

**Modelling Acoustic Signals in
the Calibration of Underwater
Electroacoustic Transducers in
Reverberant Laboratory Tanks**

S P Robinson and P M Harris

March 1999

Modelling Acoustic Signals in the Calibration of Underwater Electroacoustic Transducers in Reverberant Laboratory Tanks

Stephen P Robinson
Centre for Mechanical and Acoustical Metrology

and

Peter M Harris
Centre for Information Systems Engineering

ABSTRACT

In the calibration of an underwater electroacoustic transducer, the stepped-sinusoidal signals that are commonly used are contaminated by transients due to the resonant behaviour of the transducer, as well as noise. In addition, when calibrations are undertaken in reverberant laboratory tanks of finite size, reflections from the tank boundaries often arrive before the steady-state response of the transducer can be observed directly. However, by use of suitable signal models, it is possible to predict the steady-state response from the initial transient-dominated part of the waveform.

The approach considered here is to model the free-time response of the device by a function consisting of a sum of complex exponential terms which are used to describe both the steady-state and resonant behaviour of the device. In this report, we review two classes of estimation method: linear prediction methods (such as Prony's method and its variants), and nonlinear least-squares methods. Standard linear prediction methods are shown to be statistically biased and inefficient. This observation motivates the development of a linear prediction method in which proper account is taken of the error structure in the data. A nonlinear least-squares algorithm is also described based on a safeguarded Gauss-Newton algorithm that uses regularisation to address the ill-conditioning that is a property of the underlying problem.

Results are presented of using these methods to analyse simulated data generated to represent the measured response of a device. The effect of varying in a systematic manner the properties of the simulated data and device under test is investigated. Results are also presented for data obtained from measurements of a real transducer, and it is shown how these results may be used to establish the free-field sensitivity of the device. Finally, results are presented of using the information gained from the modelling of the transducer resonant behaviour to predict the transducer frequency response at frequencies other than those under test.

This report constitutes the deliverable of project 2.1d(ii) of the 1995-98 NMS Acoustics Metrology Programme.

© Crown Copyright 1999
Reproduced by permission of the Controller of HMSO

ISSN 1369-6785

National Physical Laboratory
Queens Road, Teddington, Middlesex, TW11 0LW

Extracts from this report may be reproduced provided the source
is acknowledged and the extract is not taken out of context.

Approved on behalf of the Managing Director, NPL
by Dr G.R.Torr, Head, Centre for Mechanical and Acoustical Metrology.

TABLE OF CONTENTS

1 INTRODUCTION	1
1.1 BACKGROUND.....	1
1.2 CONVENTIONAL SIGNAL ANALYSIS METHODS.....	2
1.3 BREAKDOWN OF CONVENTIONAL TECHNIQUES.....	3
1.4 EXTRAPOLATION USING SIGNAL MODELLING	4
1.5 ORGANISATION OF THE REPORT.....	5
2 MODELS OF THE FREE-TIME RESPONSE	6
2.1 COMPLEX EXPONENTIAL PARAMETRISATION.....	6
2.2 PRONY PARAMETRISATIONS.....	8
3 ESTIMATION	10
4 LINEAR ESTIMATION METHODS	13
4.1 LEAST-SQUARES PRONY METHOD	13
4.2 IMPROVING PRONY'S METHOD	14
5 NONLINEAR ESTIMATION METHODS	16
6 DESCRIPTION OF DATA PROCESSING ALGORITHM.....	17
6.1 ACQUIRE DATA	17
6.2 REMOVE ANY NON-ZERO DC LEVEL IN THE ACQUIRED DATA.....	17
6.3 IDENTIFY FREE-TIME WINDOW	18
6.4 APPLY LINEAR PREDICTION METHODS.....	18
6.5 APPLY NONLINEAR METHODS.....	20
6.6 VALIDATE ESTIMATES OF THE MODEL PARAMETERS	22
7 RESULTS FOR SIMULATED DATA	23
7.1 VARYING THE TIME-WINDOW	25
7.2 VARYING THE DATA SAMPLING RATE.....	25
7.3 VARYING THE NOISE LEVEL	25
7.4 VARYING THE DRIVE FREQUENCY.....	26
7.5 VARYING THE Q-FACTOR.....	26
7.6 SUMMARY	27
8 RESULTS FOR EXPERIMENTAL DATA.....	28

9 MODELS OF THE TRANSMITTING VOLTAGE RESPONSE.....	30
10 SCOPE FOR FURTHER WORK.....	33
10.1 FURTHER INVESTIGATION OF SOURCES OF ERROR.....	33
10.1.1 Evaluation of uncertainties	33
10.1.2 Instrumentation	33
10.1.3 Devices to be calibrated	33
10.2 EXTENSIONS TO THE METHOD	33
10.2.1 Combining data obtained at a range of frequencies	33
10.2.2 Inclusion of extra terms to represent echoes.....	33
10.2.3 Use of other novel techniques.....	34
10.3 INCORPORATION AND USE WITHIN CALIBRATION SOFTWARE.....	34
10.4 TRANSDUCER MODELLING.....	34
11 CONCLUSION	35
12 ACKNOWLEDGEMENTS	36
13 REFERENCES	37
APPENDIX A:NONLINEAR LEAST-SQUARES ESTIMATION.....	39
A.1 THE GAUSS-NEWTON ALGORITHM.....	39
A.2 LINE SEARCH	40
A.3 REGULARISATION	41
APPENDIX B: LINEAR DAMPED HARMONIC OSCILLATORS.....	43
B.1 FREE, MECHANICAL OSCILLATION.....	43
B.2 FORCED, MECHANICAL OSCILLATION.....	43
B.3 FORCED, ELECTRICAL OSCILLATION.....	44

1 INTRODUCTION

1.1 BACKGROUND

Ideally, the calibration of underwater electroacoustic transducers at kilohertz frequencies requires a free-field environment. For continuous wave fields, this requires the use of a large volume of water such as a lake, reservoir or ocean in order that reflections from the medium boundaries are sufficiently attenuated by propagation losses. The impracticality and expense of utilising such volumes has led to the use of laboratory tanks for measurements, with gated signals and time-windowing to isolate reflections. Laboratory tanks have the advantage of providing more controlled experimental conditions, but their use introduces further measurement problems, especially if the tank dimensions are small when measured in acoustic wavelengths.

During laboratory tank calibrations of transducers, a minimum of two transducers, one acting as a transmitter or projector and the other as a receiver, are suspended in a tank of finite size filled with water. The projector is driven by a signal consisting of a discrete-frequency tone-burst of finite duration. The signal detected by the receiving hydrophone is contaminated by (i) transients due to the resonance of the devices, (ii) reflections of the transmitted signal from the tank walls, floor and the water-surface, and (iii) noise (potentially from both acoustic and electrical sources). The extra distance travelled by the reflected signals causes them to arrive later in time than the direct path signal. The available time to observe the direct signal before the reflections arrive is termed the echo-free time, and depends on the size of the tank and the exact position of the transducers within the tank. Figure 1 is a schematic diagram of a projector and receiver in a water tank showing the main sources of reflections (ignoring multiple reflections).




Figure 1: Schematic diagram of a projector and receiver in a water tank showing the main sources of reflections (ignoring multiple reflections).

The received signal is sampled at uniformly spaced times by use of an analogue-to-digital converter, and measurements are made of the amplitude of the signal. In Figure 2 are some example waveforms obtained in this way. These relate to an ITC1001 device (used as a projector) with a resonance at approximately 18 kHz, and the two recorded waveforms correspond to drive voltages with frequencies of, respectively, 2 kHz and 20 kHz. We can identify in these figures the various “phases” in the response: the turn-on of the device followed immediately by an oscillation at its resonance frequency of approximately 18 kHz. After between three and four cycles of this oscillation, the resonant behaviour is sufficiently damped to observe the steady-state response of the device which takes the form of an undamped oscillation at the frequency of the drive voltage. Finally, the response is contaminated by a combination of the turn-off of the device and the arrival of the first reflections. These two effects are most clearly identified in the response to the 20 kHz drive voltage.




Figure 2: Measured data for the ITC1001 transducer corresponding to a 2 kHz drive voltage (left), and a 20 kHz drive voltage (right).

1.2 CONVENTIONAL SIGNAL ANALYSIS METHODS

In the case of the signals shown in Figure 2, there is enough free-time before the arrival of the first echoes for the steady-state signal to be observed directly. Typically, a time-window or gate is applied so that only the portion of the steady-state signal within the window is made available for analysis. There are a number of ways of measuring the amplitude of the steady-state component that are in common usage. These include:

- (i) Direct measurement of the peak voltage, either by a peak detector or by measuring the maximum and minimum of the digitised signal.
- (ii) Calculating the root-mean-square (RMS) voltage (by squaring, averaging and square rooting the digitised signal). Ideally, to avoid bias in the result, this is done using an integer number of cycles of the sinusoidal signal.
- (iii) Performing a fast Fourier transform (FFT) of the signal and taking the amplitude of the spectrum at the drive frequency, again using an integer number of cycles.
- (iv) Performing a “narrow-band” discrete Fourier transform (DFT) of the signal, calculating only the amplitude of the component at the drive frequency, again using an integer number of cycles.

- (v) Performing a least-squares fit of a sine-wave of the appropriate frequency and taking the amplitude of the fitted sine-curve.

All the above methods have strengths and weaknesses. The peak measurement of (i) is simple to implement but inaccurate in the presence of any noise, distortion and residual amplitude fluctuations of the waveform envelope. The other methods all require a little more processing, perhaps with dedicated software algorithms. However, methods (ii) to (v) will provide some effective averaging of small amplitude variations if many cycles are contained within the time window. Methods (iii) to (v) will also provide some discrimination against noise and provide information on phase as well as amplitude. However, errors may occur with methods (ii) to (iv) if an integer number cycles is not used for the analysis. Also, with the FFT of method (iii), an error may occur if the spectrum does not contain a point at the exact frequency of excitation. At NPL, methods (ii), (iv) and (v) are generally used depending on the number of cycles available for analysis. Both methods (iv) and (v) can still be useful even when analysing only half a cycle of signal.

1.3 BREAKDOWN OF CONVENTIONAL TECHNIQUES

In the examples from Figure 2, the steady-state signal can be observed directly. As might be expected, fewer cycles of the steady-state signal are available for analysis at lower frequencies since the period increases with decreasing frequency. If the echo-free time for a particular tank with transducers optimally positioned is denoted by T , then the number of cycles available for analysis before the arrival of reflections is equal to the product fT , where f is the frequency of excitation.

If these cycles are to be used for analysis, steady-state conditions must be reached if one of the methods of Section 1.2 is to be employed. Since the electroacoustic transducers used in underwater acoustics are resonant devices behaving as damped harmonic oscillators of quality-factor Q , it will take approximately Q cycles of the resonance frequency before the initial turn-on transients have died away. For situations where $Q > fT$, so that steady-state is not reached within the free-time available, *all* of the methods described in Section 1.2 will lead to inaccurate values for the steady-state amplitude. This situation is illustrated in Figure 3.




Figure 3: An example of a normalised waveform for a moderately high- Q transducer driven at its resonance frequency. Note that steady-state is not reached within the available time window.

