

Università degli Studi di Bologna

FACOLTÀ DI INGEGNERIA

Corso di Laurea Specialistica in Ingegneria Elettronica

METODOLOGIE DI PROGETTAZIONE HARDWARE E SOFTWARE LS

**DESIGN OF AN ACTIVE SENSING PLATFORM
FOR WIRELESS STRUCTURAL HEALTH
MONITORING**

Tesi di Laurea di:

DANIELE MUSIANI

Relatore:

Chiar.mo Prof. Ing. LUCA BENINI

Correlatori:

Chiar.ma Prof.ssa TAJANA SIMUNIC ROSING

Chiar.mo Prof. Ing. BRUNO RICCÒ

Anno Accademico 2005/2006

Introduzione

La realizzazione di sistemi di monitoraggio in tempo reale dell'integrità di strutture è oggetto di ricerche di intensità crescente. I benefici di tali sistemi sono molteplici e di enorme portata: la costruzione di strutture intelligenti, in grado di valutare continuamente le proprie condizioni potrebbe prevenire rotture e conseguenti catastrofi e ridurre significativamente i costi di manutenzione ed ispezione. Uno dei metodi più promettenti per la realizzazione di strutture con tali capacità è basato sull'integrazione nelle strutture stesse di materiali "intelligenti", che permettano di effettuare il monitoraggio per via elettronica. Tra i materiali "intelligenti", quelli piezoelettrici hanno di recente attirato l'attenzione dei ricercatori in ragione della loro versatilità, bassa invasività e costo contenuto. I materiali piezoelettrici sono caratterizzati da un accoppiamento tra le loro proprietà elettriche e meccaniche, per cui possono essere impiegati tanto per rilevare quanto per produrre vibrazioni in una struttura. I materiali "intelligenti" -e quelli piezoelettrici in particolare- aprono la strada all'implementazione su larga scala di tecniche di monitoraggio non invasive, che determinano il grado di integrità di una struttura per mezzo della misurazione e dell'analisi della risposta dinamica della struttura stessa ad uno stimolo ambientale o *ad hoc*.

Al fine di monitorare una struttura con un alto livello di accuratezza, si rende di regola necessario installare su di essa sensori ad alte densità. L'installazione di una rete di sensori interconnessi mediante fili può rivelarsi proibitivamente costosa, potendo richiedere modifiche alla struttura stessa –quali perforazioni- e molti mesi per il completamento, in particolar modo su strutture di grandi dimensioni. Le reti di sensori wireless ovviano al problema, garantendo al contempo l'ottenimento delle funzionalità di monitoraggio desiderate. Esse offrono infatti la possibilità di installare sensori sulle più diverse strutture e raccogliere dati in tempo reale in maniera economica e rapida, senza costosi interventi e cablaggi.

Il lavoro svolto in questa Tesi consiste nella progettazione, prototipizzazione e verifica di un nodo sensore wireless per applicazioni di monitoraggio attivo di

strutture. Il dispositivo è stato sviluppato per essere interfacciato ad una griglia di elementi piezoelettrici da installare o integrare nella struttura sotto osservazione e valutarne il grado di integrità secondo metodi di analisi basati su onde di Lamb. Il nodo sensore può pilotare gli elementi piezoelettrici con forme d'onda di tipologia configurabile, rilevare con elevata accuratezza la risposta della struttura alla perturbazione così indotta ed eseguire sui dati acquisiti algoritmi di analisi di complessità elevata, grazie al DSP in dotazione. Il nodo è altresì dotato di un modulo radio ZigBee per garantirne la massima integrabilità in infrastrutture wireless esistenti.

Il nodo sensore è stato inoltre dotato di funzionalità di *energy harvesting*, che lo rendono in grado di funzionare grazie alla sola energia accumulata per mezzo di pannelli solari e immagazzinata in super-condensatori. Questa caratteristica rende il nodo sensore completamente autosufficiente dal punto di vista energetico e, estendendone l'attesa di vita grazie all'eliminazione delle batterie, comporta una riduzione ulteriore dei costi di manutenzione del sistema, grazie al venir meno dell'esigenza di un periodico intervento di rinnovo delle fonti di alimentazione. Per fare fronte ai vincoli di consumo posti dall'uso della sola luce solare per l'alimentazione, si sono implementate soluzioni circuitali *ad hoc* miranti all'ottimizzazione delle risorse energetiche. In particolare, un circuito di *radio triggering* e un sistema di controllo selettivo dell'accensione dei vari componenti.

Le verifiche sul funzionamento del prototipo effettuate hanno dimostrato come esso svolga correttamente ed efficacemente tutte le operazioni richieste. Il nodo sensore oggetto di questa Tesi rappresenta quindi tanto un concreto passo in avanti verso l'implementazione di reti di monitoraggio di strutture, quanto un esempio della coniugabilità di tecniche di *energy scavenging* con alti carichi computazionali.

La Tesi, scritta in lingua Inglese essendo il lavoro stato svolto presso l'Università della California, San Diego, si articola in cinque capitoli principali. Nel primo si trattano aspetti teorici del monitoraggio di strutture. Nel secondo si delinea lo stato dell'arte delle reti di sensori, con specifica menzione alle tecniche di *energy harvesting*. Nel terzo e nel quarto si descrive l'attività di sviluppo hardware e software del nodo sensore. Nel quinto, infine, si raccolgono i risultati dei test effettuati sul prototipo realizzato.

Table of contents

Table of contents	5
List of figures	9
List of tables	13
Abstract	15
1 Structural health monitoring	17
1.1 Motivations for SHM	18
1.2 Damage detection philosophy	19
1.3 The monitoring process	20
1.4 Damage detection techniques	22
1.4.1 Natural frequencies	23
1.4.2 Time-frequency analysis	25
1.4.3 Hilbert transform	27
1.4.4 Impedance-based methods	28
1.4.5 Lamb waves	29
2 Wireless sensor networks	33
2.1 Wireless sensor nodes	34
2.1.1 Mica and Spec motes	35
2.1.2 Tmote Sky	37
2.1.3 Ecomote	38
2.2 Wireless platforms for SHM applications	39
2.2.1 A Stanford wireless node for SHM	40
2.2.2 Duranode	41
2.2.3 A miniaturized wireless node for impedance-based SHM	42
2.2.4 Wisden: a wireless sensor network for SHM	43
2.3 Energy harvesting	44
2.3.1 Solar energy	45
2.3.2 Vibrations to electricity conversion	47

2.3.3	Technologies for energy storage: batteries	50
2.3.4	Technologies for energy storage: capacitors	52
2.4	Energy scavenging solutions	56
2.4.1	Heliomote	57
2.4.2	Prometheus	58
2.4.3	Everlast	59
2.4.4	Ambimax	61
2.4.5	PicoBeacon	62
2.4.6	Other platforms	62
3	Design of a wireless node for active sensing in SHM	65
3.1	Requirements	65
3.2	Architectural design	67
3.3	Actuation circuit	70
3.4	A/D signal conditioning circuit	71
3.5	Energy management circuit	73
3.5.1	Power distribution	76
3.6	Radio triggering circuit	77
3.7	Digital core	80
3.7.1	Microcontroller	80
3.7.2	DSP and DAC	83
3.7.3	Memories	86
3.7.4	ZigBee radio	89
4	Software development	91
4.1	Microcontroller's code	91
4.1.1	Communication protocol	93
4.2	DSP's code	95
4.3	PC-side GUI	97
5	System prototyping and validation	99
5.1	Prototyping	99
5.1.1	Prototype's cost	102
5.2	System validation and characterization	103
5.2.1	Energy management circuit	103
5.2.2	Signal conditioning circuits	105
5.2.3	Radio triggering circuit	107
5.2.4	Digital core	109

5.2.5 Power consumption	111
6 Conclusions	115
Appendix A – Schematics	117
Appendix B – Prototype’s pictures	131
Appendix C – Communication protocol	137
Bibliography	141
Acknowledgements	153

List of figures

Figure 2.1 – Mica2 motes' architecture.	36
Figure 2.2 – A Mica2Dot mote.	37
Figure 2.3 – A Tmote Sky wireless node.....	38
Figure 2.4 – Ecomote wireless sensor node.....	38
Figure 2.5 – Evolution of wireless platforms for SHM.	40
Figure 2.6 – Architecture of the node for SHM by Lynch et al.	41
Figure 2.7 – The Duranode module.	42
Figure 2.8 – Continuous power/cm ³ vs. life for different power sources.	45
Figure 2.9 – Equivalent electrical model and typical I/V characteristic of a solar cell.	47
Figure 2.10 – Scheme of principle of vibrations conversion based on moving inductors.....	48
Figure 2.11 –Charge pump with variable capacitor.....	48
Figure 2.12 – MEMS variable capacitors.	49
Figure 2.13 – Vibrations to electricity conversion by means of piezoelectric materials.	50
Figure 2.14 – Comparison between technologies for energy storage.	51
Figure 2.15 – Life of a battery vs. depth of discharge.	52
Figure 2.16 – Internal structure of supercapacitors.....	54
Figure 2.17 – Life expectancy for supercapacitors.	54
Figure 2.18 – General architecture of energy harvester modules.	56
Figure 2.19 – The energy harvesting module Heliomote.....	57
Figure 2.20 – Heliomote board.	58
Figure 2.21 – Prometheus module.	59
Figure 2.22 – Battery charging circuitry on the Prometheus platform.	59
Figure 2.23 – Block diagram of the Everlast platform.....	60
Figure 2.24 – Buck converter circuit used in the Everlast module.	60
Figure 2.25 – Block diagram of the Ambimax platform.....	62
Figure 2.26 – PicoBeacon wireless module with energy harvesting capabilities.	62

Figure 2.27 – Energy harvesting solution developed at the University of Bologna. .	63
Figure 3.1 – Monitoring system’s organization.....	66
Figure 3.2 – Node’s architecture.....	69
Figure 3.3 – Actuation path’s amplifier.....	70
Figure 3.4 – Actuation path’s selection circuit.....	71
Figure 3.5 – Anti-aliasing filter.	72
Figure 3.6 – A/D amplifier.	73
Figure 3.7 – Thin-film solar panels by PowerFilm Inc.....	73
Figure 3.8 – Operation of the LT3704 DC-DC inverter.	75
Figure 3.9 – Block diagram of the energy circuit.....	76
Figure 3.10 – Transformer-based radio triggering circuit.	78
Figure 3.11 – Radio triggering circuit.	79
Figure 3.12 – Functioning of the radio triggering circuit.	79
Figure 3.13 – The radio-triggering circuit with a cascaded low-threshold MOSFET.	80
Figure 3.14 – Atmel ATmega128 in TQFP package.....	81
Figure 3.15 – TMS320F2812’s functional block diagram.	85
Figure 3.16 – Typical total power consumption vs. clock frequency for the TMS320F2812.	85
Figure 3.17 – DAC902’s functional diagram.	86
Figure 3.18 – 25A640 SPI EEPROM pinout.....	87
Figure 3.19 – CY7C1021CV33 pinout.....	88
Figure 3.20 – XBee radio in the available antenna configurations.....	89
Figure 4.1 – Flow chart of the behavior of the node.....	92
Figure 4.2 – Packet’s structure as defined in the custom protocol.	94
Figure 4.3 – The optimized actuation mechanism.	96
Figure 4.4 – Source code of the optimized actuation cycle and of the timer’s interrupt routine.	97
Figure 4.5 – Screenshot of the PC-side GUI.	98
Figure 5.1 – The node’s structure.	100
Figure 5.2 – Protomat S62 milling machine.....	100
Figure 5.3 – Copper particles produced by the milling machine.....	101
Figure 5.4 – Imperfections caused by milling tools’ deterioration.....	102
Figure 5.5 – The prototype in its final version.....	102
Figure 5.6 – DSP pins after soldering.....	102

Figure 5.7 – Charging of a 100F supercapacitor.....	103
Figure 5.8 – 3.3V (blue) supply’s output after turn-on.....	104
Figure 5.9 – +5V (blue) and -5V (pink) supplies’ output after turn-on.....	104
Figure 5.10 – +10V (blue) supply’s output after turn-on.....	105
Figure 5.11 – -10V (blue) supply’s output after turn-on.....	105
Figure 5.12 – Characteristic of the anti-aliasing filter.....	106
Figure 5.13 – Setup for damage detection tests.....	106
Figure 5.14 – Result of tests on A/D conditioning circuit.....	107
Figure 5.15 – SFDR of output amplification circuit.....	107
Figure 5.16 – SFDR of input amplification and filtering circuit.....	107
Figure 5.17 – The radio-triggering circuit and the antennas used during tests.....	108
Figure 5.18 – The output of the radio triggering circuit.....	108
Figure 5.19 – Results of actuation-acquisition loop. Actuation at 300KHz.....	109
Figure 5.20 – Results of an actuation-acquisition loop. Actuation at 600KHz, damage case.....	110
Figure 5.21 – Results of actuation-acquisition loop. Actuation at 600KHz, damage-free case.....	110
Figure 5.22 – Difference between beam’s responses in damage and damage-free cases.....	110
Figure 5.23 – A data series plotted by Matlab.....	111
Figure 5.24 – 1F 5V supercapacitor’s discharge during actuation at 100KHz, 500KHz and 1MHz.....	112

List of tables

Table 2.1 – Comparison of energy scavenging sources.....	45
Table 2.2 – Characteristics of most common rechargeable batteries.....	52

Abstract

A major focus of the structural engineering community is to develop systems and structures that can monitor their own structural integrity in real time. Besides preventing catastrophic failures, on-line damage detection would reduce costs by minimizing maintenance and inspection cycles. One of the most promising means of developing these self-monitoring structures is through the integration of smart materials into the structures themselves. Among smart materials, piezoelectric materials have drawn the attention of the research community because of their effectiveness, unobtrusiveness and low cost. These materials show a coupling between their mechanical and electrical properties and they can be used both for sensing and producing vibrations in a structure. Smart materials open the door to the implementation of non-destructive evaluation techniques, which assess the structure integrity by measuring and analyzing the dynamic response of a system to an environmental or an ad hoc stimulus.

In order to monitor structures with a high degree of accuracy, it is often required to instrument the structure with a high density of sensors. The installation of a wired sensor network can be prohibitively expensive, since it may require changes to the structure itself (e.g. drilling holes in walls) and take several months in buildings and infrastructure. This issues are addressed by wireless sensor networks, which constitute a cheap and effective way to instrument a structure and collect data without the need for expensive interventions.

In this work, an *ad-hoc* wireless sensor node for active sensing in structural health monitoring applications is designed, built and tested. The node is interfaced to a grid of piezoelectric devices to be embedded in the structure under test and it is provided with all the capabilities needed to perform an integrity assessment by means of the Lamb waves-based method. The node can drive the external piezoelectric elements with custom waveforms, sense the response of the structure with high fidelity and run complex analysis algorithms on the collected data thanks to the incorporated DSP. External communication is insured by a ZigBee radio,

which makes the node ready to be integrated in large-scale sensor networks.

In addition, the node is designed to harvest energy from the environment by means of solar panels. The energy is then stored on super-capacitors. This feature makes the node completely self-sufficient from the energy standpoint, extends its life cycle and reduces the maintenance costs, by eliminating a periodical human intervention to renew its power source. As a consequence of the power consumption constraints placed by energy harvesting, the node is provided with enhanced power management solutions at the hardware level, namely a radio triggering circuit and a selective power distribution network.

The tests run on the prototype show that the node can actually be used to perform damage detection analysis with the required accuracy. Its computational capabilities, effectiveness and low cost make the node described in this work a major contribution toward the implementation of large-scale, unobtrusive wireless sensor networks for structural health monitoring applications.

1 Structural health monitoring

Structural Health Monitoring (SHM) is the process of observing a structure over time, identifying a damage sensitive feature in the observations and performing a statistical analysis of these features to determine the health of the observed structure. Structures include examples from mechanical, civil and aerospace engineering systems. The observations consist of periodically sampled dynamic response measurements from an array of sensors deployed on the structure. Damage sensitive features are extracted through modeling of baseline structural observations, while the subsequent statistical analysis leads to quantitative information about the current state of the observed structure. SHM can be used in one of two modes. In the first mode a SHM system periodically checks the health of a structure and assesses the ability of the structure to perform an intended function as the inevitable effects of aging and degradation set in. In a second mode of operation, a SHM system is used for rapid assessment of structural state after an extreme event, such as an earthquake, to provide near real time information about structural integrity.

The introduction of smart structures into real world applications is one of the most challenging engineering problems today. The problem of SHM is not “one size fits all”. Each structure, component, geometry, or material must often be approached differently. Many different aspects of a complete SHM system must be investigated before such a project can be confidently employed. Once a SHM system is in place, continuous or routine interrogation of the desired structure should yield valuable information into the current state of operation for the configuration.

One of the common challenges facing SHM researches is coming up with a robust damage detection scheme for comprehensive structural evaluation. Many detection methods and algorithms are specifically configured to detect damage in an area very local to the sensor. Conversely, some methods are very good at providing indications of damage in a very broad, or global, area. A combination of these confined and expansive sensing techniques may be necessary for deployment of SHM on a large bridge, aircraft, or edifice. Certain types of defects may cause a

SHM system to respond in a different manner. For this reason, the sensitivity of the sensor and detection algorithms to various kinds of damage should be known. The operator must know what type of signal to expect for a specific kind, or amount, of damage. Without some sort of system calibration, the SHM output signal may indicate little or no damage for a very significant change in the structural condition.

Each individual component of a SHM system, hardware, sensors, and even software, will have to undergo extensive laboratory testing before it can be used in a critical application. After the fabrication of the complete system, experimentation must be performed to ensure the hardware and damage indication results are not adversely affected by any harsh environmental conditions. A mature SHM system has the potential to be fully integrated into the operational hardware of a host structure. If a defect is sensed, the system should have the intelligence to determine what the best course of actions should be, whether it is to shut the entire assembly down, or proceed with caution until the area can be inspected.

1.1 Motivations for SHM

People depend upon a vast amount of aerospace, civil and mechanical infrastructure of which little is known about the structural state. In the civil engineering sector there are over 600,000 highway bridges in the United States all in varying structural states. There are buildings that can undergo seismic loading that can potentially incur damage in moment resisting steel joints. In the aerospace industry airplanes undergo maintenance on a time-based schedule. Composites, which are hard to diagnose visually, are becoming the building material of choice for aerospace and some civil structures. In each of these areas monitoring the health of the system in a more analytical way and on a more regular basis could lead to improved life-safety benefits and economic savings.

On May 23, 2004, after only a year of operation, the vaulted roof on the newest terminal of Paris' Charles de Gaulle airport collapsed. The cause of the collapse was attributed to weak concrete being perforated by the supporting struts. Apparently, the concrete slowly deteriorated over time.

In a 2001 report, the Federal Highway Administration stated that 5% of the United States' Interstate bridges are structurally deficient, while 16% are functionally obsolete. Structurally deficient bridges are defined as bridges which may have

vehicle weight restrictions or may need to be completely closed for repairs before being reopened. Functionally obsolete bridges may be structurally sound, but do not match the conditions of the surrounding interstate. As the class of roads decreases from Interstate to smaller thoroughfares, the percentages of deficient or obsolete bridges increases. Pyles (2003) reports that, unlike in previous generations, the US Air Force can no longer rely on the retirement of aircraft to decrease maintenance costs. Many of the aircraft in operation are being utilized well beyond their designed service life. In addition, the capital required to train maintenance technicians and to acquire the proper repair materials is ever increasing. Finally, it may take 20 to 35 years to fully replace a fleet of aging aircraft when you include the lead times for productions. These are only three of the many examples where an effective structural health monitoring system can decrease maintenance costs and downtime. With a SHM system installed on aging aircraft, scheduled maintenance could be reduced to a minimum, or only be performed when the monitoring system indicates cause for inspection. Bridges can be continuously monitored for defects, and damage can be repaired in incipient stages before the bridges are deemed structurally deficient. Monitoring the integrity of buildings can prevent unnecessary catastrophes leading to loss of human life, or simply hasten post-earthquake inspections. Most importantly, the safety of the people operating or occupying structures with SHM systems will be significantly improved.

1.2 Damage detection philosophy

Rytter (1993) introduced a damage state classification system which has been widely accepted by the community dealing with damage detection and SHM. Following these lines, the damage state is described by answering the following questions (Sohn et al. (2003)):

- Is there damage in the system? (existence)
- Where is the damage in the structure? (location)
- What kind of damage is present? (type)
- How severe is the damage? (extent)
- How much useful life remains? (prognosis)

Generally, identification of the damage type and extent require prior knowledge of the structural behavior in the presence of each of the possible expected

failure modes for future correlation with experimental data, which is normally achieved by resorting to analytical models. For example, in operational monitoring, the modal parameters of the damaged structure must be compared to the parameters of the structure in its undamaged state, in what is called global diagnostics.

Once damage existence is detected, the use of a model of the structure in a damaged state may be used to determine the damage location, in what is called local diagnosis. Some difficulties may be encountered in the practical application of modal models, such as the knowledge of excitations and loads during machine operation, with several sources and with unknown distribution along the system. However, several output-only modal identification techniques can be found in the literature; for example Guillaume et al. (1999), Brincker et al. (2003), Rodrigues et al. (2004) and Galvin and Dominguez (2005). Prognosis, which is traditionally related to fracture mechanics and fatigue, is starting to be brought up by the modal analysis community as a field of interest, as seen in Farrar et al. (2003).

1.3 The monitoring process

The monitoring process involves the observation of a system over long periods of time, using samples of experimental data acquired periodically with adequate sensors. It also involves sensitivity feature extraction and statistical correlations to determine the actual 'health' of the system. Farrar et al. (2001) and Sohn et al. (2003) acknowledge the problem of damage detection in the context of a statistical pattern recognition paradigm. Along these lines, one may describe this paradigm as a four-step process:

- 1) Operational evaluation.
- 2) Data acquisition, fusion, and cleansing.
- 3) Feature extraction and information condensation.
- 4) Statistical model development for feature discrimination.

Operational evaluation answers questions related to the damage detection system implementation, such as economic issues, possible failure modes, operational and environmental conditions and data acquisition related limitations. For example, Aktan et al. (2000) present and discuss the prerequisite issues in creating a successful

monitoring program. Ruotolo and Surace (1997) were among the first to explicitly address the possibility of having multiple baseline configurations, proposing a technique to distinguish between changes caused by working conditions and resulting from damage, based on the singular value decomposition (SVD). More recently, Vanlanduit et al. (2005) also used a SVD based system to detect damage in structures subjected to different operational conditions, including different working excitation levels, geometrical uncertainties and surface treatments.

Data acquisition is concerned with the quantities to be measured, the type and quantity of sensors to be used, the locations where these sensors are to be placed, sensor resolution, bandwidth, and hardware. This part of the process is application specific and heavily dependent on economic factors. Another consideration is the periodicity of data acquisition. For instance, if the growth of a fatigue crack is to be monitored, then the data must be measured in an almost continuous way. On the other hand, if measurements are to be made under varying operational and environmental conditions, data normalization helps distinguish signal changes caused by operational and environmental conditions from those caused by damage. Sohn et al. (2001a) observed that the natural frequency of a bridge over the Alamosa canyon in the state of New Mexico in the USA suffered 5% deviations over a 24 hour period due to temperature variations. Sohn et al. (2003) acknowledged that data fusion first appeared as a result of defense organizations attempting to formalize procedures for integrating information from diverse sources, with the purpose of determining battlefield situations and preventing threats. Data fusion, as a discipline of SHM, is the ability to integrate data acquired from the various sensors in the measurement chain. Data cleansing is the process of selecting significant data from the multitude of information, i.e., the determination of which data is necessary (or useful) in the feature selection process.

Within the paradigm defined by Farrar et al. (2001) and Sohn et al. (2003), the field of damage detection which has received the highest attention is feature extraction. Feature extraction is the process of identifying damage sensitive properties, which allow one to distinguish between the damaged and undamaged structural states. Information condensation becomes increasingly advantageous and necessary as the quantity of data increases, particularly if comparisons are to be made between sets of data obtained over the life cycle of a system. Condensation may be seen as a form of data fusion. The evolution of the philosophy of damage detection and SHM over the last ten years can be best understood by comparing the

surveys of Doebling et al. (1996) and Sohn et al. (2003). In the first survey, the statistical pattern recognition paradigm, around which the second investigation was based, is not even mentioned. By the time of the latter, on the other hand, the concepts of operational evaluation, data acquisition, fusion and cleansing and statistical model development for features discrimination were matters whose relevance was not deemed to be noteworthy of exclusive and directed research, and were addressed only as parts of the feature extraction process.

Statistical model development for feature discrimination is the part of the SHM process that had received the least attention prior to the end of the last century. When data are available from both the undamaged and damaged structure, the statistical pattern-recognition algorithms fall into the general class referred to as supervised learning. Unsupervised learning refers to the class of algorithms that are applied to data not containing examples from the damaged structure (obtained, for instance, from models) (Sohn et al. (2003)). An important issue in the development of statistical models

is to establish the model features sensitivity to damage and to predict false damage identification. In this context, one may have either false-positives, in which damage is identified though it did not happen, or false-negatives, in which damage is not identified though it is present. Despite the fact that both these situations are undesirable, the second one can obviously lead to more severe consequences.

1.4 Damage detection techniques

Currently, an immense number of techniques exist for the identification and location of damage. Because all the techniques have their own advantages and disadvantages, there is no general algorithm that allows the resolution of all kinds of problems in all kinds of structures. Every technique tends to have damage related sensitivities; i.e., a very sensitive technique may produce false-positives, while a less sensitive technique may lead to false-negatives, the latter case being more problematic. Generally, only damage above a certain size can be detected. The quantification of damage and prediction of the remaining lifetime are beyond any doubt the most difficult issues, as well the most desirable capabilities. A high-level classification of damage detection techniques can be done based on the following distinctions between approaches: linear vs. non-linear, output only vs. input/output,

on-line vs. off-line, time domain vs. frequency domain, parametric vs. non-parametric and time varying vs. time-invariant.

In the next paragraphs the main damage detection techniques are reviewed.

1.4.1 Natural frequencies

The development of modal analysis techniques for damage detection and SHM arose from the observation that changes in the structural properties have consequences on the natural frequencies. Nevertheless, the relatively low sensitivity of natural frequency to damage requires high levels of damage and measurements made with high accuracy in order to achieve reliable results. Moreover, the capacity to locate damage is somewhat limited, as natural frequencies are global parameters and modes can only be associated with local responses at high frequencies.

Methods based on natural frequency shifts often fall into one of two categories: the forward and the inverse problem. The forward problem consists in determining what the natural frequency changes due to a known damage case (which may include its location, extension and type) will be. Typically, damage is modeled numerically and the natural frequencies are measured experimentally and compared to those related to each of the damage cases initially predicted. The inverse problem consists of determining damage parameters, such as crack length or location, from changes in the natural frequencies.

Lifshitz and Rotem (1969) present what may be the first journal article to propose the use of vibration measurements for damage detection. They search for changes in the dynamic moduli, which can be related to shifts in the natural frequencies, to detect damage in elastomers.

Stubbs and Osegueda (1990) developed damage detection methods based on modal changes of specific structural components such as beams, plates and shells. The method relates frequency shifts to changes in member stiffness using a sensitivity relationship. Stiffness reductions can be located solving an inverse problem, since damage is defined as a stiffness reduction of one of the elements forming the structure. They concluded that it is possible to locate multiple damage sites, at least in a beam. False positives occurred, however, although generally at a far smaller number of sites than those where damage was actually present. Moreover, this method has difficulties when the number of modes is lower than the number of

the damage parameters.

Sanders et al. (1992) use the frequency sensitivity method developed by Stubbs and Osegueda (1990a), combined with an internal-state-variable theory, to detect, locate and quantify damage in CFRP beams. This method includes parameters which indicate two types of damage: matrix micro-cracking (identified by changes in the extensional stiffness), and transverse cracks in the 90-degree plies (by changes in the bending stiffness). Despite the sensitivity equations used only being valid for viscous damping, they argued that the damping had little influence in the application of the methodology. Due to the fact that damage was distributed uniformly along the beam length, it was not possible to evaluate the ability of the method to locate damage.

Messina et al. (1996) showed that it is possible to locate damage more accurately if the frequency shifts are normalized relative to the undamaged structural frequencies, so that one may weigh the contribution of each vibration mode, since higher frequencies tend to suffer more pronounced shifts. The multiple damage location assurance criterion (MDLAC), proposed by Messina et al. (1998), is an extension of the DLAC to detection of multiple damage sites. Using the same principles, the damage state is indicated by search of a damage vector which maximizes the MDLAC. This formulation allows for relative quantification of the damage in each location, but not for absolute quantification. However, since the experimental frequency shift vector can be calculated, it is possible to estimate a scalar severity coefficient C , in such a way that gives the actual damage present. This scaling constant C can be estimated by either a first or second order approach, but these authors note that, although the second order approximation provides better results, for routine monitoring purposes where precise knowledge of the defect size is less important than its location, the first order approximation is likely to be an adequate choice due to its relative simplicity.

Boltezar et al. (1998) devised a method for locating transverse cracks in flexural vibrations of free-free beams by employing an inverse problem. The method is based on the assumption that the crack stiffness does not depend on the frequency of vibration (i.e., the values of the crack stiffness, which is modeled as a linear torsional spring, must be the same at the crack position for all of the measured natural frequencies). As a result, by plotting the relative stiffness along the length of the beam for at least two distinct natural frequencies, the crack location can be identified by the intersection of these curves. These authors mentioned that this idea

had already been proposed by Adams et al. (1978) for axial vibrations, and was extended in their article to include flexural vibrations in one-dimensional beams, allowing better results to be obtained. Since this model is based on the Bernoulli-Euler beam theory, the authors pay special attention to the problem of accurately knowing the values of the material properties, in particular the Young's modulus, which can be overcome by calculating its 'effective' value as suggested by Adams et al. (1978).

Sampaio et al. (2003) and Sampaio and Maia (2004) propose the detection and relative damage quantification indicator (DRQ), based on the use of the frequency domain assurance criterion (FDAC), as an effective damage indicator, capable of distinguishing a positive occurrence from a false alarm. The column vector j of the receptance matrix is the operational deflection shape (ODS) which describes the shape (in space) exhibited by the structure at each excitation frequency, given by the responses normalized by the applied forces. When a structure is damaged, its stiffness and damping change and, in consequence, so does the receptance matrix. Thus, it is reasonable to assume that the smaller the degree of correlation between the column vectors (ODSs) of and , where the superscript d stands for damaged, the larger is the damage.

1.4.2 Time-frequency analysis

Zabel (2004) applied a wavelet-based damage indicator to damage detection on reinforced concrete structures, using several case-studies, including the Z24-bridge in Switzerland. The indicator is based on the analysis of signal energy components in discrete time-scale domain. Normally, the wavelet coefficients of the impulse response function (IRF) make different contributions to the system response under consideration. Thus, if the excitation force is known, it is easy to normalize the response relative to the excitation force. Another normalization possibility is to relate each response to the response in a reference location, thus leading to the transmissibility function. According to this author, it is possible (and adequate) to consider the transmissibility functions and their wavelet decompositions in the context of SHM when the ambient excitation is unknown. However the analysis so far is based on the assumption of similar excitations for all the observations, which means that further research is needed to evaluate how this restriction can be

surpassed without the need for a known excitation.

Savov and Wenzel (2005) used a wavelet approach in order to locate damage in civil structures, based on the system's acceleration time history responses. A finite element (FE) model of a three story shear-resisting frame excited by white-noise ground acceleration and a pre-stressed reinforced concrete test beam under impact loading conditions were used as test structures. The first level fast wavelet decomposition of the measurements for both the undamaged and damaged structures was applied using the Haar mother wavelet, and the approximation and detail coefficients were obtained. Reconstruction of the actual system response was accomplished by a cross combination of the approximation coefficients of the undamaged structure with the detail coefficients of the actual measurement.

The sensitive feature is defined as the standard deviation of the error between this reconstructed signal and the actual measured signal. This method is based on the premise that the detail coefficients carry information about the local structural integrity in the time history response at the damaged sites. The advantage of this method is that one does not need to know the excitation mechanism, since the method is entirely based on the measured responses. However, damage location is compromised by the sensor mesh refinement. Extension of this approach to quantification of damage, to prediction of remaining structural life and to evaluation of more complex structures, varying operational and environmental conditions and non-linear damage phenomena are issues to which these authors will give further attention.

Li et al. (2005) present a methodology for determining crack location and size in a beam, taking advantage of wavelet finite element methods (WFEM). The idea of the WFEM is to discretize a body into wavelet finite elements which are interconnected at nodal points on the element boundaries. Some of the advantages of WFEM over traditional FE methods for modal analysis of crack problems are pointed out. The natural frequencies of the beam with various crack locations and sizes are obtained, through the use of a WFEM, and used as features in the damage detection process. The accuracy of the method was tested on the experimental data of a beam studied by Silva and Gomes (1990). Finally, these authors suggest that this procedure could be easily extended to complex structures with multiple cracks.

1.4.3 Hilbert transform

Testa (2005) argues that the Hilbert-Huang transform (HHT) is very useful for non-stationary, non-zero mean and non-linear real signals. The HHT makes an Empirical Mode Decomposition (EMD) of the time signal into narrow band components with zero mean, in which each component is called an Intrinsic Mode Function (IMF). Unlike their equivalents in the FFT, these components do not have a specific analytical representation, but each component can be associated with a physical meaning, i.e., the components are related to the mode shapes and existence of damage.

Lin et al. (2005) use the HHT technique on the phase I IASC₁-ASCE₂ benchmark problem for SHM. This benchmark has been developed to simplify the side-by-side comparison of various analysis techniques for damage identification in structures on a common basis. The scale-model structure is a 3.6 m high four-story steel frame with 2.5×2.5 m cross-sectional area. Johnson, E. A. et al. (2004) present two analytical models of the structure, both finite element based: one is a 12-DOF shear-building model, the other a 120-DOF model.

Based on the noisy acceleration data (due to ambient excitation), Lin et al. (2005) confirm that HHT is capable of identifying the natural frequencies, damping ratios, mode-shapes, stiffness matrix and damping matrix of a structure with a reasonable accuracy. Damage has been identified by comparing the stiffness of each floor before and after introducing damage. Simulation results show that the HHT technique can detect, locate and quantify damage with good accuracy. Finally, these authors acknowledge that, despite the fact that normal modes are assumed to exist in the current benchmark model, the technique can also be used on linear structures with complex eigenproperties, as demonstrated by Yang et al. (2003).

Zhang et al. (2005) presented the implementation of a method for nonlinear, non-stationary data processing, namely HHT in traditional vibration-based approaches to characterize structural damage. Also, the HHT is compared to the Fourier transfer function technique in detecting local damage with a computer model and experiments in two pile foundations of the Trinity River Relief bridge in Texas, USA. The HHT analysis showed a more significant frequency downshift than Fourier-based approaches for measurements made near the damage location. These authors observed one point which, although only applicable to very specific structures such as bridges, remains of interest: by selecting two or more similar

structural members (e.g., two of four columns with the same size, cross section, and construction in a bridge), the need for data from a reference state may be avoided if only one of these members has experienced damage.

Shi and Law (2005) present a HHT-based technique for stiffness and damping parameter identification for linear time-varying discrete dynamic systems and nonlinear systems respectively. Both procedures are tested using analytical models in free vibration. The researchers identify the non-linear system using a so-called skeleton linear model (SLM), which has similar dynamic behavior to the corresponding nonlinear system.

1.4.4 Impedance-based methods

Park et al. (2003) present an overview of piezoelectric impedance-based health monitoring where the hardware and software issues are summarized, including a discussion of future research areas and the path forward. Park et al. (1999) introduced an impedance-based method which uses a piezoelectric transducer (PZT) as both actuator and sensor. It can be shown that the electrical impedance of the PZT is directly related to the structure's mechanical impedance (Liang et al. (1994)). Thus, a relationship can be found to allow determination of the structural properties; specifically, the mass, stiffness and damping. Any change in the mechanical impedance, which could be caused by the presence of damage, will show up in the electrical impedance of the PZT. The damage state is identified when a defined metric, given by the sum of the squared differences of the impedance between the current and reference states, increases above a predefined threshold. These authors point out that this technique possesses the twin advantages of not requiring numerical models and the fact that its high-frequency excitation nature makes it very sensitive to local structural modifications.

Grisso et al. (2004) used an impedance-based SHM system to detect the onset of transverse matrix cracking in cross-ply graphite/epoxy composite. Ceramic PZT patches are attached to the composite samples (in this case, beams) to simultaneously excite the structure with high-frequency excitations (two frequency ranges, from 10 to 20 kHz and from 40 to 60 kHz) and monitor any changes in structure's mechanical impedance. In order to detect damage severity, these authors use the root mean square deviation (RMSD), using the squared differences of the impedance between

the undamaged and actual states. These authors conclude that acoustic emissions testing showed a good correlation only for thicker specimens, while the impedance-based method gave better correlation for thinner specimens.

Moura and Steffen (2004) tried to identify the best frequency bands for impedance-based structural damage identification in flexible structures. Damage was simulated in a cantilevered aluminum beam by adding a mass to the free end of the beam. By observing the results, these authors agreed that, at least for the conditions used, the best low frequency value is 15 kHz, while the best bandwidth value is to the intermediate one (4 kHz).

Peairs et al. (2005) developed a new modeling technique for impedance-based SHM, combining the spectral element method (SEM) with electric circuit analysis. These authors argue that SEM has several advantages over other conventional FE techniques, for example because SEM more accurately models higher frequency vibrations because the mass is modeled ‘exactly’ and incorporates higher order models more easily. Sensor multiplexing for transfer impedance and high frequency modeling were also investigated. However, modeling at higher frequencies is more complex because of the presence of the PZT resonances.

Simmers et al. (2005) applied an impedance-based method to detect and quantify the onset and growth of pre-crack surface corrosion. Three metrics were used: one similar to the one used by Park et al. (1999), the RMSD and a crosscorrelation. The experimental results on an aluminum beam indicate that the impedance-based method is effective for corrosion detection and tracking, and that there is a relationship between the metrics and corrosion location, surface coverage and pit depth.

1.4.5 Lamb waves

One approach that has received much attention in the last few years is the use of Lamb waves. Lamb waves are a type of elastic perturbation that can propagate across large areas of a free-free solid plate with low dispersion of energy, even in materials with a high attenuation ratio. This type of wave was first described in theory by Lamb (1917), although he never tried to produce them. Alleyne and Cawley (1992) were among the first to discuss interaction of Lamb waves with defects for non-destructive testing. Saravanos et al. (1994) presented a procedure for

delamination detection in composite materials using Lamb Waves and embedded piezoelectric sensors. Kessler et al. (2002) maintain that techniques using Lamb waves have proven to provide more information about damage type, severity and location than previously tested methods using FRFs, since Lamb waves are more sensitive to local structural defects. Piezoceramic patches were used to excite the first anti-symmetric Lamb wave (A_0 mode). The PZT actuators were chosen because of their high force output at relatively low voltages and also due to their good response qualities at low frequencies. Kessler et al. (2002) explored the optimization of Lamb wave methods for damage detection in composite materials, covering the problems of choosing the appropriate actuating frequency, pulse shape and sensor geometry for Lamb wave application. The results were compared by performing a wavelet decomposition using the Morlet wavelet, and plotting the magnitude of the coefficients at the driving frequency. Although Lamb waves display great capabilities in damage detection and localization, these authors point out that the major disadvantage of this method is that active sensing is necessary for its implementation; i.e., it requires a voltage supply and a generated signal. The high data acquisition rate needed to gain useful signal resolution is also an awkward requirement. Finally, the Lamb wave method should most likely be placed into a SHM system in conjunction with other passive detection methods, such as frequency response function (FRF) method, in order to conserve power and data storage space and because the Lamb wave data can be more difficult to interpret.

Su et al. (2002) state that noteworthy efforts have been made since the 1980s in the field of wave propagation based identification methods. Nevertheless, it is known that the propagation characteristics of Lamb waves are relatively complicated because of the dispersion phenomenon. These authors studied the interaction between the fundamental Lamb wave modes and delamination in carbon fiber-reinforced laminates. Models (FE and experimental) of a plate with a delamination were tested, using PZT wafers as sensors/actuators. The calculated and measured dynamic responses were processed using wavelet transform-based analysis in the time-frequency domain for the purpose of damage evaluation, which has proven to be very effective in diminishing the influence of broadband noises and structural vibration patterns and also showed that the response in the time-frequency space is noticeably sensitive to the delamination location.

Lee and Staszewski (2002) discussed acoustic and ultrasonic wave propagation modeling techniques, focusing on two-dimensional wave interactions

with defects in metallic structures.

Ricci et al. (2004) presented a methodology for automatic damage identification and localization in composite structural components. The damage is inflicted in the form of an impact, and identification is based on the calculation of a set of damage correlation indices obtained from the frequency analysis of the signals, obtained from PZT sensors, from the reference and damaged state of the structure. The main idea behind this approach is that elastic waves, propagating from the site where they are generated to the location where the signals are sensed, carry information about the portion of the structure through which they have traveled. Elastic waves can be very sensitive to small local defects. These authors argued that this damage index approach can be used for detection and, under certain conditions, characterization of degradation in aircraft, aerospace and civil structures and that the analysis of the waveform signals would provide information on the location and nature of small defects.

Sohn et al. (2004) and proposed a multi-scale structural health monitoring approach for detecting defects in composite structures by combining Lamb wave propagation, impedance-based methods and time reversal acoustics using a common active sensing system for local nondestructive evaluation. These authors mention that the ultrasonic research community has studied Lamb waves for the nondestructive evaluation of plates since the 1960s. They also define Lamb waves as mechanical waves corresponding to vibration modes of plates with a thickness of the same order of magnitude as the wavelength. In the study presented, the wave attenuation feature, identified using a wavelet based damage index, is used to locate the region of damage (Sohn et al. (2004)). Damage was introduced by firing a small projectile at different locations on a graphite fiber/epoxy matrix square plate with varying velocities. It was found that the three methods studied can be complementary because they can use the same sensors/actuators, which in this case were commercially available thin films with embedded PZT sensors. For instance, while it was observed that the Lamb wave propagation method is effective for thin plates, the impedance method is more suitable for detecting damage near structural joints or connections.

Toyama and Okabe (2004) studied the effects of tensile strain and transverse cracks in the propagation velocity of Lamb waves in cross-ply FRP laminates, since the waves' propagation velocity is sensitive to changes in the in-plane stiffness of the laminates. It is noted that some composite materials, such as CFRP, exhibit non-

linear stress-strain responses, which makes Lamb wave velocity based damage detection methods more complex under external loading. The Lamb wave propagation velocity was measured in GFRP (glass fiber-reinforced plastics) and CFRP laminates during tensile tests and the elastic behavior was studied. These authors make two interesting observations: First, that it seems that measuring the first symmetric Lamb wave (S_0 mode) velocity is a more accurate means of evaluating the elasticity of laminates than the tensile test. Secondly, they noted that the wave velocity depends on both damage and the laminate elasticity. Finally, they state that one needs to know the laminate stiffness as a function of strain when under external loadings, for damage detection purposes.

Sundararaman et al. (2005) use sparse phased sensor/actuator arrays for online damage detection and localization on heterogeneous and homogeneous plates and discuss the theory of narrowband and broadband beam-forming of propagating plate waves in damaged elastic media. According to these authors, a beam-former is a spatial-temporal filter that can be used to 'look' in the direction of transmitted signal while eliminating interference that cannot be removed through temporal filtering or carrier demodulation alone. Beam-formers consist of sensor arrays that are used to extract directional damage signatures. The principle on which this method is based is equivalent to an acoustic impedance discontinuity in the solid medium, so that damage, in the form of debonding, delamination, fiber breakage, matrix cracking, or voids, can be detected, localized and quantified. Experimental evaluation of this procedure was carried out on glass/epoxy woven composite, steel and aluminum plates.

2 Wireless sensor networks

In recent years, advances in miniaturization, low-power circuit design, simple, low power, reasonably efficient wireless communication equipment; and improved small-scale energy supplies have combined with reduced manufacturing costs to make possible the implementation of wireless sensor networks (WSN), a completely new technological vision.

These networks combine simple wireless communication, minimal computation facilities, and some sort of sensing of the physical environment into a new form of network that can be deeply embedded in our physical environment. Fuelled by the low cost and the wireless communication facilities, wireless sensor network have a much larger numbers of entities than hardwired ad hoc networks. Typical sensing tasks for such a device can be temperature, light, vibration, sound and radiation. The elements composing a wireless sensor network are called sensor nodes. The hoped-for size is a few cubic millimeters and the target price range less than \$1, including radio front end, microcontroller, power supply and the actual sensor.

While these networks of sensor nodes share many features with existing ad hoc network concepts, there are also a number of very differences and specific challenges. First of all, they are application specific. The large number of conceivable combinations of sensing, computing and communication technology, open the door to the development of a high number of different sensor nodes. However, it is unlikely that a sensor node designed for a particular application scenario can be effectively and efficiently employed in a different network. Secondly, since these networks have to interact with the environment, their traffic characteristics can be expected to be very different from other, human-driven forms of networks. WSNs are likely to exhibit very low data rates over a large time scale, but can have very intense traffic when something happens. Thirdly, a major constraint for WSNs is represented by energy scarcity. Energy consumption is a primary metric to be considered in the design of a WSN. Disposable batteries do not

offer a sufficient operational autonomy to the sensor nodes, and the need to prolong their lifetime has a deep impact on the system and networking architecture. Finally, WSNs are likely to be required to self-configure into connected networks with optimized datapaths toward a gateway. This necessity is due to the possibility of deploying the network in an environment subject to modifications or to the necessity to mount sensors on object which change their position. Geographical awareness is a further capability which can be required of WSNs.

