Low Cost Non-Intrusive Home Energy Monitoring

DISSERTAÇÃO DE MESTRADO

Amâncio Lucas Sousa Pereira
MESTRADO EM ENGENHARIA INFORMÁTICA

UNIVERSIDADE da MADEIRA
A Nossa Universidade
www.uma.pt

Setembro | 2011

UMa

R Low
1

T/M UMa
004
PER Low
Ex.1

72282

Low Cost Non-Intrusive Home Energy Monitoring

DISSERTAÇÃO DE MESTRADO

Amâncio Lucas Sousa Pereira
MESTRADO EM ENGENHARIA INFORMÁTICA

UNIVERSIDADE DA MADEIRA
SECTOR DE DOCUMENTAÇÃO
E ARQUIVO

ORIENTAÇÃO
Nuno Jardim Nunes

Low Cost Non-Intrusive Home Energy Monitoring

MASTER THESIS

LUCAS PEREIRA

September 2011

Thesis supervisor:

Prof. Dr. Nuno Jardim Nunes
Madeira-ITI and University of Madeira

Madeira Interactive Technologies Institute
University of Madeira
(Portugal)

“Pedras no caminho? Guardo todas, um dia vou construir um castelo...”

- *Fernando Pessoa*

“Stones in my way? I keep them all, one day I will build myself a castle…”

- *Fernando Pessoa*

Abstract

This thesis presents a low cost non-intrusive home energy monitor built on top of Non-Intrusive Load Monitoring (NILM) concepts and techniques. NILM solutions are already considered low cost alternatives to the big majority of existing commercial energy monitors but the goal here is to make its cost even lower by using a mini netbook as a whole in one solution. The mini netbook is installed in the homes main circuit breaker and computes power consumption by reading current and voltage from the built-in sound card. At the same time, feedback to the users is provided using the 11'' LCD screen as well as other built-in I/O modules. Our meter is also capable of detecting changes in power and tries to find out which appliance lead to that change and it is being used as part of an eco-feedback platform that was build to study the long terms of energy eco-feedback in individuals.

In this thesis the steps that were taken to come up with such a system are presented, from the basics of AC power measurements to the implementation of an event detector and classifier that was used to disaggregate the power load. In the last chapter results from some validation tests that have been performed are presented in order to validate the experiment.

It is believed that such a system will not only be important as an energy monitor, but also as an open system than can be easily changed to accommodate and test new or existing non-intrusive load monitoring techniques.

Keywords: Energy monitoring; Non-Intrusive Load Monitoring; AC Power; Machine Learning; Eco-feedback

Table of Contents

1 **Introduction** **1**
1.1 Motivation 1
1.2 Vision 3
1.3 Document organization 4
2 **State of the art** **6**
2.1 Background research 6
2.1.1 Power theory 6
2.1.2 Smart meters 10
2.1.3 Non-Intrusive Load Monitoring 14
2.2 Literature review 17
2.3 Challenges 20
3 **SINAIS Power meter** **22**
3.1 Profile 22
3.2 Software architecture 24
3.3 Data acquisition 26
3.4 Power calculations 28
3.5 Median filter 31
3.6 Event detection 32
3.6.1 GLR event detector 33
3.7 Feature extraction 35
3.8 Event classification 37
4 **Validation** **39**
4.1 Experimental setup 39
4.2 Power calculations 42
4.2.1 Results 43
4.2.2 Discussion 44
4.3 Event detection 45
4.3.1 Results 45
4.3.2 Discussion 47
4.4 Load disaggregation 48
4.4.1 Feature extraction 48

4.4.2 Event classification 49
4.4.3 Discussion 50
5 **Conclusions and future work** **52**
5.1 Low cost home energy monitoring 52
5.2 Project expansion 52
References **54**
Referenced Web Resources **56**

List of Figures

Figure 2.1: Power, current and voltage graph of a resistive load 7
Figure 2.2: Power, current and voltage graph of a purely reactive load 7
Figure 2.3: Power, current and voltage graph of a partially reactive load 8
Figure 2.4: AC Power triangle 9
Figure 2.5: Current and voltage affected by harmonic content 10
Figure 2.6: a) Kill a Watt P4400; b) Plugwise; c) TED; d) Wattson; e) EnerSure [w8] 12
Figure 2.7: Non Intrusive Load Monitoring workflow 15
Figure 3.8: Difference between netbook soundcards: (a) Toshiba NB300-100; (b) Macbook 13' 23
Figure 3.9: Toshiba NB300 netbook - a low cost non-intrusive home energy monitor 24
Figure 3.10: "Pipe-and-Filter" software architecture 25
Figure 3.11: Low cost Non-Intrusive Home Energy Monitor software architecture 25
Figure 3.12: Power meter installation: (a) sensors installed in the main circuit breaker; (b) sensors connected to the netbook built-in sound card 27
Figure 3.13: Current and voltage at 50 Hz with a sampling rate of 8 kHz 28
Figure 3.14: Real time Voltage RMS, Current RMS and Real Power during 1 minute 31
Figure 3.15: Microwave transients at 50 Hz: (a) turn-on; (b) turn-off 35
Figure 4.16: Local NILM installation: (a) monitor; (b) appliances being monitored 40
Figure 4.17: Simulation 1 power draw and events 41
Figure 4.18: Simulation 2 power draw and events 41
Figure 4.19: Current validation chart 43
Figure 4.20: Voltage validation chart 43
Figure 4.21: Real power validation chart 44
Figure 4.22: Representation of the windows from where the features will be extracted 49

List of equations

Equation 2.1: Apparent power 8
Equation 2.2: Relationship between real, reactive and apparent powers 9
Equation 2.3: Real Power 9
Equation 2.4: Reactive Power 9
Equation 2.5: Power factor 9
Equation 2.6: Power factor for both linear and non-linear loads 10
Equation 3.7: Real Power definition 29
Equation 3.8: Real Power definition in discrete time 29
Equation 3.9: Voltage root-mean-Square definition 29
Equation 3.10: Reactive Power definition (time delay method) 30
Equation 3.11: Reactive Power definition (time delay method) in discrete time 30
Equation 3.12: Median filter definition 31
Equation 3.13: GLR detection statistics 32
Equation 3.14: Maximum likelihood calculation for the detection window l-k 34
Equation 3.15: Likelihood minimum change of mean of interest 34
Equation 3.16: Normalization definition 36
Equation 3.17: Average power change definition 36
Equation 3.18: N^{th} degree polynomial equation 37
Equation 4.19: Percentage error 44

List of Tables

Table 4.1 : Appliances used for validation 40
Table 4.2 : Simulation 1 appliances and transitions (100 seconds) 40
Table 4.3 : Simulation 2 appliances and transitions (200 seconds) 41
Table 4.4 : Single appliance simulations event count 42
Table 4.5: Consumption results comparison 44
Table 4.6: Event detection results for simulation data 1 46
Table 4.7: Event detection results for simulation data 2 46
Table 4.8: Event detection results for simulation data 1 (test 2) 47
Table 4.9: Event detection results for simulation data 2 (test 2) 47
Table 4.10 : Features extracted from one of the microwave events 49
Table 4.11 : K-Nearest Neighbor classification results 50

1 Introduction

1.1 Motivation **1**
1.2 Vision **3**
1.3 Document organization **4**

1.1 Motivation

Back in 1992 the world leaders got together in Rio de Janeiro for the United Nations Conference on Environment and Development (UNCED) [w1]. Some of the issues addressed were, among others, alternative sources of energy to replace the use of fossil fuels and the growing scarcity of water.

In almost 20 years many actions have been taken to face those issues, with a big focus on improving and creating alternative sources of energy. The building industry has also been shifting gears towards a more environmental friendly World. Green building is the buzzword, and brings with it techniques and practices with the goal of reducing the impact of new buildings in the environment.

Energy and water efficiency are the two key points of green building with solutions like effective window placement (providing more natural light hence reducing the need for electric lighting during the day) and water conserving features (like ultra-low flush toilets and low-flow shower heads). More technical solutions are also available, for instance the on-site generation of renewable energy (e.g. solar power and wind power) or the on-site recycling of *greywater* (wastewater generated from domestic activities such as laundry and bathing) that could be used for tasks like irrigation.

Yet and although well intentioned, green building is still expensive to the average homeowner. On-site power generation for instance carries a very high installation cost and

does not pay off immediately. **In fact, the savings in money come from a more efficient use of the house utilities and not from the building itself**.

But are humans ready to start using natural resources in a more efficient and sustainable manner? The short answer is NO. The truth is that despite humans being at the center of concerns for sustainable development they are not really aware of how their actions and behaviors can affect sustainability.

Electricity is a paradigmatic example of this lack of awareness and this is proved by a series of studies that show significant contradictions between consumer perceptions and their knowledge of energy efficiency. For example in [1] authors show that most humans have a wrong perception of the most effective thing to do when being energy efficient. While there is strong evidence that generally efficiency-improving actions save more than reducing the usage of inefficient equipment, only 11.7% of participants refer to the former while 55.2% pointed out the later. Another important conclusion drawn from the same study is that home users slightly overestimate small energy savings, e.g. replacing an incandescent light bulb with a compact fluorescent lamp (CFL), and largely underestimate when potential savings are high, e.g. users estimate that central air conditioner uses only 1.3 times more energy than a room air conditioner, when in fact it uses 3.5 times as much energy.

The invisible nature of electricity has been the focus of Chisiks' work [2]. Which focus on understanding how people perceive electricity given the fact that it is not easily quantifiable. The findings are quite informative about the lack of perception regarding how much electricity is consumed by a particular device, which users tend to associate with the frequency and duration as well as with the size of the devices. For example, when presented with devices with low to moderate consumption levels respondents would refer size and duration of usage most of the time with statements like *« the television is on all the time so it must be consuming lots of electricity »* and *« the kids use the computer all the time ».*

This working hypothesis that most people lack awareness and understanding of how their everyday behaviors affect the environment is the base for eco-feedback technology which is defined as technology that provides feedback on individual or group behaviors with a goal of reducing environmental impact [3]. Eco-feedback technology has been around for more than 40 years and literature shows that providing feedback to the consumers, even at a low level of disaggregation, may result in savings between 10% and 15% [4]. However, there are studies that show that this effect is not long lasting, and that consumers tend to return to previous consumptions values is a few weeks [5].

As of today there is a renewed interest in eco-feedback technology that is mostly driven by the advances in sensing technology and in means of interaction, that create a unique opportunity to develop novel eco-feedback services that can promote sustainability. The ability to disaggregate power consumption at a low-cost, combined with the widespread of internet based social networks and lower cost smartphones or other handheld devices opens the potential of eco-feedback to millions of households that will be able to view their home resource consumption data in each of their appliances.

Nonetheless, the human half of this human-computer interaction must not be neglected once again, as the transition to sustainability must involve a personal behavior change that will only happen when technology is able to respond to peoples' changing needs, drives and motivations.

Sustainable Interaction with Social Networks, context Awareness and Innovative Services (SINAIS) is a research project run by a multi-disciplinary team working in the HCI domain, under the Carnegie Mellon Portugal partnership who's main focus of research addresses the creation and discovery of methods for people to achieve self-fulfillment sustainability [w2].

The SINAIS project has a very wide scope, involving two main research areas namely advanced infrastructure systems (AIS) and human-computer interaction (HCI). Regarding the AIS side of the project the goal is to create and deploy non-intrusive home energy monitors in several houses. This will involve the use of sensors and machine learning to monitor and make sense of human activities as well as electricity consumption. In the HCI domain, the collected data will be used along with some aspects of motivation theory to create innovative visualization techniques to provide feedback to users. These techniques will then be continuously evaluated, in the search for opportunities to maximize energy conservation and motivate behavior change.

1.2 Vision

This thesis reflects part of the work done in the AIS track of the SINAIS project regarding the home energy monitoring. Presented here are the steps that lead to the creation of a low-cost home energy monitoring device that is capable of monitoring the electricity consumption, and at the same time provide feedback to the consumer using several communication channels.

The energy monitor being created must be capable of monitoring the whole house consumption simply by analyzing the current and voltage in the building we are measuring. Also, the monitor must add the possibility to disaggregate the load. The process of measuring

the whole load being used by a house, as well as disaggregating it into individual loads is called Non-Intrusive Load Monitoring and is the central subject of this thesis.

NILM is not a new subject and it has been around for quite a while now. It is built on top of the premise that the study of how the power consumption changes over time can lead us to information about which appliance has lead to that change.

The key point of this load monitor is low-cost, and this is mainly due to two reasons: 1) the solution needs to pay-off in a short period of time (savings in the electricity bill are not reflected in big numbers, for example, an average consumer may save up to 5 Euros a month, 60 Euros in one year, meaning that the consumers will not be willing to spend a lot of money in a energy monitoring solution), ii) the meter will be deployed in at least 30 houses and this can yield a very high cost, as so final system needs to be very cost effective.

To sum up, one envisions a low cost non-intrusive home energy monitor that is capable of measuring the aggregate consumption, but and also have the ability to disaggregate the load in terms of individual appliances. This system will be installed in number of houses along with several graphical user interfaces, that will provide feedback to the householders. The system also needs to be flexible enough to be used for more technical aspects such as testing new NILM techniques and algorithms.

It is believed that this study will fill a very important gap in the future of NILM as no relevant studies have been made regarding the field of human-computer interaction with real world scenarios and users. In fact until now NILM has been very appliance focused, and simple questions like: *« What information are the users willing to have about their consumption? »; « Will disaggregated information improve how humans perceived electricity consumption? »; « How should that information be displayed? »,* and more than anything: *« Are humans willing to modify they behaviors regarding energy consumption? »* remain unanswered.

1.3 Document organization

This document is organized in the following way. In chapter 2 (State of the art) the most important theoretical aspects behind AC power systems are presented. In the follow up there is a review of smart meters, as well as other electricity monitoring projects. The remaining of this chapter is devoted to non-intrusive load monitoring. The NILM process workflow is defined as well as a supportive literature review. The chapter finished with a review of the major challenges faced in a project of this nature. In chapter 3 (SINAIS Power meter) the low

cost non-intrusive home energy monitor to be developed is described. First an overall view of the solution is given, followed by the definition of the software architecture that will support it. In the remains of the chapter all the modules that make part of the implementation are thoroughly defined. In chapter 4 (Validation) the meter is validated. First the power calculations are validated using ground truth data collected with the help of a professional power meter. Next some tests are perform to see how efficient the system is regarding the event detection and load disaggregation process. This desertation finishes with some conclusions and future work perspectives in chapter 5 (Conclusions and future work).