As an example, in the 5.5 m diameter by 5 m deep measurement tank at NPL, the echo-free time of about 2.5 ms allows 2.5 cycles of 1 kHz signal. For even a moderately resonant transducer of $Q = 5$ and resonant frequency of 1 kHz, the steady-state can never be observed

directly and classic tone-burst and time windowing methods cannot be used successfully. Of course, for any given size of tank there will be a frequency (depending on the transducer Q -factor) where the accurate use of classical techniques becomes impossible. This observation begs the question of what other methods might be employed to obtain the maximum information from whatever signal is available before the arrival of the echoes.

1.4 EXTRAPOLATION USING SIGNAL MODELLING

The difficulties outlined in Section 1.3 may at least partly be overcome if greater use can be made of the initial transient-dominated part of the waveform. This observation motivates the work summarised in this report.

If we consider the response of an electroacoustic transducer at frequencies at or below its first resonance frequency, it is reasonable to assume the behaviour to be that of a damped harmonic oscillator. This behaviour corresponds to the regime where the device may be modelled with a so-called “lumped-parameter” model. The behaviour of a damped harmonic oscillator is governed by a linear constant coefficient differential equation of the kind familiar from the analysis of systems of masses, springs and dampers (or, in an electrical analogue, a resonant LCR circuit). The solutions of such a differential equation are also well known to be combinations of damped sinusoids or, expressed in an alternative way, as a sum of complex exponentials.

Such functions are used in the work described here to model the signal observed at the start of the tone-burst before the steady-state is reached. This modelling is carried out by fitting the model to the echo-free data available, forming a best fit to the data. Since the signal *model* is not linear (it is composed of exponential functions with unknown time constants), the fitting is performed by use of a nonlinear least-squares algorithm. This approach has significant advantages over linear prediction estimators, such as the classic method of Prony. However, since a reasonable initial estimate is required for the nonlinear least-squares method to be accurate, a modified form of Prony’s method has been used to provide this, with a so-called forward-backward routine and singular value decomposition used to discriminate against spurious model components generated by noise in the data.

Since the model is a physical one (rather than a purely empirical exercise in curve-fitting), an extrapolation may be made from the fitted section of the waveform to obtain the steady-state waveform as it would appear in the absence of boundary reflections. This technique allows measurements to be made on transducers at frequencies lower than that at which the techniques of Section 1.2 may be successfully used. In addition to providing information on the steady state amplitude, the model includes parameters relating to phase, and the phase response of the transducer may also be extracted.

The results in this report indicate that for moderately resonant transducers, errors of less than 1 dB will result if data representing approximately one cycle of the resonance frequency is available for the analysis. For the example from Section 1.3 of the measurement tank at NPL, it can be seen that transducers with resonance frequencies as low as 400 Hz may be calibrated with acceptable accuracy. However, the methods described here become progressively less accurate as the data becomes more noisy and as the Q of the transducer is increased.

For the approach described above, it does not matter whether any of the poles (damped resonances) observed in the data are due to elements in the measurement system other than the projector under test (eg, the receiver, amplifiers, filters, etc). However, in practice with such measurements, it can easily be arranged for the poles due to the measurement system to be at considerably different frequencies to the projector resonance frequency (by use of a small hydrophone receiver, wideband amplifiers, no narrowband filtering, etc). In such a

case, the resonance behaviour observed in the data and reflected in the fitted model will be that due solely to the projector. An estimate of the resonant behaviour of the projector is then obtained directly from the model parameters which can be used to predict the response over a range of frequencies. A summary of this work is given in Section 9.

1.5 ORGANISATION OF THE REPORT

In Section 2 we develop models for describing the response of the device to be calibrated. Models are presented for the free-time response of the device when driven at a single (known) frequency. Throughout this presentation it is assumed that, at any given frequency, the system behaves as a linear damped harmonic oscillator.

In Sections 3, 4 and 5 we present methods for estimating from measured data the values for the parameters defining these models. In Section 3 the estimation problem is introduced in general terms, and the ill-conditioning of this particular estimation problem is illustrated. Estimation methods are described in Sections 4 and 5. These include linear prediction methods (Section 4), and a nonlinear estimation method (Section 5). The latter is based on a safeguarded Gauss-Newton algorithm that uses regularisation to address the ill-conditioning.

In Section 6 we describe the data processing algorithm for deriving from measurements made at a single drive frequency an estimate of the steady-state amplitude of the device at that frequency and the device's resonant behaviour. The algorithm comprises stages of pre-processing the data, windowing the data, applying the estimation methods indicated above, and approaches to validating the results.

In Sections 7 and 8 we present the results of applying the data processing algorithm. In Section 7 results are given for simulated data generated to represent the response of a simple device. The performance of the algorithm is investigated by varying in a systematic manner properties of the simulated data and the device under test. In Section 8 results are given for experimental data relating to an ITC1001 transducer.

In Section 9 we develop a model for the transmitting voltage response of the device regarded as a function of frequency, and we indicate how values for the parameters in this model may be derived from measurements made at a single drive frequency. Estimates of the transmitting voltage response for the ITC1001 obtained using this model are presented. As in Section 2 it is assumed that the system behaves as a linear damped harmonic oscillator.

Finally, in Sections 10 and 11, we discuss scope for further work and we present our conclusions. Acknowledgements and references are given in Sections 12 and 13.

2 MODELS OF THE FREE-TIME RESPONSE

In order to achieve good *prediction* from the model fitted to the data (which is important when direct measurement of the steady-state is limited), we consider here *physical* models. We assume that the measured signal $y(t)$ corresponds to the output of a system that behaves as a linear damped harmonic oscillator, and consequently $y(t)$ solves a linear constant coefficient differential equation of the form

$$\sum_{k=1}^{p+1} c_k D^{k-1} y(t) = 0, \quad (1)$$

where the operator D represents differentiation with respect to time t , i.e.,

$$D^k y(t) \equiv \frac{d^k}{dt^k} y(t).$$

Solutions to this differential equation include damped and undamped sinusoids and real exponential functions. With appropriate initial conditions, (1) defines completely the output $y(t)$.

In Section 2.1 we present a model for $y(t)$ based on deriving solutions to (1). Using *a priori* knowledge about $y(t)$, this model is written in the particular form

$$y(t) = A_0 \sin\{2\pi f_0 t + \phi_0\} + \sum_{k=1}^{n_r} A_k e^{d_k t} \sin\{2\pi f_k t + \phi_k\}, \quad (2)$$

where the first term is used to describe the steady-state behaviour of the device, and the remaining terms its resonant behaviour. In Section 2.2 we present alternative parametrisations for $y(t)$ based on deriving discrete versions of the differential equation (1).

2.1 COMPLEX EXPONENTIAL PARAMETRISATION

We define the *characteristic polynomial* $p(z)$ corresponding to (1) by

$$p(z) = \sum_{k=1}^{p+1} c_k z^{k-1} = \prod_{k=1}^p (z - \beta_k), \quad (3)$$

where β_k , $k = 1, \dots, p$, are the zeros of the polynomial. Then, (1) may be written in the form

$$\prod_{k=1}^p (D - \beta_k I) y(t) = 0, \quad (4)$$

where I is the identity operator, i.e.,

$$I y(t) \equiv y(t).$$

Provided the roots β_k are distinct, the solution to (1) is given by

$$y(t) = \sum_{k=1}^p \alpha_k e^{\beta_k t}. \quad (5)$$

(A repeated root β defines terms of the form $\alpha t^{k-1} e^{\beta t}$, $k = 1, \dots, r$, where r is the multiplicity of the root, in the solution $y(t)$. We assume that the response to be modelled is not characterised by such terms and, consequently, that all roots are distinct.)

Equation (5) defines the *complex exponential* parametrisation for the output $y(t)$. The variables β_k , $k = 1, \dots, p$, are known as *poles*, and are related to the coefficients c_k , $k = 1, \dots, p+1$, through (3), i.e., the poles $\{\beta_k\}$ are the roots of the characteristic polynomial defined by $\{c_k\}$. The variables α_k , $k = 1, \dots, p$, are known as *residues*, and are determined by the initial conditions for $y(t)$. The residues and poles define completely the output $y(t)$. Notice that the parameters in this model are divided into those (the residues) that appear *linearly* in the model, and those (the poles) that appear *nonlinearly*. This observation has important implications for algorithms for fitting the model to measured data.

Since $y(t)$ is a *real-valued* signal, the residues and poles are real or occur in complex conjugate pairs. Let $\beta = \beta_1 + i\beta_2$ be a pole of the system with corresponding residue $\alpha = \alpha_1 + i\alpha_2$, where i is used here to denote $\sqrt{-1}$. If β is real, α must be real also, and the component $\alpha e^{\beta t}$ defined by α and β represents *exponential* behaviour. Alternatively, if β is complex, its complex conjugate β^* is also a pole of the system with corresponding residue α^* . Without loss of generality, we define β to be the pole from this pair for which $\beta_2 > 0$. Then,

$$\alpha e^{\beta t} + \alpha^* e^{\beta^* t} = A e^{dt} \sin(2\pi f t + \phi) = S e^{dt} \cos(2\pi f t) + C e^{dt} \sin(2\pi f t), \quad (6)$$

where

$$A = \sqrt{(2\alpha_1)^2 + (2\alpha_2)^2}, \quad d = \beta_1, \quad f = \frac{\beta_2}{2\pi}, \quad \tan \phi = \frac{\alpha_1}{\alpha_2},$$

and

$$S = A \sin \phi, \quad C = A \cos \phi.$$

The parameters in (6) are the amplitude A , frequency f , damping factor d and phase angle ϕ , and (6) defines *sinusoidal* behaviour that is undamped or damped according to whether β_1 is zero or non-zero. Note that the poles of the system define the frequencies and damping factors, and the residues define the amplitudes and phase angles.

Consequently, an alternative way of writing the model (5) that uses real-valued parameters only is

$$y(t) = \sum_{k=1}^{n_s} A_k e^{d_k t} \sin(2\pi f_k t + \phi_k) + \sum_{k=1}^{n_e} \alpha_k e^{\beta_k t}, \quad (7)$$

where n_s is the number of sinusoidal components and n_e is the number of exponential components, and

$$p = 2n_s + n_e.$$

In much of this work, we will write (7) in the particular form

$$y(t) = A_0 \sin\{2\pi f_0 t + \phi_0\} + \sum_{k=1}^{n_r} A_k e^{d_k t} \sin\{2\pi f_k t + \phi_k\}, \quad (8)$$

where

- a) we have included no exponential term, i.e., $n_e = 0$, and
- b) we have written the ($n_s = p/2$) sinusoidal terms as a single undamped sinusoid to represent the steady-state behaviour of the system together with n_r damped sinusoids to represent the resonant behaviour. In (8), A_0 is the steady-state amplitude, f_0 is the drive frequency, f_k , $k > 0$, are the resonance frequencies for the system, and d_k , $k > 0$, are the damping factors for the device that are related to its Q -factors according to

$$Q_k = -\frac{\omega_k}{2d_k} = -\frac{\pi f_k}{d_k}, \quad (9)$$

where $\omega_k = 2\pi f_k$ is the angular frequency corresponding to f_k .

In some circumstances we may wish to include *a priori* knowledge about the system output. How this is done depends on the model used and the approach chosen to fit the model to the data. For example, accurate knowledge of the frequency of the signal driving the system fixes the frequency parameter f_0 for the undamped component. Furthermore, if we have determined elsewhere the frequency and Q -factor for a resonance of the system, this gives us information about the frequency and damping factor parameters for one of the damped components. This information can be incorporated as above by replacing, for example, f_1 and d_1 by their known values. Alternatively, if these values arise from measurement and are known with an associated uncertainty, the information can be incorporated in the estimation problem by including additional observation equations. The latter approach is discussed in Section 3.

