



City Research Online

City, University of London Institutional Repository

Citation: Romare, Dario (1998). The Application of Adaptive Linear and Non-Linear Filters to Fringe Order Identification in White-Light Interferometry Systems. (Unpublished Doctoral thesis, City, University of London)

This is the accepted version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/19993/>

Link to published version:

Copyright: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

City Research Online:

<http://openaccess.city.ac.uk/>

publications@city.ac.uk

**The Application of Adaptive Linear and
Non-Linear Filters to Fringe Order Identification
in White-Light Interferometry Systems**

Dario ROMARE

A Thesis Submitted for the Degree of Doctor of Philosophy

CITY UNIVERSITY

Department of Electrical, Electronic and Information Engineering

October 1998

Contents

Acknowledgements	vii
Declaration	viii
Abstract	ix
1 White-Light Interferometric Systems	1
1.1 Introduction	1
1.2 Optical Interferometry	2
1.3 Interferometric Systems as Sensors	3
1.4 Temporal Coherence and Coherence Length	4
1.5 White-Light Interferometry	7
1.6 Fibre-Optic Interferometric Sensors	9
1.7 Digital Signal Processing	10
1.8 Aims and Objectives of the Thesis	13
1.9 Summary	14
2 Time Series Modelling of WLI Systems	18
2.1 Introduction	18
2.2 Autoregressive-Moving Average Models	18
2.3 Parametric Modelling of WLI Systems	20
2.4 Inverse Filtering and the Wiener Solution	21
2.5 Estimation of the Model Parameters	23
2.6 Least Squares Filtering and Linear Prediction	27
2.7 Conclusion	28
3 Adaptive Filtering of WLI Fringe Patterns	37
3.1 Introduction	37
3.2 Adaptive Finite Impulse Response Filters	37

3.2.1	The Standard Least Mean Square (LMS) Algorithm	40
3.2.2	The Recursive Least Squares (RLS) Algorithm	42
3.2.3	The Kalman Algorithm	45
3.3	Simulation Results	47
3.4	Ill-Conditioning and Finite Precision Effects	52
3.5	Conclusion	55
4	Convergence and Tracking Problems in WLI Filtering	60
4.1	Introduction	60
4.2	Identification Rate and Filter Parameters	60
4.3	Time-Evolution of MSE and Filter Weights	64
4.4	Time-Evolution of Filter Output	69
4.5	Convergence and Tracking Aspects	74
4.6	Choice of Filter Parameters	77
4.7	Discussion	80
5	A New WLI Central Fringe Identification Scheme	85
5.1	Introduction	85
5.2	Towards Faster Convergence	85
5.3	Threshold Pre-processing	88
5.4	Simulation Results	90
5.5	Enhanced LMS Algorithms	92
5.5.1	Fixed Step-Size LMS	92
5.5.2	Variable Step-Size LMS	95
5.6	A Modified Forward-Backward LMS	98
5.6.1	Simulation Results	99
5.6.2	Experimental Evaluation	103
5.7	Matched Filter Detection	105
5.8	Comparison of Methods	107
5.9	Discussion	117
6	Alternative LMS and RLS Schemes for WLI	126
6.1	Introduction	126
6.2	Algorithms for AR Modelling	126
6.3	Algorithms for ARMA Modelling	128

6.4	Modelling with Coloured Noise	129
7	Non-Gaussian and Non-Linear Modelling and Filtering	139
7.1	Introduction	139
7.2	Non-Gaussian Modelling	139
7.2.1	Gaussianity Tests	140
7.2.2	Linear Non-Gaussian Filtering	145
7.3	Non-Linear Modelling	148
7.3.1	Linearity Tests	149
7.3.2	Non-Linear Volterra Filtering	154
7.4	Discussion	161
8	Summary and Directions for Future Work	170
8.1	Summary	170
8.2	Directions for Future Work	172
A	Glossary of Terms	180
B	List of Publications	185
C	List of Software Tools	187

List of Tables

3.1	Condition numbers of data matrix used by batch algorithms	54
3.2	Condition numbers of data matrix used by adaptive algorithms . . .	54
5.1	Computer time for sub-fringe identification	113
7.1	Success rate in white Gaussian noise	145
7.2	Success rate in broad-band Gaussian noise	145
7.3	Success rate in narrow-band Gaussian noise	145
7.4	Success rate with a $ \sin(x)/x $ profile	147

List of Figures

1.1	Output intensity with a He-Ne laser source	5
1.2	Typical spectrum of an AlGaAs LED	6
1.3	Output intensity with a LED source	6
1.4	Schematic diagram of a WLI system	8
1.5	Simulated output intensity of a WLI system with noise at 20 dB . .	11
1.6	Success rate from direct fringe visibility	12
2.1	ARMA process driven by white noise	19
2.2	Spectrum of the fringe pattern in Fig. 1.5	25
2.3	ACS and PACS of the fringe pattern in Fig. 1.5	26
3.1	FIR adaptive transversal filter	40
3.2	Flowchart for on-line prediction with the LMS algorithm	41
3.3	Flowchart for on-line prediction with the RLS algorithm	44
3.4	Flowchart for on-line prediction with the Kalman filter	46
3.5	Success rate with batch and adaptive algorithms	48
4.1	Success rate with the LMS algorithm against SNR	61
4.2	Maximum success rate against filter order	61
4.3	Success rate against μ and SNR	62
4.4	Success rate against central fringe position	62
4.5	LMS MSE across the CCD array with a SNR of 32 dB	64
4.6	LMS MSE with a SNR of 20 dB	64
4.7	LMS MSE with a SNR of 10 dB	65
4.8	LMS weights with a SNR of 32 dB	66
4.9	LMS weights with a SNR of 20 dB	67
4.10	LMS weights with a SNR of 10 dB	67
4.11	RLS MSE with a SNR of 20 dB	68

4.12	RLS weights with a SNR of 20 dB	68
4.13	Experimentally obtained white-light fringe pattern	70
4.14	Fringe pattern filtered by the LMS ($\mu = 0.1$)	71
4.15	Fringe pattern filtered by the LMS ($\mu = 0.01$)	72
4.16	Fringe pattern filtered by the RLS ($\lambda = 1.0$)	72
4.17	Fringe pattern filtered by the RLS ($\lambda = 0.8$)	73
4.18	Success rate against coherence length	79
5.1	Success rate with the thresholded technique	90
5.2	Success rate against central fringe position	91
5.3	Position of AR(2) model poles against SNR	93
5.4	Diagram showing difference between FB-LMS and MFB-LMS	98
5.5	Success rate with the thresholded MFB-LMS	100
5.6	MSE with the thresholded MFB-LMS ($\mu = 10^{-3}$)	102
5.7	MSE with the thresholded MFB-LMS ($\mu = 10^{-6}$)	102
5.8	Success rate with the centroid method against central fringe position	109
5.9	Success rate with the centroid and covariance methods	110
5.10	Success rate with the LMS predictors over 992 pixels	110
5.11	Success rate with the LMS predictors over 960 pixels	111
5.12	Success rate with the RLS predictor	112
5.13	Success rate with the matched filter	112
7.1	Density of the central 26 fringes of Fig. 4.13	141
7.2	Two common ADC non-linearities	149
7.3	Fringe pattern distorted by a quadratic non-linearity	150
7.4	Fringe pattern distorted by a cubic non-linearity	150
7.5	Fringe pattern modified by chromatic aberrations	153
7.6	AR model signal generator and FIR adaptive non-linear filtering . .	155
7.7	Performance of linear and Volterra filters	157
7.8	Performance of linear and Volterra filters	158

Acknowledgements

I wish to thank Dr. M. Sabry-Rizk and Prof. K. T. V. Grattan for their excellent supervision and encouragement throughout the whole PhD research, and Dr. A. W. Palmer and Dr. Y. N. Ning for the many valuable suggestions.

I am also indebted to Dr. K. Weir, now at Imperial College, University of London, for the patient help and technical assistance provided during the acquisition of experimental data.

Finally, but not least, the three-year financial assistance of the Science and Engineering Research Council (now Engineering and Physical Sciences Research Council), was a determinant factor for the success of this work, as was the emotional support of my girlfriend, Encarna Sanchez.

Declaration

I grant powers of discretion to the University Librarian to allow this thesis to be copied in whole or in part without further reference to me. This permission covers only single copies made for study purposes, subject to normal conditions of acknowledgement.

Abstract

Conventional optical interferometry systems driven by highly coherent light sources have a very short unambiguous operating range, a direct consequence of the flatness of the interference fringes visibility profile at the output of the system.

The range can be extended by using a white-light interferometer (WLI), which is driven by a low-coherence source and produces a Gaussian visibility profile with a unique maximum in correspondence of the central fringe.

Due to system and/or measurement noise, however, the position of the maximum (from which an accurate measurement of the measurand - displacement, temperature, pressure, flow, etc. - can be derived) is not easily detectable, and can lead to large measurement errors. This is especially true in a multiplexing scheme, where the source power is distributed evenly among various sensors, with a corresponding drop in the overall signal-to-noise ratio. The inclusion of a signal processing scheme at the receiver end is thus a necessity.

As the fringe pattern at the output of a WLI system is basically a noisy sine wave amplitude modulated by a Gaussian envelope, it can be classified as a non-stationary, narrow-band, linear but non-Gaussian signal. So far, no attempt has been made to apply digital filtering techniques, as understood in the signal processing community, to the output signal of a WLI system. This thesis constitutes a first step in that direction.

Since the only measurable information given by the system is contained in the output signal, the system is modelled as a "black box" driven by the system and measurement noise processes and containing an unknown set of parameters. Standard least squares techniques can then be applied to estimate the parameters of the model, as is usually done in the field of system identification when only noisy output measurements are available.

It is shown that identification of the model parameters is equivalent to finding a set of coefficients for an inverse filter which takes the WLI signal at its input and delivers the unknown noise process at the output.

The non-stationarity of the signal is accounted for by allowing for time variations of the model parameters; this justifies the use of adaptive filters with time-varying coefficients. A new central fringe identification scheme is proposed, based on a modification of the standard least mean square (LMS) adaptive filtering algorithm in combination with amplitude thresholding of the fringe pattern. The new scheme is shown to offer considerable improvement in the identification rate when tested against current schemes over comparable operating ranges, while retaining the computational simplicity and operational speed of the standard LMS. Its performance is also shown to be largely independent of the step-size parameter controlling the rate of convergence and tracking in the standard LMS, which is known to be the main obstacle for a successful application of the algorithm in a practical setting.

The non-Gaussianity of the signal is explored and an attempt is made to apply higher-order statistics (HOS) algorithms to central fringe identification. The effectiveness of Gaussianity tests on pilot Gaussian data is seen to depend not only on the number and length of records available but, perhaps more importantly, on the bandwidth of the process. Violation of the stationarity assumption is shown to lead to mis-classification of a seemingly non-Gaussian signal into a Gaussian one, as the visibility profile may alter the distribution of the underlying sinusoid making it appear Gaussian, even when beam diffraction and wavefront aberrations combine to produce a non-Gaussian profile. HOS-based adaptive algorithms may still be of some benefit, however, if processing is confined to that region of the fringe pattern where sufficient non-Gaussianity is allowed to develop.

Non-linear adaptive filters based on the Volterra theories are finally applied to compensate for possible non-linearities introduced by mismatches in optical components, chromatic aberrations, and analogue-to-digital converters. It is shown that although a Volterra filter is able to reproduce the low-amplitude distortions of the fringe pattern better than a linear filter does, the identification rate does not improve. Reasons are given for such behaviour.

Chapter 1

White-Light Interferometric Systems

1.1 Introduction

This chapter explains the two main properties of light which are exploited in optical measurement systems, i.e., interference and coherence. The advantages of white-light over monochromatic sources which are responsible for the recent interest in all-optical-fibre white-light interferometric systems are also described, before examining the central fringe identification problem and the techniques that have been proposed to ease it.

In Chapter two it is shown how a physical system, in this case an optical interferometer, can be approximated by a statistical model consisting of a "black-box" driven by an unknown random process and containing an arbitrary number of parameters. Identification of the central fringe is then formulated as an inverse filtering problem, and known off-line or batch schemes are presented to solve for the filter coefficients.

Adaptive filters, which give rise to on-line schemes, are introduced in Chapter three. A preliminary evaluation of off and on-line filtering algorithms for central fringe identification is presented, using simulated data. An appreciation of ill-conditioning and finite-precision effects caused by round-off errors is also included.

Chapter four discusses the effects of filter order and properties of the measured output fringe pattern (such as degree of non-stationarity and additive noise level) on the convergence and tracking performance of two of the most widely used adaptive

filtering algorithms, the standard least mean squares (LMS) and the recursive least squares (RLS).

A novel scheme using a modified version of the LMS is presented in Chapter five, which offers a much higher identification rate than that possible with the standard version, at no extra computational cost. The novel scheme is shown to approach the performance limit imposed by the matched filter, which is the ideal solution for the detection of a known signal in additive white noise, and makes the choice of the step-size parameter which controls the convergence and tracking rates in the LMS practically redundant. A comparison between the most commonly used methods for white-light central fringe identification and some of the schemes presented here ends the chapter, with greater emphasis being put on performance and computational complexity.

Chapter six contains a round-up of adaptive filtering algorithms that have not been considered in this thesis but may nevertheless be capable of improving the identification rate offered by the novel LMS scheme, on condition that they are modified along the lines of the novel LMS.

Linear filters based on higher-order statistics are examined in the first part of Chapter seven. Non-linear filters based on the Volterra theories follow, to account for non-linearities introduced by the optical system and by data-acquisition or recording instruments. It is shown that there is at present no reason to prefer a non-linear to a linear filter in current white-light interferometry measuring systems, and the advantage of using higher-order statistics for linear filtering is also doubtful.

Chapter eight is a summary of what has been achieved in the field of white-light interferometry by this thesis, and suggests possible extensions that may be worth of further study.

1.2 Optical Interferometry

Optical interference is the phenomenon which can be observed when coherent light from a source is divided into two beams which are then superposed. In the region of superposition, the resultant intensity at different points varies between maxima which exceed the sum of the component intensities, and minima which may be zero. If the light source is divided by passage through apertures placed side by side we

have division of wave-front; if, instead, the beam is split at one or more surfaces, at which part of the light is transmitted and part reflected, we have division of amplitude. The first method is only useful with small aperture sources, hence the second method is in general preferable since it gives greater interference effects [1].

Irrespective of the method used, optical interferometry exploits the fact that, for an ideal monochromatic source with wavelength in air λ_0 , the phase difference between the two beams in the region of superposition is

$$\theta = k_0 \delta L \quad (1.1)$$

where $k_0 = 2\pi/\lambda_0$ is the wave number and δL is the difference between the optical paths through which the two beams have travelled before being recombined ¹.

If I is the intensity of the light source, the resultant intensity after recombination can be expressed as [1]

$$I_{res} = \frac{I}{2}(1 + \cos \theta) \quad (1.2)$$

Hence, when δL is an integral multiple of the source wavelength λ_0 , θ is an integral multiple of 2π and I_{res} reaches its maximum value; conversely, when δL is an odd multiple of half the wavelength, I_{res} goes through minima equal to zero.

The variation of the output intensity I_{res} with θ is often referred to as a *fringe pattern* because when observed visually it appears as a succession of evenly spaced white and dark fringes.

1.3 Interferometric Systems as Sensors

Optical interferometers come in various forms and shapes, but they all share the same principle.

Those based on amplitude division usually consist of an extended source which is divided into two beams of equal intensities at a beam splitter. The two beams are recombined after reflection at two plane mirrors, and sent to a detector which responds to I_{res} . If both mirrors are fixed, δL is constant and so is I_{res} , but if one of the mirrors is allowed to move, δL varies and I_{res} with it.

The Michelson interferometer is most often used for the accurate measurement of displacements. The movable mirror is attached to the measurand, and if a displace-

¹The wave number is defined as $1/\lambda_0$ in [2], and is the number of waves/cm path in vacuum, ranging from 15000 to 25000 cm^{-1} in going from red to violet.

ment of the fringe pattern by m orders has occurred, the movement has introduced an optical path difference (OPD) equal to $m\lambda_0$ in air, or $m\lambda_0/n$ in a medium such as an optical fibre with refractive index n . The displacement of the measurand, d , is given directly as

$$d = \frac{1}{2} \frac{m\lambda_0}{n} \quad (1.3)$$

Displacements of up to one-fiftieth of a fringe can be detected, making it possible to perform measurements with an accuracy of one-hundredth wavelength, corresponding to 5 nm for green light [2].

The Jamin, Mach-Zehnder, and Rayleigh interferometers may be used to measure variations of density in gas flows, exploiting the dependence between gas density and refractive index. If t is the thickness of the gas flow traversed by the beam in one arm of the interferometer and the refractive index of air is taken to be unity, $(n - 1)t/\lambda_0$ extra waves are introduced by the passage of the gas. Hence, if a displacement of the fringe pattern by m orders has occurred, n is obtained from

$$(n - 1)t = m\lambda_0 \quad (1.4)$$

from which the gas density can be derived.

The main advantage of optical sensors over electrical sensors is their immunity to electromagnetic interference. Their use is increasing all the time and includes such diverse fields as air temperature monitoring, torque measurements, and inhomogeneity observations in glasses. A detailed account is beyond the scope of this thesis and is widely available in the literature (see, e.g. [3]).

1.4 Temporal Coherence and Coherence Length

In a practical setting, the theory of interference effects as described by Eq. 1.2 is still valid, but one has to allow for some variability, due to the impossibility of realising a perfectly monochromatic light source with a single wavelength λ_0 . When more than one wavelength is present, at zero δL all the component wavelengths from one arm of the interferometer interfere constructively with the corresponding wavelengths from the other arm, resulting in maximum output intensity. As the OPD increases, however, the different wavelengths produce fringes of slightly different spacing, so that the fringe visibility is gradually reduced and finally disappears when they no longer overlap.

Another way of explaining this phenomenon is to think of the source as consisting of wave trains of finite length. If, after division at the glass plate the OPD between the two halves of each wave train is greater than this length, there is no interference because the two halves being combined are no longer derived from the same wave train, and have lost any correlation to each other [2].

The need for partially monochromatic light is called the *temporal coherence* requirement for interference [4], and the maximum δL over which interference effects can be observed is called *coherence length*. It turns out that the sharper the line width of the source, $\delta\lambda$, the more monochromatic the light, and the greater the coherence length L_c , according to the equation

$$L_c = \frac{\lambda_0^2}{\delta\lambda} \quad (1.5)$$

where λ_0 is the mean wavelength [1].

What is more, the fringe profile is the Fourier transform of the spectrum of the light source [5], and this gives a quick tool for predicting interference effects.

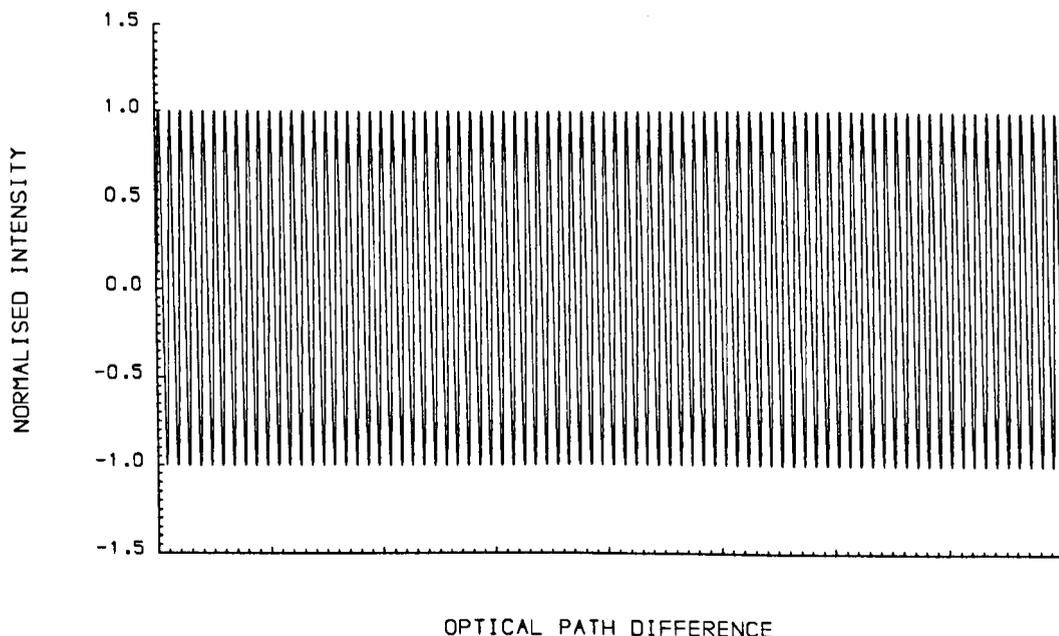


Fig. 1.1 Output intensity with a He-Ne laser source. $L_c = 1000\text{ m}$, corresponding to 1.58×10^9 fringes.

In single-mode He-Ne gas lasers the spectrum consists of a very sharp line at a wavelength of 632.8 nm [6], and the coherence length may be about 1000 metres [7]. When δL between the two arms of the interferometer reaches $\pm L_c/2$, the fringe

visibility reduces to $1/e$ of its maximum [8]. Hence, the visibility as a function of δL decays very slowly, giving an output intensity as in Fig. 1.1, which in practice can be approximated by Eq. 1.2.

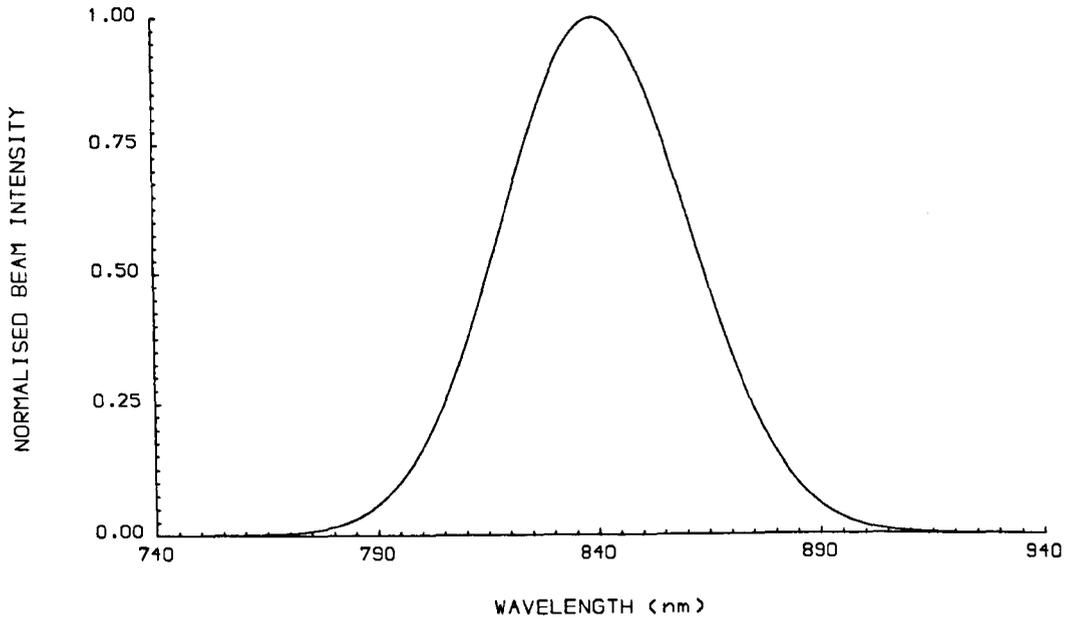


Fig. 1.2 Typical spectrum of an AlGaAs LED, emitting at a central wavelength of 840 nm with a spectral half-width of 49.4 nm .

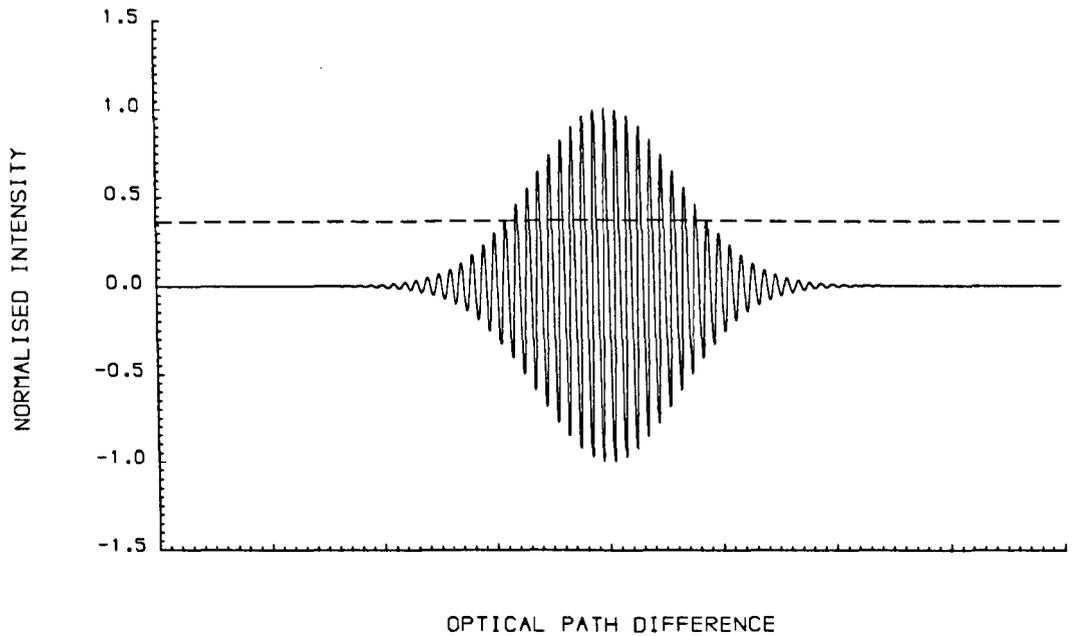


Fig. 1.3 Output intensity of an interferometer with a LED source, obtained by adding 201 wavelengths spaced by 1 nm and in phase at $\delta L = 0$, with intensities as in Fig. 1.2. $L_c = 14.3\text{ }\mu\text{m}$, corresponding to 17 fringes. The dashed line delimits the $1/e$ intensity points.

In contrast, light-emitting diodes (LEDs) and multimode laser diodes operated below threshold are polychromatic sources with a broad-band Gaussian spectral distribution and a spectral half-width ranging from 20 to 80 nm [9], as in Fig. 1.2. As the Fourier transform is a linear operation, the fringe visibility is also a Gaussian function, as in Fig. 1.3, and the coherence length may only reach a few μm .

1.5 White-Light Interferometry

The apparent inconvenience caused by polychromatic light sources can be turned to advantage by exploiting the following limitations associated with the use of high-coherence laser devices [10]:

- Unambiguous operating range corresponding to only one wavelength, which in the case of position and distance measurements allows a maximum movement of the measurand by half a wavelength (Eq. 1.3).
- The inability to identify the interference order when the interferometer is switched off and on.

These limitations are due to the long coherence length of a laser source, which generates a very flat visibility profile and makes it extremely difficult to monitor the movement of the fringes, unless complex and expensive fringe-counting methods are used [11].

When using a source with a short coherence length, the input spectrum is approximately Gaussian, hence the output fringe pattern is the cosine function in Eq. 1.2 modulated by a Gaussian visibility profile. There is then a central white fringe corresponding to the monochromatic fringe of order zero, with a few coloured maxima and minima on either side. The central fringe will remain identified if the power supply is suspended temporarily, either accidentally or on purpose.

A white-light interferometer (WLI) is an interferometric system which uses a polychromatic source and usually consists of two interferometers in series [12, 13]. Referring to Fig. 1.4, the interferometer on the left (called the sensor) can be preset to introduce a path difference δL_1 between its two arms much greater than the coherence length L_c of the source S , by adjusting the movable mirror M_2 . With this, the two output wave trains from the beam-splitter B_1 do not overlap and

cannot interfere. Upon reaching the interferometer on the right (the reference or scanner) wave train L_1 is split into L_{13} and L_{14} at the beam-splitter B_2 , while L_2 is split into L_{23} and L_{24} . Although the two wave trains reflected from the fixed mirror M_3 cannot interfere, and neither can the two wave trains from the movable mirror M_4 , L_{13} can interfere with L_{14} and L_{23} can interfere with L_{24} if the path difference δL_2 between the two arms of the reference is within L_c . The real advantage of using such an arrangement, however, is that L_{13} can interfere with L_{24} (or L_{23} interfere with L_{14}) when δL_2 is equal to δL_1 to within L_c . This allows to operate over a range much larger than that imposed by L_c , and which is now dictated only by the scanning range of the reference interferometer.

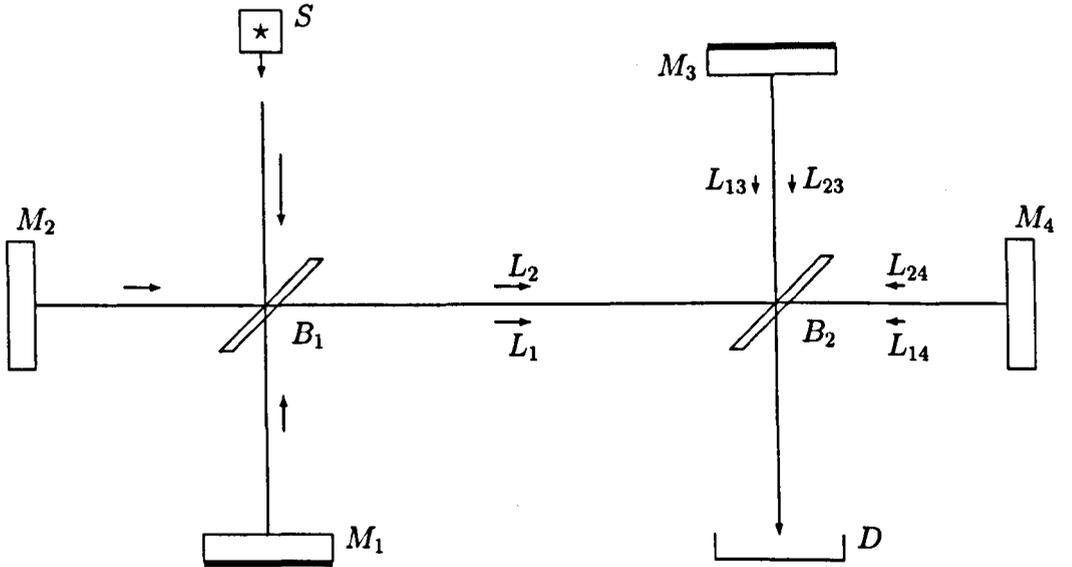


Fig. 1.4 Schematic diagram of a WLI system consisting of two Michelson interferometers.

During the calibration phase, the reference is adjusted to match δL_2 to δL_1 . During the measurement or sensing phase, the interference fringes will disappear from the detector D whenever the change in δL_1 induced by the measurand in the sensing interferometer is greater than L_c . By scanning the reference, the change in δL_1 can be matched by an equal change in δL_2 , thus restoring the interference pattern.

The system can be viewed as an optical delay line in which the two incoherent signals in the two arms of the sensing interferometer are brought together and become coherent.

The scanning can be either temporal, using a mechanical or piezo-electrical device, or spatial, using an electronic scanner [14]. With the latter, the reference interferometer produces the fringes by expanding and overlapping the beams from the two arms at an angle on a charge-coupled device (CCD) array. Compared with mechanical scanners, electronic scanners are much faster, more accurate and stable, and have smaller size. Compared with piezo-electric scanners they have a much larger operating range and do not need a high voltage driver [15]. Since no mechanical moving parts are present, their compactness and rigidity makes them especially useful for applications in harsh environments [8].

1.6 Fibre-Optic Interferometric Sensors

An additional advantage of white-light interferometry is that only the sensing unit needs to be in the sensing area, as long as the optical signal can be transmitted back to the reference interferometer. This makes it possible to take measurements in hazardous and hostile environments, and is one of the reasons for the recent interest in fibre-optic sensors, which offer more safety, reduced weight, and are more resistant to high-temperatures [9].

Such sensors exploit the fact that the OPD in a fibre is affected by its temperature, and also changes with pressure and stretching, or when an electric or magnetic field is applied. Fibre optic sensors using white-light interferometry have been developed, among others for the measurement of absolute displacement [13, 16], refractive index [17], pressure [9], and temperature [18]. In particular, with an all-fibre arrangement, where the beam splitters are replaced by optical-fibre couplers, high sensitivity can be obtained, as it is possible to have very long paths in a small space. A considerable reduction in noise can also be achieved because of the immunity of the system to perturbations in the transmission medium [10].

The peak wavelengths emitted by semiconductor light sources (LEDs and laser diodes) correspond to the minimum attenuation wavelengths of commercially available fibres, allowing easy coupling of the light to the transmission link and eliminating the need for high driving voltages and currents typical of high-coherence sources [6]. In addition, an all-fibre arrangement gives the possibility of making multipoint measurements by sending light to different sensing interferometers through a com-

mon optical busbar and back onto the same busbar to one reference interferometer, where the individual interference patterns can be recovered by coherence multiplexing [19, 20]. If the optical delays from each sensor are sufficiently spaced so that no mixing of the light beams occurs in the busbar, in fact, each individual signal can be recovered in turn by scanning the reference over its full range.

1.7 Digital Signal Processing

The normalised fringe pattern at the output of a WLI system can be expressed, after sampling, as follows [8]

$$s[n] = \exp \left[- \left(\frac{kn - \theta_s}{\pi L_c} \right)^2 \right] \cos(kn - \theta_s) \quad (1.6)$$

where n denotes a sample point, θ_s and kn are the phase differences in the sensing and reference interferometers, respectively, with k being equal to $2\pi/b$, where b is the number of samples per fringe or fringe width, and L_c is the number of fringes within the coherence length.

When $kn = \theta_s$, a perfect matching between the OPDs of the two interferometers is achieved, giving maximum interference contrast.

For a given system setup, L_c and b are fixed parameters; on the other hand, θ_s may change from one scan to the next according to the measurand, like a random variable uniformly distributed across the scanning range. Hence, the output at an arbitrary point n is a particular realisation of a random process $\{s[n]\}$. Nevertheless, having observed the sample function $s[n]$ one has no difficulty in determining the unknown θ_s from Eq. 1.6; all that is needed is to look at the sample point corresponding to the maximum value, and equate θ_s with kn . In this respect, Eq. 1.6 provides a deterministic description of the system.

In the presence of noise, however, the explicit mathematical relationship between system output and phase difference works but with a certain level of error, and it is the uncertainty in deriving the latter that will justify referring to the WLI system as a stochastic system in this thesis.

In addition to shot noise and oscillations in the electrical supply characteristic of LED sources, in laser diodes resonance effects and variations of the optical length in the laser cavity produce frequency variations in the light source which appear as amplitude variations of the visibility profile [1].

Misalignment and offset between the axis of the measurement transducer and the target displacement, optical mixing between the polarisation states of the source caused by imperfect or tilted components, diffraction and wavefront aberrations, changes in refractive index induced by temperature and pressure [21], and shot and thermal noise at the detector all contribute to make identification of the central fringe a difficult task.

These problems are exacerbated in fibre-optic systems because of the low power-coupling between the source and the fibre. Furthermore, vibrations during scanning bring the overall signal-to-noise ratio (SNR) down to 40-60 dB in temporally scanned systems [8]. For electronically scanned systems, on the other hand, the limited spatial coherence of the source reduces the visibility profile as the two beams are brought together at an angle on the CCD array, pushing the SNR further down to 20-40 dB [20].

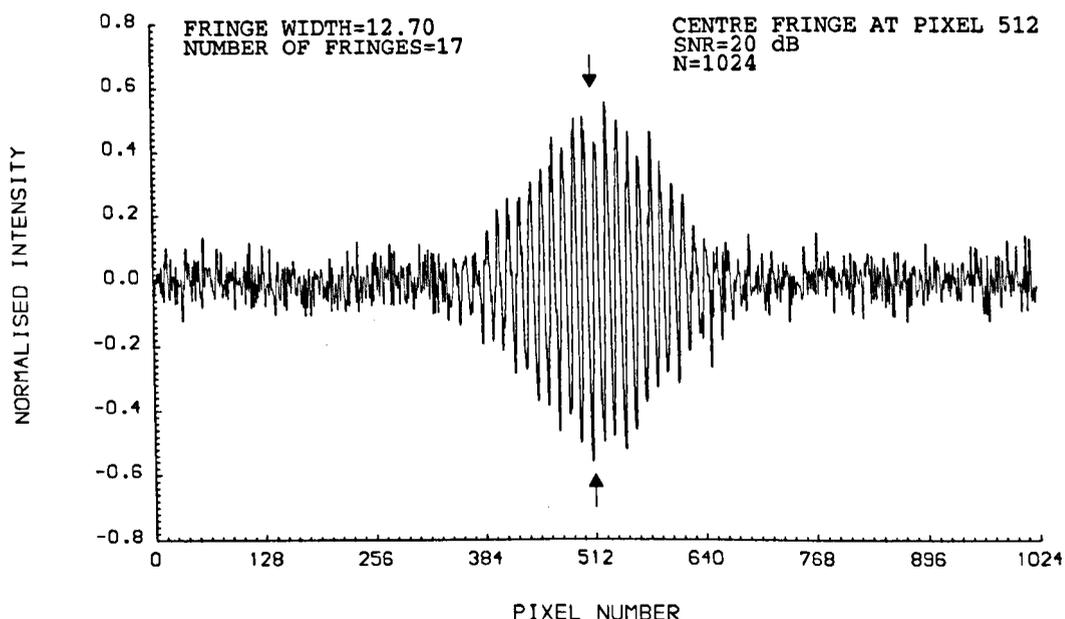


Fig. 1.5 Simulated output intensity of a WLI system with additive noise at 20 dB. The arrows point at the central fringe.

In the following, the SNR (in dB) is defined as $20 \log(A/\sigma)$, where A is the amplitude of the noise-free central fringe and σ is the root mean square of the noise, assumed stationary [22]. A SNR of 20 dB, for example, means that noise and central fringe amplitude are in the ratio 1:10.

Fig. 1.5 shows a simulated fringe pattern corrupted by pseudo-random white

Gaussian noise at 20 dB. The coherence length corresponds to 17 fringes and a CCD array consisting of 1024 pixels was assumed for the detector.

Fig. 1.6 shows the percentage of successful sub-fringe identifications from direct observation of the visibility profile, out of 300 computer simulations with the central fringe in the middle of the scanning range. The sampled intensity at pixel n is compared to that at pixel $n + 1$ in the computer, and the global maximum in the fringe pattern is reported at the end of the scan. The central fringe has been identified to sub-fringe level when the pixel corresponding to maximum intensity in the noise-free central fringe is recovered correctly in the presence of noise.

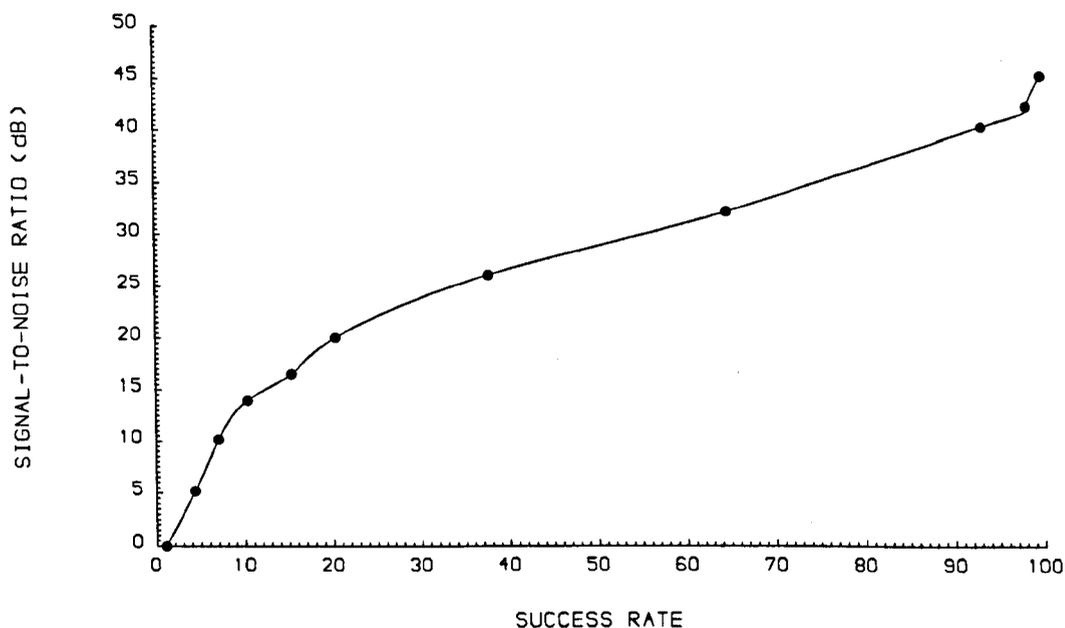


Fig. 1.6 Success rate from direct observation of the visibility profile for various SNRs.

Clearly, the identification rate degrades rapidly as the SNR falls below 40 dB. Down to around 26 dB the fringe and sub-fringe rates were the same, but below 26 dB the fringe rate was 1.5-2 times higher than the sub-fringe rate. If using a coherence multiplexing scheme or a light source with a longer coherence length, the identification rate would be even lower than that shown in Fig. 1.6.

Several signal processing operations have been proposed recently in an attempt to increase the success rate [8, 22, 23, 24, 25, 26, 27, 28, 29].

The methods reported in [8, 22, 23] use centre of gravity or centroid algorithms which exploit the symmetry of the pattern on the two sides of the central fringe. Although very simple to implement, these methods fail when the central fringe

moves just enough from the middle point on the CCD detector that the symmetry is lost. Asymmetries in the spectral profile of the source and any system misalignment will also increase the failure rate.

Multiple-wavelength combination sources, in which two [24, 25, 26] or three [27] low-coherent sources with different central wavelengths are superposed incoherently to produce a pattern with an enhanced central fringe, have shown considerable promise, but the cost and system alignment difficulty are high. This is because more optical components are needed, and techniques for finding the optimum wavelength combination of the sources have to be used [28, 29].

A simple solution for noise reduction would be to average together successive patterns, which can be easily done with the help of a portable computer. If the number of averaged traces is N , the root mean square of the noise is reduced by the factor \sqrt{N} [30]. However, such approach increases the system delay time by N , and relies on the quite unrealistic assumption that the central fringe remains still throughout.

1.8 Aims and Objectives of the Thesis

The main aims and objectives of this thesis can be listed as follows:

- Model the output of a white-light interferometric system by a suitable time-series model with a finite set of parameters.
- Exploit the connection between model identification and inverse filtering to evaluate the behaviour and performance of existing batch and recursive filters in processing the output fringe pattern.
- Improve on the central fringe identification rate of the centroid method by devising a novel scheme based on an adaptive digital filtering algorithm.
- Analyse the effects of non-linearities and evaluate the need for higher-order statistics and non-linear filtering in central fringe identification.

1.9 Summary

In this chapter white-light interferometry (WLI) was introduced as a means of overcoming the operating range limitation imposed by classical optical interferometers which use highly monochromatic sources. The problem of central fringe identification in the presence of noise was then explained, before concluding with a brief description of the techniques currently available to ease it.

In the next chapter the WLI pattern will be treated as the output of a "black-box" model consisting of a finite set of unknown parameters. Methods for the estimation of the model parameters will be described, and the central fringe identification task will be formulated as an inverse filtering problem, where the coefficients of the filter can be related to the parameters of the model.

Bibliography

- [1] M. Born and E. Wolf, *Principles of Optics: Electromagnetic Theory of Propagation, Interference and Diffraction of Light*, 4th ed. Oxford: Pergamon, 1970.
- [2] F. A. Jenkins and H. E. White, *Fundamentals of Optics*, 4th ed. New York: McGraw-Hill, 1976.
- [3] J. W. Blaker and W. M. Rosenblum, *Optics: An Introduction for Students of Engineering*. New York: MacMillan, 1993.
- [4] D. Casasent and H. J. Caulfield, "Optical data processing," in *Topics in Applied Physics*, vol. 23. Heidelberg: Springer-Verlag, 1978.
- [5] M. Francon, *Optical Interferometry*. New York: Academic, 1966.
- [6] J. T. Verdeyen, *Laser Electronics*. Englewood Cliffs, NJ: Prentice-Hall, 1981.
- [7] J. Wilson and J. F. B. Hawkes, *Optoelectronics: An Introduction*, 2nd ed. Cambridge: Prentice-Hall Int. (UK), 1989.
- [8] S. Chen, A. W. Palmer, K. T. V. Grattan, and B. T. Meggitt, "Digital signal-processing techniques for electronically scanned optical-fiber white-light interferometry," *Appl. Opt.*, vol. 31, no. 28, pp. 6003-6010, Oct. 1992.
- [9] G. Beheim, K. Fritsch, and R. N. Poorman, "Fiber-linked interferometric pressure sensor," *Rev. Sci. Instrum.*, vol. 58, no. 9, pp. 1655-1659, Sept. 1987.
- [10] Y. N. Ning, K. T. V. Grattan, and A. W. Palmer, "The use of low coherence light sources in fibre-optic interferometric systems," in *Sensors, vol. 6: Optical Sensors*, E. Wagner, R. Dändliker, and R. Spenner, ed. Weinheim, Germany: VCH Verlagsgesellschaft mbH, pp. 529-550, 1992.

- [11] P. Hariharan, "Interferometry with Lasers," in *Progress in Optics XXIV*, E. Wolf, ed. New York: Elsevier Science, 1987.
- [12] K. T. V. Grattan and B. T. Meggitt, *Optical Fiber Sensor Technology*. London: Chapman and Hall, 1995.
- [13] T. Bosselman and R. Ulrich, "High accuracy position sensing with fiber-coupled white-light interferometry," *Proc. of OFS 2*, Stuttgart, pp. 361-364, 1984.
- [14] A. Koch and R. Ulrich, "Displacement sensor with electronically scanned white-light interferometry," *Int. Congress on Optical Science and Engineering*, The Hague, Netherlands, Mar. 12-15, 1990.
- [15] S. Chen, B. T. Meggitt, and A. J. Rogers, "A large dynamic range electronically-scanned white-light interferometer with optical-fibre Young's structure," in *Fiber Optic Sensors: Engineering and Applications*, A. J. Bruinsma and B. Culshaw, eds. (*Proc. Soc. Photo-Opt. Instrum. Eng.*), vol. 1511, pp. 67-77, 1991.
- [16] A. Koch and R. Ulrich, "Fiber-optic displacement sensor with 0.02 μm resolution by white-light interferometry," *Sensors and Actuators A*, vol. 25-27, pp. 201-207, 1991.
- [17] R. Bohm and R. Ulrich, "High-accuracy fiber-optic reflectometer for fluids," *Proc. 7th Optical Fiber Sensors Conf.*, Sidney, Australia, pp. 353-356, 1990.
- [18] C. Marillier and M. Lequime, "Fiber-optic white-light birefringent temperature sensor," *Proc. Soc. Photo-Opt. Instrum. Eng.*, vol. 798, pp. 121-130, 1987.
- [19] J. L. Brooks, R. H. Wentworth, R. C. Youngquist, M. Tur, B. Y. Kim, and H. J. Shaw, "Coherence multiplexing of fiber-optic interferometric sensors," *J. Lightwave Technol.*, vol. 3, no. 5, pp. 1062-1072, Oct. 1985.
- [20] S. Chen, B.T. Meggitt, and A.J. Rogers, "Novel electronic scanner for coherence multiplexing in a quasi-distributed pressure sensor," *Electron. Lett.*, vol. 26, no. 17, pp. 1367-1369, Aug. 1990.
- [21] N. Bobroff, "Recent advances in displacement measuring interferometry," *Meas. Sci. Technol.*, vol. 4, pp. 907-926, 1993.

- [22] S. Chen, A. W. Palmer, K. T. V. Grattan, and B. T. Meggitt, "Fringe order identification in optical fibre white-light interferometry using centroid algorithm method," *Electron. Lett.*, vol. 28, no. 6, pp. 553-555, Mar. 1992.
- [23] R. Dändliker, E. Zimmermann, and G. Frosio, "Noise-resistant signal processing for electronically scanned white-light interferometry," *Proc. 8th Optical Fiber Sensors Conf.*, Monterey, CA, pp. 53-56, 1992.
- [24] S. Chen, K. T. V. Grattan, B. T. Meggitt, and A. W. Palmer, "Instantaneous fringe-order identification using dual broad-band source with widely spaced wavelengths," *Electron. Lett.*, vol. 29, no. 4, pp. 334-335, Feb. 1993.
- [25] Y. J. Rao, Y. N. Ning, and D. A. Jackson, "Synthesized source for white-light sensing systems," *Opt. Lett.*, vol. 18, no. 6, pp. 462-464, Mar. 1993.
- [26] Y. J. Rao and D. A. Jackson, "Improved synthesised source for white light interferometry," *Electron. Lett.*, vol. 30, no. 17, pp. 1440-1441, Aug. 1994.
- [27] D. N. Wang, Y. N. Ning, K. T. V. Grattan, A. W. Palmer, and K. Weir, "Three-wavelength combination source for white-light interferometry," *IEEE Photon. Technol. Lett.*, vol. 5, no. 11, pp. 1350-1352, Nov. 1993.
- [28] D. N. Wang, Y. N. Ning, K. T. V. Grattan, A. W. Palmer, and K. Weir, "The optimized wavelength combinations of two broadband sources for white light interferometry," *J. Lightwave Technol.*, vol. 12, no. 5, pp. 909-916, May 1994.
- [29] D. N. Wang, *White Light Interferometric Sensor Systems*. PhD thesis, Dept. Electrical, Electronic and Information Engineering, City Univ., London, Feb. 1995.
- [30] T. H. Wilmshurst, *Signal Recovery from Noise in Electronic Instrumentation*, 2nd ed. Bristol: IOP Publishing, 1990.

Chapter 2

Time Series Modelling of WLI Systems

2.1 Introduction

In this chapter the WLI system will be represented by a parametric linear model driven by an unobservable white noise sequence. It will be shown that the recovery of the central fringe can be treated as either a system identification problem or, equivalently, as an inverse filtering problem. Current techniques for estimating the model order and parameters will also be described.

2.2 Autoregressive-Moving Average Models

Parametric modelling of a time series is based on the assumption that the measured data under investigation evolves from a stochastic process that can be represented by a selected model with a suitable set of parameters.

The initial motivation for parametric models was the possibility of obtaining power spectrum estimates with higher frequency resolution than those produced by the more classical methods [1, 2] while reducing sidelobe leakage caused by spectral smoothing or windowing [3, 4].

Due to the fact that a continuous frequency spectrum can be approximated closely by a suitable rational function with a large number of parameters [5], the autoregressive-moving average (ARMA) model [6] has received the most attention in the literature.

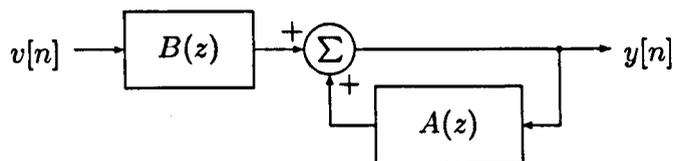
The underlying idea is that a random stationary time series $y[n]$ can be expressed in terms of its past p values and the present and past q values of a sequence of uncorrelated shock or disturbance terms $v[n]$ drawn randomly from a fixed probability distribution with zero mean ¹ [7]. Successive values of the series are related through the linear difference equation

$$y[n] = - \sum_{k=1}^p a_k y[n-k] + v[n] + \sum_{l=1}^q b_l v[n-l] \quad (2.1)$$

Taking the discrete z transform the ARMA equation in the z domain is obtained

$$y[n] = \frac{B(z)}{A(z)} v[n] = \frac{1 + b_1 z^{-1} + \dots + b_q z^{-q}}{1 + a_1 z^{-1} + \dots + a_p z^{-p}} v[n] \quad (2.2)$$

The time series $y[n]$ is thus the output of a linear discrete dynamical system driven by a white noise input and possessing a rational pole-zero transfer function $H(z) \equiv B(z)/A(z)$. Fig. 2.1 is a block diagram of an ARMA process.



$$y[n]A(z) = v[n]B(z)$$

$v[n]$ = white noise sequence
AR model : $B(z) = 1$

$y[n]$ = output sequence
MA model : $A(z) = 1$

$$\begin{aligned} A(z) &= 1 + a_1 z^{-1} + \dots + a_p z^{-p} \\ B(z) &= 1 + b_1 z^{-1} + \dots + b_q z^{-q} \end{aligned}$$

Fig. 2.1 ARMA process driven by white noise.

Setting the a_k parameters to zero gives an all-zero or moving average (MA) model, whereas setting the b_l parameters to zero gives an all-pole or autoregressive (AR) model. The MA model corresponds to a finite impulse response (FIR) filter [8], whereas the AR and ARMA models correspond to infinite impulse response (IIR) filters [9].

The spectral density of an ARMA process at a discrete frequency f can be expressed as a function of its parameters as follows

$$P_{ARMA}(f) = T\sigma_v^2 \left| \frac{1 + \sum_{l=1}^q b_l \exp(-j2\pi l f T)}{1 + \sum_{k=1}^p a_k \exp(-j2\pi k f T)} \right|^2 \quad (2.3)$$

¹Such a sequence is also known as white noise.

where T is the sampling period and σ_v^2 is the variance of the driving noise.

Eq. 2.3 provides a useful tool at the model selection stage. Since AR models are able to reproduce sharp spectral peaks at those frequencies where the denominator term approaches zero, they may be particularly useful for the modelling of processes with narrow-band spectra, whereas MA models may be better at representing broad-band processes. The more general ARMA models can account for both narrow- and broad-band behaviour, a feature which may make them preferable when dealing with processes with mixed spectra.

