TÉCNICO LISBOA

Development and implementation of a DSP-based three-phase power quality measurement system

Luís Filipe Coelho da Silva Teixeira de Magalhães

Thesis to obtain the Master of Science Degree in

Electronics Engineering

Supervisor: Prof. Pedro Miguel Pinto Ramos

Examination Committee

Chairperson:	Prof. Jorge Manuel Torres Pereira
Supervisor:	Prof. Pedro Miguel Pinto Ramos
Member of the Committee:	Prof. Gil Domingos Marques

October 2014

Acknowledgments

First of all, I'd like to thank my parents and all my family for all the sacrifices they made to make sure I would finish my studies and get a chance for a better life and Joana for all the love and support.

I would also like to thank my supervisor, Professor Pedro Ramos for the opportunity to work on a project of a subject that I enjoy and for all the support, guidance and, above all, patience offered to make sure the project works as intended.

One person who I also must mention is Mr. Pina dos Santos, who is always ready and available to help and without whom the prototyping phase would take twice as long. I also have to mention my colleagues and friends, especially those in room 1.4.42 for all the help, encouragement and fun moments, making the university feel less like a workplace and more like a home.

There are no words to describe the motivation and positive feelings and thoughts these people give me every day to keep working and to keep improving in every aspect of my life. To all of you, thank you.

Resumo

Hoje em dia, o número de aparelhos electrónicos ligados à rede eléctrica é cada vez maior. Paralelamente, o número de problemas operacionais relacionados com a energia proveniente da rede também aumentou. O facto destes aparelhos interferirem electromagnéticamente com a rede e serem também mais sensíveis a este tipo de fenómenos, faz com que haja uma necessidade da monitorização da energia recebida por parte do fornecedor.

Estes fenómenos electromagnéticos que perturbam a qualidade de energia entregue pelo fornecedor podem ter várias fontes: podem ser um problema na própria rede, podem ter origem num cliente que afecta a rede de tal forma que reduz a qualidade da energia dos clientes que estão ligados perto de si, tal como podem ser sentidos e serem originados pelo mesmo cliente, sem chegar a perturbar a rede. Pequenos produtores são um exemplo de perda de qualidade de energia eléctrica, pois representam um maior número de locais onde tem de ser controlada a injeção de energia eléctrica com um menor número de recursos.

Com a monitorização da qualidade de energia, o fornecedor é capaz de perceber a fonte dos problemas. Desta forma pode, então, proceder à sua correcção, conseguindo, assim, um serviço de qualidade. Isto porque estas perturbações têm efeitos nefastos nos elementos que as recebem, dos quais se podem destacar a redução do tempo de vida útil, perda de informação ou mesmo danificação do equipamento. O target desta aplicação não é o cliente doméstico, pois este não produz muitas perturbações, mas sim o cliente industrial, que, devido aos aparelhos utilizados na sua função, tem maior probabilidade de originar este tipo de perturbações.

Esta tese tem como objectivo desenvolver, implementar e caracterizar um sistema electrónico baseado em DSP capaz de monitorizar em tempo real a qualidade de energia de um sistema trifásico. O sistema deve detectar todo o tipo de eventos definidos nas normas internacionais referentes à qualidade de energia, registar os instantes em que esses eventos começam e terminam e ainda todos os parâmetros que caracterizam cada tipo de evento. Os dados recolhidos devem posteriormente ser enviados por ligação ethernet para um servidor web que recolhe a informação e a disponibiliza em rede através de uma página da web. A utilização da ligação ethernet nesta aplicação trás várias vantagens, como, por exemplo, o facto do sistema de monitorização ficar disponível na rede, podendo ser acedido por qualquer computador na rede, não só aquele a que está ligado, podendo mesmo ser ligado a um router com ligação wireless e estar acessível a partir da internet.

Palavras-chave: Monitorização da qualidade de energia, processador de sinal digital, deteção e registo de eventos, ethernet.

Abstract

Nowadays, the number of electronic devices connected to the power grid is increasing. With it, the number of operational problems related to the energy provided by that same power grid has also increased. This is because these devices cause electromagnetic interference and are also sensitive to these kinds of phenomena, referred to as power quality disturbances or events. As such, there's an increased need to monitor the power quality on the provider side.

These power quality disturbances or events have various possible sources: they can happen in the transmission lines, due to lightning strikes or accidents, or they can caused by an end-user, which will feel the impact of those events on his side and can also possibly transmit that disturbance to the grid, affecting other users connected to the same power grid. An example of this would be costumer generation, which represents more places where the injection of electrical energy in the power system needs to be controlled, with fewer resources than in other sectors.

With power quality monitoring, the provider is able to detect the source of these events. With this information, there's a possibility to fix the problem, thus improving the overall service provided. These events have a negative impact, such as power outages, equipment malfunction or destruction, loss of information, among others. The target of this application would be industrial energy consumers, which, due to the equipment they use, have a higher probability to generate power quality events, unlike the residential energy consumer.

The objective of this thesis is to develop, implement and characterize a DSP-based electronic system that is able to monitor, in real-time, the power quality of a three-phase power grid. This system should be able to detect every event described in international standards of power quality, to register the duration of such events, recording the starting and final instant and also to register every relevant parameter of the event. After that, the retrieved data should be sent to a PC through an Ethernet connection, which will afterwards send the data to a web server. The Ethernet connection brings several advantages and grants great flexibility and mobility for the system, because it connects the equipment to a Local Area Network (LAN), which makes it accessible for any PC in that LAN and it can always be connected to a router with wi-fi connection and the system will be accessible through the internet. All these advantages justify the use of Ethernet over the use of USB, for example, which is another common method to connect systems to a PC.

Keywords: Power quality monitoring, DSP, power quality event detection and characterization, Ethernet.

Contents

List of Figures

List of Tables

Acronyms

AC	*Alternating Current*
AC-DC	*Alternating Current-Direct Current*
ADC	*Analog to Digital Converter*
CCITT	*Comité Consultatif International Téléphonique et Télégraphique*
CRC	*Cyclic redundancy check*
CWT	*Continuous Wavelet Transform*
DAI	*Digital Applications Interface*
DC	*Direct Current*
DC-DC	*Direct Current-Direct Current*
DDR2 SDRAM	*Double Data Rate Synchronous Dynamic Random-Access Memory*
DFT	*Direct Fourier Transform*
DHCP	Dynamic Host Configuration Protocol
DMA	*Direct Memory Access*
DPI	*Digital Peripheral Interface*
DSP	*Digital Signal Processor*
EHV	*Extra High Voltage*
exFAT	*Extended File Allocation Table*
FAT	*File Allocation Table*
FFT	*Fast Fourier Transform*
FIR	*Finite Impulse Response*
FTDI	*Future Technology Devices International*
GB	*Gigabytes*
GPIO	*General Purpose Input/Output*
GPRS	*General Packet Radio Service*
HV	*High Voltage*
I/O	*Input/Output*
I^2C	*Inter-Integrated Circuit*
IC	*Integrated Circuit*
IEC	*International Electrotechnical Comission*
IEEE	*Institute of Electrical and Electronics Engineers*
IIR	*Infinite Impulse Response*
IP	*Internet Protocol*
kB	*Kilobytes*
LAN	*Local Area Network*
LV	*Low Voltage*
MB	*Megabytes*
MFLOPs	*Million Floating Point Operations Per Second*
MV	*Medium Voltage*
NAND	*Not AND*
Op-Amp	*Operational Amplifier*
PCG	*Precision Clock Generators*
PHP	*Hypertext Preprocessor*
PQ	*Power Quality*
PQDIF	*Power Quality Data Interchange Format*
pu	*per-unit*
PWM	*Pulse Width Modulation*

RAM	*Random Access Memory*
RMS	*Root Mean Square*
RTC	*Real Time Clock*
SDHC	*Secure Digital High Capacity*
SDSC	*Secure Digital Standard Capacity*
SDXC	Secure Digital Extended Capacity
SDRAM	*Synchronous Dynamic Random Access Memory*
SLG	*Single line-to-ground*
SPI	*Serial Peripheral Interface*
SPORTs	*Serial Ports*
SRU	*Signal Routing Unit*
STFT	*Short-time Fourier Transform*
TB	*Terabyte*
TCP	*Transmission Control Protocol*
THD	*Total Harmonic Distortion*
UART	*Universal Asynchronous Receiver/Transmitter*
UDP	*User Datagram Protocol*
UPS	*Uninterruptible Power Supply*
USB	*Universal Serial Bus*
VISA	*Variable Instruction Set Architecture*

1. Introduction

1.1 Motivation and purpose

Power quality is a term that refers to a wide variety of electromagnetic phenomena that characterizes the voltage and current at a given time and at a given location on the power system [1]. These phenomenons are also referred to as power quality events and they include, but are not limited to voltage sags and swells, harmonics and transients, which will be described in more detail on chapter 2. When these events occur, the quality of the energy degrades, meaning that power quality refers, in fact, to the lack of quality of the power available.

In recent years, with the increasing use of electronic and electric devices, the power quality problem has grown, because these devices are sensitive to power quality events. In some cases, they can even be the cause of such events, because they represent a non-linear, asymmetrical and/or non-steady load [2]. Examples of these kinds of loads are electric power motors, frequency conversion machines and arc furnaces [3]. Electric power commutation can also cause these events.

Flickering lights [4], power outages and the malfunction of electric devices, or even their destruction, are common consequences of this lack of energy quality. With this, the need to monitor, diagnose and prevent/correct problems in the power systems has also increased, because there's an interest of the client to get good energy quality and, of course, an interest of the provider as well to provide a good service. The provider should take action to provide a good service, either with preemptive measures such as the monitoring of the energy in the power grid and installing equipment to protect the lines from overvoltages or to compensate certain power quality events, to minimize their effects, or with reactive measures, such as quickly fixing energy faults due to fallen poles or an impaired transformer. On the other end, the end-user wants a guarantee of a good service, because, for example, in a manufacturing facility or a factory, a power outage or disturbance can halt the production of goods, which translates into lost revenue [5] [6].

But the power quality problem is not just economical. One particular application is with medical equipment [7]. Those devices are, like all electronic devices, sensitive to the quality of the energy they receive. Others also cause some power quality events, which degrade the quality of the energy in the power system. For example, x-ray machines and magnetic resonance imaging machines produce high amounts of harmonics and voltage transients on the power supply [8]. This decrease in energy quality can affect other equipment such as heart monitors and pulmonary ventilators, which have microprocessors in them, that are sensitive to these events [9]. Sometimes, these devices are the only thing to keep people alive and are invaluable tools to diagnose health problems. The companies that distribute this kind of equipment often remove modules of power supply devices for these devices to work under conditions of bad power quality in guarantee period. This is because many of the power networks of hospitals are obsolete and don't have any kind of overvoltage protection.

Because of all of this, there is a need to develop efficient, intelligent and cheap systems that can monitor the power quality [10]. These systems will give information about the quality of the energy to the provider by detecting disturbances which can help to identify their origin, so that those situations

can be dealt with. End-users can impact the quality of power of the grid and also of themselves [11]. Some of the most common causes of power quality events are external, such as voltage sags caused by lightning strikes, accidents or equipment failure on the transmission and distribution grid. But, for example, a factory can cause power quality events that will affect its own facilities. Voltage sags resulting of turning on large Alternating Current (AC) motors or significant voltage harmonics as the result of the use of non-linear solid-state power conversion equipment are some examples. These can transit to the distribution grid that is feeding the plant, affecting other users being fed by the same grid.

With the advance of technology, especially of networking and the internet of things [12] [13], these systems have less limitations regarding the way to transmit the data obtained. The implementation of the Ethernet communication in this work is important, because it allows multiple instances of this system to be placed in different nodes of the power system, gathering data regarding the power quality events and relaying that information, pre-processed, to prevent a huge dataflow that could diminish the quality of the connection and affect other users on the same network, to a central database where those events could be studied thoroughly. Information like the relationship between voltage and currents in different voltage levels and its propagation in the network can be obtained and it could also be possible to get a better understanding of voltages and currents behavior under different event conditions. For example, the design and realization of a regional power quality monitoring system can be found in [14]. This system uses power quality modules that have Universal Serial Bus (USB) 2.0, Ethernet and optical fiber as communication interfaces. These modules send their data to a web server, which uploads it to the internet for any user who would like to check it. It also uses General Packet Radio Service (GPRS) to transmit the data to other monitoring stations. This system is an example of how these systems benefit with the use of Ethernet and internet related communication schemes.

1.2 Objectives

Real-time power quality monitoring systems are very demanding in terms of acquisition and processing due to the events to detect. These are very diverse, have very different amplitude ranges and very different duration times. In this project, the monitoring is limited to the voltages of the power grid, therefore ignoring current and power measurements. The processing unit will be a Digital Signal Processor (DSP) of the SHARC family from Analog Devices (ADSP-21489).

The objective of this thesis is to develop, implement and characterize a DSP-based electronic system that is able to monitor, in real time, the power quality of a three-phase system. This system should:

- Detect every kind of event defined in national and international standards;
- Register the duration of the events, by logging the initial and final instant of the occurrence (using a Real Time Clock (RTC)), and also all the parameters that characterize that same event;
- Send the data retrieved to a web server using Ethernet.

1.3 State of the Art

Power quality measurement systems need to evaluate time power quality parameters, in which case they don't need complex signal processing because this can be extracted directly from the samples retrieved, and frequency power quality parameters, which are extracted based on a spectrum analysis of the signal [15]. The most common method to transform the signal from the time-domain to the frequency-domain is the Fast Fourier Transform (FFT) [16].

To evaluate time power quality parameters, power quality measurement and monitoring systems usually calculate the Root Mean Square (RMS) value of the input [17] [18] [19] [20]. This is the least demanding way of processing the signal, from a computational point of view. Although you can calculate it whenever you get a new sample, this value is usually calculated every cycle or half cycle. While this is a simple way of measuring power quality, it's not sufficient on its own. Events such as sags, swells and interruptions are easy to detect, given the fact that they have a big impact on the magnitude of the signal. But transients don't affect much the RMS value of the signal, so the use of this technique alone is not enough to characterize all of the events. In addition to transients, short time variations are difficult to detect given the time window of a cycle used to calculate the RMS value, and even harmonics. Basically, not all frequency power quality parameters can be evaluated with this method alone. These kinds of events are detected in a spectrum analysis.

In order to do a spectrum analysis, measurement systems need to transform the time-domain signal in a frequency-domain signal. To do this, these systems need to employ an algorithm, which can be of varying complexity. The most used algorithms are the Direct Fourier Transform (DFT), for asynchronous sampling machines and the FFT for synchronous sampling machines [15]. Each of these algorithms have their advantages, for example the FFT needs a larger number of samples, but takes less computations than the DFT [21]. But these two algorithms are flawed because they should only be applied to periodic waveforms. One way to overcome this problem is to introduce a time dependent parameter of the Fourier Transform, with an appropriate window applied to the non-stationary signal. This technique is known as the Short-time Fourier Transform (STFT). But, since the window has a fixed width and is chosen *apriori*, the frequency-time resolution is also fixed, making it difficult to analyze a sinusoidal signal of low frequency, such as the one coming from the power supply, affected by a high frequency disturbance, such as a transient. There are other techniques like the wavelet transform, which has properties like limited effective time duration, band pass spectrum, waveform similar to disturbance and orthogonality. It also allows locating information in time and frequency domains. This means that it's possible to obtain high correlation when disturbances occur and decompose these events into different components without energy aliasing. However, it has low resolution for low-frequency components, the decomposition frequency bands are fixed and it's sensitive to noise [20] [22].

Since every method has its advantages and disadvantages, the best course of action is to either use an improved version of the algorithm focused on certain areas, like the All Phase FFT [23], which is more precise on detecting higher harmonics and analyzing power quality than the classic FFT or the S-Transform [24] [25] [26], which is a hybrid of STFT and Continuous Wavelet Transform (CWT) techniques. The methods and techniques were developed and improved because they allow the

automatic classification of power quality events when used together with pattern recognition systems, knowledge based expert systems or any logic to discriminate different types of events. Some examples of these systems or techniques are fuzzy logic [27] [28] [29], artificial neural networks [30] [31] [32] or support vector machines [33] [34] [35]. The advantage of using artificial neural networks is its capability to handle the noisy data that is present in real-time measurements. However, the main drawback is the large number of training cycles and the requirement of retraining the entire neural network for every new power quality event, as demonstrated in [36] and [37]. Fuzzy logic doesn't require any training process, since it uses simple "IF-THEN" conditions.

As for commercial applications, there're a lot of products on the market that do power quality analysis. Two examples of these systems are the AEMC PowerPad Model 8335 Three-Phase Power Quality Analyzer [38] and the Fluke 1750 Three-Phase Power Quality Recorder [39]. They're both very similar, as they measure voltage ratings, current ratings, power ratings, frequency, among other things. They both have SD-card memory, in the case of PowerPad to record trend data and some alarms. The main difference is how the information is given to the user. In the case of the AEMC PowerPad, the product itself has a display in which the results of the analysis can be seen. The Fluke 1750 has a separate tablet that will show what the equipment is currently recording. They both have a way to connect to a PC, giving the user the possibility to analyze the results and to store them. On the PowerPad this is done via USB, using software called DataView, which displays and analyzes real-time data on the PC and also prints reports using standards. The Fluke 1750 transmits data over a 100BaseT high-speed Ethernet connection to the PC, which can be analyzed with the Fluke Power Analyzer Software that will generate reports with automated EN 50160 reporting and compliance. It has also a wireless control interface, which uses the Bluetooth SPP protocol to transmit the data wirelessly. On top of this, when there's no connection with the equipment, it will automatically download the data to the SD card.

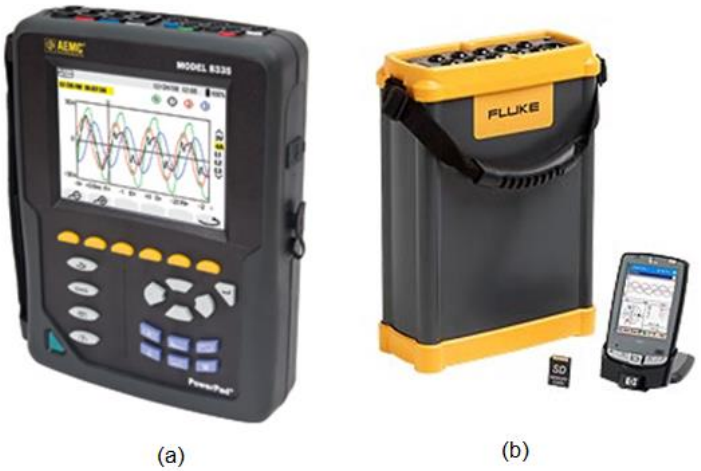

Figure 1: AEMC PowerPad 8335 (a) and the Fluke 1750 (b).

1.4 Electrical Standards

For the past few years, much effort has been made in defining power quality phenomena and their effects on electrical and electronic equipment. The recent efforts have been focused on harmonizing standards between Institute of Electrical and Electronics Engineers (IEEE) and the International Electrotechnical Comission (IEC). But there's still a need for significant additional research to establish the relationship between power quality/reliability levels and the various characteristics of the supply system. These characteristics need also to be defined in a more statistical manner to allow more effective risk assessments by end users using statistical techniques. The goal is to improve the economical side of power quality management. A roadmap for power quality standards, how they are now, how they need to be and the way to get there is described in [40].

Some of the power quality events have indices they need to abide to. In the case of Portugal, the standards used are the EN 50160 and the IEC 61000. These standards define the voltage characteristics that costumers should expect from their power supply. IEEE and IEC standards have different values than EN 50160 for these voltage characteristics. This is an important factor to notice because if the provider of electrical energy doesn't comply with the standard that the costumer is expecting, there will be compatibility problems between power system and electric devices.

The EN 50160 standard gives the main characteristics of the supply voltage, as well as the margins of error for those characteristics, for LV and MV. This is useful information for the clients using such services, because they want to know what their contract covers and if that contract with the provider is effectively being met. It's also useful information for the developers of electrical and electronic devices, because they have to make sure their equipment is compatible with the power system. It should also be important for the providers of electrical energy because if these requirements aren't met, the client is the one being harmed, since some, if not all, of the consequences of the power quality events are harmful to electrical and electronic devices. As far as responsibility goes, the client is also responsible for the quality of the energy in the power system, because he represents a load for the power system, a potential harmful one if he generates power quality events. There is always a need to make sure the impact in the power system is as low as possible. For example, if the client has large AC motors, these can generate voltage sags when they start up. It should be the clients responsibility to mitigate as much as possible this effect so it won't affect other devices he might have and the power system.

In Table 1, a comparison is made between the EN 50160 and EMC 61000. Table 2 is a complimentary table with the values regarding harmonics. These tables show the difference between the characteristics of the voltage waveforms of the two standards. There are two main reasons for this difference:

- The EMC standards concern the utility voltage, according to IEC 038, while EN 50160 deals with the supply voltage. The differences are due to voltage drops in the installation and disturbances originating from the network and from other equipment supplied from the installation.
- EN 50160 gives only general limits, which are technically and economically possible for the supplier to maintain in public distribution systems. When more rigorous conditions are

required, a more rigid and detailed agreement between costumer and supplier must be negotiated.

Table 1: Comparison between EN 51060 standard and EMC standard 61000 (source [30]).

Parameter	Supply voltage characteristics according to EN 50160	Low voltage characteristics according to EMC standard EN 61000	
		EN 61000-2-2	Other parts
Power frequency	LV, MV: mean value of fundamental measured over 10 s ±1% (49.5 - 50.5 Hz) for 99.5% of week -6%/+4% (47- 52 Hz) for 100% of week	2%	
Voltage magnitude variations	LV, MV: ±10% for 95% of week, mean 10 minutes rms values		±10% applied for 15 minutes
Rapid voltage changes	LV: 5% normal 10% infrequently Plt ≤ 1 for 95% of week MV: 4% normal 6% infrequently Plt ≤ 1 for 95% of week	3% normal 8% infrequently Pst < 1.0 Plt < 0.8	3% normal 4% maximum Pst < 1.0 Plt < 0.65 (EN 61000-3-3) 3% (IEC 61000-2-12)
Supply voltage dips	Majority: duration <1s, depth <60%. Locally limited dips caused by load switching on: LV: 10 - 50%, MV: 10 - 15%	urban: 1 - 4 months	up to 30% for 10 ms up to 60% for 100 ms (EN 61000-6-1, 6-2) up to 60% for 1000 ms (EN 61000-6-2)
Short interruptions of supply voltage	LV, MV: (up to 3 minutes) few tens - few hundreds/year Duration 70% of them < 1 s		95% reduction for 5 s (EN 61000-6-1, 6-2)
Long interruptions of supply voltage	LV, MV: (longer than 3 minutes) <10 - 50/year		
Temporary, power frequency overvoltages	LV: <1.5 kV rms MV: 1.7 Uc (solid or impedance earth) 2.0 Uc (unearthed or resonant earth)		
Transient overvoltages	LV: generally < 6kV, occasionally higher; rise time: ms - µs. MV: not defined		±2 kV, line-to-earth ±1 kV, line-to-line 1.2/50(8/20) Tr/Th µs (EN 61000-6-1, 6-2)
Supply voltage unbalance	LV, MV: up to 2% for 95% of week, mean 10 minutes rms values, up to 3% in some locations	2%	2% (IEC 61000-2-12)
Harmonic voltage	LV, MV: see Table 3	6%-5th, 5%-7th, 3.5%-11th, 3%-13th, THD <8%	5% 3rd, 6% 5th, 5% 7th, 1.5% 9th, 3.5% 11th, 3% 13th, 0.3% 15th, 2% 17th (EN 61000-3-2)
Interharmonic voltage	LV, MV: under consideration	0.2%	

Table 2: Harmonic content for LV and MV on EN 50160 (source [30]).

Odd harmonics				Even harmonics	
Not multiples of 3		Multiples of 3			
Order h	Relative voltage (%)	Order h	Relative voltage (%)	Order h	Relative voltage (%)
5	6	3	5	2	2
7	5	9	1.5	4	1
11	3.5	15	0.5	6…24	0.5
13	3	21	0.5		
17	2				
19	1.5				
23	1.5				
25	1.5				

1.5 Document organization

All the theory related with the power quality events, their description, causes, consequences and even a brief summary of a possible way to prevent them is presented in Chapter 2. Problems regarding the consensus of the power quality standards will also be presented in this chapter.

In Chapter 3 the proposed system architecture is presented together with a description of the main blocks of the system that was developed.

In Chapter 4 the power quality event method used in this project is presented, along with experimental results obtained.

Lastly, Chapter 5 presents some conclusions about this system, its advantages and future work.

2. Power Quality

In the next subchapters, a detailed description of each power quality event is presented along with the different origins of these events on different sectors of the power system, possible causes, consequences and ways to diminish their effect.

2.1 Origin of Power Quality (PQ) Events

Power quality events can occur in different sectors of the power system and can propagate to other sectors or have immediate consequences on that same sector. There are four main sectors: power plants, where the electrical energy is generated; the transmission grid, that guides the generated energy to the clients; the distribution grid, that receives the incoming energy from the transmission grid and lowers the voltage level to an adequate one for the clients; and the clients themselves that are seen as loads for this system. How power quality events are generated in these different sectors of the power system is described in the next subchapters.

2.1.1 Power plants

Electrical energy is a product and, as such, needs to have a minimum quality level to satisfy producers and end-users. This quality level can be affected by voltage and current magnitude deviation from nominal values, frequency deviation and other phenomena. These events might be enough to, for example, halt the production of a plant or damage an electronic device that was connected to the grid. So, the quality of service of electrical energy is composed of two parameters:

- Power Quality, regarding the quality of the voltage and current waveforms. The characteristics (amplitude, frequency, etc.) should be as close to the nominal level as possible;
- Power Reliability, related to the continuity of the service. This means that the service of electrical energy shouldn't stop and should be of quality.

Most of the electrical energy produced in Portugal, like in most parts of the world, is produced in power plants. The equipment in these power plants and all the adjacent power lines are carefully designed so that the system can handle an increase in load without creating voltage sags. Also, in these power plants the existing loads are linear, therefore they won't generate any power quality events. These areas are the ones with the highest levels of power quality. But from these areas to where the electrical energy is consumed by the end-user, the power quality index can drop due to disturbances on either the transmission grid, on the distribution grid or disturbances given by the loads connected to the power grid.

2.1.2 Transmission grid

In the points where power plants connect to the transmission grids, energy is injected with voltage levels ranging from 100 kV to 400 kV (High Voltage [HV] or Extra High Voltage [EHV]) to optimize the transmission of energy to the points where the transmission grid connects to the distribution grid. In this journey, the electrical energy is exposed to numerous factors that can degrade its quality.

Some of the most important factors are:

- Inductance variation of the power lines, given by the lateral oscillation of the posts which hold the lines, caused by the strong winds. This effect can be mitigated but there's always a negative impact due to random oscillations in the line terminations.
- Lightning strikes on the lines. This happens because these lines are at a high altitude and are made of conductive material. When a lightning strike occurs, it originates a transient with very high magnitude. In the transmission grid this isn't worrisome because the voltage level is already high. However, in the distribution grid lightning strikes can produce very dangerous overvoltages. This can also damage the protection equipment in the grid, originating interruptions of variable duration proportional according to the magnitude of the transient.
- Voltage drops on the lines. The impedance of the transmission lines isn't null, causing a variation of the magnitude of the voltage waveform of 10% between situations of no load and full load;
- Physical aggressions, such as fires, bird nests or line rupture.

2.1.3 Distribution grid

In the distribution grid, the voltage needs to be adapted to the end-user. This means that this voltage is fed to substations that will lower the voltage level from the EHV or HV to Medium Voltage (MV) or Low Voltage (LV). Some clients purchase the electrical energy in MV and then lower it inside of their own facility. Other customers, such as residential end-users purchase their electrical energy in LV.

In cases where the reactive energy is too high, there is a need to install capacitor banks to compensate the power factor. This will add more impedance, which will increase the probability of power quality events.

The distribution grid is also exposed to external factors that degrade the quality of the electrical energy. Since they're close to the end-users, they are subjected to a lot more physical aggressions such as accidents, animals and trees. And, since they are the connection to the costumers, they exist in larger numbers, making it the most probable source of power quality events. A lightning strike is also a much more severe occurrence in this system, because of its proximity to the clients. Even though the effects can be mitigated with protection equipment, the probability of asset damage of the clients is much higher. It can also originate interruptions that vary in duration with the magnitude of the transient event caused by the lightning.

2.1.4 Loads

Non-linear loads are known to be one of the major factors for the poor power quality and the effect that end-users have on the power grid and are also the biggest victims of the lack of power quality. This non-linearity comes with the use of devices such as diodes, transistors and microcontrollers. These devices are known to have a better efficiency, in terms of energy consumption, and better control. But their use affects the power grid which they're connected to, because they cause the distortion of the voltage waveform.

An example of non-linear loads is a switching power supply. These devices have numerous advantages regarding linear power supplies such as higher efficiency, lower consumption, less weight and reduced dimensions. But the commutation process causes a lot of power quality events because the current they use isn't sinusoidal. Thus, the voltage and current waveforms will be distorted.

2.2 PQ Events

Table 3 presents the classification of electromagnetic phenomena by the International Electrotechnical Commission [1]. The terms high-frequency and low-frequency on this table are not referred to a particular frequency range, but to the relative difference in principal frequency content of the phenomena.

Table 3: IEC's classification of the principal phenomena causing electromagnetic disturbances (source **[1]**)

Conducted low-frequency phenomena	Harmonics, interharmonics
	Signal systems (power line carrier)
	Voltage Fluctuations
	Voltage dips and interruptions
	Voltage imbalance
	Power-frequency variations
	Induced low-frequency voltages
	DC in AC networks
Radiated low-frequency phenomena	Magnetic Fields
	Electric Fields
Conducted high-frequency phenomena	Induced continuous wave voltages or currents
	Unidirectional transients
	Oscillatory transients
Radiated high-frequency phenomena	Magnetic fields
	Electric fields
	Electromagnetic fields
	Continuous waves
	Transients

The various events presented in Table 3 also need to be characterized. This characterization differs from one event to another due to the nature of its cause. Regarding steady-state phenomena, the characterization could consist of:

- Amplitude;
- Frequency;
- Spectrum;
- Modulation;

- Source impedance;
- Notch depth;
- Notch area;

Whereas for non-steady state phenomena, the following attributes can be used:

- Rate of rise;
- Amplitude;
- Duration;
- Spectrum;
- Frequency;
- Rate of occurrence;
- Energy potential;
- Source impedance;

In the next sub-chapters, the most common power quality events will be discussed.

2.2.1 Transients

Transients are rapid changes in the RMS value of voltage or current waveforms. This events have very small duration times (µs to ms) and can reach a magnitude of thousands of volts, in case of a voltage transient. These transients have two major types, depending on the polarity of the disturbance: if the disturbance is unidirectional in polarity, being it positive or negative, then it is an impulsive transient; if the disturbance is changing its polarity rapidly, then it's called an oscillatory transient.

An impulsive transient is a sudden, non-power frequency change in the steady-state condition of voltage, current, or both that is unidirectional in polarity. They are characterized by their amplitude, spectral content, rise and decay times. Since their duration time is so small, these transients are quickly damped by resistive circuit components, so they don't spread throughout the circuit. How fast they're damped and how effectively depends on the circuit itself, so different parts of the circuit will be affected differently by the same kind of transient. One example that has already been mentioned is in the case of a lightning strike which will cause an impulsive transient. These impulsive transients can excite power system resonance circuits and create another type of voltage/current transients called oscillatory transients.