New types of applications become possible thanks to WSNs. Applications include environmental control such as fire fighting or marine ground floor erosion but also installing sensors on bridges or buildings to understand earthquake vibration patterns; surveillance tasks of many kinds like intruder surveillance in premises; deeply embedding sensing into machinery where wired sensors would not be feasible, e.g., because wiring would limit flexibility, be too costly, represent a maintenance problem and could not reach the deeply embedded points. Also classes of applications include car-to-car or in-car communication. The possibilities offered by sensor networks could potentially offer societal changes when the basic size and cost problems are solved. Thus, wireless sensor networks have recently received a lot of attention in the research community.

2.1 Wireless sensor nodes

Wireless sensor nodes' development represent a real challenge because of the numerous constraints associated with them. A wireless sensor node must incorporate a computing unit, a radio communication device and some sensors. In most cases, the energy constraints make an energy harvesting circuit necessary. In addition, everything must be assembled into a very small form factor device. Thus, the design effort is directed to the choice of low-power, small size devices. Since most of the WSNs paradigms require the nodes to wake up and perform some operations periodically, a key characteristic of all node's components is the stand-by consumption, which has to be as low as possible. This is the reason why microcontrollers with a good power modes' management are largely preferred. The radio module choice is dictated by the application requirements, such as transmission range and data volume. Often, radio communication is done in high frequency bands, because of the benefits from the point of view of the antenna's size. Energy

management circuit and energy storage devices have a big impact on the node's cost and size.

As for software, since sensor nodes must be small and energy is scarce, the operating and networking software must be kept orders of magnitude simpler as compared to today's desktop computers. This simplicity may also require to break with conventional layering rules for networking software, since abstractions typically cost time and space. Some operating systems have been developed specifically for WSNs. The most famous one is TinyOS. Unlike most other operating systems, TinyOS is based on an event-driven programming model instead of multithreading. TinyOS programs are composed of event handlers and tasks with run-to-completion semantics. When an external event occurs, such as an incoming data packet or a sensor reading, TinyOS calls the appropriate event handler to handle the event. Event handlers can post tasks that are scheduled by the TinyOS kernel some time later. Both the TinyOS system and programs written for TinyOS are written in a special programming language called nesC, an extension to the C programming language, designed to detect race conditions between tasks and event handlers.

The following paragraphs contain a review of the most important wireless sensor nodes developed to date.

2.1.1 Mica and Spec motes

The Mica motes were developed by a research team at the University of California, Berkeley and they are now manufactured by Crossbow Technology Inc. Mica motes are now at the fourth stage of their evolution: after the original version, Mica2 motes, Mica2Dot motes and MicaZ motes have been developed. All these versions are characterized by a common architecture (Figure 2.1), which is based on a microcontroller, a Flash memory chip for persistent data storage, a radio transceiver and on-board power regulation circuitry. Beside the microcontroller, a co-controller is used to manage the wireless reprogramming of the node. This feature allows the mote to be easily reconfigured, without the need for expensive field maintenance. A connector is incorporated in the mote, so that custom sensor boards can be connected to it. The motes are also provided with an ID tag, thanks to which each mote can be identified.

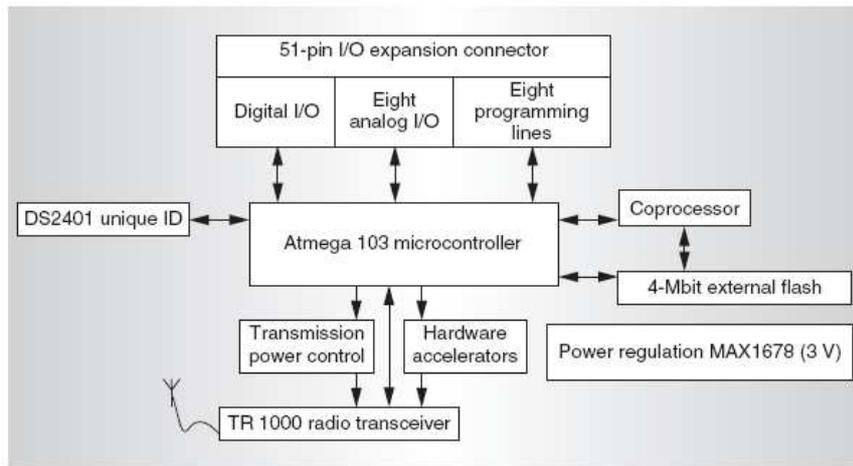


Figure 2.1 – Mica2 motes' architecture.

In the Mica2 version, the microcontroller is an ATmega103L or an ATmega128L, the Flash memory's size is 4Mb and the radio communication relies on a TR1000 transceiver, operating in the 916MHz band. The form factor is 3.2x5.76cm. The I/O equipment of the node is given by the microcontroller peripherals: the module has up to 48 GPIO lines, 8 A/D channels with 10-bit resolution, a UART port and an SPI bus. The ATmega128L has a maximum throughput equal to 8MIPS when powered with 3.3V and the maximum sampling rate of its internal ADC is 15Ksps for a 10-bit resolution and 76.9Ksps for an 8-bit resolution. The current consumption ranges from 5uA in power-down mode to 10mA in active mode at 8MHz, 3.3V. The radio module consumption is 12mA in transmission mode and 0.7uA in sleep mode.

MicaZ motes differ from Mica2 in the radio band. MicaZ motes operate in the 2.4GHz band. Mica2Dot motes are a miniaturize version of Mica2s (Figure 2.2). It is as big as a 2.5cm coin. Because of the smaller dimensions, the node has reduced interconnection capabilities, but it retains the same computational power and memory size.

After the development of Mica2Dot motes, the researchers at UCB pushed the miniaturization of their wireless node much further with the Spec motes. These devices are single chip motes integrating an AVR-like RISC core, 3KB of memory, an 8-bit ADC, FSK radio transmitter operating in the 902.4MHz band, a communication protocol accelerator, a SPI interface, and a UART port in a very small surface of approximately 2x2.5mm. The mote is a major step toward the realization of the so-called "smart dust"



Figure 2.2 – A Mica2Dot mote.

2.1.2 Tmote Sky

Tmote sky is a commercial wireless sensor platform developed by Moteiv, a company founded by some UC Berkeley alumni. Tmote Sky delivers high configurability and interoperability in a small board, about 7x3cm (Figure 2.3). Its core components are a MSP430 ultra-low power microcontroller by Texas Instruments, a 1MB Flash chip and a Chipcon 802.15.4-compatible transceiver. The node also integrates some light, humidity and temperature sensors. The I/O capabilities of the node are extended by a dedicated USB controller, which allows the node to directly communicate with a PC. A UART port and an I2C bus make the node ready for connection with other integrated chips. As for the radio communication, its transmission range is up to 125m.

The MSP430 microcontroller can run up to 8MHz. From the power consumption standpoint, its performance is outstanding. The chip draws only 500uA/MHz in active mode at 3.3V and the consumption drops to less than 1uA while in power down mode. The analog to digital conversion is performed by the microcontroller itself, thanks to the integrated ADC. The available channels are 8 and the resolution is 12 bits. The sampling period is lower than 10us. Due to the energy performance of the microcontroller, Tmote Sky shows a consumption of around 22mA when fully active.

The module runs a modified version of the TinyOS operating system called Boomerang. The customized operating system is designed to enhance the networking functionalities of the node and the deployment of reliable wireless solution. The node

has been demonstrated in real-life applications, like the real-time monitoring of a fire-fighters team during rescue operations.



Figure 2.3 – A Tmote Sky wireless node

2.1.3 Ecomote

Ecomote is an ultra-small wireless sensor node developed by Park, C. and P. Chou at the University of California, Irvine. The mote is not a single chip solution, however it achieves a remarkable result in terms of form factor. The mote is contained in a 13x10x8mm volume, including battery (Figure 2.4). The size reduction is achieved mainly by the choice of a chip integrating both a 8051 computational unit and a radio transceiver, operating in the 2.4GHz band. The MCU offers the standard I/O interfaces (i.e. UART, SPI and I2C), which are made available to external devices by ultra-small size connectors. The developers managed to fit in the small volume various sensors. In particular, a 3-axis accelerometer, a light sensor and a temperature sensor. This makes the sensing capabilities of the node completely equivalent to the Tmote Sky's ones. The memory of the MCU is expanded by an external 32KB SPI EEPROM chip. The MCU incorporates a 9-channel ADC with a resolution ranging from 6 to 12 bits and a maximum sampling rate of 100Ksps.

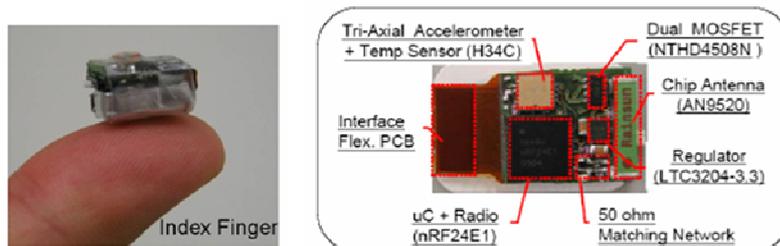


Figure 2.4 – Ecomote wireless sensor node.

Thanks to its remarkably small size, Ecomote makes the implementation of motion capture body networks possible. Anyway, the energy limits turn out to be significant. With a 40mAh battery, the node can run only for 1h in continuous listening mode and for 4h if the duty cycle is set to 50%. In addition, due to the short communication range and limited memory size, a more complex module, which acts as data aggregator, is necessary to collect the data from the motes and transmit it to some elaboration station.

2.2 Wireless platforms for SHM applications

Structural monitoring systems entail the installation of a variety of sensors in civil, mechanical and aerospace structures to monitor their environment and their response to loading. The insight into the performance of the structure helps optimizing maintenance, improving the understanding of the structures' behavior and refining mathematical models. Traditional structure health monitoring technology has employed wire-based systems to collect structural data. However, the installation of these wire-based systems can be expensive in labor, time and price. For example, a twelve-channel wire-based system may cost about \$50000, with half of the expense associated with its installation, including labor, cabling, etc. The cost of installing a monitoring system in the Tsing Ma suspension bridge was reported over \$27000 per sensing channel. Moreover, the installation of the wired systems can consume about 75% of the total testing time for large structures.

In order to reduce these monetary and time expenses for the installation of wire-based systems, new technologies in embedded systems and wireless communication have been adopted in academic and industrial research for wireless sensing and monitoring. The use of wireless communication for SHM data acquisition was firstly illustrated by Straser and Kiremidjian (Figure 2.5). Their work demonstrated the potential and cost-effectiveness of wireless monitoring systems. More recently, Lynch *et al.* extended the work by embedding damage identification algorithms into a wireless sensing unit; their work harnessed the unit's computational power for decentralized data interrogation.

For applications in civil infrastructures, a wireless monitoring system is expected to provide the capability for relatively long-distance communication within the span of the structure (usually from tens of meters to hundreds of meters), as well

as sufficient capability for local data analysis and processing. On the other hand, since wireless sensing units will most likely operate on portable batteries with finite energy, their power consumption must be recognized by potential end users. Because long-distance communication and local data interrogation capability usually demand more power, balancing the conflict between the requirements for higher data processing and communication capacity and lower power consumption becomes one of the major challenges in designing a wireless structural sensing and monitoring system. Another major challenge for a wireless structural monitoring system is the reliability and accuracy of data acquisition. A data acquisition process includes analog-to-digital conversion of sensor signals, the temporary storage of digital data, and the transfer of digitized signals by wireless communication. Problems such as circuit noise and occasional wireless communication failure should be handled properly for reliable and accurate data acquisition.

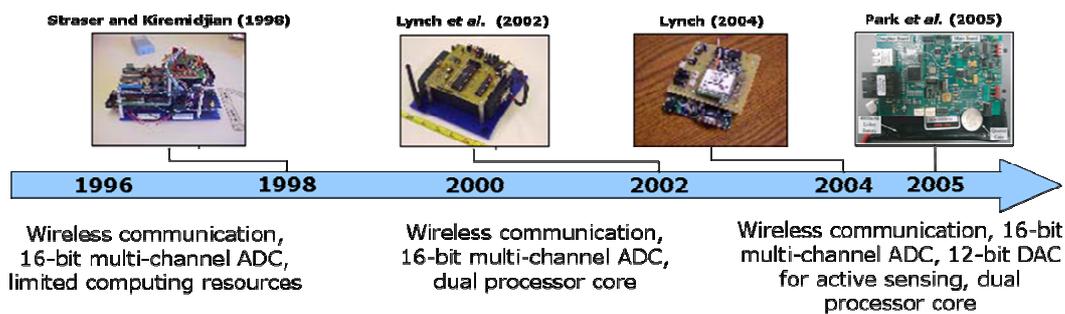


Figure 2.5 – Evolution of wireless platforms for SHM.

2.2.1 A Stanford wireless node for SHM

Lynch *et al.* developed their first wireless node for SHM in 2002. An improved version of the device was presented in 2004. The node is a fairly low-cost solution, with good acquisition and computing capabilities. The system architecture is simple (Figure 2.6). An ATmega128 microcontroller is interfaced to an ADC via a SPI port, a 128KB SRAM chip featuring a word-wide parallel I/O bus and to a 900MHz radio transceiver via a UART port. The ADC is a 4-channel 16-bit device (ADS8341) which can operate up to 100Ksps. Thanks to the architecture's simplicity, the node form factor is 10,2x6,5x4cm, including the five AA batteries powering it.

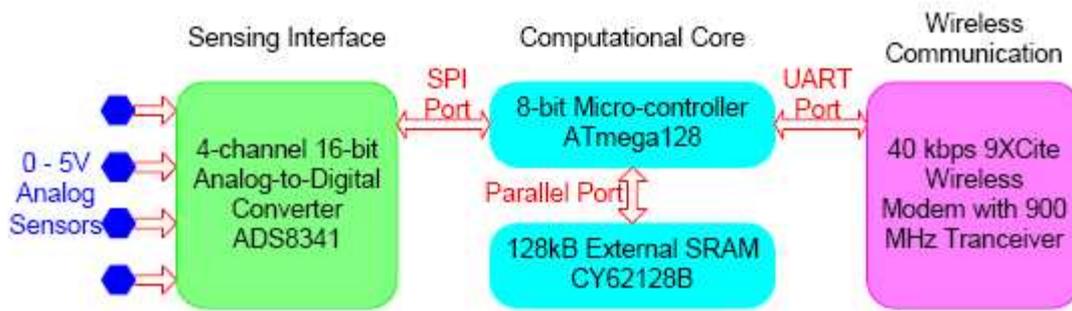


Figure 2.6 – Architecture of the node for SHM by Lynch et al.

The characteristics of the ADC mounted on the module allow the node to be interface to a wide range of sensors, such as MEMS accelerometers, gyroscopes and others. The node also has good storage capabilities: the SRAM chip can contain up to 64K 16-bit samples. The authors of the work also demonstrated the calculation of a 4096-point FFT in about 18 seconds. Basing on this information, the authors claim that, considering the radio bandwidth limitation from the wireless modem, in case a 50Hz sampling frequency is applied at each sensor and a 0 to 10Hz FFT spectrum is requested to be transferred back to central server, the system can support real-time non-stopping data collection and FFT results collection from up to 11 wireless sensing units simultaneously.

As for the energy consumption, the node draws a maximum current equal to 77mA. The stand-by consumption is 100uA. The life expectancy of the node, assuming it wakes up 10 minutes a day to perform data acquisition and analysis, is 190 days.

2.2.2 Duranode

Duranode was developed by Park, C. et al. with the University of California, Irvine in 2005. It is one of the most recent proposals about hardware solutions for wireless structural health monitoring. The node is designed to monitor vibrations in the structure under test. The sensors it includes are three MEMS accelerometers and a gyroscope. The developers chose to provide the module with high connectivity and computing capabilities, at the expense of the size optimization (Figure 2.7). The node can be used to deploy both wireless and wired network, thanks to a 802.11b wireless card and a Fast Ethernet controller and an optical transceiver. The choice is motivated by the need for a smooth transition between existing SHM systems, mostly

hardwired, and future ones, requiring no cables. The computational core is comprised of three microcontrollers. Two of them belong to the PIC family (PIC18F8680), the third one is produced by Freescale (MC9S12NE64). One of the PIC microcontroller acts as the supervising unit and it manages power consumption, data acquisition and the system memory. Its maximum throughput is 10MIPS at 4.5V. The power consumption at the maximum throughput exceeds 40mA, while using an external clock generator. The second PIC is interfaced to the 802.11b card and basically, it runs the firmware necessary for the communication through this device. Finally, the Freescale controller's function is communicating via the Ethernet connection.

Due to the complexity of the system, the power management is a complex task, both in the hardware and software domains. The node can be powered by both a battery and an AC input. Several switching regulators are required to produce the different supply voltages required by the various components out of the two power sources.

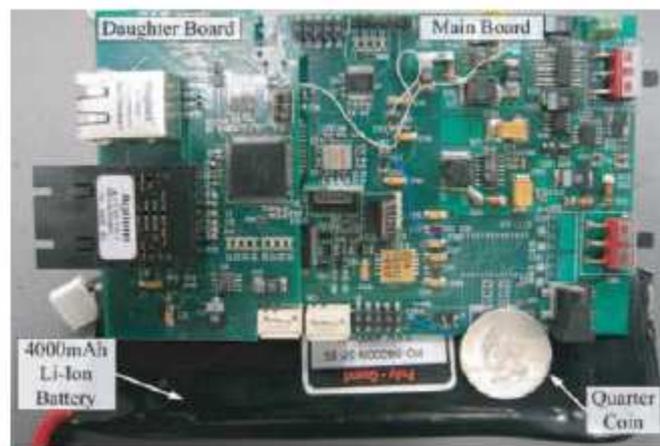


Figure 2.7 – The Duranode module.

The node can be provided with up to 64KB shared SRAM memory for data storage. The PIC microcontroller, which performs the data acquisition, has an integrated 16-channel 10-bit ADC. The sampling rate can be up to 30Ksps.

Despite the presence of the computational units, the node does not perform any on-board analysis of the data, which is simply routed to a base station.

2.2.3 A miniaturized wireless node for impedance-based SHM

In 2006, D. Mascarenas *et al.* developed a miniaturized wireless node capable

of performing an impedance-based analysis on bolt joints.

The key feature of the wireless node is the integration of an impedance measurement chip from Analog Devices (AD5933). This chip is connected to a PZT patch, glued on the structure under test or embedded in it. After being configured, the chip is fully autonomous in generating a sine wave of desired amplitude and frequency, driving the PZT with the sine wave, sense and digitize the response of the structure to the excitation caused by the PZT's vibration. The digital-to-analog conversion is performed with 24-bit resolution and the output frequency can span the 10KHz-100KHz range. The chip's analog-to-digital conversion rate is 1MSPS. Samples are produced with 12-bit resolution.

The rest of the node is composed by an ATmega128L microcontroller (8 MIPS at 3.3V), a ZigBee radio and a multiplexer, which allows up to 8 PZTs to be interfaced to the node. The FFT capabilities of the AD5933 make the calculation of a 1024-point FFT feasible with no external memory chips. The node consumption is about 65mA in active mode and 150uA in sleep mode.

A wireless power beaming system was tested as an alternative to including batteries in the module. The authors managed to charge a 0,1F supercapacitor by sending to the node a radio signal, converted to a DC signal by an on-board circuit. Given the power consumption of the node, the energy stored in the supercapacitor is sufficient to transmit 256 bytes of data.

2.2.4 Wisden: a wireless sensor network for SHM

In addition to hardware solutions, research is being conducted in the direction of effective network of sensors for SHM applications. Depending on the monitoring method and requirements, several aspects of an SHM network's implementation can turn out to be challenging.

A study concerning such a problem was conducted by Xu et al. at the University of Southern California in 2004. The researchers deployed a network of sensors over a testing structure represented by a frame ceiling in 1:1 scale ratio. The network was based on Mica2 motes connected to a custom vibrations measurement card, including a microprocessor, some accelerometers and an EEPROM chip. The focus of the study was on the network organization, data synchronization and transmission and global system reliability. The authors managed to acquire data from

tens of nodes deployed over the test structure with a 100% reliability, with a packet rate per node up to 1 packet/s, every packet being composed by 80 bytes. The result was achieved by the implementation of an adaptive self-configuration behavior of the network, that is by making nodes self-determine the best path to route packets to the base station. The nodes' firmware also features a data compression algorithm, by which the information is stored in non-volatile memory after a run-length compression. Wavelet compression is then used before data transmission. Wavelet compression a lossy operation, but, in case the received data series is considered critical by the user, the uncompressed samples can be retrieved from the nodes' EEPROM.

For consistent data visualization at the base station, all packets in the network are time-stamped. The time-stamping method consists in appending to the packet a field which is updated by every node with the time spent by the packet in that node. Along with knowledge of the network configuration allows data to be chronologically ordered at the base station.

2.3 Energy harvesting

The progressive reduction in size and power consumption of CMOS circuitry has led to a large research effort based around the vision of ubiquitous networks of wireless communication nodes. As the networks, which are usually designed to run on batteries, increase in number and the devices decrease in size, installing batteries in every module is not practical. Batteries with acceptable form factor and cost constraints do not yield the life times desired by most applications. Moreover, the replacement of depleted batteries can be costly and time-consuming. Therefore, the goal of researchers is implementing effective methods of scavenging ambient power for use by low power wireless electronic devices, in an effort to make the wireless nodes and resulting wireless sensor networks indefinitely self-sustaining.

Energy can be harvested from the environment in various ways. Light, vibrations, and air flow are some the most easily exploitable environmental power sources. For some of them, the state-of-art of collectors is advanced, while for others further research is necessary to improve the technology. However, the main difference between power sources is represented by their power density. The solar source is characterized by the highest power density, which is hundreds of times

bigger than the second best alternative. Of course this power source can be relied on only in places where solar radiation arrives. So, it is basing on the specific application, on its requirements in terms of power consumption and environmental conditions that the power source has to be chosen. Table 2.1 and Figure 2.8 compare the principal power sources and their power density.

Source	Power density ($\mu\text{W}/\text{cm}^3$)	
	1 year life time	10 years life time
Solar (outdoors)	150-150000	150-150000
Solar (indoors)	6 (office desk)	6 (office desk)
Vibrations	100-300	100-300
Acoustic noise	0.003 @ 75 dB 0.96 @ 100dB	0.003 @ 75 dB 0.96 @ 100dB
Daily temp. variation	10	10
Temperature gradient	40 @ $\Delta=10^\circ\text{C}$	40 @ $\Delta=10^\circ\text{C}$
Shoe inserts	330	330
Li batteries, non-recharg.	89	7
Li batteries rechargeable	13.7	0
Gasoline (micro engine)	403	40.3
Fuel cells (methanol)	560	56

Table 2.1 – Comparison of energy scavenging sources.

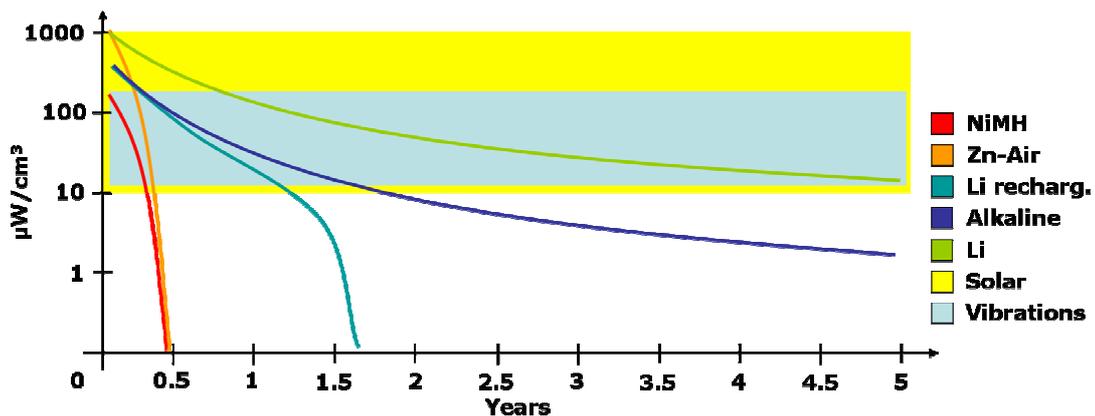


Figure 2.8 – Continuous power/cm³ vs. life for different power sources.

2.3.1 Solar energy

If outdoor sunlight, or relatively intense indoor light is available, solar cells

appear to be the best alternative for environmental energy harvesting. In sunny conditions, the power density of a solar source can be one hundred times higher than the second best choice, represented by vibrations. In addition, solar cells, which are the devices utilized to harvest solar energy, are a mature technology.

Solar cells working principle is based on the photovoltaic effect. This effect was first recognized in 1839 by French physicist A. Becquerel. The first solar cell was built by C. Fritts in 1883, who coated the semiconductor selenium with an extremely thin layer of gold to form the junctions. The device was only around 1% efficient. The modern age of solar power technology arrived in 1954 when Bell Laboratories, experimenting with semiconductors, accidentally found that silicon doped with certain impurities was very sensitive to light. This resulted in the production of the first practical solar cells with a sunlight energy conversion efficiency of around 6%.

Modern solar cells can be divided into three categories, corresponding to three different generations produced by technological advancement. The solar cells of the first generation consist of a large-area, single layer p-n junction diode, which is capable of generating usable electrical energy from light sources with the wavelengths of sunlight. These cells are typically made using a silicon wafer. By contrast, the second generation of photovoltaic materials is based on the use of thin-film deposits of semiconductors. These devices were initially designed to be high-efficiency, multiple junction photovoltaic cells. Later, the advantage of using a thin-film of material was noted, reducing the mass of material required for cell design. This resulted in a great reduction of costs for thin-film solar cells, which show top efficiencies above 15%. Currently, there are different technologies/semiconductor materials under investigation or in mass production, such as amorphous silicon, polycrystalline silicon, micro-crystalline silicon, cadmium telluride, copper indium selenide/sulfide.

Solar cells belonging to third generation are very different from the other two, broadly defined as semiconductor devices which do not rely on a traditional p-n junction to separate photogenerated charge carriers. These new devices include photoelectrochemical cells, polymer solar cells, and nanocrystal solar cells. The state-of-art of the third generation of solar cells is not very satisfactory, since the efficiency of most of the mentioned technologies is still under 10%.

Even if the third generation solar cells could be beneficial for some applications where mechanical flexibility and disposability are important, most of

current research effort is still focused in improving first and second generation devices. In particular, new internal structures of the solar cells and new material are investigated. Indeed, the record efficiency to date has been obtained with multi-junction cells, which rely on the p-n junction paradigm. Essentially, these multi-junction cells consist of multiple thin films produced using molecular beam epitaxy. Every film is made of a different light-absorbing material. A triple-junction cell, for example, may consist of GaAs, Ge, and GaInP₂, which are semiconductors. Each type of material has a characteristic band gap energy which causes it to absorb electromagnetic radiation over a portion of the spectrum. The semiconductors are carefully chosen to absorb nearly all of the solar spectrum, thus generating electricity from as much of the solar energy as possible. GaAs multi-junction devices are the most efficient solar cells to date. In late 2006 the 40% efficiency limit was broken just by these kind of solar cells. Unfortunately, they are also some of the most expensive cells per unit area (up to 40\$/cm²).

From the electrical point of view, solar cell can be modeled as a current source in parallel with a diode. An improved version of the model includes two shunt resistors to account for non-idealities (Figure 2.9). A solar cell can work over a wide range of voltages and currents, depending on the resistive value of its load. A key aspect in the employment of solar cells is making them work at their maximum power point (MPP), that is in the operating point in which the product of the output voltage by the output current is maximum.

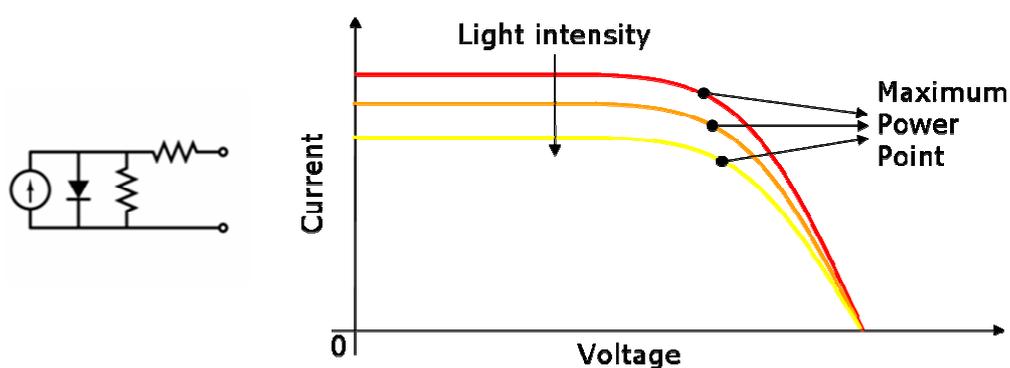


Figure 2.9 – Equivalent electrical model and typical I/V characteristic of a solar cell.

2.3.2 Vibrations to electricity conversion

In some cases wireless nodes have to be deployed in places not reached by sun light. For example, inside machinery or in shadowy parts of a building. For these

environments, vibrations to electricity conversion is an attractive alternative to solar energy harvesting, although the available power density is smaller.

As opposed to solar energy, for which the only collectors are solar panels, vibrations to electricity conversion can be performed in different ways. The most common and effective ones are three: by means of displaceable inductors, by means of variable capacitors and by means of piezoelectric materials. Displaceable inductive elements can be used to convert vibrations to electricity as well. Here, the basic principle is magnetic induction. The inductors – usually a coil- is made move in a magnetic field by vibrations (Figure 2.10). The movement produces a current which flows through the coil's wire.

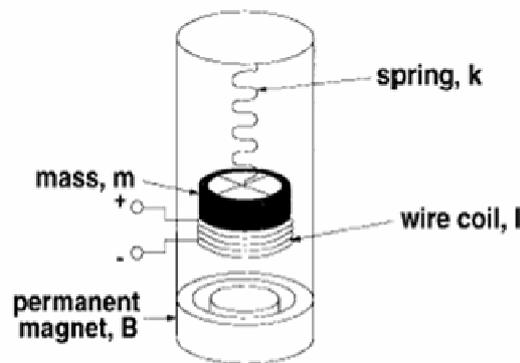


Figure 2.10 – Scheme of principle of vibrations conversion based on moving inductors.

The second way to convert vibrations to electricity is based on variable capacitors. Precisely, these solutions use capacitors the geometrical characteristics of which are modified by vibrations. The variations in the geometrical characteristics result in changes of the capacitance of the device. This can be exploited to convert mechanical work in electrical work, by inserting the capacitors in proper circuits. Basically such circuits work as charge pumps, which transfer energy from a source to the load (Figure 2.11). The energy increases during the transfer because of the mechanical work done on the variable capacitor.

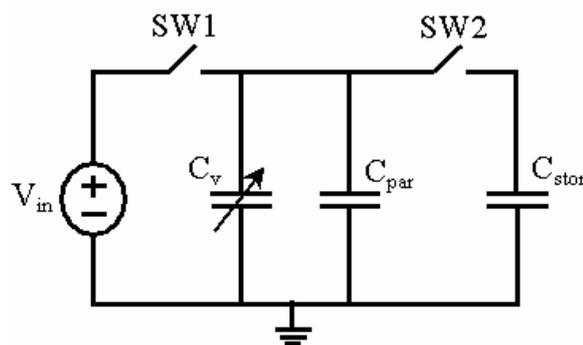


Figure 2.11 – Charge pump with variable capacitor.

The energy produced by the system is given by

$$E = \frac{1}{2} V_{in}^2 (C_V^{MAX} - C_V^{MIN}) \left(\frac{C_V^{MAX} + C_{par}}{C_V^{MIN} + C_{par}} \right)$$

where C_V^{MIN} and C_V^{MAX} are the minimum and maximum values of the variable capacitance.

The success of these solutions is due to the possibility of embedding variable capacitors in chips with relatively small effort thanks to MEMS technology. To date three main MEMS configurations have been implemented and tested. They are based on the modification of different geometrical properties of the capacitor. In the first configuration, capacitor is composed by two arrays of metal plates. One is fixed, the other one is displaceable. The vibrations make the non-fixed array of plates move. This changes the overlapping area of the plates of the two arrays and global capacitance of the device. In the second configuration, two similar arrays are used, but in this case the non-fixed array moves in such a way that the distance between the plates vary. Finally, in the third solution, a large metal plate is suspended over another one and the vibrations modify their distance (Figure 2.12). The maximum power density of all configurations is around $100\mu\text{W}/\text{cm}^3$.

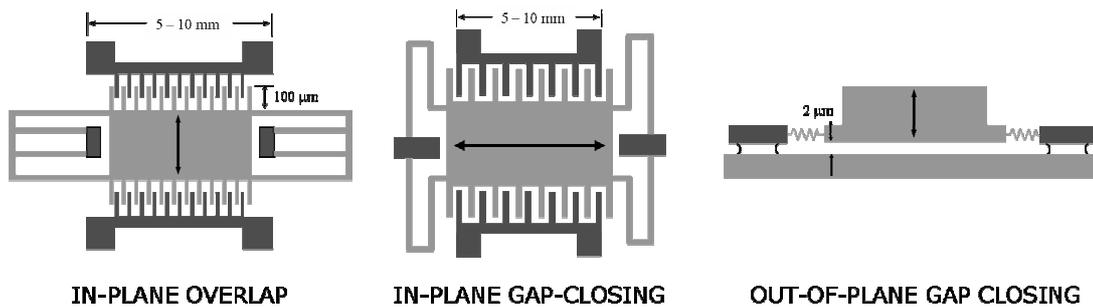


Figure 2.12 – MEMS variable capacitors.

Among vibrations to electricity conversion systems, the ones based on piezoelectric devices are characterized by the highest power density. Piezoelectric materials can be useful for vibrations-oriented energy scavenging because of the coupling between the mechanical and electrical domains they show. If a piezoelectric device is made vibrate, it converts the mechanical stress in an electrical signal (Figure 2.13). The amplitude and frequency of the vibration determine the power

output of the piezoelectric element. Current research efforts focus on realizing piezoelectric conversion devices in MEMS technology, too. The objective is embedding the piezoelectric generators in chips the same way variable capacitors are, with the benefit of a higher power density.

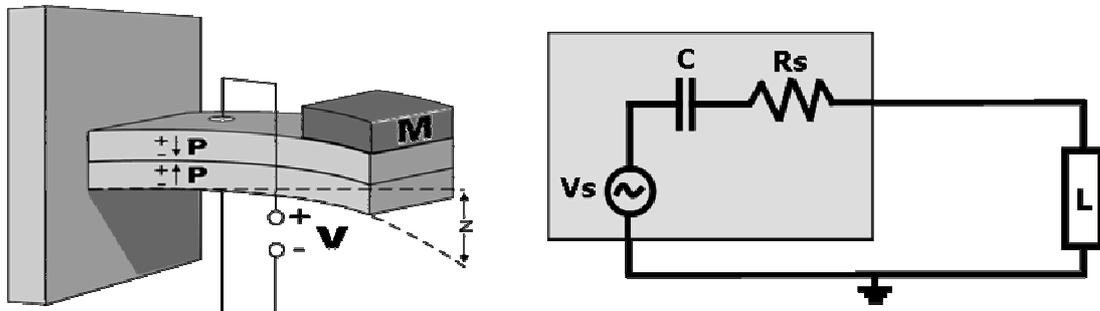


Figure 2.13 – Vibrations to electricity conversion by means of piezoelectric materials.

2.3.3 Technologies for energy storage: batteries

A battery is a device that stores chemical energy and makes it available in an electrical form. Batteries consist of electrochemical devices such as two or more galvanic cells, fuel cells or flow cells. The modern development of batteries started with the voltaic pile, announced by the Italian physicist A. Volta in 1800. Formally, an electrical "battery" is an interconnected array of similar voltaic cells. However, in many contexts it is common to call a single cell used on its own a battery. Each cell is composed of two half cells connected in series by a conductive electrolyte and has a positive terminal and a negative terminal. These do not touch each other but are immersed in a solid or liquid electrolyte. In a practical cell the materials are enclosed in a container, and a separator between the electrodes prevents them from touching. The electrolyte conducts current by allowing the passage of ions between the two electrodes. Such reactions are called faradaic, and are responsible for current flow through the cell. Non-charge-transferring (non-faradaic) reactions also occur at the electrode-electrolyte interfaces. Non-faradaic reactions are one reason that voltaic cells (particularly the lead-acid cell of ordinary batteries) get depleted when sitting unused. The electrical potential across the terminals of a battery is known as its terminal voltage.

Batteries are usually divided into two classes: primary batteries, which irreversibly transform chemical energy to electrical energy; secondary batteries, which can have the chemical reactions reversed by supplying electrical energy to the

cell, restoring their original composition. Of course, the batteries of interest for energy scavenging applications are the latter.

The voltage produced by a cell depends on the chemicals used in it, which have different electrochemical potentials. For example, alkaline and carbon-zinc cells both can output about 1.5 volts, due to the energy release of the associated chemical reactions. Because of the high electrochemical potentials of Li compounds, Li cells can provide as much as 3 V or more. So, an important research field is the investigation of materials to be employed in batteries' production. The most recent kind of cells include nickel-cadmium (NiCd), nickel metal hydride (NiMH), and lithium-ion (Li-Ion) cells, because of the increase of power density these materials bring (Figure 2.14). Anyway, the more electrolyte and electrode material in the cell, the greater the capacity of the cell.

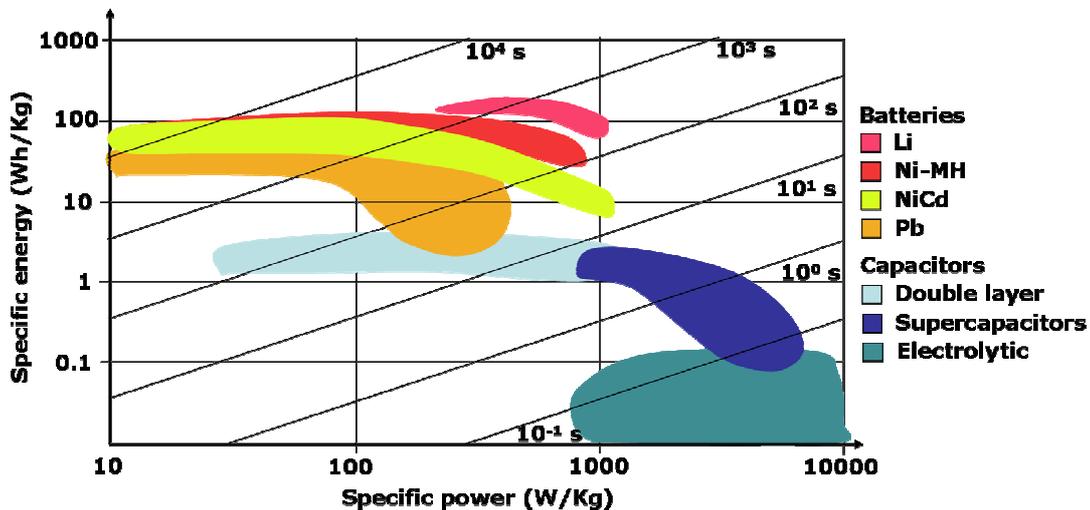


Figure 2.14 – Comparison between technologies for energy storage.

Because of the chemical reactions within the cells, the capacity of a battery depends on the discharge conditions such as the magnitude of the current, the duration of the current, the temperature, and other factors. If a battery is discharged at a relatively high rate, the available capacity will be lower than expected.

A major issue with batteries can be their self-discharge. This phenomenon is due to non-current-producing side chemical reactions, which occur within the cell even if no load is applied to it. The rate of the side reactions is reduced if the batteries are stored at low temperature, although some batteries can be damaged by freezing. High or low temperatures reduce battery performance. Disposable (or primary) batteries can lose 2% to 25% of their original charge every year. Rechargeable batteries self-discharge even more rapidly than disposable ones; up to

3% a day. In addition to the short-term effect represented by self-discharge, all batteries undergo some long-term deterioration due to the aging of the materials they are made of.

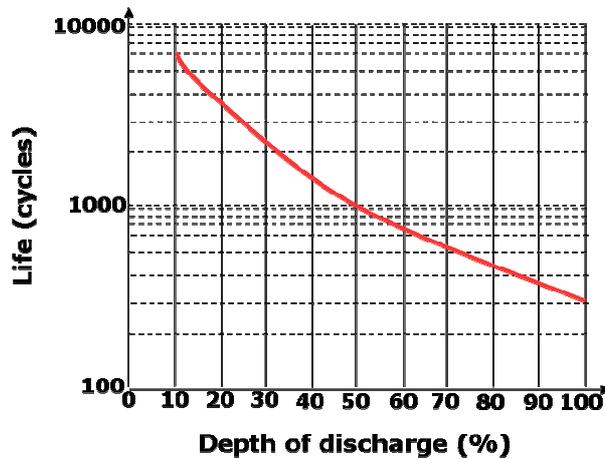


Figure 2.15 – Life of a battery vs. depth of discharge.

Finally, a batteries’ property to take into account is the resilience to the “memory effect”. The expression memory effect means the variation of the capacity of battery depending on the depth of its discharge (Figure 2.15 – Life of a battery vs. depth of discharge.). For example, NiMH and NiCd batteries can be charged several hundred times and they both can be completely discharged and then recharged without their capacity being damaged or shortened. By contrast, NiCd batteries are severely prone to the memory effect (Table 2.2 – Characteristics of most common rechargeable batteries).

	NiCd	NiMH	Li-ion
Self-discharge (%/day)	0.67	1.00	0.33
Cycle life to 80% capacity	1500	500	100
Aging effects	Low	Low	High
Memory effects	High	Low	None

Table 2.2 – Characteristics of most common rechargeable batteries.

2.3.4 Technologies for energy storage: capacitors

An alternative to batteries for energy storage are capacitors, in particular supercapacitors, a recently developed family of capacitors which show outstanding properties, in particular from the point of view of the quantity of charge they can

store. Supercapacitors have been created as an improvement of electrolytic capacitors.

The electrolytic capacitor was developed in the 1930s by the Cornell-Dubilier Electric Corporation in New Jersey. The company's scientists introduced a new way of designing capacitors featuring three major enhancements. First of all, an expanded surface area, thanks to the etching of one Al electrode by acid. This operation leaves the electrode roughened and increases the surface area available to accumulate charge. Secondly, a reduced insulation thickness between electrodes. After the etching, the electrode was oxidized, and the very thin insulating layer produced by this reaction (Al_2O_3) was used as the positive and negative charges separator. Thirdly, a liquid electrode, constituted by an electrolyte in which the first electrode was immersed. The electrolyte consisted of a paste made when boric acid dissolves in and reacts with glycol. So, the electrolytic capacitors achieved a much higher energy density than their predecessors (i.e. electrostatic capacitors) by dramatically reducing the distance between the charge layers to the micrometer range and increasing the electrodes' surface.

Supercapacitors (or ultracapacitors) are the result of another round of innovations. The charge-separation distance in supercapacitors has been reduced to the dimensions of the ions themselves within the electrolyte. Here, charges are separated by just a few nanometers. Coupling the ultra-small separation distance with a relatively vast surface area, in supercapacitors the ratio of available surface area to charge-separation distance has grown to an amazing 10^{12} . This characteristic is at the basis of their unprecedented storage capabilities. The development of supercapacitors originated in the work of Standard Oil of Ohio Research Centre (SOHIO) in the early 1960s. SOHIO researchers discovered that two pieces of activated carbon immersed in an aqueous electrolyte solution and connected across the terminals of a battery acted as a capacitor. Later, SOHIO's scientists explored the use of organic electrolytes, but at the time (early 1970s) there was really no market for such devices and little understanding of what was happening in them. SOHIO licensed its double-layer capacitor technology (Figure 2.16), as it came to be known, to NEC in 1971. During the 1980s Matsushita Electric Company patented a method of manufacturing supercapacitors having improved electrodes. In a short time, applications proliferated, especially for the coin cell types of supercapacitors such as those manufactured in Japan by Nippon Electric Company (NEC), Elna/Asahi Glass, and Matsushita. Throughout the 1980s and '90s, manufacturing of supercapacitors was

primarily an art. Recently, automated assembly techniques have replaced the labor-intensive aspects of supercapacitor manufacturing and costs have decreased substantially. For instance, around 1985 a 2.3V supercapacitor rated at 470F and manufactured by Panasonic (Matsushita Electric) cost roughly \$2/F. Today, that same supercapacitor would 10⁻²\$/F, and costs continue to decrease rapidly as ongoing automation replaces hand assembly.

