Chapter 4

The Streaming Current Meter

4.1 Introduction

4.1.1 Motivation

A streaming current (SC) meter is an instrument for measuring the charge that exists on small, suspended particles in liquid. A streaming current meter (SCM) is the only online instrument that can be used to measure coagulated particle stability for the feedback control of coagulant dosage. Therefore any attempt to improve the online control of coagulant dosing should begin with the SCM.

SCMs generally have a relatively slow and noisy response to coagulant dosage change. This chapter's main contributions lie in its approaches to improving and explain this.

The benefits of having a faster response to changes in SC are clear when an optimiser is being used, as set point changes can be tracked faster and more accurately. This will allow more jar tests to be carried out in the same amount of time (for example using multiple jars) and result in less variation between jar test results. When no optimiser is being used, a faster SC response still results in better disturbance rejection, as well as more robust controller stability in the presence of process gain changes.

4.1.2 Chapter Overview

Sections 4.2-4.5 cover the theory behind SC, its relationship to the more fundamental zeta-potential, the ways of measuring it, its use in drinking water treatment and summarises the limited theoretical work done on the measurement. This treatment of the background to SC measurement that is given here is the most complete and accessible of which the author is aware.

In section 4.6-4.8, this thesis proposes the use of signal processing technique based on the Fourier transform as a way to improve the SCMs speed of response and to improve the applicability of the result. Experimental results demonstrating its effectiveness are also presented. This algorithm was integrated into the signal processing electronics of the Accurate Measurement Accufloc SCM.

In section 4.9, a simple theoretical model for the dynamic response of the instrument is derived. No previous dynamic models of the SCM have been published. Although very limited, this model provides some insight into the model structure required as part of a model of the larger dosing system.

4.1.3 Nomenclature

The streaming current meter (SCM) has also been referred to as the streaming current detector (SCD) or streaming potential detector (SPD) as well as streaming current monitor and streaming current analyser (SCA) in various works.

The word ‘detector’ is used in earlier literature instead of the words ‘meter’, ‘monitor’ or ‘analyser’ because earlier SCMs were not considered to provide a useful quantitative measurement, only a qualitative indication of the presence and sign of charge. The words 1. ‘Meter’, 2. ‘Monitor’ and 3. ‘Analyser’ all mean the same thing in this context, except that the implication of sophistication increases in this order.

The difference between streaming potential and the more useful streaming current will be explained in the following section.

This thesis uses the more modern and descriptive term streaming current meter for all of these instruments.

4.2 Zeta-Potential and Coagulation

4.2.1 Colloidal Particles

Turbidity is caused by suspended particles in water in the size range of approximately 0.01 to 100 um in size. The larger fraction can easily be removed by settling. The smaller particles, with sizes of less than 5 um are referred to as colloidal particles (or colloids) and have extremely slow settling velocities and so cannot be practically removed by settling.

The behaviour of colloidal particles in water is strongly influenced by their electrostatic charge. This colloidal charge comes about because of the uneven surface characteristics of the particles and in most solids is negative, particularly the alumino-silicate clays typically suspended in surface water. This means that the negative charges on each particle will repel others and prevent effective agglomeration and flocculation. Neutralising this charge is the main purpose of coagulation.

However simply considering this charge as an electrostatic surface charge is an over simplification. As the whole solution must have a neutral charge, a layer of water containing ions of opposite charge surrounds each colloid. Consideration of colloidal charge purely in terms of the colloidal particles by themselves has no benefit, as the surrounding ions are, for all practical purposes, inseparable from them.

4.2.2 The Double Layer

The double layer model is used to explain the distribution of ions around each colloidal particle. This is a long standing conventional approach to colloid analysis, a more detailed description is available in such works as [Zeta-Meter 1993] and [].

For illustrative purposes we will assume a negatively charge colloid. Closest to the negative surface of the colloid there is a layer of strongly bound positive ions – this is know as the Stern layer. Further positive ions are still attracted to the colloid but are repelled by the Stern layer; likewise, negative ions are attracted to the positive ions but repelled by the colloid. A dynamic equilibrium of negative and positive ions forms

outside the Stern layer, known as the diffuse layer. The concentration of positive ions in the diffuse layer gradually decreases as the distance from the colloid increases until beyond a certain distance the ion concentrations are the same as the equilibrium in the water. This is illustrated in figure 4.1.

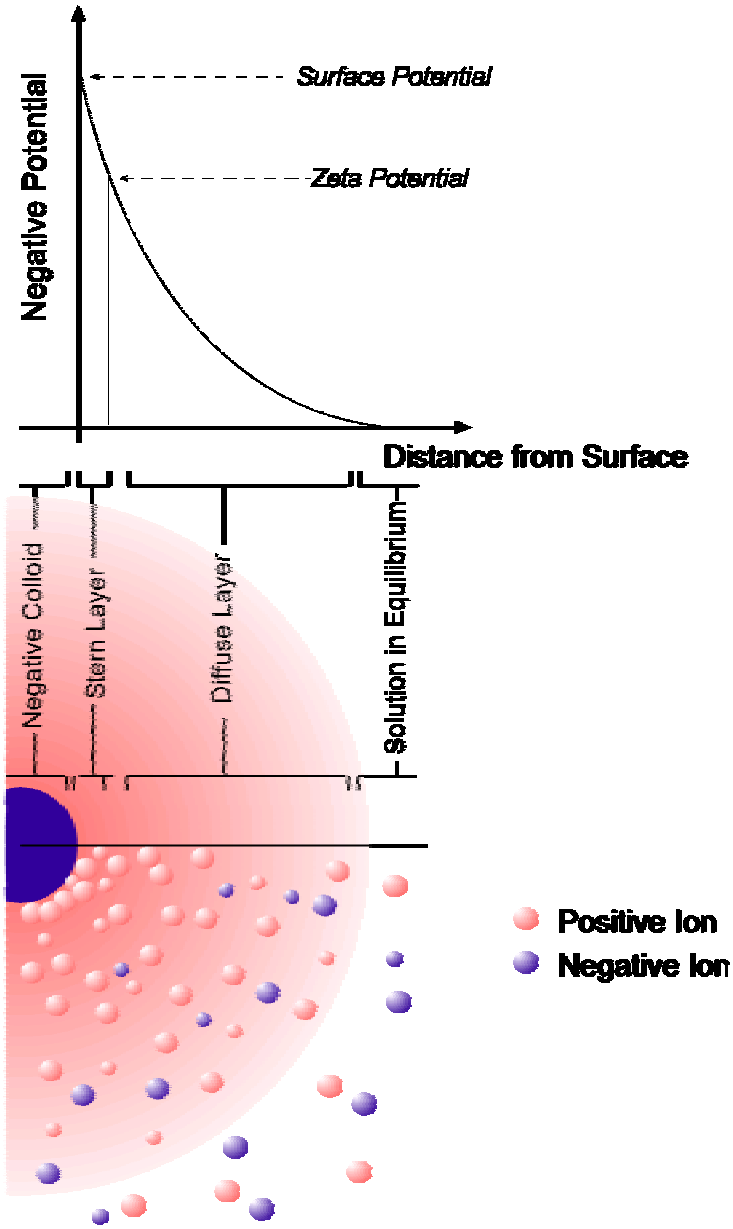

Figure 4.1 Ways to visualise the Double Layer model and the potential as a function of distance from the colloid's charged surface. The graph is actually shown upside down for a typical colloid with negative surface potential.

The strongly held positive ions near the surface and the charged layer surrounding this is where the name double layer comes from. The point just outside the stern layer is often called the shear plane as motion of the particle through the water shears the diffuse layer away.

The thickness of these layers depends on the concentration of ions in solution. At any distance from the surface the charge density is equal to the difference between the concentration of positive and negative ions at that distance. This results in an electrical potential (which is a voltage, caused by separation of charges) to exist across the layers. This potential is greatest near the surface and decreases to zero as the distance from the colloid increases. A graph of this potential curve is useful because it indicates the distance at which the interaction between colloids will occur.

The potential at the boundary between the Stern layer and diffuse layer is called the zeta potential, as shown in figure 4.2. Zeta-potential is useful, as it is a direct indication of the amount of energy required to bring separate particles together.

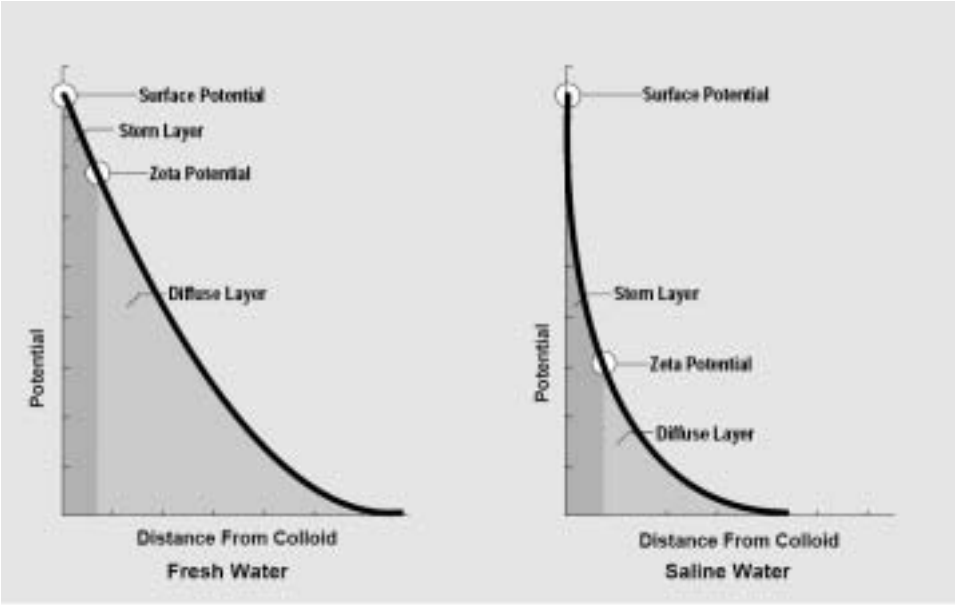

Figure 4.2 Potential vs Distance from Colloid. The higher concentration of ions in saline water causes compression of the diffuse layer. Source [Zeta-Meter 1993]

In fresh water suitable for treatment the diffuse layer is large compared to the stern layer and so zeta potential gives a good indication of surface potential, which is directly related to surface charge.

The zeta potential of a colloid's surface can be measured quite easily. A microscope is used to observe turbidity particles inside a thin chamber called an electrophoresis cell. An electric field is applied along the cell, if the particles are negatively charged they will move towards the positive end of the cell, carrying their stern layer ions with them. The average speed with which these particles move depends on the zeta-potential. It is this that determines the net electrical force operating on the particle and its stern layer ions.

In practice the relationship between zeta-potential and this average speed is not always straightforward to calculate, as it is dependant on the fluid's viscosity, dielectric constant, conductivity and temperature. As a result, zeta potential, which should have units of mV, is often expressed in terms of electrophoretic mobility. This has rather confusing units in velocity per electric field strength, typically um/s per V/cm.

Other disadvantages of measuring zeta-potential in this way include:

- It can only be determined for particles that are large enough to be detected and tracked through a microscope.

- It requires a lab technician to operate the apparatus and to make the observe the particles through a microscope. Zeta potential measurement therefore cannot be used online.

- It is not very accurate near zero, which unfortunately is the point of greatest interest because of the difficulty in tracking the particle motion.

Zeta-potential measurement has applications in laboratory studies of coagulation and flocculation processes.

4.3 The Streaming Current Sensor

4.3.1 The Streaming Current Effect

The streaming current, or streaming potential, effect is a reverse manifestation of the electrophoresis effect. This is apparent when a particle is mechanically moved through the fluid, or the fluid moved past the particle resulting in a separation of charges, causing a potential to exist.

A common situation that frequently occurs is that water with negatively charged particles is forced through a filter. The negatively charged particles become lodged in the filter, while the mobile positively charged ions are swept downstream. This separation of charges causes an electrical potential to exist across the filter, called the streaming potential. If electrodes are inserted upstream and downstream of the filter then the electrical potential can be measured. The upstream electrode is negative and the downstream electrode is positive. (Electrical potential is also referred to as voltage, as it is measured in unit called Volts.)

