# Chapter 4

# Data Processing

## 4.1 Introduction

This chapter describes software written to manipulate sampled data collected from the wind tunnel model. The primary tasks were to carry out dynamic calibration of force transducers and to convert timeseries recorded from transducer accelerometers and strain bridges into timeseries of acceleration, velocity and aerodynamic forces, suitably calibrated. The software was written and run on the same minicomputer used to collect the data and was coded in Pascal.

Besides the programmes to be discussed here, others of a more conventional nature were developed during the course of the research. These included programmes for computing data properties such as probability density, auto and cross-spectrum, frequency response function (FRF) and coherence, and for plotting graphs. Data collection software, produced in the Department of Mechanical Engineering, Monash University, was already available.

Techniques were developed to deal with cancellation of the inertial component of force transducer response, compensation for time lags in data collection, correction of analogue equipment FRFs, and integration of acceleration signals to give velocities. The individual techniques will be discussed first, followed by their application in dynamic calibration of force transducers and processing of data collected from the wind tunnel model.

The source code for computer programmes and shared procedures which used the techniques described here are presented in appendix C. Code is supplied for the programme `dyncal`, used to estimate the calibration constants of the force transducers, the sub-programme modules which made up the programme `adcal`, used to carry out the initial processing of data recorded from the cylinder transducer during wind tunnel runs, and for procedures `included` in the source files for `dyncal` and `adcal`.

## 4.2 Basic Techniques

### 4.2.1 Cancellation of Inertial Signal

#### Background

Different approaches to the problem of measuring lift forces on moving structures have been developed by different researchers. The problems involved if the forces acting on some length of the cylinder are directly measured are different to those if the pressure distribution is used to estimate force. Direct measurement of lift acting on sections of an oscillating cylinder entails some correction of force transducer output to account for the inertial component of the response due to the mass and acceleration of the exposed section, so that the forces due to the fluid can be studied separately. If the pressure distribution is measured, the pressures at a large number of tappings must be measured simultaneously for a good estimate of lift force. In addition, the pressure signals may also need correction for inertial effects. Correction for inertial effects is important, since in studies with oscillating cylinders it is the interaction between oscillation and fluid dynamic force which is of major interest. Poor correction can leave forces at the frequency of motion which will be incorrectly assigned an aerodynamic origin.

**Previous methods** If the lift forces acting on a section of the cylinder are obtained by estimating the weighted integral of the pressure distribution, then it is possible to avoid any correction for inertial effects. For this to be so, the pressure sensors must be insensitive to acceleration and mounted at the pressure tappings. Outputs of the separate transducers need to be given sinusoidal weighting before summation, depending on their circumferential location, to account for the resolution of the lift component of pressure forces. The method requires a number of pressure transducers for good accuracy, and hence tends to be expensive to implement. For example Szechenyi and Loiseau (1975) used 20 transducers at each section of the cylinder. In an earlier publication (Loiseau & Szechenyi 1972), they presented a comparison between lift measurements made directly with a force transducer attached to a $0.25\,D$ long section of a fixed cylinder, and measurements made using the pressure transducers.

Alternatively, it might be possible to approximate the lift force by manifolding tappings from the upper and lower halves of the cylinder section into separate chambers, and then sensing the pressure differential between the two, using a transducer located on the cylinder centre-plane. Only one transducer is then required per section, but there are three significant disadvantages to this method. One is that the circumferential distribution of tappings needs to be weighted so that the pressures at the shoulders of the cylinder have more influence, corresponding to the sinusoidal weighting required in the previous method. The accuracy of this spatial weighting method is open to question, and requires further investigation and comparison with other methods. Secondly, transducer output needs to be corrected for inertial effects. These are caused by the pressure gradient between the outside and centreline of the cylinder due to the acceleration of the fluid in the pressure lines. Bearman and Currie (1979) discuss such corrections applied to single-tapping pressure measurements, but extension of the method to deal with manifolded, distributed tappings is difficult, since tappings located

near the centre-plane of the cylinder have a smaller acceleration-induced response than those further away. Finally, the design of the tubing system is constrained by the need to ensure that (a) all pressure signals reach the transducer at the same time, and that (b) the system has an acceptable frequency response to pressure fluctuations (see e.g. Holmes & Lewis 1987 a, b). No results of measurements of cross flow force made using such a system have yet been published.

The most common of the published methods is to measure the forces applied to some length of an oscillating cylinder, but to subtract a signal corresponding to the component of the force due to the mass and acceleration of the cylinder. A conceptually simple, but mechanically complex way of carrying this out is to cancel the inertial forces directly, by use of counterweights. Fung (1960) used this method, with weights mounted on flexures inside the cylinder. All the other published methods use a transducer to measure acceleration of the cylinder in some way; its signal is then amplified and subtracted from the force transducer output. The amount of amplification is adjusted until the inertial force signal produced by oscillating the cylinder in the absence of flow is cancelled. Variations on the theme have been employed by Bishop and Hassan (1964 b), Jones, Cincotta and Walker (1969), Protos, Goldschmidt and Toebes (1968), Sarpkaya (1978), Howell and Novak (1979), and Steckley, Vickery and Isuymov (1989).

A final option is to ignore the inertial response of the measurement system if the inertial response is small compared to the fluid dynamic force. This may be possible if the fluid is dense enough in comparison to the cylinder. For example, this "method" may have been used by Staubli (1983), who carried out forced oscillations of a cylinder in water, since he did not discuss any cancellation scheme.

**Experimental equipment** During the manufacture of the equipment, care was taken to produce transducer outer segments which were as light as possible, in order to reduce the ratio of inertial to aerodynamic forces. The finished mass of a segment was 16 gm, compared to a mass of displaced fluid of about 0.7 gm. At the maximum amplitudes of oscillation achieved (3% $D$), the inertial component of force transducer output could be as much as ten times the component due to forces due to the flowing fluid, although usually the ratio was smaller. A typical example of signals collected from one transducer in the absence of flow, with the cylinder oscillating at an amplitude of 3% of diameter, is illustrated in figure 4.1. The presence of a large inertial component out of phase with the accelerometer signal is evident in the strain bridge signal. It is this inertial component, also present when flow is applied, that must be cancelled out before aerodynamic forces can be studied. It will be shown that the method cancellation employed was capable of cancelling the inertial component of signal to a residual RMS level of about one analogue-to-digital conversion increment.

**Added Mass** As discussed in § 1.6.6, bodies forced to oscillate in still fluid experience forces exerted by the fluid. If the amplitude of motion is small enough that flow separation does not occur, these forces consist of two types.

The first type, known as *added mass* force, is linearly related to the acceleration of the body. Physically, it is the force required to accelerate a mass of fluid which is

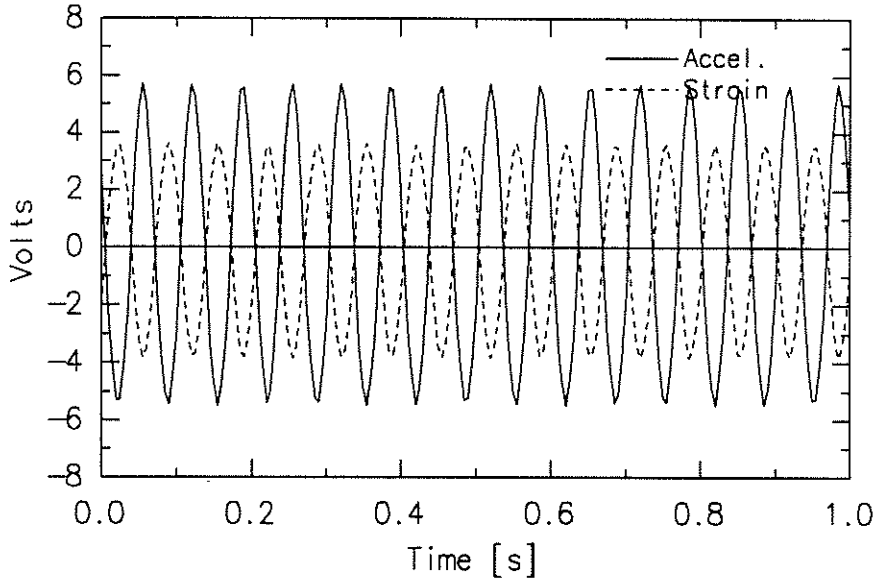


FIGURE 4.1: Acceleration and strain signals produced by oscillating the cylinder in the absence of flow.

moved from one location to another by the passage of the body (Milne-Thompson 1960 § 9.21–2). In potential flow, the added mass of a circular cylinder equals the mass of fluid it displaces. Experimental results (Sarpkaya 1978, Lundgren et al. 1979, Moe & Verley 1980), for circular cylinders with similar oscillation amplitudes and Reynolds numbers to those used here, show that the added mass is close to that predicted for potential flow.

The second type is viscous force, linearly related to cylinder velocity, caused by boundary layer shear stresses. Batchelor (1967, p. 357) presents a boundary layer method for the estimate of this viscous force. Analysis carried out using his method showed that the viscous forces are of the order of one hundred times smaller than the added mass forces for the parameters of this experimental programme (see ch. 1, § 1.6.6).

The point here is that if, to account for inertial forces, the output signals of the force transducers are adjusted to zero while the cylinder is oscillating in still fluid, added mass and viscous drag forces are also removed. Before turning to examine in detail the method by which the no-flow transducer forces were adjusted to zero, note that the added mass forces were subsequently computed and added to the processed force transducer signal, to account for this cancellation. Viscous forces were not computed since they would be below the resolution of the analogue-to-digital converter used for data collection.

**Digital cancellation** The methods described so far which cancel the inertial component of force signal have used analogue equipment. An alternative, employed here, is to use digital cancellation. Two different methods were developed and tried. In one, cancellation was done in the frequency domain; in the other, time domain methods were used. Both methods relied on fitting relationships between transducer accelerometer and strain gauge signals recorded while the cylinder was forced to oscillate in the absence of flow. The relationships were then used to predict the inertial component of

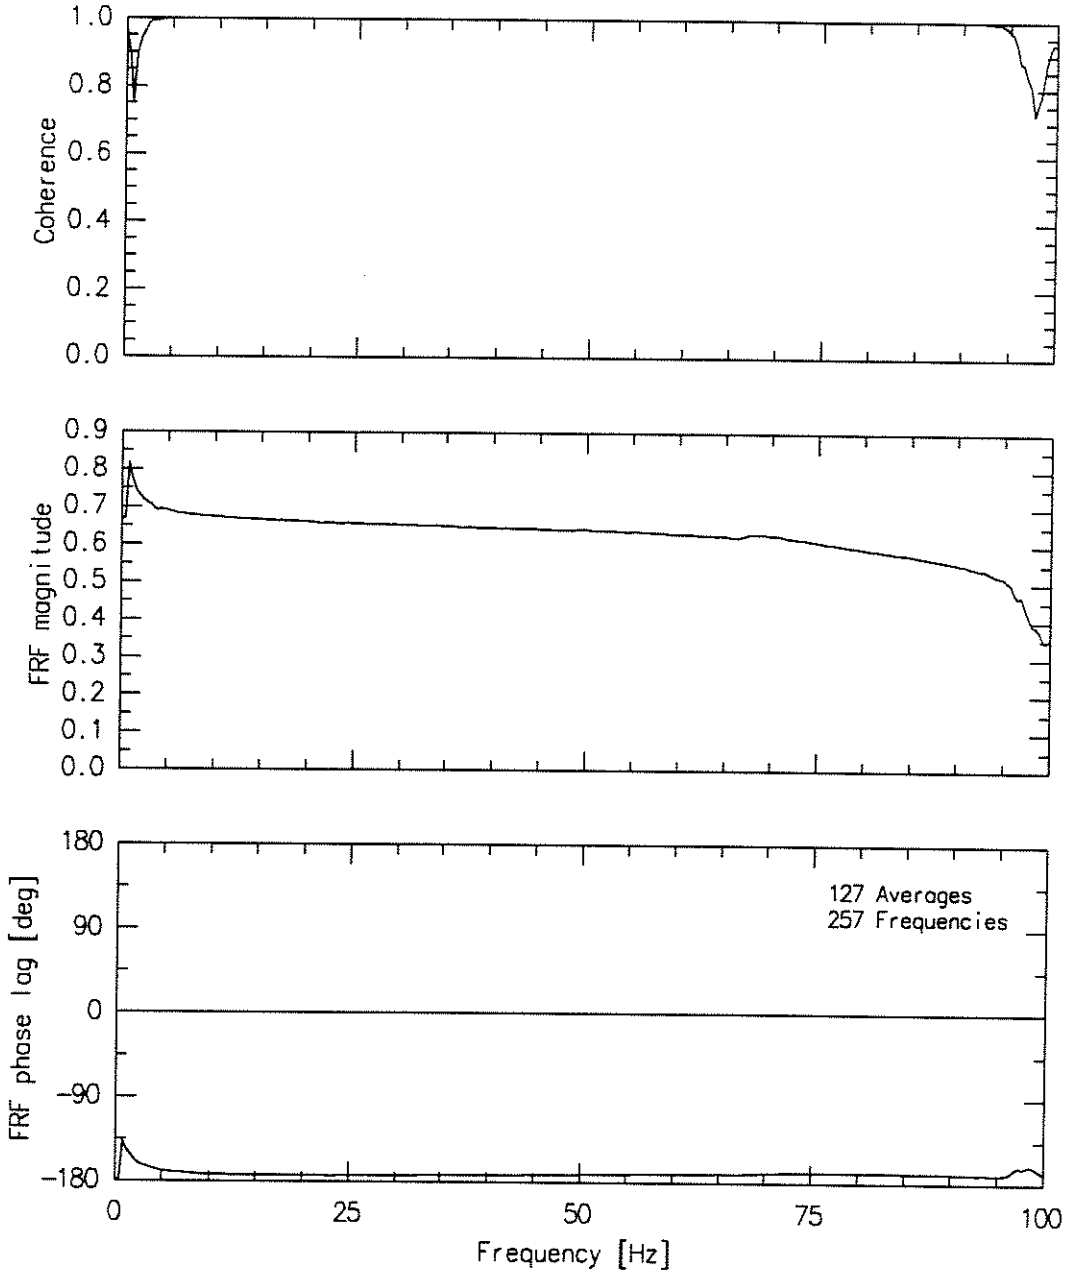


FIGURE 4.2: Frequency response of transducer strain bridge to acceleration.

strain bridge signal when flow was applied, given the output from the accelerometer. The predicted component was subtracted from the total, leaving a residual strain bridge signal representing the forces applied by the flowing fluid alone.

In the frequency domain method, the fitted relationship was an estimate of the transducer Frequency Response Function (FRF) between acceleration and strain, produced from recordings of the response of accelerometers and strain gauges to random excitation of the model applied by the shaker. An example FRF is shown in figure 4.2. It would be expected that the magnitude of the response would be approximately flat over this frequency range, given that the fundamental frequency of the transducer outer segment on the internal beam springs is around 400 Hz. The curvature in the magni-

tude response at low frequencies can be attributed to the lack of coherence between the signals, produced by low signal levels. The lack of flatness at higher frequencies is a result of higher vibration modes of the complete model assembly, leading to rotations that affect the acceleration and strain transducers differently. The behaviour at high frequencies had little or no effect on the measured aerodynamic forces, provided the relationship between acceleration and strain is modelled correctly.

The frequency domain method is described briefly in the paper by Blackburn and Melbourne (1989 b), presented in appendix B. It was less successful than the time domain method, principally because the characteristics of the transducer varied slightly with the amplitude of oscillation of the model[1]. Since the random excitation needed to compute the FRF estimates produced a range of oscillation amplitudes, it was difficult to estimate just what the characteristics were at any particular amplitude of sinusoidal oscillation to be used in wind tunnel running. Detailed investigation failed to isolate the cause of the variation. This failure resulted from conflicting evidence, but it is most likely that the variation was caused by the characteristics of the polyurethane film used to seal gaps between sections of the model.

The time domain method was based on fitting a digital Impulse Response Function (IRF) that described the relationship between accelerometer and strain bridge timeseries recorded with the cylinder oscillating in the absence of flow. The IRF was convolved with the signal from the transducer accelerometer recorded in the presence of flow to produce an estimate of the timeseries of inertial component of force transducer response.

