaluate the effectiveness of a switching waveform it is then necessary to consider the statistical properties of the error, r_α , averaged over many experimental realizations.

We begin by limiting ourselves to purely noisy backgrounds, that is backgrounds for which the expected value of β_j is $\mathbb{E}(\beta_j) = 0$. We consider backgrounds where this limitation does not hold in section 4. For these purely noisy backgrounds it is easy to see from equation (3) that $\mathbb{E}(r_\alpha) = 0$, which is to say that our estimates of α will be unbiased. They will, though, be noisy as β_j changes from one experimental realization to the next. To quantify this noise let us calculate the standard deviation, $\sigma(r_\alpha)$, which gives the statistical uncertainty in our measurement of α . This is given by

$$\sigma(r_\alpha)^2 = \mathbb{E}(r_\alpha r_\alpha^*) = \mathbb{E} \left(\frac{1}{N} \sum_{k=0}^{N-1} (\mathcal{F}\beta)_k (\mathcal{F}^{-1}c)_k \frac{1}{N} \sum_{l=0}^{N-1} (\mathcal{F}\beta)_l^* (\mathcal{F}^{-1}c)_l^* \right). \quad (7)$$

In the first step here, we have used the fact that r_α has zero mean.

It is common to characterize noise by a power spectrum, which defines the (real) amplitude $\tilde{\beta}_k$ of each Fourier component, assumed to be constant for each experimental realization. The phase ϕ_k of each component must vary in some way between realizations for this to be called noise. Let us make the simplifying assumption that the phases for each component are independently drawn from a uniform distribution over $[0, 2\pi)$ for each experimental realization. Although this noise model is quite restricted, it does provide a reasonable approximation to

typical laboratory noise and it simplifies equation (7) in an instructive way. Inserting $(\mathcal{F}\beta)_k = \tilde{\beta}_k e^{i\phi_k}$ into equation (7), and writing the expected value operator explicitly as a sum, we obtain

$$\begin{aligned}\sigma(r_\alpha)^2 &= \lim_{M \rightarrow \infty} \frac{1}{M} \frac{1}{N^2} \sum_{n=1}^M \sum_{k=0}^{N-1} \sum_{l=0}^{N-1} \tilde{\beta}_k e^{i\phi_{k,n}} (\mathcal{F}^{-1}c)_k \tilde{\beta}_l e^{-i\phi_{l,n}} (\mathcal{F}^{-1}c)_l^* \\ &= \frac{1}{N^2} \sum_{k=0}^{N-1} \sum_{l=0}^{N-1} (\mathcal{F}^{-1}c)_k (\mathcal{F}^{-1}c)_l^* \tilde{\beta}_k \tilde{\beta}_l \left(\lim_{M \rightarrow \infty} \frac{1}{M} \sum_{n=1}^M e^{i(\phi_{k,n} - \phi_{l,n})} \right),\end{aligned}$$

where $\phi_{j,n}$ is the n th realization of the stochastic variable ϕ_j . The quantity in parentheses is zero unless $k = l$, hence the sum simplifies to

$$\sigma(r_\alpha) = \frac{1}{N} \sqrt{\sum_{k=0}^{N-1} |(\mathcal{F}^{-1}c)_k|^2 \tilde{\beta}_k^2}. \quad (8)$$

We see that the individual frequency contributions sum in quadrature for this class of noise.

3. Multichannel modulation and demodulation

In section 2 we have considered how to find the modulation strength α when one parameter is modulated by control parameter c . Now we extend this to the case of multiple modulations. To help establish the notation, let us consider the signal when there are two control parameters, $c^{(1)}$ and $c^{(2)}$

$$s_j = \beta_j + \alpha^{(1)} c_j^{(1)} + \alpha^{(2)} c_j^{(2)} + \alpha^{(1,2)} c_j^{(1)} c_j^{(2)}. \quad (9)$$