2.2 PRONY PARAMETRISATIONS

Another class of parametrisation for $y(t)$ is based on deriving discrete versions of the differential equation (1). These parametrisations are important because the discrete versions of (1) on which they are based underpin the linear prediction methods (or Prony methods) described in Section 4. The parametrisations involve replacing the poles β_k by so-called Prony parameters γ_k or δ_k . These are described below.

Let Π be the *forward shift operator*,

$$\Pi y(t) = y(t + \delta t),$$

and Δ the *divided difference operator*,

$$\Delta y(t) = \frac{y(t + \delta t) - y(t)}{\delta t}, \quad \Delta = \frac{(\Pi - I)}{\delta t}.$$

Then,

$$(\Delta + \zeta_j I)e^{\beta_j t} = 0, \quad \zeta_j = \frac{(1 - e^{\beta_j \delta t})}{\delta t}, \quad (10)$$

and, consequently, $y(t)$ satisfies the *difference equation*

$$\prod_{j=1}^p (\Delta + \zeta_j I)y(t) = 0,$$

which can be written as

$$\sum_{k=1}^{p+1} \gamma_k \Delta^{k-1} y(t) = 0.$$

The parameters γ_k , $k = 1, \dots, p+1$, are called *difference form Prony parameters*. The poles β_k are recovered from the difference form Prony parameters in the following way: given γ_k , $k = 1, \dots, p+1$, we form the characteristic polynomial

$$p_\gamma(z) = \sum_{k=1}^{p+1} \gamma_k z^{k-1} = \prod_{j=1}^p (z + \zeta_j),$$

and determine its roots $-\zeta_j$, $j = 1, \dots, p$. The poles are then obtained using (10).

Furthermore,

$$(\Pi - \rho_j I)e^{\beta_j t} = 0, \quad \rho_j = e^{\beta_j \delta t}, \quad (11)$$

and, consequently, $y(t)$ satisfies the *recurrence equation*

$$\prod_{j=1}^p (\Pi - \rho_j I)y(t) = 0,$$

which can be written as

$$\sum_{k=1}^{p+1} \delta_k \Pi^{k-1} y(t) = 0.$$

The parameters δ_k , $k = 1, \dots, p+1$, are called *recurrence form Prony parameters*. The poles β_k may be recovered from the recurrence form parameters in a similar way to that described for the difference form parameters.

3 ESTIMATION

We are concerned with fitting models $y(t)$ of the forms discussed in Section 2 to data $\{(t_i, y_i): i = 1, \dots, m\}$ measured at equally spaced times t_i . If ε_i is the measurement error for the i th data value y_i , and the values t_i are known accurately, we wish to determine the function $y(t)$ that satisfies the *observation equations*

$$y_i = y(t_i) + \varepsilon_i, \quad i = 1, \dots, m. \quad (12)$$

If we assume the errors ε_i are uncorrelated samples from a Gaussian probability distribution with mean zero and standard deviation σ , unbiased and efficient estimates of the parameters defining $y(t)$ are obtained by solving

$$\text{minimise } \sum_{i=1}^m \{y_i - y(t_i)\}^2 \quad (13)$$

with respect to the parameters of $y(t)$. The residuals $e_i = y_i - y(t_i)$ evaluated at the solution provide estimates of the errors ε_i , and an estimate of σ is given by the root-mean-square error s where

$$s = \sqrt{\frac{1}{m-n} \sum_{i=1}^m e_i^2}, \quad (14)$$

and n is the number of parameters defining $y(t)$.

If *a priori* knowledge of any of the model parameters is available, this information can be represented by additional observation equations, and the estimation problem is modified accordingly. For example, if f_r and d_r are the frequency and damping factor for a resonance of the system, the estimation problem becomes

$$\text{minimise } \sum_{i=1}^m \{y_i - y(t_i)\}^2 + u^2 \{f_r - f_1\}^2 + v^2 \{d_r - d_1\}^2, \quad (15)$$

where u and v are “weights” that are used to reflect the relative accuracy between the measured data and the *a priori* knowledge.

Although we assume the device behaves linearly, the model $y(t)$ is a nonlinear function of its parameters, and problems (13) and (15) are *nonlinear least-squares* problems. Algorithms for solving these problems are discussed in Section 5. Furthermore, because of the particular type of model considered, we will see below that these problems are difficult to solve.

There is a large body of work in the literature concerned with *linear prediction* methods, also known as *Prony* methods [1, 2]. These methods have the following advantages: (a) they involve the solution of (simpler) linear problems, and (b) they are direct methods, i.e., unlike algorithms for the nonlinear problem, they do not require that initial estimates of the parameters are provided. However, because the methods do not reflect properly the error structure associated with the data, they do not provide unbiased and efficient estimates of the model parameters. Nevertheless, linear prediction methods are important because they

provide a means of generating starting estimates for algorithms for the nonlinear problem. Linear prediction methods are discussed in Section 4.

It is well-known that the estimation problem defined by (13), where $y(t)$ is replaced by a sum of exponential terms, is ill-conditioned. We can illustrate this ill-conditioning in the following way. Suppose we are given data (t_i, y_i) , $i = 1, \dots, m$, corresponding to a drive voltage of known frequency f_0 . Let us choose a model comprising $p = 4$ complex exponential terms (see (5)) in which the first two terms correspond to an undamped sinusoidal component representing the steady-state response of the device, and the second two terms to a damped component representing its resonant behaviour, i.e.,

$$y(t) = \sum_{k=1}^4 \alpha_k e^{\beta_k t}, \quad (16)$$

where

$$\beta_1 = i2\pi f_0, \quad \beta_2 = -i2\pi f_0, \quad \beta_3 = d_1 + i2\pi f_1, \quad \beta_4 = d_1 - i2\pi f_1,$$

and d_1 and f_1 are the resonance damping and frequency. If we fix d_1 and f_1 , the model is *linear* in the residue parameters α_k , $k = 1, \dots, 4$, and so the problem (13) reduces to a linear least-squares problem. Having determined the residues as the solution to this linear problem, we compute the value of the residual sum of squares

$$S(d_1, f_1) = \sum_{i=1}^m \{y_i - y(t_i)\}^2.$$

We may repeat the process for a number of choices of d_1 and f_1 and display the residual sum of squares surface S as a function of d_1 and f_1 . Four such surfaces using data obtained from measurements of an ITC1001 device, and corresponding to various choices of the drive frequency f_0 and the number of points m , are shown in Figure 4. (Here, $m = 56$ points corresponds to a time-window spanning between one and two cycles of the resonance frequency of the device.) In Table 1 we give the position of the minimum value for each surface expressed in terms of its resonance frequency and Q -factor (see (9)).

The ‘‘shape’’ of all four surfaces suggests that the residual sum of squares is insensitive to changes in the resonance damping in the sense that a large change in the value of d_1 produces only a small change in the value of the residual sum of squares. Furthermore, compared with Figure 4(a), the sensitivity to the resonance frequency is reduced for a shorter time-window (Figure 4(b)), and for measurements corresponding to drive frequencies that are different (Figures 4(c) and 4(d)) from the resonance frequency of the device (which is approximately 18 kHz). We expect the former because a shorter time-window means that less information about the resonant behaviour is available. The latter may be an indication that the model describes the data less well as the drive frequency moves away from the resonance frequency.

The surfaces illustrate the inherent ill-conditioning of the problem of finding a minimum of the residual sum of squares function, and suggest that the problem of finding accurately the values of the parameters defining this minimum is inherently ill-posed. Thus *any* numerical procedure for doing this can be expected to experience difficulties. The fact that the minimizers listed in Table 1 are different, particularly in terms of the resonance damping, means that to obtain physically meaningful and consistent solutions the inclusion of *a priori* information is important.

Finally, if we include additional damped components in the model (16), the (hyper-) surfaces that show the dependence of S on the parameters of $y(t)$ become much more complicated and typically possess a number of local minima. This fact degrades further the performance of numerical algorithms for solving (13), and means that the solution to which these algorithms converge can be sensitive to the choice of starting estimates.

<i>Drive frequency (kHz), number of points</i>	<i>Resonance frequency (kHz)</i>	<i>Q-factor</i>
(a) $f_0 = 18, m = 112$	17.796	3.01
(b) $f_0 = 18, m = 56$	17.633	2.89
(c) $f_0 = 11, m = 112$	17.878	2.75
(d) $f_0 = 25, m = 112$	17.796	3.06

Table 1: Resonance frequency and Q -factor defining the minimum of the residual sum of squares surfaces shown in Figure 4.




Figure 4: Residual sum of squares as a function of the resonance frequency and damping factor for the ITC1001 transducer and the cases (a) $f_0 = 18$ kHz, $m = 112$ (top-left), (b) $f_0 = 18$ kHz, $m = 56$ (top-right), (c) $f_0 = 11$ kHz, $m = 112$ (bottom-left), and (d) $f_0 = 25$ kHz, $m = 112$ (bottom-right). The functions are displayed as 3D surfaces above corresponding contour plots.

4 LINEAR ESTIMATION METHODS

4.1 LEAST-SQUARES PRONY METHOD

Linear estimation methods are based on the observation that the model values $y(t_i)$ satisfy a set of recurrence equations, with undetermined parameters $\delta_1, \delta_2, \dots, \delta_{p+1}$:

$$\delta_1 y(t_i) + \delta_2 y(t_{i+1}) + \dots + \delta_{p+1} y(t_{i+p}) = 0, \quad i = 1, \dots, m - p. \quad (17)$$

Replacing the values $y(t_i)$ by the measured values y_i , the recurrence equations are no longer satisfied exactly, but instead

$$\delta_1 y_i + \delta_2 y_{i+1} + \dots + \delta_{p+1} y_{i+p} = e_i, \quad i = 1, \dots, m - p, \quad (18)$$

for some errors e_i .

Figure 5 illustrates the way the recurrence equations are formed for the case $p = 4$ (four poles and residues defining two sinusoidal components) and a data set composed of 21 equispaced samples ($m = 21$). The numbers beside a data sample indicate the use of that sample within particular recurrence equations. For example, the first data sample appears in the first recurrence equation, the second data sample appears in the first and second recurrence equations, and the third data sample is used in the first three recurrence equations. In total there are $m - p = 17$ recurrence equations to be solved for $p + 1 = 5$ Prony parameters δ_i .




Figure 5: Formation of the recurrence equations in the application of the least-squares Prony method for the case $p = 4$ and $m = 21$.

The equations (18) constitute a set of linear equations for the parameters $\delta_1, \delta_2, \dots, \delta_{p+1}$. If we set $\delta_{p+1} = 1$ (note the parameters are determined only up to a scale factor), we determine estimates of the parameters $\mathbf{d} = (\delta_1, \delta_2, \dots, \delta_p)^T$ by solving the least-squares problem

$$\text{minimise } \mathbf{e}^T \mathbf{e}, \quad \mathbf{e}^T \mathbf{e} = \sum_{i=1}^{m-p} e_i^2, \quad (19)$$

with respect to \mathbf{d} . The poles β_k for the system are then recovered from the parameters $\delta_1, \delta_2, \dots, \delta_{p+1}$. Finally, the residues α_k are obtained by fitting the model (5), now regarded as a function of α_k only, to the data: this is another linear least-squares problem.