### Adaptive Digital Filters

The rest of this section is devoted to description of the time domain method. Problems due to change in transducer characteristics with oscillation amplitude were effectively overcome by this method, since the IRFs were estimated using a short record of data taken in the absence of flow before each tunnel run, with the model oscillating at the amplitude and frequency of interest.

The theory of digital filters has been extensively developed over the last three decades and a number of textbooks which treat the subject at a variety of levels are now available. The following points should be sufficient for this discussion; the reader is referred to the books by Stearns (1975), Oppenheim and Schafer (1975), Kuc (1988), and Bellanger (1989) for more detail.

1. A digital filter is broadly similar to an analogue filter, in that the input signal is convolved with the system function in the time domain to produce an output signal. (The word filter is used here in a general sense to describe any system which operates in this way.)

2. Digital filters are different from analogue ones in that the signals are only defined at discrete times, and convolution takes place by summation rather than integration.

[1] See figure 2 of the paper.

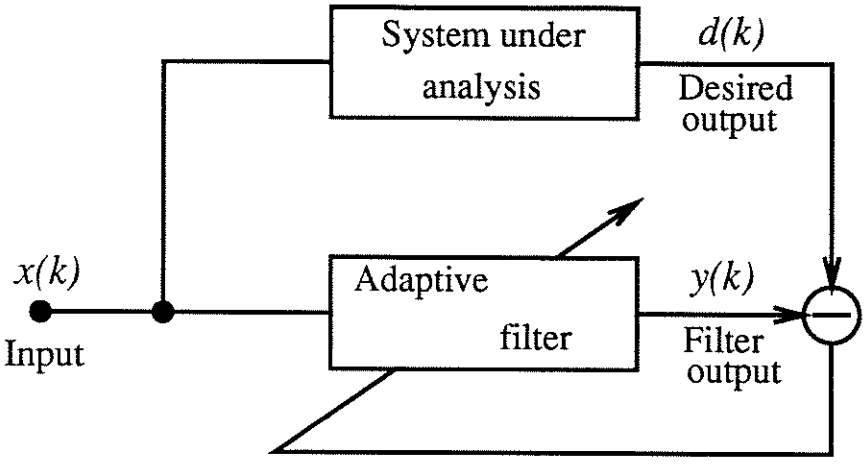


FIGURE 4.3: Schematic diagram of the use of an adaptive digital filter in system identification.

3. The characteristics of a digital filter are completely specified by its IRF, which is the sequence of values produced by applying the unit sample (impulse) to the filter.

4. Digital filters are categorized into two groups, depending on whether the IRF terminates after a finite number of time-steps, or never terminates; the filter is described as being Finite Impulse Response (FIR) or Infinite Impulse Response (IIR), respectively. FIR filters have no linear continuous-time equivalent, and bear a close relationship to the Discrete Fourier Transform (DFT) (Kuc 1988 p. 318).

FIR filters were used in this application: the impulse response can be visualized as a sequence of numerical weights as in figure 4.5.

An *adaptive* digital filter is one in which the IRF weights can be adjusted between one timestep and the next. This adjustment is usually carried out in order to minimize the error between the output of the filter and some reference signal. One situation in which an adaptive filter may be used is in the problem of *identification*, where the input and output of some physical system are available and one wishes to construct an estimate of a digital system which relates the two. In this case, the reference signal is the output of the physical system, and the IRF weights are adjusted to minimize the difference between the output of the digital filter and that of the physical system (fig. 4.3). This adaption is not confined to the case of stationary signals, and the development of the theory of adaptive digital filters has taken place partly in response to the need to model real systems which change over time.

Here, the physical system to be modelled was the force transducer, the input signal was the output of the accelerometer amplifier, and the output desired of the filter that of the strain bridge amplifier when the cylinder was forced to oscillate in the absence of flow. Signals recorded with the cylinder oscillating at the appropriate frequency and amplitude before each wind tunnel run were used to compute the IRF weights for the transducers.

## Theory

In the application to the present problem, the adaptive capability is not itself important, since the system to be modelled does not change over time. The set of equations which the method then solves is the *Yule-Walker*, or *Normal* set. The theory which lies behind these equations has been known for some time, and forms part of the theory of time series analysis (Chatfield 1984). Solutions of the equations can be obtained using standard matrix techniques, under some circumstances. Adaptive methods were chosen instead mainly because the algorithms themselves are concise (since they have been developed to run in hardware), and FORTRAN codes for them have been published (by Bellanger 1987, and Triechler, Johnson & Larimore 1987). Rather than describe the theory which lies behind the adaptive methods, the development of the equations they solve is given, and the reader is referred to these texts for more detail. The presentation here follows that of Triechler, Johnson and Larimore.

An FIR digital filter is a set of numerical weighting factors $\{w_l\}$ which define the filter response to a unit impulse. The filters to be discussed here are *causal*, that is, they deal only with the current sample and past values of the input samples $x(k), x(k-1), x(k-2) \ldots x(k-L)$. FIR filters have by definition a finite set of weights which can be conveniently thought of as a vector

$$\mathbf{W}^t = [w_0\, w_1\, w_2\, w_3 \ldots w_{N-1}]\,, \tag{4.1}$$

with subscripts running from 0 to $N-1$ for an impulse response length, or *order* of length $N$. It is also convenient to regard the subset of the input sequence $x(k) \ldots x(k-L)$ being operated on by the filter as a data vector of length $N$ $(L \geq N)$

$$\mathbf{X}(k) = [x(k)\, x(k-1)\, x(k-2)\, x(k-3) \ldots x(k-N+1)]^t\,. \tag{4.2}$$

The output sequence $y(k)$ is produced by convolving the input sequence $x(k)$ with the filter impulse response $\{w_l\}$

$$\begin{aligned} y(k) &= \sum_{l=0}^{N-1} x(k-l) w_l \\ &= \mathbf{W}^t \mathbf{X}(k). \end{aligned} \tag{4.3}$$

The sequence of samples of the *desired* waveform, $d(k)$, is used as a reference. The problem is to adjust the impulse response weights $\mathbf{W}$ so that $y(k)$ matches $d(k)$ as closely as possible. In this application, $x(k)$ is the sequence of samples of the output of an accelerometer mounted in a transducer, and $d(k)$ is the sequence of samples of the strain bridge amplifier output, both recorded in the absence of flow. The filter weights are determined so as to minimize the sum of squared errors

$$\begin{aligned} J_{ss} &= \sum_{k=N-1}^{L-1} |d(k) - y(k)|^2 \\ &= \sum_{k=N-1}^{L-1} \left|d(k) - \mathbf{W}^t \mathbf{X}(k)\right|^2 \\ &= D_{ss} - 2\mathbf{W}^t \mathbf{P}_{ss} + \mathbf{W}^t R_{ss} \mathbf{W}\,, \end{aligned} \tag{4.4}$$

where

$$D_{ss} = \sum_{k=N-1}^{L-1} |d(k)|^2 \ , \tag{4.5}$$

and

$$\begin{aligned} \mathbf{P}_{ss} &= \sum_{k=N-1}^{L-1} x(k-i) \cdot d(k), \quad 0 \le i \le N-1 \\ &= \sum_{k=N-1}^{L-1} d(k)\mathbf{X}(k) \end{aligned} \tag{4.6}$$

is an estimator of the cross-correlation vector $\mathbf{P}$ between the input and desired sequences;

$$R_{ss} = \sum_{k=N-1}^{L-1} \mathbf{X}(k)\mathbf{X}^t(k) \tag{4.7}$$

is an estimator of the input sequence autocorrelation matrix $R$, where the $i,j$th term of $R_{ss}$ is given by

$$r_{ij} = \sum_{k=N-1}^{L-1} x(k-i) \cdot x(k-j) \ . \tag{4.8}$$

To find the weights which minimize the sum of squared prediction errors, we set

$$\nabla_{\mathbf{W}} J_{ss}|_{\mathbf{W}=\mathbf{W}_{ss}^o} = \left.\frac{\partial J_{ss}}{\partial w_j}\right|_{\mathbf{W}=\mathbf{W}_{ss}^o} = \mathbf{0}, \quad 0 \le j \le N-1 \tag{4.9}$$

in equation 4.4, and find

$$R_{ss}\mathbf{W}_{ss}^o = \mathbf{P}_{ss} \ . \tag{4.10}$$

These are the Normal equations, which are to be solved for the optimal set of filter weights $\mathbf{W}_{ss}^o$. If $R_{ss}$ is non-singular it may be inverted to provide the solution:

$$\mathbf{W}_{ss}^o = R_{ss}^{-1}\mathbf{P}_{ss}. \tag{4.11}$$

If $R_{ss}$ is singular there may be infinitely many solutions of 4.10, but the inversion in 4.11 cannot be carried out.

Since it is often desired to realize the computation of the filter weights in hardware, and in some applications to produce filters which are capable of adapting filter weights as the system being modelled changes, the emphasis in digital signal processing applications has been towards solving equation 4.10 iteratively for the filter weights, that is, producing successive estimates of $\mathbf{W}_{ss}^o$ which converge to an optimal solution. There are several advantages of operating in this way, besides the adaptive capability just mentioned. Iterative schemes are often able to provide *a* solution to equation 4.10, even if $R_{ss}$ is singular, and are better behaved than direct schemes if it is nearly singular. Also the numerical schemes are usually more simple than the direct schemes since

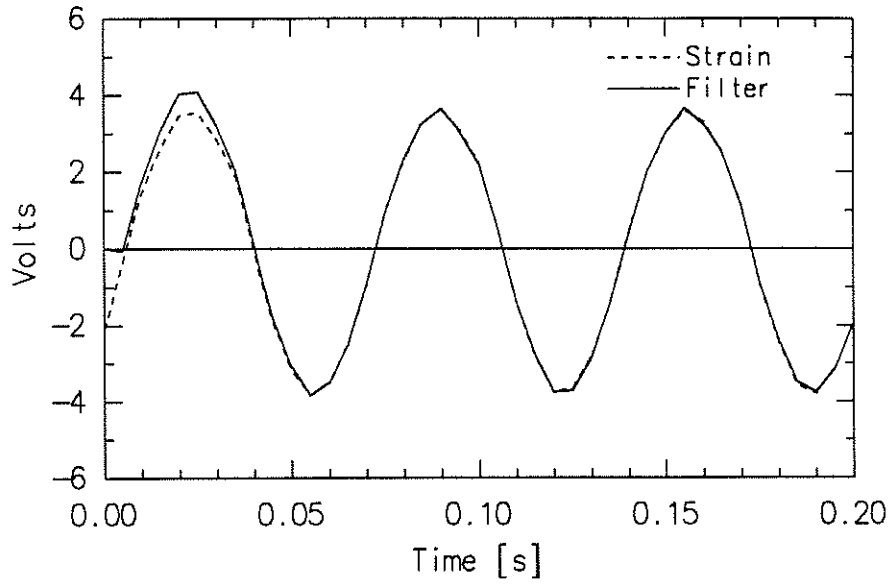


FIGURE 4.4: Demonstrating the adaption phase of 32-weight FIR digital filter.

they must be implemented in hardware. The method used here was the Fast Recursive Least Squares algorithm, which achieves speed and simplicity by avoiding matrix operations completely. Bellanger (1987, p. 209) presents a FORTRAN subroutine for this algorithm; the Pascal procedure `adaptgain`, given in appendix C, is a translation of his subroutine `FLS2`.

### Example

As a demonstration of the action of the adaption of the filter weights, consider figure 4.4, which shows strain signals obtained from one transducer while the model was oscillating in the absence of flow. Also shown in the figure is the output obtained from a 32-weight FIR filter as the filter weights are adapted so that the output of the filter matches the strain bridge signal, using the accelerometer signal as input. Filters with 32 weights were used throughout the data reduction. This length was chosen after finding that even short filters of 10 weights could perform the task of cancelling the inertial signal adequately in the absence of flow—32 weights were used to be conservative. As can be seen in figure 4.4, there is little observable difference between the measured and predicted transducer strain signals after 100 samples.

An example of an impulse response produced by the procedure is shown in figure 4.5.

### Errors

To examine what degree of cancellation of inertial component of transducer strain bridge signal could be achieved, tests were run in which the model was oscillated in the absence of flow. First, a set of data was recorded for adaption purposes, then at the same amplitude and frequency another set of data was collected to examine cancellation. This procedure replicated the experimental conditions as far as possible, with the exception

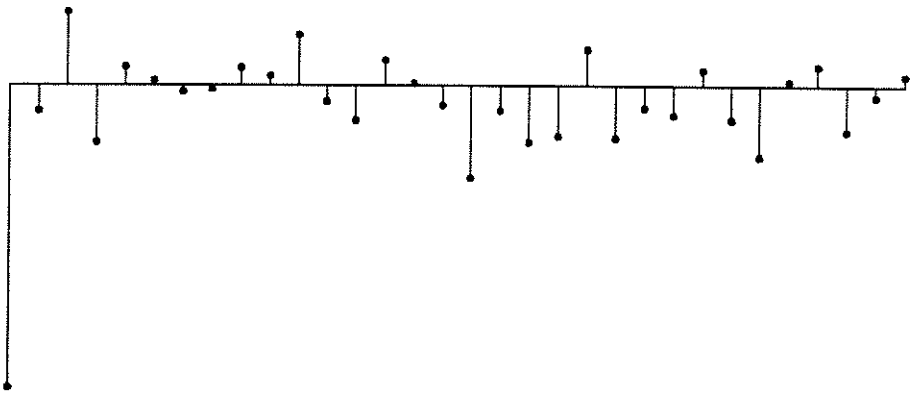

FIGURE 4.5: A typical force cancellation FIR filter impulse response.

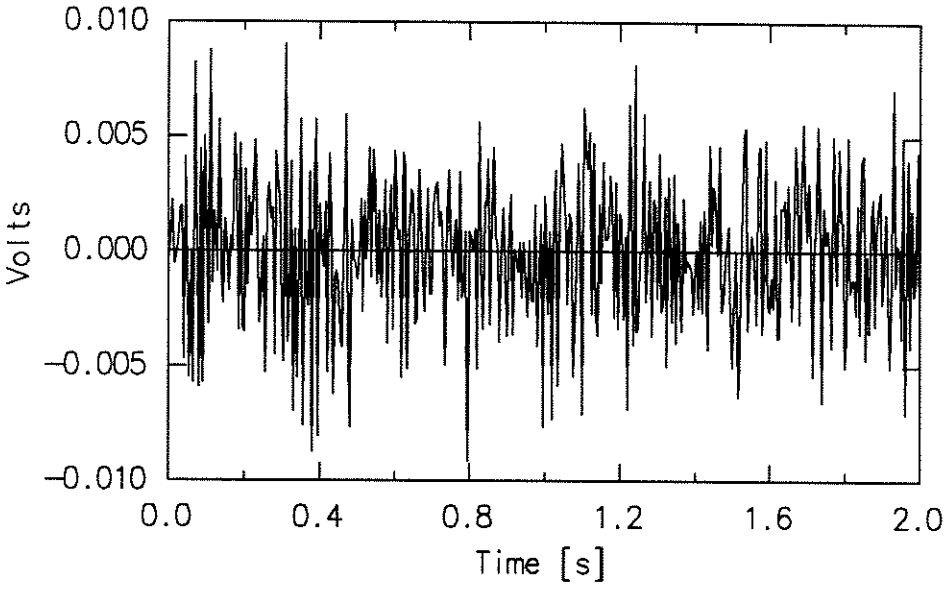


FIGURE 4.6: Demonstrating the degree of cancellation achieved. Strain bridge signal after cancellation.

that there was no flow during the second data collection. An example of the residual after cancellation can be examined in figure 4.6. The input signals were in this case the same as those shown in figure 4.1. The significance of the results shown in figure 4.6 is that the resolution of the analogue-to-digital conversion is 5 mV; the RMS size of the residual is of the same order. It is unlikely that better cancellation can be achieved.

A spectrum of the error signal is given in figure 4.7, where it can be seen that the error is concentrated at the line frequency (50 Hz) and harmonics of the shaking frequency (15 Hz).

