Politecnico di Milano
Scuola di Ingegneria Industriale e dell'Informazione
Corso di Laurea Magistrale in Ingegneria Informatica
Dipartimento di Elettronica, Informazione e Bioingegneria

A spectrum-based adaptive sampling mechanism for energy conservation in Cyber-Physical Systems

Relatore: Prof. Manuel ROVERI
Correlatore: Prof. Cesare ALIPPI

Tesi di Laurea di: Ilaria Scarabottolo
Matricola: 819927

Anno Accademico 2015-2016

To my parents,
who will always inspire me.

Abstract

The world of Cyber-Physical Systems is rapidly evolving and the applications of these systems are increasing in everyday life. With the growing size of these systems, built from the integration of computational algorithms and interacting physical components, energy management becomes a crucial issue for guaranteeing QoS applications of these technologies. Indeed, trying to manage energy supply for Cyber-Physical Systems without an accurate energy-saving policy is simply not sustainable, since their diffusion is still growing and, as a consequence, previously-employed energy-management approaches are quickly becoming obsolete.

Cyber-Physical Systems (CPS) interact with the surrounding environment through sensing units, which are embedded systems acquiring data from the physical world. Sensing units are subject to memory and computational power limitations, as normally they are small devices spread over a large geographical area. Nevertheless, they often need to perform complex operations and, therefore, guaranteeing a sufficient energy supply to these units may not be simple and energy management techniques become necessaries.

The purpose of this thesis is to explore an innovative solution for energy management in CPS, belonging to the category of adaptive sampling techniques. These techniques are aimed at scheduling the optimal sampling frequency for each CPS sensing unit, thus avoiding the acquisition of useless data and the unnecessary use of valuable resources (e.g. energy). The main idea of the proposed approach is to monitor changes in the maximum frequency contained in the acquired signal without reconstructing its whole spectrum.

The novelty of the proposed approach is based on the fact that, for the first time, the aliasing phenomenon is exploited to detect possible changes in the maximum frequency of a signal. Another important novel aspect is that, in this technique, signals are studied in the time domain instead of the frequency domain. This is a significant advantage in terms of energy savings, since monitoring the signal without reconstructing its frequency spectrum avoids the large computatioal load introduced by the spectrum analysis.

The proposed solution has been tested on several datasets and compared to other existing approaches, showing its validity and efficacy in reducing energy consumption for CPS sensing units. Finally, the designed technique has been ported on a real embedded system, in order to have a demonstration of its applicability in the real-world CPSs.

Sommario

Il mondo dei Sistemi Ciberfisici (o CPS, Cyber-Physical Systems) sta evolvendo rapidamente e le applicazioni di questi sistemi sono sempre più diffuse nella vita di tutti i giorni. Con l'aumentare delle dimensioni di questi sitemi, nati dall'integrazione di algoritmi ed elementi fisici dotati di capacità computazionale, la gestione del consumo energetico diventa un problema cruciale per garantire un'applicazione ottimale di queste tecnologie e il massimo sfruttamento delle loro potenzialità. Infatti, gestire l'approvvigionamento energetico di un sistema ciberfisico senza adottare una politica energetica volta al risparmio delle risorse non è sostenibile, poiché la diffusione di questi sistemi sta aumentando e, di conseguenza, gli approcci di conservazione dell'energia utilizzati in precedenza diventano rapidamente obsoleti.

I Sistemi Ciberfisici interagiscono con l'ambiente circostante attraverso unità di rilevazione, o *sensing units*, ossia sistemi integrati che acquisiscono dati dal mondo fisico. Le sensing units devono sottostare a limiti di memoria e di potenza di calcolo, dal momento che normalmente sono dispositivi di dimensioni ridotte distribuiti su una vasta area geografica. Ciononostante, molto spesso devono compiere operazioni complesse e, quindi, garantire un apporto di energia sufficiente a queste unità potrebbe non essere semplice e tecniche di risparmio energetico diventano necessarie.

L'obiettivo che si prefigge questa tesi è quello di esplorare una soluzione innovativa al problema della gestione energetica nei Sistemi Ciberfisici, soluzione appartente alla categoria delle tecniche di campionamento adattativo. Con il termine «campionamento adattativo» si intende la modifica della frequenza di campionamento di un dato segnale (che potrebbe essere generato da qualunque fenomeno fisico), seguendone i momenti di attività più debole o più intensa. Lo scopo di fondo è quello di evitare l'acquisizione di campioni inutili e, di conseguenza, l'utilizzo non necessario di risorse preziose come, appunto, l'energia. L'idea di base dell'approccio proposto in questa tesi è di monitorare i cambiamenti nella frequenza massima contenuta nel segnale osservato senza ricostruirne l'intero spettro di frequenza.

La frequenza massima contenuta in un segnale è un parametro importante per la corretta ricostruzione del segnale stesso, ma spesso questa frequenza non è nota a priori e, soprattutto, può evolvere nel tempo. Un approccio banale per garantire che il segnale venga perfettamente ricostruito è quello di campionare a frequenza molto più alta di quella massima contenuta nel segnale. Questo approccio però è assolutamente svantaggioso dal punto di vista energetico e, per alcune applicazioni, il dispendio di risorse sarebbe tale da renderlo insostenibile. Qui entrano in gioco le tecniche di campionamento

adattativo, che puntano a ridurre la frequenza di campionamento ogni qualvolta le dinamiche del sistema lo consentano (ovvero quando la grandezza fisica osservata non risente di variazioni sostanziali e il numero di campioni acquisito nel tempo può essere ridotto), tra le quali si colloca la soluzione proposta in questa tesi.

L'aspetto innovativo dell'approccio proposto si basa sul fatto che, per la prima volta, il fenomeno dell'*aliasing* (o ribaltamento dello spettro) viene sfruttato per identificare possibili cambiamenti nella frequenza massima contenuta nel segnale. In particolare, grazie al fenomeno dell'aliasing, è possibile riscontrare un cambiamento nella distribuzione delle frequenze e, di conseguenza, attivare ulteriori analisi per analizzare la natura del cambiamento ed eventualmente modificare la frequenza di campionamento per far fronte all'evoluzione del fenomeno fisico monitorato. Un altro importante aspetto innovativo è che, nella tecnica proposta, i segnali sono analizzati nel dominio del tempo invece che nel dominio delle frequenze. Ciò rappresenta un vantaggio significativo in termini di risparmio energetico, dal momento che monitorare il segnale senza ricostruirne lo spettro di frequenza evita la necessità di effettuare calcoli molto pesanti dal punto di vista computazionale.

La soluzione proposta consiste in una serie di fasi successive di trattamento dei dati, che comprende il filtraggio in bande di frequenza e il calcolo dell'energia contenuta nel segnale. I dati così ottenuti sono sottoposti a un Change Detection Test, un test statistico che rivela cambiamenti nella dinamica del processo sotto esame: ogni volta che viene segnalata una variazione, viene messo in atto un meccanismo di adattamento della frequenza di campionamento, basato sull'energia contenuta nelle diverse bande di frequenza, volto a gartantire che non vi sia perdita di informazione durante l'acquisizione del segnale.

La soluzione proposta è stata testata mediante simulazioni su diversi dataset e confrontata con altri approcci esistenti, rivelandone la validità e l'efficacia nella riduzione del consumo energetico nelle sensing units dei Sistemi Ciberfisici. Infine, la tecnica sviluppata è stata portata su un sistema integrato reale, composto da una scheda programmabile e da alcuni sensori in grado di misurare temperatura, umidità e pressione. Il sistema integrato in esame è soggetto ai limiti di memoria, potenza di calcolo e durata della batteria tipici di questi dispositivi, in modo da fornire una dimostrazione dell'applicabilità della soluzione proposta nei veri Sistemi Ciberfisici.

Ringraziamenti

Questa tesi ha rappresentato per me un'evoluzione dal punto di vista accademico, ma anche personale. L'idea di fondo ha preso forma poco alla volta, grazie all'esplorazione di strade nuove. Ringrazio sentitamente i miei professori, vere e proprie guide, che hanno permesso a questa idea di realizzarsi appieno e trasformarsi in un lavoro di cui, ora, sono orgogliosa. Grazie a Manuel Roveri per il sostegno continuo, per gli aiuti risolutori e per la prontezza di intervento nei momenti in cui ho avuto bisogno di lui. Grazie a Cesare Alippi per i consigli preziosi sulla mia tesi e sul mio futuro professionale.

Non ringrazierò mai abbastanza i miei genitori, senza i quali il mio percorso accademico non sarebbe certo stato lo stesso. Grazie per avermi sempre sostenuta e guidata in modo incondizionato. Grazie a Ottavia, che negli anni è diventata un modello e un punto di riferimento, nonché un'inesauribile fonte di risate. Ringrazio Lorenzo per avermi appoggiata e piacevolmente distratta durante questa importante fase dei miei studi.

Ringrazio tutti i miei amici, anche quelli d'Oltralpe, che mi hanno accompagnata per tutti questi anni e che hanno attraversato le mie stesse tappe e affrontato i miei stessi dubbi.

Mi piace vedere questa tesi come una degna conclusione di un bellissimo percorso, ma anche come un punto di partenza per un altro che seguirà e che, spero, sarà altrettanto bello.

Contents

List of Figures

List of Algorithms

List of Tables

1 Introduction

1.1 Cyber-Physical Systems

The term *Cyber-Physical Systems* (CPS) refers to interacting networks of physical and computational components, where physical entities are controlled or monitored by computer-based algorithms. Embedded computers and networks monitor and control the physical processes, usually with feedback loops, where physical processes affect computations and vice-versa. Applications of CPS have reached by now a very wide diffusion, they include medical devices and systems, assisted living, traffic control and safety, advanced automotive systems, process control, energy conservation, environmental control, avionics, critical infrastructure control (electric power, water resources, and communications systems for example), defense systems, manufacturing, and smart structures [Lee08].

Some specific technological instances of CPS are Internet of Things (IoT) and Wireless Sensor Networks (WSN). IoT is a Cyber-Physical System whose diffusion arouse recently, and it indicates the network of physical objects (like devices, vehicles or buildings) embedded with sensors and software that enable these objects to communicate, collect and exchange data. A Wireless Sensor Network, in turn, consists of sensor nodes deployed over a geographical area for monitoring physical quantities like temperature, humidity, vibrations, seismic events, etc [DRAP15].

CPS interact with the surrounding environment through sensing units. The most common model for a sensing unit includes three basic components: a sensing subsystem for data acquisition, a processing subsystem for local data processing and storage, and a wireless communication subsystem for data transmission. Sensing units might also have actuation subsystems. In addition, a power source supplies the device with the energy needed to perform all these tasks. This power source usually consists of a battery with limited duration, which could be unpractical or even impossible to recharge or substitute, because units may be deployed in a hostile environment (e.g., networks monitoring environmental parameters in difficult access areas, like underwater sensors or avalanche monitoring applications [AAG+07]). At the same time, the CPS should have a lifetime long enough to fulfill the application requirements.

Even if sensing units are supplied by an external power source (e.g., a solar panel) it is still necessary to manage power consumption carefully, since the external power source could not assure continuous availability and, in general,

the aim is reducing power consumption as much as possible to guarantee a longer lifetime for the system.

From a general point of view, there are two ways to reduce the energy consumption in CPS: network-level optimization (carried out through algorithms and protocols that enable more efficient communication) and single unit level optimization (managing the single sensor activity to reduce the amount of acquired data and consumed energy). This work will concentrate on the single unit level, but, if properly combined, these two approaches can be conjunctively managed to further improve CPS efficiency.

The most extensively adopted CPS model in literature, which is reported in Figure 1.1 for clarity, consists of one *sink node*, or *base station*, and a possibly very large number of *sensing nodes,* which could be deployed over a geographic area, called *sensing field*, or consist in mobile units like cars, people's smartphones, etc. Data are transferred from sensor nodes to the sink through a multi-hop communication paradigm [ZZSF11], but there may be other protocols as well.

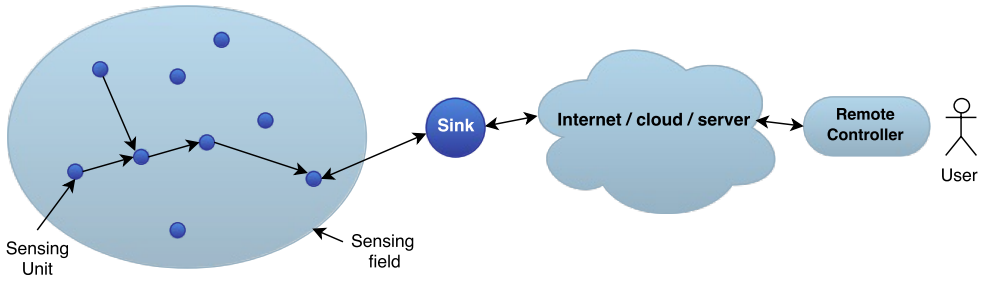

Figure 1.1: Cyber-Physical System scheme from [DRAP15]

A detailed taxonomy of energy conservation schemes currently available is presented in [DRAP15]. The most relevant subset of these approaches, closely related to this work, will be illustrated in Section 1.2, in order to put the current work in the proper perspective when compared to the general framework.

1.2 Energy conservation at CPS unit

There are different aspects of energy consumption at the single unit level on which it is possible to operate, and, of course, all of them could be taken into account simultaneously.

Duty cycling is a technique that consists in putting the electronic components and the radio transceiver in the low-power sleep mode whenever processing, acquisition or communication are not required. Ideally, the radio should be switched off as soon as there is no more data to send or receive,

and it should be resumed as soon as a new data becomes ready to be transmitted/received. Similarly, sensing units should be awakened immediately before acquiring a new sample and put back to sleep immediately after. This technique strongly reduces the energy amount spent for communication and data processing, as the unit operates less time than in case of full-time working. However, this approach is completely independent from data that are sampled by sensor units.

Quite often a-priori information about data distribution can be exploited to optimize data acquisition and further reduce the energy consumption. Usually, acquisition gathered from a real environment present strong autocorrelation and it might not be necessary to transmit every single sample to the base station. Therefore, data-driven approaches come into play. Actually, the sensing subsystem has two main energy consumption factors: one is given by the amount of data acquired and, consequently, processed and/or transmitted. In this sense, it is important to avoid the acquisition of unneeded samples which have strong spatial and/or temporal correlation ([VA04]). The second factor is the power supply needed by the sensing subsystem itself: when the sensor unit is power hungry, because it needs a lot of energy to accomplish a single acquisition (e.g., it might need to perform several operations to acquire a single sample), reducing processing and communication by exploiting data correlation may not be enough to preserve residual energy from expiring too early. *Data-driven* techniques are designed to reduce the amount of sampled data without reducing the accuracy in the observation of the phenomenon of interest, and among these techniques lays *adaptive sensing.*

Adaptive sensing is a widely employed mechanism aimed at decreasing the number of acquired samples according to the real dynamics of the physical phenomenon under monitoring and it can be implemented with three different approaches ([AADFR09]): *hierarchical sensing*, *model-based sensing* and *adaptive sampling*. The overall organization of these approaches is shown in Figure 1.2.

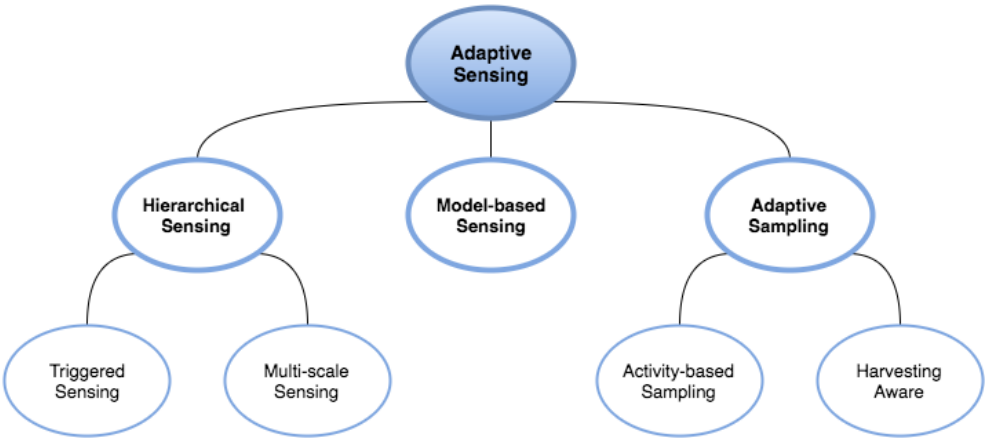

Figure 1.2: Energy-efficient data acquisition scheme from [AADFR09]

Hierarchical sensing is, in turn, divided into *triggered sensing* and *multi-scale sensing*. In both cases, these techniques assume that several sensors are installed on the sensing unit and observe the same event at a different resolution: simple sensors, which are energy-efficient, provide low resolution about the monitored event, while advanced sensors provide more accurate measurements but are more energy demanding. Normally, simple sensors are active all-of-the-time, while advanced sensors can be activated when needed, thus introducing duty cycling adaptation. In triggered sensing, advanced sensors are turned on once the low-resolution sensors detect some activities within the sensed area. For example, in security surveillance systems camera modules configured to provide low-resolution images work continuously but, if a target is detected, cameras are reconfigured into a fine-grained, high-quality image and provide images with high resolution.

Multi-scale sensing instead refers to spatial correlation and identifies areas within the monitoring field that require a more accurate observation and high-resolution data. An example is fire emergency management [TWCH07]. The sensor field is equipped with static sensors which monitor the environment. When a specific area presents an anomaly (e.g. the sampled temperature is above a predefined threshold), static units ask the base station for more specific data. As a consequence, the base station sends a mobile sensor unit to visit the potentially critical location.

In model-based sensing a forecasting model of the monitored phenomenon is built through an initial set of sampled data and used to forecast the next samples. Predictions are verified over time with actual readings and a measure of accuracy is calculated. Whenever the requested accuracy is not satisfied, the model is updated to better fit the physical phenomenon under observation. The effectiveness of this approach is bounded by the accuracy of the model and the nature of the process to be monitored.

Finally, adaptive sampling techniques consist in changing the sampling rate

according to correlations between the sensed data and information related to the available energy. If the signal under control evolves slowly with time, so that subsequent samples do not differ significantly, the sampling rate can be reduced thanks to this temporal correlation. It is also very likely that measurements taken by sensor nodes that are spatially close to each other do not differ significantly, therefore this spatial correlation can be used to reduce the energy-sensing consumption. Adaptive sampling approaches are divided in *activity based* and *harvesting aware.*

Activity-based adaptive sampling exploits temporal or spatial correlation among signal's values to further reduce the number of samples to be acquired, while it increases the sampling frequency as soon as a higher variability in data is detected (an example is the snow monitoring application of [AAG+07]).

The sampling rate can also be dynamically adjusted with harvesting-aware adaptive sampling that depends on the available energy, optimizing power consumption at the unit level by using the known residual energy and the forecasted energy coming from a harvester module ([AADFR09]). This approach is, however, more delicate, because generally some calculation is needed to compute the maximum frequency of a signal and adapt the sampling frequency consequently. Without this key step, the risk of missing some important event increases.

Each specific algorithm or technique among those presented above can concentrate more on one aspect or another, but generally they could be all considered when designing a mechanism to reduce energy consumption in a Cyber-Physical System.