2.3 Parametric Modelling of WLI Systems

Second order differential equations of the form

$$\ddot{s}(t) + \alpha\dot{s}(t) + \beta s(t) = 0 \quad (2.4)$$

are often used to describe an oscillating system. If $s(t)$ represents the amplitude of an oscillation at time t , Eq. 2.4 with $\alpha = 0$ describes simple harmonic motion, whereas with $\alpha > 0$ it describes damped harmonic motion [10].

In discrete time n , a sampled process consisting of one harmonic can be expressed by the second order difference equation [11]

$$s[n] = -a_1s[n-1] - a_2s[n-2] \quad (2.5)$$

A sampled process consisting of $p/2$ harmonic components can similarly be expressed as [12]

$$s[n] = -\sum_{i=1}^p a_i s[n-i] \quad (2.6)$$

Given initial conditions $s[1], \dots, s[p]$, Eq. 2.6 represents a deterministic process, since future behaviour is known with certainty from present and past values. However, in the presence of an external or internal noise source $v[n]$, the additive process

$$y[n] = s[n] + v[n] = -\sum_{i=1}^p a_i s[n-i] + v[n] \quad (2.7)$$

is random (or stochastic), since there is some degree of uncertainty before it actually occurs. Rewriting Eq. 2.7 as

$$y[n] = -\sum_{i=1}^p a_i y[n-i] + v[n] + \sum_{i=1}^p a_i v[n-i] \quad (2.8)$$

this represents an ARMA process with identical AR and MA parameters.

As shown in Section 1.4, the pattern at the output of a WLI system is derived from the interference effect of a continuous spectrum of frequencies lying within the bandwidth of the light source. Thus, the sampled intensity $s[n]$ expressed by Eq. 1.6 in Section 1.7 groups together an infinite number of harmonic components, and the fringe width b times the sampling period T can be viewed as the discrete equivalent of the mean wavelength λ_0 of the source. The frequency spectrum of $s[n]$ is therefore a Gaussian function with width inversely proportional to the coherence length of the source and centred at the mean frequency $1/(Tb)$.

In practice, it is not necessary to consider an infinite number of harmonics, as Fig. 1.3 demonstrated. If the sampled fringe signal can be represented quite accurately by the sum of a finite number of frequency components plus an additive noise term, the WLI system can be well approximated by an ARMA(p, p) process with identical $A(z)$ and $B(z)$ terms, with the measurement noise providing the driving sequence.

Given that the noise is physically found at the output end, this may seem somewhat odd. What one is trying to fit here, though, is not a physical model but a stochastic one. By regarding the WLI system as a "black-box" driven by an unobservable random process, one is dispensed with the impossible task of unravelling the hidden and complex mathematical relationships governing the system, and can concentrate on finding a suitable set of parameters which account for the internal system behaviour and its output in statistical terms.

2.4 Inverse Filtering and the Wiener Solution

Having chosen a suitable model for the generation of the fringe pattern, the next step is to find the set of parameters which provide the best fit between the system and the model. From Eq. 2.8, a reasonable estimator for $s[n]$ would be

$$\hat{s}[n] = -\sum_{i=1}^p \hat{a}_i y[n-i] + \sum_{i=1}^p \hat{a}_i \hat{v}[n-i] \quad (2.9)$$

where $\hat{v}[n-i] = y[n-i] - \hat{s}[n-i]$ would be a reasonable estimator for $v[n-i]$. The most often used measure of fit is the mean square error (MSE), originally formulated for the Wiener filter [13]. If $x[n]$ is a zero-mean stationary input sequence to a linear

system and $d[n]$ is a desired response, the filter which minimises the function

$$J[n] = E\{e^2[n]\} \quad (2.10)$$

where $e[n]$ is the error signal between $d[n]$ and the filter output, and E denotes expectation, is the Wiener filter, whose impulse response \mathbf{h}_0 is given by the Wiener-Hopf equations [14]

$$\mathbf{R}_{\mathbf{x}\mathbf{x}}\mathbf{h}_0 = \mathbf{R}_{\mathbf{d}\mathbf{x}} \quad (2.11)$$

where $\mathbf{R}_{\mathbf{x}\mathbf{x}}$ is the autocorrelation matrix of the filter input and $\mathbf{R}_{\mathbf{d}\mathbf{x}}$ is the cross-correlation vector between input and desired response ².

Letting $y[n]$ be the desired response and $\hat{s}[n]$ the filter output it follows that

$$J[n] = E\{v^2[n]\} + E\{(s[n] - \hat{s}[n])^2\} + 2E\{v[n](s[n] - \hat{s}[n])\} \quad (2.12)$$

Since $v[n]$ is uncorrelated with $s[n]$ and with past values of the desired response and error signals, the last term on the right-hand side of Eq. 2.12 is zero ³. Minimising $J[n]$ is thus equivalent to minimising the second term on the right-hand side. As this term is quadratic, $J[n]$ has a unique minimum with value $E\{v^2[n]\}$. If this minimum is achieved, the filter is said to be optimum in the MSE sense, because its output and error signals become equal to the noise-free fringe and measurement noise sequences, respectively, as $s[n] - \hat{s}[n] = e[n] - v[n]$.

ARMA modelling of the WLI system can therefore be seen as either a system identification from only output data problem [15], where it is assumed that the system and the IIR filter that models it are excited in parallel by the same noise sequence, and the task is to estimate the system parameters or transfer function in order to reduce the mismatch between system and model output, or as an inverse filtering problem, where the system and an inverse IIR filter are connected in series and the task is to find the inverse filter that is able to reproduce the driving noise sequence at its output end ⁴.

In the following, the WLI system will be treated as an ARMA(p, q) model with $B(z)$ in general different from $A(z)$. This is to take account not only of computational errors which may arise during the estimation phase, but especially of model

²In this thesis, correlation is used for both normalised and unnormalised covariance.

³ $E\{xy\} = E\{x\}E\{y\} = 0$ if x and y are zero-mean uncorrelated random variables.

⁴Such a filter is also called a *whitening* filter because it decorrelates its input [16].

order misspecification and parameter mismatches that may be caused by inaccurate modelling assumptions (eg., can multiplicative internal system noise, if present, be translated and add to the measurement noise at the system output, and is the noise itself white or coloured).

2.5 Estimation of the Model Parameters

The various methods used to estimate the parameters of an ARMA process all start from an auto- and cross-correlation formulation of Eq. 2.1. Multiplying both sides by $y[n - m]$ and taking expectations one obtains

$$r[m] = - \sum_{k=1}^p a_k r[m - k] + \sum_{l=0}^q b_l r_{vy}[m - l] \quad (2.13)$$

where $r[m]$ and $r_{vy}[m]$ are the autocorrelation of the output and the cross-correlation between input and output at lag m , respectively.

Eq. 2.13 for various lags gives a set of non-linear Wiener-Hopf equations, known as the ARMA Yule-Walker (ARMA Y-W) equations [7, 17]. From here one of two sets of techniques is chosen: optimum or sub-optimum.

The first use an iterative approach based on maximum likelihood estimation to solve the equations directly [18, 19, 20, 21], but they are computationally demanding and may converge to the wrong solution [22].

The second reduce the problem to a linear one by estimating the AR and MA parameters separately [23, 24, 25, 26], and as they allow to keep the computational load to levels compatible with real-time processing, they will be considered next.

To evaluate the AR parameters use is made of the fact that for a causal ⁵ system the second summation term in Eq. 2.13 drops out for $m > q$. Hence, a set of p linear equations in p unknowns, commonly referred to as the modified Y-W equations, can be formed for $q + 1 \leq m \leq q + p$ and solved by Gaussian elimination.

With N data, this would require $O(Np)$ operations to estimate the autocorrelation terms, plus $O(p^3)$ operations to invert the autocorrelation matrix. As this is Toeplitz, however ⁶, it is possible to invert it with only $O(p^2)$ operations using the Levinson-Durbin [27, 28] or the Schur [29] algorithms, either in their original or

⁵A system is causal if the output does not depend on future input values.

⁶A matrix is Toeplitz if all the elements along any of its diagonals are identical.

more computationally efficient split-forms [30, 31], with the added advantage that the storage space is reduced from $O(p^2)$ to $O(p)$ memory cells.

To decrease parameter hypersensitivity, an over-determined system with $t > p$ equations can be formed [32], resulting in a product-of-Toeplitz autocorrelation matrix which can be solved with $O(t^2)$ operations [33].

To complete the ARMA modelling it is necessary to estimate the MA component. This can be done by first computing the *residual time series*, defined as

$$\tilde{y}[n] = y[n] + \sum_{k=1}^p \hat{a}_k y[n-k] \quad (2.14)$$

for $n = p + 1, \dots, N$, using the AR parameters just estimated. Since

$$\tilde{y}[n] = \hat{A}(z)y[n] \approx \hat{B}(z)v[n] = v[n] + \sum_{l=1}^q \hat{b}_l v[n-l] \quad (2.15)$$

approximate estimates of the MA parameters can be obtained by solving the system in Eq. 2.13 for $0 \leq m \leq q$, with $p = 0$ and y replaced by \tilde{y} . Rather than using computationally difficult spectral factorisation techniques [34] a preferred approach is to approximate the MA(q) residual process with a high order AR(r) process [35], with $r \gg q$, exploiting the Wold decomposition theorem [6], which states that a finite-order MA process can be represented as a unique AR model of possibly infinite order.

The solution to the high order AR approximation requires $O(Nr + r^2)$ operations to estimate the autocorrelation matrix and invert it with the Levinson-Durbin algorithm. A further $O(rq + q^2)$ operations is needed to derive the MA from the AR parameters using the same algorithm [36].

The Wold decomposition theorem can be taken one step further to approximate the whole ARMA process by an AR process of higher order. The quality of the MA estimate from the residual time series, in fact, depends heavily on the accuracy of the AR estimate from the modified Y-W equations, which is of poor quality if the process contains spectral regions with small values [37], or if the choice of the AR order is incorrect [36]. Using an over-determined system of equations, the variance of the AR estimate is reduced but its quality is now dependent on the location of the process poles and zeros, with maximum degradation when the poles are away from the unit circle and the zeros are close to it [38].

If the WLI output can be regarded as a harmonic process in additive white Gaussian noise, its poles and zeros will tend to cluster on the unit circle [39]. The

narrow-band nature of the process will also ensure that spectral regions with very small content abound. Hence, use of the modified Y-W equations may not result in good AR estimates.

Fig. 2.2 compares the power spectral density of the noisy fringe pattern of Fig. 1.5 with that of a computer-generated zero-mean white Gaussian process. The data were normalised to have unit variance and the estimates were computed using the classical Blackman-Tukey approach [1].

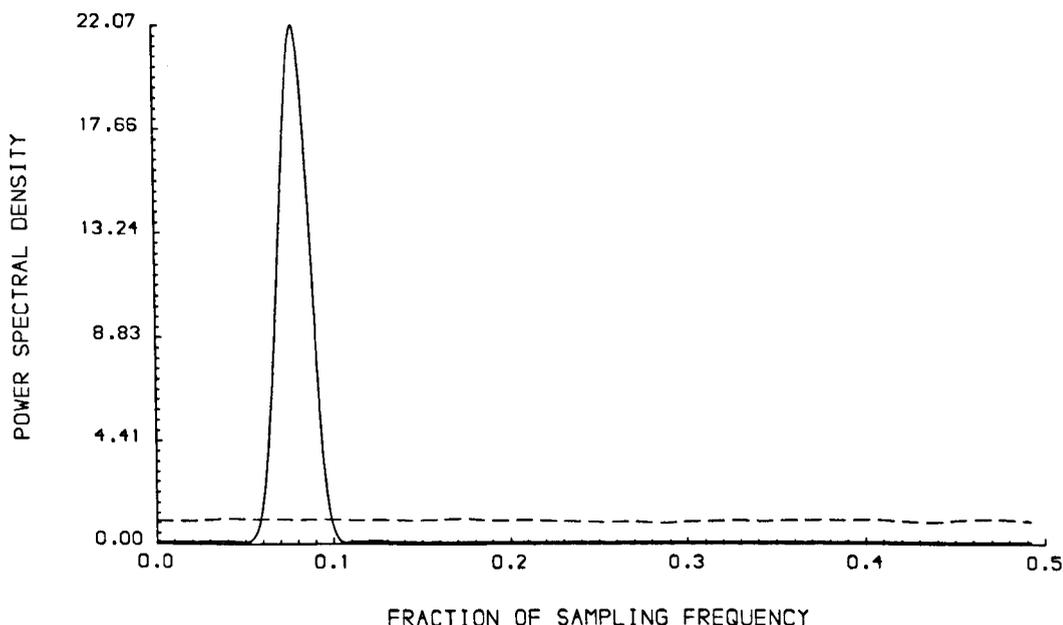


Fig. 2.2 Estimated spectrum of the central 500 samples of the fringe pattern in Fig. 1.5, smoothed with the Parzen window [40] and averaged over 30 independent realisations. The broken line represents the spectrum of a pseudo-random white Gaussian process.

The narrow-band property of the interferometric signal is clearly evident, and does not seem to require the mixed-spectrum representation of the more flexible ARMA model.

Fig. 2.3 shows estimates of the autocorrelation sequence (ACS) and the partial autocorrelation sequence (PACS) of the fringe pattern. The PACS estimate was computed with the Durbin method [28].

The ACS of an $AR(p)$ process and the PACS of an $MA(q)$ process are mixtures of damped exponentials and/or damped sine waves, whereas the PACS of an $AR(p)$ process and the ACS of an $MA(q)$ process cut off after lags p and q , respectively. Moreover, the ACS and PACS of an $ARMA(p, q)$ process are mixtures of

damped exponentials and/or damped sine waves after the first $q - p$ and $p - q$ lags, respectively [34].

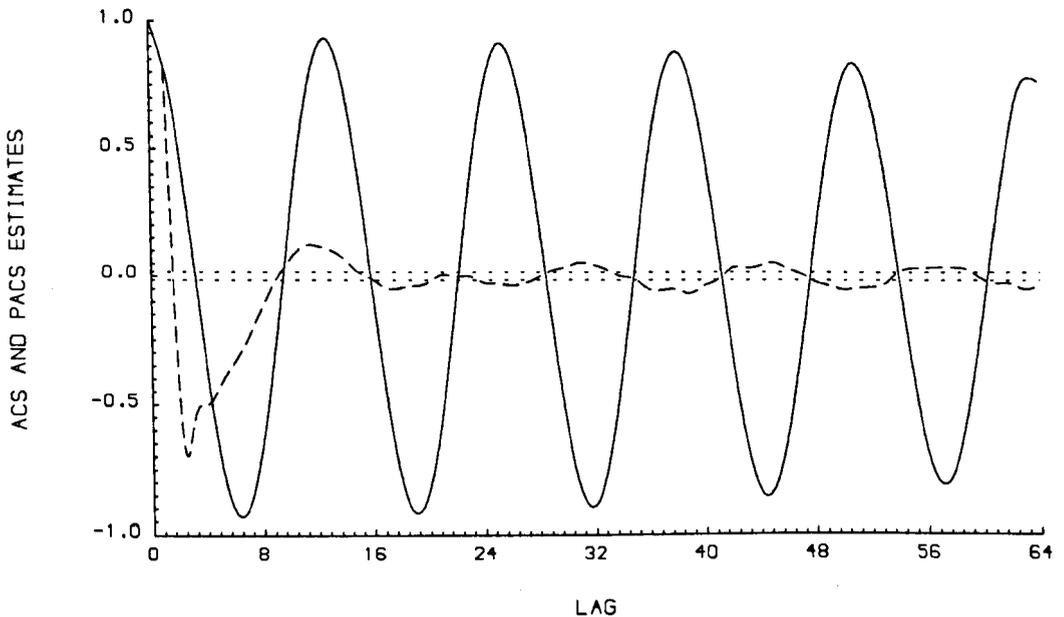


Fig. 2.3 Estimated ACS (continuous line) and PACS (broken line) of the central 500 samples of the fringe pattern in Fig. 1.5, averaged over 30 independent realisations. The dashed lines denote the 95 % confidence limits of the estimates [41, 42, 43].

From Fig. 2.3 it appears that the WLI output of Fig. 1.5 can be represented either by an ARMA process of equal but unknown AR and MA orders, or by an AR process of order between 10 and 16. Deciding on the model order from a simple visual examination of the ACS/PACS estimates is not very sensible though, as the confidence bounds rely on the large sample, stationarity, and whiteness assumptions.

Many criteria have been proposed for order selection, the final prediction error [44, 45], the information theoretic criterion [46, 47, 48], the Bayes information criterion [49, 50], the autoregressive transfer function [51], and the minimum description length [52, 53] to name but a few.

These are all closely related to earlier χ^2 -tests of significance for AR models, where the order p is tested against $p + 1$ [54, 55]. Although shown to work well for simulated processes, results have been mixed when applied to actual data. There is a general tendency to underestimate the model order [56, 57, 58, 39, 59] or to overestimate it [60]. Among the alternatives one has to mention testing for whiteness of the error sequence [61, 62] or checking the rank of the autocorrelation matrix [63, 64, 65] using the singular value decomposition [66].

One major motivation for preferring AR models is that they provide natural inverses for FIR filters, with the poles corresponding to the zeros of the filter. FIR filters do not introduce phase distortion and satisfy the bounded-input bounded-output property, whereas IIR filters become unstable when one or more of the poles stray outside the unit circle.

Although the Levinson-Durbin algorithm forces the zeros of the ARMA model (and hence the poles of the inverse filter) to lie inside the unit circle, the poles may cross to the other side in a digital implementation because of coefficient quantisation and round-off errors, which increase in severity as the bandwidth of the filter is reduced [67].

The choice between an AR, MA, or ARMA representation for a particular process is still a very complex one [68, 22].

2.6 Least Squares Filtering and Linear Prediction

The estimation of the parameters of an AR process follows similar lines to that of the AR parameters of an ARMA process. One can solve the AR Yule-Walker (AR Y-W) equations

$$r[m] = - \sum_{k=1}^p a_k r[m-k] \quad 1 \leq m \leq p \quad (2.16)$$

derived from Eq. 2.13 with $q = 0$ and $1 \leq m \leq p$. The Levinson algorithm can be used to invert the resultant Toeplitz autocorrelation matrix with $O(p^2)$ operations, after estimating the autocorrelation terms with $O(Np)$ operations.

An alternative derivation of Eq. 2.16 is provided by linear prediction analysis, which is a special case of least squares (LS) filtering [69] and gives a natural inverse for AR models by focusing directly on the available data sequence rather than on the statistics of the underlying process.

The LS filter minimises the sum of squares of the errors over the given data, replacing the Wiener ensemble averaging over all possible realisations of the process with time averaging across the observed realisation, leading to a solution which depends on the number of data.

The impulse response of the LS filter is obtained from the least squares Wiener-Hopf equations

$$(\mathbf{X}^T \mathbf{X}) \mathbf{h} = \mathbf{X}^T \mathbf{d} \quad (2.17)$$

where \mathbf{X} is a $N \times p$ matrix whose columns are shifted versions of the data vector.

The linear prediction problem is to find an estimate of the current sample of a random process from only previous samples, i.e.

$$\hat{y}[n] = - \sum_{k=1}^p w_k y[n-k] \quad (2.18)$$

for $n = p+1, \dots, N$. Defining the *forward error*⁷ as

$$e_f[n] = y[n] - \hat{y}[n] = y[n] + \sum_{k=1}^p w_k y[n-k] \quad (2.19)$$

and minimising the resultant sum of squares

$$\epsilon_f = \sum_{n=1}^N e_f^2[n] \quad (2.20)$$

by setting its partial derivative with respect to the w_k coefficients to zero, the least squares Wiener-Hopf equations for the linear predictor are obtained.

Since the elements of $\mathbf{X}^T \mathbf{X}$ are of the form

$$\sum_{n=1}^N y[n-k]y[n-m] \quad 0 \leq m-k \leq p \quad (2.21)$$

these equations are structurally identical to the AR Y-W equations, with the unnormalised estimates in Eq. 2.21 replacing the autocorrelation terms in Eq. 2.16. In particular, if the summation range in Eq. 2.20 is chosen as $1 \leq n \leq N+p$ and biased estimates are used in Eq. 2.16, the two sets become identical⁸.

Choosing instead the range $p+1 \leq n \leq N$ in Eq. 2.20, AR estimates with lower variance may be obtained, since only the available data $y[1], \dots, y[N]$ are used to construct the data matrix \mathbf{X} . This is known as the *covariance* method of least squares linear prediction [69, 71]. As $\mathbf{X}^T \mathbf{X}$ is a product of two Toeplitz matrices, the fast algorithm for the solution of the over-determined system in Section 2.5 can be used to keep the total operation count down to $O(Np + p^2)$.

2.7 Conclusion

In this chapter it was shown that the fringe pattern of a white-light interferometric system in the presence of noise may be modelled as the output of an autoregressive-

⁷*Forward* in the sense that the prediction for the current sample is a function of previous samples only.

⁸Biased estimates are preferred to unbiased ones because as the lag increases the larger bias is more than compensated for by a smaller variance [70].

moving average (ARMA) process driven by unobservable white noise. Several estimation schemes were briefly described, before showing that an autoregressive (AR) approximation to the full ARMA process may offer several advantages.

If an AR model is chosen, its order is possibly infinite (the Wold decomposition theorem). The closer the zeros of the moving average (MA) polynomial are to the unit circle, the larger the order of the AR model has to be for it to be a good approximation to the ARMA process. The estimated parameters will not converge to the true process parameters, as they have to compensate for the missing section.

The level of compensation required will also depend on the level of the observation noise. For high SNRs the best choice is usually a low-order AR model, but as the SNR decreases the order must necessarily increase to achieve adequate flexibility. This is one reason why signal processing applications make use of a large number of filter weights [72]. Indeed, virtually every time series encountered in practice can be approximated by a finite AR model of sufficiently high order [5].

Thus, it may be more sensible to consider the recovery of the noise-free pattern as an inverse filtering rather than as a system identification problem, in view of the fact that the error function minimises the deviation of the model error signal from the system driving noise, instead of the deviation between model and system parameters.

In the next chapter adaptive finite impulse response filters will be considered, which provide inverses for AR models whose parameters need to be re-estimated as frequently as possible. The performance of such filters in central fringe identification will be compared with that of the batch-processing filters introduced in this chapter.

The impact of the statistical properties of the fringe pattern data on the numerical accuracy of the estimates and the dependence of these on SNR and model order will also be examined.

Bibliography

- [1] R. B. Blackman and J. W. Tukey, *The Measurement of Power Spectra from the Point of View of Communications Engineering*. New York: Dover, 1959.
- [2] R. H. Jones, "A reappraisal of the periodogram in spectral analysis," *Technometrics*, vol. 7, no. 4, pp. 531-542, Nov. 1965.
- [3] J. P. Burg, *Maximum Entropy Spectral Analysis*. PhD thesis, Stanford Univ., Stanford, CA, May 1975.
- [4] E. Parzen, "Multiple time series modelling," in *Multivariate Analysis-II*, P. R. Krishnaiah, ed. New York: Academic, 1969.
- [5] L. H. Koopmans, *The Spectral Analysis of Time Series*. New York: Academic, 1974.
- [6] H. Wold, *A Study in the Analysis of Stationary Time Series*. Uppsala, Sweden: Almqvist and Wiksell, 1938.
- [7] G. U. Yule, "On a method of investigating periodicities in disturbed series, with special reference to Wölfer's sunspot numbers," *Philos. Trans. Royal Soc. London Ser. A*, vol. 226, pp. 267-298, July 1927.
- [8] H. E. Kallmann, "Transversal filters," *Proc. IRE*, vol. 28, no. 7, pp. 302-310, July 1940.
- [9] B. Gold and C. D. Rader, *Digital Processing of Signals*. New York: McGraw-Hill, 1969.
- [10] M. B. Priestley, *Spectral Analysis and Time Series*, vol. 1. London: Academic, 1981.

- [11] E. A. Robinson, "A historical perspective of spectrum estimation," *IEEE Proc.*, vol. 70, no. 9, pp. 885-907, Sept. 1982.
- [12] T. J. Ulrych and T. N. Bishop, "Maximum entropy spectral analysis and autoregressive decomposition," *Rev. Geophys. and Space Phys.*, vol. 13, no. 1, pp. 183-200, Feb. 1975.
- [13] N. Wiener, *Extrapolation, Interpolation and Smoothing of Stationary Time Series, with Engineering Applications*. New York: Wiley, 1949.
- [14] B. Widrow, "Adaptive filters I: fundamentals," *Tech. Rep. SU-SEL-66-126*, Stanford Electron. Lab., Stanford, CA, Dec. 1966.
- [15] G. C. Goodwin and R. L. Payne, *Dynamic System Identification: Experiment Design and Data Analysis*. New York: Academic, 1977.
- [16] A. Papoulis, "Maximum entropy and spectral estimation: a review," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 29, no. 6, pp. 1176-1186, Dec. 1981.
- [17] G. Walker, "On periodicity in series of related terms," *Proc. Royal Soc. A*, vol. 131, pp. 518-532, 1931.
- [18] P. Newbold, "The exact likelihood function of a mixed autoregressive moving-average process," *Biometrika*, vol. 61, no. 3, pp. 423-426, Dec. 1974.
- [19] G. F. Ansley, "An algorithm for the exact likelihood of a mixed autoregressive moving-average process," *Biometrika*, vol. 66, no. 1, pp. 59-65, Apr. 1979.
- [20] H. Akaike, "Maximum likelihood identification of Gaussian autoregressive moving-average models," *Biometrika*, vol. 60, no. 2, pp. 255-265, Aug. 1973.
- [21] N. K. Gupta and R. K. Mehra, "Computational aspects of maximum likelihood estimation and reduction in sensitivity function calculation," *IEEE Trans. Autom. Contr.*, vol. 19, no. 6, pp. 774-783, Dec. 1974.
- [22] S. L. Marple, Jr., *Digital Spectral Analysis with Applications*. Englewood Cliffs, NJ: Prentice-Hall, 1987.

- [23] W. Gersh, "Estimation of the autoregressive parameters of mixed autoregressive moving-average time series," *IEEE Trans. Autom. Contr.*, vol. 15, no. 5, pp. 583-588, Oct. 1970.
- [24] D. Graupe, D. J. Krause, and J. B. Moore, "Identification of autoregressive moving-average parameters of time series," *IEEE Trans. Autom. Contr.*, vol. 20, no. 1, pp. 104-107, Feb. 1975.
- [25] M. Kaveh, "High resolution spectral estimation for noisy signals," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 27, no. 3, pp. 286-287, June 1979.
- [26] J. F. Kinkel, J. Perl, L. Scharf, and A. Stubberub, "A note on covariance-invariant digital filter design and autoregressive-moving average spectral estimation," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 27, no. 2, pp. 200-202, Apr. 1979.
- [27] N. Levinson, "The Wiener rms (root mean square) error criterion in filter design and prediction," *J. Math. Phys.*, vol. 25, no. 4, pp. 261-278, Jan. 1947.
- [28] J. Durbin, "The fitting of time series models," *Rev. Int. Inst. Statist.*, vol. 28, no. 3, pp. 233-244, 1960.
- [29] T. Kailath, "A theorem of I. Schur and its impact on modern signal processing," in *I. Schur Methods in Operator Theory and Signal Processing*, I. Gohberg, ed., pp. 9-30. Basel, Switzerland: Birkhäuser Verlag, 1986.
- [30] P. Delsarte and Y. V. Genin, "The split Levinson algorithm," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 34, no. 3, pp. 470-478, June 1986.
- [31] P. Delsarte and Y. V. Genin, "On the splitting of classical algorithms in linear prediction theory," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 35, no. 5, pp. 645-653, May 1987.
- [32] R. K. Mehra, "On-line identification of linear dynamic systems with applications to Kalman filtering," *IEEE Trans. Autom. Contr.*, vol. 16, no. 1, pp. 12-22, Feb. 1971.
- [33] M. Morf, B. Dickinson, T. Kailath, and A. Vieira, "Efficient solution of covariance equations for linear prediction," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 25, no. 5, pp. 429-435, Oct. 1977.

- [34] G. E. P. Box and G. M. Jenkins, *Time Series Analysis: Forecasting and Control*, rev. ed. San Francisco, CA: Holden-Day, 1976.
- [35] D. J. Krause and D. Graupe, "Identification of autoregressive moving-average predictor models," in *Proc. 2nd Symp. on Nonlinear Estimation Theory and Its Applications*, H. W. Sorenson, chairman, pp. 174-182, 1971.
- [36] S. M. Kay, *Modern Spectral Estimation : Theory and Application*. Englewood Cliffs, NJ: Prentice-Hall, 1988.
- [37] H. Sakai and H. Tokumaru, "Statistical analysis of a spectral estimator for ARMA processes," *IEEE Trans. Autom. Contr.*, vol. 25, no. 1, pp. 122-124, Feb. 1980.
- [38] M. Kaveh and S. P. Bruzzone, "Statistical efficiency of correlation-based methods for ARMA spectral estimation," *IEE Proc.*, vol. 130, pt. F, no. 3, pp. 211-217, Apr. 1983.
- [39] T. J. Ulrych and R. W. Clayton, "Time series modelling and maximum entropy," *Phys. Earth Planet. Int.*, vol. 12, pp. 188-200, Aug. 1976.
- [40] E. Parzen, "Mathematical considerations in the estimation of spectra," *Technometrics*, vol. 3, pp. 167-190, 1961.
- [41] R. L. Anderson, "Distribution of the serial correlation coefficients," *Ann. Math. Statist.*, vol. 13, pp. 1-13, 1942.
- [42] M. S. Bartlett, "On the theoretical specification and sampling properties of autocorrelated time series," *J. Royal Statist. Soc. Suppl.*, vol. 8, pp. 27-41, 1946.
- [43] M. H. Quenouille, "Approximate tests of correlation in time series," *J. Royal Statist. Soc. Ser. B*, vol. 11, pp. 68-84, 1949.
- [44] L. D. Davisson, "The prediction error of stationary Gaussian time series of unknown covariance," *IEEE Trans. Inform. Theory*, vol. 11, pp. 527-532, 1965.
- [45] H. Akaike, "Fitting autoregressive models for prediction," *Ann. Inst. Statist. Math.*, vol. 21, pp. 243-247, 1969.

- [46] H. Akaike, "A new look at the statistical model identification," *IEEE Trans. Autom. Contr.*, vol. 19, no. 6, pp. 716-723, Dec. 1974.
- [47] H. Akaike, "Use of an information theoretic quantity for statistical model identification," in *Proc. 5th Hawaii Int. Conf. System Sci.*, pp. 99-101, 1972.
- [48] H. Akaike, "Information theory and an extension of the maximum likelihood principle," in *Proc. 2nd Int. Symp. Inform. Theory*, B. N. Petrov and F. Csaki, ed. Budapest: Akademiai Kiado, pp. 267-281, 1973.
- [49] H. Akaike, "A Bayesian analysis of the minimum AIC procedure," *Ann. Inst. Statist. Math. A*, vol. 30, no. 1, pp. 9-14, 1978.
- [50] H. Akaike, "A Bayesian extension of the minimum AIC procedure of autoregressive model fitting," *Biometrika*, vol. 66, pp. 237-242, 1979.
- [51] E. Parzen, "Some recent advances in time series modelling," *IEEE Trans. Autom. Contr.*, vol. 19, no. 6, pp. 723-730, Dec. 1974.
- [52] G. Schwarz, "Estimating the dimension of a model," *Ann. Statist.*, vol. 6, no. 2, pp. 461-464, 1978.
- [53] J. Rissanen, "Estimation of structure by minimum description length," *Circ., Syst., Signal Process.*, vol. 1, pp. 395-406, 1982.
- [54] P. Whittle, *Hypothesis Testing in Time Series Analysis*. PhD thesis, Uppsala, Sweden: Almqvist and Wiksell, 1951.
- [55] P. Whittle, "Tests of fit in time series," *Biometrika*, vol. 39, p. 309, 1952.
- [56] R. H. Jones, "Autoregression order selection," *Geophysics*, vol. 41, pp. 771-773, Aug. 1976.
- [57] J. G. Berryman, "Choice of operator length for maximum entropy spectral analysis," *Geophysics*, vol. 43, pp. 1384-1391, Dec. 1978.
- [58] H. Tong, "Autoregressive model fitting with noisy data by Akaike's information criterion," *IEEE Trans. Inform. Theory*, vol. 23, pp. 409-410, May 1977.

- [59] R. L. Kashyap, "Inconsistency of the AIC rule for estimating the order of autoregressive models," *IEEE Trans. Autom. Contr.*, vol. 25, no. 5, pp. 996-998, Oct. 1980.
- [60] R. Shibata, "Selection of the order of an autoregressive model by Akaike's information criterion," *Biometrika*, vol. 63, pp. 117-126, 1976.
- [61] J. C. Chow, "On the estimation of the order of a moving average process," *IEEE Trans. Autom. Contr.*, vol. 17, no. 3, pp. 386-387, June 1972.
- [62] J. C. Chow, "On estimating the orders of an autoregressive-moving average process with uncertain observations," *IEEE Trans. Autom. Contr.*, vol. 17, no. 5, pp. 707-709, Oct. 1972.
- [63] R. Kumaresan and D. W. Tufts, "Singular value decomposition and spectral analysis," *1st IEEE ASSP Workshop on Spectral Estimation*, McMasters Univ., Hamilton, Ontario, Canada, pp. 6.4.1-6.4.2, Aug. 1981.
- [64] D. W. Tufts and R. Kumaresan, "Estimation of frequencies of multiple sinusoids: making linear prediction behave like maximum likelihood," *Proc. IEEE*, vol. 70, no. 9, pp. 975-989, Sept. 1982.
- [65] J. A. Cadzow, "Spectral estimation: an overdetermined rational model equation approach," *Proc. IEEE*, vol. 70, no. 9, pp. 907-938, Sept. 1982.
- [66] G. H. Golub and W. Kahan, "Calculating the singular values and pseudo-inverse of a matrix," *J. SIAM Numer. Anal. Ser. B*, vol. 2, no. 2, pp. 205-224, 1965.
- [67] B. Friedlander, "Lattice filters for adaptive processing," *Proc. IEEE*, vol. 70, no. 8, pp. 829-867, Aug. 1982.
- [68] S. Treitel, P. R. Gutowski, and E. A. Robinson, "Empirical spectral analysis revisited," in *Topics in Numerical Analysis*, J. H. Miller, ed. New York: Academic, pp. 429-446, 1977.
- [69] J. Makhoul, "Linear prediction: a tutorial review," *Proc. IEEE*, vol. 63, no. 4, pp. 561-580, Apr. 1975.

- [70] G. M. Jenkins and D. G. Watts, *Spectral Analysis and its Applications*. San Francisco, CA: Holden-Day, 1968.
- [71] J. D. Markel and A. H. Gray, Jr., "Linear prediction of speech," in *Communications and Cybernetics*, vol. 12. Berlin: Springer-Verlag, 1976.
- [72] L. Ljung and T. Söderström, *Theory and Practice of Recursive Identification*. Cambridge, MA: MIT Press, 1983.

Chapter 3

Adaptive Filtering of WLI Fringe Patterns

3.1 Introduction

In Chapter 2 a general autoregressive-moving average (ARMA) model and an autoregressive (AR) approximation were proposed for describing the generation of the WLI system output.

In this chapter three popular FIR adaptive filtering algorithms will be introduced, whose coefficients are recursively updated as soon as new information becomes available.

Simulations will be carried out on a typical WLI system, in order to compare their performance with that of the batch filters described in Chapter 2.

A final section on numerical accuracy and stability ends the chapter.

3.2 Adaptive Finite Impulse Response Filters

With the filtering techniques described in Chapter 2 the coefficients of the filter are chosen by an algorithm from a given set of data to minimise the error between a desired response and the filter output. With the arrival of a new set of data the coefficients are automatically recomputed in order to adjust the centre frequency and bandwidth of the filter, if necessary, so as to keep the error down to a minimum. This is in stark contrast to a constant-coefficient filter design, where the designer first specifies a set of fixed filter characteristics, and then proceeds with an

implementation in analog or digital form, on the assumption that the statistics of future data will be the same as those upon which the design is based.

Nevertheless, in the literature the term "adaptive" is usually reserved for the sort of estimation in which the filter coefficients need to be updated within the same data set. The terms "recursive", "sequential", or "on-line" are also often used. Hence, the techniques of Chapter 2 will be referred to as "block", "batch", or "off-line" from now on.

Assuming for simplicity that the characteristics of the light source (mean wavelength, spectral width etc.), those of the optical system (refractive index of medium, angular alignment and physical properties of the mirrors, etc.), and the measurement conditions do not change over time, the statistics of the fringe pattern remain the same from scan to scan and one would be tempted to conclude that the measured signal is stationary. It is evident, though, that between the appearance of one pattern and the next there are periods in time (or in space, if an electronic scanner is used) when no interference occurs and what is observed and recorded is just noise; this situation is analogous to that encountered with electrocardiogram (ECG) signals, where the timing of the main R wave changes with the heart rate, and its centre frequency and bandwidth are very different from those of the signal between two R waves.

If stationarity can be rejected on this basis, the WLI output can at least be classified, like the ECG output, as cyclo-stationary, i.e., exhibiting statistical behaviour which is repetitive but not necessarily with a fixed period [1]. On this basis, if filtering was restricted to the interference region, the average of the least squares solutions of a block-processing filter over an increasing number of patterns would ultimately converge to the Wiener solution, as would the updated sample-to-sample solution of an adaptive filter that was processing one pattern after the other in a continuous manner.

However, both ECG and WLI outputs are non-stationary when a smaller time scale than beat-to-beat or peak-to-peak are considered. With the ECG, the centre frequency of the P and T waves is much smaller than that of the R wave in between; if a constant-coefficient filter is used with a high cut-off frequency, not enough noise may be filtered, making it difficult to extract the P and T waves. Reducing the cut-off frequency will reduce the noise content but distortion of the ECG may occur

as the R wave starts to be filtered [2]. An adaptive filter, on the other hand, may be able to track all of the P, R, and T waves if the degree of non-stationarity (DNS) is not too high, i.e., if the frequency variations are sufficiently slow relative to the sampling rate and the speed of adaptation of the filter coefficients. A block-processing filter will instead make a compromise and deliver a solution which tries to recover the P, R, and T waves contemporaneously.

A similar, although different, situation occurs with the WLI output considered over the interference region. Here, the centre frequency and bandwidth of the interference signal do not change, but the bandwidth of the output gradually increases moving away from the central fringe because, as the signal becomes weaker and weaker, the noise contribution to the output spectrum becomes larger and larger.

It was shown in [3] that the bandwidth of the Wiener filter depends not only on the bandwidth of the signal to be estimated but also on the SNR. As this decreases the frequency response of the filter, $H(f)$, shrinks so as to restrict the amount of noise passed through, according to the formula $H(f) = P_s(f)/(P_s(f) + P_n(f))$, where $P_s(f)$ and $P_n(f)$ are the spectra of signal and noise, respectively. A constant-coefficient filter has a frequency response which is fixed at the design stage and cannot be changed afterwards. A block-processing filter, instead, can handle scan-to-scan variations of the SNR (arising from changes in signal and/or noise power between scans) by adjusting its frequency response accordingly, if the approximate Wiener solution computed for the latest pattern is not averaged with the solutions from previous patterns but is used on its own. An adaptive filter can handle not only scan-to-scan but also intra-scan variations if these are not too fast, i.e., if the DNS, defined as the rate of change of the SNR along the fringe pattern, is not too high.

Two important advantages of adaptive algorithms over batch algorithms are fast operational speed and limited memory requirements. Instead of having to store the whole data set before the estimation phase can start, estimation can be undertaken while a new data sample is being acquired and the old data discarded.

Adaptive filtering of WLI fringe patterns is attractive not only in temporally scanned systems (due to the large number of samples to be processed) but also in spatially scanned systems, as the CCD clock can transform spatial information into a pulse train suitable for on-line processing.

Fig. 3.1 is a block diagram of a digital adaptive transversal finite impulse response (FIR) filter [4].

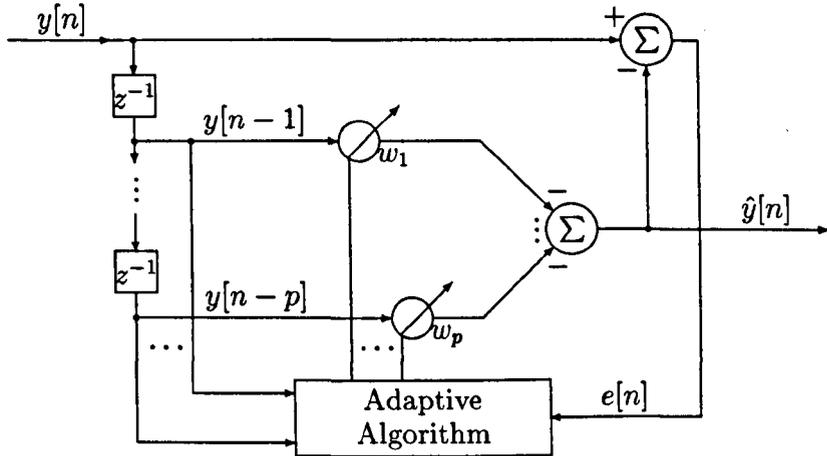


Fig. 3.1 FIR adaptive transversal filter.

At sample n the filter uses the input vector $\mathbf{y}[n-1]$ containing the previous p samples $-y[n-1], \dots, -y[n-p]$ from a tapped-delay line to produce the *predicted* output $\hat{y}[n] = \mathbf{y}^T[n-1]\mathbf{w}[n]$, where $\mathbf{w}[n]$ is the vector containing the filter coefficients w_1, \dots, w_p ; the error term $e[n]$ between the current sample $y[n]$ and its prediction $\hat{y}[n]$ is fed to an adaptive algorithm which updates $\mathbf{w}[n]$ in order to minimise a specified error function, usually the MSE criterion of Eq. 2.10.

A *filtered* output may be computed at this point using the coefficients just updated, before proceeding to the prediction for $y[n+1]$. Although the difference between the two estimates is in general negligible in environments requiring low adaptation rates, it may be worth considering the latter when dealing with a highly non-stationary process, where the filter has to respond quickly to the rapid time variations of the process parameters, although the computational cost increases by p multiplications and additions.

3.2.1 The Standard Least Mean Square (LMS) Algorithm

Rewriting Eq. 2.10 as

$$E\{e^2[n]\} = E\{(y[n] - \mathbf{y}^T[n-1]\mathbf{w}[n])^2\} \quad (3.1)$$

the MSE becomes a quadratic function of the filter weights, with a unique minimum found by setting the gradient

$$\frac{\delta E\{e^2[n]\}}{\delta \mathbf{w}} = -2E\{\mathbf{y}[n-1](y[n] - \mathbf{y}^T[n-1]\mathbf{w}[n])\} \quad (3.2)$$

to zero. This gives the Wiener-Hopf equations

$$\mathbf{R}_{y-1y-1} \mathbf{w} = \mathbf{R}_{yy-1} \quad (3.3)$$

where \mathbf{R}_{y-1y-1} is the autocorrelation matrix of the filter input vector and \mathbf{R}_{yy-1} is the cross-correlation vector between the current fringe sample $y[n]$ and the filter input vector.

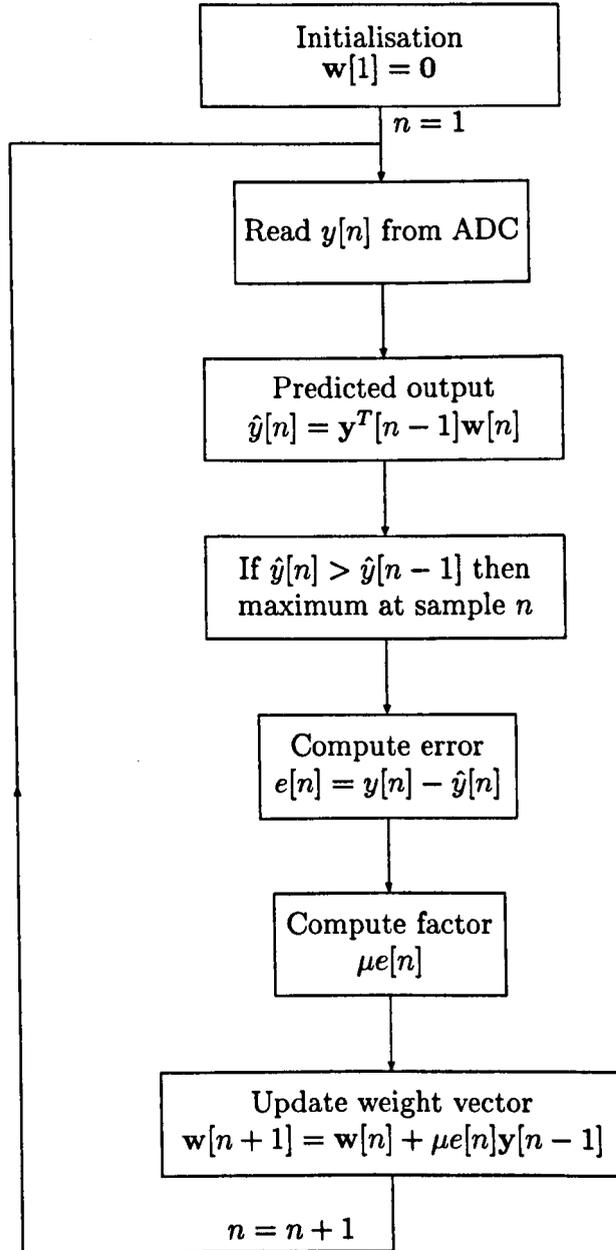


Fig. 3.2 Flowchart for on-line prediction of the WLI output with the LMS algorithm.

The LMS algorithm [5, 6] is a stochastic implementation of the method of steepest descent [7, 8] which is a gradient search technique in which the weights at $n + 1$

are equal to the weights at n plus a change proportional to the negative gradient. This amounts to descending along the p -dimensional error surface in a zig-zag fashion towards its bottom, where the Wiener solution lies. Choosing $\delta e^2[n]/\delta \mathbf{w}$ as an unbiased estimate of the unknown gradient the weight update equation becomes

$$\mathbf{w}[n+1] = \mathbf{w}[n] + \mu e[n] \mathbf{y}[n-1] \quad (3.4)$$

with initial condition $\mathbf{w}[1] = \mathbf{0}$. The constant μ is an amplification factor which controls the rate of descent or convergence.

A flowchart for on-line prediction of the fringe pattern with the LMS algorithm is shown in Fig. 3.2. At sample n the predicted output $\hat{y}[n]$ is formed, and if this is greater than the predicted output at sample $n-1$, n and $\hat{y}[n]$ are recorded and $n-1$ and $\hat{y}[n-1]$ are discarded. If the fringe position and intensity are sampled together by the analogue-to-digital converter (ADC), the sample number retained at the end of the scan gives the estimated central fringe position directly.

The total number of operations amounts to $2p+1$ multiplications and additions per iteration, plus $p-1$ shifts of the elements of the input vector.

The performance of the LMS is heavily dependent on the choice of the step-size μ , which sets a compromise between rate of convergence of the weights and excess MSE or misadjustment after convergence [9]. With μ small the descent towards the bottom of the error surface is slow but smooth, leaving a small amount of noise in the weights and a small gradient error after convergence. With μ large the rate of descent is faster but leaves a higher gradient error; if μ is too large the algorithm may also become unstable and diverge [10].

In a non-stationary environment where the statistics of the signal or system under examination may change continuously, the orientation and curvature of the error surface may move, and the weights must be able to track these variations after convergence. Again, a large μ value means faster adaptation, reducing a second contribution to the excess MSE, the lag error, caused by the time lag between the onset of the variation and the response of the weights.

3.2.2 The Recursive Least Squares (RLS) Algorithm

With the steepest descent method convergence is linear and independent of the starting approximation. With Newton's method, instead, quadratic convergence can be obtained when the iterates are getting close to the minimum [11].

Newton's method uses extra information from the curvature of the error function to adapt the step length at each iteration; computation and inversion of the second derivative matrix, or Hessian,

$$\frac{\delta^2 E\{e^2[n]\}}{\delta \mathbf{w} \delta \mathbf{w}^T} = \mathbf{R}_{y-1y-1} \quad (3.5)$$

which defines the curvature requires, however, $O(p^3)$ operations at each step.

The RLS algorithm [12, 13] is a Gauss-Newton procedure, where the exact inverse of the Hessian is substituted by an approximation $\mathbf{P}[n]$ propagated in time by means of the matrix inversion lemma [14], thus replacing the formal inversion with a simple scalar division. This makes it possible to contain the total number of operations to $O(p^2)$ at the expense of super-linear rather than quadratic convergence, and of propagation of round-off errors from one iteration to the next. Introducing a forgetting factor $\lambda < 1$ into the recursions allows to give greater importance to recent error terms than to older ones, effectively replacing the uniform LS error criterion in Eq. 2.20 with an exponentially weighted criterion of the form $\sum^n \lambda^{N-n} e^2[n]$.

A flowchart for on-line prediction of the fringe pattern using the RLS approach is shown in Fig. 3.3. In order to avoid having to compute the inverse of $\mathbf{X}^T \mathbf{X}$ as soon as this becomes invertible (i.e., at sample $n = 2p + 1$) the approximated inverse Hessian is generally set equal to $c\mathbf{I}$ at initialisation, where \mathbf{I} is the identity matrix and c is a positive constant larger than $100/\sigma_y^2$, with σ_y^2 being the variance of the data, so as to guarantee that $\mathbf{P}[n]$ remains positive-definite during repeated applications of the matrix inversion lemma [15].

If c is large the initial value of the time-varying gain vector $\mathbf{G}[n]$ will also be large, leading to fast but noisy convergence, whereas if c is small convergence will also be slow but smoother.

The factor λ allows to track slow time variations and hence reduce the lag error in a non-stationary environment after convergence, by keeping $\mathbf{G}[n]$ larger than in the case of no forgetting. Since the length of data effectively used during the estimation phase is reduced, this amplifies the accumulation and propagation of numerical errors, and with it the noisiness of the predictor weights. Hence, the compromise between weight and lag errors characteristic of the LMS algorithm applies also to the RLS algorithm, although the reason for the weight error is different.

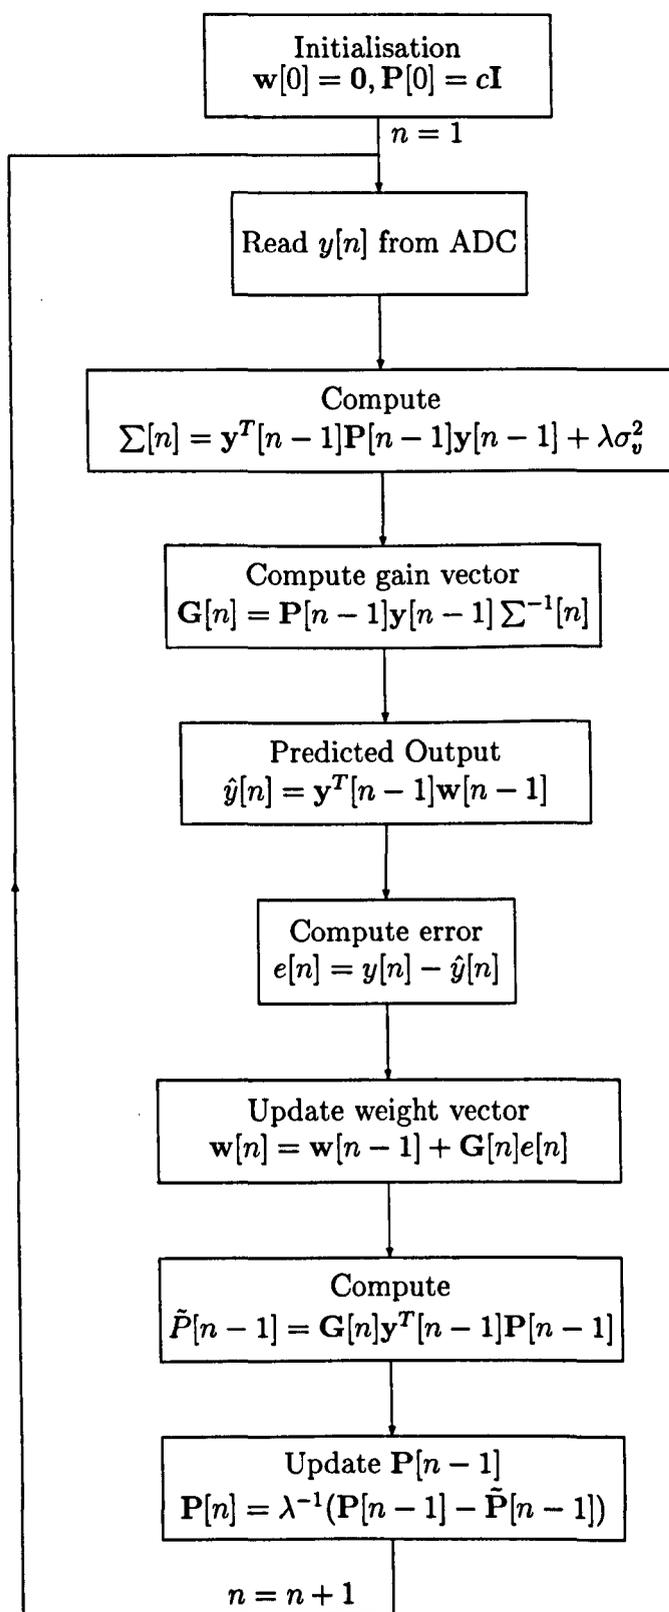


Fig. 3.3 Flowchart for on-line prediction of the WLI output with the RLS algorithm.

The parameter σ_v^2 is an estimate of the measurement noise variance, which if not available can be set to one.

A full derivation of the algorithm is found in, e.g. [16].

3.2.3 The Kalman Algorithm

The Wiener formulation leads to the optimum MSE predictor for stationary processes. Extension of the Wiener filter theory to non-stationary processes led to the development of the Kalman filter [17, 18].

