Detection of multiple oscillations in control loops

N.F. Thornhill[a,*], B. Huang[b], H. Zhang[c]

[a]Centre for Process Systems Engineering, Department of Electronic and Electrical Engineering, University College London,
Torrington Place, London WC1E 7JE, UK
[b]Department of Chemical and Materials Engineering, University of Alberta, Edmonton, Canada T6G 2G6
[c]Matrikon, Suite 1800, 10405 Jasper Avenue, Edmonton, Alberta, Canada T5J 3N4

Abstract

This article addresses the detection of oscillations in measurements from chemical processes including the case when two or more oscillations of different frequency are present simultaneously. The presence of oscillations in selected frequency ranges is determined using a new method based on the regularity of the zero crossings of filtered autocovariance functions. The work is motivated by and illustrated with industrial data that exhibit multiple plant-wide oscillations. Issues of practical implementation in an automated tool are discussed. © 2002 Elsevier Science Ltd. All rights reserved.

1. Introduction

Recent reviews [1,2] highlight the detection and diagnosis of plant-wide disturbances as a key issue facing the process industries. Oscillations are a common type of plant wide disturbance and the detection and diagnosis of oscillatory behavior in a chemical process is of importance because process variability has an impact on profit [3,4].

For industrial control consultants and engineers the detection of oscillating control loops is a basic activity. Features of interest are:

- Detection of the presence of one or more oscillations indicated by a regular pattern in the data;
- Determination of the periods of the oscillations, which can be used to infer the presence of a plant-wide oscillation;
- Determination of the magnitudes of the oscillations. A large magnitude oscillation needs to be addressed because it increases process variability while an oscillation of small magnitude may not be a cause for concern.

A tool to help the engineer should therefore automatically bring oscillatory loops to his or her attention, characterize them and highlight the presence of plant-wide oscillations.

The contribution of this paper is the detection of distributed oscillatory disturbances where the oscillation appears in several measurements. It utilizes zero-crossings of the autocovariance function (ACF) to detect multiple plant-wide oscillations and is illustrated using a case study in which many measurements in a plant participated in a plant-wide oscillation. It also deals with the case when oscillations at more than one frequency are present simultaneously in one measurement.

The following features of an oscillation were assessed from the autocovariance function:

- T_p, the period of oscillation;
- r, the regularity of the oscillation;
- $P\%$, the power of the oscillation.

The insights achieved were:

- detection of an oscillation if $r > 1$;
- determination of the number of independent oscillations present in a plant;

* Corresponding author. Tel.: +44-20-7679-3983; fax: +44-20-7388-9325.

E-mail address: n.thornhill@ee.ucl.ac.uk (N.F. Thornhill).

- detection of a plant-wide disturbance when many measurements in the plant have an oscillation of the same period;
- detection of two or more oscillations with different periods present simultaneously in one measurement.

The paper is laid out as follows. Section 2 describes detection and characterization of multiple oscillations. Section 3 introduces two case studies, one using experimental pilot plant data and the other using an industrial data set. Section 4 presents the results with a focus upon detection of multiple plant-wide oscillations. Section 5 addresses issues of practical implementation in an automated tool. The paper ends with a conclusions section.

2. Detection and characterization of oscillations

2.1. Background and context

2.1.1. Oscillation detection and diagnosis

Kedam [5] formulated a theory for the determination of oscillatory behavior from the zero crossings of a time series. Several authors have addressed the detection of oscillatory measurements in process data. Hägglund [6] detected zero crossings of the error signal in a control loop and calculated the integrated absolute error (IAE) between successive zero crossings. An oscillatory time trend has larger IAE values than a random one. Thornhill and Hägglund [7] inspected the regularity of the zero crossings of a time trend. The presence of an oscillation was inferred when the regularity was high and an estimate given of the signal to noise ratio of the oscillation. The method of Forsman and Statin [8] also used zero crossings and IAE. By distinguishing between positive and negative deviations they were able to detect asymmetrical oscillations.

The autocovariance function (ACF) can be used for oscillation detection. For example, if the absolute value of the ACF at the first minimum exceeds a threshold then the possibility of an oscillation is inferred. The patented method of Miao and Seborg [9] utilized additional cycles of the oscillatory autocovariance function. The authors demonstrated that the method could distinguish a decaying oscillation from a sustained oscillation.

Ettaleb et al. [10] considered diagnosis of oscillations in a cascade control loop and showed that the origin of an oscillation could be isolated to either the master or slave loop using a describing function analysis. Taha et al. [11] gave a diagnostic procedure based on valve characteristics while [7] distinguished limit cycle oscillation caused by non-linearity in a control loop from linear effects such as oscillatory tuning. Horch [12] showed that the covariance of the controller input and the process input would be an odd function in the presence of a

valve with hysteresis. Xia and Howell [13,14] gave an analysis of loop status that determined if an individual loop was affected by long or short term transients or a slow or fast oscillatory disturbance. Fast oscillatory disturbances were attributed to tuning or non-linearity. Thornhill et al [15] gave a non-linearity test capable of identifying the source of a plant-wide oscillation due to limit cycling and demonstrated its application in an industrial example.

It is necessary to find all the measurements or control loops in a plant having the same disturbance because the root cause will be among that group. Pryor [16] highlighted the usefulness of autocovariance functions and spectra for such a purpose. Harris et al. [17] reported plant-wide control loop assessment in which spectral analysis was useful.

2.2. Techniques

2.2.1. Use of zero crossings of the autocovariance function

The intervals between zero crossings of an oscillatory time trend are regular and can be exploited for off-line detection of oscillations. In earlier work it was found, however, that noise in the time trend caused spurious zero crossings [7].