4.2 IMPROVING PRONY'S METHOD

The procedure described above is the *least-squares Prony method* described in [1, 2, 3, 4]. It is noted in [5] that although the method is consistent as $\sigma \rightarrow 0$, it is inconsistent as $m \rightarrow \infty$. (In other words, in the presence of noise, the method gives estimates that do not converge to the true solution as the number of sampled points increases.) Consequently, the method is only useful for low noise levels regardless of how many measurements are made. This has led to the development of a number of variations on the basic method, as follows:

- a) the *modified least-squares Prony method* in which, instead of adjacent samples, every l th sample is used to satisfy the recurrence equations [1],
- b) the use of a large number p of poles with *forward* and *backward prediction* to help in distinguishing between “true” system poles and poles associated with measurement noise [2],
- c) the *Prony-SVD* method in which a large number p of poles is used with the singular value decomposition [6]. (An alternative to the singular value decomposition is to use complete orthogonal factorization [7].)

There has also been work concerned with including *a priori* information. In [2], this is done by applying a filter to the data to remove known system poles before applying a Prony method; in [7, 8], recurrence equations relating samples of the input *and* output signals are used to model the system's transfer function.

It is a property of least-squares estimation that we expect that the solution of (19) will provide unbiased and efficient estimates of the parameters \mathbf{d} in the case that the errors e_i are samples of random variables that are independent and identically distributed. Suppose that ε_i is the error associated with the i th measurement as in (12). It is reasonable to assume that the errors ε_i are samples of random variables that are independent and identically distributed, i.e., their variances and covariances are, respectively,

$$E(\varepsilon_i^2) = \sigma^2 \quad \text{and} \quad E(\varepsilon_i \varepsilon_j) = 0, \quad i, j = 1, \dots, m, \quad i \neq j. \quad (20)$$

We show below that this property is *not* inherited by the e_i , and consequently the least-squares Prony method cannot provide unbiased and efficient estimates of the parameters \mathbf{d} .

Substituting (12) into (18), we obtain

$$e_i = \sum_{k=1}^{p+1} \delta_k \{y(t_{i-1+k}) + \varepsilon_{i-1+k}\} = \sum_{k=1}^{p+1} \delta_k \varepsilon_{i-1+k}. \quad (21)$$

It follows that the variance of e_i is given by

$$E(e_i^2) = \sum_{k=1}^{p+1} \delta_k^2 E(\varepsilon_{i-1+k}^2) = \sigma^2 \sum_{k=1}^{p+1} \delta_k^2, \quad (22)$$

and the covariance of e_i with e_{i+j} , $j = 1, \dots, p$, is given by

$$E(e_i e_{i+j}) = \sum_{k=1}^{p+1-j} \delta_{j+k} \delta_k E(\varepsilon_{i+j-1+k}^2) = \sigma^2 \sum_{k=1}^{p+1-j} \delta_{j+k} \delta_k. \quad (23)$$

Since we do not expect in general the right hand side of (23) to be zero, we conclude that the errors e_i are not independent.

To obtain unbiased and efficient estimates of the parameters \mathbf{d} , we solve the *weighted* least-squares problem

$$\text{minimise } \mathbf{e}^T \mathbf{V}^{-1} \mathbf{e}, \quad (24)$$

with respect to \mathbf{d} , where V is the covariance matrix for the errors e_i defined by (22) and (23). This is the *weighted least-squares Prony method*. The problem posed in (24) correctly accounts for the error structure for the original measured data. However, because the elements of V depend on the unknown parameters \mathbf{d} , it is necessary to apply an iterative scheme in which at each iteration V is formed using the estimates of \mathbf{d} obtained at the previous iteration. The procedure is begun by setting V to be the identity matrix which corresponds to the original least-squares Prony method. The covariance matrix V needs to be further modified if the data in the above analysis comes from the filtering of the original measured data to remove known components.

5 NONLINEAR ESTIMATION METHODS

The estimation problems (13) and (15) given in Section 3 are nonlinear least-squares problems. Standard algorithms exist (see, for example, [9]) for solving this type of problem including the Gauss-Newton and full-Newton methods. These are iterative methods that at each iteration take a step towards the minimum by solving a linear least-squares problem. The algorithms differ in the amount of information that needs to be supplied about the model and how this information is used: the Gauss-Newton method requires that the first derivatives of the model with respect to its parameters are available, whereas in a full-Newton method the second derivatives are also required.

In addition to choosing an algorithm specific to least-squares problems, we can exploit structure in the model. The parameters naturally separate into those (the poles or equivalently frequencies and damping factors) that appear nonlinearly in the model, and those (the residues or equivalently amplitudes and phases) that appear linearly. The use of variable projection methods as described in [7, 10, 11, 12, 13] exploits this structure in the solution of the nonlinear least-squares problem. Yet another approach is presented in [5], involving a reparametrisation of the model and the solution of a nonlinear eigenvalue problem.


It is usual to use the solution from a linear prediction method to provide starting estimates for solving the nonlinear estimation problem. In Section 4 we noted that in order to obtain good fits using linear prediction methods, it may be necessary to choose a model for which the number n_r of damped sinusoidal components exceeds what is believed to be the number of “true” resonances. We illustrated in Section 3 that the nonlinear estimation problem (15) is inherently ill-conditioned, i.e., large changes in the frequency and damping factor parameters may produce small changes in the residual sum of squares function that we wish to minimise. These facts can make it difficult for standard algorithms to converge satisfactorily to a solution.

To address these issues, and because generally we expect to be solving “small-residual” problems, we use a Gauss-Newton algorithm, safe-guarded with a line search algorithm (Appendix A.1), with *a priori* information incorporated as described in Section 3, equation (15). Furthermore, if additional damped sinusoidal components are used to define the initial model fit, these are either explicitly removed prior to applying the Gauss-Newton algorithm or *regularisation* is used to replace the estimation problem to be solved by one that is better conditioned. A description of the Gauss-Newton algorithm for solving a general nonlinear least-squares problem is given in Appendix A.1, and its modification to include a line search and regularisation are discussed, respectively, in Appendices A.2 and A.3.

6 DESCRIPTION OF DATA PROCESSING ALGORITHM

Presented below is a flowchart for the data processing algorithm used to derive estimates of the model parameters from data measured at a single drive frequency. The steps in this flowchart are described in detail in Section 6.1 to 6.6. We illustrate the application of the data processing algorithm using data relating to an ITC1001 transducer.

Flowchart for Data Processing Algorithm



6.1 ACQUIRE DATA

An ITC1001 transducer, with a resonance at approximately 18 kHz, is driven with a discrete-frequency tone-burst signal of 5 kHz. Figure 6(a) shows the recorded waveform.

6.2 REMOVE ANY NON-ZERO DC LEVEL IN THE ACQUIRED DATA

Prior to the turn-on of the device, the recorded waveform will be essentially constant. This part of the signal will record any background dc level contaminated by measurement noise. The models presented in Section 2 assume that this level is zero, and if this is not the case it

is necessary to correct the data. This correction is made by estimating the dc level using the mean of the data values constituting that part of the signal prior to the turn-on of the device, and subtracting this value from the complete signal. It is important that only data prior to the turn-on is processed in this way, otherwise a poor estimate of the dc level may be obtained. Figure 6(b) shows the result of correcting in this way the data from Figure 6(a). In this example the first 15 points of the waveform were used to provide the estimate of the dc level, and the correction is such that visually there is little difference between the waveforms shown.

An alternative to the processing described above is to augment the signal model (8) with a constant term. However, the linear prediction methods described in Section 4 that are based on Prony parametrisations of the model do not permit this more general form of model.




Figure 6: Data for the ITC1001 transducer corresponding to a 5 kHz drive voltage: (a) acquired data (left), and (b) after removal of background dc level (right).

6.3 IDENTIFY FREE-TIME WINDOW

Time values are chosen, respectively, immediately after the turn-on of the device and immediately before the turn-off of the device or the arrival of the first reflection. That part of the signal between these times defines the free-time response of the device, and is that part of the signal that we model using the methods described in Sections 4 and 5.

The “crosses” marked in Figure 6(b) identify particular time-values used in the following examples. The left-most cross marks the location of a time t_0 immediately after the turn-on of the device. The remaining crosses mark the locations of time-points used to simulate various choices of the free-time response. The second cross at $t = t_1$ defines the end of a time-window comprising $m = 56$ points and between one and two cycles of the resonant response of the device. The third cross at $t = t_2$ defines a time-window comprising $m = 112$ points and between three and four cycles of the resonant response. Finally, the first and fourth crosses define a time-window that contains a sufficient part of the steady-state of the response to allow direct measurement of the steady-state voltage to be made.

6.4 APPLY LINEAR PREDICTION METHODS

The drive frequency f_0 is specified, together with the frequency f_r and Q -factor Q_r , defining any known resonant behaviour for the device. (Recall that f_r and Q_r may be used to define the damping factor d_r for the resonance (equation (9)) as well as the poles β_k for these components.) Using this information, the data defining the free-time response of the device is filtered to remove the specified drive and resonant components. An algorithm for doing this

is described in [2]. The filtered data samples are then processed to give estimates of the remaining unknown poles. These estimates are combined with the poles defining the known drive and resonant components, and finally the residues α_k are determined as in the standard application of Prony's method (Section 4.1).

In the same way that the least-squares Prony method described in Section 4.1 is sensitive to the presence of noise, we find that this is true also of the filtering operation. Motivated by the *modified* least-squares Prony method that was introduced in Section 4.2 as a means of improving the performance of the basic least-squares Prony method, the filter may be applied to every l th sample, rather than to adjacent samples. In Figure 7 (note the different scales used for the voltage axis) we show the results of this filtering operation using (a) $l = 1$, and (b) $l = 10$, applied to the data contained within the time-window $[t_0, t_2]$ comprising $m = 112$ points. (In this, and subsequent figures, we have shifted the time axis so that t_0 coincides with the origin.) We notice that the resonant behaviour is much clearer in the second of these figures. Furthermore, if we estimate that there are 3.25 cycles in this figure taking 0.18 milliseconds, then the frequency of the resonance is approximately 18 kHz as expected.




Figure 7: Incorporating *a priori* information about the drive voltage f_0 by filtering. A filter is applied to the data in the time-window $[t_0, t_2]$ comprising $m = 112$ points with (a) $l = 1$ (left), and (b) $l = 10$ (right).

We now model the filtered data shown in Figure 7(b) using the following linear prediction algorithms (Section 4):

- a) the least-squares Prony method,
- b) the weighted least-squares Prony method, and
- c) the least-squares Prony-SVD method.

In Figure 8 we show the solutions returned by these methods when applied to the data contained within the time-window $[t_0, t_1]$ comprising $m = 56$ points and delimited by the time-origin and the vertical dashed line. We also show the measured data and the computed model evaluated beyond this window, thus allowing us to quantify the ability of each method to estimate the steady-state amplitude.

For the least-squares and weighted least-squares Prony methods, it is necessary to choose the number n of components to model the data (equivalently, the number of poles $p = 2n$). In Figures 8(a) and 8(b) (note the different scales used for the voltage axis) we show the

solutions for $n = 1$ that corresponds to choosing a model composed of a single undamped sinusoidal component with a fixed frequency of 5 kHz and a single damped sinusoidal component of unknown frequency representing the resonant behaviour. With the least-squares Prony method, the data within the given time-window is fitted very poorly, and the method fails completely to pick-out the resonant behaviour. In contrast, with the weighted least-squares Prony method we obtain good information about the resonant behaviour of the device and the steady-state amplitude.