The system under observation is represented by a general multi-dimensional dynamic model defined in *state-space* form by a process and a measurement equation

$$\mathbf{x}[n+1] = \Phi[n+1, n]\mathbf{x}[n] + \mathbf{v}_1[n] \quad (3.6)$$

$$\mathbf{y}[n] = \mathbf{C}[n]\mathbf{x}[n] + \mathbf{v}_2[n] \quad (3.7)$$

where $\mathbf{x}[n]$ and $\mathbf{y}[n]$ are the state and measurement vectors, $\Phi[n+1, n]$ is the state-transition matrix relating the states of the system at time n to the states at time $n+1$, $\mathbf{C}[n]$ is the measurement matrix relating the states at time n to the observations at time n , and $\mathbf{v}_1[n]$ and $\mathbf{v}_2[n]$ are the process and measurement noise vectors, usually assumed to be mutually independent, zero-mean and white, with covariance matrices

$$E\{\mathbf{v}_i[n]\mathbf{v}_i^T[k]\} = \begin{cases} \mathbf{Q}_i[n], & n = k & i = 1, 2 \\ \mathbf{0}, & n \neq k \end{cases} \quad (3.8)$$

A Kalman filter for the WLI system can be obtained by representing this as a dynamic AR model driven by the measurement noise

$$y[n] = \mathbf{y}^T[n-1]\mathbf{a}[n] + v[n] \quad (3.9)$$

with the non-stationarity of the output attributed to the variation of the system parameters with time, according to the first order Markov process [19]

$$a_i[n+1] = \phi_i[n]a_i[n] + \delta a_i[n] \quad i = 1, \dots, p \quad (3.10)$$

where the ϕ_i and δa_i terms account for deterministic and random fluctuations, respectively. Putting Eq. 3.10 into matrix form gives

$$\mathbf{a}[n+1] = \Phi[n]\mathbf{a}[n] + \delta\mathbf{a}[n] \quad (3.11)$$

Eqs. 3.11 and 3.9 are state-space equations for the time-varying AR process. If no prior information is available on the ϕ_i terms, these are usually set to a constant $\beta \leq 1$ for all i and n , simplifying the state-transition matrix to $\beta\mathbf{I}$.

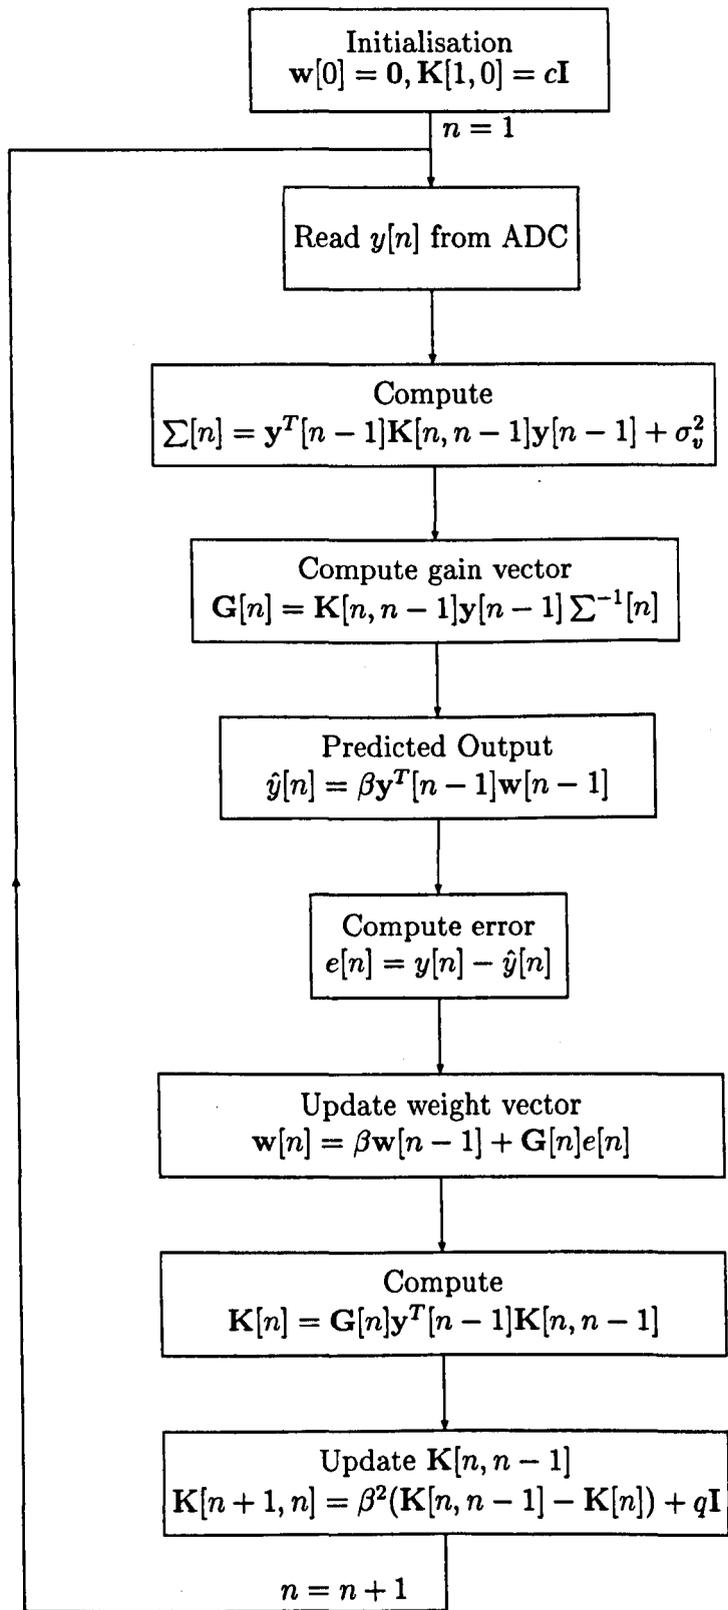


Fig. 3.4 Flowchart for on-line prediction of the WLI output with the Kalman filter.

The δa_i terms can also be assumed to be mutually independent, zero-mean white processes with constant variance q and diagonal covariance matrix $q\mathbf{I}$ [20].

A flowchart for on-line prediction of the fringe pattern with the Kalman filter is given in Fig. 3.4. $\mathbf{K}[n, n-1]$ is the $p \times p$ covariance matrix of the *predicted* estimate-error vector $\mathbf{a}[n] - \mathbf{w}[n-1]$. By minimising the trace of this matrix recursively the optimum estimate of the time-varying state vector is determined.

It can be shown that the RLS algorithm is a special case of the Kalman algorithm for the dynamic AR model considered here. The state-transition matrix and the process noise vector describe how fast the states of the system are expected to vary in time following deterministic and random changes within the system, a role which is exercised by λ in the RLS algorithm; similarly, the measurement noise vector describes the degree of confidence in the various components of the measurement vector, taking into account all sources of randomness at the measurement end, and corresponds to σ_v^2 in the RLS.

In fact, with $\lambda = 1$ and $\beta = 1, q = 0$ the two algorithms become identical for the one-dimensional case, i.e. the RLS method for a stationary AR process assumes a time-invariant state-space process equation. In a non-stationary environment the Kalman reduces to the RLS on setting $\beta = 1$ and $q \propto (1 - \lambda)/\lambda$ [16].

Having said this, the Kalman algorithm allows more flexibility as it allows to model deterministic and random variations of the system parameters independently; it is also possible to choose $\Phi[n]$ and $\delta\mathbf{a}[n]$ to have unequal entries, which may be convenient for describing different time variations of the parameters. Hence, one would expect the Kalman to perform at least as well as the RLS, if not better. Against this goes the higher computational cost, which reaches $O(p^3)$ operations per iteration.

3.3 Simulation Results

In order to obtain first-hand information on the sort of performance attainable by the adaptive filtering techniques presented in the previous section, some preliminary results from computer-generated fringe patterns will now be presented.

The simulations considered the output of an electronically scanned system with a CCD array detector consisting of 1024 pixels. To achieve a high degree of non-stationarity and thus provide a challenging situation for the algorithms, a short coherence length and fringe width were chosen, corresponding to 17 fringes and

12.7 pixels, respectively, so as to have a fast variation of the SNR along the fringe pattern. Scan-to-scan SNR fluctuations were not considered; instead, stationary white Gaussian noise was added at different levels, and for each level 500 scans were carried out with the central fringe in the middle of the scanning range. The filter weights were reset to zero at the beginning of each scan in order to test the convergence rate.

A typical output of an electronically scanned system in additive noise was shown in Fig. 1.4. The white noise assumption is a good approximation to reality in electronically scanned systems, although the noise at adjacent pixels of the CCD detector may be slightly correlated, adding a broad-band MA component. In temporally scanned systems, on the other hand, vibration of the moving scanner results in narrow-band AR noise centred at a frequency close to the mean frequency of the output pattern.

Fig. 3.5 compares the sub-fringe identification rate that can be achieved by batch and adaptive algorithms for a given filter order. An AR(10) process for the generation of the fringe pattern was assumed. The identification rate from direct fringe visibility is also shown.

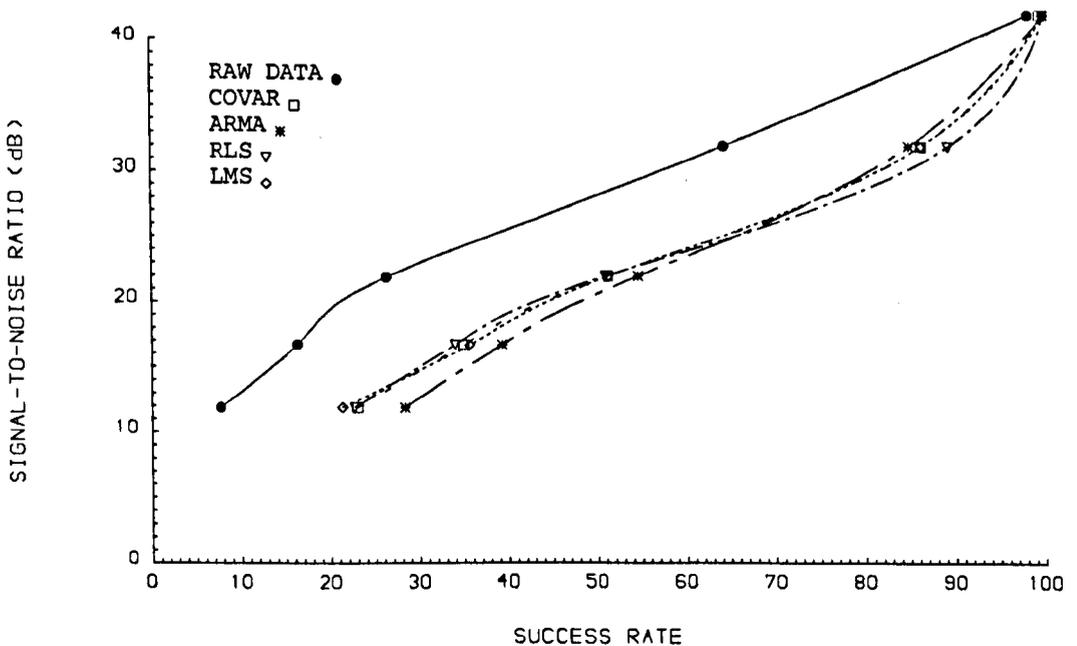


Fig. 3.5 Sub-fringe success rate with batch and adaptive algorithms out of 500 scans.

For the adaptive methods the *filtered* output of Section 3.2 was used as an estimate of the noise-free signal at pixels $n = 1, \dots, 1024$. With the *predicted* output

the performance dropped by 2-3 % on average.

The performance of the LMS filter was highly dependent on its step-size, with maximum identification rate obtained for smaller and smaller values as the SNR was decreased, suggesting that the contribution of the gradient error becomes more important as the SNR is reduced.

The noise variance σ_v^2 in the RLS filter was assumed unknown and set to one; the constant c was set to 10^4 , but smaller values produced identical results, possibly meaning that the transient phase was well over by the time the central portion of the fringe pattern was met. At high SNRs the best performance was obtained with λ between 0.9 and 1, but as the noise level was increased λ had to be kept very close to or equal to one, suggesting stationarity of the environment, although one has to remember that with $\lambda < 1$ the gain vector stays larger and induces a more noisy adaptation.

Inverting $\mathbf{X}^T\mathbf{X}$ at pixel $2p + 1$ by Gaussian elimination would allow to compute the exact LS solution, eliminating any transient phenomena. With this, the algorithm became unstable for SNRs above 50 dB, although the performance was only slightly affected for lower values.

With the fringe pattern at the centre of the scanning range the input data matrix contains only noise initially; hence, when the noise level is low the inversion of $\mathbf{X}^T\mathbf{X}$ becomes prone to round-off errors if single precision arithmetic is used, making it an ill-conditioned problem. Given that formal inversion requires $O(p^3)$ operations, it will not be considered further.

With the Kalman filter the best performance was achieved with the parameters q and β at zero and one, respectively, although small deviations from these values were not critical at high SNRs. The Kalman filter was not included in Fig. 3.5 because its success rate was identical to that of the RLS filter. In fact, with $q = 0$ and $\beta = 1$ the Kalman output was the same as the output from the RLS with $\lambda = 1$ down to 4 or 5 decimal points.

The square-root Kalman algorithm [21, 22] propagates the Cholesky square-root matrix $\mathbf{K}^{1/2}[n, n - 1]$ in place of $\mathbf{K}[n, n - 1]$, doubling the numerical precision of the Kalman filter and making it more robust to round-off errors. In fact, when computing in base i arithmetic with m significant digits, numerical difficulties can

be expected as the condition number ¹ χ of $\mathbf{K}[n, n - 1]$ approaches i^m . Since $\chi(\mathbf{K}^{1/2}[n, n - 1]) = (\chi(\mathbf{K}[n, n - 1]))^{1/2}$, using the square-root algorithm is equivalent to using double precision arithmetic.

The difference between the outputs from the standard and square-root implementations was within 0.2 % over the central portion of the fringe pattern, i.e. less than 8 % of the difference between the noise-free central and 1st-order fringe amplitudes, and clearly not enough to change the success rate. Running the standard version with double instead of single precision had the same effect. Therefore, the square-root Kalman algorithm will not be considered further, as it increases the already high computational cost ², although with longer data lengths its use in WLI may become more attractive as the chances of ill-conditioning increase.

As for the block-processing methods, the covariance algorithm for AR models of Section 2.6 was tested against the modified covariance method [25, 26], Burg method [27, 28], and the sub-optimum ARMA procedure of Section 2.5.

The modified covariance differs from the covariance method in that the sum of squares of the forward error defined in Eqs. 2.19 and 2.20 is minimised together with the sum of squares of the backward error ³

$$e_b[n] = y[n - p] - \hat{y}[n - p] = y[n - p] + \sum_{k=1}^p w_k y[n - p + k] \quad (3.12)$$

according to

$$\epsilon_{fb} = \sum_{n=p+1}^N (e_f^2[n] + e_b^2[n]) \quad (3.13)$$

so as to have twice the number of error terms involved in the minimisation task and reduce the variance of the estimates. Since the resultant autocorrelation matrix is the sum of two Toeplitz products, its inversion is possible with $O(p^2)$ operations [29, 30].

Whereas the modified covariance solves for the parameters $a_i, i = 1, \dots, p$ of the assumed AR(p) model directly, Burg method solves for the reflection or PARCOR [31] coefficients $a_i^i, i = 1, \dots, p$, of a lattice-equivalent model [32], with a_i^j being

¹An accepted measure of ill-conditioning of a matrix \mathbf{A} is its *condition number* χ , defined as the ratio between its largest and smallest singular values [23].

²With the increasing use of rotational coordinate, or CORDIC, arithmetic [24] in VLSI chips though, there may be no computational advantages in using square-root free algorithms [4].

³*Backward* in the sense that the prediction for the current data sample is a function of future samples only.

parameter a_i of an AR(j) model. Both modified covariance and Burg methods produced nearly identical coefficients and the same identification rate as the covariance method, hence they are not shown in Fig. 3.5.

Using the covariance method with backward instead of forward prediction increased the misidentification rate by 7-8 % at medium and low SNRs. The loss in performance was only due to the different data that the two filters use to produce their output; in fact, the two coefficient vectors were nearly identical, and swapping them between filters left the performance unchanged. This is in agreement with the result that the forward and backward solutions to the linear prediction problem may be different because the data read forward and backward are not in the same order, but equally valid because their statistical properties are the same [33].

For the ARMA sub-optimum procedure an ARMA(p, p) process was assumed and an AR($5p$) approximation was chosen for the MA process. The best results were obtained using an over-determined system of $t = 5p$ equations during the estimation of the AR parameters, with $p = 2$ at high and medium SNRs and $p = 4$ at low SNRs. Increasing p further the estimated error variance became negative (a sure sign of ill-conditioning) and the algorithm was exited before undertaking the MA analysis on the residual time series.

The hypothesis that the WLI system can be better represented by an ARMA model with identical AR and MA parameters was tested by setting the MA estimates equal to the AR estimates from the sub-optimum procedure. The identification rate was practically zero at all SNRs. This does not necessarily mean that the hypothesis is incorrect; what is certain is that if an optimum procedure is not adopted that solves for the two sets of parameters simultaneously, estimation errors will cause the two solution vectors to be different.

The batch techniques were also tested with the processing restricted to a central fraction of the fringe pattern. Apart from being attractive from a computational point of view, data-windowing reduces the degree of non-stationarity by restricting the SNR variation. However, the only beneficial effect of windowing was that it allowed to increase the model order for SNRs above 50 dB before ill-conditioning set in, an advantage which may not be of any practical use as far as central fringe identification is concerned, given that the success rate from the direct fringe visibility method is already as high as 100 % when employing a light source with a

typical coherence length equivalent to 17 fringes (Fig. 1.6). At lower SNRs the improvement in stationarity was apparently not enough to offset the degradation of the estimates brought about by the reduction in data availability.

3.4 Ill-Conditioning and Finite Precision Effects

Ill-conditioning is an approach towards singularity or non-invertibility [34]. There are two causes of ill-conditioning in the WLI output.

The first cause is the high degree of temporal dependence or correlation between the data points making up the zero and low-order fringes, due to the fact that the SNR is at its highest over this region, which induces near-exact relations or collinearity among adjacent columns of the data matrix, which are just shifted versions of one another and are therefore very similar. This results in uncertainty in determining how much each of the past p data points contributes separately to the present one, as there is little explanatory power unique to any of them; consequently, the estimating procedure does not have enough information to use in calculating the parameters, just as though it had a very small sample size. Any estimate based on little information cannot be held with much confidence, i.e. it will have a high variance [35].

The second cause is the insufficient dynamic range away from the interference region. Here, the interferometer output is at a constant level, and if the SNR is too high finite wordlength effects take over and result in loss of significance in the digital number representation of the computer. There is a danger that all data points are rounded to the same value, which is collinearity at its extreme.

The concept of ill-conditioning is closely related to the permanent input excitation condition. An input to a system is persistently exciting if it excites all natural frequencies or modes of the system [36]. Similarly, a sequence $y[n]$ is persistently exciting a filter with order p if the spectral density of the signal is non-zero at at least p equally spaced points [37]. In the time domain, the latter condition is satisfied if the $p \times p$ autocorrelation matrix $\mathbf{R}_{\mathbf{y}_{-1}\mathbf{y}_{-1}}$ of $y[n]$ is non-singular.

As the SNR increases the WLI interference signal becomes less persistently exciting because its bandwidth is reduced and the frequency content becomes zero at more and more points; away from the interference region the same effect can

be observed as signal quantisation takes its toll on the dynamic range and shrinks the signal down to its dc component. Similarly, increasing the filter order makes it less easy to maintain the spectral density non-zero requirement at p equally spaced frequencies.

Tables 3.1 and 3.2 show the condition numbers of the data matrix \mathbf{X} used by batch and adaptive algorithms during the computer simulations of the previous section, averaged over the 500 scans. In both cases there is a high correlation between condition numbers, SNR, and filter order.

The data matrix considered in Table 3.1 contained only the fringes within the coherence region. Expanding the window width to include the low-coherence and finally the non-coherence data, collinearity was reduced and the condition numbers with it, unless the noise level was so low that the data could not be represented accurately in the computer.

The condition numbers for Table 3.2 were calculated at several positions along the scanning range. Since the singular values of \mathbf{X} are the positive square roots of the eigenvalues of $\mathbf{X}^T\mathbf{X}$, the condition number at pixel n was computed as the square root of the eigenvalue ratio $\lambda_{max}/\lambda_{min}$ of the $p \times p$ ensemble-average autocorrelation matrix $\mathbf{R}_{y_{-1}y_{-1}} \equiv E\{\mathbf{y}[n-1]\mathbf{y}^T[n-1]\}$.

Just what is to be considered a large condition number can be an empirical matter. In econometrics, e.g., moderate to strong dependencies between the columns of the data matrix are associated with condition numbers between 30 and 100 [38].

For the error function to have a unique minimum it is essential that its Hessian be positive-definite ⁴. If the Hessian becomes negative-definite the paraboloid defining the error surface turns upside down, causing the algorithm to "blow-up". With the adaptive RLS algorithm in particular, the stabilising effect of $\mathbf{P}[0]$ is not permanent but decreases exponentially with time when $\lambda < 1$ [39]. If the data sequence ceases to be persistently exciting numerical problems will eventually arise, as by the matrix inversion lemma $\mathbf{P}[n]$ is updated as the difference of two matrices which tend to become near-singular [40]. The same problem has been observed with the $\mathbf{K}[n, n-1]$ matrix in the Kalman algorithm [41].

By effectively halving the condition number, square-root algorithms should se-

⁴The matrix \mathbf{R} is positive-definite if the quadratic form $\mathbf{x}^T\mathbf{R}\mathbf{x}$ is positive for all vectors $\mathbf{x} \neq \mathbf{0}$.

cure a search direction that points downwards when normal versions fail.

p	SNR = ∞	30 dB	20 dB	10 dB	0 dB
2	3.9	3.9	3.5	2.2	1.2
10	107.1	43.1	15.2	5.1	2.0
16	132.1	56.7	20.3	6.7	2.5
32	175.4	89.1	32.8	10.7	3.8

Table 3.1 Condition numbers of data matrix used by batch algorithms.

pixel ($p = 2$)	SNR = 30 dB	20 dB	10 dB	0 dB
250	1.0	1.0	1.0	1.0
404	11.6	3.8	1.5	1.1
436	21.2	6.8	2.3	1.2
512	32.8	10.4	3.4	1.4
($p = 10$)	30 dB	20 dB	10 dB	0 dB
250	1.0	1.0	1.0	1.0
404	21.9	7.0	2.4	1.2
436	38.1	12.1	3.9	1.6
512	65.5	20.7	6.6	2.3
($p = 16$)	30 dB	20 dB	10 dB	0 dB
250	1.0	1.0	1.0	1.0
404	27.1	8.6	2.9	1.3
436	47.6	15.1	4.9	1.8
512	85.6	27.1	8.6	2.9
($p = 32$)	30 dB	20 dB	10 dB	0 dB
250	1.0	1.0	1.0	1.0
404	34.2	10.9	3.6	1.5
436	61.4	19.4	6.2	2.2
512	123	38.9	12.3	4.0

Table 3.2 Condition numbers of data matrix used by adaptive algorithms.

3.5 Conclusion

In this chapter the LMS, RLS, and Kalman algorithms were used to update the parameters of an assumed time-varying AR model for the WLI output fringe pattern. It was shown that for this application the Kalman algorithm reduces to the RLS.

The difference between LMS and RLS is in the search direction, stochastic gradient in the former and Gauss-Newton in the latter. The Gauss-Newton direction gives near-quadratic convergence but results in an algorithm which requires at least $2p^2 + 6p$ adds and multiplies per iteration, whereas the stochastic gradient direction gives only linear convergence but the corresponding algorithm requires just $2p$ operations. Numerical problems from ill-conditioning of the data matrix also become more serious with the Gauss-Newton direction, as round-off errors accumulate and propagate from one iteration to the next.

From preliminary simulations it appears that the LMS is capable of performing as well as the RLS and batch algorithms of Chapter 2 when its step-size has been successfully optimised.

It also appears that the AR model assumption leads to an identification rate comparable to that obtainable with the more general ARMA assumption, except perhaps at low SNRs, where the absence of the MA term begins to be felt.

In the next chapter the comparison between LMS and RLS algorithms will be extended to cover convergence and tracking behaviours, and the influence of these on the identification rate will be examined together with their dependence on the filter design parameters and statistical properties of the data.

Bibliography

- [1] W. A. Gardner and L. E. Franks, "Characterization of cyclostationary random signal processes," *IEEE Trans. Inform. Theory*, vol. 21, no. 1, pp. 4-14, Jan. 1975.
- [2] M. S. Woolfson, X. B. Huang, and J. A. Crowe, "Time-varying Wiener filtering of the fetal ECG using the wavelet transform," *IEE Colloquium on Signal Processing in Cardiography*, pp. 11/1-11/6, Savoy Place, London, Mar. 1, 1995.
- [3] C. Anderson, E. Satorius, and J. Zeidler, "Adaptive enhancement of finite bandwidth signals in white Gaussian noise," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 31, no. 1, pp. 17-28, Feb. 1983.
- [4] S. Haykin, *Adaptive Filter Theory*, 2nd ed. Englewood Cliffs, NJ: Prentice-Hall, 1991.
- [5] B. Widrow and M. E. Hoff, Jr., "Adaptive switching circuits," *IRE WESCON Convention Record*, New York, pt. 4, pp. 96-104, 1960.
- [6] B. Widrow, "Adaptive filters," in *Aspects of Network in Systems Theory*, R. E. Kalman and N. DeClaris, eds. New York: Holt, Rinehart and Winston, pp. 563-587, 1971.
- [7] R. V. Southwell, *Relaxation Methods in Engineering Science*. New York: Oxford, 1940.
- [8] D. J. Wilde, *Optimum Seeking Methods*. Englewood Cliffs, NJ: Prentice-Hall, 1964.
- [9] B. Widrow, J. M. McCool, M. G. Larimore, and C. R. Johnson, Jr., "Stationary and nonstationary learning characteristics of the LMS adaptive filter," *Proc. IEEE*, vol. 64, no. 8, pp. 1151-1162, Aug. 1976.

- [10] B. Widrow, J. Glover, J. McCool, J. Kannitz, C. Williams, R. Hearn, J. Zeidler, E. Dong, Jr., and R. Goodin, "Adaptive noise cancelling: principles and applications", *Proc. IEEE*, vol. 63, no. 12, pp. 1692-1716, Dec. 1975.
- [11] R. Fletcher, *Practical Methods of Optimization, Vol. 1: Unconstrained Optimization*. New York: Wiley, 1980.
- [12] R. L. Plackett, "Some theorems in least squares," *Biometrika*, vol. 37, p. 149, 1950.
- [13] K. J. Åström, "Lectures on the identification problem - the least squares method," *Rep. 6806*, Division Autom. Contr., Lund Inst. Technol., Lund, Sweden, 1968.
- [14] M. Woodbury, "Inverting modified matrices," *Mem. Rep. 42*, Statistical Research Group, Princeton Univ., Princeton, NJ, 1950.
- [15] L. Ljung and T. Söderström, *Theory and Practice of Recursive Identification*. Cambridge, MA: MIT Press, 1983.
- [16] G. C. Goodwin and R. L. Payne, *Dynamic System Identification, Experimental Design and Data Analysis*. New York: Academic, 1977.
- [17] R. E. Kalman, "A new approach to linear filtering and prediction problems," *Trans. ASME. J. Basic Eng., Ser. D*, vol. 82, pp. 35-45, Mar. 1960.
- [18] R. E. Kalman and R. S. Bucy, "New results in linear filtering and prediction theory," *Trans. ASME. J. Basic Eng., Ser. D*, vol. 83, pp. 95-108, Mar. 1961.
- [19] Qi-tu Zhang and S. Haykin, "Tracking characteristics of the Kalman filter in a nonstationary environment for adaptive filter applications," *Proc. Int. Conf. Acoust., Speech, Signal Process.*, Boston, MA, pp. 671-674, Apr. 1983.
- [20] D. N. Godard, "Channel equalization using a Kalman filter for fast data transmission," *IBM J. Res. Develop.*, vol. 18, pp. 267-273, May 1974.
- [21] J. E. Potter, "New statistical formulas," *Memo. 40*, Instrum. Lab., MIT, Cambridge, MA, 1963.

- [22] P. G. Kaminski, A. E. Bryson, Jr., and S. F. Schmidt, "Discrete square-root filtering: a survey of current techniques," *IEEE Trans. Autom. Contr.*, vol. 16, no. 6, pp. 727-735, Dec. 1971.
- [23] A. S. Householder, *The Theory of Matrices in Numerical Analysis*. Waltham, MA: Blaisdell, 1964.
- [24] J. E. Volder, "The CORDIC trigonometric computing technique," *IRE Trans. Electron. Comput.*, vol. 8, no. 3, pp. 330-334, Sept. 1959.
- [25] T. J. Ulrych and R. W. Clayton, "Time series modelling and maximum entropy," *Phys. Earth Planet. Int.*, vol. 12, pp. 188-200, Aug. 1976.
- [26] A. H. Nuttall, "Spectral analysis of a univariate process with bad data points via maximum entropy and linear predictive techniques," *Tech. Rep. TR-5303*, Naval Underwater Systems Center, New London, CT, Mar. 1976.
- [27] J. P. Burg, "Maximum entropy spectral estimation," in *Proc. 37th Ann. Int. Meet. Soc. Explor. Geophys.*, Oklahoma City, OK, Dec. 1967.
- [28] J. P. Burg, "A new analysis technique for time series data," *NATO Adv. Study Inst. Signal Process.*, Enschede, Netherlands, pp. 12-23, Aug. 1968.
- [29] S. L. Marple, Jr., "A new autoregressive spectrum analysis algorithm," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 28, no. 4, pp. 441-454, Aug. 1980.
- [30] S. L. Marple, Jr., "Efficient least squares FIR system identification," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 29, no. 1, pp. 62-73, Feb. 1981.
- [31] J. D. Markel and A. H. Gray, Jr., "Linear Prediction of Speech," in *Communications and Cybernetics*, vol. 12. Berlin: Springer-Verlag, 1976.
- [32] F. Itakura and S. Saito, "Digital filtering techniques for speech analysis and synthesis," *Proc. 7th Int. Conf. Acoust.*, Budapest, Hungary, pp. 261-264, 1971.
- [33] C. W. Therrien, *Discrete Random Signals and Statistical Signal Processing*. Englewood Cliffs, NJ: Prentice-Hall, 1992.
- [34] A. S. Deif, *Advanced Matrix Theory for Scientists and Engineers*, 2nd. ed. Montreux, Switzerland: Gordon and Breach, 1991.

- [35] P. Kennedy, *A Guide to Econometrics*. Oxford: Robertson, 1979.
- [36] K. J. Åström and T. Bohlin, "Numerical identification of linear dynamic systems from normal operating records," in *Proc. IFAC Symp. on Theory of Self Adapting Control Systems*. New York: Plenum, Sept. 1965.
- [37] L. Ljung, "Characterization of the concept of 'persistently exciting' in the frequency domain," *Report 7119*, Division Autom. Contr., Lund Inst. Technol., Lund, Sweden, 1971.
- [38] D. A. Belsley, E. Kuh, and R. E. Welsch, *Regression Diagnostics: Identifying Influential Data and Sources of Collinearity*. New York: Wiley, 1980.
- [39] J. M. Cioffi and T. Kailath, "Fast, recursive-least-squares transversal filters for adaptive filtering," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 32, no. 2, pp. 304-337, Apr. 1984.
- [40] M. H. Verhaegen and P. Van Dooren, "Numerical aspects of different Kalman filter implementations," *IEEE Trans. Autom. Contr.*, vol. 31, no. 10, pp. 907-917, Oct. 1986.
- [41] F. M. Hsu, "Square-root Kalman filtering for high-speed data received over fading dispersive HF channels," *IEEE Trans. Inform. Theory*, vol. 28, no. 5, pp. 753-763, Sept. 1982.

Chapter 4

Convergence and Tracking Problems in WLI Filtering

4.1 Introduction

In the previous chapter adaptive algorithms were seen to be capable of achieving the same identification rate as batch algorithms, provided the step-size parameter was chosen appropriately.

This chapter concentrates on adaptive algorithms for the obvious advantages they offer (see Section 3.2), and especially on the LMS because of its simplicity of operation and fast computational speed, which are the reasons for it having become the benchmark against which all other adaptive algorithms are tested.

The sensitivity of the success rate to the step-size μ will be linked to the convergence and tracking properties of the algorithm and to the nature of the data at hand, and a comparison with the RLS algorithm will be made throughout using both simulated and experimental data.

4.2 Identification Rate and Filter Parameters

Fig. 4.1 shows that there exists an inverse relationship between noise level and optimum step-size μ_0 for maximum identification with a given filter order. The central fringe was kept in the middle of the CCD array throughout the simulations.

Fig. 4.2 shows that there exists a direct relationship between filter order and

success rate, but an inverse one between filter order and optimum step-size.

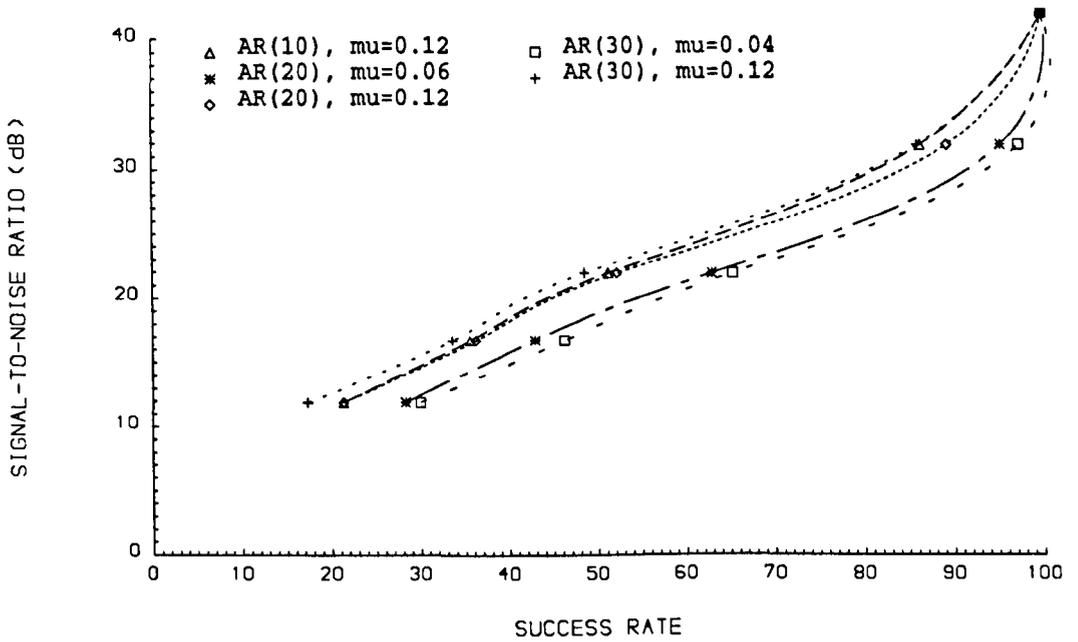


Fig. 4.1 Sub-fringe success rate (out of 500 trials) with the LMS algorithm.

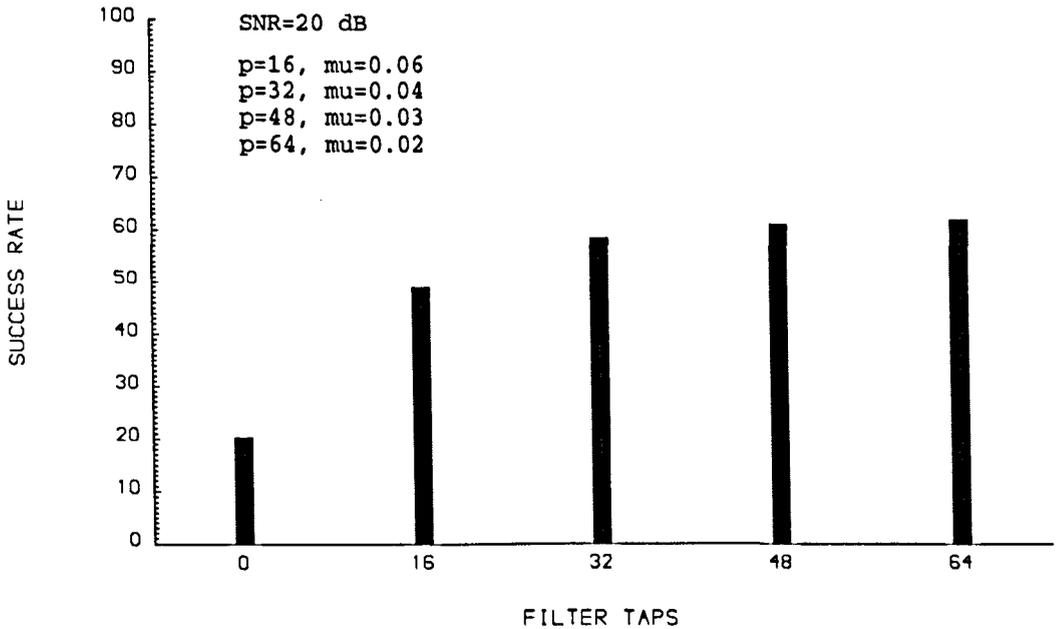


Fig. 4.2 Maximum success rate (out of 500 trials) against filter order at 20 dB.

As shown in Fig. 4.2, the identification rate improved by an absolute 9% (13%) when the filter order p was increased from 16 to 32 (64). For comparison, the success

rate with the covariance method went up 3 % (down 2 %). Just by increasing the filter order it is thus possible to beat the batch schemes, assuming μ_0 is available.

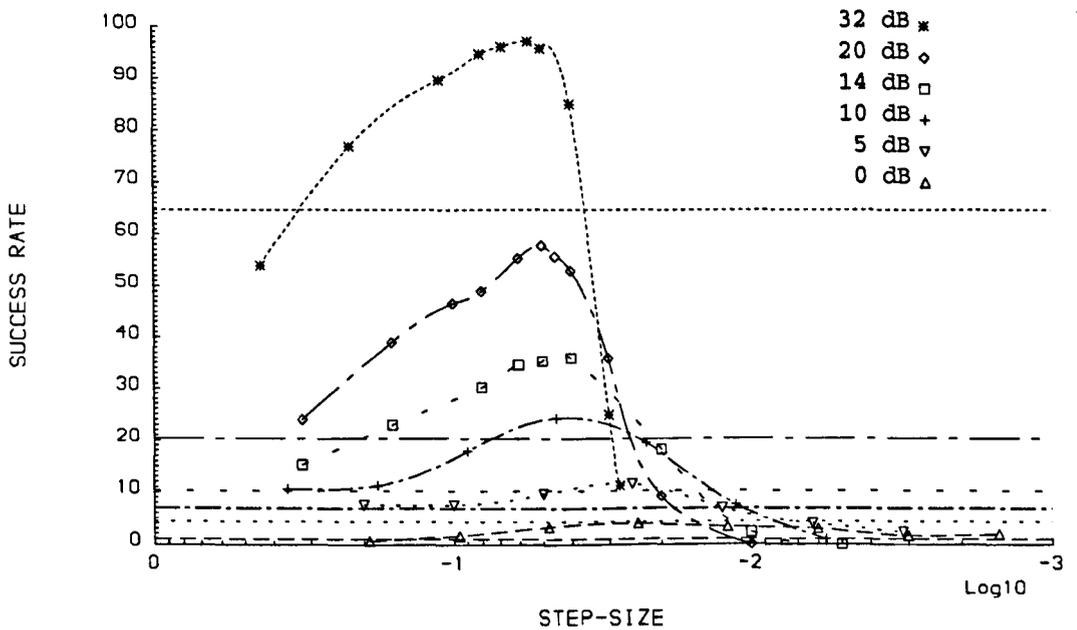


Fig. 4.3 Success rate (out of 500 trials) against μ and SNR, with $p = 32$. The horizontal lines refer to the direct visibility method.

Fig. 4.3 shows the identification rate as a function of the step-size for six values of the SNR in the range 0-32 dB.

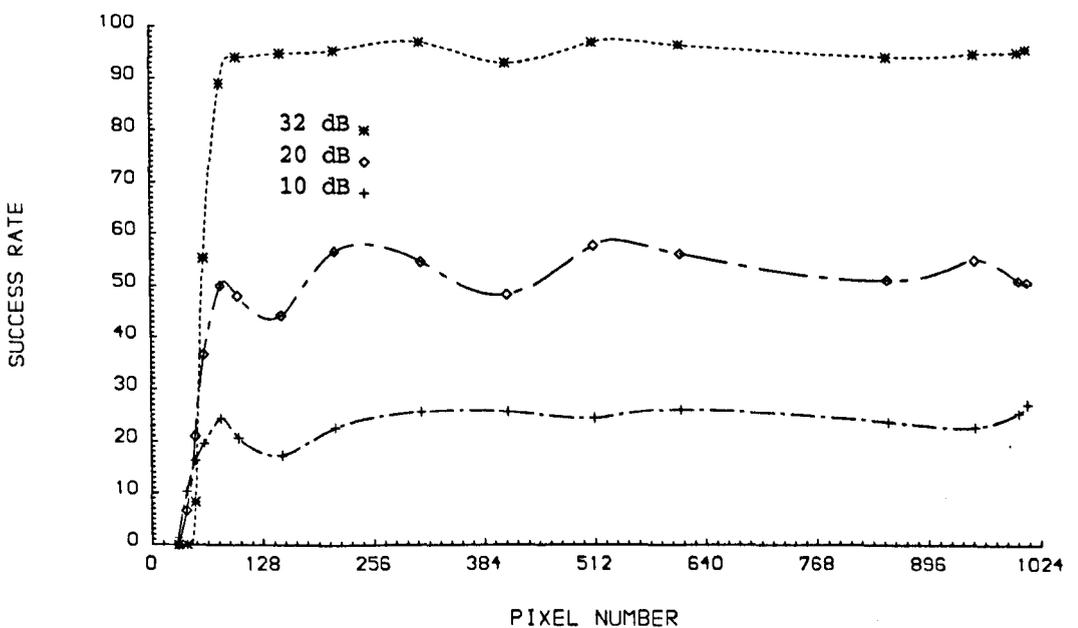


Fig. 4.4 Success rate (out of 300 trials) with the LMS algorithm against central fringe position, with $p = 32$ and μ optimised for the central pixel on the CCD array.

The sub-fringe identification rate was between one-third and one-half of the fringe identification rate for SNRs below 10 dB, otherwise the two rates were equal, as opposed to the 26 dB limit of the direct fringe visibility method (see Section 1.7).

Whereas with the latter method the rate did not change when the central fringe moved to the bottom end of the scanning range, that of adaptive algorithms was badly affected, as shown in Fig. 4.4 (fluctuations at pixels away from the bottom of the range were only caused by the randomness of the local noise amplitude, and should die out as the number of trials increases).

A set of simulations was performed with the central fringe moving continuously from pixel 262 through to pixel 761 during 500 consecutive scans. The identification rate of the direct visibility and covariance methods remained the same in the two halves of the scanning range separated by the central pixel, whereas that of the LMS was comparable to that in Fig. 4.2 only in the second half, being more than 10 % lower in the first half. Hence, optimising μ when the central fringe is in the middle of the array does not guarantee the best performance at other positions.

Another set of simulations was performed with the central fringe fixed at pixel 100 and a SNR of 20 dB. The optimum step-size was still an inverse function of the filter order but was on average 2.5-3 times higher than when the central fringe was kept in the middle of the array. Even then, the identification rate shown in Fig. 4.2 dropped by a relative 13 % (18 %) with $p = 16$ (32 or 64). With the covariance method the rate remained at the previous level with $p = 16$ and increased by as much as 6 % with $p = 32$ and 64. The rate with the direct visibility method also increased by 4 %, meaning that the performance of the LMS should not have suffered but, if anything, should have improved as a result of the fluctuations in the noise amplitude.

With the step-size optimised at pixel 100 and the central fringe at the centre of the array, the success rate also dropped by 8-9 % (15 %) with $p = 16$ or 64 (32).

In conclusion, the performance of the LMS filter in a practical setting will be influenced not only by the algorithm parameters p and μ but also by the operating range. A large μ accelerates convergence and may help to improve the performance when the operating range includes the lower portion of the scanning range, although a lower value may lead to smaller misadjustment over the central portion.

4.3 Time-Evolution of MSE and Filter Weights

Figs. 4.5-4.7 show the square of the LMS algorithm error signal $e[n]$ (as from flowchart in Fig. 3.2) during convergence and tracking, averaged over 500 consecutive scans with the central fringe in the middle of the array detector.

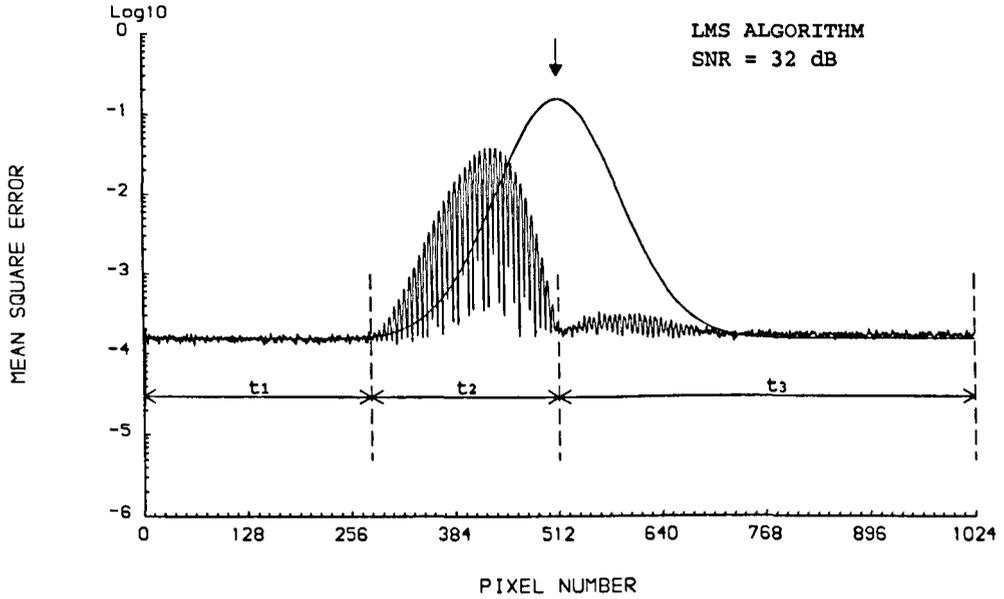


Fig. 4.5 LMS MSE across the CCD array with a SNR of 32 dB. The noise variance was -3.8 on a log scale. The superimposed fringe visibility profile is not to scale.

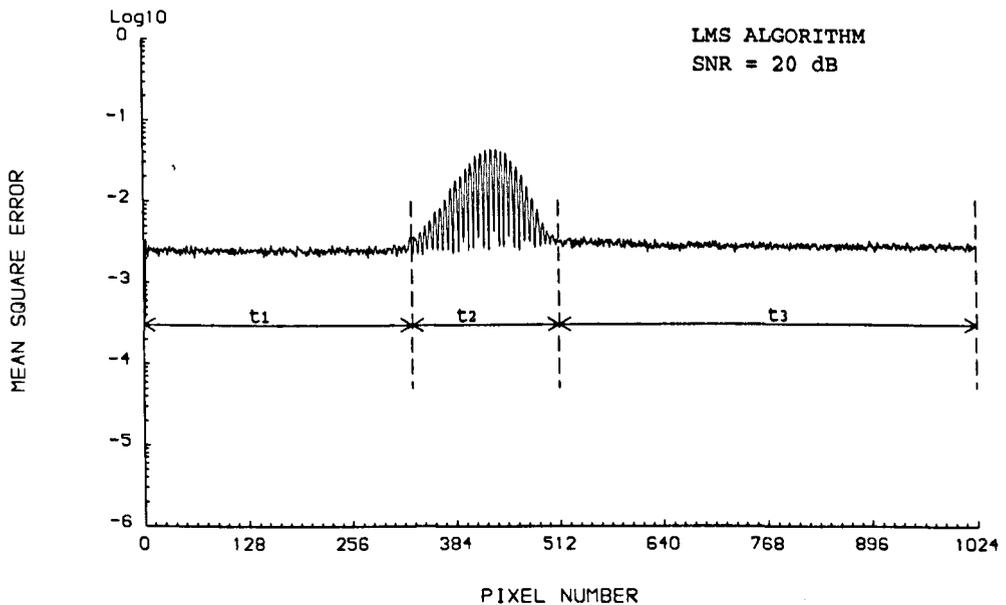


Fig. 4.6 LMS MSE with a SNR of 20 dB. The noise variance was -2.6 on a log scale.

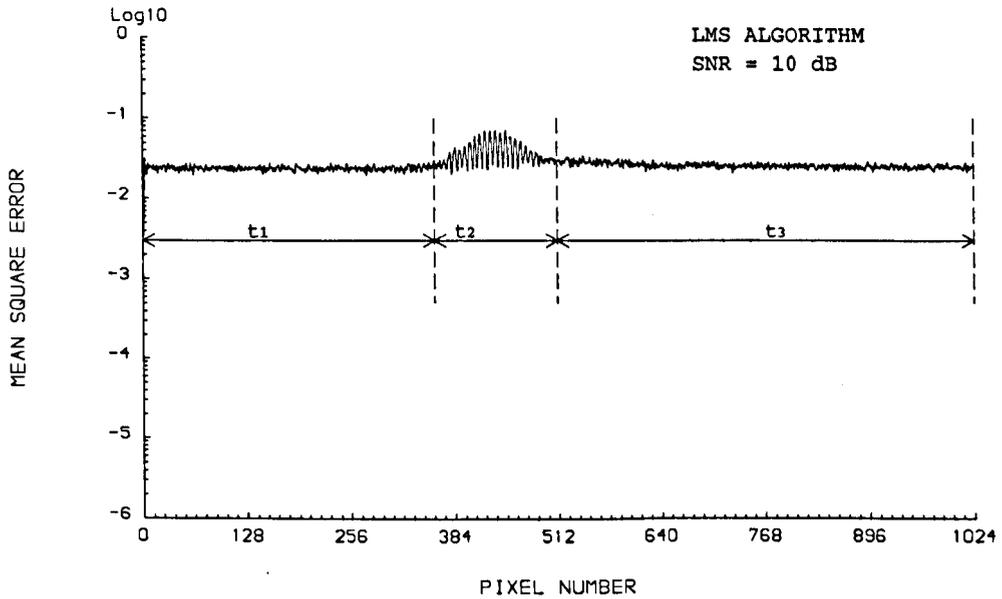


Fig. 4.7 LMS MSE with a SNR of 10 dB. The noise variance was -1.6 on a log scale.

With the amplitude of the central fringe set arbitrarily at 0.5, a noise level of 2.5 % (32 dB SNR) corresponds to a noise variance of 1.56×10^{-4} , or -3.8 on a log scale; similarly, noise levels of 10 % (20 dB) and 31 % (10 dB) correspond to noise variances of 2.5×10^{-3} and 2.5×10^{-2} , or -2.6 and -1.6 on a log scale, respectively. The step-size was optimised for each SNR but the filter order was kept fixed at sixteen.

The region delimited by t_1 in Figs. 4.5-4.7 corresponds to that portion of the array where no interference takes place and the normalised output signal from the detector, $y[n]$, consists solely of system and measurement noise $v[n]$. Here, the weights remain at zero because the algorithm cannot find any sufficiently strong correlation between successive samples.

Let's consider the updating of the weight vector \mathbf{w} in Eq. 3.4. In the mean,

$$E\{\mathbf{w}[n+1]\} = E\{\mathbf{w}[n]\} + \mu E\{e[n]\mathbf{y}[n-1]\} \quad (4.1)$$

At the beginning of the scan, the predictor output $\hat{y}[n]$ is zero because \mathbf{w} is zero, and the error signal $e[n]$ is thus equal to $y[n]$, which contains no useful signal but only $v[n]$. As the scan proceeds, the filter input vector $\mathbf{y}[n-1]$ becomes filled with past $v[n]$ values, and the last term on the r.h.s of Eq. 4.1 remains at zero because $v[n]$

is uncorrelated with its past. Hence, there is no weight adaptation, the predictor output stays at zero, and the MSE remains equal to its minimum possible value, the noise variance.

The region delimited by t_2 is where convergence takes place, and includes the effective fringe pattern up to the central fringe; the width of t_2 shrinks in passing from Fig. 4.5 to Fig 4.7 because as the noise level increases more of the small fringes remain buried under the noise. In this region, the weights start adapting since both error signal and input vector contain signal components, but the MSE initially increases with the signal level because of the inertia of the weights, which causes $\hat{y}[n]$ to lag behind $y[n]$. As convergence proceeds, however, the error signal starts to decrease, although the signal level is still increasing. An optimum value of the step-size will ensure that convergence is completed just before the central fringe is met, so that all is left to do after that is to track possible movements of the error surface caused by the non-stationarity of the fringe pattern.

Smaller values may extend the convergence region beyond the central fringe, while larger values may upset the balance between gradient and lag errors during tracking, producing a higher misadjustment over the fringes.