The electrical potential results in a current being conducted upstream through the water to remove the charge separation, this is the streaming current. The electrical potential therefore depends on the conductivity of the liquid. If the electrodes are connected together by a path with a much higher conductivity than the water, such as a wire, then the current will flow mostly through this path.

Streaming current is a more fundamental quantity than streaming potential, as it does not depend on the water's conductivity.

Streaming current can be measured whenever the water with suspended charged particles is forced through thin capillaries or other barriers to the particle's motion. However, useful measurements are extremely difficult to directly make from this effect, as the current produced is extremely small and easily obscured by potentials that exist for other reasons and electrical paths through the environment. Difficulties such as the dc offset in the electronics and electrode drift make this simple approach virtually impossible to use in a practical sensor.

Strictly speaking, zeta-potential is a property of a surface, rather than of a group of particles. A variety of surfaces can produce zeta-potential (and streaming current) in water. These, normally insignificant, effects are well recognised in materials science.

This work, however, uses those terms in the manner typical of discussion focused on water treatment, and considers them to be bulk properties of suspended colloids. Thus a sample of turbid water is said to have a certain streaming current, when correctly it is the average zeta-potential of the surfaces of the suspended colloidal particles that produce a certain streaming current in a certain sensor[1].

4.3.2 A Practical SCM Sensor Design

The practical streaming current meter (SCM) was invented around 1966 and credited to F.W Gerdes. Gerdes' approach [Gerdes 1966a] [Gerdes 1966b] overcomes many of the limitations of direct streaming current measurement and is the basis for all modern SCMs.

The SCM is based on the effect where the walls of the capillaries through which the colloidal material flows quickly gain a coating of particles and take on the surface charge characteristics of these particles.

The SC sensor consists of a piston and a close-ended chamber. A narrow gap or annulus, approximately 200-500um wide, exists between the piston and the walls of the chamber The piston is driven up and down at a fixed frequency, typically 4-5 strokes per second, forcing sample water in and out of the chamber through the annulus.

A typical inline SC sensor is shown in figure 4.3.

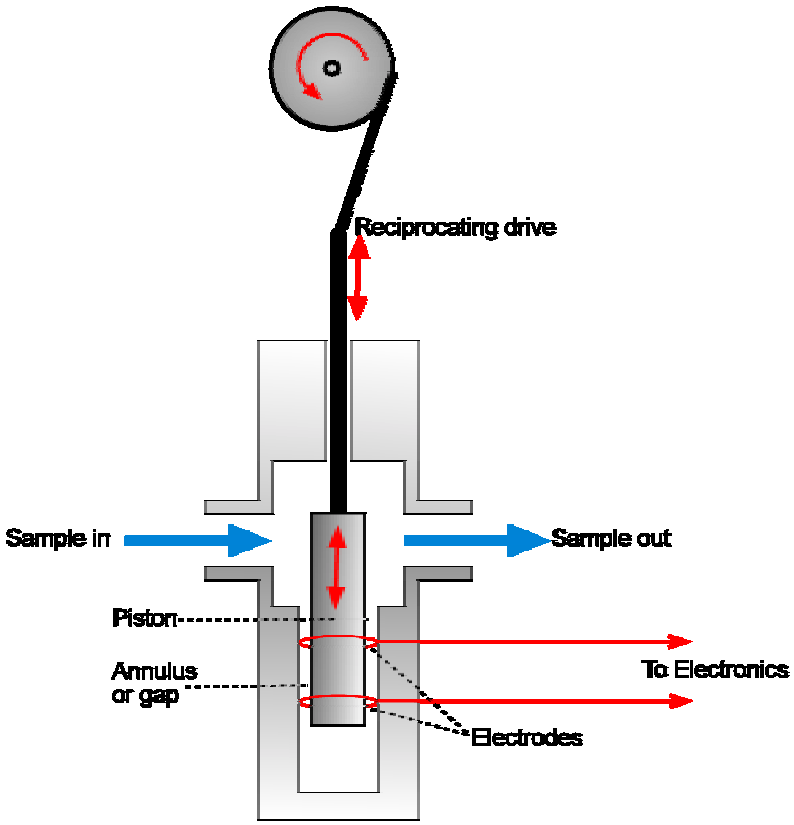

Figure 4.3 Streaming current sensor.

[1] It is clear why the former terminology is preferable in this context.

As the piston and chamber surfaces have become coated with charged particles, the water flowing rapidly up and down through the annulus results in displacement of the counter-ions. The SC signal measured by electrodes in the annulus is proportional to the water velocity and therefore alternates in time with the piston. This signal is typically in the range of 0.05uA to 5uA depending on the particular conditions.

Measurement of SC in a closed chamber has several advantages compared to measurement directly in a flowing stream:

- The closed end is electrically isolated and removes problems caused by large potentials in the process stream from other sources.

- The signal is alternating at the frequency of the piston. This allows it to be separated from external noise and offset caused by electrode drift and dissymmetry.

- Practical aspects of instrumentation, such as that the closed chamber can be shielded from electromagnetic interference and cleaned easily.

4.3.3 Mathematical Description

Unfortunately, a complete, quantitative and verified mathematical description of the sensor does not exist.

Several attempts have been made at modelling the SC sensor. The original approach by Gerdes [Gerdes 1966a] was based on a simplified triangular fluid velocity profile within the annulus. Later work [Elicker et al 1992] used more appropriate fluid profile model. The most complete work to date [Walker et al. 1996] also considers inertial effects within the annulus and compares the solutions with those found by earlier methods. Walker et al found that for the geometries used by typical SCMs the earlier approximate solutions produce very similar results to their more complete model.

All the theoretical models have the following form:

$I = k.s.\omega.\varepsilon.\zeta.f(r,R)$

Where I=average current magnitude, s=piston stroke length, ω=motor cycles per second, ε=dielectric constant of solution, ζ=zeta potential, r= piston radius, R=chamber radius, k=gain constant, f()=function which depends on the model used.

The models assume that the surface of the piston is completely and uniformly coated by the colloidal particles and by any coagulant in solution in a representative manner. This is not necessarily a valid assumption [Barron et al. 1994] [Dentel and Kingery 1989] [Dentel et al. 1989a] and probably goes a long way towards explaining the poor predictive abilities of this theory.

Gerdes acknowledged the poor predictive abilities of this type of model, and acknowledged the existence of double-equilibrium theory. This states that the charge on a particle surface is the result of an equilibrium between the ions on the surface and in the liquid. The charge on the piston surface depends on another but similar equilibrium between the same ionic concentrations in the liquid and a different adsorbed density on

the piston surface. The charge on the two different types of surfaces, which are both in equilibrium with the same liquid, will be related to one another in a way that depends on their respective equilibrium constants. Using this theory, it is only possible to predict that any change in the colloidal system that changes zeta-potential will change SC in the same direction.

The modern colloidal-science theory papers implicitly reject this equilibrium-based theory providing no discussion of how, or at what rate, the measurement surfaces take on the zeta-potential characteristics of the sample. If there exists an equilibrium between the charge on the particles in solution and the measurement surfaces then it would influence the rate at which measured SC responds to changes in coagulant dosage applied to an otherwise constant sample. Indeed in this chapter section 3.3 shows that by starting with this assumption it is possible to develop a simple dynamic model which is a good fit for the sensor's observed dynamic response.

However, even if a perfect theoretical model of the SC sensor was possible and did exist, it would still be of limited practical use. This is because it would require complete characterisation of the piston surface and the precise geometry of the annulus. Surface characteristics and the tight tolerances are easily altered by physical wear and tear, even over relatively short time frames. It is quite possible to calibrate an SCM so that it reads directly in units of zeta-potential over a limited range [Elicker et al 1992] [Dentel et al 1989a] or in units matching another SCM, however it will not stay calibrated like this when in continuous use.

The unit of SC should properly be the unit of current (the Ampere). However, the lack of a consistent relationship between this and any useful physical or chemical parameters means that SCMs normally use arbitrary units. Their scales are designed to be adjusted for convenience rather than to directly display current units and can be adjusted by their users. This thesis will follow this convention and use the term 'SC' to refer to the arbitrary units used to measure SC as well as physical streaming current.

4.3.4 Usage and History

The most useful way to apply a SCM in a typical water treatment plant is for the feed-back control of coagulant dosage. The sample is taken from after the rapid mixing stage and the SC continuously measured so that a controller can adjust the coagulant dosing rate to maintain a set SC. This is illustrated in Figure 4.4. The SC that corresponds to the desired plant operation is traditionally determined by manual jar tests. Automated jar tests (as describe in chapter 3) are also likely to be able to do this.

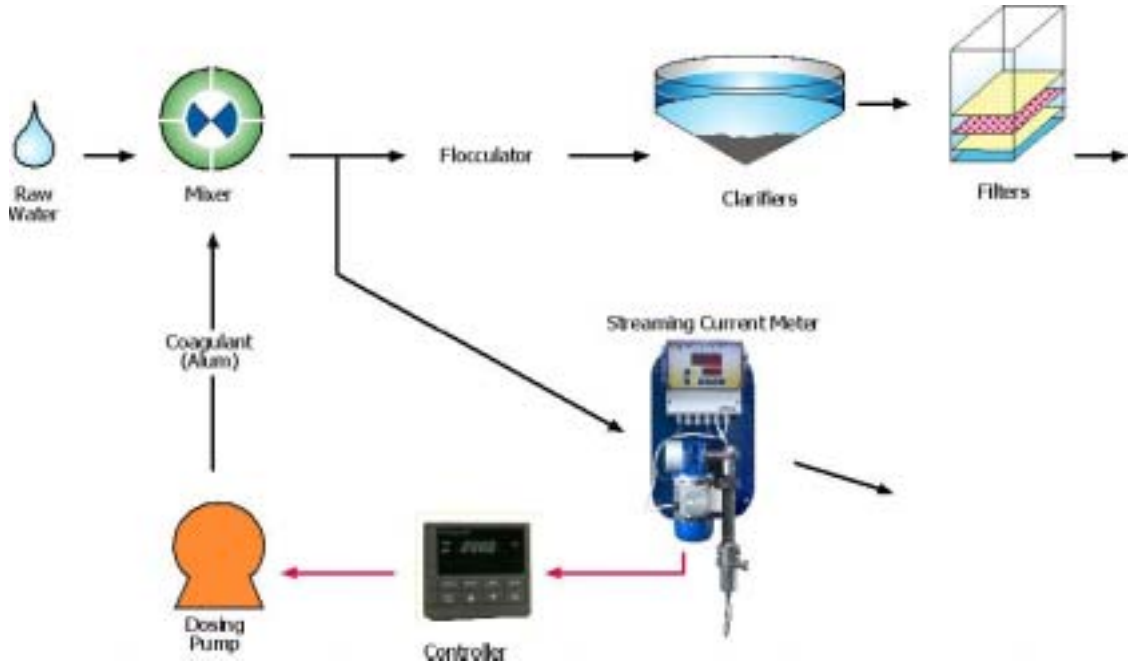

Figure 4.4 Diagram of SCM Usage in a Typical Plant

Caution and a lack of understanding kept the SCM from being used in most plants until the early 1980s. The study most widely cited, and credited with bringing widespread acceptance for use of the SCM in water treatment is the work of Dentel and Kingery, most notably [Dentel and Kingery 1989]. This includes the result that SCM use resulted in generally improving water quality and consistency, as well as decreased chemical usage by 12% during stable operation and 23% during changing raw water conditions. A survey of plant's using an SCM showing 80% of them regarded it as useful.

It is likely that these percentages would be higher today, even though many more SCMs are in use, due to better understanding by both users and manufacturers.

A problem frequently experienced include clogging of sample lines, which can normally be solved by appropriate sample point and sampling design.

Another significant problem is wear and tear on measurement surfaces and the mechanisms. SCMs are notorious for being unstable. Among some water treatment plant operators, the successful operation of a SCM has developed a (in the author's opinion, unjustified) reputation for being more of an art than a science.

