Mechanical Systems and Signal Processing 23 (2009) 987–1036

Contents lists available at ScienceDirect

Mechanical Systems and Signal Processing

journal homepage: www.elsevier.com/locate/jnlabr/ymssp

Invited Review

Cyclostationarity by examples

Jérôme Antoni *

Laboratory Roberval of Mechanics (UMR CNRS 6253), University of Technology of Compiègne, 60205 Compiègne Cédex, France

A R T I C L E I N F O

Article history:
Received 3 May 2008
Received in revised form
15 October 2008
Accepted 29 October 2008
Available online 14 November 2008

Keywords:
Cyclostationary signals
Cyclostationarity
Cyclic power
Cyclic modulation spectrum
Spectral correlation density
Spectral coherence density
Wigner-Ville spectrum
Pure and impure cyclostationarity
Cyclostationary modelling
Machine diagnostics
System identification
Source separation.

A B S T R A C T

This paper is a tutorial on cyclostationarity oriented towards mechanical applications. The approach is voluntarily intuitive and accessible to neophytes. It thrives on 20 examples devoted to illustrating key concepts on actual mechanical signals and demonstrating how cyclostationarity can be taken advantage of in machine diagnostics, identification of mechanical systems and separation of mechanical sources.

Contents

* Tel.: +33 3 44 23 45 25; fax: +33 3 44 23 44 77.
E-mail address: antoni@utc.fr

0888-3270/$ - see front matter
doi:10.1016/j.ymssp.2008.10.010

1. Introduction

1.1. Preliminaries

Cyclostationarity pertains to the new wave of signal processing techniques that is currently revolutionising the field of mechanical signature analysis. Briefly stated, a cyclostationary signal is one that exhibits some hidden periodicity of its energy flow. Such a definition includes a broad class of mechanical signals such as vibrations or acoustic measurements. This is so because many mechanical systems—e.g., reciprocating mechanisms, gears, fans, electrical motors—sustain periodic motion of their components which in turn periodically modulate the vibration or noise they radiate. The cyclostationary property covers a rich statistical typology of signals (including periodic signals and stationary random signals as particular cases) which is particularly appealing since the same "cyclostationary" formalism will apply generally and independently of the object of study.

However, the formalism of cyclostationary signals may seem difficult to grasp at first because it is quite different from the stationary logic which most of us have been trained to work with. It is the aim of this paper to introduce cyclostationarity from an intuitive approach that proceeds essentially from generalising our common experience gained from stationary signals. In particular, because it is intuitive and physically meaningful, the approach relies as much as possible on the key concept of the *energy* conveyed by a signal: this will make easier the introduction of the various types of

Nomenclature

t	time variable
τ	time-lag variable
f	(spectral) frequency variable
α	cyclic frequency variable
$\gamma_x^\alpha(f)$	spectral coherence density
$h(t; f, \Delta f)$	impulse response of a band-pass filter of bandwidth Δf and central frequency f
$H(\nu; f, \Delta f)$	frequency response of a band-pass filter of bandwidth Δf and central frequency f
P_x	signal power
P_x^α	cyclic power
$P_x(t)$	mean instantaneous power
$P_x(f)$	power spectral density
$P_x(t; f, \Delta f)$	instantaneous power spectrum
$P_x^\alpha(f; \Delta f)$	cyclic modulation spectrum
$\mathscr{P}_0\{\bullet\}$	time-averaging operator for extracting a constant value
$\mathscr{P}_\alpha\{\bullet\}$	time-averaging operator for extracting a periodic component with frequency α
$\mathscr{P}\{\bullet\}$	time-averaging operator for extracting all periodic components
$\mathscr{R}\{\bullet\}$	residual operator
$R_x(t, \tau)$	instantaneous autocorrelation function
$R_x^\alpha(\tau)$	cyclic autocorrelation function
$SC_x^\alpha(f)$	spectral correlation density
$x_{\Delta f}(t; f)$	filtered signal in the frequency band $[f-\Delta f/2; f+\Delta f/2]$
$WV_x(t, f)$	Wigner-Ville spectrum

densities that characterise cyclostationary signals in the time, the frequency and the cyclic domains. Moreover, it was decided not to make any use of the mathematical expectation operator nor of the Dirac distribution, which although useful for simplifying the mathematics, both require an additional abstraction effort that hinders intuition. Finally, a special effort has been made to illustrate every concept on real-world signals, mainly coming from acoustic and vibration measurements: this will provide the reader with another route to comprehend the subject. These precautions should make the paper accessible to most readers, and in particular to engineers or young researchers who are looking for a first introduction to cyclostationarity.

1.2. The antinomy of nonstationarity

Why should one invest so much effort in going beyond the traditional and well-established "stationary" approach, and consider an alternative such as cyclostationarity? It all relies on the fact that stationarity is more a matter of convenience than of actuality.

By definition, stationary signals are representative of physical phenomena that maintain a constant statistical behaviour in time. Yet, this property is hardly met by mechanical systems consisting of some machinery or rotating parts which, by nature, undergo a nonstationary operation. Even under constant operating conditions (speed, torque, temperature), a succession of phenomena usually takes place within the machine cycle so as to release energy on a rhythmic basis: meshing of teeth in gears, combustion of gas in internal combustion (IC) engines, inversion of forces in reciprocating or cam mechanisms, admission and exhaust of fluids in pumps, turbulence around fan blades, and so on. Such phenomena typically produce transient signatures in mechanical signals, which in turn are likely to carry critical information on the operating condition of the machine. Let us insist on the assertion that nonstationarity—as evidenced by the presence of transients—is intimately related to the concept of *information*. This is completely analogous to speech or music signals that can carry a message or a melody onlt because they consist of a succession of nonstationarities. Similarly, the identification of noise sources in an IC engine requires tracking their chronology and localising their temporal occurrences in the engine cycle.

However, when brief in time, transients are extremely difficult to track inside the machine cycle. For many decades the traditional approach has been to regard mechanical signals as if they were stationary—i.e., exempt from local information—because signal processing was not developed enough to proceed otherwise. Although time–frequency techniques have now become quite popular in industry, they are mainly "analysis" tools—as opposed to "processing" tools—and in any case they are unable to propose a versatile methodology that applies to all mechanical signals in all circumstances. This is because nonstationarity is intrinsically a non-property, defined by opposition to stationarity. For all these reasons, it is often reassuring to rely on the stationary *assumption* and to benefit from the many "on-the-shelf" tools offered by that framework. Because it is a choice by default, it is also a poor choice that deprives the user of the information conveyed by nonstationarities. This is what is referred here to as the "antinomy of nonstationarity".

Cyclostationarity comes to the fore at this juncture. Because it defines a certain type of nonstationarity, it is a *property* endowed with a rigorous framework of signal processing tools liable to apply to a broad class of mechanical signals.

1.3. State of the art

The construction of the theory of cyclostationarity is closely linked to that of modern signal processing. Nowadays the subject has probably reached its state of maturity, as evidenced by the two excellent bibliographies in Refs. [24,25], and the very complete states of the art in Refs. [17,23].

Pioneering works in cyclostationarity date back to the early sixties, but it is truly since the eighties that cyclostationarity has become a subject of active research. Most of the precursory works pertain to the field of communications, where it was recognised that the process of modulating a signal for Hertzian transmission naturally led to a cyclostationary behaviour. A great tribute is due to William A. Gardner who first established many of the theoretical foundations, laid down the currently used terminology, and also foresaw many applications. As a matter of fact, W. A. Gardner was probably the first to recognise that the cyclostationary framework is appropriate for any physical phenomenon that gives rise to data with periodic statistical characteristics: "in mechanical-vibration monitoring and diagnosis for machinery, periodicity arises from rotation, revolution, and reciprocating of gears, belts, chains, shafts, propellers, bearings, pistons, and so on" [10].

Despite this encouraging perspective, actual contributions of cyclostationarity to mechanics have remained very limited. The bibliographical compilation [24] listed 52 references related to mechanics out of a total of 1559 until 2005, and [25] listed 29 out of 786 until the same year. In both cases, this is no more that three per thousand. Most of these references—plus some others—are listed in our reference section. They will be referred to throughout the text and in particular in the last part of the paper concerned with cyclostationary applications.

1.4. Organisation of the paper

Writing a tutorial paper on cyclostationarity is a difficult task that requires finding the right balance between the temptation to provide a maximum of knowledgeable material on the subject and the necessity to follow a pedagogical—and hence simplistic—presentation. Therefore, the mathematical level was purposely kept as low as possible for the benefit of a more intuitive presentation of the key concepts that are underlying the theory of cyclostationary processes. It is hoped that this approach will facilitate the interested reader to persevere with more theoretical presentations of the subject such as [8,13,23], as well as providing him/her with the prerequisites to foresee potential applications in his/her field of interest. The key concepts that have been deemed essential towards that purpose are:

- the concept of hidden periodicities in a random signal,
- the equivalence between nonstationarity in the time domain and correlation in the frequency domain,
- the interpretation of cyclic and spectral frequencies as modulation and carrier frequencies, respectively,
- the interrelations between temporal, spectral and cyclic decomposition of the energy displayed by a signal,
- the distinction between first, second, and higher-order cyclostationarity,
- the implication of the uncertainty principle when choosing between various cyclostationary tools.

The paper is constructed around two questions and an assertion, each of which deserve a different part. The first part addresses the question "What is cyclostationarity?" and consists of the aforementioned key concepts. Whenever possible, these are introduced from first principles and illustrated on real-world examples. The second part of the paper addresses the question "How to implement cyclostationary tools?" and introduces some general guidelines as to how to construct a cyclostationary-based estimator and make it reliable. Finally the last part of the paper illustrates the assertion "Don't ignore cyclostationarity: use it to advantage" by means of several case studies concerned with the diagnostics and identification of mechanical systems, and the separation of mechanical sources.

2. What is cyclostationarity?

Two examples of real-world signals are introduced from the onset which will provide a common basis to illustrate the various concepts to be discussed later in this section.

2.1. Two introductory examples

2.1.1. First example

Fig. 1(a) displays the time history of a vibration signal captured on a 45 kW centrifugal pump operating at 2950 rpm with a delivery of 250 m^3/h. Sampling frequency is 50 kHz. From a first visual inspection, the vibration signal looks typically random and stationary. However, due to the operation of the eight blades, modulation of the vibrations by the blade-pass frequency is expected to occur. Such a hidden periodicity, if it exists, seems impossible to detect at this stage without further processing.