For the least-squares Prony-SVD method, we choose a large value for n and specify the number $r \leq n$ of these components to be associated with the unknown resonant behaviour of the device. In Figure 8(c) we show the solution for the choice $n = 5$ and $r = 1$. The data within the given time-window is not fitted as well as in Figure 8(b), but a good estimate of the steady-state amplitude is obtained.

Finally, we can apply forward and backward versions of the least-squares Prony-SVD algorithm to help distinguish between “true” system poles and poles associated with measurement noise. This, in turn, is used to validate the choices made for r and n . In Figure 8(d) we show the locations within the complex plane of the zeros of the characteristic polynomials obtained from forward prediction (using the “+” symbol) and backward prediction (using the “o” symbol). We note that the zeros occur in conjugate pairs with each pair defining a sinusoidal component, and those obtained from forward prediction lie within the unit circle (which is also marked). Furthermore, we expect the zeros obtained from backward prediction that are associated with “true” resonances to lie outside the unit circle, whereas those associated with noise will remain inside. The results shown in Figure 8(d) support, therefore, the use in this example of a model involving a single resonant component.

6.5 APPLY NONLINEAR METHODS

The results obtained from the linear prediction undertaken in Step 4 provide initial parameter values for the non-linear estimation algorithm described in Section 5. Again, the drive frequency f_0 is specified, together with the frequency f_r and damping factor d_r defining any known resonant behaviour for the device. This time, however, *a priori* information about a resonance is regarded as measurement information, and is “weighted” accordingly using knowledge of the accuracy of the *a priori* information relative to the measured data. The drive frequency is assumed to be known sufficiently accurately for this parameter to be fixed.

It is assumed that the model for fitting the data has the same number n of components as used for the linear prediction in Step 4. If n exceeds the number r of “true” resonances, we can either choose (a) explicitly to remove those additional components, or (b) to set a regularisation factor λ to be used to constrain the damping factors associated with these components (Section 5 and Appendix A.2).

In Figure 9 we show the solutions returned by the non-linear estimation algorithm using different sets of initial parameter values. For Figure 9(a), starting values were provided by the weighted least-squares Prony method and the corresponding initial fit is that shown in Figure 8(b). The fitted function comprises a single undamped and a single damped sinusoidal component. For Figure 9(b), starting values were provided by the least-squares Prony-SVD method and the corresponding initial fit is that shown in Figure 8(c). In this case, the four additional components associated with the initial fit were explicitly removed so that the fitted functions shown in Figure 9 have the same number of components.




Figure 8: Estimates of the voltage output obtained by fitting the data in the time-window $[t_0, t_1]$ comprising $m = 56$ points using (a) the least-squares Prony method (top-left), (b) the weighted least-squares Prony method (top-right), and (c) the least-squares Prony-SVD method (bottom-left). Details of the models used are provided in the text. We show (bottom-right) the positions of the zeros of the characteristic polynomials arising from forward and backward prediction using the least-squares Prony-SVD method.




Figure 9: Estimates of the voltage output obtained by fitting the data in the time-window $[t_0, t_1]$ comprising $m = 56$ points using non-linear estimation and (a) the weighted least-squares Prony method (left), and (b) the least-squares Prony-SVD method (right), to provide starting estimates.

6.6 VALIDATE ESTIMATES OF THE MODEL PARAMETERS

The computed estimates of the model parameters are validated by comparing the results obtained by applying:

- a) different estimation methods to the same data set, and
- b) the same estimation method to different amounts of data.

We consider the first of these in Table 2 where we list the computed values for the steady-state amplitude and the resonant behaviour obtained using the various analysis methods presented above. We also give the root-mean-square residual value for each fitted function which measures the accuracy to which the data (within the given time-window) is fitted. We see that there is good agreement between the various estimates of the steady-state amplitude. However, there is greater variability in the estimates of the parameters defining the resonant behaviour of the device, in particular, its Q -factor. Indeed, the estimates of the Q -factor are considerably different from the expected value of between 3 and 4. These differences may be a result of the ill-conditioning of the estimation problem that is discussed in Section 3, or may be due to a deficiency in the model.

The second validation technique is considered in Section 8 where we present estimates of the transmitting voltage response for the ITC1001 transducer obtained by analysing data taken from different time-windows.

It is also useful to compare the results obtained from the analysis of data corresponding to different drive frequencies. In this case, we expect the values of the parameters defining the resonant behaviour of the device to be similar. Again, this is considered in Section 8 where we present estimates of the transmitting voltage response for the ITC1001 transducer obtained from knowledge of its resonant behaviour estimated at different drive frequencies.

<i>Method</i>	<i>Steady-state amplitude (mV)</i>	<i>Resonance frequency (kHz)</i>	<i>Q-factor</i>	<i>Root-mean-square residual σ (mV)</i>
(a)	0.9567	17.636	2.55	0.0735
(b)	0.9557	17.570	2.61	0.0823
(c)	0.9561	17.475	2.75	0.0962
(d)	0.9537	17.402	2.03	0.0487
(e)	0.9537	17.403	2.03	0.0487

Table 2: Estimates of steady-state amplitude and resonant behaviour produced by (a) weighted least-squares Prony method (Figure 8(b)), (b) (forward) least-squares Prony-SVD method (Figure 8(c)), (c) (backward) least-squares Prony-SVD method, (d) non-linear estimation (Figure 9(a)), and (e) non-linear estimation (Figure 9(b)).

7 RESULTS FOR SIMULATED DATA

In this section we present results obtained from applying the data processing algorithm described in Section 6 to simulated data used to represent the measured response of a device. The purpose of these simulations is to investigate how well the data processing algorithm performs, and to understand its limitations in terms of various key properties of the data and the simulated device. Simulated data is used in preference to real measurement data because there is comparatively little cost in generating such data and we have complete control (and knowledge of) the properties of the data.

A number of simulations are described in which the following properties of the simulated data or device are varied in a systematic manner: (a) the time-window containing the free-time response, (b) the data sampling rate, (c) the noise level (signal-to-noise ratio), (d) the drive frequency, and (e) the Q -factor for the device. First, we describe the data on which the simulations are based.

Figure 10 shows simulated data representing the response of a device when driven at its resonance frequency. We have chosen

- a) the resonant behaviour to be defined by a frequency of 2 Hz, a Q -factor of 3, and an amplitude of 1 V, and
- b) the steady-state behaviour to have a frequency of 2 Hz and an amplitude of 1 V.

In the absence of measurement error, the response of the device during a 5 second period is shown in Figure 10. This period covers ten cycles of the resonance and the steady-state, and it is clear that the steady-state amplitude is attained within this period.

Figure 11 shows simulated data representing the response of the same device but driven at different frequencies. For the response shown on the left, the steady-state behaviour is defined by a frequency of 1.5 Hz and an amplitude of 0.1 V, and for that on the right the frequency is 2.5 Hz and the amplitude is 0.8 V. As in Figure 10, we note that the steady-state amplitude is reached within the time period shown.

Figure 12 shows simulated data representing the response of devices with different Q -factors. The devices are driven at their resonance frequency of 2 Hz (as in Figure 10), but the response on the left is generated using a Q -factor of 5, and that on the right using a Q -factor of 10. With these larger Q -factors it takes longer for steady-state to be reached and, indeed, for the case $Q = 10$, steady-state is not reached within the time period of 5 seconds shown.

The results of the simulations are given in Sections 7.1 to 7.5. For each set of simulated data, the data processing algorithm was applied in the same way. The data was filtered to remove the steady-state component defined by the known drive frequency, and the Prony-SVD method then used to provide an initial fit to the data. The number of components or poles in this fitted model was chosen on the basis of the plot of the zeros of the characteristic polynomials from forward and backward prediction (see Section 6.4, Figure 8). For some of the more difficult data sets, a large number of poles was necessary to give a sensible initial fit to the data. The components in this model associated with the noise in the data (identified as those giving zeros inside the unit circle from both forward and backward prediction) were explicitly removed, and a model comprising a single undamped sinusoidal component and single damped component fitted to the data using the nonlinear estimation algorithm.




Figure 10: Simulated data representing the response of a device when driven at its resonance frequency.




Figure 11: Simulated data representing the response of a device at frequencies below (on the left) and above (on the right) its resonance frequency.




Figure 12: Simulated data representing the response of devices with different Q -factors and driven at their resonance frequency.

7.1 VARYING THE TIME-WINDOW

Random measurement error sampled from a Gaussian distribution with mean zero and standard deviation 0.01 (representing a (peak) signal-to-noise ratio of 100) is added to the data shown in Figure 10. In Table 3 we give the results obtained from analysis of data taken from various time-windows where we assume that the data is sampled at the same rate. Results are presented for time-windows spanning (a) four cycles of resonance (and $m = 400$ points), (b) two cycles of resonance (and $m = 200$ points), (c) one cycle of resonance (and $m = 100$ points), and (d) half a cycle of resonance (and $m = 50$ points).

<i>Time-window (cycles)</i>	<i>Estimated steady-state amplitude (V)</i>	<i>Estimated resonance frequency (Hz)</i>	<i>Estimated Q-factor</i>	<i>Root-mean-square residual σ (V)</i>
4	1.000561	2.001449	2.999448	0.009755
2	1.005501	1.999270	3.032290	0.009692
1	1.024263	1.989188	3.111726	0.010634
0.5	0.273679	3.235722	2.569133	0.009631

Table 3: Estimates of steady-state amplitude and resonant behaviour for data simulated within the following time-windows: (a) 4, (b) 2, (c) 1, and (d) half a cycle of resonance.

7.2 VARYING THE DATA SAMPLING RATE

Random measurement error sampled from a Gaussian distribution with mean zero and standard deviation 0.01 (representing a (peak) signal-to-noise ratio of 100) is added to the data shown in Figure 10. In Table 4 we give the results obtained from analysis of data taken from a fixed time-window (spanning two cycles of resonance or 1 second) containing various numbers of points representing different sampling rates for collecting the data. Results are presented for (a) $m = 400$, (b) $m = 200$, (c) $m = 100$, and (d) $m = 50$.

<i>Sampling-rate (samples/sec)</i>	<i>Estimated steady-state amplitude (V)</i>	<i>Estimated resonance frequency (Hz)</i>	<i>Estimated Q-factor</i>	<i>Root-mean-square residual σ (V)</i>
400	0.995925	1.995694	2.972751	0.010363
200	1.005072	1.998376	3.018845	0.009758
100	0.981232	2.007462	2.837397	0.010000
50	0.995773	2.011301	2.943078	0.009144

Table 4: Estimates of steady-state amplitude and resonant behaviour for data simulated in a fixed time-window with various sampling rates: (a) 400, (b) 200, (c) 100 cycle, and (d) 50 samples/second.

7.3 VARYING THE NOISE LEVEL

Random measurement error sampled from a Gaussian distribution with mean zero and standard deviation σ (representing a (peak) signal-to-noise ratio of $1/\sigma$) is added to the data shown in Figure 10. In Table 5 we give the results obtained from analysis of data taken from a fixed time-window (spanning two cycles of resonance or 1 second) and a fixed sampling rate for collecting the data ($m = 200$ points). Results are presented for (a) $\sigma = 1/100 = 0.01$, (b) $\sigma = 1/75 \approx 0.013$, (b) $\sigma = 1/50 = 0.02$, and (d) $\sigma = 1/25 = 0.04$.