However, the instrument's usefulness is no longer in question. For example, the New Zealand Ministry of Health requires plants supplying more than 5000 people to have automatic control of their coagulant dosage to in order to achieve either of the highest 2 grades (to be more than merely 'satisfactory') [MOH 2003 pp 34-41]. Currently, an SCM is the only technology that can provide this.

4.4 The Need for Signal Processing

The SC sensor produces an extremely small, alternating current signal. To produce a useful measurement the signal must be amplified and processed into a relatively noise free, stable, useful, representative reading. The complete streaming current meter (SCM), is therefore a SC sensor with a set of electronics to produce this final reading.

The theory predicts that the instantaneous measured current will be directly proportional to instantaneous water velocity. As inertial effects can be neglected for the annulus dimensions used by SCMs, the instantaneous water velocity is determined by the instantaneous piston velocity. This will be a sinusoid in phase with the piston's motion, with its magnitude determined by size of the zeta-potential of the sample and its orientation (whether it is inverted or not) determined by the zeta-potential's sign.

In the idealised case of a perfect sinusoid the wide range of methods available for converting this alternating current (ac) signal to direct current (dc) reading are all equally valid, and the electronics and signal processing design is trivial.

Unfortunately, a number of problems complicate this:

- The small ac signal has a dc offset that is significantly larger than itself. Gerdes attributed this to electrode asymmetry. This author feels that potential difference between the sensor and the electronics are a more likely explanation. However isolation and shielding does not fully remove it, so the fact remains that the ac signal has a dc offset that can be several orders of magnitude larger than itself.

- The signal is polluted with noise. This takes two forms: 1. Higher frequency electrical noise picked up from nearby power conductors at 50Hz or 60Hz and other electromagnetic interference from operating electronics. 2. Sudden step-like shifts in the signal, possibly caused disturbance of the piston surface by the movement of larger individual particles through the sensor from the sample water.

- This signal is related to the piston's motion in a way that is not well explained by the available theory. This includes such effects as a phase shift that is particularly prominent near zero SC, and consistent distortion of the sinusoidal shape of the signal. Nether is explained by inertia or other mechanical phenomena. These effects are examined in more detail in section xxx.

4.5 The Classical Approaches

The signal processing proposed by Gerdes in 1966 is widely considered the best approach, and is used by a majority of modern SCMs [Walker et al. 1996]. A simplified diagram of this is shown in figure 4.5.

The input capacitor, C1, is used to remove the input offset to prevent it saturating the amplifiers. The current to voltage conversion circuit also amplifies the signal to useable levels. The synchronous rectifier consists of a switch that is controlled by a sensor that detects the position of the piston. This position sensor is commonly a rotating disc on the motor that drives the piston. When the piston is going up the switch is one way and when it is going down it is the other. This converts the ac signal to a dc signal that is positive or negative depending on the orientation of the ac signal. A final low pass filter removes the cycles at the piston frequency as well as other high frequency disturbances.

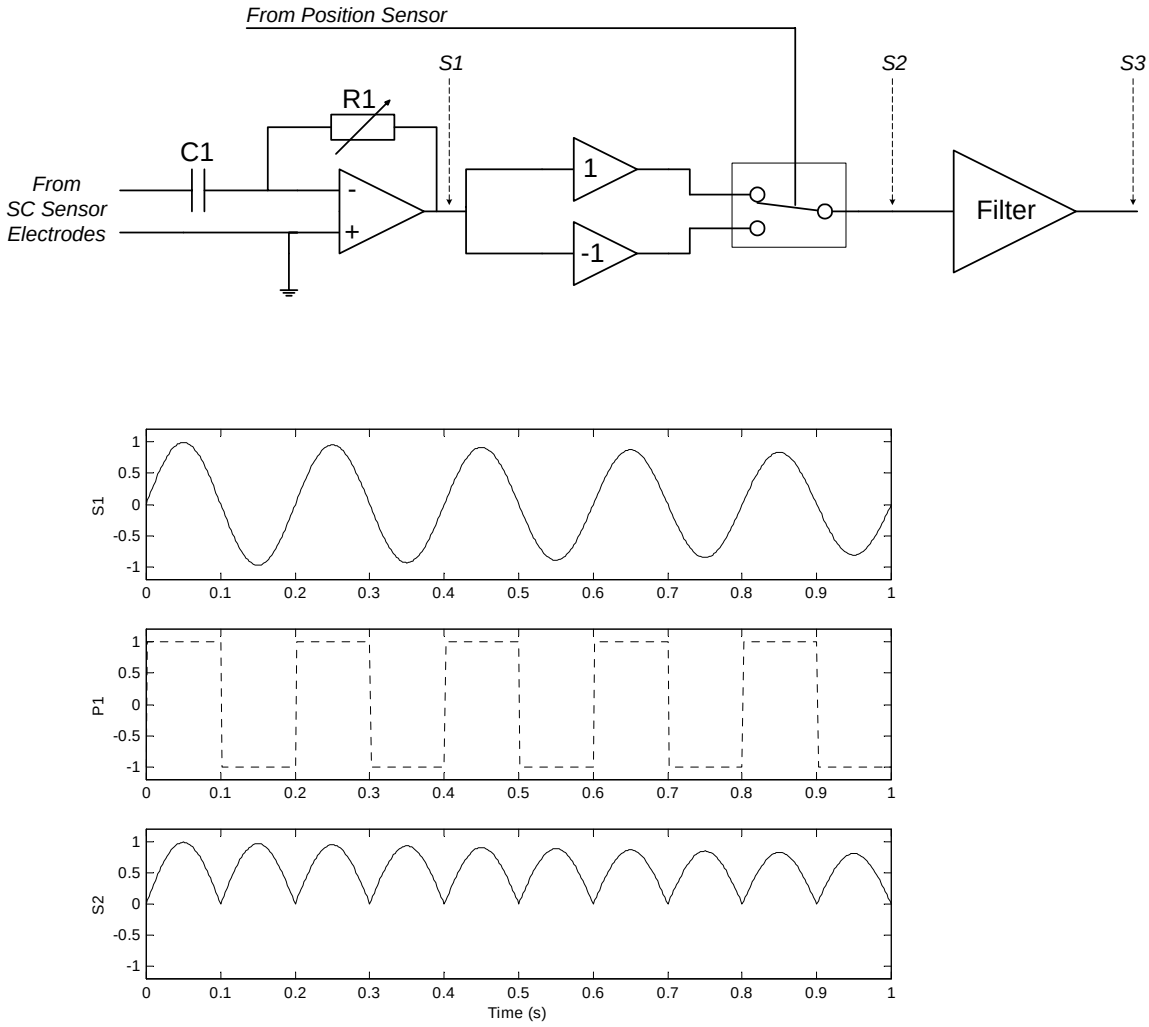

Figure 4.5 Top: Synchronous rectification signal processing circuit. Bottom: Waveforms at the marked points.

In a real SCM, the structure and order of the amplification, filtering and switching will be considerably more complicated, however the same basic approach is used.

An even simpler method, used by earlier SCMs, is to use sample and hold circuitry. This is shown in figure 4.6. The current to voltage conversion is the same. The position sensor now activates a switch briefly in the middle of the (say) downward part of each piston stroke. The signal level from this instant is then held until the next stroke. This approach has the advantage of needing less filtering at the end and therefore has a quicker response, but is much more susceptible to noise and distortion as it is only looking at a small part of the signal.

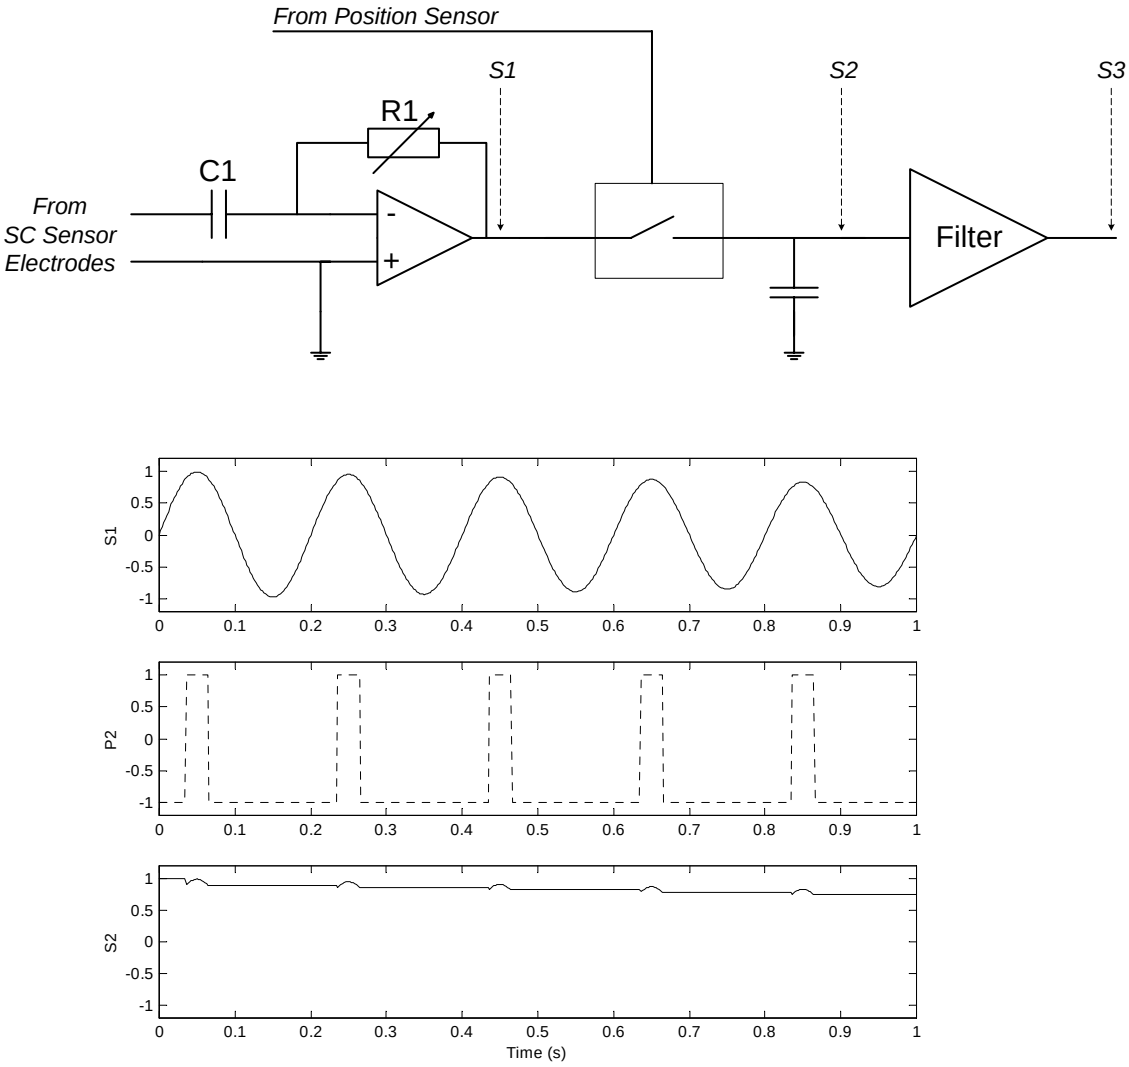

Figure 4.6 Top: Sample and hold rectification signal processing circuit. Bottom: Waveforms at the marked points.

4.6 The Fourier Transform

Let us consider, for a moment, what the signal processing is intended to accomplish. It needs to establish the strength and phase of a sinusoidal signal at the piston's frequency from a background of noise and disturbances. Fortunately a mathematical technique exists for doing exactly that – the Fourier transform.

Modern microprocessor technology allows a level of computing power unheard of in 1966 to be embedded in the signal processing electronics for less than 1/10 of the cost of the sensor manufacturing alone. Therefore there are no barriers to the implementation of this method.

The Fourier transform is one of the fundamental tools of digital signal processing. It is a generalisation of the Fourier series widely used in mathematics and engineering and is presented in a wide range of textbooks, for example [Attenborough 1994 pp 503-516].