This work is based on the activity-based adaptive sampling approach as well as duty cycling. Actually, the two approaches are complementary and often used in combination: the operating system powers the sensors on and off accordingly to the desired sampling frequency, which is, in turn, continuously adapted through an adaptive sampling algorithm. The algorithm presented here is, again, activity-based and, more specifically, it exploits the temporal correlation among signal's energy values. As a matter of fact, most changes in the system dynamics result in a variation of the signal's energy. Hence, this quantity can be used to determine whether an adaptation of the sampling frequency is needed. A major point of the proposed solution is that this can be done independently from the system's model or the network hierarchy.

In every system aimed at monitoring an evolving phenomenon, the sampling frequency plays a cardinal role, since the working principle of every sensor is to monitor a continuous signal and convert it into a discrete digital one. Cyber-Physical Systems units can monitor any kind of physical quantity: light brightness, electrical tension, acceleration or temperature are just a few examples. However, to ensure an accurate reconstruction of the observed signal, the sampling frequency has to be properly selected.

A milestone theorem about generic signals sampling is the Nyquist-Shannon theorem ([Pra10]). It states that, in order to reconstruct completely a signal whose frequencies don't exceed a given frequency f_{max}, the sampling frequency must be at least twice that frequency. The minimum sampling frequency $f_N = 2f_{max}$ is called the Nyquist frequency. Often, in real life applications, this frequency is a-priori unknown or may vary over time. A conservative approach to guarantee that, despite a possible variation of the Nyquist frequency, the signal is accurately reconstructed, is to sample at much higher frequency (in most application, it is normal to sample at 4 or 5 times f_N). This solution is clearly disadvantageous in terms of energy consumption and overall costs, and in most cases it turns out to be unsustainable when ported on the nowadays available technology, because the power consumption needed would be too large.

Several adaptive sampling algorithms bypass the Nyquist limit by forecasting the next signal's value through a wide variety of statistical models. This is an interesting approach that sacrifices complete reconstruction in favor of a significant reduction in the number of acquired samples. The forecasted sample is compared to the actual one over time and, if the discrepancy is too large, the sampling frequency is increased and the model is updated.

On the other hand, there is a family of adaptive sampling algorithms (to which belongs the one presented in this work) whose purpose is to guarantee that the Nyquist theorem is not violated and no information is lost. The solution proposed in this thesis falls into this category. The foremost advantage of this approach is that it doesn't rely on the chosen model's accuracy (this concept will be expanded in Chapter 2). By now, it is important to underline that the Nyquist theorem will be the basement above which the whole proposed solution is developed.

A final remark on energy saving techniques at single unit level in CPS is that these approaches impact substantially on network-level consumption as well. By lowering the amount of acquired data, actually, also transmission and processing will be reduced, and the risk of network congestion would be significantly decreased.

1.3 Original contribution and structure of the thesis

In a nutshell, the activity-based adaptive sampling mechanism proposed in this thesis includes the following steps: acquiring the signal at a properly-defined frequency, filtering the signal on different frequency bands, calculating energy values of the filtered signals over sliding windows, analyzing the obtained values through a change detection test to spot changes and, in case of aliasing observation, adapting the sampling frequency to guarantee that the Nyquist limit is respected.

The filtering phase is necessary because, as anticipated before, the signal is evaluated through its energy values. We emphasize that the overall amount of

energy in a given signal could remain constant, but its internal distribution over the spectrum might change. This is the reason why a filtering step is introduced: it allows the observer to identify frequency shifts inside the frequency ranges of the signal, which would result in a different distribution of the bands energies. The determination of the number of employed filters and correspondent frequency ranges will take place during the training phase, where the spectrum is analyzed to obtain the proper band configuration. A detailed description of this analysis will be given in Section 4.2.

Filtered signals are then processed to compute their energy values, which are calculated over sliding windows (i.e., non overlapping). The idea of monitoring the observed signal in the time domain instead of the frequency domain is justified by the Parseval Theorem ([Pra10]), which states the equivalence between the integral of the squared signal over time and the one of the absolute value of the Fourier Transform in the frequency domain:

$$\int_{-\infty}^{+\infty} |x(t)|^2 \, dt \; = \; \int_{-\infty}^{+\infty} |X(2\pi f)|^2 \, df \; .$$

In fact, computing the FFT to reconstruct the signal's spectrum is quite expensive from the computational (and, hence, energy) point of view; this approach avoids this heavy computation by analyzing only the energy carried by the signal in the time domain. However, by integrating energy over time, the case of a simple shift in a frequency value lying within the frequency range could not be detected; this is the reason why a preliminary filtering step is introduced.

Next, computed energy values are submitted to a change detection test. Change Detection Tests are algorithms based on statistical hypothesis tests aimed at detecting a change in the statistical behavior of data (here, the estimated energy on each band). They can be parametric, which means that they need prior information about the probability density function of the process generating the data, or nonparamentric, which are more flexible and require weaker assumptions on the input signal ([Ali14]). In order to be independent from a specific configuration, only nonparametric tests have been used in this thesis.

In the change detection phase, the occurrence of a higher frequency value, due to the evolving dynamics of the signal, could be spotted through the aliasing phenomenon. This is a major contribution of this work since, for the first time, the aliasing phenomenon (which will be illustrated in detail in Section 3.2) is exploited to monitor possible variations of the Nyquist frequency f_N that exceed the current Nyquist limit at a given time instant.

Whenever a change in the maximum frequency is detected, the sampling frequency is adapted in order to capture the new signal's dynamics. During the adaptive phase, the frequency band division of the spectrum is exploited to identify the new maximum frequency f_{max}. All bands are split into two

parts, to obtain a finer granularity. Then, the band in which the change occurred is identified and the sampling frequency is increased of a factor c that is defined by the user. Since the sampling frequency is increased, the spectrum is wider and new bands with higher frequencies are added. The band interested by the change is recalculated and the result is compared with the previous one: if no variation is detected, the new frequency value is the result of a real frequency peak and, hence, the maximum frequency has not increased and the sampling frequency can be brought back to the original value. On the other hand, if the new band is different from the one of the previous step, we are observing an aliased peak and the sampling frequency is increased again, until no more increasing is needed and the Nyquist limit is respected.

The band-based mechanism has been designed in to avoid the calculation of the Fast Fourier Transform, which is a widely employed tool in signal's analysis when a complete reconstruction of the signal spectrum is needed. The Fast Fourier Transform (FFT) calculation would immediately give the maximum frequency component of the spectrum, but it's extremely expensive in terms of computational power. We emphasize that it would be totally unrealistic to calculate it on each new samples window, but also a calculation of the FFT for each detected change turns out to be strongly penalizing. A complete description of this process will be given in Section 4.6.

To support the theoretical intuitions, experiments and simulations have been carried on three different datasets:

1. A synthetic dataset: the analyzed signal consists in the sum of fifteen sinusoidal waves to which random noise is subsequently added. A frequency shift is inserted artificially in the middle of the dataset;

2. Blue Whale Call: a real dataset with a synthetic change. These data are the result of the Bioacoustics Research Program of the Cornwell University [Pro08];

3. Electrocardiogram of a subject affected by ventricular flutter episodes: a real dataset with real changes. Ventricular flutter is a type of tachycardia affecting the ventricles. It is characterized by a sinusoidal form of the ECG instead of the typical ECG shape. Data can be found on the Physionet data bank (more specifically the records come from the MIT-BIH Malignant Ventricular Arrhythmia Database).

Moreover, a simplified version of the whole algorithm has been carried on a real embedded system, being part of a CPS: a programmable board produced by STMicroelectronics, along with its sensor expansion. This device is provided with three sensors measuring temperature, humidity and pressure, and also an accelerometer. For this experiment only temperature values have been taken into account. Despite the fact that some simplifications of the original code have been done, the basic principle of the algorithm running

on the embedded system is exactly the same presented above. The purpose of this experiment was, in fact, to have a confirmation of the possibility to transport the proposed approach on a real system and, in this sense, this has definitely proved its feasibility.

Results of these experiments have been collected and compared with an existing algorithm, showing that the proposed approach can represent an efficient mechanism for reducing energy consumption while guaranteeing a high acquisition accuracy. To evaluate its performance, several figures of merit have been taken into account: mean detection delay, which represents the ability of the application to promptly detect a change in the environment; the sampling fraction, which is the rate between the number of acquired samples in the proposed approach over the one of a fixed-rate approach; execution time; false positive and false negative rates. Clearly, requirements on false negative rate are stricter than those on false positive rate, as a missed detected change could represent a serious issue in many applications. False negative rate should ideally be equal to zero. Considering all these parameters, the results are satisfying as they show a significant improvement in energy saving and leave several points open to future development.

The thesis is organized as follows: Chapter 2 presents the related work available in the literature, Chapter 3 provides some theoretical background on signal filtering and aliasing; Chapter 4 details the contribution, presents the ideas of this work and the conceptual considerations behind it; Chapter 5 illustrates the experiments with their results and a comparison with other existing approaches; Chapter 6 details the porting of the proposed solution on the embedded system by STMicroelectronics; while Chapter 7 presents the conclusions and possible ideas for future development.

2 State of art

This chapter reviews the most popular adaptive sampling techniques available in the literature, emphasizing characteristics, strengths and witnesses.

2.1 Adaptive sampling techniques

Adaptive sampling techniques aim at modifying the sampling frequency during signal acquisition, following the signal's dynamics, in order to reduce the sampling frequency whenever possible and, consequently, reduce the energy amount consumed by the sensing unit.

Unlike Section 1.2, whose purpose was to provide the reader with a complete frame of energy management techniques available nowadays for Cyber-Physical Systems, this section will concentrate on the subset of approaches which is most related with the one presented in this work. Therefore, activity-based adaptive sampling techniques present in the literature will be illustrated.

Some of these techniques exploit measured samples correlation to reduce the amount of data acquired from the transducer. More specifically, data may change slowly with time, thus temporal correlations (i.e. the fact that subsequent samples do not differ a lot between each other) may be exploited to reduce the number of acquisitions. A similar approach can be applied when the investigated phenomenon does not change sharply between areas covered by neighboring nodes and, clearly, both temporal and spatial correlations may be jointly exploited.

Other techniques do not consider data correlation, but their purpose is to guarantee that the Nyquist-Shannon theorem is respected. Hence, they introduce other mechanisms, like change detection tests, to monitor possible changes in the maximum frequency of the signal and adapt the sampling rate accordingly.

Among algorithms that exploit data correlation, some rely on the estimation of the next sample's values ([AHY15]), the actual reading of a sample after a given number of skipped ones, and a comparison between the actual value and the forecasted one. If the error exceeds a predefined threshold, the sampling frequency is increased. Thus, these are error-based methods and they differ in the statistical method chosen for the estimate or in the metrics chosen to evaluate algorithm performance. Section 2.2 reviews the most popular ones. Though they might prove to be effective in reducing energy consumption, they have two serious limitations: firstly, they do not

assure to monitor f_{max} properly, as a sudden change in its value wouldn't be spotted if the change duration is short enough or if the actual sample doesn't drift distinctly from the forecast one. Moreover, an inaccurate estimate of the model employed to forecast samples would compromise the efficacy of the whole approach, and this is a strong drawback for systems monitoring a nonstationary physical environment, where physical phenomena may vary over time.

Some other works ([ME10, TAGL15]) set up filter systems to sample at a frequency that violates the Nyquist theorem, but they are more focused on hardware aspects and they are, usually, application specific.

Frequency-based adaptive sampling approaches (Section 2.3) are, in turn, focused at guaranteeing that the whole signal spectrum is covered and they don't rely on a statistical model to describe the observed phenomenon. Therefore, they don't exhibit the same limitations of the error-based methods and prove to be at least as efficient as those ones, or more. This thesis belongs to this second category of mechanisms.

2.2 Error-based adaptive sampling

[JC04] proposes a Kalman-Filter (KF)-based estimation technique where each sensor can use the KF estimation error to adaptively and autonomously adjust its sampling rate within a given range. When the desired sampling rate at the sensing unit level violates the range, a new sampling rate is requested from the server, which allocates new sampling rates (provided that enough resources are available), in order to minimize KF estimation error over all the active streaming sensors. At each sensor, the Kalman Filter estimator is employed to predict the future values of a signal based on those seen so far; the authors support the utilization of the Kalman Filter because it gives satisfactory results even when the process cannot be modeled accurately, it can be easily customized to provide good results on a wide range of streaming sensor data and, finally, it produces unbiased estimates even when the incoming data have high variance.

At the sensor unit level, their model considers SI_i, the current sampling interval (i.e. the number of time units between two consecutive measurements) at source S_i; SIR_i, the range within which the sampling interval can be adjusted by the source without any intervention from the server, and SI_i^{last}, the latest value of sampling interval received from the server. $SI_i^{desired}$ denotes the desired sampling interval based on the KF prediction error; sensor S_i doesn't need to contact the server for additional bandwidth provided that $(SI_i^{last} - SIR_i/2) \leq SI_i^{desired} \leq (SI_i^{last} + SIR_i/2)$. If $SI_i^{desired}$ satisfies this equation, then SI_i takes the value of $SI_i^{desired}$. This scheme helps the source to capture unexpected data trends immediately without waiting for the server intervention, which could be delayed due to network congestion or unavailability of resources. The authors evaluated their performance through

an *effective resource utilization* metric defined as $\xi = \eta \cdot m$, where m is the fraction of messages exchanged between the source and the server over the total number of samples read by the source, and η is the mean fractional error between the actual trajectory of the monitored time series and that generated by interpolation. Experiment simulations shown that, when the input parameters of this model are properly tuned, ξ can be reduced considerably with respect to uniform sampling.

Another adaptive sampling strategy is proposed in [ZR07], where a sensor network is employed to read water level for a flood warning system. The adaptive sampling behavior of their system is specific to this application and it is the result of a co-design exercise with environmental experts; nevertheless, it is an interesting example of application. It also uses a Kalman Filter coupled with a stochastic numerical hydraulic model to predict new values of water level and, when the model-based probability of the water level exceeding a predefined threshold is less than a given value, the requirement for data transmission from the sensor node at the sink is lowered. Otherwise, the requirement for transmission will arise.

In [LCJ+09], an online adaptive sampling algorithm is proposed. Here, the authors develop a previous work from Chatterjea and al. ([CH08]) and provide a theoretical framework for it employing the Box-Jenkins approach in time series analysis, which is again an online method to adapt the sampling frequency based on the values forecast through this statistical model, and again shares the limitations of this category of algorithms. Its working principle is reported in 2.1.

The underlying idea is that if a given number x of readings have already been skipped and the next reading is close to the next forecast, then the next $x + 1$ readings can be skipped (provided that the number of skipped readings doesn't overcome a predefined upper limit); otherwise, it is necessary to resume acquiring every reading until the reading and the forecast are again close enough to each other. The key contribution is to use confidence intervals to evaluate the accuracy of the forecast values: if the confidence interval on the Box-Jenkins forecast is less than the desired value, the forecast is considered accurate enough. This can be obtained provided that the underlying statistical model describing the process is well identified which, as stated before, may not be easy to guarantee in many real life applications. Anyways, experiments simulation have shown that the number of acquired samples can be reduced remarkably.

[GSBH11] describes the so called Exponential Double Smoothing-based Adaptive Sampling (EDSAS) and stresses the fact that the proposed technique has a scarce computational load and does not need an offline training phase. Its principle is quite similar to that of [LCJ+09], because it still relies on a forecast value which is compared to the real one and, based on the

Algorithm 2.1 Algorithm 0 of [LCJ+09]. $CSSL = CurrentSkipSampleLimit$, $SS = SkipSamples$, $MSSL = MaximumSkipSamplesLimit$ and ϵ is a user-specified error tolerance threshold.

```
1: Collect b samples
2: CSSL = SS = 0
3: while (1) do
4:     Acquire 1 reading
5:     Use this new reading and the previous one to interpolate samples
       skipped in the previous round, if any
6:     Make 1 forecast
7:     if (|reading-forecast|<ε) then
8:         SS = CSSL = min(CSSL + 1, MSSL)
9:     else
10:        SS = CSSL = 0
11:    while (SS > 0) do
12:        Skip 1 reading
13:        SS = SS − 1
```

result of this comparison, the step size between two consecutive samples is updated until it reaches a maximum value S_{max} representing the maximum delay tolerated by the specific application. The statistical method employed to forecast samples is, this time, the Wright's extension to Holt's method. In addition to this process, once S_{max} has been reached, a change detection mechanism based on exponentially weighted moving averages (EWMA) [WaJR+05] is introduced to minimize the false negative rate. The method has been tested on real urban traffic CO_2 pollution level dataset and compared with previous approaches ([LCS06]), showing that better sampling fraction against false negative ratio can be achieved through this process.

Another interesting work which presents the adaptive sampling problem in an original approach is [ZRDH15], where an analogy between a WSN and the human endocrine regulation system is built. Exactly like a human body receives hormone information to react to environmental changes, EASA, the proposed adaptive sampling algorithm, produces different kind of hormone information to regulate dynamically sensor nodes working status and sampling frequency. When the observed phenomenon changes quickly, the nodes make trophic hormone and increase sampling frequency. On the other hand, when high sampling rate is useless due to scarce activity, inhibitory hormone is produced and the sampling frequency is lowered. The method to determine whether there is a rapid change in the environment conditions still relies on next sampling forecasts but this time linear regression is employed and, if the relative incremental ratio between the forecast value and an average of n previous samples exceeds a given threshold, sensor activity

must be increased. Afterwards, a redundant-nodes dormancy mechanism is set into place, completing the energy management framework. Simulations show that EASA performs very well in terms of network lifetime and data accuracy when compared to fixed-rate approaches.

2.3 Frequency-based adaptive sampling

In [AAG+07], a frequency-based adaptive sampling approach is proposed. It is based on the online adaptation of the sampling frequency provided that a change in the maximum frequency of the monitored phenomenon occurs. Change detections methods have been widely discussed in the literature and here a modification of the CUmulative SUM (CUSUM) change detection test [BN93] is employed. As we have seen in Chapter 1, the maximum frequency contained in a signal determines the minimum sampling frequency necessary to reconstruct it properly. Frequency f_{max} isn't always available and changes over time in nonstationary environments. Hence, the Nyquist frequency f_N changes as well and it has to be adjusted to avoid oversampling. The proposed algorithm initially estimates f_{max} through a Fast Fourier transform by using the first W acquired data, which are assumed to be generated by a stationary process. The initial sampling frequency is $f_c = c\overline{f}_{max}$ with c larger than 2, according to Nyquist theorem. Then, maximum increment and decrement of $\overline{f}_{max}$ which can be tolerated before changing the sampling frequency are defined as follows:

$$f_{up} = \min\left\{(1+\delta)\cdot\overline{f}_{max}, \frac{f_c}{2}\right\};\ f_{down} = (1-\delta)\cdot\overline{f}_{max}$$

$\delta \in \mathbb{R}^+$ is a confidence parameter representing the minimum detectable frequency change desired by the user. A change is detected if the current maximum frequency f_{curr} overcomes one of these thresholds for h consecutive samples:

$$th_{up} = \overline{f}_{max}(1+\delta/2);\ th_{down} = \overline{f}_{max}(1-\delta/2)$$

An example of the mechanism working principle is shown in Figure 2.1, while 2.2 illustrates the complete mechanism. When a change is detected the sampling frequency is modified as follows: "if ($\left|f_{curr} - f_{up}\right| < \left|f_{curr} - \overline{f}_{max}\right|$) for h consecutive samples or if ($|f_{curr} - f_{down}| < \left|f_{curr} - \overline{f}_{max}\right|$) for h consecutive samples, then the new sampling frequency is $f_c = cf_{curr}$ ". Basically, the two defined frequencies f_{up} and f_{down} allow to detect drifts within their range and enlarge the spectrum observed by the application. A frequency peak appearing above f_{up} or below f_{down} could not be detected by this method.