An oscillatory transient refers to a voltage or current whose instantaneous values are quickly changing polarity. As with impulsive transients, the main characteristics of this event are the spectral content, duration and magnitude. For example, an oscillatory transient with 500 kHz and a typical duration of microseconds is considered a high-frequency oscillatory transient. This is an important characterization, because high-frequency oscillatory transients are often a result of a local system response to an impulsive transient. Power electronic devices produce oscillatory transients as a result of commutation and RLC snubber circuits. Back-to-back capacitor energization also results in oscillatory transients. This phenomenon occurs when a capacitor is energized in close electrical proximity to a capacitor bank already in service. Figure 2 and Figure 3 show examples of captured transient events.

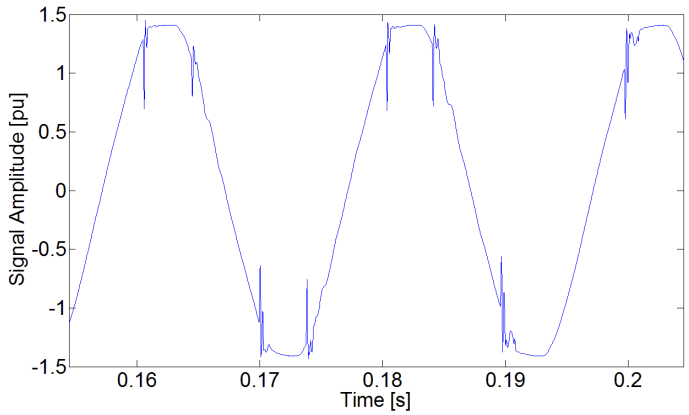

Figure 2: Example of a captured transient event **[41]**.

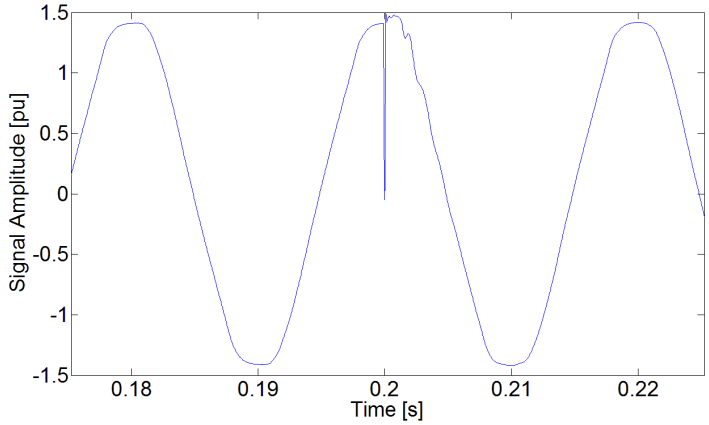

Figure 3: Another example of a captured transient event **[41]**.

According to [40], a well-known standard in the field of transient overvoltage protection is ANSI/IEEE C62.41-1991, IEEE Guide for Surge Voltages in Low Voltage AC Power Circuits. This standard defines the transient environment that equipment may encounter and provides specific test waveforms that can be used for equipment withstanding testing. The transient environment is a function of the equipment or surge suppressor location within a facility as well as the expected transients from the supply system.

2.2.2 Interruptions

An interruption is defined as a decrease in supply voltage or load current to unacceptable values for a period of time. In [1], the value is less than 0.1 per-unit (pu) for a period of time not exceeding 1 minute, for a short-duration interruption. For a sustained interruption, it would be the decrease to zero of the supply voltage for a period of time in excess of 1 minute.

In [42], the comparison of EN 50160 and EN 61000 is made. In EN 50160, short interruptions of supply voltage are only given in regard to the duration and the amount permitted per year. In LV/MV, a few hundred interruptions are permitted per year, but the duration of 70% of those events

must be under one second, even though they can go up to three minutes. A long interruption of supply voltage is defined as an interruption for more than 3 minutes, allowing ten to fifty a year. In EN 61000-6-1, 6-2, a short interruption of supply voltage is given by a reduction of 95% of the RMS nominal value for five seconds. A long interruption of supply voltage is not defined. In conclusion, for different standards, the definition of short-term and long-term interruptions differs, because the requirements and applications are also different.

Interruptions can be the result of power system faults, equipment failures and control malfunctions. Some of these interruptions may be preceded by a voltage sag when these interruptions are due to faults on the source system. Figure 4 shows an example of an interruption event. The consequences of these events depend on the equipment characteristics, duration of the interruption, etc.

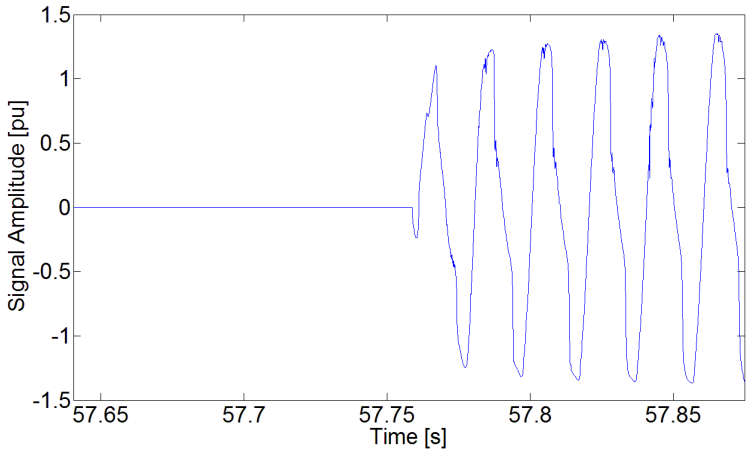

Figure 4: Example of a captured interruption event **[41]**.

There are several ways to mitigate the effects of interruptions. Some of them are at the distribution and transmission grid level, others are at the end-user level. One important way to counteract interruptions is to keep the power source close to the load. This way, the grid connecting both is shorter and less vulnerable to external physical aggressions. At the end user level, the solution is the use of an UPS (Uninterruptible Power Supply). In residential homes, hospitals, laboratories and data centers, there's usually an emergency power system which is a standby generator that can include lighting, electric generators, fuel cells, UPSs and other apparatus to provide backup power resources in a crisis or when regular systems fail.

At the distribution and transmission grid level, there are three main objectives to counteract interruptions:

- Reducing the number of failures, for example by replacing aerial lines for underground lines and increasing the isolation of the conductive materials;
- Reducing the duration of the failures, for example with the use of current limiting fuses;
- Changing system topology, for example through the implementation of redundancy on the power system.

2.2.3 Sags

Sags can be defined as a reduction of the RMS value of a voltage or current waveform. Usually, the term sag refers to short duration event and undervoltage to a longer duration event. In Portugal, the EN 50160 standard is used to define these events and in that standard, a sag has a duration of 10 ms to 1 s and it's characterized by the reduction of the RMS supply voltage value to under 90% of its nominal value. This RMS value may not be enough to keep certain equipments working in ideal conditions.

Voltage sags can be caused by system faults, such as short-circuits, switching of heavy loads or even starting of large motors. These situations will increase the amount of necessary current of the system, leading to a voltage drop. If the current magnitude is large relative to the system available fault current, the resulting voltage sag can be significant. Figure 5 shows an example of a captured sag event.

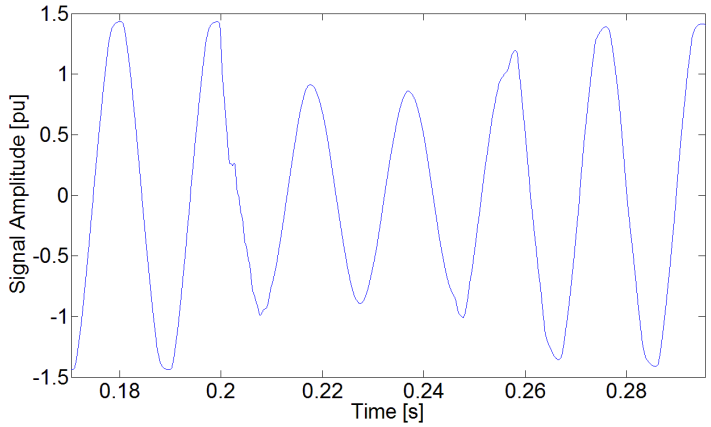

Figure 5: Example of a captured sag event **[41]**.

The consequences of voltage sags can be disastrous to certain equipments and productive processes. Some examples of these consequences are:

- In induction motors, the binary is proportional to the square of the supply voltage. If the value of the supply voltage has a small reduction, that will translate into a big reduction in the motor speed.
- In synchronous motors, the binary is proportional to the supply voltage. If the supply voltage value drops, the motor speed will drop as well, but not as much as in induction motors.
- Electronic devices usually need an AC-DC converter when they're connected to the power system. Nowadays, with the use of switching mode power supplies, the effect of voltage sags isn't as critical as with the use of linear power supplies. But, if the supply voltage value drops significantly, these power supplies may not be able to continue to keep the device working, having the same impact as a voltage interruption. In the case of domestic devices, the consequences can be as minimal as information loss, but in the case of industrial devices, the consequences can have high financial costs.

There are several ways to mitigate the effects of voltage sags, some in the transmission and distribution grids and some at the end-user level. At the transmission and distribution grid level, the solutions could be:

- Increase of the short-circuit power of the power system;
- Decrease the time to fix the defect;
- Isolate consumers with higher sensitivity to these effects.

At the consumer level, there are devices that can mitigate the effects of voltage sags such as: transformers that can automatically regulate the number of winding turns, UPSs, etc.

2.2.4 Swells

A voltage swell is the opposite of a voltage sag. It can be defined as an increase of the RMS value of a voltage or current waveform. According to [1], typical duration is between 0.5 cycles to 1 minute and typical magnitudes are between 1.1 and 1.8 pu. In EN 50160, a swell is characterized by having a duration of 10 ms or more and a magnitude greater than 1.1 pu. Usually, this magnitude won't go above 1.5 kV AC. A long duration swell is also known as an overvoltage.

Like sags, swells are usually associated with system fault conditions, but they are far less common. They can occur due to a single line-to-ground fault on the system, resulting in a temporary voltage rise on the unfaulted phases. They can also be caused by switching off a large load or switching on a large capacitor bank. Figure 6 shows an example of a swell event.

The consequences of voltage swells are typically failure or malfunction of electronic devices. Other circuit components like cables, transformers and rotational machines can have their lifetime shortened with the occurrence of these events.

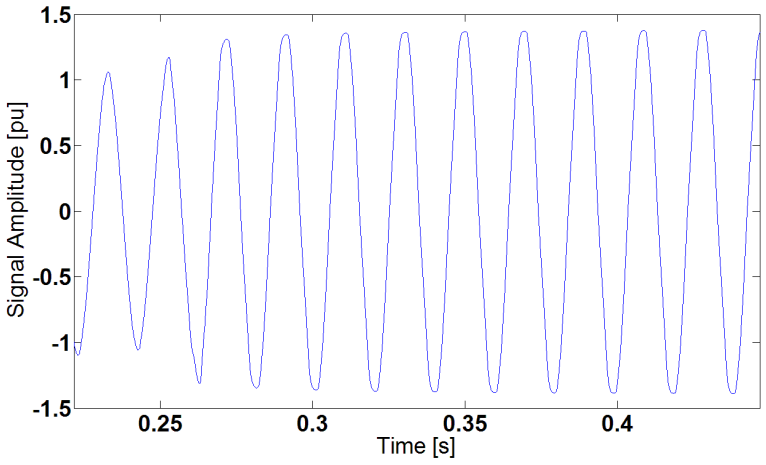

Figure 6: Example of a swell event.

2.2.5 Voltage Imbalance

Voltage imbalance is defined as the ratio of the negative or zero sequence component to the positive sequence component. The negative or zero sequence voltages in a power system generally result from unbalanced loads causing negative or zero sequence currents to flow [1]. Figure 7 shows an example of a one-week trend of imbalance measured at one point on a residential feeder.

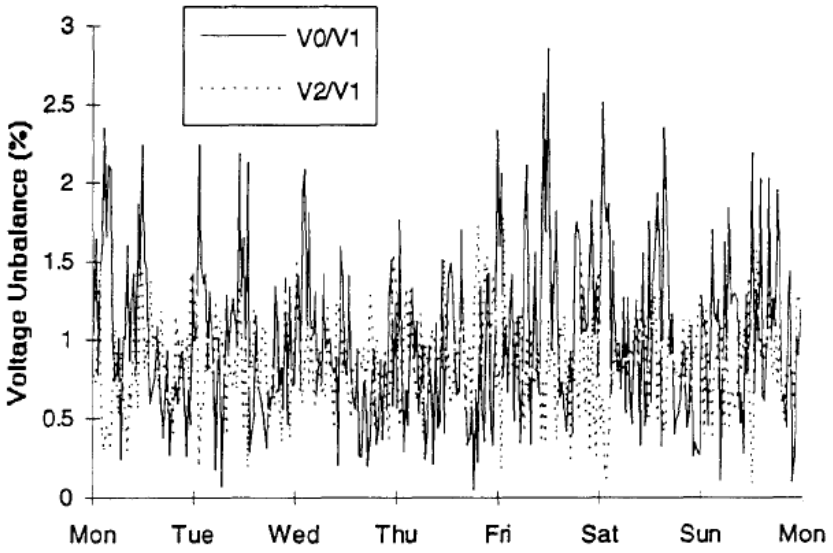

Figure 7: One-week imbalance trend on a residential feeder (source **[1]**).

One simple equation to calculate the voltage imbalance factor is

$$U_u = \frac{max\ deviation\ from\ average\ voltage}{average\ voltage} \times 100\% \text{ [1].} \quad (1)$$

This method is quite easy to apply given that the amplitudes of the voltages of the three phases are measured.

The major source of voltage imbalance is the uneven distribution of loads to the different phases of the power system. This uneven distribution will lead to an unbalanced current system, even when the voltage system is balanced. When the currents circulate through the system, they'll cause voltage drops due to the impedance of the lines, which will be different in each phase, thus unbalancing the voltage system.

Some examples of the consequences of voltage imbalance are:

- In induction motors, the temperature will rise, leading to a reduction in lifetime of these motors, since the isolation material suffers an increased deterioration when the windings have a high temperature.
- In synchronous machines the negative sequence current circulating through the stator will create a magnetic field that rotates at the same speed as the rotor speed, but on the opposite direction of the rotation given by the positive sequence. As such, the induced currents and voltages will have twice the network frequency, increasing significantly the losses in the rotor.
- Rectifiers, under normal operating conditions, create characteristic harmonic currents. When the system is unbalanced, they start creating 3rd order harmonics and its multiples. The presence of these harmonics is extremely prejudicial to the power system, because it can cause unforeseen resonance events, damaging devices attached to it.

Some examples of possible mitigation of voltage imbalance are:

- Redistribution of system loads;
- Increase of short-circuit power on the connection point;
- Power electronics;

2.2.6 Direct Current (DC) Offset

DC offset refers to the existence of a DC voltage or current in an AC power system. This can happen due to geomagnetic disturbance or to the effect of a half-wave rectification. Direct current in alternating current networks can be detrimental due to an increase in transformer saturation, additional stressing of insulation and other adverse effects.

2.2.7 Harmonics

Ideally, the sinusoidal voltage waveform of the power system would be composed of just its fundamental harmonic, at working frequency. So, the spectral content of this waveform would all be at 50 Hz, in the case of Portugal. But, in reality, the non-linear loads connected to the power system create harmonics. These harmonics are voltage and current waveforms with frequencies that are integer multiples of the working frequency. These combine with the fundamental voltage or current and produce waveform distortion. Usually, the magnitude of these harmonics is represented in a percentage of the fundamental waveform.

There are two ways to characterize the harmonic content of a waveform: either analyzing the spectral content of the signal, being able to determine the magnitude and phase angle of each individual harmonic; or using a single quantity, the total harmonic distortion, as measure of magnitude of harmonic distortion. The total harmonic distortion is given by

$$THD = \sqrt{\frac{\sum_{n=2}^{+\infty} A_n^2}{A_1^2}} \ . \tag{2}$$

Figure 8 shows the effect of harmonics on a voltage waveform and Figure 9 shows the effect of harmonics in a current waveform and the signal decomposition in its spectral content.

Harmonic distortion is a growing concern for end-users and for the overall power system due to the increasing use of power electronics equipment. The harmonic distortion problem gets even bigger with the increasing use of capacitors to compensate the power factor. These capacitors are placed in parallel with the reactance of the power system, which can produce unwanted electromagnetic resonance, at a frequency given by

$$f = \frac{1}{2\pi}\left(\frac{1}{\sqrt{LC}}\right) . \tag{3}$$

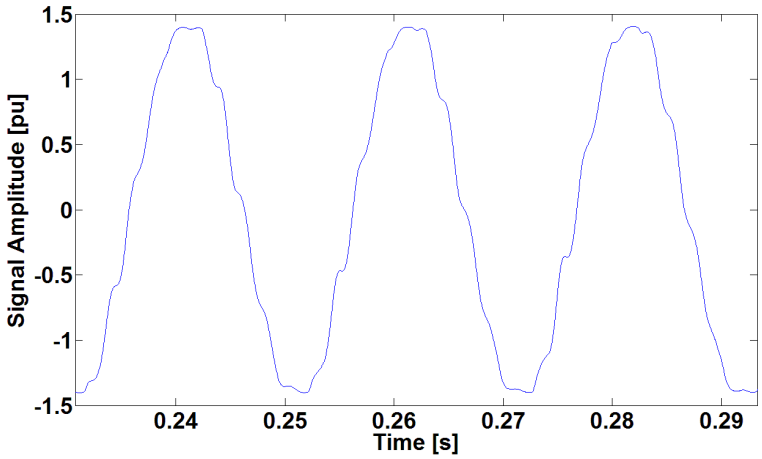

Figure 8: Impact of harmonics on the voltage waveform **[41]**.

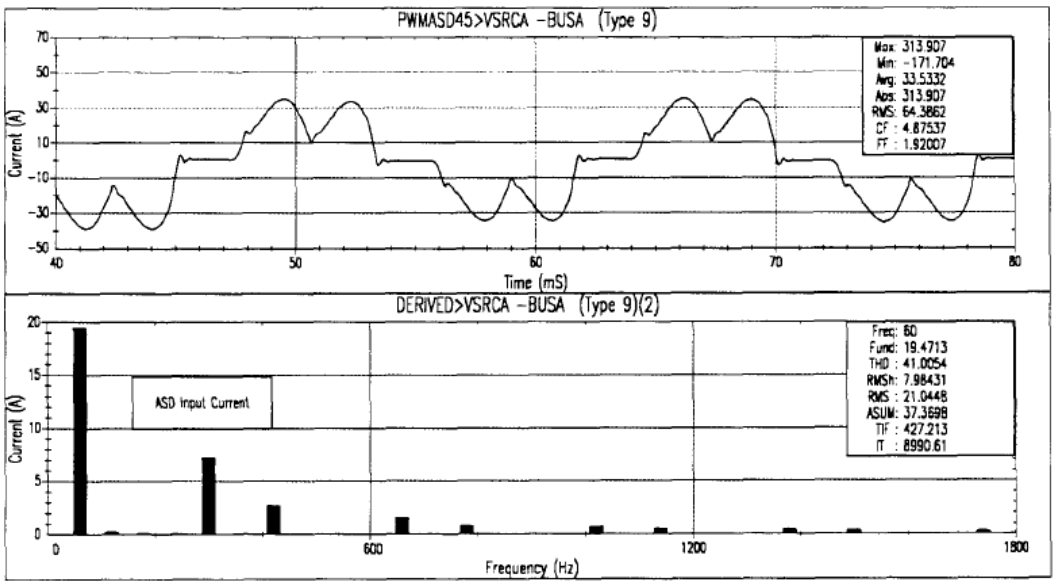

Figure 9: Impact of harmonics on a current waveform and its spectral content (source **[1]**).

In case of the current being drawn from the source having harmonic content close to the oscillation frequency, oscillatory currents with high magnitude can be occur, which can damage the capacitors and create high magnitude voltage harmonics. The most common devices that generate harmonic content are switched mode power supplies, which are used in PCs, battery chargers, consumer electronics, basically anywhere where there's a need for an efficient, low weight, size and cost solution for the power supply, and UPS.

Some examples of how to mitigate harmonic distortion are:

- Separation of harmonic content creating loads and sensitive to harmonic distortion loads;
- Use of active, passive and/or hybrid filters.

2.2.8 Interharmonics

Interharmonics are voltage and current waveforms whose frequency is not an integer multiple of the working frequency. These components can appear as discrete frequencies or as a wide-band spectrum. They are fairly common and can be found in networks of all voltage classes. The main sources of interharmonics are frequency converters, induction motors and arcing devices.

2.2.9 Notching

In the normal operation of power electronic devices, when the current is commutated from one phase to another, there is a momentary short-circuit between the two phases, originating a voltage disturbance. This disturbance is periodic and referred to as notching. This high frequency disturbance occurs continuously, so it can be characterized through the harmonic spectrum of the voltage affected. One of the most important causes of voltage notching are three-phase converters that produce continuous DC current. Figure 10 shows an example of voltage notching.

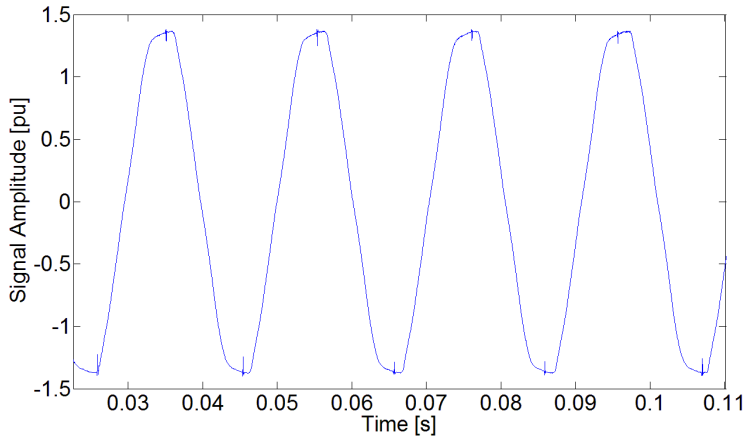

Figure 10: Example of a captured notching event **[41]**.

2.2.10 Noise

Noise consists of any unwanted waveform distortion that cannot be classified as harmonic distortion or transient. In power systems, noise can be caused by power electronic devices, control circuits, arcing equipment, loads with solid-state rectifiers and switching power supplies. Noise problems are aggravated by improper grounding.

Noise disturbs electronic devices such as microcontrollers, which are present in automatically controlled products and devices, such as automobile engine control systems, implantable medical devices, remote controls, office machines, appliances, power tools, toys and other embedded systems. The noise problem can be mitigated by using filters, isolation transformers and some line conditioners.

2.2.11 Voltage Fluctuations

Voltage fluctuations can be described as a systematic variation of the voltage envelope or a series of random voltage changes. These variations are faster than sags and swells, and their magnitude is usually between 0.9 and 1.1 pu. This voltage range is given by EN 50160. According to [1], the voltage range should be between 0.95 and 1.05 pu.

Any load that has significant current variations can cause voltage fluctuations. These voltage variations are, sometimes, erroneously referred to as flicker. The term flicker is derived from the impact of the voltage fluctuation on lighting intensity. Voltage fluctuation is the response of the power system to the varying load and light flicker is the response of the lighting system as observed by the human eye.

Some of the sources of voltage fluctuations are:

- Arc furnaces;
- Welding equipment;
- Other equipment such as x-ray machines, photocopiers, laser printers, microwave ovens and washing and drying machines.

Voltage fluctuations main consequences are oscillation in potency and binary of electrical machines, reduced efficiency of electrical equipment and flicker.

Some examples of how to mitigate voltage fluctuation disturbances are:

- Increase the power of short-circuit in the connection of the power system with the load;
- Reduce the reactive energy on the network;
- Separation of equipment that generates voltage fluctuations from equipment that is sensitive to this kind of disturbance;

2.2.12 Power Frequency Variations

The working frequency of the power system is given by the rotational speed of the generators. This rotational speed is set for a determined load, which creates a dynamic balance between the two variables. When this dynamic balance is affected, small changes in frequency occur. The size of the frequency shift and its duration depend on the load characteristics and the response of the generation system to load changes. In EN 50160, the acceptable frequency shifts are 1% of the nominal frequency value for 99.5% of the week and -6%/+4% for 100% of the week. In EN 61000-2-2, this value is 2%.

The main causes of power frequency variations are failures in the transmission and distribution grids, quick desynchronization of groups of large source of generation or a large block of load being disconnected.

These frequency variations don't have much impact in the power networks. They affect the speed of rotational machinery. In asynchronous systems, these disturbances are frequent and affect clocks that derive their timing from the 50 Hz frequency and also affect the performance of electrical motors. It can also cause harmonic filters to start working irregularly.

2.3 Summary

In this chapter the different power quality events, their possible causes, consequences and ways to possibly mitigate their effect were discussed. This will set the foundation for the developed work being discussed in subsequent chapters. Without the understanding of power quality phenomena and how it affects the voltage waveform, which is the interest of this work, it would be impossible to design a system able to detect and characterize such events.

3. System Architecture

The fully working power quality measuring system developed for the purpose of this project is independent of any other modules. The only exception is the router to be able to send reports about the captured events to a remote server. This means that it needs to be able to sample the signal coming from the power grid, adjust it to the proper voltage levels, analyze it using the various methods that will be described further on, store and send the results through the Ethernet connection. The main source of power for this system is the power grid, taking advantage of the fact that the voltage of the three phases is available. If and when this option fails, there's a backup battery to keep the system running. There's also the possibility of powering it with a 5V/2A transformer, although it should only be used in debug situations. The system architecture is the one shown in Figure 11. This system is composed of six main blocks: the signal acquisition and conditioning block, the DSP block, the Ethernet or communication block, the RTC block, the memory block and the redundant power supply block. A brief explanation of each block will be presented in the next subchapters.

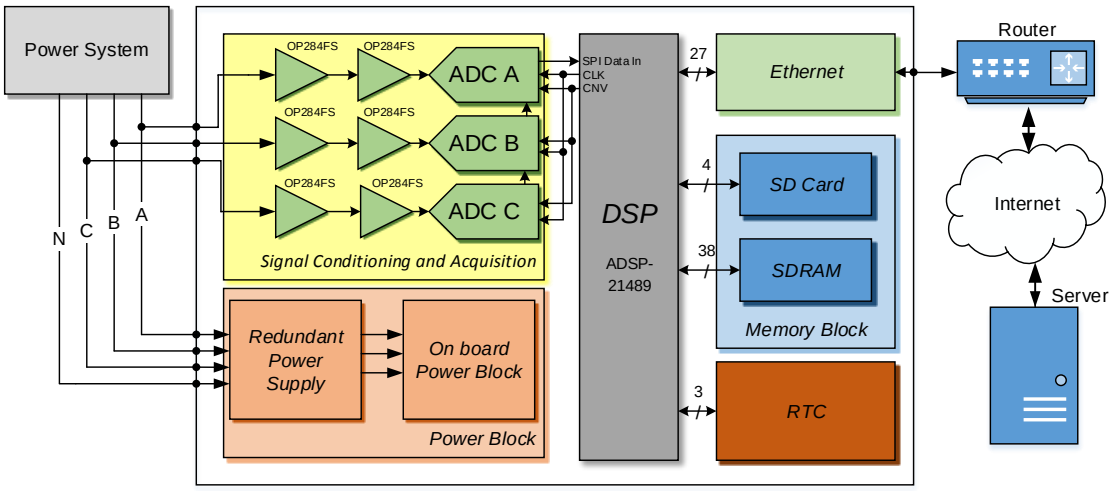

Figure 11: Block diagram of the proposed system architecture.

3.1 Signal Acquisition and Conditioning

This block comprises the circuitry for the signal conditioning and acquisition. Its purpose is to receive the input signals, in this case the three phases of the power system, adjust their voltage level to an appropriate one for the ADC, which will sample these input signals and send the samples to the DSP.

The first stage of this block is the interface between the input waveforms coming from the power system and the board. The purpose of this first stage is to reduce the voltage level of the input waveforms. There are different ways to implement this interface, but two were considered and compared. The first one would be using a voltage sensor, based on the Hall effect, which would sense the input waveforms and output a voltage level proportional to the one on its input, and the second one would be using a voltage divider. The voltage sensor solution requires an auxiliary circuit to define the voltage level at the output and it's more expensive. However, the input and output are isolated, which is a necessary feature when dealing with the power system, since the disturbances could affect and possibly damage components on the board. The voltage divider solution is less complex and

cheaper than the voltage sensor based on the Hall effect, although it doesn't isolate the input from the output. However, this problem can be solved with the use of optocouplers, which will isolate the board from the power system.

A consideration needed to be taken into account is the fact that the input signal is bipolar. Since the chosen Analog to Digital Converter (ADC) only translates positive voltage values into a digital word, there is the need for a stage that converts this waveform from bipolar to unipolar. Transforming it into a unipolar signal can be achieved with a circuit like the one shown in Figure 12.

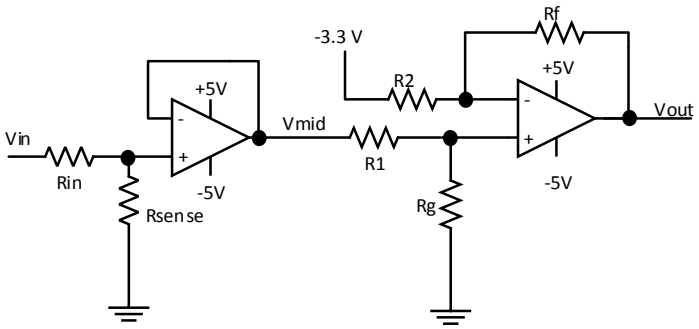

Figure 12: Circuit to transform the input signal into a unipolar signal.

The design and determination of the resistor values of this circuit takes into account the output voltage range, which is between 0 and 3.3 V, voltage level range of the ADC, and twice the nominal input voltage coming from the power system, which amounts to a peak value of 650 V or 460 V_{ef}. This is to protect the circuitry from potential high magnitude events and also to allow a bigger measurement range. The first Op-Amp of this circuit is a buffer, to offer even more protection, since it isolates the current from the input to the rest of the circuit. This circuit is replicated three times, one for each phase of the power system being sampled.