It is based on the fact that any periodic signal can be expressed as the sum of sinusoidal components. If one period of this periodic signal is sampled at N evenly spaced discrete time steps, then the discrete signal, x(n), may be represented as:

$$x(n) = \sum_{k=0}^{N-1} X(k) e^{j 2\pi kn/N} = \sum_{k=0}^{N-1} X(k)(\cos 2\pi kn/N + j.\sin 2\pi kn/N)$$

where

$$X(k) = \frac{1}{N} \sum_{n=0}^{N-1} x(n) e^{-j 2\pi nk/N}$$

$$j = \sqrt{-1}$$

X(k) is the Fourier transform and is a complex function. The Fourier transform has a far wider range of forms, theorems and uses than this. More information can be found in any textbook on modern digital signal processing, for example [Bellanger 2000].

As x(n) is strictly real, it follows that x(k) and X(N-k) are complex conjugates, therefore :

$$x(n) = \frac{1}{2} \sum_{k=0}^{N} \left[\mathrm{Re}\left(X(k)\right) \cos 2\pi kn/N + \mathrm{Im}\left(X(k)\right) \sin 2\pi kn/N \right]$$

In summary this means that any periodic signal can be broken down into a sum of sinusoids at frequencies that are a multiple of its frequency. X(0) is the constant dc part, Re(X(1)) is the cos component at the base frequency, Im(X(1)) is the sin component at the base frequency, Re(X(2)) is the cos component at twice the base frequency and Im(X(1)) is the sin component at twice the base frequency, and so on.

For the ideal waveforms predicted by the theory it is only necessary to compute the cos component of the base frequency component, Im(X(1))[2]. As in this situation, the frequency of the wanted signal is known and all dc, off-frequency and out of phase signals are to be rejected. The online signal processing does not therefore have to perform a full Fourier Transform to obtain the desired component, it can be calculated by Fourier Analysis which involves only one multiplication per sample step.

Note that the mean of X(1) over m stoke cycles can be calculated either by calculating it over each of the m cycles and averaging, or by calculating it over all m cycles.

The use of Fourier transform (FT) or synchronous rectification (SR) has two purposes:

1. Rejection of interference (noise). Interference can be defined as the signals that are not related in any way with stroke and therefore are from an outside source. These signals at not at a multiple of the stroke frequency, and are mostly considerable higher frequencies.

[2] It is unfortunate that the most relevant component of the signal is called the imaginary part. This is because the switch that synchronises the sampling is triggered at the top of a stroke cycle.

2. Rejection of distortion of the SC signal. In this chapter distortion is defined as any consistent alteration of the shape of the signal from what is desired. Exactly what is distortion and what is a useful part of the signal is not clear and will be examined later. However distortion is distinct from interference as it is, of course, at a multiple of the stroke frequency. (Making it present in the same way in each stroke)

The following sections compare the FT and SR approaches.

4.7 Noise Rejection

The raw signal coming from the electrodes will need some filtering before processing by either method, in order to prevent saturation of later amplification stages by large, high frequency interference signals, such as received radio signals and switching transients. This is ignored as these filters will be designed to have an effect only a frequencies significantly higher than the stoke rate (for example 10-100 times higher).

In the following experiments the stroke period is 206 ms (a frequency of 4.9Hz) which is typical. It is assumed that both are implemented digitally with a sample rate of 500Hz. This was chosen so as not to place unreasonable demands on the anti-aliasing filter required by the Analogue to Digital Converter. The results will be very similar for other sampling frequencies, or in an analogue implementation of the SR.

In order to use the normal signal processing analysis techniques, the two processing methods were modelled as Finite Impulse Response (FIR) filters. The only difference from conventional filters is that they are phase sensitive and so must be synchronised to the piston motion. In simulation this is easy to ensure, simply by designing the filter for the correct period (103 steps in this example) and by starting the filter at the correct point.

Figure 4.7 shows the FIR coefficients for both filters. In a real implementation it is necessary to alter the length and coefficients of the FIR filter to keep it correctly synchronised to the piston sensor. For the SR this is trivial and can be implemented in analogue electronics.

Five consecutive strokes (about 1 sec) are averaged to simulate the use of a rolling average filter. A rolling average post-filter is particularly elegant here, as it can be implemented and analysed as the result of more repetitions of the same filter co-efficients. The rolling average filter is widely used as a post filter because it is also simple and intuitive for end users who often adjust its length to suit the level of noise in their application. However, even if a more effective method is used to post-filter the per-stroke results, it will have a similar effect on both signal processing methods.

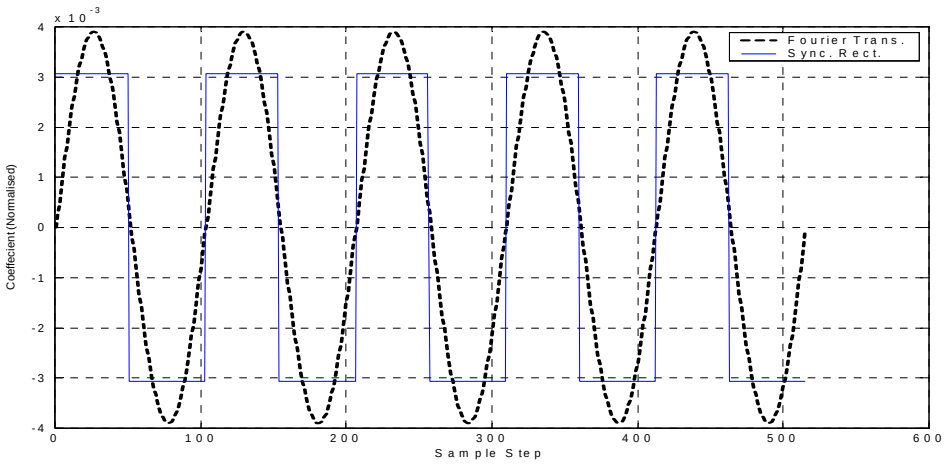

Figure 4.7 Coefficients of the FIR filters. Normalised to unit response to the ideal input.

It can be seen that the two methods are quite similar. The SR has the advantage of being easier to implement, while the FT has a more sound theoretical foundation.

Figure 4.8 shows the frequency response of the FT and SR at frequencies above the stroke rate. The FT approach is clearly better, having at least 30 dB[3] of attenuation more at the critical power supply frequencies of 50 and 60 Hz, which is a major source of interference.

Figure 4.9 shows part of figure 4.8 in the region of the stroke frequency. This indicates the FT method also has better selectivity and a generally flatter response.

[3] The definition of dB of attenuation is: x dB=$10^{(x/10)}$, so -30dB = 10^{-3}

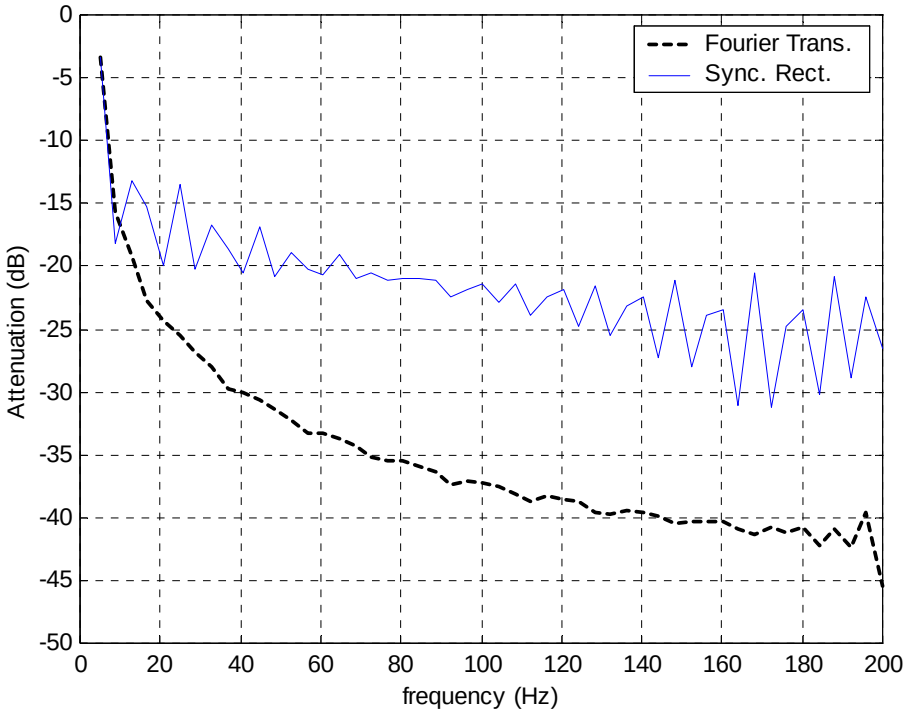

Figure 4.8 Magnitude frequency response of FT and SR methods with 5 consecutive strokes averaged. (Mean magnitude response per 10 Hz)

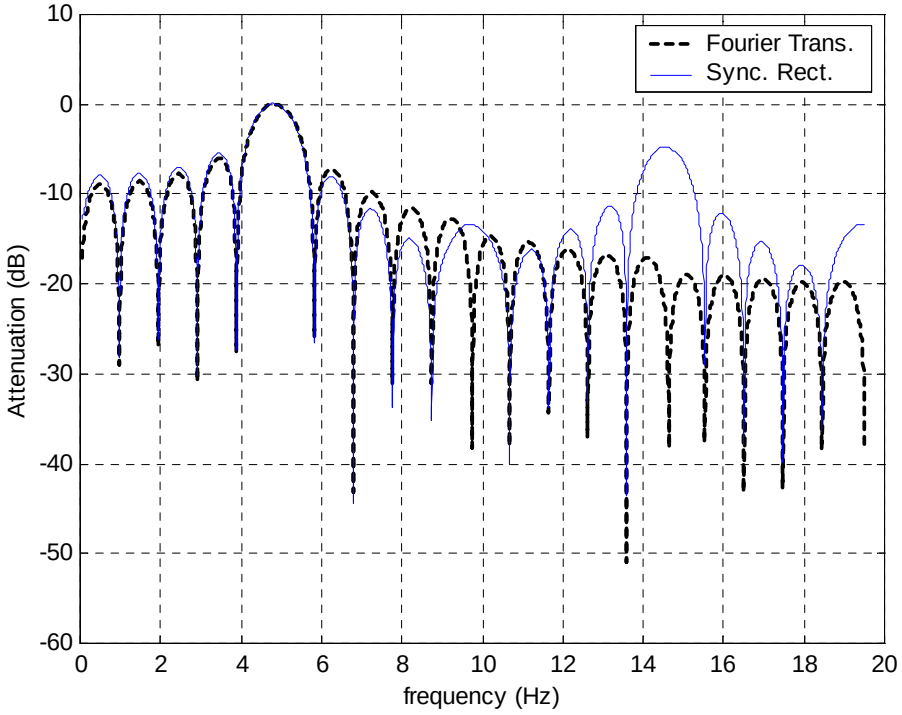

Figure 4.9 Magnitude frequency response of the FT and SR methods with 5 consecutive strokes averaged.

It is, of course, possible to filter the result of the FT or SR processing to arbitrarily increase the noise attenuation. The most commonly implemented filter is a rolling average over an adjustable time frame. This results in a, very approximate, increase in attenuation of 20dB per factor of 5 over the relevant ranges. The disadvantage of doing this is that the measurement bandwidth is decreased.

This can be illustrated by the use of a rolling-average filter to increase the attenuation of the SR method to that of the FT method. The frequency response is shown in figure 4.10. The averaging is over 5 strokes (1.03s) for the FT and 120 strokes (24.6s) for the SR. While the SR method with filter has a considerably better roll-off at low frequencies, the attenuation in the critical 50/60 Hz region and above is similar at 30-35 dB.

The signal generated by the sensor is extremely small, it is impossible to completely shield the sample chamber, and the instrument is intended to operate in proximity to large pumps and solenoid valves; therefore, 30dB of attenuation at 50/60hz can be considered a minimum.