The squares of the magnitudes of the discrete Fourier transform DFT give an estimate of the two-sided power spectrum (numerical issues related to the method are discussed in [18]). The autocovariance function (ACF) is the inverse Fourier transform of the two-sided power spectrum (the Wiener-Khinchin Theorem, [18]). The ACF of an oscillating signal is itself oscillatory with the same period as the oscillation in the time trend. The benefit the ACF for oscillation detection is that the impact of noise is reduced because white noise has an ACF that is theoretically zero for lags greater than zero. The pattern of zero crossings of the ACF therefore reveals the presence of an oscillation more clearly than the zero crossings of the time trend.

Oscillations may be present in the measurements from indicators, controlled variables (process variables), controller errors, controller outputs or in the set point. The proposed method can be used with any of these time trends. The case studies in this paper used process variables and measurements from indicators.

2.2.2. Data sets with more than one oscillation

The zero crossings of the ACF may not be regular if more than one oscillation is present. Fig. 1 shows the time trend and ACF from an industrial plant measurement having two superimposed oscillation of different periods. The bottom panel in the figure marks the positions of the zero crossings. The intervals between zero crossings of the autocovariance function reflect neither oscillation accurately because the zero crossings of the

fast and slow oscillations each destroy the regularity of the other's pattern.

The problem is addressed by frequency domain filtering. A filtered *ACF* is calculated from the inverse Fourier transform of the filtered power spectrum. Frequency domain filtering is preferred to a wavelet filter bank [19] because oscillatory time trends are global. There is thus no requirement for localization in the time domain while there is a need for precise location of the filter in the frequency domain.

The frequency-domain filter sets the power in unwanted frequency channels to zero. Because of the Nyquist sampling theorem, the upper half of the power spectrum is an aliased mirror image of the lower half. The alias of the *m*'th channel is channel $N - m + 2$, where N is the number of samples in the time trend (N is even in this formulation). The frequency domain filter operates on the required channels and also on the aliased channels.

The filter is an approximate realization of a Wiener filter [18]. A true Wiener filter also requires an estimate of the noise power within the wanted frequency channels which would then be subtracted from those channels. The trade-off for the simplification is that the power present in the wanted frequency band cannot be attributed solely to any oscillation present. The implications of this point are discussed later.

2.2.3. Assessment of period and regularity of oscillation

Each oscillation has two zero crossings and hence the intervals between zero crossings are:

$$\text{interval} = \frac{1}{2}\left(\overline{T_p} \pm \Delta T_p\right)$$

where $\overline{T_p}$ is the mean period and ΔT_p a random variation in the period. Thus $\overline{T_p}$ is twice the mean value of the intervals and the standard deviation of the period is $\sigma_{T_p} = 2 \times \sigma_{\text{intervals}}$.

The 10 intervals between the first eleven zero crossings are used for calculation of $\bar{T}_p$ and σ_{T_p}. The interval from lag zero up to the first zero crossing is excluded from the calculation because it corresponds to only one half of a completed deviation. Ten intervals are used because except in the case of a very persistent oscillation, the magnitudes of the autocovariance functions have decaying profiles and some spurious zero crossings may be found at large lags. Fewer intervals (e.g. eight) can be used for the regularity assessment when the period of oscillation is long but it is not recommended to use fewer than four intervals because the estimates of $\bar{T}_p$ and σ_{T_p} become unreliable.

An oscillation is considered to be regular if the standard deviation of the period is less than one third of the mean value. The statistic used is:
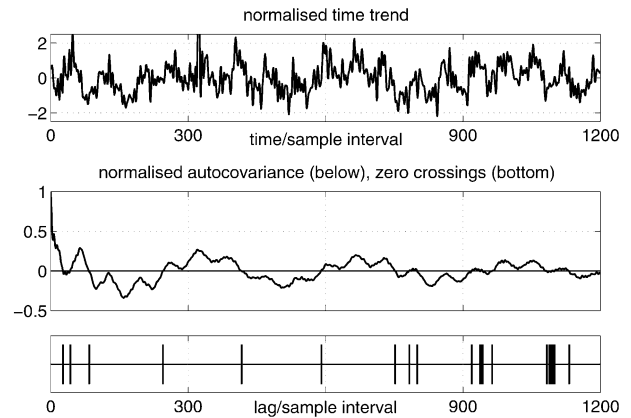

Fig. 1. An example of multiple oscillations present simultaneously (industrial tag 20).

$$r = \frac{1}{3} \times \frac{\overline{T_p}}{\sigma_{T_p}}$$

Values of $r > 1$ are taken to indicate a regular oscillation with a well defined period. The threshold value for r has a basis in statistics. If the zero crossings were random with an equal probability of arrival in each sampling interval then the intervals between zero crossings would have an exponential distribution [20]:

$$f(T_p) = \alpha e^{-\alpha T_p}$$

The standard deviation of an exponential distribution is equal to its mean. Therefore the null hypothesis of random arrivals would be $\overline{T_p} = \sigma_{T_p}$. The condition $r > 1$ rejects the hypothesis of random intervals when the observed value of the standard deviation is less than one third of its expected value under the null hypothesis, i.e:

$$\sigma_{T_p} < \frac{1}{3}\overline{T_p}$$

2.2.4. Assessment of the strength of an oscillation

The power spectrum is normalized before filtering, therefore $\sum P_i = 1$, where P_i is the power in the *i*'th frequency channel. Thus, when the spectrum is filtered it follows that the summed spectral power in the selected frequency channels gives the power as a fraction of the total. These power values include both oscillatory components and noise because the simplified Wiener filter does not subtract noise from the wanted frequency band. However, if there is a regular oscillation having $r > 1$ the summed power in the filtered spectrum would be dominated by the power of the oscillatory signal and may be utilized as an approximation to the power of the oscillation.