2.1.2. Second example

Fig. 1(b) and (c) display the noise radiated 1 m away by a 1.5 L four-cylinder diesel engine together with one of the four-cylinder pressure signals that partly produces it. The operating speed is 850 rpm and the sampling frequency of 20480 Hz. The acoustical signal is filtered in the 4–8 kHz octave band because this is where noise quality is of concern. Contrary to the previous example, the periodic modulation of the signal energy is clearly visible here. Indeed, while the cylinder pressure signal is essentially periodic with period equal to the engine cycle, the acoustical signal is more random in nature with

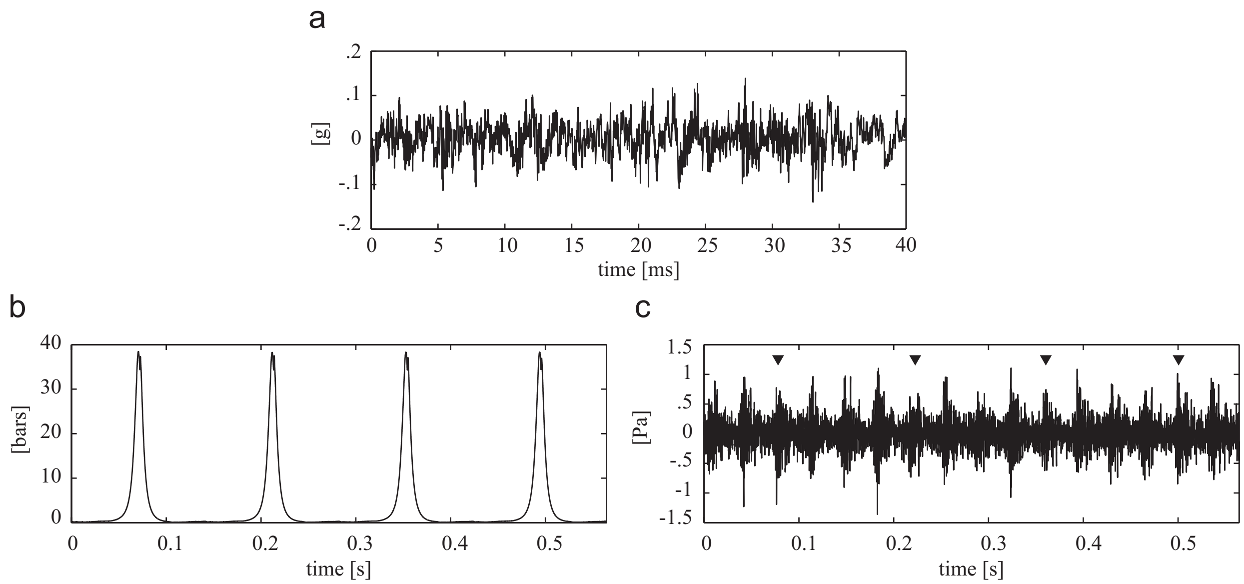

Fig. 1. (a) Vibration signal from a centrifugal pump, (b) cylinder pressure signal and (c) corresponding acoustical signal of a four-cylinder diesel engine. The sudden pressure rises in (b) are indicative of the firings in one of the four cylinders; the corresponding radiated noise is marked by arrows in (c). The other bursts of energy in (c) are due to the noise radiated by the other three cylinders whose pressure signals are not shown here.

repetitive bursts of energy synchronised on the cylinder pressure signal; since this is a four-cylinder engine, there are four bursts per cycle corresponding to the noise radiated by the explosions taking place in each cylinder. Of great interest here is to characterize the mechanism giving rise to such strong a periodical pattern in a random signal.

2.2. In search of hidden periodicities

The vibration of the pump and the noise of the diesel engine illustrated above are two typical examples of mechanical signals generated by periodic mechanisms. However, as exemplified by the pump signal, any periodic behaviour may be hard to distinguish from the visual inspection of the temporal waveforms. As we shall see now, conventional spectral analysis may not help either in the endeavour to reveal hidden periodicities when the signal of interest is mainly random in nature. Yet, one solution to the search for hidden periodicities will turn out to investigate how the signal energy—rather than the amplitude—is propagating with time.

2.2.1. Linear decomposition of energy flow: extraction of constant trends

In order to proceed from known results, let us first review the basics of stationary random signals. By definition, a *random* signal cannot be described by an equation or, equivalently, cannot be predicted exactly. The only way to describe it is by means of some *statistical* laws. For a *stationary* random signal, such statistical laws are constant with respect to time. For instance, the mean value—the direct current (DC) component—of a stationary random signal indicates the amplitude around which it oscillates on the average. For vibration signals issuing from speed or acceleration measurements the mean value is always zero unless the centre of gravity of the object under study experiences some motion (this further requires that the velocity or acceleration transducer transmits DC). The same observation holds for acoustic signals which correspond to a variation around static pressure. Of great interest therefore are the oscillations around the mean value—the alternating current (AC) component—the intensity of which is returned by the mean-square value (or its square-root, the RMS value); this is an indication of the mean rate of energy released by the signal, that is its "power". Power is a constant for all stationary signals. The following definition specifies how the mean value and the power are extracted from a stationary random signal.

Definition 1. Let us name as $\mathscr{P}_0$ the operator that extracts the time-averaged value (DC component) of a signal. The time-averaged value m_x and time-averaged power P_x of a signal $x(t)$ are then computed as

$$m_x = \mathscr{P}_0\{x(t)\} = \lim_{T\to\infty}\frac{1}{T}\int_T x(t)\,\mathrm{d}t,$$
$$P_x = \mathscr{P}_0\{|x(t)|^2\} = \lim_{T\to\infty}\frac{1}{T}\int_T |x(t)|^2\,\mathrm{d}t, \qquad (1)$$

where $\int_T(\cdot)\mathrm{d}t$ means the summation over an interval of length T. The reason of the notation $\mathscr{P}_0$ will become clear later.

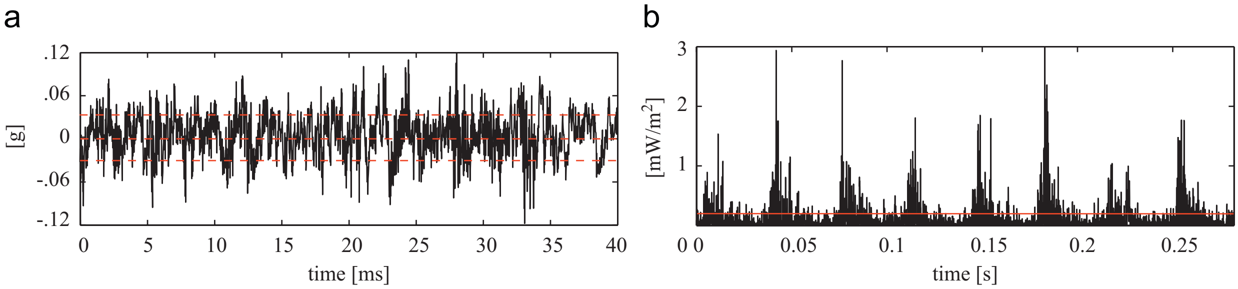

Fig. 2. (a) Vibration signal from the pump together with its time-averaged value m_x plus and minus the square-root of its time-averaged power P_x and (b) squared acoustical signal from the diesel engine together with its time-averaged power, both scaled by the air characteristic impedance ρc_0 so as to measure the radiated sound intensity.

Example 1. Fig. 2(a) displays the time-averaged value m_x plus and minus the square-root of the time-averaged power P_x computed on the pump signal. This provides a crude indication of the vibration levels of the pump by means of a "position" and a "dispersion" parameter, provided the signal does not depart too much from the assumption of stationarity.

The time-averaged power P_x of the acoustical signal from the diesel engine is shown in Fig. 2(b). In this case, proper normalisation of P_x by the air characteristic impedance ρc_0—the product of the fluid density and the sound velocity—has the physical interpretation of a "sound intensity"—i.e., the rate of energy per unit of time and of area expressed in W/m^2—radiated by the engine at the measurement point. Clearly, the stationary assumption is irrelevant here since a constant sound intensity is unable to account for the series of explosions the radiated noise consists of.

Relying only on the time-averaged value m_x and the time-averaged power P_x for characterising the statistical behaviour of a stationary random signal is surely insufficient in many instances. In particular, it is well-known that the complex structure of mechanical signals is best revealed in the frequency domain. Briefly stated, the frequency analysis of a stationary random signal consists in measuring the average flow of the energy it conveys through an infinitely narrow frequency band. Namely, let $x_{\Delta f}(t;f)$ be the filtered version of signal $x(t)$ through a frequency band of width Δf centred on frequency f. Using the previously defined $\mathscr{P}_0$-operator, the average energy flowing in that frequency band is returned by $P_x(f;\Delta f) = \mathscr{P}_0\{|x_{\Delta f}(t;f)|^2\}$. Clearly, for a random signal with no pure tone at frequency f this quantity tends to zero with the bandwidth Δf of the filter. A density $P_x(f)$ of energy flow per hertz can be measured, however, by taking the limit of the ratio $P_x(f, \Delta f)/\Delta f$:

$$P_x(f) = \lim_{\Delta f\to 0} \frac{P_x(f;\Delta f)}{\Delta f} = \lim_{\Delta f\to 0}\lim_{T\to\infty} \frac{1}{T\cdot\Delta f}\int_T |x_{\Delta f}(t;f)|^2\,\mathrm{d}t, \tag{2}$$

where the order of the two limits cannot be interchanged. This quantity defines the power spectral density (PSD) and it is to be understood as a representation of how the energy conveyed by the signal is distributed in the frequency domain—see Fig. 3). For instance, if $x(t)$ is the vibration velocity measured at a certain point of a structure, the vibration power at that point is $P_x(f)$ times twice the real part of the structural impedance thereto.

As well-known, the PSD of a random signal which does not contain any periodic waveform will be a (piecewise) *continuous* function of frequency. On the other hand, the presence of a pure tone of amplitude A at frequency f will produce a peak at that frequency with infinite magnitude $\lim_{\Delta f\to 0} A^2/(4\Delta f)$ indicating that an unbounded quantity of energy is necessary to produce it. Hence, checking for the presence of peaks in the PSD is the common practice to investigate whether the signal of interest contains periodic waveforms or whether it is purely random.

Example 2. Fig. 4 displays the three power spectral densities measured on the signals of Fig. 1. It is seen that both the vibration signal from the pump and the acoustical signal from the diesel engine have their power continuously distributed in the frequency domain, thus indicating they are essentially random in nature. This is in contrast with the "discrete" PSD of the cylinder pressure signal where energy is fully concentrated at the engine cycle frequency at 7 Hz and all its multiples, which is strongly indicative of a periodic behaviour in time. Moreover, the comparison of the power spectra in Fig. 4(b) and (c) suggests that there is little hope for the radiated noise to comprise the same periodic components as the cylinder pressure—essentially spreading over the [0–2 kHz] range—since all its frequency components below 2.5 kHz have been filtered out.