First off, the voltage divider needs to be designed. Since it must to be able to handle as much as twice the nominal voltage level, the input signal given by

$$V_{in} \simeq 650\cos(2\pi \times 50t) \tag{4}$$

needs to be transformed into

$$V_{sense} = 3.3\cos(2\pi \times 50t) \ . \tag{5}$$

This will make further calculations easier as it will be shown below. V_{sense} is given by

$$V_{sense} = \frac{R_{sense}}{R_{in} + R_{sense}} V_{in} \ . \tag{6}$$

So, in order to get (5), the ratio between the resistance values needs to be

$$\frac{V_{sense}}{V_{in}} = \frac{3.3}{650} \simeq 0.005 \ . \tag{7}$$

To get this ratio between voltage levels and also to have high input impedance, required to reduce the input current, the resistor values are

$$\frac{R_{sense}}{R_{sense}+R_{in}} = 0.005 \simeq \frac{5\,k\Omega}{5\,k\Omega + (10\ x\ 100\,k\Omega)} \ . \tag{8}$$

The voltage divider at the input of the circuit can be designed as a standalone circuit because of the high input impedance of the operational amplifier. The buffer setup won't interfere with the voltage divider, so the voltage level at R_{sense} will ideally be equal to (6). One additional note here is that R_{in} is not a single resistor, but a series of resistors. This is because the resistors will be SMD and there is no SMD resistor that can or should handle an effective voltage value of 230 V, much less 460 V. This would damage the resistor and possibly the board. With several resistors in series, the voltage is divided through all of them, reducing the voltage stress on each resistor. There is also a power concern, since one resistor couldn't dissipate the power applied to it without becoming damaged. That being said, R_{in} is composed of ten 100 kΩ resistors. The voltage level across each of them is

$$V = \frac{0.995 \times 650}{10} = 64.675\ V \ . \tag{9}$$

Since the voltage rating of the resistors is 150 V, this is an acceptable value for the voltage across each individual resistor. Also, the maximum RMS current is

$$I_{in} = \frac{V_{ef}}{R} = \frac{460}{1.005 \times 10^6} = 0.458\ mA \ . \tag{10}$$

With (10), it's possible to calculate the maximum power that each resistor will have to dissipate and make sure that value doesn't go over the power rating of the resistors. The power dissipated by R_{sense} is

$$P_{sense} = \frac{(460 \times 0.005)^2}{5000} \simeq 1\,mW \tag{11}$$

and the maximum power dissipated by each of the resistors of R_{in} is given by

$$P_{in} = \frac{64.675^2}{100000} \simeq 42\ mW \ . \tag{12}$$

Since the power rating of 0805 case SMD resistors is above 100 mW, these resistors can be used for this system.

For the second stage of this block, the signal needs to be converted from unipolar to bipolar. The signal given by (5) needs divided in half and then a DC offset needs to be added to make it unipolar. That offset will be equal to half its maximum amplitude. So, the signal at the end of this block will be

$$V_{out} = 1.65\cos(2\pi \times 50t) + 1.65 \ . \tag{13}$$

The stage responsible for this is the differential amplifier. Its transfer function is

$$V_{out} = \left(\frac{R_2 + R_f}{R_2}\right)\left(\frac{R_g}{R_1 + R_g}\right)V_+ - \left(\frac{R_f}{R_2}\right)V_- \ . \tag{14}$$

Looking at this transfer function it's easy to understand the reason for (5) and for (13). The simplest way to turn V_{out} into a unipolar signal is to have the V_+ signal in a voltage range between [-x,x] and then adding x to have a signal with a voltage range between [0,2x]. Since the chosen ADC will read values between 0 and 3.3 V, the input signal should be adjusted to a voltage range of [-1.65, 1.65] V and, to that signal, a DC offset of 1.65 V can be added to turn it into a signal between [0, 3.3] V. These

adjustments are made with the differential amplifier circuit. Knowing that the 3.3 V voltage level is available on the board, used to feed the reference pin of the ADC, it can also be used to make the wanted DC offset, defining the ratio between R_f and R_1 as 1/2. So, (14) turns into

$$V_{out} = \frac{3}{2}\left(\frac{R_g}{R_1 + R_g}\right)V_+ + 1.65\,. \tag{15}$$

The next step is to calculate the ratio between R_g and R_2. Knowing that the voltage level range of the input signal coming from the voltage divider is between [-3.3, 3.3] V it needs to be divided in half and so the ratio between R_g and R_2 can be determined as

$$\frac{3}{2}\left(\frac{R_g}{R_1 + R_g}\right)3.3 = 1.65 \rightarrow \left(\frac{R_g}{R_1 + R_g}\right) = \frac{1}{3}\,. \tag{16}$$

Knowing the ratios between the resistances, the value of those components can now be determined. This concludes the design stage of this circuit. To make sure the design is correct, the circuit was tested and simulated with PSIM. The result of the PSIM simulation is shown in Figure 13 and Figure 14.

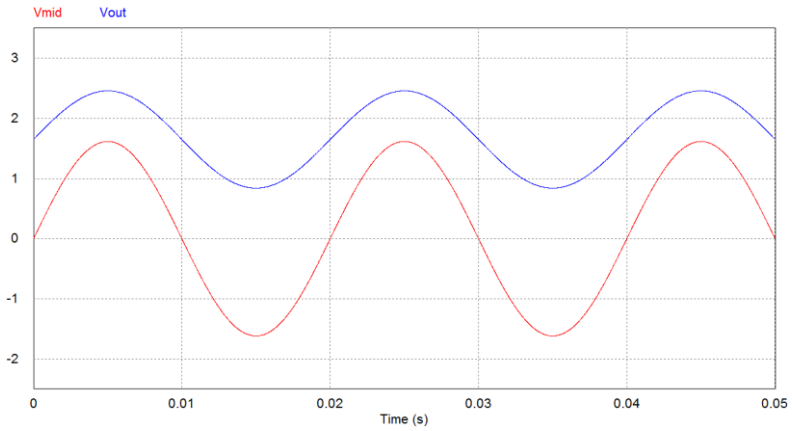

Figure 13: Output voltage (blue) for nominal input voltage (red).

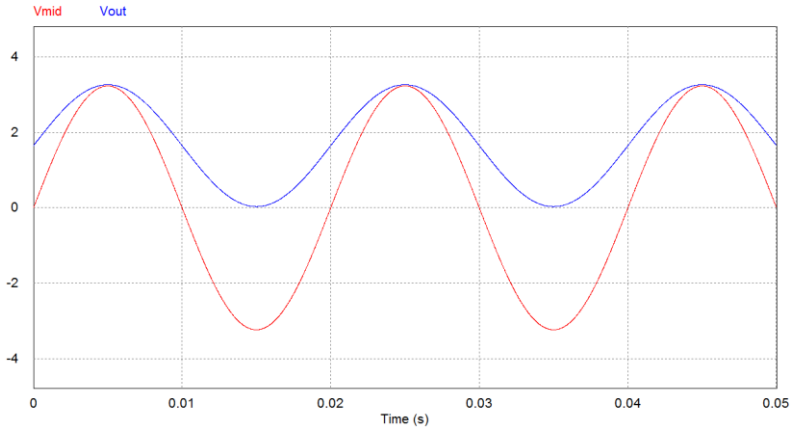

Figure 14: Output voltage (blue) for twice the nominal input voltage (red).

The experimental results are shown in Figure 15 and Figure 16.

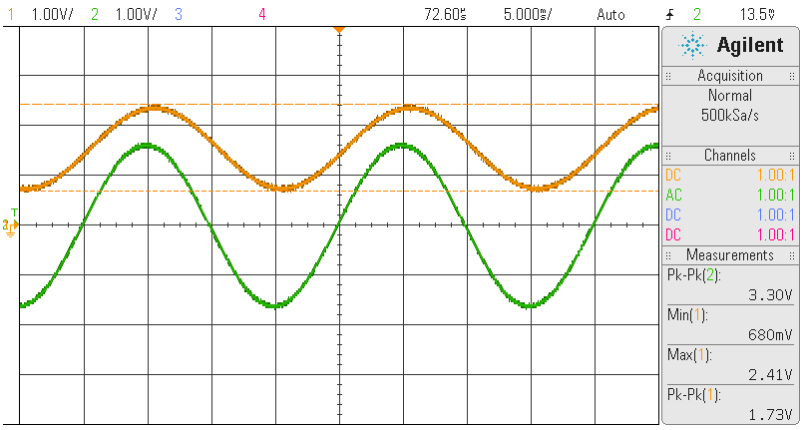

Figure 15: Experimental output voltage (yellow) for the nominal input voltage (green).

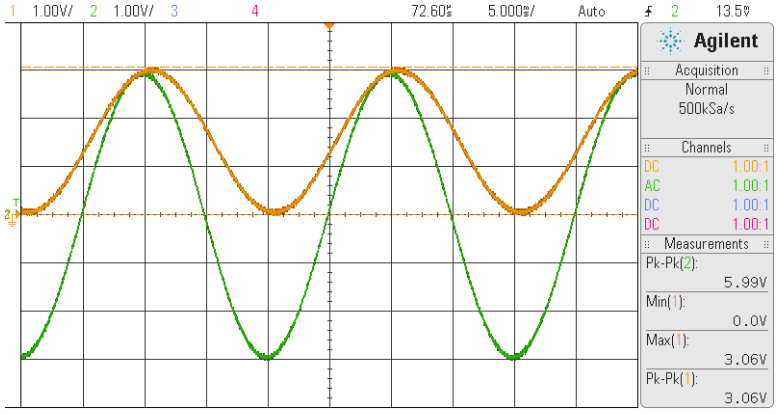

Figure 16: Experimental output voltage (yellow) for twice the nominal input voltage (green).

These experimental results were obtained using a function generator, which simulated the input voltage after the voltage divider. In Figure 16, the input voltage doesn't have the same value as the output voltage of the voltage divider when the input voltage is twice the nominal value. This is because the output voltage of the block was already at 0 V. Since the ADC doesn't take negative voltage levels at its input, the decision was not to increase further the input voltage level. This is due to variations in values of circuit components.

Another experimental result obtained was the frequency response of the signal conditioning circuit, which is shown in Figure 17 and Figure 18. These results are important to determine how the transfer function of the circuit varies with the frequency of the input signal. As shown in Figure 17 and Figure 18, the transfer function of the system remains constant until an input frequency of approximately 1 MHz. This means that power quality events with small duration, like transients, can be correctly detected and handled by the circuit.

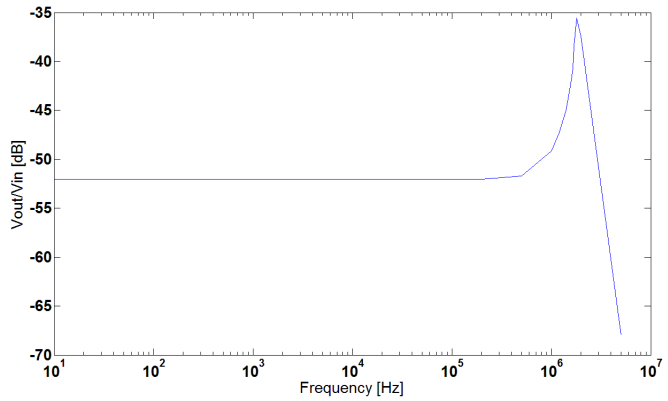

Figure 17: Magnitude plot of the circuit's frequency response.

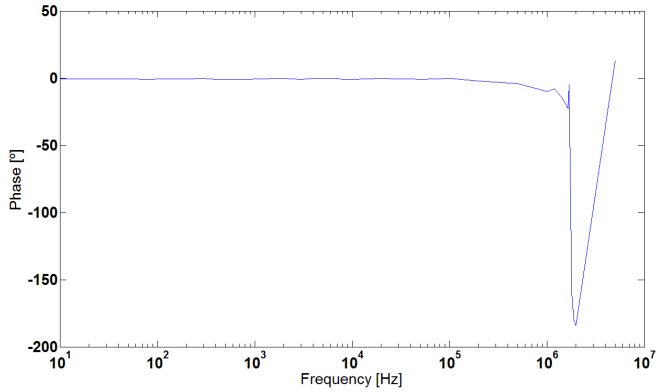

Figure 18: Phase plot of the circuit's frequency response.

The next stage of this block is the ADC. The chosen ADC is the AD7980 from Analog Devices [43]. This ADC was chosen because of its high sampling rate and the ability to chain multiple instances, reducing the number of connections to the DSP and because it uses the SPI communication protocol, which the DSP is prepared for, having dedicated pins. Since there are three phases to be sampled, there are three ADCs setup up in daisy chain. The connection scheme for the ADCs is shown in Figure 19.

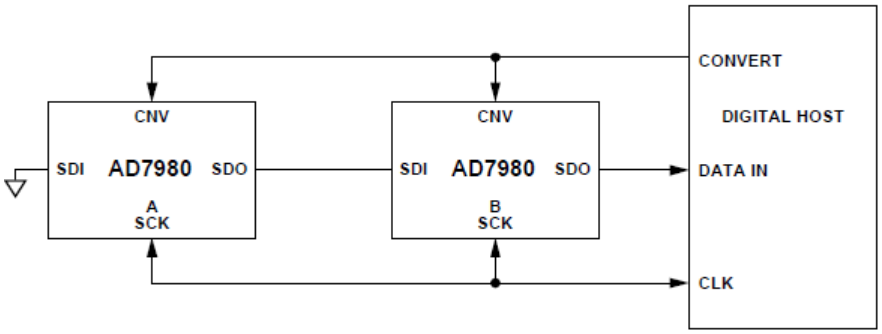

Figure 19: Chain mode without busy indicator connection diagram.

As shown in Figure 19, this connection scheme has only three lines connected to the DSP: The convert, clock and data lines. With SCK low, a rising edge on CNV initiates a conversion, selects the chain mode and disables the Busy indicator. In this mode, CNV is held high during the conversion

phase and subsequent data readback. When the conversion is complete, the MSB is output onto SDO and the AD7980 enters the acquisition phase and powers down. The remaining data bits stored in the internal shift register are clocked by subsequent SCK falling edges. For each ADC, SDI feeds the input of the internal shift register and is clocked by the SCK falling edge. Each ADC in the chain outputs its data MSB first, and 16 x N clocks are required to readback the N ADCs. The data is valid on both SCK edges. In this work, there is one ADC per phase, meaning three ADCs total, so 48 clock ticks are needed. Figure 20 shows the timing diagram of the communication protocol with the AD7980 using the connection diagram shown in Figure 19.

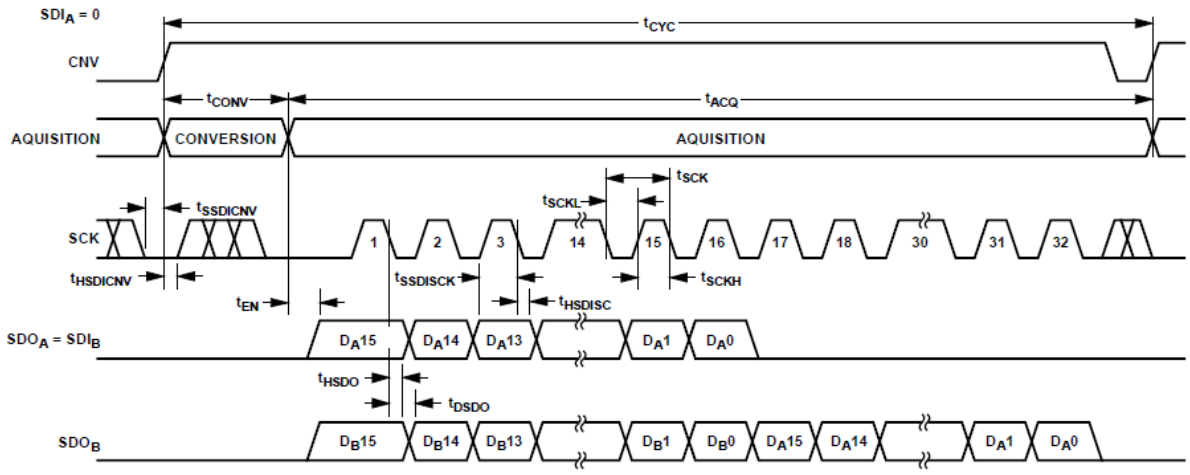

Figure 20: Chain mode without busy indicator timing diagram.

There is another important factor to keep in mind as to why the ADCs should be chained like this and that is because the sampling should be synchronized. In other words, when there's a rising edge of CNV, all ADCs start to sample at the same time, meaning that the samples will belong to the same instant in time of the three phases. This is a very important factor, since a fair comparison between the various phases can only be done if the samples belong to the same instants in time.

3.2 Redundant power supply

This block is responsible for supplying all the voltage levels required by the circuitry and is also responsible to keep the system powered when there's an interruption in supply voltage, with the help of a battery.

This block was first designed to be integrated in the same board as the DSP and other devices, but, due to the size of the of the AC-DC converters, a second board was designed to hold these components. This second board has the three AC-DC converters, the LI-Ion battery used to feed the system when there is no input voltage and a coin cell battery to feed the RTC. Figure 21 shows the simplified connection scheme between this second board and the system. A more detailed connection scheme is presented in Appendix B.

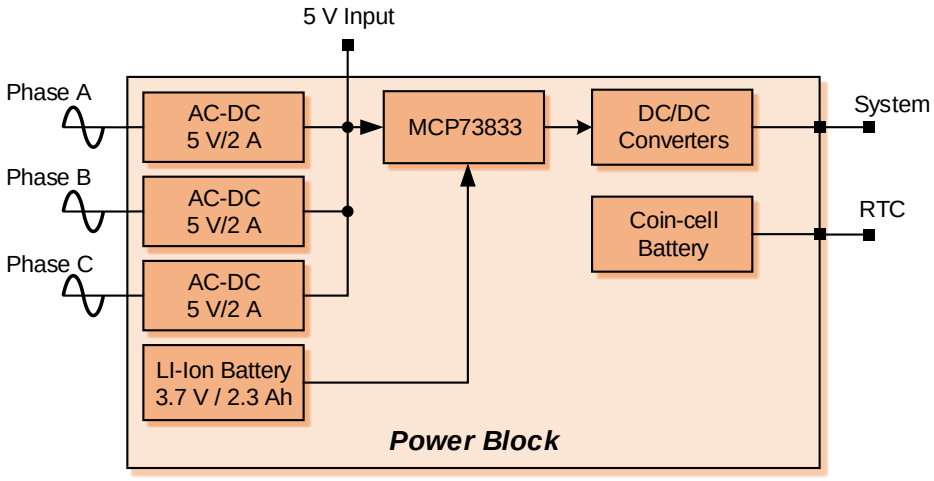

Figure 21: Simplified connection scheme of the power board to the main system.

The reason three AC-DC converters are used in this project is to guarantee that as long as one phase is working (has voltage), that phase alone can supply the entire system. Each AC-DC converter has an output of 5 V and a power rating of 10 W, which is necessary to feed the circuit. The calculations for the energy consumption of the system are presented in Appendix B, in which the worst case scenario was assumed, to make sure that each AC-DC converter could feed the system single handedly. One other aspect that was taken into consideration was the fact that the LI-Ion battery will provide 3.7 V and not the 5 V that the phase converters are giving. This means that some DC-DC converters on the system board need to be step-up converters in order to keep their DC output even when the input voltage drops below that value. In circumstances where one phase would be severely affected, a three-phase to DC converter would not work properly and the battery would be used a lot more often, instead of just using it in situations where it is absolutely needed. With the use of three AC-DC converters, each one connected to a different phase and the output voltages combined to feed the system, the load is evenly distributed along the phases, which is also necessary when using a three-phase power system.

The part of this block that was kept on the system board is the management circuit. This circuit is responsible for selecting the power source used to feed the system. It does this by checking if there's enough voltage coming from the three phase power system. If there isn't, it automatically changes to the battery. This is very important to make sure, for example, that the last reports of events are sent through Ethernet and that the files created in the SD Card are correctly closed; basically it's important to make sure everything functions properly and to make sure that an interruption doesn't cause a system failure. Another important function of this circuit is to charge the battery when it isn't being used, using the voltage coming from the power system.

The power management is done in the system board using the MCP73833 IC (Integrated Circuit) from Microchip [44]. This IC is usually used as a battery charger, but with some additional circuitry, it can also be used as a power management device [45]. Figure 22 shows the connection scheme of this IC.

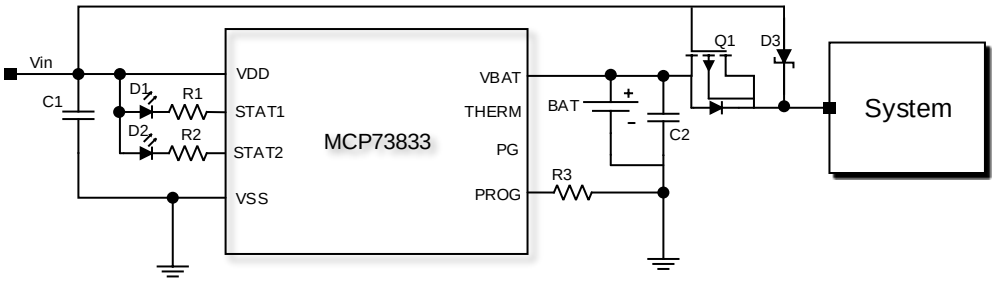

Figure 22: Schematic implementation of the battery management IC.

This circuit serves the two purposes: when there is energy present at the input of the system, the transistor Q1 will turn off and stop current flowing from the battery to the system, effectively disconnecting the battery. This means the system will be powered by the input power source, while the battery is being charged, if necessary. If the battery is already at full capacity, the IC will shut itself down, waiting for the condition to charge the battery again, which is when the voltage across the VBAT pin drops below the recharge threshold, leading to the beginning of another charge cycle. When there is no voltage present in the VDD pin, the IC will shut itself off too, Q1 will be turned on and current will flow from the battery to the system. The D1 diode will prevent this current to flow back to the power source. There is also a very useful aspect to this IC, which is the Power Good (PG) pin. This pin, when connected to the DSP, allows the program to determine if it's being powered by the power system or the battery, allowing it to save crucial data and powering down correctly instead of continuing its program and being shutdown unexpectedly. It also has status pins that are connected to two LEDs in the board, which give useful information to the user regarding the battery status, e.g. if it's charged or charging.

There is a need for additional components, besides only using the IC, because it isn't recommended to connect directly the battery to the system load. If the system were to be connected directly to the battery, the charge cycle could never end. Most Li-Ion battery chargers are based on Constant Current and Constant Voltage modes. The termination is based on the ratio of charge current and preset constant current (Fast Charge). If the system draws current from the battery, the charge current will never meet the termination value. This causes the non-termination of the charge management circuit. Also, the total system current would be limited by the charge current because the charger would deliver total system and battery charging current through the output pin. This value could not be enough to supply the whole system and it is also not recommended for applications that don't run on constant current, such as the one in this project.

3.3 DSP

Digital signal processors can perform computations for a large quantity of data samples, and use faster sampling frequencies. It is possible to create embedded measurement systems which do not require industrial-sized PCs as the main computational unit. This fact will then reduce the final cost of these measurement systems.

The pre-selected DSP for this work is the ADSP-21489 [46], which is from the SHARC family of DSPs of Analog Devices. This processor is based on a Super Harvard architecture [47] which has two separate program and data buses. It has an Input/Output (I/O) processor allowing direct interfacing with the processing core and the internal memory. There are 5 Mbits of on-chip Random Access Memory (RAM), a 16-bit wide Synchronous Dynamic Random-Access Memory (SDRAM) external memory interface and a Direct Memory Access (DMA) engine. It has a 400 MHz core clock speed, up to 2700 Million Floating Point Operations Per Second (MFLOPs) and features 32-bit fixed point and 32-/40-bit floating point arithmetic formats. In addition to its high core performance, Finite Impulse Response (FIR), Infinite Impulse Response (IIR) and FFT accelerators are provided to increase the total performance of the system. It also has a feature called Variable Instruction Set Architecture (VISA) that allows the code size to be decreased by 20% to 30% and increase the memory size availability.

The ADSP-21489 also has a total of 34 General Purpose Input/Output (GPIO) ports, consisting of 20 Digital Applications Interface (DAI) and 14 Digital Peripheral Interface (DPI) ports. These GPIO pins are internally multiplexed, using its Signal Routing Unit (SRU), to several hardware peripheral interfaces, such as generic Serial Ports (SPORTs), Precision Clock Generators (PCG), Peripheral Timers and communication interfaces, including Serial Peripheral Interface (SPI), Universal Asynchronous Receiver/Transmitter (UART) and I^2C. Routing these peripherals to specific GPIO ports is done in software and provides advantages in board layout as well as using any of the processor's peripherals as needed. It also has the ability to connect to faster external memory by providing a glueless interface to 16-bit wide Double Data Rate SDRAM (DDR2 SDRAM).

3.4 RTC

A real-time clock is a clock, in the form of an IC, which keeps track of the current time and allows the DSP to link the initial instant of a power quality event to a specific time and date. The only concern for the DSP is to process all the sampling data it can get and when it detects a power quality event, it will get the necessary information about the time and date from the RTC IC. The utilization of this IC is very important in the stage where it has to present the data to the user. Knowing that a 100 ms voltage sag with a magnitude of 0.75 pu happened is important, but knowing it happened at a certain time of the day is also important for the user to be able to link that information with something that happened in the power system to be able to detect the source of the PQ event.

This block was implemented with the BQ32000 [48] IC from Texas Instruments. This IC counts from seconds to centuries, has an automatic switchover to backup supply, which is a coin-cell battery, a serial interface to provide communication with the processor, an oscillator fail flag, which allows the processor to determine the status of the RTC oscillator and a STOP bit that allows the processor to disable the oscillator, if necessary. Figure 23 shows a how this IC is routed on this project.

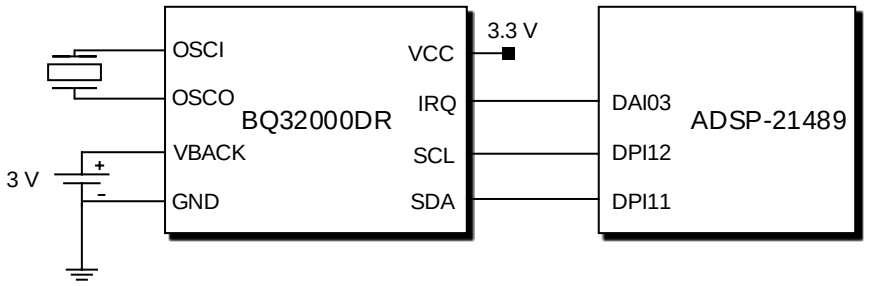

Figure 23: Schematic implementation of the RTC IC.

As it can be observed in Figure 23, this RTC has three lines connected to the processor: SDA, SCL and IRQ. These lines are used to communicate with the processor through a serial communication protocol known as I^2C. With this communication, a lot of interactions are available between RTC and processor, but only two are used in this project which are: reading the time and date and setting a new time and date.

These interactions consist of reading and writing registers within the RTC IC. Getting the time and date is done by reading the seconds, minutes, hours and so on registers of the RTC. The communication protocol to read and write registers is presented in Figure 24 and Figure 25, respectively.

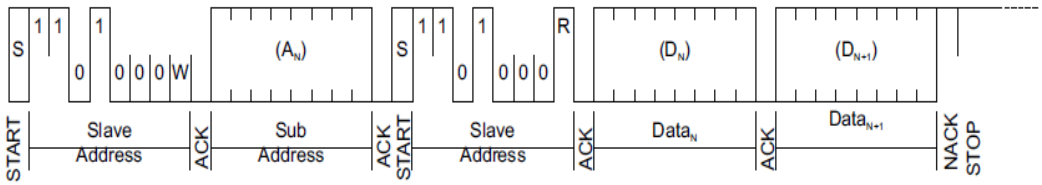

Figure 24: I^2C data read protocol.

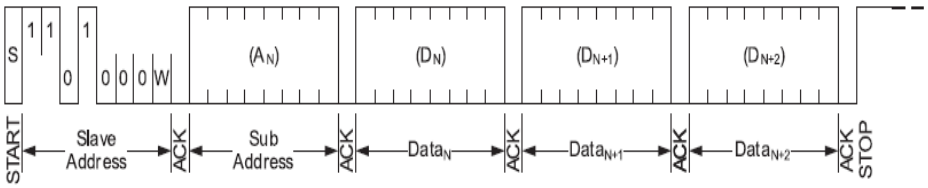

Figure 25: I^2C data write protocol.

The communication between processor and RTC is controlled by the master, which, in this case, is the processor. It is responsible to start and stop the communication and to generate the clock. The sequence to write and read registers starts in the same manner. First, a start bit is sent, followed by the slave address, which is constant and assumes the value 0x68, and the write bit, which in I^2C is 0. After that, the address of the first register to be read or written is sent. The reason as to why the address is of the first register, and not just the register of interest, is that every register is eight bits long. And, as it can be observed in Figure 24 and Figure 25, there is more than eight bits of data being

exchanged. This means that the writing and reading of registers is sequential and that the first eight bits of data received are, in fact, of the register of interest, and, if the communication is not stopped after those bits are received, the next eight bits of data will be of the next address register. Since the date and time registers have consecutive addresses (for example the seconds register has the address 0x00 and the minutes register has the address 0x01) this feature allows reading and setting the entire time and date with just one interaction, instead of multiple interactions. This interaction is stopped when a stop bit is sent.

The IRQ pin is useful when there's a need to keep something synchronized. For example, if the processor needs to be interrupted every one second, this IRQ pin can be set to produce a square wave with a one second period. In other RTCs, this pin is usually used when the IC can also work as an alarm. The processor sets a time and date to be interrupted and when that time comes, this pin becomes active and an interrupt is generated, effectively alerting the processor. In this project, that feature is not present nor it is necessary and this pin is used mostly for debug purposes, to make sure the communication between processor and RTC is working.

3.5 Ethernet communication

This system is able to communicate with a server through an Ethernet connection. Since the DSP doesn't have Ethernet connection, there is a need to use a module to make the communication possible. The use of the Ethernet connection brings many advantages like having multiple systems in different parts of the power system all connected to the same server, sending pre-processed data; the ability to connect to one of these systems remotely, since they're available in a LAN; and the ability to have longer cables than in other solutions e.g. USB.

This block is implemented with the use of the W5100 [49] IC from WIZnet. This IC is an Ethernet controller designed for embedded applications, with a fully hardwired TCP/IP stack and integrated Ethernet MAC and PHY. This TCP/IP stack supports many important protocols, such as TCP, IPv4, ICMP and ARP, in a way that is transparent for the user. This means that the user only has to worry about packet contents and not the headers for the various MAC layers present in the internet. It is also necessary to have some socket programming to create, manage and destroy connections between the processor and the internet. The use of this IC with an Ethernet port that has the magnetic circuit integrated makes the connection between the two simple and requiring no extra components. This IC is used in commercial applications such as the Arduino Ethernet shield, it's very well documented and has been a go-to reference in terms of Ethernet connectivity. The routing of this IC to the processor is shown in Figure 26. This IC has three possible ways to be connected to the processor: Direct, Indirect and SPI. In this project, the direct mode was used, taking advantage of the glueless memory interface provided by the ADSP-21489. This means that the serial communication alternative will not be used, in detriment of the parallel communication alternative. This alternative is the best for real-time applications, such as this project, because of the larger bandwith and overall speed of exchanging data. To compare, a write cycle of one word, or eight bits of data, takes a minimum of 70 ns using parallel communication. With serial communication, each bit of data takes a minimum of 7 ns of setup time and another 28 ns minimum of hold time. For every word transmitted

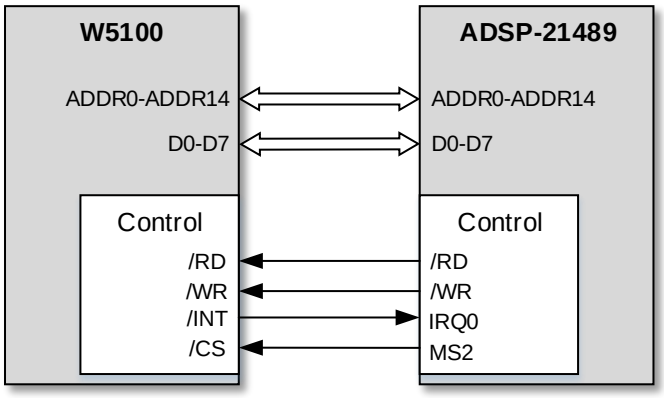

Figure 26: Schematic implementation of the Ethernet IC.

using parallel communication, two bits of data would be sent with serial communication. The disadvantage of this choice is the high number of connections between the processor and the IC. With serial communication, only four lines would be needed, while with parallel communication, twenty seven lines are required, with the direct mode. The indirect mode also uses parallel communication and reduces the number of address lines from the fifteen used in the direct mode to two. However, with this reduction in address lines comes an increase in complexity with address calculation, which isn't an appealing alternative. So, despite having a larger number of connections to the processor, the direct mode is the chosen one for this project, for its increased speed, when compared with the SPI, and easiness of use, when compared to the indirect mode.