2.2.5. Comment—oscillation detection from power spectra

An oscillatory *ACF* with a period $\overline{T_p}$ has a spectral peak at $1/\overline{T_p}$ on the frequency axis. Visual inspection
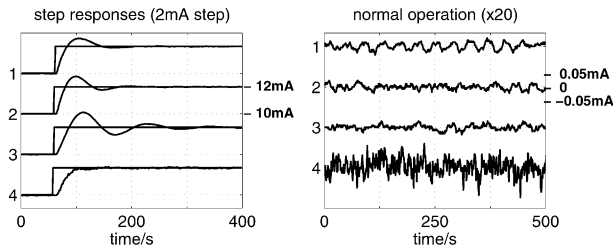
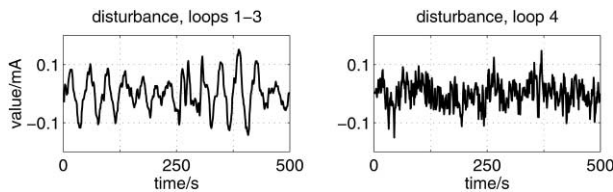Fig. 2. Step responses and normal operation of pilot plant tags 1–4.

Fig. 3. Time trends of pilot plant disturbances.

of spectra is helpful because strong peaks can be easily seen, but determination of period and regularity from the spectrum is not recommended. The ratio between the position of a peak and its bandwidth (known as the *Q*-factor) gives a measure of the regularity of the oscillation, but the presence of noise in the same frequency channels as are occupied by the spectral peak causes difficulties with the determination of bandwidth. The additional steps required to derive filtered *ACF*s are worthwhile because regular oscillations can be characterized even in the presence of noise.

3. Case studies

3.1. Data sets

3.1.1. Pilot plant data

Three time trends from a level control loop and one from a temperature loop in a university pilot plant were studied. The level loop had various tuning settings giving the step responses shown as tags 1–3 of Fig. 2. Tag 1 had a rapid response and a slight overshoot while tags 2 and 3 were slower and oscillatory with periods of about 90 and 130 s. Tag 4 was a temperature control loop. It was tuned for fast set point response and no overshoot. Fig. 2 also shows a magnified plot of mean centered data from normal running. The measurements were on a 4–20 mA scale and the sampling interval was 1 s.

The manipulated variable in the level control loop was a cold water flow which was subject to an oscillatory disturbance with a period of about 40 s that originated elsewhere in the building. The only disturbance to the temperature loop was random noise. Samples of these disturbances were captured during tests in which the cold water flow and the temperature were monitored at constant settings with no feedback control (Fig. 3).
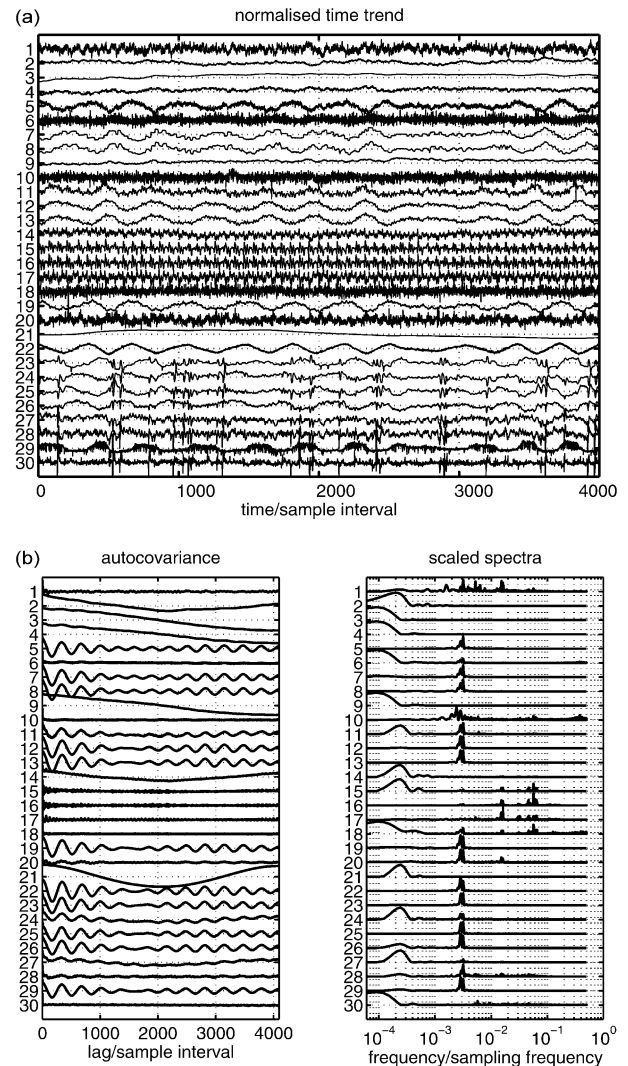

Fig. 4. (a) Time trends of 30 tags from an industrial data set normalised to the same standard deviation. (b) Autocovariance and power spectra of 30 tags from an industrial data set. Autocovariances are all unity at zero lag. Spectra show absolute magnitudes on a linear axis scaled to the same maximum peak height.

3.1.2. Industrial data

Industrial data tags 1–30 were from a data set provided courtesy of the Eastman Chemical Company, Kingsport, Tennessee. The time trends, autocovariances and power spectra are shown in Fig. 4(a) and (b). The sampling interval was 20 s and the total data record of 8192 samples represented 45.5 h of operation. The time trend plot shows just part of the data for better visualization. The frequency axis in the right hand panel of Fig. 4(b) is normalized to the sampling frequency. For example, 0.01 on the frequency axis represents 100 samples per cycle. The spectra stop at 0.5 on the frequency axis because oscillations must be sampled at least twice per cycle (i.e. the aliased part of the spectrum from 0.5 to 1.0 on the frequency axis is not shown).

Many time trends in Fig. 4(a) were oscillatory and the spectra in Fig. 4(b) display spectral peaks corresponding

to those oscillations. Some had more than one oscillation present. Therefore the challenge was the detection and accurate characterization of these oscillations and discovery of other oscillations not discernible by visual examination.