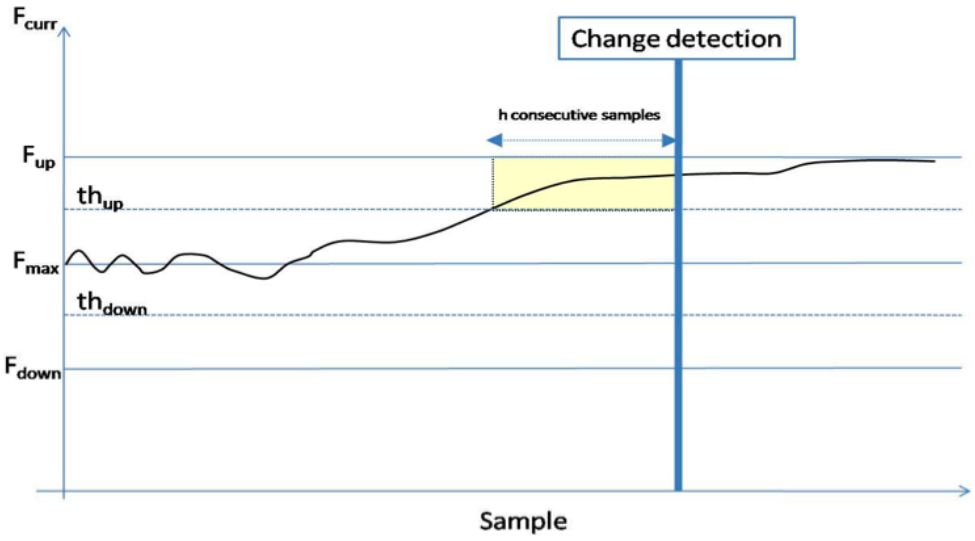

Figure 2.1: [AADFR09] Detection of a change in the maximum frequency

Algorithm 2.2 Adaptive Sampling Algorithm of [AAG+07] (c, δ, h)

1: Store the W initial samples coming from the process in $Dataset$
2: Estimate $\overline{f}_{max}$ on $Dataset$ and set $f_c = c\overline{f}_{max}$
3: Define $f_{up} = \min\left\{(1+\delta)\cdot\overline{f}_{max}, \frac{f_c}{2}\right\}$; $f_{down} = (1-\delta)\overline{f}_{max}$
4: $h_1 = 0$ and $h_2 = 0$; $i = W + 1$
5: **while** (1) **do**
6: Acquire the i-th sample and add it to $Dataset$
7: Estimate the current maximum frequency f_{curr} on the sequence $Dataset$(i-W+1,i)
8: **if** $(|f_{curr} - f_{up}| < |f_{curr} - \overline{f}_{max}|)$ **then**
9: $h_1 = h_1 + 1; h_2 = 0$
10: **else if** $(|f_{curr} - f_{down}| < |f_{curr} - \overline{f}_{max}|)$ **then**
11: $h_2 = h_2 + 1; h_1 = 0$
12: **else**
13: $h_1 = 0; h_2 = 0$
14: **if** $(h_1 > h) \,\|\, (h_2 > h)$ **then**
15: $f_c = cf_{curr}$
16: $f_{up} = \min\left\{(1+\delta)\cdot\overline{f}_{max}, \frac{f_c}{2}\right\}$
17: $f_{down} = (1-\delta)\overline{f}_{max}$
18: $\overline{f}_{max} = f_{curr}$

A snow monitoring application for avalanches forecast is employed as a case study. This is a very good example of power-hungry sensors deployed in a hostile environment, where battery replacement is quite difficult and ex-

pensive. The sensor is a "multi-frequency capacitive measuring unit designed to be embedded in a remote wireless measuring system. It is composed of a probe, a main multi-frequency injection board capable of measuring capacity at different frequencies and a wireless unit to be left on the mountain (...). At each sampling cycle the snow sensor provides measurements of snow capacitance at two different frequencies that have been proved to differentiate water from air, snow and ice. At the same time a second sensor measures the ambient temperature. All acquired information is packed in a single message and sent over the wireless channel. For each measurement the electronic injection board of the snow sensor makes several procedures (calibration, electrode pre-charging, charge sharing) in a cyclic way to obtain a reasonably stable and reliable measure. This activity makes the sensor very energy consuming. When the duty cycle mechanism substitutes the fixed sampling approach an immediate energy saving arises.".

The experiments conducted have proved that, following this approach, the number of measured samples can be reduced to a great rate with respect to the traditional fixed-rate algorithm.

In [BWCW13] the "event-sensitive adaptive sampling and low-cost monitoring (e-Sampling)" is presented, which focuses on decentralized computing in WSNs by avoiding transmission to the sink as much as possible. The field of application is structural health monitoring, in which high frequency events are not rare. Their idea is to approach energy consumption reduction at the sensor level in two different "stages", which are to some extent comparable with the four steps of this work, as it will be explained in Chapter 4. Firstly, each sensing unit has "short and recurrent "bursts" of high-rate sampling, and samples at a much lower rate at any other time.", as explained below.

Basically, four parameters are defined: D_h, the duration of each burst of high-rate sampling, R_h, high sampling frequency, D_l and R_l (duration and sampling of low-frequency or adjusted sampling, respectively). $D_h + D_l$ give the duration of the so called "Rate and Interval Adaptation" cycle. If, during D_h, the high frequency content F_h reveals to be important (which means that it overcomes a predefined threshold ω), then the minimum required rate is increased. If, in some successive adaptation cycle, the percentage of high frequency has decreased, the sampling rate is set back to the lower value.

At each Adaptation cycle, Wavelet Packet Decomposition is employed to estimate F_h([Mit01]). This methods recursively applies low-pass filters and high-pass filters, building a tree of sub-signals that sum up to the original one. At each tree level r, $s_r^1(t)$, ..., $s_r^b(t)$ signals are obtained, with $b = 2^r$, and used to estimate b_h, the maximum frequency contained in the signal. Here, band division is employed to estimate f_{max} only, there is no mechanism of change detection over the bands. 2.3 illustrates how this is managed (R_c is the current sampling rate resulting from the selection).

In e-Sampling the indication of the presence of an event is studied at the

sensor level and data are transmitted to the sink only if a change actually occurs. Experimental evaluation outlines that, when both stages (adaptive sampling and decentralized event indication) are used, e-Sampling outperforms other algorithms in term of energy consumption, included [ZR07].

Algorithm 2.3 Sampling Rate Selection of [BWCW13]

1: $U = 0$ ▷ There is an update of sampling rate;
2: $s_0^0(t) = S(t)$
3: **while** $(r < R \&\& U = 0)$ **do**
4: Compute $s_{r+1}^{2b}(t); s_{r+1}^{2b+1}(t)$
5: **if** $(thereisF_h ins_{r+1}^{2b+1}(t)) > \omega$ for any t **then**
6: $s_r^b(t) = s_{r+1}^{2b+1}(t)$ ▷ there is possibly a situation of the presence of a physical event
7: **else if** $(s_{r+1}^{2b}(t)) > \omega$ for any t **then**
8: $s_r^b(t) = s_{r+1}^{2b}(t)$ ▷ there is possibly a situation of the presence of a physical event
9: **else**
10: $U = 1$
11: $b = b_h$
12: Compute $R_m = c(R_h/2^r)(b_h + 1)$
13: Set $R_c = R_m$

2.4 Summary of the available approaches

A summary of the reviewed methods is reported in Table 2.1.

This thesis proposes a frequency-based adaptive sampling mechanism which addresses power-hungry sensors without being application-specific. This mechanism has several points in common with previous approaches, nevertheless, up to our knowledge, there is no other approach in the literature which considers all the elements presented in this work at once; that is, exploiting aliasing to spot changes in the maximum frequency, dividing the spectrum in different bands, estimating energy values on each band and monitoring this measure through a change detection test.

Table 2.1: Adaptive Sampling algorithms summary, EB: Error-Based, FB: Frequency-Based

Name	**Family**	**Assumptions**	**CDT**	**Ref**
Kalman filter	EB	input parameters must be properly tuned	No	[JC04]
FloodNet	EB	co-designed for hydraulic application	No	[ZR07]
Box-Jenkins	EB	underlying statistical model must be accurately identified	No	[LCJ+09]
EDSAS	EB	underlying statistical model must be accurately identified	Yes	[GSBH11]
EASA	EB	sensors seen as endocrine system	No	[ZRDH15]
AliRove	FB	frequency doesn't exceed range (f_{down}, f_{up})	Yes	[AAG+07]
eSampling	FB	frequency doesn't vary too quickly during low-frequency sampling	No	[BWCW13]

3 Theoretical Background

A brief theoretical explanation of the aliasing phenomenon is reported here to provide the reader with all the necessary information about the topic. This phenomenon is of particular importance to fully understand the solution proposed in this thesis, as it constitute one of the main aspects of differentiation from the other approaches existing in the literature. In fact, we concentrate on the aliasing phenomenon because it allows to spot a change in the maximum frequency of a signal without necessarily calculating the Fast Fourier Transform. The remarkable impact on the overall performance entailed by avoiding the Fast Fourier Transform calculation will be detailed in Chapter 5.

3.1 Aliasing and peak reflection

A universally known theorem on sampling theory is the Nyquist-Shannon theorem:

Theorem: In order to sample correctly a generic band-limited signal, it is sufficient to sample it at a frequency f_s at least larger than twice as the maximum frequency preset in the signal. This frequency is also called the Nyquist frequency and indicated as f_N.

A generic signal spectrum is symmetric, and usually it is represented up to half of the Nyquist frequency: $f_{max} = \frac{f_N}{2}$.

What happens if the sampling frequency is lower than the indicated one? The well-known phenomenon of aliasing occurs: when the Fourier Transform is calculated and the spectrum figure analyzed, peaks representing frequencies higher than the Nyquist one are "reflected" in the spectrum as if a lower frequency was found in the signal.

The position of the aliased peak at the reflected frequency f_r can be found through this formula, given the original frequency f_i larger than $f_N/2$:

$$f_r = |NearestIntMult(f_i, f_s) - f_i| \tag{3.1}$$

where $NearestIntMult(f_i, f_s)$ stands for the integer multiple of f_i which is nearest to f_s. For example, if $f_s = 30$ Hz and $f_r = 97$ Hz, $NearestIntMult(f_i, f_s) = 90$ Hz and $f_r = 7$ Hz.

If we want to know whether the frequency of a peak is real or aliased, we need to increase the sampling frequency: if we are observing a reflected peak, once the sampling frequency is increased, this peak at f_r will move rightward. Otherwise, if the peak is a real one, it will keep its position.

This phenomenon is strictly related to sampling frequency reduction: we can imagine a situation where a dynamic system operates most of the time under a given frequency, but sometimes a higher frequency comes into play and the corresponding peak appears (as, for example, during tachycardia episodes, as it will be illustrated in Section 5.1). To reconstruct properly the whole spectrum with a fixed-sampling frequency approach, f_s should be at least as large as twice the Nyquist one. On the other hand, high sampling frequency would be useful only the for a short amount of the time: it would be much better in terms of energy consumption to sample at the lower frequency, and observe the appearance of higher frequency peaks through aliasing. If needed, the sampling frequency could then be increased to find the actual value of the peak.

When the sampling frequency is increased, if a peak maintain its position, we have measured a frequency component in the signal below $f_N/2$ (which could appear due to environment non-stationarity), otherwise, in case of aliasing, the peak will move rightward.

3.2 Aliasing: an example

As introduced in the previous section, aliasing is the phenomenon also known as *spectrum reflection*. When a signal is analyzed in order to reconstruct its spectrum, it has to be sampled at least at the Nyquist frequency f_N, which is twice the maximum frequency present in the signal.

Since the spectrum is symmetric with respect to half of the Nyquist frequency, only the first half is represented. An example of the spectrum of a generic sine wave combination signal is shown in Figure 3.1. The original signal is

$$x(t) = \sin(2\pi f_1 t) + \sin(2\pi f_2 t)$$

with $f_1 = 10$ Hz and $f_2 = 18$ Hz and sampling frequency $f_s = 50$ Hz (the Nyquist frequency is, hence, $f_N = 25$ Hz).

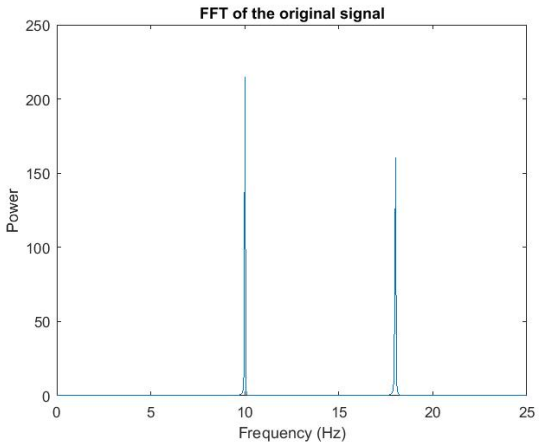

Figure 3.1: The Fourier Transform of a very simple signal composed of two sine waves

If a new sine wave is added with frequency included in the interval $(0, f_N)$, the corresponding peak will appear, as shown in Figure 3.2a (here the peak was set to 6 Hz). But what happens if the frequency of the new sine wave belongs to the interval (f_N, f_s) or $(f_s, +\infty)$?

Recalling the formula 3.1, the position of the reflected peak can be easily found.

In this example, two aliased peak are shown:

1. f_i = 44 Hz, which belongs to the interval (f_N, f_s). Its position is, thus, again 6 Hz. The corresponding Fourier transform is shown in Figure 3.2b.

2. f_i = 106 Hz, which belongs to the interval $(f_s, +\infty)$. Its position is exactly the same as the former one: 6 Hz. Figure 3.2c reports this peak, which is indistinguishable from the previous cases.

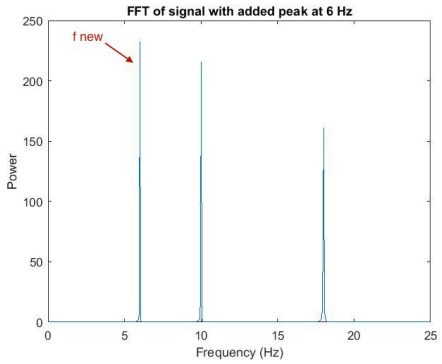

(a) FFT with peak at 6Hz

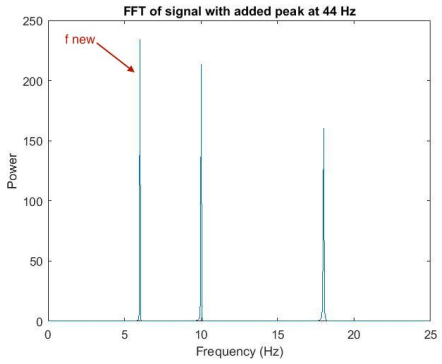

(b) FFT with peak at 44Hz

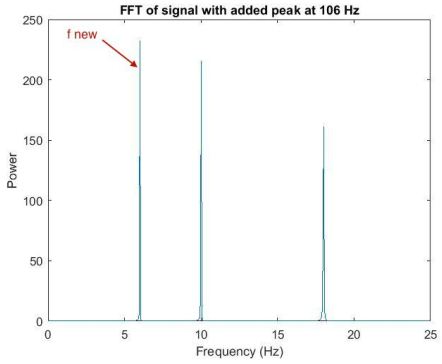

(c) FFT with peak at 56Hz, visible at 6 Hz due to aliasing

Figure 3.2: FFT of three signals with different added peaks, only the firs one is a real peak

A reflected peak originated by aliasing cannot be distinguished from a real one unless other operations are carried out on the signal. More specifically, sampling frequency has to be increased until the aliased peak stops moving rightward, as shown in figure Figure 3.3, where sampling frequency has been

brought to 100 Hz.

These preliminary analysis are in total agreement with real datasets illustrated in Chapter 5.

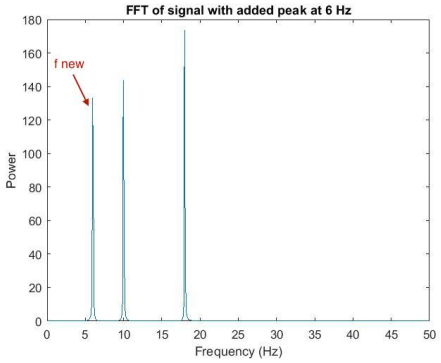

(a) FFT with peak at 6 Hz

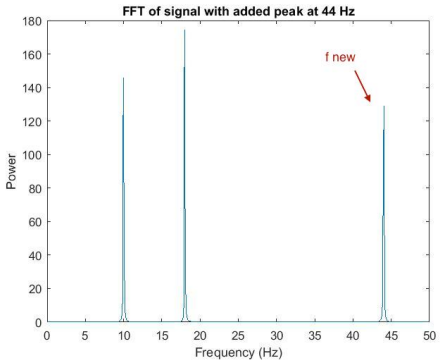

(b) FFT with peak at 44 Hz

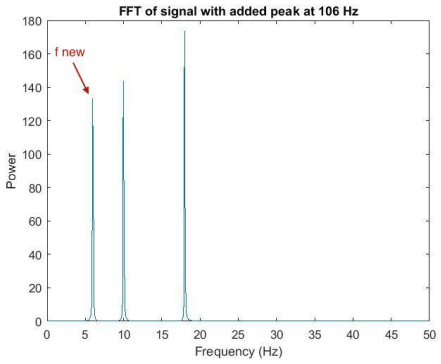

(c) FFT with peak at 56 Hz

Figure 3.3: FFT of the three signals sampled at 100 Hz. Peak from Figure 3.3a is still at 6 Hz, peak of Figure 3.3b has moved to its actual position, while peak of Figure 3.3c is still aliased.

4 Proposed approach

4.1 The general idea

This thesis presents an activity-based adaptive sampling mechanism whose purpose is to reduce the energy consumption in sensing units of Cyber-Physical Systems. It differentiates from other approaches present in the literature because, instead of forecasting sample's values and then comparing the prediction to the actual values, it samples at a rate that satisfies the Nyquist theorem and monitors possible changes through a change detection test, eventually exploiting aliasing to identify the maximum frequency of the signal.

The whole mechanism proposed by this work can be summarized into the following steps: there is a preliminary training phase, in which the spectrum is analyzed in order to obtain a complete representation of the acquired signal (maximum frequency and energy distribution over the spectrum range) and to train the change detection test. This first step is fundamental to set some of the most important parameters for the experiment, in particular, the number of bands in which the spectrum will be divided and their bandwidth. More details will be given in Section 4.2.

