Non-Intrusive Load Disaggregation based on Digital Signal Processing for Microcontroller Application

Maximiliano E. Véliz, and Alejandro D. Otero

***Abstract*—This paper presents a low complexity non-intrusive load monitoring (NILM) approach for residential electric power based on digital signal processing. The aim is to identify the real operating frames of each household appliance from the aggregated current signal frames. In the methodology, two detection methods are derived. The first method, named *direct method*, identifies the active frames of each device by identifying the most probable combination between devices in the aggregated signal. The second method, termed *indirect method*, identifies the active frames of a particular device by means of a projection of the aggregated signal onto a Fourier subspace representing the characteristic footprint of the device. The methodology is tested on 4 datasets collected in Argentina and high performance metrics are achieved. A pilot test is carried out with an ATSAMD21G18 microcontroller on the Itsy Bitsy M0 Express.**

Link to graphical and video abstracts, and to code: https://latamt.ieeer9.org/index.php/transactions/article/view/9695

***Index Terms*—Energy Disaggregation, Low complexity NILM Methology, Signal Processing, Household Appliances.**

I. Introduction

THE first non-invasive load sensing system for residential buildings was proposed by Hart in the 1990s [1]. Since then, energy disaggregation has been improved with further proposals [2]–[5]. However, according to a recent survey, many challenges remain in relation to a general approach to Non-Intrusive Load Monitoring (NILM) [6]. The main issues include: the need to train the algorithm locally as a strategy to reduce the classification error, the standardisation of NILM performance evaluation, the detection of multiple simultaneous device switching, the correct estimation in a practical scenario with noisy data, the detection of small plug-in electrical loads, the privacy of user data and the challenge of implementing a residential NILM approach on a massive scale with low-cost equipment in developing countries where no public datasets are available.

In terms of algorithmic evolution, Hart's method was based on segmenting the aggregate power signal by quantising active and reactive powers. However, this method could not detect multiple devices connected simultaneously [1], [7]. Since Hart's contribution, the problem of non-intrusive unbundling of loads has been extensively studied in the literature. A review of the state of the art [4] identifies multiple approaches to the problem, including the study of steady and transient state behaviour of electrical parameters, investigation of high- and low-frequency characteristics, study of voltage, current and power waveforms, analysis of harmonics in the frequency domain and combination of parameters to obtain new characteristics. These approaches give rise to different techniques for tackling the problem.

In recent years, advances in the field have led to a boom in machine learning approaches [8], [9].

The authors of [10] present a method for disaggregating energy based on neural networks. In the review [11], the use of different classifiers and their requirements on public datasets such as UK–DALE [12], REDD [13] and BLUED [14] is examined.

In order to design, test, and evaluate the performance of energy disaggregation algorithms, researchers need access to data on the consumption of equipment and appliances. It is evident that there is no universally adopted criterion for generating a database for the study of the non-invasive disaggregation of loads (see for example [12]–[14]). One reason for this is that, according to the approach to the subject adopted, the type of data to be evaluated changes considerably, depending on whether the macroscopic or microscopic characteristics of the loads under study are considered, and also whether the signal is sampled at high or low frequency, as described in [15]. In addition, it should be noted that the characteristics of household appliances vary from country to country, making it difficult to universally train machine learning algorithms, as shown in [16].

In Argentina, there are currently no specific implementations of techniques and algorithms for identifying and disaggregating electrical energy. Furthermore, there is no publicly available dataset containing local household appliance consumption data on which to test conventional machine learning approaches. In this context, we present an NILM approach that has been developed and validated using a local dataset. This digital signal processing-based methodology requires a training sample (frame) of each device to be disaggregated, and aims to identify the active nominal operating frames of

The associate editor coordinating the review of this manuscript and approving it for publication was Mario R. Arrieta Paternina (*Corresponding author: Maximiliano Véliz*).

We would like to thank the Institute of Industry at the National University of General Sarmiento for funding the project 'Design and manufacture of embedded systems for instantaneous reporting of electricity costs, voltage threshold detection and efficient energy management (Rector's Resolution No. 31474/2025)' and the National University of Hurlingham for funding the project 'Cortex-based data acquisition system with open-source hardware for teacher training in technology (Rector's Resolution 201/2024)'.

Maximiliano Véliz is with the Institute of Industry at the National University of General Sarmiento and at the Laboratory for Applied Research on Production and Labour at the National University of Hurlingham (e-mail: mveliz@campus.ungs.edu.ar).

Alejandro D. Otero is with the College of Engineering, University of Buenos Aires & Computational Simulation Center for Technological Applications, Buenos Aires, Argentina (e-mail: aotero@fi.uba.ar).

each device within the aggregated current signal. Compared to existing NILM methods, this is the first method based on constructing a home appliance detection indicator from the projection of signal frames obtained from a single training sample of each device. Our experimental study was carried out on datasets with unique characteristics in a local setting in Argentina, involving the acquisition and analysis of synchronous current signals at 256 samples/cycle to study on-off events of up to three different appliances alongside the aggregated signal. The dataset used in our proposal contains the ground truth of events synchronised with the voltage signal.

The literature contains research articles that demonstrate the implementation of NILM techniques on low-spec computers, such as the Raspberry Pi, or in the cloud, due to the high computational cost required for supervised and unsupervised machine learning and neural networks [17]. In particular, the work of [18] presents a measurement system that uses a NILM algorithm embedded in a Texas Instruments CC3200 chip, which is a microcontroller with integrated Wi-Fi. The algorithm used requires a large data set (at least 1000 samples of the circular buffer) for its calculation and a current sampling of 10 kHz.

The main objective of this work, which addresses some of the current challenges, is to present a novel low complexity methodology for the disaggregation of electrical energy consumption in households. Two methods are developed, based on the transformation of the signal into the frequency domain and its projection onto the Fourier subdomains associated with the individual appliances to be detected, which make it possible to detect whether they are active or not. The advantages of the proposed method are that it requires a current sampling rate of at least 1 kHz, it only needs one training sample per device state for the application, and the classification can be performed directly on a low-cost development board with microcontroller.

The following sections present the common definitions and the basis of the methods, describe four real Argentinean household datasets used to test the implemented methodology, report the results of appliance detection in these datasets, introduce a pilot application with low-cost hardware based on the Itsy bitsy M0 Express, and draw conclusions.

II. NILM Methodology

Despite the numerous NILM approaches proposed in the state of the art [19], we found at least two challenges for mass implementation in developing countries: one is the cost of the smart meters, which usually include a personal computer or cloud computing services to process the data, and the other is the availability of local datasets, as most public datasets come from developed countries [20].

The NILM methodology proposed here is based on the construction of indicators derived from the current waveform which make it possible to identify the active or inactive periods of each device from the aggregated current signal. The approach has the advantage that only one nominal training frame is required for devices with a single operating state. Two identification methods based on the orthogonal projection of signal frames are presented. The first method, named *direct method*, identifies the active frames of each device by identifying the most probable combination between devices in the aggregated signal. The second method, called *indirect method*, identifies the active frames of a particular device by means of a projection of the aggregated signal onto a Fourier subspace representing the characteristic footprint of the device.

The methods are introduced below and their computational cost will be presented later to evaluate the possibility of integrating them into microcontroller-based systems. The orthogonal projection mentioned above is based on splitting the signals into fixed length frames, which are then Fourier transformed. Both techniques rely on the scalar product of Fourier coefficient vectors, which represents the convolution of signals in the time domain. However, indirect projection also uses Gram-Schmidt orthogonalisation to progressively eliminate the contribution of devices one at a time.