2

State of the art

2.1 Background research **6**
2.1.1 Power theory 6
2.1.2 Smart meters 10
2.1.3 Non-Intrusive Load Monitoring 14
2.2 Literature review **17**
2.3 Challenges **20**

2.1 Background research

2.1.1 Power theory

As the complexity of electrical appliances increases it is not expected that they will all interact with the electric system in the same way. In fact we can find three different types of loads inside a house, purely resistive, purely reactive and partially reactive.

Appliances that consume all the energy given to them are called resistive loads. Examples of these appliances are irons, kettles and electric cookers. In a purely resistive load current and voltage are in phase, with a resulting waveform similar to the one in *Figure 2.1*.

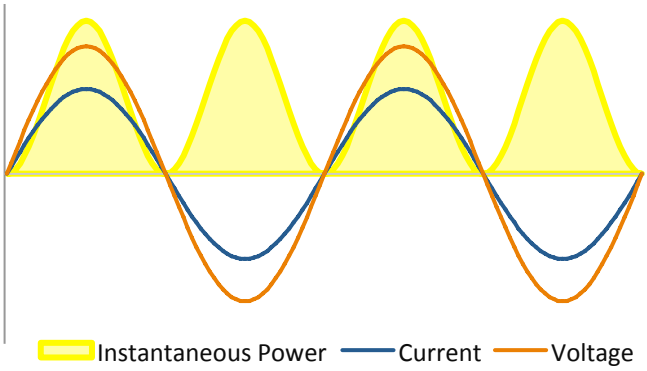

Figure 2.1: Power, current and voltage graph of a resistive load

In this kind of appliances one can see that the instantaneous power (the product of current and voltage at a given time) is always positive indicating that all the energy is flowing to the load. On the other hand appliances that consume no energy are called purely reactive loads. In a purely reactive load no power is dissipated. The case here is that power is alternately absorbed from and returned to the source. In a purely reactive load current and voltage are out of phase exactly ± 90º making the power alternate equally between cycles of positive and negative instantaneous values, which is shown in *Figure 2.2*.

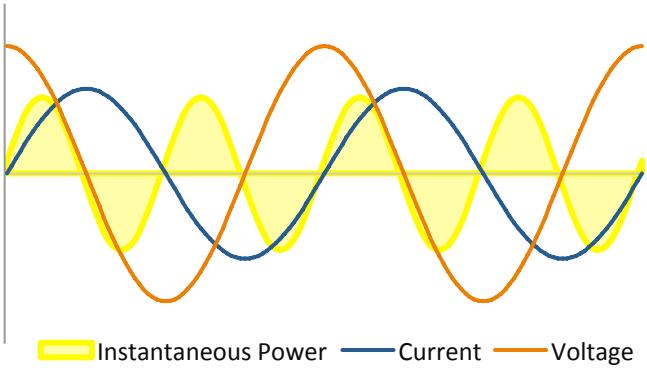

Figure 2.2: Power, current and voltage graph of a purely reactive load

Because of the constant back and forth of energy in the load, purely reactive loads will end up lowering the electric grid efficiency. Fortunately the ± 90º-phase angle between current and voltage is a "worst case scenario" that does not normally happen in real word settings. What one normally get are appliances that take in a certain amount of energy and then release some of that energy back into the supply, like for example fridges, washing machines, pillar drills and arc-welders. These are all partially reactive loads and it can be spoted in *Figure 2.3* how the instantaneous power goes from negative and then positive for a longer period of time, representing the energy flowing from the load and back. These partially reactive loads can be

either inductive, when the current lags the voltage (e.g. appliances with motors) or capacitive, when the current leads the voltage (e.g. arc welders).

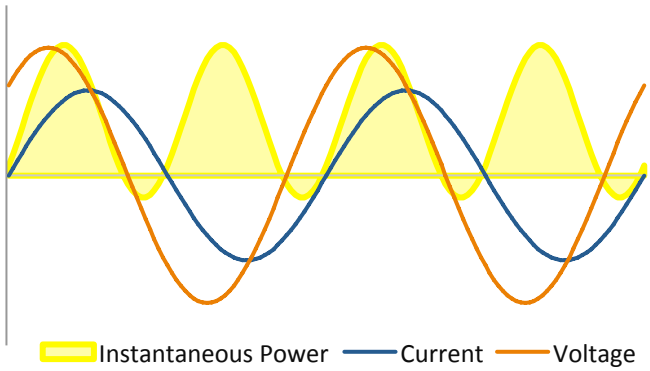

Figure 2.3: Power, current and voltage graph of a partially reactive load

Apparent, Real and Reactive Powers

The direct implication of having current and voltage out of phase, as in *Figure 2.2* and *Figure 2.3*, is that the power provided by the supply is no longer the product between current and voltage. And this happens because the current and voltage applied to the load is not equal to the current and voltage being used by the electric equipment. Instead the apparent power, which is given by the product between the voltage and current effective values over one or more full periods, is calculated:

$$S = V_{RMS} \times I_{RMS}$$

Equation 2.1: Apparent power

The root mean square (RMS) value is the effective value of a varying voltage or current. It is the equivalent steady DC (constant) value, which gives the same effect.

However apparent power is only a representation of the voltage and current being sent to the device and not the actual power being used by the load, referred to as real power (or active power), which is the power going to the load minus the power going back to the supply. The portion of power due to stored energy, which returns to the source in each cycle, is known as reactive power.

*Even though Watts = Volts * Amps, apparent power is measured in VA to differentiate it from true power.*

Real, reactive and apparent powers are very well related to each other and this relationship can be express by representing these three quantities as vectors. Real power (P) is represented as a horizontal vector and reactive power (Q) is represented as a vertical vector. The apparent

power (S) vector is the hypotenuse of a right triangle formed by connecting the real and reactive power vectors.

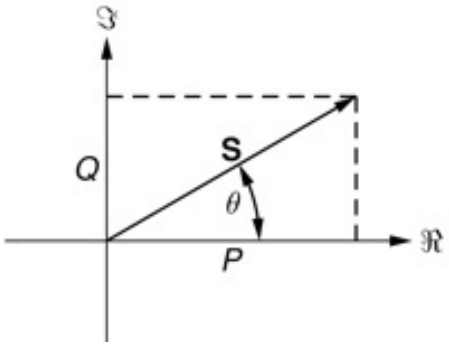

Figure 2.4: AC Power triangle

This representation is often called the power triangle. The relationship among the three quantities can then be expressed using the Pythagorean theorem:

$$S^2 = P^2 + Q^2$$

Equation 2.2: Relationship between real, reactive and apparent powers

For ideal loads, when the current and voltage are both sinusoids, real and reactive power can also be calculated from the apparent power:

$$P = S \times \cos(\phi)$$

Equation 2.3: Real Power

$$Q = S \times \sin(\phi)$$

Equation 2.4: Reactive Power

Where theta (ϕ) is the phase angle between voltage and current and $\cos(\phi)$ is the Power factor, a measure of the mains efficiency:

$$PF = \cos(\phi)$$

Equation 2.5: Power factor

If the phase angle (ϕ) is zero the real power will be the same as the apparent power, and the reactive power will be zero, with all the power being used by the load. If, on the other hand, the phase angle (ϕ) is ± 90º the real power will be zero while the reactive power will be the same has the apparent power, meaning that all the power is returned to the source.

However the power factor is also affected by the harmonic contents of the input current and voltage, which are not taken into account in *Equation 2.5*. When the currents and voltages have non-sinusoidal forms, as seen in *Figure 2.5*, the power factor is generalized to include

the effects of harmonic distortion. The ratio of real power to apparent power is called power factor and is valid for both linear (AC electric loads where the current and voltage are sinusoidal) and non-linear loads (those AC loads where the current is not proportional to the voltage):

$$PF = \frac{P}{S}$$

Equation 2.6: Power factor for both linear and non-linear loads

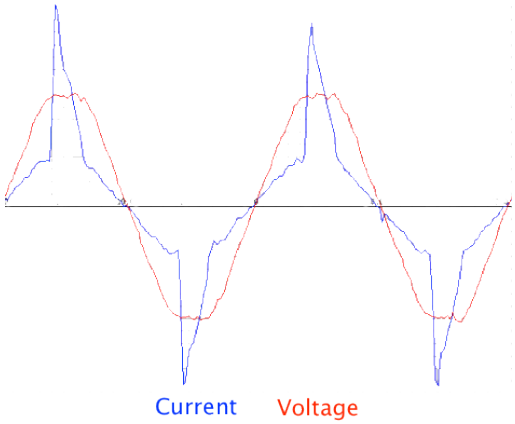

Figure 2.5: Current and voltage affected by harmonic content

2.1.2 Smart meters

Electricity exists and that is a fact, but how many of us have seen it? Electricity is there, but it is invisible and most of the times it is only noticed when suddenly none of our appliances is working any more.

As seen in the introductory chapter, humans' misunderstand energy usage, and the questions that are immediately raised are: i) how much of this invisibility is responsible for those misperceptions? ii) How can one save something that we cannot see or quantify?

Smart meters are part of the answer to the second question raised above, in the sense that they are able to quantify electricity. Broadly speaking, the term smart meter usually refers to *electric meters that record the electric energy consumption in pre-defined intervals and communicates the measured results back to the utility, also in pre-defined intervals, for monitoring and billing purposes.*

Smart meters can be defined according to two dimensions:, The time and space. Time dimension refers to the meters capability of sampling the waveforms and computing those measurements (it is important not to confuse this with the time intervals at which feedback is provided to the user). Whereas the space dimension refers to which points in the power

distribution they are measuring. The most common points are the single point, multiple point and circuit panels. These two dimensions are of big importance when deciding which kind of meter to use. For example, if we want to provide simple aggregated consumption data of a single appliance there is no need for a very high sampling rate, probably sampling at 1Hz with a single point meter is more than enough. On the other hand, if we want to study the loads of a specific home division there will the need for a circuit level meter with multiple channels and a sampling rate in the order of kHz.

In an effort to simplify the process of selecting the right meter, next some examples of commercial smart meters are presented. To a very good summary of commercial appliances please refer to [6].

Single point (plug-level) meters are probably the easiest to find and their mode of operation is very simple. Basically, the appliance is connected to the meter that in turn is connected to the outlet. Most of these meters provide basic and detailed information about the connected appliance power demands, e.g. consumption in kilowatt-hour, Voltage, Amperage, Line frequency and Power Factor. In the less expensive models feedback is usually provided through a built-in LCD display [w3], while in the most advanced cases we can find meters that are able to communicate with each other, forming a metering mesh [w4].

Multi point (whole house) meters provide measurements at the service entrance and have extra channels to track sub-panels or larger electrical loads. These are installed in the main entry feed, and the feedback to the user can be provided in several ways, e.g. portable displays, http via built-in webservers or even online services (TED [w5] used to offer integration with the Google power meter that was recently abandoned). Other examples of these whole house meters are Wattson [w6] and Current Cost [w7]. All these consist of sensors that are installed at the main entry feed and some portable display than can be placed anywhere in the house. The communication between the sensors and the displays is done using some wireless protocol of their choice.

Circuit level meters are the most powerful and expensive metering solution. These monitors offer the capability of measuring each circuit in the house, up to 12 or 16 circuits in each meter. They can even be combined to form a unique measurement mesh that will comprise each circuit in the home. Each circuit level meter provides its own way of accessing the data, but most of them fall in the following categories, serial port, Ethernet or wi-fi. These are considered, by far, the most difficult to install and the reason can easily be spotted in *Figure 2.6 (e)*.

(a) (b) (c) (d) (e)

Figure 2.6: a) Kill a Watt P4400; b) Plugwise; c) TED; d) Wattson; e) EnerSure [w8]

To increase the value of this research academic and open-source projects related to smart metering were also surveyed. The findings show several solutions that go from hacking existing meters to the installation of multi-sensor systems.

Tweet-a-watt [w9] was an attempt to build a wireless whole-house power energy monitoring system based on the cheapest versions of the kill a watt meter. *XBee* transmitters were added to each meter and a main *XBee* receiver was plugged to a central computer, responsible for the power calculations and communication with the Twitter service.

Some more advanced projects are offering web-based services to the consumer. Flukso [w10], for example, is a web-based community metering application. Consumers have to buy and install the *FluxoMeter* (a whole house monitor) that costs no more than 60 Euros. The sensed data is then sent to online repositories using the existing Internet infrastructure of the house. Once online, consumers can see their consumption patterns and compare with those from people all around the world. A similar, but yet much more advanced project is the Open Energy Monitor [w11]. This is an attempt to build an end-to-end open source platform for energy monitoring, control and analysis. Such a system integrates different kinds of sensors, e.g. temperature and electricity. All sensed data is also sent to both local and online repositories for online and offline analysis by the consumers. One of the biggest innovations / challenges of this project is the fact that they want to use actuators to respond to changes in the environment, for example, turning the lights off if daylight is detected.

It was also possible to find some interesting projects academic wise. The PowerBox [w12], developed at Cornell, is a single-point smart meter that provides all the common power metrics, and also offers the chance to remotely turn on or off the outlet being measured. Another feature is the auto shutdown that happens when a certain wattage threshold is reached. A custom C# application is used to access the power data in real time, and like in the tweet-a-watt it was also developed as an extension of the *kill a watt* smart meter. ACme [w13] is an open source hardware and software platform that enables wireless energy/power measurement and control of AC devices and was developed at Berkeley. It is an attempt to fill

the gap between inexpensive plug-level meters and expensive networked enterprise energy monitors. ACme power meters provide real, reactive, and apparent power measurements, with optional control of the attached appliance.

Although the term smart meter often refers to electricity meters, It is becoming very common to find smarts meters that measure other resources like water or natural gas.

Smart meters and energy savings

So far some of the existing solutions for smart energy metering have been presented. But how do smart meters help users save energy? The quick answer is that showing real time energy usage would help consumers reduce their energy use, hence cutting their energy bills and carbon emissions. But this soon proves to be the wrong answer. The truth is that *smart meters themselves will not necessarily mean smaller energy bills. The real potential for cutting bills lies with the consumer*. But this is not bound to happen as existing metering solutions completely fail to engage the consumer in several ways: i) they are still expensive, ii) not easy to install, specially circuit level meters, iii) feedback is either local (at the meter position in the house) or web-based, iv) they answer the wrong type of questions.