In operation, it was typical that data collections of 1024 samples were taken with the flow turned off, for adaption purposes, followed by 8192 samples with flow. The short records were used for computation of the FIR filter coefficients, which were then

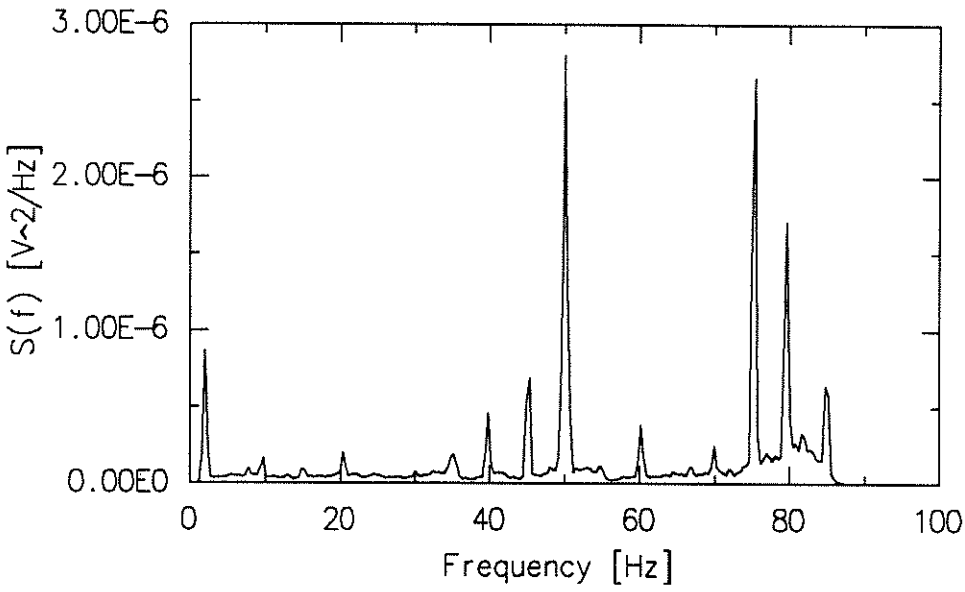


FIGURE 4.7: Spectrum of error signal.

applied to the task of cancelling the inertial component of the records recorded with flow. While it is difficult to directly assess the accuracy with which cancellation is achieved in the presence of flow, the no-flow results are encouraging. Comparisons, presented in chapter 6, of motion-correlated forces measured in smooth flow with those found by other workers indicate that the remaining forces at the frequency of cylinder motion are aerodynamic in origin and are not the result of imperfect cancellation of inertial forces.

#### Transient Effects

Transient effects associated with the start-up of the filters corrupted the beginning of the processed force timeseries (see fig. 4.22). This was accounted for in the calculation of spectra and probability densities of forces by ignoring the first 50 corrected data.

### 4.2.2 Compensation for Time Lag in Data Collection

#### Cause of Time Lag

When data were collected from the model in normal tunnel running, twelve data lines were connected to the computer for analogue-to-digital conversion, so that information about forces acting on the six transducers could be gathered at the same time. The conversion unit contained only one processor. This meant that during each clock cycle from the sampling clock, the converter was switched rapidly between one data channel and the next, so the digitizing was not simultaneous at each channel. A numerical method was devised to correct the results for this lack of simultaneity. The time delay between each channel produced an apparent phase *lead* from one channel to the next, as will be explained below.

To see the effects of a sampling time lag, consider the sketch of figure 4.8(a), which

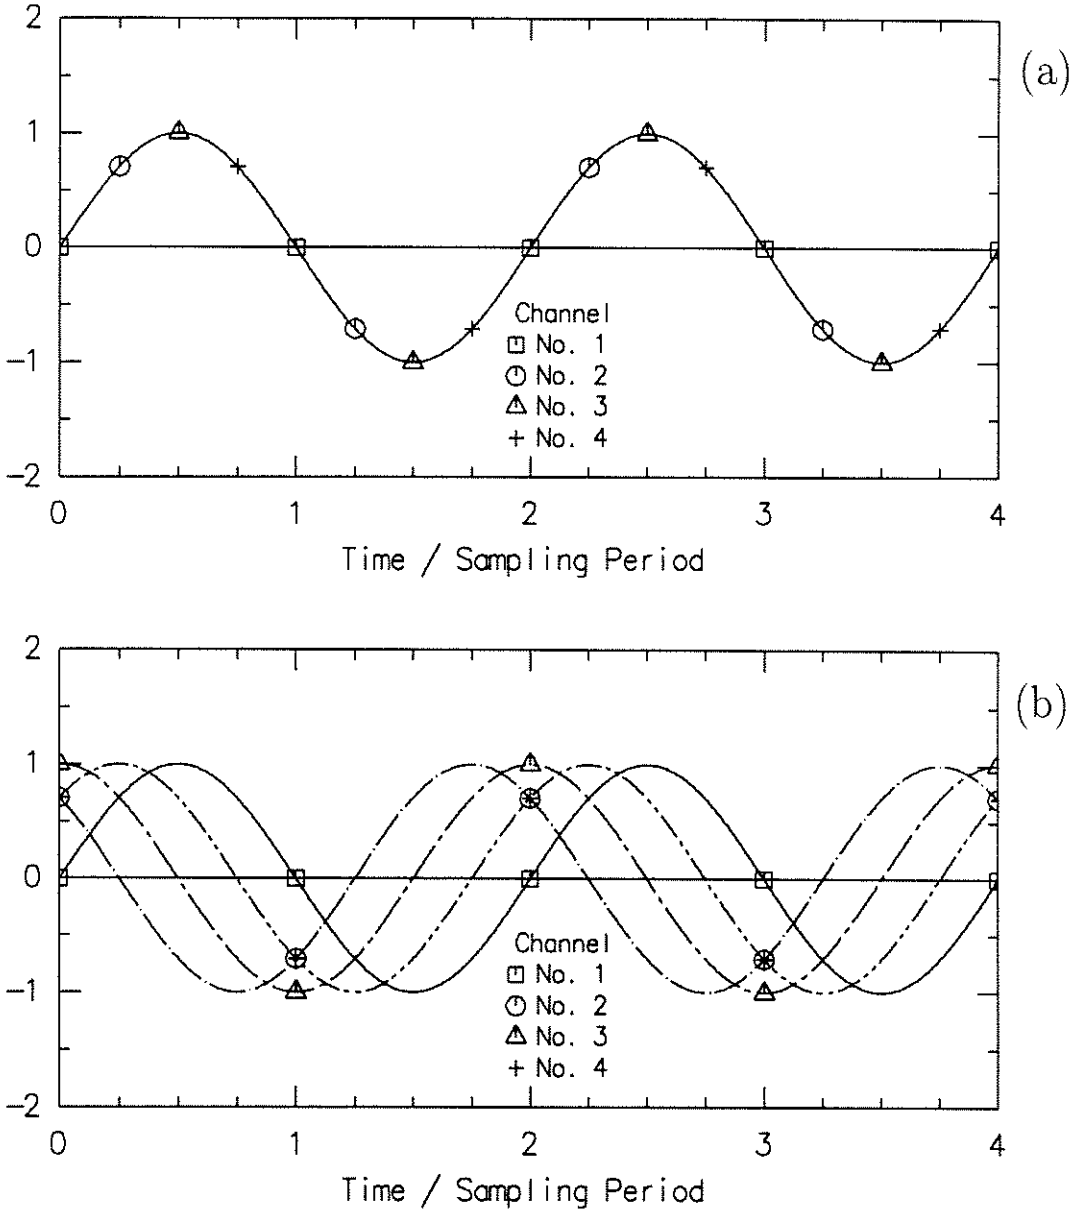


FIGURE 4.8: Illustrating the effects of sampling time lag on a sinusoid which oscillates at the Nyquist frequency. (a); Four-channel sampling of the same signal. (b); Interpolating sinusoids for the four sets of sampled data.

shows a sinusoidal signal that oscillates at the Nyquist frequency. In this hypothetical situation, the sinusoid is used as input to four data channels. In figure 4.8(b), samples from the four channels and their interpolating sinusoids are drawn, with the samples placed at the same points in time, as they would be represented in the output of the computer. It can be seen that each sinusoid has a phase lead, relative to the previous one, of $\pi/4$.

### Compensation by Phase Shifting

The above example relates to sinusoids at the Nyquist frequency, but the amount of phase shift at other frequencies is easily determined. The amount of phase shift is linearly related to the amount of time delay and the frequency, as expressed by the time delay/phase shift duality of the Fourier transform (Brigham 1988, p. 35). For continuous-time signals, a delay in the time domain is exactly equivalent to a phase shift in the frequency domain that varies linearly with frequency. This is expressed in the relationship (Otnes & Enochson 1978, p. 365)

$$\phi(f) = -2\pi f\tau \tag{4.12}$$

where $\phi(f)$ is the phase difference at frequency $f$, and $\tau$ is the time delay. The duality suggests that an appropriate way to account for the time lag between channels is to perform a phase shift on each point of the DFT of the data, for each channel after the first. The effect is equivalent to interpolating the sampled data to obtain estimates for the signals that would have been present on all channels each time the first channel was sampled.

As a consequence of the Sampling Theorem (Stearns 1975, p. 37), the phase shift in the frequency domain can successfully interpolate the data between sampling points if the original signal was frequency limited to below the Nyquist frequency and of infinite length. In practice, even though the original analogue signal may have been appropriately low-pass filtered, the digital interpolation cannot work exactly due to time-domain wrap-around of the interpolating function (see e.g. Kuc 1988, p. 341) and end-effects (*leakage & aliasing*). These effects are both due to the fact that a limited number of sampled data are available for interpolation, rather than an infinite sequence. Provided the recordings of data are long enough however, the resulting inaccuracies are insignificant away from the ends of the data records.

What remains is to compute the amount of time lag present between successive data channels. This is dependent on the way the sampling scheme is implemented. Naturally, all channels of data were sampled between one impulse and the next from the external clock source used to trigger the sampling. The way in which this was carried out was to space the sampling for the channels equally between the external clock signals, provided the number of channels requested was a power of two. If the number of channels requested was not a power of two, the sampling was carried out as though the number required was the next higher power of two, but only data from the channels requested were recorded. For example, the majority of the sampling carried out in this project used 12 channels. This meant that sampling was spaced out between channels as though 16 had been required, but the data files were written with only 12 channels, as requested. Provided the number of channels stored was known, it was then relatively easy to compute the delay between channels as a proportion of the sampling period, and it is only this proportion which is needed in the digital correction. Procedure `delay` in appendix C implements the numerical method just described.

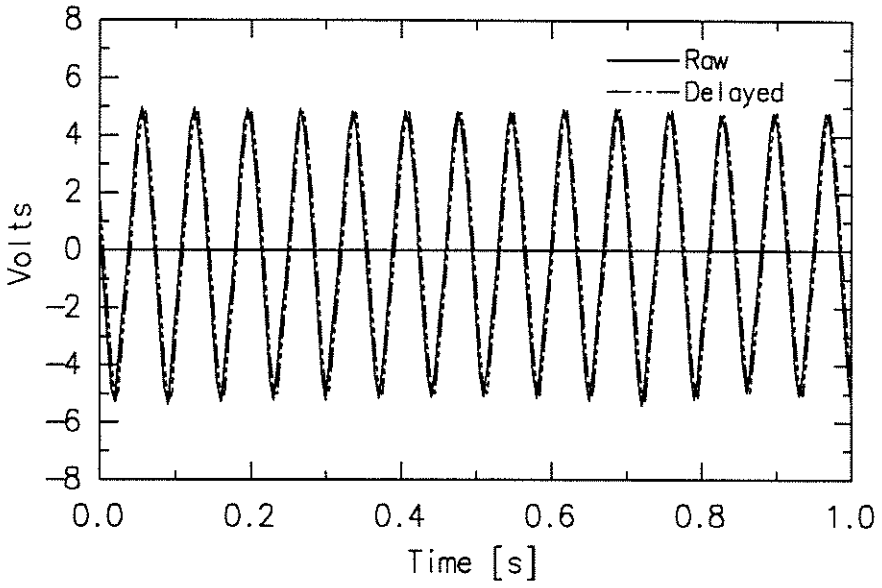


FIGURE 4.9: Sampling lag correction applied to signal from the accelerometer of the sixth transducer.

### Example

An example of the sampling lag correction is shown in figure 4.9, where signals from the accelerometer of the sixth transducer (hence the largest delay) are shown as collected, and after correction for time lag effects.

## 4.2.3 Correction of Analogue Frequency Response Functions

### Reasons for Correction

In addition to correcting signals for inter-channel lag caused by data collection hardware, further digital correction of signals was done.

All channels of data were corrected to account for the differences between low-pass filters used on each channel, because as previously explained in chapter 2, 12 matched filters were not available, so the time-lag produced by filtering was in general different from channel to channel. To show the extent of the difference, FRFs of two filters are presented in figure 4.10. (FRFs were estimated using random data analysis techniques, after applying a noise signal to the input of the filters and recording the inputs and outputs.) Using equation 4.12 in the low frequency ranges of figure 4.10, the difference in lags produced by the two filters is about 0.55 ms, which is 11% of the sampling period.

The signals produced by the accelerometers were further corrected because it was considered that the low-frequency characteristics of the accelerometers had undesirable defects. The accelerometers used in the model were calibrated against a Sundstrand servo-accelerometer which had a flat frequency response down to zero frequency. Static calibration of the servo-accelerometer consisted of aligning the accelerometer with the Earth's gravity field, then inverting it, and recording the difference between voltages