A. Signal Frames

The signals to be analysed, belonging to individual appliances and aggregated consumption, are sampled at a fixed rate to obtain their discrete counterparts. These discrete signals are then divided into frames of fixed length at regular intervals. Taking N as the number of samples in each frame, and a time-limited windowing function $w[n]$ (null outside the interval $(0, N)$ and unitary inside), each signal frame is computed as

$$x_l[n] = w[n]x[n + lH] \tag{1}$$

where $n \in \{0, \ldots, N-1\}$ is the local time index (i.e. relative to the start of the sliding extraction window), N is the window length, l is the frame index, and H is the hop size, i.e. how many samples to slide to the right when a new frame is taken [21], [22]. For each device to be analysed, a nominal signal frame is selected that is representative of its steady state operating condition.

B. Fourier Decomposition

Let $x[n]$ with $0 \leq n \leq N-1$, a discrete sample of a periodic signal of period T in time (we omit the frame index here for clarity), be decomposed into a sum of N harmonically related complex exponentials, i.e. multiples of the fundamental frequency [23]:

$$x[n] = \frac{1}{N} \sum_{k=0}^{N-1} a_k \cdot e^{jk\left(\frac{2\pi}{T}\right)n}. \tag{2}$$

The complex coefficients a_k are obtained from the discrete Fourier transform (DFT) of $x[n]$ as:

$$a_k = \sum_{n=0}^{N-1} x[n] \cdot e^{-jk\Omega_0 n}, \quad k = 0, \ldots, N-1. \tag{3}$$

C. Appliance Characteristic Vectors

Consider the case where a household with P appliances is to be analysed, the first method proposed here requires consideration of all possible combinations of these appliances.

Thus, there will be $2^P - 1$ possible scenarios corresponding to the power set (excluding the empty set) of the set of devices. In this case, the nominal frame of the combination of two or more appliances is constructed by agregating the individual nominal frames in the discrete time domain.

For each appliance ($1 \leq i \leq P$) or combination of them ($P + 1 \leq i \leq 2^P - 1$), if $k_{max} \leq N$ discrete Fourier modes are retained in the analysis, a vector of real-valued coefficients,

$$s_i = \{|a_k|\}, \text{ with } k = 0, \ldots, k_{max}, \tag{4}$$

can be formed. These vectors represent the harmonic footprint of each device or combination to be identified in the aggregated signal. The first method proposed here uses the s_i vectors directly to detect frames where the device or combination is switched on.

The second approach aims to remove the harmonic footprint of all but one device in the aggregated signal in order to detect frames where the device is active. This procedure starts by arranging, for each device i ($1 \leq i \leq P$), a matrix S_i with P rows that contain the k Fourier coefficients of the devices nominal signal frames, $S_i \in \mathbb{R}^{P \times k}$. The rows of S_i are ordered so that the last row corresponds to the device of interest s_i^*,

$$S_i = [r_j, 1 \leq j \leq P], \quad \text{with} \quad r_j = s_j, j \neq P \text{ and } r_P = s_i^*. \tag{5}$$

The essence of the method is to successively remove the signal contributions by orthogonal projections following a Gram-Schmidt process [24]. The resulting orthogonal (not orthonormal) basis $\bar{S}_i$ is then formed as

$$\bar{S}_i = \left\{ \begin{matrix} \phi_1^i \\ \vdots \\ \phi_P^i \end{matrix} \right\}. \tag{6}$$

Thus we define the characteristic vector with the distinctive footprint of device i as the last member of the base $\bar{S}_i$, which is renamed as $\hat{\phi}_i = \phi_P^i$.

D. Direct and Indirect Detection of Active Frames

The scalar product of Fourier coefficient vectors, which represents the convolution of signals in the time domain, quantifies the projection (contribution) of one signal onto another [25].

The first detection method requires considering the projection of the nominal frames s_i of each of the appliances and all possible combinations of them (including also that with all appliances turned on). Let q_l be the vector with the first N Fourier modes of the l-th frame of the aggregated signal y_l. Thus, we define the *direct active frame detection coefficient* of each of the combinations i on the l-th frame of the aggregated signal y_l as the orthogonal projection of s_i on q_l normalised by the length of q_l.

$$\alpha_{i,l} = \frac{\langle s_i, q_l \rangle}{\langle q_l, q_l \rangle}, \tag{7}$$

with $\langle \cdot, \cdot \rangle$ the usual inner product in $\mathbb{R}^N$.

This coefficient can be considered as a measure of the similarity of the two signals, so the ideal condition for detecting an active frame of combination i in frame l of the aggregated signal is $\alpha_{i,l} = 1$. A real-world application will face situations where the nominal frames of the devices are not orthogonal, there is noise or other appliances affecting the aggregated signal, etc. Therefore, in the direct method an appliance or combination i ($1 \leq i \leq P$) is selected as active if its α is the closest to 1 among all combinations:

$$i \quad \text{such that} \quad \alpha_{i,l} = \min\{|\alpha_{j,l} - 1| : 1 \leq j \leq P\}. \tag{8}$$

In order to identify frames where all devices are off, a threshold is defined such that if $\alpha_{i,l} \leq \varepsilon_\alpha$ all devices are assigned an inactive state.

Similarly, for the second detection method, let us define the *indirect active frame detection coefficient* as the orthogonal projection of $\hat{\phi}_i$ onto q_l normalised by the length of $\hat{\phi}_i$,

$$\beta_{i,l} = \frac{\langle \hat{\phi}_i, q_l \rangle}{\langle \hat{\phi}_i, \hat{\phi}_i \rangle}. \tag{9}$$

Thus, the condition to determine that appliance i is active in frame l is

$$\beta_{i,l} \geq \varepsilon_\beta, \tag{10}$$

where ε_β acts as a threshold to account for non-ideal effects in the system and frame projection errors, given that the ideal detection value would be $\beta_{i,l} = 1$. ε_β should be as close to 1 as possible to increase detection accuracy.

E. Stages of Application

The application of this new methodology involves four stages: training, testing, detection and disaggregation. The training stage consists of acquiring and tagging a nominal operating frame for each of the appliances connected to the circuit. Fig. 1 shows an example of the selection of the nominal operating frame of a particular appliance (in this case a microwave oven). As an example, the resulting individual nominal frames of the first dataset used in this work are also shown in Fig. 5.

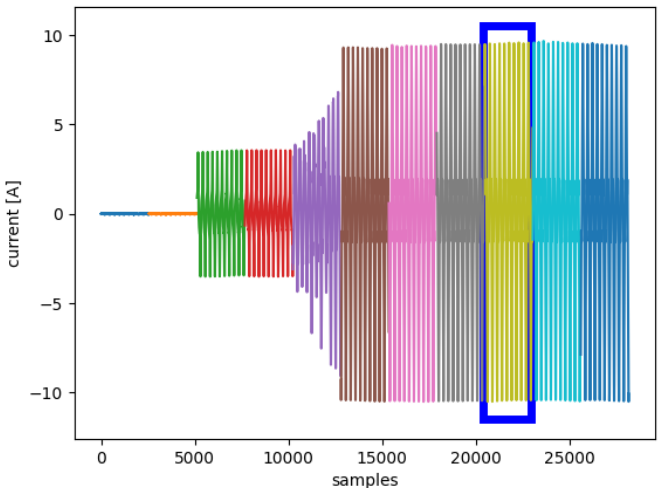

Fig. 1. Current signal of a microwave during power-on, segmented into fixed-length frames. One frame corresponding to the nominal operating state is selected for analysis.