How much *vs* Why I this much

It is strongly believed that the ability to reason on top of power consumption data would be of great interest for consumers and would be of great help in the process of engaging the consumers into having a more energy efficient behavior. *« Does my new microwave spends more than my previous one? »; « Why do I spend so much electricity at night if I am sleeping? »; « How much do I spend cooking dinner? »; « Should we eat out more often? »* are just a few examples of possible questions that consumers would like to see answered by their smart meters. But this is not what smart meters do because, although they can provide several different power metrics, the level of disaggregation is not enough to answer such questions. In short, smart meters were built to (and are good at) answering the « *How much? »* side of the question above, e.g. *« Right now you are spending 2 kWatt »; « This month you will spend 40 Euros of electricity »,* but completely fail when addressing the « *Why this much? »* aspect of the question.

The ability to answer the why side of the question lies within the level information disaggregation that we can get from our metering system and how we can re-aggregate that information with other streams of data.

2.1.3 Non-Intrusive Load Monitoring

The origins of NILM go back to the early 1980s when the term was coined by George Hart [7]. Hart realized that it was possible to monitor what was happening in a home, by visually analyzing consumption values retrieved at 5-second intervals from a single point in the house. These findings lead him to formalize a solution that would allow computers to perform the same type of analysis:

"A non-intrusive appliance load monitor determines the energy consumption of individual appliances turning on and off in an electric load, based on detailed analysis of the current and voltage of the total load, as measured at the interface to the power source. (...) It is called non-intrusive to contrast it with previous techniques for gathering appliance load data, which require placing sensors on individual appliances, and hence an intrusion onto the energy consumer's property."

NILM represents two major changes in the load-monitoring paradigm, first of all intrusive sensors like microphones and video cameras, are replaced by one single non-intrusive sensor, bringing with it lower costs and ease of installation and management. The second is less obvious and is related to the fact that while in previous intrusive techniques the hardware sensors had to perform all of the energy calculations, and then send the results through a specific communication channel to a central point where, a simple software was responsible for merging the data from all the connected channels. In NILM the sensor is only responsible for sensing and sending the sensed data to the central software. Since the sensed data is representing the whole current and voltage flowing into the house, the software is responsible for separating the measured load into separate components.

Appliance models and appliance signatures

With the increase of consumer appliances, each one with its own characteristics and working modes, it becomes unfeasible and impractical to look at them in a personalized way. In an attempt to look at theses appliances as a whole, four groups of appliances (or appliance models) were defined:

I. *Permanent consumer devices.* Devices that are consuming 24 hours a day, 7 days a week, for example a smoke-alarm or a fridge.
II. *On-Off appliances.* These represent most of the household appliances such as toasters and light bulbs.

III. *Finite state machines (FSM).* Consumer devices that pass trough several definite switching states. Example: washing machines and clothes dryers.
IV. *Continuously variable consumers.* Devices that have an infinite number of states with a variable power draw. Examples of such devices are light dimmers and power tools.

Appliance signatures, on the other hand, can be defined as a set of measurable parameters of the total load that give information about the nature and operating state of an individual appliance in the load. Appliance signatures are represented in a vector format in which each element is a feature extracted from the sensor data in the vicinity of the detected power event. Signature features can be extracted from steady-state conditions or transients. Real and Reactive power (P and Q) are the most frequently used measurements for load behavior that can be extracted from steady-state conditions. Additional information can also be extracted from steady state conditions by examining the current and voltage harmonic (features in the frequency domain) generated by the appliances. On the other hand, features can also be extracted from the switching transients. In very broad terms, a transient can be seen as the changes in current that happen in between two steady-state conditions. Although this is not true for some appliances, for example a kettle goes from zero to a steady-state value without any overshot, transients can be easily noticed in other home utilities, like light bulbs and televisions. When these are turned on a sharp rise to a peak value, followed by a drop to a steady-state value is easily noticed. A transient can be defined in terms of the shape, size, duration and time constant.

Transients are rarely noticed when shutting down appliances, meaning that they will only be useful when used in conjunction with other steady-state features.

NILM workflow

In high-level terms, the NILM workflow can be defined in six consecutive steps, as presented in *Figure 2.7*.

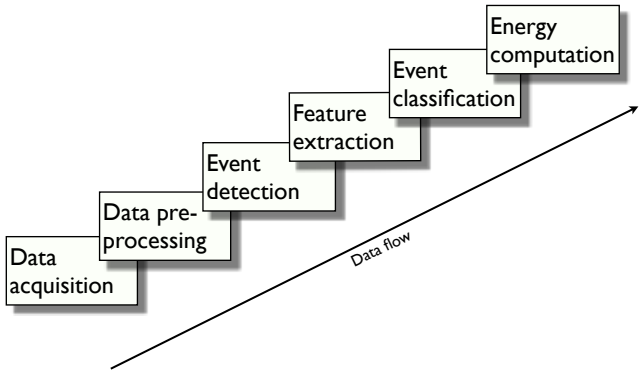

Figure 2.7: Non-Intrusive Load Monitoring workflow

1. **Data Acquisition.** The process of sampling signals that measure real world physical conditions (e.g. the magnetic field generated around a copper wire) and converting the resulting samples into digital numeric values is known as data acquisition. The NILM data acquisition module (DAQ) is responsible for measuring the characteristics of the current and voltage flowing into the house. Once the measurements are done the results must be made available to the data pre-processing module. A typical NILM DAQ consists of both a current and voltage sensor and an analog-to-digital converter.

   In this phase, sensors measure characteristics of the current and voltage signals. These measurements are taken from the main circuit breaker and then sent to the data pre-processing module.

2. **Data Pre-processing.** Only in rare occasions data is good enough to be processed as it is collected. For example, in NILM processing current and voltage separately is not enough, as these two quantities need to be combined in order to get other valuable power metrics. The same is to say that raw data needs to be pre-processed before it can be used by another processing procedure.

   At this stage, the acquired voltage and current signals are converted into traditional power metrics (active power, reactive power and power factor).

3. **Event Detection.** The building block of NILM is that a change in the power being consumed indicates a change in the state of the appliances within that house, e.g. a kettle being turned on. Hence being able to detect when the changes happen is of major importance for NILM. There are several signal processing techniques for this purpose such as filtering, differentiating and edge detection [7]. Probabilistic models are also available [6] [9].

   At this stage, changes in certain power metrics are detected and flagged as events for further processing in the proceeding modules.

4. **Feature Extraction.** In order to be able to learn to recognize the load events we need to focus on a set of generalized and critical features that can mathematically characterize them. These are extracted from the samples surrounding the event of interest. Examples of these features are the size of the total power step change, current harmonics values and transient shape or size.

   Features are extracted from the samples surrounding the detected event; together these features form an individual signature that describes the event.

5. **Event Classification.** When a load event is detected its respective signature is extracted and fed to a previously trained machine-learning algorithm that will try to find a pre-classified event that holds a similar signature. The most common approaches include supervised learning algorithms that rely on previously classified events to learn new ones and non-supervised learning approaches, for example using clustering algorithms that group events with similar power signatures.

   Previously trained machine-learning algorithms are applied to unclassified event signatures to obtain a classification.

6. **Energy Calculation.** Once we are able to keep track of all the load events as well as their associated power levels and appliances, it is possible to compute how much energy each appliance is using.

   By keeping track of all the events that occur and their associated power levels is possible to estimate how much energy each appliance is using.

2.2 Literature review

Attempts to monitor and disaggregate electricity usage in the home go back to the early 1980s. Hart et al [7] were pioneers in this field, and their approach consisted of analyzing current and voltage as it was measured from the main circuit breaker of the houses they wanted to monitor. Back then, and due to limitations in the digital sensors, they where only considering the fundamental frequency (60Hz in the US). The classification of appliances was done applying a clustering analysis algorithm, in a two-dimensional signature space of real and reactive power changes, as extracted from the steady-state power draw. Positive and negative clusters of similar magnitude were paired or matched as being the same appliance. Clusters and events with no match were paired with existing or new clusters according to a best likelihood algorithm. Finally, the pairs of clusters were associated with known power consumption levels to determine the operating schedule of individual loads. Hart ended up coining this process with the name of Non-Intrusive Load Monitoring.

An extension of [7] was proposed to extend it to commercial buildings [8]. In their work the authors apply a median filter to remove abrupt peaks in the raw signal, which are very prune to happen in commercial environments. They have also noticed that for some appliances the reactive power was relatively small, which invalidated the use the algorithms from [7] as they required non-zero reactive power. To compensate the lack of reactive power in certain appliances they considered the transient shape as an additional feature.

With the advances in sensing technology most researchers saw an opportunity to reach higher accuracy in the detection and monitoring of appliances by using « microscopic » features of the electric signal such as harmonics and waveforms.

« Macroscopic » features refer to the power changes and transient shapes that are obtained with very low sampling rates, e.g. 1Hz.

In [9] authors have used harmonics as complementary features, in addition to changes in real and reactive power. They found that using harmonics in the transient could help to distinguish loads with overlapping clusters in the $\Delta P - \Delta Q$ signature space presented in [7]. They showed that otherwise indistinguishable loads in the $\Delta P - \Delta Q$ signature space, like the turn-on and turn-off events of a personal computer or a bank of incandescent lamps, could be easily separated by looking at the third harmonic current, that is characteristic in typical personal computer power supplies.

Other harmonic approaches consist of using their spectral envelope, which is a vector of the first several coefficients of the short-time fast Fourier transform (ST-FFT) of the signal. In [10] the authors fitted a least squares procedure to each of the previously stored and labeled spectral envelops. A distance metric, Euclidian distance, was then applied to the results and the one with the shortest distance was selected as the correct classifier. Berges et al. [11] [6] have presented an end-to-end NILM solution, built with a custom data acquisition system, based on a commercial oscilloscope. In their work they applied supervised machine learning algorithms, e.g. *K*-NN and decision trees, to classify the loads, based on features that were extracted from the samples surrounding the detected events. The features used were real, reactive power and the shape of the transient profile in any of the first seven odd harmonic power coefficients, in terms of linear regression coefficients using different basis functions (polynomial, gaussian, radial and fourier).

In an attempt to improve the load disaggregation process Marchiori et al. [12] presented the circuit-level non-intrusive load monitoring. In this approach the NILM is applied individually to each circuit of the house. Authors claim that they can overcome the inability of whole-house NILM to monitor small or variable devices, because there are fewer devices in each circuit and high-powered devices (such as a stove, air conditioner or clothes dryer) each receive dedicated circuits and will not overshadow lower-power devices (such as TVs, radios or cell phone chargers).

While these approaches make use solely of power characteristics, e.g. real power and harmonic content, other researchers have come up with completely different alternatives to detect and classify changes in the power load that rely on different signal characteristics. For

example Patel et al. [13] proposed that by monitoring electric noise in a socket for transient signals they could detect most appliances that were connected to other sockets in the house. However this approach suffers from such problems as the fact that it is very computational expensive to capture and analyze transient noise, and most of all because the obtained signatures are very dependent on the household wiring which makes it impossible to create a solution that could be ported across households.

The same authors also presented ElectriSense [14], a system that focuses on sensing very high frequency (36-500 kHz) electromagnetic interference (EMI), which is constantly generated by the switch mode power supplies (SMPS) present in most modern consumer electronics, as well as fluorescent lightning. Although this is a promising approach, with a reported accuracy of 93.8%, there are some very important issues that need to be tackled, for instance the simple fact that not all modern appliances are equipped with SMPS may turn this alternative down, and also the fact that this approach (and this also applies to [13]) is merely for the detection and classification of events, as it does not seem to offer the possibility to compute the energy which is the main objective of NILM.

In [15] the authors present other approaches as well as some future directions for NILM.

In the future work section of [6] and in [16] Bergés et al. talk about using sensor fusion that will add the possibility of using other streams of data, e.g. light intensity levels, motion and sound levels, to help increase NILM performance. In [17] the authors went a step further, and presented a contactless electromagnetic field (EMF) sensor that was able to detect appliance power consumption state changes within close proximity, based on magnetic and electric field fluctuations. Its usage was recommended to aid in the training of NILM systems working as a ground truth collector. The authors admit the existence of three main limitations: i) knowing which state transitions to look at, ii) some cross-talk may exist if two or more sensors are in close proximity and iii) this approach requires additional hardware. The authors also state that such solutions will only be needed until the NILM system reaches the point where it no longer needs to be calibrated.

Given its high potential, non-intrusive approaches are being applied to domains other than electricity. For example Cohn et al. [18] developed GasSense, a single-point sensor for gas home environment. They used a single acoustic sensor placed on the home gas infrastructure, to sense gas flowing into the house. Gas usage events are detected using a sliding window step detector. After the step detector locates a step in the signal certain features are extracted from it: (1) the relative magnitude of the step change in terms of flow, (2) the slope of the change-in-flow, (3) and the rise or fall time of the unknown event. Feature vectors are created

for each detected event and then, used to build a *k*-nearest neighbor (KNN) model, that was used to automatically determine the source of the gas events.

Researchers also saw an opportunity in the field of water supply, mostly driven by the scarcity of this natural resource. In [19], HydroSense, a single-point pressure sensor that is installed in the house water infrastructure, was presented by Froehlich et al. This system allows the identification of any individual water feature activity, as well as estimates the amount of water used. The building block for this technique is that when a valve is opened or closed (be it a bathroom faucet or a mechanical valve in a dishwasher), a pressure change will occur and a pressure wave is generated in the plumbing system that can be sensed anywhere in the house.

2.3 Challenges

Despite the best efforts from researchers, the fact that no complete set of robust, widely accepted appliances features has yet been found, makes it impossible to come up with a complete NILM solution that would work for every possible appliance. This is probably the biggest challenge: selecting a set of features that will allow the quick creation of a system that offers a certain level of disaggregation that can be easily deployed in peoples homes.

Additionally, the growing number of everyday appliances and their increasing complexity makes NILM a little more complex than before, because the load composition is now far more complicated, better resembling that of a commercial building in the early days of NILM.

On the hardware side one could expect better news as a result of recent advances in sensing technology, however most of the commercially available solutions are either expensive or fail to offer the level of data granularity required by NILM. Although internally they have a sampling frequency high enough to deal with short start-up transients their external reporting rate is much slower (less or equal than 1 Hz) [23].