Then, accordingly to the results of the previous step, the spectrum is divided into the defined number of frequency bands, in order to gain visibility on small changes and to monitor intra-spectrum peak shifts. The signal coming from each band results into an independent signal, whose energy is calculated over sliding windows according to the following formula, where W indicates the window's width, j is the window index, $b = 1, ..., B$ is an index indicating the band number and $\widehat{x}_{fil,b}(t)^2$ is the signal filtered on the b-th band:

$$E_{j,b} = \int_W \widehat{x}_{fil,b}(t)^2 dt \tag{4.1}$$

This is the second step, in which a matrix E_j of B energy values $E_{j,b}$ is obtained. The purpose of this phase is to evaluate sequence $E_j = [E_{j,1}, ..., E_{j,B}]$ over time and monitor any possible change. This is achieved through a change detection test, which can be univariate (i.e., it considers each band independently from the others) or multivariate (it jointly considers all bands). Whenever a change is detected, the signal has to be analyzed in order to find out whether a change in the maximum frequency f_{max} has occurred or not. Therefore, the band division is exploited to identify which frequency range

is responsible for the detection and the possible presence of the aliasing phenomenon. If the maximum frequency has actually increased or decreased, the sampling frequency is adapted, otherwise, its value is not changed.

The final and complete framework that we obtain is summarized in Figure 4.1.

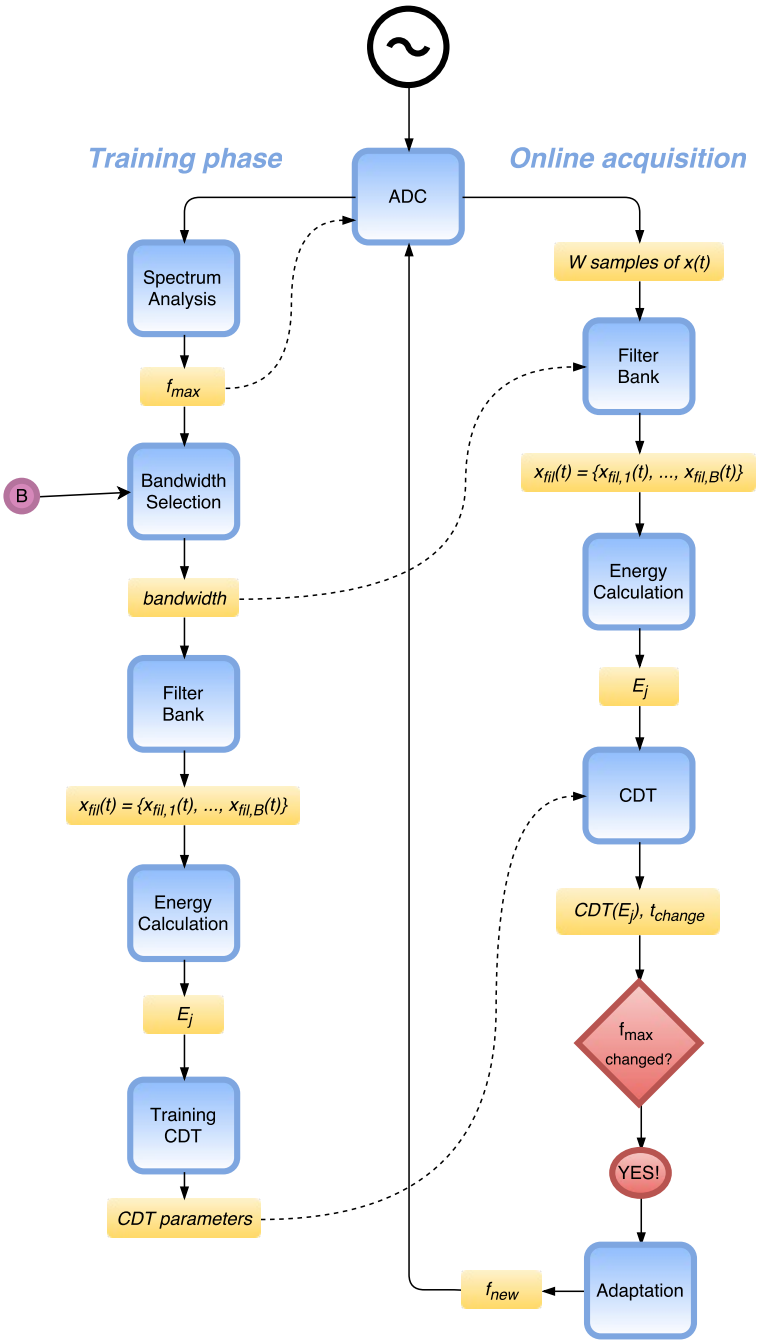

Figure 4.1: The general idea of the adaptive sampling mechanism

4.2 Training phase

A preliminary training phase is necessary to characterize of the signal's spectrum distribution. In fact, its purpose is to find the value of f_{max} and to choose the band number B along with the band configuration. This phase must take place in stationary conditions, which means that no change is occurring in the signal and the observed phenomenon is not evolving, because the energy values obtained through the filters will be used to train the considered CDT. Thanks to this, the CDT will be able to spot any variation from stationary conditions, as the picture of the signal's behavior obtained here will be held as a reference throughout the whole process.

Hence, the first Z samples are acquired at a very high frequency (if the physical phenomenon under observation has already been studied before, a frequency value of four or five times larger than the theoretical Nyquist one can be considered), and the Fast Fourier Transform is calculated. This is the only FFT calculation that will be carried out in the whole process. Once the spectrum is reconstructed, the initial value of the sampling frequency is set to $f_s = cf_{max}$, with $c \geq 2$, in order to respect the constraint given by the Nyquist theorem.

At this point, the best possible band division for each specific application-scenario has to be determined. Again, the spectrum division into frequency bands has been introduced to avoid the continuous calculation of the FFT for the detection of any possible signal change. In fact, as explained in Chapter 3, signal transmitted by a sensor in a non-stationary environment can present changes of different types. For instance, a new frequency could appear, and its value could be observed as a new peak (real or reflected), but also a simple frequency shift could occur. The latter is the most critical case: due to the fact that an integral over a window of fixed size is calculated, if a frequency peak simply shifts without changing its amplitude nor moving outside the spectrum boundaries, this change is invisible to the observer. Nevertheless, dividing the spectrum into frequency bands leverages this phenomenon because, if the bands are properly distributed, a moving peak is likely to leave a band and enter another one.

Of course there will be a trade-off effect in the choice of the number of bands: by increasing the number bands, we obtain a finer granularity and a pure shift of a peak is more easily detectable. On the other hand, a large number of bands entails a possible increase of costs and false positives, as it will be illustrated by the experiments of Chapter 5. In the current version of the proposed mechanism, the number of bands (from now on, parameter B) isn't computed automatically, but its choice is left to the user and will depend on the signal's spectrum amplitude and distribution observed in this training phase.

Another important parameter that has to be chosen carefully is the *bandwidth:* a first raw band division would simply split the spectrum range into

B equally wide bands. This choice would be reasonable for a signal whose frequencies are uniformly distributed on the spectrum range, like those of Figure 4.2, but it wouldn't be accurate for a signal whose frequencies are, for example, concentrated in the first half of the spectrum. Therefore, bandwidth has to be chosen such as the energy is uniformly distributed among bands. In fact, it wouldn't be advantageous to have one band which contains most of the signal, because this would invalidate the possibility of detecting frequency shifts. Hence, bands do not necessarily need to have the same dimension, as those of Figure 4.3. Here, a division into equally wide bands like the one of the previous example would lead to a band containing only noise (30 - 40 Hz), thus, a non-uniform division could be considered.

Otherwise, if the spectral energy is concentrated into the first part of the spectrum, another intelligent division would split the energy-denser part into equally wide bands and leave a large, almost empty band for higher frequencies. Actually, for every practical case of real world physical phenomena, a new frequency possibly appearing at a given time instant would likely be of the same order of magnitude of the other frequencies contained in the signal. Therefore, having a single, large empty band for high frequencies rather than several high-frequency empty bands should guarantee the detection of variations and avoid the unnecessary introduction of more bands.

Furthermore, the bandwidth might change along the process, following signal's evolution and possible redistribution of the signal energy: thus, an adaptive mechanism intervenes in this phase as well once a variation in f_{max} has been detected.

Once the initial sampling frequency f_s is identified and the band division has been designed, the online acquisition will begin, obtaining the sampled acquisition $\widehat{x}(t)$, $t = 1, 2, ...$

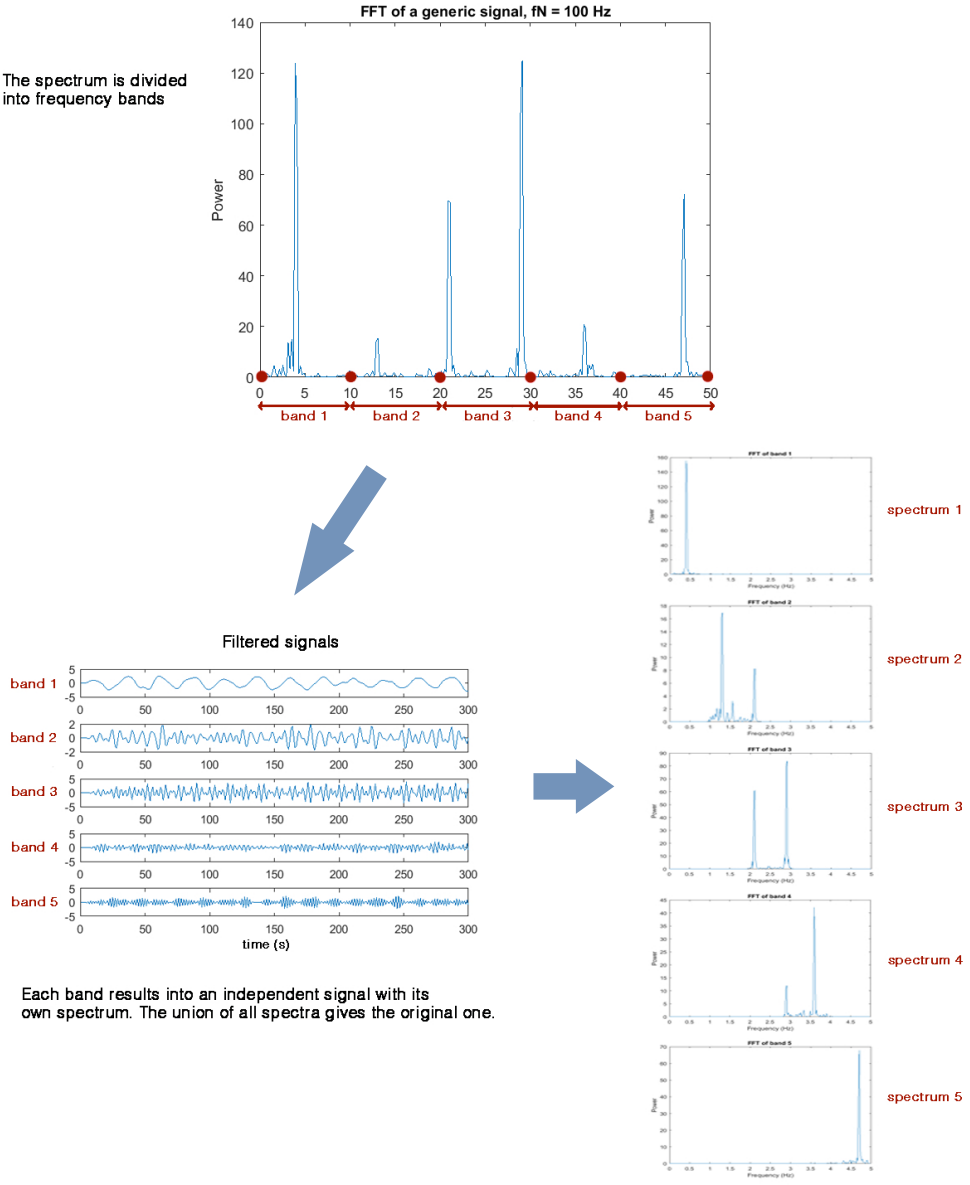

Figure 4.2: Band division and filtering process.

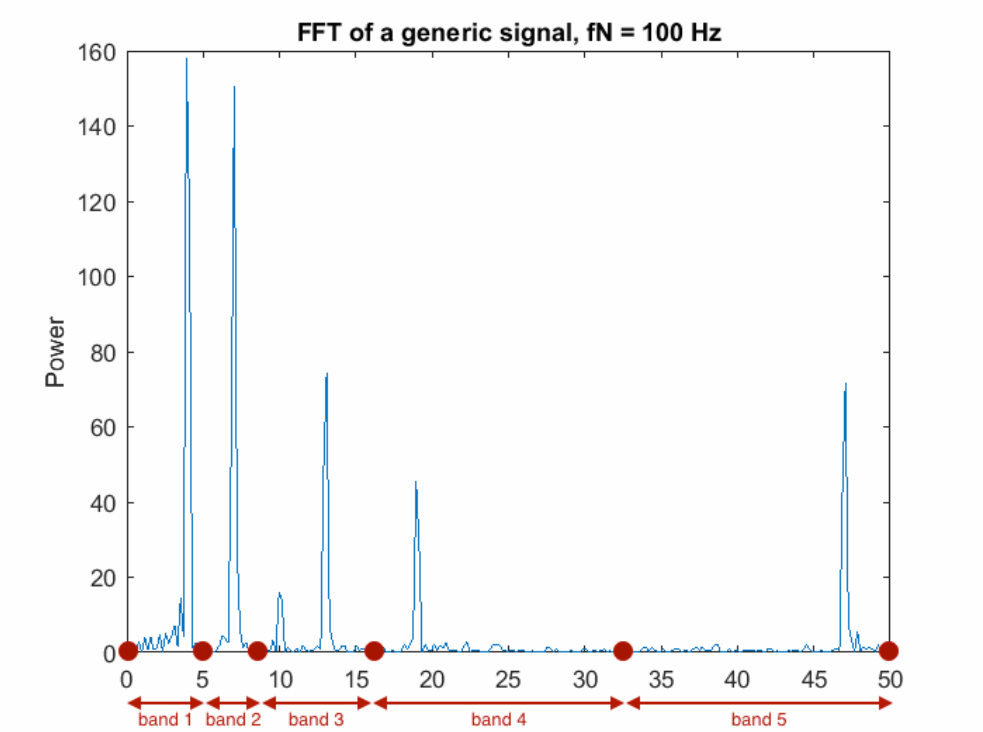

Figure 4.3: Non-uniform band division.

4.3 Filtering the spectrum

Once online acquisition begins, the signal undergoes the filtering phase, where B filtered signals are obtained, according to the filters bank designed during the training phase.

Given the sampled sequence $\widehat{x}(t)$ and B bands, the signal has to be filtered and, subsequently, integrated over sliding windows of width W. Hence, each W samples the signal will be filtered obtaining a matrix $\widehat{x}_{fil}(t) = \{\widehat{x}_{fil,1}(t), \ldots, \widehat{x}_{fil,B}(t)\}$ of B sequences $\widehat{x}_{fil,b}(t) = \{\widehat{x}_{fil,b}(1), \ldots, \widehat{x}_{fil,b}(W)\}$, whose dimension is $B \times W$, which will be used to compute the energy values of the signal.

As introduced above, this step is introduced to avoid the FFT calculation and monitor the signal in the time domain instead of the frequency domain. This approach is justified by the Parseval Theorem, which states the equivalence between the sum (or integral) of the square of a function and the sum (or integral) of the square of its transform:

$$\int_{-\infty}^{+\infty} |x(t)|^2 \, dt = \int_{-\infty}^{+\infty} |X(2\pi f)|^2 \, df$$

where $X(\omega) = \mathcal{F}_\omega\{x(t)\}$ is the continuous Fourier Transform of the signal $x(t)$ and $\omega = 2\pi f$ indicates the frequency in radians per second.

4.4 Energy computation

Recalling Equation 4.1, in order to calculate energy values of the signal, its squared power has to be integrated over fixed-width sliding windows of amplitude W. The length of the window is expressed in number of samples that are involved in each integral. Hence, each element $\widehat{x}_{fil,i}(t)$ of vector $\widehat{x}_{fil}(t)$ obtained in Section 4.3 is integrated on its W values ($i = 1, ..., B$ indicates the band number).

The result is the vector $E_j = \{e_j(\widehat{x}_{fil,1}(t)), ..., e_j(\widehat{x}_{fil,B}(t))\}$ of B elements, where j is the window index (meaning that we have already acquired $(j+1)W$ samples), and

$$e_j(\widehat{x}_{fil,i}(t)) = \int_{jW}^{(j+1)W} \left|\widehat{x}_{fil,i}(t)\right|^2 dt.$$

Vector E_j will be submitted to the change detection test. The change detection test will receive a new element E_j every time a new window of W samples of the original sequence is available, and it will test it against the previous values E_k, with $k = 1 \dots j-1$. Hence, the CDT operates on a matrix of size $j \times B$, where each energy vector E_j is stored: $E = \{E_1, ..., E_j, ...\}$.

Figure 4.4 clarifies these steps for signal of band 5 of Figure 4.2.

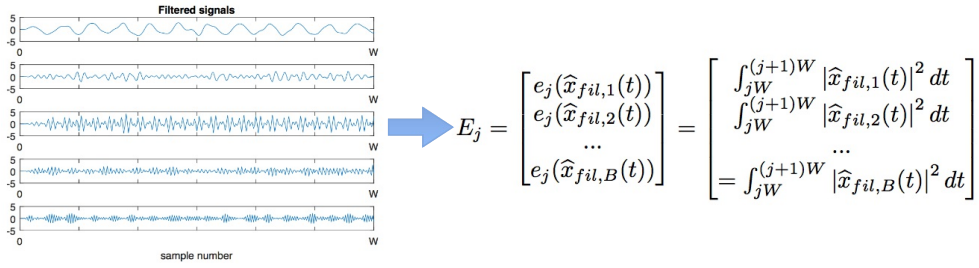

Figure 4.4: Integration process

4.5 Change Detection Test

Once the energy is calculated, a change detection test is applied to the sequence $E = \{E_1, ..., E_j\}$ obtained through the previous steps. Each W samples, a new element E_j of the sequence will be taken into account.

Change detection tests (CDTs) are powerful tools for concept drift detection, i.e., the detection of a change in the stationary conditions of the

environment. Actually, during the training phase described in Section 4.2, the signal is analyzed in stationary conditions and the computed energy values represent the stationary state. Those values are used to train the CDTs and, subsequently, to detect any possible variation from that situation.

There exist several different implementations of change detection tests ([Ali14]). In this work, two specific change detection tests have been considered. The first one is the univariate Intersection of Confidence Intervals Change Detection Test described in ([ABR11c, ABR11a, ABR11b]) for the experiments with synthetic data and real data with synthetic change point (for details, see Section 5.1). This test assumes that submitted data stream is composed of i.i.d and Gaussian distributed features, hypothesis that is satisfied provided that integration intervals, or windows, are large enough to flatten the dynamics of the signal, which often present oscillating dynamics in these kind of application.

In addition, for the experiment conducted with a real dataset, a multivariate CDT has been used ([Tac15]). The main difference between the previous test and this one, is that the former is a univariate test whereas the latter is multivariate. Multivariate tests consider all the features at once, instead of running a single test for each feature independently from the others. Features here are the energy values calculated on each band and the use of a multivariate test should allow us to reduce false positive rate. Moreover, this test is trained to work even on small input sequences and this test was particularly appropriate for the real dataset where some sequences are actually very short.

Therefore, for each window j, the new value E_j is tested against vector $E = \{E_1, \ldots, E_{j-1}\}$ looking for changes. There are two possible outputs:

$$\begin{cases} CDT(E_j) = 1 & \rightarrow \quad change\ detected \\ CDT(E_j) = 0 & \rightarrow \quad no\ change\ detected \end{cases}$$

If the output is equal to one, a change has been detected in the j-th window and further inspection is necessary to understand the nature of the change. Otherwise, the sampling mechanism continues with the acquisition of the next samples window.