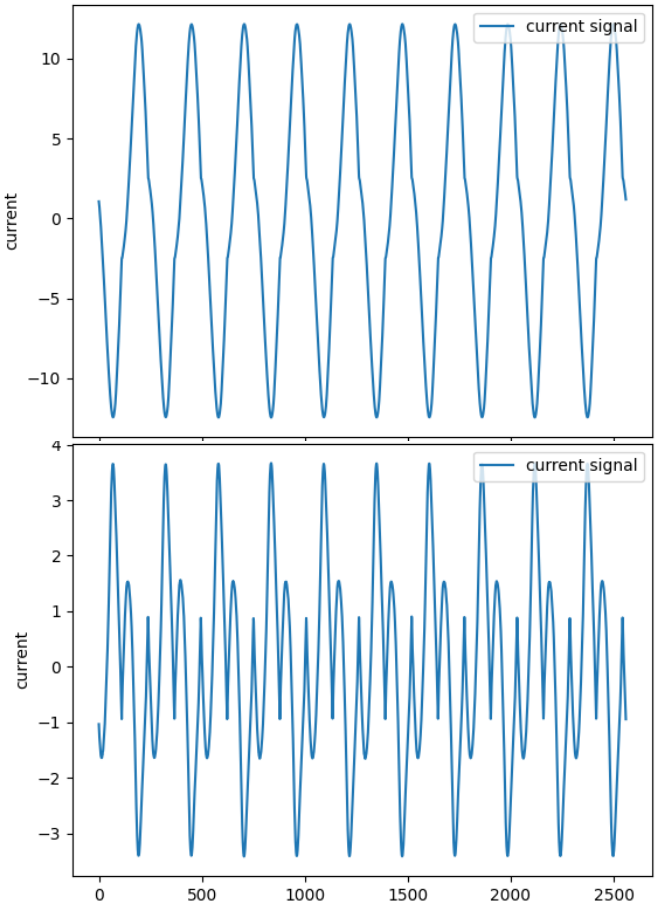

Fig. 2. Nominal combination of frames in the training stage: microwave + toaster + fridge. Top: signals added with correct polarity, bottom: polarity of microwave reversed from that shown in Fig. 5.

Once the nominal frames of each appliance have been defined, it is necessary to add them together to form the nominal frames of the combinations of two or more of them, in order to obtain the $2^P - 1$ possible scenarios to be tested. To do this, the frame alignment condition is required before adding the nominal frames of two or more devices. The phase angle between the nominal frames voltages must be zero and the polarity must be the same. Fig. 2 shows an example of the sum of frames with correct phase and polarity and the sum of frames with incorrect polarity. In the case of synchronous signal acquisition, it is only necessary to ensure the correct polarity between the frames, but in any other case where synchronisation is not guaranteed, it is possible to use techniques such as Dynamic Time Warping (DTW) [26] to perform the correct alignment between frame voltages. This ensures that the current signal of each device is summed with the correct phase shift.

In the case of the indirect method, the Gram-Schmidt procedure described in section II-C must be followed after the nominal frames of the devices have been identified.

During the test phase, the aggregated signal is divided into fixed length frames and the Fourier coefficients are obtained for each of them. The $\alpha_{i,l}$ or $\beta_{i,l}$ indicators are then calculated according to the detection method, direct or indirect, to be used.

Finally, in the disaggregation stage, the active frames of each appliance are quantified directly or indirectly from the values of $\alpha_{i,l}$ or $\beta_{i,l}$, and the energy consumed by each appliance is estimated from the number of active frames detected.

III. Experimental Datasets

In order to validate the methodology, four datasets were used. These were associated with different study scenarios in residential homes connected to the Argentine electricity distribution network.

A Dranetz Power Guia® 440S three-phase energy and power quality analyser was used to acquire the signals. It has a sampling rate of 256 samples/cycle and 8 independent channels: 4 for voltage and 4 for current.

This allows pure current samples to be taken synchronously from all four channels, considering a single voltage channel as a reference (phase A/aggregated voltage).

In a real-world operational setup, the signal acquisition of the individual devices will be performed in the training stage and probably sequentially, while in the test stage only the aggregated signal will need to be acquired, resulting in a single channel device being sufficient for the entire methodology.

In all four datasets, channel A corresponds to the aggregated current signal, while channels B to D correspond to the individual appliances whose consumption is to be disaggregated as described below.

Dataset 1 - DS1: The first dataset consists of 309760 samples. The appliances connected to the current channels are:

- Channel B: microwave (M),
- Channel C: toaster (T),
- Channel D: fridge (F).

The on/off event sequence (Fig. 3) was as follows: the fridge was switched on, followed by the microwave and then the toaster. The microwave was then switched off, and finally, the toaster was switched off. For DS1 the possible combinations to be considered in the direct method are: M only, T only, L only, M+T, M+F, T+F, M+T+F.

Dataset 2 - DS2 (194560 samples, Fig. 4): Includes a small plug-in electric load (PEL), an LED TV (L), for detection and also consists of a microwave (M) and a toaster (T). The on/off sequence was as follows: microwave on, toaster on and TV on; then microwave off, toaster off, TV off.

Dataset 3 - DS3 (199680 samples): Composed of three small plug-in electric load (PEL) for detection, including a variable load notebook charger (N), apart from a cellular phone charger (C) and LED TV (L).

Dataset 4 - DS4 (130560 samples): Consisting of a toaster (T), a drill (D) chosen to represent an intermittent load and a variable load notebook charger (N).

IV. Application

The training stage consists of recording and selecting the current signal frame, referred as nominal, of each equipment/appliance to be detected. The frame length adopted in this analysis is 2560 samples, i.e. each $x[n] \in \mathbb{R}^{2560}$,

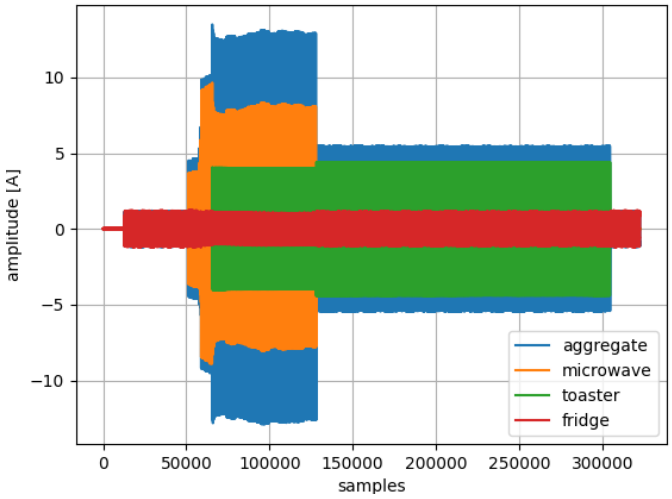

Fig. 3. DS1: Recording of four synchronous current signals, representing the aggregated current signal and those from three appliances: toaster, microwave, and fridge.

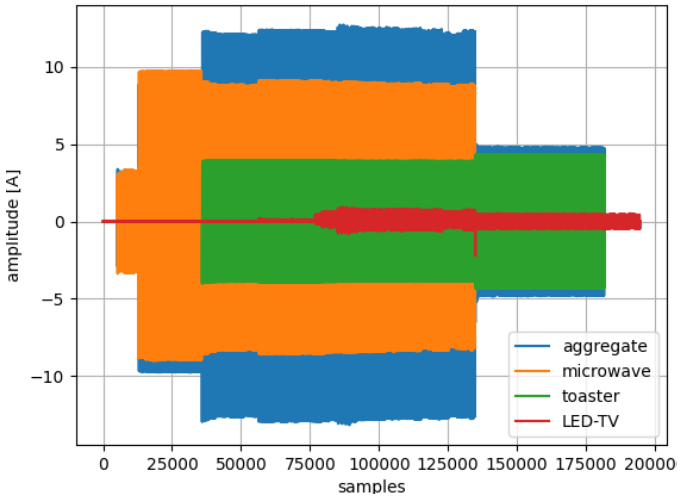

Fig. 4. DS2: Recording of four synchronous current signals, representing the aggregated current signal and those from three appliances: microwave, toaster, LED-TV.

corresponding to 10 signal periods. For example, Fig. 5 shows the nominal current frames for the appliances in DS1.