For this the reason the option was to go with a totally custom solution that would cost as little as possible. However a lower cost was not achieved without a price, and even though these were the two main aspects that greatly reduced the costs of this solution, the truth is that sound cards are far from being the ideal data acquisition system, and using an AC transformer as a voltage sensor may not be the ideal solution, specially as the loads get more complicated and we need to access the reactive component of the power.

Other challenges have to do with the deployment and maintenance of such a solution, as one needs to rely in aspects other that the meter itself, for example, the need for help from a technician to install the sensors in the breaker box, and the process of roling out a wide are wireless network to make sure that every single deployed meter can be easily tracked and updated.

Another expected challenge has to do with how people will accept the system, if they will be willing to cooperate in the different studies, or if the meter will end up being just another appliance around the house. User participation will also be important, relating the process of learning appliances. In fact, one of the biggest questions that remains unanswered about NILM is if people are willing to teach the system to correctly classify their appliances.

3

SINAIS Power meter

3.1 Profile **22**
3.2 Software architecture **24**
3.3 Data acquisition **26**
3.4 Power calculations **28**
3.5 Median filter **31**
3.6 Event detection **32**
3.6.1 GLR event detector 33
3.7 Feature extraction **35**
3.8 Event classification **37**

3.1 Profile

For research purposes the energy monitor needs to commit to a set of requirements: i) it has to sample both current and voltage in almost real time, ii) the data needs to be available for both offline and online analysis, iii) it has to allow different representations of the measured energy trough different kinds of feedback, iv) it should sense when humans are exposed to feedback and v) it haves to be non-intrusive.

In previous research [21] it was found that none of the existing commercial solutions offered enough flexibility. While some are relatively inexpensive, they do not offer all the data that is needed, nor allow enough granularity. For instance, almost all of the solutions sample current only and assume a constant voltage at 1 Hz frequency. Another drawback of the existing solutions is that some of them don't even allow us to gain direct access to the results either by connecting a PC or a wireless connection.

Given these requirements, from a hardware standpoint, the following need to be made available: i) a data acquisition module, capable of reading the current and voltage signals, ii) processing and storage elements capable of calculating and storing the energy and also the

events from the signals, iii) a visualization component to enable some form of visual and audio feedback, iv) additional sensors capable of detecting human activities e.g. motion, and v) a connection element to enable exporting data to an online repository [21].

One possible solution would be to buy each of these separately and integrate them in one final solution. However that could easily cost between 500 and 1000 Euros depending on the costs of integration, hence invalidating this solution. After some attempts with custom hardware the final decision was to use a mini notebook to sample current and voltage using its built-in sound card as an ADC. Additionally, the netbook provides all the above-mentioned elements in one compact package, which would never cost more than 300 Euros.

After testing several different netbooks the choice felt on the Toshiba NB300-100 (*Figure 3.9*). This selection had to do with the fact that this was the only netbook with a sound card able to read voltage up to the maximum peak amplitude sensed by the used voltage transformer (around 0.9 volts). *Figure 3.8* shows the results of the mini netbook when compared with a 13'' Macbook with 0.5 V maximum peak amplitude. It is easily noticed that the later does not perform well, cutting off all the values that exceed the maximum accepted amplitude.

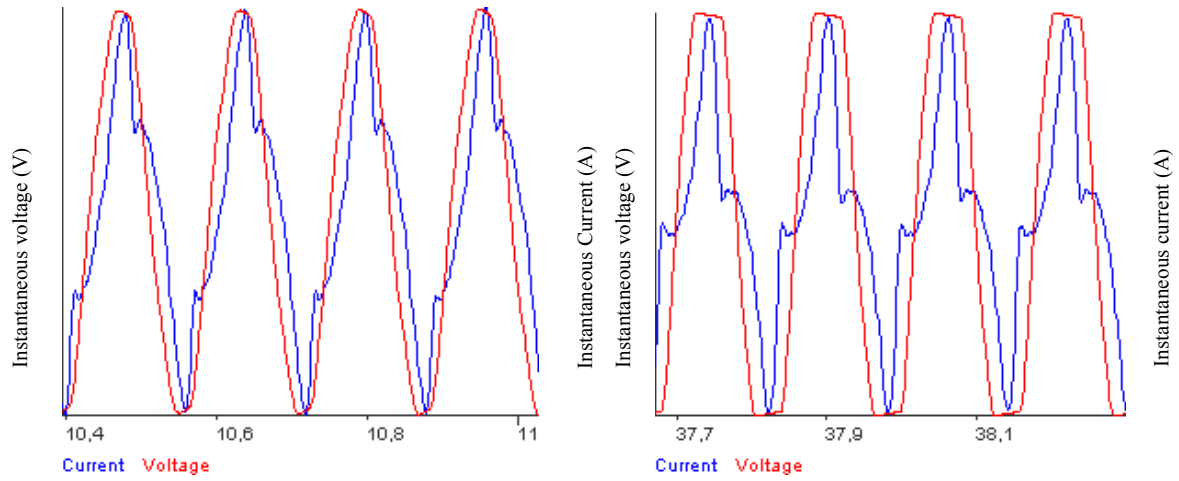

(a) *(b)*

Figure 3.8: Difference between netbook soundcards: (a) Toshiba NB300-100; (b) Macbook 13''

The notebook is equipped with a 1.66 Intel® Atom™ processor, 1 GB RAM, 250 GB hard drive, 10.1'' monitor, built-in webcam, integrated speakers, Ethernet port and Wi-Fi.

Figure 3.9: Toshiba NB300 netbook - a low cost non-intrusive home energy monitor

The laptop audio input Analog to Digital Converter (ADC) is used to sample current and voltage, the mini display and the speakers to provide the interactivity, while the Wi-Fi card enables communication over the internet and the built-in camera and microphone act as low-cost sensors for human activity.

Power measurements are taken from the main power feed to the house, covering the entire house consumption. These measurements will be used for event detection, event classification and, ultimately, information breakdown into individual appliances. In the meantime power consumption and power event data are being stored in a local database for future reference.

3.2 Software architecture

Software wise, the system mode of operation, represented in *Figure 3.11,* is based on the « pipe-and-filter » software architecture.

Pipe-and-filter is a very simple, yet powerful, architecture that consists of any number of components (filters) that transform or filter data, before passing it on via connectors (pipes) to other components. The **filter** transforms or filters the data it receives via the **pipes** with which it is connected. The **pipe** is the connector that passes data from one filter to the next. The **pump** or producer is the data source and the **sink** or consumer is the data target.

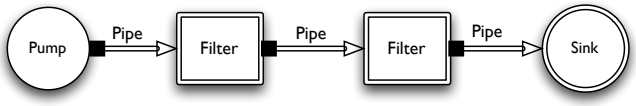

Figure 3.10: "Pipe-and-Filter" software architecture

A filter can have any number of input and output pipes and can be either active or passive. Active filters drive the data flow on the pipes and in any P&F architecture there has to be at least one of these. On the other hand a passive filter is driven by the data flow on its pipes.

Because all the filters are working *at the same time* this architecture is very suitable for systems where data transformations need to be done in (virtually) real time, which is our case.

Current and voltage are continuously sensed and sent to the *data acquisition* filter to be sampled. As current and voltage are sampled they are sent to the *power calculations* filter. This filter is responsible for doing the power calculations and drives the resulting data the splitter, which is an active filter that is responsible for sending the power samples to the filters that are connected to it. In the current implementation there are two filters and two sinks connected to the *splitter*. The *GUI* is a sink, and is responsible for plotting the power as it is being calculated (see *Figure 3.14*). The other sink is a *Request / reply socket server* that provides real time information about the system measurements in the Internet.

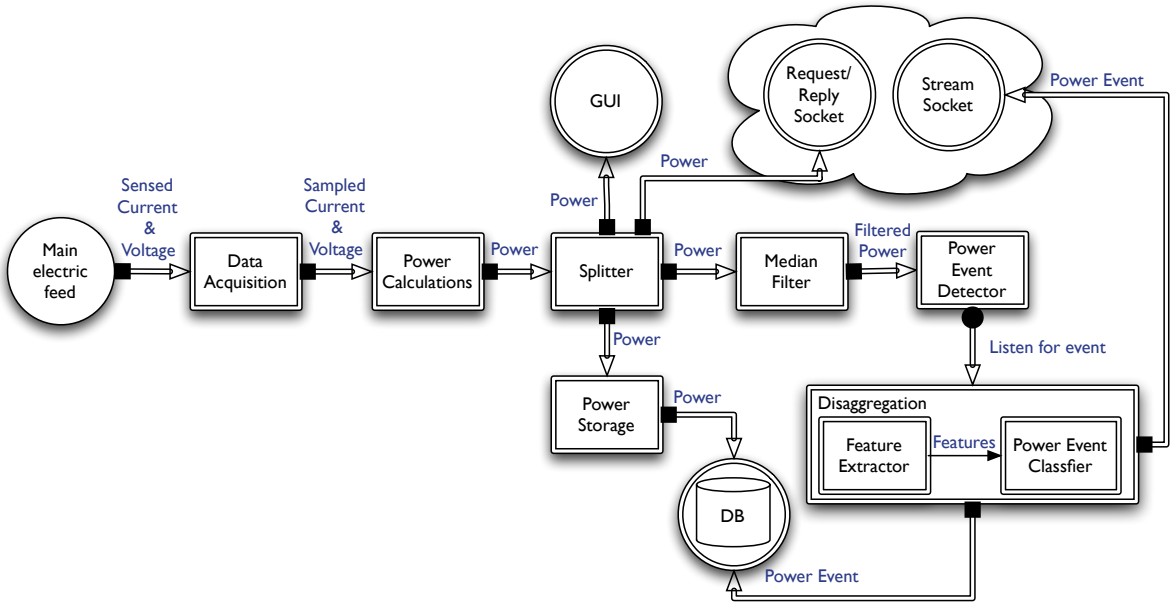

Figure 3.11: Low cost Non-Intrusive Home Energy Monitor software architecture

The *power storage* filter role is to average the power samples, based on a predefined number of samples value, and drives the resulting sample to the *database* sink. The *median filter* is used to apply a median filter to the power samples, also based on a predefined window size,

and sends them to the *power event detector* filter. The *power event detector* filter applies a detection algorithm, which will be presented in 3.6.1, to the filtered power samples, and fires a programmable event when any power event is detected.

At this stage is it important to distinguish between a power event and a programmable event. A power event refers to a change in the power load that is detected by the systems' event detector. On the other hand a programmable event consists of a "message" that is sent by the software system in response to a change of state.

The *disaggregation filter* is a composite filter that captures the events fired by the *power event detector* and it is composed by two filters that work together to disaggregate the load. The *feature extractor* is used to extract the features that will be used by the *power event classifier* to classify the power event. The event is then sent to the *database* and *stream socket* sinks. The *stream socket* sink will then stream the event (in XML format) to its clients.

Since all the filters are executed in parallel there is this sense of real time. However in such a system the filters have to be very well synchronized in order to avoid buffer overflows that may lead to loss of data and eventually system deadlocks.

3.3 Data acquisition

A Data Acquisition System (DAS or DAQ) is composed of two main components (1) sensors that convert physical parameters to electric (analog) signals and (2) an analog-to-digital converter (ADC) to convert analog sensor signals to digital values that the computer can understand.

Sampling is the process of reducing a continuous signal, like the one we obtain from the used sensors, to a discrete signal. In very broad terms the amplitude of the analog continuous wave is measured periodically at a certain rate, resulting in a set of discrete samples that represent the signal in its digital format. The accuracy of the digital approximation of the analog signal will depend on its resolution in time (the sampling rate), which in our case, despite a rule of thumb for time domain analysis suggesting about 20 points for each fundamental period [15], was set to the minimum allowed (8 kHz)) and its quantization, or resolution in amplitude (the number of bits used to represent each sample), that we set to 16 bits.

If required, and it is not the case here, some signal conditioning circuitry can be placed in between the two principal components. This component will manipulate the sensed signals before they are sent to the ADC, for example, the analog signal could be amplified before being converted to digital format.

In Portugal most residential building have single-phase electric circuits that are fed by one single Alternating Current (AC), along with a 230 V voltage, with a 50 Hz frequency. Hence, in order to measure the power being consumed we will need only two sensors installed at the main breaker circuit, as depicted in *Figure 3.12 (a):* one split-core current transformer to be placed around the cable that carries the current and one voltage transformer to be connected to one of the voltage sources.

Most current sensors work by measuring a magnetic field that is created around the copper ire that carries the current. The measured value is then transformed into an analog representation of the current.

A voltage sensor works by stepping down the input voltage until we get an output voltage that the system is able to handle.

In its essence, an ADC is a simple device that converts a continuous quantity, typically an analog voltage or current, to a discrete time digital representation proportional to the magnitude of the input signal.

These two sensors are then connected to the netbook built-in sound card using an audio splitter jack as shown in *Figure 3.12 (b)*. Custom made software is used to sample the acquired (analog) signal using the sound cards' ADC.

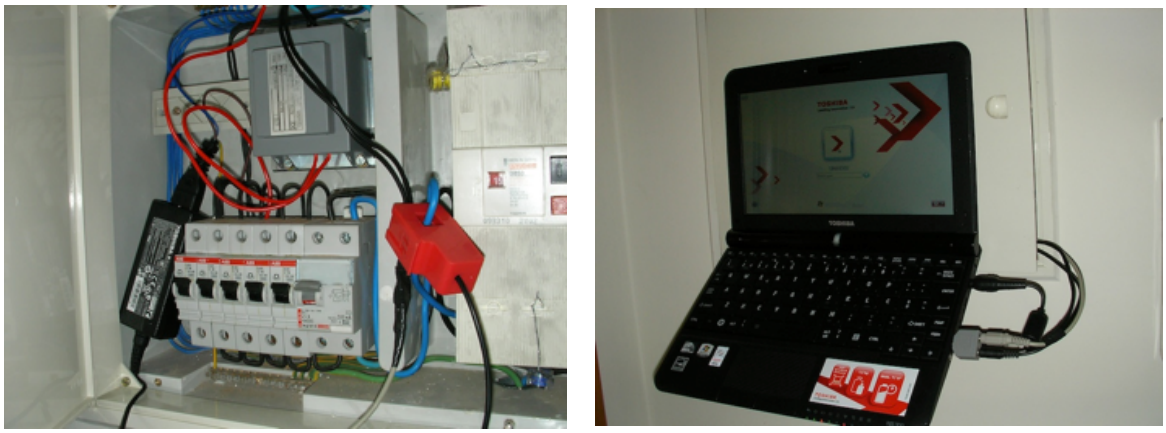

(a) *(b)*

Figure 3.12: Power meter installation: (a) sensors installed in the main circuit breaker; (b) sensors connected to the netbook built-in sound card

To be able to use the sound card as part of the data acquisition module it was necessary to make sure that both voltage and current would not mix with each other during the sampling process. Instead, each signal should be confined to one of the two available channels. This was achieved by turning the audio jack right channel into the output channel for the voltage transformer while the left channel was used for current.