The interaction between the processor and the IC is mostly done through writing and reading registers. For example, setting network information, such as IP address and gateway address, and socket management is done exclusively done by interacting with registers. But, reading and writing data that was received or that is meant to be sent, respectively, is done by reading and writing data in the W5100s buffers. However, the pointers to the position of the buffers where the data is present or where the data should be written to are also stored in registers. These registers are used in the same fashion that the RTC IC registers are: they all have constant addresses and to read or write to any given register, it is necessary to read or write to that address. The registers are eight bit long, the same as data bus width, but some of them have sixteen bit long information. Thus, to get the full information, it is necessary to read or write two consecutive addresses.

To access any given register, the first thing needed to do is to calculate the physical address. The address translation done in the ADSP-21489 allows for an easy way to interact with the W5100 registers. Knowing the IC is connected to a memory bank of the processor, the physical address of the register is

$$pRegAddr = baseBankAddr + baseICAddr + RegAddr, \quad (17)$$

where baseBankAddr is the base address of the memory bank used, baseICAddr is the base address of the W5100 and RegAddr is the address of the register. For example, the physical address of the

Interrupt Mask Register that will enable (or disable) the occurrence of interrupts coming from the various sockets is

$$0x8000000 + 0x8000 + 0x15 = 0x8008015 . \quad (18)$$

Now that the physical address is calculated, it can be used in the program to read that register and check which sockets are eligible to generate interrupt requests or write to that register to make one socket eligible (or not) to generate interrupt requests.

In order to function correctly, there are a couple of network related information, such as the IP address, Gateway IP address and Subnet mask that need to be filled before using this IC. On a debug level, the user can set this information himself to test on known conditions [50]. If the network information is filled up correctly, the chip will respond to pings automatically. But, in this project, this network information is not constant, because it isn't known where the system is going to be connected. There is a need to get this information in an automated fashion, because the system needs to have connectivity with the internet in every location possible. To do so, the DHCP client function was added and will be described in Appendix A [51]. There is also a need to allocate memory to each socket that will be used. The W5100 has a 16kB buffer, which is divided in two, a received data buffer and a send data buffer. Each of those buffers has 8kB of capacity and that capacity needs to be distributed along the four possible sockets. In this project, only two sockets are used: one for the DHCP client function, the other for the connection with the web server.

After the network information has been filled, it is possible to create connections with peers using sockets. These sockets can be UDP, TCP or IP Raw sockets, based on the protocol that they'll use. UDP, which stands for User Datagram Protocol, is a rather simple transport layer protocol. It is unreliable and gives no guarantees that the information gets to the destination, but it's faster than other protocols, since it doesn't wait for an acknowledge by the destination. There is also the possibility of sending one packet to one IP address and sending another right after to another IP address using the same socket. It's mostly used in video and audio streaming, where losing a couple of packets doesn't influence the user experience all that much. However, in applications where every packet is important, it shouldn't be used, since there is no guarantee that critical information is received correctly. TCP, which stands for Transmission Control Protocol, is another transport layer protocol, albeit more complex than the UDP. This protocol creates connections between two peers and every data packet is verified and acknowledged by the destination. If it isn't acknowledged in time, a timeout occurs and the packet is sent again. After a few timeouts, the connection is determined as dead. This connection between two peers is only terminated if one decides to terminate it or if the connection fails, which is detected by one of the peers when the packets aren't acknowledged. Even if there's no useful information flowing, the peers will keep sending keep-alive packets, to maintain their connection. Since every packet is acknowledged by the destination, there is a guarantee that the data is received correctly. This makes it the right protocol to be used in data sensitive applications, such as downloads and uploads of files. For these reasons, the TCP is the chosen protocol for this project. The IP Raw option allows the user to control every aspect of the packets used and is useful if there's a need to use a protocol that is not supported by the TCP/IP stack of the W5100. This isn't the case of this project.

Having decided on which protocol to use, there's a need to decide whether this system will be a TCP client or server. Since there's no need to have information flowing from the web server to the system, because it wouldn't add any value nor the system should be taking any requests, there's no point in it being a TCP server. Since the board is constantly searching for power quality events and there's only information to be sent when one is detected, the system should be the TCP client, because it can connect to the server only when there's useful information to send.

Another part of this project which is not directly involved with the system development is the web server that the system will communicate with. As it was already determined, this server will have to be a TCP server which the system will connect to when needed. More details about the web server will be given in Appendix A.

3.6 Memory module

Another block that is also important for this system is the memory block. This memory block is composed of two kinds of memories, volatile and non-volatile.

Volatile memory is computer memory that requires power to maintain the stored information, meaning that this kind of memory is only useful when the system is running. The DSP has only 5 Mbits of internal memory, which is not enough to process sampled data and to generate appropriate power quality event reports. So, there's a necessity to expand this memory space using external memory and, for this purpose, the MT48LC16M16A2 256 Mb SDRAM [52] from Micron was chosen. The choice for this IC was the easiest and safest for the design of this system, because it's the IC used by Analog Devices on their demonstration kit. This means that the circuit is already tested and working, there's a base of comparison if something is not working well with the IC in the system, there's already a fully working driver developed in the test programs given by Analog Devices and example codes to work with it. This IC takes advantage of the glueless connection provided by the processor, much like the Ethernet IC, meaning there's no need for any additional circuit components for it to function correctly. Figure 27 shows how the SDRAM is connected to the processor.

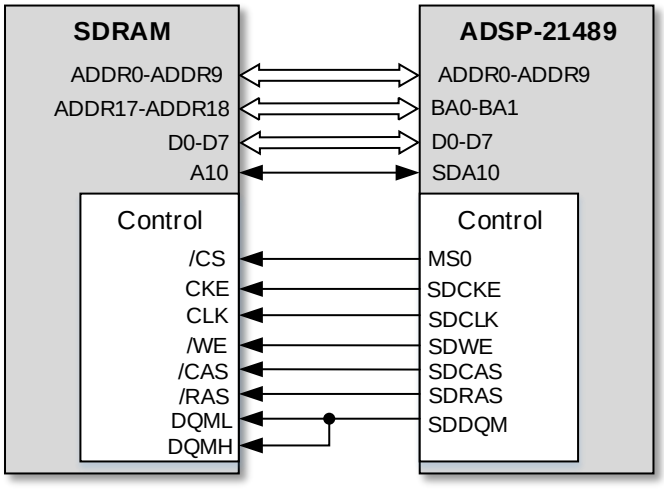

Figure 27: Schematic implementation of the SDRAM IC.

With this SDRAM present in the system, there's the possibility to store more data and to process everything safely, without the danger of overwriting data. There is also a possibility to have a bigger stack and heap, which would be allocated in this external memory, but, it was determined experimentally that it wasn't a good solution. This change would increase the execution time of all routines about ten times and, since this system is a real-time application, that isn't acceptable. But, even if it's used just to store data, mostly sampled and processed data, there's always a time penalty when using it. This penalty, however, can be reduced if the proper way of handling data is chosen, in other words, if DMA is used. It is more complex than getting or storing data using, for example, the memcpy routine, but it can reduce up to three times the time spent getting or storing data. Not only that, but since the DMA is a separate mechanism that doesn't need the processor to execute its job, it's possible to program the DMA to fetch a big chunk of data from the external memory, while the processor is running other useful code in parallel, which is important in time sensitive applications. When the DMA finishes the transfer, it generates an interrupt signaling the processor so it can process the newly received data, or just acknowledge that the data has been stored correctly. There's also the need to make sure that one DMA transfer is finished before another can start. This is pretty obvious, but can result in some waiting time if the DMA channel is already being used.

Non-volatile memory is computer memory that can get keep stored information even when not powered. This type of memory usually costs more and has a poorer performance, in terms of speed, than volatile memory. However, having this type of memory on this system is important due to the fact that it keeps the information even when the system isn't powered. In the case of a power outage, the system will still be powered by a battery, as it was described before, which will allow it to capture the possible PQ event that lead to the power outage. But, even when the battery runs out of energy and the system finally shuts down, the information regarding that PQ event, and others that couldn't be sent through Ethernet, will still be stored in this memory. It also allows gathering information when the system is not connected to any device via Ethernet, although it is not designed to do it, and will not have the same performance as with the Ethernet connection, as it will be clear further by the end of this chapter. It can be used as backup where most of the PQ event information is stored, in case the Ethernet connection suddenly fails or the server that it was connected to shuts down for some reason. All the information that is stored in the non-volatile memory can be accessed at a later time to determine the reasons for the system failure. There are many possible ways to incorporate non-volatile memory in this system, but the chosen one was the SD Card. The connection scheme of the SD Card is presented in Figure 28.

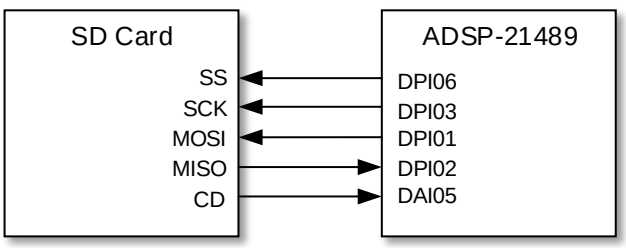

Figure 28: Interface between the DSP and the SD Card.

The CD pin shown in Figure 28 is not part of the data communication protocol of the SD card, but a pin that will assume a value of 3.3 V when the card is inserted and a value of 0 V when there is no card present in the connector. With this, it is possible for the DSP to determine if there is a card present in the system.

The SD Card was chosen in detriment of the other possible choices, not only for its advantages, but also taking in consideration the other choices already made for the rest of the system. Knowing that the Ethernet IC and the SDRAM would both use the parallel interface, it would be hard to integrate another NAND or NOR memory module in the board, at the design level, because of the high number of connections. The USB pen-drive and SD Card both use serial communication, which has a lower number of connections, which is better at design level, but worst at performance level, since this type of communication is relatively slower. But, the USB pen-drive option needs additional circuitry, basically a serial to USB communication translator, besides the connector, while the SD Card only needs the connector. Not only that, but the SD Card option can create text files in a removable media that can be opened in a computer, without the need for the system to be working. The NAND or NOR memory would certainly have a better performance, but could only be useful when the system was turned on again and the first thing it would do would be send that backup information through Ethernet. In the end, none of the choices were ideal, but the SD Card was the best choice overall.

The SPI standard defines the physical link of the communication protocol with the SD Card but not the complete data transfer protocol [53]. The use of the SPI protocol means that the processor will be the master in exchanges of data with the SD Card, being responsible for generating chip select and clock signals. Also, only the master can start the communication between the two, meaning the SD Card will never interrupt the code being executed. All SPI messages consist of command, response and data tokens. Figure 29 presents the example of a single block read operation.

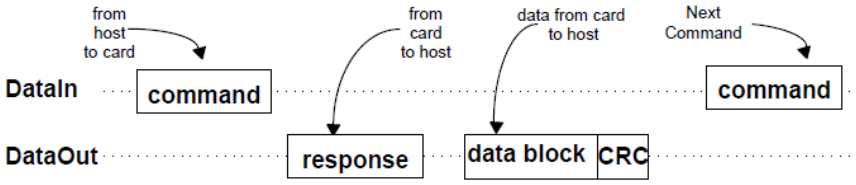

Figure 29: Single block read operation.

Upon reception of a valid read command, the card will respond with a response token followed by a data block. A valid data block is suffixed with a 16-bit CRC generated by the standard CCITT polynomial $x^{16} + x^{12} + x^5 + 1$. In the case of Standard Capacity Memory Cards, a data block can be as big as one card write block and as small as a single byte, while when using SDHC and SDXC cards, the block length is fixed to 512 bytes. So, partial block read/write operations are not allowed, which might prevent storing smaller data like the PQ event data packets in raw mode. Figure 30 presents the example of a single block write operation.

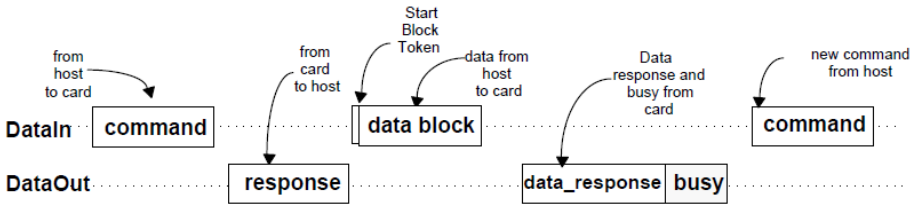

Figure 30: Single data write operation.

Upon reception of a valid write command, the card will respond with a response token and will wait for a data block to be sent from the host. CRC suffix, block length and start address restrictions are identical to the read operation. Every data block needs to have a prefix of Start Block token. After the data block has been received, the SD Card will respond with a data response token. If the data has been received successfully, it will start saving it in internal memory and it will continuously send busy tokens to the host (which is done by holding the DataOut line low). After the data has been saved to internal memory it will stop sending the busy tokens and the host should check the result of that process by sending the SEND_STATUS command. Some errors like address out of range or write protect violation are only detected when the SD Card attempts to save the data in its internal memory. The validation check done by the SD Card on the received data blocks is the CRC and general Write Error indication.

These two examples were of single block transfers and were shown to illustrate the data communication protocol with the SD Card. Usually, in applications with the SD Card, multiple data block reads and writes are used. The difference is the command used to start the data exchanged and the additional data stop block sent by the host.

When the SD Card is powered on, it is in the SD mode. This is another mode that the SD Card has available, which has a better performance than the SPI mode, but has more connections and a more complex interface. In order to change it to SPI mode, it is necessary to send a reset command (CMD0) while asserting the chip select line. The full SPI mode initialization flow is represented in Figure 31. After changing to the SPI mode, the host sends the SEND_IF_COND command, which will verify the SD Memory Card interface condition. With this command, the host informs the SD Card of its supply voltage information and asks the card whether it supports that voltage level or not. If the card responds with an illegal command, it is either a legacy card or not an SD Card at all. If the card does support CMD8 and can operate on the supplied voltage, the response echoes back the supply voltage that was set in the command argument. If one of the data token bits returns with 0, the card cannot operate on the supplied voltage. After the CMD8, the READ_OCR command is sent by the host.

The OCR is a register within the SD Card which holds the information about the voltage ranges the SD Card can work with. It also has other status bits, such as the CCS, which is the Card Capacity Status, useful to identify if the card is a SDSC card or SDHC/SDXC card. The reason why it is important to differentiate the different kinds of cards is because the SDSC cards have a capacity

range between 128 MB to 2 GB and use, by default, the FAT16 filesystem, while SDHC cards offer up to 32 GB of capacity and use, by default, the FAT32 filesystem. The SDXC cards have capacities between 64 GB and 2 TB and use, by default, the exFAT filesystem. These are just some of the differences, since the more recent and higher capacity cards also offer higher data rates.

After receiving the response of the READ_OCR command, if the host does not accept the voltage range of the SD Card it shall not proceed further with the initialization sequence. The next step is to send the SD_SEND_OP_COND command (ACMD41) to start initialization and to check if the card has completed initialization. If the host is ready to interface with a SDHC or SDXC memory card, it should set the HCS bit in the argument of ACMD41 to 1. At this point, the host will repeatedly send the ACMD41, waiting for the 'in_idle_state' bit of the response token to assume the value 0. When this bit has the value 1, it means that the card is initializing. When the response to the ACMD41 has this bit with the value 0, the card is initialized and ready to be interacted with. If the card has a 2.0 or later version, the READ_OCR command should be sent again, to check the value of the CCS bit. As it was stated before, the value of this bit will help identify the type of card present in the system, and along with it the capacity and data rate it can assume.

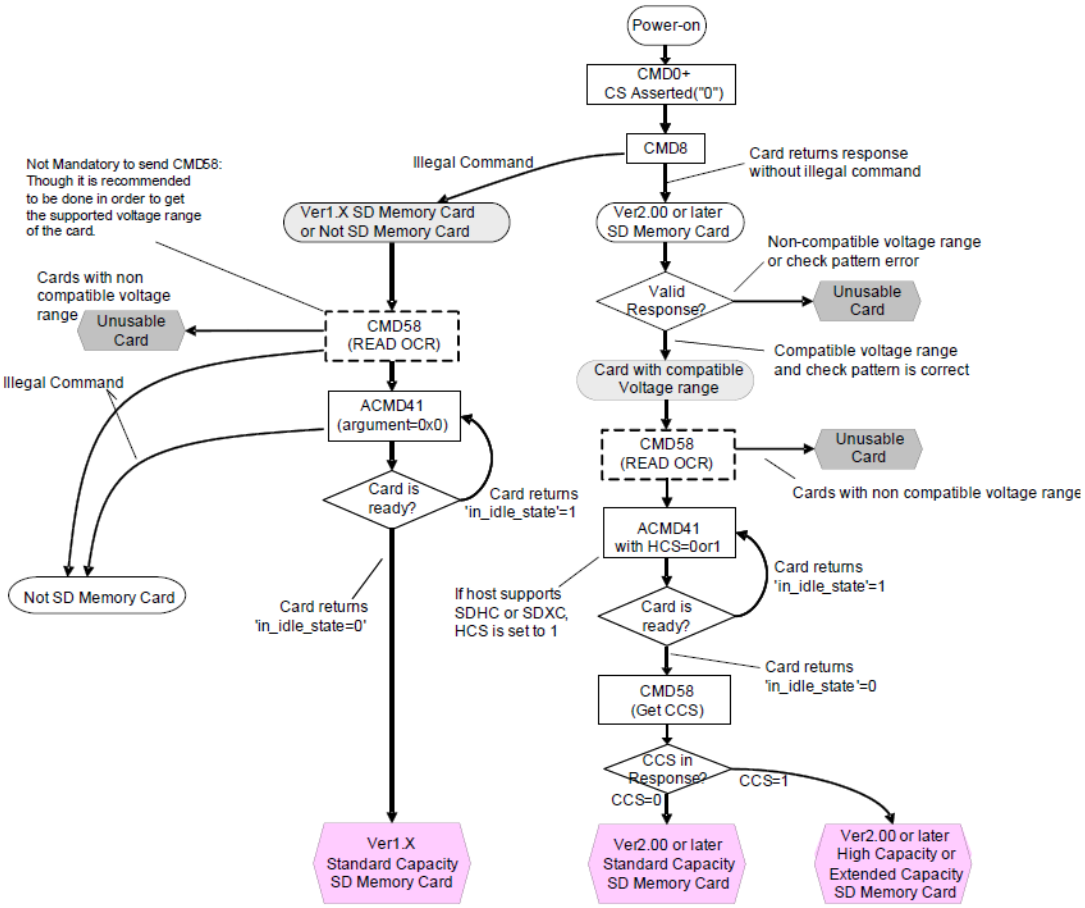

Figure 31: SD Card initialization in SPI mode flowchart **[53]**.

After the initialization of the SD memory card has been completed, it can finally be interacted with, being it to store data or retrieve previously written data. But, the data cannot be stored as individual blocks of bytes, otherwise the computer, or any other device that is able to interface with the memory card, will not be able to read the data correctly, since these devices don't read individual blocks of memory space. An abstraction layer is then needed, above this memory space, in order for all these devices to understand what is stored in the memory card. That abstraction layer is known as the filesystem that divides the memory space into clusters and sectors, which can be interpreted to determine which memory blocks hold information relative to a certain file. If the data is stored in a raw manner, as in just storing the bytes that resulted from the signal analysis into the SD card memory, it can only be read by the system that stored that information, or a similar embedded system that is prepared to retrieve data in the same raw manner. If the card is inserted into a computer or any other device that has an operating system and drivers prepared to handle a filesystem, a message will appear saying that the data is corrupted and it will ask the user if he wants to format the card. Since the aim of this project is to create text files that the user can consult when using the SD card, there is a need to implement this file system.

The FAT (File Allocation Table) file system works with, as the name says, allocation tables. The memory space is divided in clusters and sectors, where a cluster is made of several sectors. The number of bytes in each sector and cluster depend on the version of the file system used and the size of the storage device used. For example, for a storage device with 64 MB to 128 MB, FAT16 uses 2 KB clusters, while the FAT32 file system uses clusters with 1 KB of capacity. These clusters are the smallest memory allocation unit for a file or directory. Meaning that if a 512 B file is stored in a storage device with a cluster size of 1 kB, only one cluster will be used to hold the file, but the rest of the 512 B of the cluster cannot be used to store any other file, leading to unused memory space. So, the file system is more efficient, memory wise, with smaller capacity clusters. The way files and directories are stored in memory is rather simple: the root directory data cluster holds the information about the location of every starting cluster of the files it contains. So, to access a file contained in the root directory, the root directory cluster needs to be accessed to get the number of the starting cluster of the file. Once that address is obtained, the file retrieval can begin. To retrieve the file data, all the clusters allocated for that file need to be read, in order. The first data cluster holds not only the data itself, but also the pointer for the next cluster where the data continues. The data clusters need to be read continuously until an end-of-file indicator is read. Figure 32 shows an example of how three files can be read.

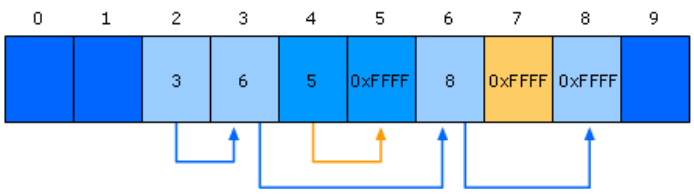

Figure 32: File reading example.

In this example, ten clusters of memory, numbered from 0 to 9, contain three files. One file occupies the clusters 2, 3, 6 and 8, another file occupies the clusters 4 and 5 and the final file occupies the cluster number 7. The clusters 0, 1 and 9 represent free memory space. The first thing to notice is that the files aren't always stored sequentially, so even after knowing the starting cluster of the file, there's a need to search that cluster for the information about the next cluster where the file is held and not read the memory space continuously until reaching an end-of-file indicator. So, to read the first file, the data from the first cluster needs to be retrieved. In that cluster, there's no end-of-file indicator, but there's a pointer to the next data cluster, which is the cluster number 3, meaning it is necessary to retrieve another data cluster to get the full data related with that file. After reading the cluster number 3 and seeing that there's still no indicator that the file is finished, it would be necessary to retrieve the next data cluster, which is the number 6 and so on. Another example, to update a file, as in adding data to that file, it's necessary to first check if the last data cluster has enough memory to hold the new data. If it does, the data is simply added, maintaining the indicator that the file ends in that cluster. But, if it doesn't have enough memory allocated, the directory data cluster needs to be consulted again to retrieve the address of the next available cluster. Once that is figured out, the end-of-file indicator is removed from the current cluster, being replaced with the address for the next data cluster that will be used to store the additional data. In conclusion, the file system adds complexity to the way data is stored and handled by the devices that interact with it, but makes it "understandable" for all of them.

The implementation of the file system was done with an open source code, with the name FatFS [54], developed as a personal project by ChaN [55]. FatFs is a generic file system module developed for small embedded systems. This module is written in C and completely separated from the disk I/O layer. This makes it completely independent of the platform it is used on, because every platform will recognize the programming language. It has a lot of features, being the most important ones the easy portability, the compatibility with Windows FAT file system and the small footprint for code and work area. This means that this module can be used to interface with SD cards, hard drives and any other type of storage devices which use the FAT file system. This module is the necessary abstraction layer for this project, because it handles the memory interactions discussed previously of clusters and sectors, taking that responsibility out of the programmers hands. This abstraction layer sits between the application level and low level memory drivers, such as the individual block read and write operations that were discussed earlier. It provides high level routines, such as f_open which will open an existing file or create and open a new file, f_write to write data on that file and f_close to close the files. These routines will handle the calculation of cluster addresses, manage end-of-file indicators and other necessary information automatically. To port this module to the project, only the low level memory drivers needed to be modified. Those drivers are the initialization driver, the memory write and read drivers and other control drivers that were discussed previously in this chapter. After this modification, the system has become capable of creating, editing and removing files and directories.

The use of the SD Card memory in this project is limited, since the card can only be interacted when the sampling and processing aren't active. The reason is the slow performance of the serial communication protocol used to interface with the card and the additional complexity of having a file

system. Writing data to a file takes too much time and cannot be done via DMA, since the data isn't flowing in just one direction. There's always a need to check the response coming from the SD Card to make sure the data was received correctly. It can and is used, however, when the system boots, to change valuable information such as the web server IP, the equipment identifier and the server password, without changing source code, and to store unsent data packets when a power outage is detected.

4. Software

In this chapter the contents of the project software, namely the program developed for the ADSP-21489 to detect and classify power quality events is discussed.

4.1 Power Quality Event Detector

There are different ways to detect power quality events. Some implementations use RMS calculations, other use mathematical transforms to represent the signal in the frequency domain and analyze its components to derive the power quality event. One thing to keep in mind is that there are many different types of power quality events and they're not easily detected using just one method. For example, RMS calculations give enough information to detect voltage sags, swells and interruptions, but aren't enough to detect fast transients, because these don't affect significantly the RMS value of the signal. On the other hand, mathematical transforms to represent the signal in the frequency domain excel in detecting power quality events that affect its harmonic content. However, they cannot detect efficiently some power quality events that affect the RMS value of the signal. If these occur gradually, there won't be any significant response in the frequency analysis of the signal.

With these considerations in mind, the proposed method for detecting and classifying power quality events is shown in Figure 33.

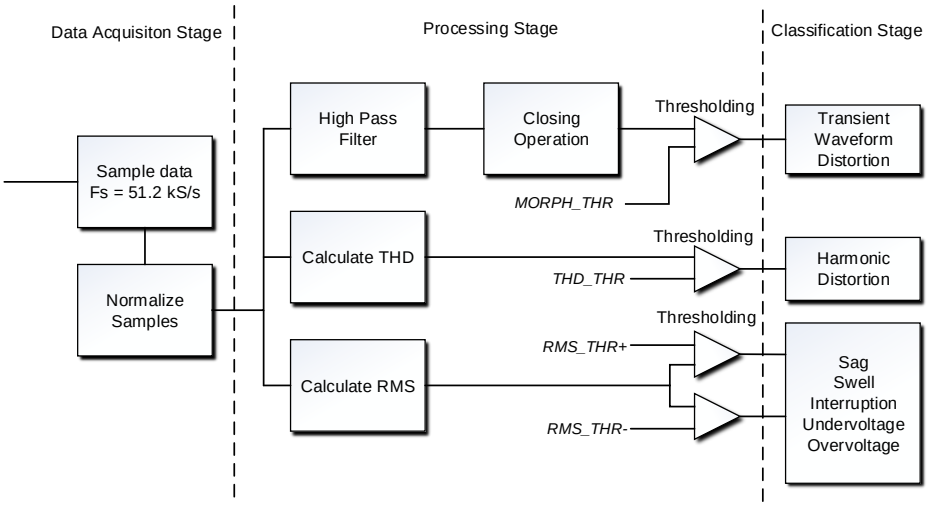

Figure 33: Detection and classification of power quality events method.

4.2 Data Acquisition Stage

As it's clear to see, the proposed method uses both time and frequency analysis of the signal to detect and classify power quality events. The time-domain analysis of the signal, which are the calculation of the RMS value and the high-pass filter, along with the closing morphological operation, are used to detect voltage sags, swells, interruptions, transients and other waveform distortion events. The frequency-domain analysis of the signal, in other words, the THD calculation is used to detect harmonic distortions.

The first step of this method is to sample the various input signals, which are the three phases of the three-phase power system. To do this, a PWM (Pulse Width Modulation) signal is generated, with a frequency of 51.2 kHz, connected to the CONV pin of the ADCs. The ADC conversion will start on the rising edge of this signal and the samples will be acquired while the pulse assumes the logical value high. The sampling frequency assumes the value of 51.2 kHz due to limitations with processing algorithms, which will be explained later in this chapter. The period while the pulse assumes the logical value low corresponds to the minimum time interval the ADCs need to initiate a new sample acquisition. When the period of this signal ends, a new sampling interval begins. The PWM signal was chosen as the CONV signal for the ADCs because it guaranteed that the signal frequency was the chosen and because it would make it independent of the execution of the rest of the program. If the program needed to reach a certain area of the code to initiate a new sampling interval, the sampling frequency would be compromised, since the code would not always take the exact same time to reach the next sampling interval. Instead, the PWM option is configured when the system starts, and it's constantly outputting that signal, with the chosen frequency, without core intervention. The SPI communication protocol is also automated: whenever 16 bits of data are received by the SPI mechanism, they're stored in an internal buffer and an interrupt is generated to the core. The core serves that interruption to read the contents of that buffer to internal memory, clearing the SPI interrupt. This means that the core is being constantly interrupted to receive the samples when they're ready to be transferred to internal memory, but if the interrupt time is small, there are no problems with the execution of the remaining code, which is mainly the processing of those samples. Not only that, but the whole sampling mechanism is automated and independent of code execution, which is a key factor to maintain a constant sampling frequency and not miss any samples, keeping a continuous acquisition.

Once the samples are transferred to internal memory, they will be translated to the real voltage value they had in the input. This is an important stop where the small differences between the ideal voltage signals and the acquired signals are corrected. Even though the three input signals are being sampled with three similar circuits, there are always differences between the circuit components. Low tolerance resistors were used to reduce this difference, but some corrections are still needed on the software level.

The equation to transform the digital word from the ADC into the real voltage value of the input signal is

$$V_{real} = \left(\frac{65535}{3.3}\right) \times \left(\left(sample\_value - DC\_OFFSET \right) \times AMP\_MOD \right), \tag{19}$$

where, the DC_OFFSET is the value to turn the signal into a bipolar signal again and the AMP_MOD value is to restore the amplitude of the signal. According to (13), the DC_OFFSET value should be 1.65 and the AMP_MOD value should be

$$AMP\_MOD = 2 \times \frac{1005}{5} = 402 . \tag{20}$$

In theory, these should be the values to correctly restore the real voltage value coming from the input signals. In reality, the experimental acquisition of samples demonstrated that some corrections need

to be made to DC_OFFSET and AMP_MOD values. To determine the correction of the DC_OFFSET value, a 0 V signal was placed on the input of the three conditioning and acquisition circuits and 100 samples were acquired with the ADCs. After that, the mean value of those samples was calculated and is used to determine the correct the DC offset of the three phase's signals. After correcting this offset, the amplitude modifier was also calculated. A signal with known amplitude was used at the input of the conditioning and acquisition circuits and that signal was sampled for long enough to detect the maximum and minimum of the input signal. After that, the new amplitude modifier was calculated that would match the maximum and minimum of the input signal with the maximum and minimum of the real voltage values. Table 4 shows the different calibration values used for each of the three phases.