In addition to the responses $\alpha^{(1)}$ and $\alpha^{(2)}$, there is also the bilinear response $\alpha^{(1,2)}$ to the product of control parameters. Put another way, this part of the signal is correlated with the *relative sign* of the two control parameters. Similarly, with three modulations, there is a trilinear term as well as the bilinear and linear ones. When P parameters are modulated by switching waveforms $c^{(k)}$, where $k = 1, \dots, P$, the signal generalizes to

$$s_j = \beta_j + \sum_{\substack{q \subseteq \{1 \dots P\} \\ q \neq \emptyset}} \alpha^q \prod_{k \in q} c_j^{(k)} = \beta_j + \sum_{\substack{q \subseteq \{1 \dots P\} \\ q \neq \emptyset}} \alpha^q w_j^q. \quad (10)$$

Here q labels all the subsets of $\{1 \dots P\}$, i.e. all the linear, bilinear, trilinear, etc terms. For convenience, the last equality renames the product of control parameters in term q as w_j^q (q being a label, not a power), which we sometimes refer to as an *analysis waveform*.

With several experimental parameters k , each modulated by a different waveform $c^{(k)}$ to produce a string of N signals s_j , the response $\alpha^{(k)}$ to each individual modulation is estimated by the generalization of equation (2):

$$\hat{\alpha}^{(k)} = \frac{1}{N} \sum_{j=0}^{N-1} s_j c_j^{(k)}. \quad (11)$$

This relies on the use of switching patterns $c^{(k)}$ that are mutually orthogonal. As before we will require the switching waveforms to be balanced, so that constant backgrounds are rejected.

The same strategy also applies to the multi-linear responses, allowing the further generalization

$$\hat{\alpha}^q = \frac{1}{N} \sum_{j=0}^{N-1} s_j w_j^q, \quad (12)$$

provided the analysis waveforms w^q are also mutually orthogonal. In order that these multi-linear responses also reject constant backgrounds we will require the analysis waveforms to be balanced, which does not automatically follow from the switching waveforms being balanced. We will call a set of switching waveforms *completely balanced* if all products of them (i.e. all analysis waveforms) are balanced.

HPS [6] have presented a simple scheme for generating a set of mutually orthogonal, completely balanced switching waveforms, and a compact notation for naming them. They start with a basic set of square waves

$$b(m)_j = (-1)^{\lceil \frac{j+1}{2^m} \rceil + 1}, \quad (13)$$

where $\lceil \cdot \rceil$ indicates the ceiling function. As in section 2, the integer j specifies the time step and runs from zero to $N - 1$, where N is the total number of steps in the switching sequence, constrained to be an even integer power of 2. The integer label m ranges from zero to $\log_2(N/2)$. For instance, when $N = 8$, we have three basic square waves

$$b(0) = \{+1, -1, +1, -1, +1, -1, +1, -1\}, \quad (14)$$

$$b(1) = \{+1, +1, -1, -1, +1, +1, -1, -1\}, \quad (15)$$

$$b(2) = \{+1, +1, +1, +1, -1, -1, -1, -1\}. \quad (16)$$

Switching waveforms c are constructed by multiplying any combination of basic square waves together. It is convenient for these to be identified by a binary code that indicates which basic square waves are included, with the n th least significant bit indicating whether the n th basic waveform is included. For example,

$$c(010) = b(1) = \{+1, +1, -1, -1, +1, +1, -1, -1\}, \quad (17)$$

$$c(101) = b(2) \times b(0) = \{+1, -1, +1, -1, -1, +1, -1, +1\}, \quad (18)$$

$$c(110) = b(2) \times b(1) = \{+1, +1, -1, -1, -1, -1, +1, +1\}. \quad (19)$$

We will find it useful later to introduce an optional, over-all inversion, indicated by an asterisk. For example, $c(101^*) = -c(101)$.

The basic waveforms are all balanced. The symmetry of the HPS waveforms also ensures that multiplying any waveform by a basic waveform flips an equal number of positive and negative entries, leaving the result balanced. In this way, all non-trivial waveforms generated with the HPS technique are balanced³. Note that the trivial pattern of constant +1 is a valid HPS waveform, which we label $c(0)$, and while it is not balanced, it is useful to include in the discussion as it can be used to extract the average value of the background.