With such configurations the system will be will be acquiring 16000 samples every second, 8000 for each waveform which at 50 Hz means that each current and voltage period will be characterized by 160 samples. This is shown in *Figure 3.13* for three periods.

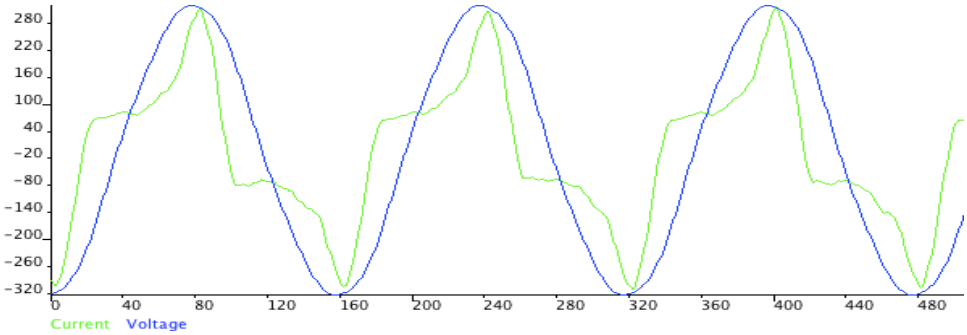

Figure 3.13: Current and voltage at 50 Hz with a sampling rate of 8 kHz

Persistence module

Although the default behavior of energy metering systems is real time processing, the possibility of storing current and voltage in files, for offline processing, was also added. This option proved to be very useful when testing detection, feature extraction and classification algorithms with different parameter selection (it will be discussed in section 4 of this thesis). The acquired signals are stored in the Audio Interchange File Format (AIFF) format [w15]. It is an uncompressed and lossless audio format, and this fact makes it perfect for storing current and voltage, as none of the initial signal characteristics will be lost. However, such a format uses a lot of disk space, about 10 MB for one minute of stereo data, at a sample rate of 44.1 kHz and a sample size of 16 bits, or over 14 GB if storing a full day of current and voltage (only feasible because the mini netbook being used has 250 GB hard drive, but not available in any commercial solution, at least a low cost one).

The fact that we are using a sound card to sample our analog signals is very important to keep the project at a low cost; still, using the sound card is only possible because in Portugal we have a single-phase system. If we had a dual-phase for example, we would have to find a different solution. Using one sound card to sample the voltage only while the other would sample current in both phases could be a possible solution

3.4 Power calculations

AC systems contain a lot of variables that can be measured, each one with its uses, for example, if one is doing an electric audit, current, voltage, harmonics, power factor and line frequency would be of very big importance. On the other hand, for household metering

applications real power is likely to be the most important value, as it tells how much power is actually being billed.

In this section it is shown hot to calculate, in the time domain, apparent power, real power, reactive power, power factor, RMS voltage and RMS current from instantaneous voltage and current measurements of a single-phase AC electric circuit.

Real Power

Real power (also known as active power) is defined as the power used by a device to produce useful work. Mathematically it is the definite integral of the voltage $v(t)$ multiplied the current $i(t)$, as follows:

$$P = \frac{1}{T}\int v(t) \times i(t)\, dt \equiv V_{RMS} \times I_{RMS} \times \cos(\phi)$$

Equation 3.7: Real Power definition

The discrete time equivalent is:

$$P \equiv \frac{1}{N}\sum_{n=0}^{N-1} v(n) \times i(n)$$

Equation 3.8: Real Power definition in discrete time

Where $v(n)$ is the sampled instance of $v(t)$, $i(n)$ is the sampled instance of $i(t)$ and N is the number of samples. Real power is then calculated simply by averaging N voltage-current products over one or more full periods.

RMS Voltage and Current

An RMS value is defined as the square root of the mean value of the square of the instantaneous value of a periodically varying quantity, averaged over one, or more, complete cycles. The discrete time equation for calculating voltage RMS is as follows:

$$V_{RMS} = \sqrt{\frac{\sum_{n=0}^{N-1} v^2(n)}{N}}$$

Equation 3.9: Voltage root-mean-Square definition

Current RMS is calculated using the same equation, only replacing voltage samples, $v(n)$, with current samples, $i(n)$.

Apparent Power and Power factor

Apparent power and power factor are calculated according to *Equation 2.1* and *Equation 2.6* respectively.

Reactive Power

Reactive power can be calculated using the power triangle that was seen in *Equation 2.2*. However noticeable errors will appear in the presence of harmonics, as the power triangle assumes that only ideal loads are present. A better approach is to introduce a time delay in one of the waveforms by shifting it by a quarter cycle [22]. The reactive power is the result of integrating the product of the normal and the shifted waveforms as follows:

$$Q = \frac{1}{T}\int v\left(t + \frac{T}{4}\right) \times i(t)\, dt \equiv V_{RMS} \times I_{RMS} \times \sin(\phi)$$

Equation 3.10: Reactive Power definition (time delay method)

And the discrete time equivalent is:

$$Q \equiv \frac{1}{N}\sum_{n=0}^{N-1} v\left(n + \frac{N}{4}\right) \times i(n)$$

Equation 3.11: Reactive Power definition (time delay method) in discrete time

The result of computing the values in this way can be observed in *Figure 3.14* where we show a graph with the real time values for current (0,33 A), voltage (231 V), apparent power (~ 100 VA), real power (~100 W), reactive power (~ 0 VAR) and power factor (0.99) at the fundamental frequency (50 Hz) during roughly 1 minute.

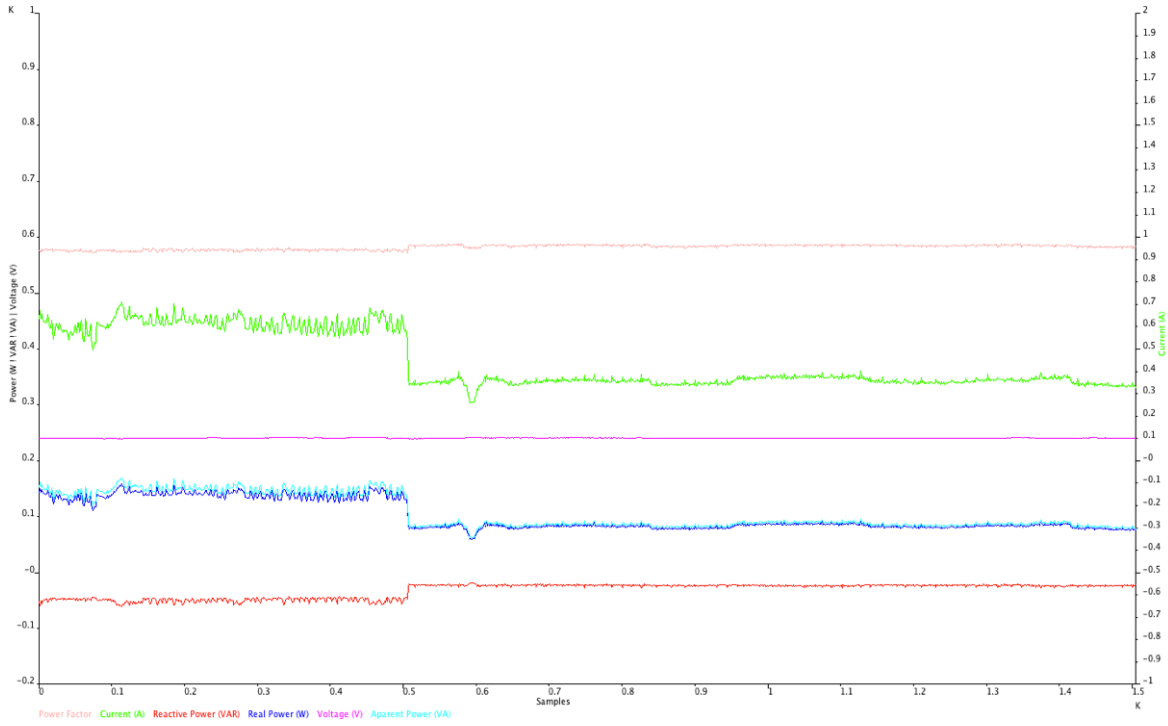

Figure 3.14: Real time Voltage RMS, Current RMS and Real Power during 1 minute

3.5 Median filter

As previously seen, in [8] authors applied a median filter to the signal before performing event detection. The median filter is a nonlinear digital filtering technique, often used to remove noise while preserving edges. And it is this edge preservation property that makes this filter very appealing to this solution. It works by sliding a symmetrically placed window across the data, point by point, and outputting the median of the data in the window at the current time.

Given a discrete set of data points, de median filter is defined as in *Equation 3.12*:

$$y[n] = median\,(x[n-k], \dots, x[n], \dots, x[n+k])$$

Equation 3.12: Median filter definition

Where $L = 2k + 1$ is the filter window length, that should be made large to increase the smoothing but at the same time should be small enough to preserve the signals resolution.

The superior performance of the median filter when applied to this kind of problem arises from the fact that median filters are geometrically oriented, shaping the local form of the signal. Also, the fact that the median filter always selects an actual data point from the input

signal (as opposed to, for example, the average filter) will help preserve important characteristics of the original signal.

In short, median filters will « pass » signals that are locally monotonic (edges) and filter those that are not (spikes).

3.6 Event detection

As appliances are turned on or off power consumption will increase or decrease and detecting when and why these changes occur is the whole point of NILM. Event detection is about finding when the change happens, while feature extraction and event classification is about telling what lead to that change. For example, in *Figure 3.14* there is one power changes in the real power draw and the event detector must be able to trigger an event alarm when it passes through it.

There are several signal processing techniques for this purpose, such as filtering, differentiating and edge detection [7]. For example, a naïve algorithm for edge detection would be one that scans the metric values (the most commonly used metric is the real power) looking for step changes, in adjacent samples that are above a predefined threshold [6]. This is considered naïve because any spike would be considered an event, even if it was noise or if it was part of the normal functioning of an appliance.

A fairly good approach to the problem of event detection is to use probabilistic methods that take into account not only adjacent but also surrounding power measurements. One of the most well known probabilistic methods for event detection is the Generalized Likelihood Ratio (GLR) that was used in [20] and adapted by [6].

In short, a likelihood ratio test is a statistical test used to compare the fit of two models. The test is based on the likelihood ratio, which expresses how many times more likely the data are under one model than the other [w14].

Regarding event detection a GLR detection algorithm will calculate a detection statistics from the natural log of a ratio of two probability distributions of power values before and after a potential change in mean:

$$S_j^k(u_1) = \sum_{i=j}^{k} \ln \frac{P_{u1}(y_i)}{P_{u_0}(y_i)}$$

Equation 3.13: GLR detection statistics

Where y_i is the sampled variable at time i, u_0 and u_1 are the mean values of the sampled sequence before and after the event, respectively. $P_u(y_i)$ is the probability density function of the sampled sequence y_i $(i = j, \dots, k)$ about the mean value u.

The behavior of the detection statistics will define the presence or absence of a step change. In short, if some appliance is turned on or off the likelihood ratio between the pre-event and the after events samples will increase meaning that there is a power change in between the two time instants represented by the two subsets of samples.

In [9] the authors present an in-depth explanation on how to create a GLR event detector as well as a very clear explanation of the GLR and why it works as a change detector.

3.6.1 GLR event detector

In its essence the GLR works with one sliding window that will be refered to as *detection window,* that will calculate the likelihood for each sample and select the index of those that are above the pre-defined threshold. Still, and given the nature of our problem, implementing the detector as it is presented in [20] holds some drawbacks:

1. *Setting the standard deviation default value and the detection statistics threshold*. Tests have shown that the standard deviation tends to increase when more appliances are in operation and the total power increases, and the latter depends very much on the consumption level of the house. As so, setting a default standard deviation would require training this parameter individually in every installation.
2. *When in presence of an event all the adjacent samples would be indexed as an event*. Although this can be solved by selecting the index with higher likelihood from the set of adjacent samples, this is highly dependent on the size of the window being used.

To minimize these issues it was opted to follow Mario Bérges [6] approach, that was built on top of the GLR but with some improvements consisting of: i) continuously updating the standard deviation and ii) implementing a voting scheme on top of the output of the maximization of the detection statistic. On top of this one also added the possibility to reset the detector every time an event is detected.

Continuously updating the standard deviation brings the advantage of having one less parameter to train. As for the voting scheme, it will make the algorithm much less sensible to noise because only indexes with a certain amount of votes will be considered events, and also there is no need to train an additional threshold value for the detection statistics. Finally,

resetting the detector when an event is detected is important to avoid having multiple detections for the same event, which is prune to happen with the standard implementation of the GLR detector, because it will be detecting events until the detection window completely passes the on / off event.

To implement this algorithm two sliding windows are needed, one for the likelihood calculation and another for the detection procedure. The likelihood calculation window will be sliding inside the detection window adding at each iteration one vote to the sample with the highest likelihood according to *Equation 3.14*. Once the maximum number of iterations for a detection window has been reached, each sample with more votes than a predefined minimum votes threshold (V_{min}) is considered to be an event.

$$g_k = max_{l \le j \le k} \sum_{j=j}^{k} \left[\frac{\widehat{V_j} \times (y_1 - u_0)}{\sigma^2} - \frac{\widehat{V_j}^2}{2 \times \sigma^2} \right]$$

Equation 3.14: Maximum likelihood calculation for the detection window l-k

Where $\widehat{V_j}$ is the value of $u_1 - u_0$ at which g_k reaches its maximum and V_m is the magnitude of the minimum change of mean of interest:

$$\widehat{V_j} = \begin{cases} V_m & if \ \frac{1}{k-j+1} \times \sum_{i=j}^{k} |y_i - u_0| < V_m \\ \frac{1}{k-j+1} \times \sum_{i=j}^{k} |y_i - u_0| & if \ \frac{1}{k-j+1} \times \sum_{i=j}^{k} |y_i - u_0| \ge V_m \end{cases}$$

Equation 3.15: Likelihood minimum change of mean of interest

Both windows have fixed lengths. The likelihood window $[l, k]$ can been seen as having two windows, $[l, j[$ and $[j, k]$, pre-event and post-event respectively. The former is used to achieve a stable mean as the reference for coming events, while the later is intended to be very sensitive to events but yet robust to disturbances. The size of the detection window is dependable on both the size of the likelihood window and the maximum number of votes that each sample can get.