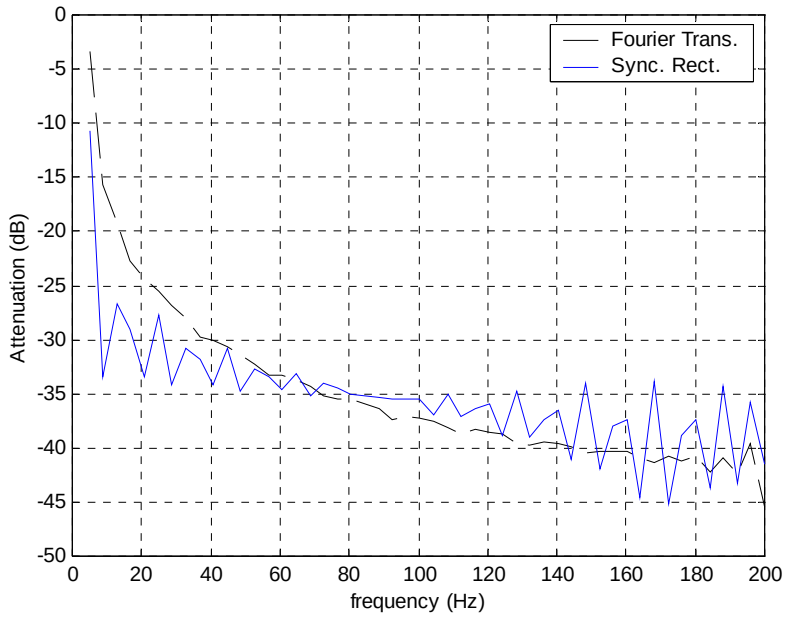

Figure 4.10 Frequency response of the FT method with 5 consecutive strokes averaged and SR with 120 averaged. (Mean magnitude response per 10 Hz)

The desired SC signal can be considered as a signal that has been modulated by the piston stokes in the same manner as the amplitude modulation used in radio transmission[4].

[4] Unfortunately, most of the mathematical tools for analysing and processing these modulated signals require a much greater separation between the baseband and modulation frequencies.

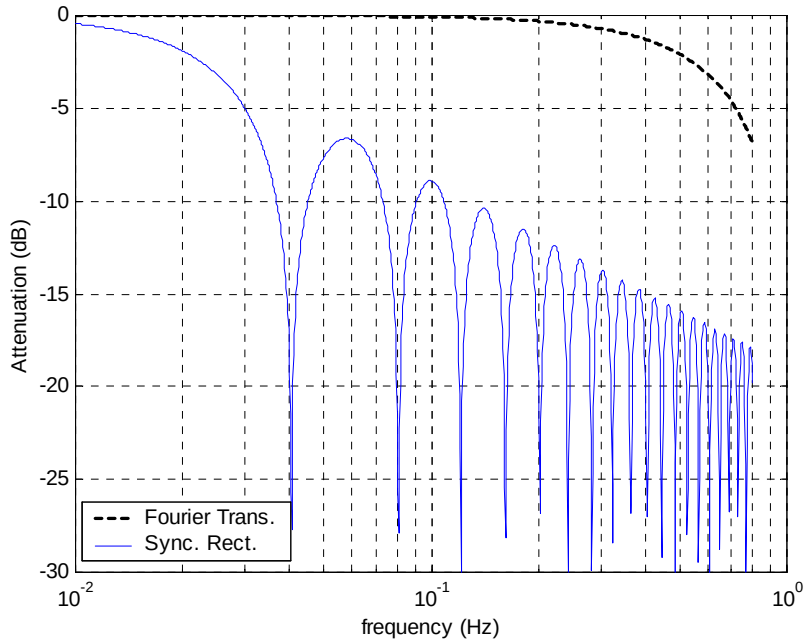

Figure 4.11 Frequency response to SC of the FT method with 5 consecutive strokes averaged and SR with 120 averaged. Note that the SR method has an identical response to the FT when 5 strokes are averaged.

With the averaging over 5 stokes both methods have very similar frequency responses, having a measurement bandwidth (as measured by the 3dB down point) of 0.59Hz. With the 120 strokes of averaging the bandwidth is 0.024Hz.

The selection of a more advanced high order post-filter would improve the SR results without having quite such an adverse effect on bandwidth. Whatever trade-off is made between noise attenuation and measurement bandwidth, the FT method is advantageous. How this signal processing contributes to the dynamic response of the complete instrument is examined in section 4.10.

4.8 Distortion Rejection

To usefully analyse the effects of distortion it is necessary to establish what sort of distortion occurs during normal operation of the SC sensor. It is also necessary to establish which effects are ‘distortion’ and which effects are useful parts of the signal. A significant amount of design goes into the physical layout of the sensor in order to minimise distortion, however the shape of the signal is never a perfect sinusoid although it can be very close.

4.8.1 Properties of the Signal from a new Piston

The following experiment was performed to examine the relevance of different components of the signal.

A sample of tap water from a bucket containing 100mg/L of Kaolin was circulated through the SCM. Kaolin is often used to simulate typical contaminants in surface waters, for example in [Bratby 1980 pp 107-109]. In this case, an accurate representation of the flocculation and settling is unnecessary, as only the charge

neutralisation effect of coagulation is being simulated. Alum was added to the bucket in 2ppm increments by a peristaltic pump. It is well known that a kaolin solution has a negative zeta-potential and the addition of aluminium sulphate makes it more positive. A signal component can therefore be considered relevant if there is a good relationship between it and the alum dose.

The digital interface to the Accufloc SCM was used to obtain signal data. This SCMs Analogue to Digital Converter (ADC) is a separate integrated circuit, which has a sampling rate of 38.4kHz, and a digital low-pass filter to reduce it to 500Hz. There is an appropriate anti-aliasing filter as part of the analogue electronics. The ADC has a 16-bit resolution and a selectable pre-amplification of up to 128 times, effectively giving it 23-bit resolution. The signal processing microprocessor receives samples at an interval of 2ms, which is a high fidelity representation of the true signal coming from the sensor. A PC can obtain this raw data over the SCM's RS-485 serial digital interface. The instrument can store only one stroke worth of signal data at a time and re-transmission to the PC takes longer than one stroke time. Data from one stoke data is downloaded and recorded every 2 seconds.

Large positive or negative readings produce a near-perfect sinusoidal shape. The most interesting readings are those close to zero, because this is where the distortion is most significant. In addition, the normal operating point of the coagulation system is quite close to zero.

Figure 4.9 shows the SC reading from this experiment plotted over time (this is Im(X(1))). The rapid SC reading increase that occur after each alum dose is caused by poor mixing in the bucket, as the location of the dosing and sampling lines invariably means that part of the newly dosed water is drawn into the SCM before mixing is complete. The drift downwards between the alum additions is caused by two factors: 1. the gradual approach to equilibrium of particles on the piston (see section xx on dynamic response) and 2. most likely the reaction of the acidic solution with air, which results in an increase in pH and thus a decrease in zeta-potential (due to double layer compression). The gradual change in the pH of the dosed water over several minutes closely matches the observed SC change.

Figure 4.10 shows the unprocessed SC waveform for each step, recorded just before the next alum dose is applied. The waveform is quite close to being the expected sinusoid, but there are significant variations.

There is often a slight asymmetry between the upward and downward strokes. A possible explanation is uneven flow in the upwards and downwards directions, caused by the fact that the piston forces the water out of the cavity, while only atmospheric pressure forces it back in again. Because water is effectively incompressible, any asymmetry caused by this effect would result in some cavitation, which is the formation of a small area of vacuum under the piston during the upstroke. Unlike most types of cavitation, this will not result in any mechanical stress to the mechanical drive, provided that the vacuum is filled before the piston starts to come down again.

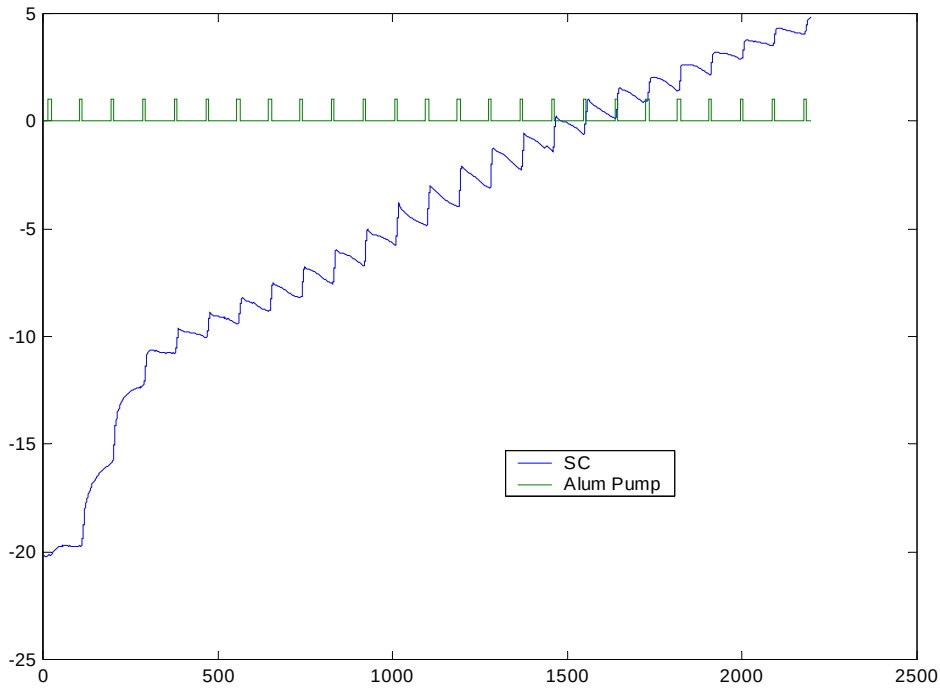

Figure 4.9 SC and alum pump speed during the experiment. Sample time = 2sec.

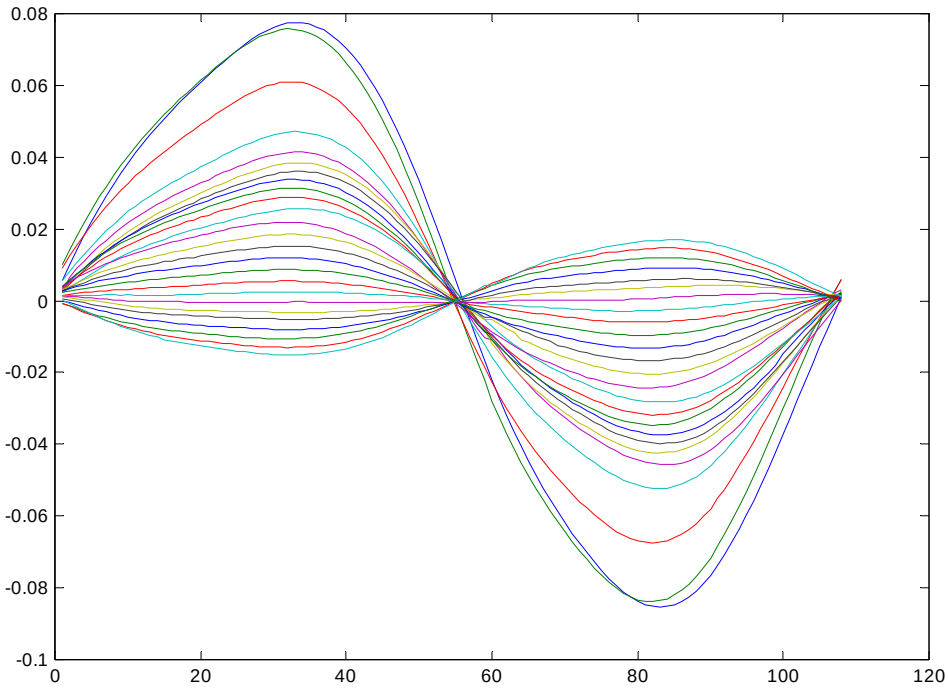

Figure 4.10 The SC waveform for each step.