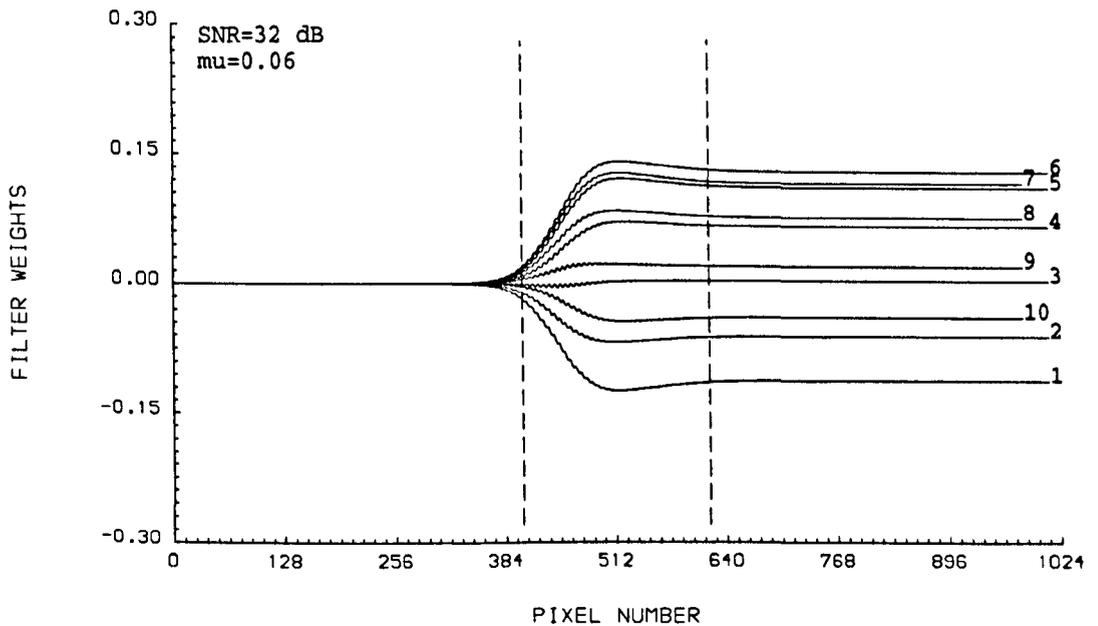


Fig. 4.8 LMS weights with a SNR of 32 dB. The dashed lines delimit the coherence region.

When the fringes are left behind, adaptation continues because past noise terms remain common to error signal and input vector in Eq. 4.1. Given sufficient time

the weights will return to zero, since in

$$E\{e^2[n]\} = E\{y^2[n]\} + E\{\hat{y}^2[n]\} - 2E\{y[n]\hat{y}[n]\} \quad (4.2)$$

the last term on the r.h.s. becomes zero and the only way to reduce the MSE further is by shrinking the second term.

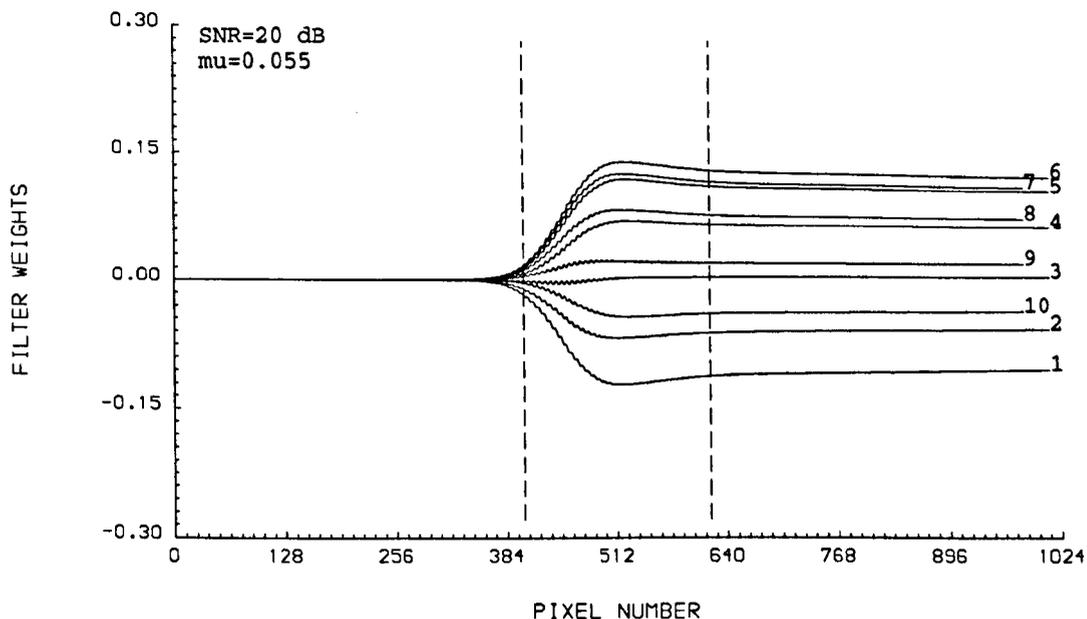


Fig. 4.9 LMS weights with a SNR of 20 dB.

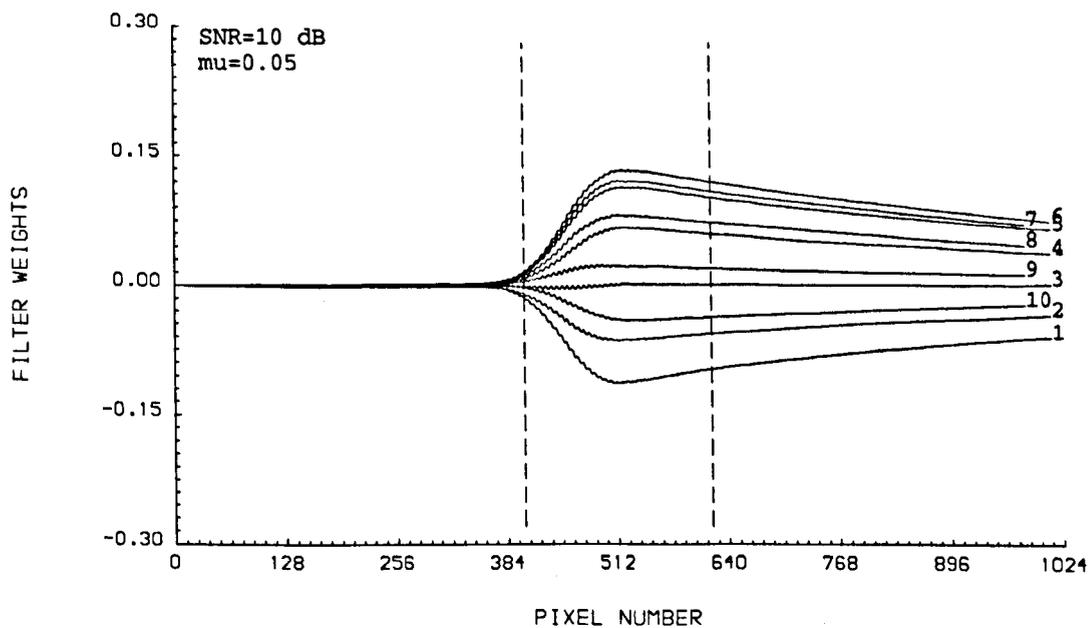


Fig. 4.10 LMS weights with a SNR of 10 dB.

Figs. 4.8-4.10 show the behaviour of the first ten weights during convergence and tracking. The μ_0 values from Fig. 4.3 were used. Figs. 4.11 and 4.12 show MSE and weight updating with the RLS filter; λ was set to one and c to forty times less than the minimum recommended in Section 3.2.2 (with σ_y^2 equal to 2.5×10^{-3} in the region delimited by t_1 , the minimum would be 40×10^3).

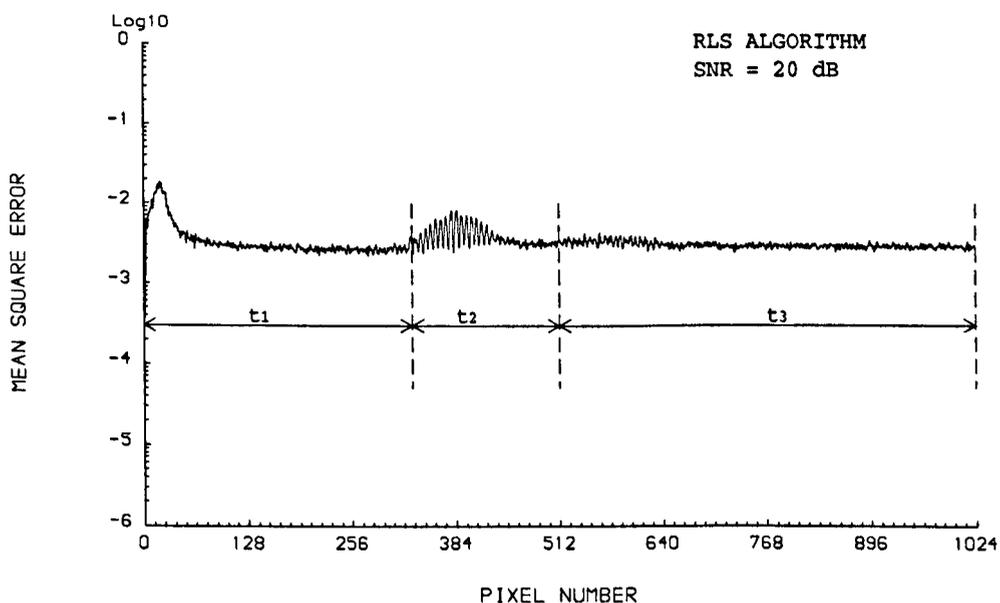


Fig. 4.11 RLS MSE with a SNR of 20 dB. The noise variance was -2.6 on a log scale.

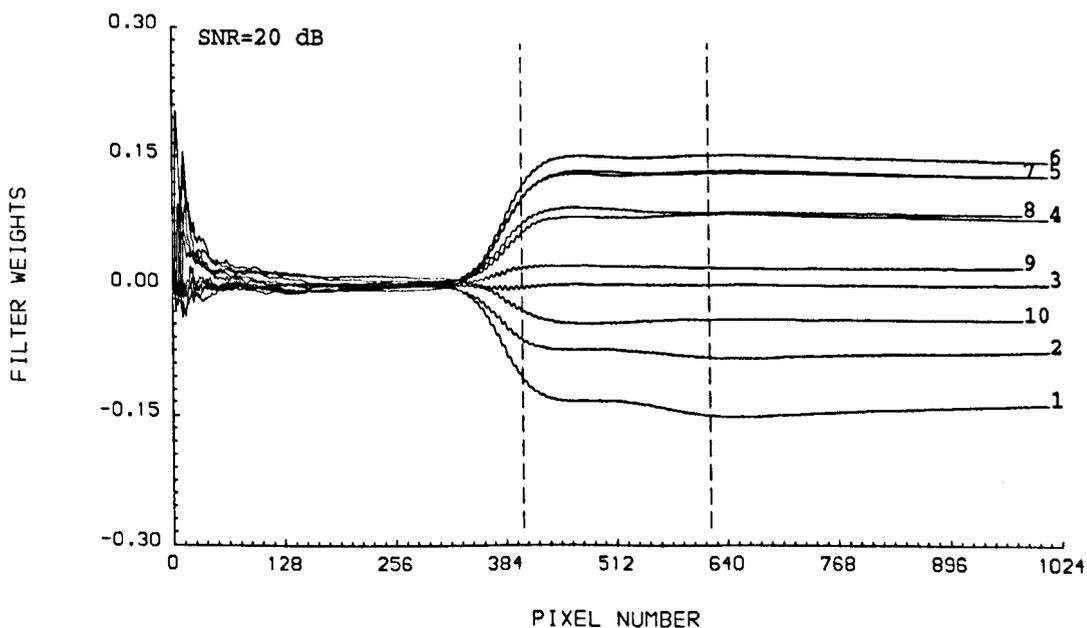


Fig. 4.12 RLS weights with a SNR of 20 dB.

Even though this choice of parameters is far from ideal when speed of adaptation is crucial, it can be seen by comparing with Figs. 4.6 and 4.9 that the RLS algorithm is far superior to the LMS in terms of convergence.

With the central fringe at pixel 100 and both μ and λ optimised, however, the identification rate of the RLS was only as high as that of the LMS with $p = 16$, and 7 % worse with $p = 32$. One reason for this may be that λ had to be reduced to 0.98 in order to make the algorithm converge in time, thus increasing the noisiness in the weights. Although the LMS suffers the same disadvantage when μ is increased, it may be that the injection of extra noise in the weights is a more serious problem with the RLS than with the LMS. This would explain why λ kept at one guaranteed the highest identification rate in the simulations of Section 3.3.

Another reason for failing to beat the LMS may be due to the influence of $P[0]$ at the lower end of the scanning range, clearly visible in Figs. 4.11 and 4.12. Although the effect of $P[0]$ decays quickly with time as λ takes over the updating role [1], the fact remains that the choice of c affects the behaviour of the RLS during this transient period. With a SNR of 20 dB, for example, central fringe and noise amplitude are in the ratio 1:10, so their variances are in the ratio 1:100; this means that the variance of the data, σ_y^2 , when the RLS comes into operation is about 100 times higher when the central fringe is at the bottom end of the range than when it is in the middle. It follows that a c value which was relatively low for the mid-range may now be too high, causing wild fluctuations in the filter weights and in the MSE before the change-over between $P[0]$ and λ takes place.

In fact, fluctuations in the MSE during the transient period were seen to increase with c . Thus, faster convergence alone is not a sufficient condition for higher identification at the bottom end of the CCD array, as the noisiness of adaptation can play a major role.

4.4 Time-Evolution of Filter Output

The fringe pattern in Fig. 4.13 is one of a set of measurements performed in the laboratory using a temporally-scanned WLI system. The light source was a LED with a nominal wavelength of 840 nm and a spectral half-width of 50 nm, corresponding to a coherence length of 14.1×10^4 nm or approximately 17 fringes.

The light beam was collimated by a $10 \times$ objective lens before being divided into two by the beam splitter of a Michelson interferometer, whose mirrors were fixed to piezo-electric transducers incorporating capacitive displacement sensors with a maximum range of $\pm 50 \mu m$. After recombination of the beams at the beam splitter, the interference signal was detected with a PIN photodiode/amplifier and displayed on an oscilloscope, before being recorded with 12 bits accuracy on a floppy disk via the serial port of a digital storage adaptor connected to the oscilloscope and a personal computer. The output from the displacement sensors was also stored digitally on the computer, but with 14 bits accuracy, so as to give an ultimate path difference resolution of $100/2^{14} \mu m$, or $6 nm$.

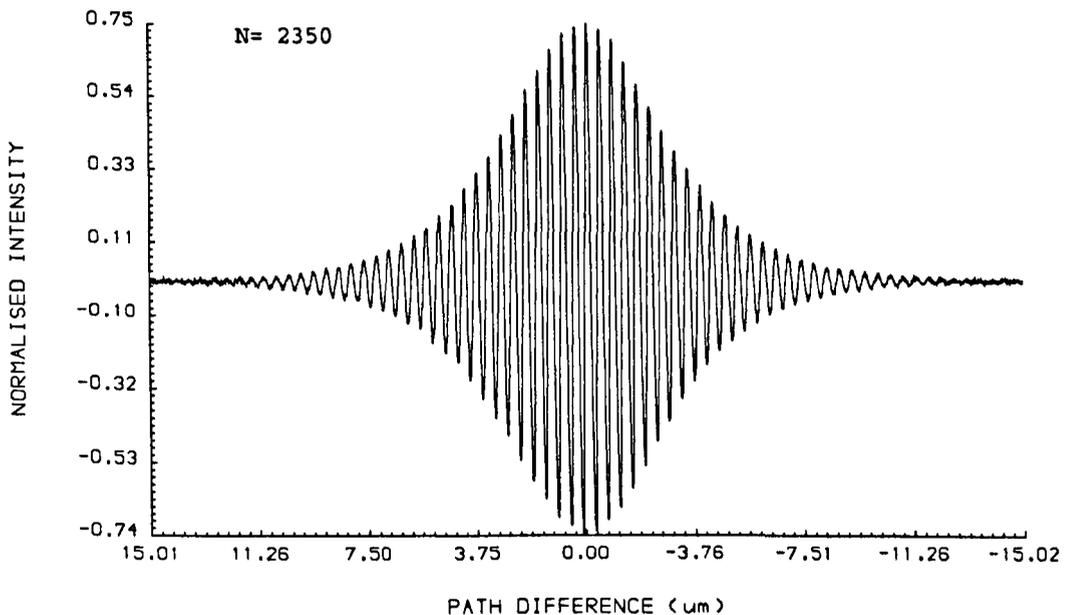


Fig. 4.13 Experimentally obtained white-light fringe pattern using a broad-band source and a mechanical scanner (after pre-processing).

The two mirrors were initially balanced to give zero OPD, and one of them was then moved to create an OPD much larger than the coherence length. The second mirror was at this point scanned through its range in steps of $\approx 13 nm$, and the output intensity was detected by the photodiode. To increase the noise level, a white-noise generator was attached to the parallel port of the storage device.

Due to vibrations induced by the movement of the scanner and the surroundings, the number of samples at the end of each scan was higher than expected. To ensure the monotonicity of the path difference, redundant information was discarded

through the use of a simple pre-processing operation incorporated into the adaptive filtering schemes, which consisted of skipping a sample if the path difference from the displacement sensors was greater than that at the previous sample.

In Fig. 4.13 only the path difference between $\pm 15 \mu m$ is shown. Over the full scanning range from $40 \mu m$ to $-40 \mu m$ a slow baseline drift was noticeable to the naked eye, possibly caused by misalignment between the optical components or by a temperature rise in the detectors during the two-minute measurement period, although the ambient temperature stayed at $25^{\circ}C$ throughout.

Comparing the signal power at 0 and $\pm 40 \mu m$ gave an estimate for the SNR at the central fringe of nearly 40 dB, a slightly higher figure than that obtained from the error variance estimated by the covariance method on the first few low-order fringes. The fringe width and coherence length of the source were found by inspection to correspond to ≈ 33.4 pixels ($0.43 \mu m$) and 17 fringes, respectively.

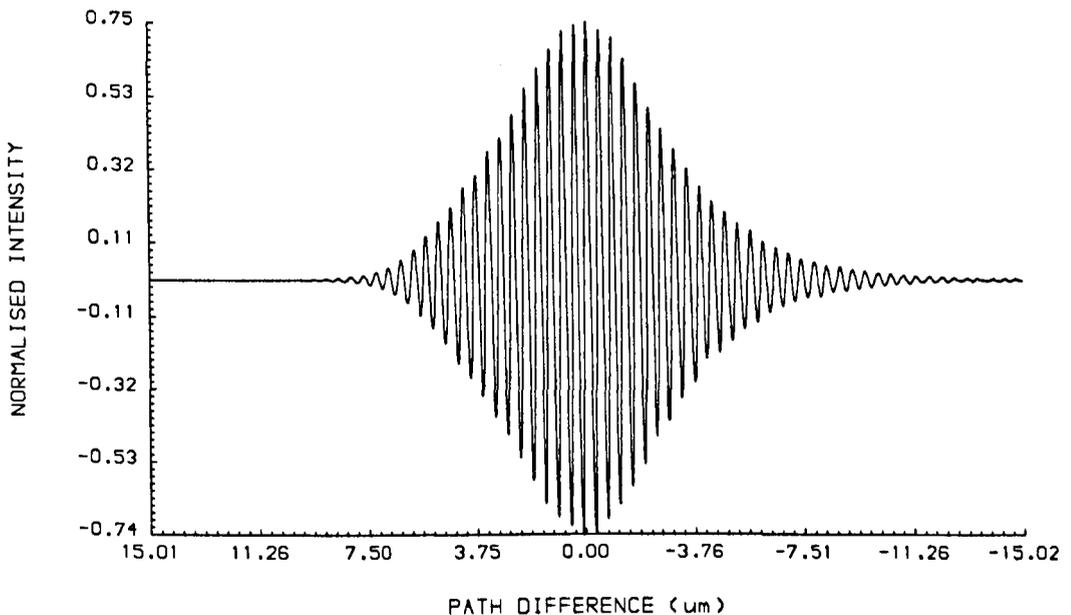


Fig. 4.14 The fringe pattern in Fig. 4.13 filtered by the LMS with $p = 32$ and $\mu = 0.1$.

Figs. 4.14 and 4.15 show the output of the LMS filter with the step-size μ equal to 0.1 and 0.01, respectively. Adaptation of the weights starts when the path difference is about $8.3 \mu m$ in the first case, but is delayed until $6.7 \mu m$ in the second case. The difference in shape of the two patterns clearly demonstrates that as the step-size is reduced convergence may become a problem, even though a relatively

large number of samples are available before the central fringe is met.

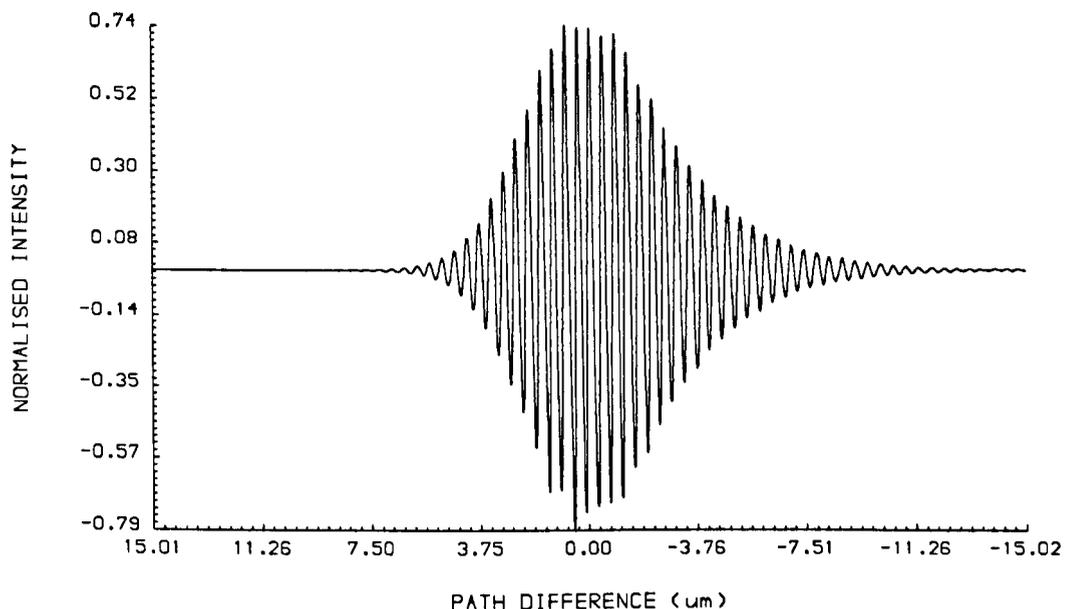


Fig. 4.15 The fringe pattern in Fig. 4.13 filtered by the LMS with $p = 32$ and $\mu = 0.01$.

Figs. 4.16 and 4.17 show the output of the RLS filter with λ equal to 1 and 0.8, respectively, and the constant c set at the recommended minimum 13×10^5 (with σ_y^2 estimated at around 7×10^{-5} at the beginning of the scan).

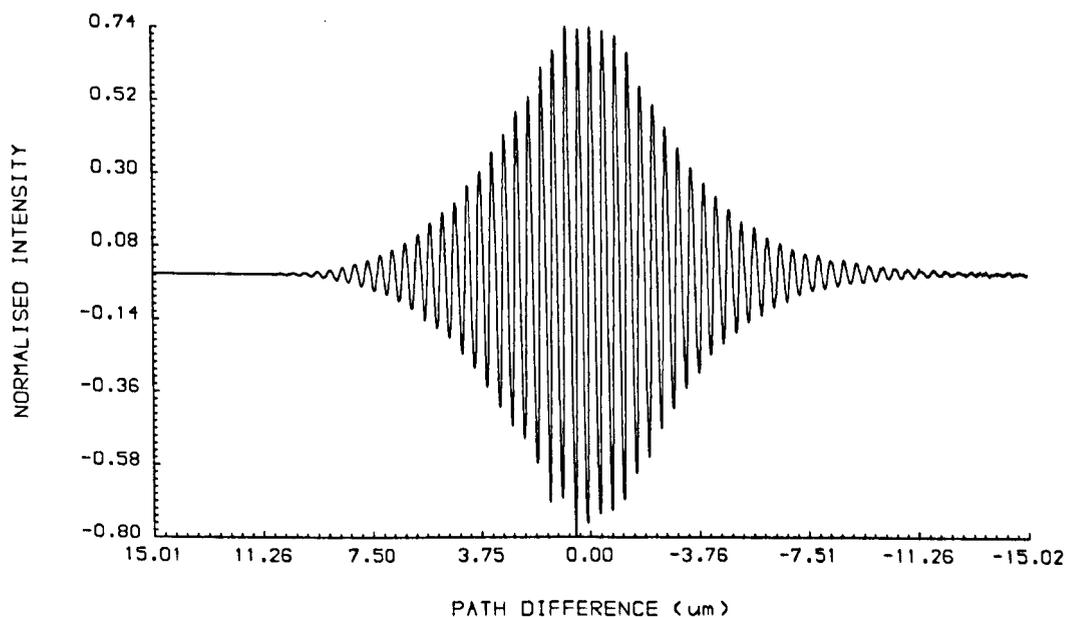


Fig. 4.16 The fringe pattern in Fig. 4.13 filtered by the RLS with $p = 32$ and $\lambda = 1.0$.

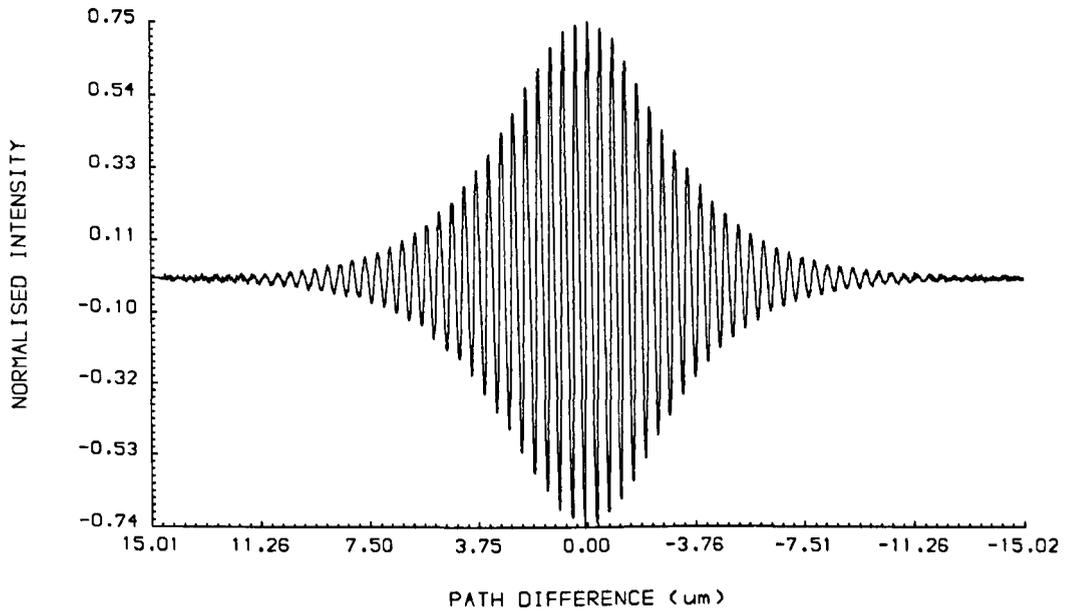


Fig. 4.17 The fringe pattern in Fig. 4.17 filtered by the RLS with $p = 32$ and $\lambda = 0.8$.

Adaptation of the weights is delayed until $9.8 \mu\text{m}$ in Fig. 4.16, but starts almost immediately in Fig. 4.17; in both cases it is faster than with the LMS algorithm, and in the latter case it is so fast that it would be difficult to spot any difference between the patterns in Figs. 4.17 and 4.13.

This leaves no doubt about the superior convergence of the RLS. As far as tracking is concerned, though, the stationary choice seems to have caused quite a large lag error after convergence, and the non-stationary choice seems not to have been able to separate the signal component from the noise component.

As in temporally scanned systems vibration noise may have a centre frequency close to that of the interference signal [2], the RLS algorithm should not be blamed if it was interpreting the noise as part of the signal and was trying to track it. However, in computer simulations where the additive noise was a pseudo-random sequence with a flat spectrum, the RLS filter was tracking both signal and noise if λ was allowed to get below a certain value. With the noise variance at 2.5×10^{-3} (SNR 20 dB) the MSE over the interference region dropped to 10^{-8} with $\lambda = 0.8$, although it was close to 2.5×10^{-3} with $\lambda = 1$. The MSE of the LMS over the fringes was also seen to decrease to 10^{-5} as μ was increased, before the algorithm became unstable.

Therefore, there is a limit on the tracking rate of both algorithms, determined by μ and λ , beyond which external noise sources can be interpreted as genuine signal variations.

4.5 Convergence and Tracking Aspects

One may ask at this point whether there is one single factor relating SNR and filter order to the behaviour of the two algorithms during convergence and tracking. The answer lies with the eigenvalues $\lambda_1, \dots, \lambda_p$ of the data autocorrelation matrix.

A necessary and sufficient condition for stability (convergence) of the steepest descent method is [3]

$$|1 - \mu\lambda_i| < 1 \quad i = 1, \dots, p \quad (4.3)$$

This follows after transforming the difference equation for the weight error vector $\mathbf{v}[n] = \mathbf{w}[n] - \mathbf{w}_o$, where \mathbf{w}_o is the optimum (Wiener) weight vector, into a set of uncoupled scalar equations [4]

$$v_i[n+1] = (1 - \mu\lambda_i)v_i[n] \quad i = 1, \dots, p \quad (4.4)$$

Hence, $\mathbf{v}[n]$ decays exponentially to zero when $0 < \mu < 2/\lambda_{max}$, irrespective of the initial conditions.

In practice, the gradient is estimated from the data and $\mathbf{v}[n]$ consists of a mixture of p noisy exponentials which are coupled together, with $\mathbf{w}[n]$ executing a random motion with amplification factor μ around the Wiener solution after convergence. Eq. 4.3 is thus referred to as the *convergence in the mean* condition for the LMS weights to the Wiener solution.

The *convergence in the mean square* condition [5]

$$\sum_{i=1}^p \frac{(\mu/2)\lambda_i}{1 - \mu\lambda_i} < 1 \quad (4.5)$$

ensures convergence of the MSE to a constant value, usually in excess of the MSE achievable by the Wiener filter by an amount equivalent to the gradient error.

Although both convergence conditions rely on the independence theory [6], so that the statistical dependence between successive gradient estimates can be ignored and the uncoupling in Eq. 4.4 may be considered valid, they have been shown to be reliable even for highly dependent data sequences [7, 8], provided μ is small enough

to make the gross temporal variation of the weight vector far slower than that of the input vector ¹. This is important because the data containing the fringe pattern are highly correlated and therefore far from being independent.

Restating Eq. 4.5 as

$$\mu < 2 \sum_{i=1}^p \frac{1 - \mu\lambda_i}{\lambda_i} \quad (4.6)$$

and since

$$\sum_{i=1}^p \frac{\lambda_i}{1 - \mu\lambda_i} > \sum_{i=1}^p \lambda_i > \lambda_{max} \quad (4.7)$$

it follows that if convergence in the mean square is ensured, convergence in the mean is also ensured. In [5] it was shown that both conditions are satisfied when $\mu < 2/(3\lambda_{max})$.

If τ_i^v is the time required for v_i to decay to $1/e$ of its initial value, then [9]

$$1 - \mu\lambda_i = e^{-1/\tau_i^v}, \quad \text{or} \quad \tau_i^v = -1/\ln(1 - \mu\lambda_i) \quad (4.8)$$

which for small μ can be approximated as

$$\tau_i^v = 1/(\mu\lambda_i) \quad (4.9)$$

The time required for the i th component of the excess MSE to decay to $1/e$ is $\tau_i^v/2$, since the MSE is a sum of squares of the v_i terms [3].

Eq. 4.9 shows that the smaller an eigenvalue is, the bigger its associated time constant τ_i^v . Eq. 4.3, on the other hand, shows that the smaller the eigenvalue, the bigger the step-size can be. This means that if p separate step-sizes could be used, the values $\mu_i = 1/\lambda_i$ would ensure convergence of all the weights in one step. With a common step-size instead, the largest eigenvalue gives an upper bound for it. Thus, since $\tau_{max}^v = 1/(\mu\lambda_{min})$ and $\mu < 1/\lambda_{max}$, it follows that $\tau_{max}^v > \lambda_{max}/\lambda_{min}$, i.e., the larger the eigenvalue ratio (or spread) of the input autocorrelation matrix the longer the weights will take to converge. The eigenvalues may be considered to be highly disparate when the ratio is greater than ten [10].

Since the eigenvalues of $\mathbf{X}^T\mathbf{X}$ are the squares of the singular values of \mathbf{X} in Table 3.2, this limit can be easily passed as the filter order and/or the SNR are increased. This conclusion can also be reached by noting that the eigenvalues can be approximated by uniformly spaced samples of the process spectrum for p sufficiently

¹The small- μ assumption, $\mu \ll 1/\lambda_{max}$, has been used extensively in the past to derive most of the results listed in this and the next sections.

high, and are bounded by the minimum and maximum values of the spectrum as p tends to infinity [11]. For a narrow-band process with spectral shape as in Fig. 2.2, decreasing the noise level decreases the spectral minimum and λ_{min} with it, whereas increasing the spectral resolution by increasing the number of parameters of the AR model increases the sharpness of the peak and λ_{max} . Adopting the spectral view it is also easy to see the connection between ill-conditioning and eigenvalue ratio, through the persistent input excitation condition of Section 3.4.

It has been shown [12, 13] that the LMS convergence in the mean square is affected less than the convergence in the mean. In other words, the filter output converges more rapidly than the weights, since those components associated with frequency regions having little power, i.e. small eigenvalues and therefore large time constants, do not contribute significantly to the excess MSE. This means that the damage caused by the eigenvalue spread during convergence is less severe in prediction and filtering than in system identification or spectral estimation, if μ is small enough to assume that the various components are uncoupled.

The RLS algorithm, on the other hand, uses a Gauss-Newton procedure, which is insensitive to the eigenvalue distribution of the data autocorrelation matrix [14]. As such, its behaviour during convergence should remain independent of the filter order, noise level, and degree of correlation of the input data.

As far as tracking is concerned, the behaviour of both algorithms is essentially the same as during convergence. The time constants of the LMS are still defined by Eq. 4.9 [10], whereas those of the RLS remain independent of the eigenvalue spread and can be approximated by $1/(1 - \lambda)$ [14]. However, the gain vector of the RLS is smaller than the initial gain $\mathbf{G}[0]$. Hence, the response of the RLS to sudden changes in the weight vector is slower during tracking than during convergence [15]. Indeed, with the step-size optimised in both algorithms, the LMS has been found to exhibit less misadjustment at steady-state in some non-stationary environments with low SNR when both weight and lag error contributions to the excess MSE are considered [16], although it is the general opinion that the RLS is able to track faster (see, e.g. [14, 17]).

4.6 Choice of Filter Parameters

Unlike with batch processing schemes, where the filter order is generally high if the noise level is high but the block width is decreased if the degree of non-stationarity is high, in adaptive schemes the different impacts that step-size, filter order, SNR, DNS, and eigenvalue spread have on the convergence rate and tracking make the choice of the filter parameters a difficult task.

Zooming in on the LMS algorithm, it follows from Eq. 4.4 that as long as $\mu < 1/\lambda_{max}$ the rate of decay of all weight error components increases with μ . Once $\mu = 1/\lambda_{max}$ the component associated with λ_{max} begins to slow down again, switching from over- to underdamping. The choice $\mu = 1/\lambda_{max}$ is therefore generally accepted as being the one that provides the fastest convergence in the mean [18], although in practice the convergence in the mean square condition $\mu < 2/(3\lambda_{max})$ may dictate otherwise.

To avoid computing λ_{max} , a rule of thumb is to use the following relationship between the trace ² and the eigenvalues of the autocorrelation matrix \mathbf{R}

$$tr(\mathbf{R}) \equiv \sum_{i=1}^p \lambda_i > \lambda_{max} \quad (4.10)$$

Since $tr(\mathbf{R}) = pr[0]$ and $r[0] = \sigma_y^2$, where σ_y^2 is the variance of the data, convergence in the mean square is guaranteed as long as $\mu < 2/(3p\sigma_y^2)$, with μ values close to this upper bound providing the fastest convergence in the mean.

Eq. 4.10 provides an approximation which is p times more conservative than the exact expression over the region delimited by t_1 in Figs. 4.5-4.7, since in this region all eigenvalues are equal. Over the fringes, however, the ratio between left and right-hand sides of Eq. 4.10 decreases quickly. With $p = 32$ and a SNR of 10 dB the ratio was 32 over the t_1 region, 3.1 at the $1/e$ intensity points, and 1.2 at the central fringe; with a SNR of 30 dB λ_{max} was so dominant over the fringes that the ratio was already 1.0 at the $1/e$ points.

In [5] the misadjustment at steady-state due to the gradient error alone was derived as

$$M = \frac{\sum_{i=1}^p (\mu/2)\lambda_i / (1 - \mu\lambda_i)}{1 - \sum_{i=1}^p (\mu/2)\lambda_i / (1 - \mu\lambda_i)} \quad (4.11)$$

²The trace of a square matrix is the sum of the elements on its leading diagonal.

which with the small- μ assumption reduces to

$$M \approx \frac{(\mu/2) \sum_{i=1}^p \lambda_i}{1 - \sum_{i=1}^p (\mu/2) \lambda_i} \approx (\mu/2) \sum_{i=1}^p \lambda_i = (\mu/2) p \sigma_y^2 \quad (4.12)$$

a result previously obtained in [3]. Eq. 4.12 shows that the misadjustment increases with the step-size and is dominated by the largest eigenvalue.

The misadjustment of the RLS at steady-state (neglecting the lag error) can instead be approximated as

$$M \approx \frac{1 - \lambda}{1 + \lambda} p \quad (4.13)$$

for λ very close to one [19], thus increasing with p and decreasing with λ but being independent of the eigenvalue spread.

To ensure fast convergence of the LMS followed by small misadjustment at steady-state, it has been suggested to choose a large μ initially, and then decrease it in discrete steps in time [18, 20]. Similarly, λ in the RLS can be made to approach one exponentially [21]. Such approach may only work well in a stationary environment, as the lag error increases linearly with μ^{-1} and with the DNS if non-stationarity causes random time variations of the system parameters [10, 14], and quadratically with μ^{-1} and the DNS if the induced variations are deterministic [22, 23, 24]. This applies also to the RLS [14, 24].

Deterministic variations can be expected when dealing with a time-varying system or signal with a trend, such as a frequency-modulated narrow-band signal whose centre frequency drifts slowly in time; in this case the DNS is constant and given by the drift rate or 'chirp'. As the amplitude of the fringe pattern is described by a well-behaved Gaussian function, it may be justified to assume deterministic variations of the WLI system parameters; in this case the DNS is given by the rate of change of the SNR within the Gaussian envelope, which is not constant even if the coherence length of the source and the fringe width are.

As for the filter order p , keeping it low reduces not only the weight but also the lag error [10, 19], since fast time variations can be tracked more easily as the memory of the filter is shortened [25]. On the other hand, in general the AR model matches the system more accurately with increased model order [26].

Attempts have been made to derive an expression for the optimum LMS step-size as a function of the filter order and some or all of signal power, noise power, eigenvalue ratio, and DNS (see, e.g. [27, 10, 28, 16, 29]). In practice, its search may need to be made by trial and error [10].

From simulations conducted on synthetic WLI fringe patterns the following results are worth mentioning:

1. On average over 8 values of the SNR in the range [0,32 dB], μ_0 was inversely proportional to p over the range of p considered.
2. μ_0 increased slightly with the SNR.
3. Increasing p the identification rate improved markedly, especially for low p , as a result of better model fitting.
4. As the coherence length was reduced to increase the slope of the Gaussian profile and the DNS with it, μ_0 also had to increase to provide faster convergence and tracking.

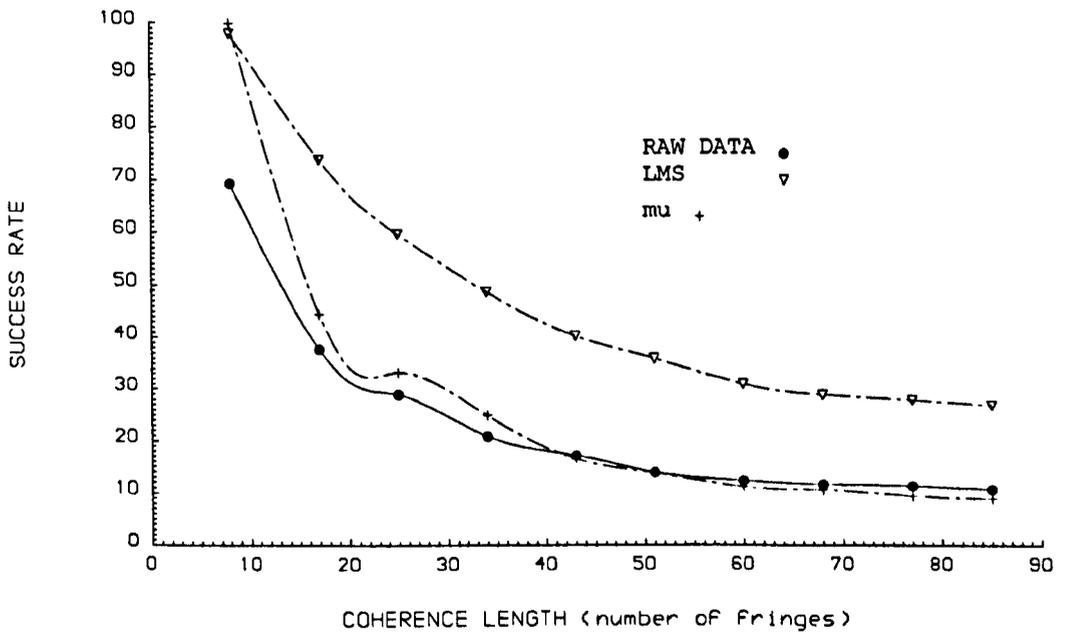


Fig. 4.18 Success rate against coherence length for a SNR of 26 dB (5 % noise) and $p = 16$. Also shown is μ_0 (amplified 555 times).

Fig. 4.18 shows that as the coherence length increases the identification rate degrades, as a consequence of the visibility profile becoming flatter around the central fringe. As the fringe pattern becomes more stationary the need for convergence and tracking decreases and so does μ_0 . It appears that μ_0 and L_c obey an inverse square or exponential law.

For the RLS algorithm, on the other hand, similar simulations have confirmed that λ is quite insensitive to changes in the filter order [14], and that its choice in general is less crucial than that of μ [1]. The identification rate is a much smoother function of λ than of μ , especially at high SNRs, where weight and lag errors seem to compensate each other during tracking. At lower SNRs instead, a reduction in lag error may be counteracted by a slightly larger increase in weight error, since the identification rate falls steadily as λ is decreased from one, although never as dramatically as with the LMS.

4.7 Discussion

In this chapter the convergence and tracking behaviour of the LMS and RLS algorithms during the filtering of the WLI fringe pattern have been examined.

It was shown that decreasing the noise level reduces the minimum value of the spectrum of the process and with it the smallest eigenvalue of the input data autocorrelation matrix. On the other hand, increasing the filter order increases the spectral resolution and with it the largest eigenvalue. In both cases the increased spectral range results in a large eigenvalue spread and slow convergence of the LMS.

The convergence of the RLS is not influenced by the eigenvalue spread, but the latter will take its toll on the numerical precision of the algorithm. This, coupled with a slower response to signal variations during tracking than during convergence, may give the LMS a slight performance advantage in non-stationary environments, on condition that its step-size is chosen appropriately.

Unfortunately, the choice of the step-size is not an easy one, as its optimum value depends not on the filter order alone but on its interplay with such environmental variables as signal power, noise power, signal bandwidth and degree of non-stationarity, some of which change continuously during scanning. The forgetting factor in the RLS is less influenced by the filter order and the SNR, as these affect the eigenvalue distribution of the data but not the convergence and tracking speed of the algorithm.

In the next chapter a simple scheme which tries to ease the convergence problem will be presented and analysed.

Alternative versions of the LMS that have been claimed to offer faster conver-

gence and/or tracking will then be assessed, before introducing a modification to the weight update equation in the standard version that leads to superior performance at no extra computational cost with next-to-nothing dependence on the step-size.

A final round-up of methods for central fringe identification in WLI systems with their pros and cons concludes the chapter.

Bibliography

- [1] L. Ljung and T. Söderström, *Theory and Practice of Recursive Identification*. Cambridge, MA: MIT Press, 1983.
- [2] S. Chen, A. W. Palmer, K. T. V. Grattan, and B. T. Meggitt, "Digital signal-processing techniques for electronically scanned optical-fiber white-light interferometry," *Appl. Opt.*, vol. 31, no. 28, pp. 6003-6010, Oct. 1992.
- [3] B. Widrow and J. McCool, "A comparison of adaptive algorithms based on the methods of steepest descent and random search", *IEEE Trans. Antennas Propag.*, vol. 24, no. 5, pp. 615-637, Sept. 1976.
- [4] B. Widrow, "Adaptive filters I: fundamentals," *Tech. Rep. SU-SEL-66-126*, Stanford Electron. Lab., Stanford, CA, Dec. 1966.
- [5] L. L. Horowitz and K. D. Senne, "Performance advantage of complex LMS for controlling narrowband adaptive arrays," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 29, no. 3, pp. 722-736, June 1981.
- [6] K. D. Senne, "Adaptive linear discrete-time estimation, *Tech. Rep. 6778-5*, Stanford Univ. Center for Systems Research, Stanford, CA, 1968.
- [7] D. C. Farden, "Stochastic approximation with correlated data," *IEEE Trans. Inform. Theory*, vol. 27, no. 1, pp. 105-113, Jan. 1981.
- [8] W. A. Gardner, "Learning characteristics of stochastic-gradient-descent algorithms: a general study, analysis and critique," *Signal Process.*, no. 6. Amsterdam, Holland: North-Holland, pp. 113-133, Apr. 1984.
- [9] L. J. Griffiths, "Rapid measurements of digital instantaneous frequency," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 23, no. 3, pp. 207-222, Apr. 1975.

- [10] B. Widrow, J. M. McCool, M. G. Larimore, and C. R. Johnson, Jr., "Stationary and nonstationary learning characteristics of the LMS adaptive filter," *Proc. IEEE*, vol. 64, no. 8, pp. 1151-1162, Aug. 1976.
- [11] U. Grenander and G. Szegö, *Toeplitz Forms and their Applications*. Berkeley, CA: Univ. of California Press, 1958.
- [12] M. J. Shensa, "Time constants and learning curves of LMS adaptive filters," *Tech. Rep. 312*, Naval Ocean Systems, San Diego, CA, Nov. 1978.
- [13] M. J. Shensa, "The spectral dynamics of evolving LMS adaptive filters," *Proc. Int. Conf. Acoust., Speech, Signal Process.*, Washington, DC, pp. 950-953, Apr. 1979.
- [14] E. Eleftheriou and D. D. Falconer, "Tracking properties and steady-state performance of RLS adaptive filter algorithms," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 34, no. 5, pp. 1097-1109, Oct. 1986.
- [15] C. F. N. Cowan and P. M. Grant, *Adaptive Filters*. Englewood Cliffs, NJ: Prentice-Hall, 1985.
- [16] N. J. Bershad and O. M. Macchi, "Comparison of RLS and LMS algorithms for tracking a chirped signal," *Proc. Int. Conf. Acoust., Speech, Signal Process.*, Glasgow, Scotland, pp. 896-899, May 1989.
- [17] J. G. Proakis and D. G. Manolakis, *Digital Signal Processing: Principles, Algorithms, and Applications*, 2nd ed. New York: MacMillan, 1992.
- [18] R. D. Gitlin, J. E. Mazo, and M. G. Taylor, "On the design of gradient algorithms for digitally implemented adaptive filters," *IEEE Trans. Circ. Theory*, vol. 20, no. 2, pp. 125-136, Mar. 1973.
- [19] R. S. Medaugh and L. J. Griffiths, "A comparison of two linear predictors," *Proc. Int. Conf. Acoust., Speech, Signal Process.*, Atlanta, GA, pp. 293-296, Apr. 1981.
- [20] G. Ungerboeck, "Theory on the speed of convergence in adaptive equalizers for digital communications," *IBM J. Res. Dev.*, vol. 16, no. 6, pp. 546-555, Nov. 1972.

- [21] T. Söderström and P. Stoica, *System Identification*. New York: Prentice-Hall, 1989.
- [22] E. Eweda and O. M. Macchi, "Tracking error bounds of adaptive nonstationary filtering," *Automatica*, vol. 21, no. 3, pp. 293-302, May 1985.
- [23] N. J. Bershad and O. M. Macchi, "Adaptive recovery of a chirped sinusoid in noise, part 2: performance of the LMS algorithm," *IEEE Trans. Signal Process.*, vol. 39, no. 3, pp. 595-602, Mar. 1991.
- [24] O. M. Macchi and N. J. Bershad, "Adaptive recovery of a chirped sinusoid in noise, part 1: performance of the RLS algorithm," *IEEE Trans. Signal Process.*, vol. 39, no. 3, pp. 583-594, Mar. 1991.
- [25] N. J. Bershad, P. L. Feintuch, F. A. Reed, and B. Fischer, "Tracking characteristics of the LMS adaptive line enhancer response to a linear chirp signal in noise," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 28, no. 5, pp. 504-516, Oct. 1980.
- [26] J. Makhoul, "Linear prediction: a tutorial review," *Proc. IEEE*, vol. 63, no. 4, pp. 561-580, Apr. 1975.
- [27] W. A. Gardner, "Nonstationary learning characteristics of the LMS algorithm," *IEEE Trans. Circ. Syst.*, vol. 34, no. 10, pp. 1199-1207, Oct. 1987.
- [28] B. Widrow and E. Walach, "On the statistical efficiency of the LMS algorithm with nonstationary inputs," *IEEE Trans. Inform. Theory*, vol. 30, no. 2, pp. 211-221, Mar. 1984.
- [29] O. M. Macchi, "Optimization of adaptive identification for time-varying filters," *IEEE Trans. Autom. Contr.*, vol. 31, no. 3, pp. 253-287, Mar. 1986.

Chapter 5

A New WLI Central Fringe Identification Scheme

5.1 Introduction

In Chapter 4 it was shown that although the LMS algorithm has a clear disadvantage over the RLS in terms of convergence, its identification rate may equal that of the RLS because of comparable, if not better, behaviour during tracking. For this to happen, however, the choice of the step-size has to be made with extra care in order to end the convergence phase as soon as possible without leaving an unnecessary high misadjustment afterwards.

In this chapter a simple scheme will be introduced, which tries to reduce the convergence disadvantage. A few alternative versions of the LMS algorithm will then be described and compared with the standard one, before moving on to an implementation which offers a very high identification rate without sacrificing computational speed and operating range.

A final assessment of the main methods that can be used at present for fringe order identification in WLI systems ends the chapter.

5.2 Towards Faster Convergence

One way to reduce the eigenvalue ratio and the chance of ill-conditioning is to add a small constant α to the elements on the leading diagonal of the input autocorrelation matrix [1]. By doing so λ_{min} is increased by a relatively larger amount than λ_{max} ,

since $\lambda_i(\mathbf{R} + \alpha\mathbf{I}) = \lambda_i(\mathbf{R}) + \alpha\lambda_i(\mathbf{I}) = \lambda_i(\mathbf{R}) + \alpha$.

The technique was introduced as a means of finding a least squares solution which could reduce the variance of the estimates at the expense of some bias in linear regression problems suffering from collinearity [2, 3], and is also known as regularisation by weight decay in the neural network community [4], as it is equivalent to adding a penalty term $\alpha\|\hat{\mathbf{w}}[n]\|^2$ to the error function in Eq. 2.10, so as to penalise large weight values. The same effect can also be obtained by adding a white noise sequence of variance α to the input time series [5].

The *leaky* LMS [6] is an adaptive implementation of this idea applied to the LMS algorithm. The weight vector is updated as

$$\mathbf{w}[n+1] = (1 - \alpha\mu)\mathbf{w}[n] + \mu e[n]\mathbf{y}[n-1] \quad (5.1)$$

where $1 - \alpha\mu$ is the leakage factor. Although it introduces some bias, the leaky LMS stabilises digital implementations; in particular, it counteracts weight drift when working with narrow-band signals by nudging the weights towards zero [7].

The identification rate in Fig. 4.3 became progressively flatter as α was increased, but unfortunately its peak decreased at the same time. Leakage might be more useful in digital implementations of the RLS or of the batch schemes in Chapter 2 when the noise level is very low, given that errors along the eigenvector corresponding to λ_{min} are increased in magnitude by the factor $1/\lambda_{min}$ during inversion of the input autocorrelation matrix [8].

A more effective way of decreasing the eigenvalue ratio is by flattening or *prewhitening* the power spectrum [9] in order to increase λ_{min} and reduce λ_{max} . This has the opposite effect of prefiltering, which enhances the SNR by suppressing frequency components outside the passband of the filter, increasing the spectral dynamic range and the eigenvalue spread [1].

Prefiltering of the fringe pattern using either low or bandpass constant-coefficient digital FIR or IIR filters resulted in some improvement in fringe identification when using the direct visibility method with SNRs in the range 0-10 dB. Above 10 dB the performance degraded rapidly as a result of amplitude and phase distortion. The effect on the LMS filter was catastrophic; the combination of signal distortion and increased eigenvalue spread caused its performance to degrade considerably throughout the range 0-40 dB.

The error sequence of a low-order predictor has been proposed for the prewhiten-

ing of narrow-band sequences [10]. The technique tries to decorrelate the signal with a low-order predictor, before the main LMS is applied. Faster convergence in a stationary environment has been observed as a result of the reduction in eigenvalue spread. Although one has to deal with two step-sizes and hence two sources of gradient error, it is claimed that the overall MSE does not increase because the eigenvalues of the prewhitened signal are smaller than those of the original signal and so smaller step-sizes can be used in both prewhitener and main LMS [11].