<i>Noise (V)</i>	<i>Estimated steady-state amplitude (V)</i>	<i>Estimated resonance frequency (Hz)</i>	<i>Estimated Q-factor</i>	<i>Root-mean-square residual σ (V)</i>
1/100	0.994144	2.003179	2.953522	0.010922
1/75	0.987408	2.010833	2.919317	0.012871
1/50	1.036989	1.997059	3.257157	0.021022
1/25	0.996407	2.007696	3.034778	0.042551

Table 5: Estimates of steady-state amplitude and resonant behaviour for data simulated with noise with the following signal-to-noise ratios: (a) 100, (b) 75, (c) 50, and (d) 25.

7.4 VARYING THE DRIVE FREQUENCY

Random measurement error sampled from a Gaussian distribution with mean zero and standard deviation 0.01 (representing a (peak) signal-to-noise ratio of 100) is added to the data shown in Figures 10 and 11. In Table 6 we give the results obtained from analysis of data taken from a fixed time-window (spanning two cycles of resonance or 1 second) and a fixed sampling rate for collecting the data ($m = 200$ points). Results are presented for a drive frequency of (a) 1.5 Hz (with steady-state amplitude of 0.1 V), (b) 2 Hz (with a steady-state amplitude of 1 V), and (c) 2.5 Hz (with a steady-state amplitude of 0.8 V).

<i>Drive frequency (Hz)</i>	<i>Estimated steady-state amplitude (V)</i>	<i>Estimated resonance frequency (Hz)</i>	<i>Estimated Q-factor</i>	<i>Root-mean-square residual σ (V)</i>
1.5	0.101186	1.995613	2.993943	0.009058
2.0	0.994609	1.993617	2.943450	0.009316
2.5	0.802332	2.001411	2.994999	0.009679

Table 6: Estimates of steady-state amplitude and resonant behaviour for data simulated with the following drive frequencies (and steady-state amplitudes): (a) 1.5 Hz (0.1 V), (b) 2 Hz (1.0 V), (c) and 2.5 Hz (0.8 V).

7.5 VARYING THE Q-FACTOR

Random measurement error sampled from a Gaussian distribution with mean zero and standard deviation 0.01 (representing a (peak) signal-to-noise ratio of 100) is added to the data shown in Figures 10 and 12. In Table 7 we give the results obtained from analysis of data taken from a fixed time-window (spanning two cycles of resonance or 1 second) and a fixed sampling rate for collecting the data ($m = 200$ points). Results are presented for a Q -factor of (a) 3, (b) 5, and (c) 10.

<i>Q-factor</i>	<i>Estimated steady-state amplitude (V)</i>	<i>Estimated resonance frequency (Hz)</i>	<i>Estimated Q-factor</i>	<i>Root-mean-square residual σ (V)</i>
3	0.999744	2.001787	2.995748	0.010379
5	0.997006	2.000979	4.986712	0.010057
10	0.892747	1.999248	8.636796	0.010789

Table 7: Estimates of steady-state amplitude and resonant behaviour for data simulated with the following Q -factors: (a) 3, (b) 5, and (c) 10.

7.6 SUMMARY

We infer the following conclusions from the results of the simulations given in Sections 7.1 to 7.5.

- a) As the time-window is reduced, the error in the values of the estimated parameters generally increases, even though the quality of the fit (as measured by σ) remains essentially the same.
- b) The accuracy of the estimated parameters and the quality of the fit (as measured by σ) are generally insensitive to the sampling rate, noise level and drive frequency.
- c) As the Q -factor of the device is increased, the error in the values of the estimated steady-state amplitude and Q -factor generally increase, even though the quality of the fit (as measured by σ) remains essentially the same. However, the accuracy of the estimated resonance frequency is generally insensitive to the choice of Q -factor.

It should be noted that in the above we have made statements about the performance of the algorithm in terms only of the *error* of the computed estimates. Note that where this error is large, we expect the standard uncertainty in the estimate to be large as well. The evaluation of uncertainties is discussed in Section 10.1.1. Important measures of an estimator's performance include its *bias* and *efficiency* expressed in terms of the mean value and the standard uncertainty of the estimated parameters. These measures may be evaluated by studying the distributions of the estimated parameters obtained by repeating many times the experiments described here with different samples of measurement error used to generate the simulated data.

8 RESULTS FOR EXPERIMENTAL DATA

In this section we present results obtained from applying the data processing algorithm described in Section 6 to data obtained from measurements made of a real device. An ITC1001 transducer, with a resonance at approximately 18 kHz, was driven with discrete-frequency tone-burst signals in the range 5 to 30 kHz in steps of 1 kHz. The acoustic signal was detected using a calibrated Reson TC4034, with a resonance at approximately 350 kHz, and the waveform was acquired using a signal analyser (12-bit ADC sampling at 10 MHz). The calibrated hydrophone and measuring equipment were chosen such that the system had no pole close to the frequency range of interest.

The transmitting voltage response for the projector at a given frequency is expressed as

$$S_P = \frac{V_H d}{V_P M_H}, \quad (25)$$

where V_H is the voltage output of the receiving hydrophone, V_P is the voltage driving the projector, d is the distance between the projector and hydrophone, and M_H is the sensitivity of the receiving hydrophone. Here, V_P and d are directly measured, and M_H is known from the calibration of the receiving hydrophone. V_H is set equal to the value of A_0 obtained from an analysis of the signal $y(t)$ measured by the receiving hydrophone.

We present in Figure 13 estimates of the transmitting voltage response for the ITC1001 projector calculated using (25) where, at each frequency, the voltage output V_H of the receiving hydrophone is estimated using the data processing procedure described in Section 6. The transmission voltage response obtained by analysing a sufficiently large time-window that includes, for each frequency, the steady-state response is shown using “small circles”. The transmission voltage responses obtained by analysing (a) data contained within the time-window $[t_0, t_2]$ comprising $m = 112$ points (between three and four cycles of resonance), and (b) data contained within the time-window $[t_0, t_1]$ comprising $m = 56$ points (between one and two cycles of resonance), are shown as, respectively, a solid line and a dashed line.




Figure 13: On the left, transmitting voltage responses from (i) free field measurements (small circles), (ii) measurements from a time-window comprising $m = 112$ points (solid line), and (iii) measurements from a time-window comprising $m = 56$ points (dashed line). On the right, the differences between (i) and (ii) (solid line), and between (i) and (iii) (dashed line).

In each case and for each frequency, the decision to accept an estimate of the steady-state response was made as objective as possible. The decision was based on how well the data within the time-window considered was fitted. A fitted model was only accepted if it satisfied the convergence criteria implemented for the nonlinear estimation problem and, wherever possible, a model was chosen to reflect the belief that only one resonance is present in the system.

9 MODELS OF THE TRANSMITTING VOLTAGE RESPONSE

In Section 2 we described models for representing the measured signal $y(t)$ that is recorded by the receiving hydrophone when the projector under test is driven by a sinusoidal drive voltage. The models developed rely on the assumption that the signal $y(t)$ corresponds to the output of a system that behaves as a linear damped harmonic oscillator. The parameter of interest in those models is the steady-state amplitude A_0 , because we can use it to determine the transmitting voltage response for the device at the given drive frequency as described in Section 8.

In this section we present a model for the transmitting voltage response as a function of the drive frequency f . This model involves the parameters defining the *resonant* behaviour of the projector rather than its *steady-state* behaviour. We indicate how the transmitting voltage response function may be derived from an analysis of a signal $y(t)$ measured at a single drive frequency.

In Appendix B we describe models describing the behaviour of linear damped harmonic oscillators, including (a) free, mechanical oscillators, (b) forced, mechanical oscillators, and (c) equivalent electrical oscillators (or LCR-circuits). In particular, for each type of oscillator, we present expressions for the steady state amplitude in terms of parameters defined by the resonant behaviour of the oscillator. Equations (37) and (38) give these expressions for the case of an LCR-circuit.

The definition of the transmitting voltage response S_p for a projector producing a spherically spreading acoustic field at a given frequency is

$$S_p = \frac{Pd}{V_p},$$

where P is the far-field acoustic pressure at a distance d from the projector when it is driven by a voltage V_p . We suppose that the device may be modelled using an equivalent electrical circuit composed of a blocked capacitance in parallel with a resonant LCR-circuit where the LCR elements represent the motional impedance of the device. The blocked capacitance is used to represent the electrical behaviour of the device at low frequencies, whereas the LCR-circuit represents its acoustic properties when driven near to its resonance. Furthermore, we assume that the pressure waveform generated by the device is a scaled replica of the current flowing through the LCR-circuit. We have chosen to use this particular equivalent circuit because if the current and pressure are directly coupled as described, the choice is consistent with the models we have used to represent the measured signals $y(t)$. We comment further on this aspect in Section 10.

If the current and pressure are coupled as above, we have

$$S_p = Kd \frac{I_0}{V_p}, \quad (26)$$

where K is an unknown scale factor and I_0 is the current. Now, given values for the resonance damping factor d and resonance frequency f for the device, (37) defines the relative sizes of the circuit elements L , $1/C$ and R , and (38) defines the amplitude response of the device. Substituting into (26), we obtain an expression for S_p as a function of the drive frequency f_0 , *viz.*,

$$S_p = \frac{Kd}{L} \frac{\omega_0}{\sqrt{((\omega^*)^2 - \omega_0^2)^2 + \omega_0^2(R/L)^2}}, \quad \omega^* = \sqrt{\frac{1}{LC}}, \quad (27)$$

where K and L are unknown. The function is completely defined by d and f (which determine R/L and $1/LC$ using (37)), and the value of S_p at a single frequency (which determines the scale factor Kd/L). Furthermore, these three pieces of information may be obtained from an analysis of the measured signal $y(t)$ at a single frequency.

We present in Figures 14 and 15 estimates of the transmitting voltage response for the ITC1001 projector calculated using (27) and estimates of the resonant behaviour of the device. For comparison, we include the (free-field) transmitting voltage response defined by the “small circles” in Figure 13. Four curves are shown in each figure corresponding to estimates of the resonant behaviour obtained by processing the data corresponding to the following drive frequencies: 10 kHz, 18 kHz, 20 kHz and 30 kHz. For Figure 14, the data analysed was contained within the time-window $[t_0, t_2]$ comprising $m = 112$ points (between three and four cycles of resonance). We show the transmitting voltage response expressed in (linear) units of Pa m/V (on the left) as well as in units of dB re 1 μ Pa m/V (on the right). For Figure 15, the data analysed was contained within the time-window $[t_0, t_1]$ comprising $m = 56$ points (between one and two cycles of resonance).

The results shown in Figure 13 indicate that good estimates of the transmitting voltage response may be obtained from an analysis of data measured in a time-window whose duration is a small number of cycles of the resonance. The value of the transmitting voltage response at a given frequency is derived from the estimate of the steady-state amplitude obtained from an analysis of the data measured at that frequency. The curves shown in the figure are a piecewise linear approximation to the transmitting voltage response generated by joining the individual values by straight-line segments.