2.2.2. Cyclic decomposition of energy flow: extraction of cyclic trends

The former examples have evidenced that the stationary approach is unable to reveal the rhythmic behaviour of cyclostationary signals. This must be so by construction, since the averaging operation embodied by the $\mathscr{P}_0$-operator smears all time structures in the signals.

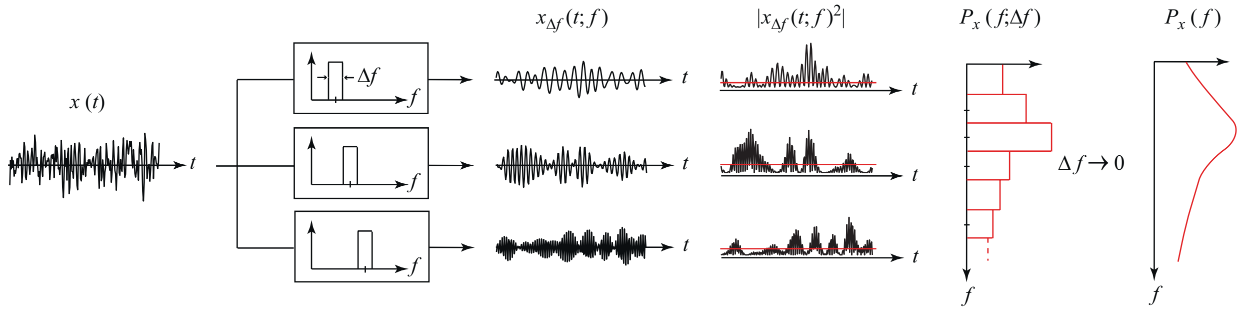

Fig. 3. Interpretation of the power spectral density as the averaged power measured at the output of a filterbank.

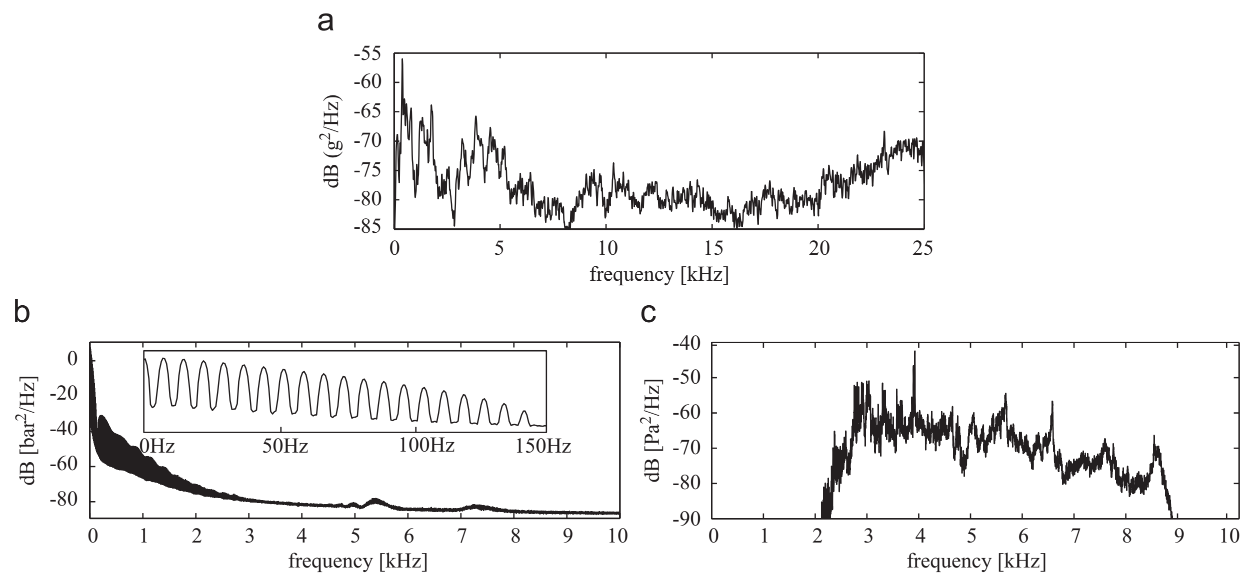

Fig. 4. (a) Power spectral density of the vibration signal from the pump ($\Delta f = 4.6$ Hz; $T \cdot \Delta f = 2$), (b) power spectral density of the cylinder pressure and (c) of the acoustical signal of the diesel engine ($\Delta f = 1.8$ Hz; $T \cdot \Delta f = 57$).

In order to remedy this shortcoming, the central idea consists of decomposing the energy flow not only into a constant trend, but also into periodic components that can depict how the energy is travelling with time. For that purpose, let us now introduce a new extraction operator.

Definition 2. Let us name as $\mathscr{P}_\alpha$ the operator that extracts the periodic component at frequency α in a time function. Specifically,

$$\mathscr{P}_\alpha\{\cdot\} = \left(\lim_{T\to\infty}\frac{1}{T}\int_T(\ \bullet)e^{-j2\pi\alpha t}\,\mathrm{d}t\right)\cdot e^{j2\pi\alpha t} = \mathscr{P}_0\{(\ \bullet)e^{-j2\pi\alpha t}\}\cdot e^{j2\pi\alpha t}. \tag{3}$$

Frequency α is commonly known as the *cyclic frequency* of the signal, and its inverse as the *cycle*.

Note that the $\mathscr{P}_\alpha$-operator first computes the Fourier coefficient $\mathscr{P}_0\{(\ \bullet)e^{-j2\pi\alpha t}\}$ at frequency α and then assigns it to that periodic component $\exp(j2\pi\alpha t)$ in order to reconstruct a pure sinusoidal signal. Note also that it reduces to the $\mathscr{P}_0$-operator formerly introduced in Definition 1 when $\alpha = 0$. In order to extract all periodic components contained in a time function, one needs to apply the $\mathscr{P}_\alpha$-operator over all candidate frequencies α contained in some set $\mathscr{A}$. This defines our last operator:

Definition 3. Let us name as $\mathscr{P}$ the operator that extracts all periodic components in a time function. Specifically,

$$\mathscr{P}\{\ \bullet\} = \sum_{\alpha\in\mathscr{A}}\mathscr{P}_\alpha\{\ \bullet\} \tag{4}$$

where the set $\mathscr{A}$ contains all cyclic frequencies α associated with non-zero periodic components.

Stated differently so as to stress the resemblance with the letter "P", the $\mathscr{P}$-operator extracts the perfectly predicable part of a time function, i.e. that part that can be fully described by an equation—here a Fourier series—and whose values can thus be computed everywhere on the time axis, including arbitrarily far future or past, in contrast to random signal (more will be said on this subject in Section 2.3). These definitions will now turn out very useful to introduce more advanced signal processing tools.

2.2.2.1. The mean instantaneous power. When applied to the energy flow $|x(t)^2|$ per unit of time at each instant t of a cyclostationary signal, the $\mathscr{P}$-operator reveals very well the repetitive bursts of energy that characterise the presence of a periodic mechanism. The resulting quantity

$$P_x(t) = \mathscr{P}\{|x(t)|^2\} \tag{5}$$

is referred to as the mean instantaneous power. In particular if the flux of energy is suspected to vary periodically, then $P_x(t)$ will have a Fourier series decomposition with non-zero coefficients, say $P_x{}^\alpha$, which we shall coin the *cyclic powers* of the signal at cyclic frequencies α; specifically,

$$P_x(t) = \sum_{\alpha\in\mathscr{A}} P_x^\alpha \cdot e^{j2\pi\alpha t}, \tag{6}$$

where the cyclic powers are obtained by use of the $\mathscr{P}_0$-operator:

$$P_x^\alpha = \mathscr{P}_0\{|x(t)|^2 \cdot e^{-j2\pi\alpha t}\}. \tag{7}$$

Example 3. Fig. 5(a) shows the square-root of the mean instantaneous power $P_x(t)$ of the vibration signal from the pump after it was high-pass filtered above 15 kHz. Here the role of the high-pass filter is to select a frequency region where the signal departs significantly from stationarity. The square-root of the mean instantaneous power may be interpreted as the envelope of the signal in that frequency band. Due to the finite length of real-world measurements, the computation of $P_x(t)$ had to be approximated by retaining in set $\mathscr{A}$ only those cyclic frequencies corresponding to the Fourier coefficients of $|x(t)^2|$ above a certain statistical level of significance, as shown in Fig. 5(b). The set $\mathscr{A}$ was found to comprise the blade-pass frequency at 393 Hz (i.e., eight times the pump rotation speed) and its first two multiples. This example illustrates quite well the capability of the cyclostationary approach to reveal the presence of a hidden periodicity in a signal, where the stationary assumption completely failed.

Fig. 5(c) displays the acoustical signal from the diesel engine together with its mean instantaneous power $P_x(t)$ normalised by the air characteristic impedance ρc_0 so as to return the instantaneous "sound intensity" radiated by the engine at the measurement point. Contrary to what is permitted by the stationary assumption, the cyclostationary approach makes it possible to track the chronology of events within the engine cycle (note the abscissa units which are

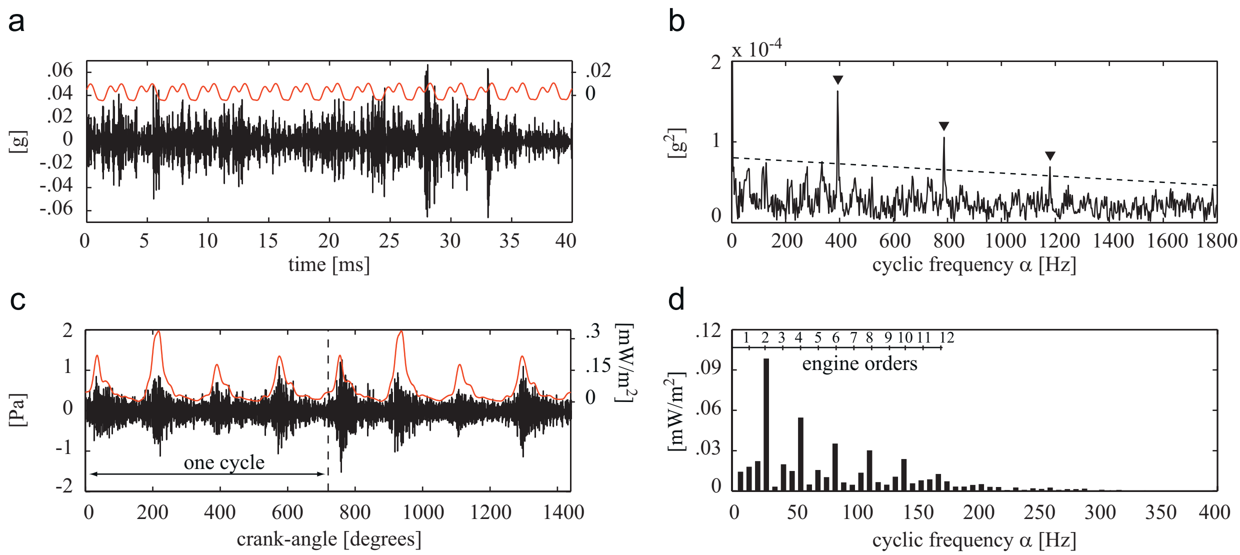

Fig. 5. (a) Vibration signal from the pump after high-pass filtration above 15 kHz together with the square-root of its mean instantaneous power $P_x(t)$, (b) amplitude spectrum of the squared magnitude of the high-pass filtered signal: the three harmonics at 393, 786 and 1179 Hz standing above the dotted threshold correspond to the non-zero cyclic powers $P_x{}^\alpha$, (c) acoustic signal from the diesel engine together with its mean instantaneous power $P_x(t)$ normalised by the air characteristic impedance ρc_0 so as to measure the radiated sound intensity and (d) the corresponding cyclic power $P_x{}^\alpha$ displayed as a function of frequency and of engine orders.