When applied to the WLI output signal, a five-weight batch prewhitener was necessary to bring down the eigenvalue ratio¹ from 400 to under 10. The eigenvalue ratio at the input of the prewhitener itself was still as high as 100. It is thus expected that a rather large step-size has to be used with a LMS prewhitener. The impact of two sources of lag error on the MSE also needs to be considered. For these reasons and the difficulty of optimising two step-sizes contemporaneously, adaptive prewhitening was not attempted.

A different idea is to restrict the use of the LMS to that region of the fringe pattern where the signal power is above a certain threshold, and save the weights at the end of each scan for the following one. This can be seen as an extreme example of the use of the time-sequenced adaptive filter for cyclo-stationary signals [12], where each of a set of adaptive filters is updated in turn, depending on the local characteristics of the signal.

As Figs. 4.8-4.10 showed, once the convergence phase is over the weights hover about their optimum settings as long as good fringe data are present, after which they start to drift at a rate proportional to the step-size and the noise level. If it can be assumed that the step-size required for optimum convergence is higher than that required during tracking, by confining weight updating to the thresholded region a smaller step-size can be used, which may increase the lag error over the fringes but hopefully reduce the gradient error by a relatively larger amount. Processing speed will also increase, allowing faster scan rates.

¹In Table 3.1, with $p = 16$ and a SNR of 20 dB.

5.3 Threshold Pre-processing

In the following, the threshold used is the signal power at the points of inflexion of the fringe visibility profile, although other values may be chosen. At these points the slope of the profile is at a maximum; hence, the rate of change of the signal power, and the DNS with it, are also maximum. During tracking, it is to be expected that the error term $e[n]$ between the signal and the LMS output increases rapidly, and that the lag error reaches its peak; this, in turn, speeds up the rate of adaptation of the weights to try and reduce the error. As one does not want to disturb the weights too much from one scan to the next, it would therefore be advisable not to update them outside the inflexion points. In practice, these points can be derived as follows.

If m and σ are the mean and standard deviation of a Gaussian process x , the underlying probability density function (PDF) peaks at m and has points of inflexion at $m \pm \sigma$. This is easily verified by setting the 1st and 2nd-derivatives of

$$\text{PDF}(x) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left[-\frac{(x-m)^2}{2\sigma^2}\right] \quad (5.2)$$

to zero. Since $\text{PDF}(m \pm \sigma) = \text{PDF}(m) \exp(-1/2)$, the density at $m \pm \sigma$ is $1/\sqrt{e}$ of its maximum value at m . Translating this to WLI, if the Gaussian visibility profile peaks at sample m , its points of inflexion are at samples $m \pm \sigma$.

The value of σ is easily found from Eq. 1.6. The cos term can be neglected because it does not enter the expression for the visibility profile; one can then equate as follows

$$\exp\left[-\left(\frac{k(m \pm \sigma) - \theta_s}{\pi L_c}\right)^2\right] = \exp\left[-\left(\frac{km - \theta_s}{\pi L_c}\right)^2\right] \exp(-1/2) \quad (5.3)$$

and as $km = \theta_s$ with zero OPD at sample m , σ is equal to $\pi L_c/(k\sqrt{2})$.

Since the amplitude of the visibility profile is $1/\sqrt{e}$ of its maximum at the inflexion points, the average signal power will be $1/e$ of its maximum. For example, with a coherence length equivalent to 17 fringes (delimited by the $1/e$ intensity points) and a sampling rate of 12.7 samples per fringe, $\sigma = 76.3$ and the inflexion points delimit the central 12 fringes or 152 samples.

As regards the choice of window used to estimate the signal power during scanning, a moving rectangular window of fixed length allows for fast updates, whereas an exponential window smooths out the effects of noise. Long windows reduce

the variance of the estimate, whereas short windows reduce the bias and the delay between the occurrence of the inflexion points and their detection. If A is the amplitude of the central fringe, $A^2/2$ its power, and σ_v^2 is the noise variance, the power at the central fringe is $A^2/2 + \sigma_v^2$ and that at the inflexion points is $A^2/2e + \sigma_v^2$. In simulations with a coherence length of 17 fringes, a sampling rate of 12.7 pixels per fringe, and a SNR of 20 dB, a rectangular window spanning two full fringes over or under-estimated the power at the central fringe and inflexion points by 2.3 % and 1.9 % on average, respectively, whereas a window length covering five fringes always under-estimated them, by 4.2 % and 4.3 % on average.

The overhead is only two multiplications, one addition and one subtraction per sample to update the current power in the moving window and compare it with the threshold. Clearly, since the threshold can only be derived from a previous estimate of the maximum signal power, its quality will depend on the stability of both signal and noise. The latter, in particular, causes the power at the inflexion points to be always greater than $1/e$ of that at the central fringe. The fractional power at the inflexion points is, in fact,

$$\left(\frac{A^2}{2e} + \sigma_v^2\right) / \left(\frac{A^2}{2} + \sigma_v^2\right) = \frac{A^2 + 2e\sigma_v^2}{e(A^2 + 2\sigma_v^2)} \quad (5.4)$$

As the noise level increases so does the fractional power. With a SNR of 20 dB it only equals $1.03/e$, but at 0 dB it is as high as $2.15/e$. If an estimate of σ_v^2 can be obtained during the calibration phase, e.g. by averaging the signal power at one or both ends of the array while the central fringe is in the middle, it can be subtracted from the maximum signal power and current power estimates to reduce the discrepancy. The average square error of the LMS algorithm at steady-state could also be used to update the σ_v^2 estimate on-line, should the noise level change during continuous measurements.

Once σ_v^2 is estimated, automatic detection and truncation of large isolated noise peaks becomes possible. A single outlier may have a dramatic influence on the short-term behaviour of the filter weights. Since the probability that $|v| > 2.6\sigma_v$ is less than $1/100$ for a random sample v drawn from a Gaussian distribution with zero mean and standard deviation σ_v , and becomes less than $1/1000$ for $|v| > 3.3\sigma_v$, any data with absolute value greater than, say, $A + 3\sigma_v$ can be treated as an outlier and truncated or substituted with its predicted value, without having to resort to an algorithm specifically designed to cope with additive impulsive noise, such as

the median LMS [13].

Slow changes in the signal mean caused by baseline drift or $1/f$ noise can instead be monitored on-line as follows

$$\bar{y}[n+1] = \bar{y}[n] + \frac{y[n+1] - y[n+1-L]}{L} \quad (5.5)$$

where $\bar{y}[n]$ is the estimated mean of the signal at sample n , computed using the present sample and $L-1$ previous samples. As an alternative, a bias weight can be added to the coefficient vector [14], making the filter capable itself of following the changes. Since with zero-mean measurement noise fringe pattern and filter output have the same mean only if the error sequence $e[n]$ has zero mean, a simple t -test on $e[n]$ would reveal whether bias compensation was needed. If $\bar{e}[n]$ and σ_e denote the mean and standard deviation of $e[n]$ computed over L samples, $\sqrt{L}\bar{e}[n]/\sigma_e$ is t -distributed with $L-1$ degrees of freedom under the null hypothesis that $e[n]$ is drawn from a Gaussian distribution with zero mean [15].

5.4 Simulation Results

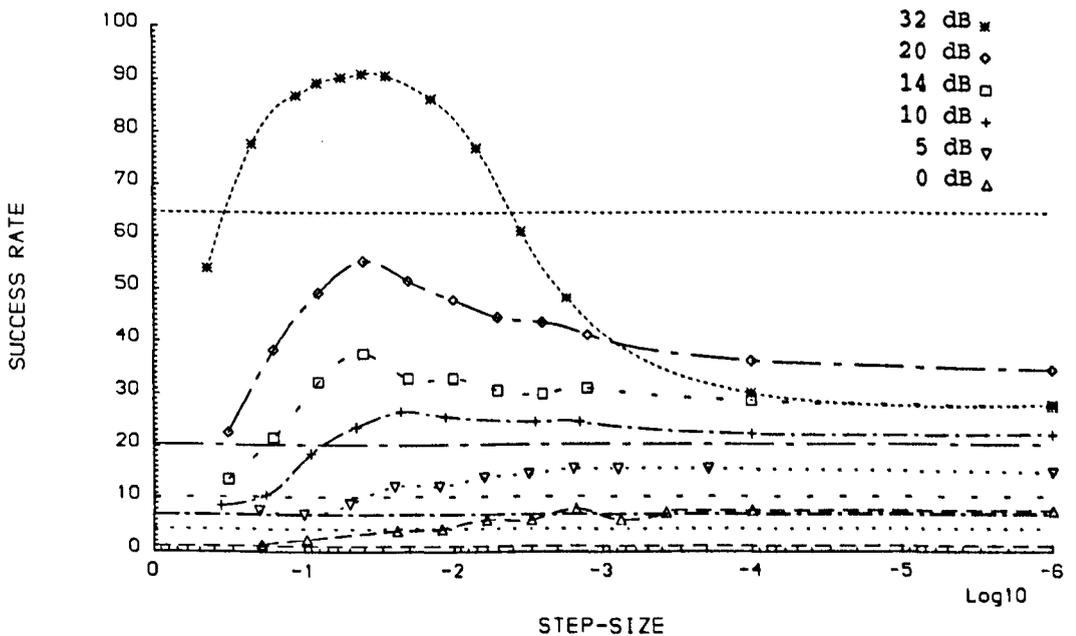


Fig. 5.1 Success rate with the thresholded technique against μ and SNR, with $p = 32$. The horizontal lines refer to the direct visibility method.

Fig. 5.1 shows the sub-fringe identification rate as a function of the step-size when the LMS algorithm with $p = 32$ weights was only used within the thresholded

region. For each value of the step-size and of the SNR considered, the following procedure was carried out.

Five preliminary scans were first taken with the central fringe in the middle of the CCD array, in order to obtain averaged estimates of σ_v^2 from the first and last 100 pixels, and of the maximum signal power, central fringe amplitude, and $1/e$ power threshold using the first 25 elements of the 32-long LMS input vector as a rectangular window. The LMS itself was not used at this stage.

Ten scans were then performed with the LMS algorithm turned on/off as soon as the signal power in the input vector went above/below the threshold, saving the weights each time round instead of resetting them to zero, thus allowing approximately 10×150 iterations in all for convergence. Finally, the identification rate was recorded over 500 consecutive trials with the central fringe at the centre of the CCD array; the LMS weights were never reset to zero. The truncation level for impulsive noise was set at $A + 4\sigma_v$, but was never exceeded during the trials.

Comparing with Fig. 4.3, which referred to the LMS without thresholding, it appears that the identification rate is a smoother function of μ for all $\mu < \mu_0$; in particular, small deviations from μ_0 have become less critical. However, μ_0 itself has not decreased significantly, and performance has improved slightly at low SNRs but has degraded by the same amount at high SNRs.

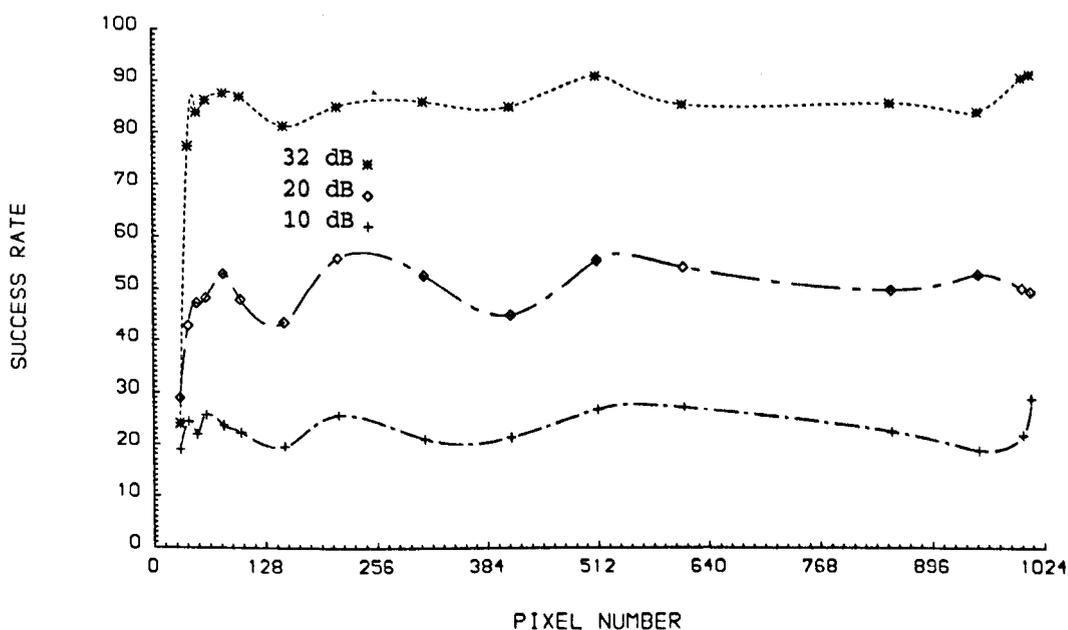


Fig. 5.2 Success rate with the thresholded technique against central fringe position, with $p = 32$ and μ optimised for the central pixel on the CCD array.

This can also be seen by comparing Fig. 5.2 with Fig. 4.4. Although the identification rate stays high for a few dozen more pixels as the central fringe approaches the bottom of the scanning range, at the very best it reaches the level of the LMS without thresholding as the central fringe moves away from it.

Therefore, thresholding may have helped the initial convergence, but in practice has done little to improve the overall performance. An additional source of non-stationarity has been introduced, caused by the sudden slope reversal of the visibility profile between scans. As a result, the DNS may even be higher at the inflexion points, and the large eigenvalue ratio there ² means that the step-size has to be kept quite high for effective tracking after convergence.

5.5 Enhanced LMS Algorithms

Several variations of the standard LMS algorithm have been proposed in an attempt to provide faster convergence with lower misadjustment and/or faster tracking capability. Both fixed and variable step-size alternatives will now be examined briefly for the case where weight updating is confined to the thresholded region.

5.5.1 Fixed Step-Size LMS

With the leaky LMS of Section 5.2 the maximum identification rate at 32 dB in Fig. 4.3 was recovered with $\mu = \mu_0$ in Fig. 5.1 and α in the range 0.01 to 0.05. At lower SNRs, however, $\alpha = 0$ remained the best choice.

In the γ -LMS [16] the current value of the weight vector in the update equation is scaled by the factor $\gamma = 1 + \mu\hat{\sigma}_v^2$, where $\hat{\sigma}_v^2$ is an estimate of the noise variance. This is equivalent to removing the noise contribution from the diagonal elements of $\mathbf{R}_{\mathbf{y}_{-1}\mathbf{y}_{-1}}$, in order to reduce the bias of the AR estimates.

As a single sinusoid in noise can be modelled by an ARMA(2,2) process with poles and zeros tending to the unit circle [17], unbiased estimates of the poles would be obtained only if the two zeros of an inverse filter were positioned, "on the average", at the pole locations. With an AR(2) approximation, the coefficients of a two-weight FIR inverse filter cannot be expected to provide unbiased estimates, as they have to compensate for the missing MA(2) term; they should, however,

²This is the square of the condition number at pixel 436 in Table 3.2.

recede on radials from the unit circle as more noise is added and the MA part becomes more significant, in order to maintain the complex conjugate property so characteristic of pseudo-periodic processes.

A non-radial trajectory from the unit circle inward has instead been observed as the SNR is reduced [18], meaning a general degradation in the estimation capability. Similar behaviour was also found here with the WLI output signal. Results with the data within the $1/e$ power threshold points are shown in Fig. 5.3. The AR estimates were computed with the covariance method, and the algorithm in [19] was used to solve for the poles, with modifications from [20].

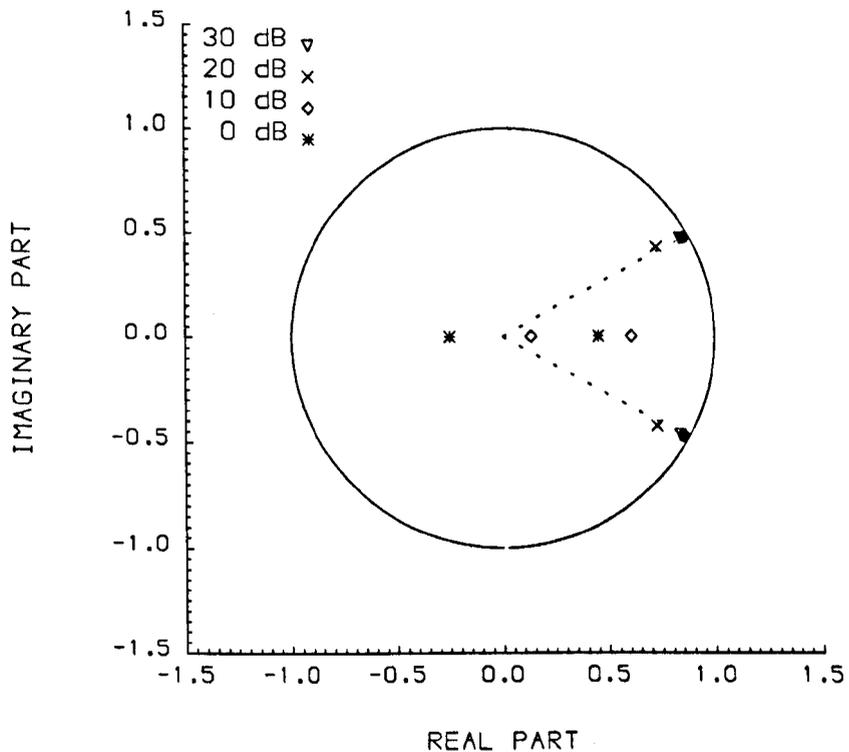


Fig. 5.3 Position of AR(2) model poles against SNR. The SNR = ∞ case is shown with bullets.

It can be seen that as the SNR falls below 20 dB the complex conjugate poles lose symmetry and start wandering about the unit circle; in this case they actually land on the real line. A positive real pole causes the autocorrelation function of the process to decay smoothly to zero, whereas with a negative one the autocorrelation function alternates in sign as it decays [21]. In both cases the pseudo-periodic behaviour is lost.

The γ -LMS may help counteract this problem, although the amount of correction needed will depend not only on the underlying SNR but also on the order of the

filter, as its zeros will position themselves so as to match the overall transfer function of the process. Convergence will also slow down as a result of the higher eigenvalue spread [22]. In fact, noise removal as used in the γ -LMS and noise injection as used in the leaky LMS are opposite operations.

Applied to the central fringe identification problem in WLI, the performance of the γ -LMS was the same as that of the standard LMS at high and medium SNRs, and slightly worse at low SNRs, which is where convergence can be affected most by changes in the noise level.

The *signed* LMS [23] replaces $e[n]$ with its sign in the weight update equation, where the sign is one if $e[n] \geq 0$ and minus one otherwise. This amounts to adopting the absolute error instead of the usual squared error criterion. One reason for considering nonmean-square cost functions is the presence of non-Gaussian noise. Since system and measurement noise are prevalently Gaussian in well-controlled WLI applications, this version was not considered.

Replacing the elements of the input vector, rather than the error term, by their sign leads to the *clipped* LMS algorithm [24], which offers fast operational speed when implemented in hardware, and fast convergence in the presence of highly deterministic signals. In fact, the optimum step-size was reduced by a factor of four at all SNRs with respect to the standard LMS, and the success rate shown in Fig. 5.2 was increased by 4.6 % on average at 32 dB, although the advantage was quickly lost at lower SNRs, with the standard outperforming the clipped LMS by 2.2 % at 20 dB and 3.7 % at 10 dB. This algorithm is therefore a simple option to be kept in mind when the SNR is high and very fast updates are needed.

The *two-sided* LMS was first proposed in [25] as the adaptive implementation of the modified covariance algorithm of Section 3.3, and was later reformulated as the *forward-backward* LMS (FB-LMS) in [26, 27]. Minimising the sum of squares of both the forward and backward errors as in Eq. 3.13 leads to the following augmented update equation for the weight vector [25]

$$\mathbf{w}[n+1] = \mathbf{w}[n] + \mu e_f[n] \mathbf{y}[n-1] + \mu e_b[n] \mathbf{y}_b[n-p+1] \quad (5.6)$$

where the backward input vector $\mathbf{y}_b[n-p+1]$ contains $-y[n-p+1], \dots, -y[n]$. Although it was shown in [27] to achieve half the misadjustment of the standard LMS for an equal rate of convergence, the success rate with the FB-LMS dropped by 8 % on average, for SNRs between 10 and 32 dB, with μ_0 values only slightly

smaller than with the standard LMS. The computational complexity also doubles, although storage requirements remain the same.

5.5.2 Variable Step-Size LMS

Since μ_0 depends, among others, on DNS, signal power, SNR, and eigenvalue ratio, none of which is a constant even within the inflexion points of the visibility profile, it would be preferable to adopt a variable step-size rather than a fixed one. Here, methods that constrain the step-size to be the same for all the weights will be reviewed before some of the more complicated alternatives that allow a separate one for each weight.

The *normalised* LMS algorithm [28, 29] divides the fixed step-size by the current signal power in the input vector. This is an attempt to avoid a gradient error amplification problem which could occur when the input vector becomes large. It also makes the algorithm more robust to outliers [30], and decreases convergence time for a given level of noise in the weights [31]. Within the inflexion points the range of the signal power is so restricted, however, that such amplification problem is very unlikely to ever occur; more interesting, perhaps, is the fact that whereas the signal power is maximum over the central fringe and minimum at the inflexion points, the DNS is maximum at the inflexion points and minimum over the central fringe, where the slope of the visibility profile is zero. Hence, the extra gradient error which is to be expected when the central fringe is approached (as predicted by Eq. 4.12) can be offset without incurring in a large penalty from the lag error.

Unfortunately, the performance of the normalised LMS was equal to that of the standard LMS at all SNRs. The strict dependence of the identification rate on μ was not removed and the curves obtained were similar to the ones in Figs. 5.1 and 5.2. The only advantage to be gained by normalisation in this application is that the algorithm is scale-invariant, which may prove useful should the signal power change between scans. In fact, normalising the data to lie in the range $[-1,1]$ instead of $[-0.5,0.5]$ caused considerable degradation in the performance of the standard LMS, unless μ was decreased by a factor of four, in which case the original performance was recovered. This would be in agreement with Eq. 4.12.

The *gradient adaptive lattice* (GAL) algorithm [32, 33] is formulated around a lattice structure [34]. Lattice filters suffer less from round-off errors arising during

computation and from quantisation errors in the filter weights [35, 36], although they require more computation and storage space; their performance is also less sensitive to the eigenvalue spread of the input autocorrelation matrix. Unlike with transversal filters, where the error term enters the update equation of each coefficient, the natural stage-by-stage decoupling which occurs in the lattice structure between the error term from each stage and the corresponding reflection coefficient means that neither the value nor the accuracy of any one coefficient estimate is affected by any of the other estimates. Moreover, the common step-size is scaled at each stage by the estimate of signal power entering that stage. Using the eigenvalue-spectrum decomposition of Section 4.5 it follows from Eq. 4.9 that the time constants at each stage are approximately equal [37].

In practice, the dependence on the input autocorrelation matrix which affects the LMS is also found with the GAL, although in a lesser degree [38, 39]. The central fringe identification rate was only as high as 73 % at 32 dB, 27 % at 20 dB, and 12 % at 10 dB. The optimum step-size also grew much faster with the SNR than in the standard LMS, questioning the ability to reduce the eigenvalue spread and increase the convergence rate.

The *momentum* LMS adds a small fraction α of the previous weight change $\delta w[n]$ to the usual correction term in the weight update equation, in an attempt to accelerate convergence and tracking. In [40] it was shown that with α negative the performance degrades because whenever $w_i[n] - w_i[n - 1]$ is positive for some i , it would be more appropriate to add a positive fraction of that amount in order to improve the convergence rate. On the other hand, a positive α is equivalent to increasing the step-size, i.e. convergence is faster but the misadjustment is larger. The authors suggest setting $\alpha = 0$ near convergence, taking one back to the standard LMS. The only advantage of including a momentum term in the weight update equation may be a smoothing effect, which can be useful in applications where error bursting is a problem.

In the *variable step size* (VSS) LMS [41] the step-size depends on the squared error as follows

$$\mu[n + 1] = \alpha\mu[n] + \gamma e^2[n] \quad (5.7)$$

where $0 < \alpha \leq 1$ provides exponential forgetting and $\gamma > 0$ controls the level of misadjustment. A large prediction error will increase μ in order to provide faster

tracking, while a small prediction error will decrease μ to reduce the misadjustment. The authors suggest typical values of 0.97 for α and 10^{-4} for γ . To prevent divergence problems and ensure a minimum tracking ability, μ is restricted to the bounded interval $[\mu_i, \mu_L]$, with $\mu_L \leq 2/(3p\sigma_y^2)$.

The VSS algorithm was very sensitive to the choice of α and γ , and its performance was at the very best equal to that of the standard LMS, provided μ_i and μ_L were chosen close enough to the optimum μ values in Fig. 5.1 and α was kept above 0.99; as for γ , this depended on the SNR, optimum values being close to 1.0 at 32 dB and 0.1 at 20 dB. There was also a tendency for μ to decrease rapidly from μ_L towards μ_i during the initial convergence phase and remain there during tracking, thus replacing the problem of finding μ_0 in the standard LMS by that of finding optimum values for α , γ , and μ_i .

Another variable step-size LMS is the VS algorithm [42]. Here, a different μ_i is assigned to each weight w_i at time n , with $\mu_i[n+1]$ set to $\alpha^{-1}\mu_i[n]$ or $\alpha\mu_i[n]$ depending on whether the sign of the i -th component of the gradient estimate, $\hat{\Delta}_i$, has changed m_0 consecutive times or has remained the same m_1 consecutive times. The rationale behind this is that when the i -th component of the MSE reaches the bottom of the error surface along the i -th dimension, $\hat{\Delta}_i$ keeps changing sign as it oscillates back and forth; μ_i should then be decreased in order to reduce the gradient error. On the other hand, away from the bottom μ_i should be increased to accelerate the rate of descent. Suggested values for m_0 and m_1 are 2 or 3, and α can be set to two for ease of hardware implementation. As with the VSS algorithm, μ_i is restricted to lie within $[\mu_i, \mu_L]$. The authors claim a reduction in convergence time from a factor of 10 at high SNRs up to 50 at low SNRs over the standard LMS, with considerable non-stationary tracking potential.

However, the performance of the VS algorithm was highly dependent on the choice of both (m_0, m_1) and (μ_i, μ_L) , and was generally inferior to that of the standard LMS. The success rate of the latter was only approached when both μ_i and μ_L were pushed close to the μ_0 values in Fig. 5.1, as was the case with the VSS algorithm. The reason for this may be due to dependencies between the components of the gradient estimates. In the standard LMS the result is a coupling between weight updates, since the correction to weight i is proportional to $\hat{\Delta}_i$, which is correlated with $\hat{\Delta}_j$ for all $j \neq i$. This may not be serious as long as the small

μ approximation is valid, as remarked in Section 4.5. In the VS LMS, however, additional coupling is present between the step-size updates, as these depend on past signs of the estimated gradient. Under these circumstances it becomes difficult to accept the independence theory and the authors' claim that the VS algorithm speeds up convergence and is less sensitive to the eigenvalue spread.

These remarks apply equally well to the VSA algorithm in [43]. Indeed, both VS and VSA are sensitive to parameters selection, and may perform poorly when the individual step sizes are chosen unwisely [44].

Two more variable-step LMS algorithms ([45, 46]) also failed to improve on the standard LMS, and achieved their best when the respective variable gain controls were kept fixed.

5.6 A Modified Forward-Backward LMS

Replacing $e_f[n]$ by $e_f[n - p]$ and $y[n - 1]$ by $y[n - p - 1]$ in Eq. 5.6 the weight update equation becomes

$$\mathbf{w}[n + 1] = \mathbf{w}[n] + \mu e_f[n - p] \mathbf{y}[n - p - 1] + \mu e_b[n] \mathbf{y}_b[n - p + 1] \quad (5.8)$$

This is equivalent to minimising the sum of square errors between $y[n - p]$ and both of its forward and backward predictions, made at time n using the past $2p$ samples, rather than the sum of square errors between $y[n]$ and its forward prediction, and between $y[n - p]$ and its backward prediction. The schematic diagram in Fig. 5.4 should make the difference clear.

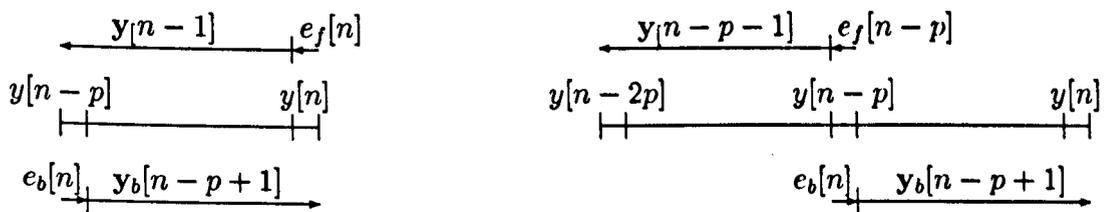


Fig. 5.4 Schematic diagram showing the difference between the original FB-LMS (left) and the modified form (right).

The two input vectors can be treated as one vector of length $2p$, containing $-y[n], \dots, -y[n - p + 1], -y[n - p - 1], \dots, -y[n - 2p]$.

Such formulation can be seen as an adaptive solution to the fixed data length smoothing problem [47]. The optimum linear smoother is the combination of two optimum linear filters, one of which works forward from the beginning of the data to some time n , while the other works from the end of the data backward to time n . It also corresponds to the delay form of a finite non-causal filter. A non-causal filter is not computationally realisable in software and not physically realisable in hardware, since its output at time n depends on future values of the input. However, by delaying the computations until the future values have occurred the filter is transformed into a computationally realisable form [48].

The computational complexity amounts to $5p + 2$ ($4p + 2$) multiplications per iteration for the filter (predictor) version, and $6p$ ($4p + 1$) additions/subtractions, plus $2p - 1$ shifts of the elements of the input vector. The forward-backward LMS requires the same number of multiplications and $5p + 1$ ($4p + 1$) additions/subtractions, plus $p - 1$ shifts, whereas the standard LMS requires $3p + 1$ ($2p + 1$) multiplications and additions/subtractions, plus $p - 1$ shifts. As for the operating range, this will be restricted to $[p + 1, N - p]$ instead of $[p + 1, N]$ of the standard and forward-backward LMS, because of the necessity of having both past and future p samples available when making a prediction for sample n .

The modified forward-backward LMS (MFB-LMS) algorithm of Eq. 5.8, coupled with the thresholding scheme introduced in Section 5.3, gave large improvements in fringe identification rate at all SNRs, with near-maximum performance being maintained for a very large range of the step-size parameter. The advantage was apparent even when the filter order p was set to half of that in the standard LMS. With this, the predictor versions of the two algorithms have the same computational complexity, and the filter version of the MFB-LMS requires less multiplications than the filter version of the standard LMS; the operating range, expressed in number of pixels, is also the same, and storage space is reduced by 25 %.

5.6.1 Simulation Results

Fig. 5.5 shows the performance of the thresholded MFB-LMS scheme with $p = 16$. Comparing with Fig. 5.1, which showed the performance of the standard LMS scheme with $p = 32$, the difference between the two algorithms is evident.

With the central fringe at ten different positions within pixels [40, 1000] the

MFB-LMS achieved a 99.7 % sub-fringe identification rate against 85.3 % at 32 dB, 67.5 % against 50.5 % at 20 dB, and 32.4 % against 24.1 % at 10 dB³. In addition, whereas the success rate of the LMS decayed exponentially as μ was reduced from μ_0 , degradation with the MFB-LMS was restricted to 2.2 % at 10 dB, 3.6 % at 20 dB, and 0 % at 32 dB, for any $\mu < \mu_0$ in Fig. 5.5.

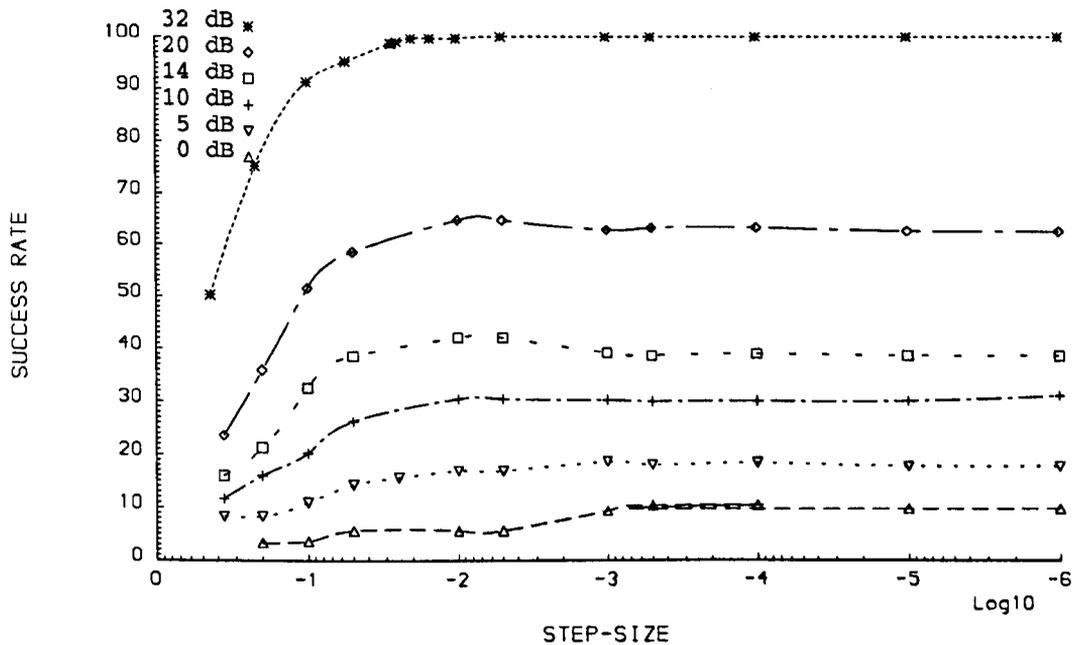


Fig. 5.5 Success rate with the thresholded MFB-LMS against μ and SNR out of 500 simulated scans with $p = 16$.

Neither the superior performance nor the insensitivity to the step-size were lost when using the predictor or the clipped-predictor versions, or when the filter or system parameters were changed. For example, at 20 dB the identification rate of the MFB-LMS filter, averaged over 3000 trials with the central fringe at ten different positions within pixels [40,985] and μ_0 derived at pixel 512, went from 67.5 % with $p = 16$ to 88.6 % with $p = 32$ and to 98.6 % with $p = 38$, with maximum degradation of 3.6, 3.0, and 0.7 %, respectively, for any value of μ less than 10^{-2} , 10^{-3} , and 10^{-4} . With the MFB-LMS predictor the rate went from 66.4 % to 88.0 % and to 98.6 %, respectively, with maximum degradation of 2.1, 1.5, and 0.7 %, whereas with the clipped-predictor it went from 63.2 % to 83.9 % and to 97.3 %, with maximum degradation of 2.0, 0.4, and 0 %.

³Out of a total of 3000 trials with $\mu = \mu_0$ in Figs. 5.1 and 5.5.

Still at 20 dB, but with a coherence length of 75 fringes and a fringe width of 8 pixels, the minimum success rate of the MFB-LMS predictor was 41 % with $p = 8$, 44 % with $p = 16$, and 50 % with $p = 32$, for all μ values equal to or below 5×10^{-2} , 5×10^{-3} , and 10^4 , respectively, whereas the maximum rate of the LMS predictor (i.e., with $\mu = \mu_0$) was only as high as 31 % with $p = 16$, 34 % with $p = 32$, and 38 % with $p = 64$.

What is the reason for the apparent insensitivity to the step-size in the MFB-LMS algorithm? When not using the thresholding technique the peak identification rate fell remarkably (e.g., from 67.5 % to 58.5 % at 20 dB with $p = 16$) in contrast to that of the standard LMS, which stayed approximately the same (e.g., 50 ± 0.5 % at 20 dB with $p = 32$). The identification rate in general also fell rapidly for $\mu \neq \mu_0$, as with the standard LMS. This, coupled with the fact that μ_0 was higher than when thresholding was used, would suggest that thresholding is critical in helping the MFB-LMS to converge.

Even assuming this to be correct, it is still puzzling to see that the identification rate remains the same for such a large spread of μ values used during convergence. Is it possible that the weights have all reached their steady-state values after the first ten passes within the $1/e$ power threshold points, irrespective of whether μ is 10^{-3} or a thousand times smaller? The answer must be negative, as the weights rose from zero to $O(10^{-1})$ or $O(10^{-2})$ when μ was 10^{-3} , but only to $O(10^{-4})$ when μ was 10^{-6} , throughout the range from zero to 32 dB.

Figs. 5.6-5.7 show the MSE across the CCD array when the MFB-LMS algorithm is operated only within the $1/e$ power threshold points (pixels 436 to 588). The noise variance was set at 2.5×10^{-3} (-2.6 on a log scale), to give a SNR of 20 dB. The MSE was averaged over 100 scans with the central fringe at the centre of the array, after allowing ten preliminary scans for convergence of the weights. The step-size was kept at 10^{-3} in Fig. 5.6 and at 10^{-6} in Fig. 5.7.

Although the algorithm is in tracking mode in Fig. 5.6 but is not in Fig. 5.7, the difference in identification rate between the two cases was only 2 %, and became zero when μ was changed from 10^{-6} to zero after the ten preliminary scans.

It was observed that the rate of increase of the MSE as μ was reduced from μ_0 was the same as in the no-threshold case, although the identification rate decreased only very slowly. This means that, as with the standard LMS, thresholding does

not help to speed up convergence but, unlike with the standard LMS, it ensures that the weights remain symmetric and invariant under scaling for all $\mu < \mu_0$.

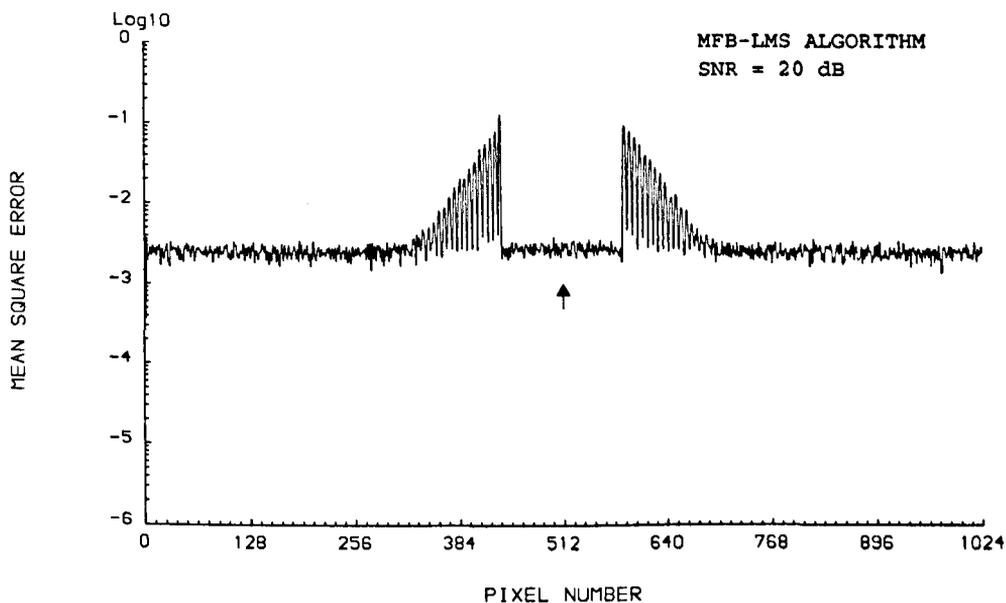


Fig. 5.6 MSE across the CCD array with the thresholded MBF-LMS algorithm and $\mu = 10^{-3}$.

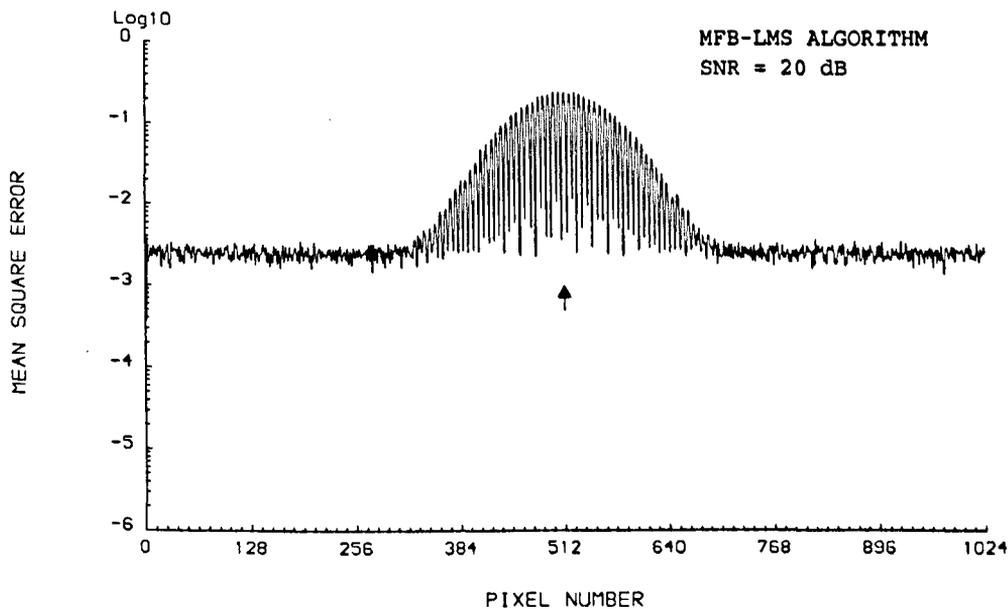


Fig. 5.7 MSE across the CCD array with the thresholded MBF-LMS algorithm and $\mu = 10^{-6}$.

With both the standard and the MFB-LMS, when $\mu = \mu_0$ the filter output is already of the same order of magnitude as the interference signal after ten scans, and

remains so afterwards; when μ is less than μ_0 , it takes longer for the filter output to grow and reach the signal level. However, whereas with the standard LMS a small step-size gives distorted patterns as long as convergence is not completed, with the MFB-LMS perfectly symmetric scaled-down patterns are produced almost immediately. Tracking is thus made a non-issue, if it is understood that tracking only starts once convergence is over, and the computational complexity can be halved by failing to update the weights altogether after a few preliminary scans, without sacrificing performance.

In conclusion, with $\mu < \mu_0$ during the preliminary scans, the coefficients and output of the thresholded MFB-LMS in the scans that follow are scaled replicas of the coefficients and output with $\mu = \mu_0$, with the scaling factor tending to unity as the measurements go on and the MSE is reduced. A step-size below the optimum value slows down convergence of all the weights uniformly, affecting the MSE but not the identification rate. Trying to induce faster convergence by decreasing the step-size steadily from 0.1 to 10^{-10} during the course of ten preliminary scans resulted in a very noisy adaptation of the weights and a low identification rate.

5.6.2 Experimental Evaluation

The comparison between standard and MFB-LMS predictors was extended to include lab measurements. The experimental set-up was the same as that described in Section 4.4, with the LED source replaced by a multimode laser diode with a nominal wavelength of 780 *nm* and an output power of 3 *mW*, being operated below threshold in order to give broad-band light with a coherence length corresponding to 75 fringes. A 1.2 metre length of multimode fibre with a core diameter of 200 μm was inserted between the source and the interferometer in order to reduce the intensity of the light falling on the photodiode. The fibre was designed for use at a wavelength of 633 *nm*, so the launch efficiency was poor and the attenuation high. The noise generator was also turned on to push the SNR down to 20 dB. Nine sets of measurements were recorded, each set consisting of 36 consecutive scans. The sampling rate was approximately 8 samples per fringe, for a total of 2049 data points per scan. Since the scans were performed at different times of the day and on different days, the position of the central fringe was checked before the start of each set using the two-wavelength method [49] with the noise generator turned off.

It was clear from the start that a reliable threshold estimate could not be obtained from a few preliminary scans, but had to be updated continuously from scan to scan. The diode driving current being supplied by a Ni-Cd battery, the source intensity was gradually reduced as the measurements went on, and both output mean and power dropped uniformly within each set, as from Eq. 1.2. The power decreased by as much as 2.7 % on average between scans, and at the end of a set the mean and central fringe amplitude could be as low as 57 % and 42 % of their initial values, respectively. As a consequence, the SNR normally dropped from 20 dB to 16 dB from the beginning to the end of a set.

Hence, the first scan of each set was reserved for the estimation of the mean and maximum power, and the battery was left to recharge at the end of each set. The first fifteen scans of the first set were also reserved for convergence of the weights; hence, the comparison between the two algorithms involved 300 scans in all. The instantaneous mean and power were updated using rectangular windows with arbitrary lengths of 128 and 64, respectively, and the mean was subtracted recursively from each sample. A bias weight was not used. The threshold for each scan was set to two-fifths of the maximum power estimated during the previous scan, limiting the weight update to forty fringes per scan on average. Since a slight drift of the central fringe was expected within each set, sub-fringe identification was not attempted.

The standard LMS predictor achieved a maximum success rate of 23 % with $p = 16$ weights, 27 % with $p = 32$, and 29 % with $p = 64$. The corresponding figures for the MFB-LMS predictor were 31 % with $p = 8$, 36 % with $p = 16$, and 38 % with $p = 32$. The identification rate of both algorithms was lower than in the simulations, and was not a smooth function of the step-size with a global maximum, as in Figs. 5.1 and 5.5; with the standard LMS the function was quite rugged and full of local maxima, whereas with the MFB-LMS the characteristic plateau beyond μ_0 was characterised by a few minor peaks and valleys.

This sort of behaviour is not surprising, given that both output power and SNR change during measurements. Even if μ was normalised by the power in the input vector, the dependence on the SNR would still remain.

Furthermore, the intensity of the diode injection current affects not only the output power but also the refractive index of the laser medium, changing the op-

tical frequency being generated and the central wavelength with it [50]. Thermal fluctuations in the cavity length also contribute to change the optical frequency; it is estimated that the central wavelength varies along with the operating temperature at a rate of $0.2 \text{ nm}/^{\circ}\text{C}$ [51]. Although this does not affect the position of the central fringe, it does affect the fringe width and the coherence length.

For all these reasons it appears that the DNS remains higher when processing measured rather than simulated data. With the thresholded MFB-LMS in particular, setting μ to zero to stop the weight adaptation after the convergence phase is over is not a good idea, as it may result in a large performance penalty.

It is envisaged that with proper offset/baseline correction and power/frequency stabilisation schemes, all of which are widely available, the success rate of both algorithms can only increase.

5.7 Matched Filter Detection

Consider the sum-of-products

$$r[n] = \sum_{i=1}^p y[n-i]z[p-i+1] \quad (5.9)$$

between the input vector and a vector of the same length containing the central fringe data. Eq. 5.9 represents a matched filter, with weight w_i equal to the noise-free and time-reversed central fringe data $z[p-i+1]$, $i = 1, \dots, p$ [52, 53].

In the absence of noise $r[n]$ is maximum when the two vectors superimpose exactly. This property is exploited in the correlation detection implementation of the matched filter. The method can be applied to WLI by computing the (unnormalised) cross-correlation $r[n]$ between the fringe pattern and a "clean" template stored in computer memory. As the template is moved across the CCD array, $r[n]$ will peak at the pixel corresponding to the best match between the two signals, and the position of the central fringe can be read off directly.

The computational complexity being only p multiplications and additions per data point, plus $p-1$ shifts of the elements of the input vector, this method can be as fast as the MFB-LMS with no weight updating after the convergence phase ⁴, unless the cross-correlation is smoothed using a window function [55, 56]; the dependency

⁴An alternative to cross-correlation is the average magnitude cross-difference [54], where multiplications are replaced by additions and subtractions.

on a step-size parameter is also removed, together with convergence and tracking issues. Its efficiency, however, depends on the quality of the template used. Since the position of the central fringe is only known exactly during calibration, this is also the only time to obtain a reliable template before the sensing phase begins.

With a template covering the two and a half central fringes ($p = 32$) averaged over ten preliminary scans in order to smooth out the noise, and processing being restricted to the region within the $1/e$ power threshold points, the technique gave an identification rate ⁵ of 100 % at 32 dB, 65.1 % at 20 dB, and 26.4 % at 10 dB. The rate at 20 dB went up to 89.1 % with $p = 64$, and to 98.1 % with $p = 76$.

From these results it is clear that a matched filter of order p is equivalent to the MFB-LMS filter of order $p/2$. In fact, the two filters are able to deliver the same performance over the same operating range and with equal computational cost. Given that a matched filter is the optimum solution for the detection of a known signal buried in additive white noise, one may ask why bother with adaptive filters.

The first point to note is that the two approaches can be combined into a single one, although the computational complexity will be the sum of the two and the operating range will be reduced by p samples. The adaptive filter can be used to remove most of the unwanted noise from the data before computing the cross-correlation. There is evidence that such schemes can be particularly effective in the presence of non-stationary noise [57, 58], as may be the case with, e.g., photoelectron noise of photosensing devices and speckle noise in laser interferometry.

The thresholded MFB-LMS scheme with $p = 16$, followed by the matched filter with $p = 32$, took the identification rate from 65.1 % of the matched filter alone to 88.3 % at 20 db, and from 26.4 % to 42.3 % at 10 dB, with only 0.4-0.8 % maximum degradation for any μ below μ_0 , and irrespective of whether the filter, predictor, or clipped-predictor version of the MFB-LMS was used. With $p = 32$ in the MFB-LMS and $p = 64$ in the matched filter, the rate was as high as 99 % for all $\mu < 10^{-3}$ at 20 dB, instead of 89.1 %. Replacing the MFB-LMS with $p = 16$ by the standard LMS filter with $p = 32$ took the identification rate from 65.1 % to only 67.2 % at 20 dB, and from 26.4 % to 31.8 % at 10 dB. Given that the dependence on μ remained as high as when using the LMS alone, there are no real advantages

⁵Out of a total of 3000 trials, with the central fringe at ten positions within pixels [40, 1000].

in combining the matched filter with the standard LMS algorithm.

The second point to note is that the quality of the cross-correlation estimate is only as good as the template from which it is calculated. If a good template cannot be obtained initially or becomes useless thereafter because one or more of sampling rate, mean wavelength, and source spectral width change during measurements, the technique loses its appeal. In this instance a multipoint measurement scheme, in which the light source is sent not to one but to two sensing units, could be used to generate an on-line template for one of the units by holding the other unit locked at the position corresponding to zero OPD. This would amount to a proper time-delay estimation problem, and a parametric approach based on the LMS or RLS algorithm [59, 60, 61] could be an alternative to the cross-correlation method.

Two disadvantages of such a scheme are, first, that the SNR over the fringes is reduced because of the sharing of the source power between the two sensors, and, second, that it does not become possible to smooth the template by averaging. Perhaps even more important is the fact that the central fringe can only move very slowly between scans for effective tracking of the time delay, and that the operating range has to be kept very low as the number of coefficients needed is twice the maximum of the integer delay.

In the next section a final assessment of central fringe identification techniques will be made in terms of success rate and speed of operation.

5.8 Comparison of Methods

The evaluation was carried out by recording the total processing time and the number of correct sub-fringe identifications at the end of a sensing phase which consisted of processing 1000 noisy patterns of the kind in Fig. 1.5, with the central fringe allowed to move randomly across the whole operating range between scans. The latter was dictated by the order p of batch and adaptive algorithms. With the fringe width used, the achievable resolution was 8 % of a fringe. A SNR of 20 dB was used for the comparison, as this is the current limit for an electronically scanned WLI system when no multiplexing is used [51].

Before the sensing phase started, five scans ⁶ were carried out to estimate the

⁶This operation was skipped for the centroid method.

threshold, followed by another ten scans ⁷ to allow for convergence of the adaptive algorithms or for the generation of the template needed by the matched filter. The reader is referred to Sections 5.4 and 5.7 for details.

Artificially generated data were preferred to real data for the following two reasons. First, unless the central fringe is kept locked throughout the measurements, there is no other way to know its precise position before each scan but to stop data acquisition, increase the SNR by increasing the output intensity, and locate the maximum point either through direct visibility alone or in conjunction with the multiple-wavelength method. Second, if an electronic scanner is not used, slow drifting of the central fringe with time means that the comparison has to be limited to fringe rather than pixel identification. Baseline and power drifts, if not corrected, will also influence the success rate of all the methods, but not necessarily in the same manner; an adaptive algorithm will always be affected in the long run, whereas the centroid will be affected only when the intra-scan variations are fast enough to tilt or shift the centre of symmetry. The degradation with an adaptive algorithm will also depend on whether or not μ -normalisation is used and, if a bias weight is not included, on the window length used for updating the mean level.

The use of a standard electronic correction scheme as suggested at the end of Section 5.6.2 may be the best way of preparing the acquired signal for further processing, as it allows to uncouple the contribution of offsets and drifts to the success/failure rate from that of the particular algorithm used for fringe identification.

The centroid method has been one of the first and most successful batch procedures for WLI fringe identification. Exploiting the symmetry of the output pattern on the two sides of the central fringe, a high success rate can be achieved by using a simple algorithm which makes extensive use of logical and comparison operators but enjoys a low arithmetic count, and which remains largely unaffected by occasional bursts of impulsive noise [51]. More specifically, the algorithm sorts through the batch of data and detects the width of each group of adjacent data that have the same sign; if this width is either greater than $b/2 + t$ or less than $b/2 - t$, where b is the fringe width and t is a tolerance, the data are not considered to constitute half a fringe, and are set to zero. The centroid or 'centre of gravity' of the sorted data is then calculated and used as an estimate for the position of the central fringe.