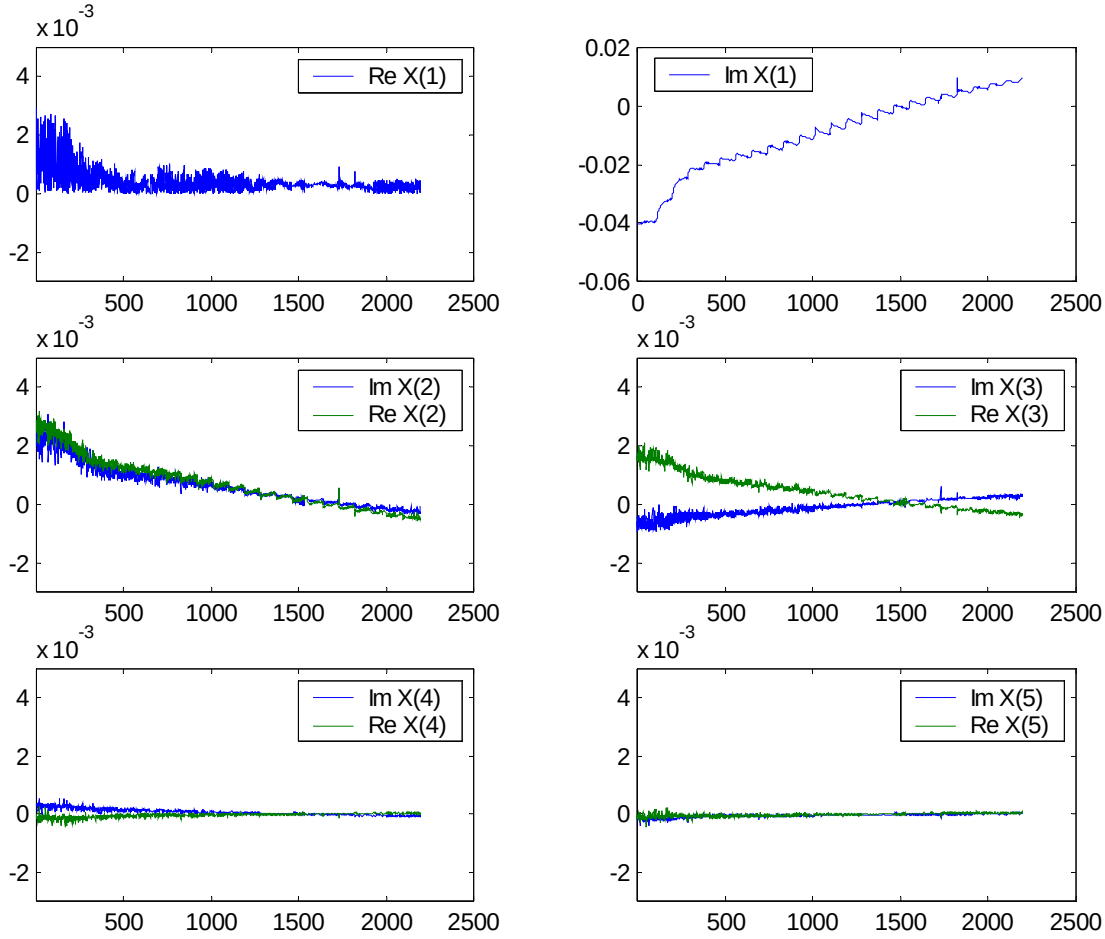

Figure 4.11. Fourier components of the signal during the experiment. Sample time = 2sec. Note that Im(X(1)) dominates the results and is shown on a different scale than the others.

Figure 4.11 is the results of a FT on the signal over time, each graph shows how a different component's real and imaginary parts change over the duration of the experiment. As the experimental conditions means that the FT is only carried out over a single period there is a greater level of noise than would be present in the reading calculated continuously online.

All the components except Im(X(1)) are removed by the FT processing, while only the components on the left of this diagram are completely removed by SR processing. It can be seen that the X(3) components have approximately 10% of the magnitude of the Im(X(1)) component. The X(3) components are reduced by a third by the SR processing, combined this brings their contribution to about 7% of the SR reading. This is not necessarily a problem, because in this case both components are well correlated with the alum dose. However, there is some evidence suggesting that the harmonic components of the signal change significantly as the condition of the sensor changes, this means that removal of the X(3) components of the signal may be more beneficial than indicated here.

The results presented here of this experiment are consistent with a large number of other experiments performed on new, defect-free, Accufloc SCMs. In fact, this test is a powerful way to predict future problems with a particular unit.

4.8.2 The Effect of Contamination on the Signal

With the new, clean, sensor tested above, the effect of the components other than Im(X(1)) can be safely ignored. However this is not possible as the sensor becomes older or if it has accumulated a coating of dirt.

Figure xx shows the result of a longer dosing experiment on a continuously flowing stream. The alum dose is varied and both the real (Re(X(1))) and imag (Im(X(1))) parts the SC signal recorded.

This data is one of the tests of the effect of alum dose on settled particle count carried out on the small-scale plant model (covered in section xxx.). In this experiment a self-tuning regulator was used to make the imag part of the SC follow a repeating pattern.

At the start of the experiment the sensor is shut down and cleaned, the signal returns to a near perfect sinusoid and the real part of the signal becomes insignificant compared to the imag part. As time progresses and dirt accumulates on the piston, the magnitude of the real part gradually increases, while the alum dosage required to control the imag part stays constant.

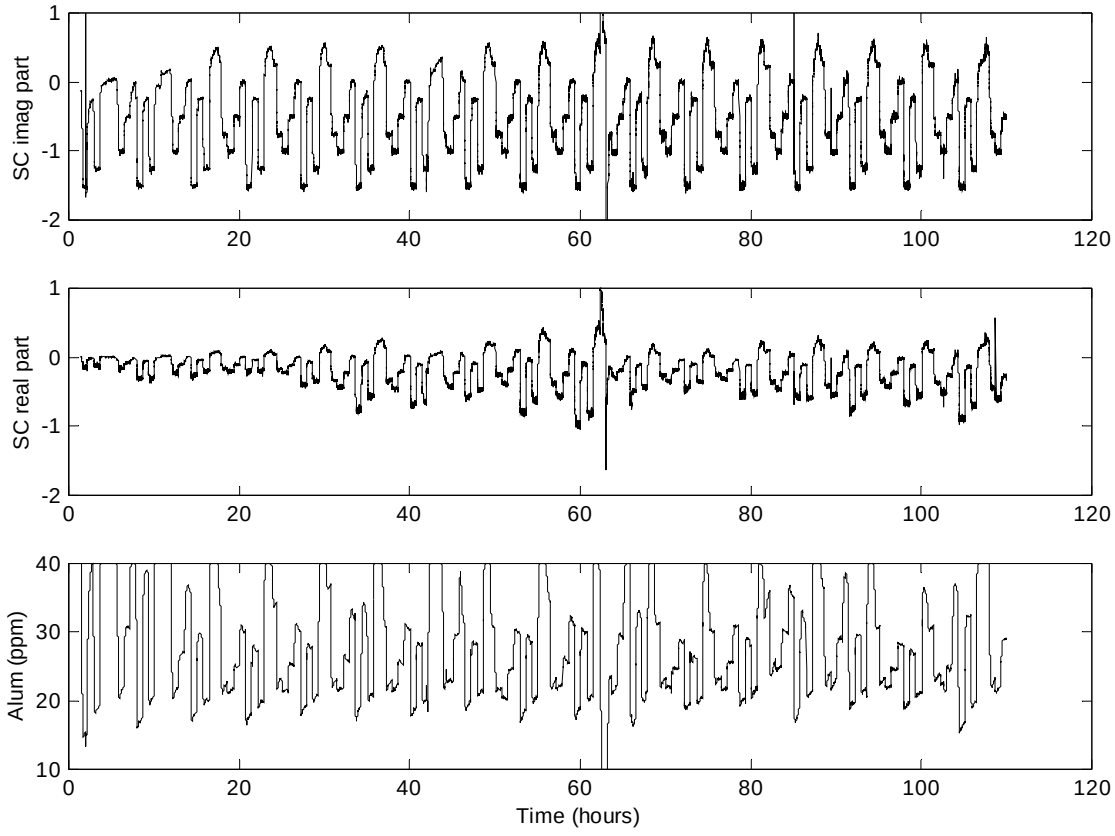

Figure 4.12 Real and Imag parts of SC signal and Alum dose over a 110hr test during a period of higher turbidity.

The results shown are broadly typical of the large number of similar results obtained. The main feature that makes test test unique is that at around 62 hours the sample was accidentally cut off by plant maintenance, bringing the flow to zero. After flow was restored shortly afterwards the imag part's response was identical, while the real part decreased and then gradually increased again. The sudden turbulent restarting of flow is likely to have had a cleaning effect on the sensor. This significantly affected the real part but not the imag part. A similar effect can be produced by manually cleaning the piston.

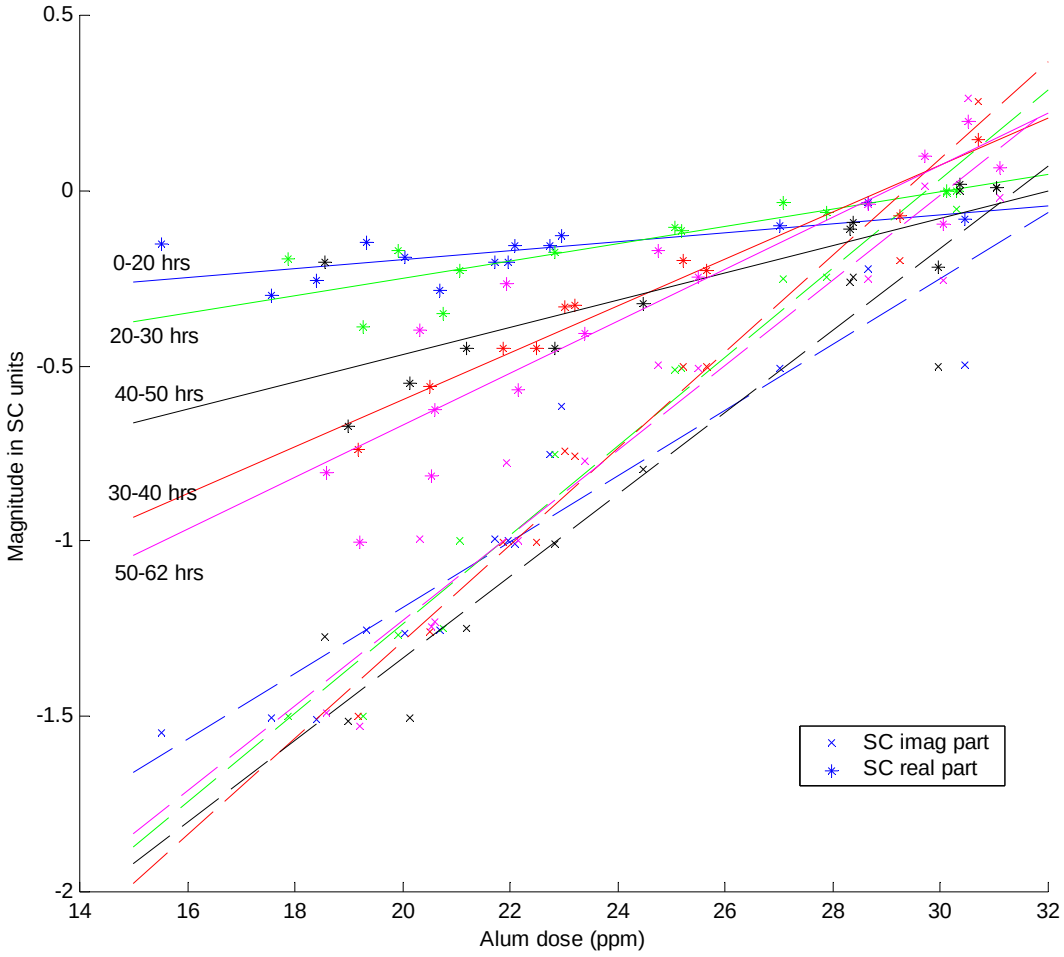

Figure 4.13 The crosses (x) and the dashed lines are the imaginary part of the signal. The asterisks (*) and solid lines are the real part.