Table 4: Values used to correct differences in the phase circuits.

	Ideal	Phase 1	Phase 2	Phase 3
DC_OFFSET	1.65	1.65141	1.68099	1.78637
AMP_MOD	402	413.379	403.075	407.215

After the real voltage value has been restored with the samples, they're normalized so that every calculation is made in pu. This ends the pre-processing stage.

Once normalized, the samples are stored in a buffer which can hold up to 1024 samples which is the number of samples per period

$$\text{num\_samples/period} = \frac{\text{fs}}{\text{f}} = \frac{51200}{50} = 1024 \ . \tag{21}$$

Whenever this buffer is full, those samples can be processed. The first thing to do is store those samples in a bigger buffer. This is because the sampling mechanism is using two buffers in a ping-pong fashion: while the buffer A isn't full, it fills the buffer A; when the buffer A is full, it begins filling buffer B, while the samples stored in buffer A are being processed. It's obvious that the sampling mechanism has to work like this, because if just one buffer was used to store the always incoming samples, the processing window would be as tiny as the time between the buffer filling up and the first sample arriving, because that sample would overwrite a previous sample and the information would be lost. With the ping-pong storage mechanism, the processing time window is now as large as one period of the input signal. One other thing to keep in mind is that there are three input signals, coming from the three phases of the power system, and they aren't in phase. This means that when one buffer of one phase fills up, the other buffers for the other phases aren't full yet, but they will be before the next period of the input signal is sampled. It becomes clear the calculation of the processing window and knowing how much time the processing has to complete is vital to avoid losing information of any of the phases. To calculate the processing window, an assumption is made: the difference of phases will always be 120º or some value close to that. This assumption is made to benchmark the processing time of the processing algorithms and make sure there are no delays. For example, the ideal scenario would be that the processing of phase A samples is finished before the processing of phase B samples starts. If this isn't guaranteed, the delay introduced by the processing algorithm

would accumulate and at a certain point data would start to be lost. Guaranteeing that the processing algorithms don't fill the processing window also allows sending event reports between phase processing. But the difference of phases isn't necessarily restrictive, since the program is designed in a way that the samples are processed as they're available. For example, if the processing of the phase A samples takes so long that the phase B samples are present already, the processing of the phase B samples will begin immediately after the phase A samples are processed. However, the processing of phase B samples must finish before phase C samples are ready and they must finish before the phase A samples are ready again. The processing window is given by the time, or the number of cycles, between one phase has finished sampling an entire period and the next phase will finish sampling an entire period. Assuming the difference of phases being equal to 120º, the zeroes of the phases will be 60º apart. With a frequency of 50 Hz, the processing window is

$$\text{proc\_window} = \frac{60\text{x}10^{-2}}{180} = 3.33 \text{ ms} \ . \tag{22}$$

Knowing that the DSP works at a 400 MHz frequency, one cycle is

$$\text{cycle} = \frac{1}{400000000} = 2.5 \text{ ns} \tag{23}$$

and so, the processing window in processing cycles is

$$\text{num\_cycles} = \frac{3.3 \times 10^{-3}}{2.5 \times 10^{-9}} = 1.32 \text{ Mcycles} \ . \tag{24}$$

Table 5 shows the values for the processing window in time and number of cycles.

Table 5: Duration of the processing window, between phases, in time and number of cycles.

Frequency (Hz)	Time (ms)	Number of Cycles (M cycles)
50	3.33	1.32
52	3.2	1.28
47	3.55	1.42

The true limitation of the processing window is the period of the input signal. As long as the processing of the samples plus the time to send event reports, if there are any to send, takes less time than the period of the input signal, then no sample will be lost. This means that the difference of phases could even be 0º and that all samples would be ready at the beginning of the processing window. The samples would be processed immediately one phase after the other and the event reports would only be sent after all the processing. Table 6 shows the values the processing window can assume in total.

Table 6: Total duration of the processing window, in time and number of cycles.

Frequency (Hz)	Time (ms)	Number of Cycles (M cycles)
50	20	8
52	19.2	7.69
47	21.3	8.51

After determining the size of the processing windows, the algorithms can be developed, making sure the sum of processing time of those algorithms does not exceed the processing window.

4.3 RMS Calculation and Detection Algorithms

The first algorithm is the RMS calculation algorithm. Since the RMS calculation is rather common, there's a library function of the DSP [56] that calculates this value, which only needs a buffer with the sampled data and the size of that buffer as inputs. The output of that function is given by

$$c = \sqrt{\frac{\sum_{i=0}^{n-1} a_i^2}{n}}\ , \tag{25}$$

where a_i is the sample and n the number of samples. The RMS is calculated over one period of the power system's voltage and refreshed every half-period. After being calculated, this value is sent to a routine that will determine if a power quality event is present. A flowchart of the algorithm used to detect the power quality event is shown in Figure 34.

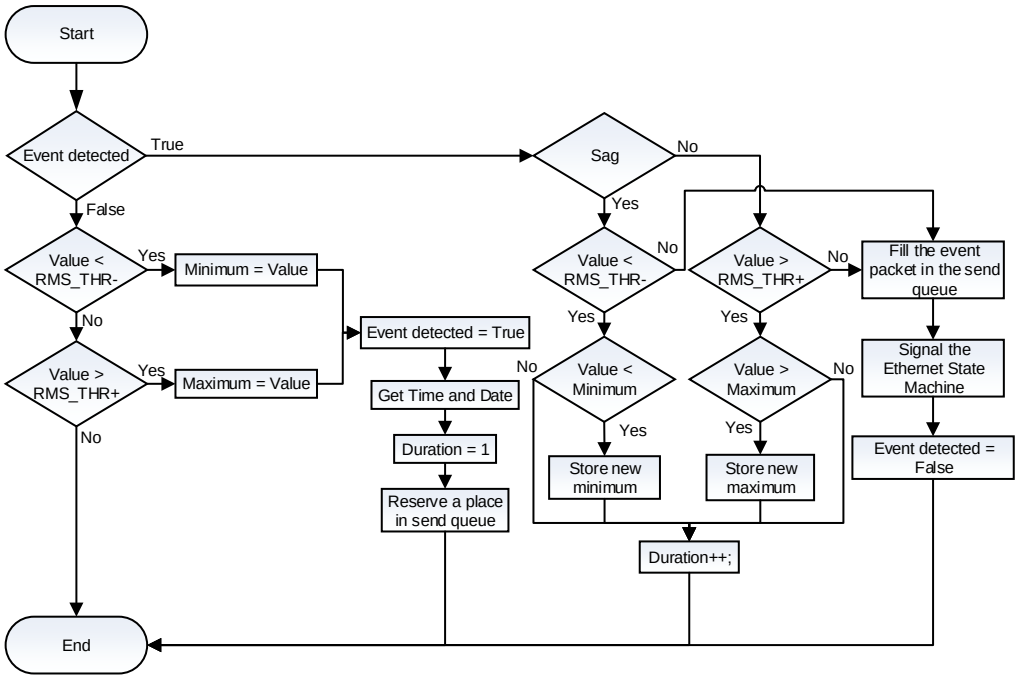

Figure 34: Flowchart of a simplified version of the RMS event detection algorithm.

The first thing this algorithm does is to compare the calculated RMS value with the pre-determined threshold values. RMS_THR+ is the maximum value the calculated value can assume without being considered a power quality event. In the same manner, the RMS_THR- is the minimum value the calculated RMS value can assume. If the calculated RMS value is higher than RMS_THR+ or lower than RMS_THR- there's a power quality event occurring. In that case, an event flag is activated and the event starts to be described, with the event type being determined (either swell or sag), the duration being set to one half period and the value being recorded as the maximum or

minimum RMS value calculated. After that, the algorithm exits, waiting for the next RMS calculation. When it happens, the algorithm checks the event flag, to see if an event had started in a prior calculation. If an event has been detected, it compares the new calculated RMS value with the threshold of that event. For example, if a swell was detected, it will compare the new value with the RMS_THR+ value. If the new RMS value calculated is still above that threshold, the algorithm tests if this value is greater than the previous maximum RMS value calculated. This makes sure that, when describing the event, the maximum RMS value that was calculated is presented. If it's higher, it is stored as the new maximum and the duration of the event is also updated. In the case of a sag event, storing the minimum value is not only useful in terms of information for the user, but it can also lead to a change of event type. If it falls down below 0.1 pu the event type needs to be changed to "Interruption". If several consecutive RMS calculations are above the threshold, the duration value will increase and if it exceeds 1 minute, the event type needs to be updated to "Overvoltage". In the case of a sag event, if the duration exceeds 1 minute the event type is updated to "Undervoltage". If the new calculated value is below the threshold value, then the event packet would be filled with all the information gathered: type, duration, magnitude, time and date. After that, an indication is sent to the Ethernet routine stating that that packet is ready to be sent and the search for a new event starts. After being fed to the algorithm, this value is then stored in a buffer that keeps a record of recent RMS values calculated.

To test this detection and classification algorithm, a simulated signal containing a RMS power quality event, made in Matlab, was fed to the system. This signal has an amplitude of 162.5 V, meaning it will trigger a sag event with a minimum value of 0.5 pu and it will affect three periods of the input signal. This event will only be applied to one of the input phases, making the power quality event a single phase sag. Figure 35 shows the value of the samples captured by the DSP and Figure 36 shows the RMS value variation over the period of those samples.

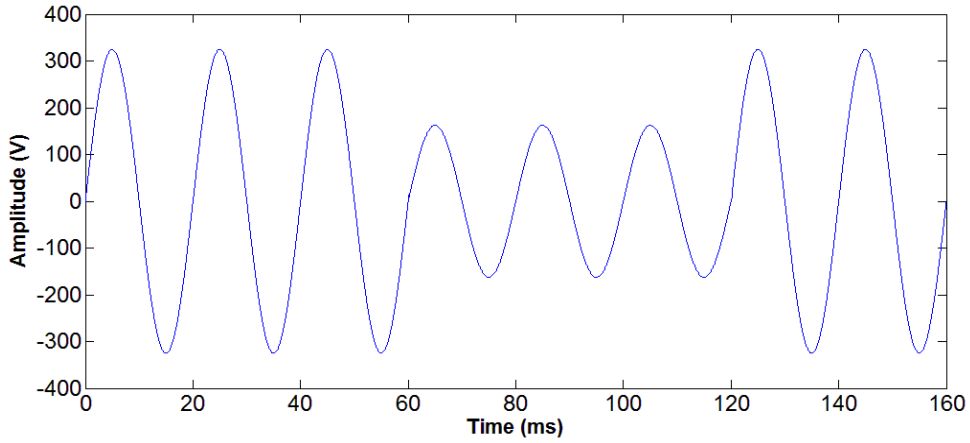

Figure 35: Sag event captured by the system.

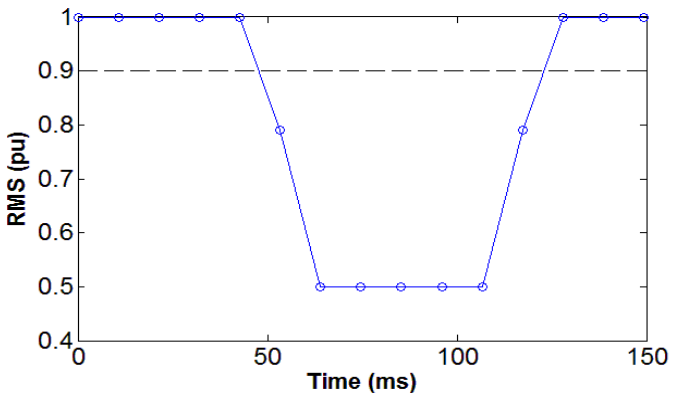

Figure 36: Variation of the RMS value over the eight periods shown in Figure 35.

As shown, the DSP is calculating a RMS value below the threshold, which will be sent to the detection algorithm that will describe the event. In Figure 36 is possible to observe that seven calculations of the RMS value are below the threshold and this will indicate the duration of the event. Five of those have the same values, being the minimum calculated value at approximately 0.5 pu. As soon as the first value is calculated below the threshold, the timestamp is acquired, so that the system can link a time reference to the beginning of the event and, since this disturbance was only applied to one phase, it's going to be described as single phase sag. Once the event is over and described, a packet will be sent to the web server to be stored in the power quality events database. A single phase event will also trigger another power quality event, which is the voltage imbalance event, since that phase will have a different RMS value than the other phases. Once the relative difference of RMS values is greater than 2%, the voltage imbalance event is triggered.

Also, when this disturbance is applied to two or even the three phases, the detected event will be classified as multi-phase sag or swell and the minimum RMS value presented will be the minimum RMS calculated in any of the three phases.

In practice, to test this algorithm, a voltage interruption was artificially caused in one of the phases, by disconnecting it from the power system. This action, though, often generated more than one power quality event: when disconnecting or connecting the plug to the power system, a transient occurred and was caught by the high-pass filter; also, it was common for the system to detect a THD event along with the transient. Figure 37 shows the captured interruption event.

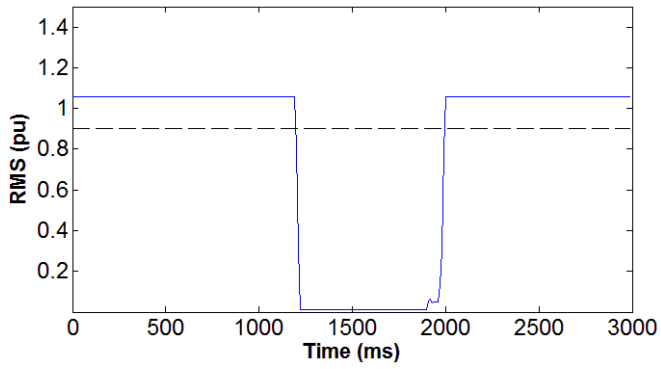

Figure 37: Interruption event registered by the developed system.

After all the algorithms were developed and tested, their execution time was determined. This is important because of the processing windows discussed earlier. Table 7 shows the processing time of the RMS related algorithms in several scenarios.

Table 7: Execution time for the RMS algorithms in all scenarios.

	Best-case scenario	Average	Worst-case scenario
Number of cycles	55447	55511	427280
Time (ms)	0.139	0.139	1.068

4.4 THD Calculation and Detection Algorithms

Another part of the processing stage is the calculation of the THD of the input signal. This calculation is the cause of the somewhat strange value for the sampling frequency. One of the objectives for this project was to have a sampling frequency over 50 kHz. So, the initial experimental trials were done using this value as the sampling frequency. But, once the THD had to be calculated, this figure needed to be readjusted. The reason is that the FFT is a function that is part of a library already given by Analog Devices and should be used because it is written in a way that takes advantage of the hardware architecture of the processor. But, the number of samples used to calculate this FFT is fixed and it must be a power of two. This means the number of input samples for the FFT calculation can only assume values which are a power of two between 8 and 65535. The calculation of the frequency bins resolution is given by

$$\Delta f = \frac{fs}{N} \ , \tag{26}$$

where fs is the sampling frequency and N the number of points of the FFT. If the 50 kHz value is used for the sampling frequency, the frequency resolution of the bins assumes a value that isn't an integer multiple of the center frequency of the input signal, nor an integer multiple of any harmonic of the input signal, leading to spectral leakage when the FFT is calculated. So, to fix this, the sampling frequency needed to be readjusted to a value that was an integer multiple of the 50 Hz frequency of the input signal, while having a frequency resolution that is also an integer multiple of the 50 Hz frequency and its harmonics. Using 4096 points and the 50 kHz sampling frequency, the frequency bin value is given by

$$\Delta f = \frac{50000}{4096} = 12.207 \text{ Hz} \ . \tag{27}$$

This value isn't an integer multiple of the 50 Hz, so the closest value that is an integer multiple is assumed. So, the new value of the bin resolution is given by

$$\Delta f = 12.5 \text{ Hz} \ . \tag{28}$$

With (23) and (25) it's possible to calculate the new sampling frequency, which is given by

$$\Delta f = \frac{fs}{N} \Rightarrow 12.5 = \frac{fs}{4096} \Rightarrow fs = 51200 \text{ Hz} \ . \quad (29)$$

The THD value is calculated using data of eight periods of the input signal, which would amount to a total of 8192 data samples but the FFT function can only take 4096 points as input. This function had to be used because there wasn't enough internal memory to allocate bigger buffers and to reduce computing time. So, instead of using 1024 data samples per period, this value was reduced to 512 data samples per period, effectively lowering the sampling frequency to half of its value, but just for the FFT/THD calculations. This is an acceptable choice, because, according to the EN 50160, only the first 40 harmonics of the power grid signal actually matter. The frequency of the 40th harmonic of the input signal is 2000 Hz. Even with 25.6 kHz of sampling frequency, the 40th can still be sampled correctly since it's below the Nyquist frequency. The THD calculation is done using (2).

After calculating the THD value, the result is fed into an algorithm to detect if a power quality event is occurring. This algorithm is a simpler version of the algorithm used for the RMS case. There's only one threshold to compare the value with, meaning there's only a power quality event when this value is above THD_THR. This threshold was set to 8%. The flowchart of this algorithm is shown in Figure 38.

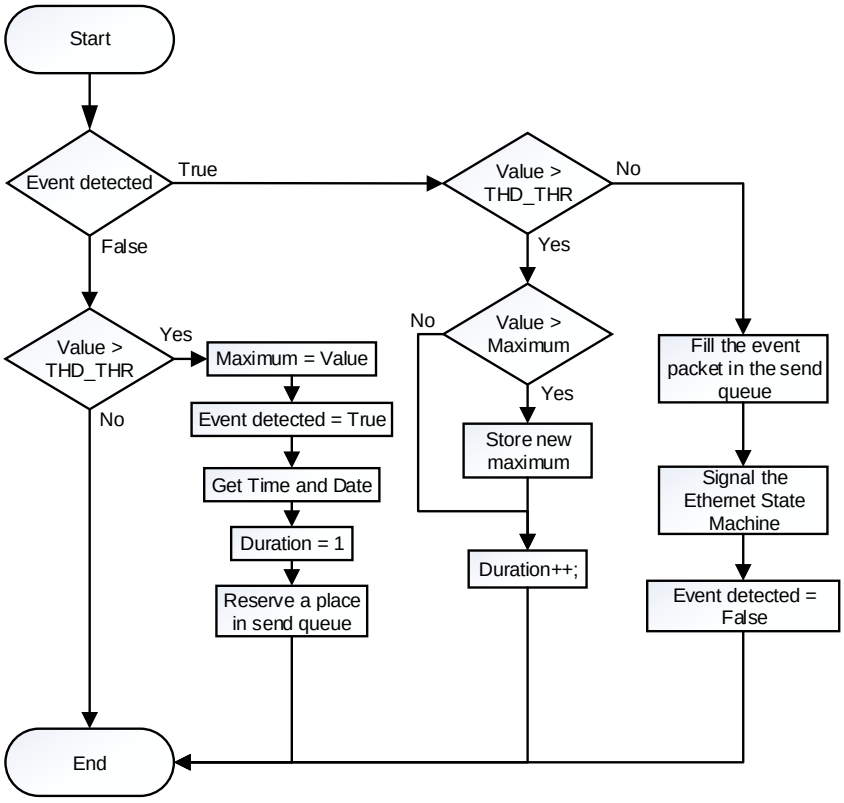

Figure 38: Flowchart of a simplified version of the THD event detection algorithm.

Once the result is over the threshold, an event flag is set up. A place in the queue of event descriptors is taken and the event type, along with the time and date is filled in that descriptor. The duration of the event is set to one and the value is stored being the maximum value calculated that was above the threshold. After that, the algorithm exits and waits for the next THD calculation. At that time, only two scenarios are possible: if the value is below the threshold, the event is finished and so

the rest of the event descriptor is filled with duration and maximum THD calculated value and a notification is sent to the Ethernet to inform that there's a package to transmit; if the value is above the threshold, the duration value is updated and the current THD value is compared with the maximum value calculated so far to update that information.

An interesting result to look at in this stage is the FFT result of the input signal. Figure 39 shows the FFT result of a real sampled signal. This signal has only one significant component, which is the 50 Hz component. This is expected since the signal coming from the power grid is a sinusoid with 50 Hz of frequency. However, when compared with a simulated signal, the real signal has higher values in the harmonic components leading to average THD values of around 1.5%. Simulated signals, on the other hand, had average THD values of 0.05%.

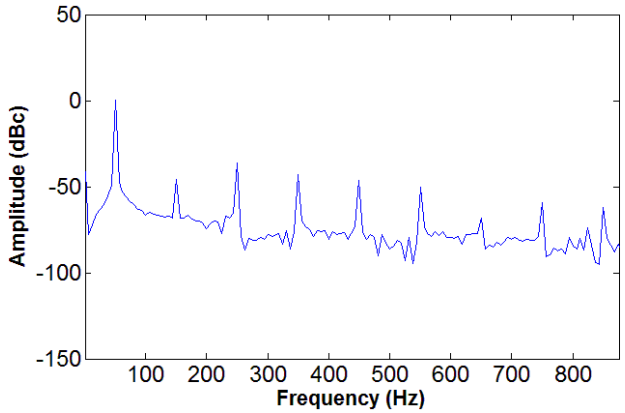

Figure 39: Spectrum of a real signal acquired by the system.

Like it was done before for the RMS algorithms, the execution time of the THD related algorithms needs to be determined. Table 8 presents the calculated execution times of these algorithms, again in several scenarios.

Table 8: Execution time for the THD related algorithms in all scenarios.

	Best-case scenario	Average	Worst-case scenario
Number of cycles	256923	257235	627886
Time (ms)	0.642	0.643	1.570

4.5 High-pass filter and Closing operation

Another part of the processing stage is the high-pass filter and the morphological closing operation [57]. The high-pass filter is used to filter out the fundamental frequency of the input signal, leaving the rest of the signal untouched. The filter output is then used to detect power quality events that have a frequency above the fundamental frequency of the input signal, such as transients and other waveform distortions.

The chosen filter is a 10^{th} order IIR Chebyshev type-II. It has a flat passband, which means that none of the harmonics of the input signal will suffer any attenuation or amplification. It has an

attenuation of 80 dB at the fundamental frequency of the input signal, effectively removing this component of the signal, with no attenuation beyond 90 Hz. It requires less computational power than an equivalent Butterworth filter, having just five biquad sections with these specifications. The coefficients of this filter were calculated with fdatool [58], an extension of Matlab to design and analyze digital filters.

The output of this high-pass filter is enough to detect events, setting a threshold and comparing the output of the filter with that threshold, much like the way RMS and THD events are detected. However, in the case of the filter and high frequency events, these may have properties that would lead to errors in detection. For example, a high frequency oscillatory transient could cross the threshold multiple times, leading to an incorrect detection of several events which are all part of the same event. Figure 40 shows an example of this situation.

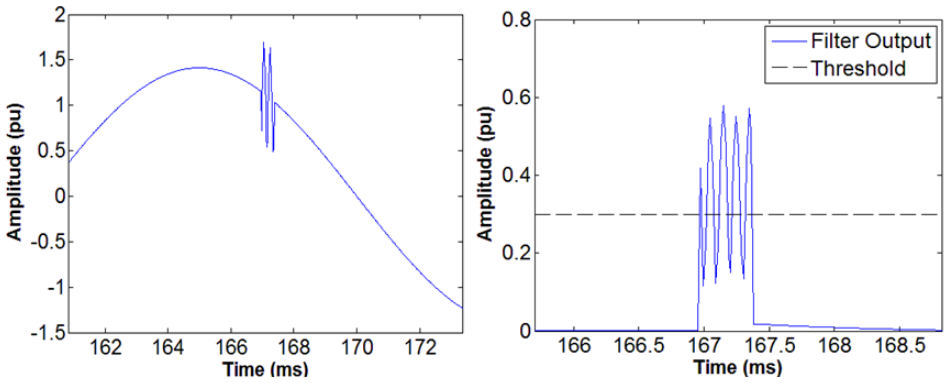

Figure 40: Example of a transient event captured and the high-pass filter output.

The morphological closing operation is commonly used in image processing to fill gaps or holes. In this project, it's used to eliminate the oscillatory behavior of the output signal of the high-pass filter, eliminating the problem with multiple crossings of the threshold. When using morphological operations, each sample of the resulting signal depends on the correspondent input sample and the samples in its neighborhood. The size of the neighborhood is determined by a structuring element which, when applied to one-dimensional vectors such as the ones in this project, is comprised of a binary vector.

The closing operation, according to its definition, is implemented using two other morphological operations: dilation and erosion:

$$u_{morph} = |u_{HP}| \bullet SE = \left(|u_{HP}| \oplus SE\right) \ominus SE \ , \tag{30}$$

with SE being the structuring element used and u_{hp} the output of the high-pass filter. As such, applying the closing operation to a signal, using a structuring element, is dilating it first and following it up with the erosion. The dilation of a function u with a binary structuring element SE is defined as

$$(u \oplus SE)(t) = \max\{u(t-\tau)\} \quad \forall SE(\tau) \neq 0, \tau \in S, t-\tau \in U \ , \tag{31}$$

where U and S are the domains of definition of the functions u and SE, respectively. This means that the value of the output sample is the maximum value of all samples in the input sample's neighborhood. On the other hand, the erosion of a function u with a binary structuring element SE is defined as

$$(u \ominus SE)(t) = \min\{u(t-\tau)\} \quad \forall SE(\tau) \neq 0, \tau \in S, t-\tau \in U \tag{32}$$

where U and S are the domains of definition of functions u and SE, respectively. This means that the erosion does the opposite of the dilation, since the value of the output sample is the minimum value of all samples in the input sample's neighborhood.

The dilation and erosion morphological operations were implemented using the van Herk-Gil-Werman algorithm [59] [60]. There are more efficient ways to implement these algorithms (e.g. [61]), but those solutions require a larger amount of available memory, which is not available in this project. Even though an external memory component is used, to provide additional memory space, the accesses to it are much slower than accesses to internal memory, decreasing the performance of those algorithms to an extent that they would exceed the time of the processing window.

Figure 41 shows an example of the morphological closing applied to the signal shown in Figure 40.

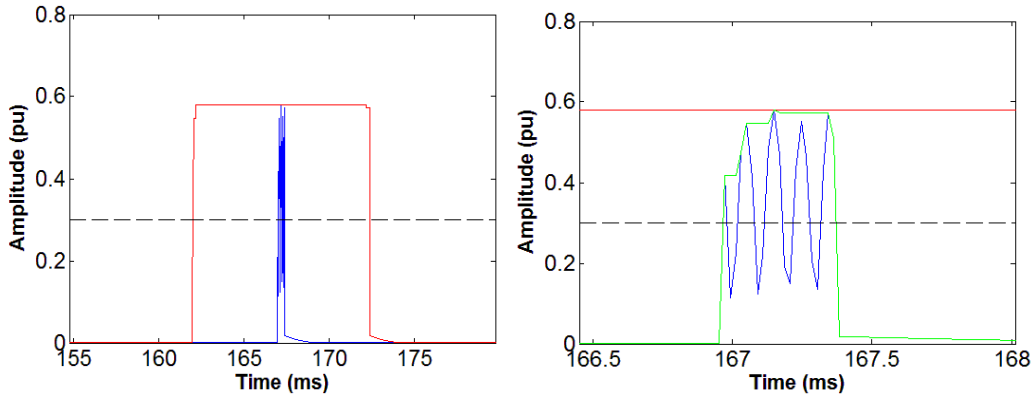

Figure 41: Result of the dilation applied to the signal shown in Figure 40 in red and the result of the erosion of the dilation in green.

As shown in Figure 41, the multiple crossing of the defined threshold is no longer a problem, making the detection and classification stage of these types of events much easier and precise. The algorithm used to detect and classify these kinds of events is the same used in the two other parts of the processing stage. Every time one output sample of the erosion operation is calculated, it is compared with the threshold value. If this calculated value is higher than the threshold level, then an event has been captured and its properties will begin to be recorded: a spot in the event queue is booked for that event, the duration is set to one, the maximum value recorded is set as the calculated value and the time and date are retrieved and added to the event packet which will be stored in the queue. From there on out, each sample above the threshold will add one more unit to the duration field and will be tested against the previous maximum value sample to determine the maximum value detected. When one sample is finally below the threshold level, the duration, maximum recorded value and event type will be stored in the event packet and a signal will be given to the Ethernet state machine to send the packet to the web server. The event type field depends on the duration value: if it does not exceed two and a half periods of the input signal, the event is classified as a transient; else, it's classified as a waveform distortion.

To test these algorithms, a simulated signal containing a transient event was fed to the system to determine if it detects and correctly classifies the power quality event. The samples are filtered from period to period, meaning that the filter input takes 1024 samples and it's possible to filter the three phases at once, using multiple channels. So, each phase samples are copied to the input filter buffer of the corresponding channel and when all the channel buffers are filled the filter will initiate. The number of samples that the filter uses is a limitation of the processor, since an IIR accelerator is being used to calculate the filter in parallel with RMS calculations.

The first thing to take into account though, is that the filter has a minimum setup time before the results are valid. This happens because only the coefficients are being fed to the filter and not the initial conditions and so these must be calculated first. Figure 42 shows the first eight outputs of the filter.

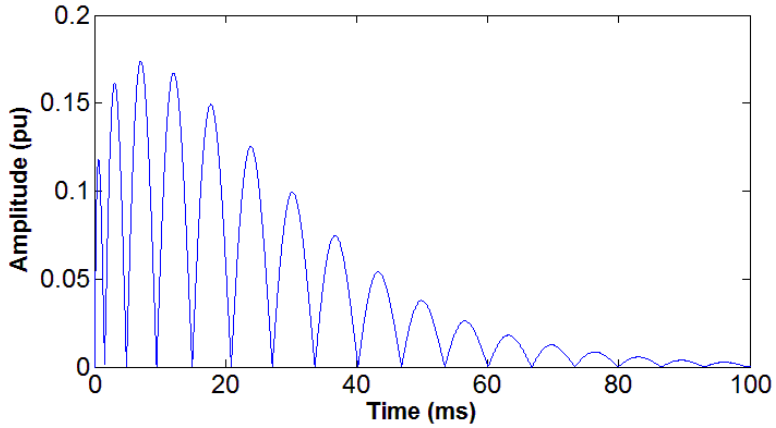

Figure 42: Digital filter setup time and first calculations of status variables.

As shown, the filter takes about 100 ms to stabilize. After the filter has stabilized, the system will start to apply the closing operation to its output, in order to remove the multiple crossings of the threshold problem described earlier. But first, some input signals with different frequencies were used to test the filter and make sure it's working correctly. Figure 43 shows the output of the used filter of different input signals with different frequencies.

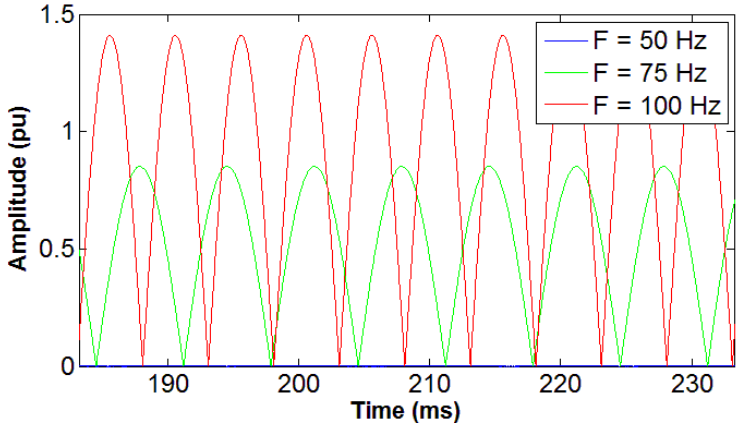

Figure 43: Filter output using input signals with frequencies between 50 and 100 Hz.