purposely displayed in degrees of rotation of the crankshaft): noise is essentially radiated around angles 20°, 200°, 380° and 560° in synchronisation with the ignitions in the four cylinders. The highest peak of sound intensity at 200° is seen to coincide perfectly with the cylinder pressure rise in Fig. 1(b)—that is the cylinder in front of which the microphone was located. The amplitudes of the cyclic powers P_x^{α} are displayed in Fig. 5(d) both as a function of frequency and of engine orders (order 1 = engine rotation speed). It nicely confirms the periodicity of the acoustic energy: the highest family of harmonics is locked on engine order 2, which is related to the firing frequency; the other harmonics at multiples of engine order $\frac{1}{2}$ are related to the engine cycle.

2.2.2.2. The instantaneous power spectrum. The mean instantaneous power $P_x(t)$ provides a global vision as how the energy of a signal is flowing with respect to time. This may not be refined enough to reveal certain types of hidden periodicities, such as in the case of the pump signal. Therefore, just as for stationary signals, it is tempting to go one step further and scrutinise how the mean instantaneous power is distributed in the frequency domain. Let us consider again $x_{\Delta f}(t;f)$, the filtered version of signal $x(t)$ through a frequency band of width Δf centred on frequency f. Now using the $\mathscr{P}$-operator, the mean instantaneous power passed through that frequency band is

$$P_x(t,f;\Delta f) = \mathscr{P}\{|x_{\Delta f}(t;f)|^2\}. \tag{8}$$

This is illustrated in Fig. 6. The quantity $P_x(t, f; \Delta f)$ describes how the energy is flowing with respect to both time and frequency and, in that sense, it is a valid time–frequency spectrum, similar to the spectrogram or Priestley's evolutionary spectrum [9]. However, just as with the latter, $P_x(t, f; \Delta f)$ is faced with the uncertainty principle which states that the frequency resolution Δf and the time resolution Δt are interrelated by the inequality

$$\Delta f \cdot \Delta t \geqslant \frac{1}{4\pi}, \tag{9}$$

so that they cannot be made arbitrarily and independently small. Stated differently, to be able to track small temporal variations of duration Δt, the filter bandwidth Δf should be at least as large as $1/(4\pi\Delta t)$. This precludes the instantaneous spectrum from being an instantaneous spectral *density* by making Δf tend to zero as in Eq. (2)—indeed it would boil down to the PSD $P_x(f)$ in that case, with complete loss of time information. Hence, it should always be remembered that the instantaneous power spectrum $P_x(t, f, \Delta f)$ explicitly depends on the user's defined Δf.

Example 4. Fig. 7(a) displays the instantaneous power spectrum of the acoustical signal from the diesel engine as a gray level image, computed over 213 consecutive engine cycles. This type of representation shows the temporal (angular) and frequency localisation of energy within the engine cycle: the successive ignitions in the four cylinders radiate sound around 20°, 200°, 380° and 560° over a wide frequency range with marked resonances around 2.9 and 5.6 kHz. Note that the energy mapping is periodic with respect to angle, so that a single image within the interval [0; 720°] is enough to describe the entire engine operation. For comparison, Fig. 7(b) displays the conventional spectrogram of the same signal computed over one engine cycle. The fact that the $\mathscr{P}$-operator was not involved in that computation clearly results in a more confused distribution of energy in the time–frequency plane. More precisely, the instantaneous power spectrum may be conceived here as the result of averaging the spectrograms of 213 engine cycles, thus reducing estimation noise to a very low level—by a factor $\sqrt{213} \approx 15$.

2.2.2.3. The cyclic modulation spectrum. The instantaneous power spectrum appears to be a very valuable tool for revealing the presence of hidden periodicities in a cyclostationary signal by displaying the energy flow jointly in time and in frequency. A natural extension is to quantify the intensity of such hidden periodicities by means of Fourier coefficients. The Fourier series expansion of the instantaneous power spectrum reads

$$P_x(t,f;\Delta f) = \sum_{\alpha \in \mathscr{A}} P_x^{\alpha}(f;\Delta f) \cdot e^{j2\pi\alpha t}, \tag{10}$$

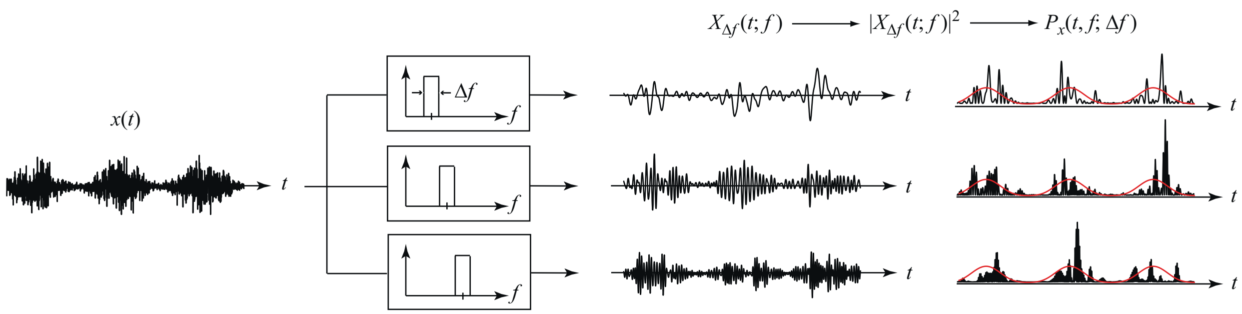

Fig. 6. Interpretation of the instantaneous power spectrum as the energy flow measured at the output of a filterbank.

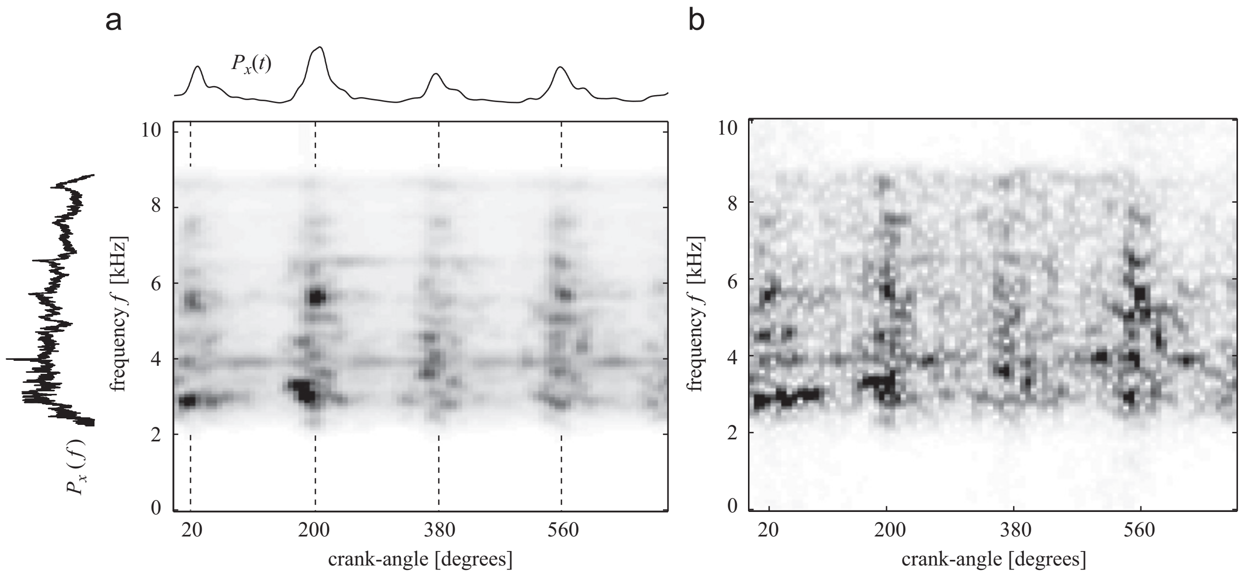

Fig. 7. (a) Instantaneous power spectrum $P_x(t, f; \Delta f)$ of the acoustical signal from the diesel engine computed over 213 consecutive engine cycles ($\Delta f = 240$ Hz and $\Delta t = 4\,\text{ms} \approx 21^\circ$). The image is displayed in gray levels where black corresponds to the highest intensity. Also shown are the mean instantaneous power $P_x(t)$ along the temporal (angle) axis and the power spectral density $P_x(f)$ along the frequency axis and (b) conventional spectrogram of the same signal computed over one cycle, with the same time–frequency resolution.

where we shall refer to the spectrum of Fourier coefficients

$$P_x^{\alpha}(f;\Delta f) = \mathscr{P}_0\{|x_{\Delta f}(t;f)|^2 \cdot e^{-j2\pi\alpha t}\} \tag{11}$$

as the *cyclic modulation spectrum*[1] of the signal. Note that the so-defined spectrum is a function of the frequency variable f and it is indexed by the frequency variable α that is the dual of time t. In order to avoid any confusion between the two frequency variables, the former is sometimes referred to as the "spectral" frequency, whereas the latter is the "cyclic" frequency as encountered in Definition 2. The representation of the cyclic modulation spectrum $P_x^{\alpha}(f, \Delta f)$ in the (f,α) plane is illustrated in Fig. 8. This type of representation (together with the spectral correlation density introduced in Section 2.4.3) is probably the most dedicated among all descriptors to analysing cyclostationary signals. First of all, it readily reveals the presence of cyclostationarity in a signal whenever a non-zero cyclic modulation spectrum is detected for any $\alpha \neq 0$—a stationary signal would have all its energy distributed along $\alpha = 0$ only. Second it makes known the complete set $\mathscr{A}$ of cyclic frequencies which the energy of the signal is travelling with as a function of time. Third it tells the magnitude of the cyclostationary component located at (f,α), where the spectral frequency f may be interpreted as the carrier frequency of the wave train transporting the energy, and the cyclic frequency α as its modulation frequency. This latter interpretation is illustrated in Fig. 9.