³ The HPS waveforms constitute only a fraction of all the balanced waves of a given length. This fraction becomes vanishingly small for long patterns but nevertheless there are enough for practical purposes. For a pattern of (even) length N , the number of non-trivial HPS waveforms is $N - 1$. By contrast, the number of balanced waveforms is $N! / (\frac{N}{2}!)^2$, that being the number of distinct ways to assign the value +1 to $N/2$ of the sites (and -1 to the rest).

In order to extract all the responses from a signal, all analysis waveforms w^q must be mutually orthogonal. The construction of the HPS waveforms ensures that all distinct waveforms are orthogonal, meaning it is sufficient to ensure that all analysis waveforms are distinct. Note well that this is a stronger constraint than ensuring that all *switching* waveforms are distinct. This point was not addressed by HPS [6], but is of great practical relevance when it comes to choosing the set of modulation waveforms. For example, let us choose three modulation waveforms $c^{(1)} = c(001)$, $c^{(2)} = c(011)$ and $c^{(3)} = c(010)$. These are mutually orthogonal, but the analysis waveform $w^{(1,2)} = c^{(1)}c^{(2)}$ is the same as $c^{(3)}$ and therefore we cannot separate the bilinear response $\alpha^{(1,2)}$ from the linear response $\alpha^{(3)}$ with this choice of modulation waveforms.

Requiring all the analysis waveforms to be distinct imposes a minimum length on the waveform, which depends on the number of parameters being modulated. Given P parameters, each modulated with a different switching pattern $c^{(k)}$, the number of independent analysis waveforms w^q is 2^P .⁴ Since the signal is made up of N numbers, s_j , the linear combinations of its elements can only represent N independent variables, and therefore a full analysis is only possible if the length N of the switching pattern is at least 2^P . This means that the number of digits in the longest binary HPS code must equal or exceed the number of parameters being modulated.

To test whether all analysis waveforms resulting from a given set of switching waveforms are distinct and orthogonal, one could simply generate them and check. In practice, however, we normally test the waveforms in a different way that is equivalent but easier to implement. With P switched parameters, the set of switches can be in any one of 2^P combined states in each time interval. If the set of switching waveforms is such that each of these 2^P states is visited an equal number of times in the N measurements, then all of the analysis waveforms are distinct and orthogonal. In the example above of $c^{(1)} = c(001)$, $c^{(2)} = c(011)$ and $c^{(3)} = c(010)$, only four of the eight combined switch states are accessed, specifically $(c^{(1)}, c^{(2)}, c^{(3)}) = \{(-1, -1, +1), (-1, +1, -1), (+1, -1, -1), (+1, +1, +1)\}$ indicating that these switching waveforms do not result in a full set of orthogonal analysis waveforms.

We will employ HPS waveforms exclusively in the rest of this paper.

4. Reproducible backgrounds

In the latter part of section 2 we considered a background β that varies purely randomly, that is without any consistent, repeatable part. We showed that although the measurement is perturbed by this background, $\hat{\alpha}$ converges to the true value α in the long-term average. In this section we allow some part of the background to vary in the same way over each set of N measurements. We will see that this gives a biased result when using the detection method described so far, and we show how a modification of the method can remove this bias.

To give a concrete example of such a background we consider a measurement recently made in our laboratory [1, 10]. In this experiment a molecular beam is subjected to a number of manipulations by static electric and magnetic fields, rf magnetic fields, and laser light, after which we measure the number of molecules in a particular quantum state. We modulate these

⁴ There are P linear responses, $P(P-1)/2!$ bilinear responses, $P(P-1)(P-2)/3!$ trilinear, etc, terminating with one response to the product of all p modulations, plus the trivial waveform $c(0)$. Summing the number of responses we obtain a total of 2^P .