3.2. Design of the filters

3.2.1. Selection of the filters

The frequency ranges for the filters were selected by inspection of the power spectra. Automation of the filter will be discussed in Section 5.

Filters for the pilot plant data were designed to separate the oscillatory disturbance in the cold water flow with a period of 40 s from the oscillations due to tuning at 90–130 s. Their specifications were:

- 2–75 samples per cycle (0.0133–0.5 on the frequency axis);
- 75 samples per cycle and longer (d.c. to 0.0133 on the frequency axis).

The spectra of the industrial data showed one group of spectral peaks at about 0.003 on the frequency axis, another at about 0.015 to 0.016 and some others at higher frequencies. Therefore three filters were specified:

- 2–35 samples per cycle (0.0286–0.5 on the frequency axis);
- 35–100 samples per cycle (0.01–0.0286 on the frequency axis);
- 100–1000 samples per cycle (0.001–0.01 on the frequency axis).

The third filter excluded disturbances having 1000 samples or more per cycle. Thus non-stationery and long term deviations were removed. For instance, a temperature indicator may detect daily temperature variations. If these were not removed then a situation like that shown in Fig. 1 could arise in which a non-stationery trend would interfere with the detection of oscillations superimposed upon the trend.

3.2.2. Spurious results

Oscillation detection can generate spurious results when the filter boundaries are placed on the frequency axis such that they split a broad spectral feature across two ranges. The remedy for false detection is to ignore detected oscillations having periods that are close to the filter boundary.

Spurious results also arise if the filter width is too small. For instance, a filter with a width of just one frequency channel would give a pure sinusoidal output even if the input signal were random noise. Guidelines for the minimum filter width are given in Appendix A.

4. Results

4.1. Pilot plant data

4.1.1. Oscillation characterization

Fig. 5 shows the autocovariances of the filtered pilot plant data, the positions of the zero crossings and the filtered spectra. Features of interest are the intervals between the zero crossings and their regularity.

Table 1 lists the oscillation statistics. Results with $r > 1$ are regular oscillations and are highlighted in the table. Thus the diagnosis detected a fast oscillation in tag 1, fast and slow oscillations present in tags 2 and 3 and no oscillations present in tag 4.

4.1.2. Discussion

The period of the slow oscillations detected in the level loop (tags 2 and 3) changed when the tuning changed. For loop 2 the period was 88 s and for loop 3 it was 134 s. These results match the periods of oscillation determined from the step responses in Fig. 2 (90 and 130 s). An oscillatory step response means that the closed loop has a resonant peak in its frequency response. Therefore, the loop responds at that frequency when excited by random noise such as is always present, for example in the signal from the level sensor or random fluctuations in the cold water flow. The higher oscillation
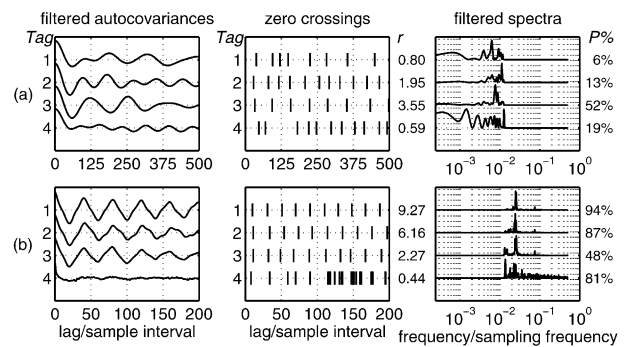

Fig. 5. Oscillation analysis of pilot plant tags: (a) low frequency (b) high frequency.

Table 1
Oscillation analysis for pilot plant tags

Tag	Period	r	Power (%)
Slow: ⩾75 samples/cycle			
1	–	0.80	–
2	88±15	1.95	13
3	134±13	3.55	52
4	–	0.59	–
Fast: 2–75 samples/cycle			
1	39±1	9.27	94
2	41±2	6.16	87
3	40±6	2.27	48
4	–	0.44	–

index *r* in tag 3 shows the regularity was higher in tag 3 which also had the most oscillatory step response.

The fast oscillations detected in tags 1–3 all had periods within one standard deviation of 40 s and therefore the disturbance was successfully detected in all three cases.

The period of the 40 s oscillation did not change when the loop tuning settings changed. However, the power in the oscillation altered with tuning settings. For example, for tag 2, 87% was associated with the disturbance and 13% with the slow oscillation while for tag 3, 48% was associated with the disturbance and 52% with the slow oscillation. A control loop like tag 2 tuned for a faster set point response has a larger closed loop bandwidth than one with a slower settling time such as tag 3. Therefore, tag 2 responded more to the fast disturbance than did tag 3.

The conclusion from the pilot plant data was that the new procedure was able to detect multiple oscillations in tags 2 and 3 and correctly reported that no oscillations were present in the non-oscillatory tag 4.

4.2. Industrial data

4.2.1. Autocovariances and zero crossings
Fig. 6 illustrates the analysis of industrial data in the range 100–1000 samples per cycle (0.001–0.01 on the

frequency axis). The spectra in the right hand panel of Fig. 6 were zero outside of that frequency range. The figure shows that tags with well defined spectral peaks have oscillatory autocovariance functions and a regular pattern of zero crossings. Analyses were also conducted in the range 35–100 and 2–35 samples per cycle. Table 2 lists the calculated oscillation statistics and highlights loops having significant oscillation with $r > 1$ and power in the oscillation above 1%.

The table of results shows, for instance, that the method detected slow and medium regular oscillations with periods of 312 and 64 samples per cycle in tag 20. Visual inspection of the spectrum [Fig. 4(b)] shows these oscillations as peaks at about 0.003 and just above 0.01 on the frequency axis. By contrast, when no filtering was used (Fig. 1) the analysis for tag 20 gave $r = 0.46$ and thus suggested no oscillation was present because the zero crossings for the medium oscillations interrupted the regularity of the zero crossings from the slow oscillation, and vice versa.