4.6 Adapting sampling frequency

After the completion of the training phase, the current value of the sampling frequency is f_s. In this configuration, the appearance of a new peak at frequency f_{new} (possibly higher than $f_s/2$) would imply a change in the energy value and, hence, it would be detected by the change detection test. The new peak could be the result of aliasing and, thus, the sampling frequency should be increased, or it could be the actual value of a new frequency component, which has no impact on the sampling frequency f_s. Therefore, a mechanism to distinguish these two scenarios is here introduced.

First of all, it is necessary to identify which band has raised the alarm of an event. In case of multiple univariate CDTs, this is pretty straightforward, because only the corresponding CDT detected a change. On the other hand, the multivariate test considers all bands at once and it is not possible to identify which feature provoked the detection. Therefore, the energy values of the last window are compared with the mean values of energy before the detection, and the band showing the maximum difference is identified as the changing one. Given that the change at time t_{change} occurred in the j-th window, we have:

$$b_{change} = \arg\max_b \left| avg(E_{0,b}, \ldots, E_{j-1,b}) - E_{j,b} \right|$$

In order to inspect more accurately the nature of the change, a heuristic approach is employed. Actually, an exact solution would again need the FFT calculation, but it can be avoided through the use of a heuristic which, here, is justified by the fact that an aliased peak will necessarily move from its position (see Chapter 3 for more details).

Algorithm 4.2 details this heuristic: after a detection, each band is split in two half, hence doubling the bands number. The energy vector E, hence, doubles its size as it will have twice the rows it had before. Then, the sampling frequency f_s is increased of a fraction h of its value, obtaining $f_{curr} = f_s + hf_s$. This increase has the effect of enlarging the spectrum and, hence, new bands are added for the higher frequencies (line 6 of Algorithm 4.2). The number and width of these new bands depend from the previous average bandwidth, in order to keep a homogeneous division. At this point, the changing band is again calculated (line 7 of Algorithm 4.2):

$$bandIndex_{new} = \arg\max_b \left| avg(E_b(1, \ldots, t_{change})) - avg(E_b(t_{change}, \ldots, t_{curr})) \right|$$

If its index doesn't vary ($bandIndex_{new} = bandIndex_{old}$) for three iterations, it means that the new peak is a real peak and we can put f_s back to its original value, which will be the one set three iterations before (thus, in line 12, f_s is put back to $f_{curr} = f_s + (k-3)hf_s$). On the other hand, if the new band index is greater than the previous one, we are observing aliasing phenomenon and the process has to be repeated until this shift in the frequency stops and a new value of the sampling frequency f_{curr} is obtained. The choice of stopping the frequency increase after three iterations instead of one has been made in order to avoid that a slowly moving aliased peak (which remains in the same band for more than one iteration even if it is actually moving) is mistaken for a real one.

4.7 The algorithm

To sum up, incoming signals are sampled at the lowest possible frequency allowed by the Nyquist limit, to save energy and reduce power consumption. Each signal is filtered through a configurable number of band-pass filters, covering the whole spectrum, and for each filtered signal the energy is computed on fixed-width windows. Once energy sequences are obtained, these are monitored by a change detection test which reveals changes in the energy values over time.

The process is illustrated in Algorithm 4.1, where line 1 to 3 refer to the training phase. Here, the initial sampling frequency is set and, thanks to the spectrum analysis, the number B of filtering bands along with the bandwidth are chosen.

Online acquisition starts with line 4. When W consequent samples are acquired (line 6), for each band (from band 1 to band B) the signal is filtered as indicated in Section 4.3 and energy values E_j are calculated according to Section 4.4 (lines 8 and 9).

In line 10, the CDT is activated and, in case of detection, the estimate t_{change} of the time instant the change started is found. If the output of the CDT is equal to one, the procedure for the adaptation of the sampling frequency is activated and its returned value is stored into f_s.

The procedure for the sampling frequency adaptation (Algorithm 4.2) works as explained in Section 4.6. Variable $flag$ serves as a condition for the while loop and indicates whether we have to keep increasing the sampling frequency or we have to stop; parameter k is used to increase the sampling frequency of the same quantity at each iteration.

Algorithm 4.1 Adaptive Sampling Algorithm (W, d)

1: Take the Z initial samples of the process and store them into $dataset$;
2: Estimate f_{max} on $dataset$ and set $f_s = cf_{max}$;
3: Define band numbert B and band-pass ranges $bandwidth$;
4: **while** (1) **do**
5: Acquire the n-th sample and add it to $dataset$;
6: **every** W **samples:**
7: **for** $(i = 1 : B)$ **do**
8: Filter the last W samples of $dataset$ on i-th band-pass range $\widehat{x}_{fil,i}(t)$;
9: Calculate the j-th energy value $E_j(b) = \int_{jW}^{(j+1)W} \left|\widehat{x}_{fil,b}(t)\right|^2 dt$;
10: t_{change} = call procedure for Change Detection Test on E_j;
11: **if** $(CDT(E_j) = 1)$ **then**
12: f_s = Adapt Sampling Frequency(E, t_{change});

Algorithm 4.2 f_{new} = Adapt Sampling Frequency (E, t_{change})

Split each band in two half;
$bandIndex_{old} = \arg\max_b \left| avg(E_b(1, ..., t_{change})) - avg(E_b(t_{change}, ..., t_{curr})) \right|$;
$count = 0, flag = 1, k = 1$;
while (flag) **do**
    $f_{curr} = f_s + khf_s$;
    $B = B + n$;
    $bandIndex_{new} = \arg\max_b \left| avg(E_b(1, ..., t_{change})) - avg(E_b(t_{change}, ..., t_{curr})) \right|$;
    if $(bandIndex_{new} = bandIndex_{old})$ **then**
        $count = count + 1$;
        if $(count = 3)$ **then** ▷ peak has not moved for three iterations
            $flag = 0$;
            $f_{curr} = f_s + (k - 3)hf_s$;
    else ▷ peak has moved from $bandIndex_{old}$ to $bandIndex_{new}, aliasing$!!!
        $bandIndex_{old} = bandIndex_{new}$;
        $count = 0$;
    $k = k + 1$;

4.8 An illustrated example

Figure 4.2 refers to the filtering process of the signal of Section 4.6 before the change point n= 500. Here, where B has been set to 5 and a uniform division has been adopted.

Figure 4.5 shows the signal spectrum before and after the appearance of a new frequency peak f_{new} at t = 500. Here, f_s is equal to 100 Hz and f_{new} seems to be around 7 Hz. If the new peak is an aliased one, its actual value can be $f = |kf_s - f_{new}|$, with $k \in \mathbb{N}$. We will see later on that we are actually observing an aliased peak.

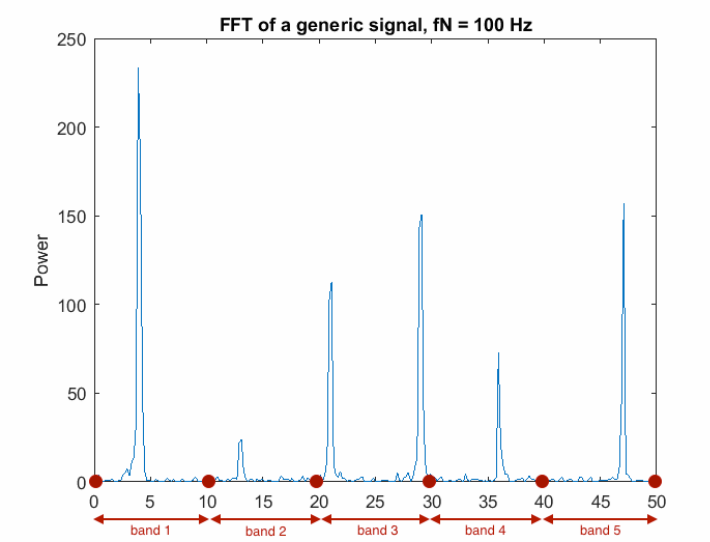

(a)

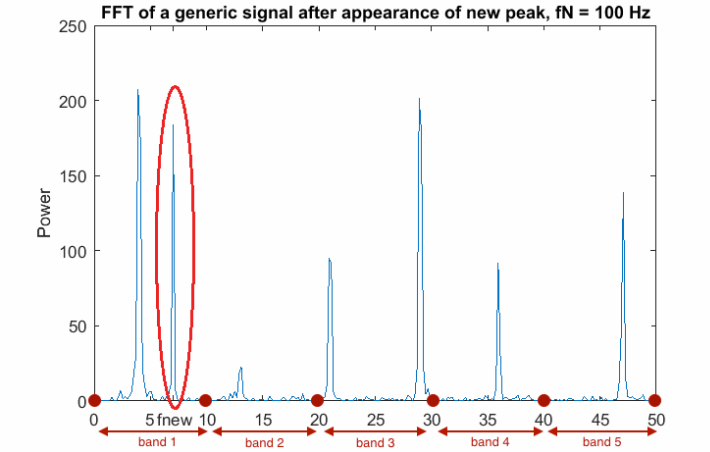

(b)

Figure 4.5: Fourier transforms of a generic signal with new peak appearance: Figure 4.5a is calculated before n=500 while Figure 4.5b reports the new peak at f_{new}

After each detection, the mechanism to distinguish real peaks from aliased ones is put into place. In the proposed example, the actual value of f_{new} is 93 Hz. Therefore, five iterations are necessary to reach the fixed point and terminate the adaptation phase. Here h has been set to 0.6 and n has been set to 3 (hence, three frequency bands are added every time the sampling frequency is increased). The values of f_{curr} for each iteration are reported in Table 4.1.

The peak at f_{new} stops moving after 5 iterations. Figure 4.6 shows the

Table 4.1: Values of f_{curr} for each iteration

k	f_{curr}
1	160
2	220
3	280
4	340
5	400

FFT for $k = 1$ and $k = 2$. It is important to notice that the FFT is not calculated in the algorithm: it has been reported here for clarity reasons. For larger values of k, the FFT is identical to Figure 4.6b.

Therefore, the current sampling frequency f_s is set to $f_{curr} = 220$ Hz, which is the first frequency of the adapting sequence to cover the Nyquist limit.

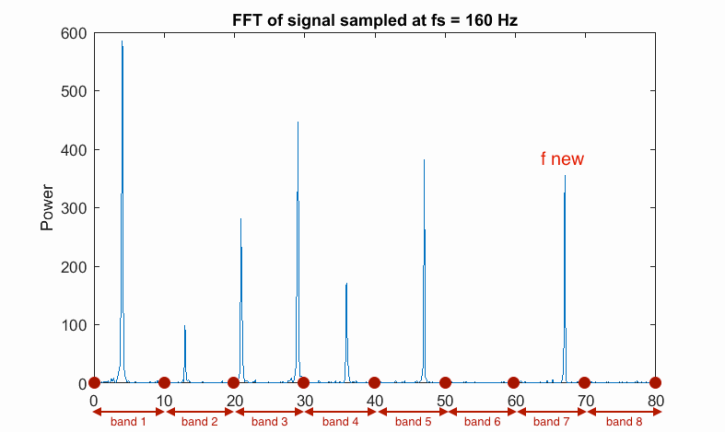

(a) Peak moves at 57 Hz

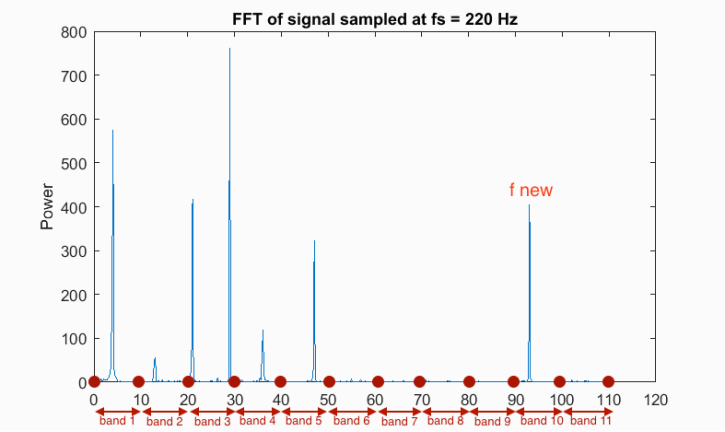

(b) Peak moves at his true position, 93 Hz

Figure 4.6: Two iterations of sampling frequency rescheduling

5 Experimental results

This chapter presents the experiments carried on the datasets presented in Chapter 1, along with their results. A description of each dataset will be provided (Section 5.1), as well as an indication of the figures of merit employed to evaluate the performance of the proposed mechanism (Section 5.2). In order to give a more complete evaluation of our approach, another adaptive sampling algorithm ([BWCW13]) has been taken into account; its working principle will be detailed in Section 5.3. Section 5.4 describes and explains the results obtain for each considered dataset, while Section 5.5 will be dedicated to general comments summarizing the work done.

5.1 Dataset description

Three datasets have been employed to evaluate the effectiveness of the proposed approach: the first is a totally synthetic dataset, with an artificial change introduced at a predefined time instant. The second dataset is a real one, but an artificial change is introduced. Finally, the third dataset is a real one, with actual changes at known time instants.

5.1.1 Synthetic dataset

The generated signal consists in the sum of fifteen sine waves with different amplitude, whose frequencies range from 5 Hz to 998 Hz. The initial sampling frequency f_s is equal to 2000 Hz and some gaussian random noise is added to the sine waves sum.

Frequency values and corresponding amplitude values are reported in Table 5.1.

Frequency	5	80	160	165	229	301	357	398
Amplitude	0.7	1	1.3	0.9	0.6	1.8	1.5	0.65

Frequency	480	561	678	680	858	980	999
Amplitude	0.4	1.18	0.98	0.5	0.28	1	1.25

Table 5.1: Frequencies and amplitudes of the synthetic dataset

The spectrum in stationary conditions is shown in Figure 5.1.

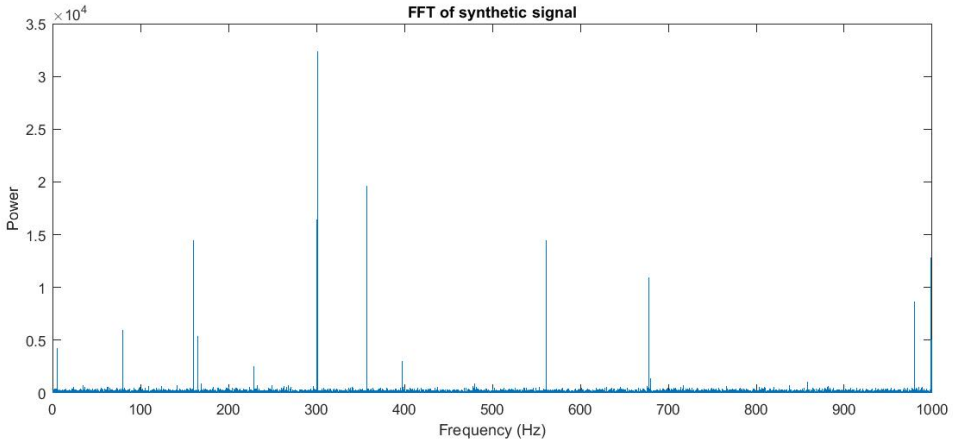

Figure 5.1: The Fourier Transform of the synthetic signal composed of fifteen sine waves

Signal consists of 10^5 points. At $t = 5 \cdot 10^4$, a change is inserted: peak at frequency 168 Hz moves to 364 Hz. This dataset models a frequency shift within the signal's frequency range, which doesn't entail a change in the sampling frequency but could not be detected without the spectrum division.

A second version of the synthetic sinusoidal dataset has been designed and employed to analyze a change in f_{max}. This time it contains only four sinusoidal components (frequencies equal to 5 Hz, 12 Hz, 21 Hz and 47 Hz), the Nyquist frequency f_N is equal to 50 Hz and the introduced change is the appearance of an aliased peak at 67 Hz. It has been used to compare the performance of the proposed approach with other existing approaches presented in Section 5.3. The length of the generated signals varies from 10^2 to 10^6, these signals have been employed to test the efficacy of the frequency adaptation phase and to measure the execution time of the considered algorithms.

5.1.2 The Blue Whale Call

This second dataset is the moan of a Blue Whale call ([Pro08]). These huge cetaceans make sounds and vocalizations for several different purposes: among possible reasons there are sexual selection and individuation of topographic features or prey location. A complete call is composed of a *trill* followed by some *moan*s. Only *moan* sounds are analyzed in the following.

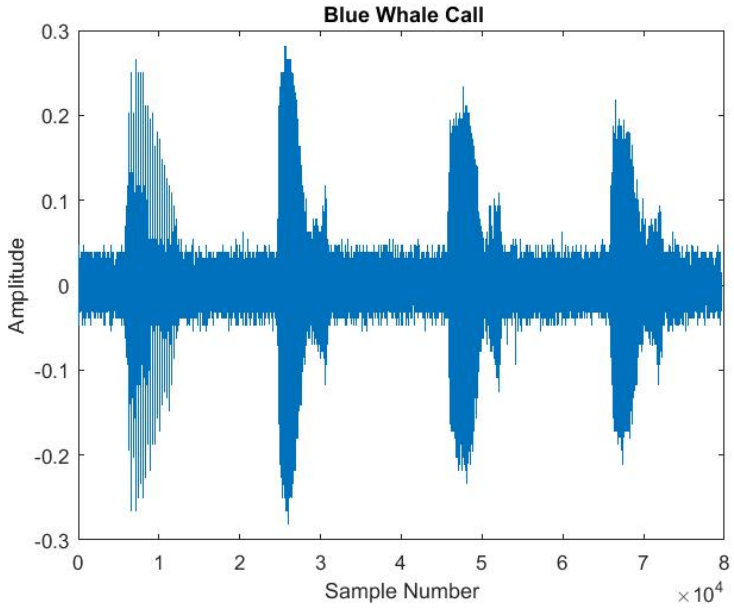

Figure 5.2: Pacific blue whale vocalization recorded by underwater microphones off the coast of California

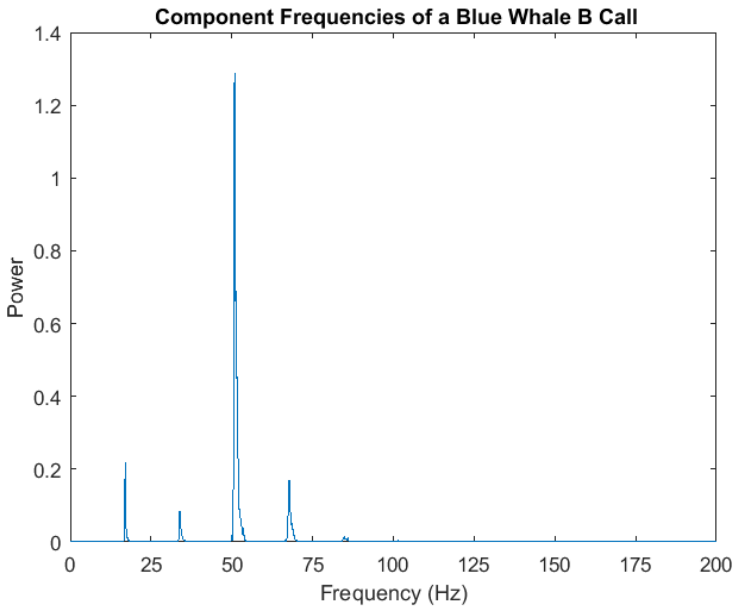

Figure 5.3: Moan from a blue whale call

Figure 5.2 shows a complete call from a Blue Whale, where the initial

sampling frequency f_s is 400 Hz. The real frequencies emitted by the whales are actually too low to be audible from human ears and hence, in the original dataset available online, they have been speeded up so that the sampling frequency passes from 4 kHz to 400 Hz.