fields in a number of ways, and in particular, we switch the signs of the static electric and magnetic fields. We seek to detect a modulation strength associated with the relative sign of the electric and magnetic fields, which would indicate that the charge distribution within electrons is not spherical. We modulate nine experimental parameters using the scheme described in this paper. Our waveforms have length $N = 4096$, and making a single *block* of 4096 measurements takes a few minutes. We then repeat this experiment, for months at a time, amassing many thousands of blocks in order to reduce the final measurement uncertainty. Noise sources include shot noise and electronic noise in the detection system, fluctuations of the beam intensity and instability of the various fields. Using the techniques described above we reduce the impact of that noise through careful choice of the modulation waveforms. However, part of the varying background is reproducible. For example, the molecular source heats up during each block of measurements, resulting in a drift of the signal s_j as the experiment progresses. There is a pause between each block where various parameters of the equipment are optimized and during this pause the source cools down again. The consequent repeating background introduces a bias that would invalidate the measurement, but we are able to suppress the bias by the method we now describe.

Let us consider a background which is the sum of a reproducible part γ and a purely random part. We define the bias ϵ^q as the error in measuring the parameter α^q , averaged over many experimental realizations. Equations (10) and (12) then give

$$\epsilon^q = \mathbb{E}(\hat{\alpha}^q - \alpha^q) = \mathbb{E} \left[\frac{1}{N} \sum_{j=0}^{N-1} \gamma_j w_j^q \right] = \frac{1}{N} \sum_{j=0}^{N-1} \gamma_j w_j^q, \quad (20)$$

where we have used the linearity of equation (12) and the fact that the random part of the background generates no bias. The last equality is intended to show the problem clearly: in each experimental realization the error due to γ is the same, producing a bias that does not average away.

To get out of this pickle we can vary the switching waveforms between experimental realizations. In this case the last equality in equation (20) does not hold, and we must write instead

$$\epsilon^q = \frac{1}{N} \sum_{j=0}^{N-1} \gamma_j \mathbb{E}[w_j^q]. \quad (21)$$

We can see that this offers a solution to the problem: the bias will vanish if the waveform, averaged over all experimental realizations, is orthogonal to the repeating part of the background. The simplest way to arrange this for all γ is to randomize the waveforms, making $\mathbb{E}[w_j^q] = 0$. Randomizing the waveforms makes sure that the error from the reproducible background is different in each measurement, in effect converting the bias into noise. There is a pleasing symmetry to this: by introducing structured waveforms we were able to reduce the impact of noisy backgrounds on our measurement, and we see that by introducing randomness into the waveforms we are able to reduce the impact of structured backgrounds.

5. Partially randomized waveforms

We wish to randomize the waveforms to reject biases from reproducible backgrounds, but we must do this carefully so as not to undermine the noise-rejecting power of the lock-in scheme.

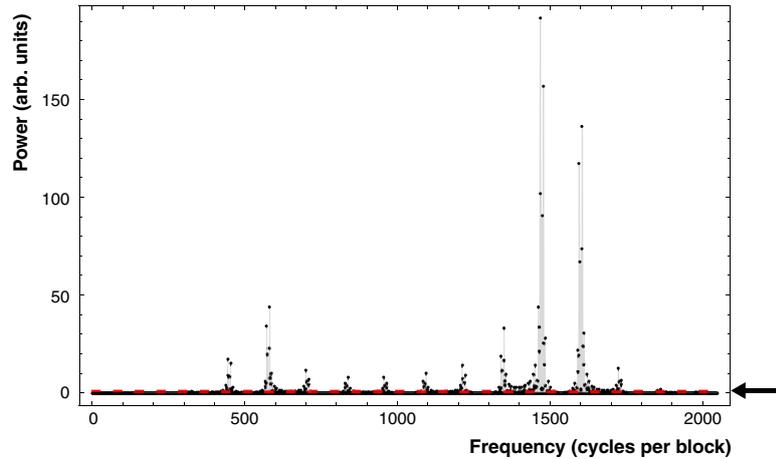


Figure 1. The power spectrum of a single switching waveform, $c(101100110111)$. The dashed line and arrow indicate the power level that would result from spreading the power equally over all Fourier components.