Figure 4.13 shows an approximation of the steady state for both real and imaginary parts against alum. It is only an approximation, as the system does not become completely steady. The lines are the least squares best fit to the data points from the time period indicated. Table 4.1 presents the parameters fitted. The distance of these points from zero alum dose and the fact that the units of SC are meaningless mean that

the slope parameters are more important than the offset parameters. The slope values are shown in bar chart form in figure 4.14.

Table 4.1. Parameters of the LS fitted lines in figure xx. (Slope and offset are defined such that SC_{linear} = slope * alum + offset)

Hours	Imag slope (SC/ppm)	Imag offset (SC)	Imag MSE (SC)	Real slope (SC/ppm)	Real offset (SC)	Real MSE (SC)	No. of samples
0-20	**0.0939**	-3.0675	0.1445	**0.0130**	-0.4572	0.0511	14
20-30	**0.1271**	-3.7782	0.0936	**0.0250**	-0.7501	0.0647	12
30-40	**0.1380**	-4.0478	0.0916	**0.0670**	-1.9376	0.0521	10
40-50	**0.1169**	-3.6726	0.1485	**0.0390**	-1.2482	0.1275	11
50-62	**0.1214**	-3.6552	0.1439	**0.0744**	-2.1578	0.1439	15
64-75	**0.1221**	-3.6955	0.1760	**0.0344**	-1.0781	0.0543	13
75-85	**0.1292**	-3.9684	0.2105	**0.0424**	-1.3580	0.1048	11
85-95	**0.1119**	-3.6205	0.0813	**0.0491**	-1.6123	0.0567	9
95-110	**0.1405**	-4.0646	0.0987	**0.0708**	-2.0523	0.0598	16

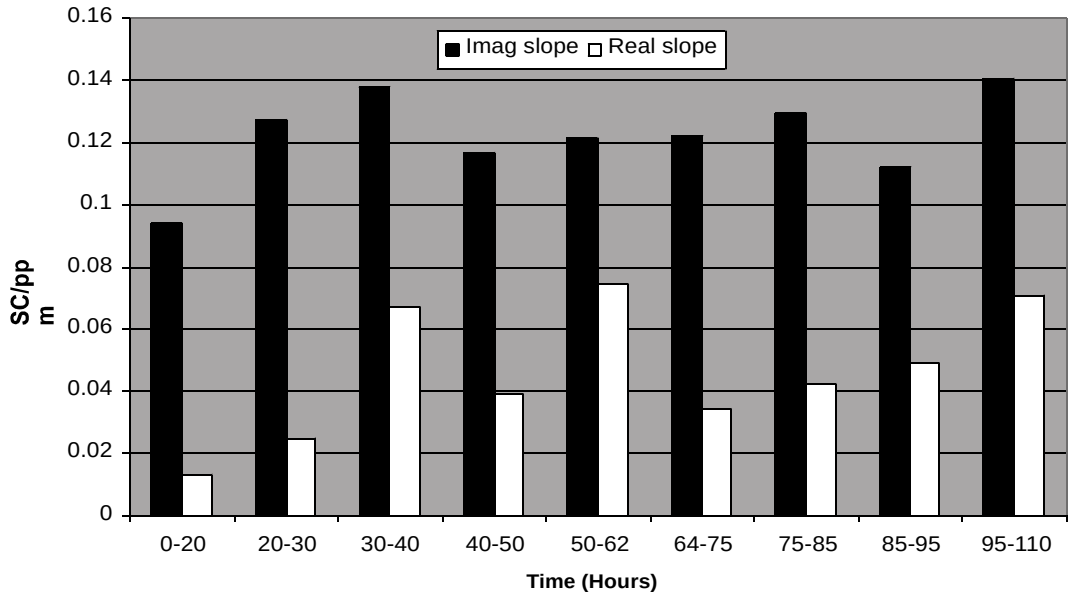

Figure 4.13 The slope values from Figure xx.

These graphs show more clearly how the slope of the relationship between alum dose and the real part changes in a fairly consistent positive direction after each even which has a cleaning effect on the sensor (such as at 0 and 62 hours), while the imag part stays relatively constant.

It is worth noting that the level of coating caused by the floc from this water source at that time, is significantly worse than normal, and the piston and sensor used was also relatively worn. This situation where large amounts of fresh and aggressive micro floc is present in the sample is very difficult to duplicate. This is unfortunate as it would be

beneficial to repeat this experiment and measure the effect of floc coating on the other components of the FT.

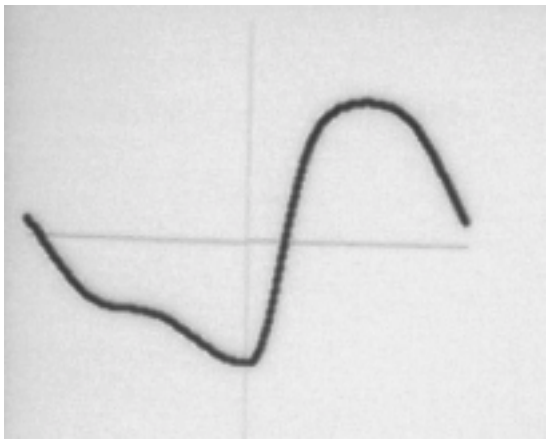

Figure 4.14. Waveform from SC sensor in the presence of significant floc contamination of the measurement surfaces.

The time domain waveform in the presence of floc build-up (show in figure xx) clearly shows that other components are present besides those recorded at the fundamental frequency. This distortion of the waveform is not predicted by the first principles theory previously presented. One possible explanation is that charged particles are being moved on and off the measurement surfaces during the stroke. This may be result from uneven flow of water caused by surface wear or damage to the measuring surfaces or mechanical misalignment caused by wear on the bearings which align the piston with the annulus.

It seems highly probable that the FT method of processing the SC signal will provide even greater benefits compared to the SR method in this situation.

It is also possible that analysis of the FT components that cause the distortion will be able to indicate the condition of the piston. This could be used to alert plant operators to the need to clean or replace the piston and avoid the common problem of drifting and unstable readings due to wear.

4.9 Dynamic Modelling

The standard theory in the previous section addresses only the steady-state output from the SC sensor. The inadequacy of the theoretical account of the way in which the colloidal particles coat the piston also means that there is no theoretical model of the sensor's response to changes in the zeta-potential of the sample, such as that caused by a coagulant dosage change. The previous examination of measurement bandwidth only considers signal processing effects and assumes that the physical changes are instantaneous.

Although this is clearly not true, experimental results do indicate that a reasonably large part of the sensor's response happens quite quickly. A typical step-response shown in figure xx, this was generated by a step change in coagulant dosage into a rapid mixer negligible detention time. Very similar results are obtained by rapidly changing the sample lines from one water source to another.

There is a transport delay, as the new water sample flows along the sample lines, then the signal changes by approximately 60% within 20 seconds of the signal beginning to change. The next 60% (25% of the total) generally takes less than another 20 seconds. Unfortunately the remaining 15% of the complete response does not fit the same exponential pattern and takes considerably more than the 10 minutes shown here to completely stabilise. This slow part can be >50% in many cases.

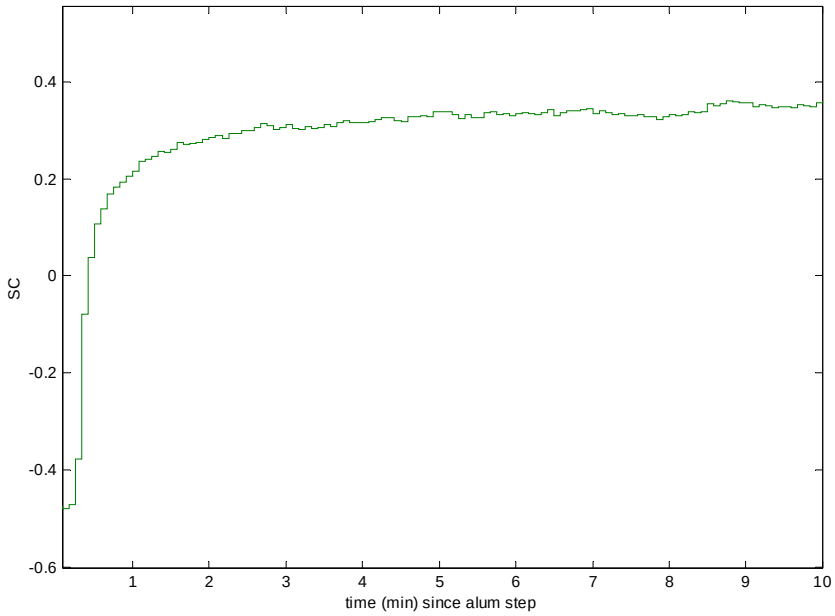

Figure 4.15 Typical SCM Step Response

A small part of this response is clearly due to back mixing in the sample lines, however only a minor increase in speed is obtained by removing the old sample from the sample lines (but not from the SCM sample chamber) first.

While the development of a complete first principles model is impossible, some further understanding can be obtained by attempting to develop a dynamic model some of the processes occurring.

It is reasonable to assume that the steady state reading produced by the SCM is directly proportional to the total charge absorbed with the particles onto the measuring surfaces of the instrument. The moving water therefore has an equal and opposite charge, which is what is actually measured.

If any given particle in the sample has a fixed probability of absorbing onto a surface, then the rate of particles accumulating onto the surfaces is proportional to the number of particles in the sample. Likewise, if any given particle on the surfaces has a fixed probability of returning to the sample, then the rate of particles leaving the surfaces is proportional to the number of particles on the surfaces. The current measured in one stroke is proportional to the sum of the charges on all of the particles absorbed onto the measuring surfaces. Because of this it is preferable to model the movement of charge in the system rather than the particles.

The rate of charge movement onto the surfaces is proportional to the total amount of charge distributed across all the particles that come into contact with the surfaces. This will depend on both the number of particles in the sample (eg turbidity) and the zeta-potential of the particles (as modified by the coagulant dosing). This total charge is the property of the sample water that it is actually measured.

The usefulness of this model can be further increased by assuming that the mixing dynamics in the sample chamber are negligible. This is reasonable at the recommended flow rate of 4L/min, as the sample chamber volume is approximately 98mL, making the retention time in the sample chamber less than 1.5 seconds.

The dynamic situation can be modelled by:

$$dq/dt = -r_{off}.q(t) + r_{on}.u(t)$$

$$S(t) = k_{sensor}.q(t)$$

Where:

q is the total charge on the measurement surfaces (C – ie measured in Coulombs)

r_{on} is the rate with which charge moves onto the surfaces from the sample (s^{-1})

r_{off} is the rate at which charge moves off the surfaces into the sample (s^{-1})

u is the transferable charge of the sample in contact with the surfaces (C)

S is the measured SC (after FT or SR processing) (in arbitrary SC units)

k_{sensor} is the gain factor of the sensor and signal processing (SC/C)

Since none of these parameters can be measured other than SC, it is beneficial to rewrite the model solely in terms of SC.

$$dS/dt = k_{sensor} (-r_{off}.q(t) + r_{on}.u(t))$$

$$= r_{off}.(v(t)-S(t))$$

where

$$v(t) = r_{off}/(k_{sensor}.r_{on}) u(t)$$

= the steady state SC reading on the sample with transferable charge u(t)

This is the equation for a standard first order dynamic response. For example, after a step change in the true SC of the solution, v(t), the difference between it and the measured SC, S(t), follows the profile of $e^{-roff.t}$.

However, the actual measured step response is clearly not first order like this. Its initial rapid movement can only be modelled by a zero in the transfer function. At first glance it may seem that this is incorrect, as the only processes that occur in the system are mixing and absorption. They are both first order processes, which can be modelled as only poles that do not ‘store energy’, and so any cascaded combination of these processes will still have no zeros.

This effect can be explained by considering that the measuring surfaces are not uniform. It can be modelled by breaking the total surface into sections. Each section has separate rates of absorption and contributes to a part of the total charge measured. With the number subscript denoting section, using the same notation as above:

$$ds_1/dt = r_{off1}.(v(t)-s_1(t))$$

$ds_1/dt = r_{off2}.(v(t)-s_2(t))$

.