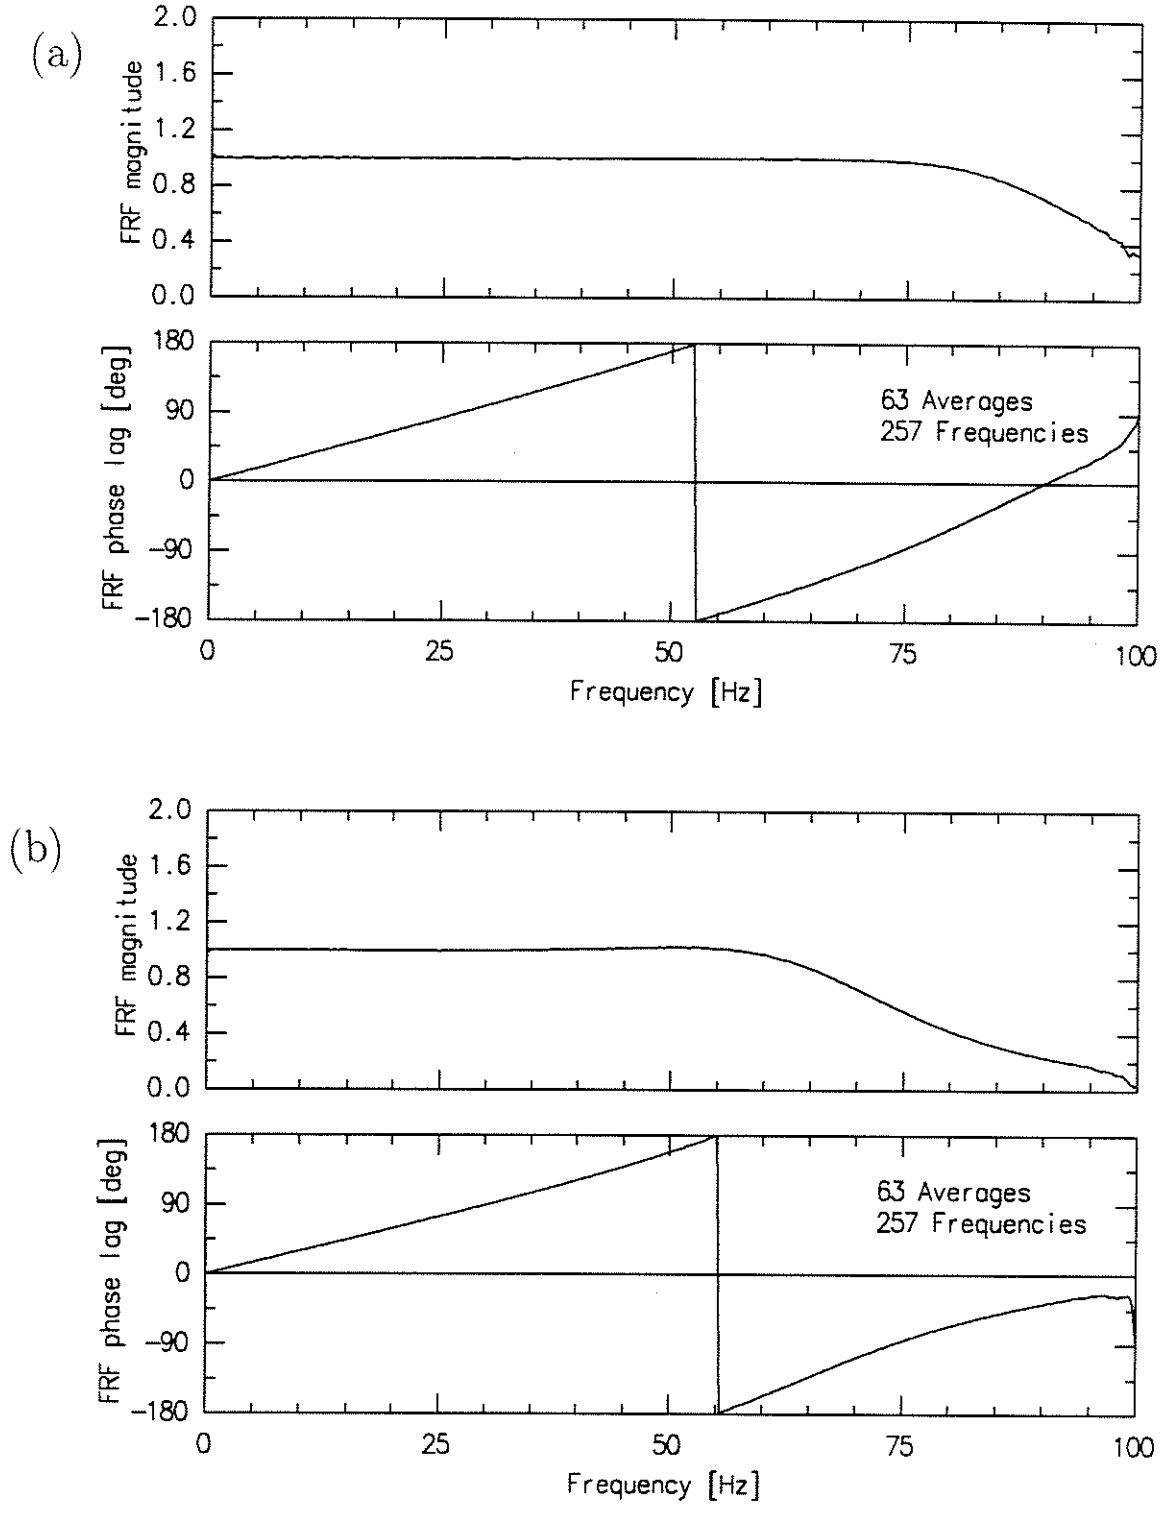


FIGURE 4.10: Frequency Response Functions for two different anti-aliasing filters used in the experiment.

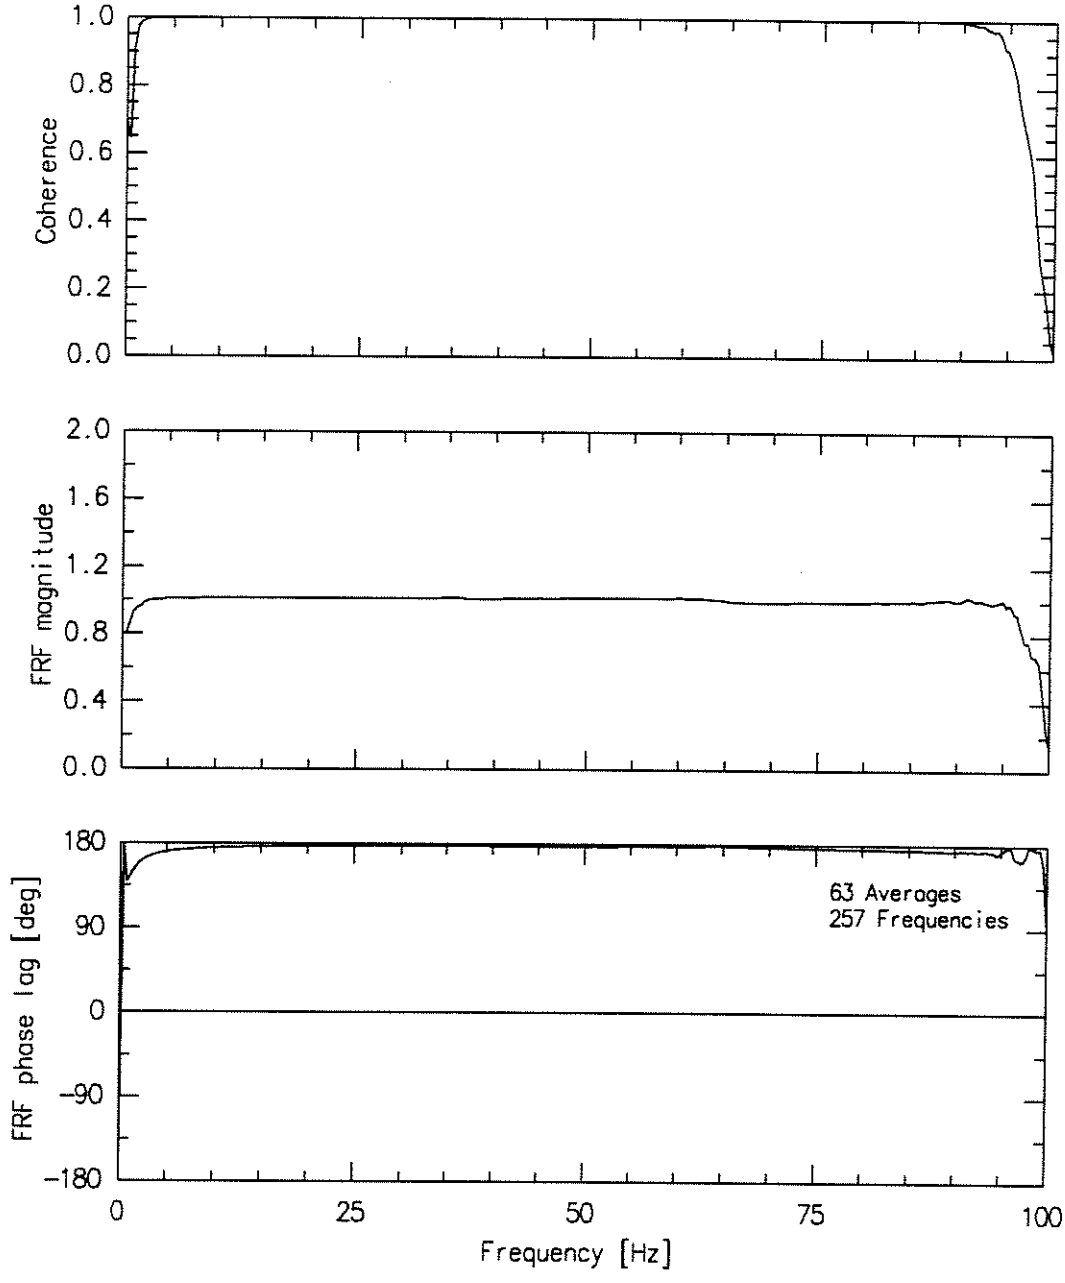


FIGURE 4.11: FRF between Sundstrand and Vibra-metrics accelerometers.

output by the accelerometer signal-conditioning unit. Relative calibrations were obtained by mounting both types of accelerometer atop the shaker used to oscillate the cylinder model, and supplying a random input signal to the shaker. The FRFs between the two sets of signals was then computed. An example of a measured accelerometer FRF is presented in figure 4.11. This shows that the magnitude and phase responses of the accelerometer installed in the model were not flat, particularly at low frequencies. The correction carried out to the accelerometer signals is intended to deal with this deficiency.

### Method

The idea behind the method is that the original physical signal—acceleration or strain—has passed through an analogue system which produced distortion. The original signal can be reconstructed by passing the distorted signal through a process which "undoes" or inverts the work of the analogue system which first produced the distortion. The processing needed is deconvolution or *inverse filtering* of the distorted signal. If the analogue system can be identified, by measuring its FRF or IRF, then the deconvolution can be carried out in either the frequency or time domains. Constructing inverting digital filters for time domain work is discussed by Kuc (1988) and by Robinson and Treitel (1980). In general, carrying out the inversion in the frequency domain is a much simpler proposition since convolution in the time domain is equivalent to multiplication in the frequency domain[2].

If the FRF of the analogue system is available, deconvolution can be carried out by dividing the DFT of the sampled data by the FRF of the analogue system. Timeseries of the deconvolved data are then produced by performing an IDFT on the result. Due to the fact that only a limited length of sampled data can be recorded, the process of convolution will produce end effects or "ringing" at the two ends of the data set, however these effects in general become negligible in the centre of the set. If the data set is long enough, the end effects will be restricted to a comparatively small part of the data and can in the most part be ignored.

In order to perform the deconvolution, estimates of the FRFs of the analogue equipment (in this case the low-pass filters and accelerometers) were needed. The estimates used in the processing were produced using a two-stage process. In the first stage, estimates of the FRF were produced by applying a random input to the transducers, recording input and output timeseries, and computing the FRF and coherence using the input and output autospectra, and the cross spectrum of the two processes (Bendat & Piersol 1971, ch. 9)[3]. FRF estimates thus produced could have been used in the calculations, essentially as look-up tables, with values at the frequencies used in the correction of timeseries found by interpolation. Instead of doing this directly, another stage of processing was added.

Experimental estimates of the FRFs had rational polynomial FRFs fitted to them using the least-squares technique derived by Levy (1959). The fitted FRFs have the form

$$H(jf) = \frac{A_0 + A_1(jf) + A_2(jf)^2 + \cdots + A_n(jf)^n}{1 + B_1(jf) + B_2(jf)^2 + \cdots + B_n(jf)^n}. \tag{4.13}$$

Typically, $n$, the order used, was six. Only experimental values of FRF for which the coherence exceeded 95% were used as a basis for the curve-fitting. A comparison of experimentally-estimated and fitted FRFs for a low-pass filter is presented in figure 4.12.

[2]An introduction to the processes of frequency domain convolution and deconvolution is given by Brigham (1988), while Press et al. (1986, ch. 12) give an overview of the field and an interesting discussion about the relative merits of time and frequency domain methods of convolution.

[3]The method of overlapping segment-averaged periodograms (Press et al. 1986, p. 248) with a *Hanning* data window was used for all spectral estimates

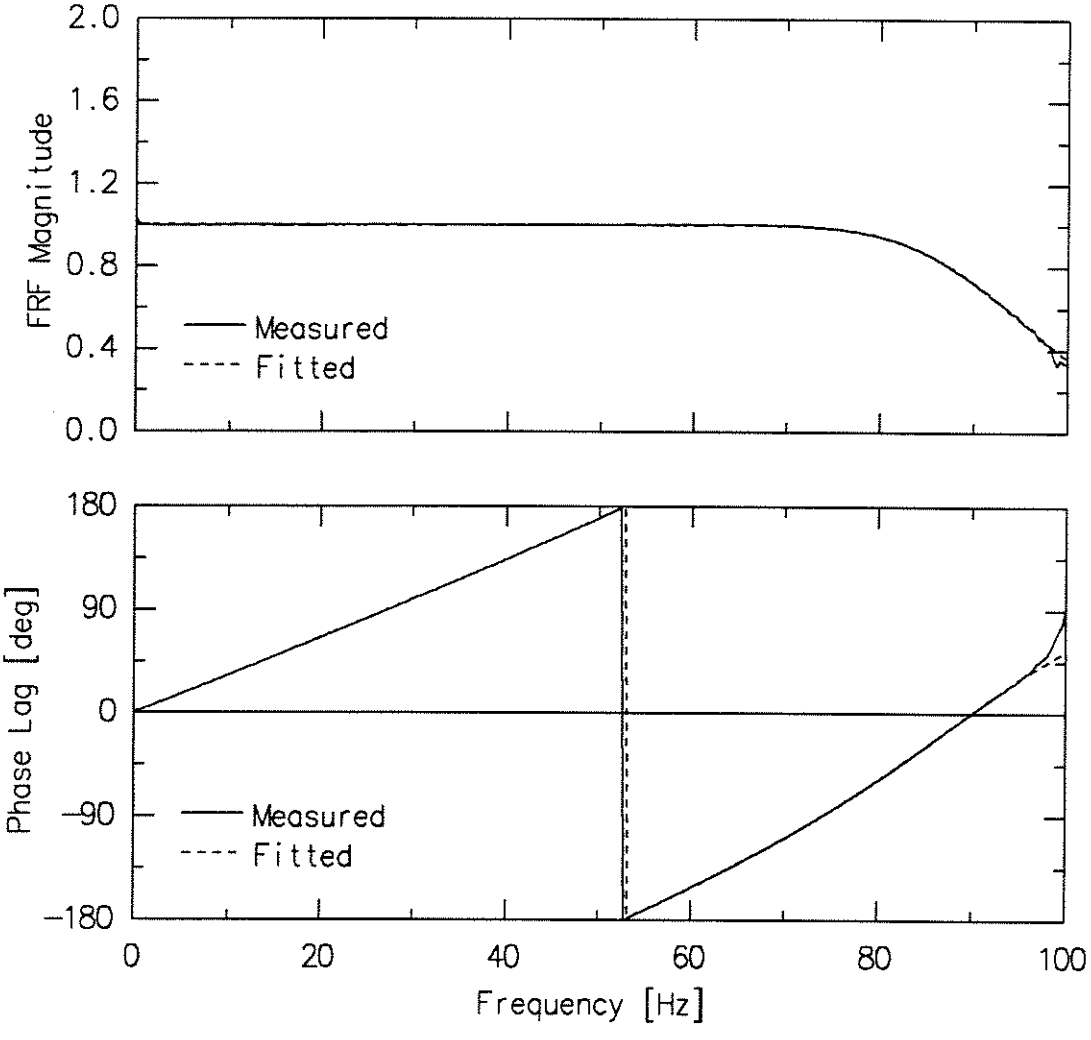


FIGURE 4.12: Comparison of measured and fitted frequency response function for a low-pass filter.

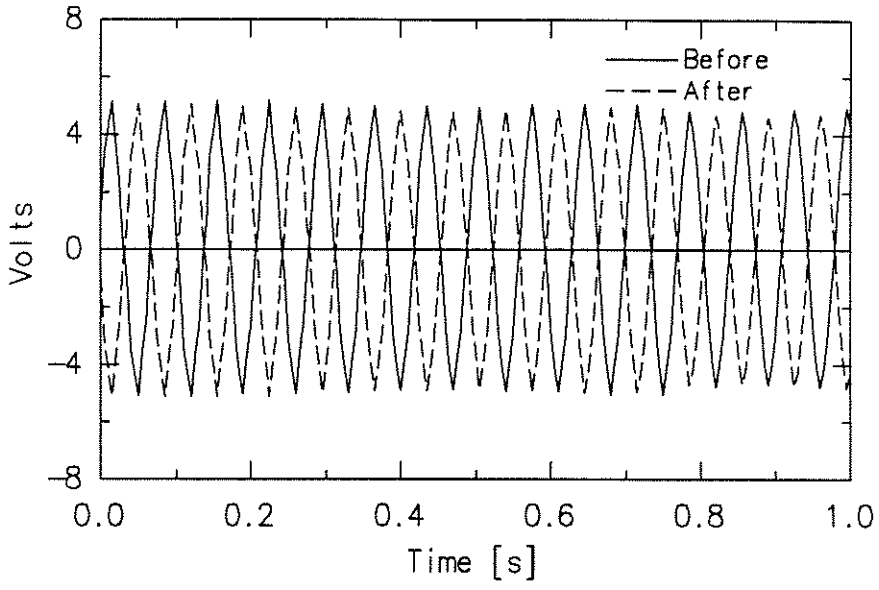


FIGURE 4.13: Accelerometer signal with and without compensation.