Our approach is to randomize the waveforms partially: enough to make sure that measurement bias from reproducible backgrounds is suppressed, while retaining the desirable properties of the lock-in scheme. The HPS waveform notation suggests a particularly simple way to do this: we generate new, partially randomized waveforms for each experimental realization by randomly generating HPS codes, but with constraints to ensure that they have the frequency properties we desire.

Let us consider again the experiment described in section 4, which uses switching waveforms of length $N = 4096$, specified by 12-bit HPS codes. To select suitable waveforms for this experiment we first generate a set of 12-bit codes, one for each switch, choosing the bits at random⁵. We then check each analysis channel of particular interest to see whether it satisfies two constraints. Firstly, its HPS code must not have more than one zero in the last four bits. This ensures that the corresponding analysis waveform will reject low-frequency noise as a result of its strong high-frequency content. Secondly, there must be no more than four zeros in the leading eight bits representing the slower basic waveforms. This ensures that the analysis waveform has a complex structure, needed to suppress the residual error due to slow variations in the background, as discussed by HPS [6]. If either test is failed, we generate a new set of codes and repeat until the tests are passed. Finally, we chose at random whether to invert each switching waveform, with probability 1/2.

This partial randomization scheme completely eliminates bias in the measurement. At any step in the sequence of N time bins, every analysis waveform is randomized, with equal probability of being -1 or $+1$ because of the randomly chosen inversions. This satisfies the requirement that $\mathbb{E}[w_j^q] = 0$. By randomizing the waveform codes between each experimental realization we also ensure that no one analysis waveform performs consistently poorly, as would be likely to happen if fixed codes were used.

To understand the impact of the partial randomization scheme, it is instructive to look at the average power spectra of some waveforms. Figure 1 shows the power spectrum of a particular

⁵ In fact, we do not randomly generate all of the codes. Two of the switch channels in our experiment must be switched with particular sequences due to hardware constraints. The procedure we describe for generating partially randomized waveforms is used for generating the other seven switching waveforms.

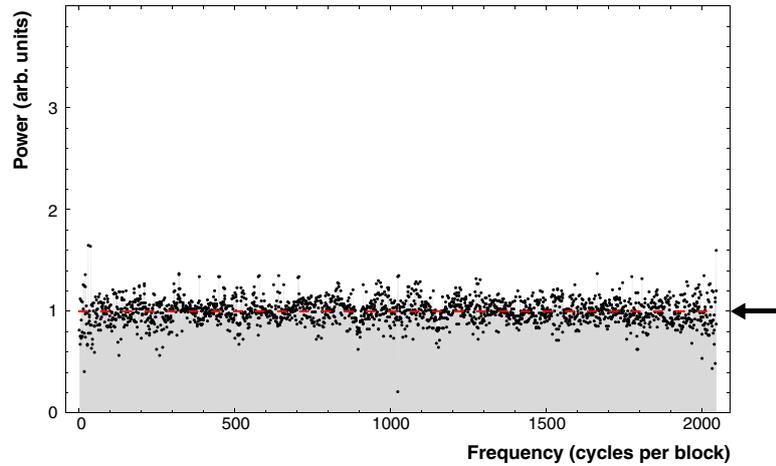


Figure 2. The average power spectrum of completely randomized HPS waveforms, with random inversion. Again, the dashed line and arrow show the power level that would result from spreading the power over all Fourier components—the white noise level—which we see approximates the resulting power spectrum well. The plot was generated by averaging 10 000 waveforms generated with randomly chosen HPS codes.