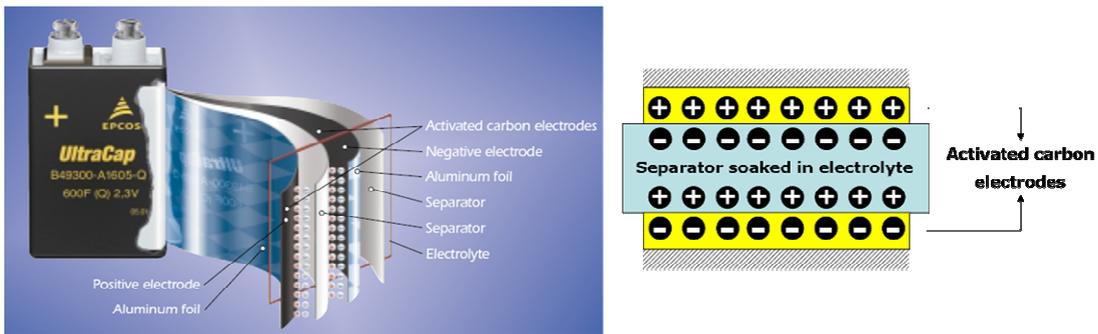


Figure 2.16 – Internal structure of supercapacitors.

Supercapacitors rely on an electrostatic effect, purely physical and highly reversible. Charge and discharge performs upon movement of ions within the electrolyte. This mode of energy storage is in clear contrast to all battery technologies, since these are based on the formation and dissolution of chemical compounds at the battery electrodes (faradic reactions). The fundamental property differences between supercapacitors and battery technologies result in extended useful life, high cycle life and virtually maintenance-free devices (Figure 2.17).

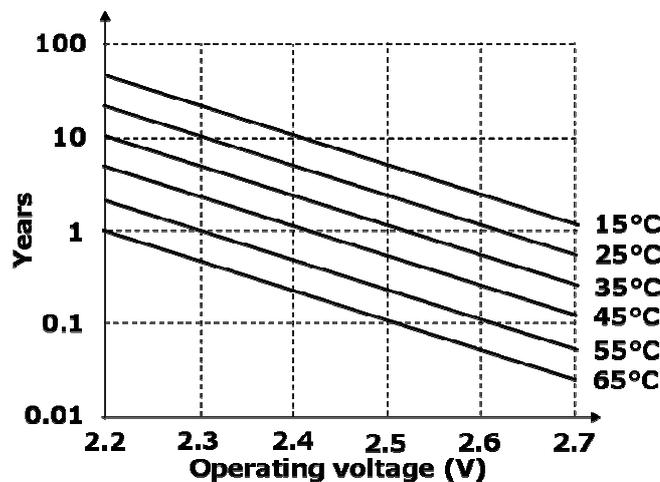


Figure 2.17 – Life expectancy for supercapacitors.

Today's supercapacitors achieve capacitances ranging up to 270F, while the

whole family of capacitors offers capacitances ranging down to picofarads (10^{-12} F). The leading manufacturers of supercapacitors today are Maxwell Technologies in the United States, NESS Capacitor Company in South Korea, Okamura Laboratory in Japan, and EPCOS in Europe. These companies manufacture carbon-carbon, or symmetric, supercapacitors. That is, both electrodes have identical construction. There are some differences in the organic salts and solvents used, and this is where supercapacitor manufacturing becomes proprietary. The supercapacitors produced by the companies listed above are rated at 2.5/2.7V, with specific capacities clustered at 5 F/g.

The outstanding performance of supercapacitors is being pushed even further by scientists. The two objectives are reducing the cost per farad and increase the capacitance. In October 2003, JEOL Ltd. in Tokyo announced an improved supercapacitor. The name coined for the new class of devices is “nanogate” or “nano-carbon” capacitor. This new component, already available for sale, has an energy density of 50/75 Wh/Kg, more than 10 times that of existing supercapacitors, and a specific capacitance of 30 F/g, about 5 times higher. The device features two carbon electrodes formed of a new, patented material whose uniqueness lies in its high porosity and accessibility for storing ions.

Even further out on the experimental edge, researchers are exploring the possibility of using carbon nanotubes for supercapacitor electrodes. The importance of carbon nanotubes lies in their uniform nanoscopic pores (about 0.8 nanometres in diameter), which could in theory store much more charge than the nanogate capacitors if the nanotubes could be properly assembled into macroscale units.

Supercapacitors are widely used in wireless nodes with energy harvesting capabilities, as described in later sections, both in conjunction with batteries and as unique energy storage devices. In both case, they are used mainly because of their extended life expectancy and lower performance degradation over time. Another key feature, which makes them suitable for some applications, is their ability to deliver peak currents much higher than batteries, without being damaged. This ability is due to the fact that supercapacitors stores the separated charges, which can be released as current very quickly. Batteries, on the contrary, store energy through chemical processes and release it the same way, so the current flow is limited by the speed of the chemical reactions. On the downside, supercapacitors have less retention capabilities. Finally, environmental friendliness also makes supercapacitors preferable to batteries, in particular in those scenarios which require wireless nodes

to be dispersed in the environment.

2.4 Energy scavenging solutions

As a consequence of the importance of energy harvesting in the implementation and deployment of effective wireless sensor networks, several research groups are busy in developing platforms dedicated to the collection of energy from the environment. Most of these platforms are designed to harvest solar energy, because of its availability and its power density. Their general architecture consists in a switching regulator, which is used to store the energy output by the solar panel in the storage device, and some circuitry which provides the supply voltage to the device being powered by the harvesting module. The complexity of the circuitry depends on the characteristics of the load device. The energy can be stored both in supercapacitors and in batteries. Actually, some harvesting modules use both of them: a supercapacitor acts as a primary buffer and power source, while the battery is used only under particular circumstances, such as low level of charge on the supercapacitor (Figure 2.18).

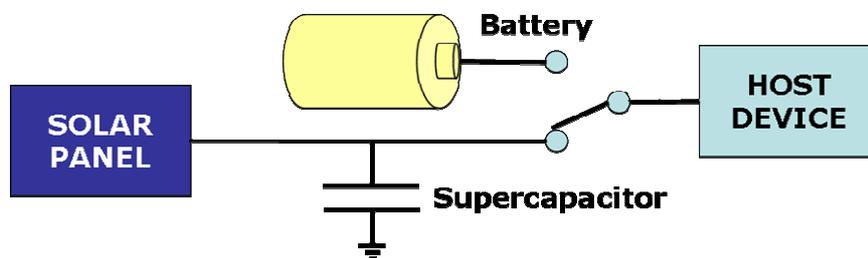


Figure 2.18 – General architecture of energy harvester modules.

The main focus in the development of such platforms is the optimization of the efficiency in the power transfer from the solar panels to energy storage device. The key problem is matching the output impedance of the solar panels to the input impedance of the cascaded circuitry. The task is not simple because of the dependence of the I/V characteristic of the solar panels on the intensity of the solar radiation. Thus, most solutions implement a dynamic MPP tracking. As described in the following sections, some of them perform the operation thanks to devices with computational capabilities, while some other ones rely on purely analog circuits.

2.4.1 Heliomote

Heliomote is an energy scavenging module developed by the University of California, Los Angeles over the last few years (Figure 2.19 and Figure 2.20). The module was initially designed to augment Mica motes sensor nodes with energy harvesting capabilities. In its most recent versions, it can be used to power a variety of sensor nodes. The module can harvest energy from the environment by means of solar panels and it uses the combination of batteries and supercapacitor to store the energy.



Figure 2.19 – The energy harvesting module Heliomote.

As in most energy scavenging solutions, the module generates the supply voltage required by the load (i.e. the wireless node) by a DC-DC boost converter. However, the module implements enhanced features aiming to maximize the efficiency of the energy collection process. First of all, the module is completely autonomous in taking decisions about the energy harvesting and energy storage. A dedicated circuit enables the module to track the maximum power point of the solar panels (MPP), thus maximizing the energy transferred from the solar panels to the storage devices. The MPP tracking is adaptive and does not depend on the specific solar panel in use. Then, an overcharge and uncharged circuit allow optimal utilization of the battery pack. Overcharge protection features are implemented with respect to the supercapacitor, too. Finally, the module is configurable in terms of the output voltage level and, basing on the battery characteristics and level of charge, it offers the possibility to connect the load directly to the storage devices, bypassing the DC-DC converter and avoiding the efficiency loss introduced by it.

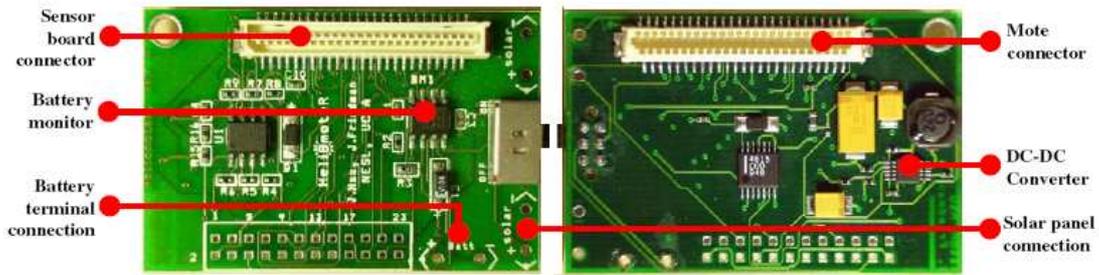


Figure 2.20 – Heliomote board.

The module is provided with an I²C interface which allows the wireless node to get information about the activity of the harvesting circuit. Heliomote can be queried to obtain data concerning the instantaneous current output of the solar panel, the level of charge of the battery pack and supercapacitors and the power consumption of the module itself. This data can be used to control performance in general and schedule tasks by adaptive duty-cycling or more complex energy aware algorithms.

The tests ran by the team which developed the module showed an energy conversion efficiency ranging from 72% to 82% depending on the load current and the supply voltage needed by the load.

2.4.2 Prometheus

Prometheus is another example of energy scavenging module, designed at the University of California, Berkeley to provide existing wireless sensor nodes with energy harvesting capabilities (Figure 2.21). Similarly to Heliomote, the module collects energy by means of solar panels. The energy is then stored in a supercapacitor, used as a primary buffer and in battery Li-Polymer battery. The supercapacitor is hardwired to the solar panel and no MPP tracking is implemented. The system usually powers the load drawing current from the supercapacitor which acts as an energy source for the battery charging, too. The battery is used only when the power output of the solar panel is not sufficient or the level of charge of the supercapacitor gets too low (Figure 2.22). This solution is used to extend the battery lifetime, which can be a main constraint for wireless systems.

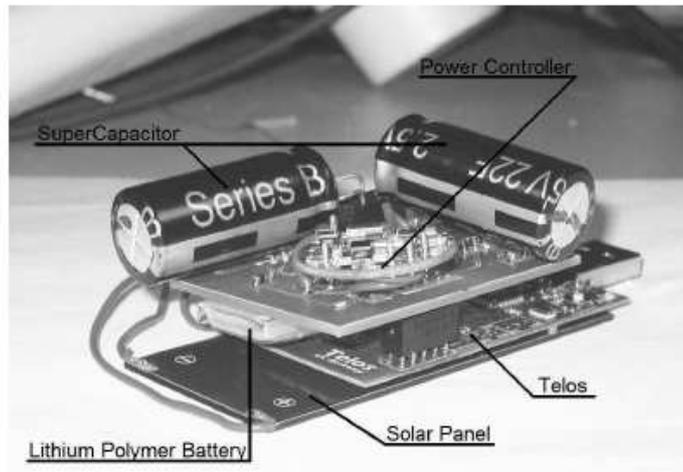


Figure 2.21 – Prometheus module.

The developers built a prototype of the module to be interfaced to Berkeley Telos motes, and wrote the software to manage the energy harvesting activity. In this case the energy management is up to the wireless node, since the energy harvesting module does not have any computational capability. The developers chose this solution to avoid the power consumption overhead introduced by active components included in the scavenging circuit. So, a small energy management routine has been appended to the normal operations executed by the mote during the wake-up phase. The behavior implemented is based on the Telos' microcontroller checking the level of charge of the supercapacitor and choosing whether to charge battery or not.

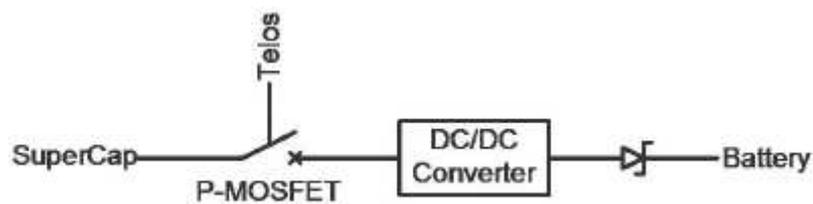


Figure 2.22 – Battery charging circuitry on the Prometheus platform.

The platform has been tested in real operational conditions and its capability to run on the harvested energy only has been demonstrated. The developers estimate a lifetime for the node ranging from 1 year to 43 years, respectively in case of a 100% duty-cycle and a 1% duty-cycle.

2.4.3 Everlast

Everlast is an energy harvesting platform designed to address the maximum

power point tracking issue (Figure 2.23). It was presented in 2006 by F. Simjee and P. Chou with the University of California, Irvine.

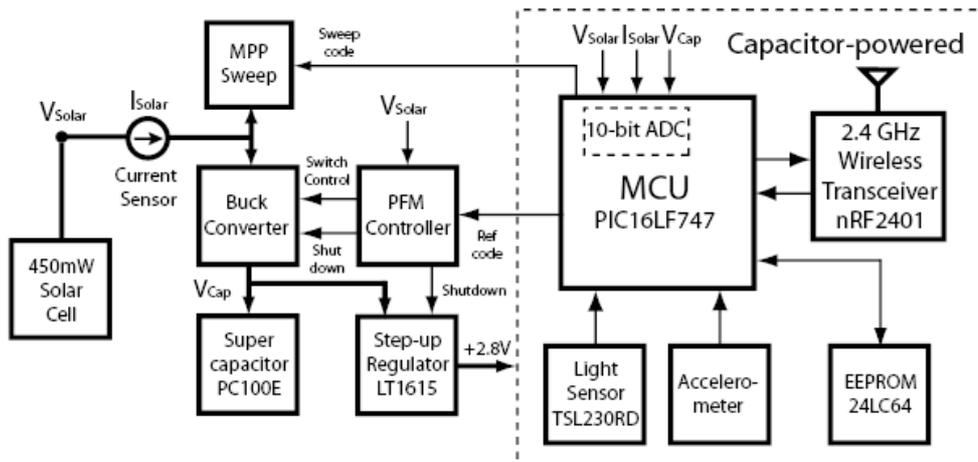


Figure 2.23 – Block diagram of the Everlast platform.

The module forgoes batteries and stores the energy in supercapacitors only. The key feature of the platform is an active tracking of MPP of the solar panel performed by a dedicated circuit, which is controlled by the microcontroller mounted on the wireless node powered by the module. The MPPT custom circuit integrated in the module aims to match the input impedance of the charging circuitry to the output impedance of the solar panel. Indeed, power transfer is maximum in case the two impedances are identical.

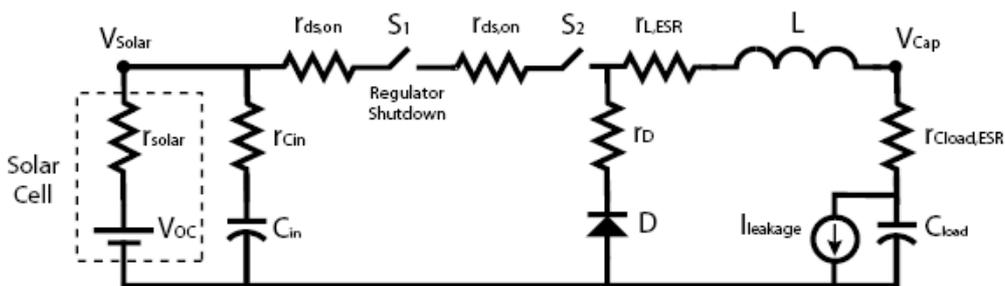


Figure 2.24 – Buck converter circuit used in the Everlast module.

As opposed to the Prometheus implementation, the supercapacitor is not hardwired to the solar panel. Instead, a buck converter gets the solar panel output and charges the supercapacitor. The impedance matching is obtained by changing the buck converter switching frequency. This results in the variation of its input impedance, similarly to what happens in switching capacitors converters (Figure 2.24). The switching frequency is determined by the uC via a PFM controller. The uC is also interface to some circuitry which can be used to profile the I/V

characteristic of the solar panel in all light conditions. This allows the MPP to be accurately determined over varying environmental conditions. Basically, the circuitry includes a DAC which controls the gate of a MOSFET connected in parallel with the solar panel. By tuning the MOSFET conductivity, the current flowing out of the solar cell can be changed and the output voltage sampled.

2.4.4 Ambimax

Ambimax is the most recent energy harvesting module presented to the research community. It has been developed by Park, C. and P. Chou at the University of California, Irvine in 2006. Ambimax is a general purpose energy harvesting platform (Figure 2.25), which can be connected to different loads and interfaced with various power sources. In particular its functioning has been demonstrated in conjunction with solar cells and a wind generator.

The module is comprised of three functional parts. The first one performs the energy scavenging, with a maximum power point tracking behavior. The second one is represented by the energy storage devices. The third one manages the connection of the load to the module.

Both supercapacitors and batteries are used for the energy storage. The energy scavenging section includes a PWM converter, of which the input impedance is varied to match the output impedance of the power source. The peculiarity of the Ambimax solution is given by the absence of any computational unit. All the MPPT is done by purely analog circuitry. Basically, the circuitry compares the voltage output of the power source with the output of a sensor, chosen to accurately reflect the expected behavior of the power source. Depending on the results of this comparison, the switching regulator is turned on and off.

As mentioned, a dedicated part of the module manages the connection of the load to the module. More exactly, it chooses whether to power the load by the supercapacitors or by the batteries. When the voltage on the supercapacitors is high enough to meet the load requirements with respect to the supply voltage, then the energy is taken directly by them. And, if it is the case, the battery undergoes charging. When the supercapacitors show a low voltage across their pins (i.e. the power source are not providing enough energy to sustain the load' activity), then the battery is connected to the load.

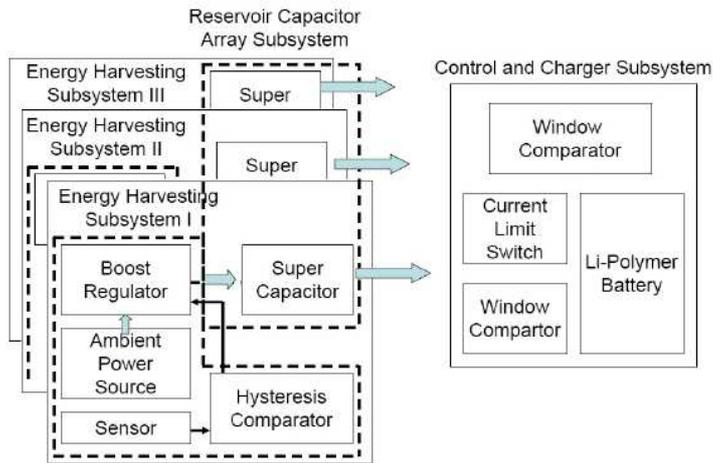


Figure 2.25 – Block diagram of the Ambimax platform.

2.4.5 PicoBeacon

The PicoBeacon platform was developed at the University of California, Berkeley in 2003 (Figure 2.26). The platform harvests solar energy from the environment to power a radio module. The remarkable feature of the platform is that it is the first fully self-contained wireless transmitter powered solely by solar energy. The radio module, computing units and energy harvesting circuits are integrated in a single chip.

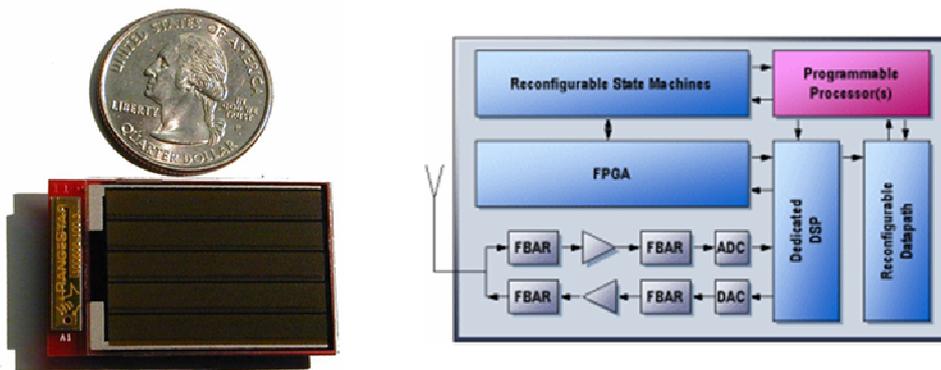


Figure 2.26 – PicoBeacon wireless module with energy harvesting capabilities.

2.4.6 Other platforms

Some energy scavenging platforms have been developed at the University of Bologna, as well. The first energy scavenging solution for a wireless node has been

developed by the author of this thesis and M. Paselli under the supervision of Prof. Benini in 2006. The platform is designed to harvest energy by solar cells and to store it in supercapacitors. The scavenging circuit is composed of a voltage regulator which transfers the energy from the solar panel to the supercapacitor and of a Schottky diode, which prevents the current from flowing back to the supercapacitor. A DC-DC boost converter produces the supply voltage required by the load. The size of the module is very small, about 2.5x2.5cm. Tests were run to demonstrate the effectiveness of the platform in powering a wireless node including a microcontroller and a radio transceiver. An adaptive duty-cycling behavior allows the whole system to achieve the maximum QoS allowed by the environmental conditions through the day-to-night cycle.

Subsequent efforts were made to address the harvesting efficiency issue. A second platform was developed by S. Raggini and D. Brunelli with Prof. Benini, featuring maximum power point tracking (Figure 2.27). This platform uses a small solar cell integrated in a DIP package as a reference. The voltage output of the pilot cell, which reflects the expected behavior of the main solar panel, is compared to the output of this latter by a comparator with hysteresis. The signal produced controls the DC-DC switching regulator charging the supercapacitor, insuring that it works at the MPP. The circuitry needed to perform the MPPT introduces a very low power consumption overhead ($<1\text{mW}$) and the global efficiency of the system is close to the nominal efficiency of the off-the-shelf switching converters mounted on the module, close to 80%.

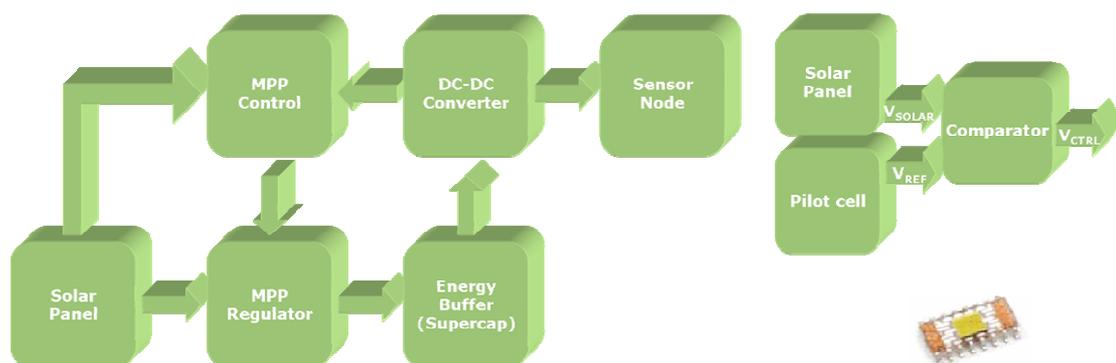


Figure 2.27 – Energy harvesting solution developed at the University of Bologna.

3 Design of a wireless node for active sensing in SHM

As stated in Chapter 1, a large amount of research has been focused on utilizing the Lamb waves method for structural health monitoring. However, the vast majority of this research has required the use of expensive and bulky instrumentation, that are not suitable for field deployment. Furthermore, the instruments used for the measurements typically contain a large number of additional features that are not utilized for structural health monitoring.

This chapter is concerned with the design, fabrication and testing of a prototype wireless sensor node for Lamb waves-based structural health monitoring. The application requirements, design choices and sensor node layout and software are presented. The issues with the prototype realization are discussed. After fabrication, the node has been tested. The results of these tests are disclosed.

From the point of view of structural health monitoring, the development of the wireless sensor node (named SHiMmer) here described is a significant step forward in moving the Lamb waves method from the laboratory to actual use. From the sensor nodes' design standpoint, Shimmer brings new ideas in the domain of energy management and shows how heavy computation can be performed at the node level, even in the case the node is powered by harvested energy only.

3.1 Requirements

The development of Shimmer is part of a joint project of the University of California, San Diego with the Los Alamos National Laboratory (LANL). The project aims to develop a wireless sensor network to be deployed over civil infrastructures and mechanical systems for SHM purposes. Bridges, industrial plants and military assets are the targets of the monitoring activity. The network is comprised of two layers of sensor nodes, characterized by different complexity and

capabilities (Figure 3.1). The system is queried by an unmanned aerial vehicle (UAV) remotely controlled or completely automated. The UAV is supposed to get close to the structure and collect a raw information on its conditions by interacting with the first network layer, composed of RFID sensing devices. Then, based on this information, it is supposed to trigger a more accurate analysis, performed by the more complex layer of the network -comprised of Shimmer nodes-, on those parts of the structure which may have damage. The project aims to execute the process daily. The assessment of the structure's integrity is usually required after events like earthquakes. However, daily monitoring can provide information about the deterioration of the structure caused by its normal use. Thus, the node is required to wake up once a day, acquire the data, run the analysis algorithm and transmit the results to the UAV.

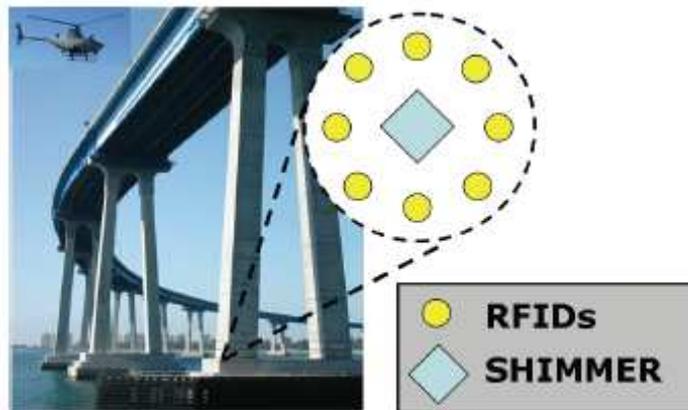


Figure 3.1 – Monitoring system's organization.

The damage detection approach used in this system is based on the Lamb Waves method. This method offers the possibility of detecting and localizing damage in the structure under test with high accuracy. The method usually uses a set of piezoelectric devices (PZT) organized in a grid, placed over a portion of the structure under test. In standard operation, two PZTs are selected in the grid. The first one works as an actuator, the second one as a sensor. The actuator PZT is made to vibrate with an appropriate input wave and the response of the structure is sensed by the second PZT. The high accuracy in the localization is achieved by sensing multiple paths between the devices in the grid, typically less than 10.

As mentioned, the network can be deployed over a variety of structures, made of different materials. The frequency and amplitude of the actuation wave and the sampling rate of the structure's response are two key factors for the success and

quality of the analysis. Most SHM methods require waves with an amplitude up to 15-20V peak-to-peak at a frequency ranging from 10 KHz to 1 MHz. Consequently, the node has to be capable of generating actuation waves characterized by an amplitude and a frequency in these ranges.

The success of the damage detection analysis is highly related to the sensing accuracy. Most SHM analysis algorithms work on data in the time domain, in which a highly accurate knowledge of the waveform is needed. Provided that the response of the structure is expected to be at the same frequency as the actuating wave, the Nyquist theorem requires a sampling rate of at least 2MHz to be used. But then, the waveform should be reconstructed using the formula

$$x(t) = \sum_{k=0}^{n-1} x(kT_s) \text{sinc}\left(\frac{t}{T_s}\right)$$

where T_s is the sampling period and $x(kT_s)$ are the samples of the signal. To avoid running a reconstruction algorithm, the wave can be oversampled. This allows the required accuracy in the time domain to be achieved, without the need for further processing of the samples. Accordingly, the desired sampling frequency has been set to 10MHz, in order to make at least 10 samples per period available through the entire frequency range. As for the data processing, the node is required to run the most common analysis algorithms described in Chapter 1.

Finally, the project aims to develop a self-sufficient network from the energy point of view, in order to reduce overall operational costs. This means the network has to be able to work for a very long time without human intervention. A direct consequence of this is the requirement of providing the node with energy harvesting capabilities and avoiding the use of batteries, which need periodical replacement.

3.2 Architectural design

The node has to accomplish several tasks: communication with the UAV, control of the PZT devices in the actuation process, sampling of the response of the structure and processing of the acquired data. These tasks involve activities both in the analog and digital domains, with very different requirements in terms of power, peak currents and supply voltages. The design of the architecture of the node has

taken place in two steps. Firstly, the components necessary to perform the required activities have been individuated. Secondly, the connections between these components have been studied and particular solutions have been adopted to meet the energy requirements.

The core activity of the node is the SHM analysis, consisting in actuation, acquisition and data processing. All these activity has to be controlled by a device with high computational power. During actuation, the device has to generate the digital samples to be converted in the analog wave driving the actuator. The maximum wave's frequency is 1MHz. To avoid distortion, a relatively high number of samples per period has to be output. Considering that the samples have to be fetched by some memory location, the device must be able to work up to several tenths of MHz. The speed requirement is made even higher by the acquisition and processing. For example, during the acquisition the response of the structure has to be sampled at 10MSPS and in the resulting sampling period the device has to store the acquired datum in some memory location. A DSP appeared to be the best solution to all these issues, since it can execute heavy computation and it is oriented to signal processing. Given the need for operations in the analog domain, a mixed-signal DSP of the TMS320 family from Texas Instruments was chosen. The members of this family can work up to 150MHz and have excellent analog-to-digital conversion capabilities thanks to the integrated ADC. The sampling rate can be as high as 12.5MSPS, thus it meets the requirements.

The signal output by a PZT in response to a mechanical strain has an amplitude of few hundreds of mV. This signal must be amplified and filtered before being sampled by the ADC. So, another block in the Shimmer's architecture must be a signal conditioning circuit which amplifies the input wave and filters it to avoid aliasing in the analog-to-digital conversion.

The desired data series length is equal to 10^4 . A sample occupies two bytes, since the desired resolution is greater than 8 bits. As a consequence, a data series would take 20KB. The maximum internal SRAM memory on the DSPs of the TMS320 family is 40KB. A FFT on N samples needs 2N memory locations. Most SHM algorithms are much more complex than a FFT. This leads to interfacing the DSP with an external SRAM chip. In addition to the SRAM, another memory chip, by instance a SPI EEPROM, is necessary to store the DSP code.

Other components involved in the node's fundamental tasks are those needed by the actuation. In particular, a DAC is necessary to convert the digital samples

produced by the DSP into an analog signal. Then, an amplification circuit must be cascaded to the DAC in order to obtain the desired wave amplitude and drive the actuator. Finally, a radio module is necessary for data transmission and reception.

Given this basic architecture consisting in a DSP, memory chips, a radio transceiver and the signal conditioning stages for actuation and acquisition, issues related to power consumption were considered. The DSPs of the TMS320 family can enter power down modes, in which their current consumption is reduced by a huge factor with respect to their active state. However, they are complex devices and their stand-by consumption is much higher than that one of low power microcontrollers. Moreover, all other components' stand-by currents would add to the DSP's, resulting in a high global stand-by consumption. This would strongly reduce the duration of the active phase.

Thus, the architecture was modified introducing a low power microcontroller with the role of supervisor and developing a power management circuit which allows all parts of the node to be selectively disconnected from the power source.

In the final architecture (Figure 3.2), the microcontroller is interfaced to the radio module and to the DSP and it controls the provision of the supply voltage to all the other components. This results in a very low power consumption in sleep mode, ideally represented by the microcontroller's own consumption, and in a very high efficiency in the energy utilization, each component is turned on only when it has some task to perform.

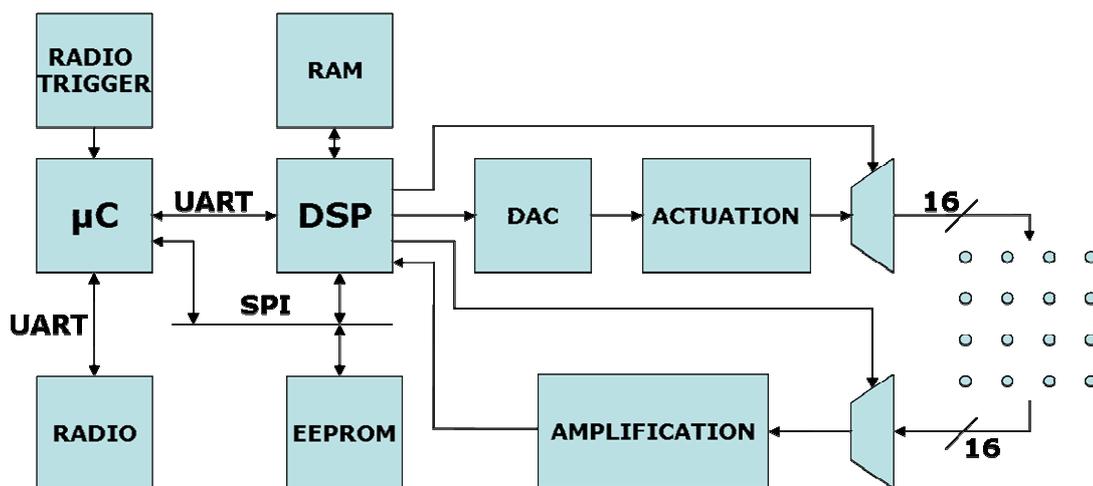


Figure 3.2 – Node's architecture.

Finally, a radio-triggering circuit is included in Shimmer to further minimize consumption during the sleep phase. The radio-triggering circuit is a completely

passive circuit which generates an interrupt signal to the microcontroller to wake it up when a proper radio signal is received. The radio-triggering circuit allows the microcontroller to enter its deepest sleep mode, while preserving a way to wake it up.

Details about all blocks composing the node’s architecture are given in the following sections.

3.3 Actuation circuit

The actuation circuit’s design has presented several challenges, due to its requirements. The actuation circuit has to drive the external actuating PZTs with the waveform output by the DAC. The PZTs devices used in Shimmer are essentially capacitive devices which show a capacitance around 6nF. The peak current drawn while driving a capacitor with a sine wave is given by

$$I_{MAX} = \pi f C V_{PP}$$

where f is the wave frequency, C is the load capacitance and V_{PP} is the peak-to-peak amplitude. Considering that the maximum actuation frequency is 1MHz, the peak current drawn during the actuation process is 380mA, assuming $V_{PP}=20V$. This high peak current and the high voltages involved in the actuation process made the choice of the components and architecture of the circuit quite challenging.

The circuit amplifying the DAC’s output was finally realized by a three-stage topology (Figure 3.3). The first stage works as a current-to-voltage converter, which gets the current output of the DAC and converts it to a voltage signal. The second stage works as an inverting adder and it shifts and amplifies the signal up to the required amplitude. The third stage is a buffer, whose purpose is driving the actuator.

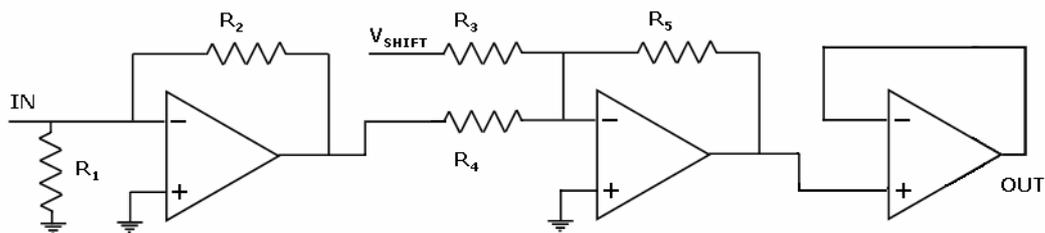


Figure 3.3 – Actuation path’s amplifier.

The first two stage are realized by a single component, integrating two opamps in the same package (Texas Instruments THS4082). The shift is made thanks to a voltage reference (MAX6037). The final stage is a high current output amplifier by Linear Technologies (LT1210).

The output of the amplifier can be used to drive any of the PZTs in the external grid, thus a selection mechanism is necessary. The first idea was to use a multiplexer, controlled by the DSP, to route the output signal to the desired PZT. This solution cannot be realized because of the characteristics of the output signal, that is high current and voltage. No commercial SMT multiplexer can bear such power dissipation. Thus, a new solution was studied and implemented making use of an array of relays (Figure 3.4). The relays can easily carry the currents involved in the actuation process and their maximum ratings in terms of voltage exceed the requirements by far. 16 relays (NEC-Tokin EF2) are mounted on Shimmer, one for each external PZT. Each relay is toggled by one of the outputs of a 4:16 multiplexer (MAX396), controlled by the DSP. In this way, the multiplexer carries only the relays' control current, which is smaller than the actuation's one by orders of magnitude. The relays cost about \$5 each.

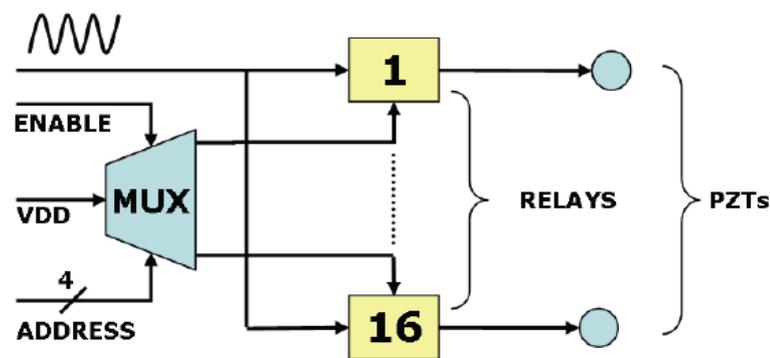


Figure 3.4 – Actuation path's selection circuit.

3.4 A/D signal conditioning circuit

The signal conditioning circuit in the analog-to-digital path has to accomplish two tasks. The first one is filtering the input signal to avoid aliasing. The second one is amplifying the signal in order to best match the input range of the ADC integrated in the DSP.

The sampling frequency is set to 10MSPS. The cutoff frequency of the anti-

aliasing filter must be less than the Nyquist frequency, which in this case is 5MHz. The input filter has been realized by two μ -power opamp from Texas Instruments (TLE2161) shipped in SOIC-8 package. These devices draw 280 μ A only. The input filter has been implemented with a 4 poles Sallen-Key topology, requiring only 4 resistors and 4 capacitors (Figure 3.5). The values of the external components was calculated by the FilterPro tool from Texas Instruments. The filter must be operated from a dual supply, since the output of the sensing PZT is expected to be a wave assuming both positive and negative values. The TLE2161s accept a supply voltage ranging from +/-3.8V to +/-18V.

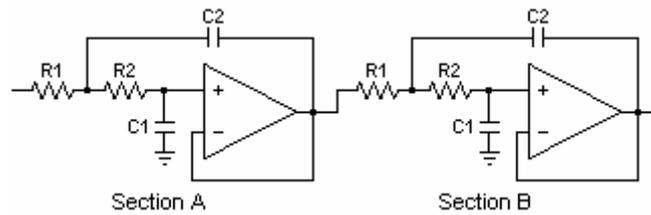


Figure 3.5 – Anti-aliasing filter.

After being filtered, the signal must be amplified. The expected amplitude of the PZT's output ranges from some tenth to a few hundreds of millivolts. The input range of the DSP's ADC is 0V-3V. So the amplifier's gain can vary from a 10 to 100, depending on the application. This gain requirement cannot be met by the TLE2161 because of their low gain-bandwidth value. Thus, the amplifier was realized by a THS4082 opamp from Texas Instruments. This opamp is characterized by a GBW=175MHz. Considering that the expected input frequency is the same as the actuation wave's one (up to 1MHz), the opamp meets the requirements. Furthermore, two opamps are integrated in the same package, so the gain can be distributed on the two of them (Figure 3.6). The current consumption is 2.9mA per amplifier at +/-5V. The amplifier topology consists in an inverting single input amplifier cascaded to an inverting adder. The adder is used to shift the signal, so that the amplifier's output has an offset equal to 1.5V. A MAX6037 voltage reference is used to provide the adder with its second input.

The amplifier's output is given by

$$\begin{cases} V_{OUT1} = -R_3 \left(\frac{V_{SHIFT}}{R_1} + \frac{V_{IN}}{R_2} \right) \\ V_{OUT2} = -\frac{R_5}{R_4} V_{OUT1} \end{cases}$$

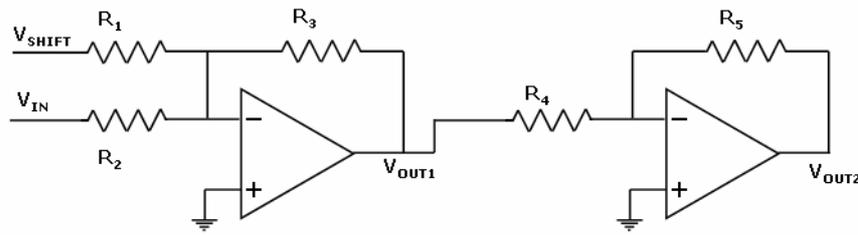


Figure 3.6 – A/D amplifier.

3.5 Energy management circuit

Shimmer is equipped with a an energy harvesting circuit which collects energy by means of solar cells and stores it in supercapacitors. No batteries are included in the node because of their lower durability and faster performance degradation. Achieving the maximum duration of the maintenance-free operational time is a key objective. The energy management circuit also includes the components needed to generate the various supply voltage required by the other parts of the node: the digital components, the analog-to-digital amplification circuit, the actuation circuit must be operated from different supplies.

The solar cells integrated in the node are produced by Powerfilm Inc. (MPT3.6-75) and they are made in thin film technology. This makes their thickness similar to a sheet of paper's one and, beside thickness, they share flexibility with paper (Figure 3.7). The solar panels sum up to an area of 100cm² and can produce up to 360mW in sunny conditions. Their cost is around \$8 each.

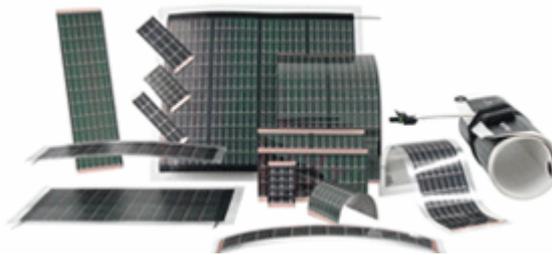


Figure 3.7 – Thin-film solar panels by PowerFilm Inc.

The supercapacitors are commercial devices belonging to the PowerStore Series manufactured by Cooper-Bussmann. Shimmer includes two different kinds of supercapacitors. The first one has a working voltage of 2.5V (B Series), the second one works at 5V (P Series). The use of two kinds of supercapacitors is motivated by

the need of obtaining clean supply voltages for the amplification circuit in the analog-to-digital part and of providing the converter generating the +10V and -10V supplies a suitable input voltage. In the first design of the energy management circuit 2.5V supercapacitors were used only. All voltage converters were connected to them. This caused two main problems. The high current peaks associated with the actuation resulted in glitches on the supercapacitors, which propagated to the other converters' outputs. Secondly, the high difference between the output voltage of the +/-10V converters and their input voltage made their functioning inefficient. As a consequence, the design was modified to accommodate two different kinds of capacitors. The 5V ones are used solely to power the +/-10V supplies, while the 2.5V provide energy for the rest of the components. In this way the +/-10V is more efficient and the actuation circuitry is completely separated from the rest of the node as for the power lines.

Two voltage regulators connected to the solar cells charges the supercapacitors. To prevent current from flowing back into the supercapacitor from the voltage regulators, a Schottky diode (BAT54 in SOT23 package) is cascaded to the voltage regulator. The diode's threshold is about 0.2V. The voltage regulator is a 2.7V regulator (TPS76618) for the 2.5V supercapacitors and a 5.2V (TPS76601) regulator for the 5V ones, so that the supercapacitors can be actually charged to their maximum operating voltage. No MPP tracking solution is implemented in Shimmer. However the energy harvesting circuit can be easily modified to accommodate circuitry dedicated to MPP tracking.

In addition to the harvesting circuit, the node includes the components needed to produce the suitable supply voltage for its different parts. The supply voltage for the microcontroller is produced by a boost converter (Maxim-Dallas MAX1674 in μ Max package) directly connected to the supercapacitor. The component was chosen because of its low minimum input voltage (0.8V) and its low current consumption (16 μ A) quiescent current and because it requires little external circuitry, consisting in an inductor and an optional diode (MBR0520). The output voltage of the MAX1674 can be set to 3.3V or 5V by connecting one its pin (FB) to ground or to the output. The DSP requires 1.8V for its digital core and 3.3V for the analog interfaces, while the signal conditioning stages require a dual supply (+/-5V and +/-10V) to generate the actuation wave and pre-amplify the signal produced by the sensing PZT. The 1.8V required by the digital core of the DSP are produced by a voltage regulator (TPS76618 in SOIC-8 package) cascaded to the boost converter producing the

microcontroller's supply voltage. The regulator was not cascaded to the 3.3V converter power the SRAM, DAC and I/O peripherals to better distribute the load current. The regulator is characterized by low dropout voltage and quiescent current (35uA). A dedicated boost converter (MAX1674) powers up the EEPROM, the RAM, the I/O interfaces of the DSP and the input and output multiplexers.