It is important to realise at this stage that to define unambiguously a modulation and a carrier frequency, the former should be smaller than the latter (the modulation should be slower than the carrier wave):

$$\alpha < f. \tag{12}$$

Moreover, by construction $P_x^{\alpha}(f; \Delta f)$ is faced with the same uncertainty principle as the instantaneous power spectrum: noting that the reciprocal of the time resolution Δt is the largest cyclic frequency $\alpha_{\max}$ that can be scanned in the filtered signal $x_{\Delta f}(t; f)$, the uncertainty principle (9) reads

$$\alpha_{\max} \leqslant 4\pi\Delta f. \tag{13}$$

A simple demonstration of this result is provided in Fig. 10. For that reason the cyclic modulation power spectrum is not eligible for being a density and it explicitly depends on the user's defined Δf.

Example 5. Fig. 11(a) displays the cyclic modulation spectrum of the acoustical signal from the diesel engine. It may be either construed as the map of Fourier coefficients of the instantaneous spectrum displayed in Fig. 7(a), or as a narrow-band spectral decomposition of the cyclic powers displayed in Fig. 5(d). In either case it provides useful additional information. First of all, the continuous distribution of energy along the spectral frequency f confirms that the signal is truly random in its waveforms, whereas the discrete distribution along the cyclic frequency α reveals that its energy is actually radiating periodically. Second, the energy cycles are seen to pertain to different mechanisms depending on the scrutinised

[1] This terminology is inspired the "modulation spectrum" used in speech processing [27].

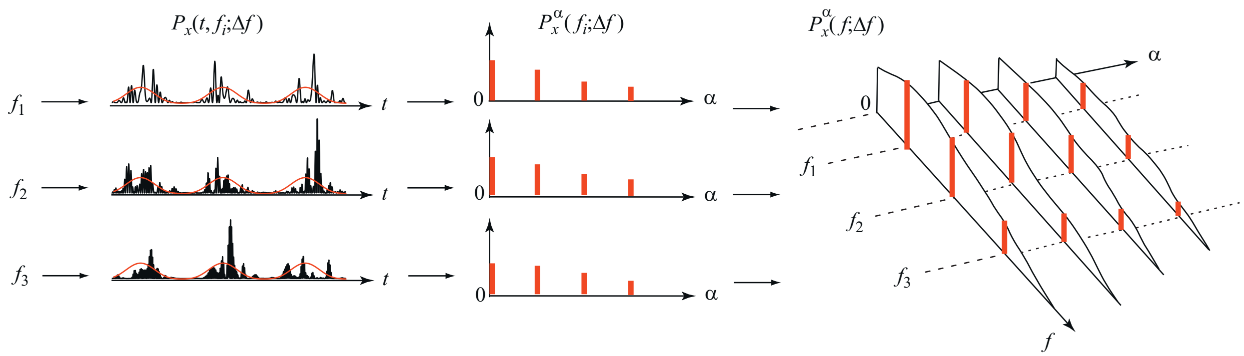

Fig. 8. Interpretation of the cyclic modulation spectrum as the cascade of Fourier spectra of the energy flow measured at the output of a filterbank.

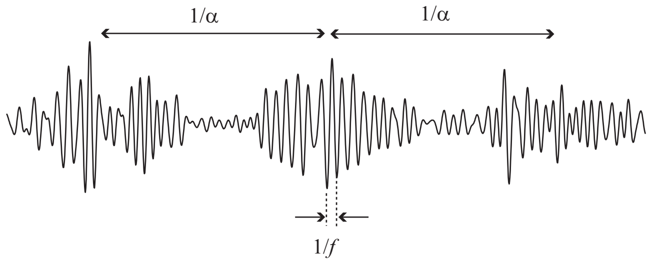

Fig. 9. Physical interpretation of the spectral frequency f and the cyclic frequency α of a cyclostationary waveform in terms of carrier and modulation frequencies, respectively.

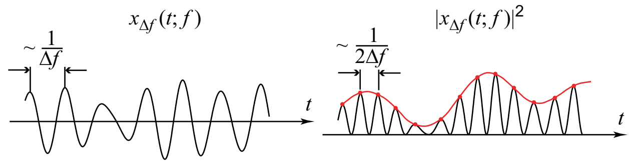

Fig. 10. Demonstration of the uncertainty principle (9) and (13) from Shannon's sampling condition. In order for the envelope of the squared signal $|x_{\Delta f}(t;f)|^2$ to be properly sampled, its maximum frequency α_{max} must not exceed half the sampling frequency $2\Delta f$, that is Δf.

frequency band: above 6 kHz, energy is essentially synchronised on engine order 2, that is the firing frequency; around the resonance at 5.6 kHz, energy is mainly synchronised on engine order 1, that is the crankshaft rotation; finally in the frequency band [2.6–3.6 kHz], energy is mostly synchronised on engine order $\frac{1}{2}$, that is the engine cycle (2 crankshaft rotations).

Fig. 11(b) displays the cyclic modulation spectrum of the vibration signal from the pump. This figure should be paralleled with the cyclic powers displayed in Fig. 5(b). The high value of the Fourier coefficient at $\alpha = 393$ Hz and over the whole f frequency range divulges the modulation by the blade-pass frequency. In contrast to Fig. 5(b), only the fundamental of the modulation can be detected here since the alpha frequency range cuts off just below the second harmonic; this is due to the uncertainty principle (13) which places an upper bound on the maximum cyclic frequency that can be scrutinised, viz $\alpha_{max} \leqslant 4\pi\Delta f = 782$ Hz. For the same reason, the presence of cyclostationarity at $\alpha = 393$ Hz can actually be detected in the cyclic modulation spectrum only if the frequency resolution is set coarser than about 30 Hz (Fig. 12)

2.2.2.4. Link with envelope analysis. The "envelope" of a signal is any function that envelops its energy fluctuations as a function of time. For it to be well-defined, it is usually computed after band-pass filtering the signal around a carrier frequency. Similarly, the "envelope spectrum" is the Fourier spectrum of the envelope function, a tool that has been extensively used in the processing of mechanical signals. These definitions are completely reminiscent of those of the mean instantaneous spectrum and of the cyclic power spectrum: $P_x(t, f; \Delta f)$ provides a surface representation of all the signal envelopes as a function of the carrier frequency f, computed in frequency bands of constant width Δf; alternatively,

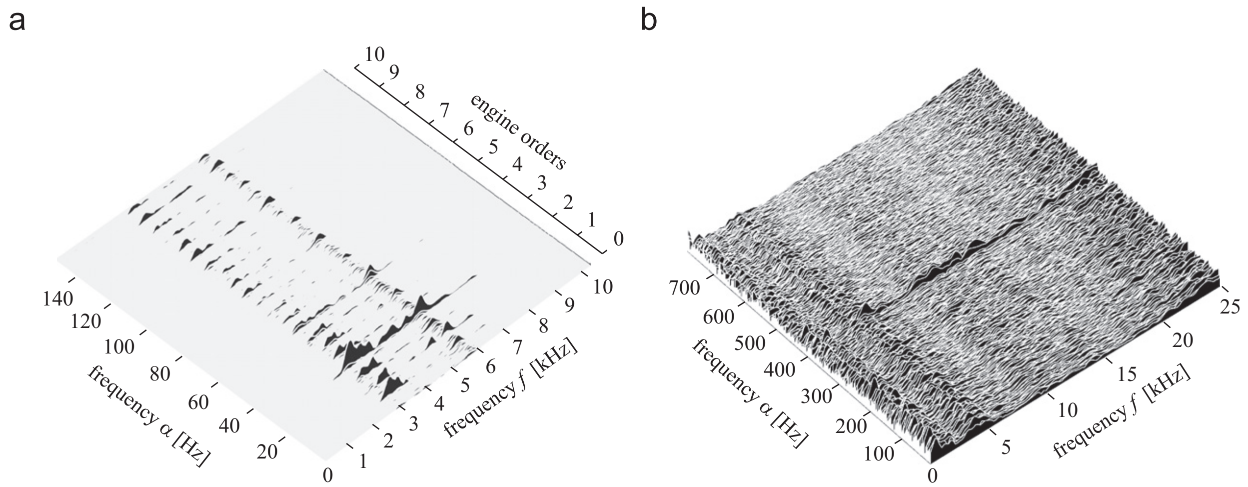

Fig. 11. (a) Cyclic modulation spectrum of the acoustical signal from the diesel ($\Delta f = 240\,\text{Hz}$; $\Delta\alpha = 0.28\,\text{Hz}$; $T\cdot\Delta f \sim 7200$) and (b) cyclic modulation spectrum of the vibration signal from the pump ($\Delta f = 62\,\text{Hz}$; $\Delta\alpha = 2.5\,\text{Hz}$; $T\cdot\Delta f \sim 65$).

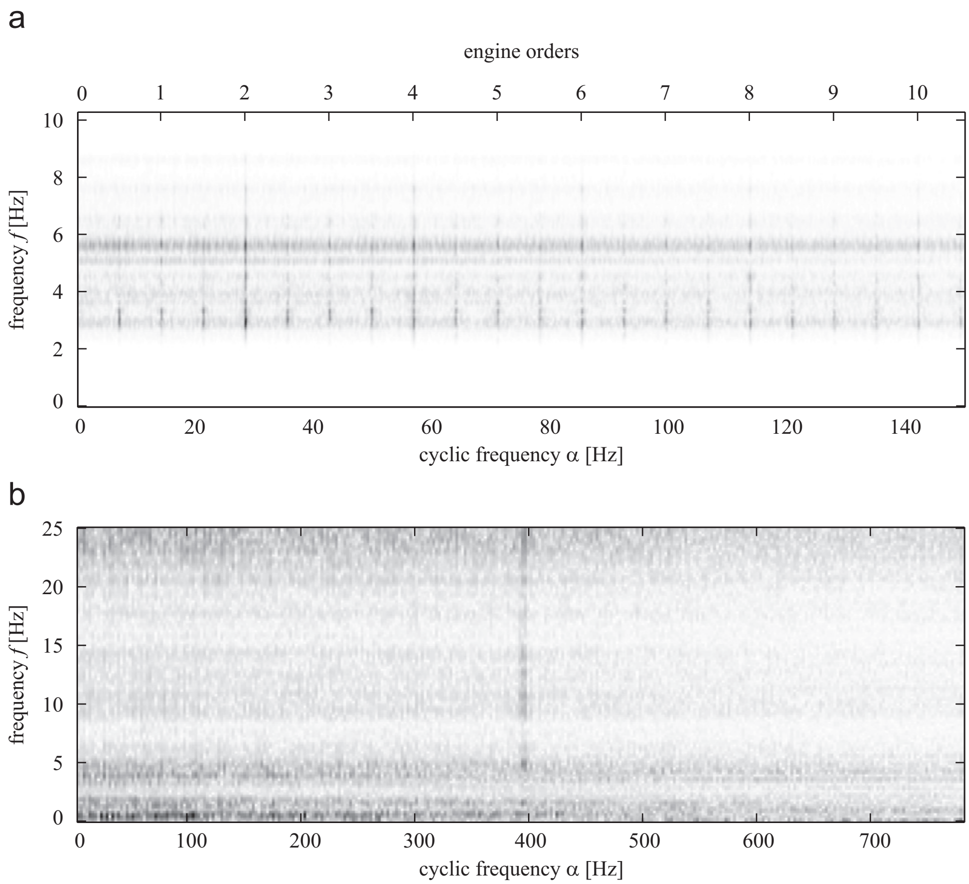

Fig. 12. Same quantities as in Fig. 11, but displayed as gray level images.