These fitted FRFs were used in processing, in favour of the measured FRF estimates from which they were derived, for two reasons. First, they provided a smooth estimate of the FRF; experimental results always have some "noise", since using the random excitation techniques the variance of the estimate produced at each frequency is never quite zero, but approaches it asymptotically as the number of segment averages increases (see e.g. Bendat & Piersol 1980, ch. 11). Second, the use of low-order fitted FRFs gives a greater degree of confidence in the regions where the coherence of the experimental estimates is low. This can be seen in figure 4.12, where at the higher frequencies the experimental estimates of the FRF deviate from a smooth curve due to low coherence produced by low signal levels. In the case of the anti-aliasing filters, the confidence in a low-order fit is increased by knowledge that the filter circuits can be described in the frequency domain using a small number of poles and zeros; effectively, carrying out the fit serves to locate them.

### Example

The correction achieved by inverting the anti-aliasing filter FRFs is difficult to observe for most of the signals dealt with, since the difference in time-lags produced by different filters is small compared to the sampling period, and the upper frequency limit of the signals recorded is usually a small proportion (typically 25%) of the Nyquist frequency. The low-frequency phase shift dealt with by inverting the accelerometer FRFs is similarly difficult to observe, as the example is given in figure 4.13 shows. Computations of frequency response, using the uncorrected and corrected records, do however reveal that the corrections have been properly applied.

### 4.2.4 Integration via the Frequency Domain

#### Application

After correction for sampling time lags and deficiencies in FRFs, the cylinder cross flow acceleration timeseries were integrated to give velocity timeseries. The motivation for this was that part of the later examination of the aerodynamic force timeseries would be to remove components correlated with cylinder acceleration and velocity. Timeseries of both acceleration and velocity were needed so that the correlation could be carried out on a least-squares basis.

#### Theory

In numerical terms the integration must use discrete equally spaced function values, which correspond to the sampled data points. Classical schemes for performing this kind of integration—such as the Trapezoidal, Simpson's and Three-Eighths Rules—are designed to provide exact integrals of polynomials of various orders, but are not necessarily suitable for the integration of timeseries, which are more naturally interpreted as the sum of sinusoids of various frequencies.

The performance of integration schemes can be examined in the frequency domain, as explained by Hamming (1986). One supplies sinusoids of different frequencies as the input to the schemes, and examines the magnitude and phase of the sinusoids produced. A perfect integration scheme will transform any input sinusoid by shifting its phase by $-\pi/2$ and reducing its magnitude by $2\pi f$ (i.e. complex division by $j2\pi f$). In general, the classical schemes can return accurate estimates of the integral, provided the significant frequency content of the sampled function is well below the Nyquist frequency. If that is not the case, the classical schemes can return very poor estimates of the integral. While it is possible to design schemes which overcome this difficulty[4], it was decided to examine a different method, based on term-by-term integration of Fourier series.

During the processing of the acceleration and force timeseries, the data were transformed to the frequency domain using a DFT, so that the correction for sampling time lags and transducer FRF imperfections could be carried out as discussed in the preceding sections. This meant that no extra numerical work was needed to produce the DFT coefficients and a method which performed time-domain integration by means of frequency domain calculations was attractive. According to the theory of Fourier transforms (Papoulis 1962, p. 40), time domain integration of the function $g(t)$ is equivalent to division of its Fourier transform $\mathcal{G}(\omega)$ by $j\omega$, and by multiplying the zero frequency value by $\pi\delta(\omega)$, where $\delta(\omega)$ is the impulse function:

$$\int_{-\infty}^{t} g(\tau)d\tau \xleftrightarrow{\mathcal{F}} \pi\mathcal{G}(0)\delta(\omega) + \frac{\mathcal{G}(\omega)}{j\omega} . \tag{4.14}$$

This suggests that division of the DFT coefficients by $j\omega$ before use of the IDFT will produce an integral of the timeseries. There is uncertainty however, about the treatment

[4]Hamming (1986, ch. 36) discusses Chebyshev design of integration schemes, where weighting coefficients are chosen to minimize the maximum error of the integration over a range of frequencies.

of the zero frequency term $\mathcal{G}(0)$, and about the correct scaling to use for the DFT coefficients. To resolve these problems, the following treatment was derived, based on the interpretation of the DFT coefficients as Fourier series coefficients of the periodic extension of the sampled section of timeseries (Stearns 1975, p. 58).

Consider the DFT-IDFT pair

$$\begin{aligned} \mathcal{G}\left(\frac{n}{NT}\right) &= \sum_{k=0}^{N-1} g(kT) e^{-j\frac{2\pi nk}{N}} \quad n = 0, 1, \ldots, N-1 \\ g(kT) &= \frac{1}{N} \sum_{n=0}^{N-1} \mathcal{G}\left(\frac{n}{NT}\right) e^{j\frac{2\pi nk}{N}} \quad k = 0, 1, \ldots, N-1 \,, \end{aligned} \tag{4.15}$$

where $g(kT)$ is obtained by sampling the continuous time function $g(t)$ $N$ times with sampling period $T$. The continuous time function $\gamma(t)$ is formed by substituting $t$, continuous time, for $kT$ in $g$. Then

$$\gamma(t) = \frac{1}{N} \sum_{n=0}^{N-1} \mathcal{G}\left(\frac{n}{NT}\right) e^{j\frac{2\pi nt}{NT}} \tag{4.16}$$

is a Fourier series reconstruction of $g(t)$. The indefinite integral of $\gamma(t)$

$$\begin{aligned} \Gamma(t) &= \int_0^t \gamma(\tau) d\tau \\ &= \int_0^t \left\{ \frac{1}{N} \sum_{n=0}^{N-1} \mathcal{G}\left(\frac{n}{NT}\right) e^{j\frac{2\pi n\tau}{NT}} \right\} d\tau \\ &= \frac{1}{N} \int_0^t \left\{ \sum_{n=0}^{N-1} \mathcal{G}\left(\frac{n}{NT}\right) \left[ \cos\frac{2\pi n\tau}{NT} + j \sin\frac{2\pi n\tau}{NT} \right] \right\} d\tau \,. \end{aligned} \tag{4.17}$$

In general the coefficients $\mathcal{G}(\frac{n}{NT})$ are complex, and, apart from the zero frequency coefficient $\mathcal{G}(0)$, are *conjugate even* around the middle (Nyquist) frequency, that is, $\mathcal{G}(\frac{N-n}{NT}) = \mathcal{G}^*(\frac{n}{NT})$, $n = 1, 2, \ldots, N/2$. Also, $\sin\frac{2\pi(N-n)t}{NT} = -\sin\frac{2\pi nt}{NT}$. Using these facts,

$$\begin{aligned} \Gamma(t) = &\frac{1}{N} \int_0^t \mathcal{G}\left(\frac{1}{2T}\right) \cos\frac{\pi\tau}{T} \, d\tau \\ &+ \frac{2}{N} \int_0^t \left\{ \sum_{n=0}^{N/2-1} \Re\left[\mathcal{G}\left(\frac{n}{NT}\right)\right] \cos\frac{2\pi n\tau}{NT} + \Im\left[\mathcal{G}\left(\frac{n}{NT}\right)\right] \sin\frac{2\pi n\tau}{NT} \right\} d\tau \,, \end{aligned} \tag{4.18}$$

where $\Re$ and $\Im$ signify real and imaginary parts.

Before carrying out the next step, presume that the zero frequency term $\mathcal{G}(0)$ can be dealt with separately. Further, we will drop all terms which produce constants in the integration—since the integration can only be performed to within a constant anyway—being content to later set the mean value of the integrated timeseries to zero. Then

$$\begin{aligned} \Gamma'(t) = &\frac{T}{N\pi} \mathcal{G}\left(\frac{1}{2T}\right) \sin\frac{\pi t}{T} \\ &+ \frac{2}{N} \sum_{n=1}^{N/2-1} \frac{NT}{2\pi n} \left\{ \Re\left[\mathcal{G}\left(\frac{n}{NT}\right)\right] \sin\frac{2\pi nt}{NT} - \Im\left[\mathcal{G}\left(\frac{n}{NT}\right)\right] \cos\frac{2\pi nt}{NT} \right\} \,. \end{aligned} \tag{4.19}$$

Interpreting $\frac{2\pi n}{NT}$ as the circular frequency $\omega$ for each of the DFT coefficients, and considering that $\sin(x) = -\sin(-x)$, the time domain integration of the Fourier series reconstruction $\gamma(t)$ can be seen to be equivalent to dividing the DFT coefficients by $j\omega$, as expected.

The zero frequency term remains to be dealt with. Since the time function $g(kT)$ is real, so is $\mathcal{G}(0)$. Physically, it is the sum $\sum_{k=0}^{N-1} g(kT)$, so $\mathcal{G}(0)/N$ is the mean of the $g(kT)$. Integrating this produces a line of constant slope $\mathcal{G}(0)/N$, which can be added on to $\Gamma'(t)$ to produce $\Gamma(t)$ within a constant:

$$\Gamma(t) = \frac{t}{N}\mathcal{G}(0) + \Gamma'(t) + C \; . \tag{4.20}$$

In the numerical implementation of the method, only the values of $\Gamma$ at discrete times $(\Gamma(kT); k = 0, 1, \ldots, N-1)$, are computed.

### Errors

The accuracy of the method depends on how well the original timeseries is matched by the Fourier series reconstruction 4.16. If the periodic repetition of the original continuous data record $g(t)$ has no frequency above the Nyquist frequency, then as a consequence of the Sampling Theorem, equation 4.16 provides an exact reconstruction (Stearns 1975), so the integration must be exact. In practice it is likely that some aliasing will be present despite low-pass filtering, because of the discontinuity produced where the periodic extensions of $g(t)$ join. This aliasing is unlikely to be important here, however, since the first harmonic of the accelerometer signal was typically less than 25% of Nyquist, so the amount of subsequent *leakage* due to periodic repetition would be small at the Nyquist frequency, and likely to be swamped by noise produced by sampling quantization in any case.

In the implementation of the method, the first term on the right of equation 4.19 was not computed. It represents the integral of the Nyquist frequency component of the DFT. The reason that it was ignored was that it did not easily fit into the data-packing scheme used in the computation. This is of little importance since all the data were low-pass filtered before return to the time domain, and the Nyquist frequency coefficient would be then set to zero in any case. The performance of the integration scheme was examined by using it to integrate samples of sums of sinusoids of different frequencies, for which the integrals could be computed analytically. The relative errors of the numerical method were always very small at all frequencies other than the Nyquist. Of course, there was always an error in the mean value too, but that was of no concern since the mean value was subsequently removed.

As a demonstration, figure 4.14 shows the analytical and computed integral of a timeseries composed of three sinusoids, and the difference between the two (error). It can be seen that the error seems to consist of a mean value with a possibility of small superimposed high frequency components (not resolvable on the scale of the plot). This is confirmed by examining the autospectrum of the error, presented in figure 4.15 (the mean value was removed before computing the spectrum). The error is concentrated at the Nyquist frequency, as expected. The resulting signal-to-noise ratio is -74.5 dB.

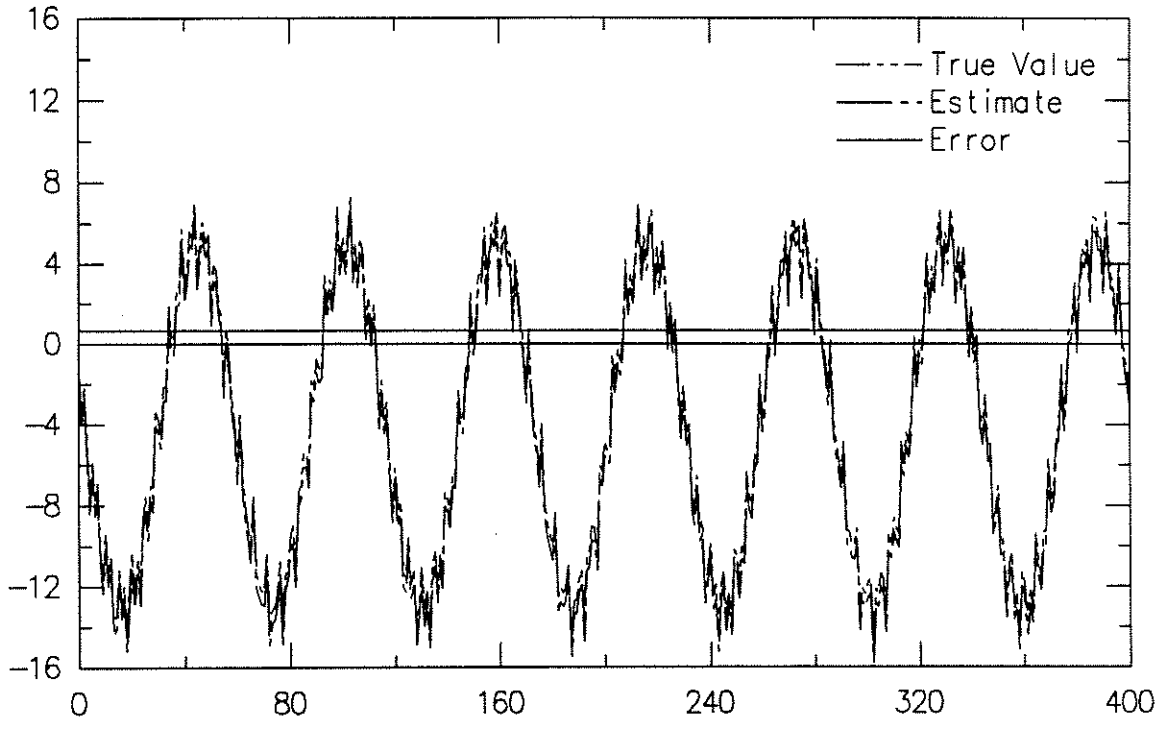


FIGURE 4.14: Analytical and computed integral of timeseries composed of three sinusoids, together with the error of the computed estimate.

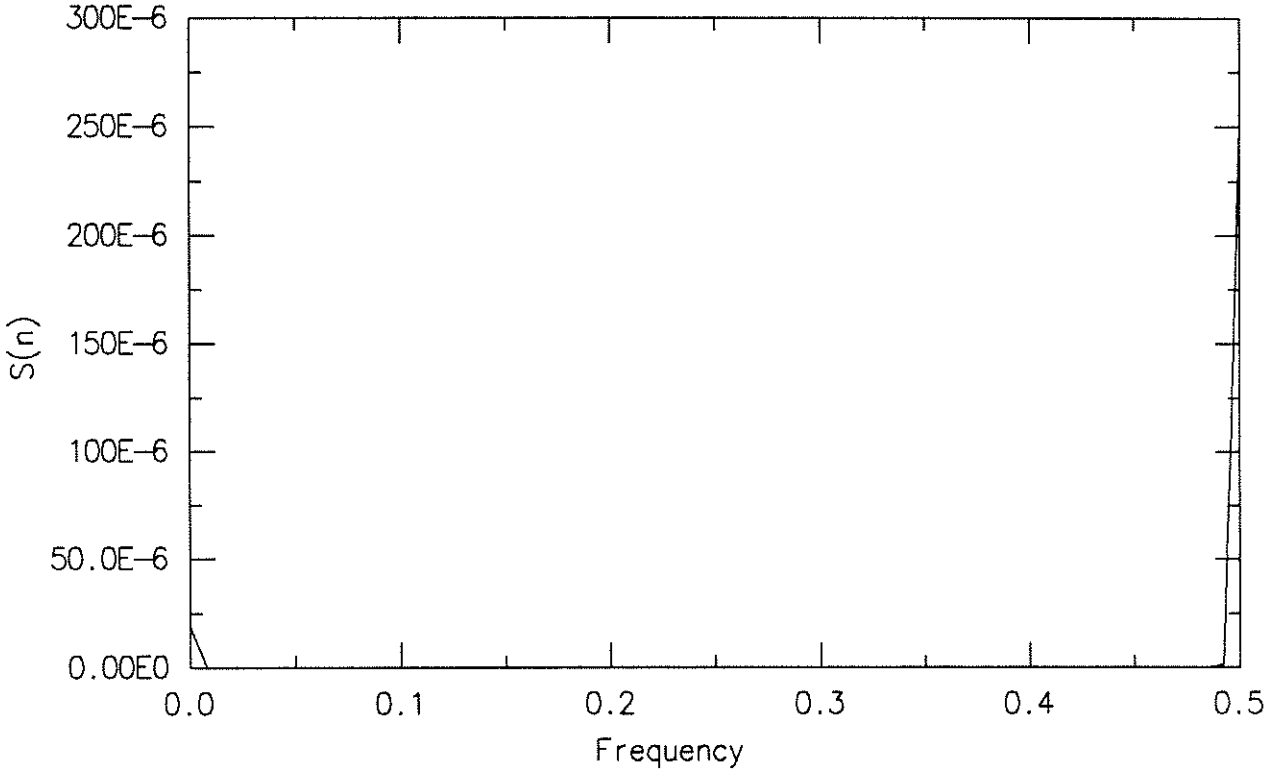


FIGURE 4.15: Autospectrum of error in fig. 4.14.