The +5V and -5V supplies are realized by a boost converter and an inverter. The boost converter is a MAX1674 in 5V configuration, while the -5V inverter is a LT3704 from Linear Technologies (Figure 3.8). The inverter is cascaded to the boost converter since its minimum output voltage is 2.5V and the component would not run if its input was connected to the 2.5V supercapacitor. The LT3704 requires some capacitors, some resistors used to set the output voltage and the switching frequency, a diode (1N5822), two inductors and a power MOSFET (FDS8880) as external circuitry.

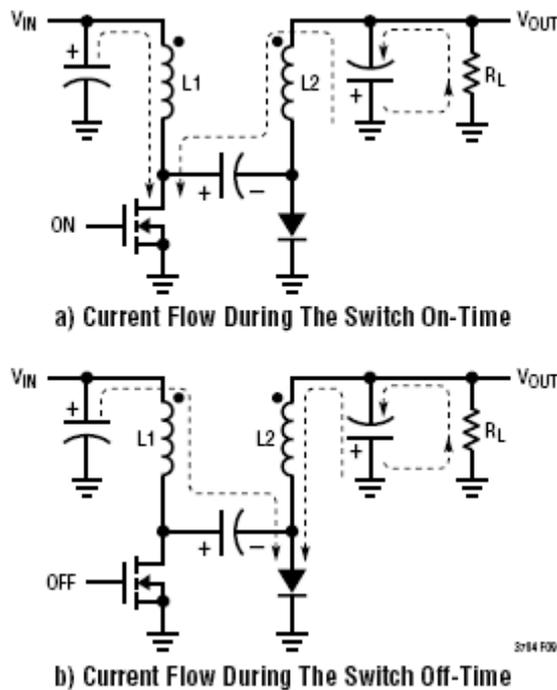


Figure 3.8 – Operation of the LT3704 DC-DC inverter.

The +10V and -10V are generated by a LT3479 and a LT3704 respectively, both from Linear Technologies. Each of them can deliver an output current more than 1A while producing the required +/-10V. This is fundamental to drive the actuation circuit, which draws very high peak currents. In this case the LT3704 is not cascaded to the positive counterpart, because the input voltage is higher.

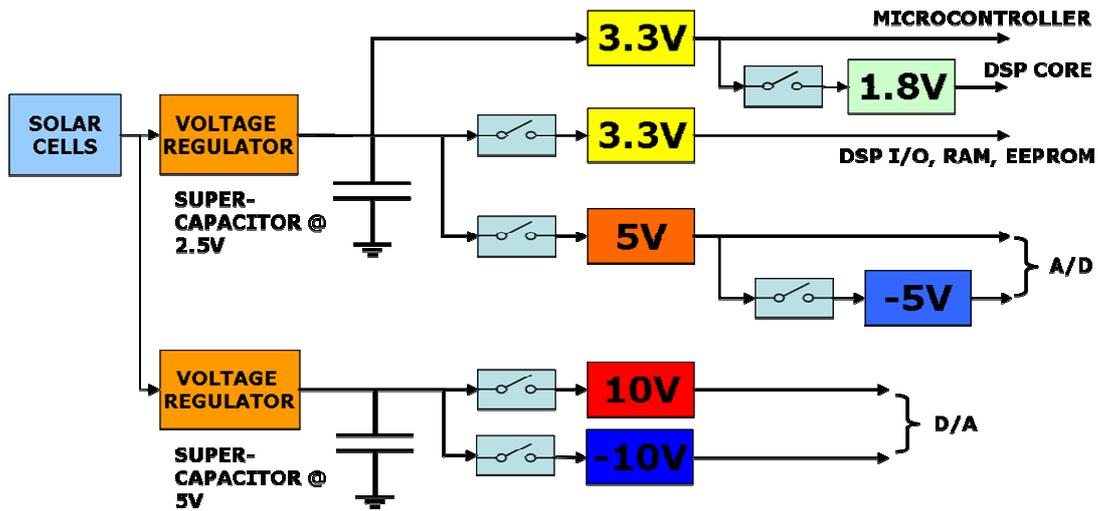


Figure 3.9 – Block diagram of the energy circuit.

3.5.1 Power distribution

The main focus in the development of the node has been on the goal of maximizing the energy available for a single session of data acquisition and processing per day. To attain this objective, the architecture of the node has been structured in such a way that the power consumption be near zero during the phase of inactivity, corresponding to most of the day. This has been achieved both by an accurate choice of components and by the implementation of power-aware solutions in the node architecture.

More exactly, a network of CMOS switches (MAX4676) have been inserted between the supercapacitors and the voltage converters. The CMOS switches allow to selectively disconnect each voltage converter from the supercapacitor, turning them off and zeroing their current consumption. They are controlled by the microcontroller supervising the node's activity by a single input signal each. The CMOS switches used in Shimmer are the MAX4626, manufactured by Maxim-Dallas Semiconductors. They are shipped in a SOT23-5 package, so their footprint is very small. They accept 3V to 5V as a supply voltage and their leakage current is in the range of 10^{-9} A, so they do not increase the global power consumption. Their very low resistance (0.5Ω) results in a small voltage drop across their terminal and low power dissipation.

3.6 Radio triggering circuit

The radio triggering circuit is an essential part of the node, which contributes to the reduction of the energy consumption during the sleep phase. Common WSNs paradigms require the sensor nodes to wake up periodically to listen for incoming queries or data packets on the radio channel. This paradigm is necessary in some applications, but in some other ones it produces a waste of energy, since the nodes' consumption in the wake up state is unavoidably higher than in sleep mode. In addition, the periodical wake up has to be triggered by a timer, which can't but keep running also during the sleep phase. Most microcontrollers can't enter the deepest of their power down modes, if one of their internal timers has to be kept on. Thus, the sleep mode consumption is not minimized. Moreover, after waking up, the node has to activate the radio module to listen for incoming messages, and radios are often the most energy-hungry devices in wireless sensors.

Eliminating the periodical wake up triggered by a timer can have significant benefits in terms of power, but it poses the problem of identifying an alternative way to wake up the node. The solution adopted by Shimmer is represented by the inclusion of a radio-triggering circuit in the sensor node.

A radio-triggering circuit is a circuit which detects radio signals emitted at a given frequency and, in case these signals reach a certain threshold, generates an interrupt signal to wake up the microcontroller. The incoming radio signals are generated by some external agent when it wants to query the node. This solution perfectly needs fit the needs of the specific application in which Shimmer is being employed. Indeed, the node is supposed to perform an analysis on the structure integrity only if the UAV requires so.

Several implementations of a radio triggering circuit can be thought of, both relying on active and on passive components. A solution employing active components could be based on an opamp amplifying the incoming radio signal, the output of which is fed into the interrupt pin of the microcontroller. Although this solution can be effective, it has two main drawbacks. Firstly, the active components would draw some current during the sleep phase, so part of the power consumption reduction obtained by putting the microcontroller in a deeper sleep state would be lost. Secondly, ultra-low power opamps should be used and these kind of device can be expensive, compared to passive components.

The amplification effect could be obtained by means of a transformer (Figure

3.10). In this way, a radio-triggering circuit based on purely passive components could be designed. Unfortunately, this solution has some problems, too, mainly related to the characteristics of the transformer. Indeed, the transformer should have a high primary-to-secondary ratio and operate at high frequencies, typically 400MHz at least, since the radio-triggering circuit is likely to share the antenna with the radio module. Commercial transformer do not offer this level of performance. A model among SURFCOIL transformers by Sprague-Goodman Electronics operate at a maximum frequency (-3dB attenuation) of 340 MHz in case of a 1:6.33 turn ratio. Another one works up to 500MHz, but in this case the turn ratio is 1:1. Furthermore, transformers are not compact devices and they could have a negative impact on the form factor of the sensor node.

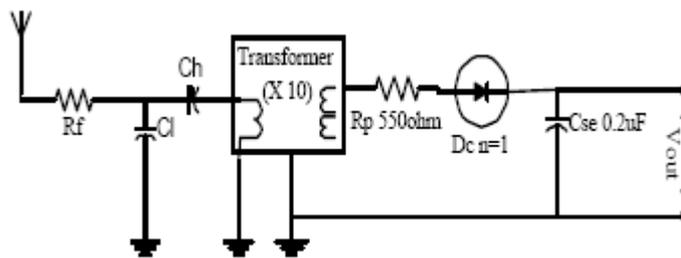


Figure 3.10 – Transformer-based radio triggering circuit.

The radio-triggering circuit embedded in Shimmer is a circuit made of purely passive components, connected to a planar antenna working in the 2.4GHz band and to an external interrupt pin of the microcontroller. The circuit consists in a network of diodes and capacitors, which rectifies the signal received by the antenna (Figure 3.11). The diodes are low barrier X-band Schottky diodes (MA4E2054B) manufactured by Macom. The MA4E2054B-287T part number corresponds to two diodes in series integrated in a single SOT-23 package. The capacitors are 100pF ceramic capacitors. The output of the circuit is a DC signal the amplitude of which is equal to the amplitude of the incoming AC signal multiplied by a factor which depends on the complexity of the circuit. Ideally, the higher the number of the stages of the network, the higher the multiplication factor. In practice, leakage currents and non-zero thresholds of the diodes impose an upper bound to the obtainable amplification.

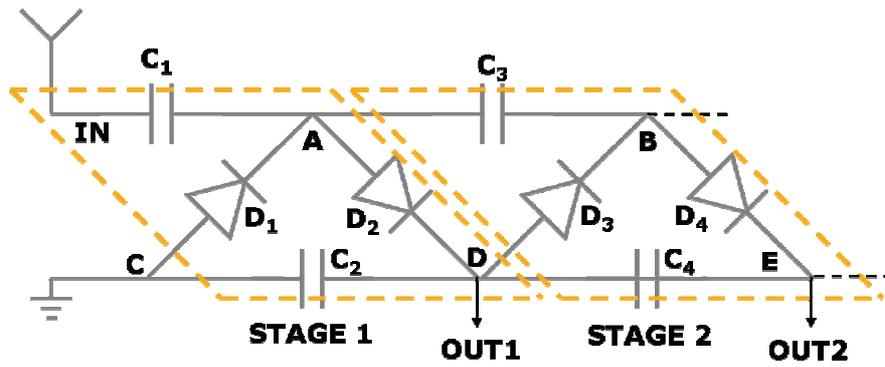


Figure 3.11 – Radio triggering circuit.

The radio triggering circuit works as an AC-DC voltage multiplier. Let V_M be the amplitude of the input AC signal. When the first negative peak arrives at the input node, capacitor C_1 is charged to V_M ($V_{IN}-V_A=V_M$) thanks to the current flowing through D_1 . When the following positive peaks arrives, capacitor C_2 is charged to $2V_M$ ($V_D-V_C=2V_M$). Then, the successive negative and positive peaks will charge C_3 to V_M ($V_A-V_B=V_M$) and C_4 to $2V_M$ ($V_E-V_D=V_M$) respectively (Figure 3.12). That is, the OUT node of the n-th stage is charged to $2nV_M$.

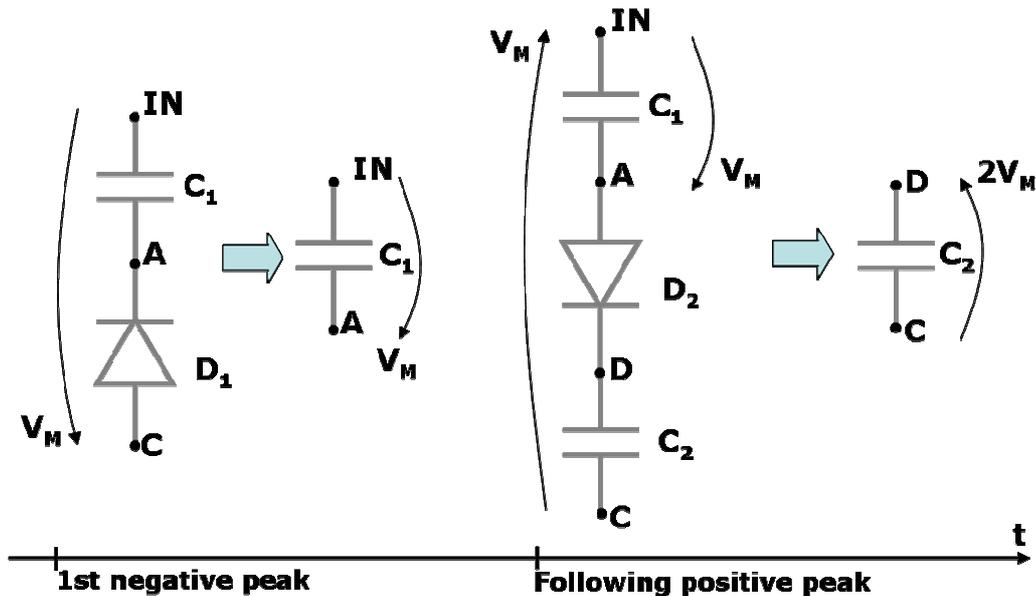


Figure 3.12 – Functioning of the radio triggering circuit.

The microcontroller requires 2.0V to trigger an interrupt. As mentioned, the output voltage depends on the complexity of the circuit. Thus, depending on operational conditions, the circuit can be tuned to obtain the required output level. The number of components can be kept low using chips integrating two diodes in the same package. In case the non-idealities do not allow to achieve the desired voltage gain, the output of the triggering circuit can be used to turn on the gate of a MOSFET

acting as a pull-up for the interrupt pin of the microcontroller (Figure 3.13). Since MOSFETs can have a threshold voltage lower than 2V, this further reduces the minimum voltage the circuit needs to output to wake-up the node. As a consequence, the power of the radio triggering signal can be reduced as well. In Shimmer a low-threshold MOSFET by Advanced Linear Devices is used (ALD110902). The device is characterized by a threshold voltage as low as 0.2V. The gate leakage current is only 30pA and the drain-source one 100pA, so the device does not have any significant effect on the current consumption of the node in sleep mode. In theory, the MOSFET's gate should be discharged after the interrupt generation, in order to turn off the MOSFET and reset the interrupt. This could be easily done by connecting the gate to a pin of the microcontroller. However, test have showed this is not necessary, because the gate gets discharged as soon as the incoming radio signal ends.

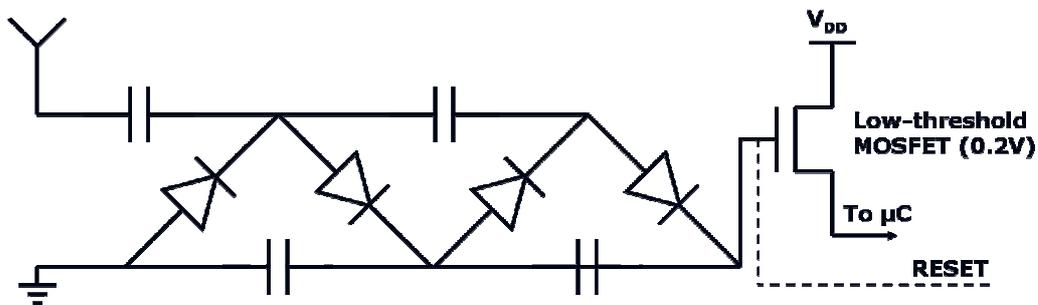


Figure 3.13 – The radio-triggering circuit with a cascaded low-threshold MOSFET.

3.7 Digital core

All parts of the sensor node described in previous paragraphs are controlled by and depend on the digital part of the sensor node. This part contains the computing components, memories and radio module. More exactly, a microcontroller with the role of node's supervisor, a DSP by which heavy SHM algorithms can be executed, a SRAM and an EEPROM chips providing the required data storage capacity and a ZigBee radio, which makes Shimmer ready for the integration in a wide range of sensor networks.

3.7.1 Microcontroller

The first component with computational capabilities is an ATMega128L

microcontroller manufactured by Atmel. The device has been chosen for the good power performances, power management features and for the reliable programming tool-chain.

The device is shipped in two different packages: TQPF and MLF. The first one has been used in the realization of the prototype, since it is much easier to solder with common laboratory tools (Figure 3.14). The microcontroller is based on an 8-bit RISC architecture and it can run up to 8MHz with a 3V supply voltage. The maximum power consumption in active mode at 8MHz is less than 20mA. Since the microcontroller does not perform heavy computational tasks, the ATmega128L clock is set to 4 MHz resulting in a current consumption of about 5 mA. The internal memory is composed of 128KB Flash memory, 4KB EEPROM and 4KB SRAM. This allows the microcontroller to run non-trivial algorithms. As for I/O capabilities, its 64 pins, organized in 5 8-bit wide I/O ports, make it ready for the interconnection with multiple external devices. The I/O ports pins can be individually multiplexed to specific functions, such as UART connection management or analog to digital conversion. In particular, the microcontroller has 2 UART ports, 1 SPI port and 8 channels for analog signals' sampling.



Figure 3.14 – Atmel ATmega128 in TQPF package.

The role of the microcontroller is supervising the activity of the node. Indeed, it is the only digital component which is powered during the sleep phase of the node. Thanks to the radio-triggering paradigm, the microcontroller can enter the deepest of its power down modes, in which all internal peripherals are turned off. This reduces the power consumption of the component to less than 5uA. All peripherals being turned off, the only way to wake up the microcontroller is by feeding a signal of at least 60% of the supply voltage into one of the external interrupt pins (INT0:7). In case the supply voltage is 3.3V, the required triggering signal's amplitude is 2.0V.

The pin chosen to be connected to the output of the radio-triggering circuit is INT0, chosen among INT0:3, which are the only suitable ones. Pins INT4:7 can be either used as edge interrupt pins or level interrupt pins. The desired behavior is edge detection, but pins INT4:7 need an I/O clock to be turned to accomplish the task and this would cause a higher power consumption. Another requirement for the wake up signal is a minimum duration of 1 μ s. The requirement is met easily, since the output of the radio-triggering circuit remains high as long as the external agent keeps transmitting the radio signal.

After waking up, the microcontroller has to power the rest of the node. This operation is preceded by a handshaking with the external agent via the radio module. The microcontroller powers the radio by closing a CMOS switch which is connected between the supply voltage line (the same used by the microcontroller itself) and the radio supply pin. The control pin of the switch is connected to pin PC0. The handshaking is necessary to avoid the node is powered up as a consequence of false positive triggering signals. The microcontroller is connected to the radio module via UART0.

Another preliminary operation is checking the charge level on the supercapacitors via A/D pins PF2:3. If the charge level is too low to perform a complete analysis, the microcontroller does not start to power up the node.

Pins PC1:5 control the CMOS switches inserted between the supercapacitors and the voltage converters providing the supply voltage to all other components of the node. The CMOS switches are closed sequentially and the microcontroller verifies that each voltage converter is producing the expected voltage by some of its A/D pins (PF0:1 and PF4:5), which are connected to the output of the 3.3V, 1.8V, +5V and +10V voltage converters respectively. This allows the microcontroller to sample the value output by the converters and detect any malfunctioning. The outputs of the +5V and +10V converters are scaled down by resistor divide.

The last CMOS switch enables the 1.8V regulator which powers the digital core of the DSP. Once woken up, the DSP handshakes with the microcontroller via UART connection (UART1 for the ATmega128) and the node's control is ceded to the DSP. During the actuation, acquisition and calculation phases, the microcontroller simply mirrors the messages sent by the DSP and received on UART1 onto UART0, so that they are transmitted to the external agent. After all operations are completed, the microcontroller turns off the CMOS switches enabling the power supplies and enters power down mode.

The ATmega128L is also connected to the on-board EEPROM chip via the SPI port. This connection has been established to make the DSP code updatable. Indeed, after the microcontroller turns on the 3.3V supply, the EEPROM content can be modified by writing the new code, which will be loaded by the DSP after its digital core is powered.

3.7.2 DSP and DAC

The computational power of Shimmer reside in a 32-bit DSP. While the microcontroller supervises the node's global activity, the DSP performs the heavy computation required by structural health monitoring algorithms. The DSP embedded in Shimmer is a mixed-signal processor (TMS320F2812) from Texas Instruments. The device is shipped in a PGF package with 176 pins. The processor was chosen because of its high speed combined with relatively low power consumption and its I/O capabilities. Another key feature determining the choice is its high speed on-chip ADC ().

The TMS320F2812 can work up to 150MHz. The clock is internally generated by an on-chip oscillator and a PLL divider which require an external crystal to work only. The clock signal is distributed to the digital core directly and to the peripherals through two clock lines, dedicated to high speed and low speed peripherals respectively. A pin (XCLK) is reserved to output a clock signal, too. This is useful in case some external peripherals needs to be provided with it.

The memory is comprised of non-volatile and SRAM blocks organized in several areas. The Flash memory capacity is 256KB (128Kx16), while the SRAM capacity is 36KB. The SRAM is divided in a M0 and a M1 (1Kx16 each), a L0 and L1 (4Kx16 each) and a H0 (8Kx16) areas. The memory space can be expanded by an external interface (EXTINF). This interface is comprised of 16 data lines and 19 address lines and allows the DSP to be connected directly to external memory chips with a capacity of up to 512KB. Five ranges of addresses are mapped on the XINTF. Reading from and writing to these five ranges activates distinct control signals among those included in the XINTF, so up to five different chips can be interfaced to the DSP. The XINTF is highly configurable to match the timing requirements of the external memories. In Shimmer, the DSP is interface to an external asynchronous SRAM whose size is 128KB.

In addition to the XINTF, a vast I/O equipment make the DSP suitable for its integration in a complex system architecture, such as Shimmer's one. The DSP can communicate through various I/O ports. UART, SPI, eCAN ports are available. All I/O pins can be individually multiplexed to one of these functions. The DSP makes use of an SPI port to communicate with an external EEPROM chip and of an UART port to communicate with the microcontroller. If no special function is enabled for an I/O pin, the pin works as a general purpose I/O pin. 21 pins are used as GPIO in Shimmer. 12 pins are used to output the value to be fed into the DAC as later described. 5 pins are connected to the actuation multiplexer and 4 pins to the sensing multiplexer.

The CPU is the point of strength of the DSP. Its 32x32-bit MAC capabilities and its 64-bit processing capabilities, enable the CPU to efficiently handle higher numerical resolution problems that would otherwise demand a more expensive floating-point processor solution. The CPU has an 8-level-deep protected pipeline with pipelined memory accesses. This pipelining enables algorithms' execution at high speeds without resorting to expensive high-speed memories. Special branch-look-ahead hardware minimizes the latency for conditional discontinuities. Special store conditional operations further improve performance.

As mentioned, an important feature of the DSP is the integrated ADC. This peripheral is essential for Shimmer operation, since it is used to sample and digitize the structure's response after actuation. The ADC can work up to 12.5MSPS and it has a resolution of 12 bits. Such specifications allow Shimmer to meet the requirements regarding acquisition in terms of both time resolution and accuracy. 16 A/D channels are available.

The DSP can be configured to fetch the code from various locations at bootstrap: internal Flash, external EEPROM, generic external device via serial or parallel port. The internal Flash memory could be used to store the code, but programming the Flash memory is a complex task. Since a desired feature is the possibility to update the algorithm run by the DSP, the solution adopted for Shimmer is storing the code in an external SPI EEPROM chip. This can be easily accessed and its content modified by the ATmega128 microcontroller.

As for power requirements and consumption, the DSP digital core is powered off a 1.8V supply voltage, while the peripherals require 3.3V. The current consumption of the digital core is around 200mA at 150MHz (Figure 3.16). Peripherals draw a current which depends on their utilization.

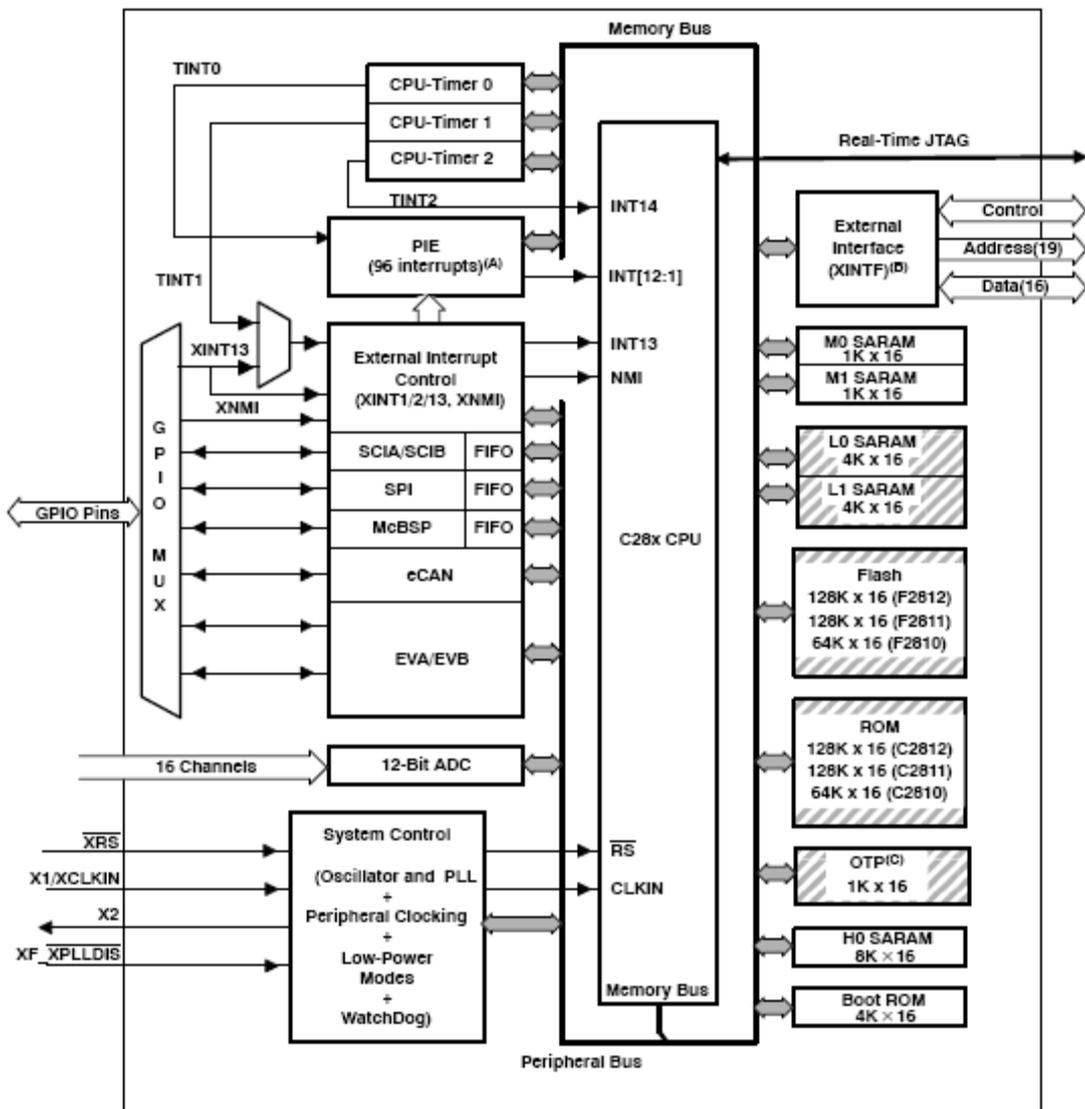


Figure 3.15 – TMS320F2812’s functional block diagram.

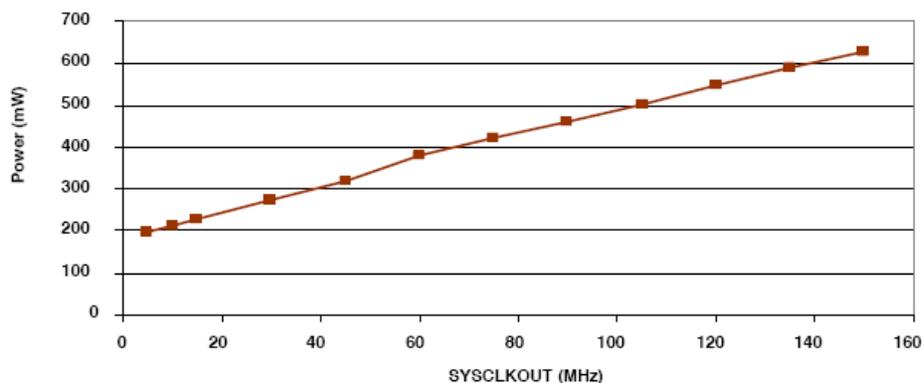


Figure 3.16 – Typical total power consumption vs. clock frequency for the TMS320F2812.

In addition to the external EEPROM and SRAM, the DSP is interfaced to a DAC from Burr-Brown (DAC902). The device is a high speed (165MSPS) digital-to-analog converter with 12-bit accuracy and output current which can be operated

from a 3.3V supply (Figure 3.17). The chip has to be provided with a clock signal. In Shimmer, this is output by the XCLK pint of the DSP. The samples to be converted are obtained through a parallel port. The device requires little external circuitry, consisting in some capacitors for decoupling purpose and a resistor which sets the range of the output current. This range can vary from 2mA to 20mA. In Shimmer, the resistor is chosen so that the 20mA range is selected. This allow to reduce the gain of the signal amplification circuit in the actuation path. The dynamic performance of the DAC is excellent: the SFDR ranges from 77dBc for a 1MHz output frequency at $f_{CK}=25\text{MHz}$ to 53dBc for a 55MHz output frequency at $f_{CK}=165\text{MHz}$.

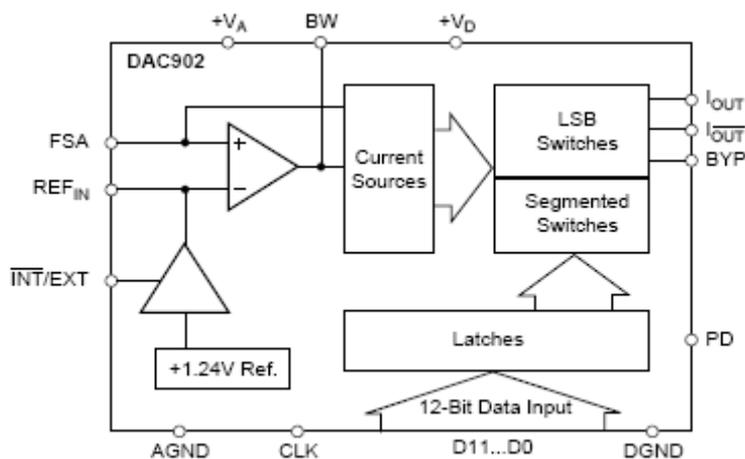


Figure 3.17 – DAC902's functional diagram.

3.7.3 Memories

Two memory chips are mounted on Shimmer. The first one is an SPI EEPROM chip (25AA640) produced by Microchip, the second one is a SRAM chip (CY7C1021CV33) from Cypress Semiconductors.

The EEPROM chip stores the DSP'S program code. The memory is interfaced to the DSP via SPI. At bootstrap, the DSP automatically fetches the code from the EEPROM by generating the required clock signal and commands. The memory chip capacity is 64Kb (8192x8) and it supports clock frequencies on the SPI port up to 1MHz. The device can operate with a 3V supply voltage and draws 3mA during write operations and 500uA during read operations. It can also enter a power down mode characterized by a current consumption as low as 500nA. It is available in many packages. The SOIC-8 was chosen (Figure 3.18).

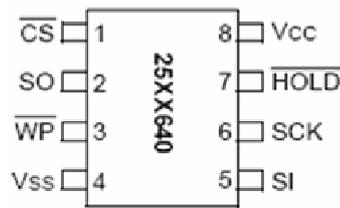


Figure 3.18 – 25A640 SPI EEPROM pinout.

The principles of operations of the SPI EEPROM are based both on hardware and software control. From the hardware point of view, the device is accessed via the SI pin, with data being clocked in on the rising edge of SCK. The CS# pin must be low and the HOLD# pin must be high for the entire operation. Data is sampled on the first rising edge of SCK after CS# goes low. If the clock line is shared with other peripheral devices on the SPI bus, the user can assert the HOLD# input and place the memory in “hold” mode. After releasing the HOLD# pin, operation will resume from the point when the HOLD# was asserted.

Software control of the device’s activity is performed by a command set, which includes read (READ), write (WRITE), write enable (WREN), write disable (WRDI), read status register (RDSR) and write status register (WRSR) commands. READ command has to be followed by the address of the location to be written. After receiving the command-address pair, the memory outputs the requested value. WRITE command must be preceded by a WREN command, otherwise the command has not effect. The command has to be followed by the destination address and by byte’s value. After a successful writing, the write enable latch is reset. The device allows a single byte or a page (32 bytes) to be written. A page can be written simply going on sending the bytes’ values, with CS# low and a valid clock signal on SCLK. The internal status register can be read to check if a writing operation is in progress or if writing is enabled. Two bits in the control register determine the level of protection of the memory content. In particular, high, low or all addresses can be protected.

Microchip produces some other SPI EEPROM chips with higher capacity which fully pinout compatible with the chip mounted in the prototype. This makes future memory expansions very easy. A memory expansion could be necessary in case custom actuation waveform have to be generated. In this case samples should be stored in a non-volatile memory and the SPI EEPROM would be the main candidate. A high number of samples is required to achieve a good accuracy in the waveform generation. Thus, the 8192B capacity of the 25AA640 could be not enough.

The second memory chip mounted on Shimmer is high speed, low power asynchronous SRAM, organized in 64K 16-bit locations (1Mbit). The chip works with a 3.3V supply voltage and is accessible via a parallel I/O port, comprised of 16 lines for data and 16 lines for addresses (Figure 3.19). Reading and writing the memory content are very easy tasks. In both cases the desired address has to be set on the input pins of the chip. If the write enable (WE#) pin is low, the writing mode is entered and the chip writes in the selected location the input value, while if it is high, the memory content pointed by the address is output on data pins. The SRAM is shipped in three different versions, corresponding to different access times. The available options are 8ns, 10ns, 12ns and 15ns. The one chosen for Shimmer is 12ns, since the DSP can work up to 150MHz, that is with a $T_{CK}=6.67ns$. An access time equal to 15ns would make the DSP wait for three clock cycles before reading the value, while the other options make a read cycle possible in only 2 DSP clock periods. Finally, the 8ns and the 10ns access times are not recommendable because they wouldn't produce any performance's improvement and, by contrast, they would increase the power consumption.

The role of the SRAM chip is storing data both during the acquisition and at runtime, when DSP is running the analysis algorithm on the acquired data. Its working speed and the EXTINT interface of the DSP allow samples to be stored in real-time during acquisition.

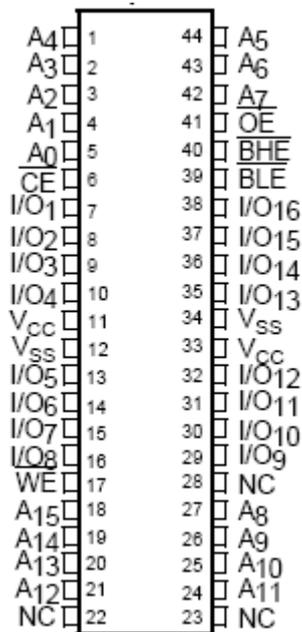


Figure 3.19 – CY7C1021CV33 pinout.

3.7.4 ZigBee radio

The radio mounted on Shimmer is a ZigBee-compliant module produced by Maxstream (Xbee OEM module). The device operates in the 2.4GHz band. The radio costs less than \$20 and it comes with a variety of antenna configurations. The radio consumes between 45mA and 50mA for receive and transmit operations, when powered by a 3.3V supply. Typical range for the radio is 90m line-of-sight and 30m in a typical indoor setting. The ease of use of the XBee radio was the main driver for its integration into the wireless sensor node: the only requirements beside the power and ground lines are the connections to the UART TX/RX pins. The microcontroller can be directly connected to the radio and communicate with external devices through a transparent serial connection. The maximum bit rate is 250kbps. The form factor, including connectors and on-chip antenna, is 2.4x2.7cm (Figure 3.20).

Radio modules with slightly lower consumption are available off-the-shelf. Their use has to be considered for the development of future versions of the node. However, the XBee radio was considered the best choice for the realization of the first prototype.

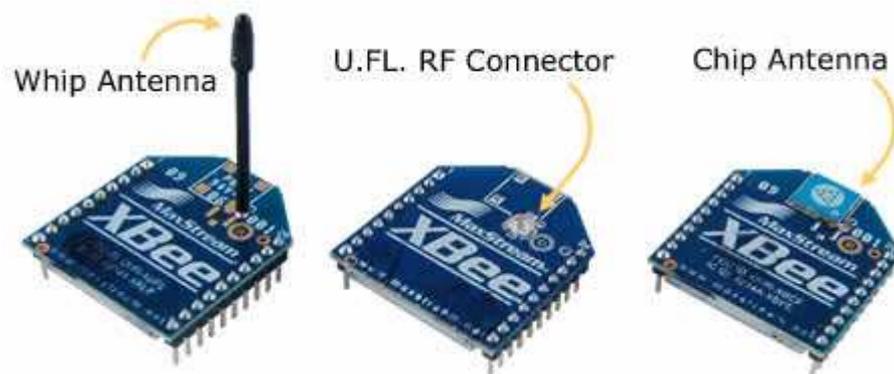


Figure 3.20 – XBee radio in the available antenna configurations.

4 Software development

In addition to the design of wireless node's architecture, the thesis work consisted in the development of the software running on the on-board microcontroller and DSP and of a graphical interface used to interact with the node by a PC. This chapter focuses on the results of such development and discusses the issues faced during it.

4.1 Microcontroller's code

The code running on the ATMega128L microcontroller mounted on Shimmer was written in the AVR Studio IDE, distributed for free by the microcontroller's manufacturer. The IDE is a complete development and programming environment which allows the developer to write the code both in Assembler and C languages, thanks to the integration with the GNU C++ toolkit. Since no particularly demanding task is required of the microcontroller, the code was entirely written in C. The microcontroller's code essentially deals with the node's state management and with the communication with the external agent (Figure 4.1).

First of all, after the wake up caused by the incoming radio-triggering signal, the microcontroller sends an handshaking packet to the external agent. The expected answer is a packet containing some settings regarding the analysis to be performed by the node. For instance, the actuation frequency, the actuator and sensor indexes and actuation duration. If the packet is not received before a preset timeout, the microcontroller simply goes back to power down mode and no operation is performed. This avoids the whole node wakes up as a consequence of a spurious radio signal received by the radio triggering circuit. If the handshake with the external agent is successful, the microcontroller start the power-up sequence, by checking the supercapacitors' level of charge and closing the CMOS switches which enable the voltage supplies. This is done sequentially and the microcontroller

performs a check after every step. If the result of any step is negative, the sequence is interrupted, reversed and the node goes back to sleep after sending an error message. For example, if the voltage on the supercapacitors is too low to perform the whole damage detection analysis, then the microcontroller does not turn on the voltage converters. Or if any supply does not work properly, the problem is signaled to the external agent and the node is powered down. All these verifications are executed by the ADC integrated in the microcontroller. The microcontroller takes care of waiting the settling time of supplies' output before checking them.

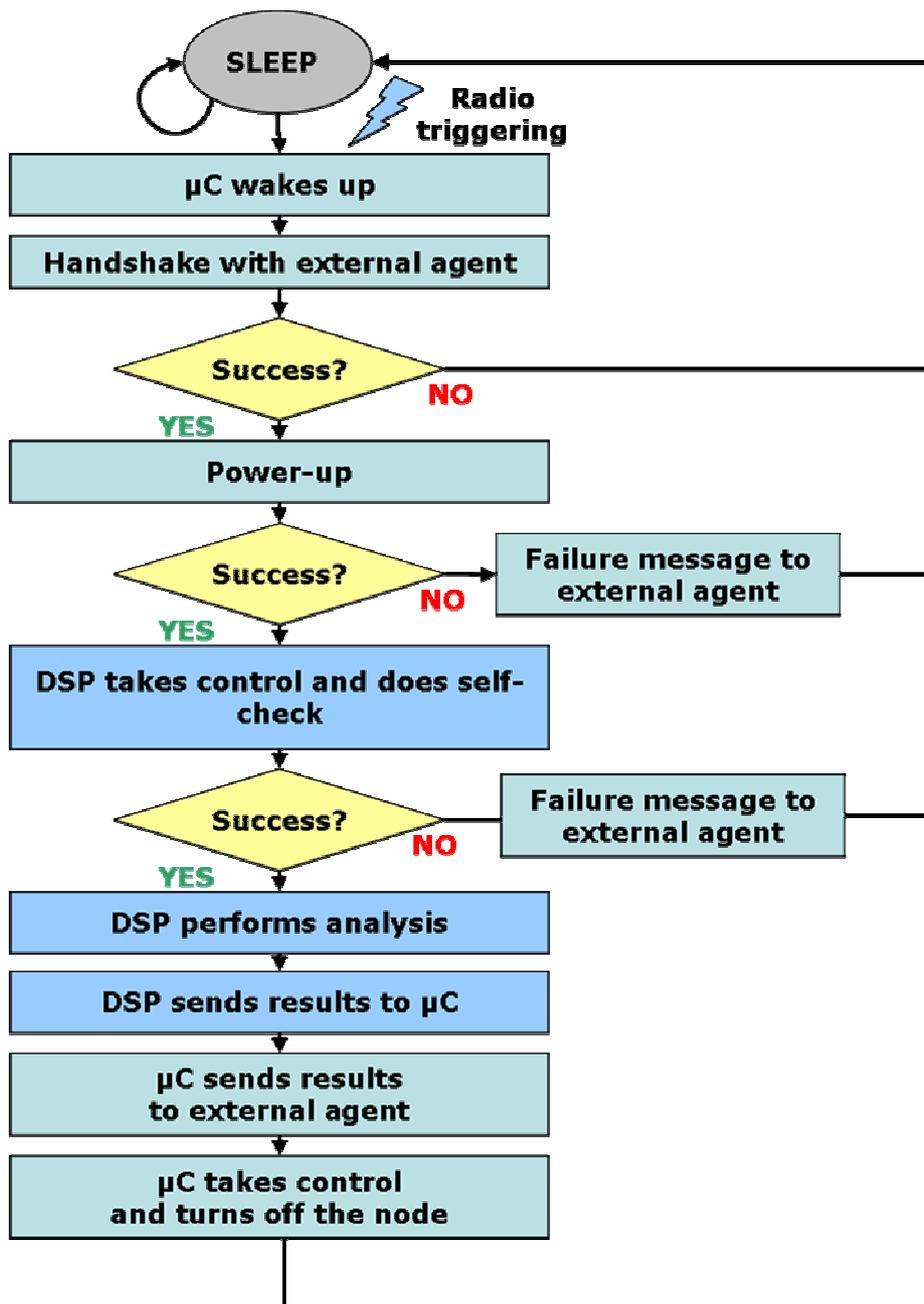


Figure 4.1 – Flow chart of the behavior of the node.

After the 3.3V converter is successfully enabled, the microcontroller checks the EEPROM and turns on the 1.8V regulator powering the DSP. At this point, an handshake is performed with the DSP and if this is successful, the microcontroller enters a state in which it simply mirrors the message sent by the DSP and received on the UART0 to the UART1 so that they are transmitted to the external agent. When the message communicating the end of acquisition is received, the microcontroller turns off the +10V and -10V supplies in order to reduce power consumption. The microcontroller goes out of this state when it receives the message signaling the end of the analysis. A software time-out prevents the node from remaining awake for ever, in case the DSP does not send the message, either because of software or hardware failure.

The development of the microcontroller's code presented no particular issues except for the management of the interrupt routines associated with the serial ports. In a first version of the code, the interrupt routines made nested calls to the functions required to parse the incoming data packets. The execution of some functions is lengthy, especially if the transmission of a packet back to the external agent is required, since it has to deal with the timing of the serial port. As a consequence, the microcontroller was unable to serve interrupts for too long time intervals. This resulted, for example in losing new incoming bytes.

To fix the problem, the code was modified by implementing a simple message queue: every macro-operation is associated with a flag; the interrupt routines set the flag corresponding to the operation which has to be performed in response to the interrupt; then, these flags are checked in the main routine and the related function is called if necessary. This solution reduces the maximum nesting level of the function calls and improves the quality of the interrupt handling. After the adoption of such mechanism, no more incoming packets went lost.

The communication protocol and management are described in the following paragraph.

4.1.1 Communication protocol

The communication between the node and the external agent is based on a custom protocol, which defines the data packets' format and a command set. Every packet exchanged by the node and the external agent is comprised of 6 fields (Figure

4.2):

- **START**. This field is 1 byte long and it contains a fixed start delimiter character.
- **OPCODE**, which is 1 byte long and contains the command transmitted.
- **DATA_LENGTH** specifies the length of the payload and it is 1 byte long, too.
- **CHECKSUM**. This field contains 1 byte calculated as the bit-wise XOR of the **OPCODE** and **DATA_LENGTH** fields. Its purpose is the implementation of a simple error detection technique.
- **DATA**. This field can have a variable size, up to 255. Its length is specified in **DATA_LENGTH**.
- **END**. A 1 byte fixed value delimiting the end of the packet.



Figure 4.2 – Packet’s structure as defined in the custom protocol.