⁷This operation was skipped for the centroid and covariance methods.

As shown in Fig. 5.8, the algorithm was quite sensitive to t . When t is large, more fringes take part in the identification process, resulting in better estimation if the symmetry of the fringe pattern is maintained. Not surprisingly, large t values gave a very high identification rate when the central fringe was in the middle of the scanning range, but very low otherwise, while small values ensured a smoother decline in performance at the expense of reduced identification in the middle.

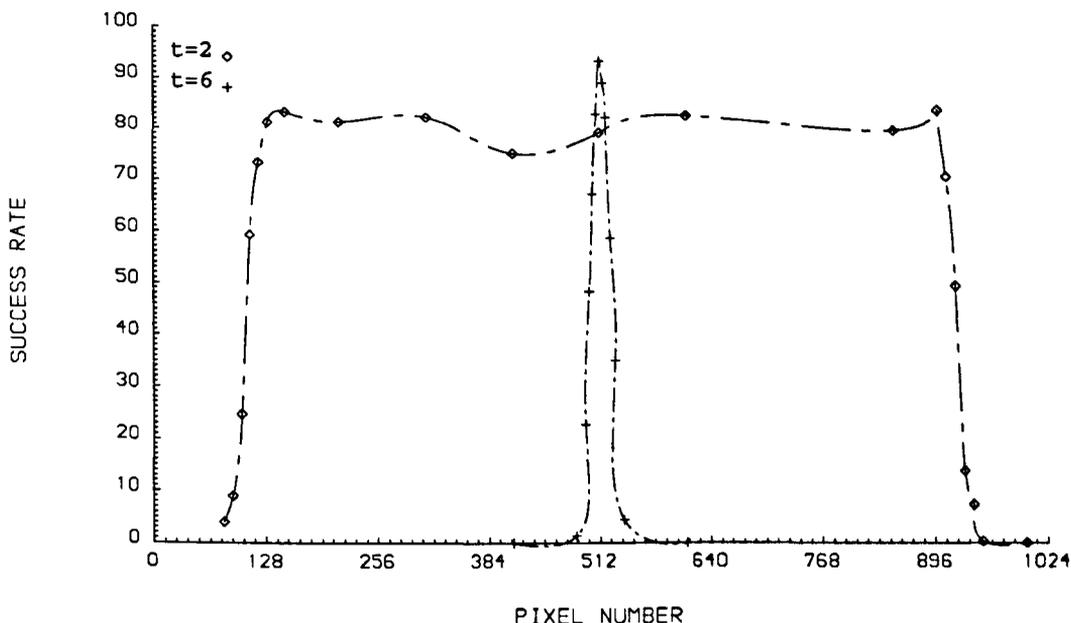


Fig. 5.8 Success rate (out of 300 trials) with the centroid method against central fringe position.

Fig. 5.9 shows the success rate achieved by the centroid method with $t = 2$ over the 1000-long trial session considered in this section. As t was increased the higher number of correct identifications around the middle of the scanning range was offset by an even higher number of failures away from the middle. The choice of a second tolerance parameter T , which should be set close to the coherence length [62], was not found to be critical.

Results using the covariance method with $p = 16$ and 32 are also shown in Fig. 5.9, for the case where forward and backward predictions for the same data sample are averaged together to produce a smoother estimate. This case is the batch equivalent of the MFB-LMS filter (the modified covariance method is, instead, the batch equivalent of the FB-LMS filter). Although the result is a higher identification rate than when either forward or backward prediction are used alone, p had to be halved in order to retain the original operating range. Upon doing this,

the performance advantage was almost completely lost, although there remained a substantial computational saving, as the number of operations per sample is $O(p^2)$.

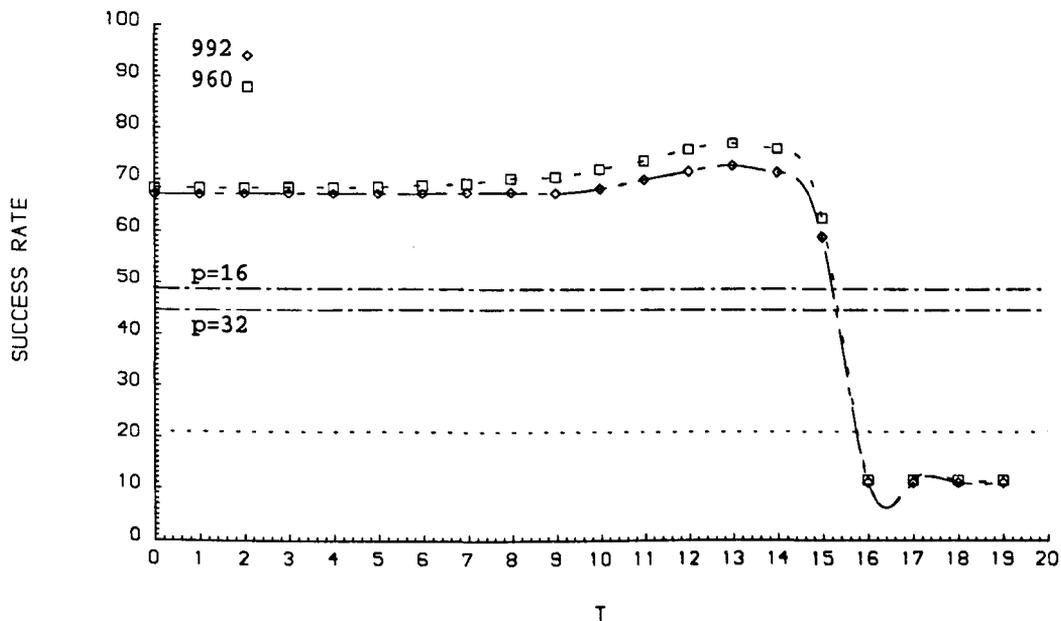


Fig. 5.9 Success rate with the centroid and covariance methods over operating ranges of 992 and 960 pixels. The dashed bottom line refers to the direct visibility method.

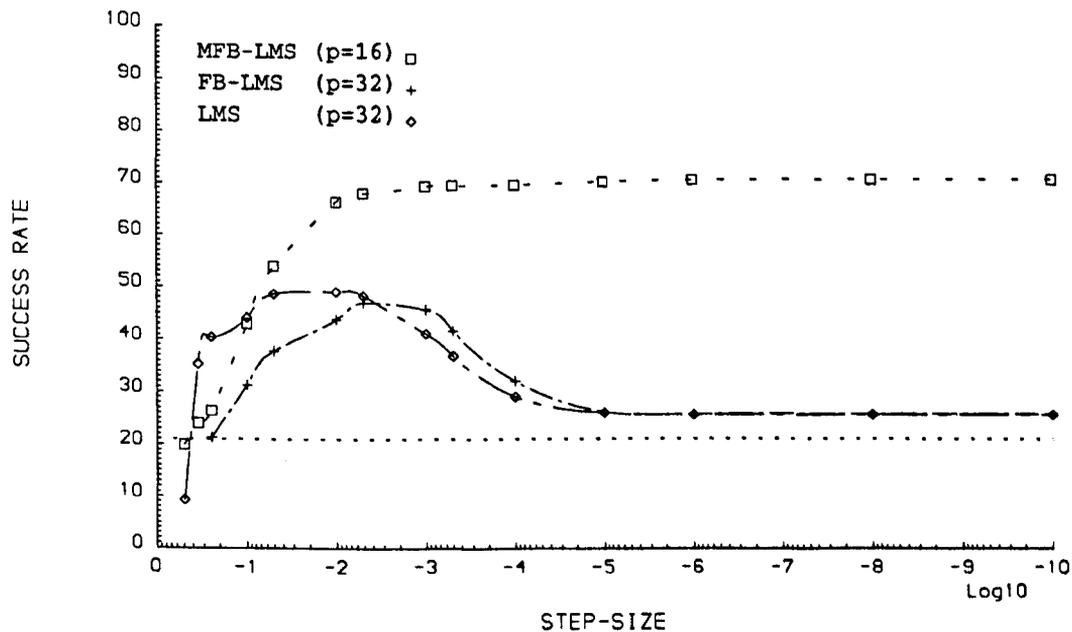


Fig. 5.10 Success rate with the LMS predictors over an operating range of 992 pixels. The dashed bottom line refers to the direct visibility method.

Fig. 5.10 shows the success rate of the LMS predictors. Whereas the standard and FB-LMS depend strongly on the step-size, the MFB-LMS gives constant, superior performance for a very large range of μ , without increasing the computational cost of the standard LMS and using only half the number of filter weights.

Fig. 5.11 shows that the MFB-LMS gets better as the filter order is increased, whereas the standard LMS does not and becomes even more sensitive to small changes of the step-size. The MFB-LMS is not affected if the weights are not updated during the sensing phase; as the operation count is halved, the algorithm becomes computationally equivalent to a matched filter.

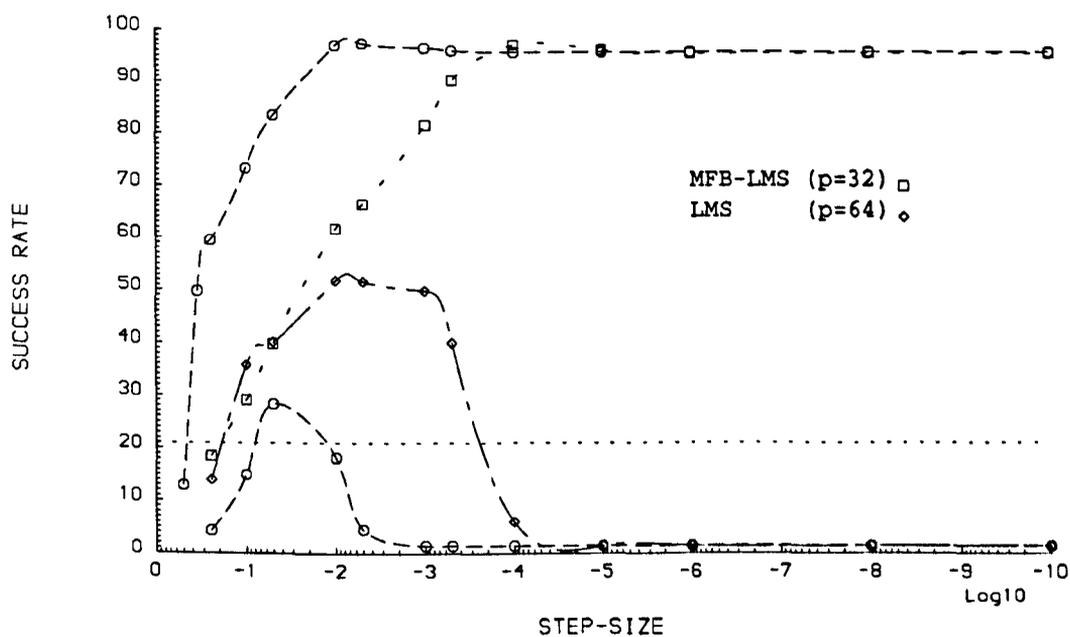


Figure 5.11 Success rate with the LMS predictors over an operating range of 960 pixels. In circles is the rate with weights fixed during the sensing phase.

The performance of the standard RLS algorithm is shown in Fig. 5.12. Comparing with the standard LMS in Figs. 5.10-5.11 it is clear that the RLS with the stationary assumption ($\lambda = 1$) offers a serious challenge to the LMS with $\mu = \mu_0$, as was the case before introducing the thresholding technique.

Apart from its computational complexity, the standard RLS without forgetting may thus be a safer choice than the standard LMS, unless the noise level is low enough to cause ill-conditioning problems. With weights fixed during the sensing phase the RLS may actually perform much better than the LMS, on condition that λ is set to a value slightly less than one during the initial convergence phase. A

modified forward-backward RLS scheme is also a possibility.

The performance of the matched filter is shown in Fig. 5.13. Also shown is the case where the template is matched against the output of the MFB-LMS predictor.

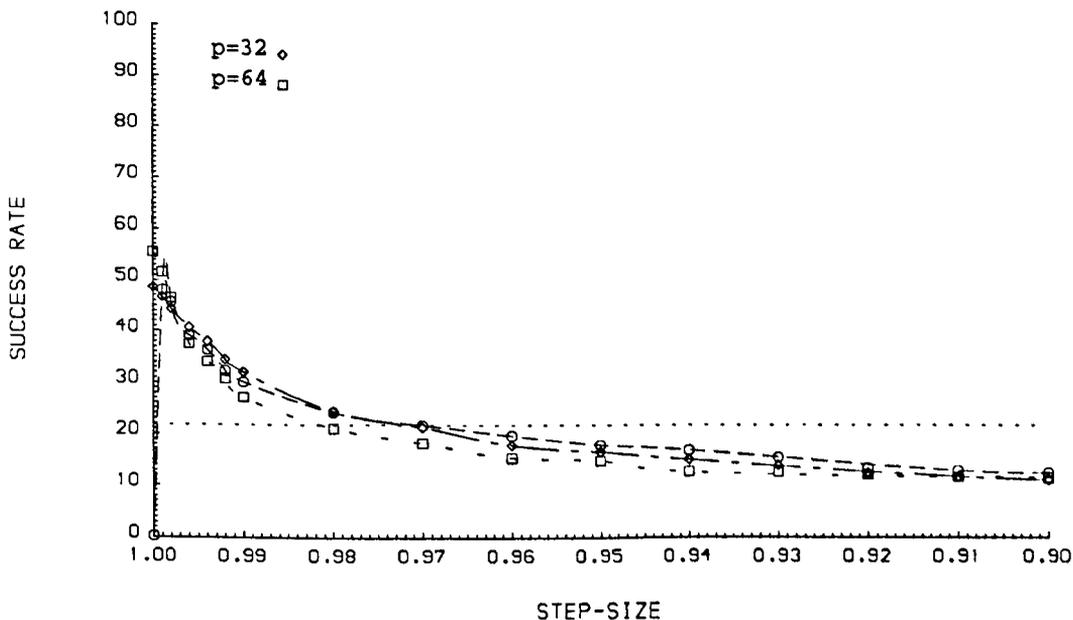


Fig. 5.12 Success rate with the RLS predictor over operating ranges of 992 pixels ($p = 32$) and 960 pixels ($p = 64$). In circles is the rate with the 64 weights fixed during the sensing phase.

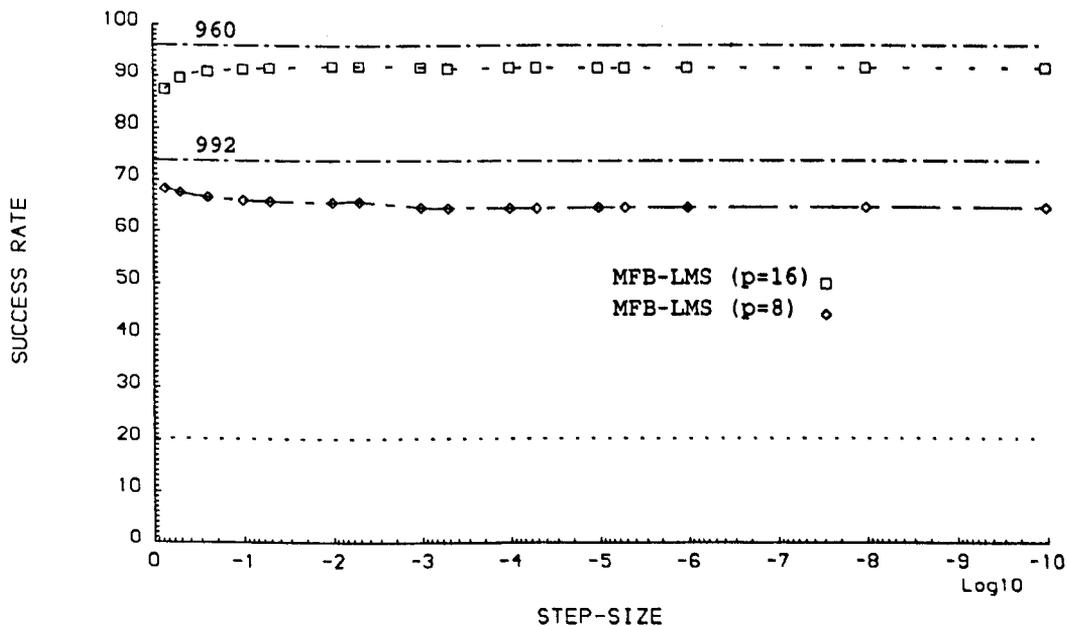


Fig. 5.13 Success rate with the matched filter over operating ranges of 992 pixels ($p = 32$) and 960 pixels ($p = 64$). The curves refer to the combined method.

In order to retain the operating range and the computational speed of either method when used alone, the filter order and the template length were halved, and the weights were kept fixed during the sensing phase.

The matched filter on its own is more successful because the length of its input vector can be double that of the combined scheme for a given operating range. The lead, however, gradually disappeared as the operating range was reduced further by increasing the template length.

Table 5.1 reports the time taken by the different algorithms for the processing of one full fringe pattern ⁸. Time for signal/noise generation and storage into a temporary data file was not included, but that for on-line retrieval of the data from the same file was. Two operating ranges were considered and, for the adaptive predictors, the fixed weights case is also given in brackets. Since on time-sharing machines the elapsed time may not be a reliable measure of the speed of an algorithm, the figures in the table are given as averages over 3000 consecutive scans.

Algorithm	992 Pixels	960 Pixels
LMS	103.3 (98.2)	111.9 (105.2)
FB	104.2 (98.8)	114.2 (105.4)
MFB	103.1 (96.7)	110.7 (102.8)
RLS	324.6 (99.7)	986.4 (105.7)
Corr	96.4	102.8
LMS+Corr	101.7 (98.7)	105.8 (101.6)
MFB+Corr	106.1 (105.5)	109.6 (108.3)
Covariance	95.1	104.5
Centroid	85.3	85.4

Table 5.1 Computer time (in *ms*) for sub-fringe identification of a WLI output with $L_c = 17$ fringes, $b = 12.7$ pixels/fringe, and $N = 1024$ CCD array data.

The centroid was the fastest of all, followed in second place by the covariance algorithm with combined forward and backward predictions and by the matched filter, which were about 12 % slower over 992 pixels and 21 % slower over 960 pixels. With the standard LMS, however, the weight update equation accounted for only

⁸ *Corr* refers to the matched filter algorithm.

4.9 % of the total processing time when $p = 32$, and for 6.0 % of the time when $p = 64$; similar figures were obtained with the two forward-backward LMS schemes. Moreover, the matched filter with $p = 32$ was 1 % slower than the covariance algorithm with $p = 16$, and only 2 % faster when the filter orders were doubled, although the latter required $O(np + p^2)$ more operations if n was the number of samples above the threshold. The FB-LMS scheme with weight updating was also just 1 % slower than the other two LMS schemes with $p = 32$, and 2 to 3 % slower with $p = 64$, although it required twice the number of operations.

It seems that with all the algorithms in the table but the centroid most of the computer time is spent on shift operations rather than on additions and multiplications. Adaptive algorithms need to shift the elements of the input vector to accommodate a new sample and discard an old one; of the two batch algorithms in the table, the covariance also needs updating of the power level if its operation is to be confined to the thresholded region. Arithmetic operations are mainly confined to the region within the points of inflexion of the visibility profile when the thresholding technique is used, whereas shifting has to be carried out at all times. It is only with the RLS algorithm that the main slowing factor becomes the weight update equation, which requires a total of $O(np^2)$ operations and accounts for over 69 % of the total processing time when $p = 32$ and 89 % when $p = 64$.

With a longer coherence length and/or fringe width more of the total time will be spent on real calculations and relatively less on shift operations. In fact, although the RLS with $p = 32$ is only 3 times slower than the LMS in Table 5.1, it was 27 times slower⁹ when processing whole temporally-scanned patterns like the one in Fig. 4.13 with $p = 16$.

As a batch algorithm can only start after all the data have been digitally stored, adaptive algorithms can offer a real speed advantage, but this is only possible if the processing time per sample remains within the scanning and analog-to-digital conversion rates. From this point of view, the advantage is easily realisable but may not be apparent when using a mechanical scanner, as it may not be important having to wait a few *ms* longer to obtain the position of the central fringe when the scanning itself has taken several seconds.

For applications requiring higher scan rates, on the other hand, an electronic

⁹The Kalman algorithm was 40 times slower than the LMS.

scanner is normally used. Read-out time of many common CCD arrays is around $100\ \mu s$, which translates in a scan time of just over $100\ ms$ for a 1024-pixel line array. Although the centroid algorithm is the fastest in the table, an adaptive algorithm can deliver its answer $85\ ms$ earlier. Obviously, in this case the advantage is more apparent but it may not be easily realisable.

For very fast and repeated measurements of critical parameters such as temperature, pressure or strain, read-out speeds of $1\ \mu s/\text{pixel}$ are available, although the read-out noise also tends to increase. At these levels none of the algorithms in the table can cope, unless implemented on a single-chip digital signal processor (DSP).

These programmable hardware devices are excellent at processing real-time signals, whereas general-purpose microprocessors found on a personal computer are much more effective at processing database and spreadsheet-type applications. Additions and subtractions can be performed as a single-cycle instruction on the latter, but multiply functions may take several hundred cycles if, as usual, they are implemented in software or micro-code as a series of add instructions. On the contrary, DSPs have an on-chip hardware multiplier capable of completing a multiply-and-accumulate operation (MAC) in a single cycle. The multicycle loop input-vector shift operations are also implemented as a single-cycle instruction on a DSP, while the inclusion of on-chip data and program memory reduces delays caused by accessing external memory.

Additionally, a Harvard-type architecture with separate program and data buses allows program fetch instructions to overlap data fetch, avoiding the bottleneck of Von Neumann-type general-purpose architectures, where in order to access some data, the instruction must be fetched first, and only then can the processor fetch the data. Pipelining is an additional method of speeding up the instruction throughput of a DSP; instructions are broken down into stages such as fetching the instruction, decoding it, fetching data, executing the instruction, and storing the result. Many DSPs also have a direct memory access (DMA) controller which relieves the CPU from transferring data to and from memory.

The increasing use of DSP systolic array architectures makes it more efficient to perform vector computations over matrices or arrays of data; in this case, a synchronous array of parallel processors consisting of multiple arithmetic logic units (ALUs) under the supervision of one control unit (CU) can handle single-instruction

and multiple-data (SIMD) streams, i.e., all the ALUs receive the same instruction from the CU but operate on different data sets at the same time.

As a result of such innovations, a 64-weight FIR filter can now be computed in 21 μs on a low-performance 16-bit fixed-point processor such as the TMS320C10 from Texas Instruments (6.25 Mips - millions of instructions per second - version), in 5.8 μs on the TMS320C20 (12.5 Mips version), in 2.6 μs on the TMS320C50 (28.6 Mips version), and in half this time on the 24-bit DSP56300 from Motorola (80 Mips version) [63, 64]; the same computation takes 4.4 μs /2.7 μs on the 32-bit floating-point TMS320C30/C40 processors (16.7 and 25 Mips versions, respectively).

Naturally, at such high speeds one should become concerned with the analog-to-digital conversion aspect, as ADCs/DACs can nowadays be a major bottleneck in many real-time DSP applications [65]. It is true that current 16 and 18-bit ADCs can only convert at sampling rates up to 48 Ksps (Kilosamples per second) corresponding to 20.8 μs /sample, and anything above a TMS320C10 could then be a waste. Similarly, most devices in the 8 to 12-bit range can convert at up to 66-75 Ksps (13-15 μs /sample). However, there is now a whole array of high-speed 8, 10 and 12-bit ADCs on the market that offer sampling rates in the Msps (Megasamples per second) region and sell for 10 US dollars or less. The 8-bit TLC5540 from Texas Instruments, for example, can convert at 40 Msps (25 ns /sample); similarly, the 10-bit TLC876 can convert at 20 Msps (50 ns /sample), and the 12-bit TLC8044 at 6 Msps (167 ns /sample).

On the DSP side itself, advances in VLSI circuit design and fabrication are continually reducing cost, space and power requirements, driving the development of higher and higher performance devices. The new C62x generation of 16-bit fixed-point DSPs from Texas Instruments is fabricated using only 550 thousand logic transistors (in contrast, Intel's Pentium requires about 5 million transistors), can deliver 1600 Mips, execute 400 million MACs per second (i.e., one MAC in 2.5 ns), perform a 1024-point complex fast Fourier transform in 70 μs , and compute 64 outputs of a 24-weight FIR filter in 3.9 μs [66], thanks to a new architecture that consists of multiple execution units running in parallel, which allow to perform up to eight instructions in one single cycle; US pricing is just over 100 dollars. Similarly, a new C67x generation of 32-bit floating-point DSPs can currently deliver 1000 Mflops (millions of floating-point operations per second), and are expected to deliver 3000

Mflops by the end of the decade ¹⁰, at a cost of less than 50 US dollars [67]. Indeed, it may not be far the day when the scan rate will be limited only by CCD technology.

5.9 Discussion

In this chapter it was shown that the slow response of the LMS algorithm, which may preclude its use in spatially-scanned WLI systems because of the small number of samples available and reduce the operating range unnecessarily in temporally-scanned systems, cannot be overcome by simply switching the algorithm off whenever the signal power falls below a preset threshold.

The thresholded MFB-LMS scheme introduced in Section 5.6 was able to achieve the performance limit imposed by the matched filter without added complexity and with very little dependence on the step-size, making it very attractive when a template is either not available initially or becomes unreliable because of slow variations of the system parameters such as wavelength or sampling rate drifts. Combining the matched filter with the MFB-LMS algorithm may also raise the identification rate of either method.

Batch filtering schemes performed worse than adaptive schemes and stopped improving quite early as the order of the filter was increased, even when forward and backward predictions were combined to obtain not only parameter estimates with lower variance but also smoother estimates of the fringe pattern.

The centroid method would be an easy choice if the central fringe was constrained to remain in the middle of the scanning range, but for continuous measurements where the central fringe can move at will from scan to scan, higher identification rates can be achieved with either the thresholded MFB-LMS algorithm, the matched filter, or the two methods combined.

Asymmetries in the output pattern caused by asymmetries of the spectral profile of the source or by difficulties in system alignment affect not only the centroid but also the adaptive filters and the multiple-wavelength method. Simulations showed that the identification rate of the MFB-LMS falls below that of the standard LMS for sharp changes of the visibility profile from the region before the central fringe to the one immediately after it caused by, e.g. chromatic aberrations. These are

¹⁰As a comparison, the TMS320C30/C40 can deliver up to 60 Mflops.

due to dispersion, i.e. to the change of the refractive index n of a material with the wavelength in air λ_0 . Since n can be approximated by an inverse function of even powers of λ_0 , according to the Cauchy equation [68]

$$n = A + \frac{B}{\lambda_0^2} + \frac{C}{\lambda_0^4} \quad (5.10)$$

where A , B and C are fixed for a given material, the symmetric distribution of the source spectrum becomes skewed in any transmissive medium which is not air (e.g., an optical fibre), as shorter wavelengths are reduced more than longer ones.

With the multiple-wavelength method asymmetries may be more acute because more than one source is involved [69]. The use of two or more widely spaced wavelengths also means that the effects of dispersion become more severe if an all-optical-fibre arrangement is used, unless the length of the fibre is limited to a few metres [70]. Another shortcoming of the method is that if the two sources do not have the same coherence length, the effects on the optimum wavelength combination and the minimum attainable SNR can be significant [71]. The numerical determination of the minimum SNR which guarantees a 100 % identification rate in [71] is also based on the assumption that the noise level in the system can be as high as the intensity difference between the central fringe in the central fringe-packet, I_{00} , and the first side fringe in the central fringe-packet, I_{01} , or the central fringe in the first side-packet, I_{10} . In practice, the probability for white Gaussian noise to be contemporaneously negatively-valued at I_{00} and positively-valued at either or both of I_{01} and I_{10} is slightly less than 0.44; this means that, on average, the assumption is too optimistic in at least 43 out of 100 scans ¹¹.

The next chapter will take a brief tour into alternative algorithms based on the LMS and RLS approaches that have not been tested in this thesis but could result useful from an implementation or a performance point of view.

¹¹If $P(A)$ is the probability that an event A occurs, \cap denotes logical AND, \cup denotes logical OR, and A , B and C are independent events, $P[A \cap (B \cup C)] = P[(A \cap B) \cup (A \cap C)] = P(A \cap B) + P(A \cap C) - P[(A \cap B) \cap (A \cap C)] = P(A) \times P(B) + P(A) \times P(C) - [P(A) \times P(B)] \times [P(A) \times P(C)]$. In our case, $P(A) = P(B) = P(C) = 0.5$.

Bibliography

- [1] J. Makhoul, "Linear prediction: a tutorial review", *Proc. IEEE*, vol. 63, no. 4, pp. 561-580, Apr. 1975.
- [2] A. E. Hoerl and R. W. Kennard, "Ridge regression: biased estimation for nonorthogonal problems," *Technometrics*, vol. 12, no. 1, pp. 55-67, Feb. 1970.
- [3] A. E. Hoerl and R. W. Kennard, "Ridge regression: applications to nonorthogonal problems," *Technometrics*, vol. 12, no. 1, pp. 69-82, Feb. 1970.
- [4] G. E. Hinton, "Learning distributed representations of concepts," in *Proc. 8th Annual Conf. Cognitive Science Soc.*, Hillsdale, NJ: Erlbaum, pp. 1-12, 1986.
- [5] J. J. Werner, "Control of drift for fractionally spaced equalizers," *US Patent 4384355*, 1983.
- [6] R. D. Gitlin, H. C. Meadors, Jr., and S. B. Weinstein, "The tap-leakage algorithm: an algorithm for the stable operation of a digitally implemented, fractionally adaptive spaced equalizer," *Bell Syst. Tech. J.*, vol. 61, no. 8, pp. 1817-1839, Oct. 1982.
- [7] W. A. Sethares, D. A. Lawrence, C. R. Johnson, Jr., and R. R. Bitmead, "Parameter drift in LMS adaptive filters," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 34, no. 4, pp. 868-879, Aug. 1986.
- [8] J. M. Cioffi, "Limited-precision effects in adaptive filtering," *IEEE Trans. Circ. Syst.*, vol. 34, no. 7, pp. 821-833, July 1987.
- [9] H. Press and J. W. Tukey, "Power spectral methods of analysis and their application to problems in airplane dynamics," *Bell Syst. Monograph no. 2606*, 1956.

- [10] S. H. Leung, B. L. Chan, and S. M. Lau, "On improving the convergence rate of LMS algorithms using prefiltering," *Proc. Int. Symp. Signal Process. Appl.*, Gold Coast, Australia, pp. 311-314, Aug. 1990.
- [11] S. H. Leung and B. L. Chan, "Fast converging LMS equaliser," *Electron. Lett.*, vol. 29, no. 10, pp. 829-830, May 1993.
- [12] E. R. Ferrara, Jr., and B. Widrow, "The time-sequenced adaptive filter," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 29, no. 3, pp. 679-683, June 1981.
- [13] T. I. Haweel and P. M. Clarkson, "A class of order statistic LMS algorithms," *IEEE Trans. Signal Process.*, vol. 40, no. 1, pp. 44-53, Jan. 1992.
- [14] B. Widrow, J. Glover, J. McCool, J. Kannitz, C. Williams, R. Hearn, J. Zeidler, E. Dong, Jr., and R. Goodin, "Adaptive noise cancelling: principles and applications", *Proc. IEEE*, vol. 63, no. 12, pp. 1692-1716, Dec. 1975.
- [15] G. W. Snedecor and W. G. Cochran, *Statistical Methods*, 7th ed. Ames, Iowa: Iowa State University Press, 1980.
- [16] O. L. Frost, III, *Private communication* to J. R. Treichler, in *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 27, no. 1, pp. 53-62, Feb. 1979.
- [17] T. J. Ulrych and R. W. Clayton, "Time series modelling and maximum entropy," *Phys. Earth Planet. Int.*, vol. 12, pp. 188-200, Aug. 1976.
- [18] J. R. Treichler, *The Spectral Line Enhancer, the Concept, and Implementation, and an Application*. PhD thesis, Stanford Univ., Stanford, CA, May 1977.
- [19] M. A. Jenkins and J. F. Traub, "Zeros of a complex polynomial," *Communications of the ACM*, vol. 15, pp. 97-99, 1972.
- [20] D. H. Withers, "Remark on algorithm 419," *Communications of the ACM*, vol. 17, p. 157, 1974.
- [21] M. B. Priestley, *Spectral Analysis and Time Series*, vol. 1. London: Academic, 1981.
- [22] J. R. Treichler, "Transient and convergent behaviour of the adaptive line enhancer", *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 27, no. 1, pp. 53-62, Feb. 1979.

- [23] C. F. N. Cowan and P. M. Grant, *Adaptive Filters*. Englewood Cliffs, NJ: Prentice-Hall, 1985.
- [24] J. L. Moschner, "Adaptive filter with clipped input data," *Rep. 6796-1*, Inform. Systems Lab., Stanford Univ., Stanford, CA, June 1970.
- [25] A. Jurkevics and T. J. Ulrych, "Representing and simulating strong ground motion," *Bull. Seism. Soc. Amer.*, vol. 68, no. 3, pp. 781-801, June 1978.
- [26] W. C. Lee, S. T. Park, I. W. Cha, and D. H. Youn, "Adaptive spatial domain forward-backward predictors for bearing estimation," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 38, no. 7, pp. 1105-1109, July 1990.
- [27] Y. C. Lim and C. C. Ko, "Forward-backward LMS adaptive line enhancer," *IEEE Trans. Circ. Syst.*, vol. 37, no. 7, pp. 936-940, July 1990.
- [28] J. I. Nagumo and A. Noda, "A learning method for system identification," *IEEE Trans. Autom. Contr.*, vol. 12, pp. 282-287, June 1967.
- [29] A. E. Albert and L. S. Gardner, *Stochastic Approximation and Nonlinear Regression*. Cambridge, MA: MIT Press, 1967.
- [30] J. A. Chambers, "Normalised order statistic LMS adaptive filtering," *IEE Colloquium*, Savoy Place, London, May 18, Digest no. 1994/125, 1994.
- [31] D. T. M. Slock, "On the convergence behaviour of the LMS and the normalized LMS algorithms," *IEEE Trans. Signal Process.*, vol. 41, no. 9, pp. 2811-2825, Sept. 1993.
- [32] L. J. Griffiths, "A continuously adaptive filter implemented as a lattice structure," *Proc. Int. Conf. Acoust., Speech, Signal Process.*, Hartford, CT., pp. 683-686, May 1977.
- [33] J. Makhoul, "Stable and efficient lattice structures for linear prediction," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 25, no. 5, pp. 423-428, Oct. 1977.
- [34] F. Itakura and S. Saito, "Digital filtering techniques for speech analysis and synthesis," in *Proc. 7th Int. Conf. Acoust.*, Budapest, Hungary, pp. 261-264, 1971.

- [35] R. E. Crochiere, "Digital lattice structures and coefficient sensitivity," *IEEE Trans. Audio Electroacoust.*, vol. 20, no. 4, pp. 240-246, Oct. 1972.
- [36] A. H. Gray, Jr., and J. D. Markel, "Digital lattice and ladder filter synthesis," *IEEE Trans. Audio Electroacoust.*, vol. 21, no. 6, pp. 491-500, 1973.
- [37] W. S. Hodgkiss and J. A. Presley, "Adaptive tracking of multiple sinusoids whose power levels are widely separated," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 29, no. 3, pp. 710-721, June 1981.
- [38] L. J. Griffiths and R. S. Medough, "Convergence properties of an adaptive noise canceling lattice structure," *Proc. Int. Conf. Decision and Control*, San Diego, CA, pp. 1357-1361, Jan. 1979.
- [39] M. L. Honig and D. G. Messerschmitt, "Convergence properties of an adaptive digital lattice filter," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 29, no. 3, pp. 642-653, June 1981.
- [40] S. Roy and J. J. Shynk, "Analysis of the momentum LMS algorithm," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 38, no. 12, pp. 2088-2098, Dec. 1990.
- [41] R. H. Kwong and E. W. Johnston, "A variable step size LMS algorithm," *IEEE Trans. Signal Process.*, vol. 40, no. 7, pp. 1633-1642, July 1992.
- [42] R. W. Harris, D. M. Chabries, and F. A. Bishop, "A variable step (VS) adaptive filter algorithm," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 34, no. 2, pp. 309-316, Apr. 1986.
- [43] J. B. Evans and B. Liu, "Variable step size methods for the LMS adaptive algorithm," *Proc. IEEE Int. Symp. Circ. Syst.*, pp. 422-425, Apr. 1987.
- [44] J. B. Evans, P. Xue, and B. Liu, "Analysis and implementation of variable step size adaptive algorithms," *IEEE Trans. Signal Process.*, vol. 41, no. 8, pp. 2517-2535, Aug. 1993.
- [45] T. J. Shan and T. Kailath, "Adaptive algorithms with an automatic gain control feature," *IEEE Trans. Circ. Syst.*, vol. 35, no. 1, pp. 122-127, Jan. 1988.

- [46] S. Karni and G. Zeng, "A new convergence factor for adaptive filters," *IEEE Trans. Circ. Syst.*, vol. 36, no. 7, pp. 1011-1012, July 1989.
- [47] D. C. Fraser and J. E. Potter, "The optimum linear smoother as a combination of two optimum linear filters," *IEEE Trans. Autom. Contr.*, vol. 14, no. 4, pp. 387-390, Aug. 1969.
- [48] E. A. Robinson and M. T. Silvia, *Digital Signal Processing and Time Series Analysis*. San Francisco, CA: Holden-Day, 1978.
- [49] S. Chen, K. T. V. Grattan, B. T. Meggitt, and A. W. Palmer, "Instantaneous fringe-order identification using dual broad-band source with widely spaced wavelengths," *Electron. Lett.*, vol. 29, no. 4, pp. 334-335, Feb. 1993.
- [50] P. Hariharan, "Interferometry with Lasers," in *Progress in Optics XXIV*, E. Wolf, ed. New York: Elsevier Science, 1987.
- [51] S. Chen, A. W. Palmer, K. T. V. Grattan, and B. T. Meggitt, "Digital signal-processing techniques for electronically scanned optical-fiber white-light interferometry," *Appl. Opt.*, vol. 31, no. 28, pp. 6003-6010, Oct. 1992.
- [52] G. L. Turin, "An introduction to matched filters," *IRE Trans. Inform. Theory*, vol. 6, pp. 311-329, 1960.
- [53] A. G. Favret, "Computer matched filter location of fetal R-waves," *Medical Biolog. Eng.*, vol. 6, pp. 467-475, 1968.
- [54] K. G. Lindcrantz and H. Lilja, "New software QRS detector algorithm suitable for real time applications with low signal-to-noise ratios," *J. Biomed. Eng.*, vol. 10, pp. 280-284, May 1988.
- [55] C. H. Knapp and G. C. Carter, "The generalized correlation method for estimation of time delay," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 24, no. 4, pp. 320-327, Aug. 1976.
- [56] G. C. Carter, "Coherence and time delay estimation," *Proc. IEEE*, vol. 75, no. 2, pp. 236-255, Feb. 1987.

- [57] P. A. Nielson and J. B. Thomas, "Effect of correlation on signal detection in arctic under-ice noise," *22nd Asilomar Conf. Signals, Systems and Computers*, Pacific Grove, CA, pp. 445-450, Nov. 1988.
- [58] J. T. Rickard, M. J. Dentino, and J. R. Zeidler, "Detection performance of an adaptive processor in nonstationary noise," *Proc. Int. Conf. Acoust., Speech, Signal Process.*, Washington, DC, pp. 136-139, Apr. 1979.
- [59] F. A. Reed, P. L. Feintuch, and N. J. Bershad, "Time delay estimation using the LMS adaptive filter - static behaviour," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 29, no. 3, pp. 561-571, June 1981.
- [60] P. L. Feintuch, N. J. Bershad, and F. A. Reed, "Time delay estimation using the LMS adaptive filter - dynamic behaviour," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 29, no. 3, pp. 571-576, June 1981.
- [61] Y. T. Chan, J. M. F. Riley, and J. B. Plant, "Modeling of time delay and its application to estimation of nonstationary delays," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 29, no. 3, pp. 577-581, June 1981.
- [62] S. Chen, A. W. Palmer, K. T. V. Grattan, and B. T. Meggitt, "Fringe order identification in optical fibre white-light interferometry using centroid algorithm method," *Electron. Lett.*, vol. 28, no. 6, pp. 553-555, Mar. 1992.
- [63] Texas Instruments, *Digital Signal Processing Product Information*.
At <http://www.ti.com/sc/docs/dsps/products.htm>.
- [64] Motorola, *Semiconductor Product Information*.
At <http://www.mot.com/SPS/DSP/documentation/index.html>.
- [65] E. C. Ifeachor and B. W. Jervis, *Digital Signal Processing: a Practical Approach*. Reading, MA: Addison-Wesley, 1993.
- [66] Texas Instruments, *Showcase Magazine*, issue 31, Nov. 1997.
- [67] Texas Instruments, *Details on Signal Processing*, Jan. 1998.
- [68] F. A. Jenkins and H. E. White, *Fundamentals of Optics*, 4th ed. New York: McGraw-Hill, 1976.

- [69] D. N. Wang, Y. N. Ning, K. T. V. Grattan, A. W. Palmer, and K. Weir, "Three-wavelength combination source for white-light interferometry," *IEEE Photon. Technol. Lett.*, vol. 5, no. 11, pp. 1350-1352, Nov. 1993.
- [70] D. N. Wang, K. T. V. Grattan, and A. W. Palmer, "Dispersion effect analysis with multiwavelength combination sources in optical sensor applications," *Optics Commun.*, vol. 127, pp. 19-24, 1996.
- [71] D. N. Wang, Y. N. Ning, K. T. V. Grattan, A. W. Palmer, and K. Weir, "The optimized wavelength combinations of two broadband sources for white light interferometry," *J. Lightwave Technol.*, vol. 12, no. 5, pp. 909-916, May 1994.

Chapter 6

Alternative LMS and RLS Schemes for WLI

6.1 Introduction

In Chapter 5 a few non-standard versions of the LMS algorithm were tried in an attempt to accelerate convergence and/or tracking while keeping the steady-state error at a minimum.

This chapter presents a brief round-up of LMS and RLS-based algorithms that were not applied here but may still be attractive for WLI central fringe identification either because they are easier to implement or execute faster than the standard version, or because they are more complex but also more stable and robust to numerical errors.

6.2 Algorithms for AR Modelling

In *block* LMS algorithms [1] the instantaneous gradient estimates are averaged over p samples before the weight vector is updated, producing a less noisy adaptation. Performance advantages have been claimed (see, e.g. [2, 3]) although the tracking speed may be affected. *Fast* LMS filters [4, 5] using the complex LMS algorithm [6] allow faster execution by replacing convolution operations in the time domain with multiplications in the frequency domain using the fast Fourier transform [7]. Exploiting the eigenvalue-spectrum decomposition, each frequency bin can be updated independently using a normalised step-size to equalise the time constants [8, 9, 10].

As with the GAL in Section 5.5.2, however, residual coupling between components may also be present in the frequency domain [11].

The *sliding window* RLS algorithm [12] is more effective than the exponentially weighted RLS in dealing with sudden changes in the statistics of the data and with impulsive noise. It may therefore recover more quickly from signal discontinuities at the power threshold points, although at a computational cost of around $5p^2 + 6p$ adds and multiplies per iteration.

QR-decomposition recursive least squares (QRD-RLS) algorithms [13, 14] replace the normal equations based on the autocorrelation matrix \mathbf{R} by a set of equations based on an orthogonal transformation of the data matrix \mathbf{X} [15]. As the condition number is square-rooted, the solution is numerically very stable and less susceptible to round-off errors. The use of matrix operations in the form of Givens rotations during the decomposition [16] allows for a high degree of parallelism, which can be exploited by a systolic array architecture [17], whose modularity and pipelined mode of operation makes it especially suitable for fast VLSI implementation [18]. Early hardware realisations [19, 20] seem to offer outstanding numerical performance.

Fast RLS algorithms update the gain vector by exploiting the shifting property of the input vector and the redundancy in the Toeplitz structure of the data matrix, reducing both computational load and memory storage from $O(p^2)$ to $O(p)$. The introduction of redundancy assumptions, however, makes them more vulnerable to round-off errors.

The RLS predictors in [21, 22] require only $8p$ operations but are unstable for $\lambda < 1$ [23, 24, 25, 26], with numerical errors increasing exponentially at the rate $1/\lambda$. The fast transversal filters (FTF) in [27, 26, 28] are relatively more stable and require from $5p$ to $9p$ operations in their predictor form, depending on the implementation; sliding-window forms have been derived in [28]. Rescuing procedures to stop the build-up of round-off errors and eventual divergence can be incorporated, such as periodic resetting of some critical internal variables to zero while updating the weights with a LMS algorithm [24, 26, 29], or compensation using error feedback to reintroduce computational redundancy with extra p operations [30, 31]. Although unlikely to raise the identification rate above that of the standard RLS, fast algorithms should be considered when computational constraints would dictate

the use of a LMS algorithm.

Fast QRD-RLS algorithms [32, 33, 34, 35] combine the numerical stability and modularity of QRD algorithms with the low operation count of fast algorithms.

Lattice filters are numerically more robust than transversal filters, and due to their modular structure they are highly suitable for fast VLSI implementation. As they yield all lower-order solutions by default [36] they can be made to self-adjust the number of weights according to the magnitude of the prediction error in real-time, without having to run a bank of filters of different orders in parallel.

Fast least squares lattice (LSL) algorithms [37, 38, 39] are the lattice counterpart of FTF algorithms. Although the derivation of the AR parameters from the reflection coefficients requires $O(p^2)$ operations, their explicit computation is not needed if a power spectrum estimate is not required; computation of the prediction errors and updating of the reflection coefficients for prediction is thus possible with twice the number of operations required by the FTF algorithms, challenging the GAL algorithm, which requires from $6p$ to $10p$ operations depending on the implementation. In fact, LSL filters have been shown to offer better performance than the GAL in some applications [40, 41].

Fast QRD-LSL algorithms [42, 43, 44] also require only $O(p)$ operations.

6.3 Algorithms for ARMA Modelling

It was shown in Chapter 2 that an ARMA model should provide at least as good a representation of the fringe pattern as the AR approximation, if not better. Although an AR polynomial of high order can compensate for the absence of the MA polynomial, allowing the latter to be different from unity may mean that the same estimation capability could be achieved with fewer parameters. In fact, rather than increasing the AR order until a white error sequence is produced, one could use the information contained in this sequence directly, as in Eq. 2.9.

All of the filtering algorithms considered in Chapter 5 were able to produce error sequences with flat power spectra and with no significant autocorrelations [45] for moderately high orders. However, increasing the filter order increases the amount of ill-conditioning and the danger of numerical problems for RLS schemes, while decreasing the speed of adaptation of the parameters for LMS schemes. With the

LMS, in particular, more scans were needed before a satisfactory white sequence could be obtained.

Adaptive algorithms that solve for the ARMA parameters include the recursive extended least squares (RELS) [46, 47] and the recursive maximum likelihood (RML) [48, 49], which can also be simplified to produce *generalised* LMS algorithms. Alternatively, the recursive instrumental variable (RIV) method [50, 51, 52, 53] provides unbiased AR estimates, but the MA parameters have to be estimated separately (see, e.g. [54]).

Lattice predictors have also been derived (see, e.g. [55, 56, 36, 57]). As a final warning, however, it remains difficult to ensure robust performance of adaptive IIR filters in many applications [58], although they provide natural inverses for ARMA models. In particular, instability during adaptation and the presence of a non-unimodal error surface are still open problems [59, 60], as are the effects of model order mismatch between true and assumed system [61].

6.4 Modelling with Coloured Noise

The white noise assumption led to the development of a stochastic ARMA(p, p) model for the generation of the observed fringe pattern in Chapter 2. Although shot and thermal noise in electronic components are believed to be white Gaussian, vibrations during scanning in temporally scanned systems are responsible for the appearance of narrow-band noise [62]. Similarly, measurement noise such as that at the CCD detector can be described as a broad-band Gaussian process.

Assuming, eg., that the noise is the MA(1) process $w[n] + c_1 w[n - 1]$, with $w[n]$ white, expanding Eq. 2.8 one obtains an ARMA($p, p + 1$) process driven by $w[n]$. This is equivalent to inserting a $C(z)$ block with transfer function $1 + c_1 z^{-1}$ between $w[n]$ and the $B(z)$ block in Fig. 2.1. In general, if the noise follows an ARMA process of order (p_1, q_1) , the output will be ARMA($p + p_1, \max(p_1, p + q_1)$) [63]. An inverse filter implemented as a whitening filter will attempt to recover the white noise creating the coloured process. Testing for whiteness of the error sequence would not be a sound diagnostic procedure in this case.

The same conclusion can be reached by noting that $E\{v[n]\hat{s}[n]\}$ in Eq. 2.12 is not necessarily zero anymore, as some of the $y[n - i]$ terms that form $\hat{s}[n]$ in Eq.

2.9 depend on those same $w[n - i]$ terms that generated $v[n]$. Hence, the MSE criterion in Eq. 2.12 is still quadratic with a unique minimum point, but reaching this point does not guarantee the equality between the noise-free fringe signal and its estimate.

Replacing white noise with the broad-band MA(3) noise process ¹ $v[n] = w[n] + w[n - 1] + w[n - 2] + w[n - 3]$, the behaviour of all the algorithms considered in Chapter 5 deteriorated. The success rate of the thresholded MFB-LMS scheme with weights fixed after ten preliminary scans went down ² from 97.8 % to 88.7 %. Similarly, the success rate of the matched filter with $p = 76$ went down from 98.1 % to 85.9 %. Even greater degradation was observed when adding the narrow-band AR(2) process $v[n] = 1.83v[n - 1] - 0.97v[n - 2] + w[n]$. This would simulate narrow-band noise centred at 60 Hz and with a 3-dB bandwidth of 5 Hz.

The RELS, RML, and RIV algorithms of the previous section can be used if one assumes that the fringe pattern is an AR process to which MA noise has been added. For an MA(m) noise process one could also replace the usual one-step prediction with multi-step prediction, a special case of the recursive instrumental variable method which exploits the fact that the optimum prediction step for noise suppression lies somewhere between the correlation distance of the noise, m , and that of the signal [64]. Replacing Eq. 2.9 by

$$\hat{s}[n] = - \sum_{i=1}^p \hat{a}_i y[n - (i + l)] + \sum_{i=1}^p \hat{a}_i \hat{v}[n - (i + l)] \quad (6.1)$$

$v[n]$ is uncorrelated with $\hat{s}[n]$ in Eq. 2.12, for any $l \geq m$. Improved frequency estimation for narrow-band signals in coloured noise has been reported in, e.g. [65, 66].

Unfortunately, multi-step prediction did not give any performance improvement over one-step prediction. The identification rate of the thresholded MFB-LMS scheme went down from 88.7 % to 76.0 % when 4-step instead of 1-step prediction was used to deal with the broad-band MA(3) noise process above. The problem by going back in time is that, although the correlation between present and past noise values fades away, that between present and past signal values also deteriorates

¹With, e.g. 12.7 samples per fringe and a sampling rate of 1 KHz, the interference signal would be centred at 79 Hz, while the noise spectrum would be approximately constant over (0,250) Hz, with a secondary lobe with cut-offs at 250 and 500 Hz.

²With $p = 38$, a SNR of 20 dB, and for all $\mu < 10^{-2}$.

rapidly when the signal is non-stationary. In addition, it may be difficult to decide a priori on the correlation distance of the noise.

The technique will also work less well with narrow-band AR noise, as in such case the autocorrelation sequence of the noise is a mixture of damped exponentials and/or damped sine waves, with no definite cut-off for the correlation distance (see Section 2.5). It is clear, however, that as the second-order statistics of the noise process become more and more similar to those of the interference signal, any filtering algorithm based on second-order statistics alone will find it increasingly difficult to discern and separate the signal from the noise.

As an aside, multi-step prediction has been shown to improve the estimation capability of an adaptive filter even in the white noise case ($m = 0$), as for complete cancellation of the sinusoidal components of a noisy signal from the error term $e[n]$ it is necessary that such components appear at the filter output with not only the same amplitude, but especially with the same phase as the corresponding components in the signal [67]. According to the theoretical analysis and practical formulae in [68, 69, 70], 4-step or 10-step prediction should be optimum for a WLI interference signal with a fringe width of 12.7 pixels per fringe. The identification rate of the LMS-based algorithms in Chapter five, however, declined progressively as the prediction step was increased from one. The highest degradation was observed at high SNRs, where the convergence and tracking capability are least and the most recent information is desperately needed.

In the next chapter the assumption that the WLI output can be completely described in terms of its second-order statistics will be questioned, and the linearity and Gaussianity constraints will be relaxed to take into account higher-order properties characteristic of linear non-Gaussian or non-linear processes.

Bibliography

- [1] G. A. Clark, S. K. Mitra, and S. R. Parker, "Block implementation of adaptive digital filters," *IEEE Trans. Circ. Syst.*, vol. 28, no. 6, pp. 584-592, June 1981.
- [2] A. Feuer, "Performance analysis of the block least mean square algorithm," *IEEE Trans. Circ. Syst.*, vol. 32, no. 9, pp. 960-963, Sept. 1985.
- [3] L. S. L. Hsieh and S. L. Wood, "Performance analysis of time domain block LMS algorithms," *Proc. Int. Conf. Acoust., Speech, Signal Process.*, Minneapolis, MN, vol. 3, pp. 535-538, Apr. 1993.
- [4] M. Dentino, J. McCool, and B. Widrow, "Adaptive filtering in the frequency domain," *Proc. IEEE*, vol. 66, no. 12, pp. 1658-1659, Dec. 1978.
- [5] E. R. Ferrara, "Fast implementation of LMS adaptive filters," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 28, no. 4, pp. 474-475, Aug. 1980.
- [6] B. Widrow, J. McCool, and M. Ball, "The complex LMS algorithm," *Proc. IEEE*, vol. 63, no. 4, pp. 719-720, Apr. 1975.
- [7] J. W. Cooley and J. W. Tukey, "An algorithm for the machine calculation of complex Fourier series," *Math. Comput.*, vol. 19, no. 90, pp. 297-301, Apr. 1965.
- [8] N. J. Bershad and P. L. Feintuch, "A normalized frequency domain LMS adaptive algorithm," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 34, no. 3, pp. 452-461, June 1986.
- [9] D. Mansour and A. H. Gray, Jr., "Unconstrained frequency domain adaptive filter," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 30, no. 5, pp. 726-734, Oct. 1982.