4.2.2. Oscillation characterization
Fig. 7 summarizes all the detected oscillations showing the oscillation index, *r*, on the horizontal axis and the period of the oscillation in samples per cycle on the vertical axis. The open circles near 35 and 100 samples
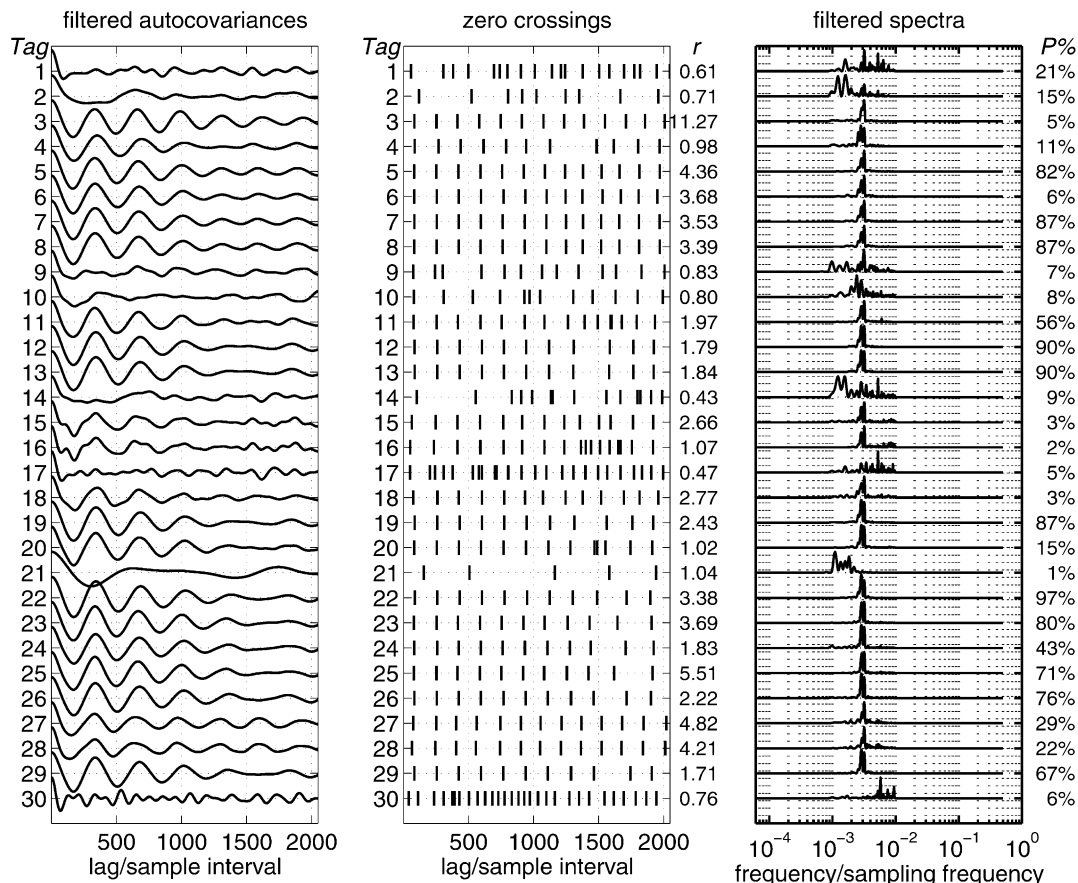

Fig. 6. Analysis of oscillations in the range 100–1000 samples per cycle in industrial data.

Table 2
Oscillation analysis for industrial data

Slow: 100–1000 sample/cycle				Medium: 35–100 sample/cycle				Fast: 2–35 sample/cycle			
Tag	Period	r	Power (%)	Tag	Period	r	Power (%)	Tag	Period	r	Power (%)
1	–	0.61	–	1	64±5.0	4.26	26	1	–	0.85	–
2	–	0.71	–	2	64±5.1	4.17	2	2	–	0.75	–
3	326±10	11.27	5	3	–	2.02	0	3	–	3.41	0
4	–	0.98	–	4	64±4.5	4.77	3	4	–	0.75	–
5	320±24	4.36	82	5	63±9.4	2.24	4	5	–	0.28	–
6	322±29	3.68	6	6	–	0.84	–	6	–	0.66	–
7	319±30	3.53	87	7	64±12	1.75	6	7	32.9±2.3	4.85	3
8	319±31	3.39	87	8	64±12	1.73	7	8	34.2±2.9	3.93	2
9	–	0.83	–	9	62±10	2.04	1	9	32.9±8.5	1.29	1
10	–	0.80	–	10	–	0.83	–	10	–	0.33	–
11	315±53	1.97	56	11	63±3.0	7.04	8	11	17.6±3.1	1.87	24
12	374±69	1.79	90	12	64±3.1	6.78	2	12	17.6±3.1	1.87	6
13	373±68	1.84	90	13	64±2.8	7.60	2	13	17.6±3.1	1.87	6
14	–	0.43	–	14	64±1.9	10.90	10	14	–	0.68	–
15	320±40	2.66	3	15	64±1.6	13.60	13	15	17.6±3.1	1.87	77
16	300±93	1.07	2	16	64±2.5	8.71	12	16	17.8±3.8	1.56	87
17	–	0.47	–	17	63±2.2	9.44	19	17	–	0.72	–
18	321±39	2.77	3	18	64±3.4	6.28	3	18	–	0.62	–
19	372±51	2.43	87	19	63±2.7	7.89	3	19	–	0.79	–
20	312±102	1.02	15	20	64±2.7	7.85	16	20	–	0.26	–
21	942±301	1.04	1	21	–	1.59	0	21	–	0.63	–
22	362±36	3.38	97	22	–	0.71	–	22	–	0.34	–
23	347±31	3.69	80	23	–	0.66	–	23	34.0±3.0	3.78	9
24	361±66	1.83	43	24	–	0.77	–	24	33.8±4.2	2.70	18
25	342±21	5.51	71	25	–	0.68	–	25	34.4±6.3	1.82	16
26	362±54	2.22	76	26	–	0.65	–	26	32.9±6.4	1.71	9
27	322±22	4.82	29	27	64±11	1.90	13	27	–	0.75	–
28	324±26	4.21	22	28	64±11	2.02	20	28	–	0.72	–
29	369±72	1.71	67	29	102±15	2.30	2	29	–	0.23	–
30	–	0.76	–	30	–	0.73	–	30	–	0.31	–

per cycle were on the filter boundaries and were disregarded, as discussed below. There were three plant-wide oscillations, as follows:

- Slow oscillation: tags {3, 5–8, 11–13, 15, 16, 18–20, 22–29}, average period 337 samples (112 min);
- Medium oscillation: tags {1, 2, 4, 5, 7–9, 11–20, 27, 28}, average period 64 samples or 21.3 min;
- Fast oscillation: tags {11–13, 15, 16}, average period 17.6 samples per cycle or 5.87 min.
- Tags {5, 7, 8, 18–20, 27, 28} had slow and medium oscillations simultaneously present and tags {11–13, 15 and 16} had simultaneous slow, medium and fast oscillations. Therefore the procedure successfully identified the presence of multiple oscillation, both plant-wide cases where the same oscillation was present in more than one tag and simultaneous oscillations where some tags had more than one oscillation.

4.2.3. False detection

An oscillation in tag 29 with a period of 102 samples per cycle was disregarded because it was near the filter boundary of 100 samples per cycle. Likewise, fast oscil-
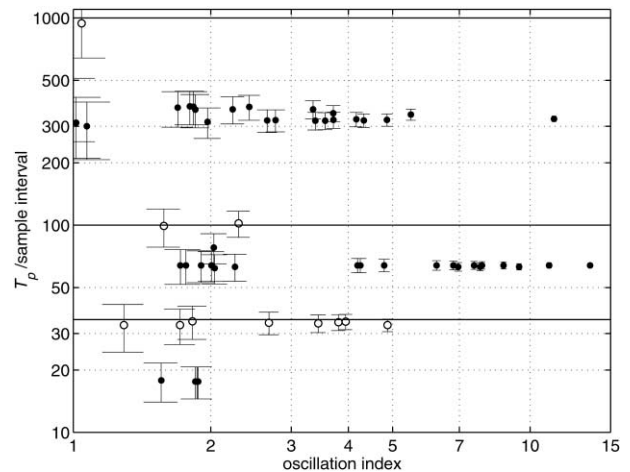

Fig. 7. Oscillations detected in the industrial data set. Open symbols are spurious oscillations lying on filter boundaries.

lations in tags {7–9, 23–26} with periods of 33 and 34 samples per cycle were ignored because they were close to the filter boundary of 35 samples per cycle. Such false oscillations arise when a tag has a broad spectral feature which falls away from the filter boundary. The trun-

cated feature resembles a spectral peak at the boundary and gives rise to an oscillatory autocovariance function.

A further oscillation analysis was run for the range 2–100 samples per cycle. No oscillations were detected at 33 or 34 samples per cycle, thus confirming that those were false oscillations. In tags 7 and 8, only the oscillations at 64 samples per cycle were detected which was the correct result. However, no oscillations were detected in tags {11–13, 15, 16} because the genuine oscillations at 64 samples per cycle and 17.6 samples per cycle interfered with one another in the manner illustrated by Fig. 1. It was concluded that in the presence of multiple oscillations it is better to use two filter ranges and to disregard oscillations with periods close to the filter boundary than to broaden the filter range.

4.2.4. Discussion

Tags {3, 5–8, 11–13, 15, 16, 18–20, 22–29} shared a common slow oscillation at 337 samples per cycle. The disturbance was a serious one because the signal power associated with that oscillation was high, more than 70% in 10 of the tags and 90% or more in three of them.

The plant-wide oscillation at 64 samples per cycle also affected many tags but the signal power was smaller, the highest being 26% in tag 1. The fast oscillation at 17.6 samples per cycle spread to five tags and the associated power was high in two of the five. The majority of the variability in the plant was therefore due to the slow oscillation at 337 samples per cycle and priority should be given to its analysis. The fast oscillation at 17.6 samples per cycle would be the second target.

4.2.5. Diagnosis

The methods described in this paper achieve the detection of tags affected by a plant-wide oscillation. Diagnosis of the root cause may be achieved using techniques described elsewhere [7,10–12,15]. A diagnosis reached on the basis of plant data can direct maintenance effort towards the likely root cause so that a definitive diagnosis using, for instance, a manual valve test can be reached promptly.

The first author together with staff of the Eastman Chemical Company identified tag 22 as the root cause of the oscillation at 337 samples per cycle using minimum variance controller performance benchmarking [21], *op-pv* maps (plots of controller output versus process variable), and non-linearity analysis [15].

4.2.6. Interpretation of the regularity, r

The regularity r of the 337 sample per cycle oscillation was highest in tags 3, 25 and 27. This does not imply, however, that one of these tags was the root cause of the disturbance because an oscillatory disturbance tends to become even more regular as it propagates through the process dynamics. For instance, a limit cycle oscillation generated by a non-linearity is periodic but not sinusoidal

and therefore contains harmonics of the fundamental frequency. The harmonics disappear as the limit cycle disturbance propagates through the plant because the plant provides low-pass mechanical filtering. Harmonics tend to reduce regularity because they slightly alter the positions of the zero crossings so an oscillation without harmonics is more regular than one with harmonics. Thus in the case when a limit cycle is the root cause the downstream tags may have oscillations with higher regularity than the root cause tag.