At the level of the microcontroller’s code, the protocol is managed by a dedicated state machine and two 258-byte buffers (a Tx buffer and a Rx buffer). The state machine transitions are triggered by the reception of a new byte on the serial port. In the reset state, the machine waits for the **START** character. When a new byte is received, if it matches the **START** value, the state machine switches to the following status. Otherwise, the byte is discarded and the state does not change. After the reception of the **START** value, the machine waits for the **OPCODE**, **DATA_LENGTH** and **CHECKSUM** bytes. If the received checksum is compared with the bit-wise XOR of the **OPCODE** and **DATA_LENGTH** calculated locally. A mismatch between them is interpreted as the result of a reception error or packet corruption and a request to re-transmit the packet is sent to the external agent. By contrast, if the two checksums match, the three bytes are copied to the Rx buffer. Then, the next bytes received on the serial port are copied to the Rx buffer and counted as long as their number equals the **DATA_LENGTH** field content. Finally then **END** character is expected. If it is not received, an error message is sent back to the external agent. In case it is received, the packet is considered well formed and a routine is called to analyze it and execute the command. Basically, the routine parses the content of the Rx buffer and calls the function which executes the command

indicated by the OP_CODE field. If the function requires some parameters, they are fetched from the DATA field in the Rx buffer.

The communication protocol allows future expansions of the commands set to be done with little effort. A new function to execute the command can be written and a new value added to the *switch{}* statement used to recognize the command and call the function.

4.2 DSP's code

The DSP's code implements a finite state machine, as well. After initialization, the DSP sends an handshake message to the microcontroller via UART0. The expected response contains the analysis settings mentioned in the previous section, i.e. actuation frequency, actuator and sensor indexes and actuation duration. After receiving the packet with these settings, the DSP starts the analysis procedure, consisting in actuation, acquisition and data processing. Finally, it sends the results to the microcontroller. At the end of every intermediate step, the DSP sends a message to the microcontroller to signal the procedure is progressing.

The communication between the DSP and the microcontroller relies on the same protocol as the one used to communicate with the external agent. The only difference is represented by a dedicated commands set.

The DSP's code development was done using a *eZdsp*TM kit. The kit includes a development board and a copy of the Texas Instruments Code Composer Studio Ideate full integration of the IDE and the board programming tools eases the debug of the code, which can even be done in real-time. The IDE can compile code written both in Assembler and in C.

As opposed to the microcontroller, the development of the DSP's code has presented significant challenges, related to the actuation and acquisition processes. During actuation, the DSP has to provide the DAC with the samples of the waveform which determines the PZT's vibration. This waveform is in the standard case a sine wave. The actuation frequency can be as high as 1MHz. Thus, the calculation of the samples is not feasible in real-time since it necessitates the evaluation of trigonometric functions in floating point and several samples per period must be output to avoid excessive distortion. The use of a LUT is mandatory. The TMS320x DSPs are shipped with an internal ROM containing the samples of several functions.

The sinusoidal waveform is one of them, so, the LUT can be built by retrieving the samples from the built-in table. Even using a LUT, the generation of the samples at a sufficient rate is not a trivial task. The LUT contains the samples of a period of the wave. Usually multiple periods must be generated during actuation. Therefore, the LUT must be parsed more than once. The first way to implement such behavior is by nesting a *for* cycle inside another one, where the inner cycle selects the sample from the LUT by increasing a pointer and the outer one counts the number of periods. This solution is not acceptable because tests showed that the overhead introduced by the branch condition of the outer cycle and the reset of the pointer being incremented by the inner one prevents a smooth connection between successive periods of the sine wave. Moreover, tests showed that the program resulted from the compilation of the C source of the algorithm did not achieve the required performance.

The next step was made on two levels: the algorithm was written in Assembler and a solution to the overhead issue was developed. The solution is based on the elimination of the outer cycle. The inner cycle is followed by an unconditional branch, which takes less time than a conditional one. In addition to that, the DSP instruction set was exploited to optimize the sample's selection, i.e. the pointer increment and reset. Two registers in the CPU are used to store the bottom and top LUT addresses. At each cycle execution, the register containing the bottom address is incremented and compared to the top address. If they match, the cycle is exited. In this way, the pointer acts as a counter, as well, and the number of operations is reduced. The following conditional branch makes the execution jump back to the cycle's initialization, which consists only in resetting the pointer to the bottom address of the LUT (Figure 4.3 and Figure 4.4).

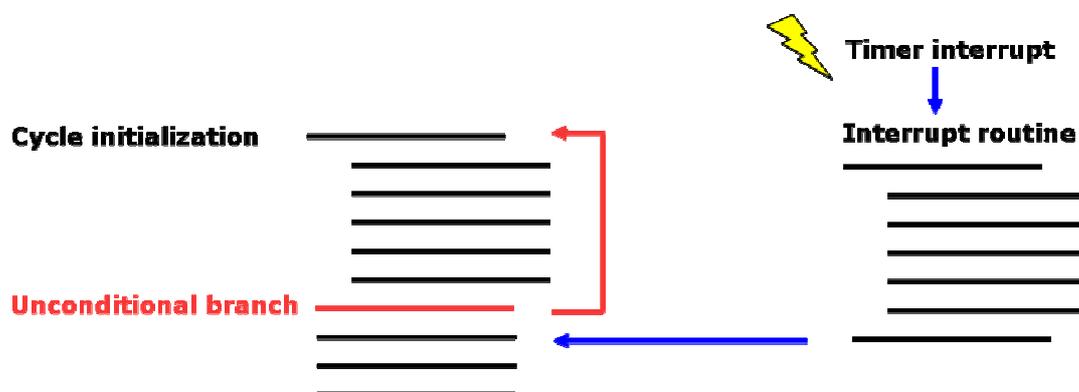


Figure 4.3 – The optimized actuation mechanism.

The proposed solution poses the problem of the actuation's interruption, which can be performed by a timeout mechanism. Before entering the infinite cycle

which parses the LUT, a timer is started and its interrupt is enabled. The timer period is set to the value sent by the microcontroller during the handshake phase. When the interrupt is triggered, the service routine is called and the actuation is stopped. To exit the cycle, the interrupt routine modifies the stack so that the return address points to the instruction following the unconditional branch instruction. In this way it is also possible to accurately control the duration of the actuation.

The same solution was adopted in the acquisition cycle. In this case, the samples are produced by the ADC and have to be stored in the SRAM memory. Depending on the number of samples, they can be stored in the internal SRAM or in the external memory chip. From the point of view of the algorithm, it makes no difference since the external memory chip is mapped in the same memory space as the SRAM and it can be accessed in a fully transparent way.

```

asm("_actuationBegin: ");
asm(" MOVL XAR1,#_actuationWave");
asm("_actuationCycle: ");
asm(" MOVU ACC,*XAR1++");
asm(" CMPR 1");
asm(" MOV *XAR2,ACC");
asm(" BF _actuationCycle, TC");
asm(" BF _actuationBegin, UNC");

asm("_actuationEnd:");
EvaRegs.T1CON.bit.TENABLE=0;

interrupt void T1_P_ISR(void){
    asm("\t PUSH XAR4");
    asm("\t PUSH XAR3");
    PieCtrlRegs.PIEACK.all = PIEACK_GROUP2;
    asm("\t SUBB SP,#8");
    asm("\t MOVL XAR4,#_actuationEnd");
    asm("\t MOVL *SP,XAR4");
    asm("\t ADDB SP,#8");
    asm("\t POP XAR3");
    asm("\t POP XAR4");
    return;
}

```

Figure 4.4 – Source code of the optimized actuation cycle and of the timer’s interrupt routine.

The software development did not include a real SHM algorithm. However, the code for a FFT on the acquired data was written and successfully executed.

The DSP code must be stored in the external EEPROM, from which it is fetched. The code has to be organized in a specific way defined in the DSP datasheets. A batch program shipped with Code Composer Studio (hex2000.exe) allows the translation of the executable file to the format required by the DSP’s bootstrap loader. After running the batch program, a common EEPROM writer was used to download the data to the EEPROM.

4.3 PC-side GUI

A PC-side graphical user interface was developed to make interaction with node possible via a computer (Figure 4.5). The GUI was developed by Microsoft Visual Studio IDE and written in Visual Basic.NET. The choice was determined by

the fact Visual Studio allows users to develop complex user interfaces with little effort and Visual Basic.NET offers built-in objects for the management of communication via serial ports. This was judged of help, since the communication with the node relies on the Xbee radio which can be interfaced to a PC via a serial link.

The GUI features both debug and real use scenario functions. It allows the user to submit standard requests to the node, such as a query about the results of the analysis, but it also allows to get detailed information about the node's state, thanks to commands which will be removed in the final version of the whole system.

As for data management, the GUI includes a data visualization panel with FFT calculation capabilities. The panel can visualize multiple waveforms. Finally, the user can also chose to export the data series to a Matlab-compatible format for further processing by the well known software.

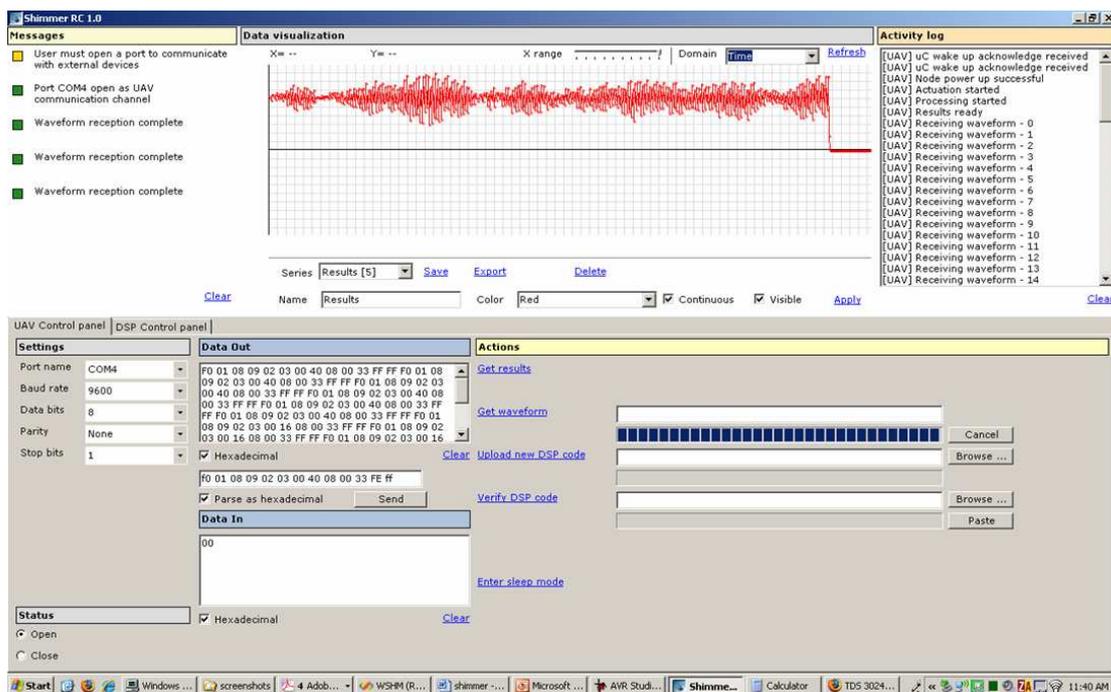


Figure 4.5 – Screenshot of the PC-side GUI.

5 System prototyping and validation

A prototype of the wireless sensor node has been built and validated. The prototyping and all tests took place in a laboratory at the California Institute for Telecommunications and Information Technology (CalIT²), located in the campus of the University of California, San Diego.

5.1 Prototyping

The prototyping of the wireless sensor node was a demanding activity. First of all, the single functional units of the node were built and tested separately. This has allowed to improve the understanding of the components' functioning. During these tests, the size of the resistors and capacitors required as external circuitry by some components was determined. The interaction between the single parts was verified and some preliminary tests were run to perform an actuation-acquisition cycle, using the DSP's development board.

Once these tests were completed, the mechanical structure of the node was designed. Basically the node is comprised of three main parts: the energy harvesting and power distribution circuit, the analog interface to the PZTs and the digital core. One of the key objectives of the whole project was to obtain a device as small as possible. The envisioned solution was to distribute all the components on three boards, corresponding to the parts just listed (Figure 5.1). The boards would be stacked vertically and joined by suitable connectors. This solution seemed to be the most convenient both practically and conceptually since it facilitates the design and build of the individual boards. Furthermore, future changes in single parts of the node would not require the whole device to be rebuilt. The resulting node's structure is shown in Figure. The distribution of the components on three boards was also driven by consideration concerning electro-magnetic interference. In particular, all switching converters were placed on the bottom side of the energy harvesting board,

so that the board itself acts as an EMI shield and prevents I/O signal conditioning functioning from being altered.

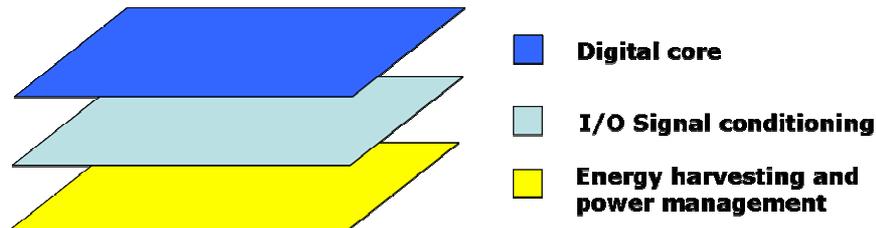


Figure 5.1 – The node's structure.

Given these guidelines, the circuit schematics and layout were prepared using the Capture and Layout tools of the OrCAD suite. The layout preparation was quite complex because of the high number of components to be placed and signals to be routed. It is at the layout definition level that space occupation was optimized. The final result, consisting of three 9x9cm boards, was achieved by performing several iterations of the components' placement and signals' routing. The task was made particularly challenging by the fact that the connectors on two adjacent boards had to be aligned. Clearly, this introduced some additional constraints and reduced the degrees of freedom. After the layout preparation, all missing components (e.g. connectors) were ordered and the prototype was built. All components and tools were bought thanks to the funds made available by the Los Alamos National Laboratory.

The major contribution to a fast prototyping was brought by a PCB milling machine (Protomat S62), manufactured by LPKF Inc (Figure 5.2). The machine can produce double layer printed circuits as big as 229x305mm in a very fast way, by milling out the signal paths on a copper-plated board. The result is similar to that obtained by the etching method, with the advantage of avoiding any acids or dangerous liquids.

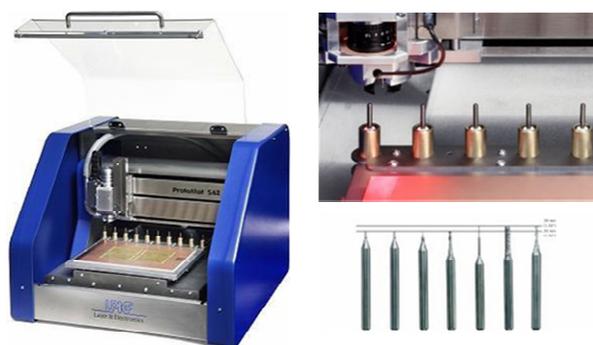


Figure 5.2 – Protomat S62 milling machine.

The signal paths are milled by means of ultra-thin drill points. The typical diameter used in the prototype's realization is 0.2mm, but points as thin as 0.1mm are available. Furthermore, the machine's resolution is as low as 0.25 μ m. The machine can be utilized to drill vias, too. The Protomat S62 comes with two control softwares: the first one (CircuitCAM) imports the output of the layout editor and translates it to a format interpretable by the milling machine; the second one (BoardMaster) features a control panel, by which all machine's operations can be supervised.

The machine has a steep learning curve and its use becomes easy after a little practice. However, all parameters must be accurately set, otherwise the quality of the milled PCB deteriorates. In particular, the signal path's contours get irregular and small copper particles remain partially attached to them, in a such a way they can detach later causing shorts in the circuit (Figure 5.3 and Figure 5.4). Therefore, a careful inspection of the PCB turned out to be mandatory.

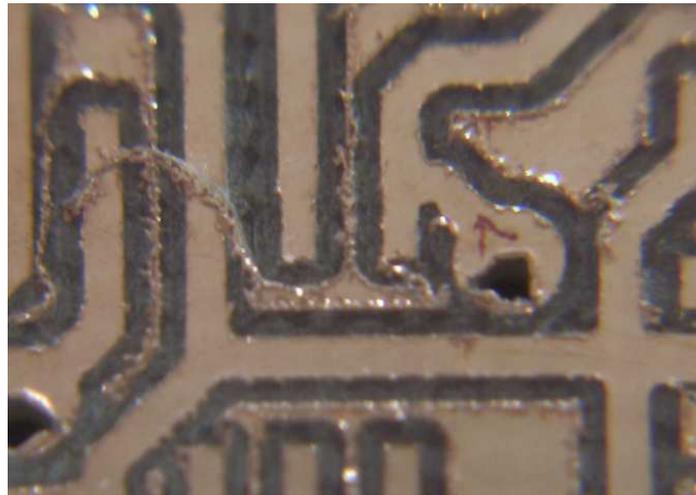


Figure 5.3 – Copper particles produced by the milling machine.

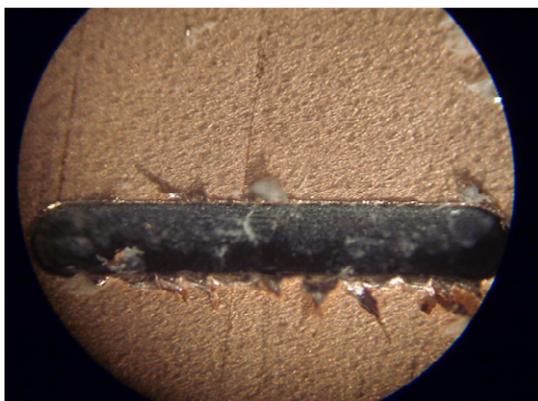


Figure 5.4 – Imperfections caused by milling tools' deterioration.

After the necessary training with the milling machine, the prototype was built. The three boards are shown in the pictures in Appendix B. The final height of the node is about 5cm (Figure 5.5).

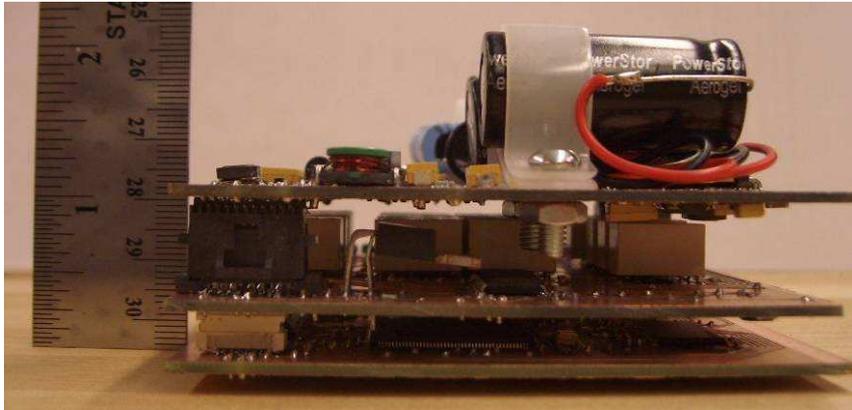


Figure 5.5 – The prototype in its final version.

The components were mounted on the boards by hand, using a bench soldering iron with an ultra-thin tip. The most challenging component was the DSP, because of the very small pin-to-pin distance, as low as 0.5mm (Figure 5.6).

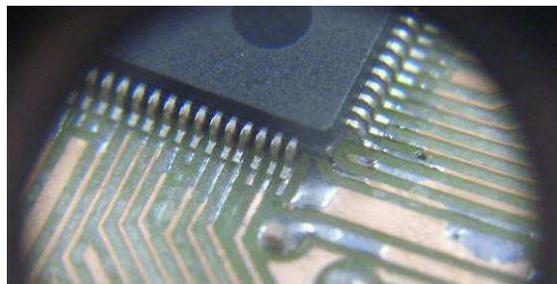


Figure 5.6 – DSP pins after soldering.

5.1.1 Prototype's cost

The cost of the prototype is about \$320. This amount is obtained by summing up all component's price, the solar panels' price and the expenses related to some materials' purchase. It is difficult to evaluate such a cost. Most of the wireless sensor nodes described in Chapter 2 are much cheaper, but their capabilities are also much simpler than Shimmer's. A benchmark which can be taken into consideration is represented by the instruments commonly used in laboratories to perform SHM

analysis. In this case the comparison shows the great contribution to the feasibility of a large scale SHM system brought by Shimmer. A National Instruments PXI Real-Time module was used in the laboratory by the other members of the research team, busy with the development of SHM algorithms. Such an instrument cost \$30000. Moreover, Shimmer's cost could be reduced by purchasing components in quantity.

5.2 System validation and characterization

Several tests aiming to validate and characterize the wireless node were run. These tests focused on the various parts and activities of the node: the energy management circuit, the A/D and the D/A signal conditioning circuits, the radio-triggering circuit and the digital core's operation.

5.2.1 Energy management circuit

The first test run on the energy management circuit aimed to characterize the supercapacitors' charging circuit. Figure 5.7 shows the charging curve of a 100F 2.5V supercapacitor by the solar panels connected to the node. The supercapacitors get fully charged in less than 1h30. The result is absolutely satisfactory, since it shows how the solar panels can harvest enough energy to fully charge the node during a typical day. This curve could be improved by the adoption of MPP solution.

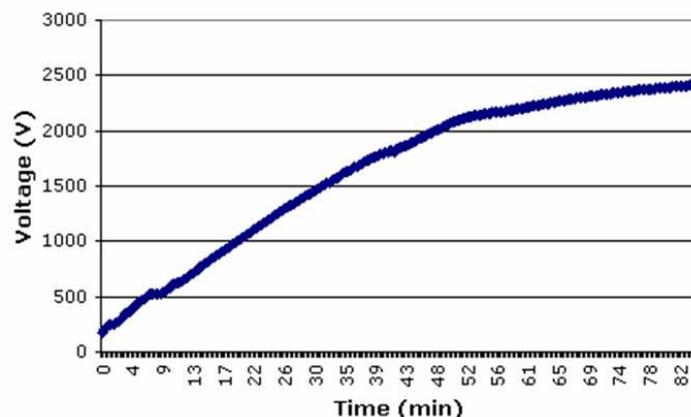


Figure 5.7 – Charging of a 100F supercapacitor.

The following pictures (Figure 5.8, Figure 5.9, Figure 5.10 and Figure 5.11) show the power-up behavior of the voltage supplies included in the node. The yellow

waveform in all pictures corresponds to the turn-on signal controlling the CMOS switches inserted between the supercapacitors and the voltage converters. The blue and pink waveform are the signals of interest.

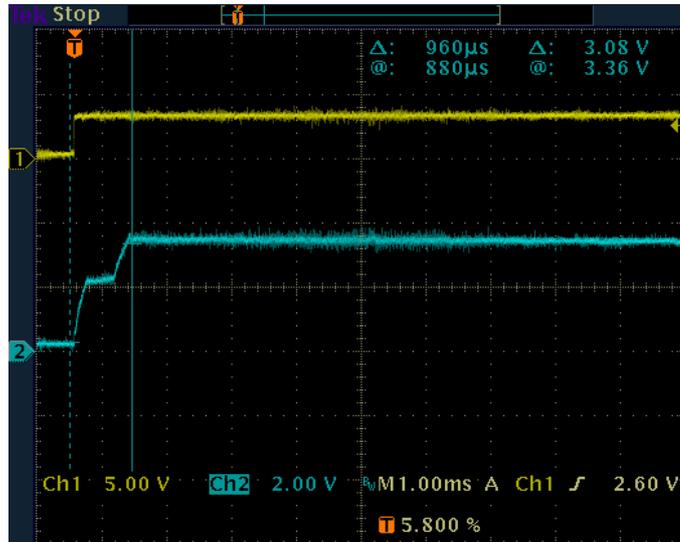


Figure 5.8 – 3.3V (blue) supply's output after turn-on.

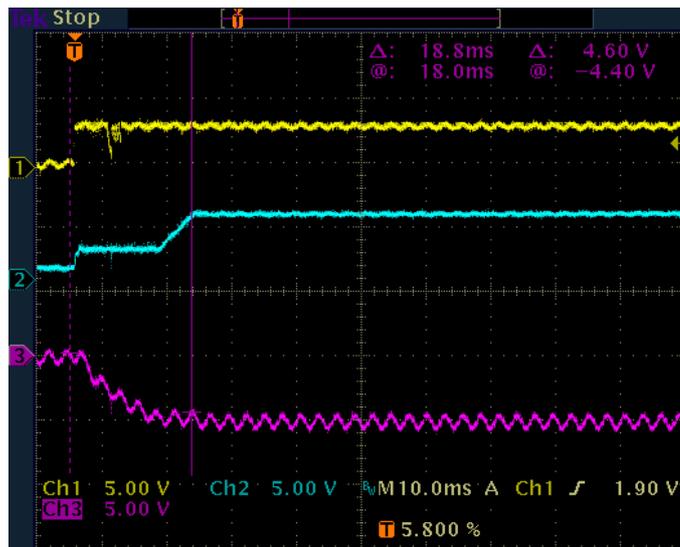


Figure 5.9 – +5V (blue) and -5V (pink) supplies' output after turn-on.

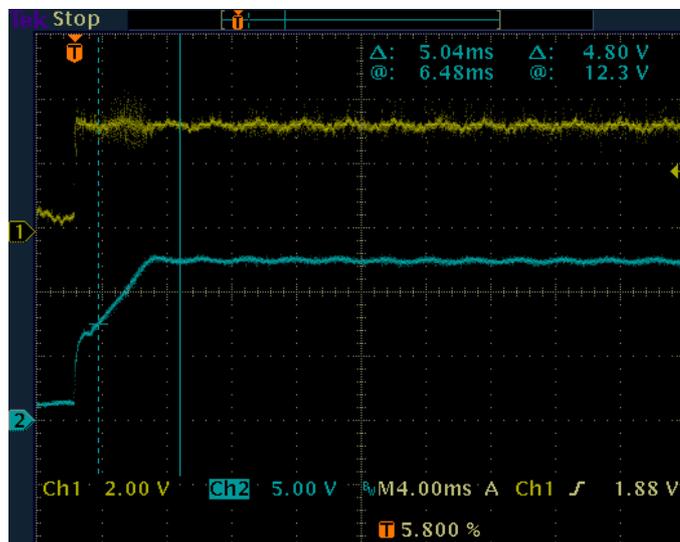


Figure 5.10 – +10V (blue) supply's output after turn-on.

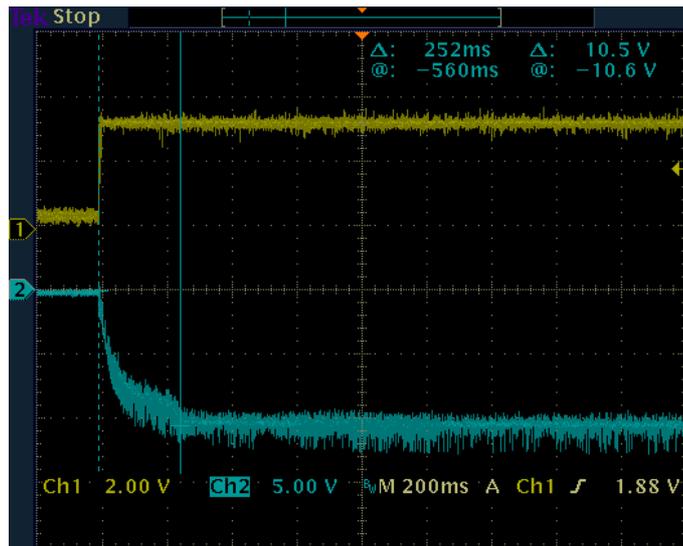


Figure 5.11 – -10V (blue) supply's output after turn-on.

As the pictures show, the slowest supply to settle its output is the voltage inverter producing the -10V. This influences the delay between the node's wake up and the beginning of the acquisition.

5.2.2 Signal conditioning circuits

The first part of the signal conditioning circuits to be tested was the anti-aliasing filter. The filter was connected to a National Instrument PXI Real-Time system, featuring a signal generator and an high-speed ADC. The filter's response to a variable-frequency input sine wave was acquired in order to evaluate the filter's characteristic. The filter showed a cutoff frequency around 1MHz, as desired (Figure 5.12).

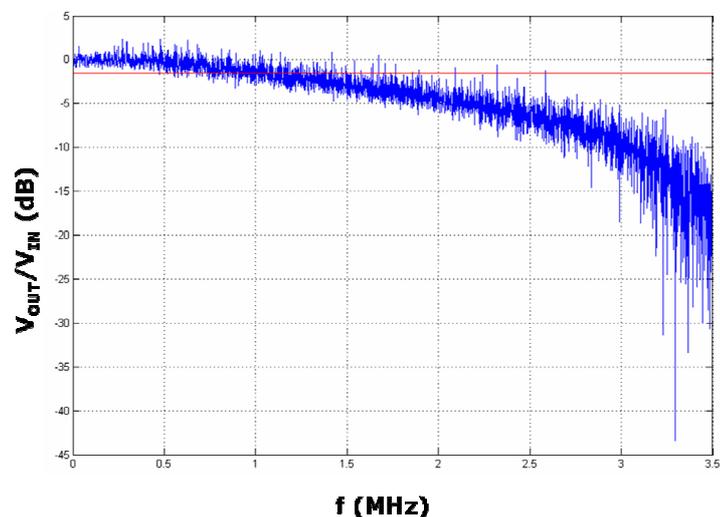


Figure 5.12 – Characteristic of the anti-aliasing filter.

By contrast, the characteristic is not as steep as expected above the cutoff frequency. This can be caused by the difference between the theoretical optimal values of the discrete components and the values actually available.

Next tests were run using the setup shown in Figure 5.13. Two PZTs were super-glued onto a steel beam. A hole was drilled in the beam and a metal pin was inserted in it. This setup was meant to simulate a real application scenario. In fact, the hole and the pin were used to simulate the damaged and damage-free conditions of the beam. By instance, the damaged condition was obtained by removing the pin, since this resulted in a discontinuity in the beam. Vice-versa, the beam's integrity was restored by inserting the pin.

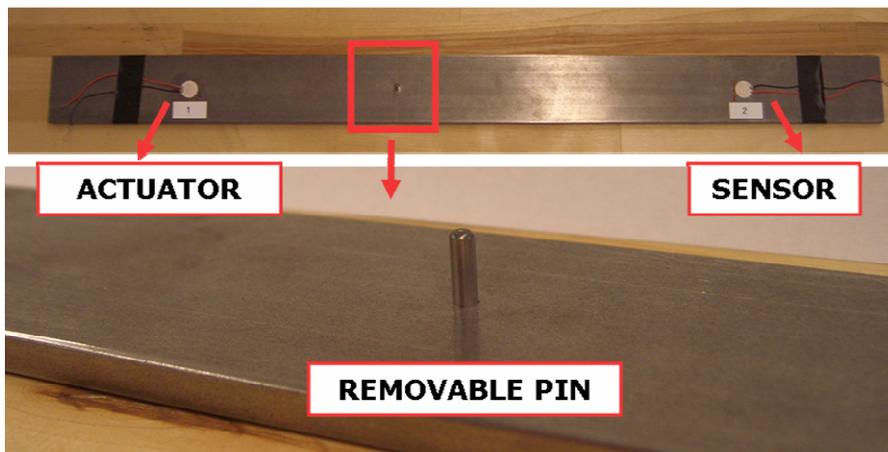


Figure 5.13 – Setup for damage detection tests.

The A/D and D/A signal conditioning circuits were fully tested by actuating the first PZT and connecting the circuit's input to the second PZT's output. Figure 5.20 shows the actuation wave (blue colored), the sensing PZT's output (pink colored) and the circuit's output (green colored). The tests allowed to conclude that both circuits were correctly designed.

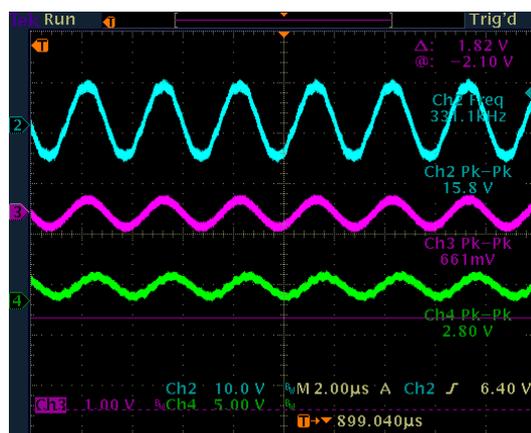


Figure 5.14 – Result of tests on A/D conditioning circuit.

Some measurements were done to characterize both input and output amplification circuits in terms of SFDR. The results are shown in Figure 5.15 and Figure 5.16. The SFDR characterization of the input amplification circuit showed a good performance: the SFDR is around 20dBc over the whole frequency range of interest. The output amplification circuit showed a minimum SFDR around 18dBc at 700KHz, with peaks of 30dBc at 400KHz and 1MHz. Future development should aim to improve the SFDR in the 700KHz band.

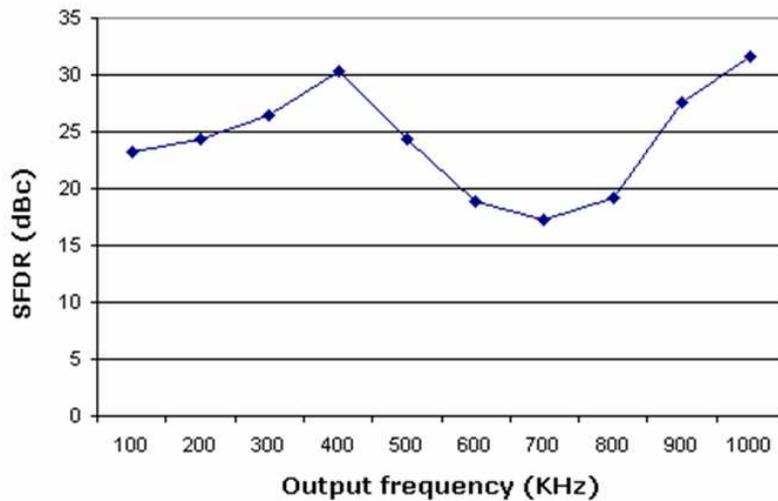


Figure 5.15 – SFDR of output amplification circuit.

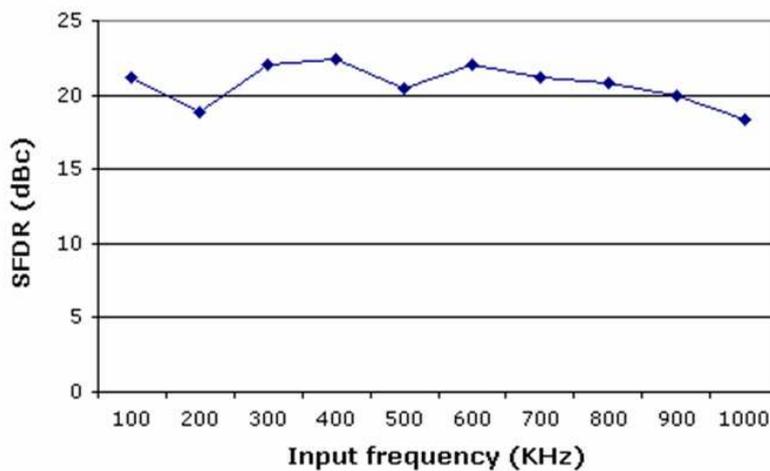


Figure 5.16 – SFDR of input amplification and filtering circuit.

5.2.3 Radio triggering circuit

Tests were performed to validate the radio-triggering circuit. Tests were run

transmitting a radio signal in the 2.4GHz band and receiving it by means of the antennas showed in Figure 5.17. The signal output by the receiving antenna was fed directly into the radio-triggering circuit. The triggering signal was produced by a microwaves amplifier connected to the transmitting antenna. The distance between the node and the origin of the signal was about 6m.

The radio-triggering circuit was demonstrated, as showed in Figure 5.18. The AC-to-DC voltage multiplier outputs the 0.2 level necessary to turn on the cascaded low-threshold MOSFET. The triggering signal's power was 100mW, below the maximum value defined by the FCC's regulations.

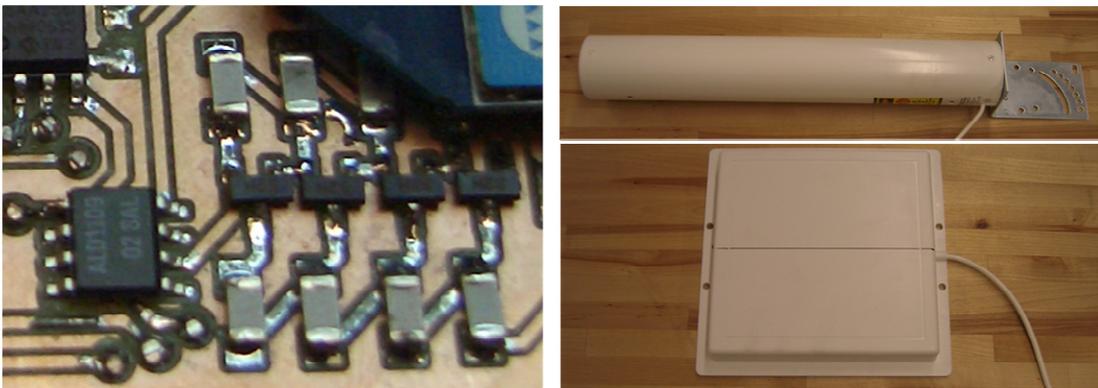


Figure 5.17 – The radio-triggering circuit and the antennas used during tests.

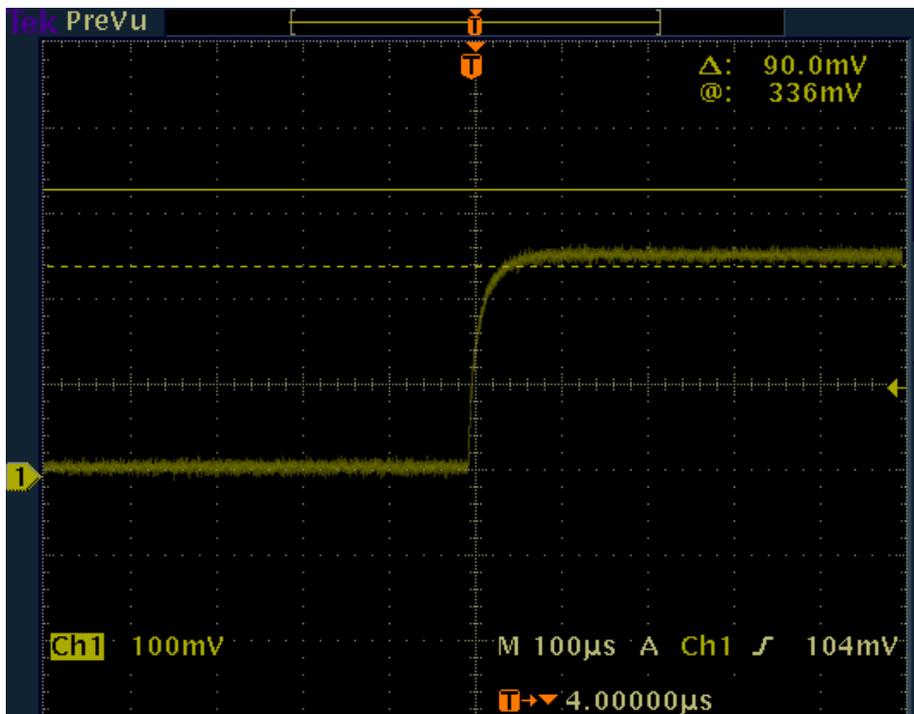


Figure 5.18 – The output of the radio triggering circuit.

5.2.4 Digital core

As a final test, a full operational test was run. The node was woken up by radio-triggering, then the required settings were sent to the node by the PC-side GUI and the actuation-acquisition loop was performed. Finally, the node was queried to get the acquired samples. Figure 5.20, Figure 5.19, Figure 5.21 show the waveforms involved in the actuation-acquisition loop in three different cases. The yellow waveform showed in all pictures is a signal used to trigger the oscilloscope, which goes high at the beginning of the actuation and goes low at the end of acquisition. The blue waveform is the actuation burst, the pink waveform is the sensing PZT's output and the green waveform is the A/D conditioning circuit's output, i.e. the signal sampled by the DSP's ADC.

The pictures show how the beam's response changes if the actuation frequency is modified. Furthermore, it is possible to see how the response changes if the pin is removed while maintaining the same actuation frequency. The pin's removal modifies the reflection paths of the signal within the steel beam. Thus, the sensing PZT is excited by the result of the superposition of different waves, as it is clearly shown in Figure 5.22.

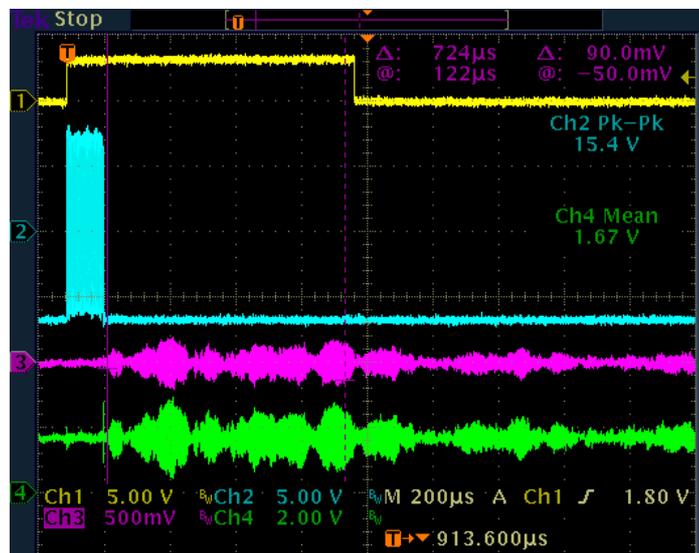


Figure 5.19 – Results of actuation-acquisition loop. Actuation at 300KHz.

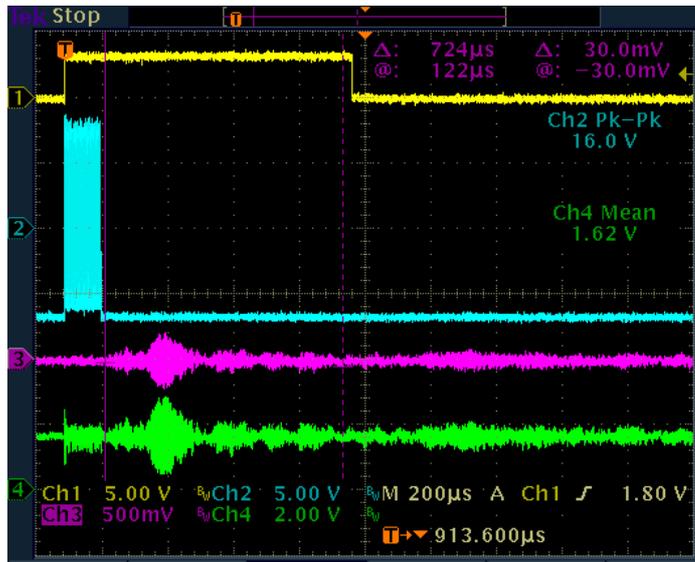


Figure 5.20 – Results of an actuation-acquisition loop. Actuation at 600KHz, damage case.

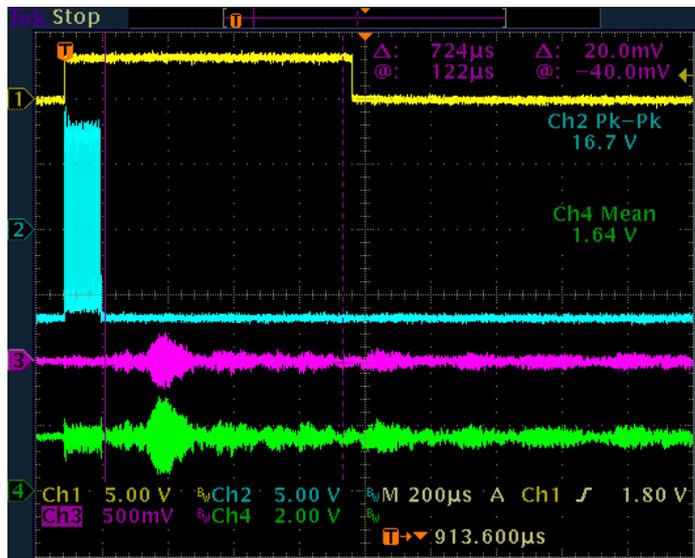


Figure 5.21 – Results of actuation-acquisition loop. Actuation at 600KHz, damage-free case.

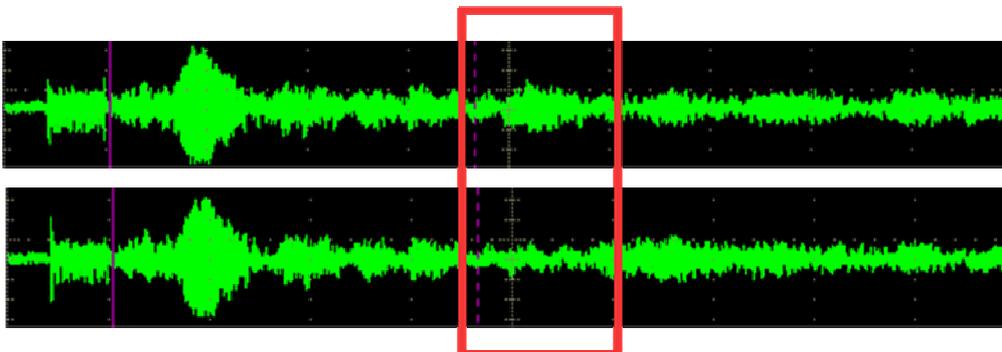


Figure 5.22 – Difference between beam’s responses in damage and damage-free cases.