- [10] G. Picchi and G. Prati, "Self-orthogonalizing adaptive algorithm for channel equalization in the discrete frequency domain," *IEEE Trans. Commun.*, vol. 32, no. 4, pp. 371-379, Apr. 1984.
- [11] J. C. Lee and C. K. Un, "Performance analysis of frequency-domain block LMS adaptive digital filters," *IEEE Trans. Circ. Syst.*, vol. 36, no. 2, pp. 173-188, Feb. 1989.
- [12] G. C. Goodwin and R. L. Payne, *Dynamic System Identification: Experiment Design and Data Analysis*. New York: Academic, 1977.
- [13] W. M. Gentleman and H. T. Kung, "Matrix triangularization by systolic arrays," *Proc. SPIE Int. Conf. Signal Process.*, vol. 298 (Real Time Signal Processing IV), pp. 298-303, 1981.
- [14] J. G. McWhirter, "Recursive least-squares minimization using a systolic array," *Proc. SPIE Int. Conf. Signal Process.*, vol. 431 (Real Time Signal Processing VI), pp. 105-112, 1983.
- [15] A. S. Householder, *The Theory of Matrices in Numerical Analysis*. Waltham, MA: Blaisdell, 1964.
- [16] W. Givens, "Computation of plane unitary rotations transforming a general matrix to triangular form," *J. Soc. Ind. Appl. Math.*, vol. 6, pp. 26-50, 1958.
- [17] C. Mead and L. Conway, *Introduction to VLSI Systems*. Reading, MA: Addison-Wesley, 1980.
- [18] S. Y. Kung, H. J. Whitehouse, and T. Kailath, eds., *VLSI and Modern Signal Processing*. Englewood Cliffs, NJ: Prentice-Hall, 1985.
- [19] B. Medoff, Nov. 1986, and T. J. Shepard, 1986, *private communications* to J. M. Cioffi, in *IEEE Trans. Circ. Syst.*, vol. 34, no. 7, pp. 821-833, July 1987.
- [20] H. Leung and S. Haykin, "Stability of recursive QRD-LS algorithms using finite-precision systolic array implementation," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 37, no. 5, pp. 760-763, May 1989.

- [21] M. Morf, T. Kailath, and L. Ljung, "Fast algorithms for recursive identification," in *Proc. Int. Conf. Decision and Control*, Clearwater Beach, FL, pp. 916-921, Dec. 1976.
- [22] L. Ljung, M. Morf, and D. D. Falconer, "Fast calculation of gain matrices for recursive estimation schemes," *Int. J. Control*, vol. 27, no. 1, pp. 1-19, Jan. 1978.
- [23] S. Ljung, *Fast Algorithms for Integral Equations and Least Squares Identification Problems*. PhD thesis no. 93, Dept. Electr. Eng., Linköping Univ., Linköping, Sweden, 1983.
- [24] D. W. Lin, "On digital implementation of the fast Kalman algorithm," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 32, no. 5, pp. 998-1005, Oct. 1984.
- [25] S. Ljung and L. Ljung, "Error propagation properties of recursive least squares adaptation algorithms," *Automatica*, vol. 21, no. 2, pp. 157-167, Mar. 1985.
- [26] J. M. Cioffi and T. Kailath, "Fast, recursive-least-squares transversal filters for adaptive filtering," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 32, no. 2, pp. 304-337, Apr. 1984.
- [27] G. Carayannis, D. G. Manolakis, and N. Kalouptsidis, "A fast sequential algorithm for least-squares filtering and prediction," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 31, no. 6, pp. 1394-1402, Dec. 1983.
- [28] J. M. Cioffi and T. Kailath, "Windowed fast transversal filters adaptive algorithms with normalization," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 33, no. 3, pp. 607-625, June 1985.
- [29] E. Eleftheriou and D. D. Falconer, "Tracking properties and steady-state performance of RLS adaptive filter algorithms," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 34, no. 5, pp. 1097-1109, Oct. 1986.
- [30] J. L. Botto and G. V. Moustakides, "Stabilizing the fast Kalman algorithms," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 37, no. 9, pp. 1342-1348, Sept. 1989.

- [31] D. T. M. Slock and T. Kailath, "Numerically stable fast transversal filters for recursive least squares adaptive filtering," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 39, no. 1, pp. 92-114, Jan. 1991.
- [32] J. M. Cioffi, "High speed systolic implementation of fast QR adaptive filters," *Proc. Int. Conf. Acoust., Speech, Signal Process.*, New York, pp. 1584-1588, Apr. 1988.
- [33] J. M. Cioffi, "The fast adaptive rotor's RLS algorithm," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 38, no. 4, pp. 631-653, Apr. 1990.
- [34] M. G. Bellanger, "The FLS-QR algorithm for adaptive filtering," *Signal Process.*, vol. 17, pp. 291-304, 1989.
- [35] I. K. Proudler, J. G. McWhirter, and T. J. Sheperd, "Fast QRD-based algorithms for least squares linear prediction," *Proc. IMA Conf. Math. Signal Process.*, Univ. of Warwick, Coventry, England, 1988.
- [36] B. Friedlander, "Lattice filters for adaptive processing," *Proc. IEEE*, vol. 70, no. 8, pp. 829-867, Aug. 1982.
- [37] M. Morf, A. Vieira, and D. T. Lee, "Ladder forms for identification and speech processing," in *Proc. Int. Conf. Decision and Control*, New Orleans, LA, pp. 1074-1078, Dec. 1977.
- [38] M. Morf and D. T. Lee, "Recursive least squares ladder forms for fast parameter tracking," in *Proc. Int. Conf. Decision and Control*, San Diego, CA, pp. 1362-1367, Jan. 1979.
- [39] D. T. Lee, M. Morf, and B. Friedlander, "Recursive least squares ladder estimation algorithms," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 29, no. 3, pp. 627-641, June 1981.
- [40] E. H. Satorius and S. T. Alexander, "Channel equalization using adaptive lattice algorithms," *IEEE Trans. Commun.*, vol. 27, no. 6, pp. 899-905, June 1979.
- [41] E. H. Satorius and J. D. Pack, "Application of least squares lattice algorithms to adaptive equalization," *IEEE Trans. Commun.*, vol. 29, no. 2, pp. 136-142, Feb. 1981.

- [42] F. Ling, "Efficient least-squares lattice algorithms based on Givens rotation with systolic array implementations," *Proc. Int. Conf. Acoust., Speech, Signal Process.*, Glasgow, Scotland, pp. 1290-1293, May 1989.
- [43] I. K. Proudler, J. G. McWhirter, and T. J. Sheperd, "QRD-based lattice filter algorithms," *SPIE Int. Conf. Signal Process.*, San Diego, CA, 1989.
- [44] P. A. Regalia and G. Bellanger, "On the duality between fast QR methods and lattice methods in least squares adaptive filtering," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 39, no. 4, pp. 879-891, Apr. 1991.
- [45] G. M. Ljung and G. E. P. Box, "On a measure of lack of fit in time series models," *Biometrika*, vol. 65, no. 2, pp. 297-303, 1978.
- [46] P. C. Young, "The use of linear regression and related procedures for the identification of dynamic processes," *Proc. 7th IEEE Symp. on Adaptive Processes*, UCLA, Los Angeles, CA, 1968.
- [47] R. L. Kashyap, "Estimation of parameters in a partially whitened representation of a stochastic process," *IEEE Trans. Autom. Contr.*, vol. 19, no. 1, pp. 13-21, Feb. 1974.
- [48] T. Söderström, "An on-line algorithm for approximate maximum likelihood identification of linear dynamic systems," *Rep. 7308*, Dept. Autom. Contr., Lund Inst. Technol., Lund, Sweden, 1973.
- [49] B. Friedlander, "A recursive maximum likelihood algorithm for ARMA line enhancement," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 30, no. 4, pp. 651-657, Aug. 1982.
- [50] L. Y. Wong and E. Polak, "Identification of linear discrete time systems using the instrumental variable method," *IEEE Trans. Autom. Contr.*, vol. 12, no. 6, pp. 707-718, Dec. 1967.
- [51] P. C. Young, "An instrumental variable method for real-time identification of a noisy process," *Automatica*, vol. 6, pp. 271-287, 1970.
- [52] B. Friedlander, "The overdetermined recursive instrumental variable method," *IEEE Trans. Autom. Contr.*, vol. 29, no. 4, pp. 353-356, Apr. 1984.

- [53] B. Porat and B. Friedlander, "The square-root overdetermined recursive instrumental variable algorithm," *IEEE Trans. Autom. Contr.*, vol. 34, no. 6, pp. 656-658, June 1989.
- [54] B. Friedlander, "Instrumental variable methods for ARMA spectral estimation," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 31, no. 2, pp. 404-415, Apr. 1983.
- [55] M. Morf, D. T. Lee, J. R. Nickolls, and A. Vieira, "A classification of algorithms for ARMA models and ladder realizations," in *Proc. Int. Conf. Acoust., Speech, Signal Process.*, Hartford, CT, May 1977.
- [56] D. T. Lee, B. Friedlander, and M. Morf, "Recursive ladder algorithms for ARMA modelling," *IEEE Trans. Autom. Contr.*, vol. 27, no. 8, pp. 753-764, Aug. 1982.
- [57] B. Friedlander, L. Ljung, and M. Morf, "Lattice implementation of the recursive maximum likelihood algorithm," in *Proc. Int. Conf. Decision and Control*, San Diego, CA, pp. 1083-1084, Dec. 1981.
- [58] C. R. Johnson, Jr., "Adaptive IIR filtering: current results and open issues," *IEEE Trans. Inform. Theory*, vol. 30, no. 2, pp. 237-250, Mar. 1984.
- [59] H. Fan and W. K. Jenkins, "An investigation of an adaptive IIR echo canceller: advantages and problems," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 36, no. 12, pp. 1819-1834, Dec. 1988.
- [60] R. L. Cupo and R. D. Gitlin, "Adaptive carrier recovery systems for digital data communications receivers," *IEEE J. Sel. Areas in Commun.*, vol. 7, no. 9, pp. 1328-1339, Dec. 1989.
- [61] C. F. N. Cowan and P. M. Grant, *Adaptive Filters*. Englewood Cliffs, NJ: Prentice-Hall, 1985.
- [62] S. Chen, A. W. Palmer, K. T. V. Grattan, and B. T. Meggitt, "Digital signal-processing techniques for electronically scanned optical-fiber white-light interferometry," *Appl. Opt.*, vol. 31, no. 28, pp. 6003-6010, Oct. 1992.

- [63] G. E. P. Box and G. M. Jenkins, *Time Series Analysis: Forecasting and Control*, rev. ed. San Francisco, CA: Holden-Day, 1976.
- [64] J. R. Zeidler, "Performance analysis of LMS adaptive prediction filters," *Proc. IEEE*, vol. 78, no. 12, pp. 1781-1806, Dec. 1990.
- [65] V. U. Reddy and A. Nehorai, "Response of adaptive line enhancer to a sinusoid low pass noise," *Proc. Inst. Elect. Eng.*, London, vol. 128, pt. F, pp. 161-166, June 1981.
- [66] A. Nehorai and M. Morf, "Enhancement of sinusoids in colored noise and the whitening performance of exact least squares predictors," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 30, no. 3, pp. 353-362, June 1982.
- [67] J. Zeidler, E. Satorius, D. Chabries, and H. Wexler, "Adaptive enhancement of multiple sinusoids in uncorrelated noise," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 26, no. 3, pp. 240-254, June 1978.
- [68] B. Egardt, T. Kailath, and V. U. Reddy, "High resolution spectral analysis using multiple interval adaptive prediction," *Circ., Syst., Signal Process.*, vol. 2, no. 4, pp. 421-443, 1983.
- [69] A. K. Gupta, "On suppression of a sinusoidal signal in broad-band noise," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 33, no. 4, pp. 1024-1026, Aug. 1985.
- [70] Y. Yorganandandum, V. U. Reddy, and T. Kailath, "Performance analysis of the adaptive line enhancer for sinusoidal signals in broad-band noise," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 36, no. 11, pp. 1749-1758, Nov. 1988.

Chapter 7

Non-Gaussian and Non-Linear Modelling and Filtering

7.1 Introduction

The modelling and filtering techniques of Chapters 2-6 were based on the customary assumption in signal processing that the observed data originate from a random process which is linear in nature and has Gaussian statistics.

In this chapter the statistical properties of the WLI fringe pattern will be analysed in order to verify the validity of such assumption, and tests based on both the second and higher-order statistics (cumulants) of the data will help to decide whether there is any advantage in using linear non-Gaussian or non-linear modelling and filtering techniques for the purpose of identifying the central fringe.

7.2 Non-Gaussian Modelling

In Chapter 2 the observed fringe pattern $y[n]$ was modelled as the output of a linear dynamical system driven by the Gaussian system/measurement noise $v[n]$. This came as a result of expressing a sampled harmonic process in difference equation form. Since a linear operation on a Gaussian random variable produces another Gaussian random variable [1], if $v[n]$ is Gaussian $y[n]$ must be Gaussian; the output and error sequences of a linear filter must also be Gaussian, as the filter input vector consists of delayed samples of $y[n]$.

In the event that $y[n]$ is non-Gaussian, it has to be assumed either that the

process generating $y[n]$ from $v[n]$ is non-linear, or that $v[n]$ is not the driving force. The first possibility will be investigated in Section 7.3. The second possibility implies that $y[n]$ results from adding $v[n]$ to the output of a linear system driven by an unobservable source. If $v[n]$ is non-Gaussian, the source may or may not be Gaussian; this case will not be considered because most of the noise in a WLI system is believed to be Gaussian ¹. If $v[n]$ is Gaussian, the source must be non-Gaussian; this case will be investigated in this section.

Non-Gaussian signals abound in the real world. Seismic reflectivities, radar returns from the ocean surface, and wind velocity data may have Bernoulli-Gaussian, Laplace, Rayleigh, Lognormal, or Weibull probability distributions [3, 4]. Communication signals passing through a fading channel may be Rayleigh or Rician [5]. Sonar signals are also non-Gaussian [6]. The property of most interest here is the fact that the cumulants of the sum of two statistically independent random processes are equal to the sum of the cumulants of the two processes [7]. Since the cumulants of a Gaussian process are identically zero (skewness and kurtosis included), cumulants should do a better job at drawing Gaussian measurement noise out of a non-Gaussian signal than second-order statistics (correlations) do, especially when the SNR is low and/or the measurement noise is coloured.

7.2.1 Gaussianity Tests

Probability density and Normal quantile-quantile plots [8] of the error sequence $e[n]$ from the batch and adaptive filtering schemes of Chapter five supported the Gaussian assumption. The Jarque-Bera (JB) test for joint skewness and kurtosis [9] was consistently passed at the 99 or 95 % confidence level. Third and fourth-order cumulant estimates were of the same order of magnitude as those of computer-generated Gaussian sequences of the same length and variance.

As for the WLI interference signal itself, the distribution of both simulated and measured patterns between the $1/e$ power threshold points exhibited the characteristic convex shape of a sinusoidal density [10] at high SNRs. Hence, the skewness was zero but the kurtosis was negative, and the data failed the JB test. As more fringes were included the convex was gradually filled by low-amplitude data until

¹The Gaussian source, non-Gaussian measurement case can be handled by an algorithm based on order statistics [2].

the distribution became peaked in the centre. Accordingly, the skewness remained at zero but the kurtosis became positive, and the JB test was failed again. Between these two extremes the kurtosis approached zero and the JB test was passed with a very high level of confidence. This happened when the data covered the $1/e$ intensity points (delimiting the coherence region) but not the whole pattern. An example of the resultant good match between fringes and Gaussian densities is shown in Fig. 7.1, which refers to the measured WLI output of Fig. 4.13 within the path difference $\pm 5.5\mu m$. As the noise level was increased the Gaussian assumption was accepted for an increasing range of fringes included in the estimates.

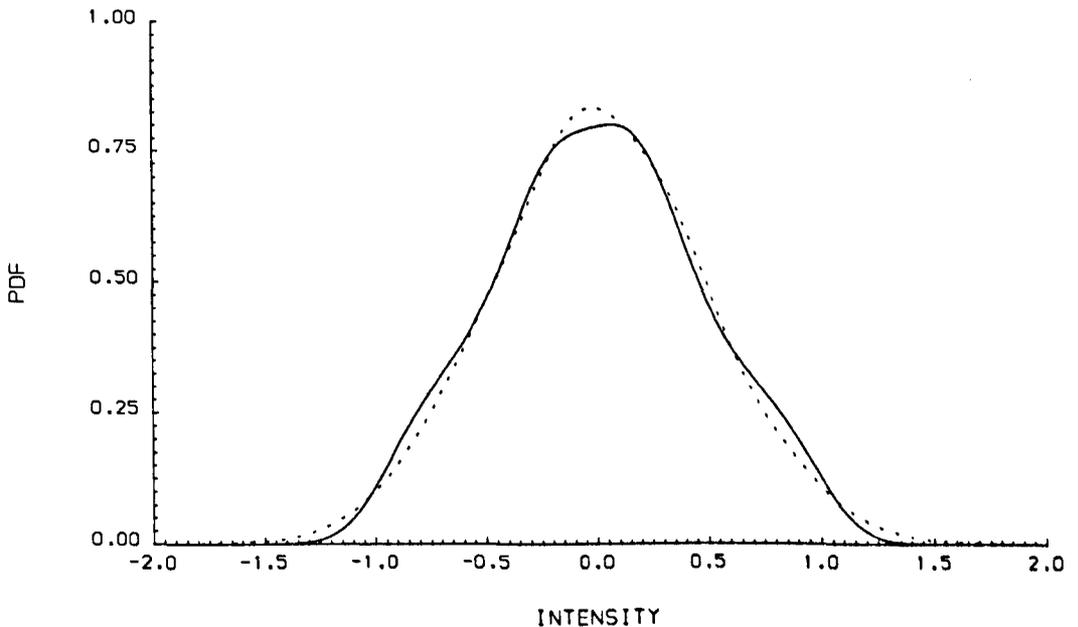


Fig. 7.1 Density of the central 26 fringes of Fig. 4.13, smoothed using a Gaussian window [11]. The dashed curve shows the density of a Gaussian sequence of the same length and variance.

Still at high SNRs, third-order sampled cumulants were several orders of magnitude smaller than those of Gaussian noise sequences of the same length, a direct consequence of the high degree of symmetry between the positive and negative portions of the fringe pattern. Fourth-order cumulants of the data between the $1/e$ power points were on average one order of magnitude larger than the corresponding cumulants of Gaussian sequences, although the latter could easily be as high as $O(10^{-1})$ for a sampled variance of $1 \pm O(10^{-2})$. Given that the kurtosis of 1000-long Gaussian sequences was also $O(10^{-1})$ but the JB test was consistently passed at the 99 % level, this should not be a concern. In fact, the sampled kurtosis of Gaussian

data fell to $O(10^{-2})$ upon increasing the sequence length ten-fold, by which time the sampled variance was within $1 \pm O(10^{-3})$.

This explains the general reluctance to the use of cumulants in signal processing; apart from being computationally more intensive to estimate than correlations, longer data sequences are needed in order to reduce the variability of the estimates, which can only be achieved by splitting the data into separate records and average over these [12]. In fact, many sampled fourth-order cumulants were smaller for the fringe pattern than for Gaussian data if averaging was not performed, and the additive property of cumulants of independent processes did not hold as well as for correlations, being totally unreliable for data lengths of $O(10^3)$ or less.

The difference between fringe pattern and Gaussian fourth-order cumulants decreased as white Gaussian sequences were replaced by narrow-band AR sequences with progressively shorter bandwidth, to end up with spectral properties similar to those of the fringe pattern. The diagonal cumulant slice [13] replicated the oscillatory behaviour of the autocorrelation function with the same period, and could be mistaken for the diagonal slice of the fringe pattern, apart from a small amplification factor (this was, e.g. less than three with the data between the $1/e$ power points and a SNR of 20 dB). Other one-dimensional slices [13] behaved similarly. It thus appears that fourth-order cumulant estimates are affected not only by the length but also by the second-order statistics of the sequence. Third-order cumulants are not affected because sampled autocorrelations are not subtracted from the sampled moments during computation, unlike in the fourth-order case.

Probability density and time-domain tests were supplemented by frequency-domain tests using higher-order spectra, as in previous work by the author [14]. The n th-order spectrum of a process is the n -dimensional Fourier transform of its n th-order cumulants [7]. As all cumulants contribute directly to its real part, whereas its imaginary part is only due to those slices not symmetric about the origin which push the process towards time-irreversibility [15], the magnitude of the higher-order spectrum of a process gives an overall measure of its departure from Gaussianity ².

Third and fourth-order spectra, known as bispectrum and trispectrum, were

²A process is time-reversible if its statistics do not depend on the direction in which it is analysed. An ARMA process is time-reversible if and only if it is Gaussian [16, 17].

estimated using both classical (non-parametric) and modern (parametric) methods of spectral analysis. The conventional direct and indirect methods [18, 7] are multi-dimensional extensions of the smoothed periodogram and correlogram [19], while the third-order recursion (TOR) method [20] calculates the bispectrum after solving for the parameters of a non-Gaussian white noise driven AR model using a set of higher-order normal equations based on the diagonal third-order slice, and can be easily adapted for the trispectrum using the fourth-order slice [21].

Two sets of WLI output data in white Gaussian noise were obtained, for a given SNR, from one hundred simulated scans between the $1/e$ power points (100-by-150 set) and from thirty simulated scans between the zero intensity points (30-by-500 set) in order to trade off variance reduction and stationarity for frequency resolution, and to test the algorithms with either negative kurtosis (100-by-150 set) or positive kurtosis (30-by-500 set). White and narrow-band AR Gaussian data sets of the same size and length were also obtained by computer simulation.

The Daniell, Parzen, and Sasaki 2-D and 3-D lag windows [22] with different widths were used in turn in the indirect method, and various sizes of Daniell (rectangular) spectral windows were employed in the direct method for smoothing in the frequency domain. Narrow windows in the time domain correspond to wide windows in the frequency domain, and give less variance at the expense of more bias and loss of resolution; the Daniell window offers the highest resolution for a fixed width, but the Parzen and Sasaki windows give lower variance and bias, respectively ³. Similarly, different model orders were considered in the TOR method. The main findings are summarised below.

With the white Gaussian data sets all three methods succeeded in producing flat and approximately zero higher-order spectra. The Daniell window was the worst, as expected, offering the highest resolution but also the highest variance, and resulting in spurious peaks along the edges of the spectral domains. Model orders between ten and twenty in the TOR method were acceptable, although it became clear that the best order for the bispectrum was not necessarily the best order for the trispectrum.

With the narrow-band AR Gaussian data sets none of the methods had difficulty

³See, e.g. [23] for 1-D lag windows design guidelines, and [24, 25] for choice of windows width in higher dimensions.

in estimating the bispectrum correctly, with the direct method being the most consistent. As for the trispectrum, the direct method was still the most consistent, giving estimates of the same magnitude as those obtained with the white Gaussian data sets. The TOR method was most successful with a model order of around twenty for the 100-by-150 set and around fifteen for the 30-by-500 set; other choices produced spurious peaks which increased the magnitude of the imaginary part, although not excessively. There was some dependency of the "best" order on the bandwidth of the simulated AR process. The indirect method was the worst of the three, giving dubious estimates of the real part for both sets. Doubling or trebling the number of records in order to decrease the variance reduced the number of peaks and their average amplitude but not dramatically.

With the WLI output data sets the bispectrum vanished, as expected. The trispectrum estimated with the indirect and TOR methods was greater than that for the AR Gaussian data sets, the magnitude difference depending on which set, window function, window width, and model order was used, but sufficient to accept non-Gaussianity comfortably down to 20 dB in all cases. However, whereas the real part of the trispectrum accounted for 90 % of its magnitude with the indirect method, the ratio between real and imaginary parts was close to one with the TOR method. If one is inclined to believe that the process generating the WLI output signal is time-reversible (this would be supported by both the general appearance of the fringe pattern and by the observed symmetry of the 1-D cumulant slices), one should also accept that the overall estimate with the TOR method was not significantly different from zero. The direct method estimate was also of the same order of magnitude or even smaller than that for the Gaussian data sets, leaving one with no clear conviction about the degree of non-Gaussianity of the interference signal of a WLI system.

Probably the best way to decide on the effectiveness of higher-order statistics in WLI is to apply them directly to the central fringe identification problem. This is carried out in the next section.

7.2.2 Linear Non-Gaussian Filtering

Tables 7.1-7.3 report the sub-fringe success rate ⁴ when batch second and higher-order algorithms were applied to simulated patterns of the kind shown in Fig. 1.5.

Method	SNR = 40 dB	30 dB	20 dB	15 dB	10 dB
Direct	98.3	64.3	24.5	16.3	7.7
2nd-order	99.7 (98.0)	85.7 (82.0)	48.4 (49.9)	35.0 (31.3)	22.7 (21.0)
3rd-order	86.7 (67.0)	48.3 (40.0)	19.5 (16.5)	13.7 (13.7)	9.7 (8.7)
4th-order	91.0 (94.0)	56.7 (50.0)	32.3 (27.6)	22.0 (22.7)	14.0 (14.3)

Table 7.1 Success rate in white Gaussian noise.

Method	SNR = 40 dB	30 dB	20 dB	15 dB	10 dB
Direct	97.3	56.7	22.0	16.7	8.7
2nd-order	99.7 (99.0)	76.7 (81.0)	27.6 (32.3)	20.0 (21.3)	13.7 (15.0)
3rd-order	83.7 (56.7)	42.7 (35.3)	19.3 (17.0)	16.3 (14.0)	13.0 (11.0)
4th-order	86.0 (87.0)	48.3 (43.7)	25.4 (22.6)	16.3 (19.3)	11.0 (12.7)

Table 7.2 Success rate in broad-band MA(5) Gaussian noise (as in [30]).

Method	SNR = 40 dB	30 dB	20 dB	15 dB	10 dB
Direct	98.0	62.3	25.4	16.7	10.3
2nd-order	99.0 (98.0)	72.0 (71.7)	29.8 (31.6)	19.3 (19.3)	13.0 (14.3)
3rd-order	82.3 (62.0)	41.0 (34.7)	20.1 (17.7)	13.7 (12.3)	9.3 (8.7)
4th-order	84.3 (86.7)	41.3 (42.3)	20.7 (20.1)	16.7 (15.3)	11.3 (12.7)

Table 7.3 Success rate in narrow-band AR(2) Gaussian noise (as in [30]).

The numbers in brackets refer to the case when the processing was restricted to the $1/e$ intensity points or coherence region (216 samples). The success rate from direct observation of the visibility profile is also shown.

⁴In percent, out of 1000 trials with the central fringe in the middle of the CCD array.

The second-order algorithm in the tables is the modified covariance method of Section 3.3. The higher-order algorithm is for a general non-Gaussian ARMA model [26], and includes singular value decomposition for AR rank-order determination, solution of an over-determined set of modified higher-order least squares Yule-Walker equations, and estimation of the MA section using the residual time series and the GM-Method in [27].

In order to avoid a possible singularity during the MA phase, the GM-Method was combined with the T-Method in [28], as previously implemented by the author in [29] following the advice in [12].

The *double C(q,k)* method in [30], which does not use a residual time series but works with the cumulants of the fringe pattern throughout, performed more poorly.

The performance of all three algorithms in the tables was affected more by coloured than by white noise, but the modified covariance remained on top in all cases, although with no noise added they all managed to identify the central fringe correctly and produced very similar sets of AR parameters.

The rank-order from the singular value decomposition was equal to two over the range 40 to 10 dB, jumping to above twenty at 0 dB; thus, the algorithms were compared for fixed model orders greater than two, as this guaranteed better performance. On the other hand, the residual time series always passed the higher-order whiteness test and the MA order was set to zero.

The tables show the identification rate for an assumed AR model order of ten only, but are quite representative of the results obtained for other model orders. In terms of speed, the third and fourth-order algorithms were slower than the second-order by factors of four and ten over 1024 data, respectively. Restricting the processing to the coherence region decreased these factors to 1.25 and 3.3.

Table 7.4 reports the success rate of second and fourth-order algorithms in the hypothetical case that the spectral distribution of the light source is rectangular instead of Gaussian; the resulting fringe visibility profile is not Gaussian but a $|\sin(x)/x|$ function [31].

Although the tests described above revealed a slightly greater departure of the fringe data within the $|\sin(x)/x|$ profile from Gaussianity, this was not enough to tip the balance in favour of the fourth-order algorithm. A bimodal Gaussian source spectrum, on the other hand, may result in a visibility profile that forces even more

Gaussianity into the fringe data, by providing a higher percentage of values close to zero. Identification rates for some typical profiles in [31] followed similar patterns to those in Tables 7.1-7.3.

Method	SNR = 40 dB	30 dB	20 dB	15 dB	10 dB
Direct	100	86.0	44.7	27.7	15.3
2nd-order	100 (100)	96.0 (92.0)	74.7 (72.3)	59.0 (54.7)	37.7 (35.0)
4th-order	100 (95.7)	79.7 (77.7)	46.7 (49.0)	35.0 (33.3)	22.7 (15.3)

Table 7.4 Success rate in white Gaussian noise with a $|\sin(x)/x|$ profile.

As for adaptive algorithms, the standard LMS and the least mean kurtosis (LMK) [32] were taken as representative of second and fourth-order methods ⁵. The LMK is a modified gradient descent method in which the negated kurtosis of the error signal is minimised, rather than the mean square. This leads to an algorithm which requires only four operations more than the LMS per iteration. The MFB-LMS of Chapter 5 was not considered because a MFB-LMK is also possible.

The LMK outperformed the LMS in Figs. 5.10-5.11 by 0.6 % with $p = 32$ and by 1.7 % with $p = 64$, in absolute terms; a slower decay of the success rate was also noticed as μ was increased from μ_0 (the latter was the same for the two algorithms) but because of this divergence occurred suddenly without the warning signs of reduced performance, a consequence of the fact that when higher-order moments of the error are considered, the degree of stability of LMS-type algorithms decreases and a stricter bound on μ must be imposed [35].

Although the performance advantage of the LMK may not be significant, the fact that it was able to achieve at least as high a success rate as the LMS may mean that long data sequences are an essential requirement when working with cumulant-based batch algorithms. Since higher-order stationarity is also implicitly assumed by these, their applicability to WLI systems is doubtful, unless a long-coherence source is used together with temporal scanning. On the contrary, cumulant-based adaptive algorithms may be beneficial both in temporally and spatially scanned

⁵The adaptive cross-correlation and parametric time-delay estimation approaches of Section 5.7 can also be implemented with fourth-order cumulants. See [33] for the first and [34] for the second approach.

systems, if fringe width and coherence length allow to restrict processing to those few central fringes where negative kurtosis can develop.

7.3 Non-Linear Modelling

The use of a linear model for the WLI system is justified only if the fringe pattern is generated by a linear process. Since a Gaussian process is linear, the Gaussianity tests of the previous section should have proven the validity of the linear assumption. In fact, in the Gaussian case the Wiener and Kalman filters provide the optimum stationary and non-stationary solutions, respectively, within the class of all linear and non-linear filters [36].

In a practical interferometric system there are many subtle sources of non-linearity, arising either in the system itself or in the associated detection and amplification circuitry. The most important are mentioned here.

Optical mixing arises from the impossibility of separating completely the optical components into the two arms of the interferometer. Angular misalignment between the source and the optics causes leakage and a non-linear, periodic relation between the phase of the interference fringes and the displacement appears, with a period of one fringe. A typical figure for the peak-to-peak phase error in a good system is 5.4° [37], corresponding to a maximum OPD error of $\pm 2.7^\circ$, which is much smaller than other errors occurring during measurements. With a fringe width of 13 pixels, for example, this error is contained to within 10 % of a pixel.

The effects of beam diffraction on linearity are difficult to identify and quantify [38], and are therefore not considered here. On the other hand, chromatic aberrations, introduced in section 5.9, result in a non-linear relation between wavelengths in air and in a refractive material like glass or an optical fibre; although the position of the central fringe remains unaffected, the visibility profile becomes skewed.

ADCs/DACs and amplifier circuitry implemented in metal-oxide semiconductor technology possess a non-linear transfer characteristic, with a systematic component due to diffusion concentration gradients and a random component due to photolithographic mismatches [39]. The random component can be translated to the output of the circuit and added to the measurement noise, but the systematic component distorts the visibility profile without skewing it.

The CCD detector could be another potential source of non-linearity, but in modern-day CCD technology this is limited to just a few parts per million.

The question to be answered in the following is thus: how strong do system and measurement non-linearities have to be in order to justify the adoption of a non-linear model for the generation of the WLI output ?

7.3.1 Linearity Tests

Two common DACs systematic non-linearities are the signed and cubic types

$$f_{NL}(x) = x + b|x| \quad \text{and} \quad ax + bx^3 \quad (7.1)$$

for a general input x , with $a=1.01333$ and $b=-0.01333$ [39]. Since one is more interested in the effects of an ADC on the analog WLI output signal prior to any digital processing operation, the non-linearities considered here will be of the quadratic or cubic type, as shown in Fig. 7.2 and expressed by

$$d[n] \equiv f_{NL}(y[n]) = ay[n] + by^2[n] \quad \text{or} \quad d[n] = ay[n] + by^3[n] \quad (7.2)$$

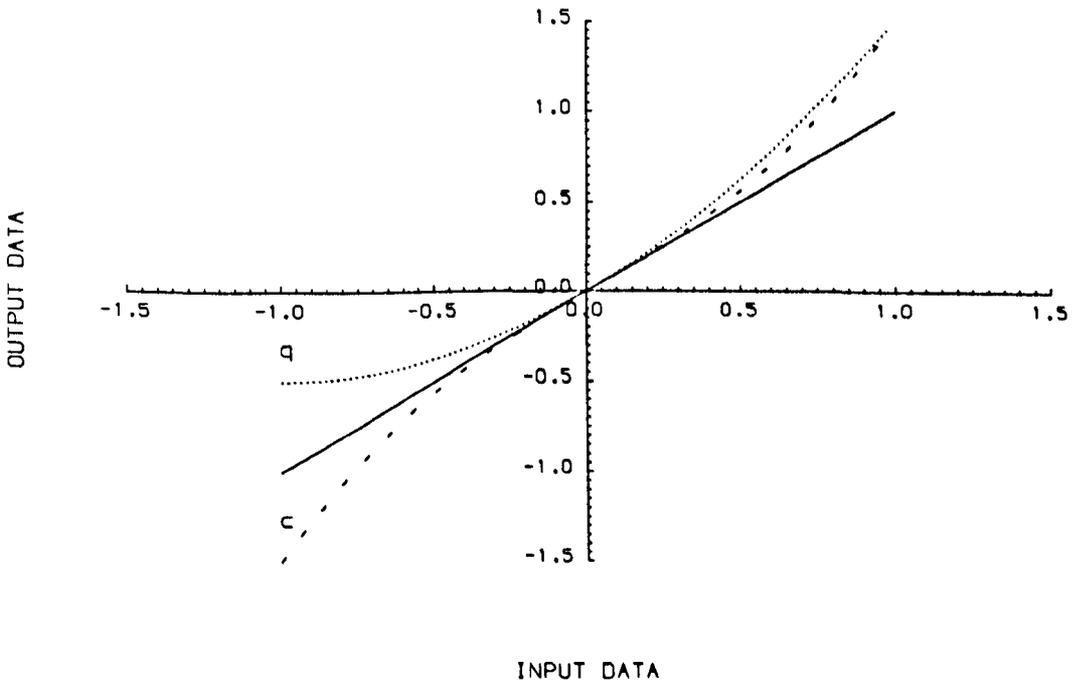


Fig. 7.2 Two common ADC non-linearities: q) quadratic; c) cubic.

Eq. 7.2 represents low-order polynomial approximations to ADCs systematic non-linearities [39], with $y[n]$ and $d[n]$ denoting the fringe pattern before and after distortion, respectively. Quadratic distortion introduces second-order harmonics in

addition to a dc term, whereas cubic distortion introduces third-order harmonics only. Their effect on a fringe pattern corrupted by Gaussian white noise can be seen in Figs. 7.3 and 7.4. The parameters a and b in Eq. 7.2 were set to 1 and 0.5.

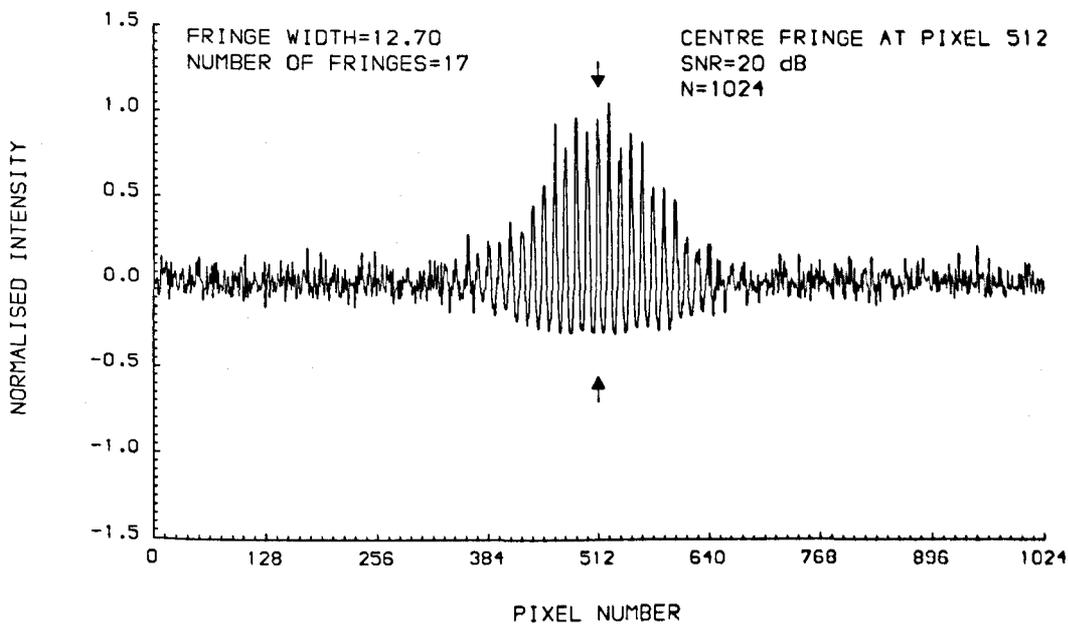


Fig. 7.3 A noise-corrupted fringe pattern distorted by a quadratic non-linearity.

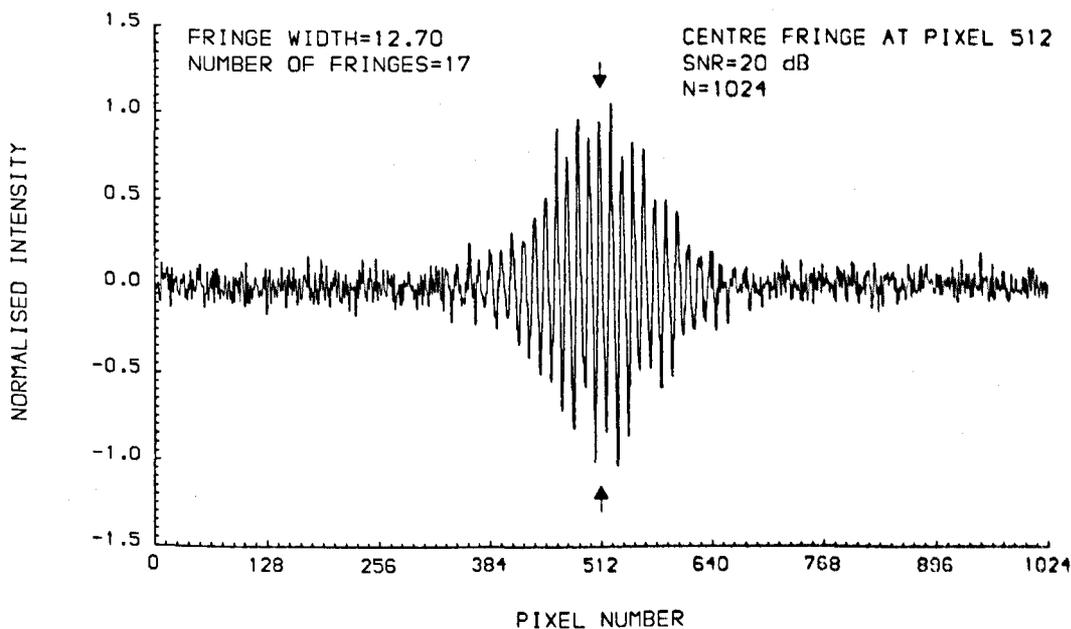


Fig. 7.4 A noise-corrupted fringe pattern distorted by a cubic non-linearity.

If $\sin(2\pi f_0 n)$ is the fundamental component of the linear fringe pattern, purely quadratic and cubic non-linearities produce the following outputs

$$\begin{aligned} d[n] &= \sin^2(2\pi f_0 n) = \frac{1}{2} - \frac{1}{2} \cos(4\pi f_0 n) \\ d[n] &= \sin^3(2\pi f_0 n) = \frac{3}{4} \sin(2\pi f_0 n) - \frac{1}{4} \sin(6\pi f_0 n) \end{aligned} \quad (7.3)$$

With typical values for a and b of 0.987 and 0.013, respectively, the transfer characteristics still looked very much linear, and neither classical nor AR power spectrum estimators were able to detect the secondary harmonics, which remained buried underneath the noise level. If A is the amplitude of the central fringe before distortion, it can easily be shown that the output power contributed by the linear term is $(aA)^2/2$ and that from the quadratic term is $(3/8)b^2A^4$, of which $(1/4)b^2A^4$ goes to the dc component. The power ratio between fundamental and second-order harmonic thus becomes $4(a/bA)^2$ in correspondence of the central fringe, and e^2 times this value at the $1/e$ intensity points. With a and b as above and A equal to one, the ratio is 23×10^3 , or 43.6 dB. Clearly, in a practical environment where the SNR is below this limit the effects of signal distortion are masked by the noise and become undetectable. The same conclusion can be reached for a cubic non-linearity.

With A and a at one and noise at 20 dB added after the distortion process, the third-order harmonic became visible only as b approached one, although the second-order harmonic was revealed sooner. Hence, the linearity tests described next were carried out with this set of values.

Third-order cumulants and bispectrum of both 100-by-150 and 30-by-500 sets of quadratically distorted patterns were non-zero, as would be expected following the loss of symmetry between the positive and negative halves of the data; however, non-Gaussianity does not necessarily imply non-linearity. The squared magnitude of the normalised bispectrum and trispectrum, also known as bicoherence and tricoherence, can be used to determine the degree of quadratic and cubic phase coupling between harmonically related components⁶. The n th-order coherence is zero for a Gaussian process and non-zero but constant over all frequencies for a linear non-Gaussian process [7]; otherwise peaks will appear where joint statistical dependencies between spectral components are introduced by non-linearities [40].

⁶Three frequencies are harmonically related if one of them is the sum of the other two.

Test statistics for zero or constant bicoherence using the sampled bispectrum alone [41] or exploiting some asymptotic properties of the bispectrum estimator [42] are available as computer programs [43, 44], but were not considered because they are based on distributional results for stationary data and do not carry over easily to the trispectrum domain ⁷. Alternatives based on suitable ratios of higher-order spectra or directly on cumulants are also possible [47, 48], as are time-domain tests for specific types of non-linearity (see, e.g. [49, 50, 51]).

Perspective views and contour maps of the bicoherence estimated with the TOR method showed a visible peak at the frequency pair (f_0, f_0) , indicating the presence of quadratic coupling as a direct result of the non-linearity, as opposed to a spontaneously excited independent mode at frequency $2f_0$. The conventional methods were also successful, although the degree of coupling was not as strong as with the TOR method, in accord with the view that conventional estimators are better as quantifiers of phase coupling, whereas AR methods are better as detectors [20].

With cubic distortion the fringe pattern remains symmetrically distributed; therefore, third-order cumulants and bispectrum should vanish like in the linear case, and no quadratic phase coupling should appear. This was indeed confirmed by the tests. Fourth-order cumulants, on the other hand, were slightly larger than in the linear case for the 30-by-500 set but slightly smaller for the 100-by-150 set; with the latter set and a SNR of 20 dB, e.g., diagonal and other 1-D slices resembled the slices of AR Gaussian processes of similar bandwidth, amplified by a factor of two and slightly distorted at the peaks and troughs. A non-zero estimate of the real part of the trispectrum with the indirect method, and of real and imaginary parts with the TOR method, were sufficient to reject Gaussianity, although the overall magnitude was slightly smaller than in the linear case with the indirect method, and either slightly smaller or larger with the TOR method, depending on the model order; the estimate was also slightly larger than in the linear case with the direct method, but still not larger than with AR Gaussian data sets.

Both the indirect and the TOR methods managed to find the third-order harmonic, producing a peak at the frequency triplet (f_0, f_0, f_0) , although the degree of coupling, as measured by the indirect method, was quite weak at 20 dB. The direct method was not successful, even when resolution and detection capabilities

⁷Confidence intervals for the sampled tricoherence have been derived only very recently [45, 46].

were pushed to the limit by progressively reducing the amount of frequency smoothing down to zero or by setting the threshold for peak detection to just above the sampled tricoherence mean.

In conclusion, quadratic and cubic distortions have to be quite strong in order to influence the statistics of the fringe pattern and make it reasonably non-Gaussian and non-linear. Cubic distortion, in particular, makes the data density appear more Gaussian by raising the kurtosis towards zero when negative, although inflating it when positive. As kurtosis and, in general, all fourth-order moments influence the real part of the trispectrum, results of any Gaussianity or linearity tests based on fourth-order statistics will depend on the proportion of the interference region used in the test. Although this may not be critical at low SNRs, where Gaussianity prevails anyway, it emphasises the importance of the stationarity assumption when implementing the tests.

Chromatic aberration effects were simulated by delaying the wavelengths by different amounts in one arm of the interferometer, in order to obtain skewed visibility profiles as in [52], an example of which is shown in Fig. 7.5.

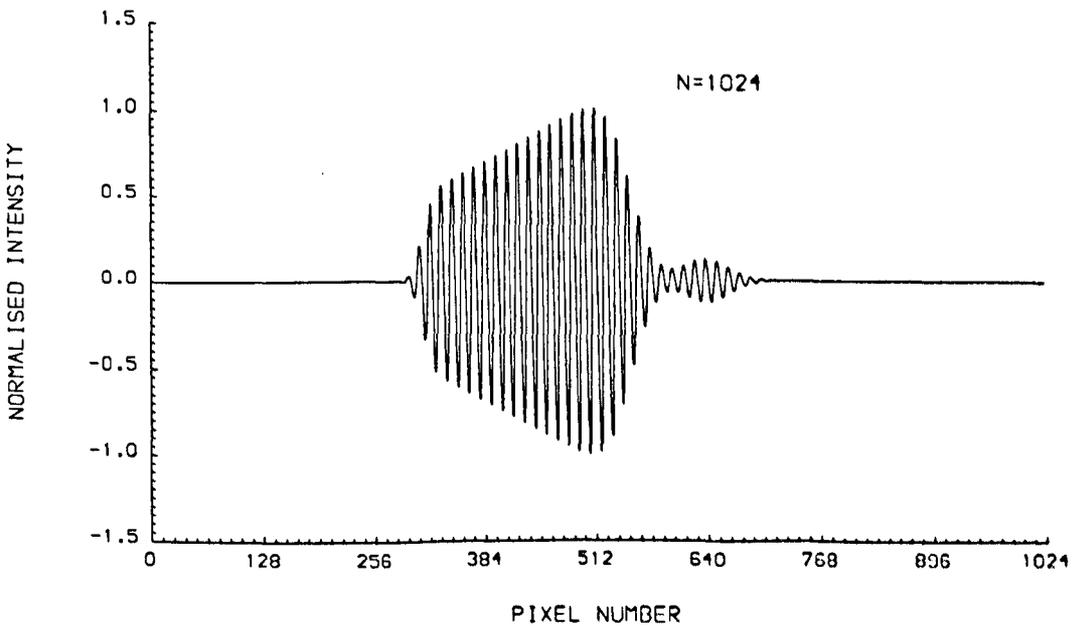


Fig. 7.5 A noise-free fringe pattern modified by chromatic aberrations.

Although the distribution of the data around their mean value remains symmetric, the process tends to become less Gaussian as the time-reversibility property is lost. Nonetheless, the spectrum and linearity tests failed to reveal the presence

of higher-order harmonics, stressing the fact that a dispersive medium does not introduce non-linear relations between the different wavelengths in the medium, although it causes a non-linear relation between wavelengths in air and the same wavelengths in the medium.

The next section presents the main conclusions drawn upon a comparison of the standard LMS predictor with non-linear versions based on a truncated discrete Volterra series expansion [53, 54].

7.3.2 Non-Linear Volterra Filtering

If $d[n]$ is the output of a linear dynamical system with rational transfer function in cascade with a zero-memory non-linearity f_{NL} as in Eq. 7.2, the overall process conforms to a particular kind of finite-memory non-linear system known as the Wiener model [55]. If f_{NL} is onto and one-to-one, the inverse function f_{NL}^{-1} can be expanded as a power series in $d[n]$,

$$y[n] = \frac{1}{a}d[n] - \frac{b}{a^3}d^2[n] + \frac{2b^2}{a^5}d^3[n] - \frac{5b^3}{a^7}d^4[n] + \dots \quad (7.4)$$

for a non-linearity of the quadratic type, or

$$y[n] = \frac{1}{a}d[n] - \frac{b}{a^4}d^3[n] + \frac{3b^2}{a^7}d^5[n] - \dots \quad (7.5)$$

for a non-linearity of the cubic type.

On the top left of Fig. 7.6 is an AR model for the generation of $y[n]$, the fringe pattern in noise. This signal goes through the non-linear function f_{NL} to produce the distorted pattern $d[n]$ shown in Figs. 7.3 and 7.4. The polynomial, or Volterra, processor on the bottom right is a non-linear adaptive FIR filter which can take not only delayed samples of $d[n]$, as a linear filter, but also their powers and all possible combinations of higher-order products.

If $s[n]$ represents the noise-free pattern and the AR approximation is valid, then

$$s[n] = -a_1y[n-1] - \dots - a_p y[n-p] \quad (7.6)$$

i.e., $s[n] \equiv f(y[n-1], \dots, y[n-p])$. Hence, from Eqs. 7.4 and 7.5 it follows that $s[n]$ can be expressed as an infinite expansion in odd and even powers of delayed $d[n]$ samples for a quadratic non-linearity, and of odd powers for a cubic non-linearity (see work by the author in [56]). Since, in practice, a is close to one and b is small,

the infinite expansion can be replaced by a third-order approximation with terms up to the cubic

$$s[n] \equiv f(d[n-1], \dots, d[n-p], d^2[n-1], \dots, d^2[n-p], d^3[n-1], \dots, d^3[n-p]) \quad (7.7)$$

for a quadratic non-linearity, or

$$s[n] \equiv f(d[n-1], \dots, d[n-p], d^3[n-1], \dots, d^3[n-p]) \quad (7.8)$$

for a cubic non-linearity. Therefore, including square and/or cubic terms at the filter input, in addition to the linear terms used by a linear filter, may help to give a better estimate of $s[n]$.

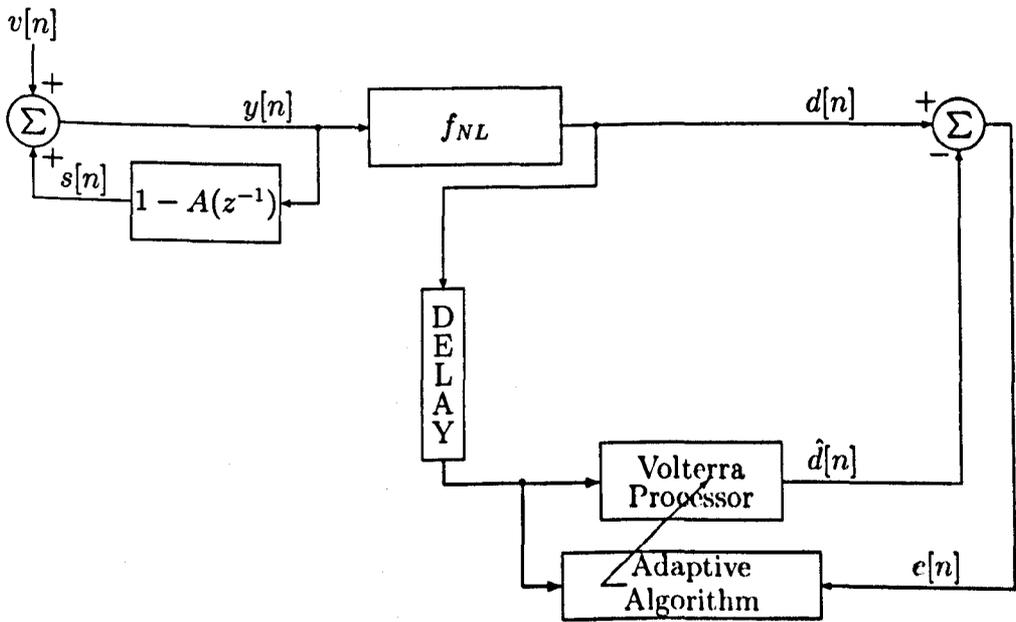


Fig. 7.6 AR model signal generator and FIR adaptive non-linear filtering.