As shown, the filter is effectively removing the 50 Hz component of the input signal, with the amplitude of the filter output being close to 0. The attenuation of the 75 Hz signal is much lower, with this component almost passing through completely. The 100 Hz component has no attenuation at all, presenting 100% of its value on the filter output. Now that the filter has been tested and proven functional, the closing operation can be applied to the input signals. A transient event was simulated in one of the phases signals and the results are shown in Figure 44 and Figure 45. Figure 44 shows the input signal that the system sampled while Figure 45 shows the results of the high-pass filter, dilation and erosion morphological operations.

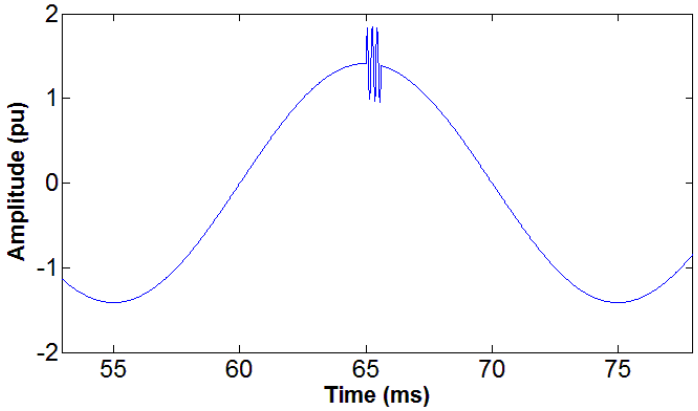

Figure 44: Simulated signal containing a transient event.

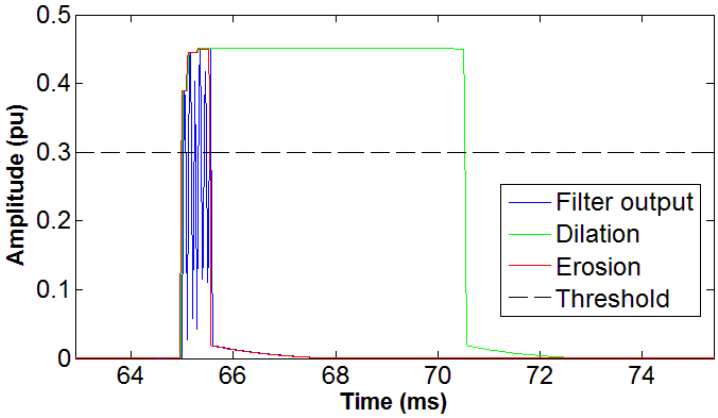

Figure 45: Results of the high-pass filter and morphological operations applied to the signal in Figure 44.

As shown, there's a significant difference between the simulated results presented in Figure 41 and the experimental results presented in Figure 45. This is because the simulated results were obtained through Matlab using algorithms done by definition. These algorithms are not meant to be efficient, memory and processing wise, meaning they were just used to test the validity of the method. The efficient version of these algorithms, developed using the van Herk-Gil-Werman algorithm, change the visual aspect of the results, but not the results themselves. As shown in Figure 45, the only difference is a shift of the central position of the event, maintaining the amplitude and duration of the event itself intact.

To classify these events, an algorithm identical to the one used in the THD case is used. That algorithm, however, is used while the erosion operation algorithm is running. This means that every time a sample is calculated using the erosion algorithm, that sample is fed to the classifying algorithm to detect events. This approach is more efficient because it avoids having a large number of memory accesses, since the value of the sample is present in internal memory and, when the erosion algorithm finishes, the events have already been detected and classified. Since the classifying algorithm is identical, it will not be reproduced here, see Figure 38 for clarification.

The last thing to determine about the filter and closing algorithms is their execution time. Since the filter is done using the accelerator, meaning the execution of the filter is in parallel with other instructions, the time for the filter to complete will not be taken into account. Only the time for the morphological operations will be taken into account. Table 9 shows the calculated execution time of these algorithms.

Table 9: Calculated execution time of the morphological operation algorithms.

	Best-case scenario	Average	Worst-case scenario
Number of cycles (M cycles)	3.361	3.362	3.743
Time (ms)	8.403	8.405	9.357

One thing to notice about the values presented in Table 9 is that even in the best-case scenario, the processing time of these algorithms is greater than the processing window between phases processing and if all phases were to calculate the closing operation on the same iteration, the sum of the three would surpass the value of the overall processing window. So, the program makes sure that only one phase calculates this morphological operation per processing window, meaning that the phase that samples eight periods and has the filter output samples first calculates this morphological operation first. The two other phases will have to wait for the next processing window. Then, the first one to reach the calculation of the morphological operation does so and makes the other one wait for the next processing window. This order is recorded and will continue to be used until the system is shutdown.

4.6 Summary

In this chapter the method for detecting and classifying the various types of power quality events was described. Along with this description, some simulated and experimental results were presented. The result of the classification of the events is presented in Appendix A. One last thing to do is to make sure that the sum of the processing time of all algorithms does not surpass the value of the overall processing window. Table 10 shows a summary of the processing time of these algorithms, focusing on the worst-case scenarios and comparing it to the processing window value.

Table 10: Summary of the worst-case scenario processing time of the algorithms used.

Algorithm	Number of Cycles	Time (ms)
RMS	427280	1.068
THD	627886	1.570
Morphological	3742618	9.357
Total	4797784	11.995
Processing window	7690000	19.2
Remaining	2892216	7.205

As shown, this system can perform all operations within the given processing window, meaning that it is a viable real time system. There's also at least 7 ms of time to send the information gathered about the power quality events captured to the web server. The processing window time can be stretched if more samples are stored before starting to process them. However this would mean finding another solution for the high-pass filter, since the accelerator only takes 1024 samples as input per channel.

5. Conclusions

The responsibility to keep a good energy quality is up to the providers. Industrial level consumers have large AC motors that could cause sags or x-ray machines that produce a lot of harmonic content in the power system and these situations need to be analyzed by the energy providers so that other users aren't affected by that equipment, meaning that the energy providers have to make sure that they provide a good quality service. However, there's still no consensus on what is considered "good" quality. This is why there are different standards with different definitions of the power quality events, which are, sometimes, complemented with a regional standard. The power quality problem is a serious one since it has many negative consequences as it was shown before. It's very important to detect these power quality events and take preventive measures so these negative consequences can be prevented or completely mitigated, ensuring that electrical and electronic devices work under good energy conditions. As such, the development of cheap, efficient and intelligent systems that can detect these power quality events is a necessity.

In this project a system that can detect and classify power quality events was developed. This system is independent of other hardware modules, with the exception of a router with internet connectivity. This system was designed to have the Ethernet communication as a core feature, and not the possibility to store the data in non-volatile memory. The innovation and major advantage of this system is its ability to automatically connect to the internet, using the DHCP, and to store that information in a web server, allowing the acquisition of data in several points of interest, such as several nodes in power systems, having multiple systems connected to a single web server. It uses a detection and classification method that is based on a method used for single phase power systems that has been proven to work successfully [62]. This method does not use a single algorithm or mathematical transform to detect multiple events, like the methods presented in the state of the art subchapter, but it uses the simplest approaches to each type of event. This lowers the complexity of the signal processing in the DSP, although more algorithms need to be used to correctly determine the different types of events that might be present in the input signals. It handles the Ethernet connectivity part automatically, without the need of user intervention, although the server IP address needs to be given with the SD Card. It's a plug and play device which is complemented with the web server application detailed in Appendix A.

For future work, some modifications could be made to improve the main system, such as:

- Choosing another Ethernet chip with power saving properties, to power down the IC when it isn't being used;
- Use another type of connection between the SD Card and the DSP, giving the system a reliable alternative to store collected data when the Ethernet connection fails;
- Adding current acquisition circuits, to allow even more data to be collected regarding power quality overall;
- Implementing a DNS client, allowing the system to retrieve the IP of the web server from the DNS server instead of being given by the user.

References

[1] The Institute of Electrical and Electronics Engineers, Inc., *"IEEE Recommended Practice for Monitoring Electric Power Quality"*. New York, 1995.

[2] F. Oliveira, H. Torre, and J. Almeida, "QUALIDADE DE ENERGIA," Universidade de Coimbra Faculdade de Ciências e Tecnologia, 2004.

[3] Z. Zou, Y. Liu, F. Zhao, and M. Peng, "Design and data processing of a real-time power quality monitoring instrument," *Electronic Measurement & Instruments, 2009. ICEMI '09. 9th International Conference on* , pp. 4-353-4-357, 2009.

[4] P. E. Issouribehere, J. C. Barbero, and F. Issouribehere, "POWER QUALITY MEASUREMENTS IN A STEEL INDUSTRY WITH ELECTRIC ARC FURNACES," *Power and Energy Society General Meeting - Conversion and Delivery of Electrical Energy in the 21st Century, 2008 IEEE*, pp. 1-8, 2008.

[5] D. O. Koval, A. U. ,. E. A. ,. C. Dept. of Electr. & Comput. Eng., R. A. Bocancea, K. Yao, and M. B. Hughes, "Frequency and duration of voltage sags and surges at industrial sites-Canadian National Power Quality Survey," *Industry Applications Conference, 1997. Thirty-Second IAS Annual Meeting, IAS '97., Conference Record of the 1997 IEEE*, vol. 3, pp. 2189-2196, 1997.

[6] A. K. ,. G. C. P. ,. S. G. K. Goswami, "Cost-Benefit Analysis of Voltage Sag Mitigation Methods in Chemical Industry," *Electrical Engineering/Electronics, Computer, Telecommunications and Information Technology (ECTI-CON), 2012 9th International Conference on*, pp. 1-4, 2012.

[7] K. Gorecki, M. Szmajda, and J. Mroczka, "Adaptive digital synchronization of measuring window in low-cost DSP power quality measurement systems," *9th International Conference Electrical Power Quality and Utilisation.*, pp. 1-4, 2007.

[8] U. Rao, S. N. Singh, and C. Kumar Thakur, "Power quality issues with medical electronics equipment in hospitals," *Industrial Electronics, Control & Robotics (IECR), 2010 International Conference on*, pp. 34-38, 2010.

[9] M. I. Buzdugan, H. Balan, and D. T. Mureşan, "An electrical power quality problem in an emergency unit from a hospital – Case study –," *Power Electronics Electrical Drives Automation and Motion (SPEEDAM), 2010 International Symposium on*, pp. 251-256, 2010.

[10] P. Xavier, "Analisador da Qualidade de Energia Baseado em DSP," Instituto Superior Técnico, 2011.

[11] R. C. Sennon, "An Overview of Pbwer Quality Standards and Guidelines from the End-User's Point-of-View," *2005 IEEE Rural Electric Power Conference*, pp. B1/1-B1/5, 2005.

[12] McKinsey & Company. The Internet of Things. [Online]. http://www.mckinsey.com/insights/high_tech_telecoms_internet/the_internet_of_things

[13] K. D. Preeuw. the internet of things | Are you ready for the internet of things?. [Online]. http://www.theinternetofthings.eu/

[14] C. Zhu, M. Hu, Z. Wu, X. Dou, and S. Zhao, "Design and realization of regional power quality monitoring system," *Electric Utility Deregulation and Restructuring and Power Technologies, 2008. DRPT 2008. Third International Conference*, 2008.

[15] M. Szmajda, K. Gorecki, and J. Mroczka, "DFT algorithm analysis in low-cost power quality measurement systems based on a DSP processor.," *9th Internation Conference Electrical Power Quality and Utilisation*, pp. 1-6, 2007.

[16] G. W. Chang, C.-Y. Chen, and M.-C. Wu, "Measuring Harmonics by An Improved FFT-based Algorithm with Considering Frequency Variations," *Circuits and Systems, 2006. ISCAS 2006. Proceedings. 2006 IEEE International Symposium*, 2006.

[17] E. Styvaktakis, M. H. J. Bollen, and I. Y. H. Gu, "Automatic Classification of Power System Events using rms Voltage Measurements," *Power Engineering Society Summer Meeting, 2002 IEEE*, vol. 2, pp. 824-829, 2002.

[18] D. L. Brooks, R. C. Dugan, M. Waclawiak, and A. Sundaram, "Indices for assessing utility distribution system RMS variation performance," *Power Delivery, IEEE Transactions on*, vol. 13, no. 1, pp. 254-259, 1998.

[19] N. Kagan, et al., "Influence of RMS variation measurement protocols on electrical system performance indices for voltage sags and swells," *Harmonics and Quality of Power, 2000. Proceedings. Ninth International Conference on*, vol. 3, pp. 790-795, 2000.

[20] V. V. Garcia, C. A. D. Gualdron, and G. O. Plata, "Obtaining patterns for classification of power quality disturbances using biorthogonal wavelets, RMS value and support vector machines," *Electrical Power Quality and Utilisation, 2007. EPQU 2007. 9th International Conference on*, pp. 1-6, 2007.

[21] V. D. K. Kasula, "Classification of Power Quality Disturbance Signals Using FFT, STFT, Wavelet Transforms and Neural Networks - A Comparative Analysis," *Conference on Computational Intelligence and Multimedia Applications, 2007. International Conference on*, vol. 1, pp. 335-340, 2007.

[22] S. De, S. Nath, P. Sinha, and S. Saha, "A virtual instrument for real time power quality measurement under nonsinusoidal conditions," *Signal Processing, Communication, Computing and Networking Technologies (ICSCCN), 2011 International Conference on*, pp. 140-144, 2011.

[23] Z. Yonghui, C. Xi, and Z. Xiyuan, "Power Quality Analysis Based on All-Phase FFT," *Intelligent System Design and Engineering Application (ISDEA), 2010 International Conference on*, vol. 2, pp. 69-72, 2010.

[24] G. Gasparesc, "Automatic classification of power quality events using S-Transform," *Environment and Electrical Engineering (EEEIC), 2011 10th International Conference on*, pp. 1-4, 2011.

[25] A. M. Gargoom, N. Ertugrul, and W. L. Soong, "Automatic Classification and Characterization of Power Quality Events," *Power Delivery, IEEE Transactions on*, vol. 23, no. 4, pp. 2417-2425, 2008.

[26] A. M. Gargoom, N. Ertugrul, and W. L. Soong, "Investigation of Effective Automatic Recognition Systems of Power-Quality Events," *Power Delivery, IEEE Transactions on*, vol. 22, no. 4, pp. 2319-2326, 2007.

[27] Y. Hai and J.-y. Chen, "Power quality evaluation based on wavelet packet decomposition and fuzzy logic," *Computer Science and Automation Engineering (CSAE), 2012 IEEE International Conference on*, vol. 3, pp. 504-507, 2012.

[28] S. K. Meher, "A Novel Power Quality Event Classification using Slantlet Transform and Fuzzy Logic," *Power System Technology and IEEE Power India Conference, 2008. POWERCON 2008. Joint International Conference*, pp. 1-4, 2008.

[29] G. Shahgholian, M. Mahdavian, A. Emami, and B. Ahmadzade, "Improve power quality using static synchronous compensator with fuzzy logic controller," *Electrical Machines and Systems (ICEMS), 2011 International Conference*, pp. 1-5, 2011.

[30] A. Cardenas, C. Guzman, and K. Agbossou, "Real-time evaluation of power quality using FPGA based measurement system," *Industrial Electronics (ISIE), 2010 IEEE International Symposium on*, pp. 2777-2782, 2010.

[31] N. T. Huang, Y. J. Zhang, D. G. Xu, X. S. Liu, and J. J. Qi, "Power quality disturbances classification by ensemble and hybrid Neural networks," *Power System Technology (POWERCON), 2010 International Conference on*, pp. 1-7, 2010.

[32] A. K. Chandel, G. Guleria, and R. Chandel, "Classification of power quality problems using wavelet based artificial neural network," *Transmission and Distribution Conference and Exposition, 2008. T&D. IEEE/PES*, pp. 1-5, 2008.

[33] S. Kaewarsa, "Detection of Power Quality Events Using DOST-Based Support Vector Machines," *Computer Science and its Applications, 2008. CSA '08. International Symposium on*, pp. 68-71, 2008.

[34] J. Susukh, S. Premrudeepreechacharn, and T. Kasirawat, "Power quality problem classification using support vector machine," *Electrical Engineering/Electronics, Computer, Telecommunications and Information Technology, 2009. ECTI-CON 2009. 6th International Conference*, vol. 1, pp. 58-61, 2009.

[35] P. Gao and W. Wu, "Power Quality Disturbances Classification using Wavelet and Support Vector Machines," *Intelligent Systems Design and Applications, 2006. ISDA '06. Sixth International Conference on*, vol. 1, pp. 201-106, 2006.

[36] J. Chung, E. J. Powers, W. M. Grady, and S. C. Bhatt, "Power disturbance classifier using a rule-based method and wavelet packet-based hidden Markov model," *Power Delivery, IEEE Transactions on*, vol. 17, no. 1, pp. 233-241, 2002.

[37] A. M. Youssef, T. K. Abdel-Galil, E. F. El-Saadany, and M. M. A. Salama, "Disturbance classification utilizing dynamic time warping classifier," *Power Delivery, IEEE Transactions on* , vol. 19, no. 1, pp. 272-278, 2004.

[38] AEMC Instruments. PowerPad Model 8335. [Online]. http://www.aemc.com/products/pdf/2136.20.pdf

[39] Fluke Corporation. Fluke 1750 Power Quality Recorder & Analyzers for Monitoring. [Online]. http://www.fluke.com/Fluke/inen/Power-Quality-Tools/Logging-Power-Meters/Fluke-1750.htm?PID=56030

[40] M. McGranaghan, "Roadmap for power quality standards development," *Power Engineering Society General Meeting, 2005. IEEE*, vol. 3, pp. 2209-2213, 2005.

[41] GIM - Instrument and Measurement Group. [Online]. http://gim.lx.it.pt/fct57708/data

[42] H. Markiewicz and A. Klajn. (2004) Voltage Disturbances Standard EN 50160 Voltage Characteristics in Public Distribution Systems. [Online]. http://www.copperinfo.co.uk/power-quality/downloads/pqug/542-standard-en-50160-voltage-characteristics-in.pdf

[43] Analog Devices, "16-Bit, 1 MSPS PulSAR ADC in MSOP/QFN AD7980," 2013.

[44] Microchip, "Miniature Single-Cell, Fully Integrated Li-Ion, Li-Polymer Charge Management Controllers," 2013.

[45] Microchip, "Designing A Li-Ion Battery Charger and Load Sharing System With Microchip's Stand-Alone Li-Ion Battery Charge Management Controller," 2008.

[46] Analog Devices, "ADSP-21489," 2010.

[47] Wikipedia. (2013) Super Harvard Architecture Single-Chip Computer. [Online]. http://en.wikipedia.org/wiki/Super_Harvard_Architecture_Single-Chip_Computer

[48] Texas Instruments, "REAL-TIME CLOCK (RTC)," 2010.

[49] WIZnet, "W5100 Datasheet Version 1.2.4," 2011.

[50] WIZnet, "W5100 Porting Guide Ver. 1.0," 2008.

[51] WIZnet, "App_note_DHCP," 2005.

[52] Micron, "Synchronous DRAM," 2008.

[53] SD Card Association, "SD Specifications - Part 1 - Physical Layer - Simplified

Specification," 2013.

[54] ChaN. FatFs - Generic FAT File System Module. [Online]. http://elm-chan.org/fsw/ff/00index_e.html

[55] ChaN. My Profile. [Online]. http://elm-chan.org/profile_e.html

[56] Analog Devices, "Visual DSP 5.0++ Run-Time Library Manual for SHARC processors," 2011.

[57] J. Serra, *Image Analysis and Mathematical Morphology*. 1982.

[58] MathWorks. (2014) Using FDATool. [Online]. http://www.mathworks.com/help/signal/ug/opening-fdatool.html

[59] M. v. Herk, "A fast algorithm for local minimum and maximum filters on rectangular and octogonal kernels," *Pattern Recognition Letters*, vol. 13, no. 7, pp. 517-521, Jul. 1991.

[60] J. ,. W. M. Gil, "Computing 2-D min, median, and max filters," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 15, no. 5, pp. 504-507, May 1993.

[61] J. Gil and R. Kimmel, "Efficient dilation, erosion, opening, and closing algorithms," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 24, no. 12, pp. 1606-1617, Jan. 2003.

[62] T. R. F. M. J. A. C. S. Pedro M. Ramos, "DSP based power quality analyzer using new signal processing algorithms for detection and classification of disturbances in a single-phase power system," *15th IMEKO TC4 Symposium on Novelties in Electrical Measurements and Instrumentation*, pp. 433-438.

[63] Analog Devices, "Estimating Power for ADSP-214xx SHARC® Processors," 2011.

Appendix A – DHCP Client and Web Server

A.1. DHCP Client

The most important feature of this system is to be completely independent of the user and the location where it's used. As such, it's extremely important for this system to be able to get the necessary network information to connect to the internet with minimal setup required from the user. This is accomplished with a protocol known as Dynamic Host Configuration Protocol (DHCP). The DHCP is a standardized network protocol used on Internet Protocol (IP) networks for dynamically distributing network configuration parameters, such as IP addresses for interfaces and services. The DHCP protocol is common nowadays, being present in pretty much every networking device. Every time people connect their devices to a public network, or even a private network in offices or schools, they don't need to configure their own IP address, gateway IP address or subnet mask. Everything is done automatically upon connection in a way that the user does not even notice. Only in specific occasions people still use static IP addresses, such as web servers. However, even in those situations, the routers have the ability to distribute IP addresses, albeit using a shorter address pool.

To have the system developed in this project connect automatically to the Internet, it needs to be connected to a router with an active DHCP server. This will allow the system to get the required network information from the router and becoming able to reach the web server that will receive the power quality event packets. It can be connected to another PC using the Ethernet port, but a lot of configurations would be needed for this setup to work. The PC would have to have a DHCP server and would need to act like a network bridge to the router present in the network. The connection with the router is plug-and-play, meaning that no other configurations are needed, except if the DHCP server is not correctly configured.

Figure 46 shows the communication required between DHCP client and DHCP server to grant the client the necessary network information for it to have Internet connectivity.

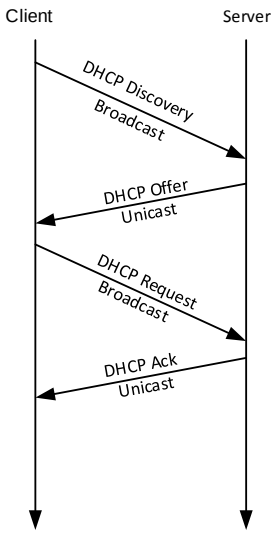

Figure 46: Example of network information retrieval using DHCP.

As it is shown in Figure 46, this protocol is based on request and reply. The exchange of information is started by the client when it joins a new network, sending a broadcast DHCP Discovery packet. Every DHCP server present in the network will answer to this request with a unicast DHCP Offer packet. This packet contains the IP address offered by the DHCP server to the client, as well as other pertinent information required by the client such as subnet mask and gateway IP address. If the client does not agree with any of the information given, being it the IP address offered or if it does not wish to connect to that specific server, it sends another broadcast DHCP Discovery packet. However, if the client agrees with everything offered by the server, it sends a broadcast DHCP Request packet that contains the information received by the DHCP server that it wishes to be connected to. That packet is finally answered with a unicast DHCP Acknowledge message, confirming the connection between the two. After this message is sent, the DHCP server will store the information regarding that device for a period of time called the lease period. During this time, the network information granted to the device is blocked. If, however, that request isn't renewed until the end of the lease time, the IP address that was given to that device is no longer assigned to it, being returned to the pool of available IP addresses that can be assigned to devices who request them from the DHCP server. Some routers implement an infinite lease time, meaning that the information is only revoked by the network administrator. Other routers implement one week long lease periods, others one day long lease periods. It all depends on the amount of devices that is usually added to that network. If it's a fairly dynamic network, with many devices being disconnected and connected, the lease time has to be low. In the case of a public area with free Internet connectivity, the lease times might not even reach the hour mark, since there's always new people joining the network, but they don't stay for long. If the network information is stale for a long period of time, the IP addresses present in the IP address pool might run out, leading to connectivity issues by the users. On the other hand, on a private network such as the ones people have at home, the lease periods can be significantly large, since there's not many users coming in or out. The lease time should always be as long as possible, because renewing network information leads to spending time and resources.

In this project, the DHCP client is implemented using a different socket than the one used to communicate with the webserver. That socket is a UPD socket that is only used to get the necessary network information and configure the Ethernet chip. After that, the socket is closed and not used again until it is necessary to renew this information. This protocol is also executed before the sampling and processing stage begins. If it fails to complete, the system will not start the sampling and processing stage since it has no way to report its findings.

A.2. Webserver communication

To show the user the results of the processing stage of the input signal, the system sends information to a web server that stores and displays it. This web server is the software complement of the system, since it's the only platform available to the user to check which power quality events happened.

The web server is implemented in PHP, a server-side scripting language designed for web development but also used as a general-purpose programming language. PHP originally stood for Personal Home Page, but now stands for Hypertext Preprocessor. This language was used because it

is easy to use and implements functions to interact with databases in an easy manner. The database is implemented using MySQL, which is the world's second most used open-source relational database management system. It has an easy graphical interface where the administrator can define the tables and fields it contains, as well as view the data that has already been entered.

The structure of the web server is shown in Figure 47.

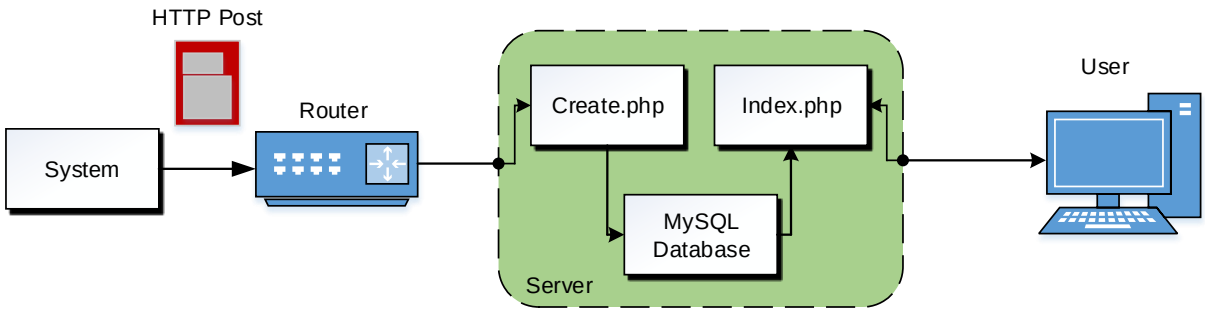

Figure 47: Communication between the system and the server and also between the server and the user.

It has two different PHP scripts running in parallel: the index.php script and the create.php script. The create script is responsible to receive the information from the system, organize it and store it adequately in the database. The system developed in this project uses a request method of the Hypertext Transfer Protocol (HTTP) to send the processed information to the web server. There are many request methods in the HTTP protocol, but the one used in this project is the Post method. The HTTP Post request method is designed to request that a web server accept the data enclosed in the request message's body for storage. It is often used when uploading a file or submitting a completed web form. The messages from the system to the web server are sent during the processing windows discussed previously. So, to minimize processing requirement, all the information is stored in byte arrays and is written in the HTTP Post message's body as a byte array, or octet stream, as well. This process is called serialization, meaning that this information has to be unserialized after being received. Upon reception, the create script turns the byte array into discernible information, filling variables with the information provided. For example, the first four bytes correspond to an integer and they're stored in a variable called password. The value of this integer is compared to a known value to determine if the message comes from the system and is a valid message. The date and time, for example, are seven separate bytes, with the first corresponding to the number of seconds, the second is the number of minutes and so on. In the end, all expected variables are filled with the information provided by the system, which is good in terms of debugging, because it allows the programmer to notice if there's something wrong with how the information is being sent, but it is also necessary when interfacing with the database, since every field of the table has to be updated individually. The last thing the script does is to create a new entry on the table of the database and store the gathered information there. This script runs every time a HTTP Post is received, meaning that this web server could handle received messages from multiple systems, allowing the idea of having multiple power quality analyzers connected to the same web server.

The index script responsibility is to retrieve data from the database and to display it to the user. It is also responsible for the layout of the website. Once the information about the power quality events is retrieved from the database, it's displayed according to the type of event detected. Figure 48 shows an example of how a power quality event is displayed. The link for the web server is http://193.136.223.175/post_test/index.php. Figure 49 shows an example of the aspect of the web page.

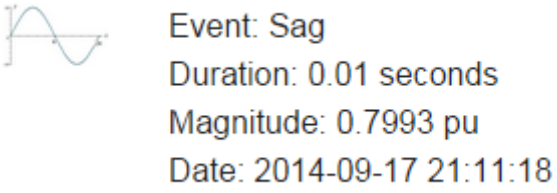

Figure 48: Example of how a power quality event is displayed.

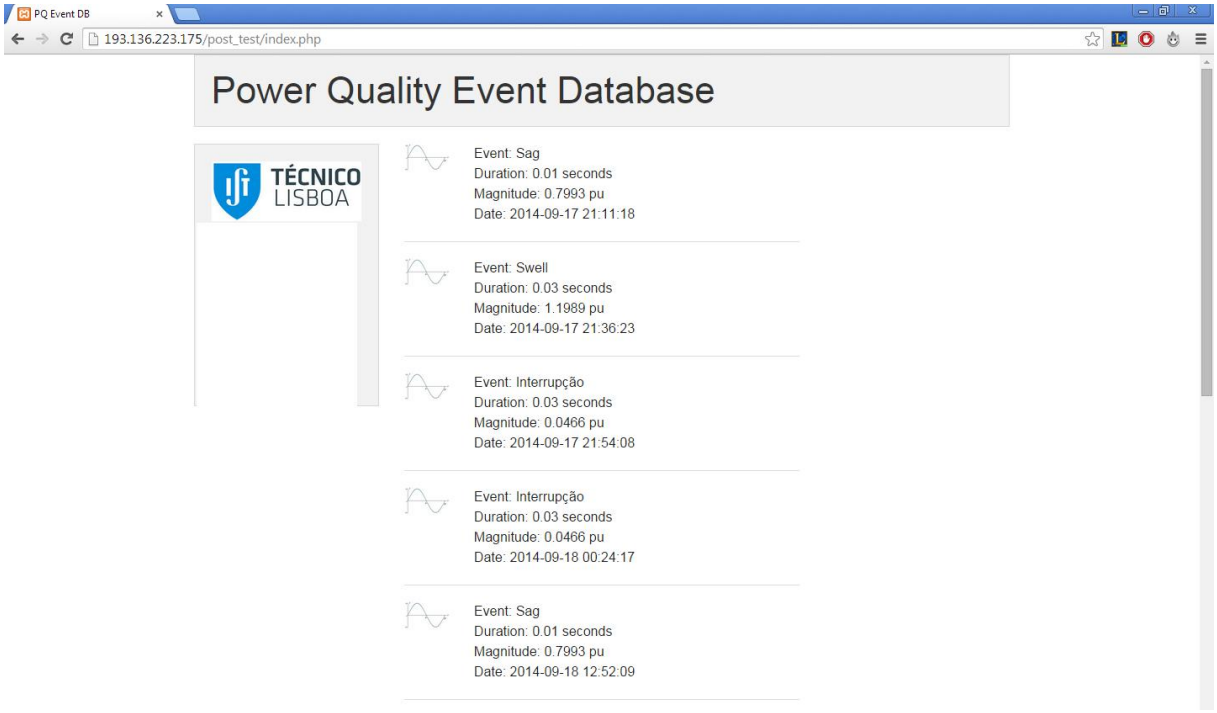

Figure 49: Example of the developed web site displaying some captured events.