The quality of the acquisition performed by the DSP was assessed by exporting the acquired data series to Matlab (Figure 5.23). As for computation, some

tests were run to collect information about the time required to run the analysis algorithm on the DSP. A FFT on 8192 samples was calculated in 510ms.

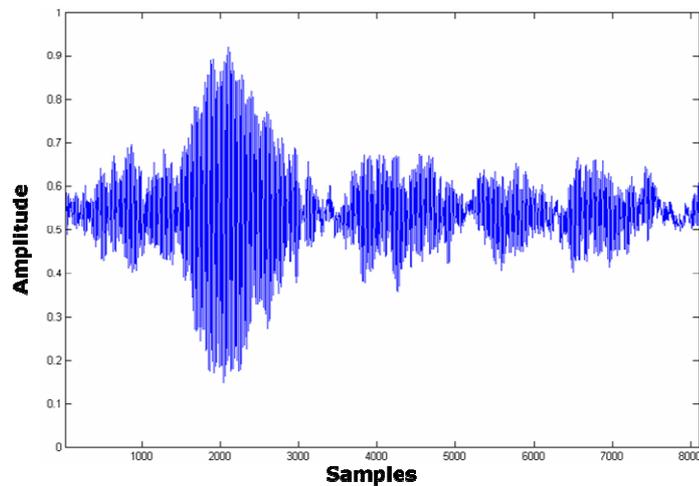


Figure 5.23 – A data series plotted by Matlab.

5.2.5 Power consumption

The characterization of the energy management circuit also included power consumption measurements. First of all, stand-by power consumption was measured. A $33\mu\text{A}$ current flowing through the 2.5V supercapacitors and a $46\mu\text{A}$ current flowing through the 5V supercapacitors were measured. This results in a sleep mode power consumption equal to $\sim 0.3\text{mW}$. The active mode power consumption was characterized in two steps: firstly, the power consumption associated with the actuation circuit, draining current from the 5V supercapacitors, was assessed, then measurements focused on the power consumption of the rest of the node.

As Figure 5.24 shows, the power consumption of the actuation circuit strongly depends on the actuation wave's frequency. At 100KHz and 500KHz, a 1F supercapacitor fully charged at 5V allows the actuation to last for about 5s. After this time interval, the voltage on the supercapacitors drops below the minimum input voltage of the positive and negative supplies powering the actuation circuit. This corresponds to a power consumption of around 2W. By contrast, at 1MHz, the actuation lasts for about 2.5s. Thus the power consumption is about the double. Such power consumptions are very high, but the typical duration of the actuation has to be taken into consideration. The usual number of sine waves periods required to sense a path between two PZTs is lower than 100. If the output frequency is 100KHz, the actuation lasts 1ms at most. Therefore, the energy stored in the supercapacitors is

sufficient to sense several paths.

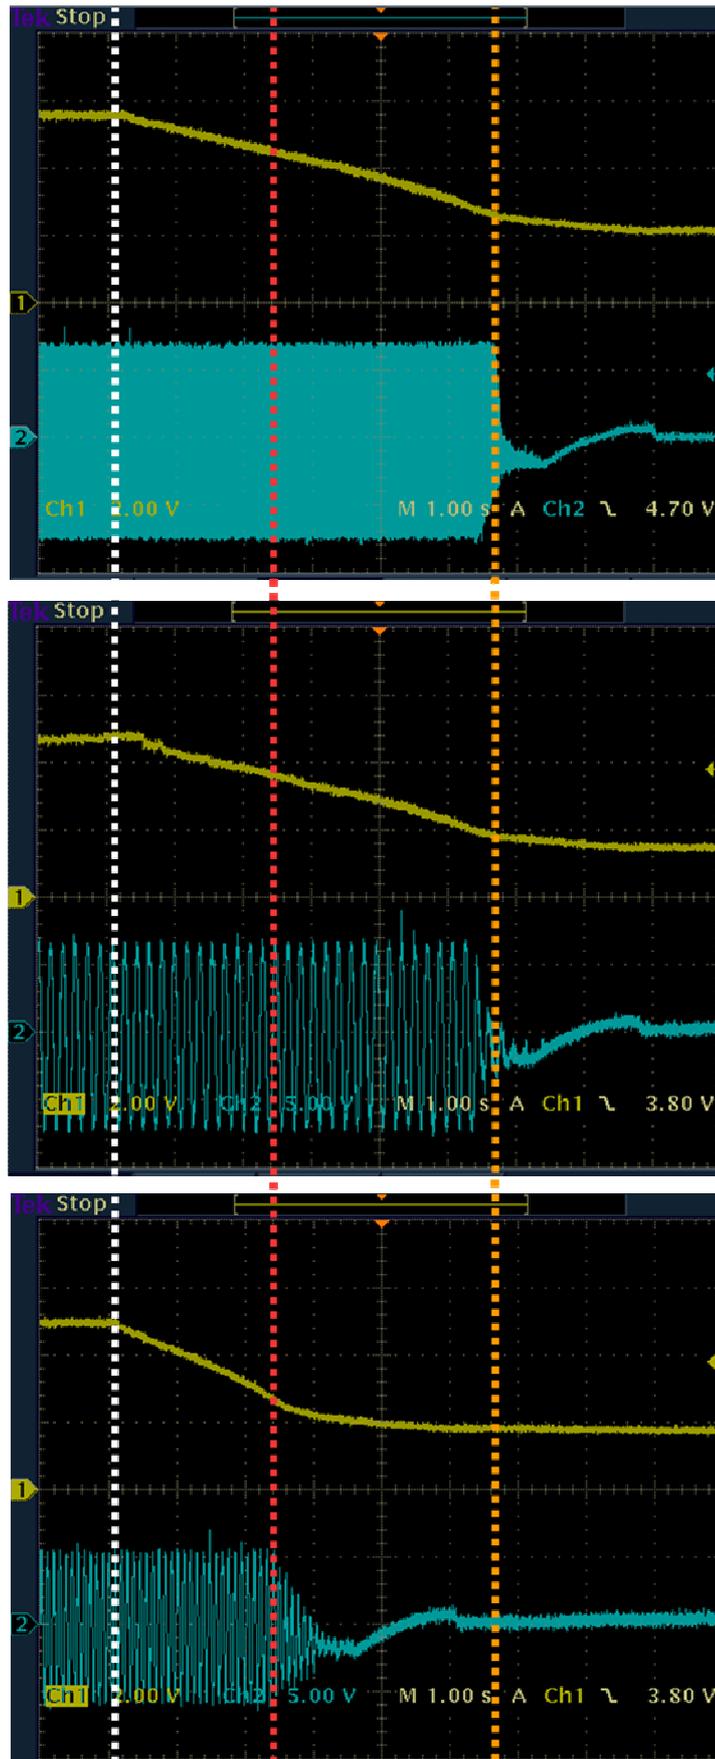


Figure 5.24 – 1F 5V supercapacitor’s discharge during actuation at 100KHz, 500KHz and 1MHz.

The power consumption during the acquisition and processing phases is around 750mW. The node can run up to 4'50'' with a 100F supercapacitor and 12'3'' with a 250F supercapacitor. The theoretical time –calculated as the ratio between the energy stored in the supercapacitors and the power consumption- is higher, but the voltage converters efficiency has to be taken into account. However, basing on the test showing that a FFT on 8192 samples can be calculated in less than 1s, it is possible to conclude that the node can clearly be used to run a typical SHM algorithm.

6 Conclusions

This thesis has outlined the development of a Lamb wave-based wireless sensor node for structural health monitoring applications.

First of all, a review of the structural health monitoring motivations, issues and techniques was given. Then the state-of-art of wireless sensor networks was analyzed from the hardware standpoint and the main energy harvesting techniques were described.

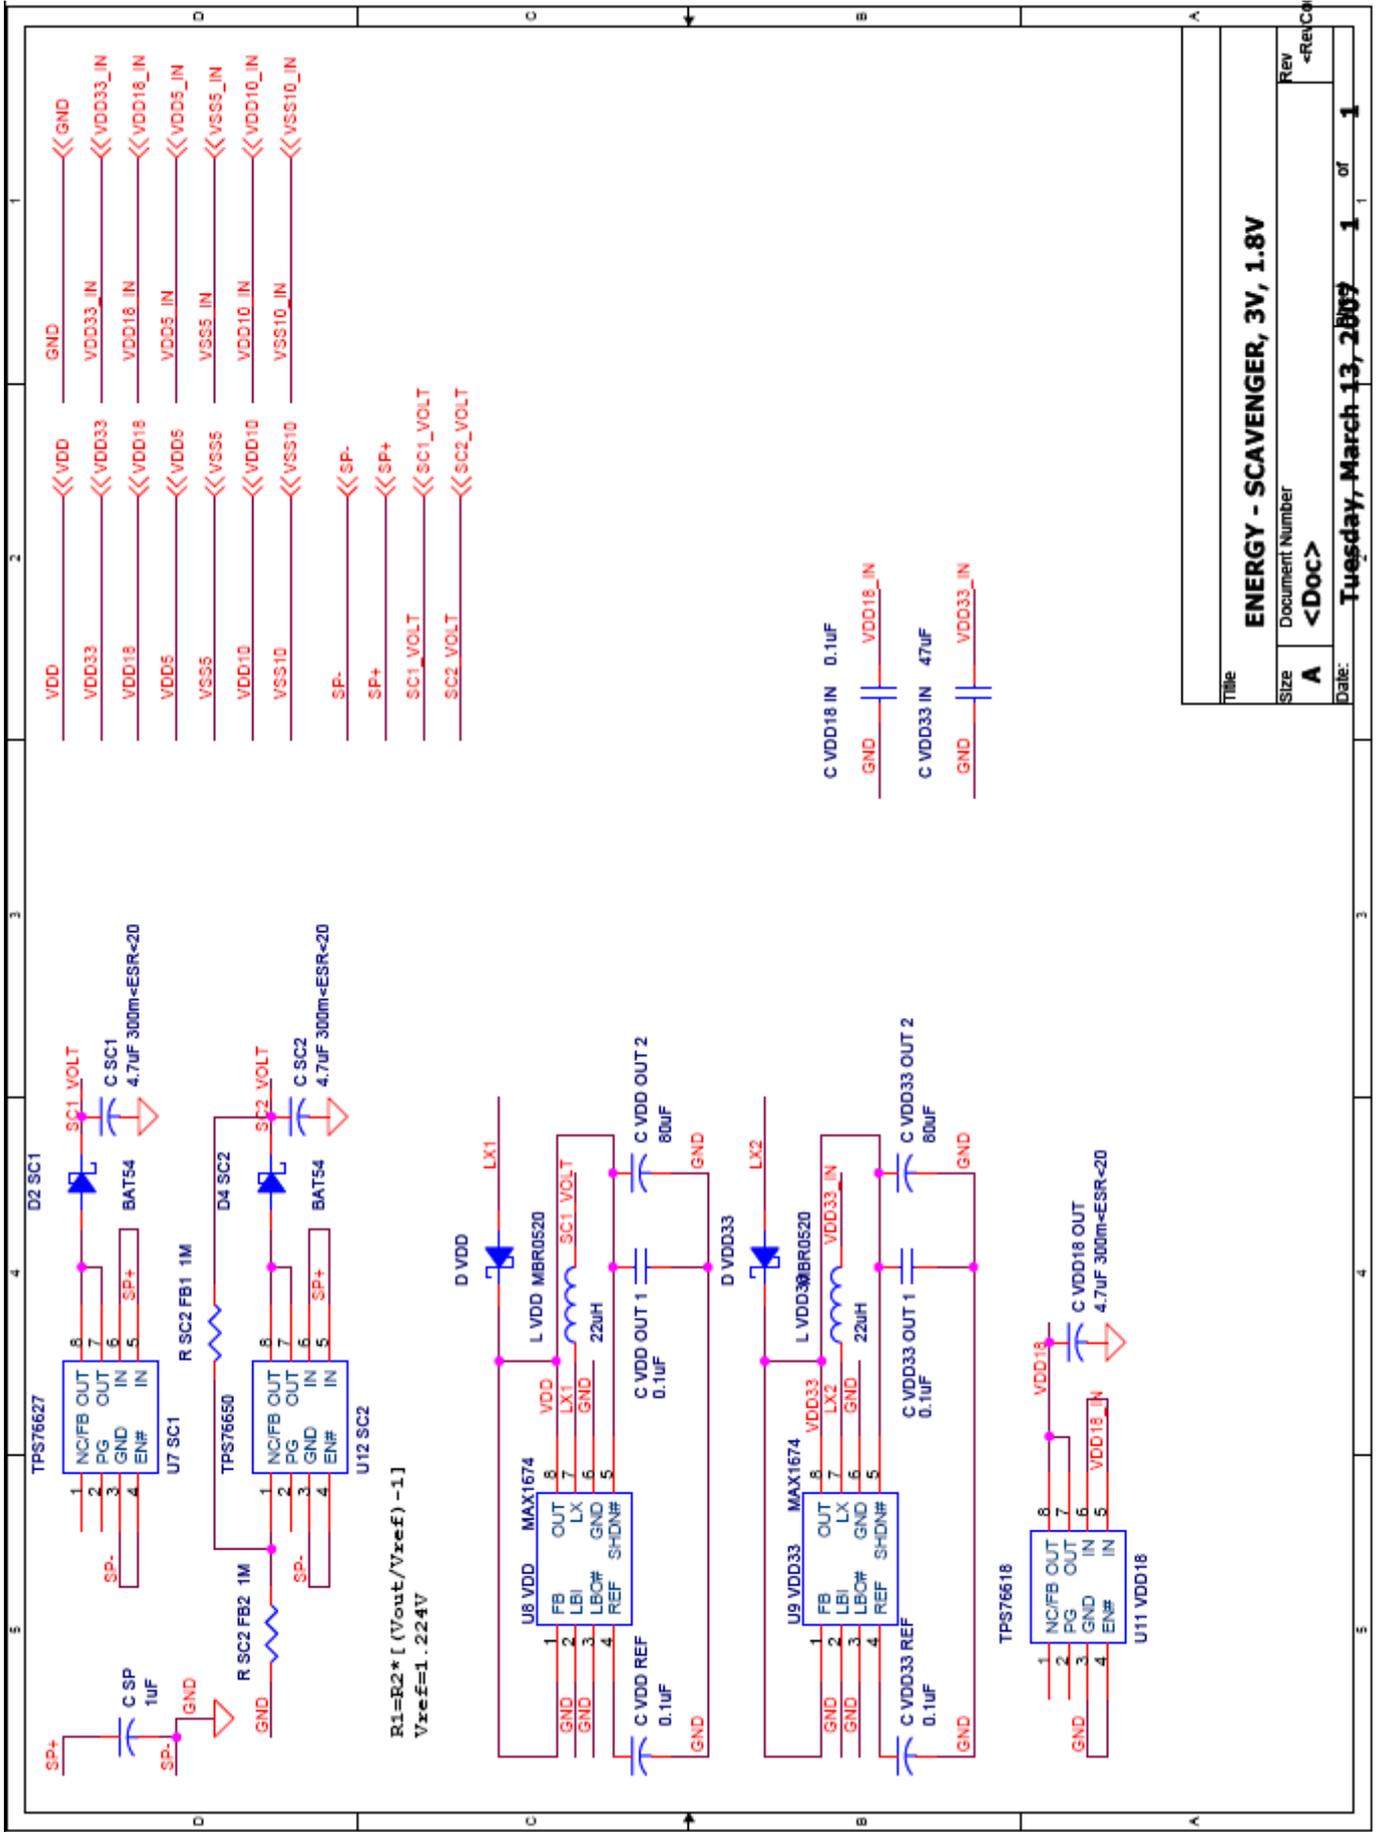
The rest of the thesis is concerned with the hardware and software design of the wireless sensor node. The node is designed to be interfaced to a grid of 16 PZTs, which can be used both as actuators and sensors. The node can sense multiple actuator-sensor paths in order to locate damage with the required accuracy. The on-board DSP also allows the node to perform heavy computation on the acquired data, enabling the execution of most common SHM algorithms. The energy harvesting circuit which the node is provided with makes the device self-sustaining from the energy point of view. The energy-aware solutions implemented in the node optimize the use of the energy and maximize the time available for computation.

The prototyping and testing results were positive and showed how the design of the sensor node represents a major contribution to the realization of an effective, low-cost and large scale wireless sensor network for structural health monitoring applications.

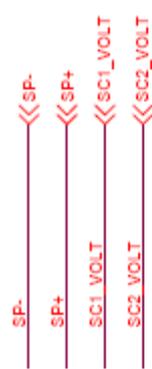
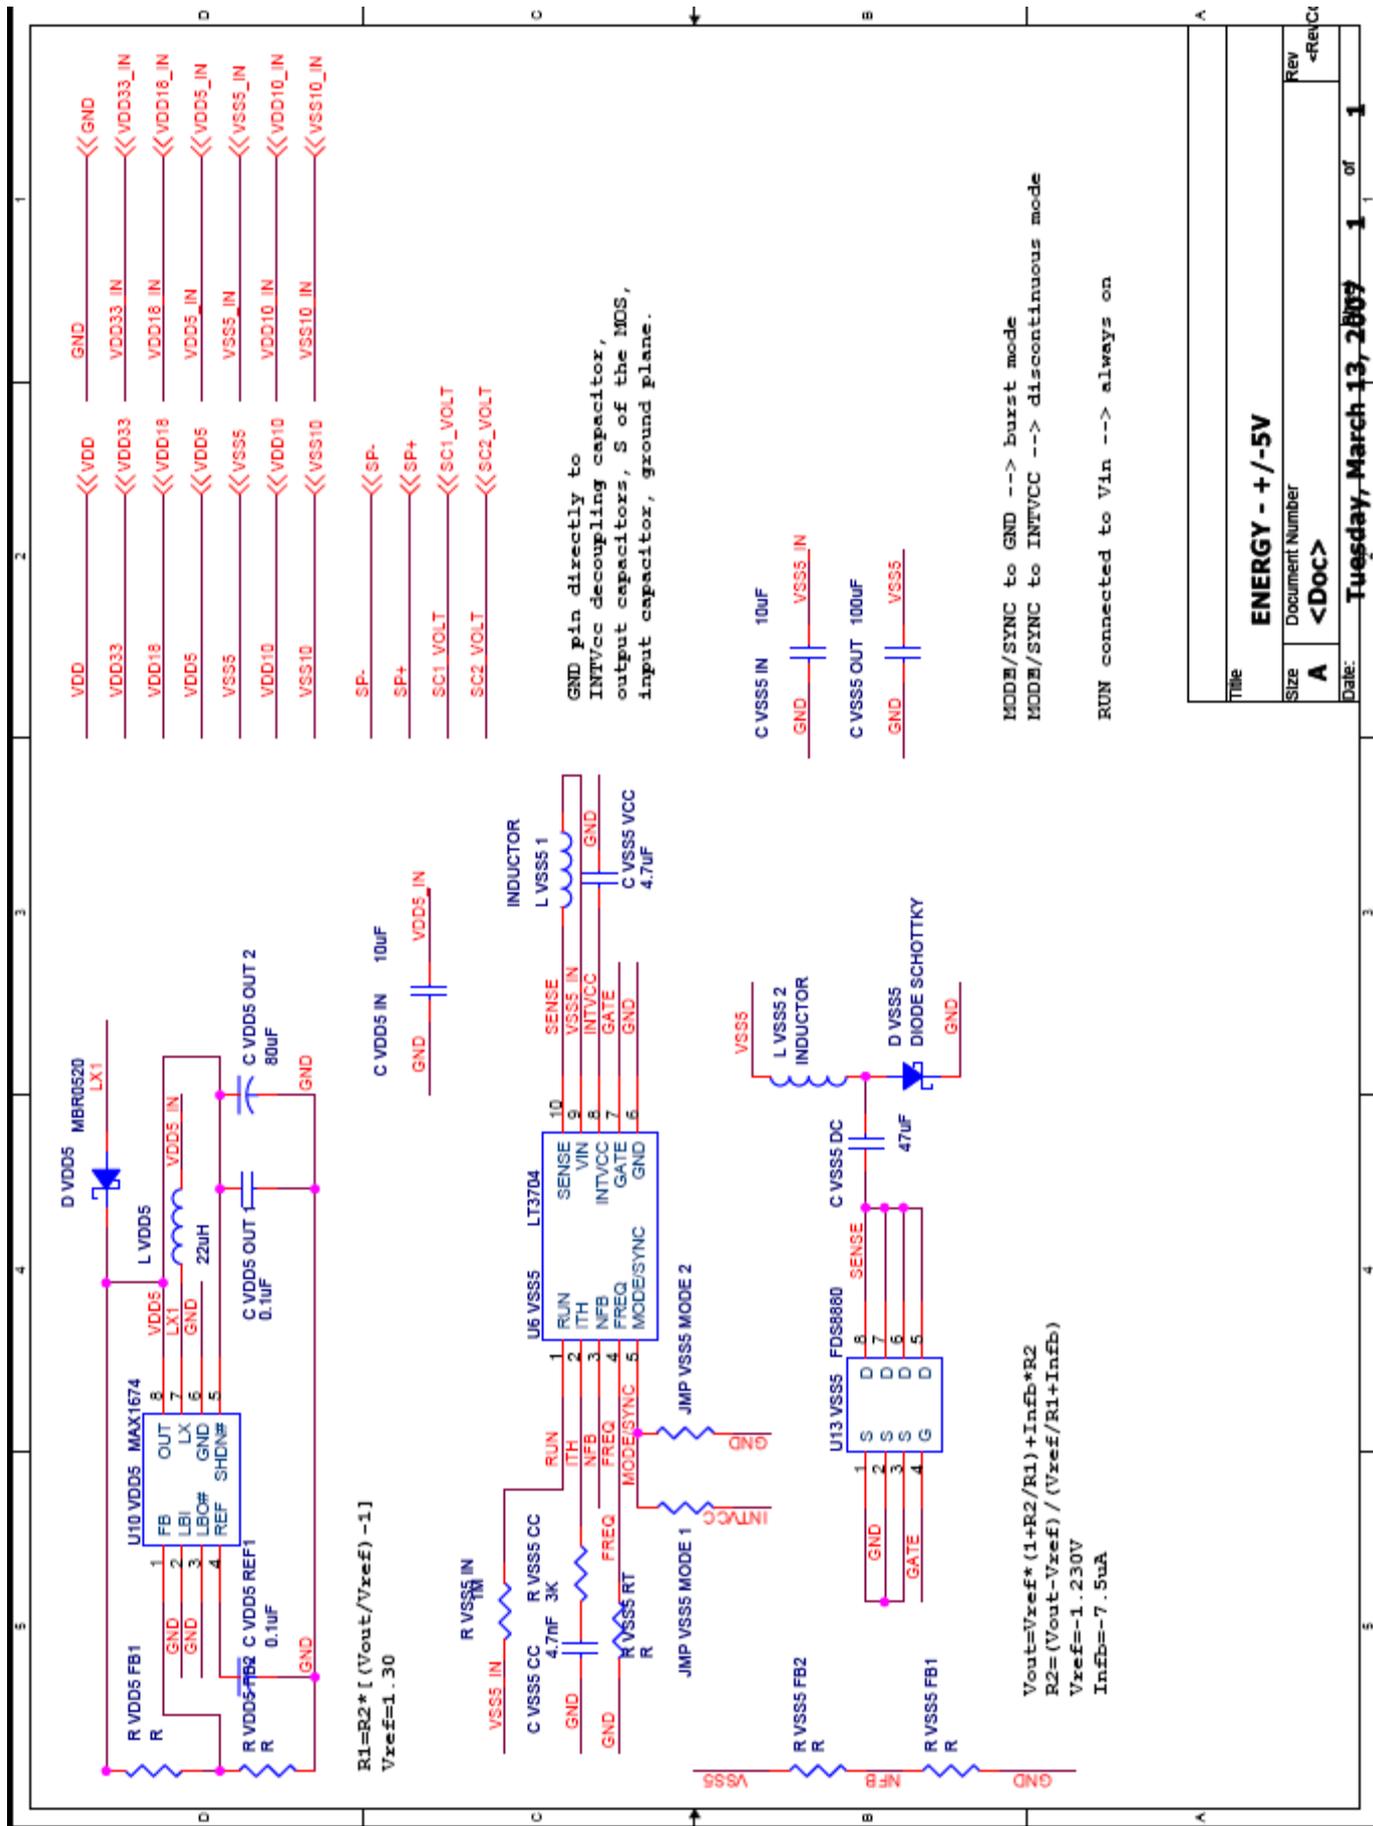
Appendix A – Schematics

Appendix A contains all schematics of the wireless sensor node. They are organized as follows:

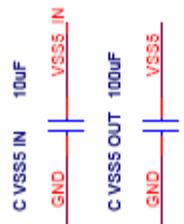
- Energy management board – harvesting circuit, 1.8V, 3V supplies;
- Energy management board – harvesting circuit, +5V, -5V supplies;
- Energy management board – harvesting circuit, +10V, -10V supplies;
- Energy management board – switches;
- Energy management board – harvesting circuit, connectors;
- I/O conditioning board – A/D circuit;
- I/O conditioning board – D/A circuit;
- I/O conditioning board – connectors;
- Digital core – DSP;
- Digital core – microcontroller;
- Digital core – radio;
- Digital core – connectors.



Title		ENERGY - SCAVENGER, 3V, 1.8V	
Size	Document Number	Rev	<RevCo
A	<Doc>		
Date:	Tuesday, March 13, 2007	1	of 1



GND pin directly to INTVCC decoupling capacitor, output capacitors, S of the MOS, input capacitor, ground plane.



MODE/SYNC to GND --> burst mode
 MODE/SYNC to INTVCC --> discontinuous mode
 RUN connected to Vin --> always on

$$R1 = R2 * [(V_{out} / V_{ref}) - 1]$$

$$V_{ref} = -1.30$$

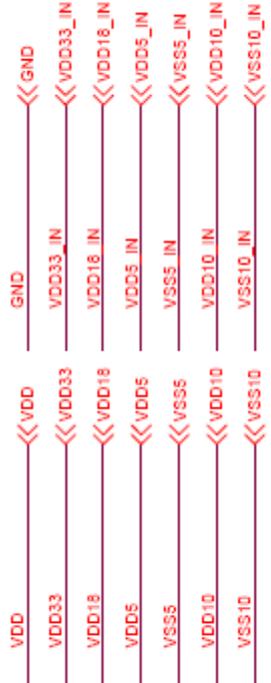
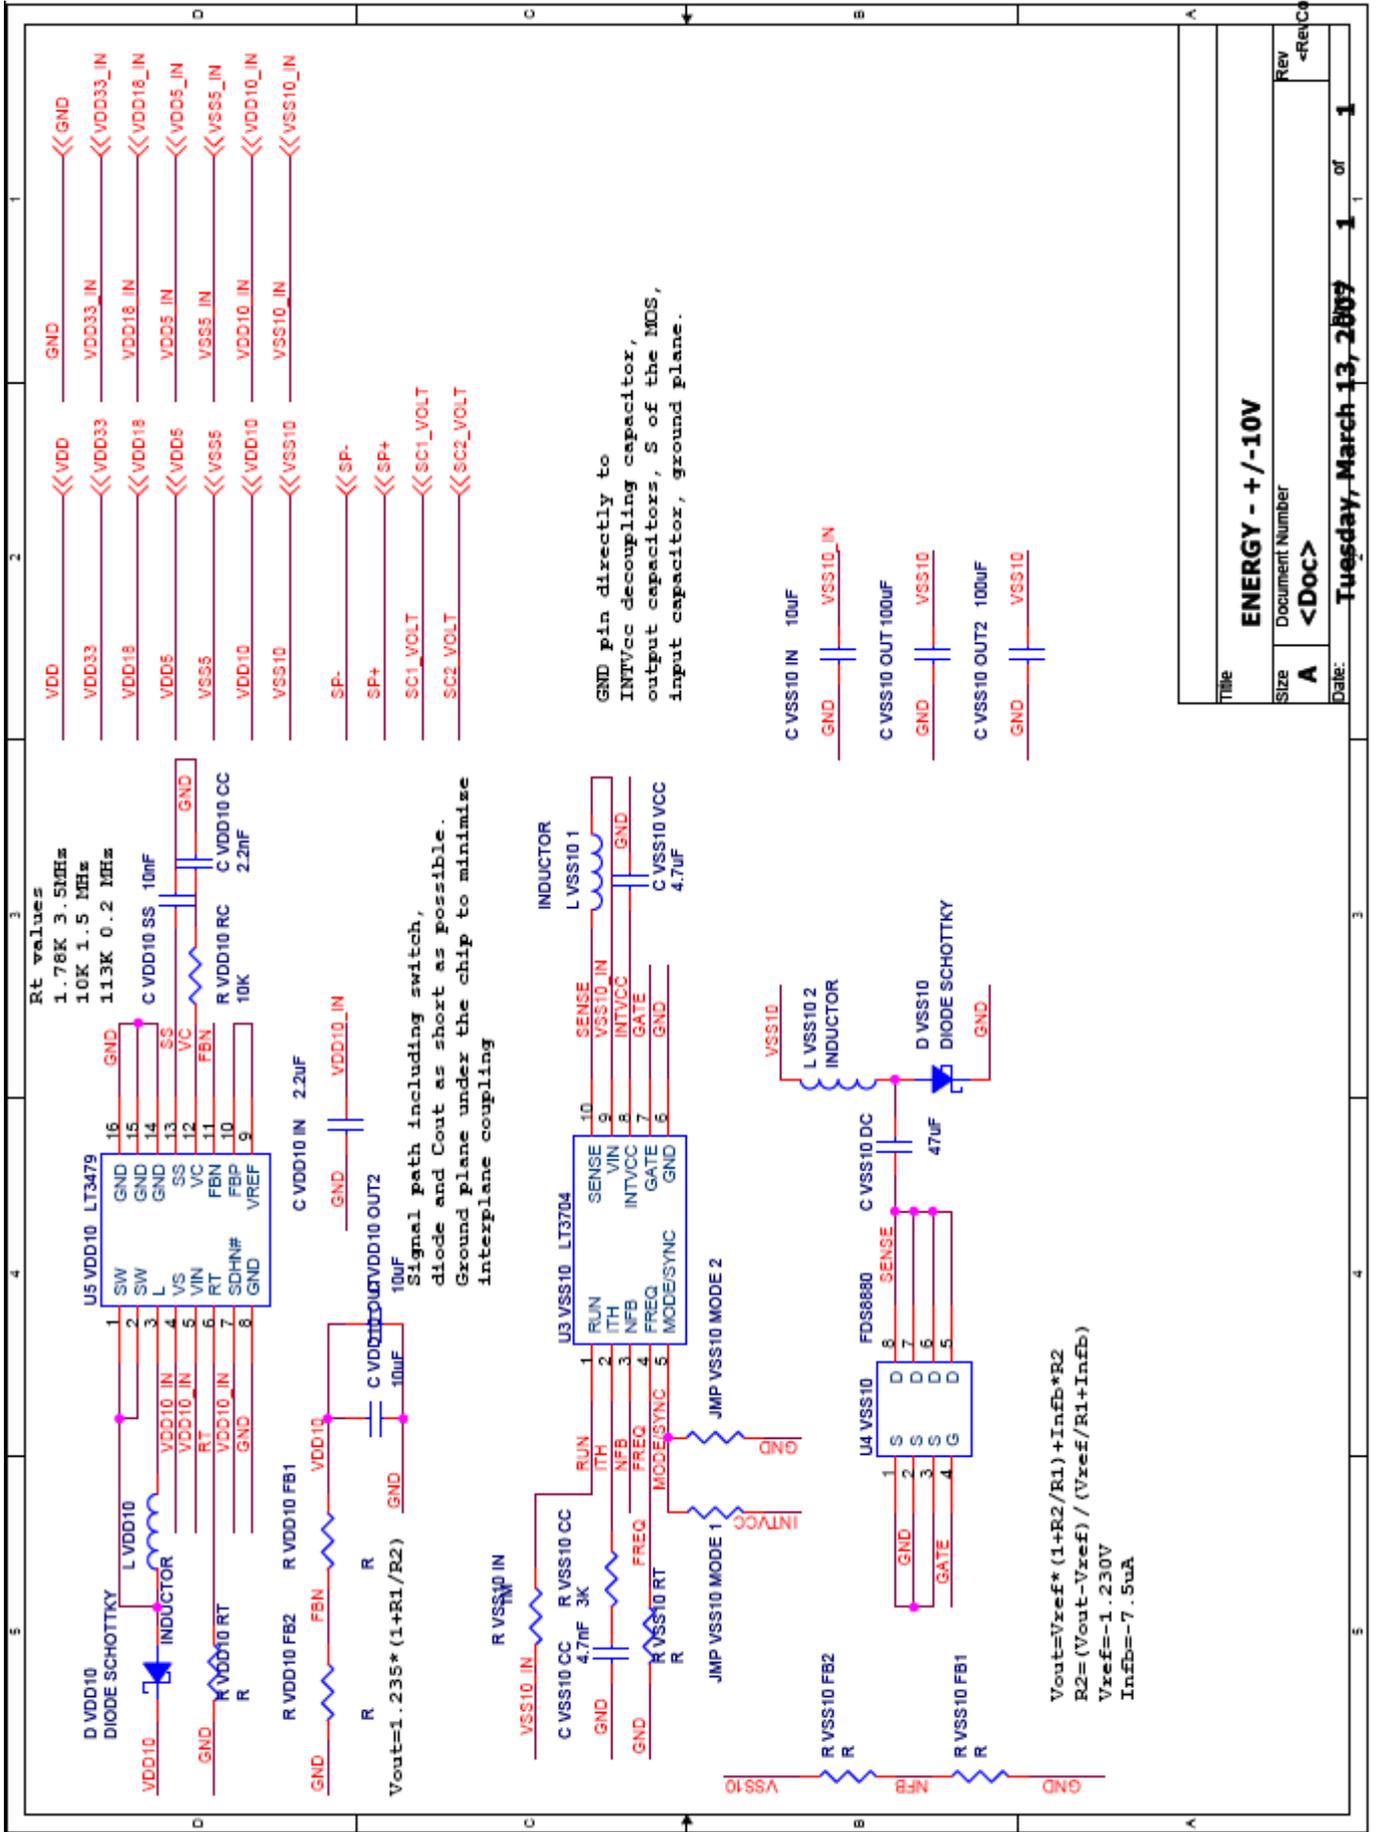
$$V_{out} = V_{ref} * (1 + R2/R1) + I_{nfB} * R2$$

$$R2 = (V_{out} - V_{ref}) / (V_{ref} / R1 + I_{nfB})$$

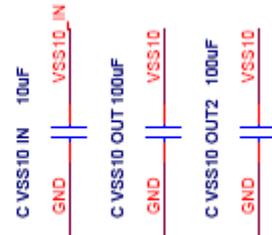
$$V_{ref} = -1.230V$$

$$I_{nfB} = -7.5uA$$

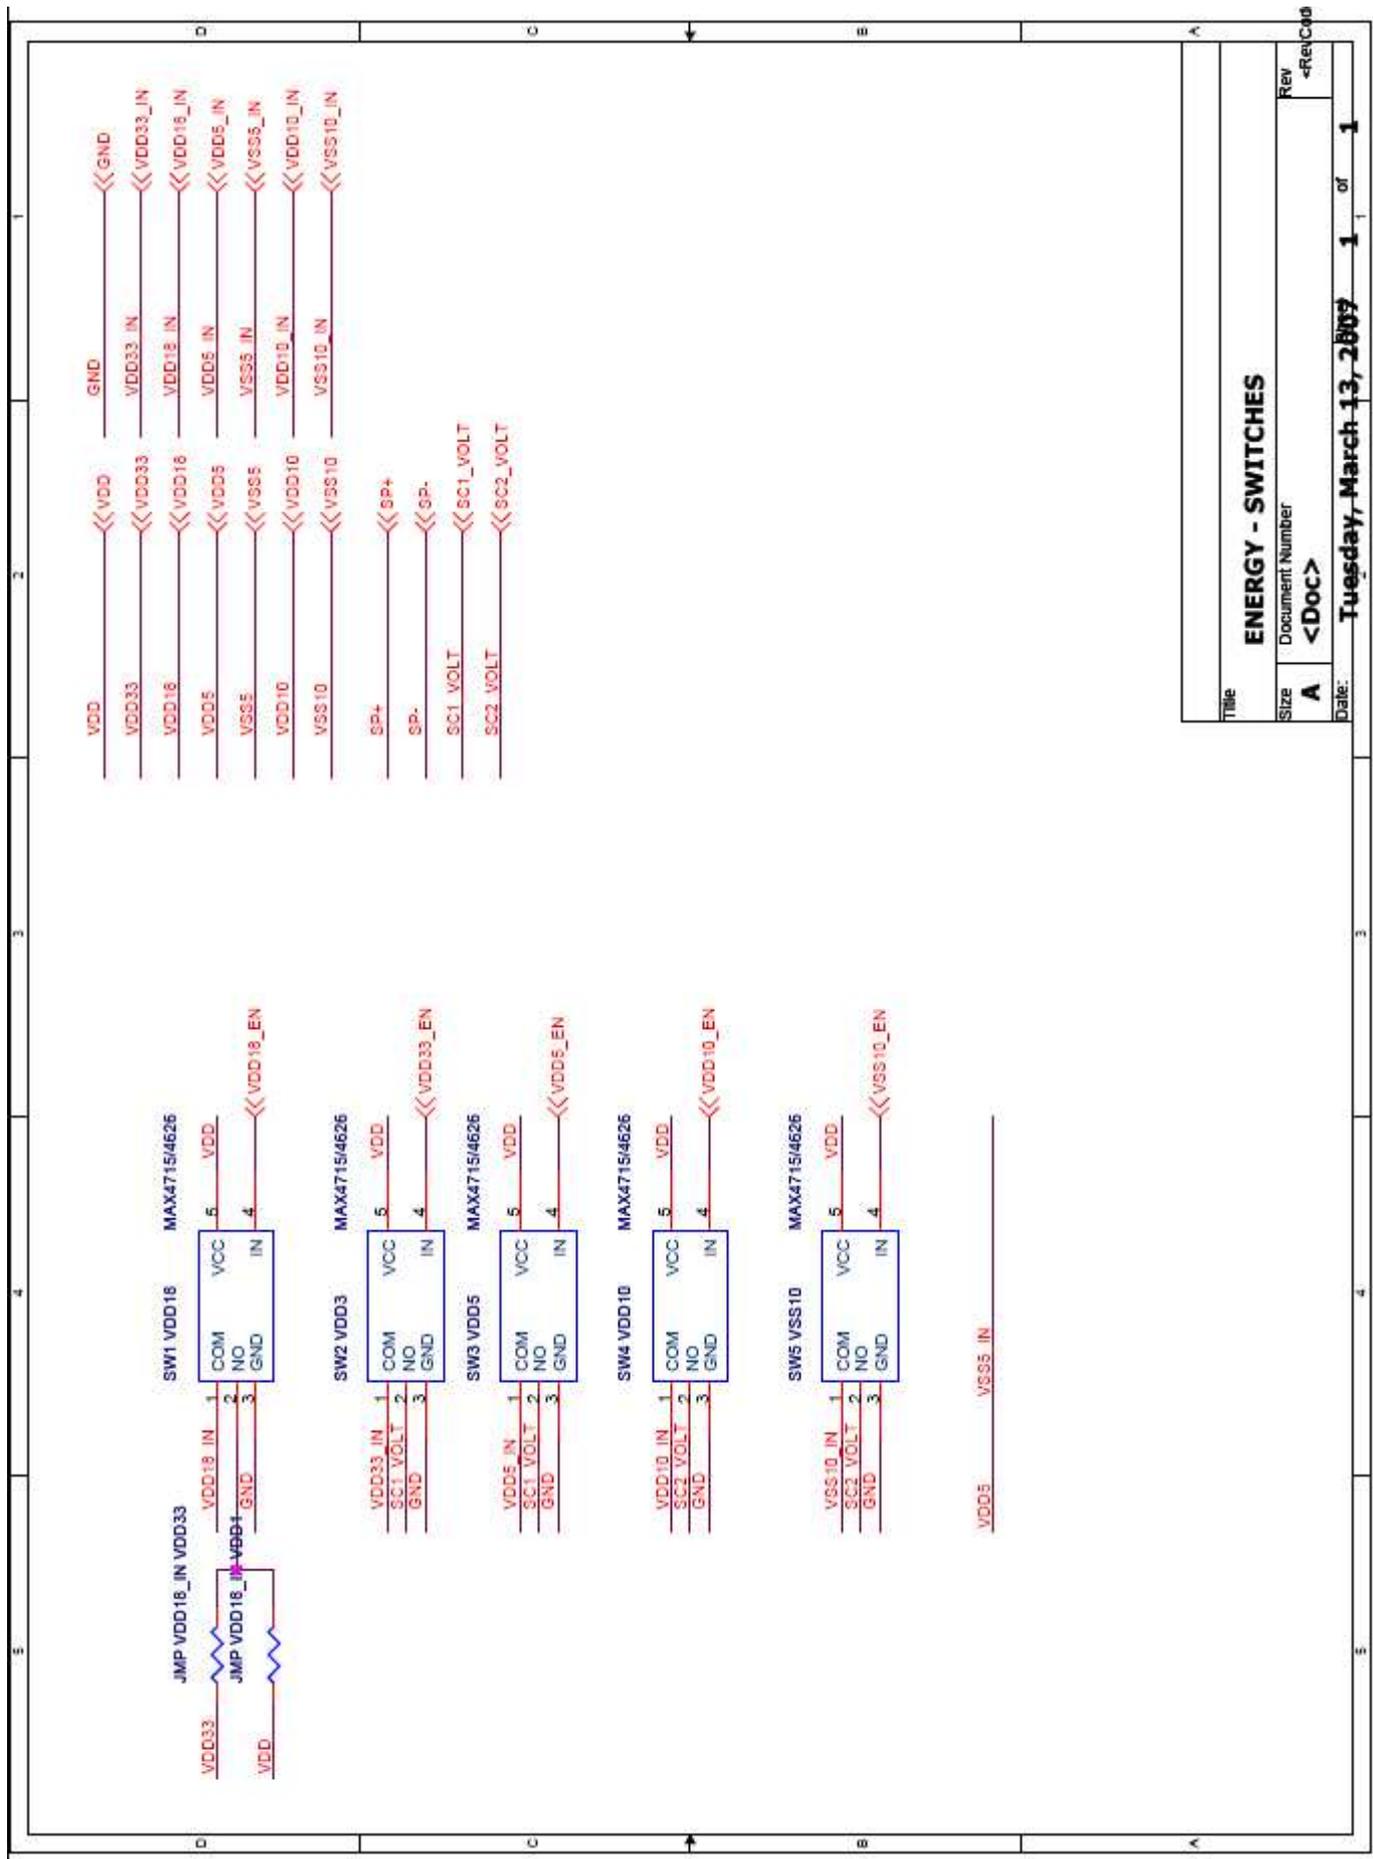
Title		ENERGY - + / -5V	
Size	Document Number	Rev	-Rev/Ct
A	<Doc>		
Date:	Tuesday, March 13, 2007	1	of 1



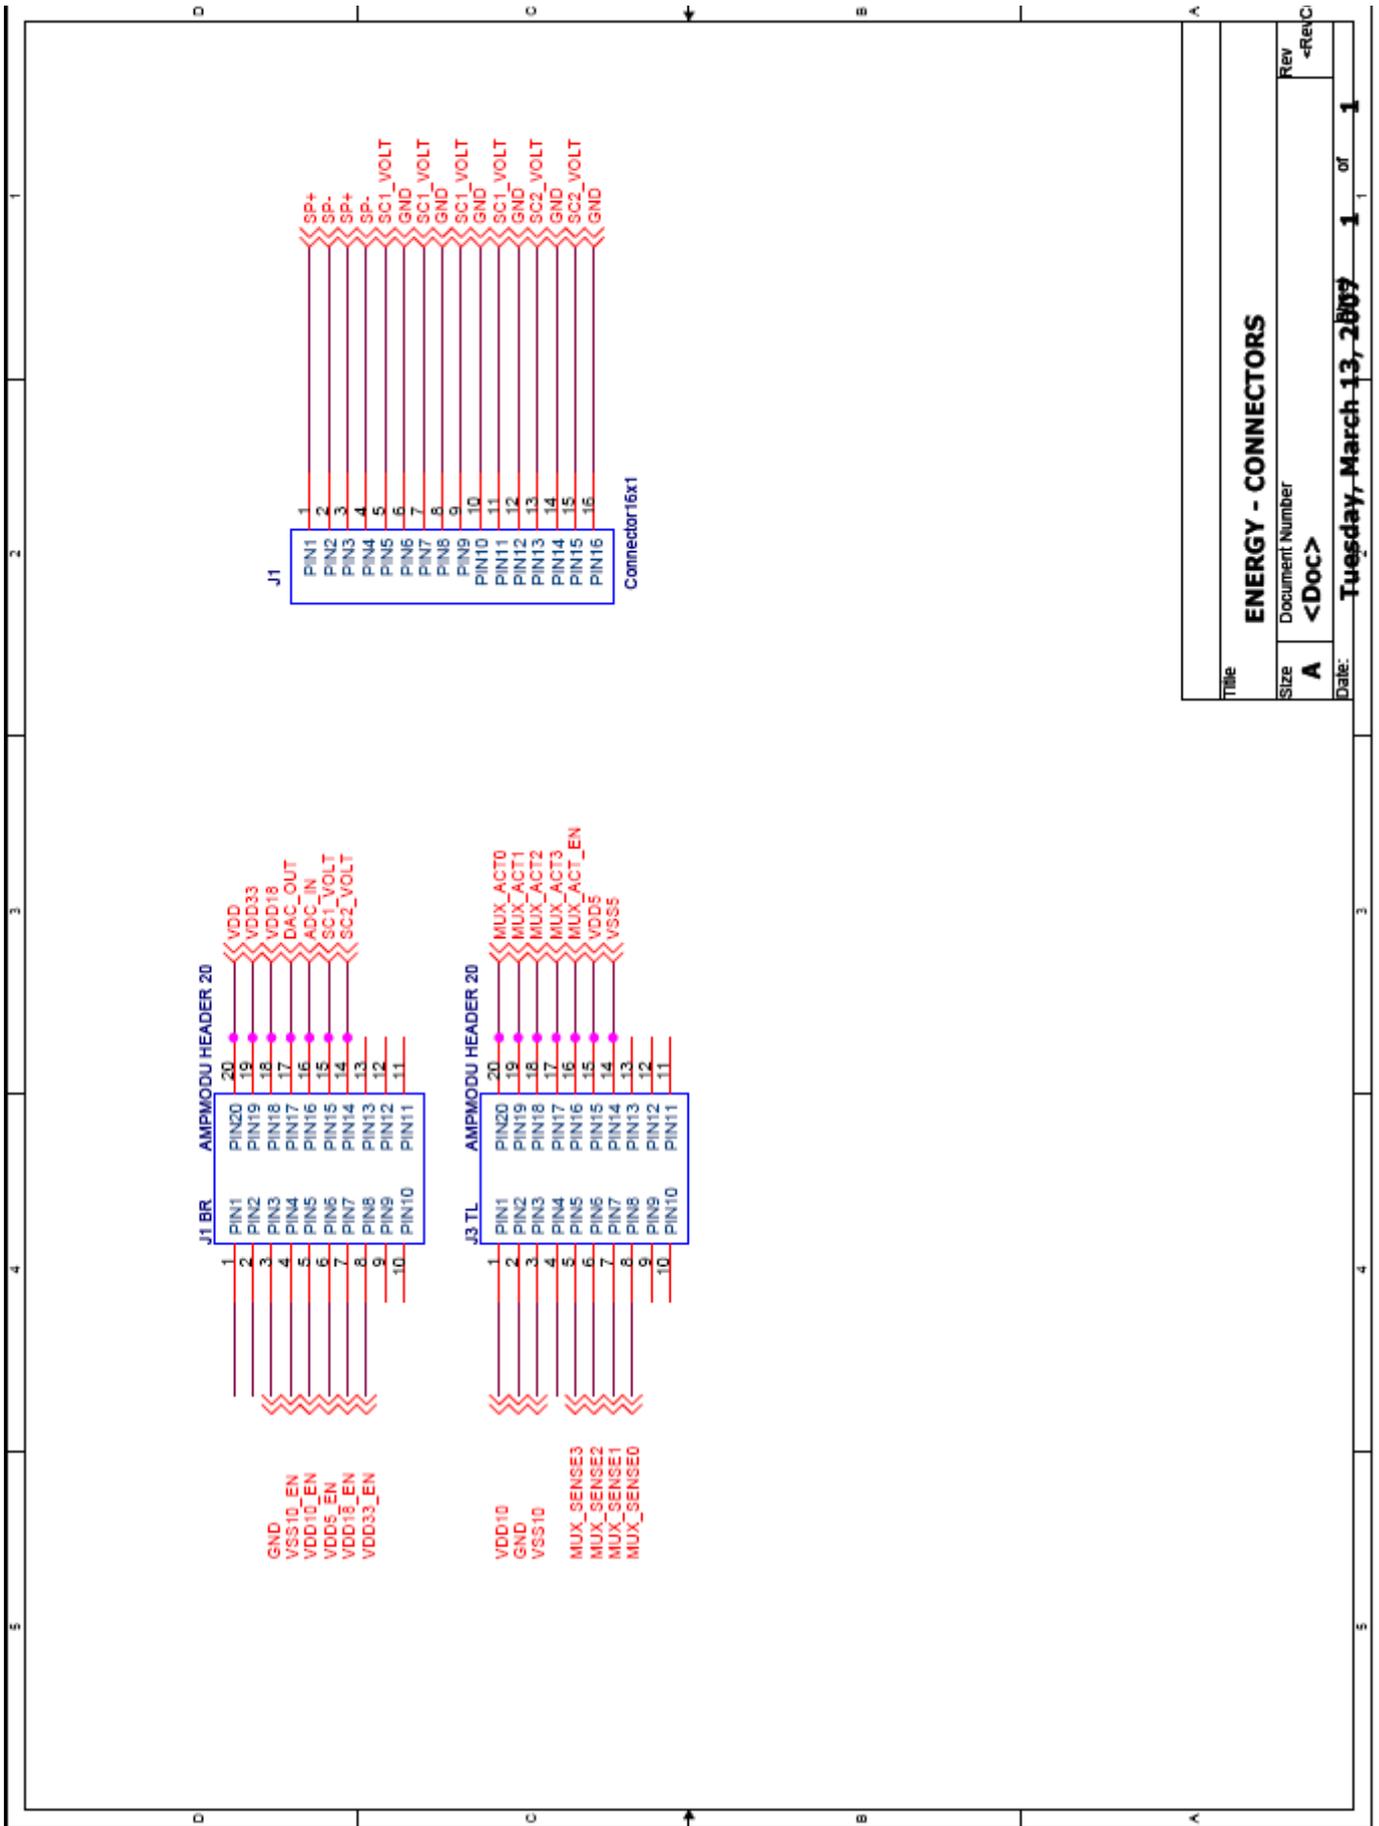
GND pin directly to INTVCC decoupling capacitor, output capacitors, S of the MOS, input capacitor, ground plane.