The process described in sections II-A to II-C is performed to obtain the two characteristic vectors s_i and $\hat{\phi}_i$ for each appliance or combination of appliances in the dataset. The length of these vectors corresponds to the number k_{max} of Fourier modes retained in the analysis.

In order to test an aggregated current signal, it must be segmented into frames of fixed length equal to that of the appliance samples. Given the sample length adopted in this study, the number of frames in each dataset is DS1: 126, DS2: 60, DS3: 78, and DS4: 51 frames. In an operational setup of this technique, aggregated signal frames are generated on-the-fly during the acquisition.

For each appliance i in each frame l, the coefficients $\alpha_{i,l}$ and $\beta_{i,l}$ are computed to apply both detection methods. In the next section, this sequence is applied to the different datasets in order to quantify the accuracy of the proposed methodology.

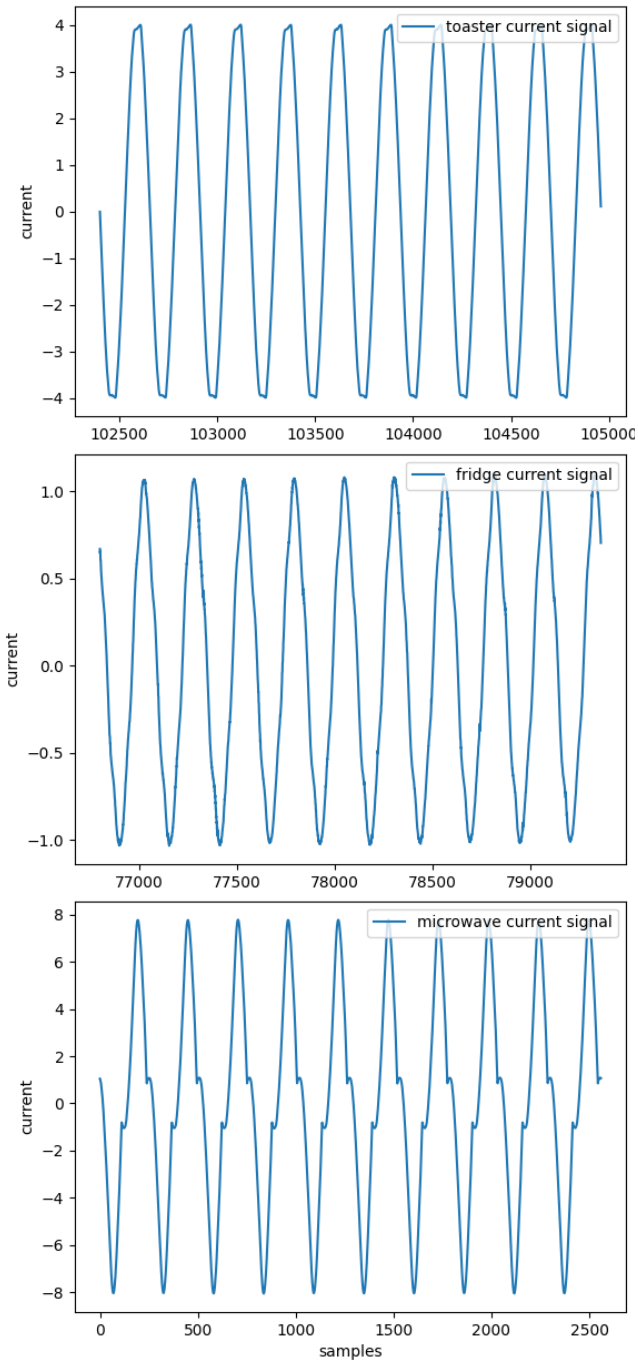

Fig. 5. Selection of nominal DS1 frames in the training phase . Upper: toaster, middle: fridge, and bottom: microwave.

V. Results

In this section, both methods are evaluated on the datasets of section III to assess their performance. This evaluation is done offline on a personal computer with AMD Radeon R7, 3.50 GHz with 8 GB RAM. The next section presents the outcomes of the pilot implementation on a microcontroller-based setup.

A. Direct Method

Fig. 6 illustrates the application of the direct detection method to DS1. The active interval of each device is also

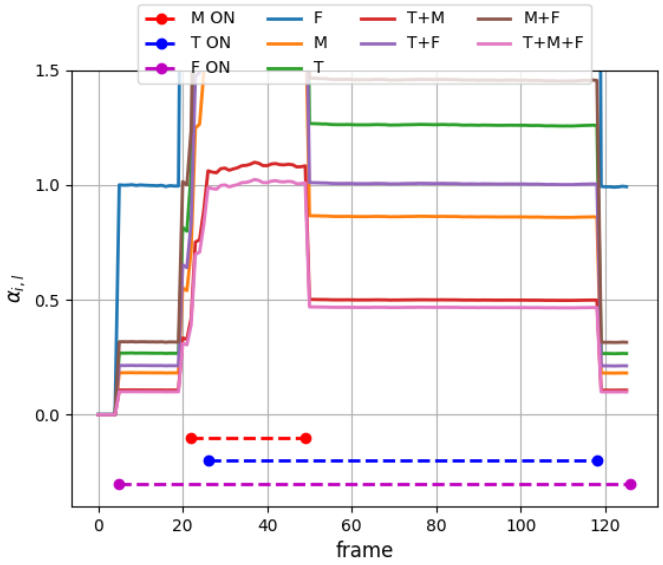

Fig. 6. Direct detection method applied to DS1. Each countinuous lines represents the result of the aggregated signal tested against a particular appliances combination. Dashed lines show the on interval of individual appliances.

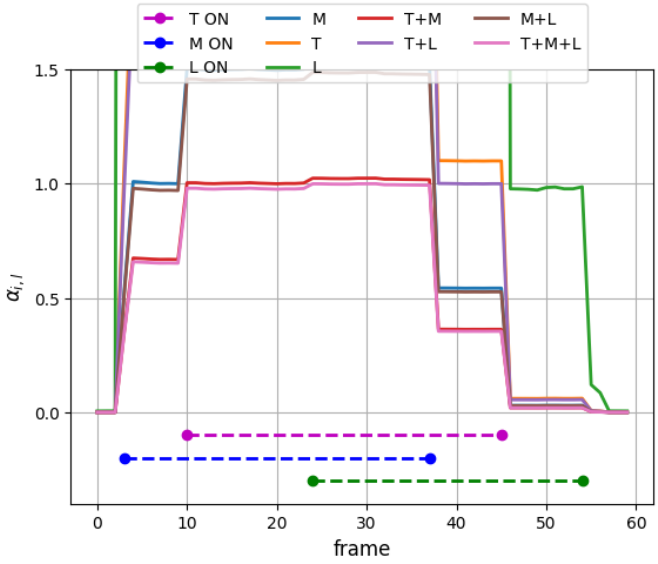

Fig. 7. Direct detection method applied to DS2. Each countinuous lines represents the result of the aggregated signal tested against a particular appliances combination. Dashed lines show the on interval of individual appliances.

indicated in the figure. Each continuous line represents the values of the direct active frame detection coefficient $\alpha_{i,l}$ for a fixed nominal combination i along the aggregated signal frames. In each frame, the curve closest to one indicates the combination positively detected. The detection appears to be unambiguous in all frames, even those occurring during a transient period. It should be noted that in this case, each frame corresponds to a duration of one-fifth of a second. Consequently, errors in a small number of frames will produce an insignificant error in the final disaggregation.