To reset the detector after an on/off has been detected the approach in [20] was replicated. This consists of purging the likelihood window $[l, k]$ in the interval $[l, \ j[$ everytime an event is detected and refill it with values that came right after the interval $[j, \ k]$. This way, when the data is purged the likelihood in *Equation 3.14* will decrease ($u_1 - u_0$ will be closer to zero) and the adjacent samples will not be tagged as events the next time the detection window slides through them.

3.7 Feature extraction

Every time an event is detected a fixed number of samples surrounding it needs to be selected in order to extract features from it.*(b)*

Figure 3.15 shows the plot of the selected samples for the real power of a microwave. In *(a)* the turn-on event, while the turned-off is shown in *(b)*:

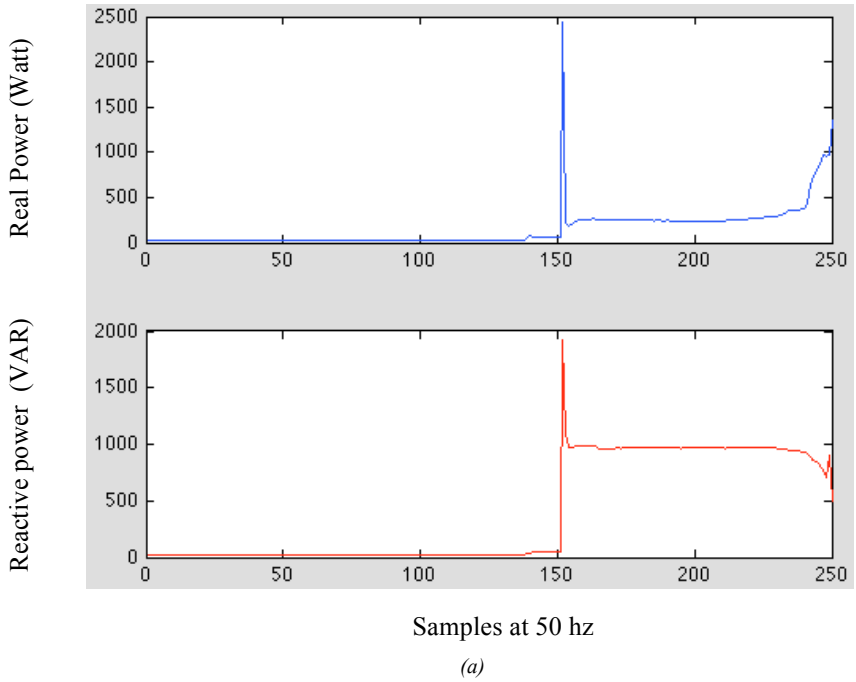

(a)

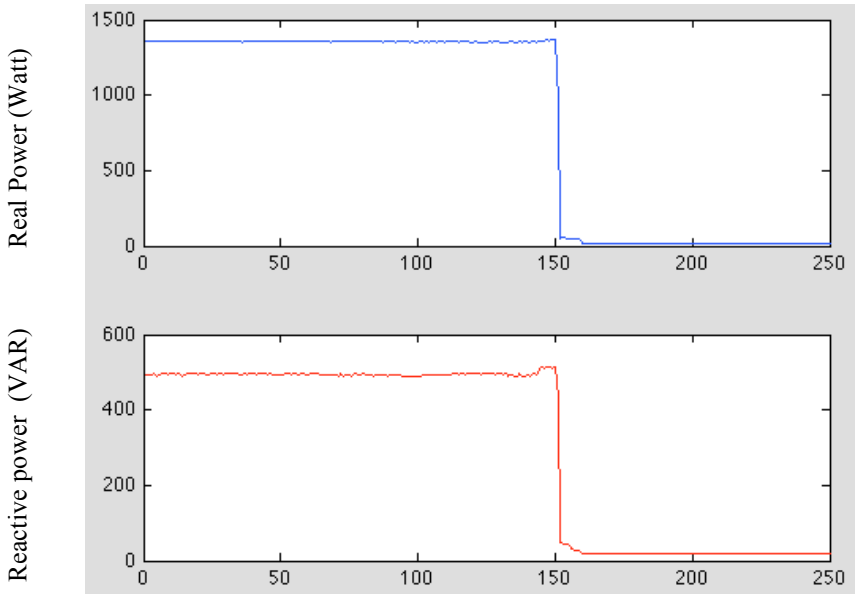

(b)

Figure 3.15: Microwave-oven transients at 50 Hz: (a) turn-on; (b) turn-off

In chapter 2 (state of the art), section 2.1.3 it was seen that in order to learn the origin of the loads one needs to mathematically characterize its changes. Such a characterization is known as a power event signature, and is composed by a set of features that are extracted from the samples surrounding the power event. In this thesis three very straightforward sets of features will be used: (1) the mean change in power generated by the event, (2) the event transient and (3) the coefficients of a polynomial fit to the power transient. But first, and before any features can be extracted, the data needs to be normalized.

Normalization

It is very important that all the parameters have the same scale before the features are extracted, especially when looking at the power transient and polynomial coefficients. Normalization is one of the well-known methods for rescaling data. In this case the data will be scaled in a range of [0, 1].

Given a discrete set of data points, a normalized data point is defined as in *Equation 3.16*:

$$y[n]_{new} = \frac{y[n] - y_{min}}{y_{max} - y_{min}}$$

Equation 3.16: Normalization definition

Where y_{max} and y_{min} are, repectively, the biggest and the smallest value among all the values that make surround the event.

Mean power change

The mean power change refers to the amount of change that just happened.

Given a set of samples that surround a power event, and the index *j* where the event happened, the average step change is given by *Equation 3.17*:

$$mean\ power\ change = mean(y[j], \dots, y[j+n]) - mean(y[j-n], \dots, y[j-1])$$

Equation 3.17: Average power change definition

Where n is the number of samples we want to average before and after the index *j*.

This feature alone should be enough to correctly classify between appliances with very distinct power consumption levels, however by itself it is not enough to classify between appliances that consume roughly the same amount.

Polynomial coefficients

Visually it can be easy to identify some appliances by looking at their start-up transient shape. Computationally a vector containing all the data points that make the transient represents the transient shape. However, and because of the big number of values that it contains it is computationaly dificult to use this as a single feature. The solution, here presented, is to perform a polynomial fit to the transient vector and use the resulting coefficients as features.

The idea behind curve fitting is that one can construct a curve, or mathematical function, that will be the best fit for a series of data points. Such a function, with degree *n*, is defined as in the following equation:

$$y = a_0 + a_1 x + a_2 x^2 + \cdots + a_n x^n$$

Equation 3.18: N^{th} degree polynomial equation

Although a « perfect » fit, when all the data points are matched, can be achieved by setting the degree of the polynomial to the number of data points minus one, this is way beyond the current needs. Instead it is believed that similar appliances will generate polynomial fits with close coefficients, and as so, it is intended to see if it is possible to use them as parameters in the classification algorithms.

Both these features can be calculated for real, reactive and harmonic powers. For instance, a step change in the reactive power can inform about the presence of a non-linear load, for example, a swimming pool motor.

3.8 Event classification

Once the power event signatures are extracted it is time to learn what that event is, or in other words, it is time to find out what appliance lead to such event. In this thesis a supervised learning method will be used for that effect.

In a supervised learning algorithm the first thing to do is to train the algorithm, which is done by presenting it with a set of previously classified data, also known as the training data set. In such algorithms the training data consists of a set of instances where the predictor features as well as the class label are already known. The learning algorithm analyzes the training data and produces a classifier that will then be used to assign class labels to future instances where the values of the predictor features are known but not the value of the class label.

K-Nearest Neighbor

The *K*-Nearest Neighbor (*K*-NN) is a very simple supervised learning algorithm and is based on the principle that the instances within a dataset will generally exist in close proximity to other instances with similar properties. Hence, if the instances have a defined class label the class of an unclassified instance can be found by observing the class labels of its *k*-nearest neighbors. In *K*-NN an instance is classified by selecting the class that is the more common among its *K* nearest neighbors.

The training phase of this algorithm consists of storing instances where the predictor features and class label are already known. In the classification phase instances without known class labels are classified by assuming the label that is most frequent among the *k*-nearest neighbors in the training dataset. The nearest neighbors are those instances of the training set that have the smallest distance (overall distance of the existing parameters) to the instance being classified. Usually the Euclidian distance is used when working with continuous variables, which is the case.

Although very simple, this « majority voting » classification scheme presents serious drawbacks when the training data is unbalanced, i.e. if one or more class labels haves more instances than the others. Also, the accuracy can be greatly degraded when in presence of noisy or irrelevant features or if the feature scales are not consistent among each other.

4
Validation

4.1 Experimental setup **39**
4.2 Power calculations **42**
4.2.1 Results ... 43
4.2.2 Discussion ... 44
4.3 Event detection **45**
4.3.1 Results ... 45
4.3.2 Discussion ... 47
4.4 Load disaggregation **48**
4.4.1 Feature extraction ... 48
4.4.2 Event classification ... 49

4.1 Experimental setup

There are two experimental setups, one to validate the power calculations and another to validate the event detection and load disaggregation modules.

The first experience was made with a commercial power meter that was installed by the local electric company.

The other experimental setup consists of a local NILM installation, as depicted in *Figure 4.16*, with a previously selected set of appliances presented in *Table 4.1*.

(a) (b)

Figure 4.16: Local NILM installation: (a) monitor; (b) appliances being monitored

Appliance	Max Power (W)
Microwave	1200
Kettle	2000
CFL (Compact Fluorescent Lamp)	10
LCD Monitor	130
Fan	50
Toaster	900
Hand blender	250
Hand mixer	250

Table 4.1 : Appliances used for validation

With these appliances two consumption simulations were made shown in *Table 4.2* and *Table 4.3*, and stored them in local audio files. These simulations were then used to test the outcomes of the system modules for the same data, but under different conditions.

Id	Appliance	Transition
1	LCD Monitor	Off – On
2	CFL	Off – On
3	Fan	On – Off
4	Fan	Off – On
5	LCD Monitor	On – Off
6	CFL	On – Off

Table 4.2 : Simulation 1 appliances and transitions (100 seconds)

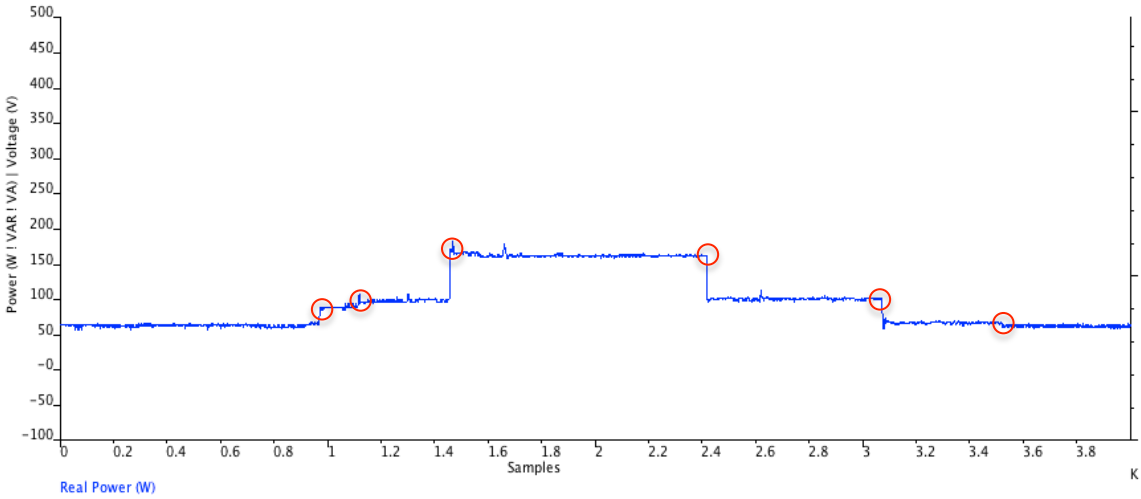

Figure 4.17: Simulation 1 power draw and events

Id	Appliance	Transition
1	Microwave	Off – On
2	CFL	Off – On
3	CFL	On – Off
4	Microwave	On – Off
5	Kettle	Off – On
6	LCD Monitor	Off – On
7	Fan	Off – On
8	LCD Monitor	On – Off
9	Kettle	On – Off
10	Fan	On – Off

Table 4.3 : Simulation 2 appliances and transitions (200 seconds)

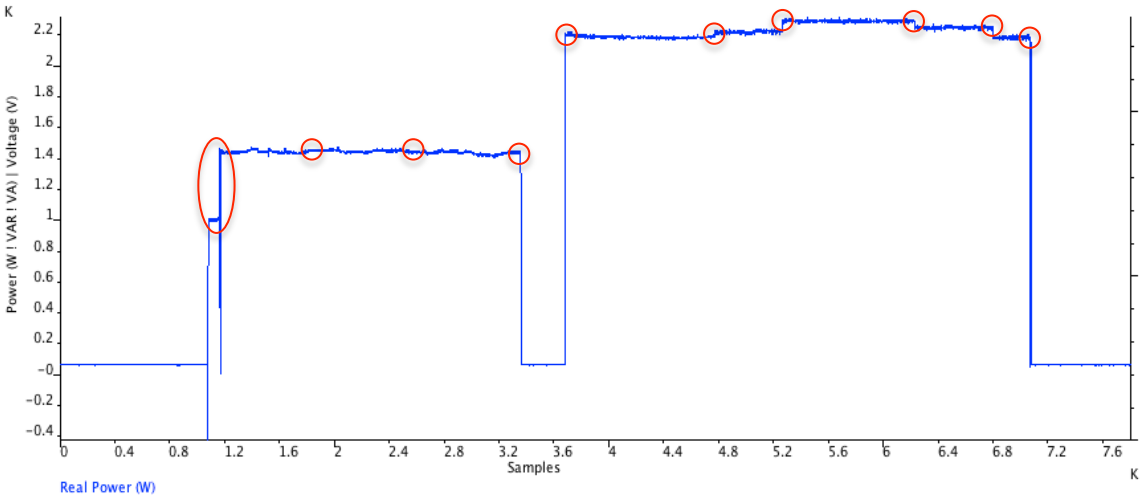

Figure 4.18: Simulation 2 power draw and events

The simulations follow a certain working logic. Simulation 1 is the most simple and is used for tests under low consumption conditions. The objective of simulation 2 is to test the cases where low consumption appliances will be working at the same time as the ones that consume more.