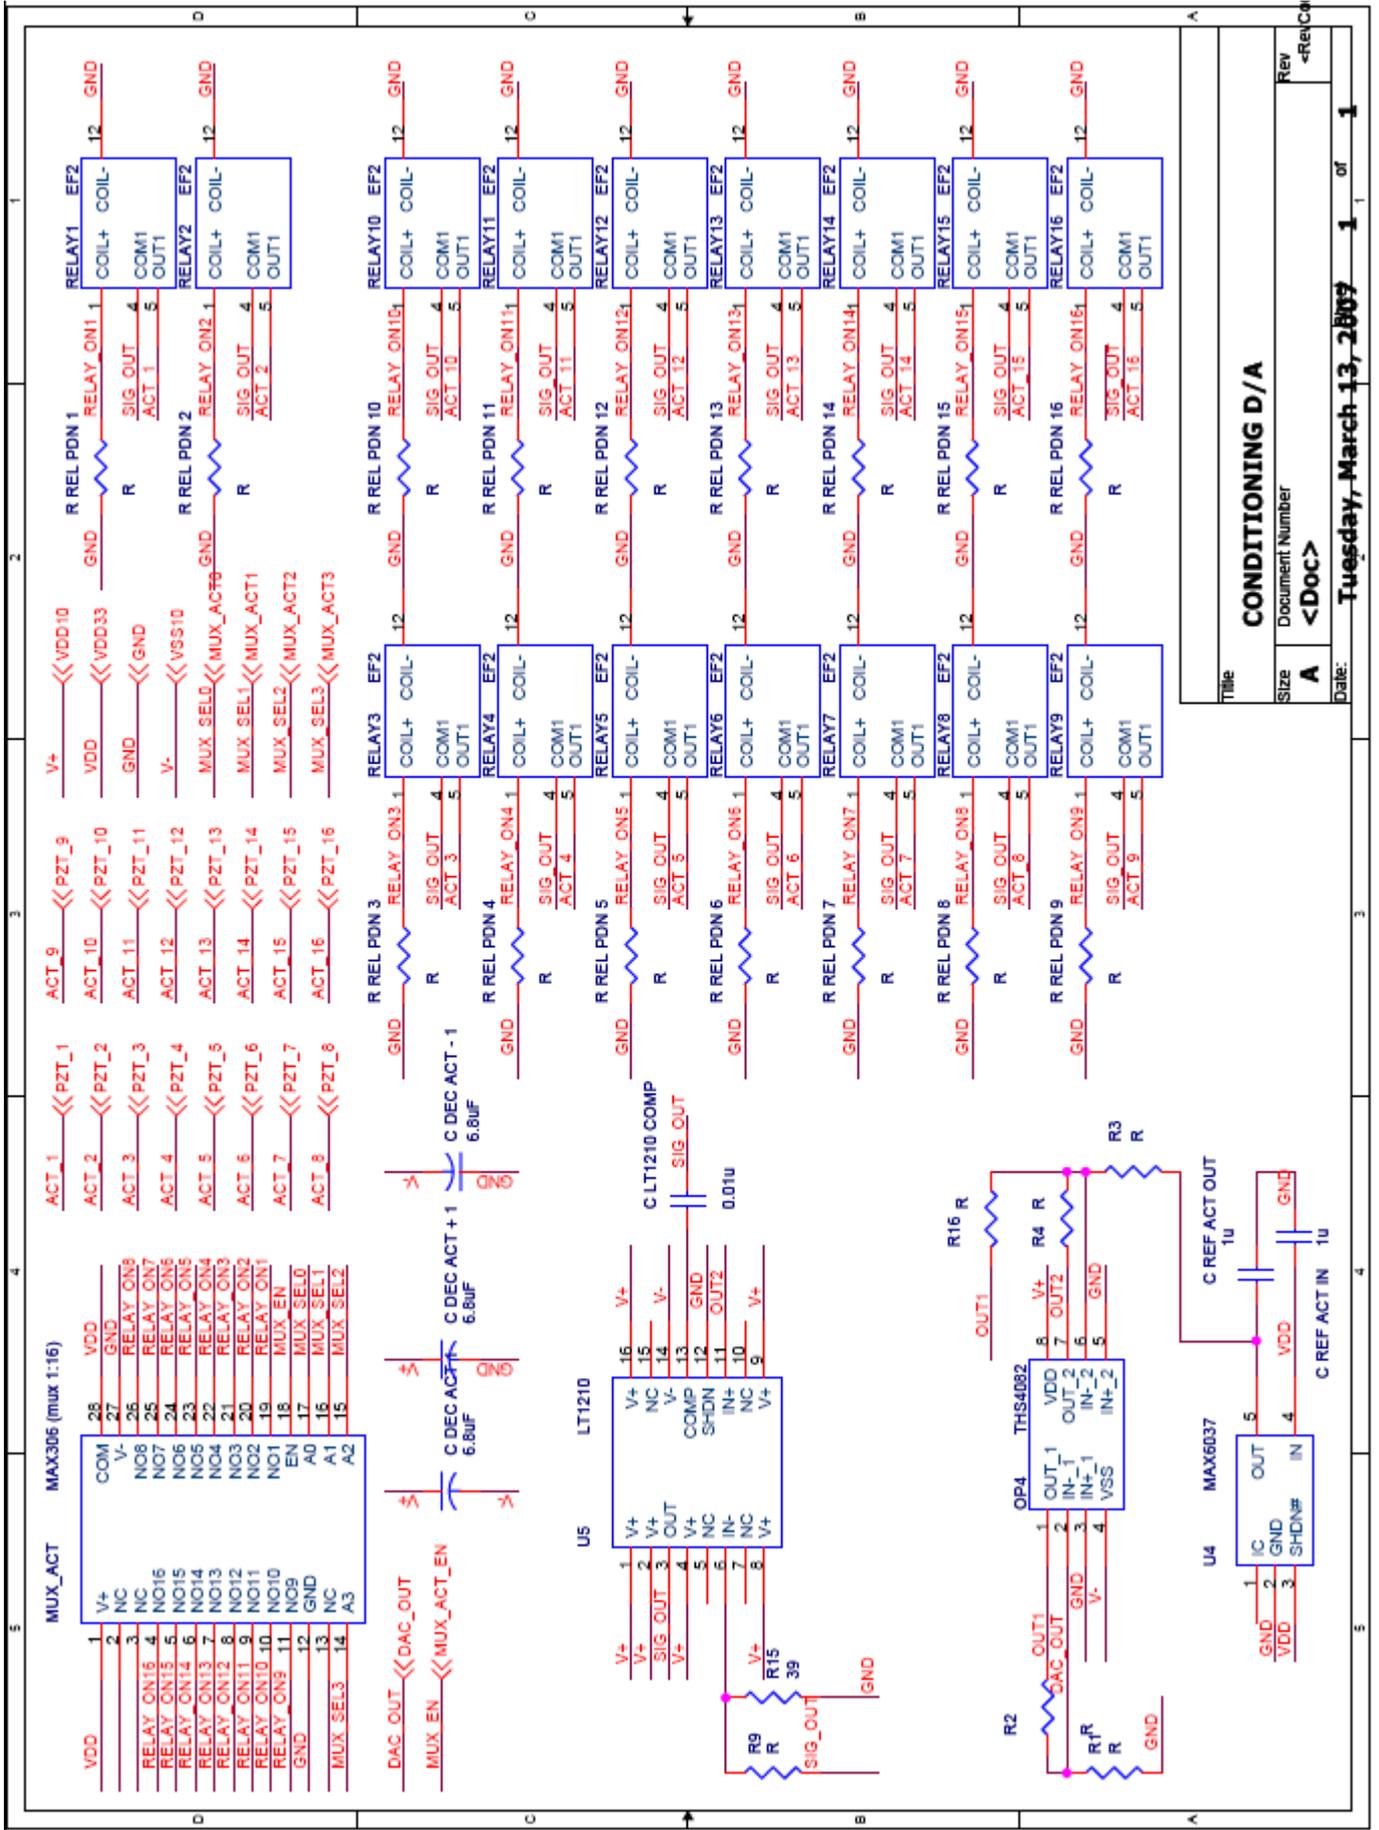
Title		ENERGY - +/-10V	
Size	Document Number	Rev	<RevC>
A	<Doc>		
Date:	Tuesday, March 13, 2007 1 of 1		



Title		ENERGY - SWITCHES	
Size	Document Number	REV	e-Ref/Doc
A	<Doc>		
Date:	Tuesday, March 13, 2007		1 of 1



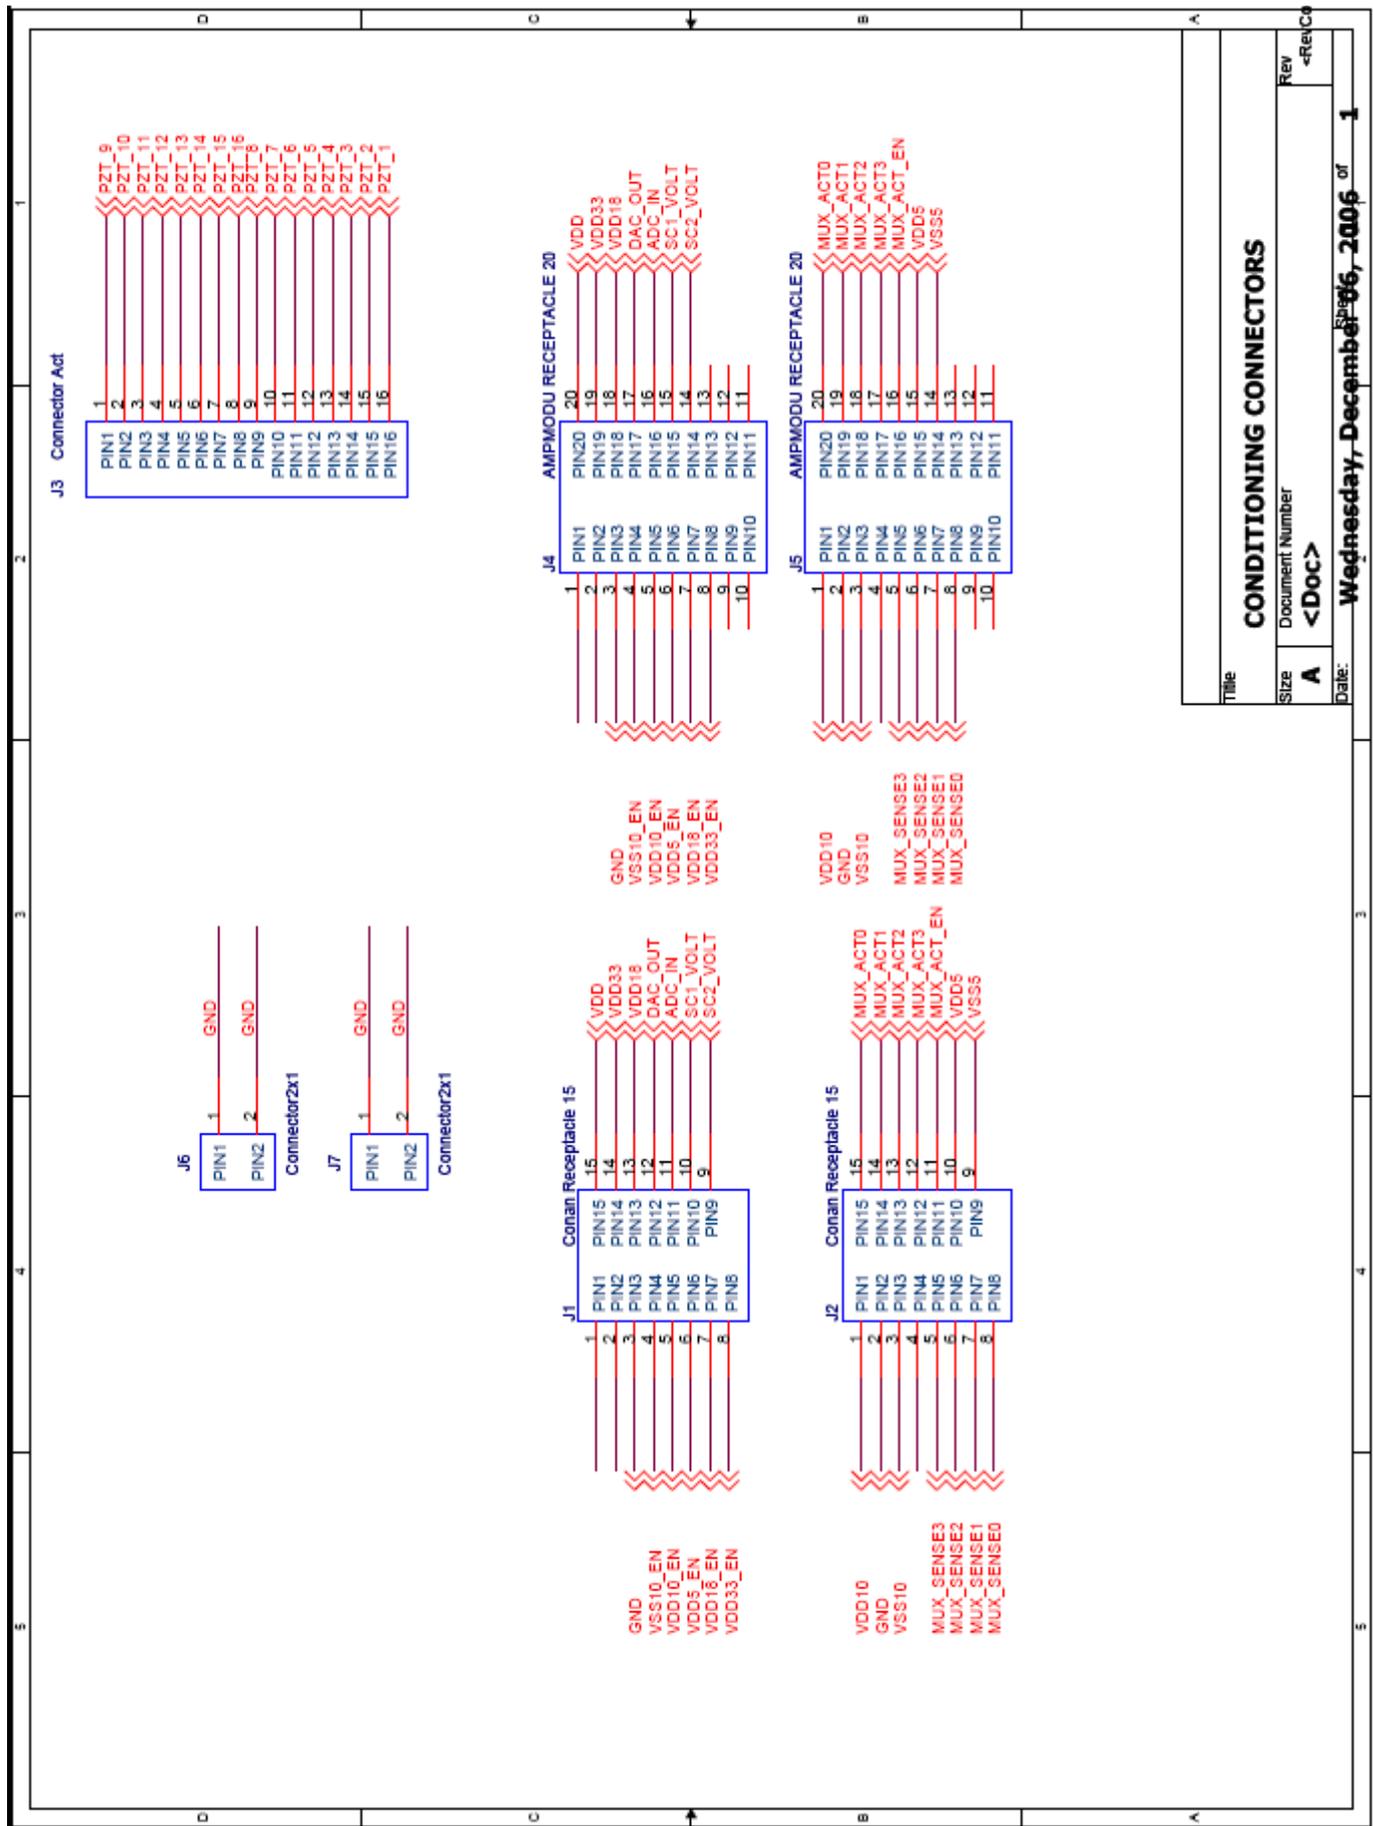
Title		ENERGY - CONNECTORS	
Size	Document Number	Rev	<RevC
A	<Doc>		
Date:	Tuesday, March 13, 2007 1 of 1		



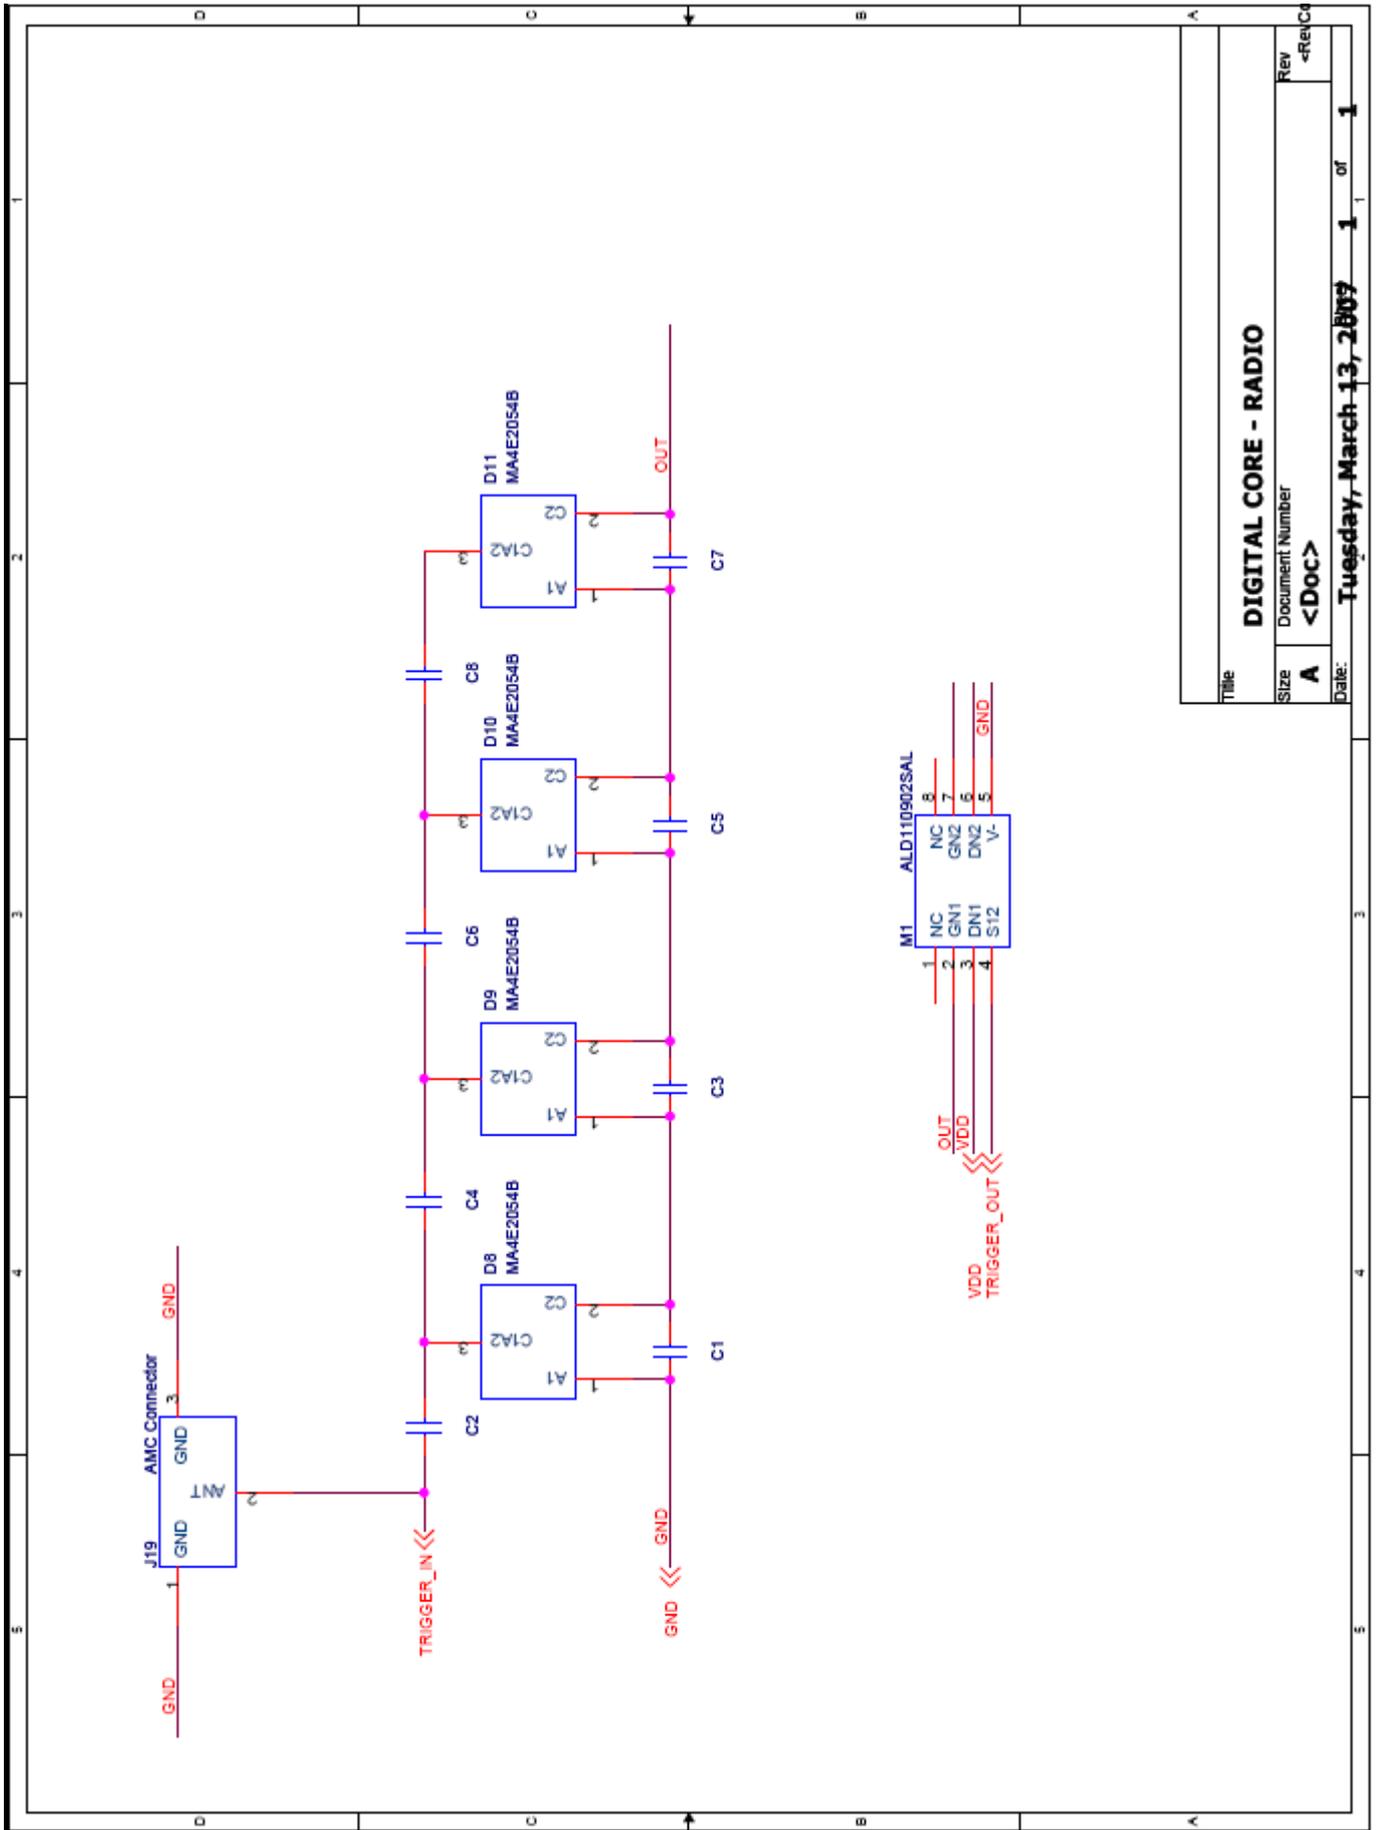
CONDITIONING D/A

Title
 Size A
 Document Number <Doc>
 Rev <RevC>

Date: **Tuesday, March 13, 2007** 1 of 1



Title	
CONDITIONING CONNECTORS	
Size	Document Number
A	<Doc>
Date:	Wednesday, December 06, 2006
Rev	1
<RevCo	



Title		DIGITAL CORE - RADIO	
Size	Document Number	Rev	<RevCo
A	<Doc>		
Date:	Tuesday, March 13, 2007	1 of 1	1

Appendix B – Prototype's pictures

Appendix B contains pictures of the final prototype.

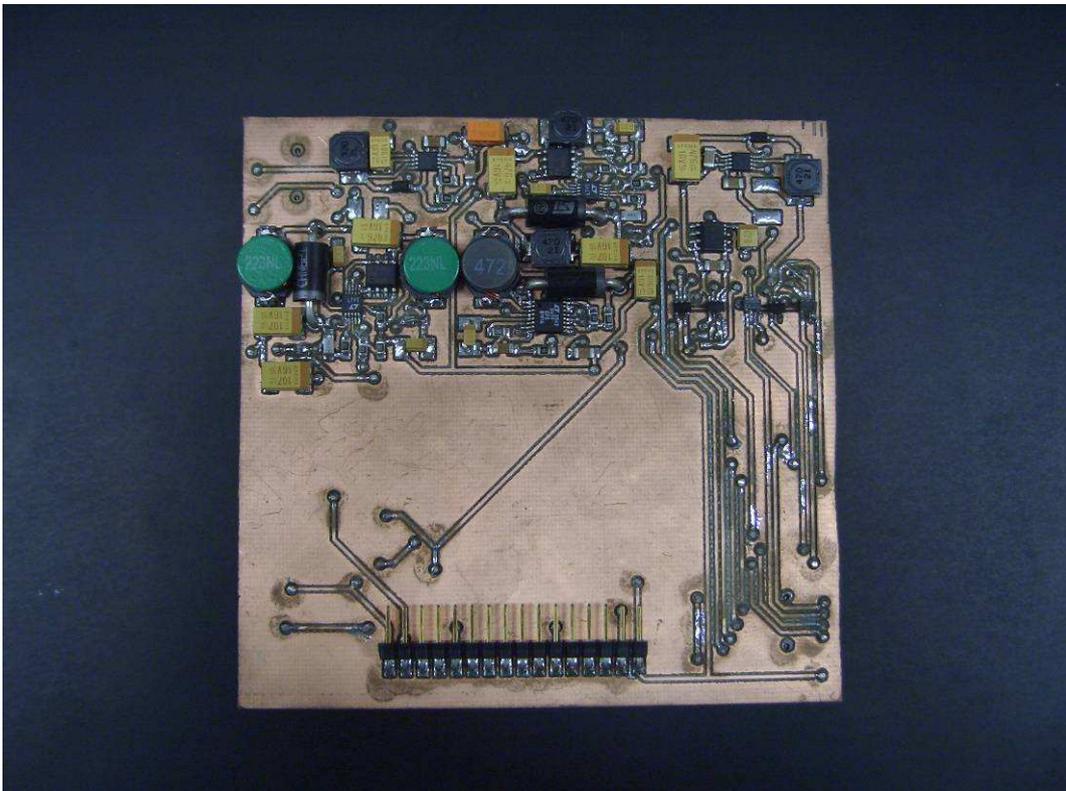


Figure B.1 – Energy board (bottom side).

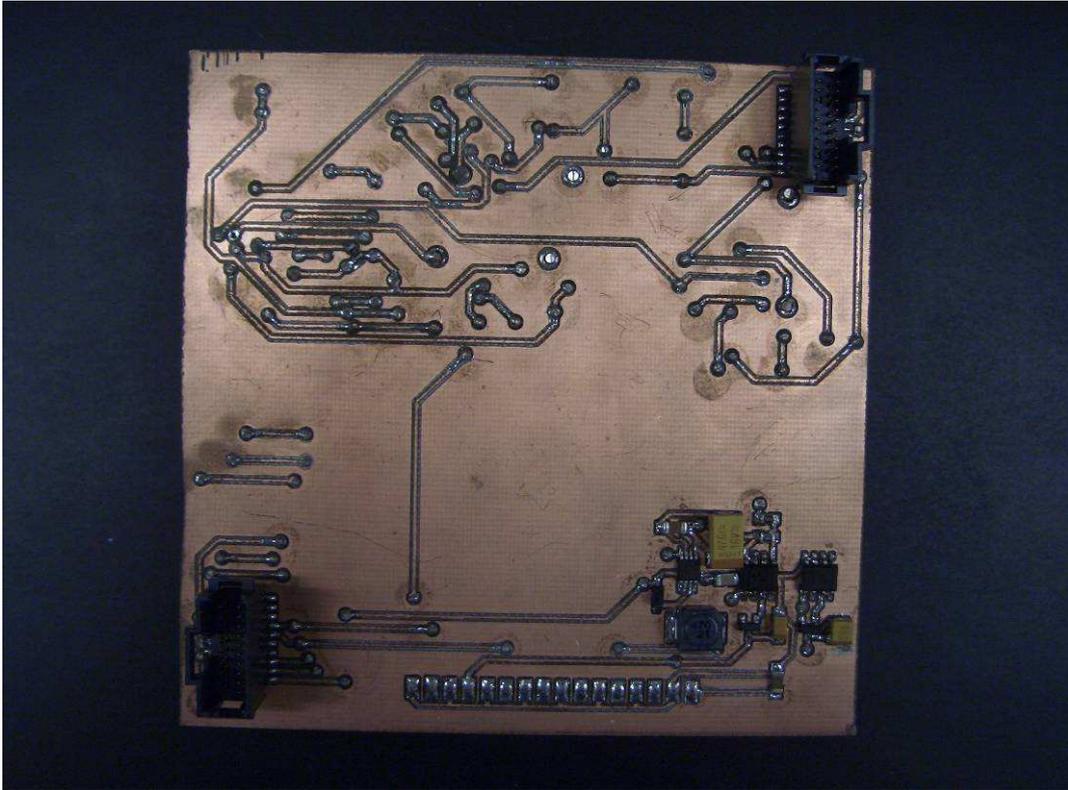


Figure B.2 - Energy board (top side).

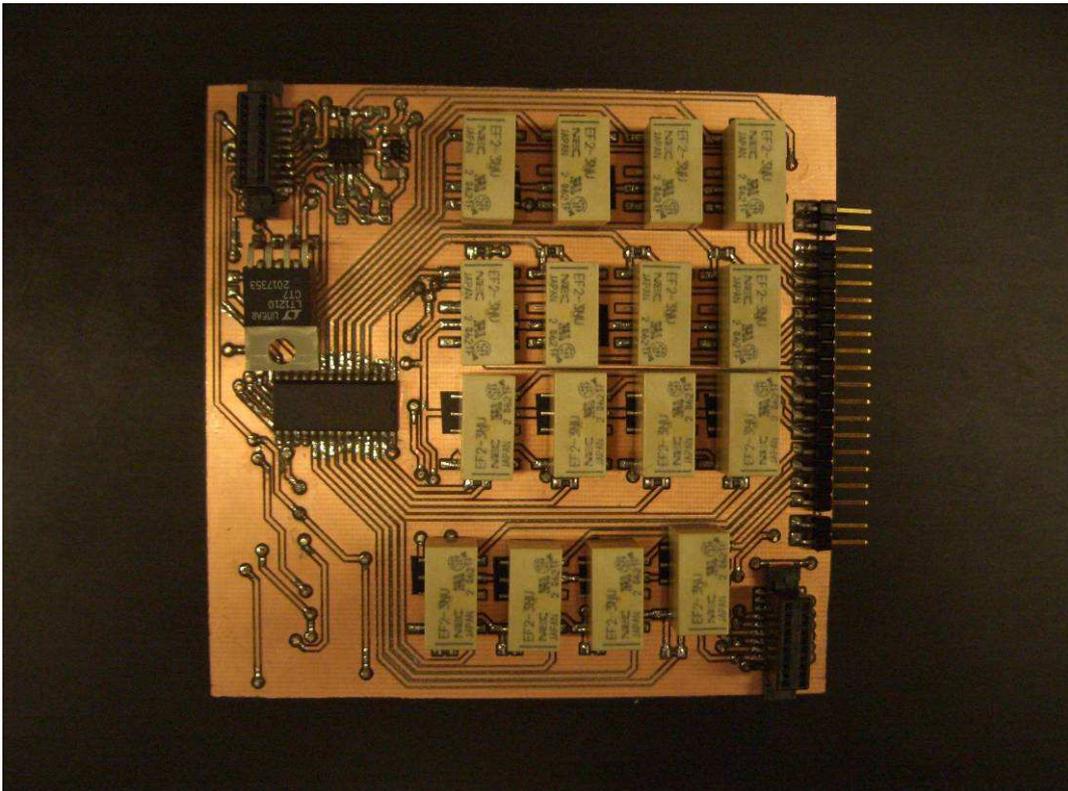


Figure B.3- I/O signal conditioning board (bottom side).

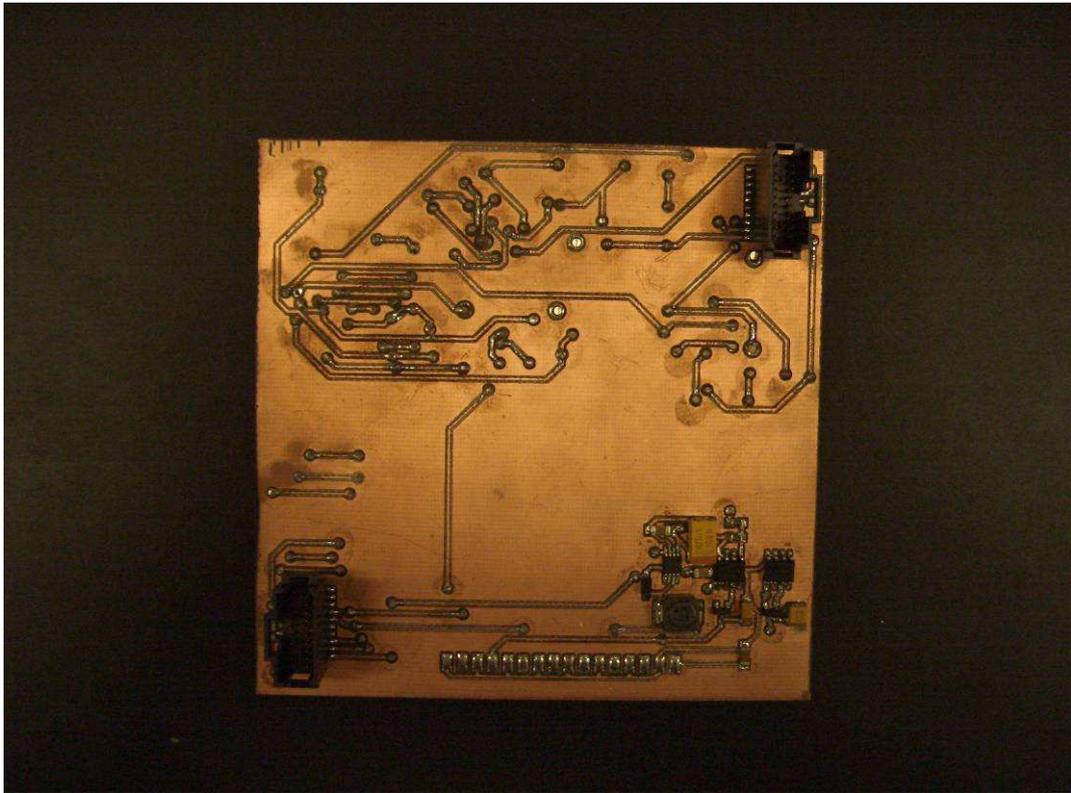


Figure B.4 -I/O signal conditioning board (top side).

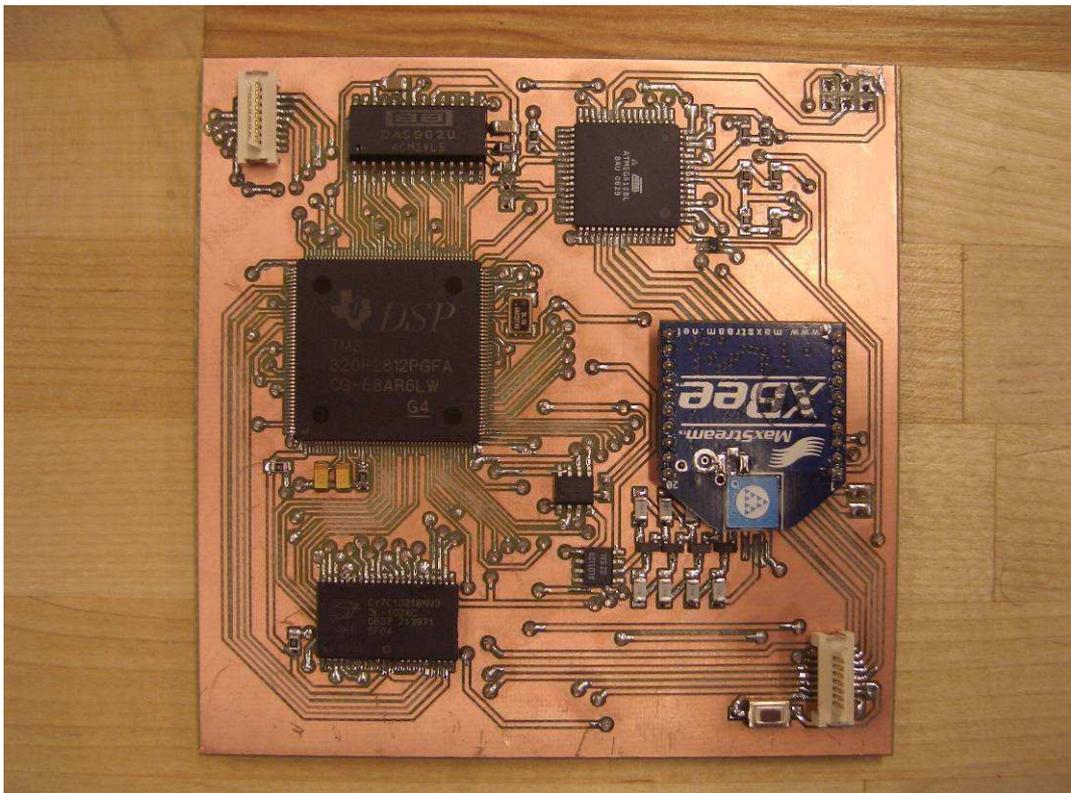


Figure B.1 – Digital core (bottom side).