Appendix B – System Power Consumption Estimation

In this appendix, an estimation of the power consumption of the developed system is provided. This estimation was used to choose the AC-DC converters used in the second board developed to power the system. This system has three main power input sources: the first is the output of the AC-DC converters, either being just one, or the three of them working together; the second, a 5 V input power connector is present if the user wants to use the system, without it being used for power quality detection (e.g. in debug or development scenarios); and the third is the 3.7 V LI-Ion battery used as backup when there's no voltage in the power grid. These three input power sources are connected to the battery management circuit, described in previous chapters, and the output of that circuit feeds several DC-DC converters which, in turn, will feed the various elements present in the system. An overview of the power distribution scheme is presented in Figure 50.

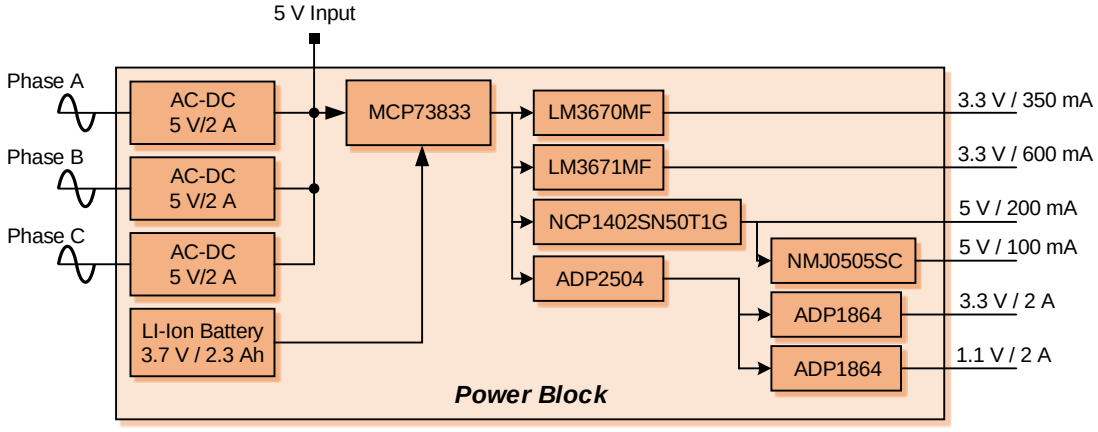

Figure 50: Block diagram of the power distribution in the system.

To correctly estimate the maximum power consumption of the system, the power that each device requires when the system is in full activity was estimated. Even though the system might not use every component at all times, this estimation will guarantee that the power supply will always be enough and the system will not fail, in terms of power consumption, at any given time. The estimation for each device was based on its datasheet specifications.

The supply current and power required by each device of this system is represented in Table 11. All the devices of the signal conditioning and acquisition circuit were condensed in a single row because they all feed of an isolated DC-DC converter, which has a maximum power rating, meaning that all those devices together cannot draw more current than the DC-DC converter can provide.

The power consumption estimate of the DSP was based on the engineer-to-engineer note EE-348 [63], from Analog Devices, and the device datasheet [46]. Again, to ensure that there are no system failures regarding the power supply, all of the indexes of this analysis were set to the highest power consumption possible, which means a 100% value on the High application activity level and 100% value on the use of every port available of the processor.

Table 11: Detailed power consumption values of the several devices in the system.

Description	Device	Voltage Supply [V]	Current [mA]	Power [mW]
Power Grid side	NMJ0505SC	5	170	850
Optocoupler	HCPL-4661	5	1.45	7.25
Memory Mangement	MCP73833	5	3	15
Ethernet	W5100	3.3	200	660
RTC	BQ32000DR	3.3	0.1	0.33
SD Card	-	3.3	200	660
SDRAM	MT48LC16M16A2	3.3	135	446
SPI Flash	M25P16	3.3	15	50
Reset	ADM706SARZ	3.3	0.2	0.66
Digital Signal Processor	ADSP-21489	3.3	134	443
Digital Signal Processor	ADSP-21489	1.1	839	923
Total				3395.24

The total system power estimation, taking into account power regulator efficiencies, is presented in Table 12.

Table 12: Total system power estimation with power regulator efficiencies.

Regulator	Output Voltage (V)	Output Current (mA)	Efficiency (%)	Input Power (mW)
NCP1402SN50T1G	5	200	85	1205
LM3670	3.3	350	88	1313
LM3671	3.3	600	82	2415
ADP2504	5	250	88	1421
Total				6354

Appendix C – Project Schematics

C.1. Signal Conditioning and Acquisition Circuit

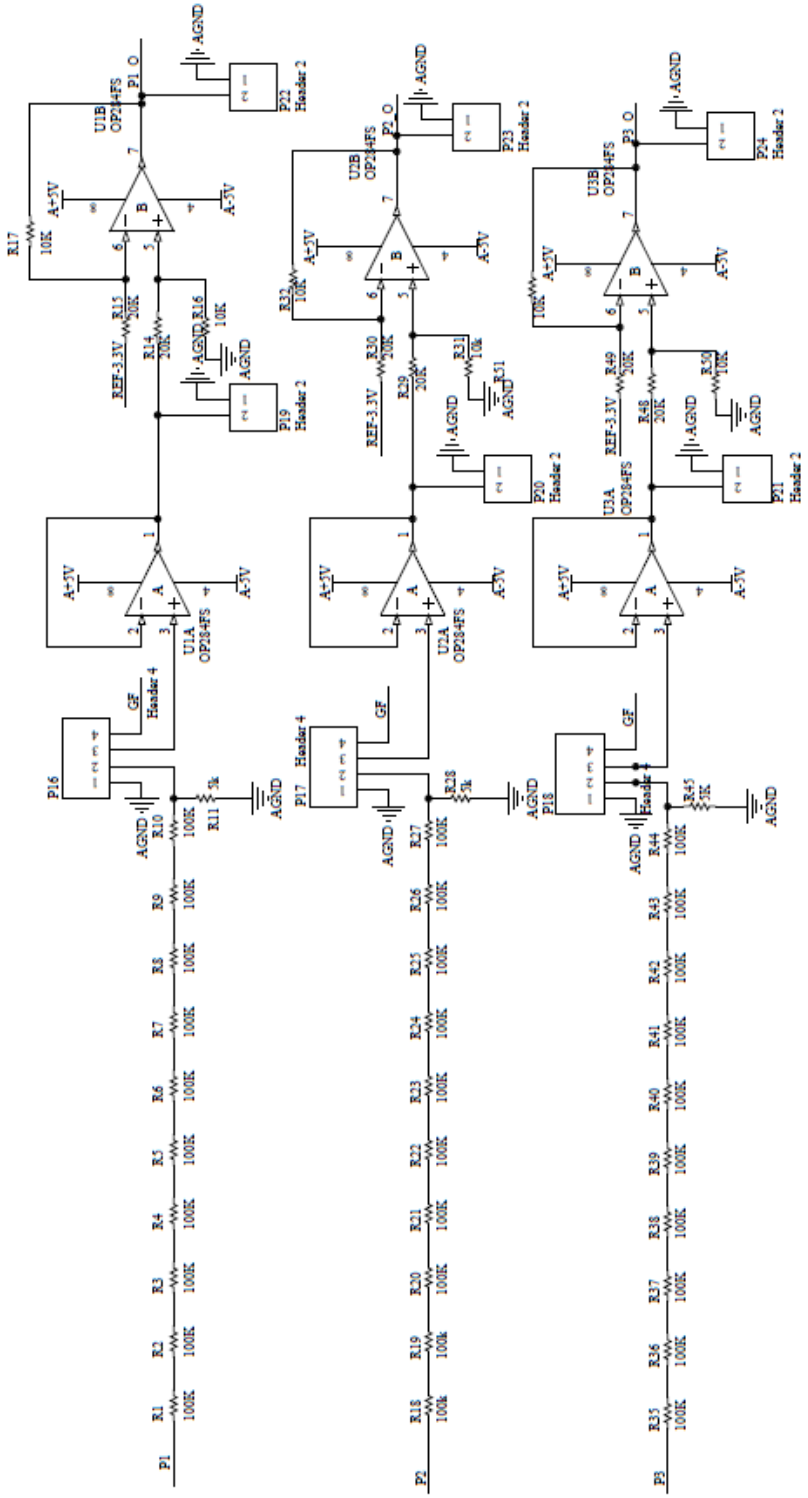

Figure 51: Signal conditioning circuit.

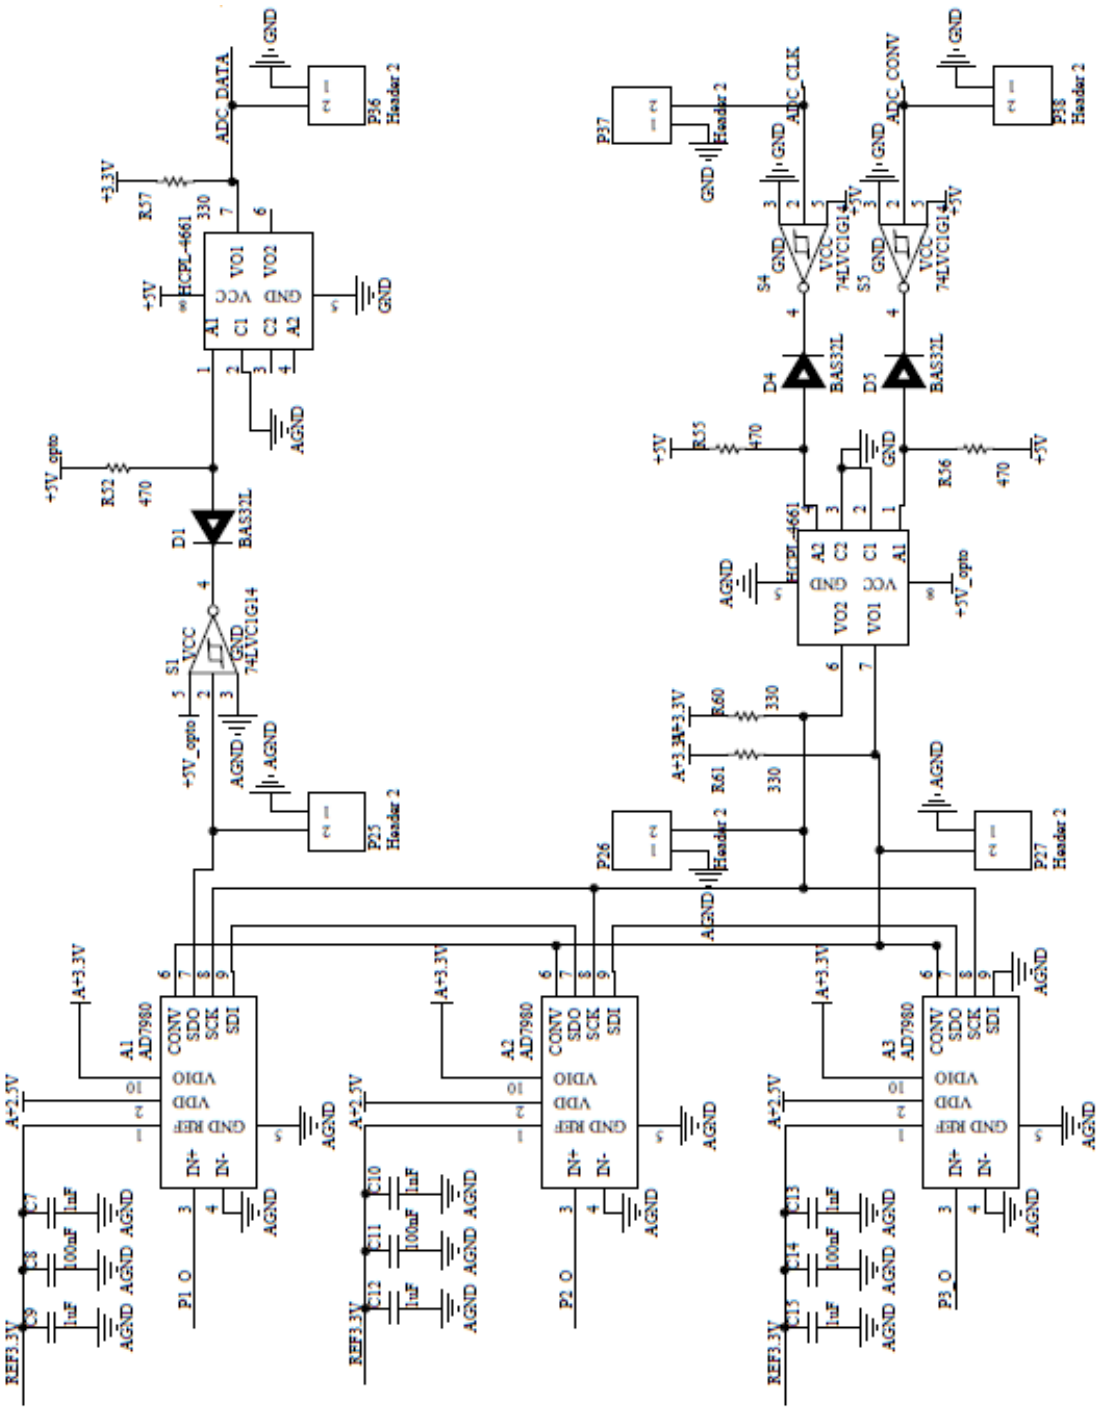

Figure 52: Signal Acquisition Circuit.

C.2. SD and RTC Schematics

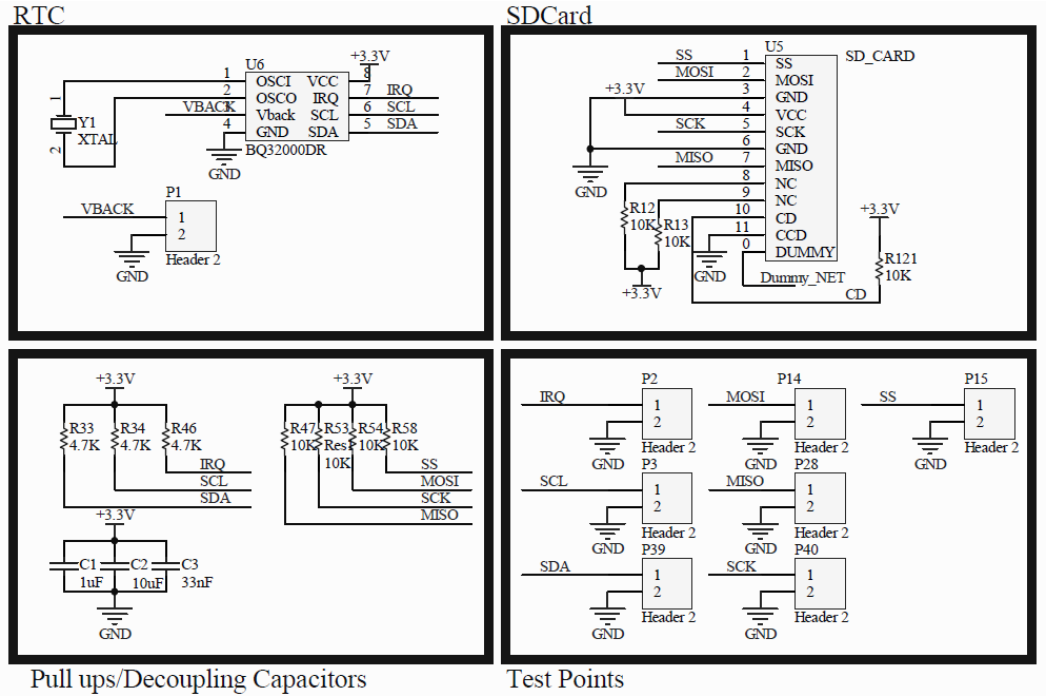

Figure 53: Schematics for the RTC IC and SD card holder, as well as decoupling capacitors, pull ups and test points.

C.3. Ethernet Schematics

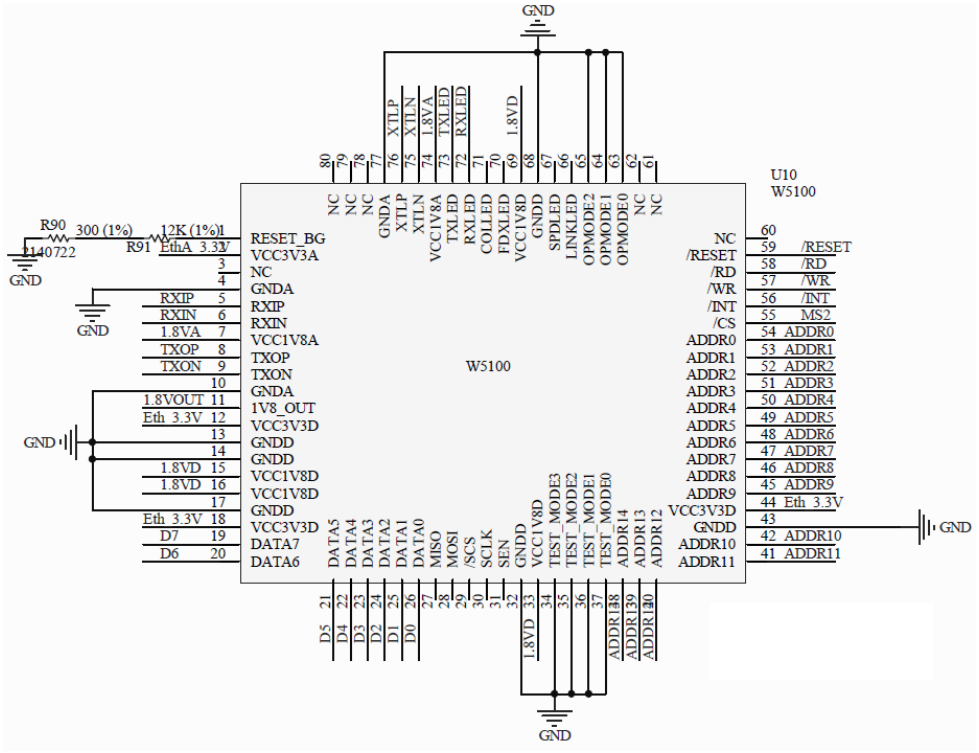

Figure 54: Schematic of the W5100 Ethernet IC.

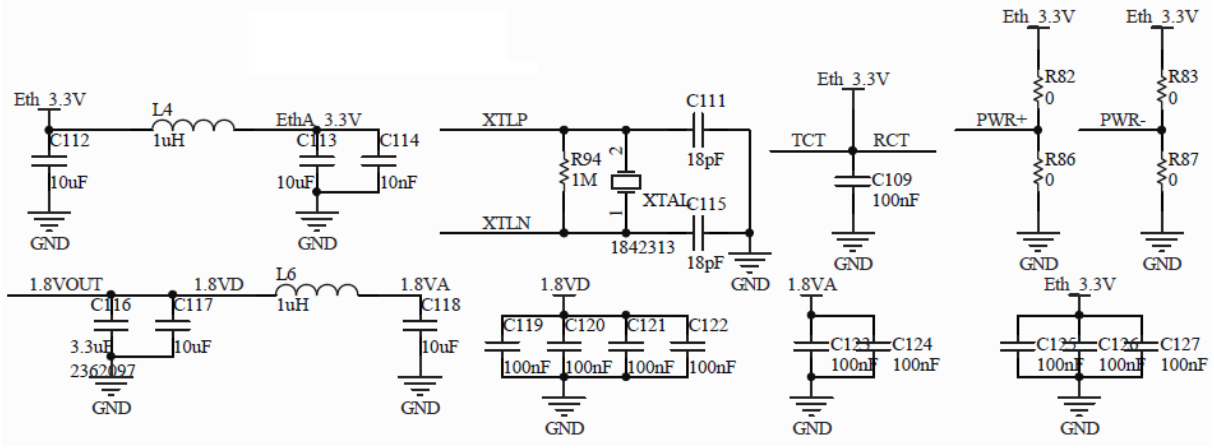

Figure 55: Decoupling capacitors, crystal, analog and digital voltage circuits.

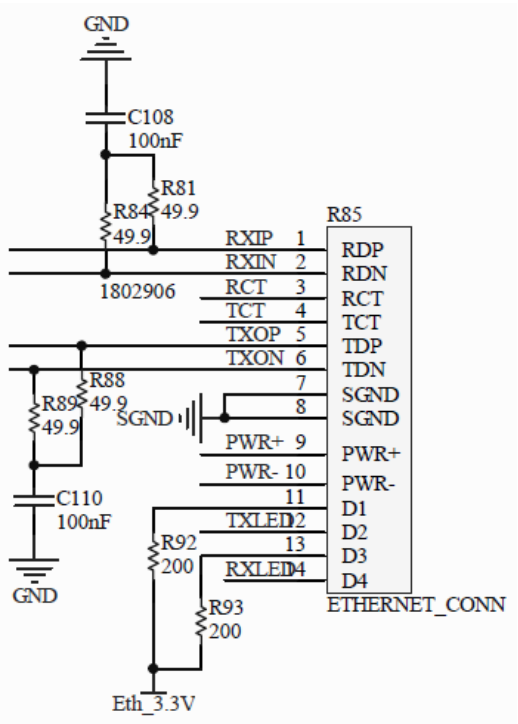

Figure 56: Schematic of the Ethernet connector.

C.4. SDRAM Schematics

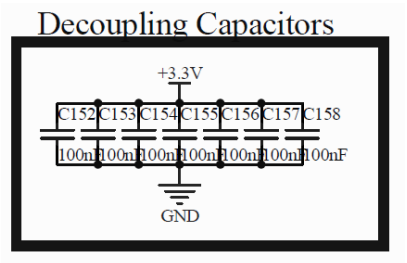

Figure 57: SDRAM decoupling capacitors.

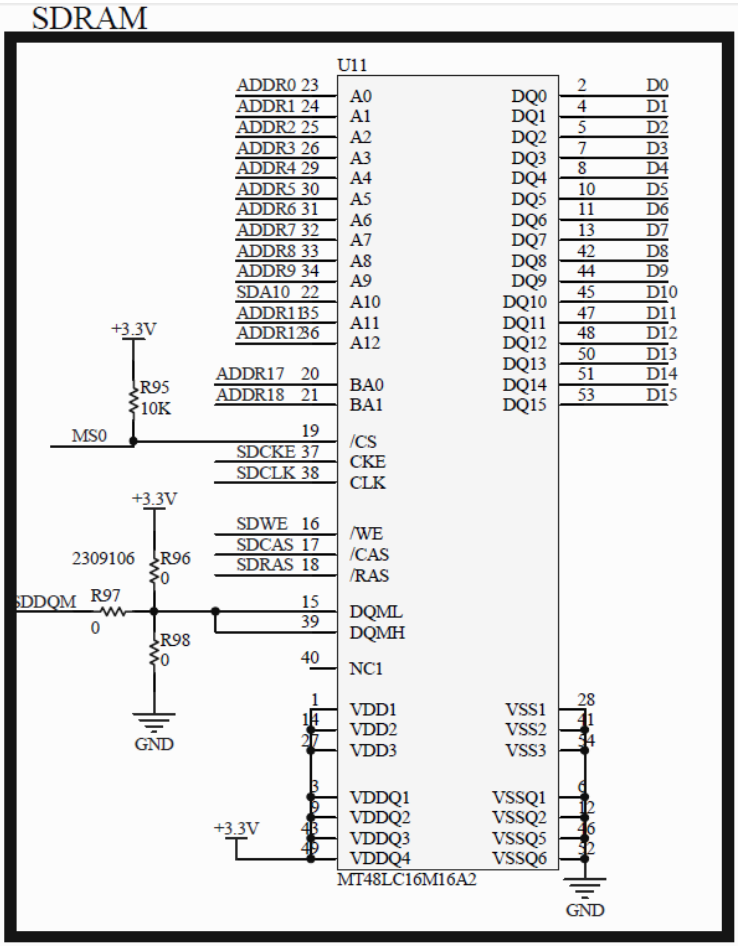

Figure 58: Schematic for the SDRAM IC.

C.5. DSP Schematics

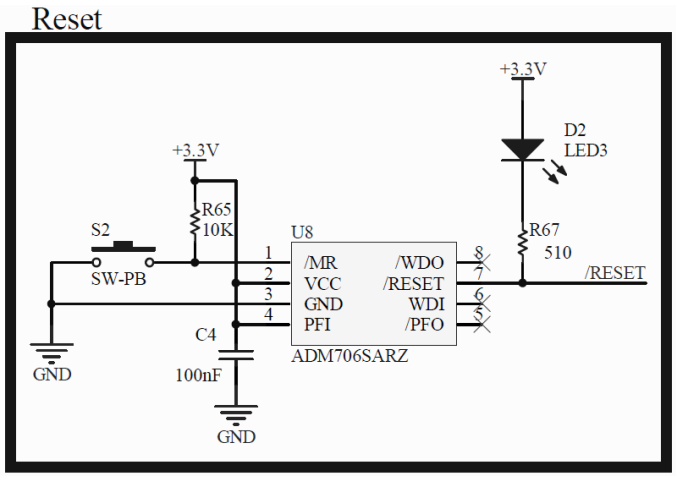

Figure 59: Reset circuit for the DSP and other ICs.

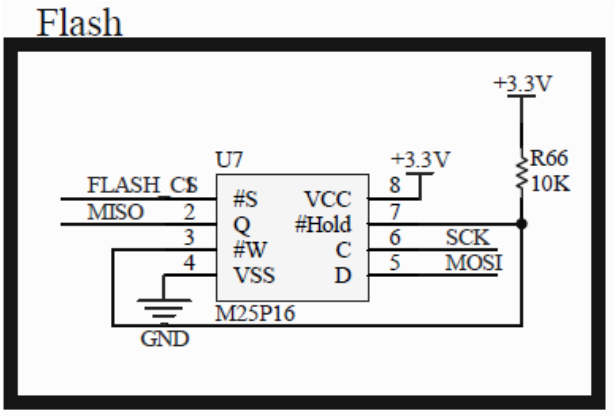

Figure 60: SPI Flash IC schematic.

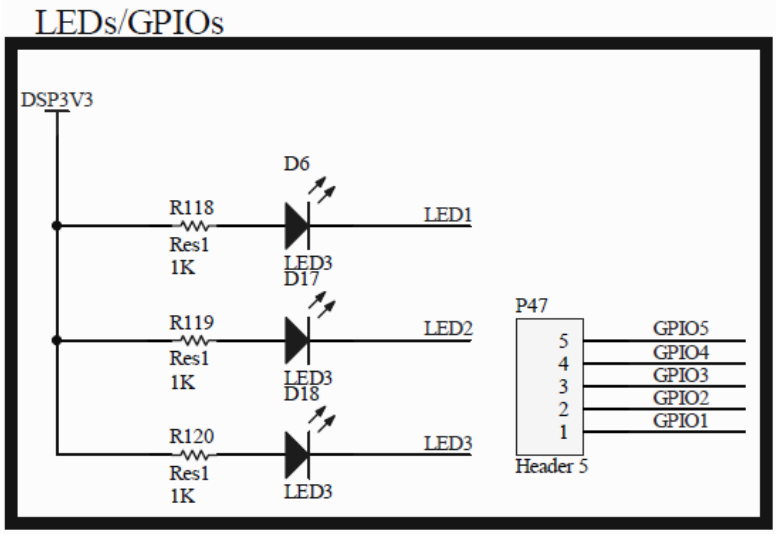

Figure 61: LEDs and GPIOs present in the system board.

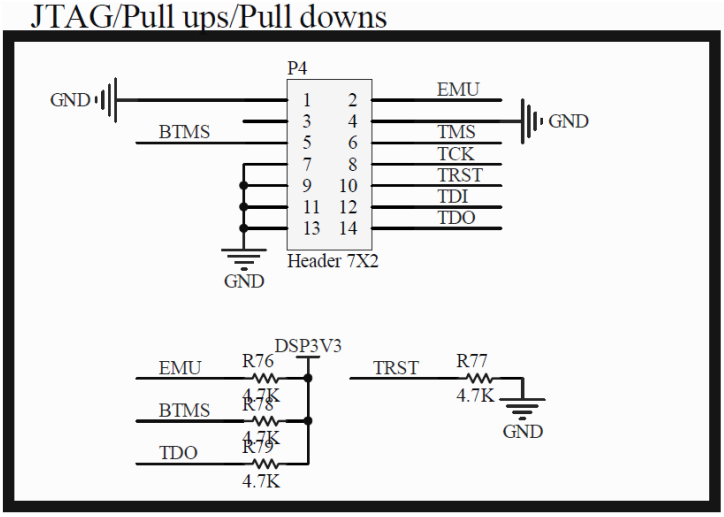

Figure 62: JTAG connector schematic and pin pull ups and pull downs.