5. Automation

5.1. Automated filter selection

The filter ranges may be designed using the following algorithmic approach which can be automated. The concept is that if an oscillation is found in any tag then all tags are inspected in a narrower frequency range around the detected oscillation.

Step 1: pre-process data by filtering out non-stationery trends having fewer than eight complete cycles in the data set.

Step 2: apply oscillation detection to the pre-processed data. Record the periods of any oscillations detected.

Step 3: place filter boundaries mid way on the frequency axis between any detected oscillations, and mid way between the highest (lowest) frequency oscillation and the highest (lowest) frequency in the pre-processing filter.

Step 4: re-do the oscillation analysis for each new frequency range. Repeat step 3 for any frequency ranges where an oscillation at more than one frequency is present to uncover cases of multiple oscillations in the same tag. Stop sub-dividing when the range contains no multiple oscillations, or when the filters become too narrow.

5.2. Worked example using industrial data

Step 1: the filter used at step 1 for the 20 s industrial data was 2–1024 samples per cycle. Any oscillations up to 1024 samples per cycle would have at least eight zero crossings in the unaliased part (i.e. lags 0–4196) of the autocovariance function. The lower bound on the filter was 2 samples per cycle (the Nyquist criterion).

Step 2: the method for counting and grouping the detected oscillations given in appendix B counted the following oscillations:

- two tags, average $\bar{T}_p$ of 17.7 samples per cycle;
- 10 tags, average $\bar{T}_p$ of 343 samples per cycle;

Step 3: two oscillations were detected so the algorithm generated two new filters to refine the search:

- 4–34 samples per cycle;

- 34–514 samples per cycle.

The following oscillations were detected:

- seven tags, average $\bar{T}_p$ of 17.7 in filter 1;
- four tags, average $\bar{T}_p$ of 63 in filter 2;
- 17 tags, average $\bar{T}_p$ of 339 in filter 2.

Step 3: the above result for filter 1 was final because only one oscillation was detected in its range. Step 2 was repeated for filter 2 because there were two oscillations in its range. The final result for the sub-divided filter 2 was:

- 19 tags, average $\bar{T}_p$ of 64;
- 21 tags, average $\bar{T}_p$ of 337.

Therefore the automated approach detected the same oscillations as the manual approach adopted earlier plus two additional oscillations in tags 18 and 1 in the 17.7 samples per cycle group. An inspection of Fig. 4(b) shows that tags 18 and 1 do indeed have sharp peaks at 17.7 samples per cycle (0.056 on the frequency axis). The automated method detected them because the range of the filter was narrower (4–34 samples per cycle) than in the manual analysis (2–35 samples per cycle).

Visual inspection shows that tags 17 and 10 also have a spectral feature in the same location but the automated analysis shows that they are not sharp enough to be oscillatory.

6. Conclusions

The work reported in this paper has introduced a procedure for oscillation detection based on the zero crossings of the autocovariance function (*ACF*). The benefit of the use of the *ACF* is the reduction of spurious zero crossings caused by noise because random noise is much reduced in the *ACF* at lags greater than 0. The method can be automated.

The autocovariance functions were computed from the inverse Fourier transform of the power spectrum of normal operating data. That route to the *ACF* allowed a computationally simple method to determine the *ACF* of components of the time trend in selected frequency ranges. Therefore fast and slow oscillations were detected independently even when they were present simultaneously.

Oscillations were characterized by their period and regularity. The magnitude was estimated approximately from the power in the selected frequency range.

Experimental data from a university pilot plant were analyzed using the new methods. Simultaneous oscillations were correctly detected in measurements from control loops that had oscillatory tuning settings and an oscillatory disturbance. No oscillation was detected for another tag that had non-oscillatory tuning and a random disturbance.

Analysis of an industrial data set found tags where more than one oscillation was present simultaneously. Also, it detected plant-wide oscillations with the same period that were simultaneously present on several measurements. It was concluded that a slow plant-wide oscillation at 337 samples per cycle was responsible for the majority of the variability in the affected measurements because the spectral power associated with that oscillation was large.

Appendix A. Minimum width for a frequency domain filter

Consider a flat power spectrum of unit magnitude filtered by a two-sided filter of width Δf centered at $\pm f_o$. The *ACF* is the inverse transform of the filtered two-sided power spectrum $P(f)$:

$$
\begin{aligned}
ACF(\tau) &= \int_{-\infty}^{\infty} P(f)\, e^{2\pi j f \tau}\, df \\
&= \left(\int_{-f_o-\Delta f/2}^{-f_o+\Delta f/2} e^{2\pi j f \tau}\, df + \int_{f_o-\Delta f/2}^{f_o+\Delta f/2} e^{2\pi j f \tau}\, df \right) \\
&= 2 \times \cos(2\pi f_o \tau) \times \frac{\sin(2\pi\tau\Delta f/2)}{\pi\tau}
\end{aligned}
\tag{A1}
$$

If the filter is narrow then Δf is small and $ACF(\tau) \approx 2\Delta f \cos(2\pi f_o \tau)$. This *ACF* is oscillatory even through the power spectrum before filtering was flat with no spectral peaks. It has zero crossings when $\tau = (2k-1)/4f_o$ where k is a positive integer. The interval between the first and eleventh zero crossing is $5/f_o$. Eleven is the key number because regularity of the oscillation was determined in this paper from the ten intervals between the first eleven zero crossings.

It is required that the *ACF* of the filtered flat spectrum not be oscillatory. Therefore the filter must be wide enough that the zero crossings of the $\sin(2\pi\tau\Delta f/2)$ term in (A1) interfere with the regularity of the first eleven zero crossings of the $\cos(2\pi f_o \tau)$ term. The $\sin(2\pi\tau\Delta f/2)$ term has zero crossings when $\tau = k/\Delta f$.