Fig. 7 illustrates a more challenging detection case due to the small PEL characteristic of the LED TV in DS2. It can be seen that the differences between combinations that differ only in whether this load is active are smaller than in the previous case where all loads were significant. Even in this case, the figure shows almost perfect detection.

B. Indirect Method

Fig. 8 demonstrates the application of the indirect detection method to DS1. The continuous line represents the values of the indirect active frame detection coefficient $\beta_{i,l}$ for each appliance. In each case, the active period is evident with coefficient values around 1. Some isolated spikes and dips occur, coinciding with the switching or transient periods of other equipment, lasting no more than a couple of frames.

Similarly, Fig. 9 shows the application of the indirect method to DS2. As with the previous method, this case is made more difficult by the presence of the LED TV. This results in two clean curves for the larger loads: toaster and microwave, but the TV leads to a more oscillatory behaviour. The active states of the first two appliances are clearly identified with coefficient values around unity. On the other hand, for the small PEL, the active period gives values around one, but with very pronounced up and down oscillations. Also, during the inactive period, the switching and transients of the other devices give rise to peaks that can exceed 0.5.

It should be recalled that in this method it is necessary to define the detection threshold ε_β. Given the results above, from now on a compromise value of 0.5 will be selected.

TABLE I
THE TABLE SHOWS THE ACCURACY OF THE DIRECT METHOD ON THE FOUR DATASETS, TAKING INTO ACCOUNT THE CLASSIFICATION OF EACH OF THE FRAMES

Dataset	T	TP	TN	FP	FN	Accuracy
1	378	240	138	0	0	100%
2	180	99	81	0	0	100%
3	234	102	117	0	15	93%
4	153	60	94	0	0	100%

C. Detection Methods Performance

In order to assess the ability of both methods in different scenarios, we apply them to the four datasets and calculate the number of true positive (TP) and negative (TN), and false positive (FP) and negative (FN) detections of each device in each case. We also adopt a definition for the accuracy of the method as the ratio

$$\text{Accuracy} = \frac{TP + TN}{T}, \tag{11}$$

where T is the total number of predictions $(TP+TN+FP+FN)$. These results are reported in Tables I and II.

As shown earlier, both methods achieve high accuracy on the first 2 datasets. The performance of the direct method drops slightly in DS3. To investigate the reasons for this behaviour, Fig. 10 shows the evolution of the $\alpha_{i,j}$ coefficient. This dataset consists of 3 small PELs with very low current levels. There are 3 critical regions where the direct method fails:

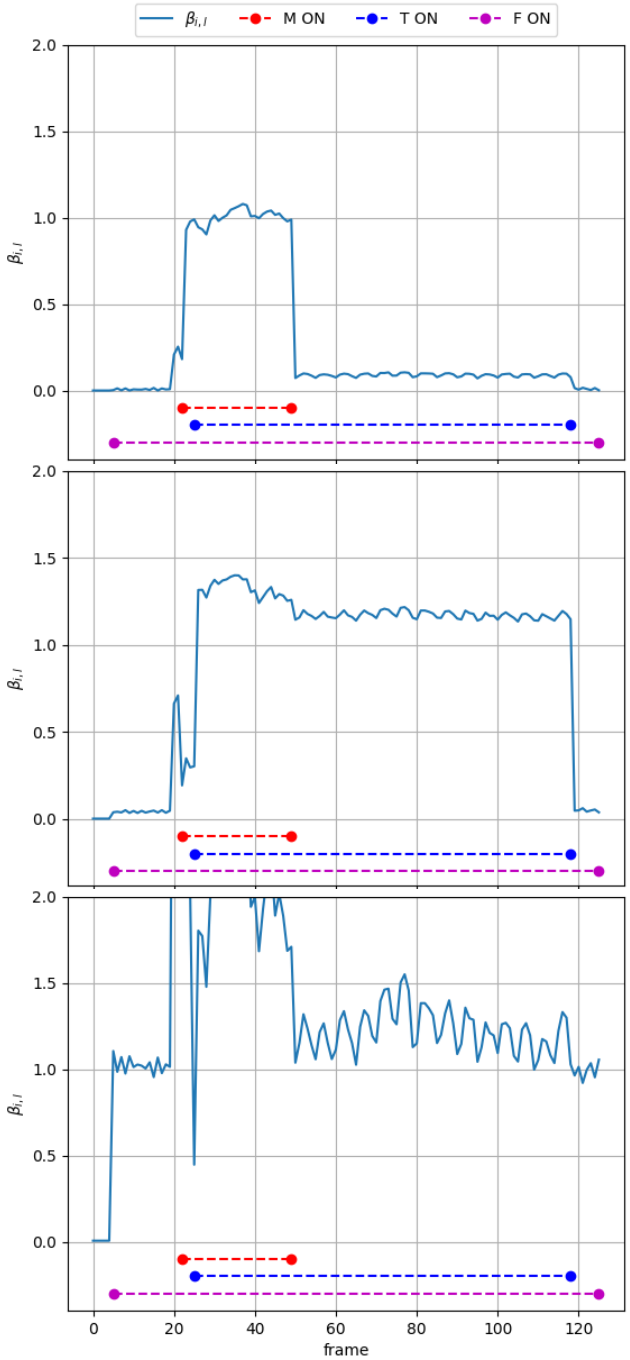

Fig. 8. Indirect detection method applied to DS1. Upper: microwave, middle: toaster, and bottom: fridge. Dashed lines show the on interval of individual appliances.

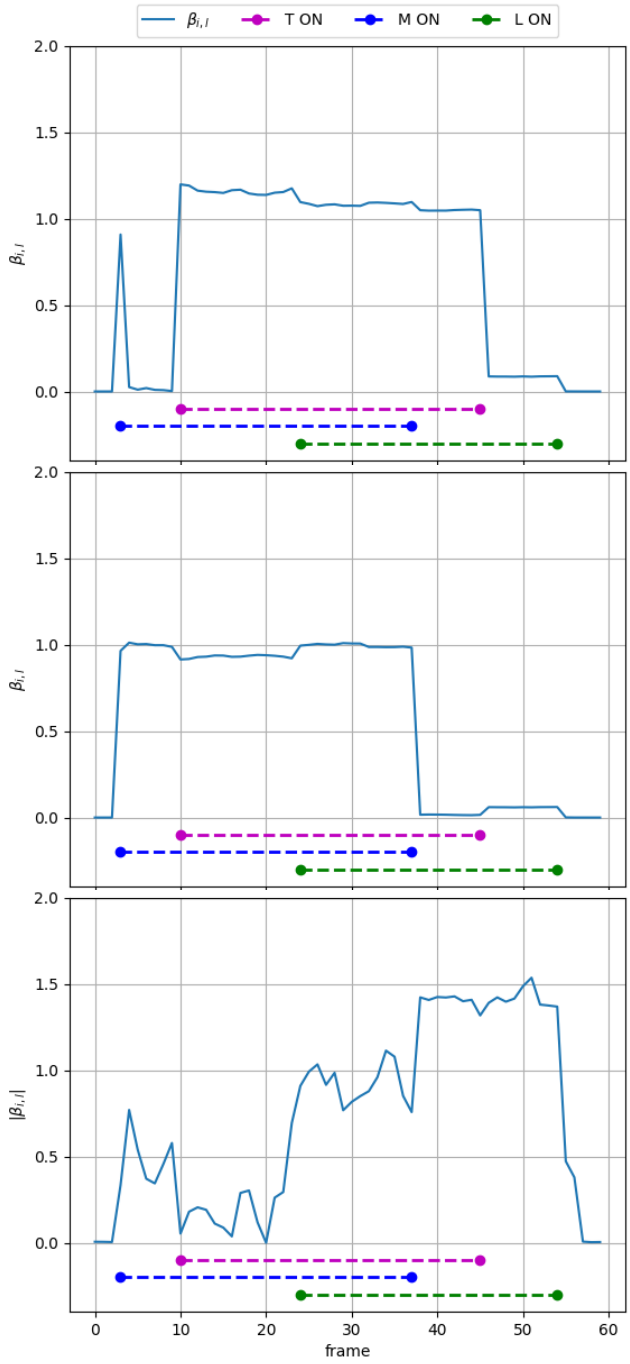

Fig. 9. Indirect detection method applied to DS2. Upper: toaster, middle: microwave, and bottom: led TV. Dashed lines show the on interval of individual appliances.