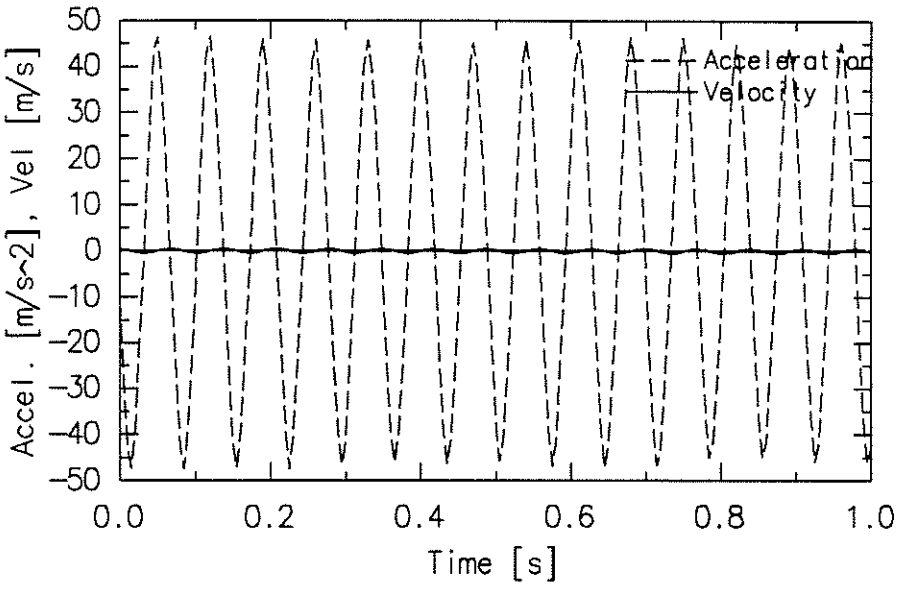


FIGURE 4.16: Timeseries of measured acceleration and computed velocity for an amplitude of 5.54 mm and a frequency of 14.24 Hz.

**Example**

An example of the use of this integration method is shown in figure 4.16, where the timeseries of acceleration and velocity for one transducer are presented. In this case, the cylinder was forced to oscillate at a frequency of 14.24 Hz and an amplitude of 5.54 mm, or 2.77% of diameter. It can be verified from the figure that the acceleration and velocity timeseries show the correct relative amplitude and phase.

### 4.2.5 Application

The *basic techniques* were applied to the central measurement tasks of the experimental programme, which were calibration of the force transducers and the processing of data recorded from the oscillating cylinder. Having dealt with the individual techniques we will now consider their implementation and use.

## 4.3 Dynamic Calibration

A dynamic method was devised for the calibration of the force transducers. While the adoption of this method was in some degree a response to problems with static calibration, as described in chapter 2, the method provided confirmation of the ability of the transducers to measure applied load when the model was oscillating. In this sense, the dynamic calibration also showed that the inertial cancellation technique worked correctly in practice. The equipment used in the dynamic calibration was described in chapter 2; the theory and signal processing are described here.

The basis of the method was to apply a fluctuating external force to the transducer segment while the cylinder was oscillating, subtract the inertial component from trans-

ducer strain bridge output, and then to compute a linear calibration factor which fitted the residual transducer signal to the applied force signal.

The fluctuating external force was applied by means of a lightweight load beam (spring), which was brought into contact with a transducer segment so that it applied a force in the cross flow direction. The beam was of cantilever type, and the end not in contact with the cylinder was held fixed. The load beam had a full bridge of strain gauges bonded to it and was itself calibrated by suspending known weights from its end. (Chapter 2 provides a more complete description of the equipment employed.) Fluctuating force was applied to the transducer by causing the cylinder to oscillate up and down, thus deflecting the beam and producing a time-varying reaction force.

The force sensed by the transducer strain gauges contained two components, one due to the mass and acceleration of the transducer outer segment, while the other was the force transmitted from the load beam. To obtain the signal from the load beam alone, the inertial component of the signal had to be calculated and subtracted out. The method described in section 4.2.1 was used to do this. Before connecting the load beam to the transducer, the cylinder was forced to oscillate and output from the transducer accelerometer and strain bridge recorded. This timeseries pair was used to produce an estimate of the FIR digital filter which related the two signals in the time domain.

With that done, the load beam was preloaded against the transducer segment and the cylinder forced to oscillate by providing input to the electromagnetic shaker. Signals were recorded from the transducer accelerometer and strain bridge amplifiers and from the load beam strain bridge amplifier. Processing of the data was done using the programme `dyncal`, presented in appendix C. The sequence of events during processing was as follows:

1. Data recorded from transducer accelerometer and strain bridge without load beam connected were read in and an FIR digital filter was fitted to relate the two timeseries (procedure `identify`).

2. Data recorded from accelerometer, transducer strain bridge and load beam strain bridge were read into storage and mean values were removed (procedures `readdata` and `demean`).

3. The FIR filter was convolved with acceleration timeseries to produce an estimate of the inertial component of force transducer response, then filter output was subtracted from transducer strain bridge timeseries (procedure `cancelforce`).

4. The force transducer strain bridge residual timeseries was transformed to the frequency domain using FFT. The effect of the antialiasing filter was removed from the result by dividing each DFT point by the complex filter magnitude (procedure `unfilter`). Then the DFT was low-pass filtered by multiplying by a filter magnitude which has the value unity at low frequencies, tapering to zero at high frequencies (procedure `bandpass`). Inverse FFT then returned the corrected force transducer strain bridge data to the time domain.

5. The load beam strain bridge timeseries was dealt with in a manner similar to that for the force transducer residual timeseries, except that it was also corrected in

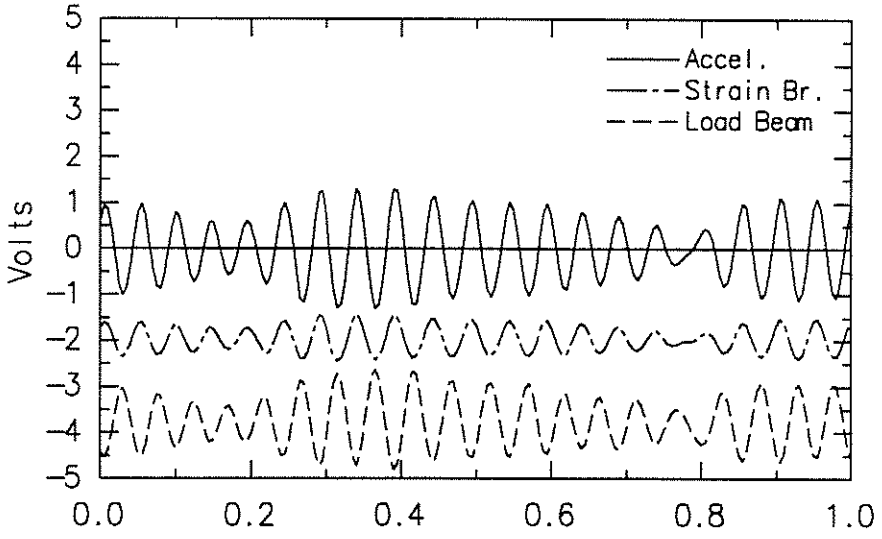


FIGURE 4.17: Timeseries from a transducer accelerometer and strain bridge, together with from the load beam, before processing.

the frequency domain for analogue-to-digital converter multiplexing time delay (procedure `delay`). This was done to bring it to zero delay relative to the force transducer strain bridge timeseries (section 4.2.2).

6. The method of least-squared error was used to fit a multiplicative constant which best related the force transducer strain bridge timeseries to the load beam strain bridge timeseries (procedure `fit`). The static calibration factor of the load beam (units Newtons/Volt) was multiplied by the fitted constant to give the calibration constant for the force transducer (written to supplementary file `dyncal.out`).

7. Timeseries of fitted and measured force were printed out for comparison, together with the difference between the two.

### 4.3.1 Example

Typical timeseries of accelerometer, transducer strain bridge and load beam strain bridge are presented in figure 4.17. This shows the narrow-band random input used during calibration. The fitted (transducer) and measured (load beam) timeseries are shown after correction in figure 4.18. The degree of agreement between the two timeseries was typical of all the calibrations obtained.

### 4.3.2 Errors

**Load beam transmissibility** One concern that can be raised about the calibration procedure is with the transmissibility of the load beam spring. If the frequency of oscillation of the cylinder assembly is near the natural frequency of the load beam, then the force transmitted to the cylinder transducer element will be much smaller than that indicated by the strain gauges on the load beam, since most of the stiffness forces in the

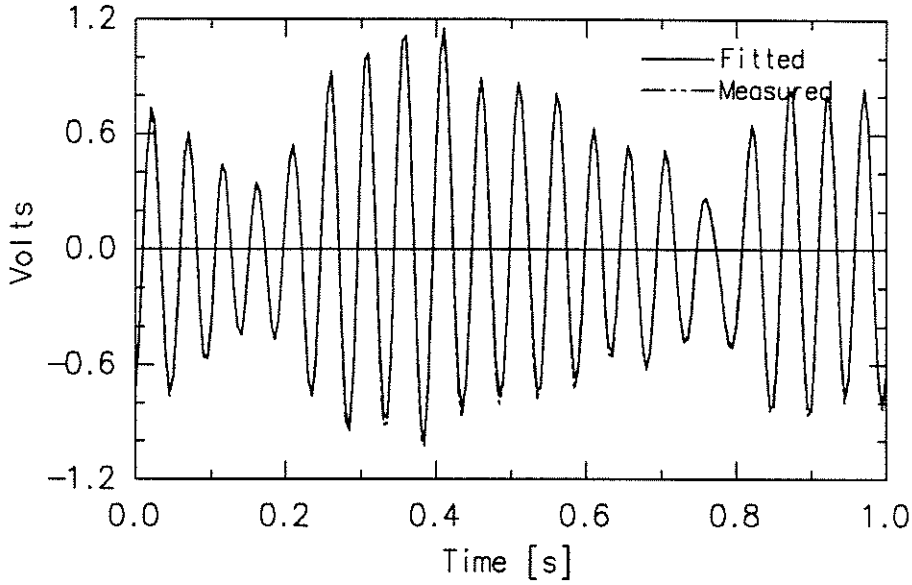


FIGURE 4.18: The fitted (force transducer) and measured (load beam) force timeseries, shown after processing.

load beam will be acting to accelerate the mass of the beam. To minimize the problem, the centre frequency of the cylinder oscillation was typically kept below 15 Hz during calibration, compared with the natural frequency of the load beam of 70 Hz, giving a frequency ratio $\omega/\omega_n = 0.21$. The magnitude of the force transmissibility is (Thomson 1972 p. 57)

$$\mathrm{TR} = \sqrt{\frac{1 + \left(2\zeta\frac{\omega}{\omega_n}\right)^2}{\left[1 - \left(\frac{\omega}{\omega_n}\right)^2\right]^2 + \left[2\zeta\frac{\omega}{\omega_n}\right]^2}} \,. \tag{4.21}$$

The worst case for TR when $\omega/\omega_n < 1$ is if $\zeta = 0$; then, with $\omega/\omega_n = 0.21$, TR $= 1.046$, implying a 4.6% error in transducer calibration. The effect would be that the forces estimated by the transducer would be too large, by a maximum of 4.6%. The phase error produced by transmissibility effects,

$$\phi = \arctan \frac{2\zeta\left(\frac{\omega}{\omega_n}\right)^3}{1 - \left(\frac{\omega}{\omega_n}\right)^2 + \left(2\zeta\frac{\omega}{\omega_n}\right)^2} \,, \tag{4.22}$$

approaches zero for $\omega/\omega_n < 1$ if the damping ratio, $\zeta$, is small (the damping ratio $\zeta$, measured for the isolated load beam, was 0.4%, although the damping provided by the contact with the cylinder could not be assessed).

**Transducer transmissibility** Another concern is with the transmissibility of the transducer itself. It was found that the resonant frequency of the transducer outer segment on the internal beam springs was 400 Hz. During calibration, the frequency ratio was $15/400 = 0.0375$, giving TR $= 1.001$, which is an insignificant error. During

measurements of aerodynamic force, the worst case was provided in high Reynolds number flows, when the maximum frequency of interest was 100 Hz, giving TR = 1.067. This frequency corresponds to a Reduced Frequency or Strouhal number $fD/U \approx 100 \times 0.2/50 = 0.4$ at the highest Reynolds numbers, with an air speed of 50 m/s. For a more typical value of $fD/U = 0.2$, TR = 1.016 at this air speed.

**Summary** The transmissibility of the load beam could produce force transducer calibration errors as high as 4.6%, in such a sense that the cross flow forces estimated using the transducers would be too large.

The errors due to transducer transmissibility are in general less significant, and errors produced by load beam and transducer transmissibilities tend to cancel each other out.

## 4.4 Processing of Cylinder Data

### 4.4.1 Programme adcal

#### Task

The core of the processing of cylinder data was performed by the programme `adcal`, the source code for which is provided in appendix C. The function of `adcal` was to convert timeseries recorded from cylinder transducer accelerometer and strain bridge amplifiers (in Volts) to timeseries of transducer acceleration, velocity, and aerodynamic force, in appropriate SI units. `Adcal` was compiled from a number of modules; the *program-block* was contained in the file `main.p` (app. C), where the operation of the programme is described in detail.

#### Procedure

As described in section 4.2.1, removal of the inertial component of transducers signal required that the cylinder be forced to oscillate at the correct frequency and amplitude for each test in the absence of flow. A short length (typically 1024 samples) of data was recorded in this condition before the flow was applied; this section of data was used to identify the FIR digital filter which related the inertial component of transducer response to accelerometer output. Then the fan pitch was advanced (with the cylinder still oscillating), and a set of data (8192 samples × 12 channels) recorded for later processing by `adcal`.

The programme was designed to operate on data recorded from the oscillating cylinder, but was just as appropriate for the processing of data from the fixed cylinder, so it was used to process all the wind tunnel data described in this thesis. For the processing of fixed cylinder data, input for the cancellation stage was provided by driving the clamped yoke assembly using band-limited noise input to the shaker.

### Precis

To carry out the processing, `adcal` used all the techniques described in section 4.2, carrying out operations in the time or frequency domains as appropriate. Discussion of the individual techniques has been presented, and the source code of appendix C is intended to be self-documenting, so only a summary of execution in the *program-block*, given in `main.p`, will be provided here.

Except for input and output stages, the programme operated on data from one transducer at a time. The sequence of data processing for one transducer was as follows:

1. Using timeseries of transducer accelerometer and strain bridge signal outputs recorded in the absence of flow, an FIR digital filter was constructed to relate on to the other (procedure `identify`).

2. Accelerometer and strain bridge timeseries recorded with flow applied were read into store (procedure `readdata`).

3. The FIR digital filter was convolved with accelerometer timeseries and result subtracted from force timeseries, leaving the component caused by aerodynamic force (procedure `cancelforce`). Mean values were removed.

4. Accelerometer and force data were transformed to the frequency domain using FFT (procedure `realFFT(...true)`).

5. DFT points were phase-shifted to introduce the time delay that brought the data to zero delay relative to first recorded channel of data (procedure `delay`).

6. The DFT was adjusted to account for the particular set of anti-aliasing filters for the two data channels, by dividing each DFT point by the complex FRF magnitude for the filter at that frequency (procedure `deconvolve`).