After isolating *moan* episodes (second, third and fourth peak in Figure 5.2), the Fourier transform is computed and the spectrum is reconstructed in Figure 5.3 : the moan call is composed of a fundamental frequency around 17 Hz and a sequence of harmonics, with the second harmonic emphasized.

A new frequency of 425 Hz is then added to the moan call. This frequency is higher than 200 Hz (i.e., $f_N/2$), so we expect to observe the aliasing phenomenon and see the new peak within the spectrum. The amplitude of the new frequency has been set to be of the same order of magnitude of the other frequencies to simulate a plausible scenario.

5.1.3 ECG during tachycardia episode

The last dataset studied in this work concerns medical applications. Two electrocardiograms of subjects affected by ventricular flutter, a form of tachycardia which affects the ventricles and could be the prelude to a sudden cardiac death, have been analyzed.

Normally, a resting adult human heart rate ranges from 60 to 120 BPM (Beats Per Minute), which is equivalent to 1-2 Hz. Values of heart rate greater than 120 indicate tachycardia, while in case of ventricular flutter it can range from 160 to 300 BPM (2.6 to 5 Hz).

Some people who experience a high risk of heart failure need constant monitoring of heart rate and other vital parameters. A lot of devices have been developed in these last few years to measure heart rate ranging, from the most accurate and sophisticated for people who are effectively at stake, to the less accurate but still working heart rate sensors on smartphones or watches. In both cases and, especially, when constant monitoring is needed, power consumption becomes an important issue. The possibility of monitoring the heart beat by sampling at a low frequency and observing abnormal phenomena through aliasing, thus, could be of great importance and could bring an effective improvement in applications which are aimed at monitoring a patient's heart up to 24h a day.

The original dataset is available at the Physiobank database ([GAG+00, Gre86]). Only records #418 and #419 have been analyzed, because the other records present also other types of arrhythmia. Figure 5.4 shows what happens in real life to the ECG when ventricular flutter arises: the usual shape of ECG (with clear definition of the QRS and T waves) is replaced by a sinusoidal waveform with smaller period.

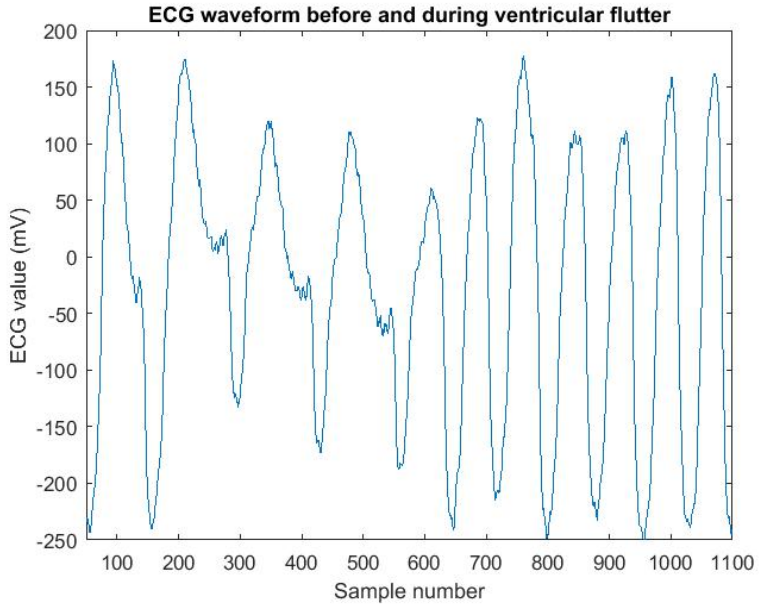

Figure 5.4: Transition between regular heart rhythm and ventricular flutter

These records last about 35 minutes each and they are sampled at 250 Hz, so that they contain about 500000 samples. In the first ECG, 60 episodes of ventricular flutter have been detected, while in the second one the episodes are 34. Table 5.2 shows the annotations available in the database, with labels indicating whether the heart rate is normal or a ventricular flutter episode is occurring.

The main signal has been divided to isolate each ventricular flutter episode plus the period of regular heart rate immediately preceding it, therefore obtaining more signals to work on. Moreover, it is better to work with one single change episode at a time and to assume that, after the first detection, the system is reset.

This dataset has been used to test the whole mechanism, including the adaptation phase. Clearly, the sampling frequency cannot be increased as data have already been recorded. Thus, all episodes have been under-sampled, which means that only one sample each n samples of the original signal has been kept to obtain a shorter signal simulating a lower sampling rate. The value of n depends on the ratio between the original sampling frequency and the desired one. Each time the sampling frequency has to be increased, more samples will be taken into account. Figure 5.5 clarifies how the sample selection is performed.

Table 5.2: Partial annotations of ECG record 418, "(N" stays for normal heart rate while "(VFL" indicates ventricular flutter

Time	Sample #	Aux
0:00.072	18	N
6:38.496	99624	VFL
6:45.996	101499	N
8:52.368	133092	VFL
8:56.152	134038	N
9:03.100	135775	VFL
9:06.512	136628	N
10:12.228	153057	VFL
10:16.460	154115	N
10:19.768	154942	VFL
10:25.164	156291	N
10:37.768	159442	VFL
10:42.064	160516	N
11:16.768	169192	VFL
11:19.228	169807	N
11:32.216	173054	VFL
11:34.692	173673	N
11:39.152	174788	VFL
11:41.612	175403	N
11:45.036	176259	VFL
11:51.472	177868	N
12:40.320	190080	VFL
12:44.996	191249	N
12:47.228	191807	VFL
12:50.780	192695	N
13:02.524	195631	VFL
13:07.176	196794	N
13:20.844	200211	VFL
13:22.536	200634	N

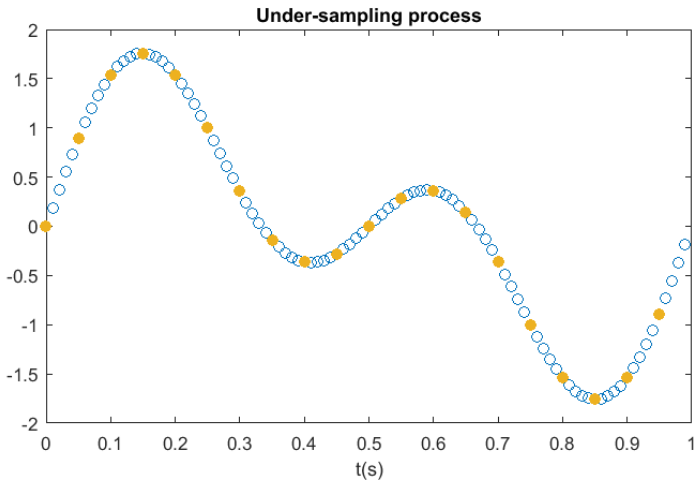

Figure 5.5: Under-sampled signal: only the orange filled samples are taken into account, while the blue ones are excluded.

5.2 Figures of merit

Giving a complete and accurate evaluation of the algorithm's performance is not an easy issue and, in particular, it is quite hard to give a precise estimation of the energy consumed by a CPS without being biased by the observed event's dynamics.

Nevertheless, it is possible to find some figures of merit that provide a good picture of the mechanism's efficacy and, most important, can be compared with other approaches.

The figures of merit employed in this work are the following:

- **False positive rate (FP)**: a false positive is a false alarm, i.e., an erroneous detection: the algorithm notifies a change that has never occurred. It is calculated as the number of false detections over the number of experiments. False positive detection represent a problem because they slow down the whole mechanism and might lead to inappropriate decisions.

- **False negative rate (FN)**: a false negative is, quite intuitively, the miss of an actual change. It is calculated as the number of missed detections again over the number of experiments. There is a trade-off effect between the values of FP and FN: an application trained to react even to small changes will result in a low FN rate, but this will automatically result in a higher FP rate, and vice versa. Usually, FN rate tends to be more dangerous than FP: in fact, a false positive usually affects performance and entails costs, but a false negative may have more serious consequences (as an example, imagine fire prevention system missing

the surge of a fire). Therefore, depending on the application criticality, it should be kept to zero. In general, it is preferred to lower the expectations on FP, rather than risking a high FN;

- **Detection delay (DD)**: all detection systems introduce a delay in the identification of a change, which is the time that elapses between the physical change in the real world and the detection notified by the system. Of course, this delay should be the lowest possible, in order to have a prompt reaction, but there is again a trade-off effect with the FP rate, because a reactive application is more likely to raise false alarms;
- **Execution time (ET)**: expressed in seconds, it's the time needed to run the code and analyze data. The execution time should be as small as possible;
- **Sampling fraction (SF)**: this is a figure of merit typical of adaptive sampling algorithms. It is defined as the rate between the number of samples acquired by an adaptive sampling mechanism over the number of samples acquired by another sampling approach over the same signal. Usually, the term of comparison is the fixed-rate sampling. Naturally, the smaller the rate, the more efficient the adaptive sampling algorithm is. It is an excellent indicator of the algorithm validity because, when considering a wide Cyber-Physical System, it is easy to recognize how heavy is the impact of a reduction in the number of samples acquired by each sensing unit. Of course, this reduction doesn't have to reduce the accuracy of the phenomenon representation, but it has an enormous influence in the energy management of the whole system.

5.3 Alternatives for the comparison

To show the effectiveness of the proposed approach and have a term of comparison, it has been tested again a fixed rate approach and another adaptive sampling approach called eSampling. The latter is an activity-based adaptive sampling algorithms which guarantees that the Nyquist theorem is respected and is described in [BWCW13].

Two datasets of Section 5.1 have been employed for the comparison: the ECG database, which is interesting because it doesn't have any artificially added perturbation, and the second version of synthetic sinusoidal dataset described in Section 5.1.1.

[BWCW13] has been chosen for a comparison because, like this work, it is an adaptive sampling algorithm based on the Nyquist theorem. In fact, it doesn't rely on the forecasting of next samples as many other approaches do, hence a comparison is more interesting. As explained in Section 2.3,

this approach couples a frequency adaptation phase, based on the individuation of the high frequency components of a signal through the Wavelet Packet Decomposition, with a change detection phase monitoring sudden variations. This second phase basically compares the difference between the most recent acquired value and the reference value q^{Ref}, computed during a training phase (performed when the observed phenomenon is stationary) and representing the reference sequence. If the relative difference is larger than a threshold, computed according to the desired confidence, a change is detected. This approach is, hence, quite similar to the one proposed in this work, although it doesn't exploit the aliasing phenomenon and, most importantly, it periodically calculates the wavelet decomposition, which is analogous to a FFT in terms of computational load.

5.4 Results description

5.4.1 Effects of the number of bands

In order to have an idea of the accuracy and effectiveness of the proposed approach, a preliminary analysis on the detection phase has been carried out and, more specifically, the effect of the number of bands on the figures of merit presented in Section 5.2 has been considered. These experiments have been done on the first two datasets, where the change was introduced artificially, hence the adaptation phase has not been considered in this first analysis.

Choosing the correct number of bands is an issue of main importance for the success of detection and plays a crucial role in determining precision and cost of the proposed mechanism.

There is evident trade-off between the algorithm's cost and its accuracy: on one hand, many bands allow a better visibility as the granularity of the experiment is finer, and shifts of smaller amplitude can be detected. On the other hand, each new band implies the calculation of two vectors of size N and S respectively, which clearly increases the complexity. Moreover, more bands induce a higher probability of having a false positive.

Concerning the synthetic dataset, the signal spectrum has been divided into 4, 5, 10 and 20 bands; W is set to 25 and t_{change} to 400 and submitted to the univariate CDT described in [ABR11c, ABR11a, ABR11b]. This test considers all bands independently and, as expected, the false positive number increases with the number of bands, as shown in Figure 5.7b. At the same time, Figure 5.7a shows that the mean detection delay decreases with the number of bands, in agreement with theoretical premises.

Again, there is a trade-off also between the number of false positives tolerated and mean detection delay, as Figure 5.8 enlightens: in order to have a small detection delay, the algorithm must be more reactive to small changes, which entails a higher probability of false positive detections. Here, the

smallest detection delay is equal to 269 samples, corresponding to 3.4 s in the original signal $\widehat{x}(t)$, while false positive rate is around 12% which is still an acceptable level.

Secondly, the Blue Whale dataset has been analyzed to see how the detection phase works.

As expected, the inserted peak is reflected back as shown in Figure 5.6, in total agreement with Section 3.2.

Clearly, in this case it is not possible to increase the sampling frequency f_s because data have been recorded already, but it is interesting to confirm that aliased frequencies can be perfectly detected by the CDT, as well as real ones.

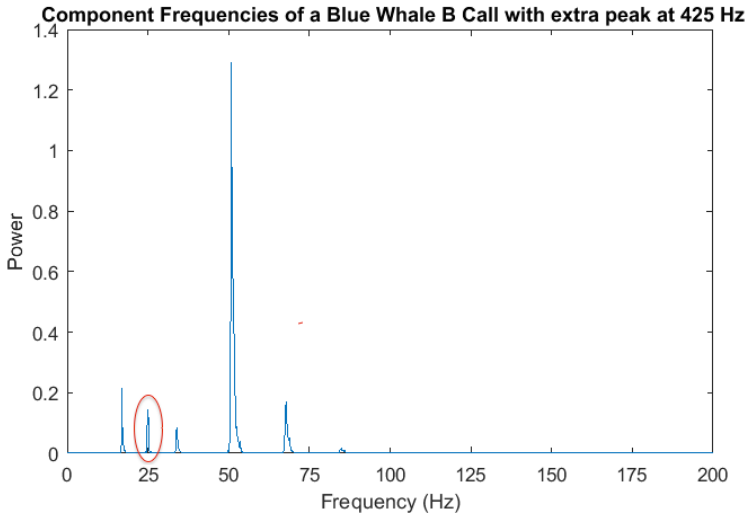

Figure 5.6: Insertion of an extra peak at 425 Hz

As for the synthetic dataset, the 200 Hz spectrum is divided into 4, 5, 10 and 20 bands of equal width. W is 15 and t is 100, and the results are in line with what expected. In this analysis there were only three available signals, each of whom corresponding to a *moan* call in the original dataset, hence mean values are calculated on three experiments.

Trends of detection delay and false positive rate are the same of the synthetic dataset, as shown in Figure 5.7. For the DD, the values in case of 10 bands and 20 bands are the same. What is different is the false positive rate (see Figure 5.7b): its value for the 20 bands case is twice the FP for 10 bands, while the DD is the same (7.3 samples, corresponding to 0.27 s). This enlightens that 20 bands are clearly too much for these experiment: they are useless in terms of detection delay and they considerably increase the PF rate.

False negative rate is, again, equal to zero independently from the number of bands.

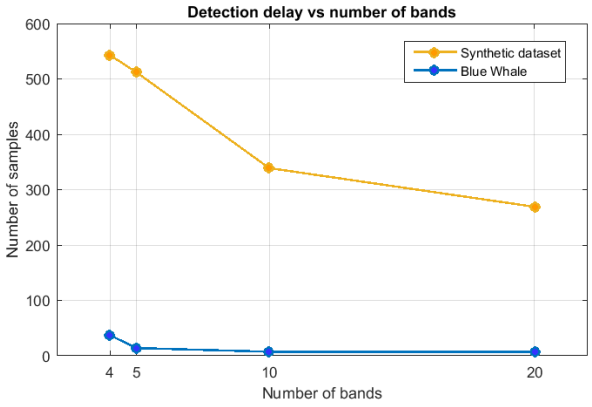

(a) Delay vs number of bands

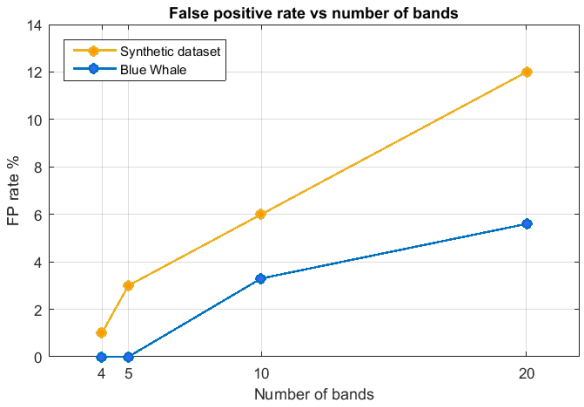

(b) False positive rate vs number of bands

Figure 5.7: Performance of hierarchical ICI CDT varying the number of bands on two different datasets.

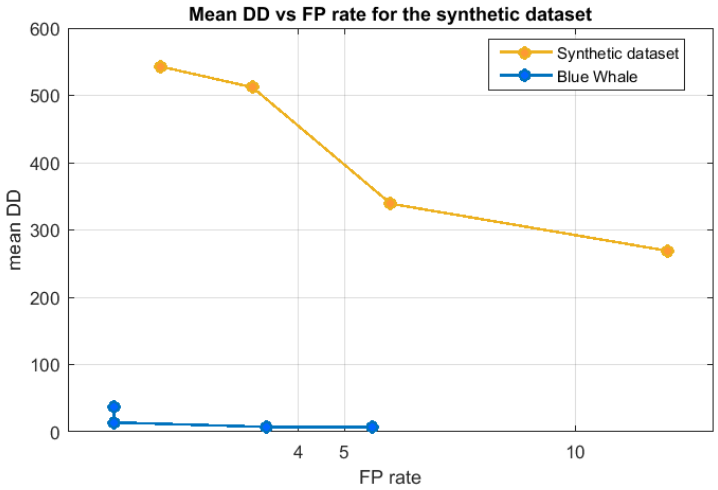

Figure 5.8: Mean detection delay over false positive rate

5.4.2 Detection and adaptation

The detection phase and the consequent adaptation phase have been deeply tested on the ECG dataset presented in Section 5.1.

The two original ECG signals have been divided into "flutter episodes": after each occurrence of the ventricular fluttering the ECG has been split, in order to have several sequences of normal rate samples followed by a fluttering period, in which the new frequency peak appears.

Figure 5.9 shows the Fourier transform of one of the longest sequence, with the most powerful peak at around 1.8 Hz and a second powerful peak at around 3.75 Hz (225 BPM). Then, the signal has been divided into normal heart rate and ventricular flutter, and the two corresponding spectrum figures are shown in Figure 5.10a and Figure 5.10b respectively. The ventricular flutter peak at around 3.75 Hz (225 BPM) is well visible and is the only one present during the flutter occurrence.

For the filtering phase of this specific application scenario, the frequency range in which there is any kind of activity is quite small, therefore, only four bands have been considered. As stated before, bands should be chosen adaptively in order to distribute energy density and to avoid the creation of a single band in which the majority of the spectral energy is concentrated. Thus, the values of the four bandwidths have been chosen as follows:

- Band 1: from 0 Hz to 0.5 Hz;
- Band 2: from 0.5 Hz to 1.5 Hz;
- Band 3: from 1.5 Hz to 2.1 Hz;
- Band 4: from 2.1 H to 3 Hz;

This frequency division takes into account the actual distribution of signal's energy. Concerning this specific experiment, frequencies of regular heart beating are well known in the medical community, as well as anomaly behaviors. Therefore, band distribution can be adapted to different situations and monitored phenomena.