Also, to get some ground truth to the classification stage, another eight simulations were created. One for each appliance in the testing suite. *Table 4.4* summarizes these eight simulations.

Appliance	On events	Off events
CFL	10	10
Fan	14	14
Kettle	9	9
Microwave	11	11
LCD Monitor	7	7
Toaster	11	11
Hand blender	11	11
Hand mixer	11	11

Table 4.4 : Single appliance simulations event count

4.2 Power calculations

To test the data acquisition and power calculation algorithms both meters were installed in the main breaker box for 24 hours, averaging RMS current and voltage , real and reactive powers, among others, every 30 seconds. The collected values in both systems were then exported to a common format and the results are show in the following section.

4.2.1 Results

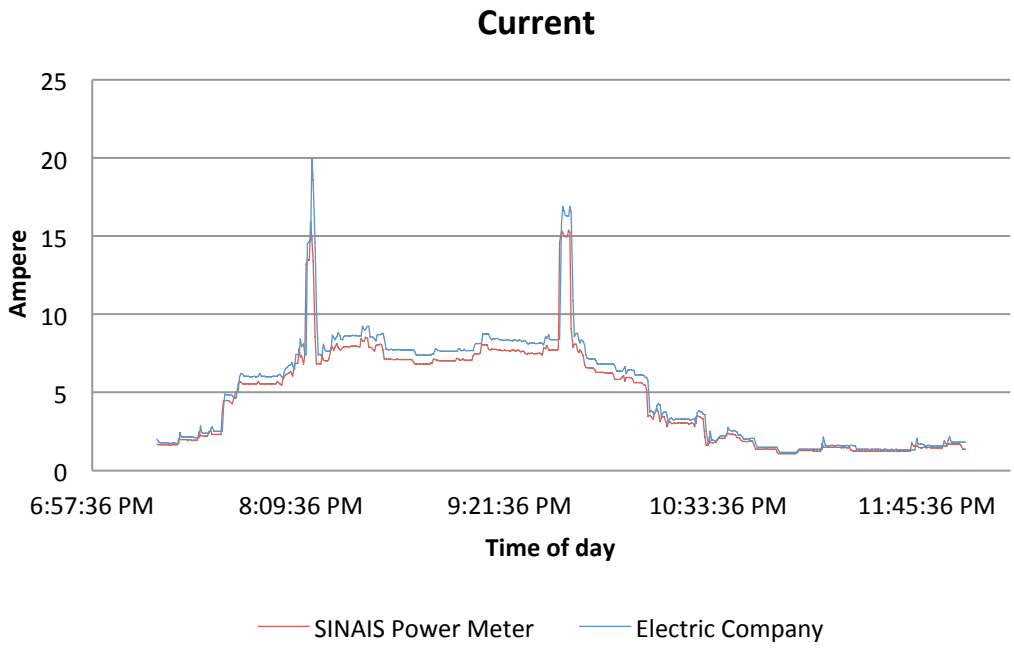

Figure 4.19: Current validation chart

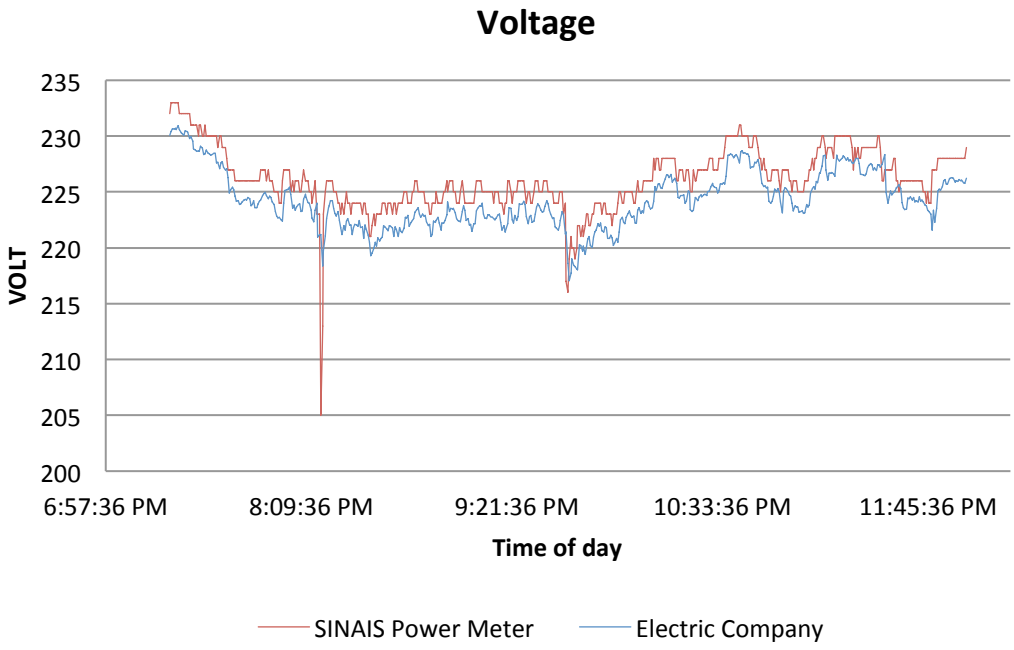

Figure 4.20: Voltage validation chart

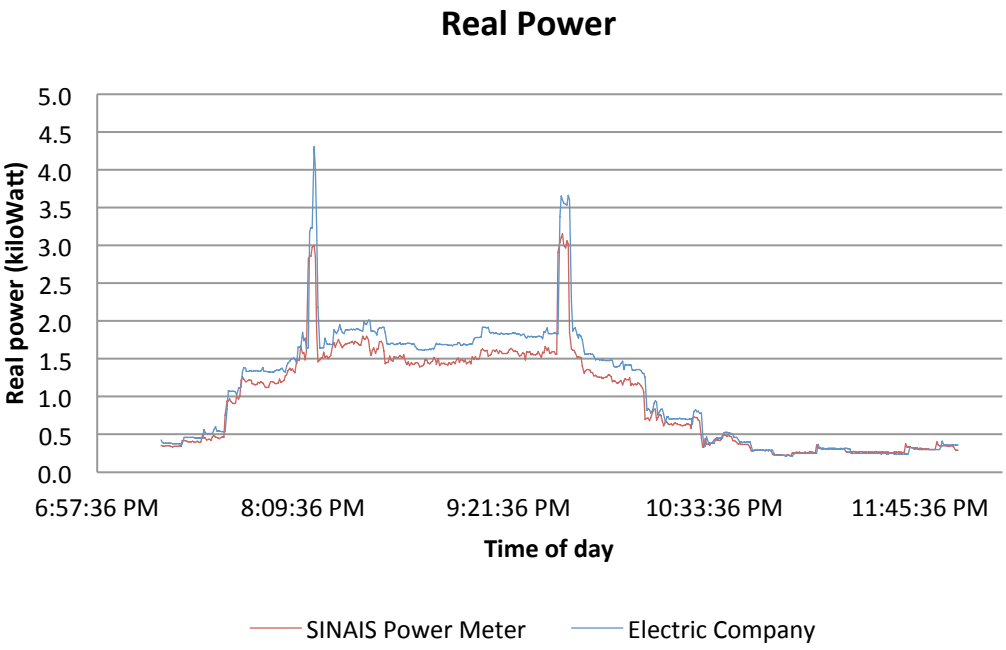

Figure 4.21: Real power validation chart

Variable	SINAIS	E.C.	% Error
Current (A)	4.68	5.10	8.24
Voltage (V)	226.17	224.27	0.84
Real Power (W)	974.68	1103.32	11.65

Table 4.5: Consumption results comparison

Assuming that the values obtained with the electric company meter are the exact values, the percentage error was calculated. This will give the difference between the approximate and exact values as a percentage of the exact value:

$$PE = \frac{|\hat{Y} - Y|}{|Y|} \times 100$$

Equation 4.19: Percentage error

Where $\hat{Y}$ is the approximate value and Y is the exact value.

4.2.2 Discussion

Current and voltage charts, *Figure 4.19* and *Figure 4.20* respectively, show how accurate the sensors are, as well as the algorithms being used to perform current and voltage root mean square calculations. And looking at the charts one can say that the results are very good, although there is a percentage error of 8.24 % in the current that is due to a calibration error.

Yet, this small error proves that the system calibration (selecting a constant value to be multiplied to the current sensor output in order to get its real value) is very important, because even a small error can produce big differences as the values increase. And the effect of a wrong current sensor calibration can be seen in *Figure 4.21*, that shows the real power for the same period of the current and voltage shown before. Again ne can see that both meter measurements follow the same trend, however the results from the custom meter are a little bit lower resulting in a 11.65 % percentage error. Still, not all of these 11.65% are due to the miss calibration of the current sensor as some of it may obviously result from the way that current and voltage are being averaged.

Unfortunately, at the time of this writing, it was not possible to compare the reactive power calculation between the two devices.

4.3 Event detection

To test the GLR event detector, the algorithm was applied, under different conditions, to the two consumption scenarios presented in *Table 4.2* and *Table 4.3*. The conditions being changed are the sizes of the pre-event and post-event windows, the minimum step change and the size of the median filter window. Because very low consumption appliances are present, the minimum power change of interest was set to 15 Watts.

The windows sizes are integer values referring to the amount of power samples. For example, at 50 Hz there are 50 samples per second, hence 150 samples represents 3 seconds. For the median filter window two sizes are used, 1 (no median filter is applied) and 11 (a value that was reached after performing cross-folding validation in other data-sets).

4.3.1 Results

In the first test the pre-event window size was set to 75 samples (1.5 seconds) and the post-event window size to 50 samples (1 second).

Simulation data 1:

Id	Appliance	Transition	Detected	
			Median: 1	Median: 11
1	LCD Monitor	Off - On	Yes	Yes
2	CFL	Off - On	Yes	Yes
3	Fan	On - Off	Yes	Yes
4	Fan	Off - On	Yes	Yes
5	LCD Monitor	On - Off	Yes	Yes
6	CFL	On - Off	No	No
False Positive			**1**	**0**

Table 4.6: Event detection results for simulation data 1

Simulation data 2:

Id	Appliance	Transition	Detected	
			Median: 1	Median: 11
1	Microwave	Off - On	Yes	Yes
2	CFL	Off - On	Yes	Yes
3	CFL	On - Off	No	No
4	Microwave	On – Off	Yes	Yes
5	Kettle	Off - On	Yes	Yes
6	LCD Monitor	Off - On	Yes	Yes
7	Fan	Off - On	Yes	Yes
8	LCD Monitor	On - Off	Yes	Yes
9	Fan	On - Off	Yes	Yes
10	Kettle	On - Off	Yes	Yes
False Positive			**24**	**22**

Table 4.7: Event detection results for simulation data 2

For the second test the pre-event window size was set to 150 samples (3 seconds) and the post-event window size to 100 samples (2 seconds).

Simulation data 1:

Id	Appliance	Transition	Detected	
			Median: 1	**Median: 11**
1	LCD Monitor	Off - On	Yes	Yes
2	CFL	Off - On	Yes	Yes
3	Fan	On - Off	Yes	Yes
4	Fan	Off - On	Yes	Yes
5	LCD Monitor	On - Off	Yes	Yes
6	CFL	On - Off	No	No
False Positive			**1**	**0**

Table 4.8: Event detection results for simulation data 1 (test 2)

Simulation data 2:

Id	Appliance	Transition	Detected	
			Median: 1	**Median: 11**
1	Microwave	Off - On	Yes	Yes
2	CFL	Off - On	Yes	Yes
3	CFL	On - Off	No	No
4	Microwave	On – Off	Yes	Yes
5	Kettle	Off - On	Yes	Yes
6	LCD Monitor	Off - On	Yes	Yes
7	Fan	Off - On	Yes	Yes
8	LCD Monitor	On - Off	Yes	Yes
9	Fan	On - Off	Yes	Yes
10	Kettle	On - Off	Yes	Yes
False Positive			**22**	**14**

Table 4.9: Event detection results for simulation data 2 (test 2)

4.3.2 Discussion

In **Simulation data 1** the accuracy is close to 100% with a very low number of false positives. This very low false positive rate is mostly due to the very low level of noise (notice that after the median filter is used the number of false positives is zero) that is present in the signal as well as the smoothness of the working mode of the appliances used. It is also possible to see that in the 4 cases the event being missed is the CFL being turned off, however if a smaller value for **Min-V** is chosen (e.g. 12 W) the system will detect all six events, even tough it will detect a few more false positives (3).

The simulation data 2 was created to see how the noise and working mode of higher consumption appliances would affect the detection of smaller appliances. And the results were as expected, one can still detect all the appliances (except the CFL being turned off), however, the number of false positives greatly increases, for example, in the second test out of 22 false positives 13 happened while the microwave was on.

A high false positive rate yields two problems, one is the unnecessary processing that will be used trying to classify something that is not an event, and second, if one of the noise events is misclassified we may end up messing up our learning signatures database, as well as the computation of the total consumption by appliance.

As one can see there is a *trade-off* in the problem of event detection: a low minimum step change will aloe the detection all the generated power events, however such a low value will dramatically increase the number of false positives. As so, it is believed that the best possible event detector will be the one that maximizes the correctly detected events while minimizing the false positives.

4.4 Load disaggregation

The load disaggregation test is done in offline mode. But first some ground truth about the appliances used in the test had to be collected. To do this individual consumption and events for each appliance were stored in audio files. The files were then loaded by the meter to do the power calculations and event detection. As told in chapter 3 (SINAIS power meter), section 3.2 the system is capable of streaming the events to a socket server, which was used to classify every event in the simulation. One the classification was done the events were loaded from the database that is being used to store all the power events after being serialized.

The following setup for the event detector was used: Pre-event window: 3 seconds, post-event window: 2 seconds, median filter: 11 samples.