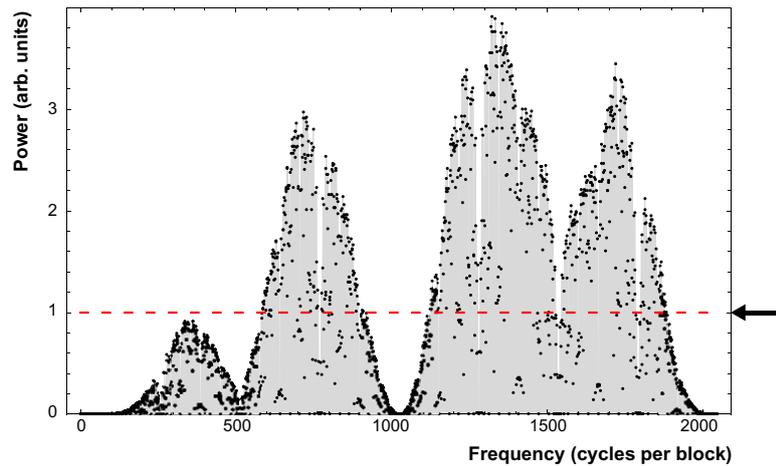


Figure 3. Average power spectrum of waveforms selected by the partial-randomization algorithm described in the text. The dashed line and arrow show the white noise spectrum. The plot was generated by averaging the power spectra of 10 000 waveforms generated using the algorithm in the text.

12-bit HPS waveform, $c(101100110111)$. We see that this has most of its power concentrated at high frequency, so it is effective in rejecting low-frequency noise. However, as we have seen, we will be susceptible to bias if we use this waveform for every measurement. At the opposite extreme, figure 2 shows the average power spectrum for 10^4 HPS waveforms chosen randomly without any constraint on the HPS code. While these waveforms will yield an unbiased result, they do not provide good noise rejection as the average power spectrum is completely white.

Figure 3 shows the average power spectrum of a waveform generated using the algorithm described above. We see that this waveform offers the best of both worlds. The average power is spread over a broad range of the higher frequencies while low-frequency components are almost

completely absent. Thinking in the time-domain, if we were to string together all the partially randomized waveforms chosen for a given analysis channel, the switching pattern would be high-frequency periodic in the short term but random and uncorrelated in the long term. It is this combination of short-term periodicity to reject noise, together with long term randomness to reject repetitive backgrounds, that makes the scheme effective.

6. Example

We give an example of the performance of the scheme by considering once again the experiment [1] described in section 4.

First we examine how well our partially randomized switching scheme rejects random noise, by comparing noise in a partially randomized analysis channel—whose waveforms were generated using the scheme described in section 5—with noise in two benchmark analysis channels. The first benchmark has the fixed waveform $c(000000000000)$, which does not switch at all. This channel measures the over-all noise in the signal, without any reduction through lock-in detection. The second uses the fixed waveform $c(000000000001)$, which is the fastest square-wave modulation available. The noise in this channel is representative of the error given by traditional, non-randomized lock-in detection. To compare the noise in each of these three cases, we take a sample of 1000 measured data blocks and demodulate the signal with the appropriate partially randomized waveform, the non-switching waveform, and the fast-switching waveform⁶. For each waveform, the standard deviation in the 1000 results provides a measure of the noise. Scaling the results so that the partially randomized technique has a noise of 1, the non-switching waveform gives a noise of 309, and the fast-switching waveform gives 0.9. We see that both the traditional, fast-switching lock-in technique and the partially randomized technique significantly reduce the noise, by more than a factor of 300. Although the partially randomized waveform gives marginally less noise suppression than the fast-switching one, this is a small price to pay for the rejection of the potential systematic errors.

To see how well the partial randomization rejects the bias coming from a reproducible background we turn to a simulation. This models the reproducible drift in the experiment as an exponentially decaying background. Figure 4 shows the signal in a channel with a partially randomized waveform, as a function of the number of blocks averaged. This is normalized to the signal detected in the fixed channel $c(101100111110)$, which meets the frequency constraints of section 5. While the bias in the fixed channel remains constant, we see that the partial randomization procedure is indeed effective in reducing the bias as we increase the number of data blocks. The data in figure 4 are well fitted by a/\sqrt{n} where n is the number of blocks, and a is a fit constant, as might be expected from the random choice of the waveform inversion for each block of data. From this scaling we estimate that any bias in the experiment of section 4, with approximately 6000 blocks of data, was reduced by a factor of order 100 by the partial waveform randomization.