$P_x^{\alpha}(f;\Delta f)$ returns the corresponding set of all envelope spectra. In this sense, cyclostationarity generalises envelope analysis by returning at once the whole information that one would obtain after band-pass filtering around all possible frequencies. But cyclostationarity is more than just this! Section 2.4.3 will present a tool—the spectral correlation density—that returns

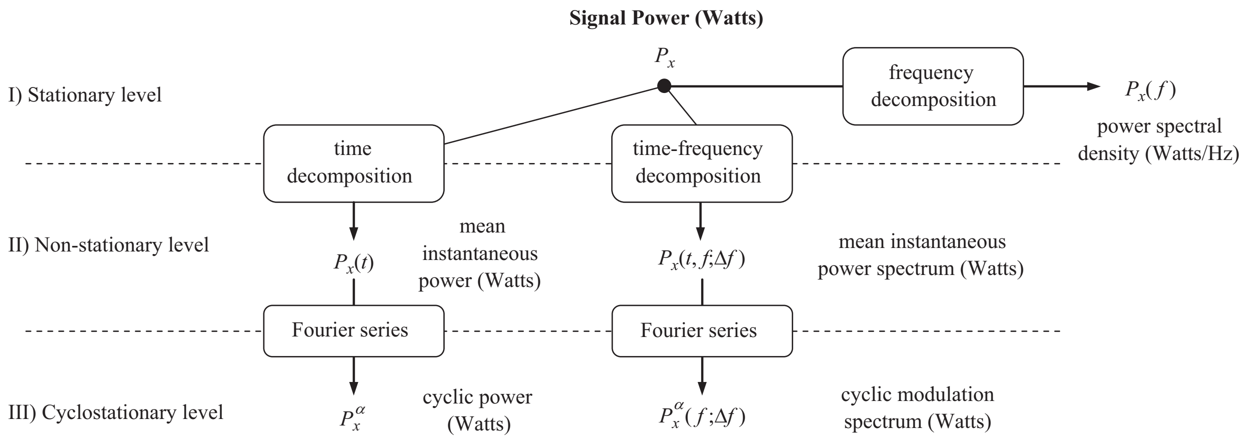

Fig. 13. Decompositions of the energy flow (by convention units of power are expressed in Watts).

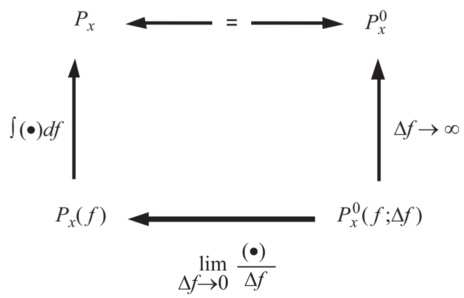

Fig. 14. Further relationships between the quantities of Fig. 13.

the same information independently of the frequency bandwidth Δf on which envelope analysis, just as $P_x(t, f; \Delta f)$ and $P_x^\alpha(f;\Delta f)$, critically depends.

2.2.3. *Synthesis on the several facets of energy*

At this juncture it has become clear that the cyclostationary property of a signal may be evidenced by decomposing its energy in one or other domain and verifying that it is flowing in a cyclic manner.

2.2.3.1. Decomposition of energy. Our approach so far has mainly consisted in considering various decompositions of the rate of energy released in a signal—i.e., the so-called *power* P_x—either in time, in frequency, or jointly in both domains. These various decompositions and their interconnections are summarised in Fig. 13. Also indicated there are the units of each quantity, and the "level" to which it pertains. Level I is dedicated to characterising stationary signals only, level II to nonstationary signals in general (assuming a generalisation of the $\mathscr{P}$-operator as discussed in Section 2.5.1), and level III to cyclostationary signals in particular. For the sake of completeness, some further relationships between the quantities of Fig. 13 are reported in Fig. 14.

2.2.3.2. Conservation laws of energy. Any proper decomposition in one or other domain should conserve the overall energy. This is not only reassuring from a physical point of view, but it is also desirable for a confident interpretation of the decomposition. The conservation of energy applies to the instantaneous power $P_x(t)$ and to the PSD $P_x(f)$; it reads:

$$P_x = \lim_{T\to\infty} \frac{1}{T}\int_T P_x(t)\,\mathrm{d}t = \int P_x(f)\,\mathrm{d}f. \tag{14}$$

Unfortunately, no similar relation applies to the instantaneous power spectrum or to the cyclic modulation spectrum in general. This is due to the fact that neither $P_x(t, f; \Delta f)$ nor $P_x^\alpha(f; \Delta f)$ are power *densities*. Care should be taken therefore when using these descriptors for *quantifying* how the power is distributed in the (t,f) or (f,α) domains. As a direct consequence of this discrepancy the integration over frequency of the instantaneous power spectrum $P_x(t, f; \Delta f)$ displayed in Fig. 7(a) does not reproduce exactly the instantaneous power $P_x(t)$ shown therein, nor does the integration over time reproduce exactly

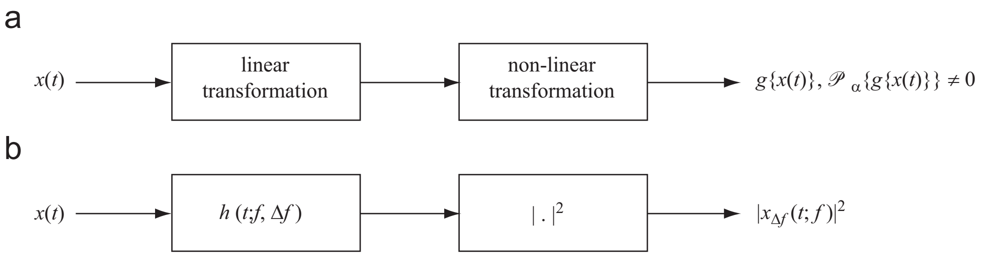

Fig. 15. (a) Scheme of the cascade of linear and non-linear transformations involved in Definition 4 and (b) particular example of such a cascade as used in the definition of the instantaneous power spectrum $P_x(f)$, where $h(t;\, f,\, \Delta f)$ refers to the impulse response of a band-pass filter of central frequency f and bandwidth Δf.

the power spectrum density $P_x(f)$. This is one reason that will motivate the introduction of another "cyclostationary" spectral descriptor in Section 2.4, namely the so-called "spectral correlation density".

2.2.4. A generic definition of cyclostationarity

Back from our tours in the time, frequency, time–frequency and frequency–frequency domains, we are now in a position to give a more generic definition of cyclostationarity.

Definition 4. A signal is said to exhibit cyclostationarity if there exists a cascade of linear and non-linear transformations that produces a periodic component. It is said to exhibit cyclostationarity at cyclic frequency α if there exists a cascade of linear and non-linear transformations that produces a sine component with frequency α. This is illustrated in Fig. 15.

In the situations explored hitherto, the linear transformation was either the identity (the "do-nothing" filter) or a pass-band filter with given bandwidth Δf, and the non-linear transformation was always a squaring devoted to producing a quantity homogeneous to energy. Definition 4 is generic in the sense that it places no constraints on the linear and non-linear transformations, so that any other variant can be envisioned. In particular it makes possible the consideration of other quantities than energy; the next subsection will briefly address this point.

2.2.5. Cyclostationary taxonomy

Definition 4 is surprisingly rich in spite of its concise formulation. First of all it makes possible the definition of several *orders* of cyclostationarity depending on the non-linear transformation it involves. Second, it allows the distinction between several *types* of cyclostationarity depending on the structure of the set $\mathscr{A}$ that contains the cyclic frequencies.

2.2.5.1. Orders of cyclostationarity. Let us consider the polynomial expansion of the non-linear transformation mentioned in Definition 4, and let us denote by n its highest degree. Any cyclostationary behaviour that can be detected by a non-linear transformation of degree n is referred to as nth order cyclostationarity. The energy-based approach we have followed so far involved a non-linear transformation of degree two (i.e., the squaring process); therefore it was implicitly restricted to the realm of second-order cyclostationarity. Another particular case is when the non-linear transformation is actually a linear one, that is when it is of degree one: it is clear from Definition 4 that a signal which exhibits first-order cyclostationarity is simply a signal which contains periodic components. This is a somewhat degenerate case since the signal processing tools it requires for its investigation do not need the specific cyclostationary apparatus developed so far; classical Fourier analysis will operate perfectly well in that case. Finally cyclostationarity of degrees higher than one and two is usually referred to as "higher-order" cyclostationarity [18]. The importance of distinguishing between the various orders of cyclostationary will be briefly discussed in Section 2.3.3.

2.2.5.2. Types of cyclostationarity. The classification of the type of cyclostationarity is made on the basis of the cyclic frequencies contained in set $\mathscr{A}$:

- repeating Definition 4, a signal is said to *exhibit* cyclostationary at cyclic frequency α if $\mathscr{A}$ contains α and its integer multiples *among other possible elements*,
- strictly speaking, a signal is said to *be* cyclostationary at cyclic frequency α if $\mathscr{A}$ contains α and its integer multiples, $k \cdot \alpha, k \in \mathbb{Z}$, *only*,
- as a particular case of a signal exhibiting cyclostationarity, a signal is said to be *poly-cyclostationary*—or quasi-cyclostationary—if $\mathscr{A}$ contains elements that are not necessarily integer multiples of each other (i.e., that are incommensurate).