4.4.1 Feature extraction

Before extracting the features one must define the samples where they will be extracted. For each feature 2 different sample sizes were defined: 100 for the medium step change (50 after the edge and 50 before it) and 75 for extracting the transient (25 before the edge and 50 after).

Figure 4.22 highlights the data samples from where the features will be extracted (in this case from a microwave) and *Table 4.10* shows the features that are extracted from this power event.

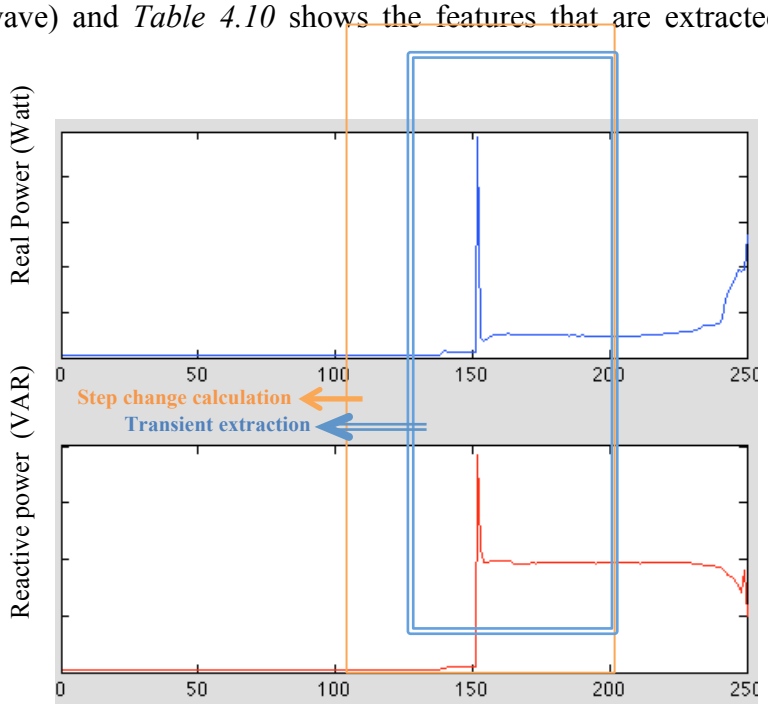

Figure 4.22: Representation of the windows from where the features will be extracted

Metric	Mean step change	Polynomial coefficients
Real Power	258.64	$y = \mathbf{-0.0074} + \mathbf{0.0194}x - \mathbf{0.0074}x^2 + \mathbf{0.00075}x^3$
Reactive Power	948.51	$y = \mathbf{-0.0288} + \mathbf{0.0849}x - \mathbf{0.0466}x^2 - \mathbf{0.0048}x^3$

Table 4.10 : Features extracted from one of the microwave events

4.4.2 Event classification

To test the classification algorithm the first step was to collect and classify 10 ON and 10 OFF power events from the appliances in *Table 4.1*, except for the CFL which was excluded from this due the difficulties in detecting its ON and OFF events. Apart from these 20 signatures, 4 others where extracted from them: 1) averaging all the points, 2) selection the median among all the points, which are referred to as “Jokers”.

In total there are 168 classified signatures, 24 for every appliance (12 ON and 12 OFF), and all these are used as learners to K-NN classifiers. In total 6 classifiers were created, 3 for each set of features, for 1, 5 and 9 nearest neighbors.

The learners were tested using the *leave-one-out* cross validation. The results of the classification process are shown in *Table 4.11*.

> *In the leave-one-out cross-validation (**LOOCV**) one single observation from the original sample is used as the validation data and the remaining observations as the training data. This is repeated such that each observation in the sample is used once as the validation data.*

Nearest neighbors	Features (P:real; Q:reactive)	Classification accuracy (%)	Features (P:real; Q:reactive)	Classification accuracy (%)
1	P	90.48	Poly P	98.02
	Q	64.29	Poly Q	98.94
	PQ	100	Poly P Poly Q	99.51
5	P	91.67	Poly P	98.40
	Q	67.26	Poly Q	99.40
	PQ	100	Poly P Poly Q	99.51
9	P	92.26	Poly P	97.78
	Q	70.24	Poly Q	99.40
	PQ	100	Poly P Poly Q	99.51

Table 4.11 : K-Nearest Neighbor classification results

4.4.3 Discussion

From the overall results the first thing to notice is that the results do not vary very much based on the number of neighbors. The second thing to notice is the low classification accuracy that is obtained when using only reactive power as a feature. This happens because the step changes in the reactive power are somehow close between different appliances.

The results for just the real power step change are better, with an overall classification rate that is above 90%. Problems, again, happen when the step change values are close between different appliances, for example: the hand mixer turning on (with a 110 Watt average step) is classified as if it was a hand blender turning on (average step of 116 Watts), and the hand mixer turning of (average step change of – 75 Watts) is confused with the fan turning off (step

change of - 72 Watts in average). The good news is that when combining real and reactive mean step changes the classification accuracy is of 100, for this set of appliances. While all the other appliances are correctly classified using only the real power step change, using the reactive power step for the hand mixer turning on (average of 75 VAR) will correctly classify it because the reactive step change of the hand blender is much smaller (23 VAR in average). The same happens between the hand mixer turning off (- 34 VAR) and the fan turning off (-5 VAR).

Regarding the polynomial coefficients the classification accuracy is very high using any of the features, and the first thing that this tells is that this solution will overcome the difficulty of appliances with close consumption levels, and probably when the same appliance is in a different consumption mode.

With a classification accuracy very close to perfection to this data set one may be tempted to believe that the event classification "silver bullet" was found, however this is not true, as the data set used in the test is very small when compared to what happens inside a house. Because of this future work in this field must include the deployment and testing of these algorisms in real world settings.

5

Conclusions and future work

5.1 Low cost home energy monitoring **52**
5.2 Project expansion **52**

5.1 Low cost home energy monitoring

Looking at the project as a whole, one must say that the final system meets the initial vision. The final system is a very cost-effective energy monitor with some load disaggregating capabilities and at the same time, a very flexible research platform for non-intrusive load monitoring.

The system presented in this thesis was deployed in 27 houses and suffered successive improvements, resulting in one peer reviewed publication [24]. Currently a version of the load monitor with disaggregation was deployed in some houses, and the expectation are to see how the availability of disaggregated information will change the way the home owners interact with the meter.

Further research also includes the possibility of sharing classified signatures across houses, as this would open a new path on load monitoring. One could for example create a system where users can upload and download signatures to and from a central repository changing forever the way that the training process of Non-Intrusive Load Monitoring.

5.2 Project expansion

NILM offers a big field of research. A recent project at University of Madeira aims at using it as part of an automatic electric extension that will detect and turn-off appliances that are found to be in stand-by mode.

Another very important project would be opening this home energy monitor to other programming languages and / or hardware systems. As seen in chapter 2 (State of the art), there are already some similar projects regarding smart meters that could be easily extended to also support NILM.

It cannot be forgotten that NILM can be easily extended to other domains, opening the possibility to create sensor networks at a very low cost. The ability to sense a whole house together with the possibility of inferring human activity will open the doors to a new kind of system where computers work in the shadow to help users in their everyday tasks. Home automation [w16] and ambient intelligence [w17], for instance, are two fields of study that can greatly benefit from NILM.,

Finally, it is also believed that there is still a very big lack of services that use of this technology, making it appealing not only for the consumers but also for the electric companies and appliance manufacturers. Hence there is an opportunity to do research in the field of service design, creating innovative services on top of low cost sensing technology.

References

[1] Attari, S. Z., DeKay, M. L., Davidson, C. I., & Bruine De Bruin, W. (2010). Public perceptions of energy consumption and savings. *Proceedings of the National Academy of Sciences of the United States of America*, 107(37), 16054-16059.

[2] Chisik, Y. (2011). An Image of Electricity: Towards an Understanding of How People Perceive Electricity. *Proceedings of the 13th IFIP TC13 Conference on Human-Computer Interaction*

[3] Froehlich, J., Findlater, L., & Landay, J. (2010). The design of eco-feedback technology. *Proceedings of the 28th international conference on Human factors in computing systems CHI 10*.

[4] Parker, D., Hoak, D., Meier, A., & Brown, R. (2006). How Much Energy Are We Using? Potential of Residential Energy Demand Feedback Devices. *Proceedings of the 2006 Summer Study on Energy Efficiency in Buildings,* -1665-06.

[5] Peschiera, G., Taylor, J. E., & Siegel, J. A. (2010). Response–relapse patterns of building occupant electricity consumption following exposure to personal, contextualized and occupant peer network utilization data. *Energy and Buildings*, 42(8), 1329-1336.

[6] Berges, M. E. (2010). A Framework for Enabling Energy-Aware Facilities Through Minimal-Intrusive Approaches. *PhD thesis, Carnegie Mellon University*.

[7] Hart, G. W. (1992). Nonintrusive appliance load monitoring. *Proceedings of the IEEE, 80(12)*, 1870-1891

[8] Norford, L. K., & Leeb, S. B. (1996). Non-intrusive electrical load monitoring in commercial buildings based on steady state and transient load-detection algorithms. *Energy and Buildings*, 24(1), 51-64.

[9] Laughman, C., Lee, K., Cox, R., Shaw, S., Leeb, S., Norford, L., & Armstrong, P. (2003). Power signature analysis. *Computer*, 1(2), 56-63

[10] Shaw, S. R., Leeb, S. B., Norford, L. K., & Cox, R. W. (2008). Nonintrusive Load Monitoring and Diagnostics in Power Systems. *IEEE Transactions On Instrumentation And Measurement*, 57(7), 1445-1454

[11] Berges, M., Goldman, E., Matthews, H. S., & Soibelman, L. (2009). Learning Systems for Electric Consumption of Buildings. *Computing in Civil Engineering*, 1-1. ASCE

[12] Marchiori, A., Hakkarinen, D., Han, Q., & Earle, L. (2011). Circuit-Level Load Monitoring for Household Energy Management. *Power,* 10(1), 40-48

[13] Patel, S. N., Robertson, T., Kientz, J. A., Reynolds, M. S., & Abowd, G. D. (2007). At the flick of a switch: Detecting and classifying unique electrical events on the residential power line. *Ubicomp07* (Vol. 4717, pp. 271-288). Springer

[14] Gupta, S., Reynolds, M. S., & Patel, S. N. (2010). ElectriSense: single-point sensing using EMI for electrical event detection and classification in the home. *Proceedings of the 12th international conference on Ubiquitous computing*, 139-148

[15] Zeifman, M., Member, S., & Roth, K. (2011). Nonintrusive Appliance Load Monitoring: Review and Outlook. *Solar Energy*, 57(1), 76-84

[16] Bergés, M., Soibelman, L., & Matthews, H. S. (2010). Leveraging data from environmental sensors to enhance electrical load disaggregation algorithms. *Proceedings of the International Conference Computing Civil and Building Engineering*

[17] Rowe, A., Berges, M., & Rajkumar, R. (2010). Contactless sensing of appliance state transitions through variations in electromagnetic fields. *Proceedings of the 2nd ACM Workshop on Embedded Sensing Systems for Energy Efficiency in Building* (pp. 19-24)

[18] Cohn, G., Gupta, S., Froehlich, J., Larson, E., & Patel, S. N. (2010). GasSense: Appliance-Level, Single-Point Sensing of Gas Activity in the Home. *Pervasive Computing*, 265-282.

[19] Froehlich, J. E., Larson, E., Campbell, T., Haggerty, C., Fogarty, J., & Patel, S. N. (2009). HydroSense: infrastructure-mediated single-point sensing of whole-home water activity. *Proceedings of the 11th international conference on Ubiquitous computing* (Vol. 9, pp. 235-244)

[20] Luo, D., Norford, L. K., Shaw, S. R., & Leeb, S. B. (2002). High Performance Commercial Building Systems Monitoring HVAC Equipment Electrical Loads from a Centralized United Technologies Corporation. *ASHRAE Transactions*, 108(1), 841-857

[21] Quintal, F., Nunes, N. J., Ocneanu, A., & Berges, M. (2010). SINAIS: home consumption package: a low-cost eco-feedback energy-monitoring research platform. *Proceedings of the 8th ACM Conference on Designing Interactive Systems* (p. 419–421)

[22] Energy, R. (2002). Measuring reactive power in energy meters. *Metering International*, 1(1), 52–54

[23] Matthews, H. S., Soibelman, L., Berges, M., & Goldman, E. (2008). Automatically Disaggregating the Total Electrical Load in Residential Buildings: a Profile of the Required Solution. *Computing in Civil Engineering*, 381-389

[24] Nunes, N., Pereira, L., Quintal, F. & Berges, M. (2011). Deploying and evaluating the effectiveness of energy eco-feedback through a low-cost NILM solution. *Proceedings of the 6th International Conference on Persuasive Technology,* Columbus, OH, USA

Referenced Web Resources

[w1] UN Conference on Environment and Development (1992). http://un.org/geninfo/enviro.html (9/2011)

[w2] SINAIS project. http://sinais.m-iti.org (9/2011)

[w3] Kill A Watt P4400. http://www.p3international.com (9/2011)

[w4] Plugwise. http://www.plugwise.com (9/2011)

[w5] The Energy Detective. http://www.theenergydetective.com (9/2011)

[w6] Wattson. http://www.diykyoto.com (9/2011)

[w7] Current Cost. http://www.currentcost.com (9/2011)

[w8] EnerSure. http://www.trendpoint.com/home/products/enersure (9/2011)

[w9] Tweet-a-Watt. http://www.ladyada.net/make/tweetawatt (9/2011)

[w10] Flukso. http://www.flukso.net (9/2011)

[w11] Open Energy Monitor. http://openenergymonitor.org (9/2011)

[w12] PowerBox. http://people.ece.cornell.edu/land/courses/ece4760/FinalProjects/s2008/cj72_xg37/cj72_xg37 (9/2011)

[w13] ACme. http://acme.cs.berkeley.edu (9/2011)

[w14] Likelihood-ratio-test. http://en.wikipedia.org/wiki/Likelihood-ratio_test (9/2011)

[w15] Audio Interchange File Format. http://en.wikipedia.org/wiki/Audio_Interchange_File_Format (9/2011)

[w16] Home automation. http://en.wikipedia.org/wiki/Home_automation (12/2011)

[w17] Ambient Intelligence. http://en.wikipedia.org/wiki/Ambient_intelligence (12/2011)