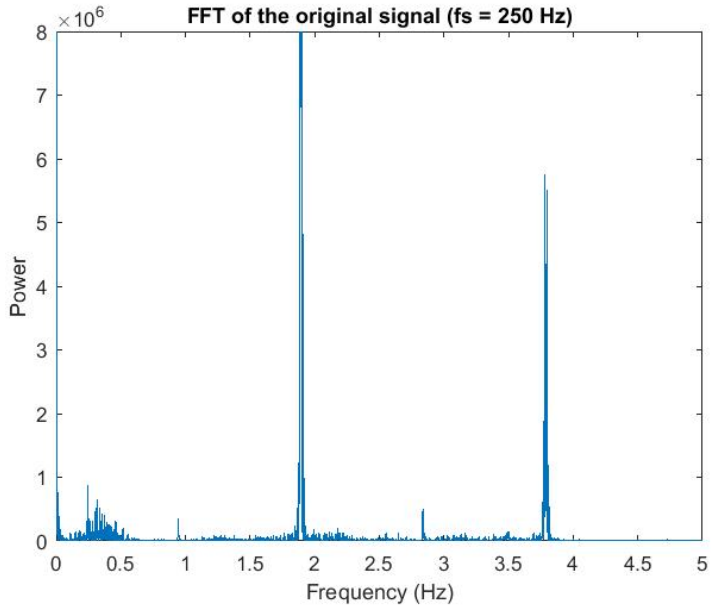

Figure 5.9: Fourier transform of the original ECG signal: two principal peaks are well visible

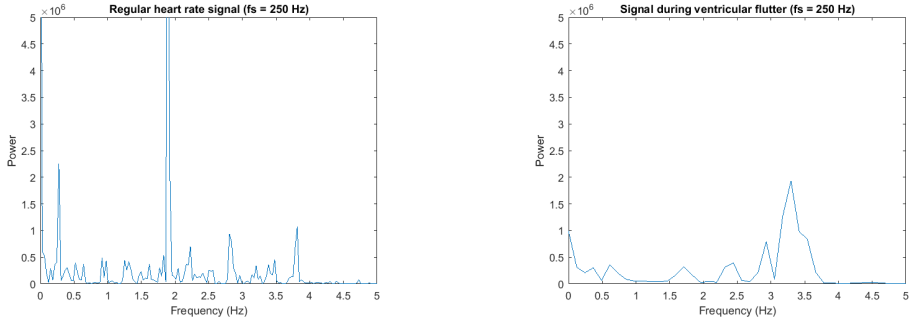

(a) **(b)**

Figure 5.10: Fourier transforms of normal heart rate in Figure 5.10a and of ventricular flutter in Figure 5.10b

These flutter episodes are sampled at 250 Hz but, in order to simulate the

actual working principle of our adaptive sampling approach, they have been under-sampled at 6 Hz. This results in a significant reduction of their length and also, considering the integration phase described in Section 4.4 (which here has been done on very small windows, $W = 5$), the episodes that remain long enough for a useful analysis are only 23. This length problem would clearly be eliminated with the collection of a dataset directly sampled at 6 Hz, in a real application.

On these remaining 23 sequences, the previously used hierarchical ICI detection test does not perform optimally, because it is not designed to work on very short sequences. Another problem is that, in most cases, flutter duration is much shorter then normal heart rate signal's one, so it is hard for the test to recognize changes that occur just before the end of the signal.

Therefore, another CDT has been used in this section: the Multivariate Change Detection Mechanism developed in [Tac15]. While working in offline mode, this test considers all samples and all four bands at once, and waits for the average value of energies to overcome a given threshold to raise the alarm state. The threshold is defined from a statistical point of view. Figure 5.11 shows an example for one of the 23 sequences, the threshold's value is depicted in red.

For the current dataset, online mode results almost always in a false positive detection: this is explained by the oscillatory dynamics of the data, which would be leveraged by wider windows for energy integrals. Unfortunately, in this case longer intervals leave only two sequences long enough to run the test. It has to be noted that, on these two long sequences, the online test performed very well and did not fail to promptly detect the change, as (Figure 5.12) illustrates.

The overall performance of this test is definitely satisfying, given the data stream intrinsic dynamics: false positive rate is around 30%, the average delay is of 4.8 samples for the energy signal, which corresponds to a delay of 24 samples in the original signal (thus, 4 s). Like the previous experiments, false negative rate is equal to zero.

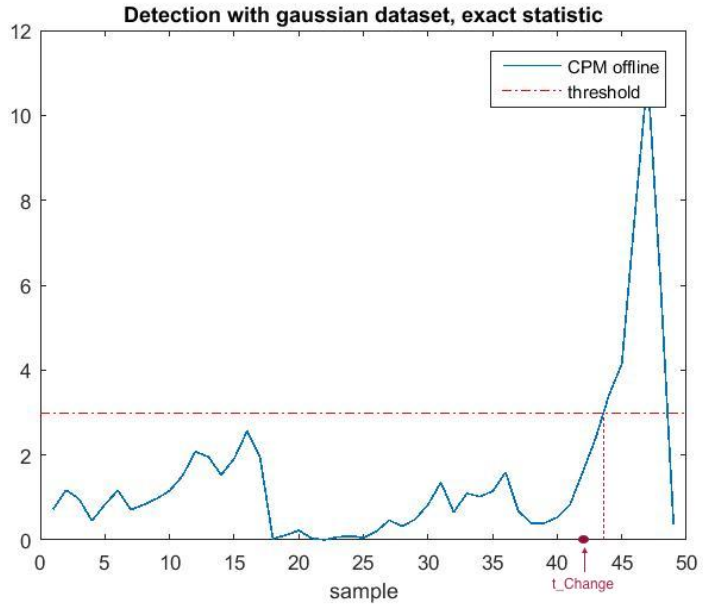

Figure 5.11: Statistic calculated by the multivariate change detection test, offline mode

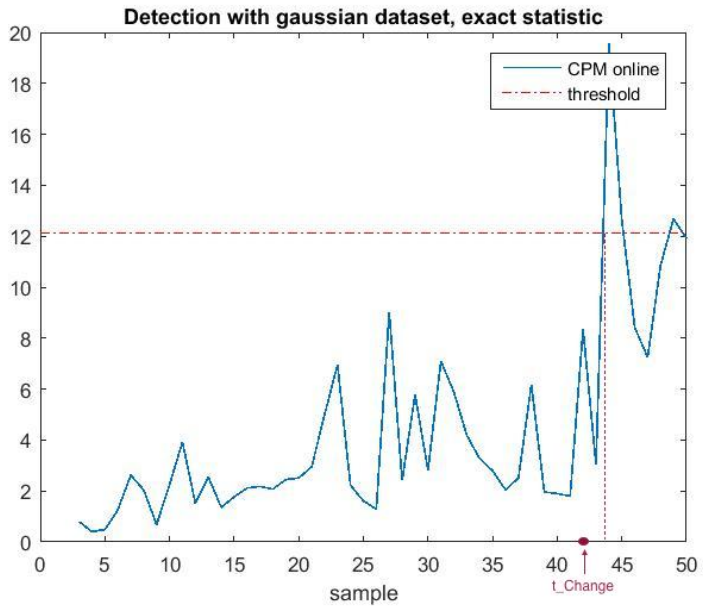

Figure 5.12: Statistic calculated by the multivariate change detection test, online mode

After a change is detected in the average value of the signal's energy, the sampling frequency has to be adapted in order to identify what kind of change has taken place.

Therefore, sampling frequency is increased of a factor $c = 0.17$ and brought to 7 Hz, as a consequence of the new peak appearance. As illustrated in Figure 5.13a, f_{new} shifts rightward and reaches the value of around 3.25 Hz, in total agreement with the expectations. Another iteration of the adapting process is carried out, and the sampling frequency is brought to $f_{s,2} = fs + 2cf_s = 8$ Hz, as Figure 5.13b shows. We know that the real peak is at 3.75 Hz and, hence, we have reached the proper sampling frequency. Anyways, another iteration is computed and, as the new frequency doesn't move, f_s is set to 8 Hz.

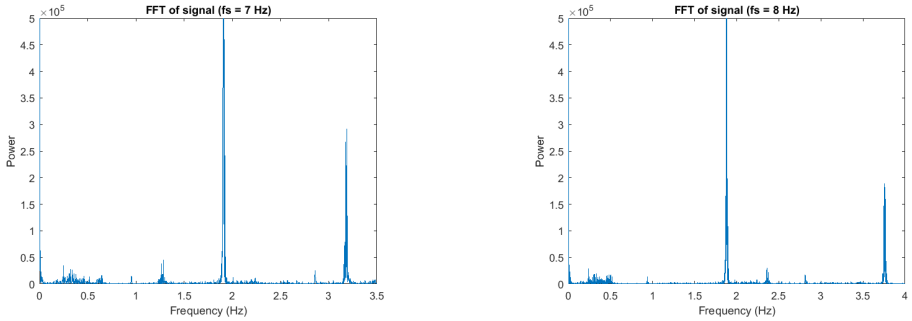

(a) **(b)**

Figure 5.13: FFT of signal sampled at 7 Hz (Figure 5.13a) and at 8 Hz (Figure 5.13b)

5.4.3 Comparison with alternative methods

As anticipated in Section 5.3, a comparison a with fixed-rate approach and an alternative approach called eSampling of [BWCW13] has been done. The comparison has been evaluated on the ECG dataset and on the second version of the synthetic dataset of Section 5.1.1.

Sampling fraction

The first figure of merit that has been considered in this comparison is the sampling fraction, which is the rate between the number of samples acquired in the proposed method and the number of samples acquired by the other mechanism employed for a comparison.

When comparing our spectrum based approach with a fixed-rate sampling one, a first consideration is that the fixed sampling frequency has to be chosen judiciously.

We have seen that the sampling frequency of heart rate in the ECG experiment can be lowered from 250 Hz to 6 Hz for our case study, which leads to

a reduction of acquired samples number of 97.6%. This result is astonishing, but we have to point out that 250 Hz is a very high sampling frequency for the original signal. Probably, it has been chosen in order to reconstruct the waveform in the best possible way, as the MIT Malignant Ventricular Arrhythmia Database has educational purposes and its ECGs serve as a term of comparison for other studies.

Hence, we can assume that a standard application samples at four times the Nyquist frequency, which is the case for most industrial applications. This leads to a reduction of 81.3% in the number of samples, still a very good result. We retain that a reasonable value to make a comparison for a fixed-rate approach is 10 Hz, which about twice the Nyquist frequency and tolerates the eventual arising of higher frequency peaks. This choice leads to an average sampling fraction of 38% for the spectrum based adaptive sampling.

eSampling alternates short periods of high frequency sampling, aimed at detecting changes in the maximum frequency, with long periods of low frequency sampling (see Section 2.3 for details). For this application, high frequency has been set to 10 Hz, and low frequency to 6Hz, which doesn't violate the Nyquist limit during normal heart-beat. When ventricular flutter arises, the low frequency is brought to 8 Hz. Clearly, the proposed spectrum based adaptive sampling algorithm outperforms eSampling, as it avoids those periods of high frequency sampling and leaves to the change detection test the responsibility to identify maximum frequency increase. Figure 5.14 reports a comparison of the three approaches. The trend of the proposed approach and that of eSampling is quite similar, this is due to the fact that for each experiment a single episode is analyzed and the frequency range is very limited. Anyways, they are not proportional to each other.

Table 5.3: Performance comparison between SpectrumBased and eSampling

	SpectrumBased		**eSampling**	
	ECG	*Synth*	*ECG*	*Synth*
mean detection delay (s)	4	1	12	29
false positive rate (%)	30	0	37	41

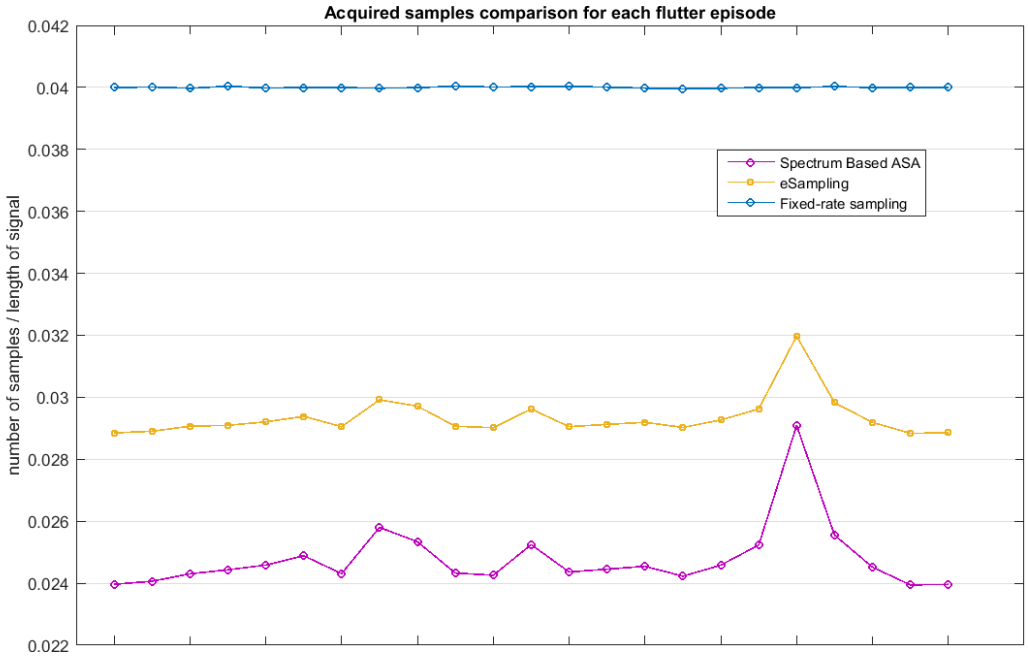

Figure 5.14: Number of samples acquired by the different approaches in each ECG flutter episode

Change detection performance

The proposed approach and eSampling have been compared also on the average detection delay and the false positive rate. Table 5.3 summarizes the results for both the real ECG dataset and the synthetic dataset created on purpose. Again, our approach proves to be more effective than eSampling.

Execution time

The execution time of both the proposed approach and eSampling algorithm has been calculated through simulations on Matlab, where the length of the original signal (expressed in number of samples) has been increased from 10^2 to 10^6. The results are summarized in Table 5.4 and show an interesting effect: for "short" signals, eSampling is quicker than SpectrumBased, but for longer signal the situation is reversed and eSampling takes way longer to terminate.

Table 5.4: Comparison between execution time of SpectrumBased and eSampling

number of samples	SpectrumBased	eSampling
10^2	0.514	0.011
10^3	0.260	0.170
10^4	0.422	1.724
10^5	6.164	26.808
10^6	22.978	98.513

Table 5.5: Execution time distinguished between pre-processing and operative phase

	SpectrumBased		eSampling	
number of samples	*pre-processing*	*operative*	*pre-processing*	*operative*
10^2	0.154	0.360	0.001	0.010
10^3	0.167	0.093	0.002	0.168
10^4	0.200	0.222	0.003	1.721
10^5	0.448	5.716	0.006	26.802
10^6	3.589	19.389	0.088	98.425

The unusual trend can be explained as follows: SpectrumBased requires an important data preprocessing phase, which is absent in eSampling. Therefore, for short signals, it is penalized by the energy calculation step and the filtering process. Moreover, errors arise in the 10^2 samples signal because of the limited number of energy values obtained. On the other hand, when the length of the signal increases, the proposed approach outperforms eSampling. One of the main contributions for this remarkable result is that SpectrumBased avoids the Fast Fourier Transform calculation to estimate the maximum frequency. Actually, eSampling computes the Wavelet Package Decomposition after each "burst" of high frequency acquisition, and this puts at disadvantage its performance on the long term.

The analysis of the partial execution times (Table 5.5), which distinguish the pre-processing phase from the operative phase, confirms the previous explanation. While the time needed to pre-process data in eSampling grows very slowly and has an almost insignificant impact on the total ET, preprocessing time fore SpectrumBased reaches the 46% of the total execution time for signals of 10^3samples and its value is always much greater than for eSampling. On the other hand, when the length of the signal increases, the pay-back of the time invested in pre-processing data becomes evident, as the total ET of SpectrumBased is almost the 25% of the one of eSampling.

5.5 General comments

Previous analyses enlighten the many positive contributions of the proposed approach, along with some other aspects that leave place to further developments.

The first analysis on the detection performance given the number of bands confirms the existence of a trade-off effect between the accuracy of the detection and the false positive rate. Actually, more bands reduce the mean detection delay, but increase false positive rate. Therefore, the choice of this parameter is still left to the user and is application-specific, because even when observing the main phenomenon an application might concentrate on having very accurate result and a potentially high false positive rate, or vice versa.

The analysis of the adaptation phase, as well as the comparison with the alternative approach, show the excellent performance of the proposed solution under many aspects. The considered figures of merit for the ECG experiment reveals to be satisfying, and they could be further improved with a dedicated monitoring system and a specifically collected dataset. In fact, the FP is quite high, but this is due to the shortness of some fluttering episodes and the impossibility to integrate on wider windows, which would leverage the signal's oscillations.

The comparison with eSampling points out that the proposed approach reveals to be more efficient in terms of acquired samples number, execution time and of promptness of detection. This is due to several factors: concerning the sampling rate SR, our approach outperforms the other one as it avoids the cyclic phases of high frequency sampling. Here, the sampling frequency is increased only when a change is detected, not in advance, as eSampling does. Promptness of detection depends on the efficiency of the CDT employed and, even if the development of the CDT itself is not the issue of this thesis, a more accurate test has been considered.

The analysis on the execution time enlightens that the proposed mechanism introduces a small overhead for short signals, for which the ET is slightly greater than eSampling. On the other hand, for longer signals, this method proves to outperform the other one. This is due to one of the main contributions of the proposed approach: actually, it avoids the heavy calculation of a FFT at each cycle, relying on the computation of the energy values instead.

We believe that an evaluation of the proposed spectrum based algorithm portability on real embedded systems was necessary, because it has been designed to work on CPS units with limited computational resources. Chapter 6 illustrates an example of application on a real CPS unit.

6 The proposed adaptive sampling mechanism on embedded systems

6.1 Implementation and porting on the dedicated system

The mechanism proposed in this work has been ported on a real embedded device, the sensing unit of a CPS. The employed device is a programmable board and a sensor expansion, which can be programmed through an online compiler provided by STMicroelectronics.

The idea behind this experiment was to test the applicability of the proposed approach on a real system, subject to memory and calculation power limitations. Despite the fact that some simplifying assumptions on the proposed algorithm have been made, the results are positive and show how the considered algorithm can actually be exploited by an embedded system being part of a CPS.

6.1.1 The considered embedded system

The components employed for this experiment are the Nucleo board STM32F411 and the sensor expansion IKS01A1. The board communicates with the PC through a USB port. The expansion is equipped with a temperature sensor, a humidity sensor, a pressure sensor and three accelerometers. In this case, only the temperature sensor has been taken into account, but any other physical quantity would have worked the same.

The sensor started sampling at a predefined frequency, collecting the first samples used for the training phase. Without any loss of generality, these samples have been analyzed with Matlab to reconstruct the signal spectrum and tune all parameters accordingly. Actually, the temperature values inside a house at daytime don't change very much and the spectrum figure outlines only very low frequencies. This is the reason why a sinusoidal component has been added artificially, in order to give some dynamics to the signal. Moreover, two different types of changes have been introduced to evaluate the system reactivity: an increase of the temperature through the emission of warm air, and the introduction of another sinusoidal component at a given instant. In both cases, the mechanism has promptly detected the change.