It is important to grasp the subtle differences in this typology. A poly-cyclostationary signal may be seen as a summation of cyclostationary signals with different cyclic frequencies: it does not produce a periodic signal after non-linear transformation, but a poly-periodic signal, that is a summation of periodic components whose common fundamental

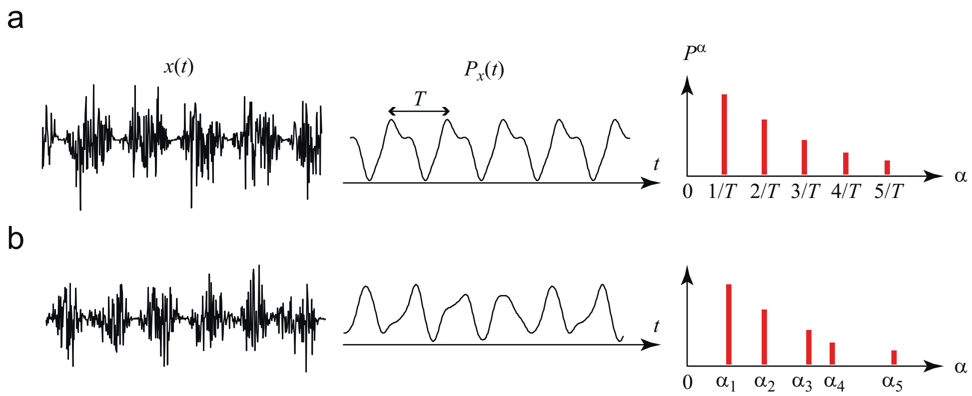

Fig. 16. (a) Example of a cyclostationary signal and (b) example of a poly-cyclostationary signal.

period is infinite in general—see Fig. 16). Note also that a cyclostationary signal does exhibit cyclostationarity, but the converse is not true in general (i.e., a signal which exhibits cyclostationarity is not necessarily cyclostationary).

2.2.5.3. Remarks. The classification of cyclostationary signals has brought to light how large the cyclostationary family is. In particular it has made clear that it includes as particular cases stationary random signals and (poly) periodic signals. The former are all members of the second and higher-order cyclostationary sets having only the zero cyclic frequency ($\alpha = 0$). The latter are members of the first-order cyclostationary set. Therefore the cyclostationary family is rich enough to encompass many of the real-life signals issuing from mechanical systems, and particularly from rotating and reciprocating mechanisms which always give rise to a mixture of several kinds of statistical behaviours.

Example 6. Because its PSD did not show any significant peak in Fig. 4(a), the vibration signal from the pump does not exhibit first-order cyclostationarity. On the other hand, the existence of non-zero coefficients in its cyclic powers and cyclic modulation spectrum in Figs. 5(b) and 11(b) proves that it exhibits second-order cyclostationarity. These coefficients being harmonically related to the fundamental blade-pass frequency 393 Hz, the signal is second-order cyclostationary at that cyclic frequency.

The cylinder pressure from the diesel engine in Fig. 1(b) is obviously first-order cyclostationary with respect to the engine cycle due to the deterministic and periodic compression/expansion part of the pressure is which is dominating the low frequencies; it is also cyclostationary on any higher order since any non-linear transformation would again produce a periodic signal.

The acoustical signal from the diesel engine mostly exhibits second-order cyclostationarity as seen from its cyclic powers and cyclic modulation spectrum in Figs. 5(d) and 11(a)—it will be shown in Section 2.3.2 that it exhibits first-order cyclostationarity as well, yet to a lesser extent. Careful inspection of its cyclic modulation spectrum reveals it is cyclostationary with a fundamental cycle equal to the engine cycle, that is engine order $\frac{1}{2}$.

Finally, a good example of a poly-cyclostationary signal is any vibration signal issuing from a complex gearbox consisting of several gears operating with a non-integer reduction ratio—e.g., see forthcoming Examples 9 and 15.

2.3. First, second and higher-order cyclostationarity

The distinction between first and second (higher) orders of cyclostationarity is of such importance that it deserves spending additional time looking into it. Indeed, it is intimately related to the notions of predictability and randomness of a signal,[2] which suggests that the two parts should be treated sequentially, with different signal processing tools.

2.3.1. The mean value of a cyclostationary signal

To begin with, let us come back to our very first approach of Section 2.2.1 where the statistical characterisation of a stationary random signal was completed in terms of its mean value and its mean power. These were two constant parameters returning a measure of position and of dispersion. Moreover, it was mentioned that for many mechanical signals the mean value is actually nil so that it can be set aside from the analysis. When moving from the stationary to the cyclostationary family of signals, this remark was voluntarily taken as granted in order to simplify the subsequent analysis. It is now time to reconsider the general situation where the mean value of a cyclostationary signal is possibly non-zero.

[2] As already mentioned after Definition 3, a predictable signal is one which can be fully described by an equation and whose values can thus be computed everywhere on the time axis, including arbitrarily far future or past. As opposed to this definition, a random signal can never be predicted exactly and everywhere on the time axis from a finite duration observation—see Ref. [9] for a rigorous presentation of these ideas wherein the words "regular" and "singular" stand for "random" and "predictable".

The mean value of a cyclostationary signal is computed following the same modus operandi as for a stationary signal, but where care is taken to replace the $\mathscr{P}_0$-operator—dedicated to extracting constant trends—by the $\mathscr{P}$-operator—dedicated to extracting periodic components. Specifically, denoting $x(t)$ the signal and $m_x(t)$ its mean value:

$$m_x(t) = \mathscr{P}\{x(t)\}. \tag{15}$$

It is of prime importance to understand that the mean value $m_x(t)$ of a cyclostationary signal is itself a *signal*—i.e., a function of time—contrary to the stationary case where it is always a constant. This is because it gathers all the periodic components in $x(t)$ that participate to the mean behaviour of the signal, i.e. to that part which is perfectly predictable on a cyclic basis (the distinction between the mean value and the time-averaged value of a signal is further discussed in Section 2.5.2).

2.3.2. The deterministic/random decomposition

As a direct consequence of definition (15), the mean value of a cyclostationary signal is that part of the signal which comprises first-order cyclostationary only. By difference, the residual part $\mathscr{R}\{x(t)\} = x(t) - \mathscr{P}\{x(t)\}$ returns the random contribution of the signal that is second and/or higher-order cyclostationary.

This decomposition is so important that it may be erected as a proposition.

Proposition 5. Any signal $x(t)$ (of presumably infinite duration) decomposes uniquely as

$$x(t) = \mathscr{P}\{x(t)\} + \mathscr{R}\{x(t)\} \tag{16}$$

where the predictable part $\mathscr{P}\{x(t)\}$ embodies all the periodic components in the signal and the residual part $\mathscr{R}\{x(t)\}$ embodies all the random components.

It is crucial to understand at this stage that the two constituents $\mathscr{P}\{x(t)\}$ and $\mathscr{R}\{x(t)\}$ in Eq. (16) may be *equivalently* and respectively construed as:

- a signal due to first-order cyclostationarity and a signal that exhibits second and/or higher-order cyclostationarity only,
- a *mean* signal and a *residual* signal—also referred to as a *centred* signal,
- a *deterministic* signal and a purely *random* signal,
- a signal having a *discrete* power spectrum and a signal having a *continuous* power spectrum.

These several equivalences are purposely highlighted here because they have been a source of constant confusion in the literature. Note that Ref. [2] is probably precursory in proposing this decomposition for analysing mechanical signals.

Example 7. Fig. 17 displays the cylinder pressure signal of the diesel engine over 4 cycles together with its decomposition into a mean value $\mathscr{P}\{x(t)\}$ and a residual random part $\mathscr{R}\{x(t)\}$. As expected from our previous investigations, the signal is found to be essentially first-order cyclostationary, i.e. with a strong periodic mean value $\mathscr{P}\{x(t)\}$ synchronised on the engine cycle. This is consistent with the discrete nature of the power spectrum observed in Fig. 4(b). It is noteworthy that the so-computed mean value is drastically different from the constant time-averaged value (DC value) of the signal which has been superposed on $\mathscr{P}\{x(t)\}$ for comparison. Finally, the residual part $\mathscr{R}\{x(t)\}$ indicates that a slight random fluctuation exists from cycles to cycles, yet it is two orders of magnitude smaller than the mean value $\mathscr{P}\{x(t)\}$; it is obviously second-order cyclostationary since squaring it would produce periodic components.

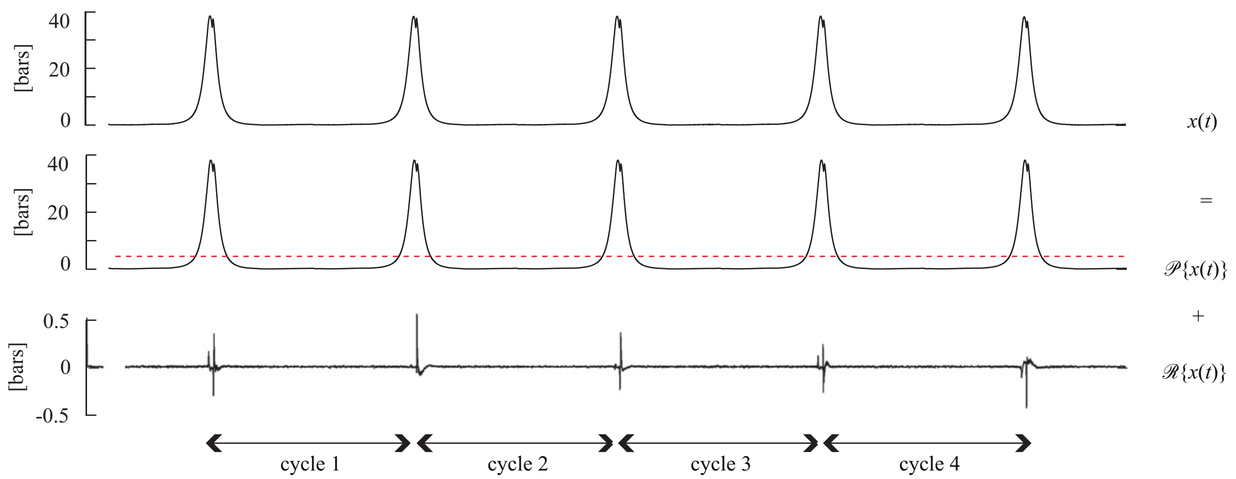

Fig. 17. Decomposition of the cylinder pressure signal from the diesel engine into its deterministic and random parts. The dotted line superposed on $\mathscr{P}\{x(t)\}$ indicates the time-averaged value (DC value) of the signal.