As the filter weights are updated by the adaptive algorithm in order to minimise the MSE between the filter output and either the quadratic signal

$$d[n] = as[n] + bs^2[n] + av[n] + bv^2[n] + 2bs[n]v[n] \quad (7.9)$$

or the cubic signal

$$d[n] = as[n] + bs^3[n] + av[n] + bv^3[n] + 3bs^2[n]v[n] + 3bs[n]v^2[n] \quad (7.10)$$

it may also be useful to include terms for the direct estimation of $s^2[n]$ or $s^3[n]$.
From Eq. 7.6

$$s^2[n] \equiv f(y^2[n-1], \dots, y^2[n-p], y[n-1]y[n-2], \dots, y[n-p+1]y[n-p]) \quad (7.11)$$

$$s^3[n] \equiv f(y^3[n-1], \dots, y^3[n-p], y^2[n-1]y[n-2], y[n-1]y^2[n-2], \dots, y^2[n-p+1]y[n-p], y[n-p+1]y^2[n-p]) \quad (7.12)$$

and from Eqs. 7.4 and 7.5

$$s^2[n] \equiv f(d^2[n-1], d^3[n-1], \dots, d^2[n-p], d^3[n-p], d[n-1]d[n-2], d[n-1]d^2[n-2], d^2[n-1]d[n-2], \dots, d[n-p+1]d[n-p], d[n-p+1]d^2[n-p], d^2[n-p+1]d[n-p]) \quad (7.13)$$

$$s^3[n] \equiv f(d^3[n-1], \dots, d^3[n-p], d^2[n-1]d[n-2], d[n-1]d^2[n-2], \dots, d^2[n-p+1]d[n-p], d[n-p+1]d^2[n-p]) \quad (7.14)$$

neglecting terms higher than the third.

Hence, a Volterra filter with a linear, quadratic, and possibly a cubic section may be able to provide all the necessary flexibility when dealing with a non-linearity of the quadratic type, while a linear plus a cubic section should suffice for a non-linearity of the cubic type.

As the squares and cubes of the delayed $d[n]$ samples are used for the simultaneous estimation of $s[n]$ and $s^2[n]$, and the cubes for the simultaneous estimation of $s[n]$ and $s^3[n]$, the linear and non-linear estimates are not separable and will remain coupled at the filter output. This was verified by computer simulations. The main objective being the detection of the central fringe and not the separation of the signal into its linear and non-linear constituents, however, such coupling may not have any negative consequences on the identification process.

Although the computational complexity becomes $O(p^2)$ for a second-order and $O(p^3)$ for a third-order LMS Volterra filter implemented as in [57], it is possible to reduce the total operation count and memory storage considerably by exploiting some symmetry conditions on the coefficients.

Each double product of the form $d[n-i]d[n-j]$ in the quadratic section, with $i \neq j$, is repeated twice; hence, only the updating of the upper triangular part of

the $p \times p$ weight matrix is needed, reducing the number of quadratic coefficients from p^2 to $p(p + 1)/2$.

Similarly, the number of triple products of the form $d[n - i]d^2[n - j]$ in the cubic section, with $i \neq j$, is $3p(p - 1)$, but each arrangement is repeated three times. The number of triple products of the form $d[n - i]d[n - j]d[n - k]$ is $p(p - 1)(p - 2)$ and could also be reduced sixfold, but are left out altogether as they are not contained in the series expansion of either $s^2[n]$ or $s^3[n]$. The overall number of cubic coefficients can be reduced from p^3 to $p + p(p - 1) = p^2$.

Extensive computer simulations on noisy patterns, with the central fringe amplitude normalised to one before the distortion process, indicated that a LMS-based Volterra filter may not be able to add any value in terms of identification rate, even when the non-linearity parameter b is as high as one and different subsets of the quadratic and cubic weight sections are considered.

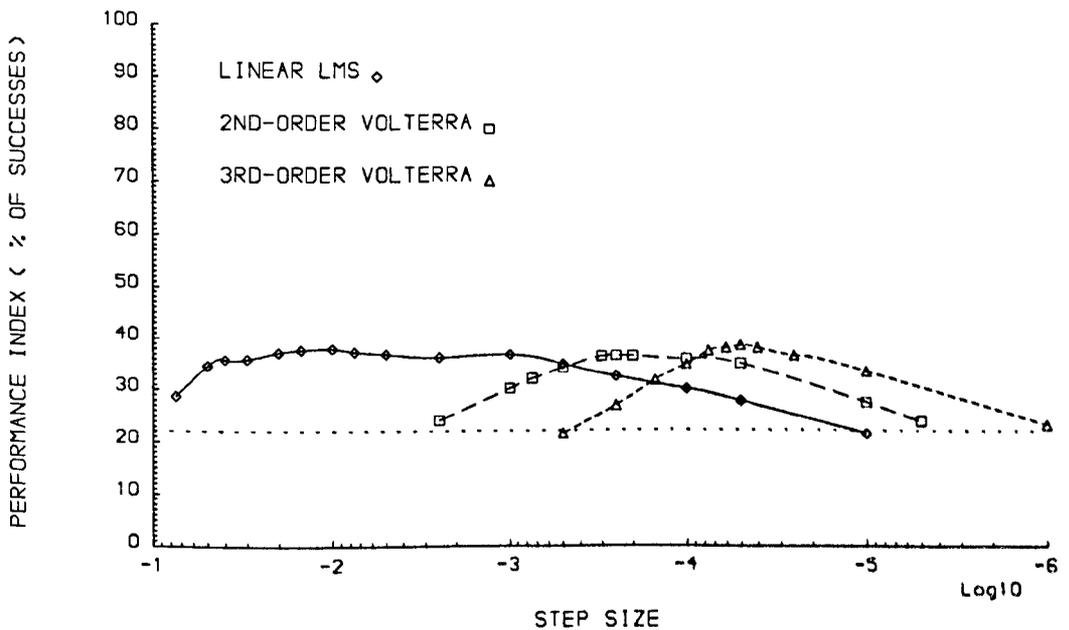


Fig. 7.7 Performance of linear and Volterra filters for the non-linear signal in Fig. 7.3.

Fig. 7.7 shows the sub-fringe identification rate of the linear LMS filter (with only the linear terms in Eq. 7.7), that of the quadratic LMS (with the additional second-order terms in Eqs. 7.7 and 7.13), and that of the quadratic-cubic LMS (with all the terms in Eqs. 7.7 and 7.13) for the non-linear signal in Fig. 7.3. The horizontal dotted line refers to the identification rate from direct fringe visibility.

The filter order p was set at 32 and, for each value of the step-size μ , the

central fringe was allowed to move at random between pixels 33 and 1024 during 1000 consecutive scans. The thresholding method of Chapter five was used to restrict the updating of the filter weights to the region within the $1/e$ power points. Although some discrepancy from the linear case would be expected, because of the non-Gaussian shape of the fringe visibility profile, the points thus calculated still delimited the central twelve fringes.

The linear filter outperformed the quadratic filter by 1.5 % in absolute terms (with $\mu = \mu_0$ in both filters) while the quadratic-cubic filter outperformed the linear filter by just 0.4 %.

Fig. 7.8 shows the identification rate of the linear LMS filter (with only the linear terms in Eq. 7.8) and that of the cubic LMS (with all the terms in Eqs. 7.8 and 7.14) for the non-linear signal in Fig. 7.4. Here, the difference between the two filters was contained to within 0.1 % (with $\mu = \mu_0$ in both filters). The inclusion of a quadratic section in the cubic filter caused a slight deterioration in performance.

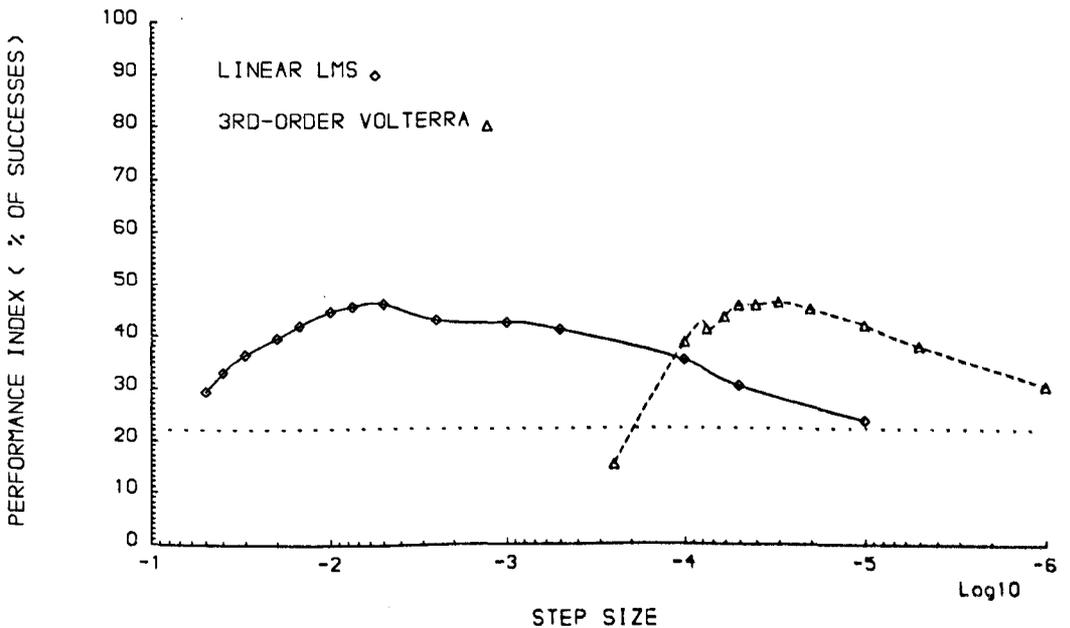


Fig. 7.8 Performance of linear and Volterra filters for the non-linear signal in Fig. 7.4.

Possible reasons for the relatively poor success of the Volterra approach are outlined below.

When both signal and additive noise $v[n]$ undergo distortion, not only $v[n]$ and its square (cube), but also multiplicative terms of the form $d[n-i]v[n]$ ($d[n-i]v^2[n]$) are propagated through a linear filter and the linear section of a Volterra filter for

a quadratic (cubic) non-linearity; multiplicative terms of the form $d^2[n - i]v[n]$ ($d^3[n - i]v[n]$) also pass through the quadratic (cubic) sections for a quadratic non-linearity, and $d^3[n - i]v^2[n]$ terms pass through the cubic section for a cubic non-linearity. Using powers of delayed $d[n]$ samples thus increases the number of noise terms interfering with the estimation process.

This possibility was investigated by adding the noise after distortion of the clean signal had taken place, as was done with the linearity tests in the previous section. Although unrealistic in practice if ADC and amplifier circuitry are the main sources of systematic non-linearities, such assumption allows to drop all the multiplicative noise terms.

The identification rate of both linear and Volterra LMS filters improved by the same amount. With a SNR of 20 dB and a cubic non-linearity parameter of one, the rate went up from 46.1 % to 60.4 % with the linear LMS and from 45.8 % to 58.2 % with the third-order Volterra without a quadratic section. The improvement may have been partly due to the absence of the multiplicative noise terms, but it cannot be ignored that the non-linear additive noise term is now also not propagated through the filters.

A second reason can be traced back to the various sources of coupling in and between linear, quadratic, and cubic sections. Non-linear operations increase the eigenvalue spread of the input autocorrelation matrix, as this becomes augmented with third and fourth (fourth and sixth) order moments sub-matrices introduced by the quadratic (cubic) section, inducing non-uniform convergence behaviour in the non-linear weights even if the input was white [58]. With correlated data the coupling between linear weights caused by statistical dependencies of the gradient estimates extends to the non-linear weights, making coupling across sections caused by dependencies between the update equations more serious. A Volterra LMS filter implemented in the frequency domain (the non-linear equivalent of the *fast* LMS filters mentioned in Section 6.2) may help decorrelate the input data and give faster convergence [59].

Coupling between linear and quadratic sections would not occur if the input data were Gaussian or, more generally, had zero third-order moments [60], and $s[n]$ alone was estimated; in this case it would be possible to optimise the two sections separately using two individual step-sizes, with the linear part of the optimum

second-order filter being equal to the optimum linear filter [61]. A second-order fast RLS Volterra filter, which would normally require $O(p^3)$ operations [62], could then be implemented in $O(p^2)$ operations [63]. Coupling between linear and cubic sections, on the other hand, would still remain [64], as fourth-order Gaussian moments are non-zero.

Another explanation for the poor behaviour of the Volterra filter, which would apply to a quadratic but not to a cubic non-linearity, is that the power series expansion for f_{NL}^{-1} is only valid as long as the mapping between $y[n]$ and $d[n]$ is one-to-one, i.e. as long as $y[n] > \frac{-a}{2b}$. Violating this condition did, in fact, affect the performance of the Volterra LMS in greater degree than the linear LMS.

Finally, a non-linear filter may be better at following the underlying distortions of the fringe pattern, but this is not essential for fringe order identification, as the ultimate aim is to recover the intensity peaks, and a linear filter may be able to model the linear signal just as well. This was tested by comparing the excess MSE of linear and Volterra LMS thresholded schemes over the second half of 1000-long sensing phases with the noise added after the distortion process and the step-size optimised for maximum identification rate. The error was defined as the difference between the predicted output and the noise-free distorted pattern, in order to consider only the MSE in excess of the noise variance.

It became apparent that a Volterra filter can indeed reproduce the non-linear pattern better than a linear filter. For example, with unit cubic distortion and additive noise at 20 dB (noise variance equal to 0.01) the excess MSE within the $1/e$ power points (innermost twelve fringes) averaged over the 500 scans was 2.73×10^{-3} (27.3 % misadjustment) for the Volterra LMS with linear and cubic sections, and 60 % larger (43.7 % misadjustment) for the linear LMS, although the identification rate over the same scans was 53.2 and 56.6 %, respectively. Over the four innermost fringes the misadjustment was 27.5 % for the Volterra LMS and 39.5 % for the linear LMS. Superimposing predicted traces to the noise-free pattern revealed that a Volterra filter is able to follow the low-amplitude distortions of the amplitude-modulated sinusoid more closely, although it may approximate the peaks less well.

The Volterra LMS was also capable of modelling the linear pattern as well as the linear LMS. With zero distortion parameter and additive noise at 20 dB the misadjustment between the $1/e$ power points was 25.7 % for the Volterra LMS with

linear and cubic sections, and 35.0 % for the linear LMS, although the identification rate was much higher for the latter (46.6 % against 35.2 %). Over the four innermost fringes the two filters had comparable misadjustment, this being 26.2 % for the linear LMS and 29.7 % for the Volterra LMS.

7.4 Discussion

This chapter has tried to answer the following question: is the fringe pattern of a WLI system evolving from a Gaussian process and, if not, can more appropriate modelling assumptions lead to higher identification rates ?

If a linear dynamical system is driven by a Gaussian process, the system output is also a Gaussian process and its higher-order statistics (cumulants) vanish. Cumulant-based algorithms thus only work for non-Gaussian processes. Even then, third-order cumulants are only useful for non-symmetric processes, otherwise fourth-order cumulants become necessary.

The symmetry and time-reversibility properties characteristic of Gaussian processes were easy to establish for the fringe pattern. Fourth-order cumulants and trispectral estimates, on the other hand, were seen to depend not only on the number and length of the data sets used but, perhaps more importantly, on the bandwidth of the process under examination, even when the process itself was Gaussian.

Moreover, the non-stationarity induced by the visibility profile was responsible for shaping the probability density of a sine-wave into a Gaussian function, even when irregularities in the emission spectrum of the light source combined to produce a non-Gaussian profile.

Similarly, the violation of the higher-order stationarity condition assumed by higher-order estimators played its part in classifying a non-stationary non-Gaussian signal as a stationary Gaussian one.

The linearity tests described in this chapter assumed that the Gaussian noise was additive at the system output. If the non-linearity is caused by the ADC or amplifier circuitry, measurement noise goes through the non-linear process together with other system noise. As the principle of superposition is not applicable, this internal noise cannot be translated to an additive, possibly coloured, output noise source as is normally assumed in linear identification [65]. Hence, it becomes much

more difficult to separate the signal from the noise contribution, as any non-linear operation on a Gaussian process produces a non-Gaussian process; narrow-band noise will also be subjected to phase coupling. These tests tend therefore to be on the optimistic side when applied to real measurements.

In addition, recent results on higher-order spectra suggest that extreme care must be exercised before drawing final conclusions from both Gaussianity and linearity tests, even in the tacitly assumed additive noise case. It has been shown that the bispectrum of a sampled signal band-limited to $[f_1, f_2]$ is zero if $2f_1 \geq f_2$ and either $f_s > 3f_2$ or $f_s < 3f_1$ are satisfied, where f_s is the sampling frequency [47]. Both conditions can easily be met in WLI systems, as Fig. 2.2 showed. The bicoherence of a linear non-Gaussian process has also been shown not to be a constant when coloured Gaussian noise is present [66]. These results will obviously carry over to the fourth-order spectrum.

The next and final chapter presents a summary of the main results obtained in this thesis, and gives the interested reader some directions for further work.

Bibliography

- [1] A. Papoulis, *Probability, Random Variables, and Stochastic Processes*. New York: McGraw-Hill, 1965.
- [2] T. I. Haweel and P. M. Clarkson, "A class of order statistic LMS algorithms," *IEEE Trans. Signal Process.*, vol. 40, no. 1, pp. 44-53, Jan. 1992.
- [3] F. W. Machell and C. S. Penrod, "Probability density functions of ocean acoustic noise processes," in *Statistical Signal Processing*, E. J. Wegman and J. G. Smith, eds. New York: Dekker, pp. 211-221, 1984.
- [4] P. A. W. Lewis and D. K. Hugus, "An analysis of 15 years of wind velocity data from ship PAPA," in *Extension of Some Models for Positive Valued Time Series*. D. K. Hugus PhD thesis, Naval Postgraduate School, Monterey, CA, 1982.
- [5] W. C. Jakes, *Microwave Mobile Communications*. New York: Wiley, 1974.
- [6] K. Sasaki, T. Sato, and Y. Nakamura, "Holographic passive sonar," *IEEE Trans. Sonics Ultrason.*, vol. 24, no. 3, pp. 193-200, May 1977.
- [7] D. R. Brillinger, "An introduction to polyspectra," *Ann. Math. Statist.*, vol. 36, pp. 1351-1374, Oct. 1965.
- [8] J. M. Chambers, W. S. Cleveland, B. Kleiner, and P. A. Tukey, *Graphical Methods for Data Analysis*. Monterey, CA: Wadsworth, 1983.
- [9] C. M. Jarque and A. K. Bera, "Efficient tests for normality, homoscedasticity and serial independence of regression residuals," *Economics Lett.*, vol. 6, pp. 255-259, 1980.
- [10] J. S. Bendat and A. G. Piersol, *Random Data: Analysis and Measurement Procedures*. New York: Wiley, 1971.

- [11] B. W. Silverman, "Density estimation for univariate and bivariate data," in *Interpreting Multivariate Data*, V. Barnett, ed. Chichester: Wiley, 1981.
- [12] J. M. Mendel, "Tutorial on higher-order statistics (spectra) in signal processing and system theory: theoretical results and some applications," *Proc. IEEE*, vol. 79, no. 3, pp. 278-305, Mar. 1991.
- [13] Y. Nagata, "Lag joint probability, higher-order covariance function and higher-order spectrum," *Bulletin de la Société franco-japonaise d'oceanographie*, vol. 8, no. 2, pp. 78-94, May 1970.
- [14] D. Romare, "Higher order statistical tests applied to ECG data," *Internal Report*, Dept. Electrical, Electronic and Information Engineering, City Univ., London, Mar. 1995.
- [15] D. R. Brillinger and M. Rosenblatt, "Computation and interpretation of kth-order spectra," in *Advanced Seminar on Spectral Analysis of Time Series*, B. Harris, ed. New York: Wiley, pp. 189-232, 1967.
- [16] G. Weiss, "Time-reversibility of linear stochastic processes," *J. Appl. Prob.*, vol. 12, pp. 831-836, 1975.
- [17] A. J. Lawrance, "Directionality and reversibility in time series," *Int. Statist. Rev.*, vol. 59, no. 1, pp. 67-79, 1991.
- [18] P. J. Huber, B. Kleiner, T. Gasser, and G. Dumermuth, "Statistical methods for investigating phase relations in stationary stochastic processes," *IEEE Trans. Audio Electroacoust.*, vol. 19, no. 1, pp. 78-86, Mar. 1971.
- [19] G. M. Jenkins and D. G. Watts, *Spectral Analysis and its Applications*. San Francisco, CA: Holden-Day, 1968.
- [20] M. R. Raghuveer and C. L. Nikias, "Bispectrum estimation: a parametric approach", *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 33, no. 4, pp. 1213-1230, Oct. 1985.
- [21] A. Swami and J. M. Mendel, "Cumulant-based approach to the harmonic retrieval and related problems," *IEEE Trans. Signal Process*, vol. 39, no. 5, pp. 1099-1109, May 1991.

- [22] K. Sasaki, T. Sato, and Y. Yamashita, "Minimum bias windows for bispectral estimation," *J. Sound Vibr.*, vol. 40, no. 1, pp. 139-148, 1975.
- [23] M. B. Priestley, *Spectral Analysis and Time Series*, vol. 1. London: Academic, 1981.
- [24] D. R. Brillinger and M. Rosenblatt, "Asymptotic theory of estimates of k-th order spectra," in *Advanced Seminar on Spectral Analysis of Time Series*, B. Harris, ed. New York: Wiley, pp. 153-188, 1967.
- [25] C. L. Nikias and A. P. Petropulu, *Higher-Order Spectral Analysis: A Non-linear Signal Processing Framework*. Englewood Cliffs, NJ: Prentice-Hall, 1993.
- [26] G. B. Giannakis and J. M. Mendel, "Cumulant-based order determination of non-Gaussian ARMA models," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 38, no. 8, pp. 1411-1423, Aug. 1990.
- [27] G. B. Giannakis and J. M. Mendel, "Identification of nonminimum phase systems using higher order statistics," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 37, no. 3, pp. 360-377, Mar. 1989.
- [28] J. K. Tugnait, "Approaches to FIR system identification with noisy data using higher-order statistics," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 38, no. 7, pp. 1307-1317, July 1990.
- [29] D. Romare, *Power Spectral Density and System Parameters Estimation - Second and Higher Order Statistics*. MSc dissertation, Dept. Electrical, Electronic and Information Engineering, City Univ., London, Oct. 1992.
- [30] G. B. Giannakis and A. Swami, "On estimating noncausal nonminimum phase ARMA models of non-Gaussian processes," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 38, no. 3, pp. 478-495, Mar. 1990.
- [31] M. Born and E. Wolf, *Principles of Optics: Electromagnetic Theory of Propagation, Interference and Diffraction of Light*, 4th ed. Oxford: Pergamon, 1970.
- [32] O. Tanrikulu and A. G. Constantinides, "Least-mean kurtosis: a novel higher-order statistics based adaptive filtering algorithm," *Electron. Lett.*, vol. 30, no. 3, pp. 189-190, Feb. 1990.

- [33] J. K. Tugnait, "Time delay estimation in unknown spatially correlated Gaussian noise using higher-order statistics," *Proc. 23rd Asilomar Conf. Signals, Systems, Computers*, pp. 211-215, Pacific Grove, CA, 1989.
- [34] H. H. Chiang and C. L. Nikias, "A new method for adaptive time delay estimation for non-Gaussian signals," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 38, no. 2, pp. 209-219, Feb. 1990.
- [35] E. Walach and B. Widrow, "The least mean fourth (LMF) adaptive algorithm and its family," *IEEE Trans. Inform. Theory*, vol. 30, no. 2, pp. 275-283, Mar. 1984.
- [36] R. C. Brown, *Introduction to Random Signal Analysis and Kalman Filtering*. New York: Wiley, 1983.
- [37] R. C. Quenelle, "Nonlinearity in interferometer measurements," *Hewlett-Packard J.*, vol. 34, p. 10, Apr. 1983.
- [38] N. Bobroff, "Recent advances in displacement measuring interferometry," *Meas. Sci. Technol.*, vol. 4, pp. 907-926, 1993.
- [39] O. Agazzi, D. G. Messerschmitt, and D. A. Hodges, "Nonlinear echo cancellation of data signals," *IEEE Trans. Commun.*, vol. 30, no. 11, pp. 2421-2433, Nov. 1982.
- [40] K. I. Kim and E. J. Powers, "Digital bispectral analysis and its applications to nonlinear wave interactions," *IEEE Trans. Plasma Sci.*, vol. 7, no. 2, pp. 120-131, Jun 1979.
- [41] T. Subba Rao and M. M. Gabr, "A test for linearity of stationary time series," *J. Time Series Analysis*, vol. 1, no. 1, pp. 145-158, 1980.
- [42] M. J. Hinich, "Testing for gaussianity and linearity of a stationary time series," *J. Time Series Analysis*, vol. 3, no. 3, pp. 169-176, 1982.
- [43] T. Subba Rao and M. M. Gabr, *An Introduction to Bispectral Analysis and Bilinear Time Series Models* (Lecture Notes in Statistics no. 24). New York: Springer-Verlag, 1984.

- [44] D. Patterson, "BISPEC: a program to estimate the bispectrum of a stationary time series," *The American Statistician*, vol. 37, no. 4, pp. 323-324, Nov. 1983.
- [45] V. Chandran, "On the computation and interpretation of auto- and cross-trispectra," *Proc. Int. Conf. Acoust., Speech, Signal Process.*, Adelaide, Australia, vol. 4, pp. 445-448, Apr. 1994.
- [46] V. Kravtchenko-Berejnoi, F. Lefeuvre, V. Krasnosel'skikh, and D. Lagoutte, "On the use of tricoherent analysis to detect non-linear wave-wave interactions," *Signal Process.*, vol. 42, pp. 291-309, 1995.
- [47] A. Swami, "Some new results in higher-order statistics," *Int. Signal Process. Workshop on Higher Order Statistics*, Chamrousse, France, pp. 135-138, July 1991.
- [48] G. B. Giannakis and M. K. Tsatsanis, "HOS or SOS for parametric modeling?," *Int. Conf. Acoust., Speech, Signal Process.*, Toronto, Ontario, Canada, pp. 3097-3100, May 1991.
- [49] A. I. McLeod and W. K. Li, "Diagnostic checking ARMA time series models using squared-residual autocorrelations," *J. Time Series Analysis*, vol. 4, no. 4, pp. 269-273, 1983
- [50] D. M. Keenan, "A Tukey nonadditivity-type test for time series nonlinearity," *Biometrika*, vol. 72, no. 1, pp. 39-44, 1985.
- [51] R. S. Tsay, "Nonlinearity tests for time series," *Biometrika*, vol. 73, no. 2, pp. 461-466, 1986.
- [52] M. J. Downs, "Problems associated with thin film measurement using double beam interference microscopy," *J. Phys. E: Sci. Instrum.*, vol. 14, pp. 24-26, 1981.
- [53] V. Volterra, *Theory of Functionals and of Integral and Integro-Differential Equations*. New York: Dover, 1959.
- [54] M. Schetzen, "Nonlinear system modeling based on the Wiener theory," *Proc. IEEE*, vol. 69, no. 12, pp. 1557-1573, Dec. 1981.

- [55] S. A. Billings, "Identification of nonlinear systems - a survey," *IEE Proc.*, vol. 127, pt. D, no. 6, pp. 272-285, Nov. 1980.
- [56] M. Sabry-Rizk, D. Romare, K. T. V. Grattan, and A. W. Palmer, "Adaptive Volterra processor for fringe order identification in white-light interferometric systems," *Proc. SBMO/IEEE MTT-S Int. Conf. Microwaves and Optoelectronics*, Rio de Janeiro, Brazil, July 24-27, IEEE Catalogue no. 95TH8080, Piscataway, NJ, ISBN 0-7803-2674-1, pp. 413-418, 1995.
- [57] M. J. Coker and D. N. Simkins, "A nonlinear adaptive noise canceler," *Proc. Int. Conf. Acoust., Speech, Signal Process.*, Denver, CO, vol. 1, pp. 470-473, Apr. 1980.
- [58] M. V. Dokic and P. M. Clarkson, "On the performance of a second-order adaptive Volterra filter," *IEEE Trans. Signal Process.*, vol. 41, no. 5, pp. 1944-1947, May 1993.
- [59] M. U. Khurram and H. M. Ahmed, "Transform domain LMS structure for nonlinear adaptive filters," in *Proc. Int. Signal Process. Workshop on Higher Order Statistics*, Chamrousse, France, July 1991. J. L. Lacoume, ed., pp. 301-301: Elsevier Science, 1992.
- [60] T. Koh and E. J. Powers, "Second-order Volterra filtering and its application to nonlinear system identification," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 33, no. 6, pp. 1445-1455, Dec. 1985.
- [61] P. Duvaut, "A unifying and general approach to adaptive linear-quadratic discrete time Volterra filtering," *Proc. Int. Conf. Acoust., Speech, Signal Process.*, Glasgow, Scotland, pp. 1166-1171, May 1989.
- [62] V. J. Mathews and J. Lee, "A fast recursive least-squares second order Volterra filter," *Proc. Int. Conf. Acoust., Speech, Signal Process.*, New York, pp. 1383-1386, Apr. 1988.
- [63] C. E. Davila, A. J. Welch, and H. G. Rylander, III, "A second-order adaptive Volterra filter with rapid convergence," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 35, no. 9, pp. 1259-1263, Sept. 1987.

- [64] S. W. Nam, S. B. Kim, and E. J. Powers, "Utilization of digital polyspectral analysis to estimate transfer functions of cubically nonlinear systems with non-Gaussian inputs," *Proc. Int. Conf. Acoust., Speech, Signal Process.*, Glasgow, Scotland, pp. 2306-2309, May 1989.
- [65] S. A. Billings and W. S. F. Voon, "Least squares parameter estimation algorithms for non-linear systems," *Int. J. Systems Sci.*, vol. 15, no. 6, pp. 601-615, 1984.
- [66] A. Swami, "Pitfalls in polyspectra," *Int. Conf. Acoust., Speech, Signal Process.*, Minneapolis, MN, vol. 4, pp. 97-100, Apr. 1993.

Chapter 8

Summary and Directions for Future Work

8.1 Summary

In this thesis a new approach to the central fringe identification problem in white-light interferometry (WLI) systems has been proposed, based on batch and adaptive, linear and non-linear digital filtering algorithm schemes.

The main motivation for the use of a digital filter in WLI is the fact that the interference signal at the output of the system has long memory (Fig. 2.3), whereas the various noise sources arising in the system or at the detector have either very short memory or no memory at all. A filter can exploit this information and reconstruct a version of the signal which is less noisy than the original.

The common assumption that the information is only linear and is fully contained in second-order statistics (correlation) was relaxed by analysing non-linear Volterra filters and linear filters based on higher-order statistics (cumulants). It was shown that an adaptive non-linear Volterra filter does not guarantee a higher identification rate than an adaptive linear filter, even when non-linear distortion of the signal has taken place (Figs. 7.7-7.8); similarly, a batch cumulant-based filter does not improve upon a batch correlation-based filter (Tables 7.1-7.4) even when the fringe visibility profile does not have the usual Gaussian shape. Reasons for failing were given, corroborated by linearity and Gaussianity tests on the fringe pattern.

Among the linear, correlation-based filters, it was found that batch algorithms

(covariance, Burg, etc.) and adaptive algorithms (LMS, RLS, Kalman) were able to achieve the same identification rate for a given filter order (Fig. 3.5). Because of the possibility of operating in real-time, it became desirable to concentrate on adaptive algorithms, and especially on the LMS, the simplest and least computationally demanding of them all. To guarantee a competitive identification rate with the LMS, its step-size had to be finely tuned according to the filter order, signal-to-noise ratio, and degree of non-stationarity (Figs. 4.1-4.3, 4.18). The forgetting factor in the RLS was much less influenced by changes in the filter order or in the external environment. Much effort was spent in analysing the LMS and RLS during convergence and tracking (Figs. 4.5-4.12, 4.14-4.17).

A thresholding technique was proposed to alleviate the slow convergence problem affecting the LMS, but failed in the task of increasing the identification rate and/or reducing the sensitivity to the step-size (Figs. 5.1-5.2). A novel forward-backward LMS algorithm, named MFB-LMS (Fig. 5.4) also failed. Combining the MFB-LMS with the thresholding technique did, however, succeed extremely well (Figs. 5.5, 5.10-5.11).

The new thresholded MFB-LMS scheme offered the following advantages:

1. The identification rate (to sub-fringe level) was much higher than with other LMS-based schemes (Figs. 5.10-5.11), RLS-based schemes (Fig. 5.12), or batch correlation-based schemes (Fig. 5.9).
2. The identification rate increased with the filter order p . Hence, the centroid method can be beaten with a moderate value of p (Figs. 5.8-5.9, 5.11).
3. The unusable scanning range was approximately $2p$ samples or pixels, translating into a usable operating range of more than 90 % for a 1024-pixel CCD array and a moderate value of p . The corresponding operating range with the centroid method may only be as high as 75 % (Fig. 5.8), as it is necessary to maintain a minimum of symmetry between the two halves of the fringe pattern separated by the central fringe.
4. Computational simplicity, limited memory requirements, and real-time mode of operation make it especially suitable for implementation on a single-chip DSP for fast and continuous measurements.

Possible disadvantages follow:

1. Although the thresholding technique offers some protection against impulsive noise, the MFB-LMS was designed to handle additive white Gaussian noise. Modifications in the form of generalised LMS or order statistic LMS algorithms would be necessary when optimising the scheme for the coloured Gaussian case or to minimise the effects of occasional bursts of impulsive noise. The centroid method, on the other hand, is alleged to resist a certain amount of coloured Gaussian and impulsive noise.
2. The addition of a standard electronic scheme for offset/baseline correction and power/frequency stabilisation of the light source is highly advisable; baseline and power drifts, in particular, introduce non-stationarities which may adversely affect the identification rate. The centroid method, on the other hand, is only affected when the variations are fast enough to tilt or shift the centre of symmetry of the fringe pattern.
3. The identification rate will also be affected, like in the centroid method, by asymmetries in the fringe pattern caused by asymmetries in the power spectrum of the source or by difficulties in system alignment.

8.2 Directions for Future Work

In addition to considering some of the alternative models and methods suggested in Chapter six, here are a few proposals which, in the author's opinion, could form an interesting extension to this thesis.

- All the simulations conducted here assumed that the central fringe in a practical WLI system is always sampled precisely at its peak, which corresponds to the position of zero OPD. This allowed to control the SNR, defined as the ratio of the central fringe amplitude to the root mean square or standard deviation of the noise. Moreover, the fringe width was never an exact integer number of samples, which meant that the first fringe on either side of the central fringe was never sampled at its peak. In practice, the central fringe will rarely be sampled at its peak; each pixel on a CCD array, for example, stores an amount of charge proportional to the average light intensity falling

inside its area. There is thus an artefact when performing computer simulations with the above assumption, concerning the difference between the central fringe and first side fringe amplitudes, whose significance for practical measurements could be analysed.

- Although shot and thermal noise in electronic components are thought to be white Gaussian, and measurement noise in general is broad-band Gaussian, vibrations during scanning in temporally scanned WLI systems can be described as a narrow-band, probably non-Gaussian process. Speckle and impulsive noise are also non-Gaussian. It would therefore be worth studying the characteristics of the noise in practical WLI systems, and to apply linear filters based on order statistics [1] when a certain amount of non-Gaussian noise corrupts the system output.
- Implementation of a digital filtering algorithm on a single-chip DSP will ultimately involve an analysis of finite precision effects, including quantisation and round-off errors. LMS-based schemes are, in general, less sensitive to finite wordlength and ill-conditioning than RLS-based schemes. LMS algorithms update the weights using instantaneous estimates of auto and cross-correlation, whereas RLS algorithms use implicit matrix inversion. Ill-conditioning of the data matrix is thus a more serious problem with the latter, as round-off errors accumulate and propagate from one iteration to the next. In one experiment, for example, the LMS algorithm could operate properly with as few as 7 bits, whereas fast RLS algorithms needed at least 10 bits [2]. Round-off errors in a 32-bit floating-point DSP should follow the same pattern as on a personal computer or workstation using the IEEE standard format for single precision representation (23-bit mantissa, 8-bit exponent, and one sign bit), which was the format used throughout this thesis, and can therefore be easily anticipated. Quantisation errors caused by limited precision and overflow brought about by a small dynamic range in a fixed-point implementation may be harder to predict.

Quantisation errors of the LMS filter weights are proportional to μ^{-1} both during convergence [3] and steady-state [4, 5]. Quantisation errors in the correction term of the update equation may also take the algorithm to the

edge of stability [5]. It may well be that such effects are hardly felt by a 24-bit or 16-bit fixed-point DSP, but it would be interesting to know the number of bits that can be excluded before the performance of the algorithm in a practical setting degrades.

- On the modelling side, one may investigate time-lag dependencies and the minimum embedding dimension (see, e.g. [6, 7]) in order to improve parsimony of representation.
- On the signal processing side, one may try de-noising with wavelets [8]. Whereas the discrete Fourier transform uses basis functions (sines and cosines) which are infinite in extent, and thus loses all reference of time, the discrete Wavelet transform uses a set of basis functions derived from a prototype function which is finite in time. Dilations and translations of this prototype allow a representation of the signal at different time-scales and frequency resolutions, which is particularly useful for the analysis of both smooth and abrupt non-stationarities. Of particular interest is the problem of choosing the type and amount of thresholding to be applied to the wavelet coefficients for optimum noise removal and signal reconstruction.

As the Wavelet transform can be used iteratively to represent finer and finer details in the data, it may also be interesting to use various decomposition levels and combine estimates from each inverse Wavelet transform into an overall estimate for the reconstructed signal. A lot of research is going on at the moment into trying to find optimum combinations of models outputs (see, e.g. [9, 10]).

- Time series prediction with Support Vector learning machines is another research possibility. The Support Vector (SV) method was discovered in 1965 for constructing separating hyperplanes that minimise the number of classification errors in pattern recognition problems, and has now become a general approach to function approximation, regression estimation and signal processing problems [11]. The basic idea is to map the low-dimensional data vectors of the input space non-linearly onto high-dimensional vectors of a feature space, and use linear regression or classification in this space to construct the optimum separating hyperplane. Vectors in the feature space that

lie closest to the separating hyperplane are called support vectors because if all other vectors were removed, the separating hyperplane would remain the same. Support vectors in regression can be interpreted as pulling the estimate towards the response, and good experimental results have already been obtained in time series prediction [12].

- A Volterra expansion is the most general way of representing a non-linear time series model in terms of powers of its input sequence [13]. The finite-parameter AR version used in Chapter seven was a special case, derived in order to provide a parsimonious approximation through the use of an FIR filter.

Another special case is the class of bilinear models [14] which constitute the simplest non-linear extension to linear AR and ARMA models, and can approximate to an arbitrary degree of accuracy any "reasonable" Volterra expansion over a finite time interval [15].

An alternative parsimonious representation is offered by the threshold autoregressive (TAR) models [16], which make use of a piecewise-linear approximation to functions which are not analytic at all points because of discontinuities, and can therefore overcome the convergence difficulties experienced by the Volterra series when non-linear systems are modelled that include saturating elements ¹.

A non-parsimonious representation is instead offered by real-time recurrent neural networks [18], which have the ability to perform highly non-linear time-varying input-output mappings by feeding the output of every neuron in the network back to the input, and are therefore well suited for non-linear adaptive filtering or prediction of non-stationary signals. To cut down on the very high computational cost, which is of $O(N^4)$, where N is the number of neurons in the network, a modular pipelined structure consisting of M smaller recurrent networks in cascade may be used [19, 20], which necessitates a much lower number N_M of neurons per module and $O(MN_M^4)$ operations per iteration.

Critics of neural networks argue that it may be extremely difficult to know precisely what has been extracted from the data, as the relationships that

¹See, e.g. [17] for a review of TAR models and related linearity tests.

take place inside the network are so complex that they are as yet not fully understood. There is also some fear among practitioners that they may not be sufficiently powerful to solve temporal tasks effectively, as the impact of changing a weight in the network - even if appropriate - is likely to be masked by other weights if their values are inappropriate, resulting in an error surface full of local minima [21]. On the contrary, the output of a Volterra filter is a linear function of its coefficients, and the MSE is a quadratic function with one global minimum.

Bibliography

- [1] T. I. Haweel and P. M. Clarkson, "A class of order statistic LMS algorithms," *IEEE Trans. Signal Process.*, vol. 40, no. 1, pp. 44-53, Jan. 1992.
- [2] F. Ling and J. G. Proakis, "Numerical accuracy and stability: two problems of adaptive estimation algorithms caused by round-off errors," *Proc. Int. Conf. Acoust., Speech, Signal Process.*, San Diego, CA, pp. 30.3.1-30.3.4, Apr. 1984.
- [3] S. T. Alexander, "Transient weight misadjustment properties for the finite precision LMS algorithm," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 35, no. 9, pp. 1250-1258, Sept. 1987.
- [4] C. Caraiscos and B. Liu, "A roundoff error analysis of the LMS adaptive algorithm," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 32, no. 1, pp. 34-41, Feb. 1984.
- [5] J. M. Cioffi, "Limited-precision effects in adaptive filtering," *IEEE Trans. Circ. Syst.*, vol. 34, no. 7, pp. 821-833, July 1987.
- [6] G. A. Darbellay and M. Slama, "How nonlinear is your time series ? A new method and a case study," *Neural Network World*, vol. 4, no. 5, pp. 483-493, May 1997.
- [7] H. Pi and C. Peterson, "Finding the embedding dimension and variable dependencies in time series," *Neural Computation*, vol. 6, pp. 509-520, 1994.
- [8] I. Daubechies, *Ten Lectures on Wavelets*. Philadelphia, PA: SIAM, 1992.
- [9] M. LeBlanc and R. Tibshirani, "Combining estimates in regression and classification," *Tech. Rep.*, Dept. Statistics, Toronto Univ., Ontario, Canada, pp. 1-21, Nov. 1993.

- [10] S. Hashem and B. Schmeiser, "Improving model accuracy using optimal linear combinations of trained neural networks," *IEEE Trans. Neural Networks*, vol. 6, no. 3, pp. 792-794, 1995.
- [11] V. Vapnik, S. Golowich, and A. Smola, "Support vector method for function approximation, regression estimation, and signal processing," in *Advances in Neural Information Processing Systems*, vol. 9, M. Mozer, M. Jordan, and T. Petsche, eds. Cambridge, MA: MIT Press, 1997.
- [12] K.-R. Müller, A. J. Smola, G. Rätsch, B. Schölkopf, J. Kohlmorgen, and V. Vapnik, "Predicting time series with support vector machines," *Int. Conf. Artificial Neural Networks*, Lausanne, Switzerland, Oct. 8-10, 1997.
- [13] N. Wiener, *Non-Linear Problems in Random Theory*. Boston, MA: MIT Press, 1958.
- [14] C. W. J. Granger and A. P. Andersen, *An Introduction to Bilinear Time Series Models*. Göttingen, Germany: Vandenhoeck and Ruprecht, 1978.
- [15] R. W. Brockett, "Volterra series and geometric control theory," *Automatica*, vol. 12, pp. 167-172, 1976.
- [16] H. Tong and K. S. Lim, "Threshold autoregression, limit cycles and cyclical data," *J. Royal Statist. Soc. Ser. B*, vol. 42, no. 3, pp. 245-292, 1980.
- [17] H. Tong, *Non-Linear Time Series: A Dynamical System Approach*. Oxford: Oxford Univ. Press, 1990.
- [18] R. J. Williams and D. Zipser, "A learning algorithm for continually running fully recurrent neural networks," *Neural Computation*, vol. 1, pp. 270-280, 1989.
- [19] S. Haykin and L. Li, "Real-time nonlinear adaptive prediction of nonstationary signals," *CRL Rep. no. 276*, Communications Research Lab., McMaster Univ., Hamilton, Ontario, Canada, 1993.
- [20] L. Li and S. Haykin, "A cascaded recurrent neural networks for real-time nonlinear adaptive filtering," *Int. Conf. Neural Networks*, San Francisco, CA, pp. 857-862, 1993.

- [21] M. C. Mozer, "Neural net architectures for temporal sequence processing," in *Time Series Prediction: Forecasting the Future and Understanding the Past*, A. S. Weigend and N. A. Gershenfeld, eds. (*Proc. NATO Advanced Research Workshop on Comparative Time Series Analysis*, Santa Fe, New Mexico, May 14-17, 1992, vol. XV). Reading, MA: Addison-Wesley, pp. 243-264, 1993.

Appendix A

Glossary of Terms

ADC analogue-to-digital converter

AR autoregressive time-series model. AR models correspond to IIR filters

ARMA autoregressive-moving average time-series model. ARMA models correspond to IIR filters

broad-band random process characterised by a smooth power spectral density. Also known as wide-band

CCD charge-coupled device detector which responds to the intensity of light

central fringe fringe with the highest intensity in a WLI interference pattern. Its highest point corresponds to zero OPD

coherence length maximum OPD within which interference effects can be observed. Often expressed in number of interference fringes after sampling

coloured noise random noise process consisting of a sequence of correlated random variables. Can be either broad-band or narrow-band

convergence phase during which the initialised weights of an adaptive filter are updated to minimise an error function

DAC digital-to-analogue converter

DNS degree of non-stationarity. For a WLI output signal, it is high/low if the variation of signal power or SNR along the scanning axis is fast/slow relative to the sampling rate or fringe width

driving noise the unobservable process which excites a discrete dynamical system

DSP digital signal processing or processor

dynamic range range of values that a device can process without overflow or distortion

error function any function of the error signal or filter weights. Also called error surface when more than one weight is involved

error signal difference between the observed signal and its estimate by a filter

excess MSE difference between steady-state MSE and MSE of an optimum filter

FB-LMS forward-backward LMS adaptive filtering algorithm

FIR finite impulse response filter. The current output is a combination of the present input and a finite number of past inputs

fixed-point number format in which a quantity is represented in computer hardware as an integer

floating-point number format in which a quantity is represented in computer hardware as a real number. Gives greater dynamic range than fixed-point but creates round-off errors

forgetting factor parameter controlling the speed of adaptation of the RLS filter weights during convergence and tracking

fringe identification the task of identifying the central fringe when the interference pattern is corrupted by noise

fringe visibility a function which provides amplitude-modulation of the light intensity at the output of an optical interferometer. Well approximated by a Gaussian profile in a WLI system

fringe width number of samples forming a fringe after sampling

gradient error the weight error in the LMS algorithm

IIR infinite impulse response filter. The current output is a combination of the present input and a finite number of past outputs, plus possibly a finite number of past inputs

ill-conditioning a problem is ill-conditioned if its solution is very sensitive to small changes in the data

inflexion points points at which the second derivative of the Gaussian fringe visibility profile becomes zero

lag error the steady-state contribution to the MSE which is additional to the weight error and is only due to the non-stationarity of the environment

LED light-emitting diode light source

linear process any process which results as a linear combination of its elementary constituents. In an ADC or DAC, linearity is the precision with which the digital output/input tracks the analog input/output

LMS least mean square adaptive filtering algorithm

LS least squares filter or error criterion

MA moving average time-series model. MA models correspond to FIR filters

MFB-LMS modified forward-backward LMS adaptive filtering algorithm

misadjustment ratio between excess MSE and MSE of an optimum filter

MSE mean square error. A quadratic function of the filter weights when the filter is a linear function of its weights

narrow-band random process characterised by a peaked power spectral density

OPD optical path difference between two light beams in an optical interferometer

power spectral density the distribution of power/energy as a function of frequency. Also called frequency, power or energy spectrum

PDF probability density or distribution function. The distribution of signal values generated by a deterministic or random process

process deterministic or random mechanism responsible for the generation of the observed signal or data sequence

quantisation error error introduced by the computer hardware when storing a continuous quantity with a finite wordlength

RLS recursive least squares adaptive filtering algorithm

round-off error error introduced by the computer hardware when performing operations between real numbers. Accumulates with increasing amounts of calculation, and determines the stability or instability of a numerical algorithm

SNR Signal-to-noise ratio. Usually expressed in decibels (dB)

stability a system is stable if its output remains bounded in response to a bounded input. An algorithm is stable if round-off errors introduced at one stage of the computation do not propagate through later stages with increasing magnitude

stationary random process whose statistics are time-invariant

steady-state phase following convergence in a stationary or non-stationary environment

step-size parameter controlling the speed of adaptation of the LMS filter weights during convergence and tracking

sub-fringe identification the task of identifying the sample closest to the highest point in the central fringe, when the interference pattern is corrupted by noise

tracking steady-state phase during which the weights of an adaptive filter have to be readjusted in response to non-stationarities

truncation error error introduced by an algorithm when calculating a discrete approximation to a continuous function

weight error steady-state contribution to the MSE of an adaptive algorithm in a stationary environment. Also called gradient error when referring to the LMS

white noise random process consisting of a sequence of uncorrelated random variables. Characterised by a flat power spectral density

WLI White-light interferometer. An optical interferometer which uses a broad-band light source

wordlength number of bits or bytes used to store a quantity in the computer hardware

Appendix B

List of Publications

1. D. Romare, *Power Spectral Density and System Parameters Estimation - Second and Higher Order Statistics*. MSc dissertation, Dept. Electrical, Electronic and Information Engineering, City Univ., London, Oct. 1992.
2. M. Sabry-Rizk, D. Romare, K. T. V. Grattan, and A. W. Palmer, "The application of super-resolution adaptive algorithms to fringe order estimation in all-optical-fibre interferometric sensors," *Proc. Int. Conf. Appl. Photon. Technol.*, Toronto, Ontario, Canada, June 21-23, 1994. In *Applications of Photonic Technology*, G. A. Lampropoulos, ed. New York: Plenum, 1995.
3. M. Sabry-Rizk, D. Romare, K. T. V. Grattan, and A. W. Palmer, "Efficient adaptive filtering scheme for fringe order identification in electronically scanned white-light interferometric sensors," submitted to *IEEE Trans. Instrum. and Meas.*
4. D. Romare, "Higher order statistical tests applied to ECG data," *Internal Report*, Dept. Electrical, Electronic and Information Engineering, City Univ., London, Mar. 1995.
5. M. Sabry-Rizk, D. Romare, W. Zgallai, K. T. V. Grattan, P. Hardiman, and J. Oriordan, "Higher order statistics (HOS) in signal processing: are they of any use?," *IEE Colloquium*, Savoy Place, London, May 22, Digest no. 1995/11, 1995.
6. M. Sabry-Rizk, D. Romare, K. T. V. Grattan, and A. W. Palmer, "Adaptive Volterra processor for fringe order identification in white-light interferometric

- systems," *Proc. SBMO/IEEE MTT-S Int. Conf. Microwaves and Optoelectronics*, Rio de Janeiro, Brazil, July 24-27, IEEE Catalogue no. 95TH8080, Piscataway, NJ, ISBN 0-7803-2674-1, pp. 413-418, 1995.
7. D. Romare, M. Sabry-Rizk, K. T. V. Grattan, and A. W. Palmer, "Superior LMS-based technique for white-light interferometric systems," *IEEE Photon. Technol. Lett.*, vol. 8, no. 1, pp. 104-106, Jan. 1996.
 8. D. Romare, M. Sabry-Rizk, and K. T. V. Grattan, "A comparison of different adaptive filtering techniques applied to white-light interferometric systems," *Proc. Int. Conf. Appl. Photon. Technol.*, Ottawa, Ontario, Canada, June 1996. In *Applications of Photonic Technology*, G. A. Lampropoulos, ed. New York: Plenum, 1997.
 9. D. Romare, M. Sabry-Rizk, and K. T. V. Grattan, "Evaluation of different adaptive filtering techniques applied to white-light interferometric sensor systems," *Sensors and Actuators A*, vol. 63, pp. 197-203, 1997.
 10. D. Romare, M. Sabry-Rizk, and K. T. V. Grattan, "Digital signal processing techniques for central fringe identification in white-light interferometry," *Inst. Physics Conf. on Applied Optics and Optoelectronics*, Brighton Centre, Brighton, UK, Mar. 16-19, 1998.

Appendix C

List of Software Tools

The covariance and modified covariance algorithms, Burg method, and the sub-optimal ARMA procedure of Section 3.3 were implemented in Fortran code as supplied in ¹.

The algorithms for Tables 3.1-3.2 and for Figs. 5.3 and 7.1 were written in S-Plus ² (a registered trademark of Mathsoft, Inc.).

All other algorithms were derived from the relevant literature (where applicable), written in Fortran by the author and run on a Sun-Sparc (a trademark of Sun Microsystems, Inc.) workstation under the Unix operating system (a registered trademark of UNIX Systems Laboratory, Inc.) in a time-sharing environment.

The text was typeset in LaTeX ³, a special version of TeX (a trademark of the American Mathematical Society).

All figures were produced with Gino for Fortran (a registered trademark of Bradly Associates Ltd.), with the exception of Figs. 1.4, 2.1, 3.1-3.4, 5.4 and 7.6, which were produced directly with LaTeX commands.

1. S. L. Marple, Jr., *Digital Spectral Analysis with Applications*. Englewood Cliffs, NJ: Prentice-Hall, 1987.
2. *S-PLUS Guide to Statistical and Mathematical Analysis*, Version 3.3. Seattle, WA: Statsci, a division of Mathsoft, Inc., 1995.
3. L. Lamport, *LaTeX: A Document Preparation System*. Reading, MA: Addison-Wesley, 1986.