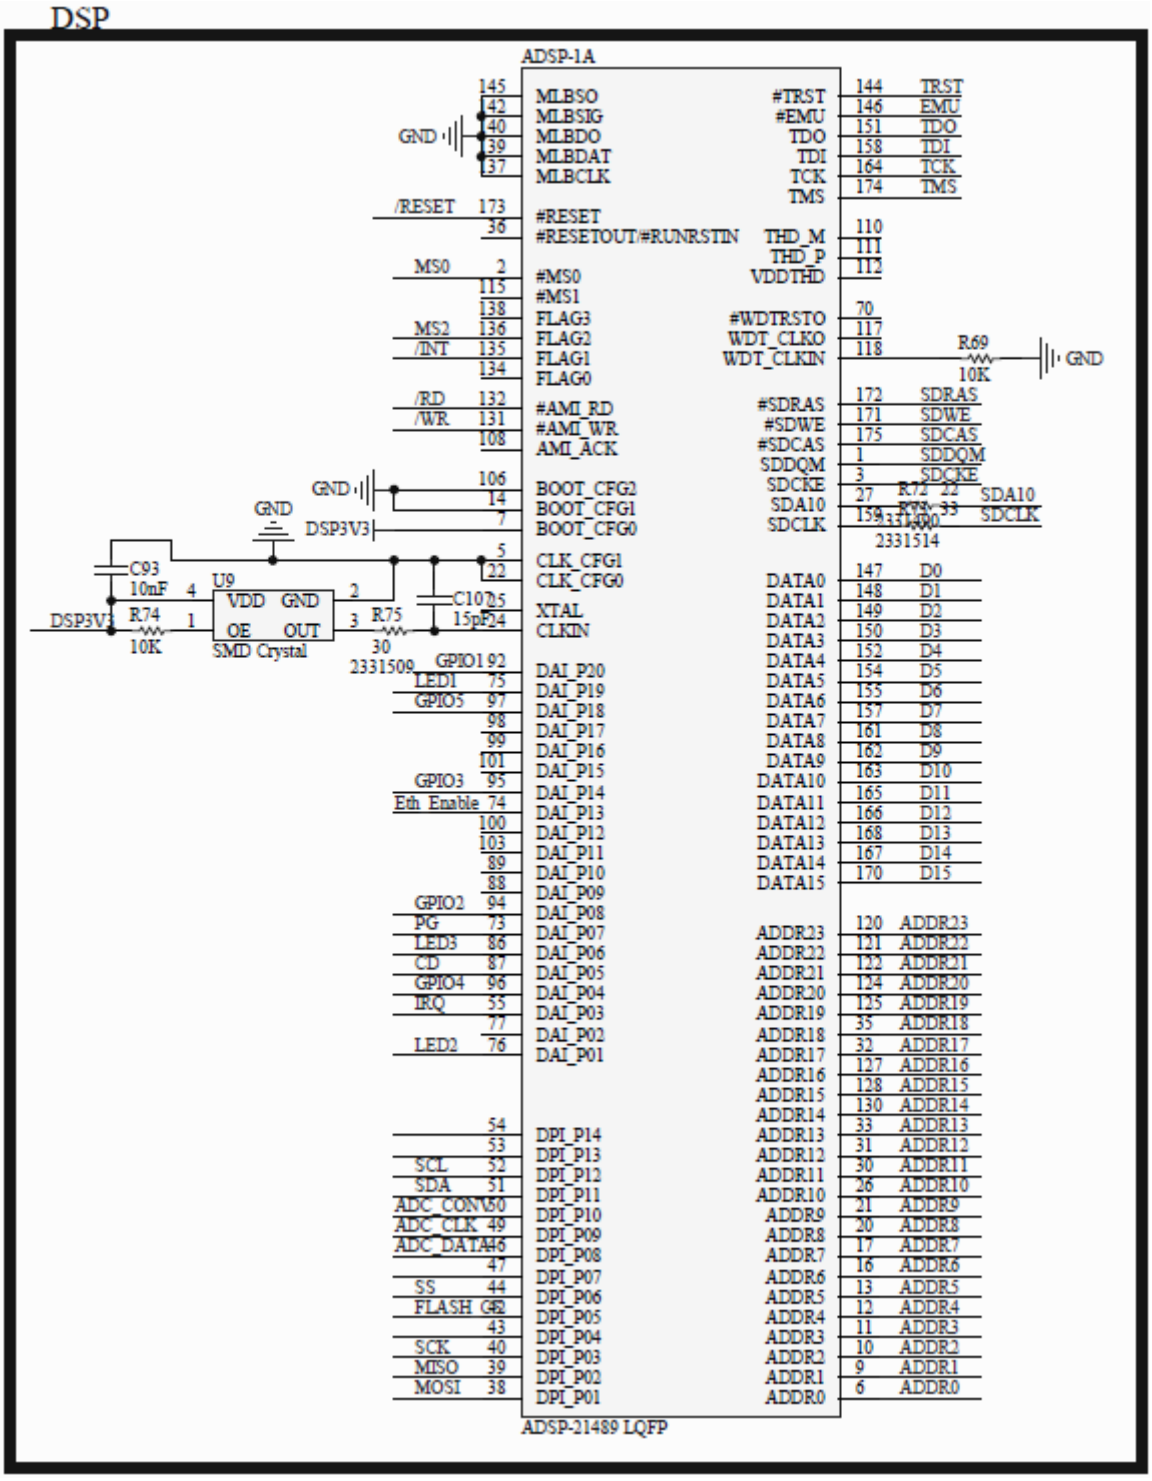

Figure 63: Schematic of the DSP IC.

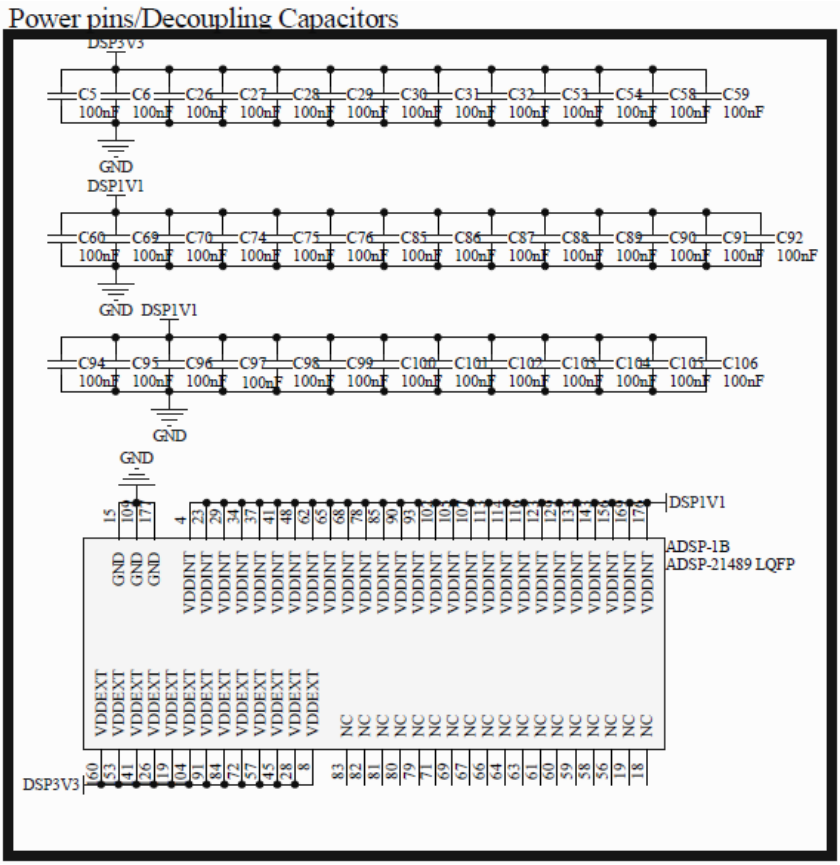

Figure 64: Decoupling capacitors and power pins of the DSP.

C.6. Power Block Schematics – Power Grid

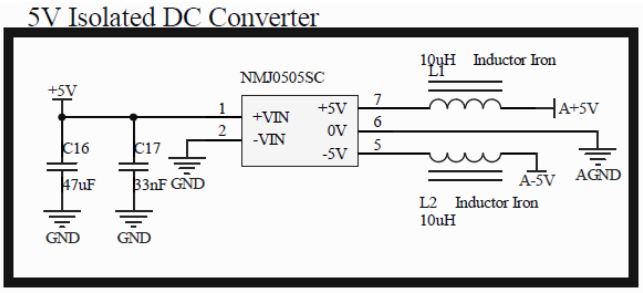

Figure 65: Schematic of the 5V isolated DC-DC converter.

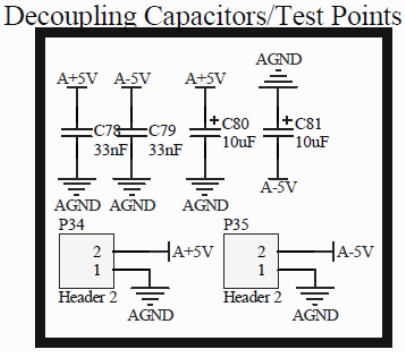

Figure 66: Decoupling capacitors and test points.

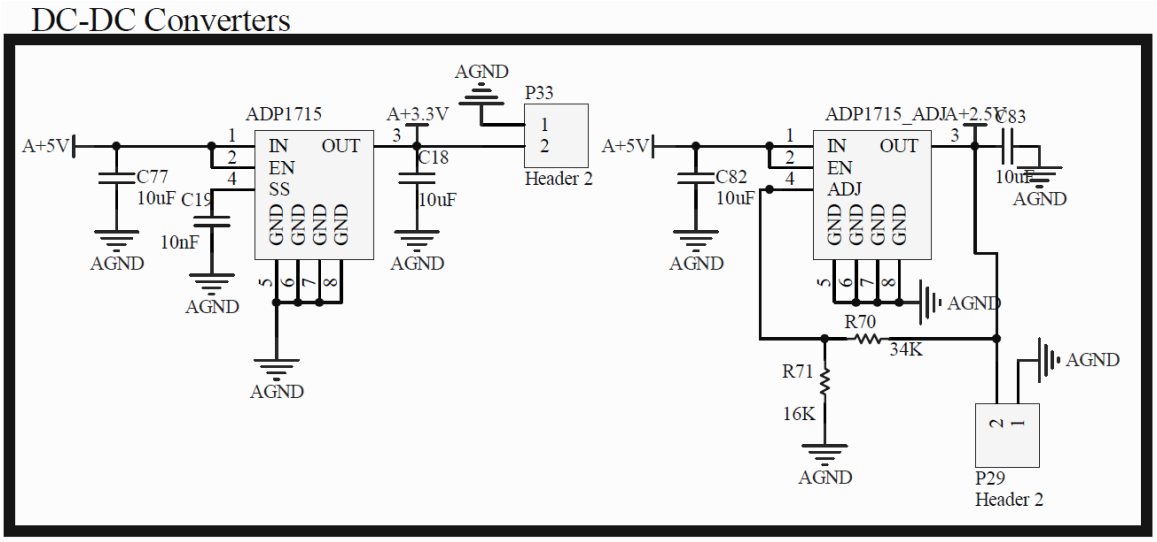

Figure 67: Schematics of the DC-DC converters used and their test points.

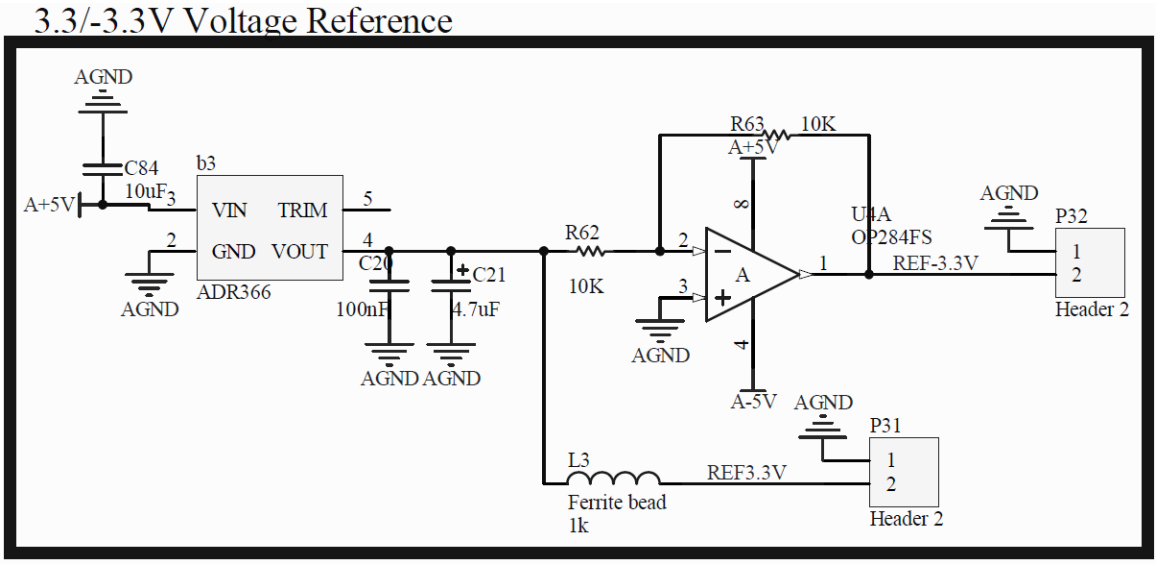

Figure 68: Schematics of the voltage references used and their test points.

C.6. Power Block Schematics – DSP

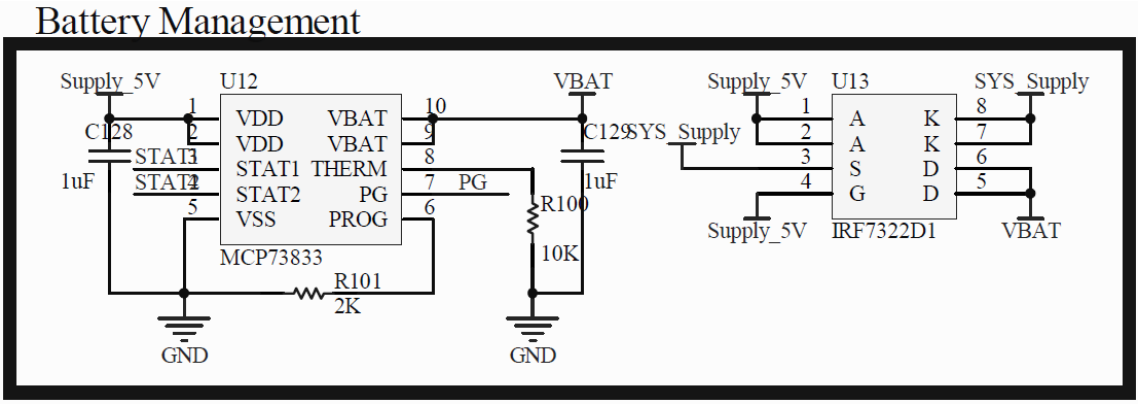

Figure 69: Schematic of the battery management circuit.

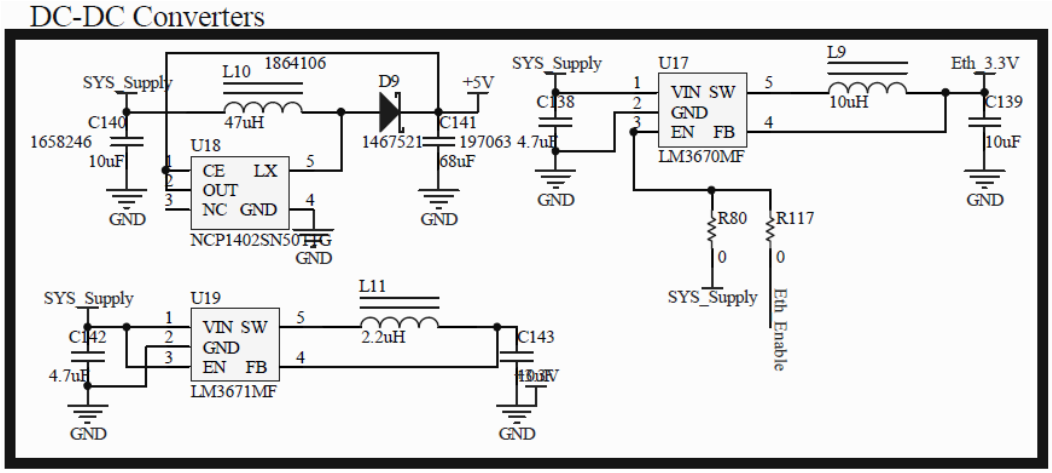

Figure 70: Schematics of the DC-DC converters used.

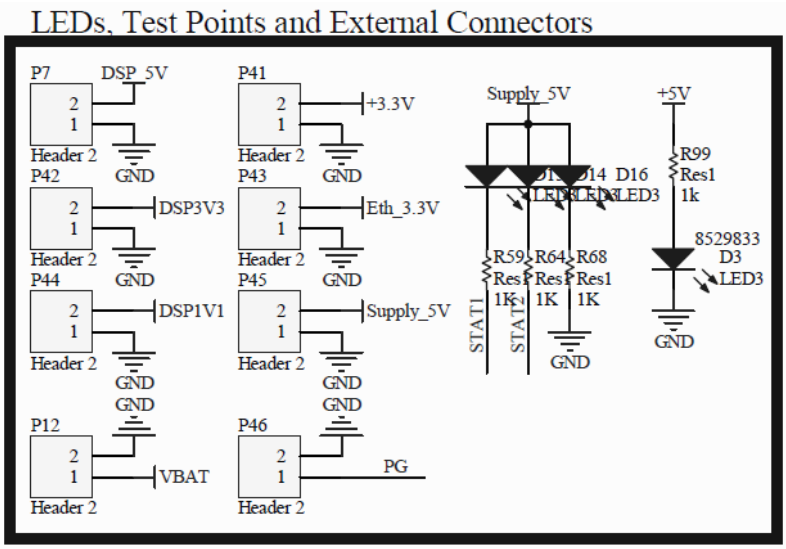

Figure 71: Test points, LEDs and external connectors.

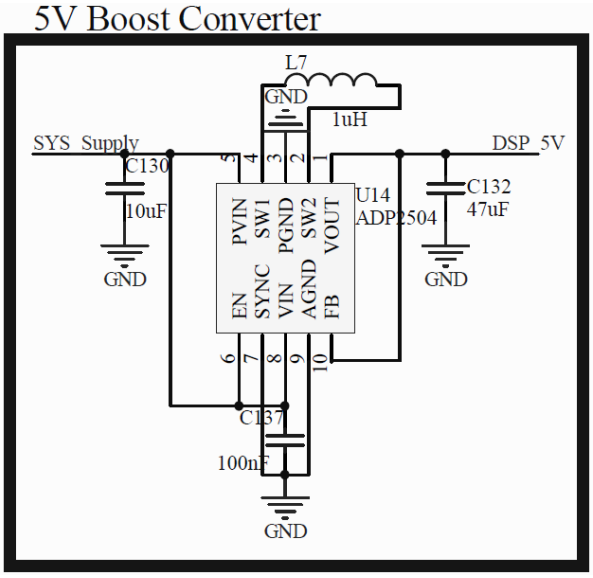

Figure 72: 5V Boost DC-DC converter used to feed the DSP DC-DC converters.

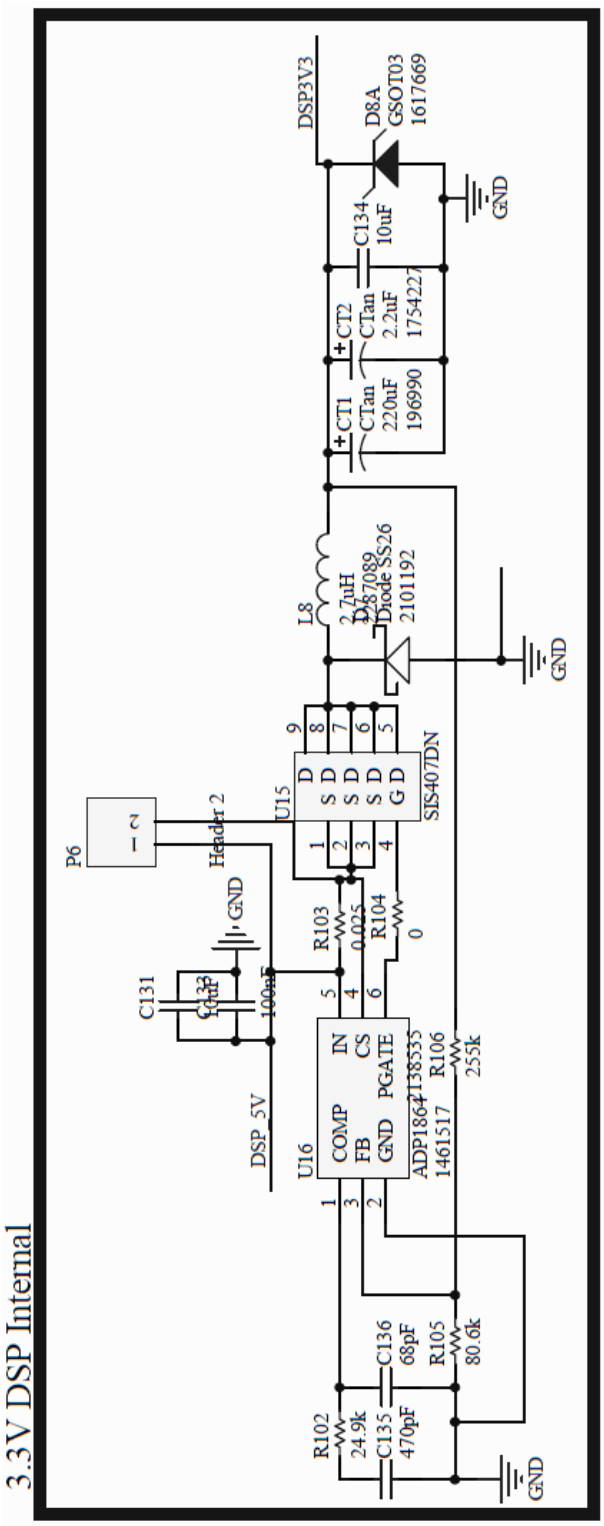

Figure 73: Schematic of the circuit to generate 3.3 V and feed it to the DSP.

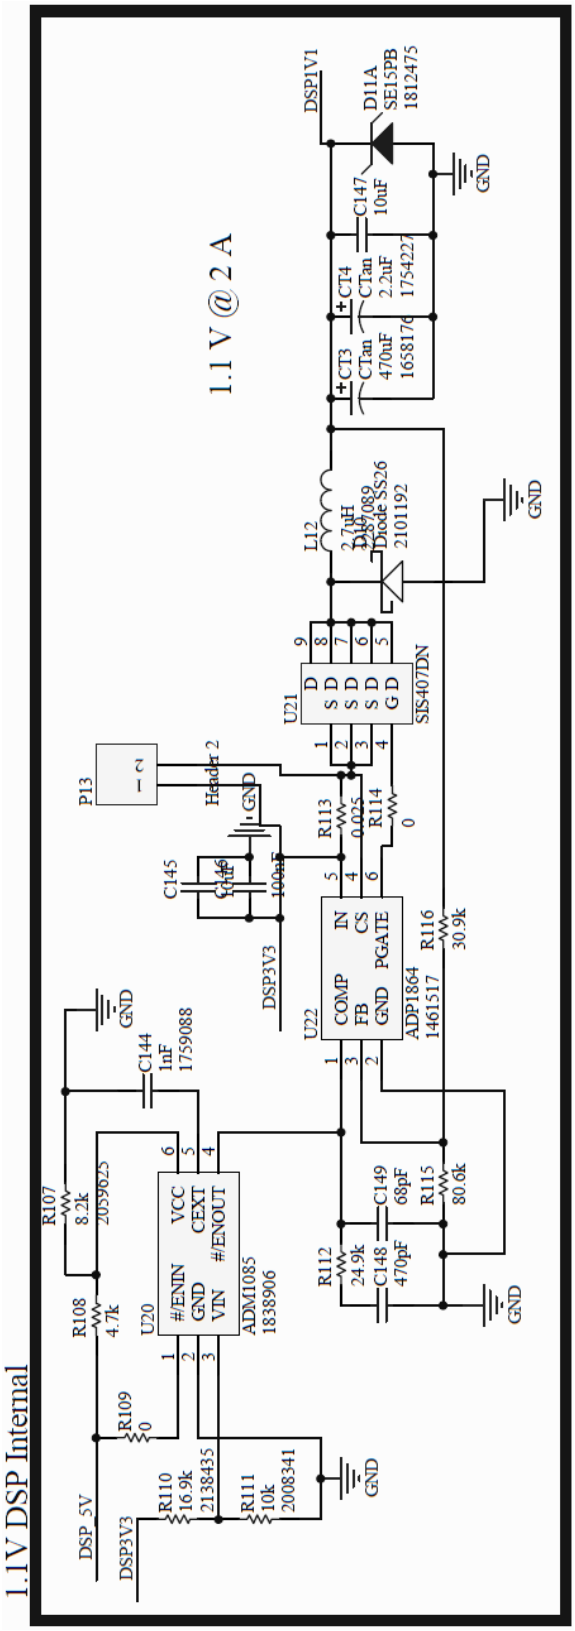

Figure 74: Schematic of the circuit to generate 1.1 V and feed it to the DSP.

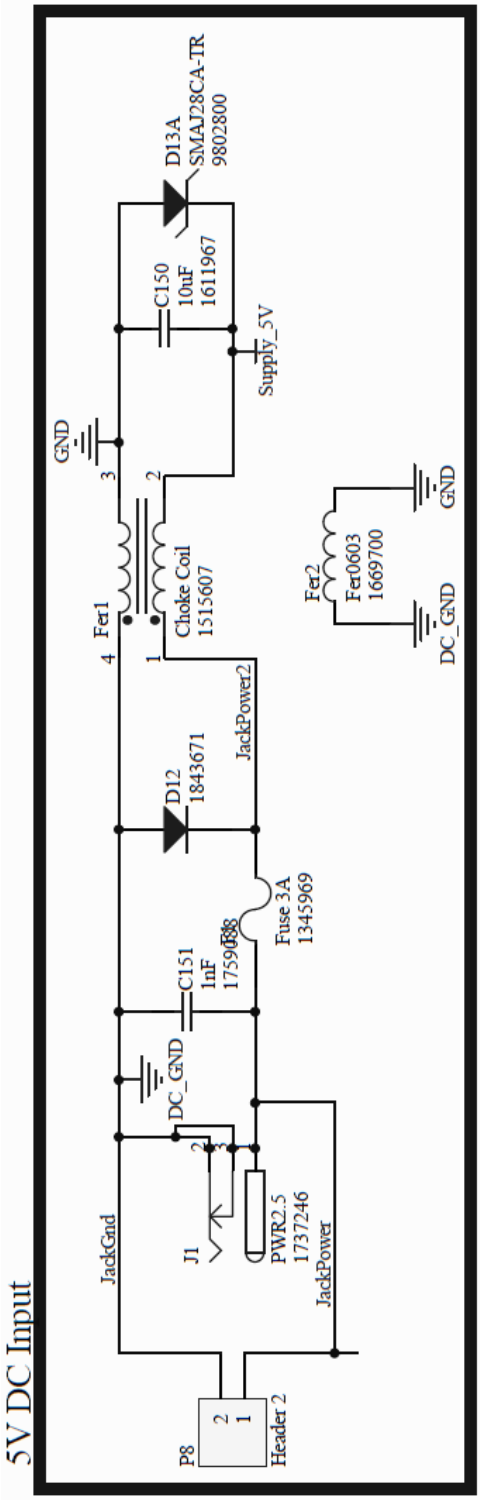

Figure 75: 5 V DC Input circuit schematics.

Appendix D – Developed System

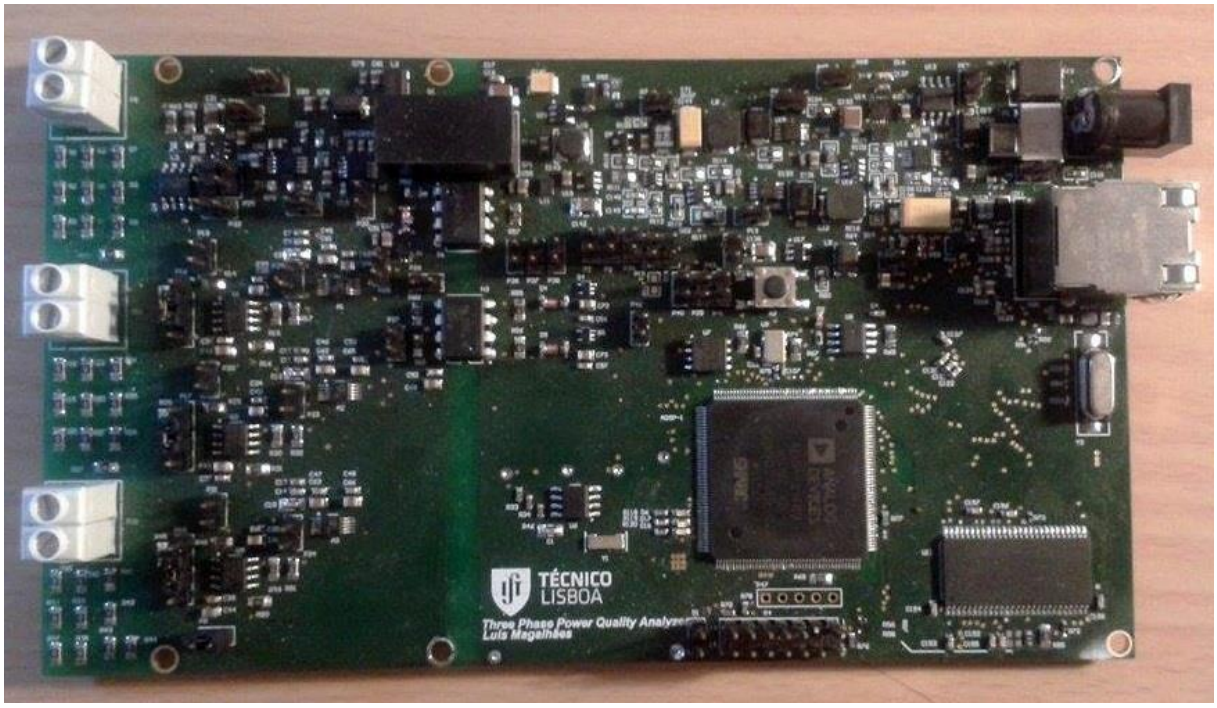

Figure 76: Top view of the main board developed.

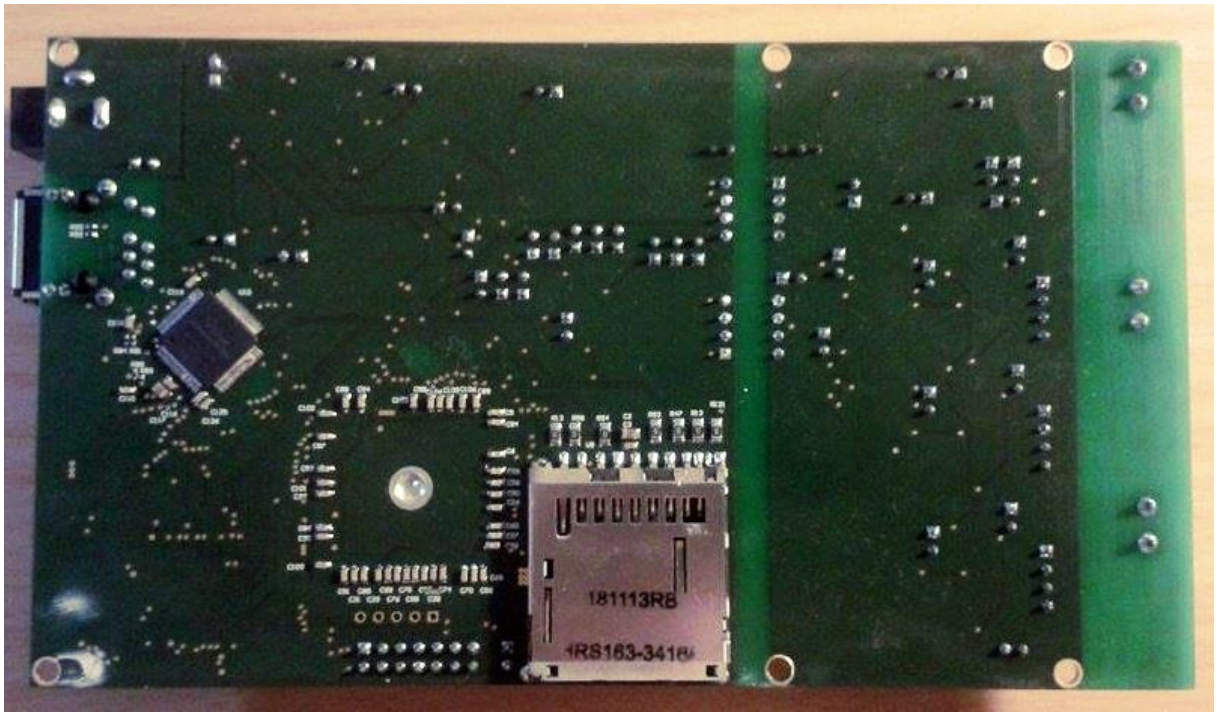

Figure 77: Bottom view of the main board developed.

Figure 78: Top view of the power board developed.

Figure 79: Top view of the complete system.

Figure 80: Another view of the complete system.