In contrast, the curves shown in Figures 14 and 15 are obtained by evaluating a *model* for the transmitting voltage response defined by parameters that depend on estimates of the resonant behaviour. Note that each curve in these figures is generated *completely* from the analysis of data corresponding to a *single* drive frequency. However, the estimates of the transmitting voltage response generated in this way are less satisfactory, particular for low frequencies. The reasons for this effect may be that (a) the model of the transmitting voltage response is deficient (because the equivalent electrical circuit from which the model is derived is too simple), and (b) the estimates of the resonant behaviour are not sufficiently accurate. Certainly, the estimated model obtained from measurements made at the lowest frequency of 10 kHz are the least accurate, and this is where the data is hardest to model. Furthermore, we know from our discussion of the conditioning of the estimation problem (Section 3) that obtaining estimates of the resonant behaviour is more difficult than estimating the steady-state amplitude.




Figure 14: Transmission voltage responses obtained from estimates of the resonant behaviour of the device. Each curve is derived from the analysis of data corresponding to a single drive frequency. The data analysed was contained within the time-window $[t_0, t_2]$ comprising $m = 112$ points.




Figure 15: Transmission voltage responses obtained from estimates of the resonant behaviour of the device. Each curve is derived from the analysis of data corresponding to a single drive frequency. The data analysed was contained within the time-window $[t_0, t_1]$ comprising $m = 56$ points.

10 SCOPE FOR FURTHER WORK

A number of topics for further work are now identified.

10.1 FURTHER INVESTIGATION OF SOURCES OF ERROR

10.1.1 Evaluation of uncertainties

In this work we have concentrated on presenting methods for estimating the *value* of various key parameters required in the calibration of an underwater transducer. For such estimates to be useful, it is necessary also to quantify their statistical uncertainty. This quantification can be done as part of the nonlinear estimation algorithm: the matrix

$$s^2(J^T J)^{-1},$$

where s is the root-mean-square error given by (14) and J is the Jacobian matrix evaluated at the solution (Section 5 and Appendix A), provides an estimate of the covariance matrix for the solution parameters, and contains the variances of the steady-state amplitude and the parameters defining the resonant behaviour.

10.1.2 Instrumentation

The influence of the instrumentation on the accuracy of the results requires further investigation. For example, it is believed that the gating unit used may have introduced spurious “glitches” at the start of the tone-bursts for some of the measurements, which might lead to increased transient behaviour at the waveform start.

10.1.3 Devices to be calibrated

Particular problems exist when the methods described here are applied to high- Q transducers since insufficient information about the transducer resonances is present in the data. Difficulties can also be found in discriminating accurately between two closely spaced frequencies, as might occur when a high- Q transducer is driven close to resonance. There is a requirement to apply the techniques presented in this report to a wide variety of devices, including high- Q transducers, driven at various frequencies.

10.2 EXTENSIONS TO THE METHOD

10.2.1 Combining data obtained at a range of frequencies

In this work we have considered modelling data obtained by measuring the response at a single drive frequency. However, we would expect benefits from modelling simultaneously a number of sets of measurement data corresponding to different drive frequencies. Such modelling would require the model to be extended to include:

- a) terms representing the resonant behaviour of the device that would be common to all the measurement sets, and
- b) terms representing the steady-state behaviour that would relate to particular measurement sets.

10.2.2 Inclusion of extra terms to represent echoes

The modelled data has been restricted to lie within a time-window containing the free-time response. This time-window is delimited by the turn-on of the device and the arrival time of the first echo. However, we would expect benefits for high- Q transducers from extending this time-window into the region containing the first few reflections. Again, this extension would

require the model to be augmented to include *delayed* terms used to model the resonant behaviour of the device arising from the arrival of subsequent reflections (and possibly its turn-off). This approach has been considered in [7].

10.2.3 Use of other novel techniques

Techniques such as wavelet analysis may have the potential to improve the methods presented here by enabling better de-noising of the signals and accurate estimation of the arrival time of reflections. In addition, there would be benefit from comparison with other techniques which may be used to overcome the problems created by the presence of boundary reflections. Such techniques include the use of transient suppression techniques [14-16], echo elimination by cepstral analysis, cross-correlation techniques and the method of time-delayed spectrometry [17].

10.3 INCORPORATION AND USE WITHIN CALIBRATION SOFTWARE

Currently, the software to execute the method requires user input to window the waveform data, and estimate the number of poles, etc. Some of these functions could be further automated and there is scope for the analysis software to be better integrated with the software which controls the instrumentation and acquires the waveforms.

The analysis software has direct application in the calibration of transducers in measurement tanks wherever the restricted echo-free time limits the lowest frequency at which measurements can be made. Such conditions also occur in the measurement of the echo reduction and transmission loss of panels made from anechoic materials.

10.4 TRANSDUCER MODELLING

Further work is required to improve the modelling of the transducer in order to estimate the response (transfer function) at frequencies other than the excitation frequency. Only a very crude model for the transducer has been assumed (in electrical terms a simple LCR circuit with a blocked capacitor in parallel), and improved results may be obtained with a fuller model that explicitly includes circuit elements representing the radiation impedance.

In addition, further work is needed to determine if the modelling can be applied to data obtained from transducers excited with other than sinusoidal drive waveforms. We have chosen not to consider, for example, driving the device with an impulse to measure its impulse response directly because of practical problems in generating such a drive signal. The use of an "imperfect" impulse is expected to give a response containing many frequency components not suitable for modelling in the way described in this work.

11 CONCLUSION

By use of suitable signal models, it has been shown that it is possible to predict the steady-state amplitude of stepped-sinusoid signals from the initial transient-dominated part of a measured waveform. The approach considered has been to model the free-time response of the device by a function consisting of a sum of complex exponential terms which are used to describe both the steady-state and resonant behaviour of the device. Two classes of estimation method have been reviewed: linear prediction methods (such as Prony's method), and nonlinear least-squares methods. Linear prediction methods have been shown to be statistically biased and inefficient. This observation motivated the development of a linear prediction method in which proper account is taken of the error structure in the data. A nonlinear least-squares algorithm has also been described based on a safeguarded Gauss-Newton algorithm that uses regularisation to address the ill-conditioning that is a property of the underlying problem.

Results have been presented of using these methods to analyse simulated data generated to represent the measured response of a device. The effect of varying in a systematic manner the properties of the simulated data and device under test has been investigated. These investigations suggest that useful results can be obtained from the analysis of a waveform that has not reached steady-state. Furthermore, the principal factors that affect the quality of the results are the length of the time-window containing the free-time response and the Q -factor of the device.

Results have also been presented for data obtained from measurements of a real transducer, and it is shown how these results may be used to establish the free-field sensitivity of the device. These results indicate that for moderately resonant transducers, errors of less than 1 dB will result if data defining approximately one cycle of the resonance frequency is available.

Finally, results are presented of using the information gained from the modelling of the transducer resonant behaviour to predict the transducer frequency response at frequencies other than those under test.

12 ACKNOWLEDGEMENTS

This report constitutes the deliverable of project 2.1d(ii) of the 1995-98 NMS Acoustics Metrology Programme and has been funded by the National Measurement System Policy Unit of the UK Department of Trade and Industry.

The authors would like to thank the staff of the Underwater Sound Reference Detachment for informative discussions during this work, Geraint Green (NPL) for undertaking some of the measurements, and Maurice Cox (NPL) for his advice and comments on this report.

13 REFERENCES

- [1] BEATTY, L.G., GEORGE, J.D and ROBINSON, A.Z. Use of the complex exponential expansion as a signal representation for underwater acoustic calibration. *J. Acoust. Soc. Am.*, 1978, **63**, 1782-1794.
- [2] MARPLE, S.L. *Digital Spectral Analysis with Applications*. Prentice Hall, Englewood Cliffs, New Jersey. 1987.
- [3] TRIVETT, D.H. and ROBINSON, A.Z. Modified Prony method approach to echo-reduction measurements. *J. Acoust. Soc. Am.*, 1981, **70**, 1166-1175.
- [4] TUFTS, D.W. and KUMARESAN, R. Estimation of frequencies of Multiple Sinusoids: Making linear prediction perform like maximum likelihood. *Proceedings of the IEEE*, 1982, **70**, 975-989.
- [5] OSBORNE, M.R. and SMYTH, G.K. A modified Prony algorithm for exponential fitting. *SIAM J. Sci. Comput.*, 1995, **16**, 119-138.
- [6] GEORGE, J.D., JAIN, V.K. and AINSLEIGH, P.L. Estimating steady-state response of a resonant transducer in a reverberant underwater environment. *IEEE International Conference on Acoustics, Speech and Signal Processing*, 1992, **5**, 1988, 2737-2740.
- [7] AINSLEIGH, P.L. and GEORGE, J.D. Signal modeling in reverberant environments with application to underwater electroacoustic transducer calibration. *J. Acoust. Soc. Am.* 1995, **98**, 270-279.
- [8] GEORGE, D.J. and GOODMAN, D.M. Application of the output error system identification method to the calibration of underwater acoustic transducers. *Underwater Acoustic Data Processing*, Kluwer Academic Publishers, 1989, 39-46.
- [9] GILL, P.E., MURRAY, W. and WRIGHT, M.H. *Practical Optimization*. Academic Press, 1981.
- [10] AINSLEIGH, P.L. and GEORGE, J.D. Variable projection methods with application to sums of exponentials in white noise. *NRL Memorandum Report 6643*, 1990.
- [11] AINSLEIGH, P.L. and GEORGE, J.D. Full Newton and constraint methods for semilinear signal modeling. *IEEE Transactions on Signal Processing*, 1993, **41**, 1689-1692.
- [12] GEORGE, J.D. and AINSLEIGH, P.L. A constrained Newton algorithm for maximum likelihood modeling of sums of exponentials in noise. *IEEE International Conference on Acoustics, Speech and Signal Processing*, 1990, **5**, 2443-2446.
- [13] GOLUB, G.H. and PEREYRA, V. The differentiation of pseudoinverses and nonlinear least squares problems whose variables separate. *SIAM J. Numer. Anal.*, 1973, **10**.
- [14] PIQUETTE, J.C. Method for transducer transient suppression, I: Theory. *J. Acoust. Soc. Am.* **92**(3), 1992.

- [15] PIQUETTE, J.C. Method for transducer transient suppression, II: Experiment. *J. Acoust. Soc. Am.* **92**(3), 1992.
- [16] PIQUETTE, J.C. Applications of the method for transducer transient suppression to various transducer types. *J. Acoust. Soc. Am.* **94**(2), 1992.
- [17] GIANGRECO, C. *Mesures acoustiques appliquées aux antennes sonar*. Lavoisier, 1997.
- [18] FLETCHER, R. A modified Marquardt subroutine for nonlinear least squares. *Report R6799, Atomic Energy Research Establishment, Harwell, England*, 1971.

APPENDIX A:NONLINEAR LEAST-SQUARES ESTIMATION

A.1 THE GAUSS-NEWTON ALGORITHM

We give here a brief description of the Gauss-Newton algorithm [9] for solving a nonlinear least-squares problem.