TABLE II
THE TABLE SHOWS THE ACCURACY OF THE INDIRECT METHOD ON THE FOUR DATASETS, TAKING INTO ACCOUNT THE CLASSIFICATION OF EACH OF THE FRAMES

Dataset	T	TP	TN	FP	FN	Accuracy
1	378	240	138	1	0	99%
2	180	99	81	1	0	99%
3	234	111	101	10	12	90%
4	153	74	79	0	21	87%

- Frames 5 – 15: when the LED TV is switched on, the method assigns a different combination. Looking at the TV's current curve, we can see that during these frames the TV draws an increased current, which then falls back to its nominal value, a kind of warm-up phase.
- Frames 19 – 25: the method does not detect that the cell phone charger is active. Again, looking at the current curve, we observe that in these frames the current drawn is much lower than the nominal current.
- Around frame 55: something similar to what happens

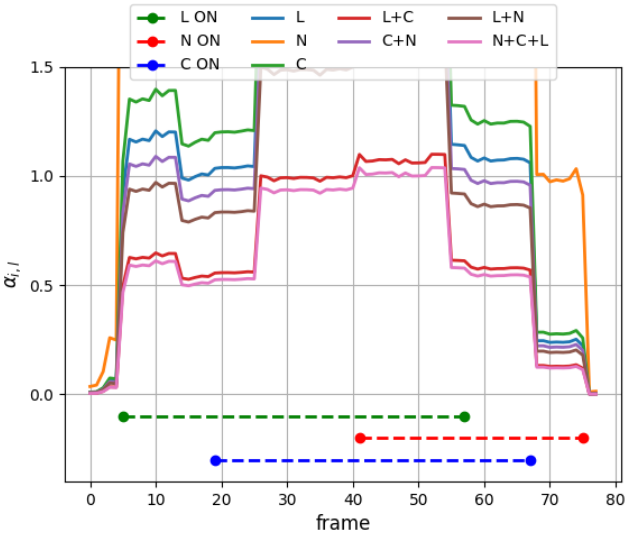

Fig. 10. Direct detection method applied to DS3. Each countinuous lines represents the result of the aggregated signal tested against a particular appliances combination. Dashed lines show the on interval of individual appliances.

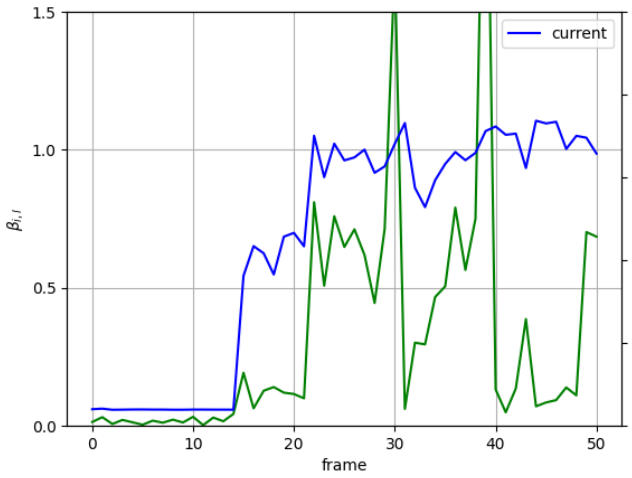

Fig. 11. $\beta_{i,j}$ coefficient (green line - left axis) and current (blue line - right axis) of the notebook in DS4.

during switching on happens, causing the LED TV not to be detected. It seems that there is a low current state produced in the switching off process, which can be found in the current curve.

The performance of the indirect method also deteriorates slightly in DS3 and DS4. In both cases, the main source of error is the detection of the notebook. An explanation for this can be found by comparing the curves for the $\beta_{i,j}$ coefficient and that for the current, see Fig. 11. It can be seen that the regions where the method becomes inaccurate correspond to the partial load operating conditions of the notebook charger (e.g. frames 14 – 21 DS4). In DS4 there are also oscillating values associated with transients in other devices.

The computational complexity of both methods depends on the number of Fourier modes retained in the analysis. In order to assess their suitability for implementation in low-cost

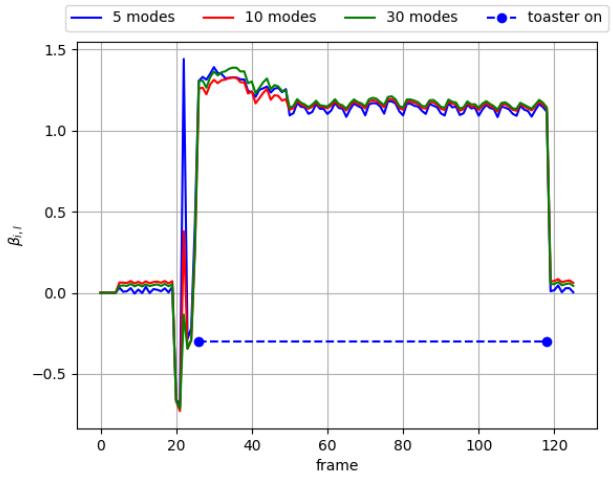

Fig. 12. The figure illustrates how the projection of β coefficient frames varies between different Fourier modes in toaster detection in DS1.

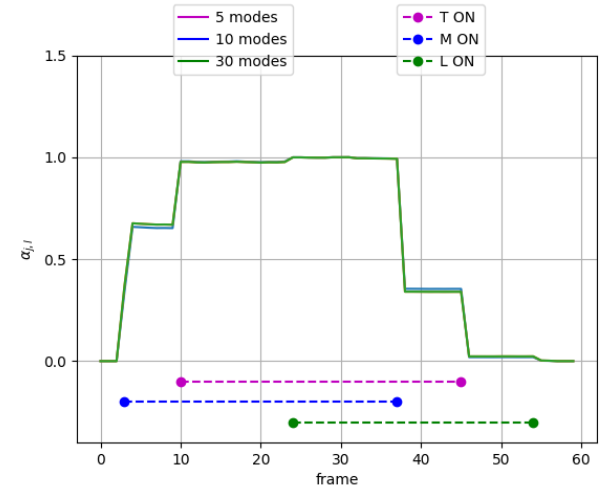

Fig. 13. The figure shows how the frame projection varies using α coefficients as the Fourier modes vary. The combination M+T+TV is tested on the DS4.

microcontroller-based devices, the trade-off between accuracy and computational cost must be considered. As an example of the changes in detection capability, the β coefficients for 5, 10, 30 modes in DS1 are shown in Fig. 12 and the α coefficients of the M+T+L combination in DS4 are shown in Fig. 13. These results are consistent with what has been observed for other detection targets in the four datasets not shown here, suggesting that in some cases good detection can be achieved with as few as 5 modes. To increase detection confidence, 10 modes are used in the microcontroller implementation.