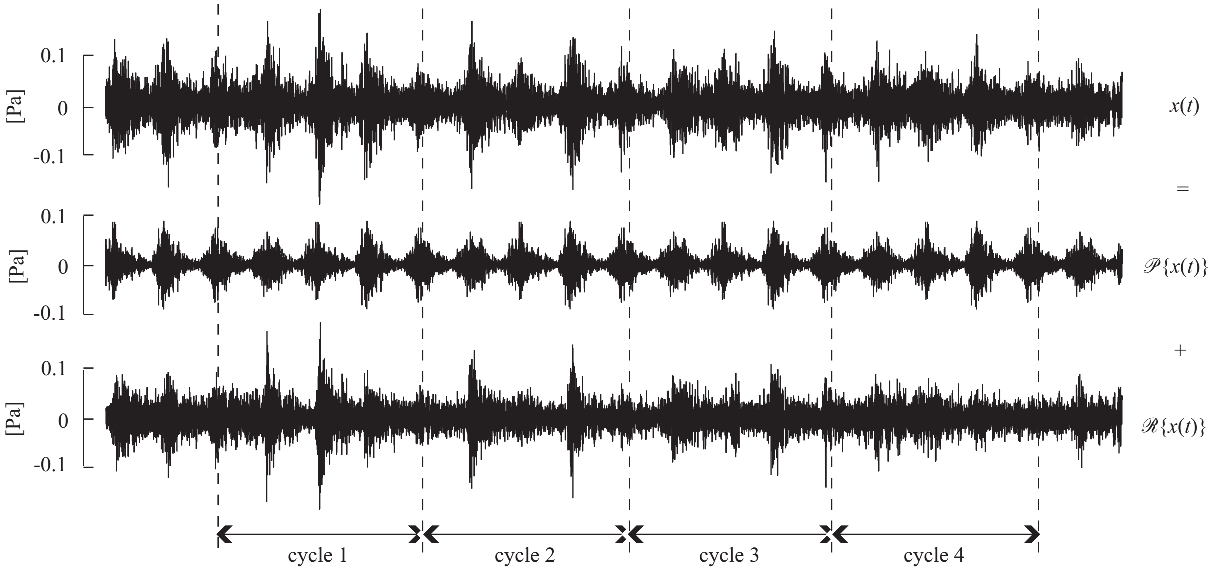

Fig. 18. Decomposition of the acoustical signal from the diesel engine into its deterministic and random parts.

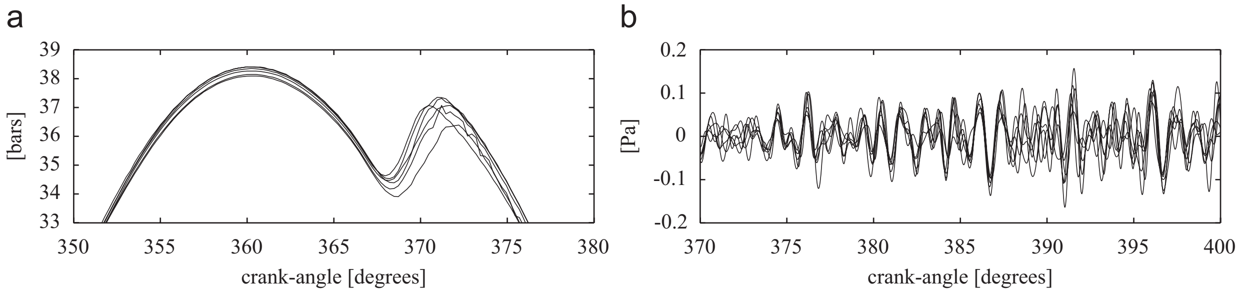

Fig. 19. Superposition of several cycles of (a) the cylinder pressure signal and (b) the acoustical signal from the diesel engine over a short angular sector.

For comparison, Fig. 18 displays the deterministic/random decomposition of the corresponding acoustical signal. In this case, the mean value $\mathscr{P}\{x(t)\}$ of the signal is found larger than expected, although the random part $\mathscr{R}\{x(t)\}$ is still predominating. Interestingly, the mean value signal seems to originate primarily from the four firings taking place within the engine cycle.

More insight can be gained into the deterministic/random behaviour of the diesel engine signals by zooming on the superposition of several successive cycles over a short angular sector, as displayed in Fig. 19.

2.3.3. Pure versus impure cyclostationarity

Let us consider again the cylinder pressure signal of Example 7. It was observed to exhibit strong first-order cyclostationarity since it gives rise to a periodic component without need of a non-linear transformation. Actually it also exhibits strong second-order cyclostationarity since squaring it would give rise to significant periodic component as well. Indeed, any non-linear transformation of arbitrary order n would give rise to strong periodic components in that case, so that the signal surely exhibits cyclostationary at all orders. On the other hand, the residual part $\mathscr{R}\{x(t)\}$ obtained after performing the deterministic/random decomposition exhibits only a low level of second-order cyclostationarity. What would happen in the limit case when that part is identically zero? Would it be right to say that a perfectly periodic signal exhibits strong cyclostationarity at any order while obviously it is all due to the first-order cyclostationary component?

In order to clarify this situation, it is convenient to distinguish between *pure* and *impure* cyclostationarity. Impure cyclostationarity of order n refers to that kind of cyclostationarity that is induced by cyclostationary components of orders less than n. For instance a periodic signal will exhibit impure cyclostationarity at any order. Pure cyclostationarity of order n refers to that kind of cyclostationarity that remains after all the cyclostationary components of lower orders have been removed. For instance a periodic signal is purely cyclostationary at order one and does not exhibit pure cyclostationarity at any higher (than one) order.

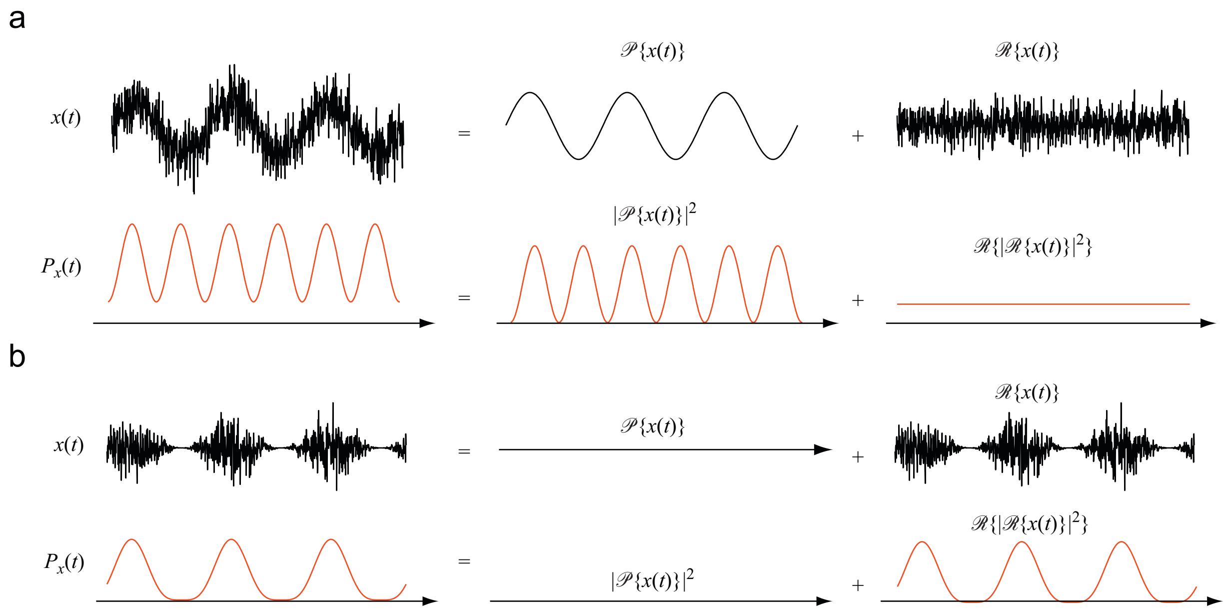

Fig. 20. Decomposition of synthetic signals into their deterministic and random parts, paralleled with the decomposition of their mean instantaneous powers and (a) a pure first-order cyclostationary signal and (b) a pure second-order cyclostationary signal.

Having said that, it appears that Definition 4 actually referred to impure cyclostationarity. The decomposition introduced in Proposition 5 will now make the definition of pure cyclostationarity perfectly precise up to order 2.

Definition 6. A signal is said *purely* cyclostationary at order 1 if its residual part $\mathscr{R}\{x(t)\}$ does not exhibit cyclostationarity at any order.

Definition 7. A signal is said *purely* cyclostationary at order 2 if its deterministic part $\mathscr{P}\{x(t)\}$ is nil and its residual part is cyclostationarity at order 2.

Pure higher-orders of cyclostationarity can be defined in the same way: for instance pure nth-order cyclostationarity refers to any hidden periodicity that can be evidenced by using a non-linear transformation of order n after all cyclostationary components of lower orders $(<n)$ have been removed. The study of pure higher-order cyclostationarity is outside the scope of this paper and the interested reader is invited to consult Refs. [18,19].[3]

Example 8. A typical example of a first-order cyclostationary signal $x(t)$ is made of a sinusoid $p(t)$ buried in stationary background noise $v(t)$ of variance σ^2: $x(t) = p(t)+v(t)$ The $\mathscr{P}$-operator then makes it possible to achieve the separation $\mathscr{P}\{x(t)\} = p(t)$ and $\mathscr{R}\{x(t)\} = v(t)$. The instantaneous power of the signal reads $P_x(t) = \mathscr{P}\{|x(t)|^2\} = |p(t)|^2 + \sigma^2$ wherein the background noise contributes to a constant term—see Fig. 20(a).

A typical example of a pure second-order cyclostationary signal is made of a stationary noise $v(t)$ of variance σ^2 modulated by a sinusoid: $x(t) = p(t) \cdot v(t)$. In this case $\mathscr{P}\{x(t)\} = 0$ and $\mathscr{R}\{x(t)\} = x(t)$. The instantaneous power of the signal is given by $P_x(t) = \mathscr{P}\{|x(t)|^2\} = |p(t)|^2 \cdot \sigma^2$—see Fig. 20(b).

2.3.4. *Mixed cyclostationarity: the motivation for a sequential approach*

Real-world systems rarely produce signals that are purely cyclostationary at a given order only; they are rather a combination of several orders of cyclostationarity, as illustrated by the acoustical signal from the diesel engine in Fig. 18. Such signals are said to exhibit *mixed cyclostationarity*. In many applications, however, restricting the analysis to the first two orders will turn out sufficient because they completely characterise the signal in terms of its (time-varying) mean position and mean dispersion.

In order to make explicit the distinction between first and second-order *pure* cyclostationarity, it is recommended to perform the fundamental decomposition of Proposition 5 where first-order cyclostationarity is extracted by the $\mathscr{P}$-operator

[3] For those readers familiar with the use of moments and cumulants, nth-order cyclic moments are devoted to characterising *impure* cyclostationarity of order n, whereas nth-order cyclic cumulants are devoted to *pure* cyclostationarity of order n. Indeed, nth-order cumulants are related to nth-order moments by removing in the latter the interactions coming from all lower (than n) order moments.