Suppose we wish to minimise the (general) function $E(\mathbf{a})$ with respect to the parameters \mathbf{a} . Let \mathbf{a}_k be an estimate of the solution \mathbf{a}^+ satisfying

$$\mathbf{a}^+ = \mathbf{a}_k + \mathbf{p}.$$

A quadratic approximation to the objective function E can be obtained by taking the first three terms of the Taylor-series expansion about \mathbf{a}_k , i.e.,

$$E(\mathbf{a}_k + \mathbf{p}) = E(\mathbf{a}_k) + \mathbf{g}_k^T \mathbf{p} + \frac{1}{2} \mathbf{p}^T G_k \mathbf{p},$$

where \mathbf{g}_k is the gradient of E evaluated at \mathbf{a}_k (i.e., the vector of first derivatives), and G_k is the Hessian matrix evaluated at \mathbf{a}_k (i.e., the matrix of second derivatives). Now, a stationary point for this quadratic approximation satisfies the linear system

$$G_k \mathbf{p} = -\mathbf{g}_k. \quad (28)$$

A minimisation algorithm in which \mathbf{p} is defined by (28) is called Newton's method, and the solution to (28) is called Newton's direction. A new estimate of the solution \mathbf{a}^+ is formed from

$$\mathbf{a}_{k+1} = \mathbf{a}_k + \mathbf{p}, \quad (29)$$

and the process is repeated using \mathbf{a}_{k+1} .

Now suppose that E takes the form of a sum of squares of nonlinear functions, i.e.,

$$E(\mathbf{a}) = \sum_{i=1}^m e_i^2(\mathbf{a}).$$

Let $J(\mathbf{a})$ be the Jacobian matrix for $\mathbf{e}(\mathbf{a}) = (e_1(\mathbf{a}), \dots, e_m(\mathbf{a}))^T$, i.e., the i th row of J contains the gradient vector for $e_i(\mathbf{a})$, and let $G_i(\mathbf{a})$ denote the Hessian matrix for $e_i(\mathbf{a})$. Then, the gradient vector and Hessian matrix for E are, respectively,

$$\mathbf{g}(\mathbf{a}) = 2J(\mathbf{a})^T \mathbf{e}(\mathbf{a}),$$

and

$$G(\mathbf{a}) = 2J(\mathbf{a})^T J(\mathbf{a}) + 2Q(\mathbf{a}), \quad Q(\mathbf{a}) = \sum_{i=1}^m e_i(\mathbf{a}) G_i(\mathbf{a}).$$

Consequently, the Newton direction \mathbf{p} is given by

$$(J_k^T J_k + Q_k) \mathbf{p} = -J_k^T \mathbf{e}_k,$$

where we use the subscript k to denote evaluation at \mathbf{a}_k .

In the Gauss-Newton method we assume that the first order term $J_k^T J_k$ dominates the second order term Q_k , and the Newton direction is approximated by the solution to the linear system

$$J_k^T J_k \mathbf{p} = -J_k^T \mathbf{e}_k. \quad (30)$$

This assumption is valid provided the residuals e_i corresponding to the estimate \mathbf{a}_k are small. Note that the solution to (30) is a solution of the linear least-squares problem

$$J_k \mathbf{p} = -\mathbf{e}_k. \quad (31)$$

This solution defines the Gauss-Newton direction.

The Gauss-Newton algorithm is an iterative method that at each iteration takes a step towards the minimum by solving the linear least-squares problem (31) and updating the current estimate using (29). The iterations are repeated until the estimates are judged to have converged. Convergence may be detected by examining (a) the change in the value of the objective function E , (b) the size of the Gauss-Newton step, and (c) the size of the gradient of E . Convergence is not guaranteed: for example, if \mathbf{a}_k is a poor estimate of the solution, the iterates generated by the algorithm may diverge. Furthermore, convergence may be slow, as can happen when the residuals e_i are large or if the Jacobian matrix J is poorly conditioned. For this reason, the Gauss-Newton algorithm is modified to improve its performance in these situations. These modifications are described in the following sections.

A.2 LINE SEARCH

We replace the update step (29) by

$$\mathbf{a}_{k+1} = \mathbf{a}_k + \mu \mathbf{p},$$

where \mathbf{p} is the Gauss-Newton direction and μ controls the length of the step.

The parameter μ is chosen so that the step $\mu \mathbf{p}$ produces a ‘‘sufficient decrease’’ in the objective function E . Following [9], a sufficient decrease is achieved when μ satisfies the condition

$$E(\mathbf{a}_k + \mu \mathbf{p}) - E(\mathbf{a}_k) \leq \rho \mu \mathbf{g}_k^T \mathbf{p} < 0, \quad (32)$$

where we set $\rho = 0.05$. In other words, the step is accepted if the change in the objective function exceeds (in absolute value) some threshold, where this threshold depends on the gradient of the objective function at the new estimate. If this condition is not satisfied, a new value μ_{new} for μ is chosen as follows.

Firstly [9], define

$$\mu_1 = \max\{\mu_m, \rho \mu\},$$

where μ_m locates the minimum of a quadratic approximation to E along the direction \mathbf{p} . Then, set

$$\mu_{\text{new}} = \begin{cases} \mu_1, & \rho_1 \leq \mu_1 \leq 1, \\ \mu/2, & \text{otherwise.} \end{cases}$$

where we set $\rho_1 = 0.002$. The motivation for this scheme is that we try to find a step that produces a sufficient decrease in E (by accepting, if possible, μ_m), but we do not want the step to be too large or small. Starting with $\mu = 1$ (the default Gauss-Newton step), we repeat this procedure until either the condition (32) is satisfied or we have reached a maximum number of iterations.

An alternative to the Gauss-Newton algorithm with a line search is to use the Levenberg-Marquardt algorithm [9, 18]. The Levenberg-Marquardt search direction is defined as the solution to

$$(J_k^T J_k + \mu_k I) \mathbf{p} = -J_k^T \mathbf{e}_k,$$

where μ_k is a non-negative scalar. A unit step is always taken along \mathbf{p} , i.e., \mathbf{a}_{k+1} is given by (29), with μ_k chosen in order to ensure a descent. An algorithm that adjusts the value of μ_k according to the relationship between the actual and predicted sum of squares is proposed in [18].

A.3 REGULARISATION

We replace the nonlinear least-squares problem (15) by

$$\text{minimise } \sum_{i=1}^m \{y_i - y(t_i)\}^2 + u^2 \{f_r - f_1\}^2 + v^2 \{d_r - d_1\}^2 + \lambda^2 \sum_{k>1} d_k^2,$$

where λ is a regularisation parameter. The function to be minimised here involves

- a) the residuals $y_i - y(t_i)$ that measure the errors between the measured data and the fitted model,
- b) the residuals $f_r - f_1$ and $d_r - d_1$ that relate to any *a priori* information about a resonance of the device, and
- c) terms that relate to constraints of the form $d_k = 0$.

The idea behind the inclusion of the terms under c) is to define a problem whose solution is close to the solution of the problem (15) that we wish to solve but which is better conditioned than that problem.

As before, the “weights” u and v are chosen to reflect the relative accuracies between the measured data and the *a priori* information. The regularisation factor λ is chosen to be as large as possible such that the residual sum of squares for the regularised problem does not exceed that for the unregularised problem by more than a given factor, say 10%. Information about the residual sum of squares for the unregularised problem may be derived from knowledge concerning the likely measurement errors in the data, or from the fits to data measured at high drive frequencies for which regularisation is not necessary.

Our experience is that regularisation is required principally in cases where many damped sinusoidal terms are included in the model. The use of regularisation is avoided therefore by

choosing a model which includes only a small number of such terms. The inclusion of many terms of this type is usually necessary to ensure that, in the presence of measurement error, linear prediction methods give a good initial fit to the data. Regularisation is then avoided by removing explicitly these terms, leaving a small number of damped sinusoidal terms to represent the resonances of the device, before applying the Gauss-Newton algorithm.

APPENDIX B: LINEAR DAMPED HARMONIC OSCILLATORS

B.1 FREE, MECHANICAL OSCILLATION

The differential equation

$$m\ddot{y}(t) + c\dot{y}(t) + ky(t) = 0$$

models the free (mechanical) oscillation of a body, where m is the mass of the body, c is the damping constant, and k is the spring constant. Provided

$$c^2 < 4mk,$$

the roots of the characteristic equation are

$$\beta = -\frac{c}{2m} \pm i \frac{\sqrt{4mk - c^2}}{2m},$$

and the motion is the damped oscillation

$$y(t) = A_r e^{d_r t} \cos(2\pi f_r t + \phi_r),$$

where

$$d_r = -\frac{c}{2m}, \quad 2\pi f_r = \frac{\sqrt{4mk - c^2}}{2m}. \quad (33)$$

A_r and ϕ_r are defined by initial conditions comprising values for y and \dot{y} at $t = 0$.

B.2 FORCED, MECHANICAL OSCILLATION

The differential equation

$$m\ddot{y}(t) + c\dot{y}(t) + ky(t) = F_0 \cos(2\pi f_0 t) \quad (34)$$

models the forced (mechanical) oscillation of a body. The solution is now

$$y(t) = A_r e^{d_r t} \cos(2\pi f_r t + \phi_r) + A_0 \cos(2\pi f_0 t + \phi_0),$$

where

$$\frac{A_0}{F_0} = \frac{1}{\sqrt{m^2(\omega_r^2 - \omega_0^2)^2 + \omega_0^2 c^2}}, \quad (35)$$

and

$$\omega_0 = 2\pi f_0, \quad \omega_r = \sqrt{\frac{k}{m}}. \quad (36)$$

Given values for m , c and k , the expressions (35) and (36) define the amplitude response of the system as a function of the drive frequency f_0 . Notice that given values for d_r and f_r that define the resonant behaviour of the system, we are unable to determine absolute values for m , c and k . Instead, we can use (33) to determine their *relative* sizes as follows:

$$\frac{c}{m} = -2d_r, \quad \frac{k}{m} = \frac{1}{4}((2d_r)^2 + (4\pi f_r)^2).$$

Then, a scaled version of the amplitude response is defined by

$$m \frac{A_0}{F_0} = \frac{1}{\sqrt{(\omega_r^2 - \omega_0^2)^2 + \omega_0^2 (c/m)^2}}, \quad \omega_r = \sqrt{\frac{k}{m}},$$

where m is unknown.

B.3 FORCED, ELECTRICAL OSCILLATION

The electrical equivalent to the forced mechanical oscillation described above is an LCR-circuit whose current response to an applied potential satisfies

$$L\dot{I} + RI + \frac{1}{C} \int I dt = E_0 \sin(2\pi f_0 t).$$

Differentiating with respect to time, we obtain

$$L\ddot{I} + R\dot{I} + \frac{1}{C} I = 2\pi f_0 E_0 \cos(2\pi f_0 t).$$

If we set $m = L$, $c = R$, $k = 1/C$, $F_0 = 2\pi f_0 E_0$ and $y = I$, this is equivalent to (34). Therefore, we deduce that the amplitude response of the system is given by

$$\frac{I_0}{E_0} = \frac{\omega_0}{\sqrt{L^2 (\omega_r^2 - \omega_0^2)^2 + \omega_0^2 R^2}},$$

where

$$\omega_0 = 2\pi f_0, \quad \omega_r = \sqrt{\frac{1}{LC}}.$$

As for the forced mechanical system, if all we know is the resonant behaviour as defined by d_r and f_r , we are unable to determine the absolute values for the parameters L , C and R . Instead, we can use (33) to determine their *relative* sizes as follows:

$$\frac{R}{L} = -2d_r, \quad \frac{1}{LC} = \frac{1}{4}((2d_r)^2 + (4\pi f_r)^2), \quad (37)$$

and a scaled version of the amplitude response is defined by

$$L \frac{I_0}{E_0} = \frac{\omega_0}{\sqrt{(\omega_r^2 - \omega_0^2)^2 + \omega_0^2 (R/L)^2}}, \quad \omega_r = \sqrt{\frac{1}{LC}}, \quad (38)$$

where L is unknown.