VI. Pilot Application

In this section we present the tests of the register implementation and the calculation of the α and β coefficients, for which we have chosen a low-cost commercial development board. The Itsy Bitsy M0 Express is an Adafruit development board, based on Microchip's ATSAMD21G18 microcontroller

[27]. The main features include: 32-bit ARM Cortex-M0+ core, clock speed up to 48 MHz, 256 KB flash memory, 32 KB RAM, and 2 MB of integrated SPI flash memory available to store files or data. The Itsy Bitsy M0 Express has a 12-bit ADC and can achieve sample rates in the 200-300 kHz range, depending on configuration. However, it has RAM limitations, so it is necessary to optimise the code to handle the signals efficiently. Signals are calculated and processed in chunks, avoiding the use of complete lists that consume too much RAM.

In this pilot setup, the itsy Bitsy M0 express is connected to a measurement board based on the M90E36A chip, which reports 1 current channel via SPI protocol, obtaining current samples at 1 kHz, according to the procedure presented in previous work [28]. All of the processing was implemented using the CircuitPython programming language [29]. Although this did not achieve operational performance, which may be the aim of a next low-level code implementation, it did allow easy testing of the setup capabilities. For the processing test, the flash memory of the Itsy Bitsy M0 express is used to hold the file samples obtained from the M90E36A integrated circuit at 1 kHz.

After successfully testing the pilot implementation, we found that it was possible to store more than 5 minutes of the 1 kHz aggregated current signal in the local flash memory of the Itsy Bitsy (1.14 MB file size). Overall, the processing, which involves separating the frames of the aggregated current signal every 200 samples (10 periods of the signal) and calculating the Fourier coefficients, is the most time-consuming processing operation. For each frame of the aggregated current signal, $\sim$66 s of processing time was required.

The outcomes of the pilot test, in terms of run time and RAM usage (given that the implementation was done using the CircuitPython programming language), are summarised below:

1) Aggregate current signal recording, 1 frame (10 signal cycles), 1 kHz: 802 bytes. file stored in the Itsy Bitsy flash memory.
2) Reading the log file in binary blocks and calculating Fourier coefficients up to the 10th harmonic in 1 frame (200 samples):
   - Run time: 66.7 s.
   - Memory used: 6048 bytes.
3) Computation of the α coefficient:
   - One-time training stage: sum of frames and calculation of Fourier coefficients up to the 10th harmonic, 7 combinations for the three appliances:
     - Run time: 3.3 s.
     - Memory used: 736 bytes.
   - For each frame of the aggregated signal analysed: alpha coefficient calculation for each appliance combination:
     - Run time: 1.5 s
     - Memory used: 2912 bytes
4) Computation of the β coefficient:
   - One-time training stage: the Fourier coefficients for the individual appliances are already computed in the previous step.
   - For each frame of the aggregated signal analysed: construction of detection bases for each appliance following the Gram-Schmidt process on each frame.
     - Run time: 0.26 s
     - Memory used: 7376 bytes

These processing times and memory requirements allow a reporting rate close to $1/min$ with this pilot implementation based on CircuitPython. Foreseeing a low-level native language implementation of this methodology will allow to go well below the minute that is sufficient for the NILM in domestic scenarios. The memory usage corresponds only to dynamic RAM used for data during program execution, as measured by the gc.mem alloc() and gc.mem free() functions in CircuitPython.

VII. Conclusions

In this work, two methods based on digital signal processing are proposed for application as a NILM technique. They are built on the common basis of high frequency sampling of the aggregated signal to be decomposed and its transformation into the frequency domain. Subsequently, these detection methods attempt to recover the signals corresponding to the individual devices by means of projections on the associated Fourier subdomains, in order to finally detect whether they are active or not. The continuous signals are divided into frames of fixed length before being analysed. The main input of both methods are the nominal signal frames of the devices to be detected. This approach has the advantage that only one nominal training frame is required for devices with a single operating state and its implementation can be carried out locally on low-cost microcontroller-based hardware. The classification rate achieved per frame was one minute on average. This could be improved by using a low-level programming language or by implementing it on a development board with higher RAM capacity.

The proposed methodology was tested on 4 household datasets consisting of different combinations of appliances. These datasets included loads with different characteristics: low and high current, variable and intermittent loads were included and combined. In general, high accuracy detection results were obtained. In particular, where loads were significant, detection was almost perfect, with only a few minor detection errors when small loads were included in the datasets. Typically, false positives and false negatives were associated with strong transients generated by the larger loads, or with the change of state of current consumption in multi-state devices. Even in the worst cases, false positives and negatives represented only one or two frames in the series. Given that these datasets were designed with more frequent state changes on the equipment than in normal use cases, it can be expected that in real operating situations each state or combination will last longer, so the accuracy of the methods would be higher. When considering very low current PELs, some inaccuracies were also found for variable load conditions. Since the consumption of these types of loads is very low, the impact of these inaccuracies could be considered negligible in real-world applications.

Regarding the differences between the two methods, there is no obvious performance advantage, but their particular characteristics could make one preferable to the other. In this sense, the direct method might be preferable in some situations as it provides an unambiguous answer, whereas the indirect method requires the threshold coefficient to be chosen. Conversely, the indirect method may offer better scalability as the number of tests in each frame scales with the number of devices, whereas the direct method requires testing all possible combinations.

In this study, scenarios consisting of up to 3 devices were analysed, since the capacity of the acquisition devices allowed the simultaneous acquisition of 4 current channels, thus providing the possibility of obtaining 3 individual signals plus the aggregated one synchronously. Future work will test this methodology on a larger number of devices. Attempts will also be made to quantify the effects of aggregated low current loads.

REFERENCES

[1] G. W. Hart, "Nonintrusive appliance load monitoring," *Proceedings of the IEEE*, vol. 80, no. 12, pp. 1870–1891, 1992. [Online]. Available: https://doi.org/10.1109/5.192069

[2] S. Darby *et al.*, "The effectiveness of feedback on energy consumption," *A Review for DEFRA of the Literature on Metering, Billing and direct Displays*, vol. 486, no. 2006, p. 26, 2006.

[3] B. Neenan, J. Robinson, and R. Boisvert, "Residential electricity use feedback: A research synthesis and economic framework," *Electric Power Research Institute*, vol. 3, 2009. [Online]. Available: https://www.epri.com/research/products/1016844

[4] P. A. Schirmer and I. Mporas, "Non-intrusive load monitoring: A review," *IEEE Transactions on Smart Grid*, vol. 14, no. 1, pp. 769–784, 2022. [Online]. Available: https://doi.org/10.1109/TSG.2022.3189598

[5] R. Gopinath, M. Kumar, C. P. C. Joshua, and K. Srinivas, "Energy management using non-intrusive load monitoring techniques–state-of-the-art and future research directions," *Sustainable Cities and Society*, vol. 62, p. 102411, 2020. [Online]. Available: https://doi.org/10.1016/j.scs.2020.102411

[6] S. Dash and N. Sahoo, "Electric energy disaggregation via non-intrusive load monitoring: A state-of-the-art systematic review," *Electric Power Systems Research*, vol. 213, p. 108673, 2022. [Online]. Available: https://doi.org/10.1016/j.epsr.2022.108673

[7] G. W. Hart, "Residential energy monitoring and computerized surveillance via utility power flows," *IEEE Technology and Society Magazine*, vol. 8, no. 2, pp. 12–16, 1989. [Online]. Available: https://doi.org/10.1109/44.31557