⁶ These data were taken with the partially randomized modulation scheme, so the fast-switching signal does not represent any real signal. The signal for the partially randomized waveform that we are comparing with has zero mean, and neither it nor the fast-switching signal are correlated to the other waveforms, so it still serves as a useful comparison of noise.

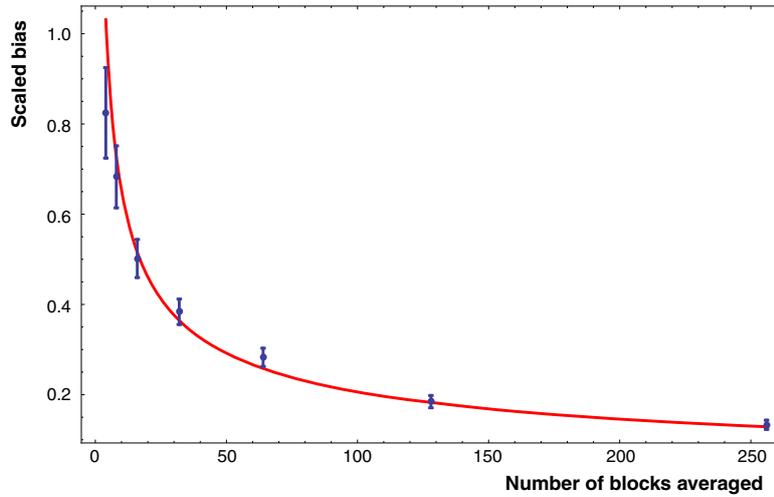


Figure 4. Bias from an exponentially decaying background, as a function of the number of blocks averaged, using the partially randomized waveform scheme. The y-axis is scaled to the bias in a particular, static waveform that meets the frequency constraints of section 5, $c(101100111110)$. Each point is the average result of 100 independent simulations, with the error bar indicating the standard error in the mean. The line shows a fit to a/\sqrt{n} , where n is the number of blocks and a is a fit constant.

7. Summary and conclusions

We have shown that the multi-channel lock-in detection scheme introduced by HPS yields information on all the responses, but only after minimum-length and completeness requirements are applied to the choice of modulation waveforms. A method is given for testing the suitability of the waveform set. We have shown how to evaluate the standard deviation of the measurement results in the presence of random noise. Finally, we have pointed out that many experiments also have a reproducing pattern of background variation that can bias the results of the measurement. We have described a scheme for partially randomizing the modulation waveforms in such a way that this bias averages to zero and merely contributes to the noise. We have demonstrated the effect of this partial randomization scheme by quantifying its impact on a recent high-precision measurement.

Acknowledgments

This work was supported by the Royal Society, the UK EPSRC and STFC and by the European Commission.

References

- [1] Hudson J J, Kara D M, Smallman I J, Sauer B E, Tarbutt M R and Hinds E A 2011 *Nature* **473** 493
- [2] Cosens C R 1934 *Proc. Phys. Soc. Lond.* **46** 818
- [3] Michels W C and Curtis N L 1941 *Rev. Sci. Instrum.* **12** 444
- [4] Michels W C and Redding E D 1949 *Rev. Sci. Instrum.* **20** 566

-
- [5] Hawrit M 1985 *Interview of Robert Dicke by Martin Hawrit (Niels Bohr Library and Archives)* (College Park, MD: American Institute of Physics) (www.aip.org/history/ohilist/4572.html)
- [6] Harrison G E, Player M A and Sandars P G H 1971 *J. Phys. E: Sci. Instrum.* **4** 750
- [7] Schmitz C H, Locker M, Lasker J M, Hielscher A H and Barbour R L 2002 *Rev. Sci. Instrum.* **73** 429
- [8] Masciotti J, Lasker J and Hielscher A 2008 *IEEE Trans. Instrum. Meas.* **57** 182
- [9] Hintenaus P M and Trinker H 2013 *IEEE Trans. Instrum. Meas.* **62** 785
- [10] Kara D M, Smallman I J, Hudson J J, Sauer B E, Tarbutt M R and Hinds E A 2012 *New J. Phys.* **14** 103051