7. Using the FRF computed between the output of the transducer accelerometer and a reference accelerometer, the accelerometer DFT was adjusted to account for imperfect frequency response (procedure `deconvolve`).

8. Bandpass filter of both data channels. The first 1% of spectral lines were set to zero, and the remainder of the DFT was low-pass filtered, starting at (typically) 82.5% of Nyquist frequency. The low-pass filtering was carried out to deal with any high frequency noise that may have been amplified during deconvolution of anti-aliasing filters (filter magnitude was near zero at Nyquist frequency) (procedure `bandpass`).

9. The frequency domain component of integration was carried out on acceleration data—the first step in computation of cylinder velocity (procedure `freqinteg`).

10. Acceleration, velocity, and force data were transformed back to the time domain (`realFFT(...false)`).

11. The time domain steps in integration of acceleration were carried out (procedures `finishinteg` and `removemean`).

12. Unit conversion of acceleration, velocity, and force data was carried out, using calibration factors to convert to SI units (procedure `calibrate`).

### Verification

It is impossible to provide direct verification of the operation of the programme, which requires another, independent, 'correct' set of results for the same experimental set-up. This state of affairs is of course the norm for any new experimental technique used in a new set-up.

Before further discussion, note that the various causes of error which could be present in measurements of lift force made with the equipment can be categorized.

*Variations in boundary conditions* may produce significant differences between the results of one experimental set-up and another. Strictly speaking, these do not produce measurement errors, but may limit the comparison of results from experiment to experiment.

*Bad transducer design* can result in transducers which distort the input signal or are unexpectedly sensitive to inputs other than those under study.

*Calibration error* will produce systematic errors in the results.

*Low signal to noise ratio* will produce results which have a comparatively large amount of signal that is not coherent with the input to the transducer.

*Poor signal processing techniques,* due either to bad choice of algorithm or errors in coding.

It is this last class, concerned with signal processing, which is under discussion here. As pointed out above, it is not possible to directly gauge the size of these errors, but indirect checks have been made. These included: checking results against previously published results, examining the individual techniques used in signal processing and results before and after the use of the various procedures discussed above, and checking output by hand calculations. These checks all indicated that the signal processing worked as intended.

The greatest concern must lie with the inertial force cancellation technique, considering that the inertial component of transducer signal is such a large proportion of total signal when the cylinder is oscillating at maximum amplitude. If this cancellation is unsatisfactory, residual forces will exist at the frequency of motion which will be falsely assigned an aerodynamic origin. It is the possiblity of error here which is most difficult to directly assess, and at the same time most important for the experimental programme as a whole. Two kinds of experimental results address this concern. The first is the demonstration (section 4.2.1) that in the absence of flow, the inertial cancellation technique performed satisfactorily, giving an RMS error signal similar to the analogue-to-digital quantization level (5 mV). The second is the comparison of the behaviour of the forces at the frequency of oscillation in smooth flow with published results (to be discussed in chapter 6). It was found that these forces showed the previously observed rapid change in phase with respect to cylinder motion as the cylinder frequency traversed the Strouhal frequency. The facts that the forces at the cylinder oscillation frequency (a) have similar magnitude from transducer to transducer, (b) show

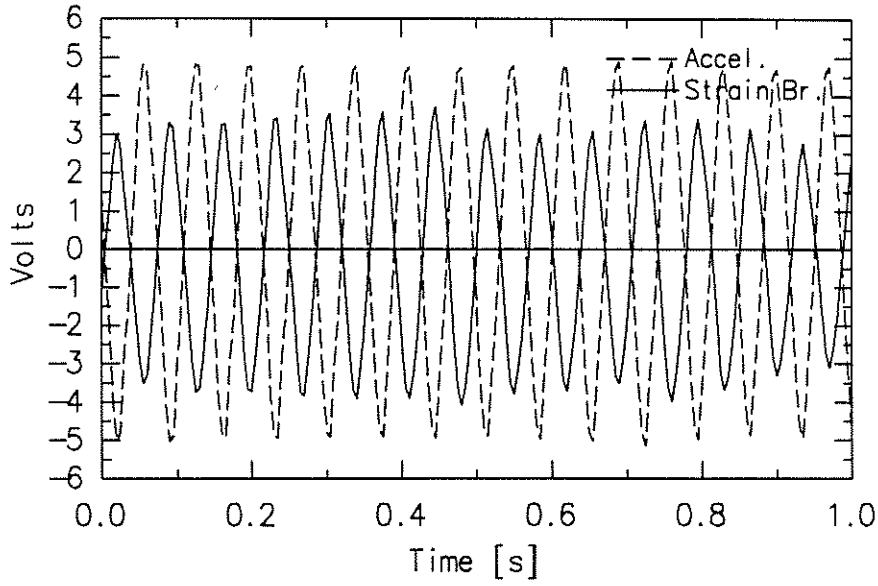


FIGURE 4.19: Data recorded from one transducer in low Reynolds number, low turbulence flow, prior to processing.

the same variation in phase, (c) are repeatable, (d) match previous results in form, show that the inertial cancellation technique also worked when flow was applied.

Of the other causes of error, *boundary condition variations* will be discussed in later chapters although the topic has already received attention in chapter 1. *Transducer design* has been addressed directly in chapter 2 and during commissioning of the equipment. *Calibration errors* were addressed in section 4.3. *Signal to noise ratio* was maximized by adjusting the magnitude of the recorded signals so that they were of the same order as the available range of the analogue-to-digital converter ($\pm 10$ V), as shown in examples presented in this chapter.

### Examples

Two examples of the operation of `adcal` are presented here. They come from the ends of the ranges of flow regime: one set of examples is from low Reynolds number, low turbulence flow; the other from high Reynolds number, high turbulence flow.

**Low Reynolds number** The first set of results were recorded at a Reynolds number of $1.6 \times 10^5$. The turbulence intensity was 0.6%, and the amplitude of cylinder oscillation was 2.77% of diameter (average from all transducers), while the frequency was 14.24 Hz. The working section air speed was 12.5 m/s.

Unprocessed data from one transducer are shown in figure 4.19. The acceleration signal was nearly periodic, showing that the cross flow forces were unable to significantly influence the forced oscillation of the cylinder. The force transducer signal contains a large component at the cylinder oscillation frequency, out of phase with acceleration. This dominance was due to the inertial component of transducer response, which was larger than the aerodynamic component.

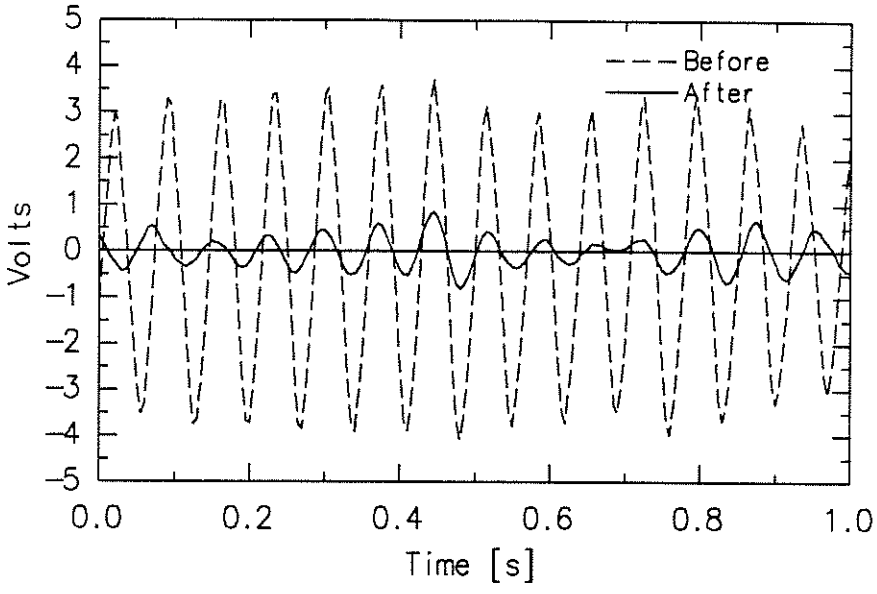


FIGURE 4.20: Low Reynolds number, low turbulence flow. Force transducer signal before and after processing (no unit conversion applied).

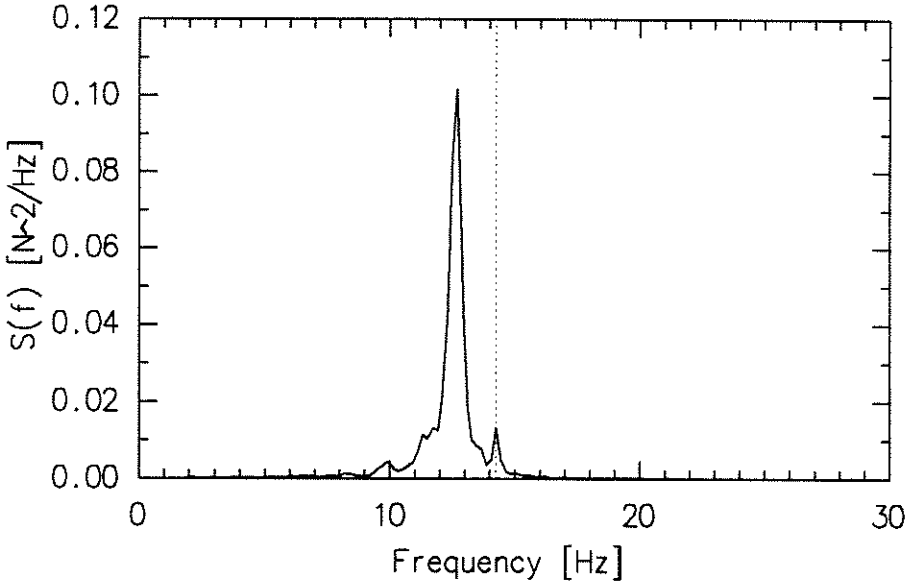


FIGURE 4.21: Low Reynolds number, low turbulence flow. Spectrum of processed force transducer output. Dotted line indicates cylinder oscillation frequency.

Figure 4.20 shows, for comparison, the force transducer signal before and after processing. No unit conversion has been performed on the processed signal (an option of `adcal`), so comparison of magnitudes in Volts shows directly the influence of the inertial component of response. Note the obvious presence of narrow-band vortex shedding.

A spectrum of the processed force transducer output is shown in figure 4.21; immediately apparent is the presence of two peaks in the spectrum—one at the cylinder oscillation frequency (14.24 Hz) and the other at the vortex shedding frequency (12.7 Hz, giving a Strouhal number $f_v D/U = 12.7 \cdot 0.2/12.5 = 0.203$).

The start-up transient is shown in the processed force signal of figure 4.22. The

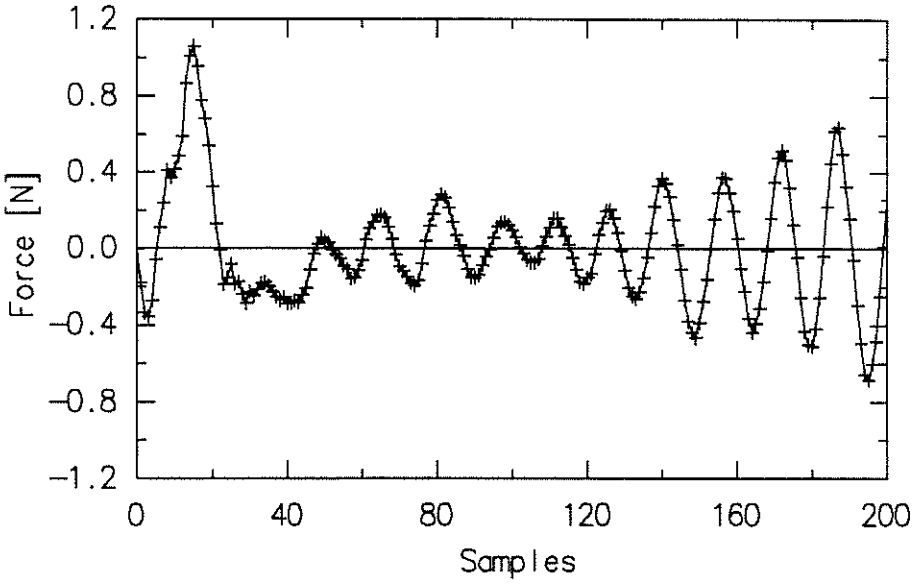


FIGURE 4.22: Low Reynolds number, low turbulence flow. Start-up transient effect at beginning of processed force transducer output.

transient was due to the start-up of the 32-weight force cancellation filter. The abscissa scale has been given as number of samples—the total length of time was the same as in figure 4.20; the sampling period was 5 ms. Two points can be made here. The first is that the total length of recording of all the data sets is 8192 points, or over 20 times the portions shown in figures 4.20 and 4.22—the shorter length is shown to emphasize detail. The second is that the start-up transient section was omitted from computations of average data properties (such as the spectrum of figure 4.21), in all the results presented in this dissertation.

Traces of acceleration (corrected and calibrated) from all six transducers are shown in figure 4.23. The mean peak level of 46 $m^2/s$ corresponds to an amplitude of oscillation of 5.75 mm (2.87% of diameter) at a frequency of 14.24 Hz. Velocities are presented in figure 4.24, again, the peak level 0.49 m/s is appropriate, and the phase is retarded by $\pi/2$ compared to the acceleration traces, as expected. In fact, the peak values indicate that the cylinder was not oscillating with a perfectly sinusoidal waveform, since they both indicate slightly different amplitudes. The amplitude of 5.54 mm quoted previously was estimated on the basis of the RMS velocity signals averaged over the six transducers. Note that it is difficult to differentiate between the six sets of acceleration and velocity signals at the scale at which figures 4.23 and 4.24 have been reproduced.

Aerodynamic forces from all six transducers are shown in figure 4.25. It can be seen that vortex shedding was well correlated along the span of the cylinder and also that the magnitude of forces at all transducers is similar.

**High Reynolds number** An analogous set of results is presented here for the high Reynolds number ($5.5 \times 10^5$), high turbulence intensity (18%) configuration. In this case the sampling period was reduced to 2.5 ms (giving a Nyquist frequency of 200 Hz); the sensitivity of the strain bridge and accelerometer amplifiers was reduced as well, to keep the signals within the $\pm 10$ V range of the analogue-to-digital converter.

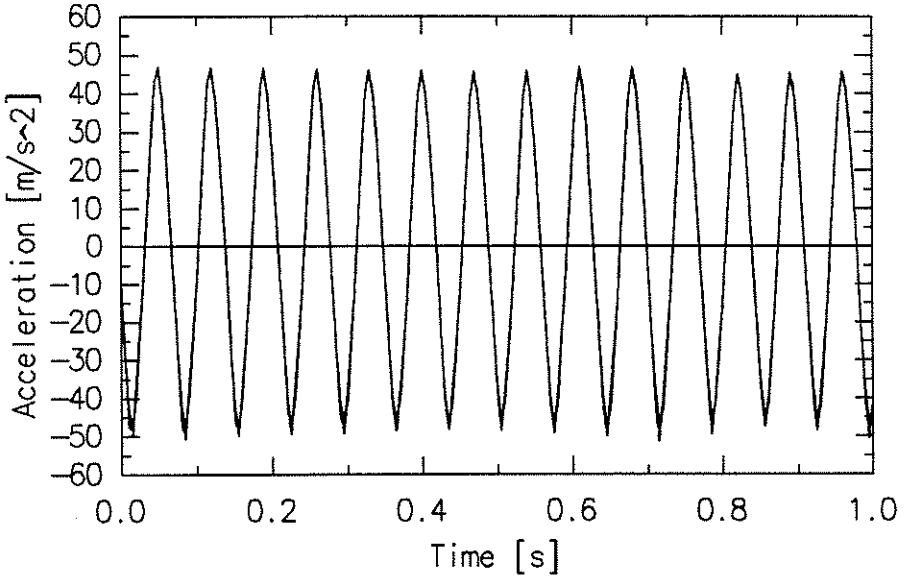


FIGURE 4.23: Low Reynolds number, low turbulence flow. Acceleration traces from all six transducers.