[8] N. Batra, J. Kelly, O. Parson, H. Dutta, W. Knottenbelt, A. Rogers, A. Singh, and M. Srivastava, "NILMTK: an open source toolkit for non-intrusive load monitoring," in *Proceedings of the 5th international conference on Future energy systems*, 2014, pp. 265–276. [Online]. Available: https://dl.acm.org/doi/10.1145/2602044.2602051

[9] P. A. Schirmer and I. Mporas, "Statistical and electrical features evaluation for electrical appliances energy disaggregation," *Sustainability*, vol. 11, no. 11, p. 3222, 2019. [Online]. Available: https://doi.org/10.3390/su11113222

[10] J. Kelly and W. Knottenbelt, "Neural nilm: Deep neural networks applied to energy disaggregation," in *Proceedings of the 2nd ACM International Conference on Embedded Systems for Energy-Efficient Built Environments*, 2015, pp. 55–64. [Online]. Available: DOI:10.1145/2821650.2821672

[11] C. Shin, S. Rho, H. Lee, and W. Rhee, "Data requirements for applying machine learning to energy disaggregation," *Energies*, vol. 12, no. 9, p. 1696, 2019. [Online]. Available: DOI:10.3390/en12091696

[12] J. Kelly and W. Knottenbelt, "The UK-DALE dataset, domestic appliance-level electricity demand and whole-house demand from five uk homes," *Scientific data*, vol. 2, no. 1, pp. 1–14, 2015. [Online]. Available: DOI:10.1038/sdata.2015.7

[13] Z. Kolter and M. J. Johnson, "REDD: A public data set for energy disaggregation research," in *Workshop on data mining applications in sustainability (SIGKDD), San Diego, CA*, vol. 25, no. Citeseer, 2011, pp. 59–62. [Online]. Available: https://zicokolter.com/publications/kolter2011redd.pdf

[14] K. Anderson, A. Ocneanu, D. Benitez, D. Carlson, A. Rowe, and M. Berges, "BLUED: A fully labeled public dataset for event-based non-intrusive load monitoring research," in *Proceedings of the 2nd KDD workshop on data mining applications in sustainability (SustKDD)*, vol. 7. ACM, 2012. [Online]. Available: DOI:10.1145/2339530.2339536

[15] W. A. de Souza, F. D. Garcia, F. P. Marafão, L. C. P. Da Silva, and M. G. Simões, "Load disaggregation using microscopic power features and pattern recognition," *Energies*, vol. 12, no. 14, p. 2641, 2019. [Online]. Available: DOI:10.3390/en12142641

[16] D. Murray, L. Stankovic, V. Stankovic, S. Lulic, and S. Sladojevic, "Transferability of neural network approaches for low-rate energy disaggregation," in *ICASSP 2019-2019 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2019, pp. 8330–8334. [Online]. Available: https://doi.org/10.1109/ICASSP.2019.8682889

[17] Q. Liu, K. M. Kamoto, X. Liu, M. Sun, and N. Linge, "Low-complexity non-intrusive load monitoring using unsupervised learning and generalized appliance models," *IEEE Transactions on Consumer Electronics*, vol. 65, no. 1, pp. 28–37, 2019. [Online]. Available: https://doi.org/10.1109/TCE.2019.2891160?utm_source=chatgpt.com

[18] M. Nardello, M. Rossi, and D. Brunelli, "An innovative cost-effective smart meter with embedded non intrusive load monitoring," in *2017 IEEE PES Innovative Smart Grid Technologies Conference Europe (ISGT-Europe)*. IEEE, 2017, pp. 1–6. [Online]. Available: https://doi.org/10.1109/ISGTEurope.2017.8260242

[19] E. Aladesanmi and K. Folly, "Overview of non-intrusive load monitoring and identification techniques," *IFAC-PapersOnLine*, vol. 48, no. 30, pp. 415–420, 2015. [Online]. Available: https://doi.org/10.1016/j.ifacol.2015.12.414

[20] H. K. Iqbal, F. H. Malik, A. Muhammad, M. A. Qureshi, M. N. Abbasi, and A. R. Chishti, "A critical review of state-of-the-art non-intrusive load monitoring datasets," *Electric Power Systems Research*, vol. 192, p. 106921, 2021. [Online]. Available: https://doi.org/10.1016/j.epsr.2020.106921

[21] M. Müller, *Fundamentals of music processing: Audio, analysis, algorithms, applications*. Springer, 2015, vol. 5. [Online]. Available: https://doi.org/10.1007/978-3-319-21945-5

[22] J. B. Allen and L. R. Rabiner, "A unified approach to short-time fourier analysis and synthesis," *Proceedings of the IEEE*, vol. 65, no. 11, pp. 1558–1564, 1977. [Online]. Available: https://doi.org/10.1109/PROC.1977.10770

[23] A. V. Oppenheim, N. S. Hamid, and A. S. Willsky, *Fourier series representation of periodic signals*. Prentice Hall, 1997.

[24] Å. Björck, "Numerics of Gram-Schmidt orthogonalization," *Linear Algebra and Its Applications*, vol. 197, pp. 297–316, 1994. [Online]. Available: https://doi.org/10.1016/0024-3795(94)90493-6

[25] B. Mulgrew, P. Grant, and J. Thompson, *Digital signal processing: concepts and applications*. Red Globe Press London, 1999. [Online]. Available: https://doi.org/10.1007/978-1-349-14944-5

[26] C. Myers, L. Rabiner, and A. Rosenberg, "Performance tradeoffs in dynamic time warping algorithms for isolated word recognition," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 28, no. 6, pp. 623–635, 1980. [Online]. Available: https://doi.org/10.1109/TASSP.1980.1163491

[27] A. Pinales and D. Valles, "Autonomous embedded system vehicle design on environmental, mapping and human detection data acquisition for firefighting situations," in *2018 IEEE 9th Annual Information Technology, Electronics and Mobile Communication Conference (IEMCON)*. IEEE, 2018, pp. 194–198. [Online]. Available: https://doi.org/10.1109/IEMCON.2018.8615022

[28] M. E. Véliz, G. E. Real, and A. D. Otero, "Flexible and configurable embedded electrical energy measurement system to acquire and process high-frequency features," *HardwareX*, p. e00539, 2024. [Online]. Available: https://doi.org/10.1016/j.ohx.2024.e00539

[29] D. K. Fisher, R. S. Fletcher, and S. S. Anapalli, "Python software integrates with microcontrollers and electronic hardware to ease development for open-source research and scientific applications," *Advances in Internet of Things*, vol. 11, no. 1, pp. 42–58, 2021. [Online]. Available: https://doi.org/10.4236/ait.2021.111004

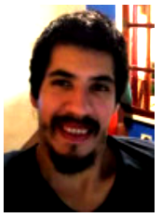

Maximiliano E. Véliz is a teaching researcher at the Institute of Industry, National University of General Sarmiento (UNGS), and the Laboratory of Applied Research on Production and Labor, National University of Hurlingham (UNAHUR). I specialize in research and technological development of embedded systems, with a strong focus on optimizing electrical energy efficiency, particularly in smart grids and distributed generation. Currently, my work centers on Non-Intrusive Load Monitoring (NILM).

Alejandro D. Otero is an Assistant Professor at University of Buenos Aires and Researcher at the CSC-CONICET (Argentina), where he leads the Renewable Energy Group. He specializes in numerical simulation techniques for solving PDE problems from fluid and solid mechanics, heat and mass transfer and multiphysics (fluid–structure interaction, thermomechanics, porous media flow, etc.) which typically appear in sciences and engineering. Lately he devoted to the implementation of those techniques in HPC environments.