Let the requirement be for at least two zero crossings of the $\sin(2\pi\tau\Delta f/2)$ function before $\tau = 5/f_o$. Therefore, the condition for the width of the filter is:

$$\frac{2}{\Delta f} \leqslant \frac{5}{f_o} \quad \text{or} \quad \Delta f \geqslant \frac{f_o}{2.5}$$

Appendix B. Grouping of detected oscillations

This Appendix shows the method used for automated counting and grouping of detected oscillations. It was adapted from a classification procedure by Chatfield and Collins [22].

Step 1: the periods of the significant oscillations (for instance, those with $r \geqslant 1.0$ and power above 1%) are arranged in a vector. Their standard deviations σ are arranged in another vector and their oscillation indexes, r, in a third. A fourth vector, w, of the same length as the others is initialized to all ones.

Step 2: scaled distances between the periods of oscillation are calculated. The scaled distances are:

$$d_{i,j} = \frac{|T_{p_i} - T_{p_j}|}{\max(\sigma_i, \sigma_j)}$$

The condition $d_{i,j} < 1$ indicates that two oscillations belong together because the difference between the periods of two oscillations is less than the larger of the two standard deviations.

The $d_{i,j}$ are arranged into a symmetric matrix. The row and column indexes (i and j) for the smallest of any $d_{i,j} < 1$ are determined. Only the upper off-diagonal elements are considered because the diagonal elements are zero and the lower off-diagonal elements are the same as the upper elements.

Step 3: the vectors are revised. T_{p_j} is deleted from the vector and the more regular of the two oscillations is recorded, i.e. T_{p_i} is kept if $r_i > r_j$ or replaced by T_{p_j} if $r_j > r_i$. Likewise, σ_j is deleted and σ_i is kept if $r_i > r_j$ or replaced by σ_j if $r_j > r_i$. In the r-vector, r_i is replaced by $\max(r_i, r_j)$ and r_j is deleted. In the w vector, w_i is replaced by $w_i + w_j$ and w_j is deleted. Thus the w vector keeps track of how many oscillations have been grouped.

Step 4: steps 2 and 3 are repeated until there are no more changes in the vector of T_p values (i.e. the $d_{i,j}$ are all greater than one). The final result is a short T_p vector giving the period of the most regular oscillation in each group of oscillations and a corresponding w vector giving the number of oscillations in the group.

An oscillation belongs to a group if its period is different from the period of the most regular oscillation in the group by less than one standard deviation. For instance, tag 25 with an oscillation period of 342 ± 21 belongs with tag 3, the most regular oscillation in the group with a period of 326 because 21 is greater than 342–326.

References

[1] S.J. Qin, Control performance monitoring—a review and assessment, Comput. Chem. Eng. 23 (1998) 173–186.

[2] L. Desborough, R. Miller, Increasing customer value of industrial control performance monitoring—Honeywell's experience. Preprints of CPC6, Tucson, AZ, 2001 p.153–186.

[3] G.D. Martin, L.E. Turpin, R.P. Cline, Estimating control function benefits, Hydrocarbon Processing (June) (1991) 68–73.

[4] J.P. Shunta, Achieving world class manufacturing through process control, Prentice Hall, NJ, 1995.

[5] B. Kedam, Time series analysis by higher order zero crossings, IEEE Press, New York, 1993.

[6] T. Hägglund, A control-loop performance monitor, Control Eng. Practice 3 (1995) 1543–1551.

[7] N.F. Thornhill, T. Hägglund, Detection and diagnosis of oscillation in control loops, Control Engng. Prac. 5 (1997) 1343–1354.

[8] K. Forsman, A. Stattin, A new criterion for detecting oscillations in control loops, European Control Conf., Karlsruhe, Germany, 1999.

[9] T. Miao, D.E. Seborg, Automatic detection of excessively oscillatory feedback control loops, in: IEEE Conf. Control. Appl., Hawaii, 1999, pp. 359–364.

[10] L. Ettaleb, M.S. Davies, G.A. Dumont, E.E. Kwok, Monitoring oscillations in multiloop systems, in: IEEE Int. Conf. Cont. Appl., 1996, pp. 859–863.

[11] O. Taha, G.A. Dumont, M.S. Davies, Detection and diagnosis of oscillations in control loops, in: IEEE Conf. Decision and Control, Kobe, Japan, 1996, pp. 2432–2437.

[12] A. Horch, A simple method for detection of stiction in control valves, Control. Engng. Prac 7 (1999) 1221–1231.

[13] C. Xia, J. Howell, Controller output based, single number statistics for control loop status monitoring, in: CHEMFAS4, Korea, 7–8 June 2001.

[14] C. Xia, J. Howell, Loop status statistics, in: CHEMFAS4, Korea, 7–8 June 2001.

[15] N.F. Thornhill, S.L. Shah, B. Huang, Detection of distributed oscillations and root-cause diagnosis, in: *CHEMFAS4*, Cheju Island, Korea, 7–8 July 2001.

[16] C. Pryor, Autocovariance and power spectrum analysis, Control Engineering (Oct) (1982) 103–106.

[17] T.J. Harris, C.T. Seppala, P.J. Jofreit, B.W. Surgenor, Plant-wide feedback control performance assessment using an expert system framework, Control Engng. Practice 4 (1996) 1297–1303.

[18] W.H. Press, B.P. Flannery, S.A. Teukolsky, W.T. Vetterling, Numerical recipes, Cambridge University Press, Cambridge, 1986.

[19] G. Strang, T. Nguyen, Wavelets and filter banks, Wellesley-Cambridge Press, Wellesley MA, USA, 1996.

[20] J.E. Freund, Mathematical Statistics, second ed., Prentice Hall, NJ, 1971.

[21] L. Desborough, T. Harris, Performance assessment measures for univariate feedback control, Can. J. Chem. Eng. 70 (1992) 1186–1197.

[22] C. Chatfield, A.J. Collins, Introduction to Multivariate Analysis, Chapman and Hall, London, UK, 1980.