$ds_n/dt = r_{offn}.(v(t)-s_n(t))$

$S = f_1.s_1 + f_2.s_2 + \ldots. + f_n.s_n$

where

f_i = the proportion of the total contributed by section i

This is, of course, still a simplification, as there is can be a significant amount of mixing between the samples and exchanging of charge among the particles.

This can be readily converted to a state space model:

$$\bar{S}'(t) = \begin{bmatrix} -r_{off1} & 0 & \cdots & 0 \\ 0 & -r_{off2} & & 0 \\ \vdots & & \ddots & \\ 0 & 0 & & -r_{offn} \end{bmatrix} \bar{S}(t) + \begin{bmatrix} r_{off1} \\ r_{off2} \\ \vdots \\ r_{offn} \end{bmatrix} v(t)$$

where

$$\bar{S}(t) = \begin{bmatrix} s_1(t) \\ s_2(t) \\ \vdots \\ s_n(t) \end{bmatrix}$$ ***/ sense? X f

There is no need to model the dynamic response of individual physical sections of the piston, as only their absorption coefficients are important. In fact, it is still possible to adequately represent the observed responses by reducing the number of sections in the model to two, a relatively fast responding section and a slow one. This 2nd order version has poles at $-r_{off1}$ and $-r_{off2}$ and a zero at $-\frac{2r_{off1}r_{off2}}{r_{off1}+r_{off2}}$.

This response can normally be modelled as a single pole and zero, as the effect of the faster pole is obscured.

The relationship between coagulant dose and SC was investigated using the small-scale plant model described in section 3.6. The streaming current response is modelled further in section 5.2. Here it is also demonstrated that one pole and one zero are sufficient to model the dynamic response. The sample time used is 5 seconds, which adequately represents the measured dynamics.

The dynamics of the rapid mixer can be ignored, as it has a detention time that of approximately 1-2 seconds, based on its volume and flow rate. There is a transport delay of 3 sample steps (15 seconds) for the water to move through the pipes from the dosing point to the SCM's sample chamber. The steady-state gain of the coagulant to SC relationship is non-linear and depends on variable water parameters, however the nature of the dynamic response is relatively consistent. Under one particular set of raw water conditions the linear continuous model, which is the best (least squares) fit to the

recorded data, has a 15 second pure delay, a pole at -0.0161, a zero at -0.0389 and a constant of 0.0642.

4.10 Benefits of Using the FT Method

4.10.1 Comparison Simulation

The greater noise rejection of the FT method allows little post-filtering to be used in order to preserve the fast component of the SC sensor's response. This can be illustrated by the following simulation experiment.

A normalised step response is calculated from the example dynamic model given above (to make the continuous model realistic a fast extra pole is added at –0.5). A SC signal from the sensor is synthesised at one second intervals. The magnitude of the sine wave is the SC reading, Gaussian white noise of variance 0.05 is added and a 'hum' with a random frequency normally distributed around 50Hz (variance 1) with random phase and a magnitude normally distributed around 2 (variance 1). An example signal is shown in figure 4.x. This is representative of the signal from the sensor with a very significant, but realistic, level of interference.

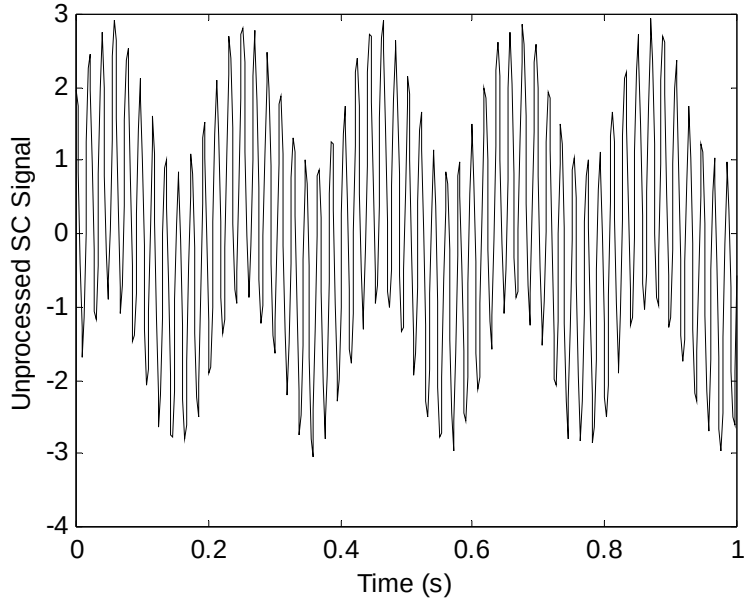

Figure 4.x. Example of unprocessed SC signal with significant interference.

Simulating the FT and SR methods over 1000 results in a mean squared error (MSE) between the true SC magnitude and the measured result of 1.134×10^{-5} for the FT and 1.52×10^{-6} for the SR result. Some post-filtering is necessary to improve the performance of the SR method.

A fifth order digital low-pass Butterworth filter with a cut-off at 0.05Hz (0.1 times the Nyquist frequency) results in a similar MSE when the signal magnitude is constant. An averaging filter over 24 seconds results in a somewhat better MSE when the signal magnitude is constant (very similar to averaging 120 strokes as described in section 4.7; it has similar attenuation at 50-60Hz to the FT method.) Table 4.1 shows the MSE for 600 seconds of a steady state signal with unit magnitude.

Table 4.1 Steady State Mean Squared Error of three signal processing methods

	FT Only	**SR + Butterworth**	**SR + Averaging**
Steady State MSE in simulation	0.126×10^{-4}	0.140×10^{-4}	0.061×10^{-4}

The Butterworth filter has a group delay of 10.2 sec and a maximally flat response in the pass-band, while the averaging filter has a group delay of 11.5 sec, no overshoot in the time domain and a poorer attenuation at higher frequencies. The results of each method on a simulated sensor step response are shown in figure 4.xx.

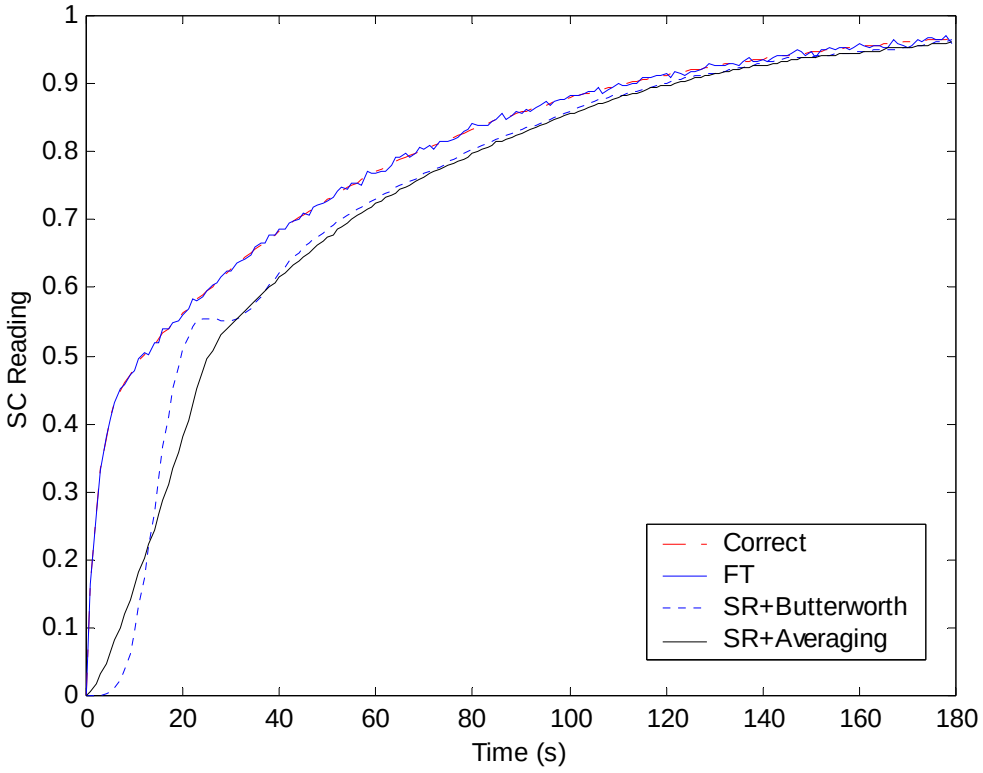

Figure 4.xx. Simulated instrument response to unit step.

The loss of bandwidth caused by the presence of the filtering has a significant effect on the controller design that is possible. For example the PID controller tuning described in section 5.10 is designed for the rapid set-point tracking and disturbance rejection required for online optimisation. It is not possible to tune a PID controller with the same level of performance with this post-filtering; the resulting system is unstable.

This indicates that the use of the FT signal processing method has significant benefits compared to the conventional SR method in situations where greater bandwidth and a rapid response to changes is required.

4.10.2 Practical Results

The Accurate Measurement Accufloc SCM implements the FT signal processing method and an adjustable rolling average filter in its signal processing electronics. This version of SCM was introduced in 2003, and since that time appears to have captured a majority of the NZ market for new SCMs. The unit has also been exported from NZ in small, but growing, quantities. [Hellier 2005].

It is worth noting that the SCM used with the optimiser in the previous chapter was an Accufloc. Without the greater bandwidth of the FT method it may not have been practical to use the jar testing optimiser with a PID controller as described in chapter 3. The response may have been too slow for the PID controller tuning to have a fast enough rise time while being non-oscillatory across all flow ranges. The quality of the more advanced control of the coagulant dose described in the following chapters benefits from the increased measurement bandwidth.

Meaningful comparing the dynamic responses of this SCM with one from a different manufacturer is difficult, because the physical design of the sensor and the condition of the piston significantly influence the signal generated. However several informal studies by plant operators who have evaluated an Accufloc SCM beside a conventional SCM indicated that, as expected, there is a relatively linear relationship between the steady-state readings and also that the Accufloc provides a noticeably quicker initial response to changes. This tends to suggest that conventional SCMs use a significant amount of post-filtering; therefore, the manufacturers expect that the level of noise rejection achieved by this would sometimes be necessary.

A quicker response also has the benefit of allowing plant operators to establish if the coagulant control loop is functioning correctly by examination of the readings in real time. This can be a considerable benefit in tuning, set up and fault finding because many plants do not have sufficient telemetry to allow these to be done from recorded data.

In addition, the greater distortion rejection of the FT Theoretically means that higher levels of degradation and contamination of the sensor surfaces can be tolerated without affecting the usefulness of the resulting measurement. However, this is extremely difficult to quantify, as it requires observations of a plant using the unit over an extended period (for example 1-2 years) and comparisons with a very similar plant using a unit of the same age with the old signal processing. Given the difficulty of measuring the effectiveness of coagulation and flocculation and the uniqueness of plants, such data is unlikely ever to be available.

4.11 Chapter Conclusions

The theoretical science of streaming current measurement available is inadequate to completely explain the observed phenomena, although there is a well known, but hard to model, relationship with zeta-potential.

It has been demonstrated how the use of the Fourier transform in SC signal processing is superior in both noise rejection and distortion rejection compared to the normal synchronous rectification method. The only advantage SR has over FT is that it can be more easily implemented.

The analysis of other components of the FT of the raw signal indicates that the relationship between the in-phase fundamental frequency component is quite constant as the piston and sensor develops a coating of floc and dirt. The other fundamental frequency component gradually changes under these conditions, and it is very likely that the harmonic components change as well. More experiments are required to understand this effect. A more complete understanding of the SC sensor and the signal it generates would allow a SC instrument to identify when it is not producing correct readings, thus removing one of the major problems with this technology.

A simple theoretical model of the dynamic response of the instrument can be developed by considering that the dynamic response of the complete sensor is the sum of the response of a number of independent first order parts, corresponding to different regions of the sensor. This provides a framework to explain the presence of a fast acting part of the step response. It is desirable to not filter out the fast part of the response, as it is a very useful indication of dosing performance and allows for better controller design.

Post-filtering which does not remove this fast part of the response but provides sufficient noise attenuation is very difficult SR. The level of performance achieved by the relatively advanced controllers examined in the follow chapter would not be possible with a conventional streaming current meter using the synchronous rectification signal processing method.