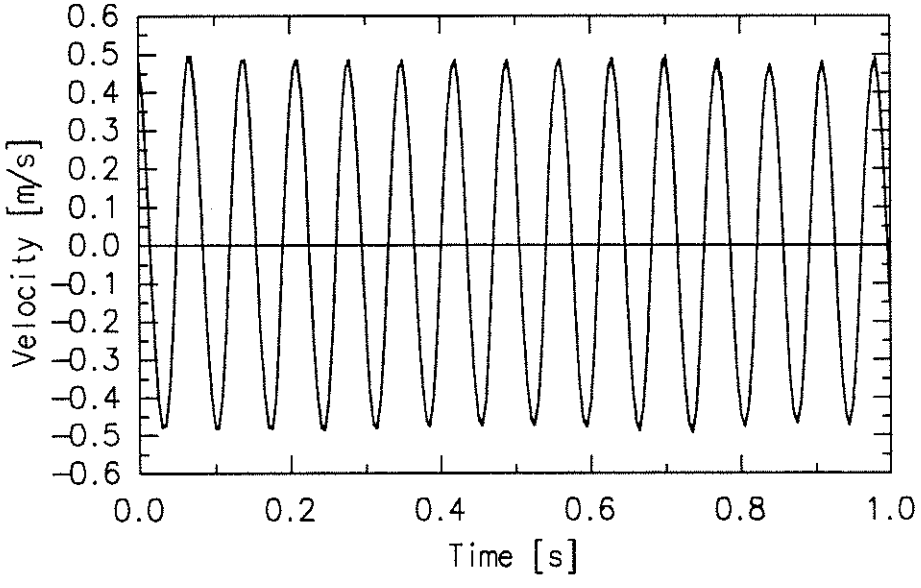


FIGURE 4.24: Low Reynolds number, low turbulence flow. Velocity traces from all six transducers.

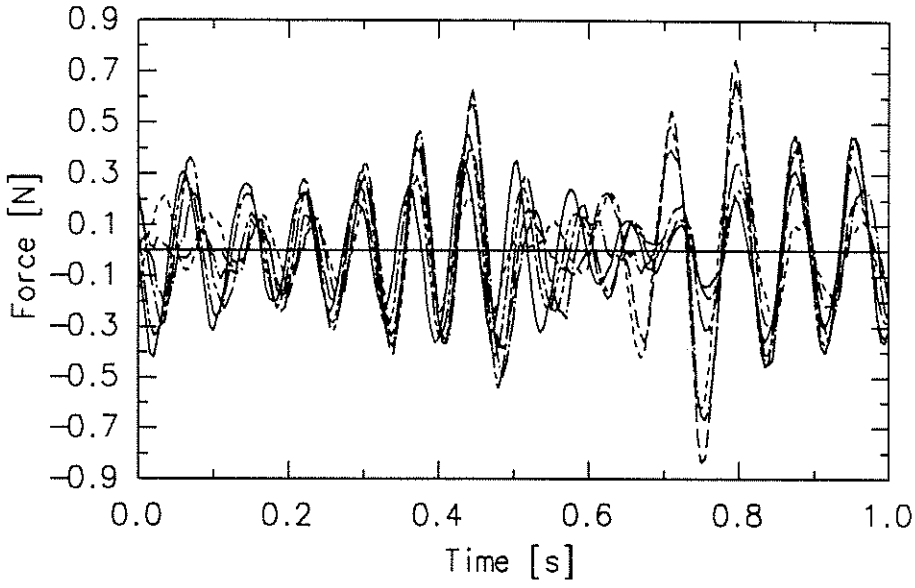


FIGURE 4.25: Low Reynolds number, low turbulence flow. Timeseries of aerodynamic force from all six transducers.

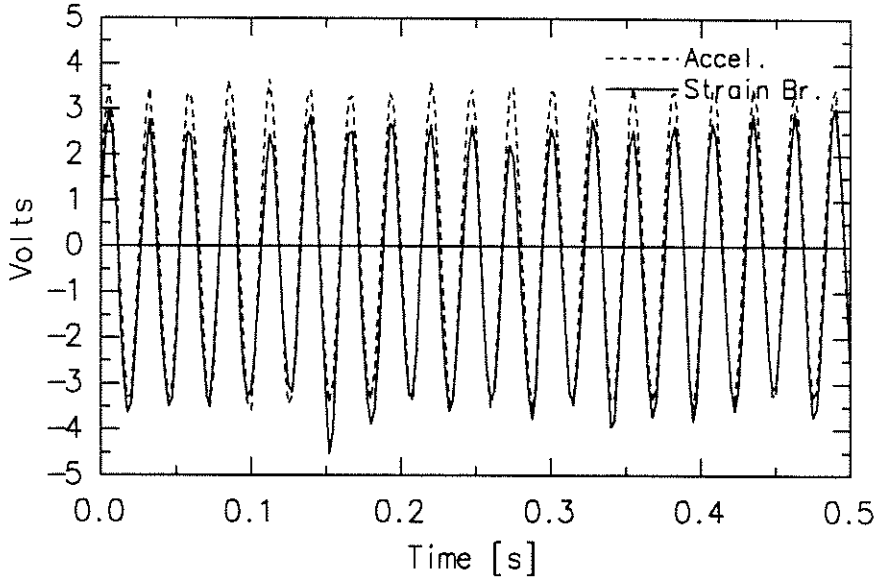


FIGURE 4.26: Data recorded from one transducer in high Reynolds number, high turbulence flow, prior to processing.

The unprocessed accelerometer and strain bridge signals from one transducer are shown in figure 4.26; again some influence of cross flow force on the motion can be seen, but the motion was still almost periodic, indicating that the cylinder was in forced oscillation.

Output from one force transducer, before and after processing, is shown in figure 4.27. Unit conversion has not been carried out, so that the relative magnitudes of the total and aerodynamic components of force are shown, as before. The residual (aerodynamic) force has a quasi-periodic nature, reflected in the spectrum of figure 4.28. As in figure 4.21, a significant peak can be seen in the spectrum at the cylinder oscil-

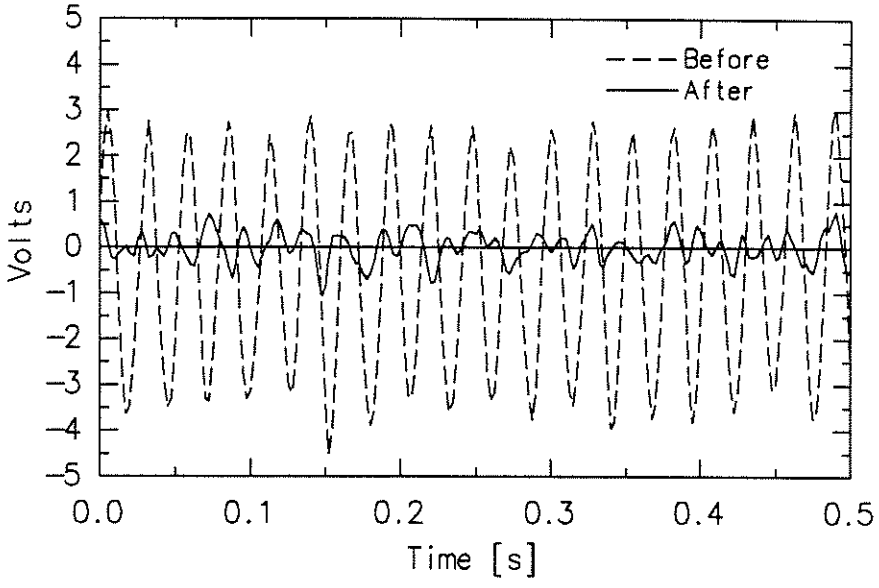


FIGURE 4.27: High Reynolds number, high turbulence flow. Force transducer signal before and after processing (no unit conversion applied).

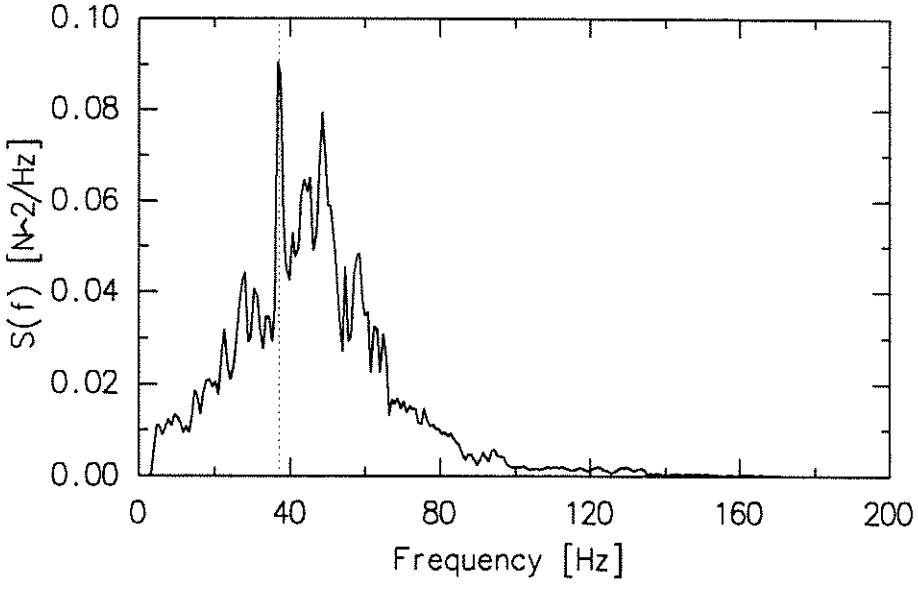


FIGURE 4.28: High Reynolds number, high turbulence flow. Spectrum of processed force transducer output.

lation frequency, 37.1 Hz. The remainder of the spectrum shows a peak at a frequency of 48 Hz, giving a Strouhal number of $48 \cdot 0.2/37.5 = 0.243$. The relative bandwidth is much greater than in the low Reynolds number, low turbulence flow.

Time traces of force from all six transducers, in figure 4.29, shows that the degree of spanwise coherence of the vortex shedding was much lower in this flow than in the low Reynolds number, low turbulence flow (fig. 4.25). Occasional well-correlated events can be seen.

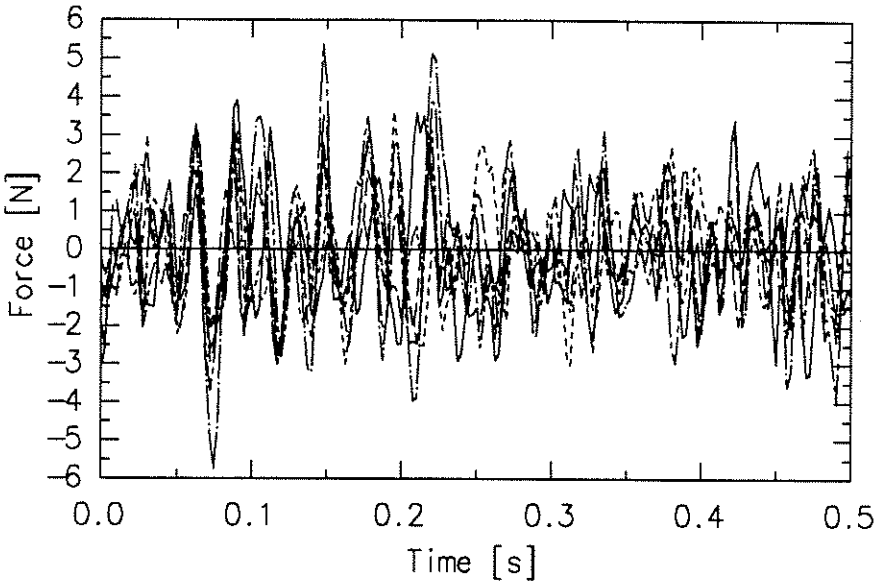


FIGURE 4.29: High Reynolds number, high turbulence flow. Timeseries of aerodynamic force from all six force transducers.

### 4.4.2 Post processing

After processing by `adcal`, and before further work, two modifications were made to the force transducer outputs.

#### Return of Added Mass force

The process of inertial signal cancellation also removed the influence of added mass forces in the absence of flow, as pointed out in section 4.2.1. Since the magnitude of the force coefficient for acceleration-related forces in the absence of flow is known, these forces were computed and added back into the output of the force transducers. The force was proportional to the mass of fluid displaced by the transducer and its acceleration

$$F_{a_0} = -C_{a_0} \rho V \ddot{y} \ . \tag{4.23}$$

Here, $\rho$ is the density of the fluid, $V$ and $\ddot{y}$ are the volume and cross flow acceleration of the transducer, and $C_{a_0}$ is the added mass coefficient at zero flow speed. As discussed in section 4.2.1, $C_{a_0} = 1$ for a circular cylinder. $F_{a_0}$ is the added mass force applied to the transducer; the sign change of equation 4.23 accounts for the fact that the force applied to the transducer was equal and opposite that applied to the fluid.

For the data in chapter 6, forces correlated with cylinder acceleration and velocity were computed and used to calculate coefficients of motion-dependent forces. The forces correlated with acceleration were reduced to a coefficient of added mass. The process just described effectively added an offset of unity to the added mass coefficient $C_a$.

The relative size of these forces is shown in figure 4.30, which is the processed output of the force trace shown in figure 4.20 Also, these forces were converted to dimensionless form, as discussed below.

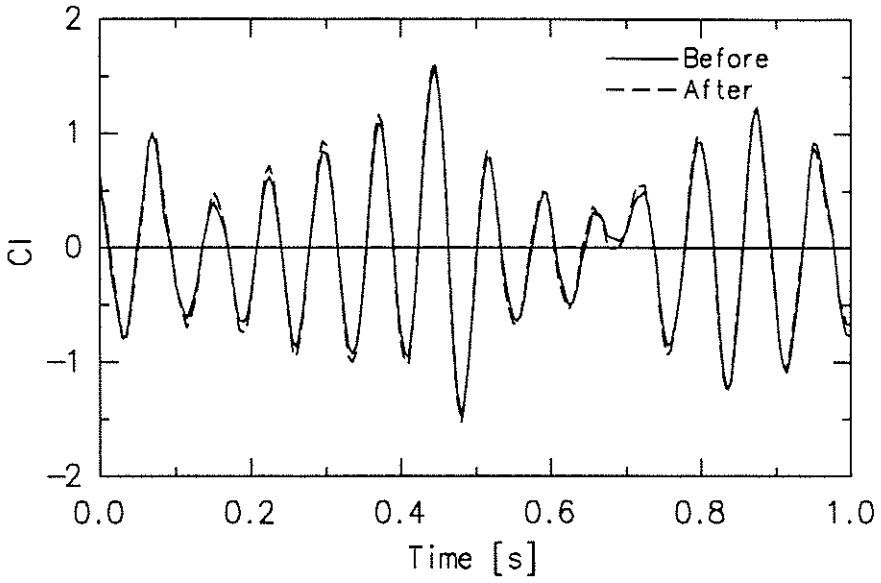


FIGURE 4.30: Example showing the relative size of forces due to the re-introduction of no-flow added mass.

### Conversion to dimensionless form

The processed cross flow forces were converted to dimensionless form:

$$C_L = \frac{l}{\frac{1}{2}\rho U_{ref}^2 K_u^2 D\Delta L}\,, \tag{4.24}$$

where $\frac{1}{2}\rho U_{ref}^2$ was the dynamic pressure provided by the reference Pitot-static tube, $K_u$ was a correction factor between velocity indicated by the reference Pitot-static tube and mean velocity found at the cylinder centreline position using a hotwire probe, and $D\Delta L$ was the projected area of the transducer. $l$ was cross flow force, giving $C_l$ as the coefficient of lift, or cross flow force, at the transducer. All forces presented subsequently are in dimensionless form.

## 4.5 Closure

The techniques used in the initial processing of data from the wind tunnel model have been described. Examples of the use of the individual techniques and of their integrated application for data processing have been presented, and the errors produced by the individual and combined techniques have been discussed, although from the experimental, rather than numerical point of view.

Some further manipulation of data will be described in the following chapters, but that will build on the methods described here, in that all subsequent processing and discussion will assume that the processing described in this chapter has previously been carried out on the data. Put another way, later work assumes that the data are available in the form of acceleration, velocity and dimensionless cross flow force (corrected for no-flow added mass), for each transducer.