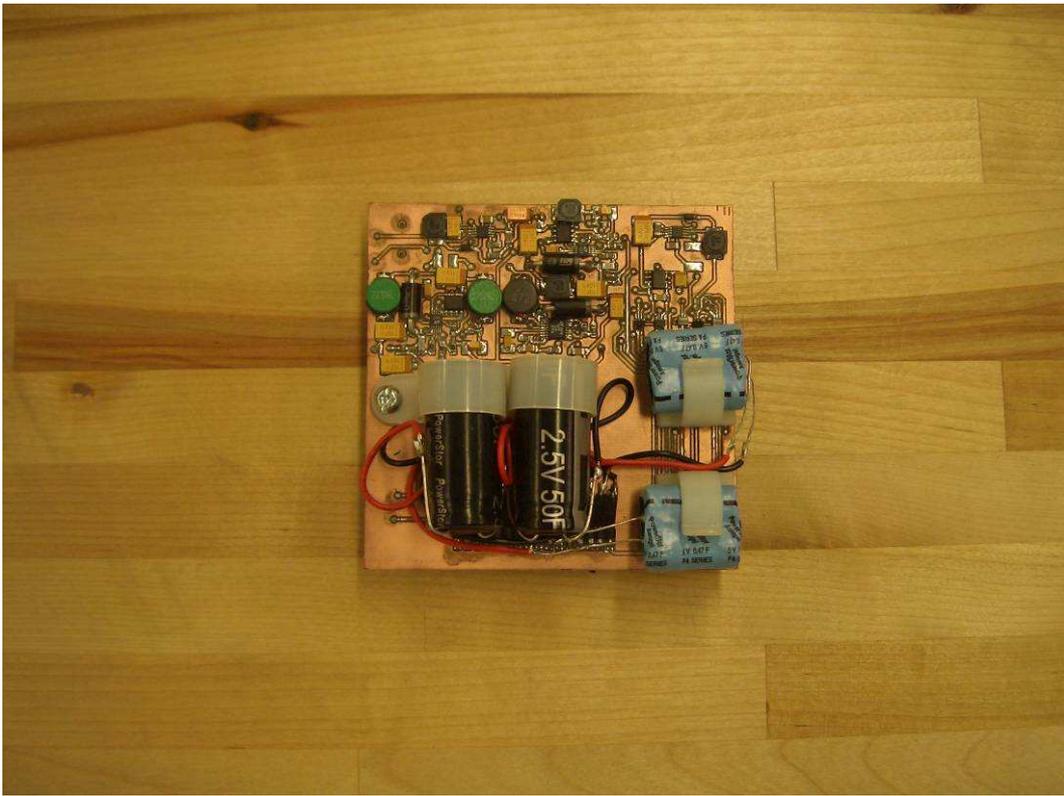


Figure B.2 – Final prototype (bottom view).

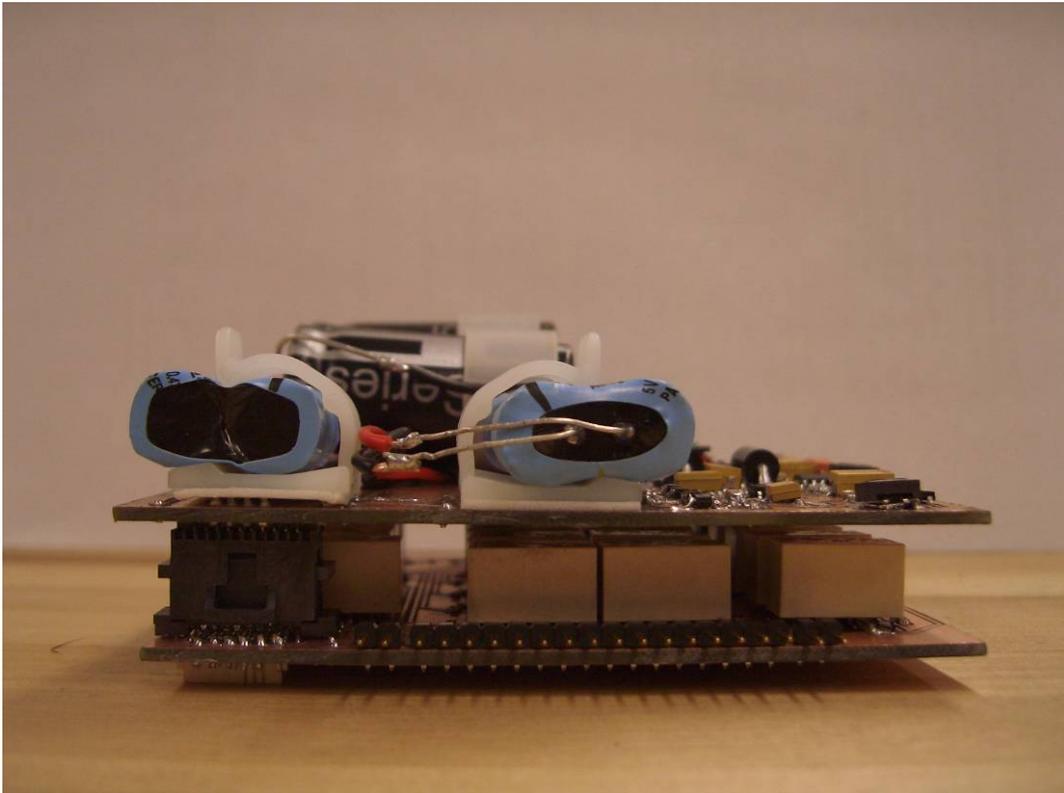


Figure B.3 – Final prototype (side view 1).

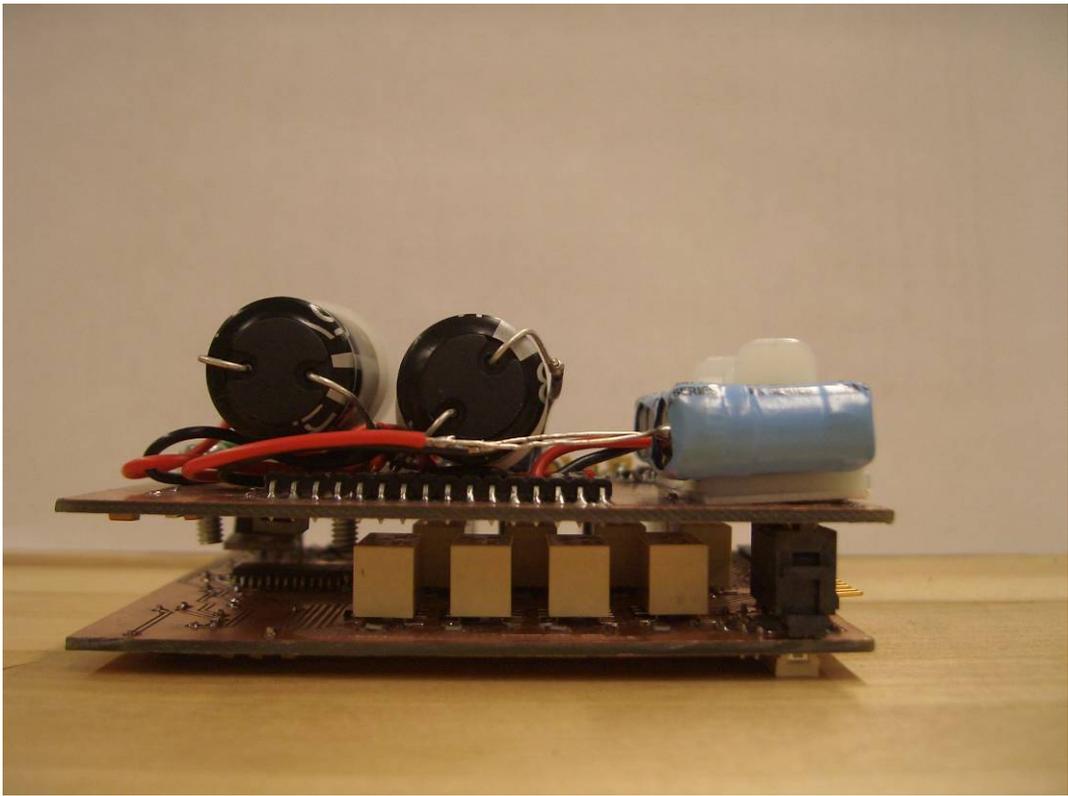


Figure B.4 - Final prototype (side view 2).

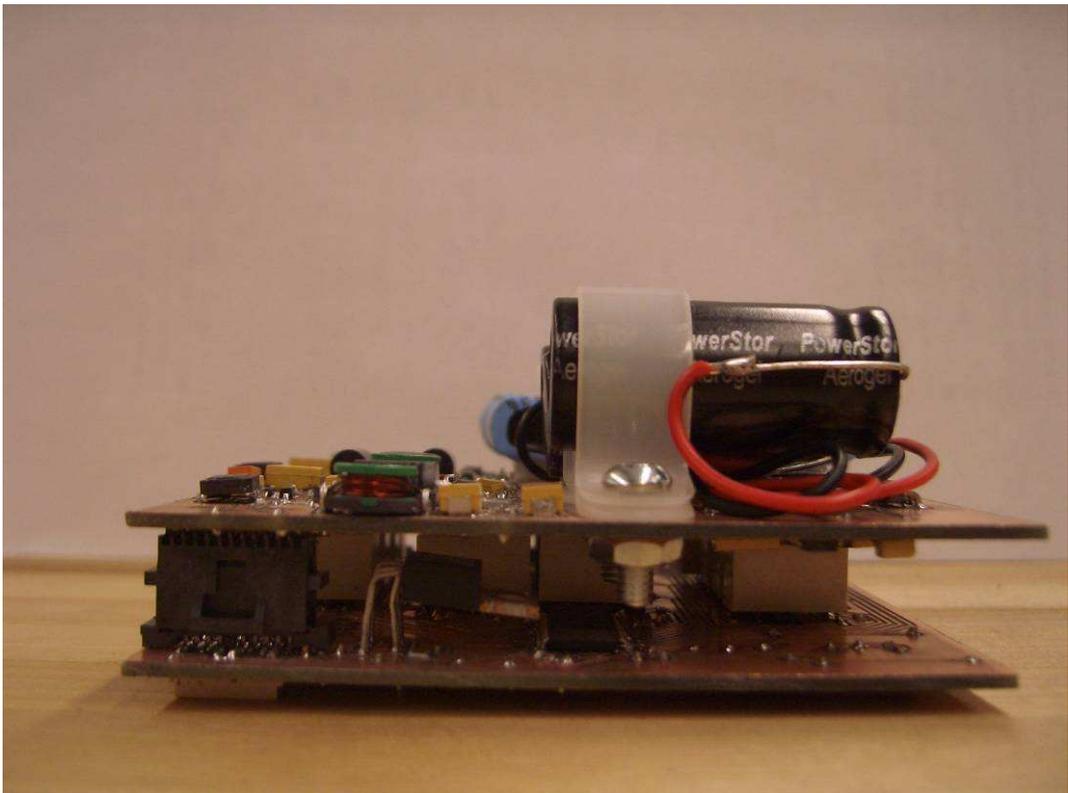


Figure B.5 - Final prototype (side view 3).

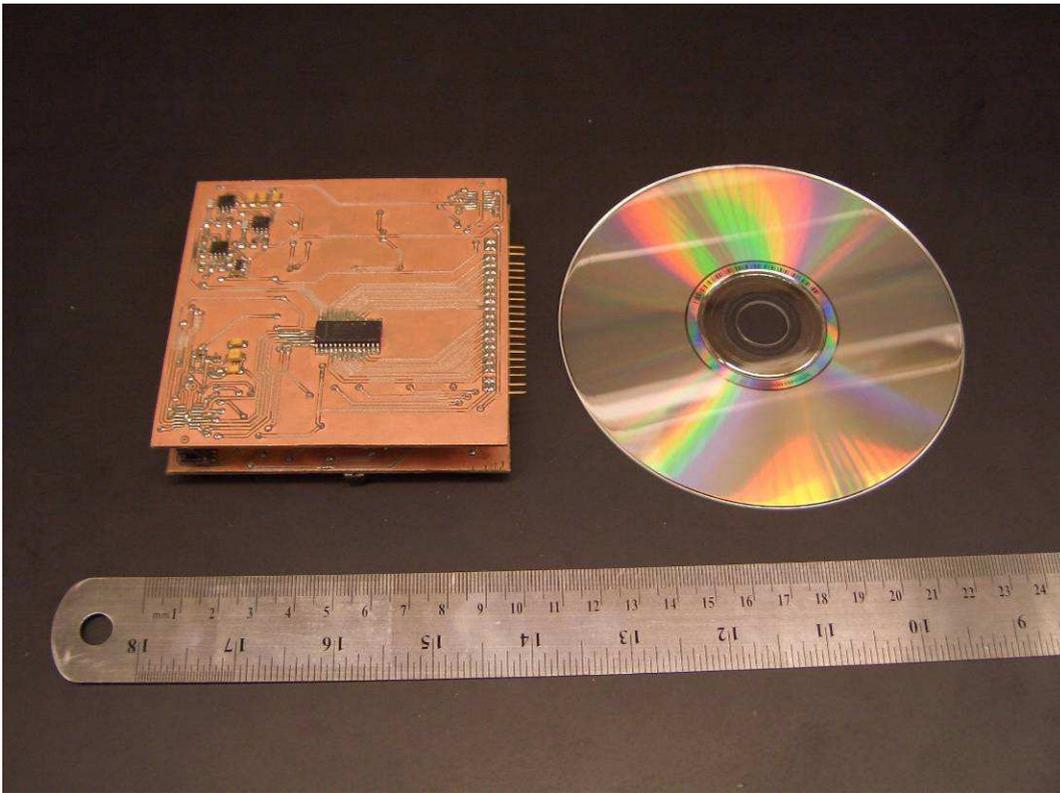


Figure B.6 –Comparison between the prototype’s and a CD’s sizes.



Figure B.7 – Prototype’s size measurement.

Appendix C – Communication protocol

This appendix contains detailed information about the communication protocol and data exchange between the node and the external agent and between the microcontroller and the DSP. The START and END fields' fixed values are 0xF0 and 0xFF respectively.



Node-UAV interaction

The following opcodes for node-UAV interaction are implemented in the current version of the microcontroller's firmware:

Opcode Name	Opcode value
OP_UAV_ERROR	0x00
OP_UAV_ACK	0x01
OP_UAV_RESULTS	0x02
OP_UAV_WAVEFORM	0x03
OP_UAV_SLEEP	0x04
OP_UC_ERROR	0x10
OP_UC_ACK	0x11
OP_UC_WAKE_UP_ACK	0x12
OP_UC_PU_OK	0x13
OP_UC_START_ACTUATION	0x14
OP_UC_START_PROCESSING	0x15
OP_UC_RESULTS_READY	0x16
OP_UC_RESULTS	0x17
OP_UC_WAVEFORM	0x18
OP_UC_SLEEP	0x19

The basic steps of the UAV-node interaction are:

Sender	Receiver	Opcode	Payload
Node	UAV	OP_UC_WAKE_UP_ACK	None
UAV	Node	OP_UAV_ACK	Settings (two bytes for actuation frequency, two bytes for actuation duration, one byte for actuator index, one byte for sensor index)
Node	UAV	OP_UC_PU_OK	None
Node	UAV	OP_UC_START_ACTUATION	None
Node	UAV	OP_UC_START_PROCESSING	None
Node	UAV	OP_UC_RESULTS_READY	None
UAV	Node	OP_UAV_RESULTS	TBD
Node	UAV	OP_UC_RESULTS	TBD

Acquired samples can be retrieved by exchanging the following messages:

Sender	Receiver	Opcode	Payload
UAV	Node	OP_UAV_WAVEFORM	Samples index (samples are organized in 254-wide groups)
Node	UAV	OP_UC_WAVEFORM	Requested samples

Microcontroller-DSP interaction

A set of opcodes dedicated to the DSP was defined. The opcodes are listed in the following table.

Opcode name	Opcode value
OP_DSP_ERROR	0x20
OP_DSP_ACK	0x21
OP_DSP_START_ACTUATION	0x22
OP_DSP_START_PROCESSING	0x23
OP_DSP_RESULTS_READY	0x24
OP_DSP_RESULTS	0x25
OP_DSP_WAVEFORM	0x26

The basic steps of the microcontroller-DSP interaction are:

Sender	Receiver	Opcode	Payload
DSP	μC	OP_UC_ACK	None
μC	DSP	OP_UC_ACK	Settings
DSP	μC	OP_UC_START_ACTUATION	None
DSP	μC	OP_UC_START_PROCESSING	None
DSP	μC	OP_UC_RESULTS_READY	None
UAV	μC	OP_UAV_RESULTS	TBD

Acquired samples can be retrieved by exchanging the following messages:

Sender	Receiver	Opcode	Payload
μC	DSP	OP_UC_WAVEFORM	Samples index (samples are organized in 254-wide groups)
DSP	μC	OP_DSP_WAVEFORM	Requested samples

Error messages

The communication protocol allows all entities to send error messages including a description of the specific error. An error message is identified by an opcode value chosen among OP_UAV_ERROR, OP_UC_ERROR and OP_DSP_ERROR, depending on the entity sending the message. The kind of error can be specified by setting the first byte of the payload to one of the values listed in the following table.

Error ID	Error value
ERROR_FORMAT	0x01
ERROR_ENERGY_LOW_MAIN	0x02
ERROR_ENERGY_LOW_AUX	0x03
ERROR_VDD33_FAIL	0x04
ERROR_VDD18_FAIL	0x05
ERROR_VDD10_FAIL	0x06
ERROR_VDD5_FAIL	0x07
ERROR_DSP_COMM	0x08
ERROR_CRC	0x09
ERROR_RESULTS_EXPECTED	0x0A

ERROR_PZT_INDEXES	0x0B
ERROR_ACK_EXPECTED	0x0C
ERROR_ACK_EXPECTED_FROM_DSP	0x0D
ERROR_ACK_EXPECTED_FROM_UAV	0x0E
ERROR_PU_KO	0x0F
ERROR_TIMEOUT_UAV_ACK	0x10
ERROR_TIMEOUT_DSP_ACK	0x11
ERROR_TIMEOUT_START_ACTUATION	0x12
ERROR_TIMEOUT_START_PROCESSING	0x13
ERROR_TIMEOUT_RESULTS_READY	0x14
ERROR_TIMEOUT_DISPATCH_RESULTS	0x15

Bibliography

1. Adams, R. D., Cawley, P., Pye, C. J. and Stone, B. J, 1978, “A Vibration Technique for Non-Destructively assessing the Integrity of Structures”, *Journal of Mechanical Engineering Science*, Vol. 20, 93–100.
2. Aktan, A. E., Catbas, F. N., Grimmelsman, K. A. and Tsikos, C. J, 2000, “Issues in Infrastructure Health Monitoring for Management”, *Journal of Engineering Mechanics*, Vol. 126, 711–724.
3. Alleyne, D. N. and Cawley, P, 1992, “The Interaction of Lamb Waves with Defects”, *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, Vol. 39, 381–397.
4. Alleyne, P., Cawley, D. “The interaction of lamb waves with defects”. In *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*,, pages 381–397, May 1992.
5. Anandarajah, A., Moore, K., Terzis, A., and I.-J. Wang. “Sensor networks for landslide detection”. In *SenSys '05: Proceedings of the 3rd international conference on Embedded networked sensor systems*, pages 268–269, New York, NY, USA, 2005. ACM Press.
6. ATMega128L <http://www.atmel.com>.
7. Backman, D., Flynn, E., Swartz, R., Hundhausen, R., and Park, G.. “Active piezoelectric sensing for damage identification in honeycomb composite panels”. In *Proceedings of 24th International Modal Analysis Conference*, 2006.
8. Boltezar, M., Strancar, B. and Kuhelj, A, 1998, “Identification of Transverse Crack Location in Flexural Vibrations of Free-Free Beams”, *Journal of Sound and Vibration*, Vol. 211, 729–734.
9. Brincker, R. and Andersen, P, 2003, “A Way of getting Scaled Mode Shapes in Output Only Modal Testing”, In *Proceedings of the 21st International Modal Analysis Conference (IMAC XXI)*, Kissimmee, Florida, USA, paper no. 141.
10. Brincker, R., Andersen, P., Kirkegaard, P. H. and Ulfkjaer, J. P, 1995, “Damage Detection in Laboratory Concrete Beams”, In *Proceedings of the 13th International Modal Analysis Conference (IMAC XIII)*, Nashville, Tennessee, USA, 668–674.

11. Brincker, R., Kirkegaard, P. H., Andersen, P. and Martínez, M. E, 1995, “Damage Detection in an Offshore Structure”, In Proceedings of the 13th International Modal Analysis Conference (IMAC XIII), Nashville, Tennessee, USA, 661–667.
12. Brincker, R., Ventura, C. E. and Andersen, P, 2003, “Why Output-Only Modal Testing is a Desirable Tool for a Wide Range of Practical Applications”, In Proceedings of the 21st International Modal Analysis Conference (IMAC XXI), Kissimmee, Florida, USA, paper no. 265.
13. Chintalapudi, K., Paek, J., Gnawali, O., Fu, T. S., Dantu, K., Caffrey, J., Govindan, R., Johnson, E., and Masri, S.. “Structural damage detection and localization using netshm”. In IPSN '06: Proceedings of the fifth international conference on Information processing in sensor networks, pages 475–482, New York, NY, USA, 2006. ACM Press.
14. Cooper-Bussman supercapacitors <http://www.cooperet.com>
15. Doebling, S. W., Farrar, C. F., Prime, M. B. and Shevits, D. W, 1996, “Damage Identification and Health Monitoring of Structural and Mechanical Systems from Changes in their Vibration Characteristics: A Literature Review”, Los Alamos National Laboratory, USA.
16. Doebling, S. W., Farrar, C. R. and Prime, M. B, 1998, “A Summary Review of Vibration-Based Damage Identification Methods”, Shock and Vibration Digest, Vol. 30, 91–105.
17. Enjeti, P., Howze, J., and Palma, L.. “An approach to improve battery run-time in mobile applications with supercapacitors”. In IEEE 34th Annual Power Electronics Specialists Conference, 2003.
18. Farrar, C. R., Baker, W. E., Bell, T. M., Cone, K. M., Darling, T. W., Duffey, T. A., Eklund, A. and Migliori, A, 1994,, “Dynamic Characterization and Damage Detection in the I-40 Bridge over the Rio Grande”, Los Alamos National Laboratory, USA.
19. Farrar, C. R., Doebling, S. W. and Nix, D. A, 2001, “Vibration-based Structural Damage Identification”, Philosophical Transactions of the Royal Society of London Series A-Mathematical Physical and Engineering Sciences, Vol. 359, 131–149.
20. Farrar, C. R., Sohn, H. and Robertson, A. N, 2004, “Applications of Nonlinear System Identification to Structural Health Monitoring”, In Proceedings of the 2nd European Workshop on Structural Health Monitoring, Forum am Deutschen Museum, Munich, Germany, 59–67.

21. Farrar, C. R., Sohn, H., Hemez, F. M., Anderson, M. C., Bement, M. T., Cornwell, P. J., Doebling, S. W., Lieven, N., Robertson, A. N. and Schultze, J. F, 2003, “Damage Prognosis: Current Status and Future Needs”, Los Alamos National Laboratory, USA.
22. FIRE. Fire project: Smokenet. <http://fire.me.berkeley.edu/smokeNet.html>.
23. Galvin, P. and Dominguez, J, 2005, “Modal Identification of a Pedestrian Bridge by Output-Only Analysis”, In Proceedings of the 23rd International Modal Analysis Conference (IMAC XXIII), Orlando, Florida, USA, paper no. 63.
24. Giraldo, D., Caicedo, J., and Dyke, S.. “Experimental phase II of the structural health monitorino benchmark problem”. In EM '03: Proceedings of the 16th ASCE Engineering Mechanics Conference, 2003.
25. Grisso, B. J., Martin, L. A. and Inman, D. J, 2005, “A Wireless Active Sensing System for Impedance-based Structural Health Monitoring”, In Proceedings of the 23rd International Conference on Modal Analysis (IMAC XXIII), Orlando, Florida, USA, paper no. 236.
26. Grisso, B. J., Peairs, D. M. and Inman, D. J, 2004, “Impedance-based Health Monitoring of Composites”, In Proceedings of the 22nd Inter26 The Shock and Vibration Digest / July 2006 national Modal Analysis Conference (IMAC XXII), Dearborn, Michigan, USA, paper no. 274.
27. Gu, L. and Stankovic, J. A.. “Radio-triggered wake-up for wireless sensor networks”. *Real-Time Syst.*, 29(2-3):157–182, 2005.
28. Guillaume, P., Hermans, L. and Van der Auweraer, H, 1999, “Maximum Likelihood Identification of Modal Parameters from Operational Data”, In Proceedings of the 17th International Modal Analysis Conference (IMAC XVII), Kissimmee, Florida, USA, 1887–1893.
29. Guillaume, P., Parloo, E., Verboven, P. and De Sitter, G, 2002, “An Inverse Method for the Identification of Localized Excitation Sources”, In Proceedings of the 20th International Modal Analysis Conference (IMAC XX), Los Angeles, California, USA, 1382–1388.
30. Guo, N. and Cawley, P. “The interaction of Lamb waves with delaminations in composite laminates”. *Acoustical Society of America Journal*, 94:2240–2246, Oct. 1993.
31. Hailu, B., Hayward, G., Gachagan, A., McNab, A., and Farlow, R.. “Comparison of different piezoelectric materials for the design of embedded transducers for structural health monitoring applications”. In *Ultrasonics Symposium*, 2000 IEEE,

pages 1009–1012, Oct. 2000.

32. Huang, J.-H., Amjad, S., and Mishra, S.. “Cenwits: a sensor-based loosely coupled search and rescue system using witnesses”. In *SenSys '05: Proceedings of the 3rd international conference on Embedded networked sensor systems*, pages 180–191, New York, NY, USA, 2005. ACM Press.
33. Intanagonwiwat, C., Govindan, R., Estrin, D., Heidemann, J., and Silva, F. “Directed diffusion for wireless sensor networking”. *IEEE/ACM Trans. Netw.*, 11(1):2–16, 2003.
34. Jiang, X., Polastre, J., and Culler, D.. “Perpetual environmentally powered sensor networks”. In *IPSN '05: Proceedings of the 4th international symposium on Information processing in sensor networks*, page 65, Piscataway, NJ, USA, 2005. IEEE Press.
35. Kessler, S. S., Spearing, S. M. and Atalla, M. J, 2002, “In-Situ Damage Detection of Composites Structures using Lamb Waves Methods”, In *Proceedings of the 1st European Workshop on Structural Health Monitoring*, Ecole Normale Supérieure, Cachan, Paris, France, 374–381.
36. Lamb, H, 1917, “On Waves in an Elastic Plate”, *Proceedings of the Royal Society of London, Containing Papers of a Mathematical and Physical Character*, Vol. 93, 293-312.
37. Lee, B. C. and Staszewski, W. J, 2002, “Modelling of Acousto-ultrasonic Wave Interaction with Defects in Metallic Structures”, In *Proceedings of the International Conference on Noise and Vibration Engineering (ISMA 2002)*, Leuven, Belgium, 319–327.
38. Legendre, S., Massicotte, D., Goyette, J., and Bose, T.. “Wavelet-transform-based method of analysis for Lamb-wave ultrasonic nde signals”. In *Instrumentation and Measurement*, *IEEE Transactions on*, pages 524–530, June 2000.
39. Lehfeldt, E. and Holler, P.. “Lamb waves and lamination detection”. *Ultrasonics*, 5:255, 1967.
40. Li, B., Chen, X. F., Ma, J. X. and He, Z. J, 2005, “Detection of Crack Location and Size in Structures using Wavelet Finite Element Methods”, *Journal of Sound and Vibration*, Vol. 285, 767–782.
41. Li, H. C. H., Weis, M., Herszberg, I. and Mouritz, A. P, 2004, “Damage Detection in a Fiber Reinforced Composite Beam using Random Decrement Signatures”, In *Proceedings of the 2nd European Workshop on Structural Health Monitoring*, Forum am Deutschen Museum, Munich, Germany, 1091–1098.

42. Liang, C., Sun, F. P. and Rogers, C. A., 1994, "An Impedance Method for Dynamic Analysis of Active Material System", *Journal of Vibration and Acoustics*, Vol. 116, 121–128.
43. Lifshitz, J. M. and Rotem, A., 1969, "Determination of Reinforcement Unbonding of Composites by a Vibration Technique", *Journal of Composite Materials*, Vol. 3, 412–423.
44. Lin, M., 1999, "Development of SMART Layer for Built-In Structural Diagnostics", *Structural Health Monitoring 2000*, Stanford University, Palo Alto, California, USA, 603–611.
45. Lin, S., Yang, J. N. and Zhou, L., 2005, "Damage Identification of a Benchmark Building for Structural Health Monitoring", *Smart Materials and Structures*, Vol. 14, 162–169.
46. Lynch, J., Law, K., Kiremidjian, A., Kenny, T., Carryer, E., and Partridge, A.. "The design of a wireless sensing unit for structural health monitoring". In *Proceedings of the 3rd International Workshop on Structural Health Monitoring*, Stanford, CA, 2001.", 2001.
47. Mascarenas, D. L., Todd, M. D., Park, G., and Farrar, C. R.. "A miniaturized electromechanical impedance-based node for the wireless interrogation of structural health". In *Proceedings of SPIE – Health Monitoring and Smart Nondestructive Evaluation of Structural and Biological Systems*, Mar. 2006.
48. Maxstream Xbee radio <http://www.maxstream.com>.
49. Messina, A., Jones, I. A. and Williams, E. J., 1992, "Damage Detection and Localization using Natural Frequency Changes", In *Proceedings of the 1st Conference on Identification*, Cambridge, England, UK, 67–76.
50. Messina, A., Jones, I. A. and Williams, E. J., 1996, "Damage Detection and Localization using Natural Frequency Changes", In *Proceedings of the Conference on Identification in Engineering Systems*, Swansea, Wales, UK, 67–76.
51. Messina, A., Williams, E. J. and Contursi, T., 1998, "Structural Damage Detection by a Sensitivity and Statistical-Based Method", *Journal of Sound and Vibration*, Vol. 216, 791–808.
52. Micrel Group, University of Bologna <http://www-micrel.deis.unibo.it/>
53. Montalvao, D., Maia, N. M., and Ribeiro, A. M.. "A Review of Vibration-based Structural Health Monitoring with Special Emphasis on Composite Materials". *The Shock and Vibration Digest*, 38(4):295–324, 2006.
54. Moura Jr., J. R. V., 2005, "Damage Detection Techniques for Aeronautic

Structures”, In Proceedings of the 23rd International Modal Analysis Conference (IMAC XXIII), Orlando, Florida, USA, paper no. 101.

55. Moura Jr., J. R. V. and Steffen Jr., V, 2004, “Impedance-based Health Monitoring: Frequency Band Evaluation”, In Proceedings of the 22nd International Modal Analysis Conference (IMAC XXII), Dearborn, Michigan, USA, paper no. 161.

56. Pakzad, S., Kim, S., Fenves, G., Glaser, S., Culler, D., and Demmel, J.. “Multi-purpose wireless accelerometers for civil infrastructure monitoring”. In Proceedings of the 5th International Workshop on Structural Health Monitoring, 2005.

57. Park, C. and Chou, P. H.. “Ambimax: Efficient, autonomous energy harvesting system for multiple-supply wireless sensor nodes”. In SECON '06: Proceedings of Third Annual IEEE Communications Society Conference on Sensor, Mesh, and Ad Hoc Communications and Networks, 2006.

58. Park, G., Cudney, H. and Inman, D. J, 1999, “Impedance-Based Health Monitoring Technique for Civil Structures”, Structural Health Monitoring 2000, Stanford University, Palo Alto, California, USA, 523–532.

59. Park, G., Rutherford, A. C., Sohn, H. and Farrar, C. R, 2005, “An Outlier Analysis Framework for Impedance-based Structural Health Monitoring”, Journal of Sound and Vibration, Vol. 286, 229–250.

60. Park, G., Sohn, H., Farrar, C. R. and Inman, D. J, 2003, “Overview of Piezoelectric Impedance-Based Health Monitoring and Path-Forward”, Shock and Vibration Digest, Vol. 35, 451–463.

61. Park, J. M., Kong, J. W., Kim, D. S. and Yoon, D. J, 2005, “Nondestructive Damage Detection and Interfacial Evaluation of Single-Fibers/ Epoxy Composites using PZT, PVDF and P(VDF-TrFE) Copolymer Sensors”, Composites Science and Technology, Vol. 65, 241–256.

62. Park, N. G. and Park, Y. S, 2003, “Damage Detection using Spatially Incomplete Frequency Response Functions”, Mechanical Systems and Signal Processing, Vol. 17, 519–532.

63. Park, Y. S., Park, H. S. and Lee, S. S, 1988, “Weighted-Error-Matrix Application to Detect Stiffness Damage by Dynamic-Characteristic Measurement”, Modal Analysis: The International Journal of Analytical and Experimental Modal Analysis, Vol. 3, 101–107.

64. Peairs, D. M., Park, G. and Inman, D. J, 2005, “Impedance-based SHM Modeling using Spectral Elements and Circuit Analysis”, In Proceedings of the 23rd

International Modal Analysis Conference (IMAC XXIII), Orlando, Florida, USA, paper no. 232.

65. PowerFilm Inc <http://www.powerfilmsolar.com>

66. Pyles, R., 2003, Aging Aircraft: USAF Workload and Material Consumption Life Cycle Patterns, RAND, Arlington, VA.

67. Raghunathan, V. and Chou, P. H.. “Design and power management of energy harvesting embedded systems”. In ISLPED '06: Proceedings of the 2006 international symposium on Low power electronics and design, pages 369–374, New York, NY, USA, 2006. ACM Press.

68. Raghunathan, V., Kansal, A., Hsu, J., Friedman, J. and Srivastava, M.. “Design considerations for solar energy harvesting wireless embedded systems”. In IPSN '05: Proceedings of the 4th international symposium on Information processing in sensor networks, page 64, Piscataway, NJ, USA, 2005. IEEE Press.

69. Ricci, F., Banerjee, S. and Mal, A. K., 2004, “Health Monitoring of Composite Structures using Wave Propagation Data”, In Proceedings of the 2nd European Workshop on Structural Health Monitoring, Forum am Deutschen Museum, Munich, Germany, 1035–1042.

70. Rodrigues, J., Brincker, R. and Andersen, P., 2004, “Improvement of Frequency Domain Output-Only Modal Identification from the Application of the Random Decrement Technique”, In Proceedings of the 22nd International Modal Analysis Conference (IMAC XXII), Dearborn, Michigan, USA, paper no. 92

71. Ruotolo, R. and Surace, C., 1997, “Damage Detection Using Singular Value Decomposition”, In Proceedings of the 2nd International Conference on Structural Damage Assessment using Advanced Signal Processing Procedures (DAMAS 97), Sheffield, England, UK, 87–96.

72. Rytter, A., 1993,, “Vibration Based Inspection of Civil Engineering Structures”, PhD Thesis, Aalborg University, Denmark.

73. Sampaio, R. P. C., Maia, N. M. M. and Silva, J. M. M., 2003, “The Frequency Domain Assurance Criterion as a Tool for Damage Identification”, In Proceedings of the 5th International Conference on Damage Assessment of Structures (DAMAS 2003), Southampton, England, UK, 69–76.

74. Sanders, D. R., Kim, Y. I. and Stubbs, N., 1992, “Nondestructive Evaluation of Damage in Composite Structures Using Modal Parameters”, Experimental Mechanics, Vol. 32, 240–251.

75. Saravanos, D. A., Birman, V. and Hopkins, D. A., 1994, “Detection of

Delaminations in Composite Beams using Piezoelectric Sensors”, In Proceedings of the 31st AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference, 181–191.

76. Savov, K. and Wenzel, H., 2005, “System Identification and Damage Detection using Wavelet Analysis: Applications in Frame Structures”, In Proceedings of the 1st International Operational Modal Analysis Conference (IOMAC 2005), Copenhagen, Denmark, 511–518.

77. Shi, Z. Y. and Law, S. S., 2005a, “Parameter Identification-Based Damage Detection for Nonlinear Systems”, In Proceedings of the 1st International Operational Modal Analysis Conference (IOMAC 2005), Copenhagen, Denmark, 503–510.

78. Shi, Z. Y. and Law, S. S., 2005b, “Parameter Identification-Based Damage Detection for Linear Time-Varying Systems”, In Proceedings of the 1st International Operational Modal Analysis Conference (IOMAC 2005), Copenhagen, Denmark, 527–534.

79. Silva, A. J. P. F., 2001, “Propagação de Delaminação em Compósitos de Matriz Polimérica”, PhD Thesis, Instituto Superior Técnico, Technical University of Lisbon, Portugal (in Portuguese).

80. Silva, J. M. M. and Gomes, A. J. M. A., 1990, “Experimental Dynamic Analysis of Cracked Free-Free Beams”, *Experimental Mechanics*, Vol. 30, 20–25.

81. Silva, J. M. M. and Gomes, A. J. M. A., 1991, “Crack Modelling Using Torsional Springs”, *Actas das Quartas Jornadas de Fractura*, Sociedade Portuguesa de Materiais, Instituto Superior Técnico, Technical University of Lisbon, Portugal (in Portuguese).

82. Silva, J. M. M. and Gomes, A. J. M. A., 1994, “Crack Identification of Simple Structural Elements through the use of Natural Frequency Variations: The Inverse Problem”, In Proceedings of the 12th International Modal Analysis Conference (IMAC XII), Honolulu, Hawaii, USA, 1728–1735.

83. Simjee, F. and Chou, P. H., “Everlast: long-life, supercapacitor-operated wireless sensor node”. In ISLPED '06: Proceedings of the 2006 international symposium on Low power electronics and design, pages 197–202, New York, NY, USA, 2006. ACM Press.

84. Simmers Jr., G., Sodano, H., Park, G. and Inman, D. J., 2005, “Impedance based Structural Health Monitoring to Detect Corrosion”, In Proceedings of the 23rd International Modal Analysis Conference (IMAC XXIII), Orlando, Florida, USA,

paper no. 56.

85. Simmers, J., Garnett E., Hodgkins, J. R., Mascarenas, D. D., Park, G., and Sohn, H. "Improved piezoelectric self-sensing actuation". *Journal of Intelligent Material Systems and Structures*, 15(12):941–953, 2004.
86. Smith, T., Mars, J., and Turner, G.. "Using supercapacitors to improve battery performance". In *Power Electronics Specialists Conference*, 2002.
87. Sohn, H. and Farrar, C. R., 2001, "Damage Diagnosis using Time Series Analysis of Vibration Signals", *Smart Materials and Structures*, Vol. 10, 446–451.
88. Sohn, H., Farrar, C. R., Hemez, F. M., Shunk, D. D., Stinemates, D. W. And Nadler, B. R., 2003, "A Review of Structural Health Monitoring Literature: 1996–2001", Los Alamos National Laboratory, USA.
89. Sohn, H., Farrar, C. R., Hunter, N. F. and Worden, K., 2001, "Structural Health Monitoring using Statistical Pattern Recognition Techniques", *Journal of Dynamic Systems Measurement and Control-Transactions of the ASME*, Vol. 123, 706–711.
90. Sohn, H., Park, G., Wait, J. R., Limback, N. P. and Farrar, C. R., 2004, "Wavelet-based Signal Processing for Detecting Delamination in Composite Plates", *Smart Materials and Structures*, Vol. 13, 153–160.
91. Sohn, H., Park, H., Law, K. H. and Farrar, C. R., 2005, "Instantaneous Online Monitoring of Unmanned Aerial Vehicles without Baseline Signals", In *Proceedings of the 23rd International Modal Analysis Conference (IMAC XXIII)*, Orlando, Florida, USA, paper no. 259.
92. Sohn, H., Wait, J. R., Park, G. and Farrar, C. R., 2004a, "Multi-Scale Structural Health Monitoring for Composite Structures", In *Proceedings of the 2nd European Workshop on Structural Health Monitoring*, Forum am Deutschen Museum, Munich, Germany, 721–729.
93. Spearing, S., and Soutis, C.. "Damage detection in composite materials using Lamb wave methods". In *Proceedings of the American Society for Composites*, Sept. 2001.
94. Stubbs, N. and Osegueda, R., 1990a, "Global Non-Destructive Damage Evaluation in Solids", *Modal Analysis: The International Journal of Analytical and Experimental Modal Analysis*, Vol. 5, 67–79.
95. Stubbs, N. and Osegueda, R., 1990b, "Global Damage Detection in Solids–Experimental Verification", *Modal Analysis: The International Journal of Analytical and Experimental Modal Analysis*, Vol. 5, 81–97.
96. Su, Z., Ye, L. and Bu, X., 2002, "Evaluation of Delamination in Laminated

Composites based on Lamb Waves Methods: FEM Simulation and Experimental Verification”, In Proceedings of the 1st European Workshop on Structural Health Monitoring, Ecole Normale Supérieure, Cachan, Paris, France, 328–335.

97. Sundararaman, S., Adams, D. E. and Rigas, E. J, 2005, “Characterizing Damage in Plates through Beamforming with Sensor Arrays”, In Proceedings of the 23rd International Modal Analysis Conference (IMAC XXIII), Orlando, Florida, USA, paper no. 249.

98. Testa, R. B, 2005, “Characteristics and Detection of Damage and Fatigue Cracks”, Advanced Course on Dynamic Methods for Damage Detection in Structures, CISM (Centre International des Sciences Mécaniques), Udine, Italy.

99. Thien, A. B.. “Pipeline structural health monitoring using micro-fiber composite active sensors”. Master’s thesis, University of Cincinnati, 2006.

100. TMS320F1812 <http://www.ti.com>

101. Toyama, N. and Okabe, T, 2004, “Effects of Tensile Strain and Transverse Cracks on Lamb-Wave Velocity in Cross-Ply FRP Laminates”, *Journal of Materials Science*, Vol. 39, 7365–7367.

102. Vanlanduit, S., Parloo, E., Cauberghe, B., Guillaume, P. and Verboven, P, 2005, “A Robust Singular Value Decomposition for Damage Detection under Changing Operating Conditions and Structural Uncertainties”, *Journal of Sound and Vibration*, Vol. 284, 1033–1050.

103. Wang, Y., Lynch, J. P., and Law, K. H.. “Design of a low-power wireless structural monitoring system for collaborative computational algorithms”. In Health Monitoring and Smart Nondestructive Evaluation of Structural and Biological Systems IV. Proceedings of the SPIE, Volume 5768, pp. 106-117 (2005)., pages 106–117, May 2005.

104. Xu, N., Rangwala, S., Chintalapudi, K. K., Ganesan, D., Broad, A., Govindan, R., and Estrin, D.. “A wireless sensor network for structural monitoring”. In SenSys ’04: Proceedings of the 2nd international conference on Embedded networked sensor systems, pages 13–24, New York, NY, USA, 2004. ACM Press.

105. Yang, J. N., Lei, Y., Pan, S. and Huang, N, 2003, “Identification of Linear Structures Based on Hilbert-Huang Transform Part II: Complex Modes”, *Earthquake Engineering and Structural Dynamics*, Vol. 32, 1443–1467.

106. Zabel, V, 2004, “A Wavelet-Based Damage Detection Indicator for Reinforced Concrete Structures”, In Proceedings of the 22nd International Modal Analysis Conference (IMAC XXII), Dearborn, Michigan, USA, paper no. 70.

107. Zhang, R., King, R., Olson, L. and Xu, Y. L, 2005, “Dynamic Response of the Trinity River Relief Bridge to Controlled Pile Damage: Modelling and Experimental Data Analysis Comparing Fourier and Hilbert- Huang Techniques.” *Journal of Sound and Vibration*, Vol. 285, 1049– 1070.
108. Zou, Y., Tong, L., and Steven, G.. “Vibration-based model-dependant damage (Delamination) identification and health monitoring for composite structures - A review”. *Journal of Sound and Vibration*, 2:357–378, 2000.

Acknowledgements

I would like to thank a number of people who contributed to make this thesis possible. First of all, I'd like to thank Prof. Luca Benini for believing in my skills and providing me with the great occasion to work at the University of California, San Diego. Prof. Tajana Simunic Rosing with the Department of Computer Science and Engineering (CSE) at UCSD deserves my thankfulness as well. She made me enter a great research project and gave me advice whenever I faced a problem.

I also feel obliged to thank all people in the California Institute for Telecommunications and Information Technology (CalIT²): David Mascarenas, Kaisen Lin and Eric Flynn. They welcomed me in their team and helped me not to miss home, by sharing endless nights in the lab and introducing me to the questionable merits of American food and life-style. They all are the most important thing my stay at UCSD gave me: real friends.

Next, I would like to acknowledge the role of all people at the CSE Department, who were very helpful since the first day I entered the department. Virginia McIlwain's and Robin Vespertino's help with administrative matters was very important. Christy Flowers and Sheila Manalo deserve a special thank for their endless patience with my packages' deliveries.

I also want to thank all other people I met in San Diego and help me enjoy my stay. In particular, my roommates Chintan Jain and Parag Dixit, who allowed me to invade most of their cupboard with my pasta, and Edoardo Regini, who shared some real Italian lunches with me.

Most of all I want to thank my family: my father Fabio, my mother Tatiana and my little sister Adele. They always have given me their support.

Finally, I will never thank Laura enough for all the beautiful moments we shared. You are the most precious person in my life.