Figure 6.1 shows the Nucleo components, with the programmable board (A) and the sensor expansion (B). The sensing device is connected to the PC through a USB cable.

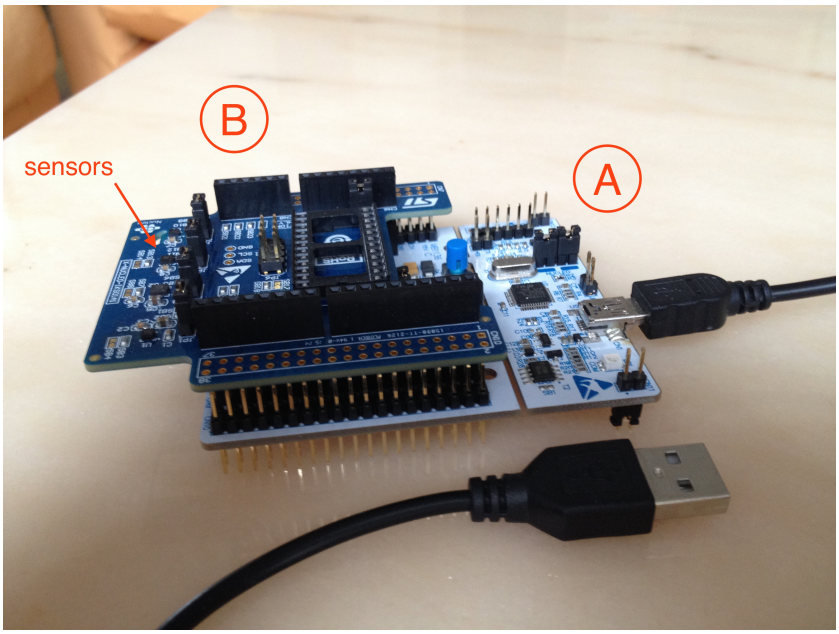

Figure 6.1: Nucleo F411 board (A) and its expansion with sensing devices (B).

6.1.2 Assumptions

Some simplifying assumptions on the proposed mechanism have been made due to the necessity of avoiding computational overhead and, secondary, for time constraints.

First of all, the signal has not been filtered and split into several bands, but it has been considered as a single frequency band instead. Actually, as the idea here was to have a proof of concept for the algorithm portability, a single band interested by a change was considered. If necessary, filter banks could be considered on the embedded device, depending on the specific application, in order to manage the frequency bands and could be activated according to the process needs.

Without any loss of generality the Fast Fourier Transform calculation of the training phase has been calculated on Matlab.

Finally, the change detection test employed is slightly different from the one used in the previous chapters, as it doesn't have an implementation in C++ and hasn't been published yet at the time of this experiment.

All these assumptions clearly simplify the overall mechanism proposed in Chapter 4, but they surely don't invalidate the experiment, as the core of the proposed mechanism is maintained: all samples are integrated over time, a change detection test runs on the energy values, and the sampling frequency is adapted with the same mechanism of the offline synthetic experiments, although the number of bands is not increased.

6.2 Results

As anticipated before, the results are encouraging for this experiment and the developed application revealed to be able to detect changes in the environment.

In the first experiments, the temperature values of the area around the sensor have been increased, and the application has immediately spotted the change. After a first increase of the sampling frequency, though, its original value has been restored, as no aliasing phenomenon has been detected (the anomaly wasn't due to a change in the maximum frequency contained in the spectrum). Figure 6.2, obtained with Matlab, shows the Fast Fourier Transform before and after the change: no significant difference can be identified, since the spectrum itself has not changed.

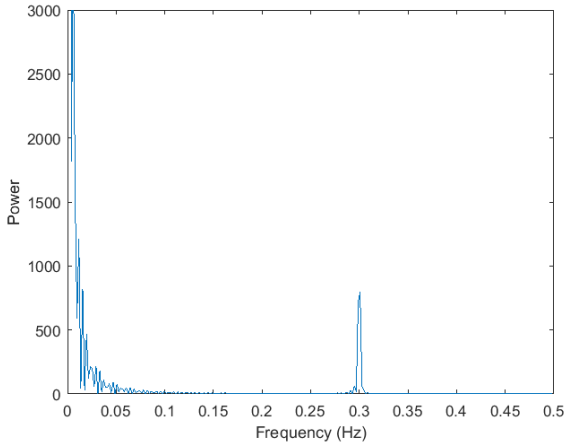

(a) FFT before t_{change}

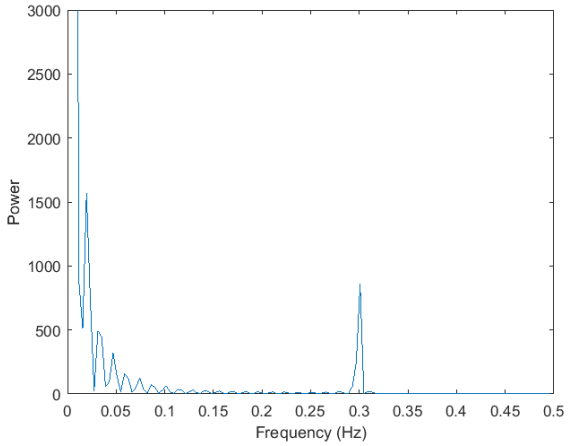

(b) FFT after t_{change}

Figure 6.2: Fourier transforms of temperature signal before change in Figure 6.2a and after change in Figure 6.2b

In the second experiments, a sinusoidal component has been introduced at $t = 200\ s$. The original sampling frequency was 1 Hz, and the lower sine wave was at f = 0.3 Hz. The high frequency peak was at 150.4 Hz, resulting in aliasing peak visible at 0.4 Hz. Figure 6.3 shows the two well visible peaks.

The detection delay depends on the window length W, as all W samples are aggregated to obtain the energy value. Hence, the minimum detection delay is W/f_s. In this experiment, the delay doesn't exceed this lower bound, hence it is almost 10 seconds in the worst case scenario for each detection.

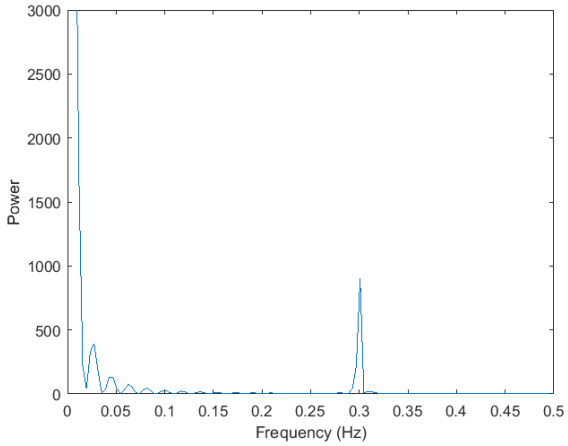

(a) FFT before t_{change}

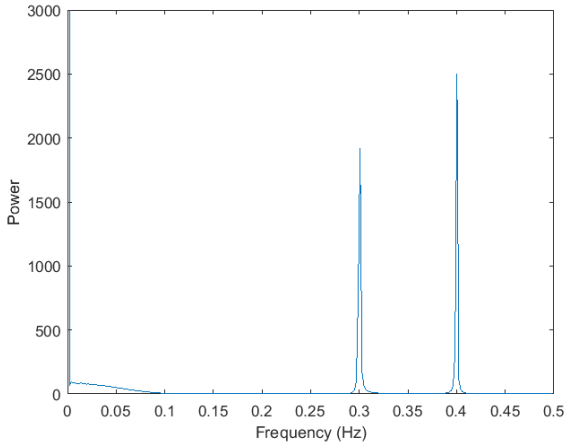

(b) FFT after t_{change}

Figure 6.3: Fourier transforms of temperature signal before change in Figure 6.3a and after change in Figure 6.3b

7 Conclusions

The reduction of the energy consumption in Cyber-Physical Systems is a crucial issue for several applications that are becoming essential in today's industrial and social reality. As these systems are expected to grow even more in the next future, it is fundamental to work on energy saving and battery preservation both from an environmental and economical point of view.

In this chapter the main characteristics of the proposed mechanism are summarized. Section 7.1 describes the contribution of our methodology together with its limits, while Section 7.2 discusses some possible future works.

7.1 Contributions and limits

Many existing energy saving approaches consist in adaptive sampling techniques at the sensing unit level of CPSs and some of them also employ a change detection test to improve performance.

For the first time, the proposed approach exploits the aliasing phenomenon to monitor possible changes in the maximum frequency contained in the signal. In addition to this, the proposed approach integrates the idea of monitoring the signal in the time domain instead of the frequency domain, through the calculation on the energy values. This approach avoids the use of statistical models aimed at forecasting the next samples, and, at the same time, avoids the heavy computational load due to the FFT.

The problem of peak's shift undetectability, introduced by the integration of the energy over time, is overcome by the division of the spectrum into frequency bands. Actually, if the proper number of bands is introduced, a peak moving within the original frequency range is likely to leave one band and entering another one, modifying the energy values of both.

Like in most of the real world applications, the maximum frequency of the signal and, consequently, the Nyquist frequency, are supposed to be unknown. If, at any time, this frequency should exceed the current sampling frequency, this approach is able to detect a change through aliasing and adapt the sampling frequency accordingly.

The whole process proposed in this work integrates all these aspects, thus providing a complete solution for adaptive sampling and energy consumption management. Experimental results have demonstrated its effectiveness compared to fixed-rate approaches and to another algorithm existing in the literature.

Moreover, the mechanism has been carried on a real embedded device, which proves that it represents a viable solution that could be employed in real world applications.

One of the principal limits of this approach is, however, the choice of some parameters which is still left to the user and is application-specific, as well as the need to perform a preliminary spectrum analysis. As stated in the previous chapters, bands need to be chosen adaptively to avoid an unbalanced energy distribution among them; therefore, an initial training phase is needed for band configuration, as well as for training the CDT. Nevertheless, this happens in most of the existing approaches and it is totally reasonable to suppose that most of the specific real world application provide a first phase of inspection in stationary conditions.

Another limitation is the trade-off effect between the detection delay and the false positive rate introduced by the number of bands in which the spectrum is divided. Actually, a high number of bands entails a high detection accuracy, because even minor shifts can be detected. On the other hand, each new band introduces higher probabilities of false alarm raising, which is an unavoidable phenomenon of every change detection mechanism.

Finally, the fact that the signal is integrated over sliding windows introduces a lower bound for detection delay: in the worst case scenario, which is a change occurring at the very first sample of a new window, the detection delay will be of W samples at least.

7.2 Perspectives for future work

The directions for future improvements are various, from both theoretical and implementative perspectives. We would like to explore:

- a more articulated mechanism to reschedule sampling frequency after change detection, possibly avoiding the activation of new frequency bands, which might be unpractical in a real installation;
- the systematic utilization of multivariate change detection tests, which should leverage the phenomenon of false alarm raising due to the increase of the band number;
- a wider exploration of the type of filter employed, which could be application-specific;
- the collection of a more appropriate dataset for ventricular flutter detection, or for any other phenomenon showing the sudden appearance of a higher frequency. We are convinced that this would confirm the effectiveness of the proposed approach and improve the results on the figures of merit.

- the design and development of an ad-hoc system to test the proposed approach on a real observed phenomenon: in this work, simulations and experiments have been carried out on existing datasets and brought on a real embedded device in a simplified version; it would be interesting to actually build a CPS to test the improvement given by the application of what has been analyzed so far in this thesis;
- the integration of the proposed approach with a network protocol to better distribute resources over the interested area. In effect, this thesis studies a method to reduce energy at the sensor level and does not consider the possible intervention of the sink or of the remote server. An exploration of this second aspect would lead to a mechanism including every aspects of a Cyber-Physical System and proposing an all-round approach.

Our work was inspired by the idea of improving the existing energy consumption management methods. An integrated approach considering several aspects of power reduction can be a valid alternative to the existing ones and contributes to the further development and employment of these arising technologies.

8 Bibliography

[AADFR09] Cesare Alippi, Giuseppe Anastasi, Mario Di Francesco, and Manuel Roveri. Energy management in wireless sensor networks with energy-hungry sensors. *Instrumentation and Measurement Magazine, IEEE*, 2009.

[AAG+07] Cesare Alippi, Giuseppe Anastasi, Cristian Galperti, Francesca Mancini, and Manuel Roveri. Adaptive sampling for energy conservation in wireless sensor networks for snow monitoring applications. In *Mobile Adhoc and Sensor Systems, 2007. MASS 2007. IEEE International Conference on*, 2007.

[ABR11a] C. Alippi, G. Boracchi, and M. Roveri. An effective just-in-time adaptive classifier for gradual concept drifts. In *Neural Networks (IJCNN), The 2011 International Joint Conference on*, pages 1675 –1682, 31 2011-aug. 5 2011.

[ABR11b] Cesare Alippi, Giacomo Boracchi, and Manuel Roveri. A distributed self-adaptive nonparametric change-detection test for sensor/actuator networks. In Timo Honkela, Wlodzislaw Duch, Mark Girolami, and Samuel Kaski, editors, *Artificial Neural Networks and Machine Learning ICANN 2011*, volume 6792. Springer Berlin Heidelberg, 2011.

[ABR11c] Cesare Alippi, Giacomo Boracchi, and Manuel Roveri. A just-in-time adaptive classification system based on the intersection of confidence intervals rule. *Neural Networks*, 2011.

[AHY15] Noura Al-Hoqani and Shuang-Hua Yang. Adaptive sampling for wireless household water consumption monitoring. *Procedia Engineering*, 119: 1356 – 1365, 2015.

[Ali14] Cesare Alippi. *Intelligence for Embedded Systems: A Methodological Approach.* Springer Publishing Company, Incorporated, 2014.

[BN93] Michelle Basseville and Igor V. Nikiforov. *Detection of Abrupt Changes: Theory and Application.* Prentice-Hall, Inc., 1993.

[BWCW13] M. Z. A. Bhuiyan, G. Wang, J. Cao, and J. Wu. Energy and bandwidth-efficient wireless sensor networks for monitor-

ing high-frequency events. In *Sensor, Mesh and Ad Hoc Communications and Networks (SECON), 2013 10th Annual IEEE Communications Society Conference on*, pages 194–202, June 2013.

[CH08] Supriyo Chatterjea and Paul Havinga. *Distributed Computing in Sensor Systems: 4th IEEE International Conference, DCOSS 2008 Santorini Island, Greece, June 11-14, 2008 Proceedings*, chapter An Adaptive and Autonomous Sensor Sampling Frequency Control Scheme for Energy-Efficient Data Acquisition in Wireless Sensor Networks, pages 60–78. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.

[DRAP15] G. Ditzler, M. Roveri, C. Alippi, and R. Polikar. Learning in nonstationary environments: A survey. *IEEE Computational Intelligence Magazine*, 10(4): 12–25, Nov 2015.

[GAG+00] Ary L. Goldberger, Luis A. N. Amaral, Leon Glass, Jeffrey M. Hausdorff, Plamen Ch. Ivanov, Roger G. Mark, Joseph E. Mietus, George B. Moody, Chung-Kang Peng, and H. Eugene Stanley. Physiobank, physiotoolkit, and physionet: Components of a new research resource for complex physiologic signals. *Circulation*, 101(23): e215–e220, 2000.

[Gre86] SD Greenwald. Development and analysis of a ventricular fibrillation detector. Master's thesis, MIT Dept. of Electrical Engineering and Computer Science, 1986.

[GSBH11] M. Gupta, L. V. Shum, E. Bodanese, and S. Hailes. Design and evaluation of an adaptive sampling strategy for a wireless air pollution sensor network. In *Local Computer Networks (LCN), 2011 IEEE 36th Conference on*, pages 1003–1010, Oct 2011.

[JC04] Ankur Jain and Edward Y. Chang. Adaptive sampling for sensor networks. In *Proceeedings of the 1st International Workshop on Data Management for Sensor Networks: In Conjunction with VLDB 2004*, DMSN '04, pages 10–16, New York, NY, USA, 2004. ACM.

[LCJ+09] Yee Wei Law, Supriyo Chatterjea, Jiong Jin, Thomas Hanselmann, and Marimuthu Palaniswami. Energy-efficient data acquisition by adaptive sampling for wireless sensor networks. In *Proceedings of the 2009 International Conference on Wireless Communications and Mobile Computing: Connecting the World Wirelessly*, 2009.

[LCS06] H. Liu, A. Chandra, and J. Srivastava. esense: energy efficient stochastic sensing framework for wireless sensor platforms. In

Information Processing in Sensor Networks, 2006. IPSN 2006. The Fifth International Conference on, pages 235–242, 2006.

[Lee08] E. A. Lee. Cyber physical systems: Design challenges. In *Object Oriented Real-Time Distributed Computing (ISORC), 2008 11th IEEE International Symposium on*, pages 363–369, May 2008.

[ME10] M. Mishali and Y. C. Eldar. From theory to practice: Sub-nyquist sampling of sparse wideband analog signals. *IEEE Journal of Selected Topics in Signal Processing*, 4(2): 375–391, April 2010.

[Mit01] Sanjit K. Mitra. *Digital Signal Processing: A Computer-Based Approach*. McGraw-Hill School Education Group, 2nd edition, 2001.

[Pra10] C. Prati. *Segnali e sistemi per le telecomunicazioni*. McGraw-Hill, 2010.

[Pro08] Cornell University Bioacoustics Research Program. Blue whale project, 2008.

[Tac15] Giada Tacconelli. Learning-and-forgetting in nonstationary environments: a multivariate change detection mechanism. Master's thesis, Politecnico di Milano, 2015.

[TAGL15] Samba Traore, Babar Aziz, Daniel Le Guennec, and Yves Louet. Adaptive non-uniform sampling of sparse signals for green cognitive radio. *Computers and Electrical Engineering*, 2015.

[TWCH07] Y. C. Tseng, Y. C. Wang, K. Y. Cheng, and Y. Y. Hsieh. imouse: An integrated mobile surveillance and wireless sensor system. *Computer*, 40(6): 60–66, June 2007.

[VA04] Mehmet C. Vuran and Ian F. Akan. Spatio-temporal correlation: Theory and applications for wireless sensor networks. *Comput. Netw.*, 45(3): 245–259, June 2004.

[WaJR+05] Geoffrey Werner-allen, Jeff Johnson, Mario Ruiz, Jonathan Lees, and Matt Welsh. Monitoring volcanic eruptions with a wireless sensor network. In *in Proceedings of the Second European Workshop on Wireless Sensor Networks*, 2005.

[ZR07] Jing Zhou and David Roure. Floodnet: Coupling adaptive sampling with energy aware routing in a flood warning system. *Journal of Computer Science and Technology*, 22(1): 121–130, 2007.

[ZRDH15] J. Zhang, L. Ren, Y. Ding, and K. Hao. Adaptive sampling algorithm with endocrine regulation mechanism for wireless sensor network. In *Intelligent Systems and Knowledge Engineering (ISKE), 2015 10th International Conference on*, pages 502–507, Nov 2015.

[ZZSF11] C. Zhang, X. Zhu, Y. Song, and Y. Fang. C4: A new paradigm for providing incentives in multi-hop wireless networks. In *INFOCOM, 2011 Proceedings IEEE*, pages 918–926, April 2011.